

A Survey of Symbolic Execution Techniques

Roberto Baldoni¹, [Emilio Coppa](#)², Daniele Cono D’Elia²,
[Camil Demetrescu](#)², and [Irene Finocchi](#)²

¹[Cyber Intelligence and Information Security Research Center](#), Sapienza University of Rome

²[SEASON Lab](#), Sapienza University of Rome

{[baldoni](#),[coppa](#),[delia](#),[demetres](#)}@dis.uniroma1.it, [finocchi](#)@di.uniroma1.it

Abstract

Many security and software testing applications require checking whether certain properties of a program hold for any possible usage scenario. For instance, a tool for identifying software vulnerabilities may need to rule out the existence of any backdoor to bypass a program’s authentication. One approach would be to test the program using different, possibly random inputs. As the backdoor may only be hit for very specific program workloads, automated exploration of the space of possible inputs is of the essence. Symbolic execution provides an elegant solution to the problem, by systematically exploring many possible execution paths at the same time without necessarily requiring concrete inputs. Rather than taking on fully specified input values, the technique abstractly represents them as symbols, resorting to constraint solvers to construct actual instances that would cause property violations. Symbolic execution has been incubated in dozens of tools developed over the last four decades, leading to major practical breakthroughs in a number of prominent software reliability applications. The goal of this survey is to provide an overview of the main ideas, challenges, and solutions developed in the area, distilling them for a broad audience.

“Sometimes you can’t see how important something is in its moment, even if it seems kind of important. This is probably one of those times.”

(Cyber Grand Challenge highlights from DEF CON 24, August 6, 2016)

1 Introduction

Symbolic execution is a popular program analysis technique introduced in the mid ’70s in the context of software testing to check whether a certain property can be violated by a program [[King, 1975](#), [Boyer et al., 1975](#), [King, 1976](#), [Howden, 1977](#)]. Aspects of interest could be that no division by zero is ever performed, no NULL pointer is ever dereferenced, no backdoor exists that can bypass authentication, etc. While in general there is no automated way to decide some properties (e.g., the target of an indirect jump), heuristics and approximate analyses can prove useful in practice in a variety of settings, including mission-critical and security applications.

In a concrete execution, a program is run on a specific input and a single control flow path is explored. Hence, in most cases concrete executions can only underapproximate the analysis of the property of interest. In contrast, symbolic execution can simultaneously explore multiple paths that a program could take under different inputs. This paves the road to sound analyses that can yield strong guarantees on the checked property. The key idea is to allow a program to take on *symbolic* – rather than concrete – input values. Execution is performed by a *symbolic execution engine*, which maintains for each explored control flow path: (i) a first-order Boolean *formula* that describes the conditions satisfied by the branches taken along that path, and (ii) a *symbolic memory store* that maps variables to symbolic expressions or values. Branch execution

```

1. void foobar(int a, int b) {
2.     int x = 1, y = 0;
3.     if (a != 0) {
4.         y = 3+x;
5.         if (b == 0)
6.             x = 2*(a+b);
7.     }
8.     assert(x-y != 0);
9. }

```

Figure 1: Warm-up example: which values of **a** and **b** make the **assert** fail?

updates the formula, while assignments update the symbolic store. A *model checker*, typically based on a *satisfiability modulo theories* (SMT) solver [Biere et al., 2009], is eventually used to verify whether there are any violations of the property along each explored path and if the path itself is realizable, i.e., if its formula can be satisfied by some assignment of concrete values to the program’s symbolic arguments.

Symbolic execution techniques have been brought to the attention of a heterogenous audience since DARPA announced in 2013 the Cyber Grand Challenge, a two-year competition seeking to create automatic systems for vulnerability detection, exploitation, and patching in near real-time [Shoshitaishvili et al., 2016].

More remarkably, symbolic execution tools have been running 24/7 in the testing process of many Microsoft applications since 2008, revealing for instance nearly one third of the bugs discovered during the development of Windows 7, which were missed by other program analyses and blackbox testing techniques [Godefroid et al., 2012].

In this article, we survey the main aspects of symbolic execution and discuss its extensive usage in software testing and computer security applications, where software vulnerabilities can be found by symbolically executing programs at the level of either source or binary code. We start with a simple example that highlights many of the fundamental issues addressed in the remainder of the article.

1.1 A Warm-up Example

Consider the C code of Figure 1 and assume that our goal is to determine which inputs make the **assert** at line 8 of function **foobar** fail. Since each input parameter can take as many as 2^{32} distinct integer values, the approach of running concretely function **foobar** on randomly generated inputs will unlikely pick up exactly the assert-failing inputs. By evaluating the code using symbols for its inputs, instead of concrete values, symbolic execution overcomes this limitation and makes it possible to reason on *classes of inputs*, rather than single input values.

In more detail, every value that cannot be determined by a static analysis of the code, such as an actual parameter of a function or the result of a system call that reads data from a stream, is represented by a symbol α_i . At any time, the symbolic execution engine maintains a state $(stmt, \sigma, \pi)$ where:

- $stmt$ is the next statement to evaluate. For the time being, we assume that $stmt$ can be an assignment, a conditional branch, or a jump (more complex constructs such as function calls and loops will be discussed in Section 2 and Section 5, respectively).
- σ is a *symbolic store* that associates program variables with either expressions over concrete values or symbolic values α_i .
- π denotes the *path constraints*, i.e., is a formula that expresses a set of assumptions on the symbols α_i due to branches taken in the execution to reach $stmt$. At the beginning of the analysis, $\pi = true$.

Depending on $stmt$, the symbolic engine changes the state as follows:

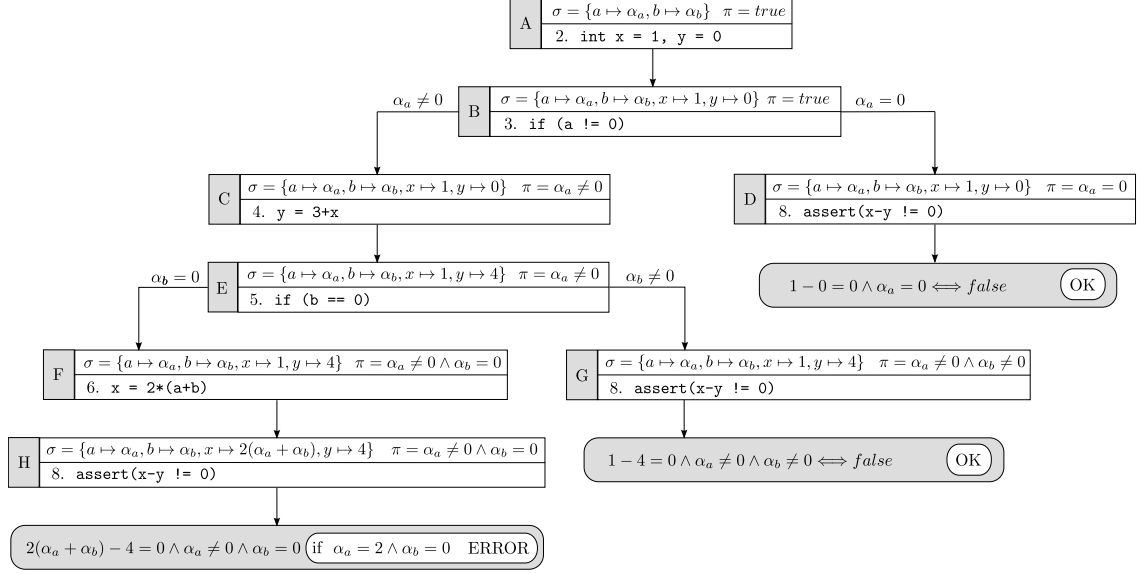


Figure 2: Symbolic execution tree of function `foobar` given in Figure 1. Each execution state, labeled with an upper case letter, shows the statement to be executed, the symbolic store σ , and the path constraints π . Leaves are evaluated against the condition in the `assert` statement.

- The evaluation of an assignment $x = e$ updates the symbolic store σ by associating x with a new symbolic expression e_s . We denote this association with $x \mapsto e_s$, where e_s is obtained by evaluating e in the context of the current execution state and can be any expression involving unary or binary operators over symbols and concrete values.
- The evaluation of a conditional branch `if e then s_{true} else s_{false}` affects the path constraints π . The symbolic execution is forked by creating two execution states with path constraints π_{true} and π_{false} , respectively, which correspond to the two branches: $\pi_{true} = \pi \wedge e_s$ and $\pi_{false} = \pi \wedge \neg e_s$, where e_s is a symbolic expression obtained by evaluating e . Symbolic execution independently proceeds on both states.
- The evaluation of a jump `goto s` updates the execution state by advancing the symbolic execution to statement s .

A symbolic execution of function `foobar`, which can be effectively represented as a tree, is shown in Figure 2. Initially (execution state *A*) the path constraints are `true` and input arguments `a` and `b` are associated with symbolic values. After initializing local variables `x` and `y` at line 2, the symbolic store is updated by associating `x` and `y` with concrete values 1 and 0, respectively (execution state *B*). Line 3 contains a conditional branch and the execution is forked: depending on the branch taken, a different statement is evaluated next and different assumptions are made on symbol α_a (execution states *C* and *D*, respectively). In the branch where $\alpha_a \neq 0$, variable `y` is assigned with `x + 3`, obtaining $y \mapsto 4$ in state *E* because $x \mapsto 1$ in state *C*. In general, arithmetic expression evaluation simply manipulates the symbolic values. After expanding every execution state until the `assert` at line 8 is reached on all branches, we can check which input values for parameters `a` and `b` can make the `assert` fail. By analyzing execution states $\{D, G, H\}$, we can conclude that only *H* can make `x-y = 0` true. The path constraints for *H* at this point implicitly define the set of inputs that are unsafe for `foobar`. In particular, any input values such that:

$$2(\alpha_a + \alpha_b) - 4 = 0 \wedge \alpha_a \neq 0 \wedge \alpha_b = 0$$

will make `assert` fail. An instance of unsafe input parameters can be eventually determined by invoking a *model checker* [Biere et al., 2009] to solve the path constraints, which in this example would yield $a = 2$ and $b = 0$.

1.2 Challenges in Symbolic Execution

In the example discussed in Section 1.1 symbolic execution can identify *all* the possible unsafe inputs that make the `assert` fail. This is achieved through an exhaustive exploration of the possible execution states. From a theoretical perspective, exhaustive symbolic execution provides a *sound* and *complete* methodology for any decidable analysis. Soundness prevents false negatives, i.e., all possible unsafe inputs are guaranteed to be found, while completeness prevents false positives, i.e., input values deemed as unsafe are actually unsafe. As we will discuss later on, exhaustive symbolic execution is unlikely to scale beyond small applications. Hence, in practice we often settle for less ambitious goals, e.g., by trading soundness for performance.

Challenges that symbolic execution has to face when processing real-world code can be significantly more complex than those illustrated in our warm-up example. Several observations and questions naturally arise:

- *Memory*: how does the symbolic engine handle pointers, arrays, or other complex objects? Any arbitrarily complex object can be regarded as an array of bytes and each byte associated with a distinct symbol. However, when possible, exploiting structural properties of the data may be more convenient: for instance, relational bounds on the class fields in object-oriented languages could be used for refining the search performed by symbolic execution.
- *Environment*: how does the symbolic engine handle interactions with the environment? Real-world applications constantly interact with the environment (e.g., the file system or the network) through libraries and system calls. These interactions may cause side-effects (such as the creation of a file) that could later affect the execution and must be therefore taken into account. Evaluating any possible interaction outcome is generally unfeasible: it could generate a large number of execution states, of which only a small number can actually happen in a non-symbolic scenario. A typical strategy is to consider popular library and system routines and create models that can help the symbolic engine analyze only significant outcomes.
- *Loops*: how does the symbolic engine handle loops? Choosing the number of loop iterations to analyze is especially critical when this number cannot be determined in advance (e.g., depends on an input parameter). The naive approach of unrolling iterations for every valid bound would result in a prohibitively large number of states. Typical solutions are to compute an underapproximation of the analysis by limiting the number of iterations to some value k , thus trading speed for soundness. Other approaches infer loop invariants through static analysis and use them to merge equivalent states.
- *State space explosion and path selection*: how does symbolic execution deal with path explosion? Language constructs such as loops might exponentially increase the number of execution states. It is thus unlikely that a symbolic execution engine can exhaustively explore all the possible states within a reasonable amount of time. In practice, heuristics are used to guide exploration and prioritize certain states first (e.g., to maximize code coverage). In addition, symbolic engines can implement efficient mechanisms for evaluating multiple states in parallel without running out of resources.
- *Constraint solver*: what can a constraint solver do in practice? Constraint solvers suffer from a number of limitations. They can typically handle complex constraints in a reasonable amount of time only if they are made of linear expressions over their constituents. Symbolic execution engines normally implement a number of optimizations to make queries as much *solver-friendly* as possible, for instance by splitting queries into independent components to be processed separately or by performing algebraic simplifications.
- *Binary code*: what issues can arise when symbolically executing binary code? While the warm-up example of Section 1.1 is written in C, in several scenarios binary code is the only




Figure 3: Concrete and abstract execution machine models.

available representation of a program. However, having the source code of an application can make symbolic execution significantly easier, as it can exploit high-level properties (e.g., object shapes) that can be inferred statically by analyzing the source code.

Depending on the specific context in which symbolic execution is used, different choices and assumptions are made to address the questions highlighted above. Although these choices typically affect soundness or completeness, in several scenarios a partial exploration of the space of possible execution states may be sufficient to achieve the goal (e.g., identifying a crashing input for an application) within a limited time budget.

1.3 Organization of the Article

The remainder of this article is organized as follows. In Section 2, we discuss the overall principles and evaluation strategies of a symbolic execution engine. Section 3 through Section 8 address the key challenges that we listed in Section 1.2. Prominent applications based on symbolic execution techniques are discussed in Section 9, while concluding remarks are addressed in Section 10.

2 Symbolic Execution Engines

In this section we describe some important principles for the design of symbolic executors as well as crucial tradeoffs that arise in their implementation. Moving from the concepts of concrete and symbolic runs, we also introduce the idea of “concolic” execution.

2.1 Concrete, Symbolic, and Concolic Execution

As shown in the warm-up example (Section 1.1), a symbolic execution of a program can generate – in theory – all possible control flow paths that the program could take during its concrete executions on specific inputs. While modeling all possible runs allows for very interesting analyses, it is typically unfeasible in practice, especially on real-world software, for a variety of reasons.

First, as extensively discussed in Section 6, the number of control flow paths to be generated and analyzed could be prohibitively large, due to branch instructions and loops. In the worst case, if the code contains an unbounded loop, symbolic execution could keep running forever, generating a potentially infinite number of paths (we refer to Section 5 for an example).

Moreover, as observed in Section 1, symbolic engines are clients of SMT solvers, which are continuously invoked during the analysis. Although powerful SMT solvers are currently available, the time spent in constraint solving is still one of the main performance barriers for symbolic engines. It may also happen that the program yields constraints that the solver cannot handle well (e.g., non-linear constraints), in spite of the fact that symbolic executors often use more than one solver in order to support as many decidable logical fragments as possible.

A standard approach to limit the resources (running time and space usage) required by the execution engine and to handle complex constraints is to mix concrete and symbolic execution: this is dubbed *concolic execution*, where the term concolic is a portmanteau of the words concrete and symbolic. The basic idea is to have the concrete execution drive the symbolic execution (see also Figure 3). Besides the symbolic store and the path constraints, a concolic execution engine also maintains a concrete store σ_c . After choosing an arbitrary input to begin with, it executes the program both concretely and symbolically by simultaneously updating the two stores and the path constraints. In order to explore different paths, the path conditions given by one or more

branches can be negated and the SMT solver invoked to find a satisfying assignment for the new constraints, i.e., to generate a new input.

Example. Consider the C function in Figure 1 and suppose to choose $a = 1$ and $b = 1$ as input parameters. Under these conditions, the concrete execution takes path $A \rightsquigarrow B \rightsquigarrow C \rightsquigarrow E \rightsquigarrow G$ in the symbolic tree of Figure 2. Besides the symbolic stores shown in Figure 2, the concrete stores maintained in the traversed states are the following:

- $\sigma_c = \{a \mapsto 1, b \mapsto 1\}$ in state A ;
- $\sigma_c = \{a \mapsto 1, b \mapsto 1, x \mapsto 1, y \mapsto 0\}$ in states B and C ;
- $\sigma_c = \{a \mapsto 1, b \mapsto 1, x \mapsto 1, y \mapsto 4\}$ in states E and G .

After checking that the `assert` conditions at line 8 succeed, we can generate a new control flow path by negating the last path constraint, i.e., $\alpha_b \neq 0$. The solver at this point would generate a new input that satisfies the constraints $\alpha_a \neq 0 \wedge \alpha_b = 0$ (for instance $a = 1$ and $b = 0$) and the execution would continue in a similar way along the path $A \rightsquigarrow B \rightsquigarrow C \rightsquigarrow E \rightsquigarrow F$.

As shown by the example, the symbolic information maintained during a concrete run can be exploited by the execution engine, for instance, to obtain new inputs and explore new control flow paths. We will further discuss this aspect in Section 6.2.

It is worth noticing that concolic execution trades soundness for performance: false negatives are indeed possible, because some program executions – and therefore possible erroneous behaviors – may be missed. In the literature, this is also regarded as an *under-approximate* form of program analysis.

Many papers exploit variants of concolic execution or different ways of mixing concrete and symbolic runs. For instance, in *execution-generated testing* (see, e.g., KLEE [Cadar et al., 2008], EXE [Cadar et al., 2006], and [Cadar and Sen, 2013], the symbolic engine always executes concretely the operations that involve only concrete values. This makes it possible to reason even over complex (e.g., non-linear) operations if they involve only concrete values. *Selective symbolic execution* [Chipounov et al., 2012] takes a different approach, by interleaving portions of code that are concretely run with fully symbolic phases. The interleaving must be done carefully in order to preserve the meaningfulness of the whole exploration. When an argument x for a function call to concretize is symbolic, the engine should convert it to some concrete value in order to perform the call, which is equivalent to corseting the exploration to a single path in the callee. When the call returns and the symbolic phase resumes, the concrete value for x becomes part of the path constraints for the remainder of the exploration. However, a large number of paths may be then excluded. S²E [Chipounov et al., 2012] presents a systematic approach to consistently cross the symbolic/concrete boundary in both directions: it describes a strategy to deal with constraints introduced on symbolic values as a consequence of concretization, and introduces a number of consistency models – where a state is *consistent* when there exists a feasible path to it from the initial state – which suit different analyses. Constraints updated to account for concrete values are marked as *soft*, and whenever a branch in the symbolic domain is disabled because of a soft constraint, execution goes back and picks a value for the concrete call that would enable that branch. Throughout this article, we will see other uses of concretization (see, e.g., Section 3 and Section 7) and of concolic execution (see Section 6).

2.2 Design Principles of Symbolic Executors

A number of performance-related design principles that a symbolic execution engine should follow are summarized in [Cha et al., 2012]. Most notably:

1. *Progress*: the executor should be able make forward progress for an arbitrarily long time without exceeding the given resources. Memory consumption can be especially critical, due to the potentially gargantuan number of distinct control flow paths.

2. *Work repetition*: no execution work should be repeated, avoiding to restart a program several times from its very beginning in order to analyze different paths that might have a common prefix.
3. *Analysis reuse*: analysis results from previous runs should be reused as much as possible. In particular, costly invocations to the SMT solver on previously solved path constraints should be avoided.

Due to the large size of the execution state space to be analyzed, different symbolic engines have explored different tradeoffs between, e.g., running time and memory consumption, or performance and soundness/completeness of the analysis.

Symbolic executors that attempt to execute multiple paths simultaneously in a single run – also called *online executors* – clone the execution state at each input-dependent branch. Examples are given in KLEE [Cadar et al., 2008], AEG [Avgerinos et al., 2011], S²E [Chipounov et al., 2012]. These engines never re-execute previous instructions, thus avoiding work repetition. However, since many active states need to be kept in memory, they put a huge burden on memory consumption, possibly hindering progress. Effective techniques for reducing the memory footprint include *copy-on-write*, which tries to share as much as possible between different states [Cadar et al., 2008]. Moreover, executing multiple paths in parallel requires to ensure isolation between execution states, e.g., keeping different states of the OS by emulating the effects of system calls.

Reasoning about a single path at a time, as in concolic execution, is the approach taken by so-called *offline executors*, such as SAGE [Godefroid et al., 2008]. Running each path independently of the others results in low memory consumption with respect to online executors and in the capability of reusing immediately analysis results from previous runs. On the other side, work can be largely repeated, since each run usually restarts the execution of the program from the very beginning. In a typical implementation of offline executors, runs are concrete and require an input seed: the program is first executed concretely, a trace of instructions is recorded, and the recorded trace is then executed symbolically.

Hybrid executors such as MAYHEM [Cha et al., 2012] attempt at balancing between speed and memory requirements: they start in online mode and generate checkpoints, rather than forking new executors, when memory usage or the number of concurrently active states reaches a threshold. Checkpoints maintain the symbolic execution state and replay information. When a checkpoint is picked for restoration, the concrete state is restored and the online exploration resumes.

2.3 Caching

Caching is a powerful technique to achieve time-space tradeoffs and is embodied in symbolic executors in different ways. Most prominently:

- *Function caching*. A function f , and more in general any part of a program, may be called multiple times during an execution, either at the same calling context or at different ones. The traditional symbolic execution approach requires to symbolically execute f at each call. [Godefroid, 2007] proposes a compositional approach that dynamically generates *function summaries*, allowing the symbolic executor to effectively reuse prior discovered analysis results. A similar idea has been also proposed in [Boonstoppel et al., 2008]. The main intuition is that, if two program states differ only for some program values that are not read later, the executions generated by the two program states will produce the same side effects. Side effects of a portion of code can be therefore cached and possibly reused later.
- *Loop summarization*. In order to avoid redundant executions of the same loop under the same program state, loop summaries can be computed and cached for later reuse, similarly to function summaries. We refer to Section 5 for details on a loop summarization strategy proposed in [Godefroid and Luchaup, 2011].
- *Constraint reuse*. In order to speed up constraint solving, different works support the reuse of constraint solutions based on syntactic or semantic equivalence of the constraints. Examples

Symbolic engine	References	Project URL (last retrieved: August 2016)
CUTE	[Sen et al., 2005]	-
DART	[Godefroid et al., 2005]	-
JCUTE	[Sen and Agha, 2006]	https://github.com/osl/jcute
KLEE	[Cadar et al., 2006, Cadar et al., 2008]	https://klee.github.io/
SAGE	[Godefroid et al., 2008, Elkarablieh et al., 2009]	-
BITBLAZE	[Song et al., 2008]	http://bitblaze.cs.berkeley.edu/
CREST	[Burnim and Sen, 2008]	https://github.com/jburnim/crest
PEX	[Tillmann and De Halleux, 2008]	http://research.microsoft.com/en-us/projects/pex/
RUBYX	[Chaudhuri and Foster, 2010]	-
JAVA PATHFINDER	[Păsăreanu and Rungta, 2010]	http://babelfish.arc.nasa.gov/trac/jpf
OTTER	[Reisner et al., 2010]	https://bitbucket.org/khooy/otter/
BAP	[Brumley et al., 2011]	https://github.com/BinaryAnalysisPlatform/bap
CLOUD9	[Bucur et al., 2011]	http://cloud9.epfl.ch/
MAYHEM	[Cha et al., 2012]	-
SYMDROID	[Jeon et al., 2012]	-
S ² E	[Chipounov et al., 2012]	http://s2e.epfl.ch/
FUZZBALL	[Martignoni et al., 2012, Caselden et al., 2013]	http://bitblaze.cs.berkeley.edu/fuzzball.html
JALANGI	[Sen et al., 2013]	https://github.com/Samsung/jalangi2
PATHGRIND	[Sharma, 2014]	https://github.com/codelion/pathgrind
KITE	[do Val, 2014]	http://www.cs.ubc.ca/labs/isd/Projects/Kite
SYMJS	[Li et al., 2014]	-
CIVL	[Siegel et al., 2015]	http://vsl.cis.udel.edu/civl/
KEY	[Hentschel et al., 2014]	http://www.key-project.org/
ANGR	[Shoshitaishvili et al., 2015, Shoshitaishvili et al., 2016]	http://angr.io/
TRITON	[Saudel and Salwan, 2015]	http://triton.quarkslab.com/
PyEXZ3	[Ball and Daniel, 2015]	https://github.com/thomasjball/PyExZ3
JDART	[Luckow et al., 2016]	https://github.com/psycopaths/jdart
CATG	-	https://github.com/ksen007/janala2
PySYMEMU	-	https://github.com/feliam/pysymemu/
MIASM	-	https://github.com/cea-sec/miasm

Figure 4: Selection of symbolic execution engines, along with their reference article(s) and software project web site (if any).

are given in EXE [Cadar et al., 2006], KLEE [Cadar et al., 2008], and [Yang et al., 2012, Visser et al., 2012]. We will further discuss this optimization in Section 7.

2.4 Tools

Table 4 lists a number of symbolic execution engines that have worked as incubators for several of the techniques surveyed in this article. The novel contributions introduced by tools that represented milestones in the area are described in the appropriate sections throughout the article.

3 Memory model

Our warm-up example of Section 1.1 presented a simplified memory model where data are stored in scalar variables only, with no indirection. A crucial aspect of symbolic execution is how memory should be modeled to support programs with pointers and arrays. This requires extending our notion of memory store by mapping not only variables, but also memory addresses to symbolic expressions or concrete values. In general, a store σ that explicitly models memory addresses can be thought as a mapping that associates memory addresses (indexes) with either expressions over concrete values or symbolic values. We can still support variables by using their address rather than their name in the mapping. In the following, when we write $x \mapsto e$ for a variable x and an expression e we mean $\&x \mapsto e$, where $\&x$ is the concrete address of variable x . Also, if v is an array and c is an integer constant, by $v[c] \mapsto e$ we mean $\&v + c \mapsto e$. A memory model is an important design choice for a symbolic engine, as it can have a significant influence on the coverage achieved by symbolic execution, as well as on the scalability of constraint solving [Cadar and Sen, 2013].

The *symbolic memory address* problem [Schwartz et al., 2010] arises when the address referenced in the operation is a symbolic expression derived from user input instead of a concrete value. In the remainder of this section, we discuss a number of popular solutions.

3.1 Fully Symbolic Memory

At the one end of the spectrum, an engine may treat memory addresses as fully symbolic. This is the approach taken by a number of works (e.g., BITBLAZE [Song et al., 2008], [Thakur et al., 2010], BAP [Brumley et al., 2011], and [Trtík and Strejček, 2014]). Two fundamental approaches, pioneered by King in its seminal paper [King, 1976], are the following:

- *State forking.* If an operation reads from or writes to a symbolic address, the state is forked by considering all possible states that may result from the operation. The path constraints are updated accordingly for each forked state.

Example. Consider the example shown in Figure 5. The write operation at line 4 affects either $a[0]$ or $a[1]$, depending on the unknown value of array index i . State forking creates two states after executing the memory assignment to explicitly consider both possible scenarios (Figure 6). The path constraints for the forked states encode the assumption made on the value of i . Similarly, the memory read operation $a[j]$ at line 5 may access either $a[0]$ or $a[1]$, depending on the unknown value of array index j . Therefore, for each of the two possible outcomes of the assignment $a[i]=5$, there are two possible outcomes of the `assert`, which are explicitly explored by forking the corresponding states.

```

1. void foobar(unsigned i, unsigned j) {
2.     int a[2] = { 0 };
3.     if (i>1 || j>1) return;
4.     a[i] = 5;
5.     assert(a[j] != 5);
6. }

```

Figure 5: Memory modeling example: which values of i and j make the `assert` fail?

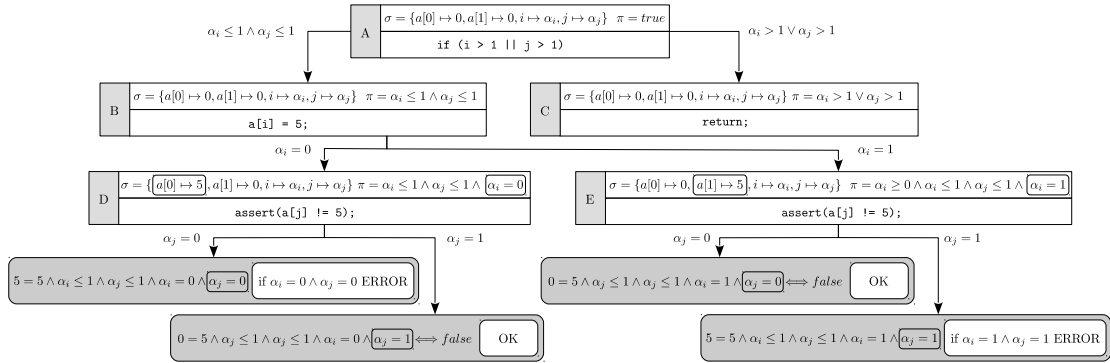


Figure 6: Fully symbolic memory via state forking for the example of Figure 5.

- *if-then-else formulas.* An alternative approach consists in encoding the uncertainty on the possible values of a symbolic pointer into the expressions kept in the symbolic store and in the path constraints, without forking any new states. The key idea is to exploit the capability of some solvers to reason on formulas that contain if-then-else expressions of the form $ite(c, t, f)$, which yields t if c is true, and f otherwise¹. The approach works differently for memory read and write operations. Let α be a symbolic address that may assume the concrete values a_1, a_2, \dots :

– reading from α yields the expression $ite(\alpha = a_1, \sigma(a_1), ite(\alpha = a_2, \sigma(a_2), \dots))$;

¹In propositional logic, the $ite(c, t, f)$ expression could be replaced with the formula $(c \wedge t) \vee (\neg c \wedge f)$.

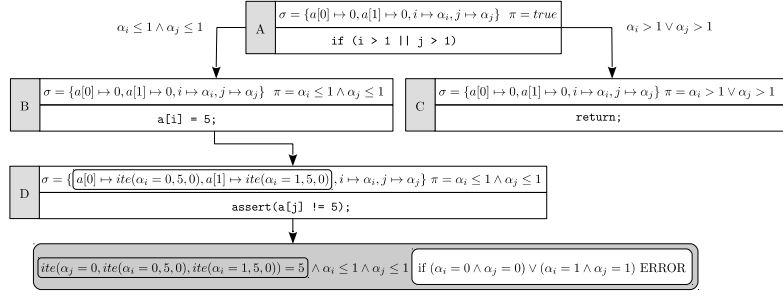


Figure 7: Fully symbolic memory via if-then-else formulas for the example of Figure 5.

- writing an expression e at α updates the symbolic store for each a_1, a_2, \dots as $\sigma(a_i) \leftarrow ite(\alpha = a_i, e, \sigma(a_i))$.

Notice that in both cases, a memory operation introduces in the store as many *ite* expressions as the number of possible values the accessed symbolic address may assume. The *ite* approach to symbolic memory is used, e.g., in ANGR [Shoshitaishvili et al., 2016] (Section 3.3).

Example. Consider again the example shown in Figure 5. Rather than forking the state after the operation `i=5` at line 4, the if-then-else approach updates the memory store by encoding both possible outcomes of the assignment, i.e., $a[0] \mapsto ite(\alpha_i = 0, 5, 0)$ and $a[1] \mapsto ite(\alpha_i = 1, 5, 0)$ (Figure 7). Similarly, rather than creating a new state for each possible distinct address of `a[j]` at line 5, the uncertainty on j is encoded in the single expression $ite(\alpha_j = 0, \sigma(a[0]), \sigma(a[1])) = ite(\alpha_j = 0, ite(\alpha_i = 0, 5, 0), ite(\alpha_i = 1, 5, 0))$.

In general, a symbolic address may reference any cell in memory, making the approaches described above intractable. Fortunately, in many practical cases the set of possible addresses a memory operation may reference is small [Song et al., 2008], as in the example shown in Figure 5 where indexes i and j range in a bounded interval.

To model fully symbolic pointers, an extensive line of research (e.g., EXE [Cadar et al., 2006], KLEE [Cadar et al., 2008], SAGE [Elkarablieh et al., 2009]) leverages the expressive power of some SMT solvers, which can model operations on arrays as first-class entities in constraint formulas using *theories of arrays* in their decision procedures [Ganesh and Dill, 2007].

3.2 Address Concretization

In all cases where the combinatorial complexity of the analysis explodes as pointer values cannot be bounded to sufficiently small ranges, *address concretization*, which consists in concretizing a pointer to a single specific address, is a popular alternative. This can reduce the number of states and the complexity of the formulas fed to the solver and thus improve running time, although may cause the engine to miss paths that, for instance, depend on specific values for some pointers.

Concretization is a natural choice for offline executors (Section 2.2) such as DART [Godefroid et al., 2005] and early SAGE releases [Godefroid et al., 2008] that concretely execute one path at a time while collecting path constraints along executed paths. Systems such as CREST [Burnim and Sen, 2008] and CUTE [Sen et al., 2005] are capable of reasoning only about equality constraints for pointers, as they can be solved efficiently, and resort to concretization for general symbolic references.

3.3 Partial Memory Modeling

To mitigate the scalability problems of fully symbolic memory and the loss of soundness of memory concretization, MAYHEM [Cha et al., 2012] explores a middle point in the spectrum by introducing

a *partial* memory model. The key idea is that written addresses are always concretized and read addresses are modeled symbolically if the contiguous interval of possible values they may assume is small enough. This model is based on a trade-off: it uses more expressive formulas than concretization, since it encodes multiple pointer values per state, but does not attempt to encode all of them like in fully symbolic memory [Avgerinos, 2014]. A basic approach to bound the set of possible values that an address may assume consists in trying different concrete values and checking whether they satisfy the current path constraints, excluding large portions of the address space at each trial until a tight range is found. This algorithm comes with a number of caveats: for instance, querying the solver on each symbolic dereference is expensive, the memory region may not be continuous, and the values within the memory range of a symbolic pointer might have structure. MAYHEM [Cha et al., 2012] thus performs a number of optimizations, including Value Set Analysis [Duesterwald, 2004] and forms of query caching (Section 7), to refine ranges efficiently. If at the end of the process the range size exceeds a given threshold (e.g., 1024), the address is concretized. ANGR [Shoshitaishvili et al., 2016] also adopts the partial memory model idea and extends it by optionally supporting write operations on symbolic pointers that range within small contiguous intervals (up to 128 addresses).

3.4 Complex Objects

[Khurshid et al., 2003] introduces a generalization of traditional symbolic execution to advanced constructs of object-oriented languages such as Java and C++. The authors describe a verification framework that combines symbolic execution and model checking to handle dynamically allocated data structures such as lists and trees.

In particular, they generalize symbolic execution by introducing *lazy initialization* to effectively handle dynamically allocated objects. Compared to our warm-up example (Section 1.1), the state representation is extended with a *heap configuration* used to maintain such objects. Symbolic execution of a method taking complex objects as inputs starts with uninitialized fields, and assigns values to them in a lazy fashion, i.e., they are initialized when first accessed during execution.

When an uninitialized reference field is accessed, the algorithm forks the current state with three different heap configurations, in which the field is initialized with: (1) `null`, (2) a reference to a new object with all symbolic attributes, and (3) a previously introduced concrete object of the desired type, respectively. This on-demand concretization enables symbolic execution of methods without requiring an a priori bound on the number of input objects. Also, forking the state as in (2) results into a systematic treatment for aliasing, i.e., when an object can be accessed through multiple references.

[Khurshid et al., 2003, Visser et al., 2004] combine lazy initialization with user-provided *method preconditions*, i.e., conditions which are assumed to be true before the execution of a method. Pre-conditions are used to characterize those program input states in which the method is expected to behave as intended by the programmer. For instance, we expect a binary tree data structure to be acyclic and with every node - except for the root - having exactly one parent. Conservative preconditions are used to ensure that incorrect heap configurations are eliminated during initialization, speeding up the symbolic execution process.

Further refinements to lazy initialization are described in a number of works, e.g., [Deng et al., 2012, Geldenhuys et al., 2013, Rosner et al., 2015], which all share the goal of reducing the number of heap configurations to generate when forking the state. [Deng et al., 2012] also provides a formal treatment of lazy initialization in Java.

Of a different flavor is the technique presented in [Shannon et al., 2007] for symbolic execution over objects instantiated from commonly used libraries. The authors argue that performing symbolic execution at the representation level might be redundant if the aim is to only check the client code, thus trusting the correctness of the library implementation. They discuss the idea of symbolically executing methods of the Java `String` class using a finite-state automaton that abstracts away the implementation details. They present a case study of an application that dynamically generates SQL queries: symbolic execution is used to check whether the statements conform to the SQL grammar and possibly match injection patterns. The authors mention that their approach

might be used to symbolically execute over standard container classes such as trees or maps. It is worth mentioning that symbolic execution is used to detect SQL injection vulnerabilities also in [Fu et al., 2007].

4 Interaction with the environment

When a program interacts with its environment – e.g., file system, environment variables, network - a symbolic executor has to take into account the whole software stack surrounding it, including system libraries, kernel, drivers, etc.

A body of early works (e.g., DART [Godefroid et al., 2005], CUTE [Sen et al., 2005], and EXE [Cadar et al., 2006]) includes the environment in symbolic analysis by actually executing external calls using concrete arguments for them. This indeed limits the behaviors they can explore compared to a fully symbolic strategy, which on the other hand might be unfeasible. In an online executor this choice also results in having calls from distinct execution paths interfere with each other.

Another way to tackle the problem is to create an abstract model that captures these interactions. For instance, in KLEE [Cadar et al., 2008] symbolic files are supported through a simple *symbolic file system*, which is private for each execution state. In particular, it consists of a directory with n symbolic files whose number and sizes are specified by the user. An operation on a symbolic file results in forking $n + 1$ state branches: one for each possible file, plus an optional one to capture unexpected errors in the operation. As the number of functions in a standard library is typically large and writing models for them is an error-prone and rather expensive process [Ball et al., 2006], models are generally implemented at system call-level rather than library level. An additional benefit of this approach is that it allows for symbolic exploration of the libraries as well.

AEG [Avgerinos et al., 2011] models most of the system environment that an attacker can possibly use as an input source, including the file system, network sockets, and environment variables. Additionally, more than 70 library and system calls are emulated, including thread- and process-related system calls, and common formatting functions to capture potential buffer overflows. Symbolic files are handled as in KLEE [Cadar et al., 2008], while symbolic sockets are dealt with in a similar manner, with packets and their payloads being processed as in symbolic files and their contents. CLOUD9 further extends support to many other POSIX libraries, allowing the user to also control advanced conditions in the testing environment. For instance, CLOUD9 is able to simulate delay, reordering, and dropping of packets caused by a fragmented network data stream.

S²E [Chipounov et al., 2012] remarks that models, other than expensive to write, rarely achieve full accuracy, and also lose it as the modeled system evolves. It would be thus preferable to have analyzed programs to interact with the real environment while exploring multiple paths. In their S²E platform the authors rely on virtualization to perform the desired analysis on the real software stack, preventing side effects from propagating across independent execution paths.

The approach followed in DART [Godefroid et al., 2005] is different, as the goal is to enable automated unit testing. DART deems as foreign interfaces all the external variables and functions referenced in a C program along with the arguments for a top-level function. External functions are simulated by nondeterministically returning any value of their specified return type. Library functions are normally not considered external functions as they are controlled by the program, but in practice the user can adjust the boundary between library and external functions to simulate the desired effects.

5 Loops

Loops are one of the main causes of path explosion: each iteration of a loop can be seen as an `if-goto` statement, leading to a conditional branch in the execution tree. If the loop condition involves one or more symbolic values, the number of generated branches may be potentially infinite.

```

1. int x = sym_input(); // e.g., read from file
2. while (x > 0) {
3.     x = sym_input();
4. }

```

Figure 8: Loop example with input read from the environment [Cadaru and Sen, 2013].

Example. Consider the code fragment of Figure 8 [Cadaru and Sen, 2013], where `sym_input()` is an external routine that interacts with the environment (e.g., by reading input data from a network) and returns a fresh symbolic input. The path constraint set at any final state has the form:

$$\pi = \left(\bigwedge_{i \in [1, k]} \alpha_i > 0 \right) \wedge (\alpha_{k+1} \leq 0)$$

where k is the number of iterations and α_i is the symbol produced by `sym_input()` at the i -th iteration.

The problem of path explosion due to symbolic execution of loops has been attacked from different sides. A first natural strategy adopted by many symbolic engines is to limit the loop exploration up to a certain number of iterations. Obviously, this may lead to missing interesting paths in the program. For this reason, some works (e.g., AEG [Avgerinos et al., 2011]) have also considered the opposite strategy, allowing the engine to fully explore some loops. To mitigate the path explosion problem, only a single instance of the symbolic executor is allowed to fully unroll a loop, while other instances conservatively explore other paths. This approach has been shown to be effective in some application contexts such as security (e.g., identification of buffer overflows) where interesting behavior may be observed at the loop boundaries.

By using static or dynamic analysis techniques, it may be possible to derive properties over a loop that can be exploited by the symbolic engine to significantly prune branching paths. For instance, knowledge of the exact number of loop iterations - or at least a constant upper bound on it - can significantly help the engine. Section 6.4 provides a more general discussion of how preconditions can help symbolic execution. Nevertheless, even symbolic execution can be used to derive loop invariants. Indeed, if a program contains an assertion after the loop, the approach presented in [Păsăreanu and Visser, 2004] works backwards from the property to be checked and it iteratively applies approximation to derive loop invariants. The main idea is to pick the asserted property as the initial invariant candidate and then to exploit symbolic execution to check whether this property is inductive. If the invariant cannot be verified for some loop paths, it is replaced by a different invariant. The next candidate for the invariant is generated by exploiting the path constraints for the paths on which the verification has failed. Additional refinements steps are performed to guarantee termination.

[Godefroid and Luchaup, 2011] presents a technique that automatically derives partial summarizations for loops. A loop summarization is similar to a function summary (Section 2.3), using a set of preconditions and a set of postconditions. These are computed dynamically during the symbolic execution by reasoning on the dependencies among loop conditions and symbolic variables. As soon as a loop summary is computed, it is cached for possibly subsequent reuse. This not only allows the symbolic engine to avoid redundant executions of the same loop under the same program state, but also makes it possible to generalize the loop summary to cover even different executions of the same loop that run under different conditions. A main limitation of this approach is that it can generate summaries only for loops that iteratively update symbolic variables across loop iterations by adding a constant, non-zero amount.

[Slaby et al., 2013] introduces a technique of a different flavor that analyzes cyclic paths in the control flow graph of a given program and produces *templates* that declaratively describe the program states generated by these portions of code into a symbolic execution tree. By exploiting templates, the symbolic execution engine needs to explore a significantly reduced number of program states. A drawback of this approach is that templates introduce quantifiers in the path




Figure 9: Pruning unrealizable paths example: (a) Fragment of code; (b) Symbolic execution of the fragment of code: the *true* branch in node D is not explored since its path constraints ($\alpha_a \leq 0 \wedge \alpha_a > 1$) are not satisfiable, i.e., there exists no assignment of variable **a** that can drive a concrete execution first toward node D and then through its *true* branch.

constraints: in turn, this may significantly increase the burden on the constraint solver.

It has also been observed that loop executions may strictly depend on input features. *Loop-extended symbolic execution* [Saxena et al., 2009] is able to effectively explore a loop whenever a grammar describing the input program is available. Relating the number of iterations with features of the program input can guide the exploration of the program states generated by a loop.

6 Path explosion

One of the main challenges of symbolic execution is the path explosion problem. Since symbolic execution may fork off a new execution engine instance at every branch, the total number of executors may be exponential in the number of branches in the program. This impacts both time and space, as a symbolic executor may need to keep track of an exponential number of pending branches to be explored. A common approach is to compute an under-approximation of the analysis that only explores a relevant subset of the state space.

6.1 Pruning Unrealizable Paths

A first natural technique for reducing the path space is invoking the constraint solver at each branch, pruning branches that are not realizable. Indeed, if the constraint solver is able to prove that the logical formula given by the path constraints of a branch is not satisfiable, then there exists no assignment of the program input values that would drive a real execution toward that path. For this reason, the symbolic engine can safely discard the path involving that branch without affecting soundness of the approach.

Example. Consider the example shown in Figure 9. Assuming that **a** is a local variable bound to an unconstrained symbol α_a , a symbolic engine would start the execution of the code fragment in Figure 9a by evaluating the branch condition $a > 0$. Before expanding both branches, the symbolic engine queries a constraint solver to verify that no contradiction arises when adding to the path constraints π the *true* branch condition ($\alpha_a > 0$) or the *false* branch condition ($\alpha_a \leq 0$). Since both paths are feasible, the engine forks the execution states B and D (Figure 9b). A similar scenario happens when the engine evaluates the branch condition $a > 1$. However, since α_a is not unconstrained anymore, some contradictions may be actually possible. The engine queries the solver to check the following path constraints: (1) $\alpha_a > 0 \wedge \alpha_a > 1$, (2) $\alpha_a > 0 \wedge \alpha_a \leq 1$, (3) $\alpha_a \leq 0 \wedge \alpha_a > 1$, and (4) $\alpha_a \leq 0 \wedge \alpha_a \leq 1$. Notice that formula $\alpha_a \leq 0 \wedge \alpha_a > 1$ does not admit a valid solution and thus the related path can be safely dropped by the engine since it is unrealizable. On the other hand, other paths admit a valid solution and can be further explored by the engine.

This approach is commonly referred as *eager evaluation* of path constraints since path constraints

Heuristic	Goal
BFS	<i>Maximize coverage</i> [Chipounov et al., 2012, Tillmann and De Halleux, 2008]
DFS	<i>Exhaust paths, minimize memory usage</i> [Cadar et al., 2006, Chipounov et al., 2012] [Tillmann and De Halleux, 2008, Godefroid et al., 2005]
Random path selection	<i>Randomly pick a path with probability based on its length</i> [Cadar et al., 2008]
Code coverage search	<i>Prioritize paths that may explore unexplored code</i> [Cadar et al., 2006, Cadar et al., 2008, Cha et al., 2012] [Chipounov et al., 2012, Groce and Visser, 2002]
Buggy-path-first	<i>Prioritize bug-friendly path</i> [Avgerinos et al., 2011]
Loop exhaustion	<i>Fully explore specific loops</i> [Avgerinos et al., 2011]
Symbolic instruction pointers	<i>Prioritize paths with symbolic instruction pointers</i> [Cha et al., 2012]
Symbolic memory accesses	<i>Prioritize paths with symbolic memory accesses</i> [Cha et al., 2012]
Fitness function	<i>Prioritize paths based on a fitness function</i> [Xie et al., 2009, Cadar and Sen, 2013, Xie et al., 2009]
Subpath-guided search	<i>Use frequency distributions of explored subpaths to prioritize less covered parts of a program</i> [Li et al., 2013]

Figure 10: Common path selection heuristics discussed in literature.

are eagerly checked at each branch and is typically the default approach adopted by most symbolic engines. We refer to Section 7 for a discussion of the possible benefits given by the opposite strategy, i.e., *lazy evaluation*, where path constraints are lazily checked in order to possibly reduce the burden on the solver.

6.2 Bounding Computational Resources

Another common approach is to limit the amount of resources symbolic execution is allowed to use. For instance, the computation may time out after a certain amount of time. Since only a fraction of paths may be explored, the search should be prioritized by looking at the most promising paths first. There are several strategies for selecting or generating the next path to be explored. We now briefly overview some of the most interesting techniques that have been shown to be effective in the literature.

Search Heuristics. Given a set of unexplored paths, a search heuristic should select the most promising path to explore. Many works have introduced novel search strategies, showing their effectiveness in specific application contexts. These heuristics have often been tailored to help the symbolic engine achieve a specific goal (e.g., overflow detection). Finding a universally optimal strategy for prioritizing path exploration remains an open problem. Table 10 provides a sample of prominent search heuristics discussed in prior works.

The most common strategies are *depth-first search* (DFS) and *breadth-first search* (BFS). DFS continuously expands a path as much as possible, before backtracking to the deepest unexplored branch. BFS explores all unexplored paths in parallel, repeatedly expanding each of them by a fixed slice. DFS is often adopted for minimizing the memory usage of the symbolic engine: since the chosen path will be sooner or later fully explored, the memory needed for keeping its state will be released as well. Unfortunately, paths containing loops and recursive calls can easily stall the symbolic engine. For this reason, some tools prefer prioritizing paths using BFS. Although memory pressure can be higher, this strategy may allow an engine to quickly explore diverse paths and possibly detecting interesting behaviors. On the other hand, if the ultimate goal requires to fully terminate the exploration of one or more paths, BFS may take a very long time.

Another popular strategy is *random path selection* that, as its name would suggest, randomly picks a path for exploration. This heuristic has been refined in several variants. For instance, KLEE [Cadaru et al., 2008] assigns probabilities to paths based on their length and on the branch arity. Namely, it favors paths that have been explored fewer times, preventing starvation caused by loops and other path explosion factors.

Several works, such as EXE [Cadaru et al., 2006], KLEE [Cadaru et al., 2008], MAYHEM [Cha et al., 2012], and S²E [Chipounov et al., 2012], have discussed heuristics aimed at maximizing code coverage. For instance, the *coverage optimize search* discussed in KLEE [Cadaru et al., 2008] computes a weight for each state and then randomly selects a state according to these weights. The weight is obtained by taking into account the minimum distance to an uncovered instruction, the call stack of the state, and whether the state has recently covered new code. Of a similar flavor is the heuristic proposed in [Li et al., 2013], called *subpath-guided search*, which attempts to explore *less traveled* parts of a program by selecting the subpath of the control flow graph that has been explored fewer times. This is achieved by maintaining a frequency distribution of explored subpaths, where a subpath is defined as a consecutive subsequence of length n from a complete path. Interestingly, the value n plays a crucial role with respect to the code coverage achieved by a symbolic engine using this heuristic and no specific value has been shown to be universally optimal.

Other search heuristics try to prioritize paths likely leading to states that are *interesting* according to some goal. For instance, the *buggy-path first* strategy in AEG [Avgerinos et al., 2011] picks paths whose past states have contained small but unexploitable bugs. The intuition is that if a path contains some small errors, it is likely that it has not been properly tested. There is thus a good chance that future states may contain interesting, and hopefully exploitable, bugs. Similarly, the *loop exhaustion* strategy discussed in AEG [Avgerinos et al., 2011] explores paths that visit loops. This approach is inspired by the practical observation that common programming mistakes in loops may lead to buffer overflows or other memory-related errors. In order to find exploitable bugs, MAYHEM [Cha et al., 2012] instead gives priority to paths where symbolic memory accesses are identified or symbolic instruction pointers are detected.

Fitness functions have been largely used in the context of search-based test generation [McMinn, 2004]. A fitness function measures how close an explored path is to achieve the target test coverage. Several papers, e.g., [Xie et al., 2009, Cadaru and Sen, 2013, Xie et al., 2009], have applied this idea in the context of symbolic execution. As an example, [Xie et al., 2009] introduces *fitnex*, a strategy for concolic execution that prioritizes paths that are *closer* to take a specific branch. In more detail, given a branch condition of the form $|a - c| == 0$ and a path that has reached the branch, *fitnex* computes a closeness equal to $|a - c|$ by leveraging the concrete values of the two variables a and c in that path. Similar fitness values can be computed for other kinds of branch conditions. The path with the lowest fitness value for a branch is selected by the symbolic engine. Paths that have not reached the branch yet get the worst-case fitness value.

Dynamic Test Generation. Traditional symbolic execution does not scale over large programs. Although search heuristics may help prioritize some interesting paths, symbolic execution may still proceed extremely slow. Indeed, the engine must simulate any instruction in the program and heavily relies on the constraint solver in order to make any progress in the execution. *Dynamic test generation*, initially introduced in DART [Godefroid et al., 2005], is a technique that can help symbolic execution scale to large programs. The main idea is to execute a program both concretely and symbolically. This kind of execution is often referred to as concolic execution (Section 2.1). Initially, a random input is generated and a concrete execution is started. In parallel, a symbolic execution is also started. Whenever the concrete execution takes a branch, the symbolic execution is directed toward the same branch and the constraints extracted from the branch condition are added to the current set of path constraints. In other words, the symbolic execution is driven by a specific concrete execution. A consequence of this approach is that the symbolic engine does not need to invoke the constraint solver to decide whether a branch condition is (un)satisfiable, since this is directly tested by the concrete execution. Other paths of the program execution can be then explored by selecting a previously taken branch and negating its constraints. Using a

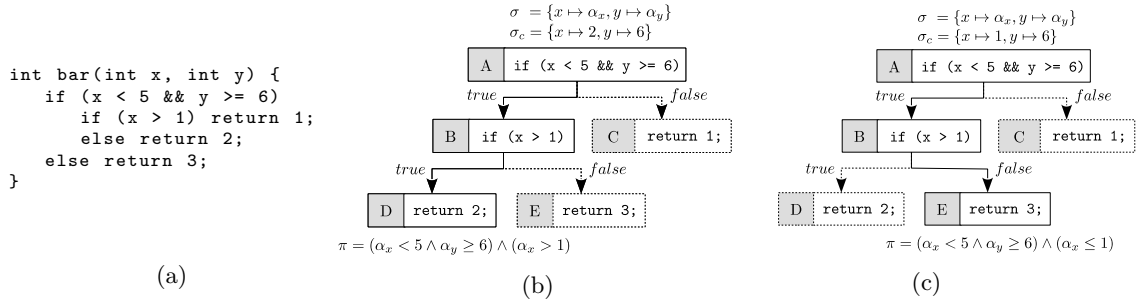


Figure 11: Dynamic test generation example. (a) Source code of function `bar`. (b) Symbolic execution tree for the function `bar`. Solid edges show the path taken by the concolic execution when $x = 2$ and $y = 6$. These input values have been randomly chosen. (c) Concolic execution when $x = 1$ and $y = 6$. These input values have been obtained using a constraint solver, after negating the path constraints of node `B` in the function `bar`.

constraint solver, the symbolic engine can generate a new (random) input that drives the concrete execution toward a new path. This strategy can be repeated as much as needed to achieve the desired coverage.

Example. An example of dynamic test generation is shown in Figure 11. Consider the function `bar` (Figure 11a) that takes two integer inputs x and y . To start a first concrete execution, a symbolic engine may initially randomly pick $x = 2$ and $y = 6$ as input values. The concrete execution induced by these inputs is presented in Figure 11b: both the first and second branch condition (nodes `A` and `B`) are satisfied, directing the execution toward the first `return` statement (node `D`). Nodes `C` and `E` are skipped since their associated branch conditions are not met by the current input values. For instance, node `E` is not executed since the condition $x > 1$ (node `B`) directs the path toward the node `D`. An engine that desires to symbolically execute a path containing the node `E` must track the constraints during the concrete execution over $x = 2$ and $y = 6$ and then negate the branch condition $x > 1$. To generate a new input, the engine then invokes a solver over the constraints $\neg(\alpha_x > 1) \wedge (\alpha_x < 5 \wedge \alpha_y > 5)$, getting, e.g., $x = 1$ and $y = 6$ (Figure 11c). Notice that since y is not involved in the branch condition that is currently negated, the engine may reuse its value and include an additional constraint $\alpha_y = 6$. This optimization may drastically reduce the solving time required to obtain a solution from the constraint solver.

Although dynamic test generation uses concrete inputs to drive the symbolic execution toward a specific path, it still needs to pick a branch to negate whenever a new path has to be explored. Notice also that each concrete execution may add new branches that will have to be visited. Since the set of non-taken branches across all the performed concrete executions can be very large, the search heuristics discussed in Section 6.2 still play a crucial role. For instance, DART [Godefroid et al., 2005] chooses the next branch to negate using a DFS strategy. Additional strategies for picking the next branch to negate have been presented in literature. For instance, the *generational search* algorithm discussed in SAGE [Godefroid et al., 2008] systematically yet partially explores the state space, maximizing the number of new tests generated while also avoiding redundancies in the search. This is achieved by negating constraints following a specific order and by limiting the backtracking of the search algorithm. Since the state space is only partially explored, the initial input plays a crucial role in the effectiveness of the overall approach. The importance of the first input is similar to what happens in traditional *black-box fuzzing* and, for this reason, symbolic engines such as SAGE are often referred as *white-box fuzzers*.

6.3 Under-Constrained Symbolic Execution

A possible approach to avoid path explosion is to cut the code to check, say a function, out of its enclosing system and check it in isolation. Lazy initialization with user-specified preconditions (Section 3.4) follows this principle in order to automatically reconstruct complex data structures.

However, taking a code region out of an application has proven to be quite difficult due to the entanglements with the surrounding environment [Engler and Dunbar, 2007].

The main issue is that errors detected in a function analyzed in isolation may be false positives, as the input may never assume certain values when the function is executed in the context of a full program. Some prior works, e.g., CHECK 'N' CRASH [Csallner and Smaragdakis, 2005], first analyze the code in isolation and then test the generated crashing inputs using concrete executions.

Under-constrained symbolic execution [Engler and Dunbar, 2007] is a twist on symbolic execution that allows for the analysis of a function in isolation by marking some symbolic inputs as *under-constrained*. Intuitively, a symbolic variable is under-constrained when in the analysis we do not account for constraints on its value that should have been collected along the path prefix from the program's entry point to the function to analyze. Under-constrained variables have the same semantics as classic symbolic variables except when used in an expression that can cause an error to occur. In particular, an error is reported only if all the solutions for the currently known constraints on the variable cause it to occur, i.e., the error is context-insensitive and thus a true positive. Otherwise, its negation is added to the path constraints and execution resumes as normal. This choice can be regarded as an attempt to reconstruct preconditions from the checks inserted in the code: any subsequent action violating an added negated constraint will be reported as an error.

Although this technique is not sound as it may miss errors, it can still scale to find interesting bugs in larger programs. Also, the application of under-constrained symbolic execution is not limited to functions only: for instance, if a code region (e.g., a loop) may be troublesome for the symbolic executor, it can be skipped by marking the locations affected by it as under-constrained.

6.4 Preconditioned Symbolic Execution

AEG [Avgerinos et al., 2011] proposes *preconditioned symbolic execution* as a novel method to drive symbolic execution toward certain subsets of the input space. A subset is determined by a precondition predicate: inputs that do not satisfy it will not be explored. The intuition behind preconditioned symbolic execution is that we can narrow down the state space we are exploring by specifying goal-oriented (e.g., exploitability) conditions as a precondition, e.g., all symbolic inputs should have the maximum size to expose buffer overflow bugs. The main benefit of preconditioned symbolic execution is simple: by limiting the size of the input state space before execution begins, we can prune program paths and therefore explore the target program more efficiently. Indeed, preconditions need to be selected carefully. If a precondition is too specific, no bug may be detected; if it is too general, almost the entire state space will be explored.

This technique enforces a precondition by adding the precondition constraints to the path predicate during initialization. Adding constraints may seem counterintuitive since there are more checks to perform at branch points during symbolic execution. However, the state space shrink caused by the precondition constraints outweighs the decision procedure overhead at branching points. When the precondition for a branch is unsatisfiable, neither further checks are made, nor a new execution engine instance is forked for the branch.

Four prominent categories of preconditions are:

- *None*: there is no precondition, thus space exploration proceeds as normal;
- *Known length*: symbolic inputs are of known maximum length, e.g., a network packet has a fixed size, or static analysis can determine the input length;
- *Known prefix*: symbolic inputs have a known prefix, e.g., a fixed header string such as the initial *magic code* in a binary, or a network packet header.
- *Fully known*: all input bytes are concrete, as in a concolic execution; it can be used, for instance, to generate a working exploit from a known crashing input.

<pre> // N symbolic branches if (input[0] < 42) [...] [...] if (input[N-1] < 42) [...] // symbolic loop strcpy(dest, input); // M symbolic branches if (input[N] < 42) [...] [...] if (input[N+M-1] < 42) [...] </pre>	<pre> 1. void foo(int x, int y) { 2. if (x < 5) 3. y = y * 2; 4. else 5. y = y * 3; 6. return y; 7. } </pre>
(a)	(b)

Figure 12: (a) Preconditioned symbolic execution example [Avgerinos et al., 2011]; (b) State merging example

Example. Consider the code fragment in Figure 12a where `input` is an array of $s \geq n + m$ bytes. The impact of preconditions on the state space size is as follows:

- *None*: the input space size is 256^s , and up to $2^n \cdot s \cdot 2^m$ execution engine instances are needed, due to $n + m$ symbolic branches and up to S loop iterations;
- *Known length*: if we assume a string length s , i.e., the first $(s - 1)$ bytes of `input` are not $\backslash 0$, the loop is concretized, and the state space size is reduced to 2^{n+m} ;
- *Known prefix*: if a prefix of $p < n$ bytes is known for `input`, the first p branches and p loop iterations are concrete, and the state space size becomes $2^{n-p} \cdot s \cdot 2^m$;
- *Fully known*: as all input bytes are concrete, the state space size is trivially 1.

6.5 State Merging

Several static program analysis techniques such as abstract interpretation merge states corresponding to different paths into a state that over-approximates them. In a precise symbolic execution, however, merging is not allowed to introduce any approximation or abstraction, and therefore can only change formulas to have them characterize sets of execution paths. In other words, a merged state will be described by a formula that represents the disjunction of the formulas that would have described the individual states if they were kept separate.

Example. Consider the function of Figure 12b and its symbolic execution tree shown in Figure 13a. Initially (execution state A) the path constraints are *true* and input arguments x and y are associated with symbolic values α_x and α_y , respectively. Line 2 contains a conditional branch and the execution is forked: depending on the branch taken, a different statement is evaluated next and different assumptions are made on symbol α_x (execution states C and D , respectively). After expanding every execution state until the `return` at line 6 is reached on all branches, the symbolic execution tree gets populated with two additional states D and E . If a symbolic execution engine desires to reduce the number of active states, then state merging can be performed. For instance, Figure 13b shows the symbolic execution DAG for the same piece of code when a state merging operation is performed before evaluating the `return` statement at line 6: D' is now a merged state that fully captures the former execution states D and E using the *ite* expression $ite(\alpha_x < 5, 2 * \alpha_y, 3 * \alpha_y)$ (Section 3.1). Note that the path constraints of the execution states D and E can be merged into the disjunction formula $\alpha_x < 5 \vee \alpha_x \geq 5$ and then simplified to *true* in D' .

Trade-Offs: to Merge or Not to Merge? Early works [Godefroid, 2007, Hansen et al., 2009] have shown that merging techniques effectively decrease the number of paths to explore, but

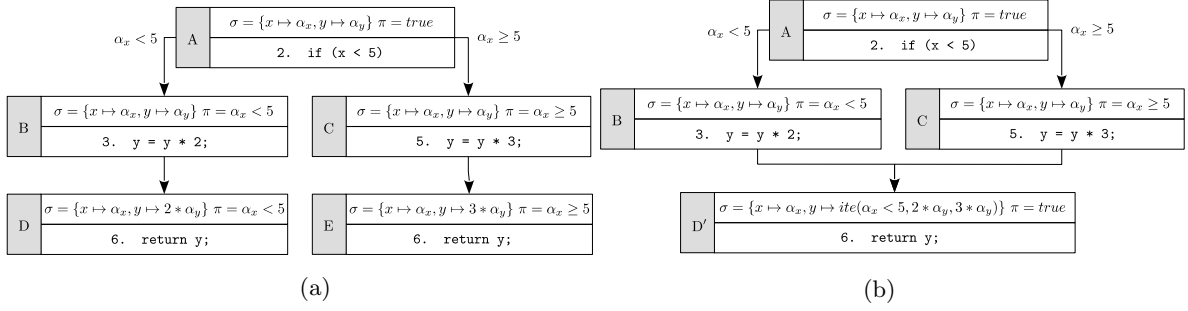


Figure 13: Symbolic execution of function `foo` of Figure 12b: (a) without state merging; (b) with state merging.

also put a burden on constraints solvers, which typically encounter difficulties when dealing with disjunction. Merging can also introduce new symbolic expressions in the code, for instance when merging different concrete values from a conditional assignment into a symbolic expression over the condition. [Kuznetsov et al., 2012] provides an excellent discussion of the design space of state merging techniques. At the one end of the spectrum, complete separation of the paths does not perform any merge and is used in search-based symbolic execution (Section 6.2). At the other end, static state merging combines states at control-flow join points, thus essentially representing a whole program with a single formula. Static state merging is used in whole-program verification condition generators, e.g., [Xie and Aiken, 2005, Babic and Hu, 2008]), which typically trade precision for scalability by, for instance, unrolling loops only once.

Merging Heuristics. Intermediate merging solutions adopt heuristics to identify state merges that can speed the exploration process up. Indeed, generating larger symbolic expressions and possibly extra solvers invocations can outweigh the benefit of having fewer states, leading to poorer overall performance [Hansen et al., 2009, Kuznetsov et al., 2012]. *Query count estimation* [Kuznetsov et al., 2012] relies on a simple static analysis to identify how often each variable is used in branch conditions past any given point in the CFG. The estimate is used as a proxy for the number of solver queries that a given variable is likely to be part of. Two states make a good candidate for merging when their differing variables are expected to appear infrequently in later queries. *Veritestng* [Avgerinos et al., 2014] implements a form of merging heuristic based on a distinction between easy and hard statements. Hard statements are those that involve system calls, indirect jumps, and other operations that are difficult to reason about statically. Static merging is performed on sequences of easy statements, whose effects are captured using *ite* expressions, while per-path symbolic exploration is done whenever a hard-to-analyze statement is encountered.

Dynamic State Merging. Ideally, in order to maximize the opportunities for merging, a symbolic execution engine should traverse a CFG so that a combined state for a program point can be computed from its predecessors, e.g., if the graph is acyclic, by following a topological ordering. However, this would prevent search exploration strategies aiming at prioritizing more “interesting” states over others. [Kuznetsov et al., 2012] introduces *dynamic state merging* to identify opportunities for merging regardless of the exploration order imposed by the search strategy. Suppose the symbolic engine maintains a worklist of states and a bounded history of their predecessors. When the engine has to pick the next state to explore, it first checks whether there are two states s_1 and s_2 from the worklist such that they do not match for merging, but s_1 and a predecessor of s_2 do. If the expected similarity between s_2 and a successor of s_1 is also high, the algorithm attempts a merge by advancing the execution of s_1 for a fixed number of steps. This captures the idea that if two states are similar, then also their respective successors are likely to become similar in a few steps. If the merge fails, the algorithm lets the search heuristic pick the next state to explore.

6.6 Leveraging Program Analysis and Optimization Techniques

A deeper understanding of a program’s behavior can help a symbolic engine to focus on promising states, e.g., by pruning uninteresting parts of the computation tree. Several classical program analysis techniques have been explored in the symbolic execution literature. Some prominent examples are listed below:

- *Program slicing* is a method that, starting from a subset of a program’s behavior, extracts from the program the minimal sequence of instructions that faithfully represents that behavior [Weiser, 1984]. This information can help a symbolic engine in several ways: for instance, given a program slice related to a target program point, symbolic exploration can be restricted to paths contained in the program slice. We discuss an example of use in Section 9.3.
- *Taint analysis* attempts to identify variables that can be modified by predefined taint sources such as user input and can be performed both statically and dynamically. Dynamic taint analysis typically yields more accurate results, and is often employed to explore symbolically only the parts of an execution that depend upon tainted values [Schwartz et al., 2010].
- *Fuzzing* is a software testing technique that randomly mutates user-provided test inputs to cause crashes or assertion failures and find potential memory leaks. Fuzzing, as discussed in Section 6.2, can be augmented with symbolic execution to collect constraints for an input and negate them to generate new inputs. On the other hand, a symbolic executor can be augmented with fuzzing to reach deeper states in the exploration more quickly and efficiently: we present two embodiments of this approach in Section 9.1.
- *Branch predication* is a strategy for mitigating misprediction penalties in pipelined executions by avoiding jumps over very small sections of code: for instance, control-flow forking constructs such as the C ternary operator can be replaced with a predicated `select` instruction. [Collingbourne et al., 2011] reports an exponential decrease in the number of paths to explore from the adoption of this strategy when cross-checking two implementations of a program using symbolic execution.
- *State matching* determines whether an abstract state is subsumed by another, and can be used to analyze an under-approximation of a program’s behavior. For instance, [Anand et al., 2006, Visser et al., 2006] explore different heap shapes in the context of test generation for data structures, using subsumption checking to determine whether a symbolic state is being revisited.
- *Type checking* can be effectively mixed with symbolic analysis [Khoo et al., 2010]; for instance, type checking can determine the return type of a function that is difficult to analyze symbolically: such information can then potentially be used by the executor to prune certain paths².

7 Constraint solving

Constraint satisfaction problems arise in many domains, including analysis, testing, and verification of software programs. Constraint solvers are decision procedures for problems expressed in logical formulas: for instance, the boolean satisfiability problem (also known as SAT) aims at determining whether there exists an interpretation of the symbols of a formula that makes it true. Although SAT is a well-known NP-complete problem, recent advances have moved the boundaries for what is intractable when it comes to practical applications [De Moura and Bjørner, 2011].

²Interestingly, [Khoo et al., 2010] discusses also how symbolic analysis can help a type checker. For instance, a symbolic engine can provide context-sensitive properties over a variable that would rule out certain type errors, improving the precision of the type checker.

Observe that some problems are more naturally described with languages that are more expressive than the one of boolean formulas with logical connectives. For this reason, satisfiability modulo theories (SMT) generalize the SAT problem with supporting theories to capture formulas involving, for instance, linear arithmetic inequalities and operations over arrays (see, e.g., Section 3.1). SMT solvers map the atoms in an SMT formula to fresh boolean variables: a SAT decision procedure checks the rewritten formula for satisfiability, and a theory solver checks the model generated by the SAT procedure.

In a symbolic executor, constraint solving plays a crucial role in checking the feasibility of a path, generating assignments to symbolic variables, and verifying assertions. The two most popular solvers used in symbolic executors are STP and Z3. STP [Ganesh and Dill, 2007, Ganesh, 2007] is an SMT solver with bitvector and array theories initially developed at Stanford and employed in, e.g., EXE [Cadaru et al., 2006], KLEE [Cadaru et al., 2008], MINESWEEPER [Brumley et al., 2008], and AEG [Avgerinos et al., 2011]. Z3 [De Moura and Bjørner, 2008] is an SMT solver developed at Microsoft with support for nonlinear arithmetic, bitvector, and array theories, and is used in, e.g., MAYHEM [Cha et al., 2012], SAGE [Godefroid et al., 2012], and ANGR [Shoshitaishvili et al., 2016]. CVC3 [Barrett and Tinelli, 2007] is another SMT solver that supports theories for linear arithmetic, bitvectors, arrays, and quantifiers, and is employed in JAVA PATHFINDER [Păsăreanu and Rungta, 2010] along with CHOCO [Prud’homme et al., 2015] for integer/real constraints and CORAL [Souza et al., 2011] for complex mathematical constraints. Modern symbolic executors can typically choose between different underlying solvers through a common API, and also resort to a native interface to a specific solver for better performance.

However, despite the significant advances observed over the past few years – which also made symbolic execution practical in the first place [Cadaru and Sen, 2013] – constraint solving remains one of the main obstacles to the scalability of symbolic execution engines.

In the remainder of this section, we address different techniques to extend the range of programs that can be handled by symbolic execution and to optimize the performance of constraint solving. Two prominent approaches consist in: (i) reducing the size and complexity of the constraints to check, and (ii) unburdening the solver by, e.g., resorting to constraint solution caching, deferring of constraint solver queries, or concretization.

Constraint Reduction. A common optimization approach followed by both solvers and symbolic executors is to reduce constraints into simpler forms. For example, the *expression rewriting* optimization can apply classical techniques from optimizing compilers such as constant folding, strength reduction, and simplification of linear expressions (see, e.g., KLEE [Cadaru et al., 2008]).

EXE [Cadaru et al., 2006] introduces a *constraint independence* optimization that exploits the fact that a set of constraints can frequently be divided into multiple independent subsets of constraints. This optimization interacts well with query result caching strategies, and offers an additional advantage when an engine asks the solver about the satisfiability of a specific constraint, as it removes irrelevant constraints from the query. In fact, real programs typically have many independent branches, which can often introduce irrelevant constraints that add up quickly.

Another fact that can be exploited by reduction techniques is that the natural structure of programs can lead to the introduction of more specific constraints for some variables as the execution proceeds. As path conditions are generated by conjoining new terms to an existing sequence, it might become possible to rewrite and optimize existing constraints. For instance, adding an equality constraint of the form $x := 5$ enables not only the simplification to true of other constraints over the value of the variable (e.g., $x > 0$), but also the substitution of the symbol x with the associated concrete value in the other subsequent constraints involving it. The latter optimization is also known as *implied value concretization* and, for instance, it is employed by KLEE [Cadaru et al., 2008].

In a similar spirit, S²E [Chipounov et al., 2012] introduces a bitfield-theory expression simplifier to replace with concrete values parts of a symbolic variable that are masked away by bit operations. For instance, if x is a 4-bit symbolic value, the most significant bit in the expression $x | 1000$ is known to be one. The simplifier can propagate information across the tree representation of an expression, and if all bits in an expression are known, it replaces the expression with the

```

1. int non_linear(int v) {
2.     return (v*v) % 50;
3. }

4. void test(int x, int y) {
5.     z = non_linear(y);
6.     if (z == x) {
7.         if (x > y + 10) ERROR;
8.     }
9. }

```

Figure 14: Example with non-linear constraints.

corresponding constant.

Reuse of Constraint Solutions. The idea of reusing previously computed results to speed up constraint solving can be particularly effective in the setting of a symbolic executor, especially when combined with other techniques such as constraint independence optimization. Most reuse approaches for constraint solving are currently based on syntactic or semantic equivalence of the constraints.

EXE [Cadaru et al., 2006] caches the results of satisfiability queries and constraint solutions in order to avoid calling the solver when possible. A cache is managed by a server process that can receive queries from multiple parallel instances of the execution engine, each exploring a different program state.

KLEE [Cadaru et al., 2008] implements an incremental optimization strategy called *counterexample caching*. A tree representation-based cache maps constraint sets to concrete variable assignments, or to a special null value when a constraint set is unsatisfiable. When an unsatisfiable set in the cache is a subset for a given constraint set S , S is deemed unsatisfiable as well. Conversely, when the cache contains a solution for a superset of S , the solution trivially satisfies S too. Finally, when the cache contains a solution for one or more subsets of S , the algorithm tries substituting in all the solutions to check whether a satisfying solution for S can be found.

Memoized symbolic execution [Yang et al., 2012] introduces a new approach for a more efficient application of symbolic execution. The work is motivated by the observation that applying symbolic execution often requires several successive runs of the technique on largely similar underlying problems, e.g., finding a bug and then examining the program again to check the validity of the fix. A trie-based data structure compactly encodes the choices taken when exploring different paths, allowing successive runs to reuse previously computed results where possible.

The Green framework [Visser et al., 2012] explores constraint solution reuse across runs of not only the same program, but also similar programs, different programs, and different analyses. Constraints are distilled into their essential parts through a *slicing* transformation and represented in a canonical form to achieve good reuse, even within a single analysis run. [Jia et al., 2015] presents an extension to the framework that supports constraint reuse based on the logical implication relations among constraints, leading to better reuse and faster execution time for the symbolic analysis.

Lazy Constraints. [Ramos and Engler, 2015] adopts a timeout approach for constraint solver queries. In their initial experiments, the authors traced most timeouts to symbolic division and remainder operations, with the worst cases occurring when an unsigned remainder operation had a symbolic value in the denominator. They thus implemented a solution that works as follows: when the executor encounters a branch statement involving an expensive symbolic operation, it will take both the true and false branches and add a *lazy* constraint on the result of the expensive operation to the path conditions. When the exploration reaches a state that satisfies some goal (e.g., an error is found), the algorithm will check for the feasibility of the path, and suppress it if deemed as unreachable in a real execution.

Compared to the *eager* approach of checking the feasibility of a branch as encountered (Section 6.1), a *lazy* strategy may lead to a larger number of active states, and in turn to more solver queries. However, the authors report that the delayed queries are in many cases more efficient than their eager counterparts: the path constraints added after a lazy constraint can in fact narrow down the solution space for the solver.

Concretization. [Cadaru and Sen, 2013] discusses limitations of classical symbolic execution in the presence of formulas that cannot be (efficiently) solved by constraint solvers.

Example. In the code fragment of Figure 14, the engine stores a non-linear constraint of the form $\alpha_x = (\alpha_y * \alpha_y) \% 50$ for the *true* branch at line 6. A solver that does not support non-linear arithmetic fails to generate any input for the program.

A concolic executor generates some random input for the program and executes it both concretely and symbolically: a possible value from the concrete execution can be used for a symbolic operand involved in a formula that is inherently hard for the solver, albeit at the cost of sacrificing soundness in the exploration. For instance, in the presence of three nested branches with only one being nonlinear, DART [Godefroid et al., 2005] starts from a random valid input for the function, and then alters it when symbolically exploring the two linear branches. The work resorts to concretization also to avoid performing expensive or imprecise alias analysis on pointers.

To partially overcome the incompleteness due to concretization, [Păsăreanu et al., 2011] suggests to consider *all* the path constraints collectable over a path before binding one or more symbols to specific concrete values. Indeed, DART [Godefroid et al., 2005] concretizes symbols based on the path constraints collected up to a target branch. In this manner, a constraint contained in a subsequent branch in the same path is not considered and it may be not satisfiable due to already concretized symbols. If this happens, DART restarts the execution with different random concrete values, hoping to be able to satisfy the subsequent branch. The approach presented in [Păsăreanu et al., 2011] requires instead to detect *solvable* constraints along a full path and to delay concretization as much as possible.

8 Symbolic execution of binary code

The importance of performing symbolic analysis of program properties on binary code is on the rise for a number of reasons. Binary code analysis is attractive as it reasons on code that will actually execute: not requiring the source code significantly extends the applicability of such techniques (to, e.g., common off-the-shelf proprietary programs, firmwares for embedded systems, and malicious software), and it gives the ground truth important for security applications whereas source code analysis may yield misleading results due to compiler optimizations [Song et al., 2008]. Binary analysis is relevant also for programs written in dynamic languages, as nowadays they are executed in modern runtimes that deeply transform and optimize their code before just-in-time compiling it.

Analyzing binary code is commonly seen as a challenging task due to its complexity and lack of a high-level semantics. Modern architectures offer complex instruction sets: modeling each instruction can be difficult, especially in the presence of multiple side effects on processor flags to determine branch conditions. The second major challenge comes from the lack of the higher-level semantics present in source code, especially when no debugging information is available. Types are not explicitly encoded in binary code: even with register types, it is common to store values as one type and read them as another. Similar considerations can be made for array bounds as well. Also, control-flow graph information is not explicitly available, as control flow is performed through jump instructions at both inter- and intra-procedural level. The function abstraction at the binary level does not exist as we intend it at source-code level: functions can be separated in non-contiguous pieces, and code may also call in the middle of a code block generated for a source-level function.

In the remainder of this section we provide an overview of how symbolic executors can address some of the most significant challenges in the analysis of binary code.

8.1 Lifting to an Intermediate Representation

Motivated by the complexity in modeling native instructions and by the variety of architectures on which applications can be deployed (e.g., x86, x86-64, ARM, MIPS), symbolic executors for




Figure 15: Lowering and lifting processes in native vs. source code processing.

binary code typically rely on a *lifter* that transforms native instructions into an *intermediate representation* (IR), also known as *bytecode*. Modern compilers such as LLVM [Lattner and Adve, 2004] typically generate IR by *lowering* the user-provided source code during the first step of compilation, optimize it, and eventually lower it to native code for a specific platform. Source-code symbolic executors can resort to compiler-assisted lowering to reason on bytecode rather than source-language statements: for instance, KLEE [Cadar et al., 2008] reasons on the IR generated by the LLVM compiler for static languages such as C and C++. Figure 15 summarizes the relationships between source code, IR, and binary code.

Reasoning at the intermediate representation level allows program analyses to be written in an architecture-independent fashion. Translated instructions will always expose all the side-effects of a native instruction, and support for additional platforms can be added over time. A number of symbolic executors use VEX, the intermediate representation format from the Valgrind dynamic instrumentation framework [Nethercote and Seward, 2007]. VEX is a RISC-like language designed for program analysis that offers a compact set of instructions for expressing programs in static single assignment (SSA) form [Cytron et al., 1991]. Lifters are available for both 32-bit and 64-bit ARM, MIPS, PPC, and x86 binaries.

ANGR [Shoshitaishvili et al., 2016] performs analysis directly on VEX IR. Authors chose VEX over other IR formats as at that time it was the only choice that offered a publicly available implementation with support for many architectures. Also, they mention that writing a binary lifter can be a daunting task, and a well-documented and program analysis-oriented solution can be a bonus. BITBLAZE [Song et al., 2008] uses VEX too, although it translates it to a custom intermediate language. The reason for this is that VEX captures the side effects of some instructions only implicitly, such as what EFLAGS are set by x86 instructions: translating it to a custom language simplified the development of BITBLAZE’s analysis framework.

The authors of S²E [Chipounov et al., 2012] have implemented an x86-to-LLVM-IR lifter in order to use the KLEE [Cadar et al., 2008] symbolic execution engine for whole-system symbolic analysis of binary code in a virtualized environment. The translation is transparent to both the guest operating system and KLEE, thus enabling the analysis of binaries using the full power of KLEE. Another x86-to-LLVM-IR lifter that can be used to run KLEE on binary code is `mcsema`³.

8.2 Reconstructing the Control Flow Graph

A control flow graph (CFG) can provide valuable information for a symbolic executor as it captures the set of potential control flow transfers for all feasible execution paths. A fundamental issue that arises when reconstructing CFGs for binaries is that the possible targets of an indirect jump may not be identified correctly. Direct jumps are straightforward to process: as they encode their targets explicitly in the code, successor basic blocks can be identified and visited until no new edge is found. The target of an indirect jump is determined instead at run time: it might be computed by carrying out a calculation (e.g., a jump table) or depend on the current calling context (e.g., a function pointer is passed as argument, or a virtual C++ method is invoked).

CFG recovery is typically an iterative refinement process based on a number of program analysis techniques. For instance, Value Set Analysis (VSA) [Duesterwald, 2004] is a technique that can be used to identify a tight over-approximation of certain program state properties (e.g., the set of possible targets of an indirect jump or a memory write). In BITBLAZE [Song et al., 2008] an initial CFG is generated by inserting special successor nodes for unresolved indirect jump targets. This choice is conceptually similar to widening a fact to the bottom of a lattice in a data-flow

³<https://github.com/trailofbits/mcsema>.

analysis. When an analysis requires more precise information, VSA is then applied on demand.

ANGR [Shoshitaishvili et al., 2016] implements two algorithms for CFG recovery. An iterative algorithm starts from the entry point of the program and interleaves a number of techniques to achieve speed and completeness, including VSA, inter-procedural backward program slicing, and symbolic execution of blocks. This algorithm is however rather slow and may miss code portions reachable only through unresolved jump targets. The authors thus devise a fast secondary algorithm that uses a number of heuristics to identify functions based on prologue signatures, and performs simple analyses (e.g., a lightweight alias analysis) to solve a number of indirect jumps. The algorithm is context-insensitive, so it can be used to quickly recover a CFG without a concern for understanding the reachability of functions from one another.

8.3 Code Obfuscation

In recent years, code obfuscation has received considerable attention as a cheap way to hinder the understanding of the inner workings of a proprietary program. Obfuscation is employed not only to thwart software piracy and improve software security, but also to avoid detection and resist analysis for malicious software [Udupa et al., 2005, Yadegari et al., 2015].

A significant motivation behind using symbolic/concolic execution in malware analysis is to get around code obfuscations. However, current analysis techniques have trouble dealing with some of those obfuscations, leading to imprecision and/or excessive resource usage [Yadegari and Debray, 2015]. For instance, obfuscation tools can transform conditional branches into indirect jumps that symbolic analysis find difficult to analyze, while run-time code self-modification might conceal conditional jumps on symbolic values so that they are missed by the analysis.

A few works have described obfuscation techniques aiming at thwarting symbolic execution. [Sharif et al., 2008] introduces a *conditional code obfuscation* scheme based on one-way hash functions that makes it hard to identify the values of symbolic variables for which branch conditions are satisfied. They also present an encryption scheme for the code to execute based on a key derived from the value that satisfies a branch condition. [Wang et al., 2011] takes a step forward by proposing an obfuscation technique that works in spite of the fact that it uses linear operations only, for which symbolic execution usually works well. The obfuscation tool inserts a simple loop incorporating an unsolved mathematical conjecture that converges to a known value after a number of iterations, and the produced result is then combined with the original branch condition.

[Hai et al., 2016] presents BE-PUM, a tool to generate a precise CFG in the presence of obfuscation techniques that are common in the malware domain, including indirect jumps, structured exception handlers (SEHs), overlapping instructions, and self-modifying code. While engines such as BITBLAZE [Song et al., 2008] typically rely on existing disassemblers like IDA Pro⁴ for obfuscated code, BE-PUM relies on concolic execution for deobfuscation, using a binary emulator for the user process and stubs for API calls.

[Yadegari and Debray, 2015] discusses the limitations of symbolic execution in the presence of three generic obfuscation techniques: (1) conditional-to-indirect jump transformation, also known as *symbolic jump problem* [Schwartz et al., 2010]; (2) conditional-to-conditional jump transformation, where the predicate is deeply changed; and (3) symbolic code, when code modification is carried out using an input-derived value. The authors show how resorting to bit-level taint analysis and architecture-aware constraint generation can allow symbolic execution to circumvent such obfuscations.

9 Sample Applications

The last decade has witnessed an increasing adoption of symbolic execution techniques not only in the software testing domain, but also to address other compelling engineering problems such as automatic generation of exploits or authentication bypass. We now discuss three prominent

⁴<https://www.hex-rays.com/products/ida/>.

applications of symbolic execution techniques to these domains. Examples of extensions to other areas can be found, e.g., in [Cadaru et al., 2011].

9.1 Bug Detection

Software testing strategies typically attempt to execute a program with the intent of finding bugs. As manual test input generation is an error-prone and usually non-exhaustive process, automated testing techniques have drawn a lot of attention over the years. Random testing techniques such as fuzzing are cheap in terms of run-time overhead, but fail to obtain a wide exploration of a program state space. Symbolic and concolic execution techniques on the other hand achieve a more exhaustive exploration, but they become expensive as the length of the execution grows: for this reason, they usually reveal shallow bugs only.

[Majumdar and Sen, 2007] proposes *hybrid concolic testing* for test input generation, which combines the ability of random search to reach deep program states with the ability of concolic execution to achieve a wide exploration. The two techniques are interleaved: in particular, when random testing saturates (i.e., it is unable to hit new code coverage points after a number of steps), concolic execution is used to mutate the current program state by performing a bounded depth-first search for an uncovered coverage point. For a fixed time budget, the technique outperforms both random and concolic testing in terms of branch coverage. The intuition behind this approach is that many programs show behaviors where a state can be easily reached through random testing, but then a precise sequence of events - identifiable by a symbolic engine - is required to hit a specific coverage point.

[Stephens et al., 2016] refines this idea and devises a vulnerability excavation tool based on ANGR [Shoshitaishvili et al., 2016], called Driller, that interleaves fuzzing and concolic execution to discover memory corruption vulnerabilities. The authors remark that user inputs can be categorized as *general* input, which has a wide range of valid values, and *specific* input: a check for particular values of a specific input then splits an application into *compartments*. Driller offloads the majority of unique path discovery to a fuzzy engine, and relies on concolic execution to move across compartments. During the fuzzy phase, Driller marks a number of inputs as interesting (for instance, when an input was the first to trigger some state transition) and once it gets stuck in the exploration, it passes the set of such paths to a concolic engine, which preconstraints the program states to ensure consistency with the results of the native execution. On the dataset used for the DARPA Cyber Grand Challenge qualifying event, Driller could identify crashing inputs in 77 applications, including both the 68 and 16 applications for which fuzzing and symbolic execution alone succeeded, respectively. For 6 applications, Driller was the only one to detect a vulnerability.

9.2 Bug Exploitation

Bugs are a consequence of the nature of human factors in software development and are everywhere. Those that can be exploited by an attacker should normally be fixed first: systems for automatically and effectively identifying them are thus very valuable.

AEG [Avgerinos et al., 2011] employs preconditioned symbolic execution to analyze a potentially buggy program in source form and look for bugs amenable to stack smashing or return-into-libc exploits [Pincus and Baker, 2004], which are popular control hijack attack techniques. The tool augments path constraints with exploitability constraints and queries a constraint solver, generating a concrete exploit when the constraints are satisfiable. The authors devise the *buggy-path-first* and *loop-exhaustion* strategies discussed in Section 6.2 to prioritize paths in the search. On a set of 14 Linux applications, AEG could construct 16 control hijack exploits, 2 of which against previously unknown vulnerabilities.

MAYHEM [Cha et al., 2012] takes another step forward by presenting the first end-to-end exploitable bug finding system working on programs in binary form. It adopts a hybrid execution model based on checkpoints and two components: a concrete executor that injects taint-analysis instrumentation in the code and a symbolic executor that takes over when a tainted branch or jump instruction is met. Exploitability constraints for symbolic instruction pointers and format strings

are generated, targeting a wide range of exploits, e.g., SEH-based and jump-to-register ones. Three path selection heuristics help prioritizing paths that are most likely to contain a bug (e.g., those containing symbolic memory accesses or instruction pointers). A virtualization layer intercepts and emulates all the system calls to the host OS, while preconditioned symbolic execution can be used to reduce the size of the search space. Also, restricting symbolic execution to tainted basic blocks only gives very good speedups in this setting, as in the reported experiments more than 95% of the processed instructions were not tainted. MAYHEM was able to find exploitable vulnerabilities in the 29 Linux and Windows applications considered in the evaluation, 2 of which were previously undocumented. Although the goal in MAYHEM is informing the user that an exploitable bug exists, the generated simple exploits can be likely transformed in an automated fashion to work in the presence of OS defenses such as address space layout randomization and data execution prevention [Schwartz et al., 2011].

9.3 Authentication Bypass

Software backdoors are a method of bypassing authentication in an algorithm, a software product, or even in a full computer system. Although sometimes these software flaws are injected by external attackers using subtle tricks such as compiler tampering [Karger and Schell, 1974], there are reported cases of backdoors that have been surreptitiously installed by the hardware and/or software manufacturers [Costin et al., 2014], or even by governments [Zitter, 2013].

Different works [Davidson et al., 2013, Zaddach et al., 2014, Shoshitaishvili et al., 2015] have exploited symbolic execution for analyzing the behavior of binary firmwares. Indeed, an advantage of this technique is that it can be used even in environments, such as embedded systems, where the documentation and the source code that are publicly released by the manufacturer are typically very limited or none at all. For instance, [Shoshitaishvili et al., 2015] proposes Firmalice, a binary analysis framework based on ANGR [Shoshitaishvili et al., 2016] that can be effectively used for identifying authentication bypass flaws inside firmwares running on devices such as routers and printers. Given a user-provided description of a privileged operation in the device, Firmalice identifies a set of program points that, if executed, forces the privileged operation to be performed. The program slice that involves the privileged program points is then symbolically analyzed using ANGR. If any such point can be reached by the engine, a set of concrete inputs is generated using an SMT solver. These values can be then used to effectively bypass authentication inside the device. On three commercially available devices, Firmalice could detect vulnerabilities in two of them, and determine that a backdoor in the third firmware is not remotely exploitable.

10 Conclusions

Techniques for symbolic execution have evolved significantly in the last decade, leading to major practical breakthroughs. In 2016, DARPA has challenged the global innovation community with a \$2M prize to build a computer that can hack and patch unknown software with no one at the keyboard. The winner, MAYHEM [Cha et al., 2012], was also the first autonomous computer system to play the Capture-The-Flag contest at the DEF CON 24 hacker convention⁵. The event demonstrated that tools for automatic exploit detection based on symbolic execution can be competitive with human experts, paving the road to unprecedented applications and the rise of start-ups that have the potential to shape software security and reliability in the next decades.

This survey has discussed some of the key aspects and challenges of symbolic execution, presenting them for a broad audience. To explain the basic design principles of symbolic executors and the main optimization techniques, we have focused on single-threaded applications with integer arithmetic. Symbolic execution of multi-threaded programs is treated, e.g., in [Khurshid et al., 2003, Sen, 2007, Bucur et al., 2011, Farzan et al., 2013, Bergan et al., 2014, Guo et al., 2015], while techniques for programs that manipulate floating point data are addressed in, e.g.,

⁵<https://www.defcon.org/html/defcon-24/dc-24-ctf.html>.

[Meudec, 2001, Botella et al., 2006, Lakhota et al., 2010, Collingbourne et al., 2011, Barr et al., 2013, Collingbourne et al., 2014, Ramachandran et al., 2015].

We hope that this survey will help non-experts grasp the key inventions in the exciting line of research of symbolic execution, inspiring further work and new ideas.

Acknowledgements. This work is supported in part by a grant of the Italian Presidency of the Council of Ministers and by the CINI (Consorzio Interuniversitario Nazionale Informatica) National Laboratory of Cyber Security.

Live Version of this Article. We complement the traditional scholarly publication model by maintaining a live version of this article at <https://github.com/season-lab/survey-symbolic-execution/>. The live version incorporates continuous feedback by the community, providing post-publication fixes, improvements, and extensions.

References

- [Anand et al., 2006] Anand, S., Păsăreanu, C. S., and Visser, W. (2006). Symbolic execution with abstract subsumption checking. In *Proceedings of the 13th International Conference on Model Checking Software*, SPIN 2006, pages 163–181, Berlin, Heidelberg. Springer-Verlag, ISBN: 3-540-33102-6, 978-3-540-33102-5, DOI: [10.1007/11691617_10](https://doi.org/10.1007/11691617_10), http://dx.doi.org/10.1007/11691617_10.
- [Avgerinos, 2014] Avgerinos, A. (2014). *Exploiting Trade-offs in Symbolic Execution for Identifying Security Bugs*. PhD thesis, <http://repository.cmu.edu/cgi/viewcontent.cgi?article=1478&context=dissertations>.
- [Avgerinos et al., 2011] Avgerinos, T., Cha, S. K., Hao, B. L. T., and Brumley, D. (2011). AEG: automatic exploit generation. In *Proceedings of the Network and Distributed System Security Symposium*, NDSS 2011. http://www.isoc.org/isoc/conferences/ndss/11/pdf/5_5.pdf.
- [Avgerinos et al., 2014] Avgerinos, T., Rebert, A., Cha, S. K., and Brumley, D. (2014). Enhancing symbolic execution with veritesting. In *Proceedings of the 36th International Conference on Software Engineering*, ICSE 2014, pages 1083–1094, New York, NY, USA. ACM, ISBN: 978-1-4503-2756-5, DOI: [10.1145/2568225.2568293](https://doi.org/10.1145/2568225.2568293), <http://doi.acm.org/10.1145/2568225.2568293>.
- [Babic and Hu, 2008] Babic, D. and Hu, A. J. (2008). Calysto: Scalable and precise extended static checking. In *Proceedings of the 30th International Conference on Software Engineering*, ICSE 2008, pages 211–220, New York, NY, USA. ACM, ISBN: 978-1-60558-079-1, DOI: [10.1145/1368088.1368118](https://doi.org/10.1145/1368088.1368118), <http://doi.acm.org/10.1145/1368088.1368118>.
- [Ball et al., 2006] Ball, T., Bounimova, E., Cook, B., Levin, V., Lichtenberg, J., McGarvey, C., Ondrusek, B., Rajamani, S. K., and Ustuner, A. (2006). Thorough static analysis of device drivers. In *Proceedings of the 1st ACM SIGOPS/EuroSys European Conference on Computer Systems 2006*, EuroSys 2006, pages 73–85, New York, NY, USA. ACM, ISBN: 1-59593-322-0, DOI: [10.1145/1217935.1217943](https://doi.org/10.1145/1217935.1217943), <http://doi.acm.org/10.1145/1217935.1217943>.
- [Ball and Daniel, 2015] Ball, T. and Daniel, J. (2015). Deconstructing dynamic symbolic execution. In *Proceedings of the 2014 Marktoberdorf Summer School on Dependable Software Systems Engineering*. IOS Press, <https://www.microsoft.com/en-us/research/publication/deconstructing-dynamic-symbolic-execution/>.
- [Barr et al., 2013] Barr, E. T., Vo, T., Le, V., and Su, Z. (2013). Automatic detection of floating-point exceptions. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL 2013, pages 549–560, New York, NY, USA. ACM, ISBN: 978-1-4503-1832-7, DOI: [10.1145/2429069.2429133](https://doi.org/10.1145/2429069.2429133), <http://doi.acm.org/10.1145/2429069.2429133>.

- [Barrett and Tinelli, 2007] Barrett, C. and Tinelli, C. (2007). Cvc3. In *Proceedings of the 19th International Conference on Computer Aided Verification, CAV 2007*, pages 298–302, Berlin, Heidelberg. Springer-Verlag, ISBN: 978-3-540-73367-6, DOI: [10.1007/978-3-540-73368-3_34](https://doi.org/10.1007/978-3-540-73368-3_34), <http://dl.acm.org/citation.cfm?id=1770351.1770397>.
- [Bergan et al., 2014] Bergan, T., Grossman, D., and Ceze, L. (2014). Symbolic execution of multithreaded programs from arbitrary program contexts. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2014*, pages 491–506, New York, NY, USA. ACM, ISBN: 978-1-4503-2585-1, DOI: [10.1145/2660193.2660200](https://doi.org/10.1145/2660193.2660200), <http://doi.acm.org/10.1145/2660193.2660200>.
- [Biere et al., 2009] Biere, A., Heule, M. J. H., van Maaren, H., and Walsh, T., editors (2009). *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, ISBN: 978-1-58603-929-5, ISSN: 0922-6389.
- [Boonstoppel et al., 2008] Boonstoppel, P., Cadar, C., and Engler, D. R. (2008). Rwsset: Attacking path explosion in constraint-based test generation. In *14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2008*, pages 351–366. DOI: [10.1007/978-3-540-78800-3_27](https://doi.org/10.1007/978-3-540-78800-3_27), http://dx.doi.org/10.1007/978-3-540-78800-3_27.
- [Botella et al., 2006] Botella, B., Gotlieb, A., and Michel, C. (2006). Symbolic execution of floating-point computations. *Software Testing, Verification & Reliability*, 16(2):97–121, ISSN: 0960-0833, DOI: [10.1002/stvr.v16:2](https://doi.org/10.1002/stvr.v16:2), <http://dx.doi.org/10.1002/stvr.v16:2>.
- [Boyer et al., 1975] Boyer, R. S., Elspas, B., and Levitt, K. N. (1975). Select: a formal system for testing and debugging programs by symbolic execution. In *Proceedings of the International Conference on Reliable Software*, pages 234–245, New York, NY, USA. ACM, DOI: [10.1145/800027.808445](https://doi.org/10.1145/800027.808445), <http://doi.acm.org/10.1145/800027.808445>.
- [Brumley et al., 2008] Brumley, D., Hartwig, C., Liang, Z., Newsome, J., Song, D., and Yin, H. (2008). *Botnet Detection: Countering the Largest Security Threat*, chapter Automatically Identifying Trigger-based Behavior in Malware, pages 65–88. Springer US, DOI: [10.1007/978-0-387-68768-1](https://doi.org/10.1007/978-0-387-68768-1).
- [Brumley et al., 2011] Brumley, D., Jager, I., Avgerinos, T., and Schwartz, E. J. (2011). BAP: A binary analysis platform. In *Proceedings of the 23rd International Conference on Computer Aided Verification, CAV 2011*, pages 463–469, Berlin, Heidelberg. Springer-Verlag, ISBN: 978-3-642-22109-5, DOI: [10.1007/978-3-642-22110-1_37](https://doi.org/10.1007/978-3-642-22110-1_37), <http://dl.acm.org/citation.cfm?id=2032305.2032342>.
- [Bucur et al., 2011] Bucur, S., Ureche, V., Zamfir, C., and Candea, G. (2011). Parallel symbolic execution for automated real-world software testing. In *Proceedings of the Sixth Conference on Computer Systems, EuroSys 2011*, pages 183–198. ISBN: 978-1-4503-0634-8, DOI: [10.1145/1966445.1966463](https://doi.org/10.1145/1966445.1966463), <http://doi.acm.org/10.1145/1966445.1966463>.
- [Burnim and Sen, 2008] Burnim, J. and Sen, K. (2008). Heuristics for scalable dynamic test generation. In *Proceedings of the 2008 23rd IEEE/ACM International Conference on Automated Software Engineering, ASE 2008*, pages 443–446, Washington, DC, USA. IEEE Computer Society, ISBN: 978-1-4244-2187-9, DOI: [10.1109/ASE.2008.69](https://doi.org/10.1109/ASE.2008.69), <http://dx.doi.org/10.1109/ASE.2008.69>.
- [Cadar et al., 2008] Cadar, C., Dunbar, D., and Engler, D. (2008). KLEE: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation, OSDI 2008*, pages 209–224, Berkeley, CA, USA. USENIX Association, <http://dl.acm.org/citation.cfm?id=1855741.1855756>.

- [Cadar et al., 2006] Cadar, C., Ganesh, V., Pawlowski, P. M., Dill, D. L., and Engler, D. R. (2006). EXE: Automatically generating inputs of death. In *Proceedings of the 13th ACM Conference on Computer and Communications Security*, CCS 2006, pages 322–335, New York, NY, USA. ACM, ISBN: 1-59593-518-5, DOI: [10.1145/1180405.1180445](https://doi.org/10.1145/1180405.1180445), <http://doi.acm.org/10.1145/1180405.1180445>.
- [Cadar et al., 2011] Cadar, C., Godefroid, P., Khurshid, S., Păsăreanu, C. S., Sen, K., Tillmann, N., and Visser, W. (2011). Symbolic execution for software testing in practice: Preliminary assessment. In *Proceedings of the 33rd International Conference on Software Engineering*, ICSE 2011, pages 1066–1071, New York, NY, USA. ACM, ISBN: 978-1-4503-0445-0, DOI: [10.1145/1985793.1985995](https://doi.org/10.1145/1985793.1985995), <http://doi.acm.org/10.1145/1985793.1985995>.
- [Cadar and Sen, 2013] Cadar, C. and Sen, K. (2013). Symbolic execution for software testing: Three decades later. *Communications of the ACM*, 56(2):82–90, ISSN: 0001-0782, DOI: [10.1145/2408776.2408795](https://doi.org/10.1145/2408776.2408795), <http://doi.acm.org/10.1145/2408776.2408795>.
- [Caselden et al., 2013] Caselden, D., Bazhanyuk, A., Payer, M., McCamant, S., and Song, D. (2013). HI-CFG: construction by binary analysis and application to attack polymorphism. In *18th European Symposium on Research in Computer Security*, ESORICS 2013, pages 164–181. DOI: [10.1007/978-3-642-40203-6_10](https://doi.org/10.1007/978-3-642-40203-6_10), http://dx.doi.org/10.1007/978-3-642-40203-6_10.
- [Cha et al., 2012] Cha, S. K., Avgerinos, T., Rebert, A., and Brumley, D. (2012). Unleashing mayhem on binary code. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, SP 2012, pages 380–394, Washington, DC, USA. IEEE Computer Society, ISBN: 978-0-7695-4681-0, DOI: [10.1109/SP.2012.31](https://doi.org/10.1109/SP.2012.31), <http://dx.doi.org/10.1109/SP.2012.31>.
- [Chaudhuri and Foster, 2010] Chaudhuri, A. and Foster, J. S. (2010). Symbolic security analysis of ruby-on-rails web applications. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, CCS 2010, pages 585–594, New York, NY, USA. ACM, ISBN: 978-1-4503-0245-6, DOI: [10.1145/1866307.1866373](https://doi.org/10.1145/1866307.1866373), <http://doi.acm.org/10.1145/1866307.1866373>.
- [Chipounov et al., 2012] Chipounov, V., Kuznetsov, V., and Candea, G. (2012). The S2E platform: Design, implementation, and applications. *ACM Transactions on Computer Systems*, 30(1):2, DOI: [10.1145/2110356.2110358](https://doi.org/10.1145/2110356.2110358), <http://doi.acm.org/10.1145/2110356.2110358>.
- [Collingbourne et al., 2011] Collingbourne, P., Cadar, C., and Kelly, P. H. (2011). Symbolic cross-checking of floating-point and simd code. In *Proceedings of the Sixth Conference on Computer Systems*, EuroSys 2011, pages 315–328, New York, NY, USA. ACM, ISBN: 978-1-4503-0634-8, DOI: [10.1145/1966445.1966475](https://doi.org/10.1145/1966445.1966475), <http://doi.acm.org/10.1145/1966445.1966475>.
- [Collingbourne et al., 2014] Collingbourne, P., Cadar, C., and Kelly, P. H. J. (2014). Symbolic crosschecking of data-parallel floating-point code. *IEEE Transactions on Software Engineering*, 40(7):710–737, DOI: [10.1109/TSE.2013.2297120](https://doi.org/10.1109/TSE.2013.2297120), <http://dx.doi.org/10.1109/TSE.2013.2297120>.
- [Costin et al., 2014] Costin, A., Zaddach, J., Francillon, A., and Balzarotti, D. (2014). A large-scale analysis of the security of embedded firmwares. In *Proceedings of the 23rd USENIX Security Symposium*, pages 95–110. <https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/costin>.
- [Csallner and Smaragdakis, 2005] Csallner, C. and Smaragdakis, Y. (2005). Check 'n' crash: Combining static checking and testing. In *Proceedings of the 27th International Conference on Software Engineering*, ICSE 2005, pages 422–431, New York, NY, USA. ACM, ISBN: 1-58113-963-2, DOI: [10.1145/1062455.1062533](https://doi.org/10.1145/1062455.1062533), <http://doi.acm.org/10.1145/1062455.1062533>.

- [Cytron et al., 1991] Cytron, R., Ferrante, J., Rosen, B. K., Wegman, M. N., and Zadeck, F. K. (1991). Efficiently computing static single assignment form and the control dependence graph. *ACM Transactions on Programming Languages and Systems*, 13(4):451–490, ISSN: 0164-0925, DOI: [10.1145/115372.115320](https://doi.org/10.1145/115372.115320), <http://doi.acm.org/10.1145/115372.115320>.
- [Davidson et al., 2013] Davidson, D., Moench, B., Jha, S., and Ristenpart, T. (2013). FIE on firmware: Finding vulnerabilities in embedded systems using symbolic execution. In *Proceedings of the 22nd USENIX Conference on Security, SEC 2013*, pages 463–478, Berkeley, CA, USA. USENIX Association, ISBN: 978-1-931971-03-4, <http://dl.acm.org/citation.cfm?id=2534766>.
- [De Moura and Bjørner, 2008] De Moura, L. and Bjørner, N. (2008). Z3: An efficient smt solver. In *Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS'08/ETAPS'08*, pages 337–340, Berlin, Heidelberg. Springer-Verlag, ISBN: 3-540-78799-2, 978-3-540-78799-0, DOI: [10.1007/978-3-540-78800-3_24](https://doi.org/10.1007/978-3-540-78800-3_24), <http://dl.acm.org/citation.cfm?id=1792734>.
- [De Moura and Bjørner, 2011] De Moura, L. and Bjørner, N. (2011). Satisfiability modulo theories: Introduction and applications. *Communications of the ACM*, 54(9):69–77, ISSN: 0001-0782, DOI: [10.1145/1995376.1995394](https://doi.org/10.1145/1995376.1995394), <http://doi.acm.org/10.1145/1995376.1995394>.
- [Deng et al., 2012] Deng, X., Lee, J., and Robby (2012). Efficient and formal generalized symbolic execution. *Automated Software Engineering*, 19(3):233–301, ISSN: 0928-8910, DOI: [10.1007/s10515-011-0089-9](https://doi.org/10.1007/s10515-011-0089-9), <http://dx.doi.org/10.1007/s10515-011-0089-9>.
- [do Val, 2014] do Val, C. G. (2014). Conflict-driven symbolic execution: How to learn to get better. MSc thesis, University of British Columbia.
- [Duesterwald, 2004] Duesterwald, E., editor (2004). *Analyzing Memory Accesses in x86 Executables*, CC 2004, Berlin, Heidelberg. Springer Berlin Heidelberg, ISBN: 978-3-540-24723-4, DOI: [10.1007/978-3-540-24723-4_2](https://doi.org/10.1007/978-3-540-24723-4_2), http://dx.doi.org/10.1007/978-3-540-24723-4_2.
- [Elkarablieh et al., 2009] Elkarablieh, B., Godefroid, P., and Levin, M. Y. (2009). Precise pointer reasoning for dynamic test generation. In *Proceedings of the Eighteenth International Symposium on Software Testing and Analysis, ISSTA 2009*, pages 129–140, New York, NY, USA. ACM, ISBN: 978-1-60558-338-9, DOI: [10.1145/1572272.1572288](https://doi.org/10.1145/1572272.1572288), <http://doi.acm.org/10.1145/1572272.1572288>.
- [Engler and Dunbar, 2007] Engler, D. and Dunbar, D. (2007). Under-constrained execution: Making automatic code destruction easy and scalable. In *Proceedings of the 2007 International Symposium on Software Testing and Analysis, ISSTA 2007*, pages 1–4, New York, NY, USA. ACM, ISBN: 978-1-59593-734-6, DOI: [10.1145/1273463.1273464](https://doi.org/10.1145/1273463.1273464).
- [Farzan et al., 2013] Farzan, A., Holzer, A., Razavi, N., and Veith, H. (2013). Con2colic testing. In *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2013*, pages 37–47, New York, NY, USA. ACM, ISBN: 978-1-4503-2237-9, DOI: [10.1145/2491411.2491453](https://doi.org/10.1145/2491411.2491453), <http://doi.acm.org/10.1145/2491411.2491453>.
- [Fu et al., 2007] Fu, X., Lu, X., Peltsverger, B., Chen, S., Qian, K., and Tao, L. (2007). A static analysis framework for detecting sql injection vulnerabilities. In *Proceedings of the 31st Annual International Computer Software and Applications Conference, COMPSAC 2007*, pages 87–96, Washington, DC, USA. IEEE Computer Society, ISBN: 0-7695-2870-8, DOI: [10.1109/COMPSAC.2007.43](https://doi.org/10.1109/COMPSAC.2007.43), <http://dx.doi.org/10.1109/COMPSAC.2007.43>.

- [Ganesh, 2007] Ganesh, V. (2007). *Decision Procedures for Bit-vectors, Arrays and Integers*. PhD thesis, https://ece.uwaterloo.ca/~vganesh/Publications_files/vg2007-PhD-STANFORD.pdf.
- [Ganesh and Dill, 2007] Ganesh, V. and Dill, D. L. (2007). A decision procedure for bit-vectors and arrays. In *Proceedings of the 19th International Conference on Computer Aided Verification, CAV 2007*, pages 519–531, Berlin, Heidelberg. Springer-Verlag, ISBN: 978-3-540-73367-6, DOI: http://dx.doi.org/10.1007/978-3-540-73368-3_52, <http://dl.acm.org/citation.cfm?id=1770351.1770421>.
- [Geldenhuys et al., 2013] Geldenhuys, J., Aguirre, N., Frias, M. F., and Visser, W. (2013). Bounded lazy initialization. In Brat, G., Rungta, N., and Venet, A., editors, *5th International NASA Formal Methods Symposium, NFM 2013*, pages 229–243, Berlin, Heidelberg. Springer Berlin Heidelberg, ISBN: 978-3-642-38088-4, DOI: [10.1007/978-3-642-38088-4_16](http://dx.doi.org/10.1007/978-3-642-38088-4_16), http://dx.doi.org/10.1007/978-3-642-38088-4_16.
- [Godefroid, 2007] Godefroid, P. (2007). Compositional dynamic test generation. In *Proceedings of the 34th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2007*, pages 47–54, New York, NY, USA. ACM, ISBN: 1-59593-575-4, DOI: [10.1145/1190216.1190226](http://doi.acm.org/10.1145/1190216.1190226), <http://doi.acm.org/10.1145/1190216.1190226>.
- [Godefroid et al., 2005] Godefroid, P., Klarlund, N., and Sen, K. (2005). DART: Directed automated random testing. In *Proceedings of the 2005 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2005*, pages 213–223, New York, NY, USA. ACM, ISBN: 1-59593-056-6, DOI: [10.1145/1065010.1065036](http://doi.acm.org/10.1145/1065010.1065036), <http://doi.acm.org/10.1145/1065010.1065036>.
- [Godefroid et al., 2012] Godefroid, P., Levin, M. Y., and Molnar, D. (2012). Sage: White-box fuzzing for security testing. *Queue*, 10(1):20:20–20:27, ISSN: 1542-7730, DOI: [10.1145/2090147.2094081](http://doi.acm.org/10.1145/2090147.2094081), <http://doi.acm.org/10.1145/2090147.2094081>.
- [Godefroid et al., 2008] Godefroid, P., Levin, M. Y., and Molnar, D. A. (2008). Automated whitebox fuzz testing. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2008*. http://www.isoc.org/isoc/conferences/ndss/08/papers/10_automated_whitebox_fuzz.pdf.
- [Godefroid and Luchau, 2011] Godefroid, P. and Luchau, D. (2011). Automatic partial loop summarization in dynamic test generation. In *Proceedings of the 2011 International Symposium on Software Testing and Analysis, ISSTA 2011*, pages 23–33, New York, NY, USA. ACM, ISBN: 978-1-4503-0562-4, DOI: [10.1145/2001420.2001424](http://doi.acm.org/10.1145/2001420.2001424), <http://doi.acm.org/10.1145/2001420.2001424>.
- [Groce and Visser, 2002] Groce, A. and Visser, W. (2002). Model checking java programs using structural heuristics. In *Proceedings of the 2002 ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2002*, pages 12–21, New York, NY, USA. ACM, ISBN: 1-58113-562-9, DOI: [10.1145/566172.566175](http://doi.acm.org/10.1145/566172.566175), <http://doi.acm.org/10.1145/566172.566175>.
- [Guo et al., 2015] Guo, S., Kusano, M., Wang, C., Yang, Z., and Gupta, A. (2015). Assertion guided symbolic execution of multithreaded programs. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015*, pages 854–865, New York, NY, USA. ACM, ISBN: 978-1-4503-3675-8, DOI: [10.1145/2786805.2786841](http://doi.acm.org/10.1145/2786805.2786841), <http://doi.acm.org/10.1145/2786805.2786841>.
- [Hai et al., 2016] Hai, N. M., Ogawa, M., and Tho, Q. T. (2016). *Obfuscation Code Localization Based on CFG Generation of Malware*, pages 229–247. Springer International Publishing, Cham, ISBN: 978-3-319-30303-1, DOI: [10.1007/978-3-319-30303-1_14](http://dx.doi.org/10.1007/978-3-319-30303-1_14), http://dx.doi.org/10.1007/978-3-319-30303-1_14.

- [Hansen et al., 2009] Hansen, T., Schachte, P., and Søndergaard, H. (2009). Runtime verification. chapter State Joining and Splitting for the Symbolic Execution of Binaries, pages 76–92. Springer-Verlag, Berlin, Heidelberg, ISBN: 978-3-642-04693-3, DOI: [10.1007/978-3-642-04694-0_6](https://doi.org/10.1007/978-3-642-04694-0_6), http://dx.doi.org/10.1007/978-3-642-04694-0_6.
- [Hentschel et al., 2014] Hentschel, M., Bubel, R., and Hähnle, R. (2014). Symbolic execution debugger (sed). In Bonakdarpour, B. and Smolka, S. A., editors, *Proceedings of Runtime Verification 2014*, RV 2014, pages 255–262. Springer, DOI: [10.1007/978-3-319-11164-3_21](https://doi.org/10.1007/978-3-319-11164-3_21), http://dx.doi.org/10.1007/978-3-319-11164-3_21.
- [Howden, 1977] Howden, W. E. (1977). Symbolic testing and the dissect symbolic evaluation system. *IEEE Transactions on Software Engineering*, 3(4):266–278, ISSN: 0098-5589, DOI: [10.1109/TSE.1977.231144](https://doi.org/10.1109/TSE.1977.231144), <http://dx.doi.org/10.1109/TSE.1977.231144>.
- [Jeon et al., 2012] Jeon, J., Micinski, K. K., and Foster, J. S. (2012). SymDroid: Symbolic Execution for Dalvik Bytecode. Technical Report CS-TR-5022, Department of Computer Science, University of Maryland, College Park, <http://www.cs.umd.edu/~jfoster/papers/cs-tr-5022.pdf>.
- [Jia et al., 2015] Jia, X., Ghezzi, C., and Ying, S. (2015). Enhancing reuse of constraint solutions to improve symbolic execution. In *Proceedings of the 2015 International Symposium on Software Testing and Analysis*, ISSTA 2015, pages 177–187, New York, NY, USA. ACM, ISBN: 978-1-4503-3620-8, DOI: [10.1145/2771783.2771806](https://doi.org/10.1145/2771783.2771806), <http://doi.acm.org/10.1145/2771783.2771806>.
- [Karger and Schell, 1974] Karger, P. A. and Schell, R. R. (1974). Multics security evaluation: Vulnerability analysis. Technical report, HQ Electronic Systems Division: Hanscom AFB, MA, <http://csrc.nist.gov/publications/history/karg74.pdf>.
- [Khoo et al., 2010] Khoo, Y. P., Chang, B.-Y. E., and Foster, J. S. (2010). Mixing type checking and symbolic execution. In *Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2010, pages 436–447, New York, NY, USA. ACM, ISBN: 978-1-4503-0019-3, DOI: [10.1145/1806596.1806645](https://doi.org/10.1145/1806596.1806645), <http://doi.acm.org/10.1145/1806596.1806645>.
- [Khurshid et al., 2003] Khurshid, S., Păsăreanu, C. S., and Visser, W. (2003). Generalized Symbolic Execution for Model Checking and Testing. In *Proceedings of the 9th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, TACAS 2003, pages 553–568, Berlin, Heidelberg. Springer-Verlag, ISBN: 3-540-00898-5, DOI: [10.1007/3-540-36577-x_40](https://doi.org/10.1007/3-540-36577-x_40), <http://dl.acm.org/citation.cfm?id=1765871.1765924>.
- [King, 1975] King, J. C. (1975). A new approach to program testing. In *Proceedings of the International Conference on Reliable Software*, pages 228–233, New York, NY, USA. ACM, DOI: [10.1145/800027.808444](https://doi.org/10.1145/800027.808444), <http://doi.acm.org/10.1145/800027.808444>.
- [King, 1976] King, J. C. (1976). Symbolic execution and program testing. *Communication of ACM*, 19(7):385–394, ISSN: 0001-0782, DOI: [10.1145/360248.360252](https://doi.org/10.1145/360248.360252), <http://doi.acm.org/10.1145/360248.360252>.
- [Kuznetsov et al., 2012] Kuznetsov, V., Kinder, J., Bucur, S., and Candea, G. (2012). Efficient state merging in symbolic execution. In *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2012, pages 193–204, New York, NY, USA. ACM, ISBN: 978-1-4503-1205-9, DOI: [10.1145/2254064.2254088](https://doi.org/10.1145/2254064.2254088).
- [Lakhota et al., 2010] Lakhota, K., Tillmann, N., Harman, M., and De Halleux, J. (2010). Flopsy: Search-based floating point constraint solving for symbolic execution. In *Proceedings of the 22nd IFIP WG 6.1 International Conference on Testing Software and Systems*, ICTSS 2010, pages 142–157, Berlin, Heidelberg. Springer-Verlag, ISBN: 3-642-16572-9,

- 978-3-642-16572-6, DOI: [10.1007/978-3-642-16573-3_11](https://doi.org/10.1007/978-3-642-16573-3_11), <http://dl.acm.org/citation.cfm?id=1928028.1928039>.
- [Lattner and Adve, 2004] Lattner, C. and Adve, V. (2004). LLVM: A compilation framework for lifelong program analysis & transformation. In *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization*, CGO 2004, pages 75–86, Washington, DC, USA. IEEE Computer Society, ISBN: 0-7695-2102-9, DOI: [10.1109/cgo.2004.1281665](https://doi.org/10.1109/cgo.2004.1281665), <http://dl.acm.org/citation.cfm?id=977395.977673>.
- [Li et al., 2014] Li, G., Andreasen, E., and Ghosh, I. (2014). Symjs: Automatic symbolic testing of javascript web applications. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE 2014, pages 449–459, New York, NY, USA. ACM, ISBN: 978-1-4503-3056-5, DOI: [10.1145/2635868.2635913](https://doi.org/10.1145/2635868.2635913), <http://doi.acm.org/10.1145/2635868.2635913>.
- [Li et al., 2013] Li, Y., Su, Z., Wang, L., and Li, X. (2013). Steering symbolic execution to less traveled paths. In *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications*, OOPSLA 2013, pages 19–32. ISBN: 978-1-4503-2374-1, DOI: [10.1145/2509136.2509553](https://doi.org/10.1145/2509136.2509553), <http://doi.acm.org/10.1145/2509136.2509553>.
- [Luckow et al., 2016] Luckow, K., Dimjašević, M., Giannakopoulou, D., Howar, F., Isberner, M., Kahsai, T., Rakamarić, Z., and Raman, V. (2016). JDart: A dynamic symbolic analysis framework. In *Proceedings of the 22nd International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, TACAS 2016, pages 442–459, New York, NY, USA. Springer-Verlag New York, Inc., ISBN: 978-3-662-49673-2, DOI: [10.1007/978-3-662-49674-9_26](https://doi.org/10.1007/978-3-662-49674-9_26), http://dx.doi.org/10.1007/978-3-662-49674-9_26.
- [Majumdar and Sen, 2007] Majumdar, R. and Sen, K. (2007). Hybrid concolic testing. In *Proceedings of the 29th International Conference on Software Engineering*, ICSE 2007, pages 416–426, Washington, DC, USA. IEEE Computer Society, ISBN: 0-7695-2828-7, DOI: [10.1109/ICSE.2007.41](https://doi.org/10.1109/ICSE.2007.41).
- [Martignoni et al., 2012] Martignoni, L., McCamant, S., Poosankam, P., Song, D., and Marniatis, P. (2012). Path-exploration lifting: Hi-fi tests for lo-fi emulators. In *Proceedings of the Seventeenth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS XVII, pages 337–348, New York, NY, USA. ACM, ISBN: 978-1-4503-0759-8, DOI: [10.1145/2150976.2151012](https://doi.org/10.1145/2150976.2151012), <http://doi.acm.org/10.1145/2150976.2151012>.
- [McMinn, 2004] McMinn, P. (2004). Search-based software test data generation: A survey. *Software Testing, Verification & Reliability*, 14(2):105–156, ISSN: 0960-0833, DOI: [10.1002/stvr.v14:2](https://doi.org/10.1002/stvr.v14:2), <http://dx.doi.org/10.1002/stvr.v14:2>.
- [Meudec, 2001] Meudec, C. (2001). Atgen: automatic test data generation using constraint logic programming and symbolic execution. *Software Testing Verification and Reliability*, 11(2):81–96, DOI: [10.1002/stvr.225](https://doi.org/10.1002/stvr.225).
- [Nethercote and Seward, 2007] Nethercote, N. and Seward, J. (2007). Valgrind: A framework for heavyweight dynamic binary instrumentation. In *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2007, pages 89–100, New York, NY, USA. ACM, ISBN: 978-1-59593-633-2, DOI: [10.1145/1250734.1250746](https://doi.org/10.1145/1250734.1250746), <http://doi.acm.org/10.1145/1250734.1250746>.
- [Pincus and Baker, 2004] Pincus, J. and Baker, B. (2004). Beyond stack smashing: recent advances in exploiting buffer overruns. *IEEE Security Privacy*, 2(4):20–27, ISSN: 1540-7993, DOI: [10.1109/MSP.2004.36](https://doi.org/10.1109/MSP.2004.36).

- [Prud'homme et al., 2015] Prud'homme, C., Fages, J.-G., and Lorca, X. (2015). *Choco Documentation*. TASC, INRIA Rennes, LINA CNRS UMR 6241, COSLING S.A.S., <http://www.choco-solver.org>.
- [Păsăreanu and Rungta, 2010] Păsăreanu, C. S. and Rungta, N. (2010). Symbolic pathfinder: Symbolic execution of java bytecode. In *Proceedings of the IEEE/ACM International Conference on Automated Software Engineering, ASE 2010*, pages 179–180, New York, NY, USA. ACM, ISBN: 978-1-4503-0116-9, DOI: [10.1145/1858996.1859035](https://doi.org/10.1145/1858996.1859035), <http://doi.acm.org/10.1145/1858996.1859035>.
- [Păsăreanu et al., 2011] Păsăreanu, C. S., Rungta, N., and Visser, W. (2011). Symbolic execution with mixed concrete-symbolic solving. In *Proceedings of the 2011 International Symposium on Software Testing and Analysis, ISSTA 2011*, pages 34–44, New York, NY, USA. ACM, ISBN: 978-1-4503-0562-4, DOI: [10.1145/2001420.2001425](https://doi.org/10.1145/2001420.2001425), <http://doi.acm.org/10.1145/2001420.2001425>.
- [Păsăreanu and Visser, 2004] Păsăreanu, C. S. and Visser, W. (2004). Verification of java programs using symbolic execution and invariant generation. In *Model Checking Software: 11th International SPIN Workshop, SPIN 2004*, pages 164–181. Springer Berlin Heidelberg, DOI: [10.1007/978-3-540-24732-6_13](https://doi.org/10.1007/978-3-540-24732-6_13), http://dx.doi.org/10.1007/978-3-540-24732-6_13.
- [Ramachandran et al., 2015] Ramachandran, J., Păsăreanu, C., and Wahl, T. (2015). Symbolic execution for checking the accuracy of floating-point programs. *ACM SIGSOFT Software Engineering Notes*, 40(1):1–5, ISSN: 0163-5948, DOI: [10.1145/2693208.2693242](https://doi.org/10.1145/2693208.2693242), <http://doi.acm.org/10.1145/2693208.2693242>.
- [Ramos and Engler, 2015] Ramos, D. A. and Engler, D. (2015). Under-constrained symbolic execution: Correctness checking for real code. In *Proceedings of the 24th USENIX Conference on Security Symposium, SEC 2015*, pages 49–64, Berkeley, CA, USA. USENIX Association, ISBN: 978-1-931971-232, <http://dl.acm.org/citation.cfm?id=2831143.2831147>.
- [Reisner et al., 2010] Reisner, E., Song, C., Ma, K.-K., Foster, J. S., and Porter, A. (2010). Using symbolic evaluation to understand behavior in configurable software systems. In *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering, ICSE 2010*, pages 445–454, New York, NY, USA. ACM, ISBN: 978-1-60558-719-6, DOI: [10.1145/1806799.1806864](https://doi.org/10.1145/1806799.1806864), <http://doi.acm.org/10.1145/1806799.1806864>.
- [Rosner et al., 2015] Rosner, N., Geldenhuys, J., Aguirre, N. M., Visser, W., and Frias, M. F. (2015). Bliss: Improved symbolic execution by bounded lazy initialization with sat support. *IEEE Transactions on Software Engineering*, 41(7):639–660, ISSN: 0098-5589, DOI: [10.1109/TSE.2015.2389225](https://doi.org/10.1109/TSE.2015.2389225).
- [Saudel and Salwan, 2015] Saudel, F. and Salwan, J. (2015). Triton: A dynamic symbolic execution framework. In *Symposium sur la sécurité des technologies de l'information et des communications, SSTIC, France, Rennes, June 3-5 2015*, pages 31–54. SSTIC, http://triton.quarkslab.com/files/sstic2015_wp_fr_saudel_salwan.pdf.
- [Saxena et al., 2009] Saxena, P., Poosankam, P., McCamant, S., and Song, D. (2009). Loop-extended symbolic execution on binary programs. In *Proceedings of the Eighteenth International Symposium on Software Testing and Analysis, ISSTA 2009*, pages 225–236, New York, NY, USA. ACM, ISBN: 978-1-60558-338-9, DOI: [10.1145/1572272.1572299](https://doi.org/10.1145/1572272.1572299).
- [Schwartz et al., 2010] Schwartz, E. J., Avgerinos, T., and Brumley, D. (2010). All you ever wanted to know about dynamic taint analysis and forward symbolic execution (but might have been afraid to ask). In *Proceedings of the 2010 IEEE Symposium on Security and Privacy, SP 2010*, pages 317–331, Washington, DC, USA. IEEE Computer Society, ISBN: 978-0-7695-4035-1, DOI: [10.1109/SP.2010.26](https://doi.org/10.1109/SP.2010.26), <http://dx.doi.org/10.1109/SP.2010.26>.

- [Schwartz et al., 2011] Schwartz, E. J., Avgerinos, T., and Brumley, D. (2011). Q: Exploit hardening made easy. In *Proceedings of the 20th USENIX Conference on Security, SEC 2011*, pages 25–25, Berkeley, CA, USA. USENIX Association, <http://dl.acm.org/citation.cfm?id=2028067.2028092>.
- [Sen, 2007] Sen, Koushik and Agha, G. (2007). A race-detection and flipping algorithm for automated testing of multi-threaded programs. In *Hardware and Software, Verification and Testing: Second International Haifa Verification Conference, HVC 2006*, pages 166–182, Berlin, Heidelberg. Springer Berlin Heidelberg, ISBN: 978-3-540-70889-6, DOI: [10.1007/978-3-540-70889-6_13](https://doi.org/10.1007/978-3-540-70889-6_13), http://dx.doi.org/10.1007/978-3-540-70889-6_13.
- [Sen and Agha, 2006] Sen, K. and Agha, G. (2006). CUTE and jCUTE: Concolic unit testing and explicit path model-checking tools. In *Proceedings of the 18th International Conference on Computer Aided Verification, CAV 2006*, pages 419–423, Berlin, Heidelberg. Springer-Verlag, ISBN: 3-540-37406-X, 978-3-540-37406-0, DOI: [10.1007/11817963_38](https://doi.org/10.1007/11817963_38), http://dx.doi.org/10.1007/11817963_38.
- [Sen et al., 2013] Sen, K., Kalasapur, S., Brutch, T., and Gibbs, S. (2013). Jalangi: A selective record-replay and dynamic analysis framework for javascript. In *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2013*, pages 488–498, New York, NY, USA. ACM, ISBN: 978-1-4503-2237-9, DOI: [10.1145/2491411.2491447](https://doi.org/10.1145/2491411.2491447), <http://doi.acm.org/10.1145/2491411.2491447>.
- [Sen et al., 2005] Sen, K., Marinov, D., and Agha, G. (2005). CUTE: A concolic unit testing engine for c. In *Proceedings of the 10th European Software Engineering Conference Held Jointly with 13th ACM SIGSOFT International Symposium on Foundations of Software Engineering, ESEC/FSE-13*, pages 263–272, New York, NY, USA. ACM, ISBN: 1-59593-014-0, DOI: [10.1145/1081706.1081750](https://doi.org/10.1145/1081706.1081750), <http://doi.acm.org/10.1145/1081706.1081750>.
- [Shannon et al., 2007] Shannon, D., Hajra, S., Lee, A., Zhan, D., and Khurshid, S. (2007). Abstracting symbolic execution with string analysis. In *Proceedings of the Testing: Academic and Industrial Conference Practice and Research Techniques - MUTATION, TAICPART-MUTATION 2007*, pages 13–22, Washington, DC, USA. IEEE Computer Society, ISBN: 0-7695-2984-4, DOI: [10.1109/taicpart.2007.4344094](https://doi.org/10.1109/taicpart.2007.4344094), <http://dl.acm.org/citation.cfm?id=1308173.1308254>.
- [Sharif et al., 2008] Sharif, M. I., Lanzi, A., Giffin, J. T., and Lee, W. (2008). Impeding malware analysis using conditional code obfuscation. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2008*. http://www.isoc.org/isoc/conferences/ndss/08/papers/19_impeding_malware_analysis.pdf.
- [Sharma, 2014] Sharma, A. (2014). Exploiting undefined behaviors for efficient symbolic execution. In *Companion Proceedings of the 36th International Conference on Software Engineering, ICSE Companion 2014*, pages 727–729, New York, NY, USA. ACM, ISBN: 978-1-4503-2768-8, DOI: [10.1145/2591062.2594450](https://doi.org/10.1145/2591062.2594450), <http://doi.acm.org/10.1145/2591062.2594450>.
- [Shoshitaishvili et al., 2015] Shoshitaishvili, Y., Wang, R., Hauser, C., Kruegel, C., and Vigna, G. (2015). Firmalice - automatic detection of authentication bypass vulnerabilities in binary firmware. In *22nd Annual Network and Distributed System Security Symposium, NDSS 2015*. DOI: [10.14722/ndss.2015.23294](https://doi.org/10.14722/ndss.2015.23294), <http://www.internetsociety.org/doc/firmalice-automatic-detection-authentication-bypass-vulnerabilities-binary-firmware>.
- [Shoshitaishvili et al., 2016] Shoshitaishvili, Y., Wang, R., Salls, C., Stephens, N., Polino, M., Dutcher, A., Grosen, J., Feng, S., Hauser, C., Krügel, C., and Vigna, G. (2016). SOK: (state of) the art of war: Offensive techniques in binary analysis. In *IEEE Symposium on Security and Privacy, SP 2016*, pages 138–157. DOI: [10.1109/SP.2016.17](https://doi.org/10.1109/SP.2016.17), <http://dx.doi.org/10.1109/SP.2016.17>.

- [Siegel et al., 2015] Siegel, S. F., Zheng, M., Luo, Z., Zirkel, T. K., Marianiello, A. V., Edenhofner, J. G., Dwyer, M. B., and Rogers, M. S. (2015). *Civl: The concurrency intermediate verification language*. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2015*, pages 61:1–61:12, New York, NY, USA. ACM, ISBN: 978-1-4503-3723-6, DOI: [10.1145/2807591.2807635](https://doi.org/10.1145/2807591.2807635), <http://doi.acm.org/10.1145/2807591.2807635>.
- [Slaby et al., 2013] Slaby, J., Strejcek, J., and Trtik, M. (2013). Compact symbolic execution. In *11th International Symposium on Automated Technology for Verification and Analysis, ATVA 2013*, pages 193–207. DOI: [10.1007/978-3-319-02444-8_15](https://doi.org/10.1007/978-3-319-02444-8_15), http://dx.doi.org/10.1007/978-3-319-02444-8_15.
- [Song et al., 2008] Song, D., Brumley, D., Yin, H., Caballero, J., Jager, I., Kang, M. G., Liang, Z., Newsome, J., Poosankam, P., and Saxena, P. (2008). Bitblaze: A new approach to computer security via binary analysis. In *Proceedings of the 4th International Conference on Information Systems Security, ICISS 2008*, pages 1–25, Berlin, Heidelberg. Springer-Verlag, ISBN: 978-3-540-89861-0, DOI: [10.1007/978-3-540-89862-7_1](https://doi.org/10.1007/978-3-540-89862-7_1), http://dx.doi.org/10.1007/978-3-540-89862-7_1.
- [Souza et al., 2011] Souza, M., Borges, M., d’Amorim, M., and Păsăreanu, C. S. (2011). Coral: Solving complex constraints for symbolic pathfinder. In *Proceedings of the Third International NASA Formal Methods Symposium, NFM 2011*, pages 359–374, Berlin, Heidelberg. Springer-Verlag, ISBN: 978-3-642-20397-8, DOI: [10.1007/978-3-642-20398-5_26](https://doi.org/10.1007/978-3-642-20398-5_26), <http://dl.acm.org/citation.cfm?id=1986308.1986337>.
- [Stephens et al., 2016] Stephens, N., Grosen, J., Salls, C., Dutcher, A., Wang, R., Corbetta, J., Shoshitaishvili, Y., Kruegel, C., and Vigna, G. (2016). Driller: Augmenting fuzzing through selective symbolic execution. In *23rd Annual Network and Distributed System Security Symposium, NDSS 2016*. <http://www.internetsociety.org/sites/default/files/blogs-media/driller-augmenting-fuzzing-through-selective-symbolic-execution.pdf>.
- [Thakur et al., 2010] Thakur, A., Lim, J., Lal, A., Burton, A., Driscoll, E., Elder, M., Andersen, T., and Reps, T. (2010). Directed proof generation for machine code. In *Proceedings of the 22nd International Conference on Computer Aided Verification, CAV 2010*, pages 288–305, Berlin, Heidelberg. Springer-Verlag, ISBN: 3-642-14294-X, 978-3-642-14294-9, DOI: [10.1007/978-3-642-14295-6_27](https://doi.org/10.1007/978-3-642-14295-6_27), http://dx.doi.org/10.1007/978-3-642-14295-6_27.
- [Tillmann and De Halleux, 2008] Tillmann, N. and De Halleux, J. (2008). Pex: White box test generation for .net. In *Proceedings of the 2nd International Conference on Tests and Proofs, TAP 2008*, pages 134–153, Berlin, Heidelberg. Springer-Verlag, ISBN: 3-540-79123-X, 978-3-540-79123-2, DOI: [10.1007/978-3-540-79124-9_10](https://doi.org/10.1007/978-3-540-79124-9_10), <http://dl.acm.org/citation.cfm?id=1792786.1792798>.
- [Trtík and Strejček, 2014] Trtík, M. and Strejček, J. (2014). *Symbolic Memory with Pointers*, pages 380–395. ATVA 2014. Springer International Publishing, Cham, DOI: [10.1007/978-3-319-11936-6_27](https://doi.org/10.1007/978-3-319-11936-6_27), http://dx.doi.org/10.1007/978-3-319-11936-6_27.
- [Udupa et al., 2005] Udupa, S. K., Debray, S. K., and Madou, M. (2005). Deobfuscation: Reverse engineering obfuscated code. In *Proceedings of the 12th Working Conference on Reverse Engineering, WCRE 2005*, pages 45–54, Washington, DC, USA. IEEE Computer Society, ISBN: 0-7695-2474-5, DOI: [10.1109/WCRE.2005.13](https://doi.org/10.1109/WCRE.2005.13), <http://dx.doi.org/10.1109/WCRE.2005.13>.
- [Visser et al., 2012] Visser, W., Geldenhuys, J., and Dwyer, M. B. (2012). Green: Reducing, reusing and recycling constraints in program analysis. In *Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering, FSE 2012*, pages 58:1–58:11, New York, NY, USA. ACM, ISBN: 978-1-4503-1614-9, DOI: [10.1145/2393596.2393665](https://doi.org/10.1145/2393596.2393665), <http://doi.acm.org/10.1145/2393596.2393665>.

- [Visser et al., 2004] Visser, W., Păsăreanu, C. S., and Khurshid, S. (2004). Test Input Generation with Java PathFinder. In *Proceedings of the 2004 ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA 2004, pages 97–107, New York, NY, USA. ACM, ISBN: 1-58113-820-2, DOI: [10.1145/1007512.1007526](https://doi.org/10.1145/1007512.1007526), <http://doi.acm.org/10.1145/1007512.1007526>.
- [Visser et al., 2006] Visser, W., Păsăreanu, C. S., and Pelánek, R. (2006). Test input generation for java containers using state matching. In *Proceedings of the 2006 International Symposium on Software Testing and Analysis*, ISSTA 2006, pages 37–48, New York, NY, USA. ACM, ISBN: 1-59593-263-1, DOI: [10.1145/1146238.1146243](https://doi.org/10.1145/1146238.1146243), <http://doi.acm.org/10.1145/1146238.1146243>.
- [Wang et al., 2011] Wang, Z., Ming, J., Jia, C., and Gao, D. (2011). *Linear Obfuscation to Combat Symbolic Execution*, pages 210–226. Springer Berlin Heidelberg, Berlin, Heidelberg, ISBN: 978-3-642-23822-2, DOI: [10.1007/978-3-642-23822-2_12](https://doi.org/10.1007/978-3-642-23822-2_12), http://dx.doi.org/10.1007/978-3-642-23822-2_12.
- [Weiser, 1984] Weiser, M. (1984). Program Slicing. *IEEE Transactions on Software Engineering*, SE-10(4):352–357, ISSN: 0098-5589, DOI: [10.1109/TSE.1984.5010248](https://doi.org/10.1109/TSE.1984.5010248), <http://dx.doi.org/10.1109/TSE.1984.5010248>.
- [Xie et al., 2009] Xie, T., Tillmann, N., de Halleux, J., and Schulte, W. (2009). Fitness-guided path exploration in dynamic symbolic execution. In *Proceedings of the 2009 IEEE/IFIP International Conference on Dependable Systems and Networks*, DSN 2009, pages 359–368. DOI: [10.1109/DSN.2009.5270315](https://doi.org/10.1109/DSN.2009.5270315), <http://dx.doi.org/10.1109/DSN.2009.5270315>.
- [Xie and Aiken, 2005] Xie, Y. and Aiken, A. (2005). Scalable error detection using boolean satisfiability. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL 2005, pages 351–363, New York, NY, USA. ACM, ISBN: 1-58113-830-X, DOI: [10.1145/1040305.1040334](https://doi.org/10.1145/1040305.1040334), <http://doi.acm.org/10.1145/1040305.1040334>.
- [Yadegari and Debray, 2015] Yadegari, B. and Debray, S. (2015). Symbolic execution of obfuscated code. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, CCS 2015, pages 732–744, New York, NY, USA. ACM, ISBN: 978-1-4503-3832-5, DOI: [10.1145/2810103.2813663](https://doi.org/10.1145/2810103.2813663), <http://doi.acm.org/10.1145/2810103.2813663>.
- [Yadegari et al., 2015] Yadegari, B., Johannesmeyer, B., Whitely, B., and Debray, S. (2015). A generic approach to automatic deobfuscation of executable code. In *Proceedings of the 2015 IEEE Symposium on Security and Privacy*, SP 2015, pages 674–691, Washington, DC, USA. IEEE Computer Society, ISBN: 978-1-4673-6949-7, DOI: [10.1109/SP.2015.47](https://doi.org/10.1109/SP.2015.47), <http://dx.doi.org/10.1109/SP.2015.47>.
- [Yang et al., 2012] Yang, G., Păsăreanu, C. S., and Khurshid, S. (2012). Memoized symbolic execution. In *Proceedings of the 2012 International Symposium on Software Testing and Analysis*, ISSTA 2012, pages 144–154, New York, NY, USA. ACM, ISBN: 978-1-4503-1454-1, DOI: [10.1145/2338965.2336771](https://doi.org/10.1145/2338965.2336771), <http://doi.acm.org/10.1145/2338965.2336771>.
- [Zaddach et al., 2014] Zaddach, J., Bruno, L., Francillon, A., and Balzarotti, D. (2014). AVATAR: A framework to support dynamic security analysis of embedded systems’ firmwares. In *21st Annual Network and Distributed System Security Symposium*, NDSS 2014. <http://www.internetsociety.org/doc/avatar-framework-support-dynamic-security-analysis-embedded-systems-firmwares>.
- [Zitter, 2013] Zitter, K. (2013). How a crypto backdoor pitted the tech world against the nsa. <https://www.wired.com/2013/09/nsa-backdoor/all/>.