Detecting Vulnerabilities in C Programs Using Trace-Based Testing

Dazhi Zhang, Donggang Liu, Yu Lei, David Kung, Christoph Csallner, and Wenhua Wang
Department of Computer Science and Engineering
The University of Texas at Arlington

Abstract

Security testing has gained significant attention recently due to frequent attacks against software systems. This paper presents a trace-based security testing *approach. It reuses test cases generated from previous testing methods to produce* execution traces. *An execution trace is a sequence of program statements exercised by a test case. Each trace is symbolically executed to produce* program constraints *and* security constraints. *A program constraint is a constraint imposed by program logic on program variables. A security constraint is a condition on program variables that must be satisfied to ensure system security. A security flaw exists if there is an assignment of values to program variables that satisfies the program constraint but violates the security constraint. This approach detects security flaws even if existing test cases do not trigger them. The novelty of this method is a test model that* unifies program constraints and security constraints *such that formal reasoning can be applied to detect vulnerabilities. A tool named* SecTAC *is implemented and applied to 14 benchmark programs and 3 open-source programs. The experiment shows that* SecTAC *quickly detects all reported vulnerabilities and 13 new ones that have not been detected before.*

1   Introduction

Software security has gained significant attention in recent years due to the huge number of security attacks that exploit vulnerabilities in software. *Security testing* is becoming an active area of research, aiming at identifying software vulnerabilities effectively. Recently, many approaches have been proposed to detect vulnerabilities in programs [25, 14, 11, 5, 22, 1, 20, 6, 23, 4, 10, 12, 26].

Static analysis has been used to scan source code for errors that crash a system or cause security problems [24, 25]. These static analysis tools use heuristics to determine if a security problem could occur; they usually approximate or even ignore runtime dynamics such as branch conditions and how buffer elements are visited. Thus, they are often imprecise, causing many false alarms.

Dynamic analysis examines program execution to detect security problems [13, 8, 9, 1, 20, 23]. These tools feed test data to a program and monitor its runtime behavior. A security vulnerability is detected if the behavior is considered abnormal, e.g., the program accessed a buffer outside its bounds. Although dynamic analysis reduces false alarm rates, it requires test inputs that actually cause security problems. This places a huge burden on testers.

Dynamic symbolic execution, also called *concolic testing*, is often used in automatic test data generation for finding errors that crash a system or cause security problems [11, 22, 21, 6, 4, 10, 12, 26]. These tools perform concrete and symbolic execution of a program simultaneously to explore as many paths as possible. They do not need inputs that can actually cause security problems. However, they are either ineffective in the sense that unguided path exploration may not cover important vulnerabilities, or do not scale well to large and complex programs.

In this paper, we propose a novel security testing approach using *trace-based symbolic execution* and *satisfiability analysis*. Trace-based symbolic execution avoids the search space explosion of conventional symbolic execution. In our approach, each existing test case is used to generate an execution trace, i.e., the sequence of exercised program statements. Symbolic execution is then applied to produce two kinds of predicates. The first predicate is called a *program constraint* (PC), which specifies a condition that program variables must satisfy after the execution of the statement. The second predicate is called a *security constraint* (SC), which specifies the condition that program variables must satisfy to ensure the security of the given program execution. A security vulnerability is detected if there is an assignment of values to program variables satisfies PC but violates SC, i.e., PC $\wedge \neg$ SC is satisfiable.

The advantages of our approach are as follows. First, As opposed to previous approaches, we can guide our search to focus on those features of the user program that are most important to the user–as indicated by the developers' willingness to write test cases for them. Our approach can detect security flaws even if these existing test cases do not trigger them. In other words, our technique can generate new inputs that trigger security problems, even if the user-supplied inputs do not. Second, we propose a test model that unifies program constraints and security constraints using logical expressions so that formal reasoning can be performed to detect security vulnerabilities. Hence, our ap-

proach can handle new types of vulnerabilities by simply formulating new security requirements for them. Third, trace-based symbolic execution also makes it possible to test programs for vulnerability in parallel. This is because analyses on different execution traces are independent from each other. We can partition the test cases into a number of disjoint subsets and analyze these subsets in parallel. This cannot be directly achieved in dynamic symbolic execution based approaches since test cases exercising different paths are generated during path exploration. Certainly, tools like DART [11] and CUTE [22] can be modified to reuse existing test cases and only test the paths exercised by these test cases. However, in this case, they lose the benefit of automatically exploring program paths.

To evaluate the effectiveness of our approach, we implemented a tool named SecTAC (A Security Testing Approach for C programs) and applied it to 14 benchmark programs given in [28] and 3 open source programs. The benchmark programs were designed to evaluate buffer overflow detection tools by simulating historic real-world vulnerabilities in server programs. Compared with the results in [28, 27, 26], SecTAC can detect every reported vulnerability as long as the vulnerability exists in the execution traces tested in our experiments. In addition, SecTAC detected 6 previously unreported vulnerabilities in the 14 benchmark programs. SecTAC also detected 7 vulnerabilities in the open-source programs that, to the best of our knowledge, have not been reported previously.

The rest of this paper is organized as follows. In the next section, we explain our basic ideas. In Section 3, we overview the SecTAC design. In Section 4, we describe the SecTAC implementation. In Section 5, we present the experiment result. In Section 6, we review related work. We discuss the limitations of SecTAC in Section 7 and draw some conclusions in Section 8.

2 Basic Ideas of Our Approach

Software systems must be tested to ensure that the required functionalities are correctly implemented. Unlike conventional software testing, our goal is to detect security vulnerabilities that exist in the software system. A program is said to be *vulnerable* if there is an execution path that can be exploited to compromise the security of the system. To detect such security vulnerability, we rely on a set of *security requirements* that must be satisfied by all execution paths of the program. An example of security requirements is that the length of the string copied to a buffer using strcpy must not exceed the capacity of the buffer.

Testing for security vulnerabilities implies the generation of test cases that can effectively detect violations of security requirements. However, it is well known that effective test case generation is both difficult and time-consuming. Therefore, it is desirable to reuse the test cases that are already generated during conventional software testing. The merit of this is twofold. First, these test cases typically ac-

complish some required coverage criteria such as branch coverage. Second, the branches covered by the test cases are deemed important by the developer. Our goal is to provide a security testing method for software developers who have access to the source program and the test cases produced by traditional functional testing.

In our approach, we use existing test cases to generate execution traces. Each execution trace is a sequence of source code statements exercised by a test case. There are no loops in execution traces since a loop in the original program will be unfolded when it is exercised by a test case. We then symbolically execute each execution trace to determine whether it contains a security vulnerability. Symbolic execution of each trace produces two kinds of predicates. The first predicate is the *program constraint* (PC), which is updated during the symbolic execution of the trace; it specifies a condition that the program variables must satisfy. In other words, the program constraint specifies the possible values of variables at each point during the symbolic execution of the trace. The second predicate is the *security constraint* (SC), which is produced at certain points during the symbolic execution of the trace; it specifies a condition that program variables must satisfy to ensure the security of the software system. A security problem will occur when the values of some variables violate the security constraint. Testing C programs for vulnerabilities is therefore equivalent to determining whether at each point in the trace, there exists an assignment of values to program variables that satisfies PC but violates SC.

Program constraints: The program constraint at a given point in the trace is determined by the program statements exercised to reach this point. These statements include declaration statements, assignment statements, branching statements, and library function calls; they impact the values of variables as follows:

- A declaration statement contains important information about the *type* and *size* of the declared program variable. These two pieces of information determine the initial program constraint on the variable. As an example, the declared size of a buffer or an array constrains the space available for holding data.

- An assignment statement constrains the value of its left expression to the result of its right expression.

- A branching statement indicates that different execution paths could be taken under different conditions. However, our execution trace is produced by running the program under a real test case. We already know which execution path is taken by the test case. Hence, we can immediately determine a condition expression that specifies a constraint between the involved variables. For example, if statement "if(i>j)" exercises the FALSE branch, we know that $i \leq j$ is a constraint between i and j.

- A library function call restricts the range of its return value if it has one. For example, the return value of func-

tion `open` is always greater than or equal to -1. In addition, some library functions have side-effects (i.e., modifying the states in addition to returning a value) that also impose constraints on variables. For example, calling function `getcwd` will change the content of the buffer specified by the parameter.

According to the above rules, symbolically executing each statement produces an expression describing the constraint between the program variables involved in the statement. To distinguish it from the program constraint (PC), we call such expression the *program constraint conjunction* (PCC). PCC may get updated during program execution. The program constraint at any given point in the trace can be expressed as the conjunction of all current PCCs.

Security constraints: Producing security constraints requires clearly-defined high-level security requirements, e.g., the length of the string copied to a buffer must not exceed the capacity of the buffer. A wide range of security vulnerabilities like buffer overflow, SQL injection, and format string, are caused by improper uses of operations such as `strcpy`, `sql.exec`, and `printf`. Correct uses of such operations can be expressed as security requirements, which can then be used to generate security constraints. For example, a security requirement for `strcpy` will be "*the length of the second argument must not exceed the capacity of the first argument*". If the trace includes a statement `strcpy(a,b)`, where `a` is a buffer and `b` is a string, we produce a security constraint: `a.space>b.strlen`, where `a.space` is the capacity of buffer `a` and `b.strlen` is the length of string `b`. We use first-order logic to express security constraints.

security-critical func.	security requirement
strcpy(dst,src)	dst.space>src.strlen
strncpy(dst,src,n)	(dst.space $\geq$ n) $\wedge$ (n $\geq$ 0)
strcat(dst,src)	dst.space>dst.strlen + src.strlen
getcwd(buf,size)	(buf.space$\geq$ size) $\wedge$ (size $\geq$ 0)
fgets(dst,size,f)	(dst.space $\geq$ size) $\wedge$ (size $\geq$ 0)
scanf(format, ...)	# formats = # parameters-1
printf(format, ...)	# formats = # parameters-1

Table 1. Security requirements for library function calls. "x.space" is the size of the memory allocated to x and "x.strlen" is the string length of x.

`SecTAC` can detect the violation of a security requirement as long as such requirement can be expressed as a condition that program variables must satisfy. In the current implementation, we support two kinds of security requirements: *pointer addition requirements* and *function parameter requirements*. The former is derived from a useful observation made in [15], i.e., the result of a pointer addition must point to the same original object. The latter is generated from *security-critical library functions*, i.e., the library functions whose parameters must satisfy a condition to ensure the security of a software system. For example,

functions `strcpy` and `printf` are both security-critical library functions. We have selected 20+ library functions that are well known to be "insecure" and formulated their security requirements. Table 1 shows some of these functions and their security requirements. Although these requirements are written by hand, in practice we have found it to be not too difficult for well-known functions.

```
1:    void foo(int a,char *s){
2:        char buf[10];
3:        if(a>0)
4:            strcpy(buf,s);
5:    }
```

Figure 1. A sample program

An example: Figure 1 shows a sample program, which copies the second argument `s` into a buffer, if the first argument is greater than 0. Assume that there is only one security requirement, i.e., the length of a string copied to a buffer using function `strcpy` must not exceed the capacity of the buffer. Furthermore, we assume that both arguments are user inputs, meaning they can be any values that are not known in advance. Now, consider a test case that includes the call `foo(x,y)` with `x=1` and `y="test"`. This test case generates an execution trace (1,2,3,4) of statement numbers. Although this test case does not trigger any security problem, we will demonstrate that our method can effectively find the vulnerability in the trace. Table 2 shows the result of symbolically executing this execution trace. The first column indicates the statement number, and the second and third columns give the program and security constraints at the respective statements.

Line#	Program Constraint	Security Constraint
1	(MIN$\leq$ a$\leq$ MAX)$\wedge$ (s.strlen$\geq$ 0)	TRUE
2	(MIN$\leq$ a$\leq$ MAX)$\wedge$ (s.strlen$\geq$ 0)	TRUE
3	(0<a$\leq$ MAX)$\wedge$ (s.strlen$\geq$ 0)	TRUE
4	(0<a$\leq$ MAX)$\wedge$ (s.strlen$\geq$ 0)	s.strlen<10

Table 2. Program and security constraints for the execution trace $(1, 2, 3, 4)$

As shown in the table, the PC at statement 1 is (MIN$\leq$a$\leq$MAX)$\wedge$(s.strlen$\geq$ 0), where [MIN, MAX] defines the range of an integer number, which is usually machine dependent, and s.strlen is a symbolic value denoting the length of string s. This is because both a and s are user inputs, i.e., a can be any integer value and s can be any string. The security constraint at statement 1 is TRUE since the statement does not include any operation that may violate any security requirement. More specifically, it does not include a call to the `strcpy` function. Statement 2 is a declaration statement of a buffer; it sets the space of the buffer to `10`. We do not include this in the program constraint. Instead, we directly update the `space` field of the

buffer, i.e., `buf.space=10`, for simplicity.

Statement 3 is a condition statement and the test case exercises the TRUE branch, which implies that `a > 0` must be TRUE. Thus, the program constraint changes from $(MIN \leq a \leq MAX) \wedge (s.strlen \geq 0)$ to $(0 < a \leq MAX) \wedge (s.strlen \geq 0)$, as shown in the third line of the table. A security constraint is produced at statement 4, as shown in the fourth line of the table. The reason is that function `strcpy` is associated with a security requirement, i.e., the string length of the second argument must be less than the space allocated to the first argument. As a result, we produce a security constraint: `s.strlen<buf.space`. Since `buf.space=10`, we have `s.strlen<10`.

A security vulnerability exists at a given point if an assignment of values to variables satisfies PC but violates SC, i.e., $PC \wedge \neg SC$ is satisfiable. At statement 4, we check the satisfiability of $PC \wedge \neg SC$, i.e., $(0 < a \leq MAX) \wedge (s.strlen \geq 0) \wedge \neg(s.strlen < 10)$. We use a theorem prover and find that `a=1` and `s="012345678910"` satisfies $PC \wedge \neg SC$. Thus a test case can be generated to uncover the vulnerability.

3 SecTAC Design

The goal of `SecTAC` is to detect security vulnerabilities in a program. As discussed, `SecTAC` reuses existing test cases for achieving high coverage and reducing testing effort. Specifically, we extract the execution trace of the program under each test case and then analyze each execution trace to determine whether it contains a security vulnerability. Figure 2 shows the workflow of `SecTAC`.
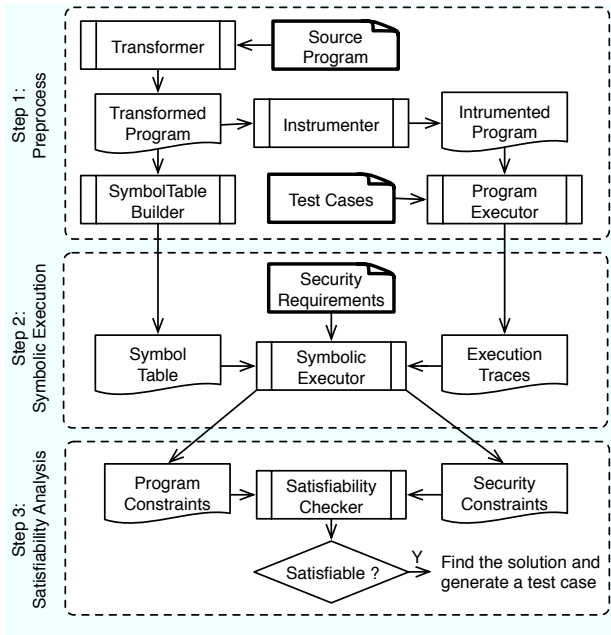

Figure 2. `SecTAC` **Workflow**

`SecTAC` performs security testing through three steps, *preprocessing*, *symbolic execution*, and *satisfiability anal-*

ysis, as indicated in Figure 2. In preprocessing, we generate execution traces from existing test cases and prepare the symbol table for tracking the state of program variables; in symbolic execution, we analyze every execution trace to extract the program and security constraints at each point in the trace; and in satisfiability analysis, we find inputs that can detect security vulnerabilities.

Preprocessing: In this step, we first use the *transformer* to transform the source program into three-address code to simply the analysis. To obtain execution traces, the *instrumenter* parses and inserts the trace-logging code into this transformed program. This transformed, instrumented program is compiled and then executed by the *program executor* using all test cases. The trace-logging code generates an execution trace for each test case.

The *symbol-table builder* constructs a *symbol table* for all program variables for effectively tracking the program constraints on them. In addition to the size and type information, each program variable is also associated with additional attributes. For example, for a pointer that points into a buffer, we introduce two attributes to track *which buffer and which position in the buffer it points into* so that we can test the out-of-bounds buffer access.

Symbolic execution: We use the *symbolic executor* to symbolically execute the trace to capture program constraints and check the pattern of each executed statement against the security requirements. Whenever a security requirement applies, e.g., a security-critical function call or a pointer addition statement is exercised, we generate a security constraint corresponding to such security requirement. The program and security constraints are predicates on the symbolic values of program variables and their attributes.

Satisfiability analysis: For each statement in the trace that generates a security constraint (SC), we get the program constraint (PC) at that statement and use a *satisfiability checker* to check if $PC \wedge \neg SC$ is satisfiable. If it is, a security vulnerability is detected. The solution given by the satisfiability checker is then used to generate test data to uncover the vulnerability. We express both program and security constraints using the SMT-LIB format [19] and use the Yices SMT-solver [7] as the satisfiability checker.

4 SecTAC Implementation

In this section, we describe the implementation of `SecTAC` in detail based on the workflow in Figure 2.

4.1 Step 1: Preprocessing

The main tasks of preprocessing are (1) *generating execution traces* and (2) *constructing the symbol table*.

4.1.1 Generating Execution Traces
In `SecTAC`, the program is transformed by CIL [18], instrumented by the Java parser generator `JavaCC`, and executed under each test case to produce the corresponding execution trace. An execution trace was previously defined in Section 2 as a sequence of source code statements exercised by a

test case. This definition facilitates the understanding of the basic ideas of our approach. However, our implementation generates execution traces consisting of sequences of expressions and special marks. Expressions are either assignment statements or library function calls. Special marks are used to indicate: (1) function call entry and exit, (2) conditional branching, (3) parameter passing, and (4) returning of values to variables from function calls.

Note that declaration statements are not included in execution traces since they are not "executed" by test cases. However, they contain important information about the *type* and *size* of program variables. SecTAC handles declaration statements in the symbol-table builder.

4.1.2 Constructing the Symbol Table

The symbol table is used to track the state of program variables; it includes information about all program variables and user-defined functions in the trace. Specifically, the symbol-table builder parses the program and creates a *data object* for each program variable and a *function object* for each user-defined function. These objects include various attributes to track the state of program variables. Next we describe the creation of the objects and discuss features added to address *pointer dependency*.

Data objects: The symbol-table builder creates a *data class* for every program variable type. A data class includes the size and type information as well as some other attributes about the data type it represents; an object of this class is created for each program variable of this data type. We have a pre-defined class for each primitive type or primitive type with qualifiers. For example, we use classes `Int` and `BCharacter` for integers and characters declared in the program, respectively. For each composite type, we create a class using its type name. We also have a pre-defined class `Pointer` for pointers, arrays, and buffers. All the above data classes are extended from a common *base class* `BaseType` that defines common attributes such as name, type, and symbolic value. It also includes a `typesize` field to record the size of the memory allocated to the variable. For example, the `typesize` field of an `Int` object is 4 in a 32-bit computer.

Function objects: SecTAC also creates a class for each user-defined function to facilitate the trace analysis, i.e., help locate objects in the scope of any user-defined function. For every function class `f`, we create objects for the parameters to the corresponding function and the local variables declared in this function. These objects are the members of this function class `f`. Other statements in the function body are not included in class `f`.

All function classes are extended from a common abstract base class `Function` that includes a `getObject` method, which can be used to locate the object representing a local variable or function parameter in the scope of a user-defined function given a name.

In C programs, the global variables or static variables declared in the file scope are not included in any function. To track these variables, SecTAC also constructs a `Global` class and a *file-scope* class for each file, and puts the variables in these classes accordingly.

Pointer dependency: It is possible that several pointer-type variables are declared and point to the same array. For example, we can declare "`char p[10]`" and define a pointer "`char *q=p+5`" in a C program. We know that both pointers `p` and `q` point into the same array. The only difference is that `p` points to the beginning of the array, while `q` points to the sixth element of the array. The pointer objects are said to be *related* or *dependent* if they point into the same array. Hence, `p` and `q` are related. We notice that the operation on a given pointer object may impact its related pointer objects. For example, if we copy a string of length 6 to `p`, then the string lengths of `p` and `q` become 6 and 1, respectively. If we immediately copy another string of length 4 to `q`, then the string lengths of `p` and `q` become 9 and 4, respectively.

To correctly analyze the impact of pointer operations on *related* pointer objects, the `Pointer` class also includes a `start` field and a `space` field. A pointer object uses `start` to record its starting position in the array, and `space` to record the size of the space from its starting position to the end of the array. Thus, we can determine how the operation on one pointer object can impact others. From the previous example, we know that the `start` fields of the objects for `p` and `q` are 0 and 5, respectively. If a string of length 6 is copied to `p`, then we immediately know that `q` is impacted and its string length should be 1.

Object locating: Object locating addresses how to determine the target object(s) of a program statement. For example, for statement "`i=j.id;`", we need to locate the objects created for variable `i` and the member `id` of the structure `j`. As discussed before, each function class provides a method `getObject` to locate the object created for variables in its scope given a name. However, when a member of a composite type variable, e.g., `j.id` in the above example, is referenced, we need to further locate the *member object* representing the member of this variable. Every class created for a composite type variable (e.g., struct or array) has a method `getObject` to locate the member object given a *name* or an *offset*.

4.2 Step 2: Symbolic Execution

Once we have the execution trace and the symbol table, we start to analyze the execution trace statement by statement to capture the program and security constraints using symbolic execution (symbolic value propagation).

4.2.1 Producing Program and Security Constraints

The program constraint will be initialized when we are building the symbol table. Specifically, when we create an object for a program variable, we produce a program constraint conjunction according to the variable declaration information. For example, statement "`int i;`" leads to the creation of an `Int` type object `i`. Thus, we produce

a program constraint conjunction (MIN $\leq$ `i.sym` $\leq$ MAX), where `i.sym` is the symbolic value of `i`. The program constraint will also be updated when a statement in the trace is symbolically executed.

- If it is an assignment statement, the attributes of the object for the right part determines the attributes of the object for the left part. In this case, we directly update the attributes of the left object instead of updating the program constraint.

- If it is a branch statement, we update the program constraint based on which branch is exercised. For example, a conditional expression "@true i>j" indicates that the TRUE branch is exercised. Thus, `SecTAC` generates a program constraint conjunction (`i.sym>j.sym`).

- If the statement calls a library function, we need to update the program constraint according to its semantics. If the return value of the library function is assigned to a variable, we generate a program constraint conjunction according to this return type. Since some library functions have constraints on their return values, a program constraint conjunction that further restricts the range of the returned value is produced. For example, the return value of `fopen` is always greater than or equal to -1, which is different from the default range of its return type. In addition, some library functions have side-effects on their parameters. Some side-effects can be considered as equivalent to updating the object attributes, e.g., for `strcpy(dst,src)`, the `strlen` field (a symbolic value that denotes the string length) of the `dst` object is updated to that of the `src` object. Some side-effects, however, impose constraints on the involved parameters. For example, after calling `getcwd(buf,n)`, the `strlen` of `buf` is less than n if the length of the current path is less than n, and unchanged otherwise. In this case, we also generate a program constraint conjunction.

A program statement in the execution trace is said to be *security critical* if it may violate a security requirement. In the current implementation of `SecTAC`, any statement involving either a security-critical function or a pointer addition is a security-critical statement. `SecTAC` produces a security constraint, i.e., a first order logic expression, at every security-critical statement.

4.2.2 Algorithm for Symbolic Executor

We now describe the detail of the `SecTAC` symbolic executor. We first create a *stack* to keep track of the current function object, i.e., the active function object in use, which is always the one at the top of the stack. `SecTAC` then processes each statement in the trace according to the following rules: (1) if it is a function entry, `SecTAC` creates a new object of this function class and pushes the object into the stack; (2) if it is a function return, `SecTAC` pops an object from the stack; (3) If it is an assignment statement, `SecTAC` performs symbolic execution on the left

and right expressions, and updates the object attributes for the involved variables; (4) if it is a conditional statement, `SecTAC` produces a program constraint conjunction that captures which branch is exercised; (5) if it is a library function call, `SecTAC` processes as follows. If the function is in the right part of an assignment statement, a new object is created according to its return type. If the function further limits its return value to a smaller range compared to its type, the program constraint on this object is updated. If the function also has side-effects, the attributes of the involved objects are updated accordingly, and the program constraint is also updated as needed. If the function is also a security-critical function, a security constraint is generated.

Symbolic Execution on Expressions: A critical part of symbolic execution is the symbolic execution on expressions. The symbolic execution procedure on a given expression `e` works as follows: (1) if `e` is a constant number or character, a new object of the class for such data type is created, and its symbolic value is set to this constant value; (2) if it is a constant string, a `Pointer` object is created, and its `strlen` field is set to be the length of this constant string; (3) if it is a variable, we will locate the corresponding object and return it; (4) if it is $*v$, we locate the object corresponding to v and return the object specified by the `point_to` field of this pointer object; (5) if it is $\&v$, we locate the object corresponding to v and create a `Pointer` object. We then set the `point_to` field of the newly created object to the object corresponding to v; (6) if it is $v.m$, we locate the object of v, then return its member object with the name m; (7) if it is e_1 *op* e_2, we recursively perform symbolic execution on expressions e_1 and e_2. Based on the types of the returned objects, we take different actions; (8) if it is a library function call, we handle it in the same way as we handle library function calls.

`SecTAC` generates a security constraint for every pointer addition to check whether the result still points to the same original object. We thus take special care of the addition between a `Pointer` object `p` and an `Int` object `i` as follows:

- If `p` points to a buffer, we create a new `Pointer` object `obj` and set its `space`, `start`, and `strlen` fields based on `p` and `i`. Specifically, `obj.space` and `obj.start` are set to `p.space-i.sym` and `p.start+i.sym` respectively. `obj.strlen` is set to the following conditional expression:

  `((p.strlen`$\geq$`i.sym) (p.strlen-i.sym) newsym)`

  This expression indicates that `obj.strlen` is set to `p.strlen-i.sym` if `p.strlen`$\geq$`i.sym`, and a new symbol `newsym` otherwise. A program constraint conjunction is also produced for the new symbol `newsym`, i.e., `newsym`≥ 0. Finally, object `obj` is returned.

- If `p` points to a composite type object, e.g. array or struct, then we need to find a member object inside this composite object through offset `i`. In this case, we use the

getObject(i) method in object p.point_to to locate and return the object.

- If p points to neither a buffer nor a composite type data, then it is a pointer arithmetic. In this case, a new object will be created in a similar way as the first case. The only difference here is that the strlen field need not be set.

Pointer analysis : We will discuss how we address the pointer dependency problem mentioned in Section 4.1.2. Specifically, when we create an object for a buffer, we also include a number of links in this object through which we can locate all Pointer objects that operate on this buffer. Let us consider a particular pointer p that points into a buffer. When we update the object for this pointer, we will need to find the object for the original buffer this pointer points into and locate all Pointer objects that operate on this buffer. Let q be a Pointer object we find. We first check p.start and q.start to decide their relative positions in the buffer. There are two cases:

- If q's position in the buffer is before that of p's, we compare q.strlen with the distance between them. If q.strlen is larger than the distance, we have to update q.strlen accordingly.

- If q's position in the buffer is after p's, we compare p.strlen with the distance between their positions. If p.strlen is larger than the distance, we have to update q.strlen accordingly.

4.3 Step 3: Satisfiability Analysis

Finally, the program and security constraints are expressed in SMT-LIB [19] format, which is recognized by many SMT solvers. We use the SMT solver Yices [7] to check the satisfiability of PC $\wedge\neg$ SC for each SC and the PC at the same point in the trace. Note that the PC at a given point in the trace may include a huge number of conjunctions. In this case, checking the satisfiability of PC $\wedge\neg$ SC could be very expensive. However, we note that a lot of PC conjunctions are actually irrelevant to SC since they only involves variable symbols that do not impact SC. Removing these irrelevant conjunctions will not change the result of satisfiability analysis. We thus use only *SC-dependent* PC conjunctions to save the cost. Two conjunctions are said to be *directly related* if they include at least one common variable symbol. Then, starting from an empty S, we first identify all PC conjunctions that are directly related to SC and put them in S. We then *repeatedly* check every PC conjunction and add it into S if it is SC-dependent, i.e., directly related to at least one conjunction in S. We stop when there are no more SC-dependent PC conjunctions. Let PC$'$ be the conjunction of all conjunctions in S. We only need to check the satisfiability of PC$' \wedge\neg$ SC instead of PC $\wedge\neg$ SC.

4.4 False Negatives and Positives

SecTAC can detect a security vulnerability in a program if (1) the vulnerability is modeled by one of the security requirements, (2) an execution path that can trigger such security problem is exercised by one of the test cases, (3) the program and security constraints are derived correctly, and (4) the theorem prover for satisfiability analysis can correctly find a solution if PC $\wedge\neg$ SC is satisfiable. In other words, there will be false negatives if one of the above three conditions is false. For example, if we derive constraints on library functions from documentation, there might be false negatives due to the inconsistency between the documentation and the actual implementation. We can use LFI to check such consistency [17]. Similarly, SecTAC will generate a false positive if (1) the theorem prover returns a solution when PC $\wedge\neg$ SC is not satisfiable or (2) the program and security constraints are extracted incorrectly. In our experiments, we did not find any false positive.

5 Experiments

To evaluate the effectiveness of our approach, we applied SecTAC on 14 benchmark programs [28], two open source http server programs, nullhttpd-0.5.1 and lancer, and an open source ftp server program bftpd-2.3. We used their latest versions in our experiment. The benchmark programs represent various kinds of memory corruption vulnerabilities in certain versions of the **Bind**, **Sendmail**, and **Wu-ftp** programs. They have been used to evaluate the effectiveness of many buffer overflow detection tools [28, 27, 26]. For each of these programs, there is a buggy version and a fixed version. We used the buggy version in our experiment. Our results show that SecTAC can detect every reported vulnerability as long as the vulnerability exists in the traces. In addition, SecTAC also detected six vulnerabilities in the benchmark programs, four vulnerabilities in nullhttpd-0.5.1, four vulnerabilities in lancer, and one vulnerability in bftpd-2.3 that, to the best of our knowledge, have not been reported previously. Next, we will report our findings in detail.

Table 3 summarizes our experimental results. The first 14 rows show the result of evaluating SecTAC on the 14 benchmark programs [28]. As shown in the last column of the table, we found new vulnerabilities in Bind 4, Sendmail 1, Sendmail 3, Wu-ftp 2, Wu-ftp 3, nullhttpd-0.5.1, lancer, and bftpd2.3 programs.

Test inputs: For each buggy benchmark program version, a specific input file or hard-coded assignment to variables is provided in [28] as the test data to trigger the vulnerability. However, a major merit of SecTAC is that it can detect vulnerabilities under test cases from functional testing that do not trigger vulnerability. Hence, in our experiments, *whenever it is possible, we construct test inputs that exercise paths containing the reported vulnerabilities but do not trigger them*. Only when it is impossible to find a test case exercising the known vulnerable path without triggering the vulnerability, do we use the test input provided in [28]. For the sake of presentation, we call a test case *normal* if it does not trigger any vulnerability. We call the test cases

Program	LOC	Input	LOT	Time(mm:ss)	#KnownBugs	#FoundBugs	#FP	Remark
Bind 1	1116	www.cnn.com	539	00:01	1	1	0	
Bind 2	1306	cnn.com	1117	00:01	1	1	0	
Bind 3	380	default	365	00:01	1	1	0	
Bind 4	645	www.nbc.com; www.cnn.com	162	00:01	1	2	0	1 new bug
Sendmail 1	537	default	6207	00:02	6(5)*	6	0	1 new bug
Sendmail 2	791	default	5509	00:03	1	1	0	
Sendmail 3	416	default	2534	00:03	1	2	0	1 new bug
Sendmail 4	485	default	1379	00:03	4	4	0	
Sendmail 5	622	default	6669	00:03	3	3	0	
Sendmail 6	390	default	129	00:01	1	1	0	
Sendmail 7	929	default	2145	00:03	2	2	0	
Wu-ftp 1	503	/tmp/aa	79	00:01	4	4	0	
Wu-ftp 2	744	/tmp/test.c	106	00:01	1	2	0	1 new bug
Wu-ftp 3	689	/tmp/aa	399	00:01	6	8	0	2 new bugs
nullhttpd-0.5.1	2328	50 test cases	12447	08:07	1	3	0	2 new bug
lancer	4261	50 test cases	118657	49:18	0	4	0	4 new bugs
bftpd-2.3	5766	10 test cases	65027	11:42	0	1	0	1 new bug

* According to the BAD marks in the program, there are 6 bugs in the trace. However, we found that one of them is not a bug.

Table 3. Experimental Results. "LOC" represents the number of lines of the code; "Input" represents the program input we use; "LOT" represents the number of lines in the execution trace exercised by the test case; "Time" represents the time that our tool used; "#KnownBugs" is the number of previously reported vulnerabilities in the execution trace; "#FoundBugs" is the number of vulnerabilities found by SecTAC; and "#FP" is the number of false positives.

provided in [28] *default* in the table.

For http server programs, we randomly generate 50 normal http requests. For the ftp server program, we manually generate 10 test cases that include basic ftp commands such as "ls", "get", and "put". We use the GCC bounds checking extension to monitor the program execution. These test cases do not trigger any out-of-bounds operation. Next we describe the test input to every benchmark program tested in the experiment.

In the Bind 1 program, buffer overflow occurs when a negative value is passed as the third argument of memcpy. In [28], a constant string "*sls.lcs.mit.edu*" is hard-coded as the second argument of strcpy to achieve this. We use "*www.cnn.com*" instead as the normal test data under which the program runs normally, and SecTAC can detect this vulnerability. Similarly, for the Bind 2 program, we use string "*cnn.com*" as the normal test input instead of the original hard-coded input "*sls.lcs.mit.edu*" that crashes the program. The Bind 3 program does not check the buffer space when calling memcpy. The provided test case is a file s3.in whose content is "9283721". However, we notice that as long as its content is not "0", the vulnerability always occurs. Thus we just use the original test case. The Bind 4 program uses sprintf without boundary checking. A string of 1072 bytes long is provided in [28] as the input to trigger the vulnerability. We do not use this input; instead, we use a normal test input as given in Table 3.

Most of the vulnerabilities in the Sendmail programs are caused by out-of-bounds pointer operations. These operations are usually in a loop where the pointer is increased by 1 for each iteration. As a result, in the execution trace, the out-of-bounds operation of a pointer only occurs when

a test case can actually trigger the vulnerability. In other words, the execution trace under a normal test case does not contain the vulnerability. Thus, we use test cases provided by the benchmark programs in our experiments.

The test input to each Wu-ftp program is a string that represents a path. For the Wu-ftp 1 program, the original test case in [28] is "*/tmp/*" followed by 24 '*a*'s. This is carefully designed to trigger the buffer overflow caused by strcpy. For the Wu-ftp 2 program, the original test case is also a specific complex path with 9 subdirectories, which triggers the vulnerability caused by strcat. For the Wu-ftp 3 program, the length of the input path is made more than 47 to trigger the vulnerability caused by strcpy. In our experiments, we use normal test inputs. Specifically, for Wu-ftp 1 and Wu-ftp 3, we use a normal input "*/tmp/aa*" that does not trigger the vulnerability. For Wu-ftp 2, we use "*/tmp/test.c*", which is the path of an existing file and does not trigger the vulnerability.

Performance: We did the experiments on a 2GHz Core 2 Desktop running Ubuntu-8.10 Linux operating system. We let the Java use a maximum of 1G heap memory during our experiments. The fifth column of Table 3 shows the execution time of SecTAC for analyzing all traces for each program. The execution time is the sum of the times needed for trace-based symbolic execution and satisfiability analysis, which increases nearly linearly with the trace size in our experiments. We can see that SecTAC can quickly analyze C programs for vulnerability.

New vulnerabilities: In addition to the known bugs, SecTAC also detected six new vulnerabilities in the 14 benchmark programs as shown in Table 4. Test cases that trigger these vulnerabilities can be directly derived

Program	Test Input that Triggers the New Vulnerability	Location of the New Vulnerability	Remarks
Bind 4	www.cnn.com; www.nbc.com	ns-lookup-bad.c:277	*nsp* out-of-bound
Sendmail 1	default	crackaddr-bad.c:460	*buflim* out-of-bound
Sendmail 3	default	mime1-bad.c:212	*infile* out-of-bound
Wu-ftp 2	a 200 bytes long string for argv[1]	call_fb_realpath.c:94	*strcpy* buffer overflow
Wu-ftp 3	/a...a (48 a's)/aa	realpath-2.4.2-bad.c:269	*where* out-of-bound
Wu-ftp 3	/a...a (48 a's)/aa	realpath-2.4.2-bad.c:257	*strcpy* buffer overflow

Table 4. New Vulnerabilities in the benchmark programs

from the solutions given by the satisfiability checker Yices [7] in SecTAC. Notably, we detected a vulnerability in code that was previously considered to be safe. The authors of [28] explicitly commented the line 257 of file realpath-2.4.2-bad.c in the Wu-ftp 3 program as a safe call. However, our experiment shows that it is not. As shown in Table 4, when the length of a directory name is long enough, the strcpy function at line 257 will overflow the destination buffer whose size is only 46 bytes.

For nullhttpd-0.5.1, SecTAC found three buffer overflow vulnerabilities at line 143 of file "http.c". The attacker can overflow three different buffers in this line of code. In addition, it also found a new vulnerability at line 58 of file "config.c", where the program uses snprintf to copy a string variable config.server_base_dir and a constant string "/bin" to buffer server_bin_dir. However, the space allocated to server_bin_dir is 255. If the string length of config.server_base_dir is 255, the buffer is not null terminated and the string "/bin" cannot be copied to the buffer, causing a configuration error. Lines 59 to 61 in the same file have the same vulnerability. For the lancer program, SecTAC found four buffer overflow problems in "handler.c" and "host.c". These problems have the same pattern: the author declared a buffer with the size of n, and used strncpy to copy at most n-1 non-zero characters to the buffer. However, the value at position n-1, which does not belong to this buffer, could be a non-zero value. Thus, it is possible that the string in the buffer is not properly null-terminated, which may cause buffer overflow. For the bftpd-2.3 program, SecTAC detected that the buffer "bu_host" (whose content is from an external input) may be not properly null-terminated. We have reported this vulnerability to the author of the program and a new version was subsequently released to fix this bug.

6   Related Work

The method in [14] detects buffer overflow using existing test cases. They do not perform symbolic analysis and ignore branch conditions, causing many false alarms. The predictive testing in [16] inserts assertions into the source program and uses a combination of concrete and symbolic execution on the given test inputs to discover assertion violations. DART [11, 10] and CUTE [22] can automatically generate test cases. However, they use concrete values for complex constraints that they cannot handle. They may miss many paths that are covered by the test cases carefully designed in traditional testing. SecTAC takes advantage of

previous test effort. In addition, DART and CUTE overlook useful information about variables and functions such as pointer dependency and function return type. SPLAT [26] improves DART by introducing a length attribute in each buffer. It also represents a fixed-length prefix of the buffer elements symbolically. Other buffer elements are represented using concrete values during execution. The limitation is that when the program visits a buffer element beyond the prefix, their symbolic execution becomes *concrete*. SecTAC generates new objects only when a buffer element is visited, which improves the precision and reduces the cost. EXE [6] and KLEE [4] were developed to achieve high branch coverage. They can detect memory overflow vulnerabilities. SAGE [12] also employs trace-based symbolic execution with satisfiability analysis. However, SAGE works on the binary level; a lot of useful information in the source code is no longer available for analysis.

7   Limitations and Suggestions

SecTAC has a number of limitations. First, we must have the test cases ready before doing the security testing. The effectiveness of SecTAC depends on the completeness of the existing test cases. In fact, the branch coverage of the test cases determines the number of paths that our method can check. Second, the size of an execution trace for large complex programs may be huge. Analyzing a large execution trace can cause many problems. For example, it may be the case that a large number of statements in the trace generate security constraints. As a result, SecTAC may invoke the SMT solver very frequently, which can slow down security testing significantly. We plan to improve SecTAC by managing program and security constraints more efficiently, e.g., by using BDDs [2, 3].

8   Conclusion and Future Work

In this paper, we proposed an approach for testing the security of C programs using trace-based symbolic execution and satisfiability analysis. We developed a tool named SecTAC to demonstrate the effectiveness of our approach. We evaluated this tool on 14 benchmark programs and 3 open source programs. The result shows that our tool quickly identified every reported vulnerability in the traces and also found 13 new vulnerabilities. In conclusion, our tool is effective and efficient in testing the security of current software systems.

We are interested in the following directions. First, although our approach can handle multi-threaded programs as long as the test cases are available, it only analyzes a

specific combination of the traces generated by different threads. We propose to identify the trace for each thread and seek effective ways to combine them to improve the detection of security vulnerabilities in multi-threaded programs. Second, we will also seek solutions to further improve the efficiency of `SecTAC` and conduct more experiments on large and complex programs to evaluate our approach.

Acknowledgment

The authors would like to thank the anonymous reviewers for their valuable comments.

References

[1] M. Abadi, M. Budiu, U. Erlingsson, and J. Ligatti. Control-flow integrity. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, pages 340–353, 2005.

[2] S. B. Akers. Binary decision diagrams. *IEEE Transaction on Computers*, C-27(6):509 – 516, June 1978.

[3] R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Transaction on Computers*, C-35(8):677–691, 1986.

[4] C. Cadar, D. Dunbar, and D. Engler. KLEE: Unassisted and automatic generation of high-coverage tests for complex system programs. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2008.

[5] C. Cadar and D. Engler. Execution generated test cases: How to make systems code crash itself. In *Proceedings of the International SPIN Workshop on Model Checking of Software*, 2005.

[6] C. Cadar, V. Ganesh, P. Pawlowski, D. Dill, and D. Engler. EXE: automatically generating inputs of death. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, pages 322–335, 2006.

[7] B. Dutertre and L. de Moura. A fast linear-arithmetic solver for DPLL(T). In *Proceedings of the Computer-Aided Verification Conference (CAV)*, pages 81–94, 2006.

[8] G. Fink, C. Ko, M. Archer, and K. Levitt. Towards a property-based testing environment with applications to security-critical software. In *Proceedings of the 4th Irvine Software Symposium*, pages 39–48, 1994.

[9] A. Ghosh, T. O'Connor, and G. McGraw. An automated approach for identifying potential vulnerabilities in software. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 104–114, 1998.

[10] P. Godefroid. Compositional dynamic test generation. In *Proceedings of the Symposium on Principles of Programming Languages (POPL)*, 2007.

[11] P. Godefroid, N. Klarlund, and K. Sen. DART: directed automated random testing. In *Proceedings of the ACM SIGPLAN conference on Programming Language Design and Implementation*, pages 213–223, 2005.

[12] P. Godefroid, M. Y. Levin, and D. Molnar. Automated whitebox fuzz testing. In *Proceedings of the Network and Distributed Systems Security (NDSS)*, pages 151–166, 2008.

[13] R. Hastings and B. Joyce. Purify: Fast detection of memory leaks and access errors. In *Proceedings of the Winter USENIX Conference*, pages 125–136, 1992.

[14] E. Haugh and M. Bishop. Testing C programs for buffer overflow vulnerabilities. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, pages 123–130, 2003.

[15] R. Jones and P. Kelly. Backwards-compatible bounds checking for arrays and pointers in C programs. In *Proceedings of the International Workshop on Automated Debugging*, 1997.

[16] P. Joshi, K. Sen, and M. Shlimovich. Predictive testing: amplifying the effectiveness of software testing. In *Proceedings of the Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering*, pages 561–564, 2007.

[17] P. Marinescu and G. Candea. LFI: A practical and general library-level fault injector. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN)*, 2009.

[18] G. C. Necula, S. McPeak, S. Rahul, and W. Weimer. CIL: Intermediate language and tools for analysis and transformation of C programs. In *Proceedings of the International Conference on Compiler Construction*, pages 213–228, 2002.

[19] S. Ranise and C. Tinelli. The satisfiability modulo theories library(smt-lib). www.SMT-LIB.org, 2006.

[20] M. Ringenburg and D. Grossman. Preventing format-string attacks via automatic and efficient dynamic checking. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, pages 354–363, 2005.

[21] K. Sen. Concolic testing. In *Proceedings of the IEEE/ACM nternational Conference on Automated Software Engineering (ASE)*, 2007.

[22] K. Sen, D. Marinov, and G. Agha. CUTE: a concolic unit testing engine for C. In *Proceedings of the joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 263–272, 2005.

[23] E. C. Sezer, P. Ning, C. Kil, and J. Xu. Memsherlock: an automated debugger for unknown memory corruption vulnerabilities. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, pages 562–572, 2007.

[24] J. Viega, J. T. Bloch, Y. Kohno, and G. McGraw. ITS4: A static vulnerability scanner for C and C++ code. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*, page 257, 2000.

[25] D. Wagner, J. Foster, E. Brewer, and A. Aiken. A first step towards automated detection of buffer overrun vulnerabilities. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, pages 3–17, 2000.

[26] R. Xu, P. Godefroid, and R. Majumdar. Testing for buffer overflows with length abstractions. In *Proceedings of the International Symposium on Software Testing and Analysis (ISSTA)*, pages 27–38, 2008.

[27] M. Zhivich, T.Leek, and R. Lippmann. Dynamic buffer overflow detection. In *Proceedings of the Workshop on the Evaluation of Software Defect Detection Tools*, 2005.

[28] M. Zitser, R. Lippmann, and T. Leek. Testing static analysis tools using exploitable buffer overflows from open source code. In *Proceedings of the ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 97–106, 2004.