

Demand-Driven Compositional Symbolic Execution

Saswat Anand^{*1}, Patrice Godefroid², and Nikolai Tillmann²

¹ Georgia Institute of Technology, saswat@cc.gatech.edu

² Microsoft Research, {pg,nikolait}@microsoft.com

Abstract. We discuss how to perform symbolic execution of large programs in a manner that is both compositional (hence more scalable) and demand-driven. Compositional symbolic execution means finding feasible interprocedural program paths by composing symbolic executions of feasible intraprocedural paths. By demand-driven, we mean that as few intraprocedural paths as possible are symbolically executed in order to form an interprocedural path leading to a specific target branch or statement of interest (like an assertion). A key originality of this work is that our demand-driven compositional interprocedural symbolic execution is performed entirely using first-order logic formulas solved with an off-the-shelf SMT (Satisfiability-Modulo-Theories) solver – no procedure in-lining or custom algorithm is required for the interprocedural part. This allows a uniform and elegant way of summarizing procedures at various levels of detail and of composing those using logic formulas.

We have implemented a prototype of this novel symbolic execution technique as an extension of Pex, a general automatic testing framework for .NET applications. Preliminary experimental results are encouraging. For instance, our prototype was able to generate tests triggering assertion violations in programs with large numbers of program paths that were beyond the scope of non-compositional test generation.

1 Introduction

Given a sequential program P with input parameters \vec{I} , the *test generation problem* consists in generating automatically a set of input values to exercise as many program statements as possible. There are essentially two approaches to solve this problem. *Static* test generation [15, 22, 7] consists in analyzing the program P statically, using symbolic execution techniques to attempt to compute inputs to drive P along specific paths or branches, but without ever executing the program. In contrast, *dynamic test generation* [16, 11, 5] consists in executing the program, typically starting with some random inputs, while simultaneously performing a symbolic execution to collect symbolic constraints on inputs obtained from predicates in branch statements along the execution, and then using a constraint solver to infer variants of the previous inputs in order to steer program

* The work of this author was done mostly while visiting Microsoft Research.

executions along alternative program paths. Since dynamic test generation extends static test generation with additional runtime information, it can be more powerful [11, 10], and is therefore used as the basis of this work.

As recently pointed out [10], automatic test generation (whether static or dynamic) does not scale to large programs with many feasible program paths, *unless* test generation is performed *compositionally*. Inspired by interprocedural static analysis, compositional test generation consists in encoding test results of lower-level functions with test *summaries*, expressed using preconditions over function inputs and postconditions over function outputs, and then re-using those summaries when testing higher-level functions. In contrast with traditional interprocedural static analysis, the framework introduced in [10] involves detailed summaries where function preconditions and postconditions are represented using logic formulas, and the interprocedural analysis (test generation) is performed using an automated theorem prover. A key component of this approach is thus *compositional symbolic execution*: how to find feasible interprocedural program paths by composing symbolic executions of feasible intraprocedural paths, represented as logic “summaries”.

In this paper, we develop compositional symbolic execution further. We present a detailed formalization of how to generate first-order logic formulas with uninterpreted functions in order to represent function summaries and allow compositional symbolic execution using a SMT (Satisfiability-Modulo-Theories) solver. Our formalization generalizes the one of [10] as it allows incomplete summaries (which correspond to only a subset of all paths of a function) to be expanded *lazily* on a *demand-driven* basis, instead of being expanded in the fixed “innermost-first” order described in [10]. With demand-driven symbolic execution, as few intraprocedural paths as possible are symbolically executed in order to form an interprocedural path leading to a specific target branch or statement of interest (like an assertion). This increased flexibility also allows test generation to adapt dynamically, as more statements get covered, in order to focus on those program statements that are still uncovered. In practice, real-life software applications are very complex, and allowing the search to be demand-driven is often key to reach a specific target in a reasonable time. It is also useful for selective regression testing aimed at generating tests targeted to cover new code embedded in old code.

We have implemented a prototype of demand-driven compositional symbolic execution as an extension of Pex [20], a general automatic testing framework for .NET applications. Preliminary experimental results are encouraging. For instance, our prototype implementation was able to generate tests triggering assertion violations in programs with large numbers of program paths that were beyond the scope of non-compositional test generation.

2 Background

We assume we are given a sequential program P with input parameters \vec{I} . *Symbolic execution* of P means symbolically exploring the tree \mathcal{T} defined by the

execution paths of the program when considering all possible value assignments to input parameters. For each execution path ρ , i.e., a sequence of statements executed by the program, a *path constraint* ϕ_ρ is constructed that characterizes the input assignments for which the program executes along ρ . Each variable appearing in ϕ_ρ is thus a program input, while each constraint is expressed in some theory T decided by a constraint solver (for instance, including linear arithmetic, bit-vector operations, etc.). A constraint solver is an automated theorem prover which also returns a satisfying assignment for all variables appearing in formulas it can prove satisfiable. All program paths can be enumerated by a search algorithm that explores all possible branches at conditional statements. The paths ρ for which ϕ_ρ is satisfiable are *feasible* and are the only ones that can be executed by the actual program. The solutions to ϕ_ρ exactly characterize the inputs that drive the program through ρ . Assuming that the constraint solver used to check the satisfiability of all formulas ϕ_ρ is sound and complete, this use of symbolic execution for programs with finitely many paths amounts to program verification.

In practice, symbolic execution of large programs is bound to be imprecise due to complex program statements (pointer manipulations, floating-point operations, etc.) and calls to operating-system and library functions that are hard or impossible to reason about symbolically with good enough precision at a reasonable cost. Whenever precise symbolic execution is not possible during dynamic test generation, concrete values can be used to simplify constraints and carry on with a simplified, partial symbolic execution [11].

Systematically executing symbolically *all* feasible program paths does not scale to large programs. Indeed, the number of feasible paths can be exponential in the program size, or even infinite in presence of loops with unbounded number of iterations. This *path explosion* can be alleviated by performing symbolic execution *compositionally* [10].

Let us assume the program P consists of a set of functions. In what follows, we use the generic term of *function* to denote any part of the program P whose observed behaviors are summarized; obviously, any other kinds of program fragments such as arbitrary program blocks or object methods can be treated as “functions” as done in this paper. To simplify the presentation, we assume that the functions in P do not perform recursive calls, and that all the executions of P terminate. (These assumptions do not prevent P from possibly having infinitely many executions paths, as is the case if P contains a loop whose number of iterations may depend on some unbounded input.)

In compositional symbolic execution [10], a function summary ϕ_f for a function f is defined as a formula in propositional logic whose propositions are constraints expressed in some theory T . ϕ_f can be derived by successive iterations and defined as a disjunction of formulas ϕ_w of the form $\phi_w = pre_w \wedge post_w$, where pre_w is a conjunction of constraints on the inputs of f while $post_w$ is a conjunction of constraints on the outputs of f . ϕ_w can be computed from the path constraint corresponding to the execution path w as described later. An input to a function f is any value that can be read by f in some of its execu-

```

int abs(int x){
    if(x > 0) return x;
    else if(x == 0)
        return 100;
    else return -x;
}

int testAbs(int p, int q){
    int m = abs(p);
    int n = abs(q);
    if(m > n && p > 0)
        assert false; //target
}

```

Fig. 1. Example program

tions, while an output of f is any value written by f in some of its executions and later read by P after f returns. To simplify the presentation, we assume in what follows that each function takes a fixed number of arguments as inputs and returns a single value.

3 Motivating Example and Overview

To illustrate the motivation for *demand-driven* compositional symbolic execution, consider the simple program in Fig. 1, which consists of a top-level function `testAbs` which calls another function `abs`. Intraprocedural execution trees for each function are shown in Fig. 2. Each node in such trees represents the execution of a program statement such that a path from the root of the tree to a leaf corresponds to an intraprocedural path. Each such path can be identified by its leaf node. Edges in execution trees are labeled with constraints expressed in terms of the function inputs. The conjunction of constraints labeling the edges of a path represents its associated path constraint as defined earlier. For example, Fig. 2(a) shows the (partial) execution tree of function `abs`, shown in Fig. 1, after the execution of `abs` with a single input $x=1$. In what follows, we call a node *dangling* if it represents a path that has not been exercised yet. For example, after executing the `abs` with input $x=1$, any path on which the input is less than or equal to 0 is not exercised. In Fig. 2(a), the sole dangling node is denoted by a circle.

The demand-driven compositional symbolic execution we develop in this work has two key properties: given a specific target to cover, it tries to (1) explore as few paths as possible (called *lazy exploration*) and to (2) avoid exploring paths that can be guaranteed not to cover the target (called *relevant exploration*). We now illustrate these two features.

Lazy Exploration Assume that we first run the program of Fig. 1 by executing the function `testAbs` with $p=1$ and $q=1$. This first execution will exercise the then branch of the first conditional statement in `abs` (node 3), as well as the else branch of the conditional statement in `testAbs` (node 10). The execution trees of `abs` and `testAbs` resulting from this execution are shown in Fig. 2(a) and (c), respectively. Suppose we want to generate a test input to cover node 11, corresponding to the assertion in `testAbs`. The search ordering described in [10] is not target-driven and would attempt to next exercise the unexplored paths in

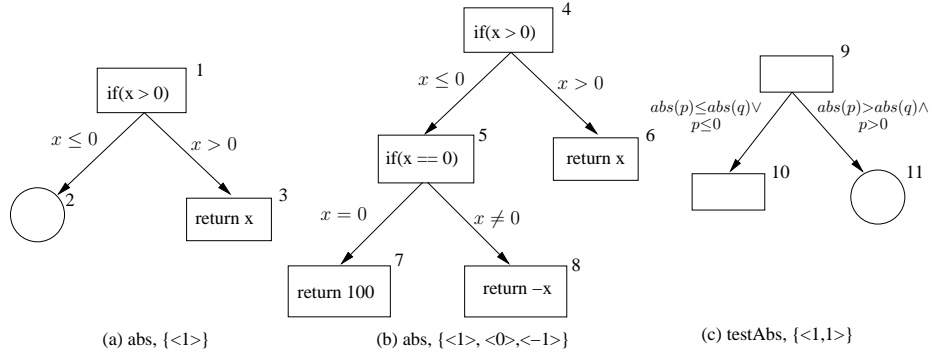


Fig. 2. Execution trees for the functions `abs` and `testAbs` from Fig. 1. Each execution tree represents paths exercised by a set of test inputs shown as vectors inside the curly braces.

the innermost, lower-level function `abs`. In contrast, the more flexible formalization introduced in the next section allows us to check whether a combination of currently-known fully-explored intraprocedural paths are sufficient to generate a new test input covering the target node. In this example, this is the case as the assertion can be reached in `testAbs` without exploring new paths in `abs`, for instance with values $p=2$ and $q=1$.

Relevant Exploration Now, assume we first execute the program with inputs $p=0$ and $q=1$. Suppose our target is again node 11 corresponding to the assert statement. From the condition guarding the assert statement, observe that any combination of input values for p and q where p has a non-positive value has no chance to cover the target. As we will see, our proposed algorithm is able to infer such information automatically from the previous execution with inputs $p=0$ and $q=1$, and will thus prune automatically the entire sub-search tree where p is not greater than 0.

4 Demand-Driven Compositional Symbolic Execution

4.1 Main Algorithm

Algorithm 1 outlines our test-generation algorithm. Given a program P , Algorithm 1 iteratively computes a set of test inputs to cover all reachable statements in P . The algorithm starts with an empty set of intraprocedural execution trees, and a random program input. It performs two steps in sequence until all reachable nodes in the program have been explored. (1) Function `Execute` executes the program with some test input, both normally and symbolically. During symbolic execution of the specific path exercised by the test input, new nodes and edges with constraint labels are added to the intraprocedural execution trees of the individual program functions being executed, while dangling nodes, used as place-holders along this specific path in previous executions, become regular

```

input : Program P
output: A set of test inputs
exTrees ← emptyExTree ;
input ← RandomInput();
repeat
  if input ≠ emptyModel then
    exTrees ← Execute(P, input, exTrees);
    OutputTest(input);
  else
    RemoveDanglingNode(n);
  end
  n ← ChooseDanglingNode(exTrees);
  if n ≠ nil then
    input ← FindTestInput(exTrees, n);
  end
until n = nil ;
return;

```

Algorithm 1: Test-input generation algorithm

nodes. (2) `ChooseDanglingNode` chooses a dangling node as the next target to be covered, using any heuristic (search strategy). If there is no dangling node remaining, the algorithm terminates. Otherwise, `FindTestInput` computes a test input to cover the target, as will be described next.

4.2 Compositional Symbolic Execution

In compositional symbolic execution, the condition under which a node in a function can be reached from the program’s entry point is the conjunction of (1) the condition under which the function’s entry node can be reached, referred to as *calling context*; and (2) the condition under which the node can be reached within its function, referred to as the *local (intraprocedural) path constraint*.

Local (Intraprocedural) Path Constraint. The local path constraint of a node n in the intraprocedural execution tree \mathcal{T}_f of function f is defined as the path constraint of the path w from the entry node of f to the statement represented by n . The local path constraint of node n , represented by $localpc(n)$, is expressed in terms of the input parameter symbols $\vec{\mathcal{P}}_f$ of f and represents a precondition $pre(w)$ for execution of the path w [10]. It is defined as follows.

$$localpc(n) := lpc_n \wedge \bigwedge_{\text{for each } g(\vec{a}) \text{ appearing in } lpc_n} D_g(\vec{a})$$

where lpc_n is the conjunction of constraints appearing on the edges of the path w from the root of \mathcal{T}_f to n , and each definition predicate $D_g(\vec{a})$ represents the (possibly partial) summary currently available for function g , called from f with \vec{a} as arguments, and mentioned in lpc_n . Definition predicates are formally defined as follows.

Definition predicate. When function f calls function g during symbolic execution, we treat the return value of the function call to g as a (fresh) symbolic input to f . We represent the return value by the expression $g(\vec{a})$, where \vec{a} are the arguments expressed in terms of $\vec{\mathcal{P}}_f$. If the return value is used in a conditional statement of f , then $g(\vec{a})$ appears in the path constraint. The function symbol g will be treated as an uninterpreted function symbol by the constraint solver, and we restrict possible interpretations by an axiom of the form $\forall x. g(x) = E[x]$, where $E[x]$ is an expression that may involve the bound variable x . As an example, for the **abs** function in Fig. 1, *abs* can be defined as follows (where *ITE* denotes the If-Then-Else construct):

$$\forall x. \text{abs}(x) = \text{ITE}(x > 0, x, \text{ITE}(x = 0, 100, -x))$$

However, return values on some paths of a function may be currently unknown since paths are explored incrementally and on-demand. In those cases, we cannot use the above encoding directly. We could use a special *undefined* value that represents the result of an unexercised path, and lift all operations accordingly. Instead, we use a *definition-predicate* D_g for each function symbol that represents the return value of a function call. We define this predicate with the axiom δ_g as follows.

$$\delta_g := \forall \vec{\mathcal{P}}_g. D_g(\vec{\mathcal{P}}_g) \Leftrightarrow \bigvee_{\text{leaf } l \text{ in } \mathcal{T}_g} \text{localpc}(l) \wedge \text{ret}(l)$$

where

$$\text{ret}(l) := \begin{cases} \mathcal{G}_l & \text{if } l \text{ is a dangling node} \\ g(\vec{\mathcal{P}}_g) = \text{Ret}_g(l) & \text{otherwise} \end{cases}$$

In the above definition, $\text{Ret}_g(l)$ represents the return value of g , which is an expression in terms of $\vec{\mathcal{P}}_g$, on the fully-explored intraprocedural path represented by l . For each dangling node d , \mathcal{G}_d represents an auxiliary boolean variable that uniquely corresponds to d ; we use these boolean variables in Sec. 4.3 to control the search by specifying whether the exploration of a new execution path through a dangling node is permissible.

For the example shown in Fig. 1, suppose we execute **testAbs** with $p = 1$ and $q = 1$. The execution trees for **abs** and **testAbs** obtained from this input are shown in Fig. 2(a) and (c) respectively. Now, the local path constraint of the node n , labeled 11 in the figure, will be as follows.

$$\text{localpc}(n) := \text{abs}(p) > \text{abs}(q) \wedge p > 0 \wedge D_{\text{abs}}(p) \wedge D_{\text{abs}}(q)$$

With the above input, since only the path where $x > 0$ has been explored in **abs**, there is a dangling node d , labeled 2, which represents the (unexplored) else branch of the conditional statement. The definition predicate D_{abs} is then defined by the following axiom.

$$\delta_{\text{abs}} := \forall x. D_{\text{abs}}(x) \Leftrightarrow \text{ITE}(x > 0, \text{abs}(x) = x, \mathcal{G}_d)$$

If all the paths of `abs` had been explored (as shown in Fig. 2(b)), its definition-predicate axiom would instead be as follows.

$$\begin{aligned} \delta_{abs} &:= \forall x. D_{abs}(x) \Leftrightarrow (x \leq 0 \wedge x = 0 \wedge abs(x) = 100) \\ &\quad \vee (x \leq 0 \wedge x \neq 0 \wedge abs(x) = -x) \\ &\quad \vee (x > 0 \wedge abs(x) = x) \\ &= \forall x. D_{abs}(x) \Leftrightarrow ITE(x \leq 0, ITE(x = 0, abs(x) = 100, abs(x) = -x), \\ &\quad abs(x) = x) \end{aligned}$$

Note that, with the specific innermost-first search order used in [10] for incrementally computing summaries, dangling and target nodes are always in the current innermost function in the call stack and the above formalization of partial summaries can then be simplified. In contrast, the formalization presented here is more general as it allows dangling nodes and target nodes to be located anywhere in the program.

Calling-context Predicate. The calling-context predicate associated with a function f describes under which conditions, and with which arguments, f can be reached. The calling-context predicate of function f , written as $C_f(\vec{a})$, evaluates to true iff on some program path f can be called with arguments \vec{a} . $C_f(\vec{a})$ is defined by the *calling-context axiom* γ_f as follows.

$$\gamma_f := \begin{cases} \forall \vec{a}. C_f(\vec{a}) \Leftrightarrow \vec{a} = \vec{I} & \text{if } f \text{ is entry function of program } P \\ \forall \vec{a}. C_f(\vec{a}) \Leftrightarrow \bigvee_{\substack{\text{for each} \\ \text{function } g \text{ in } P}} C_f^g(\vec{a}) & \text{otherwise} \end{cases}$$

with

$$C_f^g(\vec{a}) := \exists \vec{\mathcal{P}}_g. C_g(\vec{\mathcal{P}}_g) \wedge (knownC_f^g(\vec{a}) \vee unknownC_f^g)$$

where

$$\begin{aligned} knownC_f^g(\vec{a}) &:= \bigvee_{m \in \text{callsites}(\mathcal{T}_g, f)} \vec{a} = \text{args}(m) \wedge \text{localpc}(m) \\ unknownC_f^g &:= \bigvee_{\text{dangling node } d \text{ in } \mathcal{T}_g} \text{localpc}(d) \wedge \mathcal{G}_d \end{aligned}$$

We distinguish two cases in γ_f . First, if f is the entry function of the program P , then the arguments of f are the program inputs \vec{I} . Otherwise, $C_f(\vec{a})$ is true iff f can be called from some function g with arguments \vec{a} . $C_f^g(\vec{a})$ represents the condition under which g may call f with arguments \vec{a} . $C_f^g(\vec{a})$ in turn evaluates to true iff (1) g itself can be called with arguments $\vec{\mathcal{P}}_g$; and either (2.a) f can be called from g in a *known* call site denoted by $m \in \text{callsites}(\mathcal{T}_g, f)$ with arguments $\vec{a} = \text{args}(m)$, where $\text{args}(m)$ denote the arguments (in terms of $\vec{\mathcal{P}}_g$) passed to call at m ; or (2.b) f *might* be called (with unknown arguments) on a path in g , represented by a dangling node d , that has not been explored so far. In either case, the local path constraint $\text{localpc}(m)$ leading to the known call site m or $\text{localpc}(d)$ leading to a possible call site d , respectively, is appended as a condition necessary to reach the respective call site.

Consider again the program shown in Fig. 1 with `testAbs` as the top-level entry function. The calling-context predicate for `testAbs` is then defined by the following axiom.

$$\gamma_{testAbs} := \forall p, q. C_{testAbs}(p, q) \Leftrightarrow p = \vec{I}(0) \wedge q = \vec{I}(1).$$

For the function `abs`, the definition of the calling-context predicate is more complicated because `abs` can be called twice in `testAbs`. Suppose the execution trees of `abs` and `testAbs` are as shown in Fig. 2(b) and (c) respectively. For both known call-sites of `abs` in `testAbs`, where p and q are passed as arguments, `localpc` evaluates to `true`. And, there is one unknown call-site, which is represented by the dangling node d (labeled 11). For d , we have `localpc`(d) := `abs`(p) > `abs`(q) \wedge $p > 0 \wedge D_{abs}(p) \wedge D_{abs}(q)$. Now, $C_{abs}(a)$ is defined by the axiom γ_{abs} as follows.

$$\begin{aligned} \gamma_{abs} &:= \forall a. C_{abs}(a) \Leftrightarrow C_{abs}^{testAbs}(a) \\ C_{abs}^{testAbs}(a) &:= \exists p, q. C_{testAbs}(p, q) \wedge (a = p \vee a = q \\ &\quad \vee (\text{abs}(p) > \text{abs}(q) \wedge p > 0 \wedge D_{abs}(p) \wedge D_{abs}(q) \wedge \mathcal{G}_d)) \end{aligned}$$

Note that an existential quantification is used in C_f^g to limit the scope of parameter symbols \mathcal{P}_g to specific call-sites. However, this existential quantification can be eliminated by skolemization since it always appears within the scope of the universal quantifier in the definition of γ_f .

Also note that the formalization proposed in [10] does not require calling-context predicates because it only supports a fixed inner-most ordering in which intraprocedural paths are explored. Since we relax here the restriction on the exploration ordering so that paths can be explored in any order on-demand, calling-context predicates become necessary.

Interprocedural path constraint. Given a node n in the intraprocedural execution tree \mathcal{T}_f of a function f , path constraints of interprocedural paths leading to n are represented by Ψ_n , which is defined recursively as follows:

$$\Psi_n = \text{localpc}(n) \wedge C_f(\vec{\mathcal{P}}_f) \wedge \bigwedge_{C_g(\vec{a}) \text{ appears in } \Psi_n} \gamma_g \wedge \bigwedge_{g(\vec{a}) \text{ appears in } \Psi_n} \delta_g$$

Ψ_n represents the disjunction of path constraints of all interprocedural paths to target n that can be formed by joining intraprocedural paths, represented by execution trees of different functions. (Disjunctions arise from the definitions of γ_g and δ_g .) An intraprocedural path p in \mathcal{T}_f can be *joined* with an intraprocedural path q in \mathcal{T}_g , if either (1) p ends at a leaf node (possibly a dangling node) in \mathcal{T}_f , and q starts at a node in \mathcal{T}_g corresponding to a call-site of f in g ; or, (2) p ends at a node representing a call-site of g in f and q starts at the entry-node of \mathcal{T}_g ; or, (3) p ends at a dangling node, and q starts from the entry-node of \mathcal{T}_g , where g is any arbitrary function.

With compositional symbolic execution, the size of an interprocedural path constraint is linear in the sum of the sizes of the execution trees \mathcal{T}_f [10].

Examples. As our first example, suppose the execution trees for `abs` and `testAbs` are as shown in Fig. 2(b) and (c), respectively. If the target is the node labeled 11, then the interprocedural path constraint is as follows.

$$\begin{aligned} & abs(p) > abs(q) \wedge p > 0 \wedge D_{abs}(p) \wedge D_{abs}(q) \wedge p > 0 \wedge C_{testAbs}(p, q) \\ & \wedge \forall x. D_{abs}(x) \Leftrightarrow ITE(x \leq 0, ITE(x = 0, abs(x) = 100, abs(x) = -x), abs(x) = x) \\ & \wedge \forall p, q. C_{testAbs}(p, q) \Leftrightarrow p = \vec{I}(0) \wedge q = \vec{I}(1) \end{aligned}$$

As another example, suppose the execution trees for `abs` and `testAbs` are again as shown in Fig. 2(b) and (c), respectively. Now if the target is node labeled 2, the path constraint is as follows (where \mathcal{G}_{11} represents the unique boolean variable corresponding to the dangling node labeled 11):

$$\begin{aligned} & x \leq 0 \wedge C_{abs}(x) \\ & \wedge \forall x. D_{abs}(x) \Leftrightarrow ITE(x \leq 0, ITE(x = 0, abs(x) = 100, abs(x) = -x), abs(x) = x) \\ & \wedge \forall a. C_{abs}(a) \Leftrightarrow \exists p, q. C_{testAbs}(p, q) \wedge (a = p \vee a = q \\ & \quad \vee (abs(p) > abs(q) \wedge p > 0 \wedge D_{abs}(p) \wedge D_{abs}(q) \wedge \mathcal{G}_{11})) \\ & \wedge \forall p, q. C_{testAbs}(p, q) \Leftrightarrow p = \vec{I}(0) \wedge q = \vec{I}(1) \end{aligned}$$

4.3 Demand-Driven Symbolic Execution

In compositional symbolic execution, interprocedural paths are formed by combining intraprocedural paths. To allow compositional symbolic execution to be demand-driven, we allow in this work (unlike [10]) interprocedural paths to be formed by combining intraprocedural paths that end in dangling nodes. We call an interprocedural path *partially-explored* iff it goes through one or more dangling nodes; otherwise we call the path *fully-explored*. Note that a fully-explored path may end at, but not go through, a dangling node.

Algorithm 2 is used to find a feasible, interprocedural path from the entry of the program to a target node using demand-driven compositional symbolic execution. The algorithm corresponds to the subroutine `FindTestInput` in Algorithm 1. It takes as input a set of intraprocedural execution trees `exTrees`, and a dangling node `n` in one of these execution trees, which is the target to cover. It returns either (1) a designated value `emptyModel` representing the fact that the target node is unreachable, or (2) program inputs \vec{I} that exercises a path that may cover the target. The algorithm calls an SMT solver by invoking the function `FindModel(Ψ)`, which returns a model for the path constraint Ψ if it is satisfiable, or returns `emptyModel` otherwise. $G(\Psi)$ represents the set of all boolean flags that appear in the path constraint Ψ , each of which uniquely corresponds to a dangling node in `exTrees`. The algorithm first computes the interprocedural path constraint for the target node `n` in `exTrees` as presented in Sec. 4.2. Then it performs two steps, referred to as *lazy exploration* and *relevant exploration* in what follows.

Lazy Exploration In this step, the algorithm checks if it is possible to form a feasible, fully-explored, interprocedural path to `n` by combining only (fully-explored) intraprocedural paths in `exTrees`. To do so, it computes a constraint

<p>input : Set of execution trees <code>exTrees</code>, target node <code>n</code> to be covered</p> <p>output: Program inputs that <i>may</i> cover <code>n</code>, or <code>emptyModel</code> if the target is unreachable</p> <pre> $\Psi_n \leftarrow \text{InterprocPC}(n, \text{exTrees});$ input $\leftarrow \text{FindModel}(\Psi_n \wedge \bigwedge_{\mathcal{G}_d \in G(\Psi_n) \wedge d \neq n} \mathcal{G}_d = \text{false});$ if input = <code>emptyModel</code> then input $\leftarrow \text{FindModel}(\Psi_n);$ end return input ; </pre>

Algorithm 2: Demand-driven, compositional FindTestInput algorithm

that represents the disjunction of the path constraints of all such paths and checks its satisfiability. The new constraint is formed by conjoining Ψ_n with equations that set all variables but \mathcal{G}_n in $G(\Psi_n)$ to *false* so that all intraprocedural paths that end at a dangling node other than `n` are made infeasible. If the augmented constraint is satisfiable, `FindModel` returns a program test input that is guaranteed to cover the target (provided symbolic execution has perfect precision). Otherwise, we need to explore new partially-explored intraprocedural paths, which is done in the next step.

Relevant Exploration We say that a partially-explored, interprocedural path is *relevant* if it ends at the target. In other words, such a path starts at the program entry, goes through one or more dangling nodes, finally taking the path from the root node of \mathcal{T}_f to the target node `n`, where \mathcal{T}_f represents the execution tree of function f where `n` is located. In this second step, the algorithm checks if a feasible relevant path can be formed by combining all (both fully-explored and partially-explored) intraprocedural paths in `exTrees`. To do so, the algorithm checks satisfiability of Ψ_n with a second call to `FindModel`. If Ψ_n is unsatisfiable, the algorithm returns `emptyModel` representing unreachability of the target. Otherwise, it returns a program input that *might* exercise a path to the target. This time, the boolean variables in $G(\Psi_n)$ are not constrained to any specific value as is done in the previous step. As a result, the constraint solver assigns *true* to a boolean variable if the path to the corresponding dangling node is used to form the interprocedural path to the target. Such a relevant path is not guaranteed to reach the target, since the program’s behavior at dangling nodes, which may appear on a relevant path, is currently unknown.

The following theorems characterize the correctness of the above algorithms. These theorems hold *assuming symbolic execution has perfect precision*, i.e., that constraint generation and solving is both sound and complete for all program statements. (Proofs are omitted due to space limitations.)

Theorem (Relative Completeness) If Algorithm 2 returns `emptyModel`, then the target `n` is unreachable.

Theorem (Progress) If Algorithm 2 returns a program input \vec{T} (different from `emptyModel`), then the execution of the program with \vec{T} exercises a new intraprocedural path (i.e., at least one dangling node is removed from `exTrees`).

Theorem (Termination) If the program has a finite number of paths, Algorithm 1 terminates.

5 Preliminary Experiments

We have implemented a prototype of demand-driven compositional symbolic execution in Pex [20], a general automatic testing framework for .NET programs. Pex generates test inputs for parameterized unit tests [21] by performing a variation of dynamic [11] test generation using the SMT constraint solver Z3 [8]. Pex’ goal is to analyze as many feasible execution paths of a given .NET program as possible in a given amount of time. During the search, Pex picks the next target node using a scheduling algorithm that is fair between all dangling nodes. Pex is a comprehensive testing tool and framework, which has been used within Microsoft on several .NET applications and contributed to finding many bugs (including several security-critical ones) in released software and software still under development at Microsoft.

We present experiments with three programs written in C# using both non-compositional and demand-driven compositional symbolic execution. These experiments were conducted on a 3.4 GHz Pentium 4 with 2 GB memory.

HWM is program that takes a string as input, and an assertion fails if the input string contains all of the four substrings: “Hello”, “world”, “at”, “Microsoft!”. Although it is a simple program, it has hundreds of millions of feasible whole-program paths. The program has a main method that calls `contains(s,t)` four times in succession. `contains(s,t)` checks if string `s` contains substring `t`. `contains(s,t)` calls `containsAt(s,i,t)` that checks if `s` contains `t` starting from index `i` in `s`.

Parser is a parser for a subset of a Pascal-like language. The program takes a string as input, and successfully parses it if it represents a syntactically valid program in the language. An assertion is violated if parsing is successful. A valid program starts with the keyword “program” followed by an arbitrary string representing program name. Furthermore, the body of the program starts with keyword “begin” and end with keyword “end”. And the body may optionally include function definitions.

IncDec is a program that takes an integer as argument. It increments it several times and then decrements until a certain condition specified as an assertion is satisfied.

The table in Fig. 3 presents results of experiments. The three first columns represent the total number of executions, the total time taken over all executions, and the time taken per execution. (Execution time includes time taken by the constraint solver.) The last column shows whether the respective technique was able to generate an input that violates the assertion contained in each program.

Benchmark	No. of Executions		Time in sec time per execution				Exception found	
	new	old	new	old	new	old	new	old
HWM	37	maxed	65	705	1.75	0.02	yes	no
Parser	144	maxed	71	338	0.49	0.01	yes	yes
IncDec	74	1207	14	43	0.18	0.03	yes	yes

Fig. 3. Comparison between new (demand-driven, compositional) and old (non-compositional) symbolic execution techniques

In the column showing the number of executions, “maxed” denotes that non-compositional symbolic execution hits an upper bound of 20,000 executions; in those cases, total execution time represents the time taken to reach the upper bound.

We make the following observations from the table in Fig. 3. (1) The number of executions required with demand-driven compositional symbolic execution is often several orders of magnitude smaller compared to non-compositional symbolic execution. (2) The improvement in total time cannot be measured as non-compositional symbolic execution technique hits the upper bound on the number of execution in two of the three cases. (3) The time taken for each execution increases when the symbolic execution is demand-driven and compositional, as the formulas generated are more complicated and the constraint solver needs more time to solve those, although most can be solved in seconds. (4) In the case of HWM, only the search with demand-driven compositional symbolic execution is able to find the assertion violation, whereas the non-compositional search is lost in search-space due to path explosion. The other two examples have fewer execution paths, and the fair search heuristics implemented in Pex are able to find the assertion violations, even with non-compositional searches.

6 Other Related Work

Interprocedural *static* analysis always involves some form of summarization [19]. Summaries are usually defined either at some fixed-level of abstraction, e.g., for points-to analysis [17], or as abstractions of intraprocedural pre and postconditions, e.g., projections onto a set of predicates [3, 23]. Even when a SAT solver is used for a precise intraprocedural analysis [6, 23, 2], the interprocedural part of the analysis itself is carried out either using some custom fixpoint computation algorithm [4, 23] or by in-lining functions [6, 2], the latter leading to combinatorial explosion.

In contrast with prior work on interprocedural static analysis, we represent function summaries as uninterpreted functions with arbitrary pre/postconditions represented as logic formulas, and we use an SMT solver to carry out the interprocedural part of the analysis. Of course, the constraint solver may need to in-line summaries during its search for a model satisfying a whole-program path constraint, but it will do so lazily, only if necessary, and while memoizing

new induced facts in order to avoid re-inferring those later, hence simulating the effect of caching previously-considered calling contexts and new summaries inferred by transitivity, as in compositional algorithms for hierarchical finite-state machine verification [1].

How to perform *abstract* symbolic execution with simplified summary representations [14, 2, 12] in *static* program analysis is orthogonal to the demand-driven and compositionality issues addressed in our paper.

The use of automatically-generated software stubs [11] for abstracting (over-approximating) lower-level functions during dynamic test generation [18, 9] is also mostly orthogonal to our approach. However, the practicality of this idea is questionable because anticipating side-effects of stubbed functions accurately is problematic. In contrast, our approach is compositional while being grounded in testing and concrete execution, thus without ever generating false alarms.

Demand-driven dynamic test generation for single procedures has previously been discussed in [16, 13]. This prior work is based on dataflow analysis, does not use logic and automated theorem proving, and does not discuss interprocedural analysis. As discussed earlier, our work extends the compositional test generation framework introduced in [10] by precisely formalizing how to implement it using first-order logic formulas with uninterpreted functions and a SMT solver, and by allowing it to be demand-driven.

7 Conclusion

This paper presents an automatic and efficient symbolic execution technique for test-input generation, which is both demand-driven and compositional. By demand-driven, we mean that, given a target to cover, the technique aims to explore as few program paths as possible (called lazy exploration), and avoid exploring paths that can be guaranteed not to cover the target (called relevant exploration). By compositional, we mean that, instead of enumerating all interprocedural paths one-by-one, the technique finds feasible, interprocedural paths by combining intraprocedural paths. Because the technique is demand-driven, it can be very efficient when the goal is to cover a particular location in the program (e.g., an assertion). And, due to its compositionality, it can alleviate the path-explosion problem, which severely limits the scalability of automatic test-input generation. We have implemented a prototype of the proposed technique on top of Microsoft’s Pex test-generation tool. Preliminary experimental results are promising. Currently, we are extending our prototype to handle implementation issues such as summarizing side-effects through the heap. Future work includes applying the technique to a larger set of programs to further assess its effectiveness.

Acknowledgments. We thank Jonathan ‘Peli’ de Halleux, one of the Pex developers, Nikolaj Bjørner and Leonardo de Moura for the Z3 SMT constraint solver and their support, and the anonymous reviewers for helpful comments.

References

1. R. Alur and M. Yannakakis. Model Checking of Hierarchical State Machines. In *Proceedings of FSE'98*, pages 175–188, 1998.
2. D. Babic and A. J. Hu. Structural Abstraction of Software Verification Conditions. In *CAV'2007*, Berlin, July 2007.
3. T. Ball, R. Majumdar, T. Millstein, and S. Rajamani. Automatic Predicate Abstraction of C Programs. In *Proceedings of PLDI'2001*, 2001.
4. W.R. Bush, J.D. Pincus, and D.J. Sielaff. A static analyzer for finding dynamic programming errors. *Software Practice and Experience*, 30(7):775–802, 2000.
5. C. Cadar, V. Ganesh, P. M. Pawlowski, D. L. Dill, and D. R. Engler. EXE: Automatically Generating Inputs of Death. In *ACM CCS*, 2006.
6. E. Clarke, D. Kroening, and F. Lerda. A Tool for Checking ANSI-C Programs. In *Proceedings of TACAS'2004*, 2004.
7. C. Csallner and Y. Smaragdakis. Check'n Crash: Combining Static Checking and Testing. In *ICSE'2005*. ACM, May 2005.
8. L. de Moura and N. Bjørner. Z3, 2007. Web page: <http://research.microsoft.com/projects/Z3>.
9. D. Engler and D. Dunbar. Under-constrained execution: making automatic code destruction easy and scalable. In *Proceedings of ISSTA'2007*, 2007.
10. P. Godefroid. Compositional Dynamic Test Generation. In *POPL'2007*, pages 47–54, Nice, January 2007.
11. P. Godefroid, N. Klarlund, and K. Sen. DART: Directed Automated Random Testing. In *PLDI'2005*, pages 213–223, Chicago, June 2005.
12. D. Gopan and T. Reps. Low-level Library Analysis and Summarization. In *CAV'2007*, Berlin, July 2007.
13. N. Gupta, A. P. Mathur, and M. L. Soffa. Generating Test Data for Branch Coverage. In *Proceedings of ASE'2000*, pages 219–227, September 2000.
14. S. Khurshid and Y. L. Suen. Generalizing Symbolic Execution to Library Classes. In *PASTE'2005*, Lisbon, September 2005.
15. J. C. King. Symbolic Execution and Program Testing. *Journal of the ACM*, 19(7):385–394, 1976.
16. B. Korel. A Dynamic Approach of Test Data Generation. In *ICSM*, pages 311–317, San Diego, November 1990.
17. V. B. Livshits and M. Lam. Tracking Pointers with Path and Context Sensitivity for Bug Detection in C Programs. In *Proceedings of FSE'2003*, 2003.
18. R. Majumdar and K. Sen. Latest: Lazy dynamic test input generation. Technical report, UC Berkeley, 2007.
19. T. Reps, S. Horwitz, and M. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *Proceedings of POPL'95*, pages 49–61, 1995.
20. N. Tillmann and J. de Halleux. Pex, 2007. Web page: <http://research.microsoft.com/Pex>.
21. N. Tillmann and W. Schulte. Parameterized unit tests. In *ESEC-FSE'2005*, pages 253–262. ACM, 2005.
22. W. Visser, C. Pasareanu, and S. Khurshid. Test Input Generation with Java PathFinder. In *ISSTA'04*, Boston, July 2004.
23. Y. Xie and A. Aiken. Scalable Error Detection Using Boolean Satisfiability. In *Proceedings of POPL'2005*, 2005.