

Towards Denial-of-Service Memory Vulnerabilities

Tianhan Lu¹, Yu-Ju Lee¹, Wen-Wei Liao^{2, 3*}

¹ Department of Computer Science, University of Colorado Boulder, CO 80309-0430, USA.

² University of Colorado Boulder Cooperative Institute for Research in Environmental Sciences, USA.

³ International College of Semiconductor Technology, National Chiao Tung University, Taiwan.

*Corresponding author. Tel.: 303-497-3549; email: Wenwei.Liao@Colorado.Edu

Manuscript submitted May 4, 2019; accepted July 20, 2019.

doi: 10.17706/jsw.14.9.423-436

Abstract: We address the problem of verifying a program to be free of Denial-of-Service memory vulnerabilities. More specifically, we define a program to be safe from DoS attacks if its memory usage at any time during execution is linear to sizes of its inputs. We design an analysis algorithm that verifies if a program satisfies this definition, and reports code snippets in the program that may cause a nonlinear amount of memory usage in case the verification fails. We also formally prove the correctness of our algorithm w.r.t. the above definition. Our experimental results indicate that the analysis algorithm is both effective and efficient.

Key words: Program verification, software security, static analysis.

1. Introduction

Denial-of-Service memory vulnerabilities are an important category of software security bugs, especially for server programs. Successful exploits may exhaust memory resources, leaving the software unavailable to serve requests from benign users. In this paper, we address the problem of either verifying a program to be free of Denial-of-Service vulnerabilities, or labeling which code snippets might cause DoS vulnerabilities. In particular, we concretize the definition of being free from Denial-of-Service memory vulnerabilities into that, at any time during execution, the program uses a linear amount of memory usage w.r.t. sizes of its input variables.

Consider a program model that supports alias, array allocation, collection-typed variables (e.g. lists, maps, sets), mutation, and method invocation. It is challenging for static analysis techniques to produce useful results for such a program while terminate in a short amount of time. Additionally, when a verification procedure fails, it is challenging but desirable for it to return useful information, i.e. which code snippets may cause DoS vulnerabilities.

We propose an analysis algorithm that can both verify non-existence of DoS memory vulnerabilities and report suspicious code snippets back to a user if verification fails. Overall, the contributions of this paper are:

We design an analysis algorithm to verify the linear amount of memory usage of a program, and report a set of code snippets that may cause a nonlinear amount of memory usage when verification fails (Section 4).

We formally prove the analysis algorithm is sound and correct (Section 5) w.r.t. our definition of DoS memory vulnerabilities mentioned above.

The experimental results indicate that our analysis algorithm is both effective and efficient (Section 6). On average, it takes 9.5s for our tool to analyze a 21k lines-of-code program.

2. Motivating Examples

In this section, we present the results from our analysis algorithm over a few example programs that motivated our analysis algorithm. Consider Example 1 in Fig. 1. Our analysis will report verification failure on this program, and report the loop at line 2-5 as a code snippet that potentially increase memory usage for a nonlinear amount. Intuitively, we consider this program as vulnerable to DoS attacks because method argument x is controlled by a user and hence the amount of memory usage is controlled by a user.

Consider Example 2 in Fig. 2. Our analysis will report verification success on this program. Intuitively, the while loop at line 3-6 will be iterated for the length of list l2. Note that if line 5 is replaced by `cons ([1], [2])`, then our analysis will report verification failure, because the amount of memory increment caused by the loop will be nonlinear.

```

1 Unit method(List[Int] l, int x) {
2   while (x > 0)
3     cons(l, x);
4     x = x - 1;
5 }

```

Fig. 1. Example 1.

```

1 Unit method(List[Int] l1, List[Int] l2) {
2   iter = iter(l1);
3   while (hasnxt(iter)) {
4     x = next(iter);
5     cons(l2, x); // cons(l1, l2);
6   }
7 }

```

Fig. 2. Example 2.

Consider Example 3 in Fig. 3. Our analysis will report verification success on this program. Intuitively, although method fibonacci is recursive and the maximum number of times executing the method is not trivial to obtain, this method does not increase memory usage.

Consider Example 4 in Fig. 4. Our analysis will report verification failure on this program. Intuitively, the number of times executing loop at line 6-13 is difficult to obtain (in fact it is exponential to variable number) and the loop allocates an array at line 9. Therefore, the amount of memory increase is not linear to input values.

3. Definition

3.1. Language

In this subsection, we formalize the core language (see Fig. 3). The supported types are Integer type `Int`, Boolean type `Bool`, Unit type `Unit`, Array type `Arr[τ]`, List type `List[τ]`, and Iterator type `Iter[τ]`. In particular, the list type is used to model collection types such as sets, lists and maps. A program p is defined as a main method m_{main} together

with a list of methods \bar{m} . A method m consists of its return type τ , its name M , its list of arguments \bar{x} , and its body statement s . As usual, a statement can be assignment $x = e$, method invocation $x = m(\bar{x})$, array allocation $x = \text{newarr } \tau[e]$, array read $x = x'[e]$, array write $x[e] = e$ sequential composition of statements $s_1; s_2$, disjunction $\text{if}(e)\text{then } s_1\text{else } s_2$, loop $\text{while}(e) s$, return statement `retx` and skip statement `skip`. Moreover, we have statements

modeling collection operations. Statement $x = \text{iter}(x')$ creates an iterator value (which is associated with list x') and assigns it to variable x . Statement $x = \text{next}(x')$ takes the next element from iterator x' and assign it to variable x . Statement $\text{cons}(x, x')$ appends list- or non-list-typed variable x' to list-typed

variable x , thus modeling list concatenation as well as list append. Expressions can be variables, values, arithmetic and comparison operations, negation, and querying if an iterator has reached the end, i.e. $\text{hasnxt}(x)$. Values include integers, booleans, addresses, iterator values $\text{iter}(n, a)$ (where n denotes the current index of iterator and a denotes the address of the corresponding list value), array values $[v, \dots, v]$ and list values $\langle v, \dots, v \rangle$.

3.2. Operational Semantics

In this subsection, we formalize the small-step operational semantics (see Fig. 6) of the language defined in Section 3.1. The operational semantics defined here formally describes how an abstract machine executes a program written in the language defined in Section 3.1. Additionally, they will be used in Section 5 to prove that the analysis defined in Section 4 is correct, i.e. Theorem 2. Intuitively, the operational semantics defined in Fig. 6 respects the semantics of a Java or C++ program, but without the Object-Oriented feature. Scalar values (i.e. integer, booleans, addresses and iterators) reside in the stack and other values (i.e. array and list) reside in the heap.

```

1 Unit fibonacci(int number) {
2   if(number == 1 || number == 2)
3     return 1;
4   return fibonacci(number-1) + fibonacci(number -2); // recursion
5 }

```

Fig. 3. Example 3.

```

1 Unit fibonacci(int number, List[Int] l) {
2   if (number == 1 || number == 2)
3     return 1;
4   fibo1=1; fibo2=1; fibo=1;
5   i= 3;
6   while (i<= number) {
7     //Fibonacci number is sum of previous two Fibonacci number
8     fibo = fibo1 + fibo2;
9     arr = newarr Int[1];
10    arr[0] = fibo;
11    cons(1, arr);
12    fibo1 = fibo2;
13    fibo2 = fibo;
14    i = i + 1;
15  }
16 }

```

Fig. 4. Example 4.

A program state σ is defined as pair of stack frames and a heap $[\bar{\kappa}, \rho]$. Notation $\bar{\kappa}$ denotes a list of stack frames $\langle \kappa_1, \dots, \kappa_n \rangle$ where each stack frame is a mapping from variables to values, i.e. $\kappa \stackrel{\text{def}}{=} [x \mapsto v]$. Heap ρ is a mapping from addresses to values $\rho \stackrel{\text{def}}{=} [a \mapsto v]$. We sometimes denote list $\langle \kappa_1, \dots, \kappa_n \rangle$ as $\kappa^h :: \kappa^t$ where κ^h denotes the head of list (i.e. κ_1) and κ^t denotes the tail (i.e. $\kappa_2, \dots, \kappa_n$). A program's initial state σ_0 is defined as $[\kappa_0, \rho_0]$. Initial stack κ_0 is defined as $[x^0 \mapsto v^0]$ where x^0 denotes input variables and v^0 denotes input values. Initial store ρ_0 is defined as $[a^0 \mapsto v^0]$ where a^0 denotes addresses that point to input values v^0 . Note that we assume the program is already type-checked before being executed. As a result, in the rules below (see Fig. 6) each variable is typed and their types are available during execution.

To simplify the representation of rules in Fig. 6, we define a notation $\sigma[x \mapsto v]$ to denote value mutation that happens at the top frame κ^h of a program state $\sigma \stackrel{\text{def}}{=} (\kappa^h :: \kappa^t, \rho)$, i.e. $\sigma[x \mapsto v] \stackrel{\text{def}}{=} (\kappa^h[x \mapsto v] :: \kappa^t, \rho)$. Notation $\kappa[x \mapsto v]$ denotes a stack that is the same as stack κ , except that variable x is mapped to value

v . We also define a similar notation $\rho[a \mapsto v]$ to denote a heap that is the same as heap ρ , except that address a is mapped to value v . We use notation $x : \tau$ to denote that variable x 's type is τ .

Rule E-Method creates a new stack frame before invoking the method by copying values from the current stack frame. The return value is the only value that is stored in the top stack frame after method invocation, which is guaranteed by Rule Ret. Rule E-NewArray allocates a new chunk of space in the store. Rule E-ArrayW writes a value into an array value in the store. Rule E-ArrayR reads a value from an array value in the store. Rule E-Iter creates an iterator value that is associated with a list. Rule E-Next retrieves a value from a list value in the store. Rule E-ConsAppend appends a value into a list value in the store, based on the types of two arguments. Rule E-ConsConcat concatenates a list value into another list value in the store, based on the types of two arguments. The remaining rules are standard and hence not further explained. The small-step operational semantic rules for evaluating expressions are standard and hence omitted.

3.3. Size Semantics

In this subsection, we define the memory usage $|v|_\rho$ of a value w.r.t. a heap ρ , i.e. the *reachable size* of value v in the heap ρ , as well as the memory usage $|\sigma|$ of a program state (see Fig. 7). Intuitively, the size of a list or array value is the size of the reachable heap starting from the value. Similarly, the size of a state is the size of the reachable heap starting from list and array values stored in the stack. Note that we do not take into consideration the sizes of scalar values (i.e. integers, booleans, and iterator values), because our observation is that memory usage can only be significantly increased by list values and array values. Notation $\text{Dom}(\kappa)$ denotes the domain of mapping κ and $\text{Dom}(\rho)$ denotes the domain of mapping ρ .

3.4. Problem Definition

Formally, the problem that we address in this paper is defined as follows. Given a program p and an initial state $\sigma_0 = [\kappa_0, \rho_0]$ (which instantiated with input values v_1^0, \dots, v_n^0), we either verify that the program's memory usage at any

point of execution is linear to the total size of input values (see Fig. 8), or if the verification fails, output a set of statements that may cause the constraint to be violated. Note that initial state σ_0 trivially satisfies the constraint in Fig. 8 by setting all q_i to 1.

| | | | |
|------------|-----|-----|--|
| Program | p | ::= | $m_{main}; \bar{m}$ |
| Method | m | ::= | $\tau M(\bar{x})\{s\}$ |
| Statements | s | ::= | $x = e \mid x = M(\bar{x}) \mid x = \text{newarr } \tau[e] \mid x = x'[e] \mid x[e] = e'$ $\mid x = \text{iter}(x') \mid x = \text{next}(x') \mid \text{cons}(x, x')$ $\mid s_1; s_2 \mid \text{if}(e)\text{then } s_1 \text{ else } s_2 \mid \text{while}(e) s \mid \text{ret } x \mid \text{skip}$ |
| Expression | e | ::= | $x \mid v \mid e_1 \oplus e_2 \mid e_1 \bowtie e_2 \mid e_1 \vee e_2 \mid \neg e \mid \text{hasnxt}(x)$ |
| Value | v | ::= | $n \in \mathbb{Z} \mid b \in \{\text{true}, \text{false}\} \mid a \in \text{Addr}$ $\mid \text{iter}(n^{\in \mathbb{N}}, a^{\in \text{Addr}}) \mid [v, \dots, v] \mid \langle v, \dots, v \rangle$ |

(a) The core calculus. Operator \oplus and \bowtie stand for arithmetic and comparison operations. M denotes a string.

| | | | |
|-------|--------|-----|--|
| Types | τ | ::= | $\text{Int} \mid \text{Bool} \mid \text{Unit} \mid \text{Arr}[\tau] \mid \text{List}[\tau] \mid \text{Iter}[\tau]$ |
|-------|--------|-----|--|

(b) Types.

Fig. 5. Language definition.

4. Algorithm

This section presents the analysis algorithm in the form of rules (see Fig. 10). The analysis takes as input a program and apply Rule T-Program, which will apply other rules. Specifically, judgement form $\vdash s : (\Delta, \Theta)$ means that the algorithm takes as input statement s , and then outputs a set of statements Δ , such that executing any of them may cause a program state not to preserve its memory usage as a linear combination

of input values, as well as the type of memory usage Θ of statement s . Note that for all rules in Fig. 10, it is assumed that before applying a rule, the precondition is that we have Theorem 1 hold true. In other words, each rule infers the output Δ and Θ only if Theorem 1 holds before applying the rule. Also note that the analysis algorithm is parametrized by two procedures ConstProp and LoopBound. Procedure ConstProp takes as input a program and outputs the set of expressions that must be constant for any input values. Procedure LoopBound takes as input a loop and outputs the symbolic loop bound (i.e. the maximum number of times executing the loop body).

We define the type of memory usage Θ of a statement s in Fig. 9a. If a statement s is of type Θ_{NoInc} , then if starting from any program state, executing the statement will not increase the size of the program state. If a statement s is of type $\Theta_{IncConst}$, then if starting from any program state, executing the statement will increase the size of the program state by a constant. Note that we define type Θ_{NoInc} as a subtype of type $\Theta_{IncConst}$ in Fig. 9b, because when q is 0, a statement satisfying the definition of type $\Theta_{IncConst}$ also satisfies the definition of type Θ_{NoInc} . If a statement s is of type $\Theta_{IncConst}$, then if starting from any program state, executing the statement will increase the size of the program state by a linear combination of input values, together with a constant. Note that we define type $\Theta_{IncConst}$ as a subtype of type Θ_{IncLin} in Fig. 9b, because when all q_i are 0, a statement satisfying the definition of type Θ_{IncLin} also satisfies the definition of type $\Theta_{IncConst}$. If a statement does not satisfy any of the above three

definitions about the type of memory usage, we define it to be of type $\Theta_{IncArbitrary}$, denoting that executing the statement may arbitrarily increase the memory usage of a program.

Next we explain the rules. Notation $|x|$ denotes the size of a value in a concrete heap, where the value is mapped from variable x in a concrete stack. We use notation $\text{Dom}(p)$ to denote the set of variables defined in all methods of program p .

Rule T-Next, T-Iter, T-Assign, T-ArrayR, T-Ret trivially do not increase memory usage of a program state. Rule T-ConsAppend can at most increase memory usage by a linear amount, because it is assumed that variable x' will use a linear amount of memory (recall that it is assumed Theorem 1 holds true before applying each rule). Similarly, Rule T-ConsConcat can at most increase memory usage by a linear amount. Rule T-NewArrayConst applies when procedure ConstProp infers that expression e is always a constant for any input values. As a result, memory increase is a constant. Rule T-NewArrayConst applies when procedure ConstProp cannot guarantee that expression e is always a constant for any input values. As a result, we conservatively assert that the amount of memory increase is arbitrary. Rule T-ArrayW increases memory usage by a linear amount, because 1) if expression e' is a scalar value, then sizes of scalar values are constant (see Fig. 7), 2). if expression e' is a non-scalar value, then according to the assumption Theorem 1, the memory increase will be a linear amount.

Rule T-WhileLin1 applies when procedure LoopBound infers that the loop's bound is a linear expression using any variable defined in the program and when the memory usage type of the loop body is a subtype of type $\Theta_{IncConst}$, i.e.

the amount of memory increment caused by executing the loop body once is at most a constant. Rule T-WhileLin2 applies when procedure LoopBound infers that the loop's bound is a linear expression using any variable defined in the program and when the memory usage type of the loop body is a supertype of type Θ_{IncLin} . Rule T-WhileNonLin1 applies

when procedure LoopBound cannot infer that the loop's bound is a linear expression using any variable defined in the program and when the memory usage type of the loop body is a subtype of type Θ_{NoInc} . Rule T-WhileNonLin2 applies

when procedure LoopBound cannot infer that the loop's bound is a linear expression using any variable defined in the program and when the memory usage type of the loop body is a supertype of type $\Theta_{IncConst}$.

Rule T-Method applies when the invoked method M is not a recursive function. Intuitively, assigning the return value to variable can at most increase memory usage (when compared with the memory usage before invocation) by a same amount as executing method M . Rule T-MethodRec applies when the invoked method M is a (mutually) recursive function. Intuitively, we first collect the set of all possible invocation relations starting from method M , i.e. which method may invoke which method. Then we define a constraint that enforces that, if method M_i invokes M_{i+1} , then

- The set of statements (that possibly increase memory usage nonlinearly) inferred from method M_{i+1} must be included in the set inferred from method M_i , because if a statement from method M_{i+1} may cause
- vulnerabilities then method M_i should also report it (because method M_i may invoke method M_{i+1}). The memory usage type of method M_i should be at most as restrictive as the type of method M_{i+1} (because method M_i may invoke method M_{i+1}).

$$\frac{\text{E-WhileT} \quad \langle \sigma, e \rangle \rightsquigarrow \text{true} \quad \langle \sigma, s \rangle \rightsquigarrow \langle \sigma', s' \rangle}{\langle \sigma, \text{while}(e) s \rangle \rightsquigarrow \langle \sigma', s'; \text{while}(e) s \rangle} \quad \frac{\text{E-WhileF} \quad \langle \sigma, e \rangle \rightsquigarrow \text{false}}{\langle \sigma, \text{while}(e) s \rangle \rightsquigarrow \langle \sigma, \text{skip} \rangle} \quad \frac{\text{E-Assign} \quad \langle \sigma, e \rangle \rightsquigarrow v}{\langle \sigma, x = e \rangle \rightsquigarrow \langle \sigma[x \mapsto v], \text{skip} \rangle}$$

E-Method

$$\frac{\frac{\bar{x} = x_1, \dots, x_n \quad \sigma = [\overbrace{\kappa^h :: \kappa^t, \rho}^{\bar{\kappa}}] \quad \forall i \in [1, n]. \kappa^h[x_i] = v_i \quad \tau M(\bar{x})\{S\}}{\kappa_m = [x'_i \mapsto v_i]_{i \in [1, n]} \quad \text{where } x'_i \text{ is a fresh variable}}}{\langle [\kappa_m :: \bar{\kappa}, \rho], S \rangle \rightsquigarrow^* \langle \sigma', \text{skip} \rangle} \quad \sigma' = [\overbrace{x_r \mapsto v_r}^{\bar{\kappa}}, \rho']}{\langle \sigma, x = M(\bar{x}) \rangle \rightsquigarrow \langle [\bar{\kappa}[x \mapsto v_r], \rho'], \text{skip} \rangle}$$

$$\frac{\text{E- E-Ret} \quad \sigma = [\kappa^h :: \kappa^t, \rho] \quad v = \kappa^h[x] \quad \sigma' = [\overbrace{x \mapsto v}^{\kappa^t}, \rho]}{\langle \sigma, \text{ret } x \rangle \rightsquigarrow \langle \sigma', \text{skip} \rangle} \quad \frac{\text{E-SeqSkip} \quad \langle \sigma, s_1 \rangle \rightsquigarrow \langle \sigma', \text{skip} \rangle}{\langle \sigma, s_1; s_2 \rangle \rightsquigarrow \langle \sigma', s_2 \rangle} \quad \frac{\text{E-Seq} \quad \langle \sigma, s_1 \rangle \rightsquigarrow \langle \sigma', \text{skip} \rangle \quad s'_1 \neq \text{skip}}{\langle \sigma, s_1; s_2 \rangle \rightsquigarrow \langle \sigma', s'_1; s_2 \rangle}$$

$$\frac{\text{E-IfT} \quad \langle \sigma, e \rangle \rightsquigarrow \text{true}}{\langle \sigma, \text{if}(e) \text{ then } s_1 \text{ else } s_2 \rangle \rightsquigarrow \langle \sigma, s_1 \rangle} \quad \frac{\text{E-IfF} \quad \langle \sigma, e \rangle \rightsquigarrow \text{false}}{\langle \sigma, \text{if}(e) \text{ then } s_1 \text{ else } s_2 \rangle \rightsquigarrow \langle \sigma, s_2 \rangle}$$

E-NewArray

$$\frac{\langle \sigma, e \rangle \rightsquigarrow v \in \mathbb{N} \quad \rho' = \rho[a \mapsto \overbrace{[0, \dots, 0]}^{v \text{ elements}}]}{a \text{ is a fresh address}}}{\langle \sigma, x = \text{newarr } \tau[e] \rangle \rightsquigarrow \langle [\sigma[x \mapsto a], \rho'] , \text{skip} \rangle}$$

E-ArrayW

$$\frac{\rho' = [\kappa, \rho] \quad \kappa[x] = a \quad \rho[a] = [v_1, \dots, v_i, \dots, v_n] \quad \langle \sigma, e \rangle \rightsquigarrow i \quad (1 \leq i \leq n) \wedge (i \in \mathbb{N}) \quad \langle \sigma, e' \rangle \rightsquigarrow v'_i}{\rho' = \rho[a \mapsto [v_1, \dots, v'_i, \dots, v_n]]}{\langle \sigma, x[e] = e' \rangle \rightsquigarrow \langle [\kappa, \rho'] , \text{skip} \rangle}$$

E-ArrayR

$$\frac{\sigma = [\kappa, \rho] \quad \rho[\kappa[x]] = [v_1, \dots, v_i, \dots, v_n] \quad \langle \sigma, e \rangle \rightsquigarrow i \quad 1 \leq i \leq n \wedge i \in \mathbb{N}}{\langle \sigma, x = x'[e] \rangle \rightsquigarrow \langle \sigma[x \mapsto v_i], \text{skip} \rangle}$$

E-Iter

$$\frac{\sigma = [\kappa, \rho]}{\langle \sigma, x = \text{iter}(x') \rangle \rightsquigarrow \langle \sigma[x \mapsto \text{iter}(0, \kappa[x'])], \text{skip} \rangle}$$

E-Next

$$\frac{\sigma = [\kappa, \rho] \quad \kappa[x] = \text{iter}(i, a) \quad \rho[a] = \langle v_1, \dots, v_n \rangle \quad 0 \leq i < n}{\langle \sigma, x = \text{next}(x') \rangle \rightsquigarrow \langle \sigma[x \mapsto v_{i+1}], \text{skip} \rangle}$$

E-ConsAppend

$$\frac{\sigma = [\kappa, \rho] \quad x': \tau \quad x: \text{List}[\tau] \quad \kappa[x] = a \quad \rho[a] = \langle v_1, \dots, v_n \rangle \quad \rho' = \rho[a \mapsto \langle v_1, \dots, v_n, \kappa[x'] \rangle]}{\langle \sigma, \text{cons}(x, x') \rangle \rightsquigarrow \langle [\kappa, \rho'] , \text{skip} \rangle}$$

E- ConsConcat

$$\frac{\sigma = [\kappa, \rho] \quad x': \text{List}[\tau] \quad x: \text{List}[\tau] \quad \kappa[x] = a \quad \rho[a] = \langle v_1, \dots, v_n \rangle \quad \kappa[x'] = a' \quad \rho[a'] = \langle v'_1, \dots, v'_m \rangle \quad \rho' = \rho[a \mapsto \langle v_1, \dots, v_n, v'_1, \dots, v'_m \rangle]}{\langle \sigma, \text{cons}(x, x') \rangle \rightsquigarrow \langle [\kappa, \rho'] , \text{skip} \rangle}$$

E- HasNext

$$\frac{\sigma = [\kappa, \rho] \quad \kappa[x] = \text{iter}(i, a) \quad \rho[a] = \langle v_1, \dots, v_n \rangle}{\langle \sigma, \text{hasnxt}(x) \rangle \rightsquigarrow i < n}$$

Fig. 6. Small-step operational semantics.

In the end, we try to obtain a least fix point solution to the above constraint and use the solution as the output of this rule. A least fix point preserves analysis precision as much as possible.

Rule T-Seq and T-If first infer types of memory usage and sets of statements (that possibly increase memory usage nonlinearly) for each sub-statement, and then join the types (i.e. choose the least restrictive type to guarantee soundness) and union the sets.

$$\begin{aligned}
 |n|_\rho &\stackrel{\text{def}}{=} 0 \\
 |b|_\rho &\stackrel{\text{def}}{=} 0 \\
 |a|_\rho &\stackrel{\text{def}}{=} |\rho[a]|_\rho \\
 |\text{iter}(n, a)|_\rho &\stackrel{\text{def}}{=} 0 \\
 |[v_1, \dots, v_n]|_\rho &\stackrel{\text{def}}{=} n + \sum_{i=0}^n |v_i|_\rho \\
 |\langle v_1, \dots, v_n \rangle|_\rho &\stackrel{\text{def}}{=} n + \sum_{i=0}^n |v_i|_\rho \\
 |\sigma| &\stackrel{\text{def}}{=} n + \sum_{x_i \in \text{Dom}(\kappa)} (|\kappa[x_i]|_\rho) \text{ where } \sigma = [\kappa, \rho]
 \end{aligned}$$

Fig. 7. Size semantics.

$$\forall \sigma. \text{If } \langle \sigma_0, p \rangle \rightsquigarrow^* \langle \sigma, s \rangle, \text{ then } |\sigma| = n + \sum_{x_i \in \text{Dom}(\kappa_0)} (q_i \cdot |\kappa_0[x_i]|_{\rho_0}) \text{ where } q_i \text{ is a constant}$$

Fig. 8. (Problem definition) Linear memory usage

$$\begin{aligned}
 \Theta &\stackrel{\text{def}}{=} \Theta_{\text{NoInc}} \mid \Theta_{\text{IncConst}} \mid \Theta_{\text{IncLin}} \mid \Theta_{\text{IncArbitrary}} \\
 \text{S: } \Theta_{\text{NoInc}} &\stackrel{\text{def}}{=} \forall \sigma. \text{If } \langle \sigma, s \rangle \rightsquigarrow \langle \sigma', s' \rangle, \text{ then } |\sigma'| - |\sigma| = 0 \\
 \text{S: } \Theta_{\text{IncConst}} &\stackrel{\text{def}}{=} \forall \sigma. \text{If } \langle \sigma, s \rangle \rightsquigarrow \langle \sigma', s' \rangle, \text{ then } |\sigma'| - |\sigma| = q \text{ where } q \text{ is a constant} \\
 \text{S: } \Theta_{\text{IncLin}} &\stackrel{\text{def}}{=} \forall \sigma. \text{If } \langle \sigma, s \rangle \rightsquigarrow \langle \sigma', s' \rangle, \text{ then } |\sigma'| - |\sigma| = \left(\sum_{x_i \in \kappa_0} (q_i \cdot |\kappa_0[x_i]|_{\rho_0}) \right) + q_0 \\
 &\text{where } q_i \text{ is a constant and } \sigma = [\kappa_0, \rho_0] \text{ is the initial program state} \\
 \text{S: } \Theta_{\text{IncArbitrary}} &\stackrel{\text{def}}{=} \text{Otherwise.}
 \end{aligned}$$

(a) Types.

$$\Theta_{\text{NoInc}} <: \Theta_{\text{IncConst}} <: \Theta_{\text{IncLin}} <: \Theta_{\text{IncArbitrary}}$$

(b) Type lattice.

Fig. 9. Types of memory usage of a statement

5. Soundness

In this section, we present the soundness theorems that establish the correctness of the analysis algorithm defined in Section 4 w.r.t. the problem definition in Section 3.4 and the operational semantic defined in Fig. 6, and formally prove them.

Theorem 1 (Linearity). *Given program p , if $\vdash p: (\emptyset, \Theta)$ and the program starts with initial state σ_0 , then*

$$\forall \sigma. \text{If } \langle \sigma_0, p \rangle \rightsquigarrow \langle \sigma, s \rangle, \text{ then } \forall \kappa_i^{i \in [1, n]} \forall x \in \kappa_i \cdot |\kappa[x]|_\rho = \sum_{x_i \in \text{Dom}(\kappa_0)} (q_i \cdot |\kappa_0[x_i]|_{\rho_0})$$

where $\sigma = [\langle \kappa_1, \dots, \kappa_n \rangle, \rho]$ and q_i is a constant

Proof. We prove this theorem by a mathematical induction on each rule defined in the algorithm (see Figure 10).

Base case. It is trivial that the above theorem holds for any initial state σ_0 , because the left hand side and right hand side are syntactically the same.

Inductive cases. Next we prove the above theorem holds true for any state that is reachable from the initial state σ_0 .

The inductive assumption is that the above theorem is valid before applying each rule.

- Rule WhileLin1: Since the memory usage type of the loop body is a subtype of type Θ_{IncConst} , the amount of memory increment for all variables is at most a constant. Since the loop bound is a linear expression using sizes of program variables (which are linear expressions of input values according to the inductive assumption) and the total amount of memory increment of each variable is a product of the above two, the total memory usage increment for each variable is at most a linear amount.
- Rule WhileLin2: Since the memory usage type of the loop body is a supertype of type Θ_{IncLin} , the amount of memory increment for all variables can be at least a linear expression of variable sizes. Since the loop bound is a linear expression using sizes of program variables (which are linear expressions of input values according to the inductive assumption) and the total amount of memory increment of each variable is a product of the above two, the total memory usage increment for each variable can be arbitrary.
- Rule WhileNonLin1: Since the memory usage type of the loop body is a subtype of type Θ_{NoInc} , the amount of memory increment for all variables is 0. As a result, the total amount of memory increment for all variables is 0.
- Rule WhileNonLin2: Since the loop bound can be arbitrary, the total amount of memory increment for all variables can also be arbitrary.
- Rule T-Method: Any memory increment caused by invoking method M may escape to the caller, which is exhibited by outputting the memory usage type of statement s .

$$\begin{array}{c}
 \text{T-Next} \qquad \qquad \qquad \text{T-Iter} \qquad \qquad \qquad \text{T-ConsAppend} \qquad \qquad \qquad \text{T-ConsConcat} \\
 \hline
 \vdash x = \text{next}(x') : (\emptyset, \Theta_{\text{NoInc}}) \quad \vdash x = \text{iter}(x') : (\emptyset, \Theta_{\text{NoInc}}) \quad \vdash \text{cons}(x, x') : (\emptyset, \Theta_{\text{IncLin}}) \quad \vdash \text{cons}(x, x') : (\emptyset, \Theta_{\text{IncLin}}) \\
 \\
 \text{T-Assign} \qquad \qquad \qquad \text{T-NewArrayConst} \qquad \qquad \qquad \text{T-NewArray} \\
 \hline
 \vdash x = e : (\emptyset, \Theta_{\text{NoInc}}) \quad \frac{\text{ConstProp}(p) \text{ infers } e \text{ is a constant}}{\vdash x = \text{newarr } \tau[e] : (\emptyset, \Theta_{\text{IncConst}})} \quad \frac{\text{ConstProp}(p) \text{ cannot infers } e \text{ is a constant}}{\vdash x = \text{newarr } \tau[e] : (\{x = \text{newarr } \tau[e]\}, \Theta_{\text{IncArbitrary}})} \\
 \\
 \text{T-ArrayR} \qquad \qquad \qquad \text{T-ArrayW} \qquad \qquad \qquad \text{T-If} \\
 \hline
 \vdash x = x'[e] : (\emptyset, \Theta_{\text{NoInc}}) \quad \vdash x = x'[e] : (\emptyset, \Theta_{\text{IncLin}}) \quad \frac{\vdash s_1 : (\Delta_1, \Theta_1) \quad \vdash s_2 : (\Delta_2, \Theta_2)}{\vdash \text{if}(e) \text{ then } s_1 \text{ else } s_2 : (\Delta_1 \cup \Delta_2, \Theta_1 \vee \Theta_2)}
 \end{array}$$

T- WhileLin1

$$\frac{\text{LoopBound}(\text{while}(e) s) \text{ infers the loop bound is } \left(\sum q_i \cdot |x_i| + q_0 \right) \text{ where } x_i \in \text{Dom}(p) \text{ and } q_0, q_1 \text{ are constants}}{\vdash s : (\Delta, \Theta) \quad \Theta <: \Theta_{\text{IncConst}}} \quad \hline
 \vdash \text{while}(e) s : (\emptyset, \Theta_{\text{IncLin}})$$

T- WhileLin2

$$\frac{\text{LoopBound}(\text{while}(e) s) \text{ infers the loop bound is } \left(\sum q_i \cdot |x_i| + q_0 \right) \text{ where } x_i \in \text{Dom}(p) \text{ and } q_0, q_1 \text{ are constants}}{\begin{array}{c} \vdash s: (\Delta, \Theta) \quad \Theta_{\text{IncLin}} <: \Theta \\ \hline \vdash \text{while}(e) s: (\{\text{while}(e) s\}, \Theta_{\text{IncArbitrary}}) \end{array}}$$

T- WhileNonLin1

$$\frac{\text{LoopBound}(\text{while}(e) s) \text{ cannot infer the loop bound is } \left(\sum q_i \cdot |x_i| + q_0 \right) \text{ where } x_i \in \text{Dom}(p) \text{ and } q_0, q_1 \text{ are constants}}{\begin{array}{c} \vdash s: (\Delta, \Theta) \quad \Theta <: \Theta_{\text{NoInc}} \\ \hline \vdash \text{while}(e) s: (\emptyset, \Theta_{\text{NoInc}}) \end{array}}$$

T- WhileNonLin2

$$\frac{\text{LoopBound}(\text{while}(e) s) \text{ cannot infer the loop bound is } \left(\sum q_i \cdot |x_i| + q_0 \right) \text{ where } x_i \in \text{Dom}(p) \text{ and } q_0, q_1 \text{ are constants}}{\begin{array}{c} \vdash s: (\Delta, \Theta) \quad \Theta_{\text{IncConst}} <: \Theta \\ \hline \vdash \text{while}(e) s: (\{\text{while}(e) s\}, \Theta_{\text{IncArbitrary}}) \end{array}}$$

T-MethodRec

$$\frac{\begin{array}{c} \tau M(\bar{x})\{s\} \\ \text{Method } M \text{ is (mutually)recursive. Let the set of invocation relations be } \{M_i \text{ invokes } M_{i+1}\} \text{ where } M_0 = M \\ \tau M_i(\bar{x})\{s_i\} \quad \vdash s_i: (\Delta_i, \Theta_i) \quad C = (\forall \Delta_i, \Delta_{i+1} \subseteq \Delta_i) \wedge (\forall \Theta_i, \Theta_{i+1} <: \Theta_i) \\ \text{Let the least fix point solution to constraint } C \text{ be } \Delta_{lfp} = \Delta_0 \dots \text{ and } \Theta_{lfp} = \Theta_0 \dots \end{array}}{\vdash x = M(\bar{x}): (\Delta_0, \Theta_0)}$$

T-Method

$$\frac{\tau M(\bar{x})\{s\} \quad \text{Method } M \text{ is not recursive} \quad \vdash s: (\Delta, \Theta)}{\vdash x = M(\bar{x}): (\Delta, \Theta)}$$

T-Ret

$$\frac{}{\vdash \text{ret } x: (\emptyset, \Theta_{\text{NoInc}})}$$

T-Seq

$$\frac{\begin{array}{c} \vdash s_1: (\Delta_1, \Theta_1) \quad \vdash s_2: (\Delta_2, \Theta_2) \\ \hline \vdash s_1; s_2: (\Delta_1 \cup \Delta_2, \Theta_1 \vee \Theta_2) \end{array}}{\vdash s_1; s_2: (\Delta_1 \cup \Delta_2, \Theta_1 \vee \Theta_2)}$$

T-Program

$$\frac{\begin{array}{c} p = m_{\text{main}}; \bar{m} \quad m_{\text{main}} = M \quad \tau M(\bar{x})\{s\} \\ \hline \vdash s: (\Delta, \Theta) \end{array}}{\vdash p: (\Delta, \Theta)}$$

Fig. 10. Analysis algorithm.

- Rule T-MethodRec: Constraint C guarantees that if method M_i invokes M_{i+1} , then the amount of memory increment for statement s_i is at least as much as for statement s_{i+1} , thus considering the possibility that memory increment during invoking method M_{i+1} might escape to its caller method M_i . The least fix point solution preserves as much precision as possible.
- Rule T-Next, T-Iter, T-Ret, T-Assign, T-ArrayR: According to size semantics defined in Figure 7, there is no memory increment.
- Rule T-ConsAppend: According to the inductive assumption, the amount of memory usage of variable x' is at most a linear expression of input values. Hence the amount of memory usage variable x shall remain a linear expression of input values.
- Rule T-ConsConcat: According to the inductive assumption, the amount of memory usage of variable x' is at most a linear expression of input values. Hence the amount of memory usage of variable x shall remain a linear expression of input values.
- Rule T-NewArrayConst: Since procedure ConstProp infers that expression e is always a constant for any input values, the amount of memory increment is a constant.
- Rule T-NewArray: Since procedure ConstProp cannot infer that expression e is always a constant for any input values, we conservatively assume that the amount of memory increment can be

- arbitrary.
- Rule T-ArrayW: If expression e' is scalar-typed (i.e. integer, boolean, and iterator), then the amount of memory increase is 0.
Otherwise if expression e' is not scalar-typed (i.e. list, array), then according to the inductive assumption, the amount of memory usage of variable e' is at most a linear expression of input values. Hence the amount of memory usage of variable x shall remain a linear expression of input values, soundly disregarding the memory decrement by removing the pointer to value v_i (see Rule E-ArrayW)
 - Rule T-Seq: The join of memory usage types Θ_1 and Θ_2 returns the least restrictive type, i.e. a more conservative memory usage type among types Θ_1 and Θ_2 that over-approximates the amount of memory increment.
 - Rule T-If: The join of memory usage types Θ_1 and Θ_2 returns the least restrictive type, i.e. a more conservative memory usage type among types Θ_1 and Θ_2 that over-approximates the amount of memory increment.

Theorem 2 (Soundness). *Given program p , if $\vdash p: (\emptyset, \Theta)$ and the program starts with initial state σ_0 , then*

$$\forall \sigma. \text{If } \langle \sigma_0, p \rangle \rightsquigarrow \langle \sigma, s \rangle, \text{ then } |\sigma| = \sum_{x_i \in \text{Dom}(\kappa_n)} (q_i \cdot |\kappa_0[x_i]|_{\rho_0}) \text{ where } q_i \text{ is a constant}$$

Proof. The proof can be directly obtained after proving Theorem 1 and using the definition of $|\sigma|$ in Fig. 7.

6. Experiment

6.1. Implementation

We implemented the algorithm defined in Fig. 10 using T.J. Watson Libraries for Analysis (WALA)¹, targeting at Java programs. Since the Object-Oriented feature of Java is not modeled in the core calculus defined in Fig. 6 and hence not considered in the algorithm, the experimental results should be interpreted as sound and correct without considering the Object-Oriented feature.

In the implementation, procedure ConstProp adopts a simple strategy, i.e. if an expression is itself a constant, then return true and otherwise return false. Procedure LoopBound tries to syntactically identify loops that iterate over a collection- or array-typed variable. The solver used in Rule T-MethodRec makes a most conservative estimation, i.e. as long as a method may be recursive and may increase memory usage, consider it as being able to increase memory usage arbitrarily.

6.2. Experimental Results

The set of Java programs come from Space and Time Analysis for Cybersecurity (STAC). 3 of the programs contain memory exhaustion vulnerabilities and others do not. The experiments were run on a MacBook pro'15, 2017. The machine has a 4-core 6th generation Intel i7 CPU, 16GB memory and 1TB SSD. Running time was measured by the Real time (wall clock time) in the output of time command.

To evaluate our technique, we propose the following two research questions:

Effectiveness. How effective is our technique in verifying linear amount of memory usage or reporting code snippets that may cause Denial-of-Service vulnerabilities? As shown in the second column Figure 11, our technique is able to verify that only 11.41% of all loops in the benchmarks might cause nonlinear amount of memory usage. Moreover, as shown in the third column, only 0.72% of all method are recursive methods and might cause nonlinear amount of memory usage. As a result, only these reported loops and

¹ http://wala.sourceforge.net/wiki/index.php/Main_Page

methods need a secure analysis to manually look into, thus excluding the vast majority of the code.

Efficiency. How fast is our verification algorithm? As shown in Fig. 12, on average it took 9.5 seconds for our tool to analyze a 21k lines-of-code program, which we consider as indicating that our algorithm can scale to real life programs.

| Benchmark name | Number of vulnerabilities | Vulnerable loops | Recursive methods |
|------------------|---------------------------|--------------------|--------------------|
| Airplan 5 | 1 | 35/299 (11.17%) | 2/1848 (0.11%) |
| InfoTrader | 0 | 11/133 (8.27%) | 12/1041 (1.15%) |
| Linear Algebra | 1 | 14/81 (17.28%) | 0/221 (0.00%) |
| Malware Analyzer | 0 | 2/15 (13.33%) | 4/79 (5.06%) |
| Powerbroker 1 | 0 | 8/129 (6.20%) | 15/1532 (0.97%) |
| Powerbroker 4 | 0 | 8/115 (6.96%) | 14/1493 (0.94%) |
| Tweeter | 1 | 16/124 (12.90%) | 3/566 (0.53%) |
| Withmi 3 | 0 | 14/96 (14.58%) | 8/1294 (0.62%) |
| Withmi 4 | 0 | 16/95 (16.84%) | 9/1258 (0.71%) |
| Average | | (11.17%) | (0.11%) |

Fig. 11. Effectiveness: Suspicious code snippets. Column “Vulnerable loops” present the percentage of loops in all loops that may cause nonlinear amount of memory increment. Column “Recursive methods” show the percentage of recursive methods in all methods that allocates memory, which are all considered to may increase the memory usage with a nonlinear amount (see Section 6.1) for detail).

| Benchmark name | Lines of code | Time consumption |
|------------------|---------------------|------------------|
| Airplan 5 | 25869 loc | 11.165s |
| InfoTrader | 29335 loc | 8.419s |
| Linear Algebra | 3817 loc | 8.236s |
| Malware Analyzer | 2605 loc | 5.943s |
| Powerbroker 1 | 31705 loc | 10.696 |
| Powerbroker 4 | 31625 loc | 10.053 |
| Tweeter | 9896 loc | 8.128s |
| Withmi 3 | 26072 loc | 11.632s |
| Withmi 4 | 25972 loc | 11.227s |
| Average | 20766.22 loc | 9.5s |

Fig. 12. Efficiency: Running time.

6.3. Discussion and Limitations

Since the analysis algorithm defined in Fig. 10 is flow- and context-insensitive, its precision can be improved by inferring or verifying pre-conditions before applying each rule. Additionally, Rule T-Method is conservative under the program model of passing references (instead of values) back to a caller, i.e. a nonlinear amount of memory allocated in a callee method does not necessarily escape back to the caller

method.

7. Related Works

The most relevant works are in the area of resource and bound analysis using static program analysis techniques. While these works do not return useful information about resource usage if the analysis fails, our algorithm is designed with this feature. Additionally, some of these works may complement our work in the sense that, the procedure `LoopBound` defined in Figure 10 may be implemented with these works. Below are more detailed arguments supporting the above claims.

Sized-types. The works of Pedro B Vasconcelos *et al.* [15], [16] take a type-based approach by traversing program statements and generating constraints in a particular form w.r.t. variable sizes in a syntax-directed way, and finally obtaining closed-form solutions to the constraints using an external solver. The solution to the constraints is the symbolic resource usage of the given program. Compared with our work, these works define a functional core calculus that behaves differently than ours from the perspective of memory usage, i.e. while the fact that memory usage may be affected by mutation is not modeled by a functional core calculus, it is modeled by our core calculus. Since an analysis designed for such a functional core calculus [15], [16] does not take this effect into consideration, these works do not work on the core calculus defined in this paper (see Fig. 5 and 6), thus not giving useful information to secure analysts.

Cost analysis for Java programs. The works of Elvira Albert *et al.* [1]-[3] take a similar approach as Pedro B Vasconcelos *et al.* [15], [16] by generating constraints w.r.t. variable sizes in a syntax-directed way and finally obtaining closed-form solutions to the constraints using an external solver. The closed-form solution to the constraints is the symbolic resource usage of the given program. The difference here is that apart from generating constraints w.r.t. variable sizes, this line of works also path-insensitively generate constraints w.r.t variable values, thus achieving more precision. The work of Diego Esteban Alonso-Blas *et al.* [4] extends the work of Elvira Albert *et al.* [1]-[3] by defining a more expressive form of constraints w.r.t. variable sizes. The work of Antonio Flores-Montoya *et al.* [8] extends the work of Elvira Albert *et al.* [1]-[3] by excluding invalid constraints w.r.t. variable sizes (via identifying infeasible control flow paths) and hence improves analysis precision. Compared with our work, this line of works are more precise in terms of inferring resource bounds, because they collect constraints w.r.t. both variable values and sizes. However, recall that the problem addressed by this paper (see Section 1) is to either verify a program to use a linear amount of memory usage or report its vulnerable code snippets. For this purpose, this approach is not suitable because it is neither guaranteed that an external solver can return a closed-form solution (i.e. symbolic resource usage bounds), nor the solver can ever terminate.

When any of the two situations happens, this analysis will not give any helpful information to a secure analyst.

RAML. The works of Jan Hoffmann *et al.* [12], [13] targets at functional programs. It defines polynomial bound templates with unknown coefficients, and then uses a type-directed approach to collect constraints on the coefficients from a particular form of the language construct (i.e. a recursive call in pattern matching). Eventually, the set of collected constraints become a linear programming problem and is solved by an external solver. On one hand, this approach is similar to our work because both approaches take advantage of the insight that, traversing collection-typed variables is a very common coding pattern and can be used to infer resource usage of a program. On the other hand, similar to the work of Pedro B Vasconcelos *et al.* [15], [16], since this technique works for a functional core calculus, it does not consider mutation's effects on memory usage, thus not returning useful information to secure analysts.

SPEED. The works of Sumit Gulwani *et al.* [9]-[11] provide various ways to abstract program transitions such that the time bound can be directly inferred from the resultant transitions. Meanwhile, the abstraction

of transitions is generated by and is guaranteed to be sound (i.e. semantic-preserving) via Abstract Interpretation [7] based techniques. The work of Florian Zuleger [17] extends the works of Sumit Gulwani *et al.* [9]-[11] by adopting a new numeric domain for abstract interpretation, i.e. size-change abstraction. The domain is less expensive than other numerical domains, and yet at most times it is sufficient to effectively infer time bounds. This lines of work are complement to our work, because they can be used as an analysis module (i.e. procedure LoopBound) in the algorithm (see Fig. 10) to infer linear loop bounds.

Alternate runtime and size analysis. The work of Marc Brockschmidt [5] alternatively infers and refines variable sizes and time bounds for each Strong-Connected-Component in a program. This approach is complement to our work because procedure LoopBound defined in Fig. 10 may be implemented with this technique.

Compositional analysis. Similar to the work of Jan Hoffmann *et al.* [12], [13], the work of Quentin Carbonneaux [6] first defines bound templates with unknown coefficients for a given program, and then collect constraints on the coefficients. The set of constraints finally become a linear programming problem and are solved by an external solver. The difference between this work and the work of Jan Hoffmann *et al.* [12], [13] is that, this work targets at imperative programs while the other targets at functional programs. This work is not suitable for the purpose of verifying a program to use a linear amount of memory usage or reporting its vulnerable code snippets, because if a program's resource bound does not fit the bound templates, then this analysis cannot provide any useful information to a secure analyst (that is concerned with DoS memory vulnerabilities of the given program).

Simple bound analysis. The work of Moritz Sinn [14] takes advantage of ranking functions to infer time bound of a program. Ranking functions and transition abstraction [9]-[11], [17] are similar in the sense that, they both semantically abstract a program transition into a simpler form, from which a time bound is easier to infer. As a result, this work is also complement to our work because it can be used as a module (i.e. procedure LoopBound) in the algorithm (see Fig. 10) to infer linear loop bounds.

8. Conclusion

In this paper, we address the problem of Denial-of-Service memory vulnerabilities from one perspective, i.e. a program may only exhibit a memory usage linear to its input values. We present an analysis algorithm to verify a program to satisfy this definition, or report code snippets that may cause the program to use a nonlinear amount of memory if the verification fails. We then formally prove the correctness of the analysis algorithm. Our experimental results also indicate that our analysis algorithm is bot effective in verification and efficient in analysis time.

References

- [1] Elvira, A., Puri, A., Samir, G., Germán, P., & Damiano, Z. (2007). Cost analysis of java bytecode. *Proceedings of the European Symposium on Programming*.
- [2] Elvira, A., Samir, G., & Miguel, G. Z. (2007). Heap space analysis for java bytecode. *Proceedings of the 6th international symposium on Memory management*.
- [3] Elvira, A., Samir, G., & Miguel, G. Z. (2009). Live heap space analysis for languages with garbage collection. *Proceedings of the 2009 International Symposium on Memory Management*.
- [4] Diego, E. A. B., & Samir, G. (2012). On the limits of the classical approach to cost analysis. *International Static Analysis Symposium*.
- [5] Marc, B., Fabian, E., Stephan, F., Carsten, F., & Jürgen, G. (2014). Alternating runtime and size complexity analysis of integer programs. *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*.
- [6] Quentin, C., Jan, H., & Zhong, S. (2015). Compositional certified resource bounds. *ACM SIGPLAN Notices*, 50(6), 467–478, 2015.

- [7] Patrick, C., & Radhia, C. (1997). Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. *Proceedings of the 4th ACM Sigact-sigplan Symposium on Principles of Programming Languages*.
- [8] Antonio, F. M., & Reiner, H. (2014). Resource analysis of complex programs with cost equations. *Asian Symposium on Programming Languages and Systems*.
- [9] Sumit, G., Sagar, J., & Eric, K. (2009). Control-flow refinement and progress invariants for bound analysis. *ACM Sigplan Notices*, 44(6), 375–385.
- [10] Sumit, G., Krishna, K. M., & Trishul, C. (2009). Speed: Precise and efficient static estimation of program computational complexity. *ACM Sigplan Notices*.
- [11] Sumit, G., & Florian, Z. (2010). The reachability-bound problem. *ACM Sigplan Notices*.
- [12] Jan, H., Klaus, A., & Martin, H. (2011). Multivariate amortized resource analysis. *ACM SIGPLAN Notices*.
- [13] Jan, H., Ankush, D., & Shu, C. W. (2017). Towards automatic resource bound analysis for ocaml. *ACM Sigplan Notices*.
- [14] Moritz, S., Florian, Z., & Helmut, V. (2014). A simple and scalable static analysis for bound analysis and amortized complexity analysis. *International Conference on Computer Aided Verification*.
- [15] Pedro, B. V. (2008). *Space Cost Analysis Using Sized Types*. PhD thesis, University of St Andrews.
- [16] Pedro, B. V., & Kevin, H. (2003). Inferring cost equations for recursive, polymorphic and higherorder functional programs. *Proceedings of the Symposium on Implementation and Application of Functional Languages*.
- [17] Florian, Z., Sumit, G., Moritz, S., & Helmut, V. (2011). Bound analysis of imperative programs with the size-change abstraction. *International Static Analysis Symposium*.



Tianhan Lu obtained a degree of bachelor of science in computer science from Nanjing University in 2015 (Nanjing, Jiangsu, China). He received a degree of master of science in computer science from University of Colorado Boulder in 2017 (Boulder, Colorado, USA). Tianhan Lu is currently a PhD student majored in computer science in University of Colorado Boulder (started in 2015). His research interest includes static program analysis and verification, with a special focus on analyzing the amount of resources and costs that a program may use at runtime. Resource and cost analysis has a variety of applications in real world, including but not limited to analyzing software performance and verifying software to be free of cybersecurity vulnerabilities.



Yu-Ju Lee received his B.S. degree in electrical engineering from National Chung Cheng University, Chiayi, Taiwan, in 2005 and M.S. degree in electrical engineering from University of Southern California, Los Angeles, CA, USA, in 2010. He is currently working toward his Ph.D. degree in computer science at University of Colorado Boulder, Boulder, CO, USA. Before joined CU Boulder, he worked as a system architect and had several patents in Wi-Fi, Bluetooth and embedded system. His research interests are machine learning and deep learning application for system performance improvement, IoT system and networking design and analysis, and wireless coexistence system.



WenWei Liao is a research associate in the University of Colorado Boulder Cooperative Institute for Research in Environmental Sciences (CIRES), he is also a visiting associate professor in the National Chiao Tung University International College of Semiconductor Technology. His research interests span both education and computer technologies. Much of his work has been on improving the education technologies and the performance of the computer technologies. Dr. Liao is the corresponding author of the paper "Towards Denial-of-Service Memory Vulnerabilities". He is also a good principal investigator in the University

of Colorado Boulder.