Hindawi Publishing Corporation
Journal of Applied Mathematics
Volume 2014, Article ID 789213, 12 pages
http://dx.doi.org/10.1155/2014/789213

Research Article

Robustness Analysis of Floating-Point Programs by Self-Composition

Liqian Chen, Jiahong Jiang, Banghu Yin, Wei Dong, and Ji Wang

National Laboratory for Parallel and Distributed Processing, National University of Defense Technology, Changsha 410073, China

Correspondence should be addressed to Liqian Chen; lqchen@nudt.edu.cn

Received 14 February 2014; Accepted 7 April 2014; Published 20 May 2014

Academic Editor: Xiaoyu Song

Robustness is a key property for critical systems that run in uncertain environments, to ensure that small input perturbations can cause only small output changes. Current critical systems often involve lots of floating-point computations which are inexact. Robustness analysis of floating-point programs needs to consider both the uncertain inputs and the inexact computation. In this paper, we propose to leverage the idea of self-composition to transform the robustness analysis problem into a reachability problem, which enables the use of standard reachability analysis techniques such as software model checking and symbolic execution for robustness analysis. To handle floating-point arithmetic, we employ an abstraction that encompasses the effect of rounding and that can encompass all rounding modes. It converts floating-point expressions into linear expressions with interval coefficients in exact real arithmetic. On this basis, we employ interval linear programming to compute the maximum output change or maximum allowed input perturbation for the abstracted programs. Preliminary experimental results of our prototype implementation are encouraging.

1. Introduction

Uncertainty and inexactness in computing have attracted much attention in computer science. In Cyber Physical Systems (CPS), the discrete world of computation is integrated with the continuous world of physical processes. Moreover, CPS run in the open environmental context and thus have to deal with uncertain data which may come from noisy sensor data or approximate computation. Hence, inputs for programs in CPS are of intrinsic uncertainty. On the other hand, due to finite precision on computers, physical values are truncated into digital ones. In modern computers, real numbers are approximated by a finite set of floating-point numbers. Due to the pervasive rounding errors, numerical computation using floating-point arithmetic is not exact. Since many safety-critical CPS systems (such as aircrafts, automobiles, and medical devices) often involve lots of numerical computations, there is a great need to ensure that these programs are *robust* with respect to the uncertain input as well as the inexact computation.

Although *robustness* is long known as a standard correctness property for control systems [1], considering the robustness of programs is quite recent [2–5]. Intuitively, robustness of a program means that small input perturbations of the program can cause only small output changes. Much existing work on analyzing robustness of programs assumes that the analyzed program is in exact real arithmetic, although floating-point computation is pervasive in practical applications. This paper targets the analysis of robustness properties of floating-point programs.

A program using floating-point arithmetic often exhibits more robustness issues than that using exact real arithmetic, due to the misunderstandings and nonintuitive behaviors of floating-point semantics. Although floating-point arithmetic is quite different from the exact real arithmetic, most developers of floating-point programs will write programs as if computations were done in exact arithmetic. For the same input, the control flow of the program using floating-point arithmetic can be different from the one that would be taken assuming exact real arithmetic. Similarly,

for two inputs whose values are close to each other, the resulting two control flows of the same program can be different (even when following exact real arithmetic). Two different control flows may lead to very large difference in outputs.

We illustrate the robustness problem due to floating-point computation using a motivating example shown in Figure 1, which is a "classroom" example of a robustness problem frequently used in the field of geometric computations [6]. The program `Orientation` implements the 2D orientation test that decides whether a point r lies to the left of, to the right of, or on the line $\overrightarrow{pq}$ defined by the 2 points p, q, by evaluating the sign of a determinant `det` which is expressed in terms of the coordinates of the input points. Due to rounding errors, the floating-point computation of the determinant `det` may lead to a wrong result when the true determinant (via exact real arithmetic) is close to zero. From the robustness point of view, for this program, even a very small input perturbation may lead to an output change of 1 or 2. If the rounding modes for the floating-point operations are not determinate in the program, the output change can be 2 even when there is no perturbation in the inputs (by running the program in different rounding modes). This misinformation may then lead to a failure of a computational geometry application (e.g., crash or not terminate) or produce wrong results [6].

Analyzing robustness of floating-point programs is more challenging than analyzing programs assuming exact real arithmetic, since besides the input perturbations, we need to consider also the inexactness of floating-point computation. The floating-point program itself acts as if inputs were perturbated due to the pervasive rounding errors or nondeterminate rounding modes. There exist a few known pitfalls of analyzing and verifying floating-point programs [7].

In this paper, we present a robustness analysis method for floating-point programs. The key idea is to leverage the self-composition technique from the field of secure information flow to transform the robustness analysis problem into a reachability (safety) problem. Then we use standard rechability analysis techniques such as software model checking and symbolic execution to analyze the self-composed programs. To cope with floating-point arithmetic, we utilize a rounding mode insensitive abstraction method to abstract floating-point expressions into linear expressions with interval coefficients in the field of reals. On this basis, we use interval linear programming to compute the maximum output change (when given the input perturbation) or the maximum allowed input perturbation (when given the output change) for the abstracted programs. The preliminary experimental results are promising on benchmark programs.

The rest of the paper is organized as follows. Section 2 reviews the IEEE 754 floating-point arithmetic and the basic theory of interval linear systems as well as interval linear programming. Section 3 presents the robustness analysis approach via self-composition for programs (that assume exact real arithmetic). Section 4 presents the techniques to handle floating-point arithmetic. Section 5 presents our prototype implementation together with preliminary experimental results. Section 6 discusses some related work before Section 7 concludes.

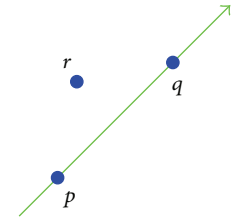


```
(1) int Orientation (float px, float py, float qx, float qy, float rx, float ry)
(2) {
(3)     float pqx = qx − px, pqy = qy − py;
(4)     float prx = rx − px, pry = ry − py;
(5)     float det = pqx ∗ pry − pqy ∗ prx;
(6)     if (det > 0) return 1;
(7)     if (det < 0) return −1;
(8)     return 0;
(9) }
```

Figure 1: Floating-point implementation for orientation test of 2D points.

2. Preliminaries

In this section, we briefly provide the background on the IEEE 754 floating-point arithmetic and the basic theory on interval arithmetics as well as interval linear programming.

2.1. The IEEE 754 Floating-Point Arithmetic. A digital computer cannot represent all possible real numbers in mathematics exactly. In computing, floating-point numbers provide an approach to represent a finite subset of the real numbers. In this paper, we focus on analyzing programs with respect to the binary formats of the IEEE 754 floating-point standard [8] which is the most commonly used floating-point representation and is followed by almost all modern computers.

In the IEEE 754 standard, the binary representation of a floating-point number x can be described as $x = (-1)^S \times M \times 2^E$, where

(i) S is the 1-bit *sign* of x, which represents that x is positive (when $S = 0$) or negative (when $S = 1$);

(ii) $E = e - \mathbf{bias}$ is called the *exponent*, where e is a biased $\mathbf{e}$-bit unsigned integer and $\mathbf{bias} = 2^{\mathbf{e}-1} - 1$;

(iii) $M = m_0.m_1m_2 \ldots m_{\mathbf{p}}$ is called the *significand*, where $f = .m_1m_2 \ldots m_{\mathbf{p}}$ represents a $\mathbf{p}$-bit fraction and m_0 is the hidden bit without need of storage.

The values of $\mathbf{e}$, $\mathbf{bias}$, $\mathbf{p}$ depend on the floating-point formats. The IEEE 754 standard supports several formats, among which the basic formats include

(i) 32-bit single-precision format, where $\mathbf{e} = 8$ (and thus $\mathbf{bias} = 127$), $\mathbf{p} = 23$;

(ii) 64-bit double-precision format, where $\mathbf{e} = 11$ (and thus $\mathbf{bias} = 1023$), $\mathbf{p} = 52$.

According to the value of e, the floating-point numbers can be divided into the following categories:

(i) *normalized* number $(-1)^S \times 1.f \times 2^{e-\mathbf{bias}}$, when $1 \le e \le 2^{\mathbf{e}} - 2$;

(ii) *denormalized* number $(-1)^S \times 0.f \times 2^{1-\mathbf{bias}}$, when $e = 0$ and $f \neq 0$;

(iii) $+0$ or -0, when $e = 0$ and $f = 0$;

(iv) $+\infty$ or $-\infty$, when $e = 2^{\mathbf{e}} - 1$ and $f = 0$;

(v) NaN (Not a Number), when $e = 2^{\mathbf{e}} - 1$ and $f \neq 0$.

Let $\mathbf{F}$ be the set of floating-point formats. For each $\mathbf{f} \in \mathbf{F}$, we define

(i) $mf_{\mathbf{f}} \stackrel{\text{def}}{=} 2^{1-\mathbf{bias}-\mathbf{p}}$, the smallest nonzero positive floating-point number;

(ii) $Mf_{\mathbf{f}} \stackrel{\text{def}}{=} (2 - 2^{-\mathbf{p}})2^{2^{\mathbf{e}}-\mathbf{bias}-2}$, the largest noninfinity floating-point number.

In general, the result of a floating-point operation may not be exactly representable in the floating-point representation, and thus the result needs to be rounded into a floating-point number. The IEEE 754 standard supports four rounding modes: toward nearest, toward $+\infty$, toward $-\infty$, and toward zero. In this paper, in order to distinguish floating-point arithmetic operations from exact real arithmetic ones, we introduce additional notations. As usual, $\{+, -, \times, /\}$ are used as exact rational arithmetic operations. The corresponding floating-point operations are denoted by $\{\oplus_{\mathbf{f},r}, \ominus_{\mathbf{f},r}, \otimes_{\mathbf{f},r}, \oslash_{\mathbf{f},r}\}$, tagged with a floating-point format $\mathbf{f} \in \mathbf{F}$ and a rounding mode $r \in \{+\infty, -\infty, 0, n\}$ (n representing rounding to nearest). We also use ? to denote arbitrary rounding mode.

Due to rounding errors, many well-known algebraic properties (such as associativity and distributivity) over the reals do not hold for floating-point arithmetic.

Example 1. Consider the following expressions in the 32-bit single-precision floating-point arithmetic:

$$\begin{aligned} \left(2^{24} \oplus_{32,?} - 2^{24}\right) \oplus_{32,?} 1 &= 1 \\ \left(2^{24} \oplus_{32,-\infty} 1\right) \oplus_{32,-\infty} - 2^{24} &= 0 \\ \left(2^{24} \oplus_{32,+\infty} 1\right) \oplus_{32,+\infty} - 2^{24} &= 2. \end{aligned} \tag{1}$$

Note that in the 32-bit single-precision format, the significand is $M = m_0.m_1m_2 \ldots m_{\mathbf{23}}$. However, to represent the exact result of $2^{24} + 1$ over the reals, we need one more bit for the significand M (say $m_{\mathbf{24}}$). Hence, rounding happens. $2^{24} \oplus_{32,-\infty} 1$ will result in 2^{24}, while $2^{24} \oplus_{32,+\infty} 1$ will result in $(1 + 2^{-23}) \times 2^{24}$.

2.2. Interval Linear Systems and Interval Linear Programming. Let $\underline{A}, \overline{A} \in \mathbb{R}^{m \times n}$ be two matrices with $\underline{A} \le \overline{A}$, where comparison operators are defined element-wise; then the set of matrices $\mathbf{A} \in \mathbb{IR}^{m \times n}$ defined by

$$\mathbf{A} = \left[\underline{A}, \overline{A}\right] = \left\{A \in \mathbb{R}^{m \times n} : \underline{A} \le A \le \overline{A}\right\} \tag{2}$$

is called an *interval matrix*, and the matrices $\underline{A}, \overline{A}$ are called its bounds. Let us define the *center matrix* of $\mathbf{A}$ as $A_c = (1/2)(\underline{A} + \overline{A})$ and the *radius matrix* as $\Delta_A = (1/2)(\overline{A} - \underline{A})$. Then, $\mathbf{A} = [\underline{A}, \overline{A}] = [A_c - \Delta_A, A_c + \Delta_A]$. An *interval vector* is a one-column interval matrix $\mathbf{d} = [\underline{d}, \overline{d}] = \{d \in \mathbb{R}^m : \underline{d} \le d \le \overline{d}\}$, where $\underline{d}, \overline{d} \in \mathbb{R}^m$ and $\underline{d} \le \overline{d}$.

Let $\mathbf{A}$ be an $m \times n$ interval matrix and b be a vector of size m. The following system of interval linear inequalities

$$\mathbf{A}x \le b \tag{3}$$

denotes an *interval linear system*, that is, the *family* of all systems of linear inequalities $Ax \le b$ such that $A \in \mathbf{A}$.

Definition 2 (weak solution). A vector $x \in \mathbb{R}^n$ is called a *weak solution* of the interval linear system $\mathbf{A}x \le b$, if it satisfies $Ax \le b$ for some $A \in \mathbf{A}$. Furthermore, the set

$$\Sigma_{\exists}(\mathbf{A}, b) = \left\{x \in \mathbb{R}^n : \exists A \in \mathbf{A}, Ax \le b\right\} \tag{4}$$

is said to be the *weak solution set* of the system $\mathbf{A}x \le b$.

The weak solution set of an interval linear system is characterized by the following theorem [9].

Theorem 3. *A vector $x \in \mathbb{R}^n$ is a weak solution of $\mathbf{A}x \le b$ if and only if it satisfies $A_c x - \Delta_A |x| \le b$.*

Let $\mathbf{A} \in \mathbb{IR}^{m \times n}$ be an $m \times n$ interval matrix, $b \in \mathbb{R}^m$ be an m-dimensional vector, and $\mathbf{c} \in \mathbb{IR}^n$ be an n-dimensional interval vector. The *family* of linear programming (LP) problems

$$f(A, b, c) = \max\left\{c^T x : Ax \le b\right\} \tag{5}$$

with data satisfying

$$A \in \mathbf{A}, \quad c \in \mathbf{c} \tag{6}$$

is called an interval linear programming (ILP) problem.

In this paper, we are only interested in computing the upper bound $\overline{f}(\mathbf{A}, b, \mathbf{c}) = \sup\{f(A, b, c) : A \in \mathbf{A}, c \in \mathbf{c}\}$. In general, according to Theorem 3, to compute the exact $\overline{f}(\mathbf{A}, b, \mathbf{c})$, in the worst case up to 2^n LP problems have to be solved, one for each orthant. Recall that a (closed) orthant is one of the 2^n subsets of an n-dimensional Euclidean space defined by constraining each Cartesian coordinate axis to be either nonnegative or nonpositive. In each orthant, we consider the following LP problem:

$$\begin{aligned} \max \quad & \sum_{j=1}^{n} c'_j x_j \\ \text{s.t.} \quad & \bigwedge_{0 \le i \le m} \sum_{j=1}^{n} A'_{ij} x_j \le b_i, \end{aligned} \tag{7}$$

where

$$c'_j = \begin{cases} \overline{c}_j & \text{if } x_j \geq 0 \\ \underline{c}_j & \text{if } x_j < 0 \end{cases}$$
$$A'_{ij} = \begin{cases} \underline{A}_{ij} & \text{if } x_j \geq 0 \\ \overline{A}_{ij} & \text{if } x_j < 0. \end{cases} \quad (8)$$

And $\overline{f}(\mathbf{A}, b, \mathbf{c})$ will be the the maximum over all the optimal values of the 2^n LP problems with one per each orthant.

3. Robustness Analysis via Self-Composition

3.1. Robustness of Programs. In this paper, we follow the definition for robustness of programs used by Majumdar and Saha [5]. Let f be a function with inputs $x_1, \ldots, x_n$ and output y; that is, $y = f(x_1, \ldots, x_n)$. The function f is said to be (δ, ϵ)-*robust* in the ith input x_i if a perturbation of at most δ in the input x_i can only cause a change of at most ϵ in the output; that is,

$$\forall x_i, x'_i \cdot \left|x_i - x'_i\right| \leq \delta \implies \left|f(x_1, \ldots, x_i, \ldots, x_n) - f\left(x_1, \ldots, x'_i, \ldots, x_n\right)\right| \leq \epsilon, \quad (9)$$

where $\delta, \epsilon \in \mathbb{R}$ are nonnegative constant parameters specified by users. Recall that we consider the perturbation over only one input at a time while assume that there is no perturbation over all other inputs at the same time.

Moreover, in practice, users may be interested in the maximum output change of y with respect to x_i and δ; that is,

$$\overline{\epsilon}_\delta \stackrel{\text{def}}{=} \max_{x,x'} \left\{ \left|y - y'\right| \mid \begin{array}{c} y = f\left(x_1, \ldots, x_i, \ldots, x_n\right) \\ y' = f\left(x_1, \ldots, x'_i, \ldots, x_n\right) \\ \left|x_i - x'_i\right| \leq \delta \end{array} \right\}. \quad (10)$$

Similarly, users may be interested in the maximum input perturbation allowed over x_i with respect to y and ϵ; that is,

$$\overline{\delta}_\epsilon \stackrel{\text{def}}{=} \max_{y,y'} \left\{ \left|x_i - x'_i\right| \mid \begin{array}{c} y = f\left(x_1, \ldots, x_i, \ldots, x_n\right) \\ y' = f\left(x_1, \ldots, x'_i, \ldots, x_n\right) \\ \left|y - y'\right| \leq \epsilon \end{array} \right\}. \quad (11)$$

Example 4. Consider the program shown in Figure 2, which implements a piece-wise linear function. When $x = 1.001$, the two branches give the same result $y = 1002.001$ in exact real arithmetic (assuming floats are reals). It is easy to see that in exact real arithmetic, this program is $(0.1, \epsilon_0)$-*robust* for all $\epsilon_0 \geq 100.1$ but is not $(0.1, \epsilon_1)$-*robust* for all $\epsilon_1 < 100.1$. This can be deduced by observing that in exact real arithmetic, given the input perturbation $\delta = 0.1$, the maximum output change of y is $\overline{\epsilon}_\delta = 100.1$; given the output change $\epsilon = 100.1$, the maximum input perturbation allowed over x is $\overline{\delta}_\epsilon = 0.1$.

```
(1) float piecewise_linear(float x) {
(2)     float y;
(3)     if(x < 1.001)
(4)         y = x + 1001.0;
(5)     else
(6)         y = x * 1001.0;
(7)     return y;
(8) }
```

Figure 2: A floating-point program `piecewise_linear`.

3.2. Self-Composition. The idea of self-composition is firstly used in the field of secure information flow [10, 11] to characterize noninterference. Let P be a program and P' be a copy of P with each variable x in P replaced by a fresh variable x'. Using Hoare triples, noninterference can be characterized as

$$\left\{L = L'\right\} \; P; P' \; \left\{L = L'\right\}, \quad (12)$$

where L denotes low-security variables. In other words, it requires that running two instances of the same program with equal low-security values and arbitrary high-security values results in equal low-security values. Hence, via self-composition, a secure information flow property of P reduces to a reachability property over single program executions of the program $P; P'$.

In this paper, we would like to leverage the idea of self-composition to reduce the robustness problem of a program P into an equivalent reachability problem over $P; P'$. Assume that program P has n input variables $x_1, \ldots, x_n$ and an output variable y. Similarly, using Hoare triples, the (δ, ϵ)-*robustness* of program P over the ith input x_i can be characterized as

$$\left\{\left|x_i - x'_i\right| \leq \delta \wedge \bigwedge_{1 \leq j \leq n, j \neq i} x_j = x'_j\right\} \; P; P' \; \left\{\left|y - y'\right| \leq \epsilon\right\}. \quad (13)$$

Example 5. Consider again the program `piecewise_linear` in Figure 2. The self-composition of the function body is shown in Figure 3. To express the robustness property, we add the assumption $|x - x'| \leq \delta$ as a precondition at the beginning of the self-composed program and add an assertion $|y - y'| \leq \epsilon$ as a postcondition at the end.

Essentially, the copied program P' has the same program code as P but uses variables with different initial values. Hence, there exists inherent symmetry and redundancy in the self-composed programs. In order to make the following analysis and verification process for self-composed programs easier, program transformations can be used to optimize the self-composed programs. In the field of secure information flow analysis, Terauchi and Aiken [12] proposed type-directed transformation to improve self-composition. The main idea of type-directed transformation is not to self-compose branch (or loop) statements when the branch (or loop) condition is only dependent on the values of low-security variables. In addition, for an assignment statement

```
(1) assume(−δ ≤ x − x′ ≤ δ)
(2)
(3) if (x < 1.001)
(4)     y = x + 1001.0;
(5) else
(6)     y = x ∗ 1001.0;
(7)
(8) if (x′ < 1.001)
(9)     y′ = x′ + 1001.0;
(10) else
(11)    y′ = x′ ∗ 1001.0;
(12)
(13) assert(−ϵ ≤ y − y′ ≤ ϵ)
```

Figure 3: Self-composition of the `piecewise_linear` program for robustness analysis.

$\{x := e;\}$, when the right-hand expression e is only dependent on the values of low-security variables, its self-composition is simplified as $\{x := e; x' := x;\}$.

With respect to robustness, a similar transformation can be applied. Intuitively, we could consider the perturbed input variable x_i as a high-security variable and all other input variables x_j's as low-security variables where $j \neq i$. Hence, similarly to type directed transformation, we do not self-compose branch (or loop) statements when the branch (or loop) condition is not dependent on the values of perturbed input variables. For an assignment statement $\{x := e;\}$, when the right-hand expression e is not dependent on the values of perturbed input variables, its self-composition is simplified as $\{x := e; x' := x;\}$.

Example 6. Consider the function *min_plus*1 shown in Figure 4, which implements $\min(x + 1, y)$ by adding 0.1 to x ten times. The optimized self-composition result of the function body after applying transformation is given in Figure 5, when we consider the perturbation over the input variable x (while assuming no perturbation over y). More specifically, since the loop condition $i < 10$ in the original program is not dependent on the value of the perturbed input variable x, we do not self-compose the loop statement and thus there is only one loop in the transformed resulting self-composed program.

3.3. Robustness Analysis of Self-Composed Programs. Via self-composition, the robustness analysis problem can be reduced to solving a standard reachability (safety) problem. The recent success of automatic analysis and verification tools (such as SLAM [13], CBMC [14], and ASTRÉE [15]) aiming at checking reachability properties in programs makes this approach promising. In the following, we will present two popular reachability analysis approaches that fit for analyzing robustness, that is, software model checking and symbolic execution.

3.3.1. Checking Robustness by Software Model Checking. Software model checking [16] provides an automatic approach to check whether a program satisfies a property by exploring

```
(1) float min_plus1(float x, float y){
(2)     float z;
(3)     int i;
(4)     i = 0;
(5)     while (i < 10) {
(6)         x = x + 0.1;
(7)         i = i + 1;
(8)     }
(9)     if (y <= x) z = y;
(10)    else z = x;
(11)    return z;
(12) }
```

Figure 4: A floating-point program `min_plus1`.

```
(1) assume(−δ ≤ x − x′ ≤ δ and y = y′)
(2)
(3) i = 0; i′ = i;
(4) while (i < 10) {
(5)     x = x + 0.1; x′ = x′ + 0.1;
(6)     i = i + 1; i′ = i;
(7) }
(8) if (y <= x) z = y;
(9) else z = x;
(10) if (y′ <= x′) z′ = y′;
(11) else z′ = x′;
(12)
(13) assert(−ϵ ≤ z − z′ ≤ ϵ)
```

Figure 5: Transformed self-composition of the `min_plus1` program for robustness analysis.

the state space of the program. For the robustness analysis problem, the property to be checked is an assertion at the end of the self-composed programs stating that the output change is bounded by ϵ, that is, `assert` $(-\epsilon \leq y - y' \leq \epsilon)$. A main advantage of using software model checking is that it will generate a counterexample when the robustness property does not hold. The counterexample shows an execution trace which violates the robustness property. A counterexample is very helpful for the users to identify the source of nonrobustness.

3.3.2. Finding Maximum Output Change (or Input Perturbation) by Symbolic Execution. Symbolic execution [17, 18] is a technique to analyze a program by executing the program with symbolic rather than concrete values as program inputs. The process of symbolic execution essentially generates and explores a symbolic execution tree which represents all execution paths followed during the process. Each tree node represents a symbolic execution state, while each edge represents a program transition between the states. At any tree node, the symbolic execution state includes a program counter, a *path condition* (PC) that encodes the constraints on the symbolic inputs to reach that node, a *path function* (PF) that represents the current values of the program variables as function of symbolic inputs when the path condition holds

true. The path condition is a boolean expression over the symbolic inputs. The path function describes the expected result of the program, under the given path condition. Due to conditional branches and loops in a program, the symbolic execution of a program will result in a set of paths, each of which is described by a pair $\langle \mathrm{PC}, \mathrm{PF}\rangle$ of the path condition PC and the associated path function PF.

We now show how to use symbolic execution to conduct a robustness analysis of program P with inputs $x_1, \ldots, x_n$ and output y. First, the analysis algorithm performs symbolic execution on the self-composed program $P; P'$. Assume that the algorithm collects, at the end of the self-composed program, a set $\mathcal{S}$ of pairs $\langle \mathrm{PC}, \mathrm{PF}\rangle$ of the path condition PC and the associated path function PF. Then for each $s \triangleq \langle \mathrm{PC}, \mathrm{PF}\rangle \in \mathcal{S}$, we compute the maximum output change $\overline{\epsilon}_\delta^s$:

$$\begin{aligned} &\max \quad \left|\mathrm{PF}_y - \mathrm{PF}_{y'}\right| \\ &\text{s.t.} \quad \left(\left|x_i - x_i'\right| \le \delta \wedge \bigwedge_{1\le j\le n, j\ne i} x_j = x_j'\right) \wedge \mathrm{PC}. \end{aligned} \tag{14}$$

Here, PF_y and $\mathrm{PF}_{y'}$ denote the symbolic expressions that the path function PF maps the variables y and y' to, respectively. Let $\overline{\epsilon}_\delta$ be the maximum element of $\{\overline{\epsilon}_\delta^s \mid s \in \mathcal{S}\}$; that is, $\overline{\epsilon}_\delta = \max(\{\overline{\epsilon}_\delta^s \mid s \in \mathcal{S}\})$. If $\overline{\epsilon}_\delta \le \epsilon$, then the original program P is (δ, ϵ)-*robust*.

Similarly, given the bound of output change ϵ, computing the maximum allowed input perturbation is reduced to solving a series of the following optimization problems for each $s \triangleq \langle \mathrm{PC}, \mathrm{PF}\rangle \in \mathcal{S}$:

$$\begin{aligned} &\max \quad \left|x_i - x_i'\right| \\ &\text{s.t.} \quad \left(\bigwedge_{1\le j\le n, j\ne i} x_j = x_j'\right) \wedge \mathrm{PC} \wedge \left|\mathrm{PF}_y - \mathrm{PF}_{y'}\right| \le \epsilon. \end{aligned} \tag{15}$$

And $\overline{\delta}_\epsilon$ will be the maximum element of $\{\overline{\delta}_\epsilon^s \mid s \in \mathcal{S}\}$; that is, $\overline{\delta}_\epsilon = \max(\{\overline{\delta}_\epsilon^s \mid s \in \mathcal{S}\})$.

4. Robustness Analysis of Floating-Point Programs

In this section, we consider the robustness analysis problem of floating-point programs. In Section 3.3, we propose to utilize software model checking and symbolic execution to perform robustness analysis of self-composed programs (in exact real arithmetic). However, most existing software model checkers and symbolic execution tools can not be directly applied to floating-point programs, since they rely on constraint solvers that often assume good algebraic properties such as associativity and distributivity over the reals which do not hold for floating-point arithmetic. To handle floating-point arithmetic, we have to resort to bit-precise modeling of floating-point arithmetic or abstracting floating-point arithmetic to real number arithmetic.

CBMC (C Bounded Model Checker) [14] is one of the few software model checkers that have considered floating-point arithmetic. CBMC employs a sound and complete decision procedure for floating-point arithmetic [19, 20]. It precisely encodes floating-point operations as functions on bit-vectors. Each floating-point operation is further modeled as a formula in propositional logic. The formula is then handled by a SAT-solver in the backend to check for satisfiability.

When we consider symbolic execution of floating-point programs, both the path condition and the path function will involve floating-point expressions. Hence, to compute the maximum output change (or maximum allowed input perturbation), we need optimization methods supporting floating-point constraints. However, as far as we know, even for linear programming, there is no available sound solver supporting floating-point constraints. To this end, in this paper, we abstract the optimization problem with floating-point constraints into an interval linear programming problem (i.e., linear programming problem with interval coefficients) over the reals. The main idea is to use the so-called *floating-point linearization* technique [21, 22] to abstract floating-point expressions into linear real number expressions with interval coefficients (in the form of $\Sigma_i [a_i, b_i] x_i$).

4.1. Floating-Point Abstraction. In this subsection, we will explain how to abstract floating-point expressions into interval linear expressions over the reals.

First, let us consider the upper bound on rounding errors due to one floating-point operation. Let $R_{\mathbf{f},r}(x)$ denote the floating-point rounding function that maps a real number x to a floating-point number (or a runtime error due to, for example, overflows) with respect to the floating-point format $\mathbf{f}$ and the rounding mode r. The amount of the rounding error due to $R_{\mathbf{f},r}(x)$ depends on the category of x.

(i) If x is in the range of normalized numbers, then $|R_{\mathbf{f},r}(x) - x| \le \varepsilon_{\mathrm{rel}} \cdot |x|$ where $\varepsilon_{\mathrm{rel}} = 2^{-\mathbf{p}}$ (wherein $\mathbf{p}$ is the number of bits of fraction in the significand of the floating-point format $\mathbf{f}$). In this case we consider the relative rounding error $\varepsilon_{\mathrm{rel}}$.

(ii) If x is in the range of denormalized number, then $|R_{\mathbf{f},r}(x) - x| \le \varepsilon_{\mathrm{abs}}$, where $\varepsilon_{\mathrm{abs}} = mf_{\mathbf{f}}$ (wherein $mf_{\mathbf{f}}$ is the smallest nonzero positive denormalized floating-point number in the floating-point format $\mathbf{f}$, which is also the gap between two neighboring denormalized numbers). In this case, we consider the absolute rounding error $\varepsilon_{\mathrm{abs}}$.

The rounding errors of these two cases can be unified as

$$\left|R_{\mathbf{f},r}(x) - x\right| \le \max\left(\varepsilon_{\mathrm{rel}} \cdot |x|, \varepsilon_{\mathrm{abs}}\right). \tag{16}$$

Since max is not a linear operation, we derive an overapproximation

$$\left|R_{\mathbf{f},r}(x) - x\right| \le \varepsilon_{\mathrm{rel}} \cdot |x| + \varepsilon_{\mathrm{abs}}. \tag{17}$$

Furthermore, when $b \ge 0$, $|y| \le b$ is equivalent to $y = [-1, 1] \times b$. Hence,

$$R_{\mathbf{f},r}(x) - x = [-1, 1]\left(\varepsilon_{\mathrm{rel}} \cdot |x| + \varepsilon_{\mathrm{abs}}\right); \tag{18}$$

that is,

$$R_{\mathbf{f},r}(x) = [1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \times x + [-\varepsilon_{\text{abs}}, \varepsilon_{\text{abs}}]. \quad (19)$$

In general, we could abstract floating-point operations into interval linear expressions in real number semantics. For example,

$$x \oplus_{\mathbf{f},r} y, \quad (20)$$

that is,

$$R_{\mathbf{f},r}(x+y) \quad (21)$$

can be abstracted into

$$[1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \times (x+y) + [-\varepsilon_{\text{abs}}, \varepsilon_{\text{abs}}]; \quad (22)$$

that is,

$$[1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \times x + [1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \times y + [-\varepsilon_{\text{abs}}, \varepsilon_{\text{abs}}]. \quad (23)$$

The advantage of this kind of rounding mode insensitive floating-point abstractions is that the result is sound with respect to arbitrary rounding modes, since $R_{\mathbf{f},r}(x)$ always satisfies $R_{\mathbf{f},-\infty}(x) \le R_{\mathbf{f},r}(x) \le R_{\mathbf{f},+\infty}(x)$ while $|R_{\mathbf{f},r}(x) - x| \le \varepsilon_{\text{rel}} \cdot |x| + \varepsilon_{\text{abs}}$ has already taken into account the extreme cases of $r = -\infty$ and $r = +\infty$. This is of practical importance, since we may not know the exact rounding mode for each floating-point operation. For example, C99 provides the `fesetround()` function to set the current rounding mode. Of course, when we know the exact rounding mode for the floating-point operation, we could make the floating-point abstraction more precise. For example, if the current rounding mode is toward nearest, then

$$R_{\mathbf{f},n}(x) = \left[1-\frac{\varepsilon_{\text{rel}}}{2}, 1+\frac{\varepsilon_{\text{rel}}}{2}\right] \times x + \left[-\frac{\varepsilon_{\text{abs}}}{2}, \frac{\varepsilon_{\text{abs}}}{2}\right]. \quad (24)$$

In addition, if we know the range of x, we may also define more precise floating-point abstractions. For example, if we know that x is in the range of denormalized numbers, then

$$R_{\mathbf{f},r}(x) = x + [-\varepsilon_{\text{abs}}, \varepsilon_{\text{abs}}]. \quad (25)$$

For the sake of generality, in this paper, we use the following rounding mode insensitive floating-point abstraction:

$$R^{\#}_{\mathbf{f},?}(x) = [1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \times x + [-\varepsilon_{\text{abs}}, \varepsilon_{\text{abs}}], \quad (26)$$

where we assume $|x| < Mf_{\mathbf{f}}$.

More clearly, we use the following abstraction for floating-point arithmetic:

$$\begin{aligned}
R^{\#}_{\mathbf{f},?}(x \oplus_{\mathbf{f},?} y) &= [1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \times x + [1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \\
&\quad \times y + [-\varepsilon_{\text{abs}}, \varepsilon_{\text{abs}}] \\
R^{\#}_{\mathbf{f},?}(x \ominus_{\mathbf{f},?} y) &= [1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \times x + [-1-\varepsilon_{\text{rel}}, -1+\varepsilon_{\text{rel}}] \\
&\quad \times y + [-\varepsilon_{\text{abs}}, \varepsilon_{\text{abs}}] \\
R^{\#}_{\mathbf{f},?}(x \otimes_{\mathbf{f},?} y) &= [1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \times x \times y + [-\varepsilon_{\text{abs}}, \varepsilon_{\text{abs}}] \\
R^{\#}_{\mathbf{f},?}(x \oslash_{\mathbf{f},?} y) &= [1-\varepsilon_{\text{rel}}, 1+\varepsilon_{\text{rel}}] \times \frac{x}{y} + [-\varepsilon_{\text{abs}}, \varepsilon_{\text{abs}}].
\end{aligned} \quad (27)$$

Specially, for a constant number c that appears in the source code, we use the following abstraction:

$$R^{\#}_{\mathbf{f},?}(c) = \left[R^{\#}_{\mathbf{f},-\infty}(c), R^{\#}_{\mathbf{f},+\infty}(c)\right]. \quad (28)$$

4.2. Symbolic Execution of Abstracted Floating-Point Programs. From Section 4.1, we see that floating-point expressions can be soundly abstracted into real number expressions with interval coefficients. Since the multiplication $x \times y$ and division x/y are not linear expressions when both x and y are not constant numbers, in order to obtain linear expressions with interval coefficients, we replace y with its interval range denoted as $[\underline{y}, \overline{y}]$. In symbolic execution, y is always an expression over the symbolic input values. We assume users provide the interval ranges for those symbolic input values. Then, all floating-point expressions can be abstracted as interval linear expressions. Therefore, the resulting path conditions of symbolic execution consist of interval linear constraints while the resulting path functions consist of interval linear expressions.

Finally, the problems of computing the maximum output change and the maximum allowed input perturbation are reduced to solving a series of interval linear programming problems. For example, computing the maximum output change requires the solutions of the following interval linear programming problems:

$$\begin{aligned}
\max \quad & \Sigma_i [\underline{a}_i, \overline{a}_i] \times x_i + \Sigma_i \left[\underline{a}'_i, \overline{a}'_i\right] \times x'_i + b \\
\text{s.t.} \quad & \left(\left|x_i - x'_i\right| \le \delta \wedge \bigwedge_{1 \le j \le n, j \ne i} x_j = x'_j\right) \\
& \wedge \bigwedge_k \Sigma_i \left[\underline{A}_{ki}, \overline{A}_{ki}\right] \times x_i + \Sigma_i \left[\underline{A}'_{ki}, \overline{A}'_{ki}\right] \times x'_i \le c_k,
\end{aligned} \quad (29)$$

where $\Sigma_i[\underline{a}_i, \overline{a}_i] \times x_i + \Sigma_i[\underline{a}'_i, \overline{a}'_i] \times x'_i + b$ denotes the abstracted output change $\text{PF}_y - \text{PF}_{y'}$ (or $\text{PF}_{y'} - \text{PF}_y$) while $\bigwedge_k \Sigma_i[\underline{A}_{ki}, \overline{A}_{ki}] \times x_i + \Sigma_i[\underline{A}'_{ki}, \overline{A}'_{ki}] \times x'_i \le c_k$ denotes the abstracted PC.

Example 7. Consider the self-composed program `piecewise_linear` in Example 5. Suppose we would like to compute the maximum output change, given the input perturbation $\delta = 0.1$ over x. The self-composed program includes four paths overall. Let us consider for example the path that takes the `else` branch in both the unprimed program P and the primed program P'. Since x, y are of `float` type, $\varepsilon_{\text{rel}} = 2^{-23}$ and $\varepsilon_{\text{abs}} = 2^{-149}$ for the 32-bit single precision floating-point format. We will have

$$\begin{aligned}
&\text{PC} : x \ge R^{\#}_{\mathbf{f},-\infty}(1.001) \wedge x' \ge R^{\#}_{\mathbf{f},-\infty}(1.001) \\
&\text{PF}_y : \left[1-2^{-23}, 1+2^{-23}\right] \\
&\qquad \times \left[R^{\#}_{\mathbf{f},-\infty}(1001.0), R^{\#}_{\mathbf{f},+\infty}(1001.0)\right] \times x \\
&\qquad + \left[-2^{-149}, 2^{-149}\right]
\end{aligned}$$

$$\begin{aligned}\mathrm{PF}_{y'} : &\left[1-2^{-23}, 1+2^{-23}\right] \\ &\times \left[R^{\#}_{\mathbf{f},-\infty}(1001.0), R^{\#}_{\mathbf{f},+\infty}(1001.0)\right] \times x' \\ &+ \left[-2^{-149}, 2^{-149}\right].\end{aligned} \tag{30}$$

Thus, we get the following interval linear programming problem:

$$\begin{aligned}\max \quad &\left[1-2^{-23}, 1+2^{-23}\right] \\ &\times \left[R^{\#}_{\mathbf{f},-\infty}(1001.0), R^{\#}_{\mathbf{f},+\infty}(1001.0)\right] \times x \\ &+ \left[-1-2^{-23}, -1+2^{-23}\right] \\ &\times \left[R^{\#}_{\mathbf{f},-\infty}(1001.0), R^{\#}_{\mathbf{f},+\infty}(1001.0)\right] \times x' \\ &+ \left[-2^{-148}, 2^{-148}\right]\end{aligned} \tag{31}$$

$$\begin{aligned}\text{s.t.} \quad &x - x' \le 0.1 \wedge -x + x' \le 0.1 \\ &\wedge -x \le -R^{\#}_{\mathbf{f},-\infty}(1.001) \wedge -x' \le -R^{\#}_{\mathbf{f},-\infty}(1.001).\end{aligned} \tag{32}$$

Solving the above interval linear programming problem by the method described in Section 2.2 will give us 100.10023889571252. After we deal with all other paths in the same way, we will find that 100.10023889571252 is the maximum output change with respect to the given input perturbation 0.1. Hence, the program `piecewise_linear` in floating-point arithmetic is at least $(0.1, 100.10023889571252)$-*robust*.

5. Implementation and Experimental Results

We have implemented a robustness analysis tool RAFP, based on the symbolic execution and floating-point abstraction techniques presented in Section 4. Given an input perturbation δ over one input variable of the program, RAFP can compute the maximum output change. Furthermore, if the user also provides a candidate output change ϵ and would like to check whether the program is (δ, ϵ)-*robust*, RAFP will check this property during the process of computing maximum output change and will stop once one path violating the property is found. Also, given an output change ϵ, RAFP can compute the maximum allowed input perturbation for floating-point programs. RAFP is built on top of Symbolic PathFinder (SPF) [23] which is a symbolic execution engine for Java programs. We use SPF to extract the path conditions together with the associated path functions. For linear programming, RAFP makes use of the Java Binding for GLPK (GNU Linear programming kit) called GLPK-Java [24].

To conduct experiments on checking robustness properties of floating-point programs via software model checking, we choose CBMC (C Bounded Model Checker) [14] which implements bounded model checking for ANSI-C programs using SAT/SMT solvers. CBMC utilizes a bit-precise modeling for floating-point operations and employs a sound and complete decision procedure for floating-point arithmetic. CBMC provides an option `--floatbv` to use IEEE floating point arithmetic and options for choosing rounding modes. However, CBMC does not support to use different rounding modes for the floating-point operations in the same program. In other words, all floating-point operations in a program are of the same rounding mode during the analysis. We use the default rounding mode `--round-to-nearest` during our experiments. Moreover, CBMC provides `__CPROVER_assume()` and `__CPROVER_assert()` statements, which are needed for robustness analysis of self-composed programs. Both statements take Boolean conditions. The `__CPROVER_assume()` statement restricts that the program traces should satisfy the assumed condition. For the `__CPROVER_assert()` statement, CBMC will check whether the asserted condition holds true for all runs of the program.

We have conducted experiments on a selection of benchmark examples using both RAFP and CBMC. Table 1 shows the comparison of performance and the resulting output changes. The column "δ_{in}" shows the considered input perturbation over one input variable of the program. The column "ϵ_{max}" shows the resulting maximum output change computed by RAFP with respect to the given input perturbation. The column "ϵ_{unr}" gives the largest possible output change that we have tried with CBMC such that the program is not $(\delta_{\mathrm{in}}, \epsilon_{\mathrm{unr}})$-*robust* with respect to the given input perturbation. The column "ϵ_r" gives the smallest output change that we have tried with CBMC such that the program is $(\delta_{\mathrm{in}}, \epsilon_r)$-*robust* with respect to the given input perturbation (Note that CBMC can be used only to check whether a program is (δ, ϵ)-*robust* and can not be used to compute the amount of output change with respect to the given input perturbation. During our experiments, we try CBMC with different candidate values of ϵ to find ϵ_{unr} and ϵ_r.). Since CBMC uses the same rounding mode for all floating-point operations in the same program during the analysis, the output change is always 0 when the given input perturbation is 0. Hence, for those rows that specify input perturbation as 0, we do not need to run CBMC and thus we mark the table entry with $\star$ in this case. Our tool RAFP utilizes rounding mode insensitive floating-point abstraction and thus in principle it holds that $\epsilon_{\mathrm{max}} \ge \epsilon_r \ge \epsilon_{\mathrm{unr}}$, which is confirmed by the experimental results.

The program `piecewise_linear` corresponds to the program shown in Example 4. `Max1`, `MorePaths` come from JPF Continuity [25]. `Max1` is a floating-point program that implements $\max(x, y)$, and thus it is $(\delta_{\mathrm{in}}, \delta_{\mathrm{in}})$-*robust*. `MorePaths` is a floating-point program that involves both a step function and a max function, and thus it is $(\delta_{\mathrm{in}}, 1.0)$-*robust* for all $\delta_{\mathrm{in}} \le 1.0$. `Orientation` (which corresponds to the program shown in Figure 1) together with `Filtered_Orientation` are extracted from the computational geometry algorithms library CGAL [26] and address robust geometric computation. `Filtered_Orientation` is an improved version of `Orientation` via static filter technique. The approximate result of computing the sign of a determinant

Table 1: Experimental results for benchmark examples.

Program	δ_{in}	RAFP		CBMC			
		$\epsilon_{\max}$	t (ms)	ϵ_{unr}	t (ms)	ϵ_r	t (ms)
piecewise_linear	0	$2.3889571220452006e-4$	29	$\star$	$\star$	$\star$	$\star$
	0.01	10.010238895712442	33	?	>1 h	?	>1 h
	0.1	100.10023889571252	34	?	>1 h	?	>1 h
Max1	0	$1.4012987984203268e-45$	44	$\star$	$\star$	$\star$	$\star$
	0.01	0.010000000000000007	55	0.0099	215	0.01	16066
	0.1	0.10000000000000006	54	0.099	227	0.1	16262
MorePaths	0	1.0000000027939686	59	$\star$	$\star$	$\star$	$\star$
	0.01	1.0000000027939686	79	0.99	279	1.0	379
	0.1	1.0000000027939686	86	0.99	314	1.0	471
Orientation	0	2.0	55	$\star$	$\star$	$\star$	$\star$
	$0.1e-8*\mathscr{E}$	2.0	68	0	1840	1.0	6845
	$0.1e-3*\mathscr{E}$	2.0	73	0	1338	1.0	13327
	$0.1e-2*\mathscr{E}$	2.0	80	1.0	14165	2.0	413
Filtered_Orientation	0	1.0	54	$\star$	$\star$	$\star$	$\star$
	$0.1e-8*\mathscr{E}$	1.0	70	0	1044	1.0	29641
	$0.1e-2*\mathscr{E}$	1.0	73	0	898	1.0	30261
	$0.1*\mathscr{E}$	2.0	75	0	719	1.0	26463
	$\mathscr{E}$	2.0	68	1.0	13206	2.0	654

is compared with a given positive filter bound $\mathscr{E}$ (rather than compared with zero). When the approximate result is in the interval $[-\mathscr{E}, \mathscr{E}]$, Filtered_Orientation gives 0. During our experiments, we set $\mathscr{E} = 1.5e-5$ (and for the sake of comparison, we express the input perturbation in terms of $\mathscr{E}$ also for Orientation although here $\mathscr{E}$ does not appear). The outputs of Orientation and Filtered_Orientation are always −1 (negative), 0 (zero), or 1 (positive). Hence, in Table 1, the resulting output changes for these two programs are always 0, 1.0, or 2.0. From Table 1, we could find that Filtered_Orientation is more robust than Orientation. For example, given the input perturbation $\delta = 0.1e-2*\mathscr{E}$, CBMC finds that for $\epsilon = 1.0$, Filtered_Orientation is robust while Orientation is not. Similarly, given the input perturbation $\delta = 0.1e-2*\mathscr{E}$, RAFP gives $\epsilon_{\max} = 1.0$ for Filtered_Orientation but gives $\epsilon_{\max} = 2.0$ for Orientation.

The column "t (ms)" presents the analysis times in milliseconds when the analyzers run on a 2.5 GHz PC with 4 GB of RAM running Windows 7. (RAFP runs further on a Java Virtual Machine (JVM) while CBMC runs further on a virtual machine VMWare running Fedora 12.) From Table 1, we could see that RAFP outperforms CBMC in time efficiency. Especially for piecewise_linear, CBMC could not even finish the analysis process in 1 hour. The low efficiency of CBMC is because that CBMC uses a sound and complete decision procedure for floating-point arithmetic. Especially, the multiplication and division floating-point operations may generate formulae that are expensive to decide and quite hard for SAT solvers to solve [27]. Hence, checking robustness properties of floating-point programs via CBMC may have limitations in scalability due to the current expensive decision procedures for floating-point logic. During our experiments, the approach via symbolic execution of abstracted floating-point programs is much more efficient. In principle, symbolic execution may suffer from the path explosion problem. However, the recent success of symbolic execution tools such as KLEE [28] on analyzing large-scale programs [17] makes this approach promising.

6. Related Work

6.1. Robustness Analysis of Programs. Robustness is a standard correctness property for control systems [1]. Robustness analysis of programs has received increasing attention in the recent years. Majumdar and Saha [5] took a first step toward analyzing the robustness of programs in control systems. They also utilized symbolic execution and optimization techniques to compute the maximum difference in program outputs with respect to the given input perturbation. However, they assumed exact real arithmetic in the program. Continuity as one aspect of robustness for software was firstly considered in [29]. Recently, Chaudhuri et al. presented logic-based mostly automated methods to determine whether a program is continuous [2] or Lipschitz continuous [3, 4], and more recently to determine whether a decision-making program is consistent under uncertainty [30]. Quite recently, Shahrokni and Feldt [31] conducted a systematic review of software robustness. However, much existing work on robustness analysis does not handle floating-point arithmetic in the program. Bushnell [25] presented a symbolic execution based approach to identify continuties and discontinuties associated with path condition boundaries for floating-point software, but it did not consider the true floating-point semantics. Besides, Gazeau et al. [32]

presented a nonlocal method for proving the robustness of floating-point programs but which needs much manual work. Recently, Goubault and Putot [33] proposed an abstract interpretation based robustness analysis method for finite precision implementations.

6.2. Safety Analysis of Floating Point Programs. Monniaux [7] described common pitfalls in analyzing and verifying floating-point programs. Abstract interpretation [34] based techniques have shown quite successful on analysis of floating-point programs. In [35], Goubault analyzed the origin of the loss of precision in floating-point programs based on abstract interpretation. Following this direction, a static analyzer FLUCTUAT [36] was developed. The abstract interpretation based static analyzer ASTRÉE [15] checks for floating-point run-time errors based on the computed set of reachable values for floating-point variables. As in ASTRÉE, we rely on the floating-point abstraction technique of [21] to soundly abstract floating-point expressions into ones over the field of reals. Chen et al. [37, 38] utilized interval linear constraints to design numerical abstract domains and to construct sound floating-point implementations [39]. Ivančić et al. [40] used bounded model checking based on SMT solvers to detect numerical instabilities in floating-point programs, based on a mixed integer-real model for floating-point variables and operations. Brain et al. [41] recently improved the bit-precise decision procedure for the theory of floating-point arithmetic based on a strict lifting of the conflict-driven clause learning algorithm in modern SAT solvers to abstract domains. Barr et al. [42] presented a method to automatically detect the floating-point exception through symbolic execution.

6.3. Self-Composition. The idea of self-composition is firstly used in the field of secure information flow [10, 11], to characterize noninterference. Terauchi and Aiken [12] proposed the type-directed transformation approach to make self-composition work in practice with off-the-shelf automatic safety analysis tools. Recently, Barthe et al. [43] proposed a general notion of product program that is beneficial to relational verification, which could be considered as the generalization of self-composition. Kovacs et al. [44] presented a general method to analyze 2-hypersafety properties by applying abstract interpretation on the self-compositions of the control flow graphs of programs.

7. Conclusion

We have proposed a self-composition based approach for robustness analysis of programs, which enables making use of off-the-shelf automatic reachability analysis tools to analyze robustness properties of programs. Then, we have shown how to use software model checking and symbolic execution techniques on self-composed programs to analyze program robustness properties. In particular, we have considered the robustness analysis problem of floating-point programs. To deal with floating-point arithmetic during symbolic execution, we have utilized a rounding mode insensitive floating-point abstraction to abstract floating-point expressions into interval linear expressions in exact real arithmetic. On this basis, the maximum output change (when given the input perturbation) or maximum allowed input perturbation (when given the input perturbation) are computed based on symbolic execution and interval linear programming for abstracted floating-point programs. Experimental results of our prototype implementation are encouraging.

It remains for future work to exploit the intrinsic symmetry of self-composed programs to reduce the number of considered paths during robustness analysis. We also plan to improve the prototype implementation and to conduct more experiments on larger realistic floating-point programs.

Conflict of Interests

The authors declare that there is no conflict of interests regarding the publication of this paper.

Acknowledgments

This work is supported by the 973 Program under Grant no. 2014CB340703, the 863 Program under Grant no. 2011AA010106, the NSFC under Grant nos. 61202120, 61120106006, and 91318301 and the SRFDP under Grant no. 20124307120034.

References

[1] S. Pettersson and B. Lennartson, "Stability and robustness for hybrid systems," in *Proceedings of the 35th IEEE Conference on Decision and Control*, pp. 1202–1207, December 1996.

[2] S. Chaudhuri, S. Gulwani, and R. Lublinerman, "Continuity analysis of programs," in *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '10)*, pp. 57–69, ACM, January 2010.

[3] S. Chaudhuri, S. Gulwani, and R. Lublinerman, "Continuity and robustness of programs," *Communications of the ACM*, vol. 55, no. 8, pp. 107–115, 2012.

[4] S. Chaudhuri, S. Gulwani, R. Lublinerman, and S. NavidPour, "Proving programs robust," in *Proceedings of the 19th ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE '11)*, pp. 102–112, ACM, 2011.

[5] R. Majumdar and I. Saha, "Symbolic robustness analysis," in *Proceedings of the Real-Time Systems Symposium (RTSS '09)*, pp. 355–363, IEEE, December 2009.

[6] L. Kettner, K. Mehlhorn, S. Pion, S. Schirra, and C. K. Yap, "Classroom examples of robustness problems in geometric computations," in *Proceedings of the European Symposium on Algorithms (ESA '04)*, vol. 3221 of *Lecture Notes in Computer Science*, pp. 702–713, Springer, 2004.

[7] D. Monniaux, "The pitfalls of verifying floating-point computations," *ACM Transactions on Programming Languages and Systems*, vol. 30, no. 3, article 12, 2008.

[8] IEEE Computer Society, "IEEE standard for binary floating point arithmetic," Tech. Rep. ANSI/IEEE Std 745-1985, 1985.

[9] J. Rohn, "Solvability of systems of interval linear equations and inequalities," in *Linear Optimization Problems with Inexact Data*, pp. 35–77, Springer, 2006.

[10] G. Barthe, P. R. D'Argenio, and T. Rezk, "Secure information flow by self-composition," in *Proceedings of the 17th IEEE Computer Security Foundations Workshop (CSFW '04)*, pp. 100–114, IEEE, June 2004.

[11] Á. Darvas, R. Hahnle, and D. Sands, "A theorem proving approach to analysis of secure information flow," in *Proceedings of the 2nd International Conference on Security in Pervasive Computing (SPC '05)*, vol. 3450 of *Lecture Notes in Computer Science*, pp. 193–209, Springer, 2005.

[12] T. Terauchi and A. Aiken, "Secure information flow as a safety problem," in *Proceedings of the International Static Analysis Symposium (SAS '05)*, vol. 3672 of *Lecture Notes in Computer Science*, pp. 352–367, Springer, 2005.

[13] T. Ball and S. K. Rajamani, "The SLAM project: debugging system software via static analysis," in *Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '02)*, pp. 1–3, ACM Press, January 2002.

[14] E. M. Clarke, D. Kroening, and F. Lerda, "A tool for checking ANSI-C programs," in *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS '04)*, vol. 2988 of *Lecture Notes in Computer Science*, pp. 168–176, Springer, 2004.

[15] B. Blanchet, P. Cousot, R. Cousot et al., "A static analyzer for large safety-critical software," in *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '03)*, pp. 196–207, ACM Press, June 2003.

[16] R. Jhala and R. Majumdar, "Software model checking," *ACM Computing Surveys*, vol. 41, no. 4, article 21, 2009.

[17] C. Cadar and K. Sen, "Symbolic execution for software testing: three decades later," *Communications of the ACM*, vol. 56, no. 2, pp. 82–90, 2013.

[18] J. C. King, "Symbolic execution and program testing," *Communications of the ACM*, vol. 19, no. 7, pp. 385–394, 1976.

[19] A. Brillout, D. Kroening, and T. Wahl, "Mixed abstractions for floating-point arithmetic," in *Proceedings of the 9th International Conference Formal Methods in Computer Aided Design (FMCAD '09)*, pp. 69–76, IEEE, November 2009.

[20] L. Haller, A. Griggio, M. Brain, and D. Kroening, "Deciding floating-point logic with systematic abstraction," in *Proceedings of the International Conference Formal Methods in Computer Aided Design (FMCAD '12)*, pp. 131–140, IEEE, 2012.

[21] A. Miné, "Relational abstract domains for the detection of floating-point run-time errors," in *Proceedings of the European Symposium on Programming (ESOP '04)*, vol. 2986 of *Lecture Notes in Computer Science*, pp. 3–17, Springer, 2004.

[22] A. Miné, *Weakly relational numerical abstract domains [Ph.D. thesis]*, Ecole Polytechnique, Palaiseau, France, 2004.

[23] C. S. Pasareanu, W. Visser, D. H. Bushnell, J. Geldenhuys, P. C. Mehlitz, and N. Rungta, "Symbolic pathfinder: integrating symbolic execution with model checking for java bytecode analysis," *Automated Software Engineering*, vol. 20, no. 3, pp. 391–425, 2013.

[24] H. Schuchardt, "GLPK for Java," 2014, http://glpk-java.sourceforge.net/.

[25] D. Bushnell, "Continuity analysis for floating point software," in *Proceedings of the 4th workshop on Numerical Software Verification (NSV '11)*, 2011.

[26] E. Fogel and M. Teillaud, "The computational geometry algorithms library cgal," *ACM Communications in Computer Algebra*, vol. 47, no. 3, pp. 85–87, 2013.

[27] D. Kroening, "The CPROVER User Manual," http://www.cprover.org/cbmc/doc/manual.pdf.

[28] C. Cadar, D. Dunbar, and D. R. Engler, "KLEE: unassisted and automatic generation of high-coverage tests for complex systems programs," in *Proceedings of the Symposium on Operating Systems Design and Implementation (OSDI '14)*, pp. 209–224, USENIX Association, 2008.

[29] D. Hamlet, "Continuity in software systems," in *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA '02)*, pp. 196–200, ACM, July 2002.

[30] S. Chaudhuri, A. Farzan, and Z. Kincaid, "Consistency analysis of decision-making programs," in *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14)*, pp. 555–568, Springer, 2014.

[31] A. Shahrokni and R. Feldt, "A systematic review of software robustness," *Information & Software Technology*, vol. 55, no. 1, pp. 1–17, 2013.

[32] I. Gazeau, D. Miller, and C. Palamidessi, "A non-local method for robustness analysis of floating point programs," in *Proceedings of the Workshop on Quantitative Aspects of Programming Languages (QAPL '12)*, vol. 85 of *Electronic Proceedings in Theoretical Computer Science*, pp. 63–76, 2012.

[33] E. Goubault and S. Putot, "Robustness analysis of finite precision implementations," in *Proceedings of the Asian Symposium on Programming Languages and Systems (APLAS '13)*, vol. 8301 of *Lecture Notes in Computer Science*, pp. 50–57, Springer, 2013.

[34] P. Cousot and R. Cousot, "Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints," in *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '77)*, pp. 238–252, ACM, 1977.

[35] E. Goubault, "Static analyses of the precision of floating-point operations," in *Proceedings of the International Static Analysis Symposium (SAS '01)*, vol. 2126 of *Lecture Notes in Computer Science*, pp. 234–259, Springer, 2001.

[36] E. Goubault, M. Martel, and S. Putot, "Asserting the precision of floating-point computations: a simple abstract interpreter," in *Proceedings of the European Symposium on Programming (ESOP '02)*, vol. 2305 of *Lecture Notes in Computer Science*, pp. 209–212, Springer, 2002.

[37] L. Chen, A. Miné, J. Wang, and P. Cousot, "Interval polyhedra: an abstract domain to inferinterval linear relationships," in *Proceedings of the International Static Analysis Symposium (SAS '09)*, vol. 5673 of *Lecture Notes in Computer Science*, pp. 309–325, Springer, 2009.

[38] L. Chen, A. Miné, J. Wang, and P. Cousot, "An abstract domain to discover interval linear equalities," in *Proceedings of the International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI '10)*, vol. 5944 of *Lecture Notes in Computer Science*, pp. 112–128, Springer, 2010.

[39] L. Chen, A. Miné, and P. Cousot, "A sound floating-point polyhedra abstract domain," in *Proceedings of the Asian Symposium on Programming Languages and Systems (APLAS '08)*, vol. 5356 of *Lecture Notes in Computer Science*, pp. 3–18, Springer, 2008.

[40] F. Ivančić, M. K. Ganai, S. Sankaranarayanan, and A. Gupta, "Numerical stability analysis of floating-point computations using software model checking," in *Proceedings of the 8th ACM/IEEE International Conference on Formal Methods and*

Models for Codesign (MEMOCODE '10), pp. 49–58, IEEE, July 2010.

[41] M. Brain, V. D'Silva, A. Griggio, L. Haller, and D. Kroening, "Deciding floating-point logic with abstract conflict driven clause learning," in *Formal Methods in System Design*, 2013.

[42] E. T. Barr, T. Vo, V. Le, and Z. Su, "Automatic detection of floating-point exceptions," in *Proceedings of the SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '13)*, pp. 549–560, ACM, 2013.

[43] G. Barthe, J. M. Crespo, and C. Kunz, "Relational verification using product programs," in *Proceedings of the 17th International Symposium on Formal Methods (FM '11)*, pp. 200–214, Springer, 2011.

[44] M. Kovacs, H. Seidl, and B. Finkbeiner, "Relational abstract interpretation for the verification of 2-hypersafety properties," in *Proceedings of the ACM Conference on Computer and Communications Security (CCS '13)*, pp. 211–222, 2013.