

Invertibility Conditions for Floating-Point Formulas^{*}

Martin Brain³, Aina Niemetz¹, Mathias Preiner¹, Andrew Reynolds²,
Clark Barrett¹, and Cesare Tinelli²

¹ Stanford University

² The University of Iowa

³ University of Oxford and City, University of London

Abstract. Automated reasoning procedures are essential for a number of applications that involve bit-exact floating-point computations. This paper presents conditions that characterize when a variable in a floating-point constraint has a solution, which we call invertibility conditions. We describe a novel workflow that combines human interaction and a syntax-guided synthesis (SyGuS) solver that was used for discovering these conditions. We verify our conditions for several floating-point formats. One implication of this result is that a fragment of floating-point arithmetic admits compact quantifier elimination. We implement our invertibility conditions in a prototype extension of our solver CVC4, showing their usefulness for solving quantified constraints over floating-points.

1 Introduction

Satisfiability Modulo Theories (SMT) formulas including either the theory of bit-exact floating-point [13] or universal quantifiers [26,34] are widely regarded as some of the hardest to solve. Problems that combine the two (i.e. universal quantification over floating-point sorts) are rare – experience to date has suggested these problems are hard for solvers and would-be users should either give up or develop their own incomplete techniques. However, progress in theory solvers for floating-point [11] and the use of expression synthesis for handling universal quantifiers [31,29] suggest that these problems may not be entirely out of reach after all, which could potentially impact a number of interesting applications.

This paper makes substantial progress towards a scalable approach for solving quantified floating-point constraints directly in an SMT solver. Building on the concept of *invertibility conditions* introduced in [28] and a state-of-the-art expression synthesis engine [31], we generate invertibility conditions for the functions and predicates of the theory of floating-point. These can be used for quantifier elimination or in local-search techniques. Specifically the paper contributes the following:

- In Section 3, we present invertibility conditions for a large subset of floating-point operators and predicates in the SMT-LIB standard for floating-points.
- In Section 4, we present a custom methodology based on syntax-guided synthesis and decision tree learning that we developed for the purpose of synthesizing the invertibility conditions presented in this work.

^{*} This work was partially supported by the National Science Foundation under award 1656926.

- In Section 5, we present a quantifier elimination procedure for a fragment of floating-point formulas that is based on invertibility conditions and give experimental evidence of its potential, based on quantified floating-point problems coming from a verification application.

Motivation Our goal in this paper is to develop techniques for solving quantified floating-point formulas. We briefly mention two potential applications of quantified floating-point constraints.

A common task in hardware design and high-performance computing [18] is to find a simple, tight, polynomial approximation of a more costly expression f over a bounded range b . Given a polynomial $p(x)$ with co-efficients c_n, \dots, c_0 , this can be stated as: $\exists c_n, \dots, c_0. \forall x. b(x) \Rightarrow -\epsilon < f(x) - p(x) < \epsilon$.

This is structurally similar to another application, the challenge of generating inductive invariants in template-based abstract interpreters [10]. For example, given a transition system over two variables with initial condition I , transition relation T , a safety property P and an invariant template $a \cdot x + b \cdot y < c$, safety checking can be reduced to solving:

$$\begin{aligned} \exists a, b, c. \forall x, y, x', y'. \quad & I(x, y) \Rightarrow a \cdot x + b \cdot y < c \\ & \wedge (a \cdot x + b \cdot y < c \wedge T(x, y, x', y')) \Rightarrow a \cdot x' + b \cdot y' < c \\ & \wedge a \cdot x + b \cdot y < c \Rightarrow P(x, y) \end{aligned}$$

Solving this formula directly using an SMT solver (rather than following the classical approach of implementing abstract domains) would allow for greater precision and on-the-fly abstraction refinement. It would also enable non-linear and polynomial abstraction templates.

Related Work To our knowledge, no previous work specifically discusses techniques for solving universally quantified floating-point formulas. In [11], Brain et al. provide a comprehensive review of decision procedures for quantifier-free bit-exact floating-point using both SMT-based as well as other approaches. They identify four groups of techniques; bit-blasting approaches that use floating-point circuits to generate bit-vector formulas [14,17,22,35]; interval techniques that use partitioning and interval propagation [9,33,25,24]; optimization and numerical approaches that work with complete valuations [3,20,6,23]; and axiomatic techniques that use partial or total axiomatizations of the theory of floating-point into other theories, often reals [16,15].

On the other hand, approaches for universal quantification have been developed in modern SMT solvers that target other background theories including linear arithmetic [31,19,7] and bit-vectors [34,29,28]. At a high level, these approaches use model-based refinement loops that lazily add instances of universal quantifiers until they reach a conflict at the quantifier-free level, or otherwise saturate with a model.

Our work is partly based on an adaptation of the approach of Niemetz et. al. [28], to the theory of floating-points, showing that invertibility conditions exist and show promise for solving quantified floating-point constraints.

2 Preliminaries

We assume the usual notions and terminology of many-sorted first-order logic with equality (denoted by \approx). Let Σ be a *signature* consisting of a set Σ^s of sort symbols and a set Σ^f of interpreted (and sorted) function symbols. Each function symbol f has a sort $\tau_1 \times \dots \times \tau_n \rightarrow \tau$, with arity $n \geq 0$ and $\tau_1, \dots, \tau_n, \tau \in \Sigma^s$. We assume that Σ includes a Boolean sort `Bool` and the Boolean constants \top (true) and \perp (false). We further assume the usual definition of well-sorted terms, literals, and (quantified) formulas with variables and symbols from Σ , and refer to them as Σ -terms, Σ -atoms, and so on. For a Σ -term or Σ -formula e , we denote the *free variables* of e (defined as usual) as $FV(e)$ and use $e[x]$ to denote that the variable x occurs free in e . We write $e[t]$ for the term or formula obtained from e by replacing each occurrence of x in e by t .

A *theory* T is a pair (Σ, I) , where Σ is a signature and I is a non-empty class of Σ -interpretations (the *models* of T) that is closed under variable reassignment, i.e., every Σ -interpretation that only differs from an $\mathcal{I} \in I$ in how it interprets variables is also in I . A Σ -formula φ is *T-satisfiable* (resp. *T-unsatisfiable*) if it is satisfied by some (resp. no) interpretation in I ; it is *T-valid* if it is satisfied by all interpretations in I . We will sometimes omit T when the theory is understood from context.

We briefly recap the terminology and notation of Brain et al. [13] which defines an SMT-LIB theory of floating-point numbers T_{FP} based on IEEE-754 2008 [2].⁴ The signature of T_{FP} includes a parametric family of sorts $\mathbb{F}_{\varepsilon, \sigma}$ where ε and σ are integers greater than or equal to 2 giving the number of bits used to store the exponent and significand, respectively. Each of these sorts contains five kinds of constants; normal numbers of the form $1.s * 2^e$, subnormal numbers of the form $0.s * 2^{-2^{\sigma-1}-1}$, two zeros ($+0$ and -0), two infinities ($+\infty$ and $-\infty$) and a single not-a-number (NaN). We assume a map $v_{\varepsilon, \sigma}$ for each sort, which gives the value that each of these constants represents in terms of $\mathbb{R}^* = \mathbb{R} \cup \{+\infty, -\infty, \text{NaN}\}$. The theory also provides a rounding-mode sort `RM`, which contains five elements $\{\text{RNE}, \text{RNA}, \text{RTP}, \text{RTN}, \text{RTZ}\}$.

Theory T_{FP} contains a full set of arithmetic operations $\{|\cdot|, +, -, \cdot, \div, \sqrt{\cdot}, \max, \min\}$ as well as `rem` (remainder), `rtoi` (round to integral) and `fma` (combined multiply and add with just one rounding). The precise semantics of these operations are given in [13] and they all follow the same general pattern: $v_{\varepsilon, \sigma}$ is used to project the arguments to \mathbb{R}^* , the normal arithmetic is performed in \mathbb{R}^* , then the rounding mode and the result are used to select one of the adjoints of $v_{\varepsilon, \sigma}$ to convert the result back to $\mathbb{F}_{\varepsilon, \sigma}$.

Theory T_{FP} further defines a set of ordering predicates $\{<, >, \leq, \geq\}$ and a set of classification predicates $\{\text{isNorm}, \text{isSub}, \text{isInf}, \text{isZero}, \text{isNaN}, \text{isNeg}, \text{isPos}\}$. In the following, we denote the rounding mode of an operation above the operator symbol, e.g., $a \overset{\text{RTZ}}{+} b$ adds a and b and rounds the result towards zero. We use the infix operator style for `isInf` ($\dots \approx \pm\infty$), `isZero` ($\dots \approx \pm 0$), and `isNaN` ($\dots \approx \text{NaN}$) for conciseness. We further use $\text{min}_n/\text{max}_n$ and $\text{min}_s/\text{max}_s$ for floating-point constants representing the minimum/maximum normal and subnormal numbers, respectively. We will omit rounding mode and floating-point sorts if they are clear from the context.

⁴ The full theory in [13] includes several additional operators which we omit from discussion here, such as floating-point minimum/maximum, equality with floating-point semantics (`fp.eq`), and conversions between sorts.

Symbol	SMT-LIB Syntax	Sort
isNorm, isSub	fp.isNormal, fp.isSubnormal	$\mathbb{F}_{\varepsilon, \sigma} \rightarrow \text{Bool}$
isPos, isNeg	fp.isPositive, fp.isNegative	$\mathbb{F}_{\varepsilon, \sigma} \rightarrow \text{Bool}$
isInf, isNaN, isZero	fp.isInfinite, fp.isNaN, fp.isZero	$\mathbb{F}_{\varepsilon, \sigma} \rightarrow \text{Bool}$
$\approx, <, >, \leq, \geq$	=, fp.lt, fp.gt, fp.leq, fp.geq	$\mathbb{F}_{\varepsilon, \sigma} \times \mathbb{F}_{\varepsilon, \sigma} \rightarrow \text{Bool}$
$ \dots , -$	fp.abs, fp.neg	$\mathbb{F}_{\varepsilon, \sigma} \rightarrow \mathbb{F}_{\varepsilon, \sigma}$
rem	fp.rem	$\mathbb{F}_{\varepsilon, \sigma} \times \mathbb{F}_{\varepsilon, \sigma} \rightarrow \mathbb{F}_{\varepsilon, \sigma}$
$\sqrt{\cdot}$, rti	fp.sqrt, fp.roundToIntegral	$\text{RM} \times \mathbb{F}_{\varepsilon, \sigma} \rightarrow \mathbb{F}_{\varepsilon, \sigma}$
$+$, $-$, \cdot , \div	fp.add, fp.sub, fp.mul, fp.div	$\text{RM} \times \mathbb{F}_{\varepsilon, \sigma} \times \mathbb{F}_{\varepsilon, \sigma} \rightarrow \mathbb{F}_{\varepsilon, \sigma}$
fma	fp.fma	$\text{RM} \times \mathbb{F}_{\varepsilon, \sigma} \times \mathbb{F}_{\varepsilon, \sigma} \times \mathbb{F}_{\varepsilon, \sigma} \rightarrow \mathbb{F}_{\varepsilon, \sigma}$

Table 1: Considered floating-point predicates/operators with SMT-LIB 2 syntax

3 Invertibility Conditions for Floating-Point Formulas

In this section, we adapt the concept of invertibility conditions introduced by Niemetz et al. [28] to theory T_{FP} . Intuitively, an invertibility condition ϕ_c for a literal $l[x]$ is the exact condition under which $l[x]$ has a solution for x , i.e., $\phi_c \Leftrightarrow \exists x. l[x]$.

Definition 1. (*Floating-Point Invertibility Condition*) Let $l[x]$ be a Σ_{FP} -literal. A quantifier-free Σ_{FP} -formula ϕ_c is an invertibility condition for x in $l[x]$ if $x \notin FV(\phi_c)$ and $\phi_c \Leftrightarrow \exists x. l[x]$ is T_{FP} -valid.

As a simple example of an invertibility condition, given literal $|x| \approx t$ where $|x|$ denotes the absolute value of x , a solution for x exists only if t is not negative, i.e., if $\text{isNeg}(t)$ holds. We introduce additional terminology for the sake of the discussion. We define the *dimension* of an invertibility condition problem $\exists x. l[x]$ as the number of free variables it contains. For example, if s and t are variables, then the dimension of $\exists x. x + s \approx t$ is two, the dimension of $\exists x. \text{isZero}(x + s)$ is one, and the dimension of $\exists x. \text{isZero}(|x|)$ is zero. A literal $l[x]$ is *fully invertible* if its invertibility condition is \top . A term e is an (unconditional) *inverse* for x in $l[x]$ if $l[e]$ is equivalent to \top . For example, the literal $-x \approx t$ is fully invertible and $-t$ is an inverse for x in this literal. We say that e is a *conditional inverse* for $l[x]$ if $l[e]$ is an invertibility condition for $l[x]$.

Our primary goal in this work is to establish invertibility conditions for all floating-point constraints that contain exactly one operator and one predicate. These conditions collectively suffice to characterize when any literal $l[x]$ containing exactly one occurrence of x , the variable to solve for, has a solution. In total, we were able to establish 167 out of 188 invertibility conditions (counting commutative cases only once) using a syntax-guided synthesis framework which we describe in more detail in Section 4. In this section, we present a subset of these invertibility conditions, highlighting the most interesting cases where we succeeded (or failed) to establish an invertibility condition. Due to space restrictions, we omit the conditions for the remaining cases⁵.

Table 2 lists the invertibility conditions for equality with the operators $\{+, -, \cdot, \div, \text{rem}, \sqrt{\cdot}, |\dots|, -, \text{rti}\}$, parameterized over a rounding mode R (one of RNE, RNA,

⁵ Note to reviewers: all omitted invertibility conditions can be found in the appendix.

Literal Invertibility Condition	
$x + s \approx t$	$t \approx (t - s) + s \vee t \approx (t - s) + s \vee s \approx t$
$x - s \approx t$	$t \approx (s + t) - s \vee t \approx (s + t) - s \vee (s \neq t \wedge s \approx \pm\infty \wedge t \approx \pm\infty)$
$s - x \approx t$	$t \approx s + (t - s) \vee t \approx s + (t - s) \vee s \approx t$
$x \cdot s \approx t$	$t \approx (t \div s) \cdot s \vee t \approx (t \div s) \cdot s \vee (s \approx \pm\infty \wedge t \approx \pm\infty) \vee (s \approx \pm 0 \wedge t \approx \pm 0)$
$x \div s \approx t$	$t \approx (s \cdot t) \div s \vee t \approx (s \cdot t) \div s \vee (s \approx \pm\infty \wedge t \approx \pm 0) \vee (t \approx \pm\infty \wedge s \approx \pm 0)$
$s \div x \approx t$	$t \approx s \div (s \div t) \vee t \approx s \div (s \div t) \vee (s \approx \pm\infty \wedge t \approx \pm\infty) \vee (s \approx \pm 0 \wedge t \approx \pm 0)$
$x \text{ rem } s \approx t$	$t \approx t \text{ rem } s$
$s \text{ rem } x \approx t$?
$\sqrt{x} \approx t$	$t \approx \sqrt{(t \cdot t)} \vee t \approx \sqrt{(t \cdot t)} \vee t \approx \pm 0$
$ x \approx t$	$\neg \text{isNeg}(t)$
$-x \approx t$	\top
$\text{rti}(x) \approx t$	$t \approx \text{rti}(t)$

Table 2: Invertibility conditions for floating-point operators (excl. fma) with \approx .

RTP, RTN, or RTZ). Note that operators $\{+, \cdot\}$ and the multiplicative step of fma are commutative, and thus the invertibility conditions for both variants are identical, and only one variant is considered. Each of the first six invertibility conditions in this table follow a pattern. The first two disjuncts are instances of the literal to solve for (for a term involving rounding modes RTP and RTN), followed by disjuncts for handling special cases for infinity and zero. Based on extensive runs of our syntax-guided synthesis procedure, we believe these solutions are close to minimal in term size. This indicates that these literals are not fully invertible, nor do they have simple conditional inverses. In particular, notice that it is not the case that either of the two disjuncts is an invertibility condition. For example, for addition with rounding mode RNE, neither $t \approx (t - s) + s$ nor $t \approx (t - s) + s$ is an invertibility condition. Neither is $t \approx (t - s) + s$ for *all* rounding modes R. The condition given in Table 2, however, exactly characterizes when a solution for this case exists.

We found that t is a conditional inverse for the case of $\text{rti}(x) \approx t$ and $x \text{ rem } s \approx t$, that is, substituting t for x is an invertibility condition. For the latter, we discovered an alternative invertibility condition:

$$|t + t| \leq |s| \vee |t + t| \leq |s| \vee \text{ite}(t \approx \pm 0, s \neq \pm 0, t \neq \pm \infty) \quad (1)$$

In contrast to the condition from Table 2, this version does not involve rem. This implies that certain applications of floating-point remainder, including those whose first argument is an unconstrained variable, can be eliminated based on this equivalence. Interestingly, for $s \text{ rem } x \approx t$, we did not succeed in finding an invertibility condition. This case appears to not admit a concise solution, we discuss further details below.

Table 3 gives the invertibility conditions for \geq . Since these constraints admit more solutions, they typically have simpler invertibility conditions. In particular, with the exception of rem, all conditions only involve floating-point classifiers.

Literal Invertibility Condition	
$x \overset{R}{+} s \geq t$	$(\text{isPos}(s) \vee \text{ite}(s \approx \pm\infty, (t \approx \pm\infty \wedge \text{isNeg}(t)), \text{isNeg}(s))) \wedge t \not\approx \text{NaN}$
$x \overset{R}{-} s \geq t$	$\text{ite}(\text{isNeg}(s), t \not\approx \text{NaN}, \text{ite}(s \approx \pm\infty, (t \approx \pm\infty \wedge \text{isNeg}(t)), (\text{isPos}(s) \wedge t \not\approx \text{NaN})))$
$s \overset{R}{-} x \geq t$	$(\text{isPos}(s) \vee \text{ite}(s \approx \pm\infty, (t \approx \pm\infty \wedge \text{isNeg}(t)), \text{isNeg}(s))) \wedge t \not\approx \text{NaN}$
$x \overset{R}{\cdot} s \geq t$	$(\text{isNeg}(t) \vee t \approx \pm 0 \vee s \not\approx \pm 0) \wedge s \not\approx \text{NaN} \wedge t \not\approx \text{NaN}$
$x \overset{R}{\div} s \geq t$	$(\text{isNeg}(t) \vee t \approx \pm 0 \vee s \not\approx \pm 0) \wedge s \not\approx \text{NaN} \wedge t \not\approx \text{NaN}$
$s \overset{R}{\div} x \geq t$	$(\text{isNeg}(t) \vee t \approx \pm 0 \vee s \not\approx \pm 0) \wedge s \not\approx \text{NaN} \wedge t \not\approx \text{NaN}$
$x \text{ rem } s \geq t$	$\text{ite}(\text{isNeg}(t), s \not\approx \text{NaN}, (t + t \overset{\text{RNE}}{\leq} s \wedge t \not\approx \pm\infty)) \wedge s \not\approx \pm 0$
$s \text{ rem } x \geq t$?
$\sqrt{x} \geq t$	$t \not\approx \text{NaN}$
$ x \geq t$	$t \not\approx \text{NaN}$
$-x \geq t$	$t \not\approx \text{NaN}$
$\overset{R}{\text{rti}}(x) \geq t$	$t \not\approx \text{NaN}$

Table 3: Invertibility conditions for floating-point operators (excl. fma) with \geq .

When considering literals with predicates, the invertibility conditions for cases involving $x + s$ and $s - x$ are identical for every predicate and rounding mode. This is due to the fact that $s - x \equiv s + (-x)$, independent from the rounding mode. Thus, the negation of the inverse value of x for an equation involving $x + s$ is the inverse value of x for an equation involving $s - x$. Similarly, the invertibility conditions for $x \cdot s$ and $s \div x$ over predicates $\{<, \leq, >, \geq, \text{isInf}, \text{isNaN}, \text{isNeg}, \text{isZero}\}$ are identical for all rounding modes.

For all predicates except $\{\approx, \text{isNorm}, \text{isSub}\}$, the invertibility conditions for operators $\{+, -, \div, \cdot\}$ contain floating-point classifiers only. All of these conditions are also independent from the rounding mode. Similarly, for operator fma over predicates $\{\text{isInf}, \text{isNaN}, \text{isNeg}, \text{isPos}\}$, the invertibility conditions contain only floating-point classifiers. All of these conditions except for $\text{isNeg}(\text{fma}(x, s, t))$ and $\text{isPos}(\text{fma}(x, s, t))$ are also independent from the rounding mode.

For all floating-point operators with predicate isNaN, the invertibility condition is \top , i.e., an inverse value for x always exists. This is due to the fact that every floating-point operator returns NaN if one of its operands is NaN, hence NaN can be picked as an inverse value of x . Conversely, we identified four cases for which the invertibility condition is \perp , i.e., an inverse value for x never exists. These four cases are $\text{isNeg}(|x|)$, $\text{isInf}(x \text{ rem } s)$, $\text{isInf}(s \text{ rem } x)$, and $\text{isSub}(\text{rti}(x))$. For the first three cases, it is obvious why no inverse value exists. The intuition for $\text{isSub}(\text{rti}(x))$ is that integers are not subnormal, and as a result if x is rounded to an integer it can never be a subnormal number. All of these cases can be easily implemented as rewrite rules in an SMT solver.

For operator fma, the invertibility conditions over predicates $\{\text{isInf}, \text{isNaN}, \text{isNeg}, \text{isPos}\}$ contain floating-point classifiers only. For predicate isZero, the invertibility conditions are more involved. Equations (2) and (3) show the invertibility conditions for $\text{isZero}(\text{fma}(x, s, t))$ and $\text{isZero}(\text{fma}(s, t, x))$ for all rounding modes R.

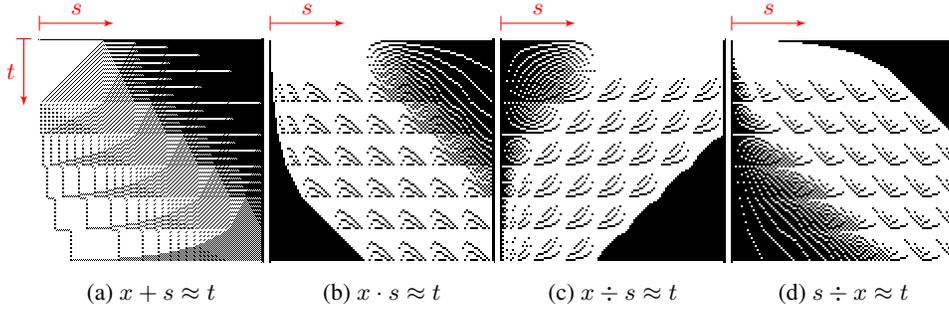


Fig. 1: Invertibility conditions for $\{+, \cdot, \div\}$ over \approx for $\mathbb{F}_{3,5}$ and rounding mode RNE.

$$\text{fma}^R(-(t \overset{\text{RTP}}{\div} s), s, t) \approx \pm 0 \vee \text{fma}^R(-(t \overset{\text{RTN}}{\div} s), s, t) \approx \pm 0 \vee (s \approx \pm 0 \wedge t \approx \pm 0) \quad (2)$$

$$\text{fma}^R(s, t, -(s \overset{\text{RTP}}{\cdot} t)) \approx \pm 0 \vee \text{fma}^R(s, t, -(s \overset{\text{RTN}}{\cdot} t)) \approx \pm 0 \quad (3)$$

These two invertibility conditions contain case splits similar to those in Table 2 and indicate that, e.g., $-t \overset{\text{RTP}}{\div} s$ is an inverse value for x when $\text{fma}^R(-(t \overset{\text{RTP}}{\div} s), s, t) \approx \pm 0$ holds.

As we will describe in Section 4, an important aspect of synthesizing these invertibility conditions was considering their visualizations. This helped us determine which invertibility conditions were relatively simple and which exhibited complex behavior.

Figure 1 shows the visualizations of the invertibility conditions for operators $\{+, \cdot, \div\}$ over \approx from Table 2 for sort $\mathbb{F}_{3,5}$ with rounding mode RNE (each of the literals is two-dimensional). We use 227×227 pixel maps over all possible values of s and t , where the pixel at point (s, t) is white if the invertibility condition is true, and black if it is false.⁶ The values of s are plotted on the horizontal axis and the values of t are plotted on the vertical axis. The leftmost two columns (resp. topmost two rows) give the value of the invertibility condition for $s = \pm 0$ (resp. $t = \pm 0$); the rightmost column (resp. bottom row) gives its value for NaN; the next two columns left of (resp. next two rows on top of) NaN give its value for $\pm\infty$; the remainder plots the values of the subnormal and normal values of s and t , left-to-right (resp. top-to-bottom) in increasing order of their absolute value, alternating positive and negative values. These visualizations give an intuition of the complexity of the behavior of invertibility conditions, which is a consequence of the complex semantics of floating-point operations.

Figure 2 gives the invertibility condition visualizations for remainder over \approx with sort $\mathbb{F}_{3,5}$ and rounding mode RNE. The visualization on the left hand shows that solving for x as the first argument is relatively easy. It suggests that an invertibility condition for this case involves a linear inequality relating the absolute values of s and t , which we were able to derive in Equation (1). Solving for x as the second argument, on the other hand, is much more difficult, as indicated by the right picture, which has a significantly more complex structure. We conjecture that no simple solution exists for the latter problem. The visualization of the invertibility condition gives some of the intuition for this: The diagonal divide is caused by the fact that output t will always have a smaller absolute value than the input s . The top-left corner represents subnormal/subnormal computation, this acts as fixed-point and behaves differently from the rest of

⁶ Notice that we consider all possible $(2^{\sigma-1} - 1) * 2$ NaN values as one single NaN value. Thus, for sort $\mathbb{F}_{3,5}$ we have 227 floating-point values (instead of $2^8 = 256$).

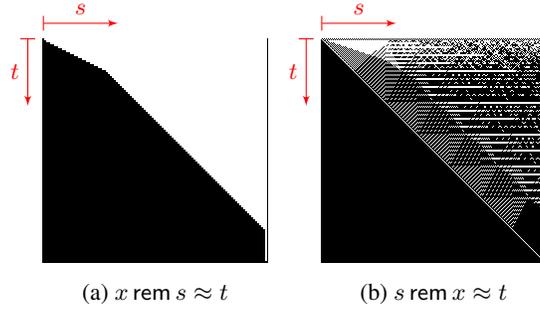


Fig. 2: Invertibility conditions for rem over \approx for $\mathbb{F}_{3,5}$.

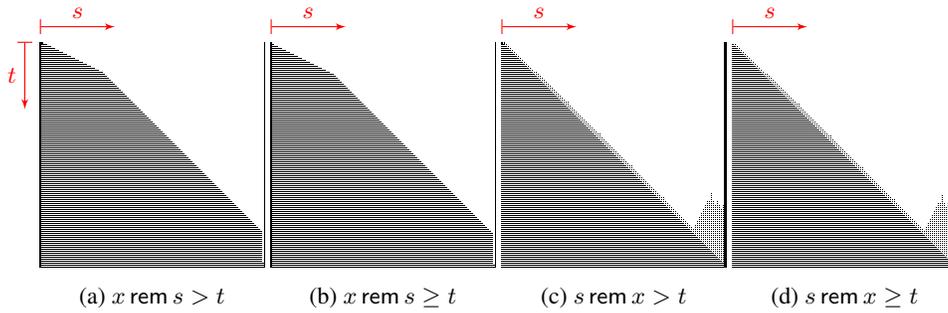


Fig. 3: Invertibility conditions for rem over inequalities for $\mathbb{F}_{3,5}$.

the function. The stepped blocks along the diagonal occur when s and t have the same exponent and thus the pattern is similar to the invertibility condition for $+$ shown in Figure 1. Portions right of the main diagonal appear to exhibit random behavior. We believe this is the result of repeated cancellations in the computation of the remainder for those values, which suggests a behavior that we believe is similar to the Blum-Blum-Shub random number generator [8].

For remainder with inequalities, we succeeded in determining invertibility conditions for \leq and \geq if x is the first argument. However, for $x \text{ rem } s$ over $\{<, >\}$, and $s \text{ rem } x$ over $\{\geq, \leq, <, >\}$ we did not. This is particularly surprising considering that the invertibility conditions for non-strict and strict inequalities are nearly identical (varying only by a handful of pixels), as shown in Figure 3. Note that for x as the first argument, all variations of the concise invertibility conditions for non-strict inequality we considered failed as solutions for the strict inequality. This behavior is representative of the many subtle corner cases we encountered while synthesizing these conditions.

Figure 4 shows visualizations for invertibility conditions involving fma . The left two images are visualizations for the invertibility conditions for isZero . The corresponding invertibility conditions are given in Equations (2) and (3) above. We were not able to determine invertibility conditions for operator fma over predicate isSub , which are visualized in the rightmost two pictures in Figure 4. Finally, we did not succeed in finding invertibility conditions for fma with binary predicates, which are particularly challenging since they are three-dimensional. Finding solutions for these cases is ongoing work (see Section 4 for a more in-depth discussion).

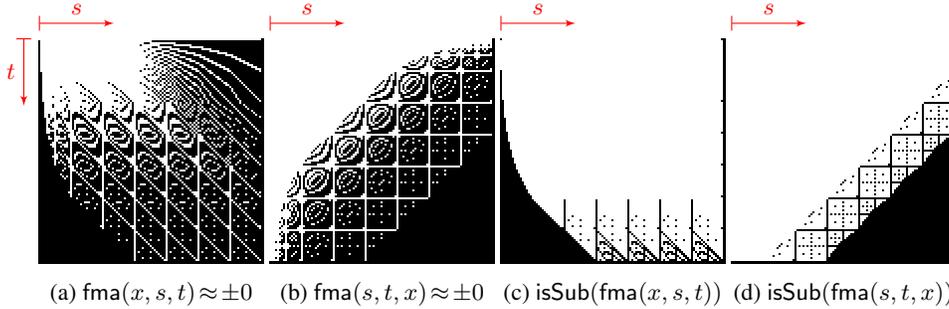


Fig. 4: Invertibility conditions for fma over $\{\text{isZero}, \text{isSub}\}$ for $\mathbb{F}_{3,5}$ and rnd. mode RNE.

4 Synthesis of Floating-Point Invertibility Conditions

In this section, we present an extension of the syntax-guided synthesis (SyGuS) paradigm we developed for the custom purpose of finding the invertibility conditions from Section 3. We use the SyGuS capabilities of the SMT solver CVC4 [31,4], which has recently been extended to support reasoning about the floating-point theory [11]. Our techniques are agnostic to theories, and can be used for synthesis problems over any finite domain. We will use the invertibility condition for floating-point addition with equality as a running example.

Establishing an invertibility condition requires solving a synthesis problem with *three* levels of quantifier alternation. In particular, for floating-point addition with equality, we are interested in finding a solution for predicate IC that satisfies the conjecture:

$$\exists \text{IC}. \forall s, t. (\text{IC}(s, t) \Leftrightarrow (\exists x. x \overset{\text{R}}{+} s \approx t)) \quad (4)$$

for some rounding mode R. In other words, this conjecture states that $\text{IC}(s, t)$ holds exactly when there exists an x that, when rounding the result of adding x to s according to mode R, yields t . Furthermore, we are interested in finding a solution for IC that holds *independently of the format* of x, s, t . Note that SMT solvers are not capable of reasoning about constraints that are parametric to floating-point formats. To address this challenge, following methodology from previous work [28], our strategy for establishing (general) invertibility conditions first solves the synthesis conjecture for a fixed format $\mathbb{F}_{\varepsilon, \sigma}$, and subsequently checks whether that solution also holds for other formats. The choice of the number of exponent bits ε and significand bits σ in $\mathbb{F}_{\varepsilon, \sigma}$ balances two criteria:

1. ε, σ should be large enough to exercise many (or all) of the behaviors of the operators and relations in our synthesis conjecture,
2. ε, σ should be small enough for the synthesis problem to be tractable.

In our experience, the best choices for (ε, σ) depended on the particular invertibility condition we were solving. The most common choices for (ε, σ) were $(3, 5)$, $(4, 5)$ and $(4, 6)$. For most two-dimensional invertibility conditions (those that involve two variables s and t), we used $(3, 5)$, since the required synthesis procedures mentioned below were roughly eight times faster than for $(4, 5)$. For one-dimensional invertibility conditions, we often used higher precision formats. Since floating-point operators like

addition take as additional argument a rounding mode R , we assume a fixed rounding mode when solving, and then cross-check our solution for multiple rounding modes.

Assume we have chosen to synthesize the invertibility condition for conjecture (4) for format $\mathbb{F}_{3,5}$ and rounding mode RNE. Notice that current SyGuS solvers [31,1] support only two levels of quantifier alternation. However, we can expand the innermost quantifier in this conjecture to obtain the conjecture:

$$\exists IC. \forall st. (IC(s, t) \Leftrightarrow (\bigvee_{i=0}^{226} i + s \approx t)) \quad (5)$$

where for simplicity of notation we use $i = 0, \dots, 226$ to denote the values of $\mathbb{F}_{3,5}$. This methodology was also used in Niemetz et al. [28], where invertibility conditions for bit-vector operators were synthesized for bit-width 4 by giving the conjecture of the above form to an off-the-shelf SyGuS solver. In contrast to that work, we found that the synthesis conjecture above is too challenging to be solved efficiently by current state-of-the-art enumerative SyGuS solvers. The reason for this is twofold. First, the smallest viable floating-point format is $3 + 5 = 8$ bits, which requires the body of (5) to have a significantly large number of disjuncts (227), which is more ten times larger than the 16 disjuncts required when synthesizing 4-bit invertibility conditions for bit-vectors. Second, floating-point formulas are much harder to solve than bit-vector formulas, due to the complexity of their bit-blasted encodings. Thus, a significantly challenging satisfiability query must be solved *for each* candidate considered within the SyGuS solver.

To address the above challenges, we perform a more extreme preprocessing step on our synthesis conjecture, which computes the input/output behavior of the invertibility condition on all points in the domain of s and t . In other words, we rephrase our synthesis conjecture as:

$$\exists IC. \bigwedge_{i=0}^{226} \bigwedge_{j=0}^{226} (IC(i, j) \Leftrightarrow c_{i,j}) \quad (6)$$

where each $c_{i,j}$ is a Boolean constant (either \top or \perp) determined by a quantifier-free satisfiability query. In particular, for each pair of floating-point values (i, j) , constant $c_{i,j}$ is \top if $x + i \approx j$ is satisfiable, and \perp if it is unsatisfiable. In practice, we represent the above conjecture as a 227×227 table, which we call the *full I/O specification* of invertibility condition IC. In our experience, computing this table for most two-dimensional invertibility conditions of sort $\mathbb{F}_{3,5}$ required 15 minutes (for $227 * 227 = 51529$ quantifier-free queries), and 2 hours for sort $\mathbb{F}_{4,5}$ (requiring $483 * 483 = 233289$ queries). This process was accelerated by first applying random sampling over possible values of x to quickly test if a query is satisfiable. For some operators, notably floating-point remainder, this required significantly more time than for others (up to a factor of 2). Due to the high cost of this preprocessing step, we generated a database with the full I/O specifications for *all* invertibility conditions from Section 3 using a cluster of 50 nodes with Intel Xeon E5-2637 with 3.5GHz and 32GB memory, and then shared this database among multiple developers. Computing the full I/O specifications for $\mathbb{F}_{3,5}$, $\mathbb{F}_{4,5}$, and $\mathbb{F}_{4,6}$ required a total of 459 days of CPU time (6.1 for $\mathbb{F}_{3,5}$, 54.7 for $\mathbb{F}_{4,5}$, and 398.5 for $\mathbb{F}_{4,6}$). Despite the heavy cost of this step, it was crucial for accelerating our framework for synthesizing invertibility conditions, which we describe in following.

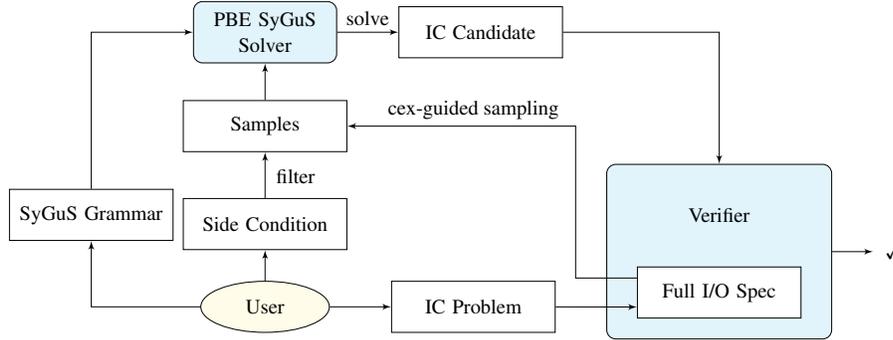


Fig. 5: Architecture for synthesizing invertibility conditions for floating point formulas.

Figure 5 summarizes our architecture for solving synthesis conjectures of the above form. The user first selects an invertibility condition problem to solve, where assume the full I/O specification has been computed using the aforementioned techniques. At a high level, our architecture can be seen as an *interactive synthesis environment*, where the user manages the interaction between two subprocedures:

1. a SyGuS solver with support for decision tree learning, and
2. a solution verifier storing the full I/O specification of the invertibility condition.

We use a counterexample-guided loop, where the SyGuS solver provides the solution verifier with candidate solutions, and the solution verifier provides the SyGuS solver with an evolving subset of sample points taken from the full I/O specification. These points correspond to counterexamples to failed candidate solutions, and are sampled in a uniformly random manner over the domain of our specification. To accelerate the speed at which our framework converges on a solution, we configure the solution verifier to generate multiple counterexample points (typically 10) for each iteration of the loop. The process terminates when the SyGuS solver generates a candidate solution that is correct for all points according to its full I/O specification.

We give the user control over both the solutions and counterexample points generated in this loop. First, as is commonly done in syntax-guided synthesis applications, the user in our workflow provides an input grammar to the SyGuS solver. This is a context-free grammar in a standard format [30], which contains a guess of the operators and patterns that may be involved in the invertibility condition we are synthesizing. Second, note that the domain of floating-point numbers can be subdivided into a number of subdomains and special cases (e.g. normal, subnormal, not-a-number, infinity), as well as split into different classifications (e.g. positive and negative). Our workflow allows the user to provide a *side condition*, whose purpose is to focus on finding an invertibility condition that is correct for one of these subdomains. The side condition acts as a filtering mechanism on the counterexample points generated by the solution verifier. For example, given the side condition $\text{isNorm}(s) \wedge \text{isNorm}(t)$, the solution verifier checks candidate solutions generated by the SyGuS solver only against points (s, t) where both arguments are normal, and consequently only communicates counterexamples of this form to the synthesis solver. The solution verifier may also be configured to estab-

lish that the current candidate solution generated by the SyGuS solver is *conditionally* correct, that is, it is true on all points in the domain that satisfy the side condition.

There are several advantages to the form of the synthesis conjecture in (6) that we exploit in our workflow. First, its structure makes it easy to divide, that is, our synthesis workflow at all times sends only a subset of the conjuncts of (6) for some (i, j) pairs. As a result, we do not burden the underlying synthesis solver with the entire conjecture at once, which would not scale in practice. A second advantage is that it is in *programming-by-examples* (PBE) form, since it consists of a conjunction of concrete input-output pairs. As a consequence, specialized algorithms can be used by the synthesis solver to generate solutions for (approximations of) our conjecture in a way that is highly scalable in practice. These techniques are broadly referred to as decision tree learning or unification algorithms. As a brief review (see [1] for a recent SyGuS-based approach), a decision tree learning algorithm is given as input a set of good examples $c_1 \mapsto \top, \dots, c_n \mapsto \top$ and bad examples $d_1 \mapsto \perp, \dots, d_m \mapsto \perp$. The goal of a decision tree algorithm is to find a predicate (often called a *classifier*) that evaluates to true on all the good examples, and false on all the bad examples. In our context, a classifier is expressed as an if-then-else tree of Boolean sort. Sampling the space of conjecture (6) provides the decision tree algorithm with good and bad examples and the returned classifier is a candidate solution that we give to the solution verifier. The synthesis solver of CVC4 uses a decision tree algorithm, which we use in our workflow. Due to the scalability of this algorithm and the fact that only a small subset of our conjecture is considered at any given time, candidate solutions are typically generated by the SyGuS solver in our framework in a matter of seconds.

Another important aspect of the synthesis solver in Figure 5 is that it is configured to generate *multiple* solutions for the current set of sample points. Due to the manner in which the SyGuS-based decision tree learning algorithm works, these solutions tend to become *more general* over the runtime of the solver. As a simple example (assuming exact integer arithmetic), say the solver is given input points $(1, 1) \mapsto \top, (2, 0) \mapsto \top, (1, 0) \mapsto \perp$ and $(0, 1) \mapsto \perp$ for (s, t) . It enumerates predicates over s and t , starting with simplest predicates first, say $s \approx 0, t \approx 0, s \approx 1, y \approx 1, s+t > 1$, and so on. After generating the first four predicates, it constructs the solution $\text{ite}(s \approx 1, t \approx 1, t \approx 0)$, which is a correct classifier for the given set of points. However, after generating the fifth predicate in this list, it returns that $s+t > 1$ is by itself a solution, and can be seen as a generalization of the previous solution since it requires no case splitting.

Since more general candidate solutions have a higher likelihood of being actual solutions in our experience, our workflow critically relies on the ability of the user to manually terminate the synthesis procedure when they are satisfied with the last generated candidate. Our synthesis procedure logs a list of candidate solutions that each satisfy the conjecture on the current set of sample points. When the user terminates the synthesis process, the solution verifier will check the last solution generated in this list. The user may further customize which candidate solution is tested by modifying this list by hand, if they have an intuition that a specific candidate is more likely to be correct.

Experience The first challenging invertibility condition we solved with our framework was addition with equality for rounding mode RNE. Initially, we used a generic grammar that contained the entire signature of floating-points. As a first key step towards

solving this problem, the synthesis procedure suggested the single literal $t \approx s + (t - s)$ as candidate solution. Although counterexamples were found for this candidate, we noticed that it satisfied over 98% of the specification, and a visualization of its I/O behavior showed similar patterns to the invertibility condition we were solving for. Based on these observations, we focused our grammar towards literals of this form. In particular, we used a function that takes two floating-points x, y and two rounding modes R_1, R_2 as arguments and returns $x \overset{R_1}{+} (y \overset{R_2}{-} x)$ as a builtin symbol of our grammar. We refer to such a function as a *residual* computation of y , noting that its value is often approximately y . By including various functions for residual computations, we focused the effort of the synthesizer on more interesting predicates. The end solution involved multiple residual computations, as shown in Table 2. Our initial solution was specific to the rounding mode RNE. After solving for several other rounding modes, we were able to construct a parametric solution that was correct for all rounding modes. In total, it took roughly three days of developer time to discover the generalized invertibility condition for addition with equality. Many of the subsequent invertibility conditions took a matter of hours, since by then we had a good intuition for the residual computations that were relevant for each case.

Invertibility conditions involving `rem`, `fma`, `isNorm`, and `isSub` were challenging and required further customizations to the grammar, for instance to include constants that corresponded to the minimum and maximum normal and subnormal values. Three dimensional invertibility conditions (which in this work is limited to cases of `fma` with binary predicates) were especially challenging since the domain of their conjecture is a factor of 227 larger for $\mathbb{F}_{3,5}$ than the others. Following our strategy for solving the invertibility conditions for specific formats and rounding modes, in ongoing work we are investigating solving these cases by first solving the invertibility condition for a fixed value c for one of its free variables u . Solving a two-dimensional problem of this form with a solution φ may suggest a generalization that works for all values of u where all occurrences of c in φ are replaced by u .

We found the side condition feature of our workflow important for narrowing down which subdomain was the most challenging for the conjecture in question. For instance, for some cases it was very easy to find invertibility conditions that held for all cases where both s and t were normal (resp. subnormal), but very difficult when s was normal and t was subnormal or vice versa.

Finally, we found that it was often helpful to visualize the input/output behavior of candidate solutions. It was often the case that the difference between a candidate solution and the desired behavior of the invertibility condition was indicative of a required modification to the grammar, as well as suggesting which parts of the domain of the conjecture we should focus on.

4.1 Verifying Conditions for Multiple Formats and Rounding Modes

We verified the correctness of all 167 invertibility conditions by checking them against their corresponding full I/O specification for floating-point formats $\mathbb{F}_{3,5}$, $\mathbb{F}_{4,5}$, and $\mathbb{F}_{4,6}$ and all rounding modes, which required 1.6 days of CPU time. This is relatively cheap compared to computing the specifications, since checking is essentially constant eval-

$\text{QE}_{\text{FP}}(\exists x. P(t_1, \dots, t_j[x], \dots, t_n))$, where $x \notin \text{FV}(t_i)$ for $i \neq j$:
 If $t_j[x] = x$, return $\text{getIC}(x, P)$.
 Otherwise, $t_j[x] = \diamond(s_1, \dots, s_k[x], \dots, s_m)$ where $m > 0$, $x \notin \text{FV}(s_i)$ for $i \neq k$.
 Let $Q[y] = P(t_1, \dots, t_{j-1}, \diamond(s_1, \dots, s_{k-1}, y, s_{k+1}, \dots, s_m), t_{j+1}, \dots, t_n)$ where y is
 a fresh variable.
 Return $\text{getIC}(y, Q[y]) \wedge \text{QE}_{\text{FP}}(\exists x. s_k[x] \approx y)$.

Fig. 6: Recursive procedure QE_{FP} for computing quantifier elimination for x in the unit linear formula $\exists x. P(t_1, \dots, t_j[x], \dots, t_n)$. The free variables in this formula and the fresh variable y are implicitly universally quantified.

uation of invertibility conditions for all possible input values. However, this quickly becomes infeasible with increasing precision, since the time required for computing the I/O specification roughly increases by factor 8 for each bit.

As a consequence, we generated quantified floating-point problems to verify the 167 invertibility for formats $\mathbb{F}_{3,5}$, $\mathbb{F}_{4,5}$, $\mathbb{F}_{4,6}$, $\mathbb{F}_{5,11}$ (Float16), $\mathbb{F}_{8,24}$ (Float32), and $\mathbb{F}_{11,53}$ (Float64) and all rounding modes. Each problem checks the T_{FP} -unsatisfiability of formula $\neg(\phi_c \Leftrightarrow \exists x. l[x])$, where $l[x]$ corresponds to the floating-point literal, and ϕ_c to its invertibility condition. In total, we generated 3786 problems ($116 * 5 + 51^7$ for each floating-point format) and checked them using CVC4 [4] (master 546bf686) and Z3 [17] (version 4.8.4).

We consider an invertibility condition to be verified for a floating-point format and rounding mode if at least one solver reports unsatisfiable. Given a CPU time limit of one hour and a memory limit of 8GB for each solver/benchmark pair, we were able to verify 3577 (94.5%) invertibility conditions overall, with 99.2% of $\mathbb{F}_{3,5}$, 99.7% of $\mathbb{F}_{4,5}$, 100% of $\mathbb{F}_{4,6}$, 93.8% of $\mathbb{F}_{5,11}$, 90.2% of $\mathbb{F}_{8,24}$, and 84% of $\mathbb{F}_{11,53}$. This verification with CVC4 and Z3 required a total of 32 days of CPU time. All verification jobs were run on cluster nodes with Intel Xeon E5-2637 3.5GHz and 32GB memory.

We encountered an issue with the invertibility condition for literal $\text{rti}(x) \approx t$, where CVC4 reported satisfiable for format $\mathbb{F}_{3,5}$ and $\mathbb{F}_{5,11}$, and Z3 reported unsatisfiable. This is due to a bug in SymFPU [12], the back-end floating-point library used in CVC4, that occurs only for certain floating-point formats.

5 Quantifier Elimination for Unit Linear Floating-Point Formulas

In this section, we present a quantifier elimination procedure for a restricted fragment of floating-point formulas, which is based on the invertibility conditions presented in Section 3. It is applicable to *unit linear* quantified formulas, that is, formulas of the form $\exists x. P[x]$ where P is *linear in x* , i.e., containing exactly one occurrence of x .

Figure 6 gives a quantifier elimination procedure QE_{FP} for unit linear floating-point formulas $\exists x. P[x]$. We write $\text{getIC}(y, Q[y])$ to indicate the invertibility condition for y in $Q[y]$, which amounts to a table lookup for the appropriate condition as given in

⁷ 116 invertibility conditions from rounding mode dependent operators and 51 invertibility conditions where the operator is rounding mode independent (e.g., rem).

Section 3. The recursive procedure returns a conjunction of conditions based on the path on which x occurs in P . If x occurs beneath multiple nested function applications, a fresh variable y is introduced and used for referencing the intermediate result of the subterm we are currently solving for. We demonstrate this in the following example.

Example 2. Consider the unit linear formula $\exists x. (x \cdot u) \dagger s \geq t$. Invoking the procedure QE_{FP} on this input yields, after two recursive calls, the conjunction

$$\text{getlC}(y_1, y_1 \dagger s \geq t) \wedge \text{getlC}(y_2, y_2 \cdot u \approx y_1) \wedge \text{getlC}(x, x \approx y_2)$$

where y_1 and y_2 are fresh variables. The third conjunct is trivially equivalent to \top . This formula is quantifier-free, its properties are given by the following theorem.

Theorem 1 *Let $\exists x. P$ be unit linear in x and P be quantifier-free, and let \mathcal{I} be a Σ_{FP} -interpretation for the free variables of P . Then, \mathcal{I} satisfies $\neg \exists x. P$ if and only if there exists a Σ_{FP} -interpretation \mathcal{J} that extends \mathcal{I} and satisfies $\neg \text{QE}_{\text{FP}}(\exists x. P)$.*

Previous work [28] presented a similar algorithm for solving unit linear bit-vector literals. In that work, a counterexample-guided loop was devised that made use of Hilbert-choice expressions for representing quantifier instantiations. In contrast to that work, here we only present a quantifier elimination algorithm. Extending our techniques to a general quantifier instantiation strategy is the subject of ongoing work. In the following, we present preliminary work in this direction.

Implementation We implemented a prototype extension of the SMT solver CVC4 that leverages the technique above for solving quantified floating-point formulas. For handling quantified formulas, CVC4 uses a basic model-based instantiation loop (see, e.g., [34,32] for instantiation approaches for other theories). This technique maintains a quantifier-free set of constraints F corresponding to instantiations of universally quantified formulas. It terminates with unsatisfiable if F is unsatisfiable at the quantifier-free level, and concludes with satisfiable if it can show that the given quantified formulas are satisfied by a model for F . For T_{FP} , the instantiations are substitutions of universally quantified variables to concrete floating-point values, e.g. $\forall x. P(x) \Rightarrow P(0)$, which can be highly inefficient in the worst case for higher precision.

We extend this basic loop with a preprocessing pass that constructs valid theory lemmas based on the invertibility conditions corresponding to literals of quantified formulas $\forall x. P$ that contain exactly one occurrence of x . As an example, given a quantified formula φ of the form $\forall x. \neg((x \cdot u) \dagger s \geq t \wedge Q(x))$, the following lemma is added:

$$(\text{getlC}(y_1, y_1 \dagger s \geq t) \Rightarrow y_1 \dagger s \geq t) \wedge (\text{getlC}(y_2, y_2 \cdot u \approx y_1) \Rightarrow y_2 \cdot u \approx y_1)$$

where y_1 and y_2 are fresh (existentially quantified) variables. Notice that this lemma preserves satisfiability since if the invertibility conditions hold for an equation, then without loss of generality we may assume that y_1 and y_2 are a solution for that equation. We then add the instantiation lemma $\varphi \Rightarrow \neg((y_2 \cdot u) \dagger s \geq t \wedge Q(y_2))$. Although x is not necessarily linear in the body of φ , if both invertibility conditions hold, then the

combination of the above lemmas implies $(y_2 \cdot u) + s \geq t$, which together with the instantiation lemma allows the solver to infer that the remaining portion of the quantified formula Q cannot hold for y_2 . An inference of this form may be more productive than enumerating the possible values of x in instantiations.

Evaluation. We considered all 61 benchmarks from SMT-LIB [5] that contained quantified formulas over floating-points (logic FP), which correspond to verification conditions from the software verification competition that use a floating-point encoding [21]. The invertibility conditions required for solving their literals include floating-point addition, multiplication and division (both arguments) with equality and inequality. We implemented all cases of invertibility conditions for solving these cases. We extended our SMT solver CVC4 (GitHub master 5d248c36) with the above preprocessing pass (GitHub cav19fp 9b5acd74), and compare its performance with (configuration CVC4-ext) and without (configuration CVC4-base) the above preprocessing pass enabled to the SMT solver Z3 (version 4.8.4). All experiments were run on the same cluster as above with a memory limit of 8GB and a 1800 second time limit. Overall, CVC4-base solved 35 benchmarks within the time limit (with no benchmarks uniquely solved compared to CVC4-ext), CVC4-ext solved 42 benchmarks (7 of these uniquely solved compared to the base version), and Z3 solved 56 benchmarks. While CVC4-ext solves significantly fewer benchmarks than Z3, we believe that the improvement over CVC4-base is indicative that our approach for invertibility conditions shows potential for solving quantified floating-point constraints in SMT solvers. A more comprehensive evaluation and implementation is left as future work.

6 Conclusion

We have presented invertibility conditions for a large subset of combinations of floating-point operators over floating-point predicates supported by SMT solvers. These conditions were found by a framework that utilizes syntax-guided synthesis solving, customized for our problem and developed over the course of this work. We have shown that invertibility conditions imply that a simple fragment of quantified floating-points admit compact quantifier elimination, and have given preliminary evidence that an SMT solver that partially leverages this technique can have a higher success rate on floating-point problems coming from a software verification application.

For future work, we plan to extend techniques for quantified and quantifier-free floating-points to incorporate our findings, in particular to lift existing quantifier instantiation approaches (e.g., [28]) and local search procedures (e.g., [27]) for bit-vectors to floating-points. We also plan to extend and use our synthesis framework for other challenging related synthesis tasks, such as finding conditions under which more complex constraints have solutions, including those having multiple occurrences of a variable to solve for. Our synthesis framework is agnostic to theories and can be used for any sort with a small finite domain. It can thus also be leveraged for solutions to quantified bit-vector constraints. Finally, we would like to establish formal proofs of correctness of our invertibility conditions that are independent of floating-point formats.

References

1. Alur, R., Radhakrishna, A., Udupa, A.: Scaling enumerative program synthesis via divide and conquer. In: Tools and Algorithms for the Construction and Analysis of Systems (TACAS). pp. 319–336 (2017)
2. Association, I.S.: 754-2008 - IEEE Standard for Floating-Point Arithmetic. <https://ieeexplore.ieee.org/servlet/opac?punumber=4610933> (2008)
3. Barr, E.T., Vo, T., Le, V., Su, Z.: Automatic detection of floating-point exceptions. SIGPLAN Not. **48**(1), 549–560 (Jan 2013)
4. Barrett, C., Conway, C.L., Deters, M., Hadarean, L., Jovanović, D., King, T., Reynolds, A., Tinelli, C.: CVC4. In: Proceedings of the 23rd International Conference on Computer Aided Verification. pp. 171–177. CAV’11, Springer-Verlag (2011)
5. Barrett, C., Stump, A., Tinelli, C.: The Satisfiability Modulo Theories Library (SMT-LIB). www.SMT-LIB.org (2010)
6. Ben Khadra, M.A., Stoffel, D., Kunz, W.: goSAT: Floating-point satisfiability as global optimization. In: FMCAD. pp. 11–14. IEEE (2017)
7. Bjørner, N., Janota, M.: Playing with quantified satisfaction. In: 20th International Conferences on Logic for Programming, Artificial Intelligence and Reasoning - Short Presentations, LPAR 2015, Suva, Fiji, November 24-28, 2015. pp. 15–27 (2015)
8. Blum, L., Blum, M., Shub, M.: A simple unpredictable pseudo-random number generator. SIAM J. Comput. **15**(2), 364–383 (1986)
9. Brain, M., D’silva, V., Griggio, A., Haller, L., Kroening, D.: Deciding floating-point logic with abstract conflict driven clause learning. Formal Methods in System Design **45**(2), 213–245 (2014)
10. Brain, M., Joshi, S., Kroening, D., Schrammel, P.: Safety verification and refutation by k-invariants and k-induction. In: Blazy, S., Jensen, T. (eds.) Static Analysis - 22nd International Symposium, SAS 2015, Saint-Malo, France, September 9-11, 2015, Proceedings. Lecture Notes in Computer Science, vol. 9291, pp. 145–161. Springer (2015)
11. Brain, M., Schanda, F., Sun, Y.: Building better bit-blasting for floating-point problems (April 2019), to appear at TACAS’19
12. Brain, M., Schanda, F., Sun, Y.: Building better bit-blasting for floating-point problems. In: Vojnar, T., Zhang, L. (eds.) Tools and Algorithms for the Construction and Analysis of Systems - 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part I. Lecture Notes in Computer Science, vol. 11427, pp. 79–98. Springer (2019). https://doi.org/10.1007/978-3-030-17462-0_5, https://doi.org/10.1007/978-3-030-17462-0_5
13. Brain, M., Tinelli, C., Rümmer, P., Wahl, T.: An automatable formal semantics for IEEE-754 floating-point arithmetic. In: 22nd IEEE Symposium on Computer Arithmetic, ARITH 2015, Lyon, France, June 22-24, 2015. pp. 160–167. IEEE (2015)
14. Brillout, A., Kroening, D., Wahl, T.: Mixed abstractions for floating-point arithmetic. In: FMCAD. pp. 69–76. IEEE (2009)
15. Conchon, S., Iguernlala, M., Ji, K., Melquiond, G., Fumex, C.: A three-tier strategy for reasoning about floating-point numbers in SMT. In: CAV. pp. 419–435. Springer (2017)
16. Dumas, M., Melquiond, G.: Certification of bounds on expressions involving rounded operators. ACM Transactions on Mathematical Software **37**(1), 1–20 (2010)
17. De Moura, L., Bjørner, N.: Z3: An efficient smt solver. In: Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 337–340. TACAS’08/ETAPS’08, Springer-Verlag (2008)

18. Drane, T.: Lossy Polynomial Datapath Synthesis. Ph.D. thesis, Imperial College London (February 2014)
19. Dutertre, B.: Solving exists/forall problems in yices. Workshop on Satisfiability Modulo Theories (2015)
20. Fu, Z., Su, Z.: XSat: A fast floating-point satisfiability solver. In: CAV. pp. 187–209. Springer (2016)
21. Heizmann, M., Chen, Y., Dietsch, D., Greitschus, M., Nutz, A., Musa, B., Schätzle, C., Schilling, C., Schüssele, F., Podelski, A.: Ultimate automizer with an on-demand construction of floyd-hoare automata - (competition contribution). In: Tools and Algorithms for the Construction and Analysis of Systems - 23rd International Conference, TACAS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings, Part II. pp. 394–398 (2017)
22. Lapschies, F.: SONOLAR, the solver for non-linear arithmetic. <http://www.informatik.uni-bremen.de/agbs/florian/sonolar> (2014)
23. Liew, D.: JFS: JIT fuzzing solver. <https://github.com/delcypher/jfs>
24. Marre, B., Bobot, F., Chihani, Z.: Real behavior of floating point numbers. In: SMT Workshop (2017)
25. Michel, C., Rueher, M., Lebbah, Y.: Solving constraints over floating-point numbers. In: International Conference on Principles and Practice of Constraint Programming. pp. 524–538. Springer (2001)
26. de Moura, L.M., Bjørner, N.: Efficient e-matching for SMT solvers. In: Automated Deduction - CADE-21, 21st International Conference on Automated Deduction, Bremen, Germany, July 17-20, 2007, Proceedings. pp. 183–198 (2007)
27. Niemetz, A., Preiner, M., Biere, A.: Precise and complete propagation based local search for satisfiability modulo theories. In: Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part I. pp. 199–217 (2016)
28. Niemetz, A., Preiner, M., Reynolds, A., Barrett, C., Tinelli, C.: Solving quantified bit-vectors using invertibility conditions. In: Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part II. pp. 236–255 (2018)
29. Preiner, M., Niemetz, A., Biere, A.: Counterexample-guided model synthesis. In: Tools and Algorithms for the Construction and Analysis of Systems - 23rd International Conference, TACAS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings, Part I. pp. 264–280 (2017)
30. Raghothaman, M., Udupa, A.: Language to specify syntax-guided synthesis problems (05 2014)
31. Reynolds, A., Deters, M., Kuncak, V., Tinelli, C., Barrett, C.W.: Counterexample-guided quantifier instantiation for synthesis in SMT. In: Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part II. pp. 198–216 (2015)
32. Reynolds, A., King, T., Kuncak, V.: Solving quantified linear arithmetic by counterexample-guided instantiation. *Formal Methods in System Design* **51**(3), 500–532 (2017)
33. Scheibler, K., Kupferschmid, S., Becker, B.: Recent improvements in the SMT solver iSAT. *MBMV* **13**, 231–241 (2013)
34. Wintersteiger, C.M., Hamadi, Y., de Moura, L.M.: Efficiently solving quantified bit-vector formulas. *Formal Methods in System Design* **42**(1), 3–23 (2013)
35. Zeljić, A., Wintersteiger, C.M., Rümmer, P.: Approximations for model construction. In: IICAR. pp. 344–359. Springer (2014)

A Omitted Floating-Point Invertibility Conditions

Literal	Invertibility Condition
$x \overset{R}{+} s \not\approx t$	$s \approx \text{NaN} \Rightarrow t \not\approx \text{NaN}$
$x \overset{R}{-} s \not\approx t$	$s \approx \text{NaN} \Rightarrow t \not\approx \text{NaN}$
$s \overset{R}{-} x \not\approx t$	$s \approx \text{NaN} \Rightarrow t \not\approx \text{NaN}$
$x \overset{R}{\cdot} s \not\approx t$	$s \approx \text{NaN} \Rightarrow t \not\approx \text{NaN}$
$x \overset{R}{\div} s \not\approx t$	$s \approx \text{NaN} \Rightarrow t \not\approx \text{NaN}$
$s \overset{R}{\div} x \not\approx t$	$s \approx \text{NaN} \Rightarrow t \not\approx \text{NaN}$
$x \text{ rem } s \not\approx t$	$t \approx \text{NaN} \Rightarrow (s \not\approx \pm 0 \wedge (\neg \text{isNeg}(s) \Rightarrow \text{isPos}(s)))$
$s \text{ rem } x \not\approx t$	$t \approx \text{NaN} \Rightarrow (s \not\approx \pm \infty \wedge s \not\approx \text{NaN})$
$\sqrt[x]{x} \not\approx t$	\top
$ x \not\approx t$	\top
$-x \not\approx t$	\top
$\overset{R}{\text{rti}}(x) \not\approx t$	\top

Table 4: Invertibility conditions for floating-point operators (excl. fma) with $\not\approx$.

Literal	Invertibility Condition
$x \overset{R}{+} s \leq t$	$(\text{isNorm}(s) \vee \text{isNeg}(s) \vee \text{isSub}(s) \vee s \approx \pm 0 \vee (t \approx \pm \infty \wedge \text{isPos}(s) \wedge \neg \text{isNeg}(t))) \wedge t \not\approx \text{NaN}$
$x \overset{R}{-} s \leq t$	$((t \approx \pm \infty \wedge \neg \text{isNeg}(t)) \vee \neg \text{isNeg}(s) \vee s \not\approx \pm \infty) \wedge s \not\approx \text{NaN} \wedge t \not\approx \text{NaN}$
$s \overset{R}{-} x \leq t$	$(\text{isNorm}(s) \vee \text{isNeg}(s) \vee \text{isSub}(s) \vee s \approx \pm 0 \vee (t \approx \pm \infty \wedge \text{isPos}(s) \wedge \neg \text{isNeg}(t))) \wedge t \not\approx \text{NaN}$
$x \overset{R}{\cdot} s \leq t$	$(t \approx \pm 0 \vee \neg \text{isNeg}(t) \vee s \not\approx \pm 0) \wedge s \not\approx \text{NaN} \wedge t \not\approx \text{NaN}$
$x \overset{R}{\div} s \leq t$	$(t \approx \pm 0 \vee \neg \text{isNeg}(t) \vee s \not\approx \pm \infty) \wedge s \not\approx \text{NaN} \wedge t \not\approx \text{NaN}$
$s \overset{R}{\div} x \leq t$	$(t \approx \pm 0 \vee \neg \text{isNeg}(t) \vee s \not\approx \pm 0) \wedge s \not\approx \text{NaN} \wedge t \not\approx \text{NaN}$
$x \text{ rem } s \leq t$	$s \text{ rem } s \leq t \vee t \text{ rem } s \leq t \vee (s \approx \pm \infty \wedge \text{isPos}(t))$
$s \text{ rem } x \leq t$?
$\sqrt[x]{x} \leq t$	$\text{isPos}(t) \vee t \approx \pm 0$
$ x \leq t$	$\text{isPos}(t) \vee t \approx \pm 0$
$-x \leq t$	$t \not\approx \text{NaN}$
$\overset{R}{\text{rti}}(x) \leq t$	$t \not\approx \text{NaN}$

Table 5: Invertibility conditions for floating-point operators (excl. fma) with \leq .

Literal	Invertibility Condition
$x \overset{R}{+} s < t$	$\text{ite}((s \neq \pm 0 \wedge \neg \text{isNeg}(s) \wedge \neg \text{isNorm}(s)), (\text{isSub}(s) \wedge (\neg \text{isNeg}(t) \vee t \neq \pm \infty)), (\neg \text{isNeg}(t) \vee t \neq \pm \infty)) \wedge t \neq \text{NaN}$
$x \overset{R}{-} s < t$	$(\text{isPos}(s) \vee s \neq \pm \infty) \wedge \text{ite}(t \approx \pm \infty, \neg \text{isNeg}(t), t \neq \text{NaN}) \wedge s \neq \text{NaN}$
$s \overset{R}{-} x < t$	$\text{ite}((s \neq \pm 0 \wedge \neg \text{isNeg}(s) \wedge \neg \text{isNorm}(s)), (\text{isSub}(s) \wedge (\neg \text{isNeg}(t) \vee t \neq \pm \infty)), (\neg \text{isNeg}(t) \vee t \neq \pm \infty)) \wedge t \neq \text{NaN}$
$x \overset{R}{\cdot} s < t$	$s \neq \text{NaN} \wedge \text{ite}(\text{isNeg}(t), (s \neq \pm 0 \wedge t \neq \pm \infty), ((s \neq \pm 0 \vee t \neq \pm 0) \wedge t \neq \text{NaN}))$
$x \overset{R}{\div} s < t$	$\text{ite}(s \approx \pm \infty, (\neg \text{isNeg}(t) \wedge t \neq \pm 0), (\neg \text{isNeg}(t) \vee t \neq \pm \infty)) \wedge s \neq \text{NaN} \wedge t \neq \text{NaN}$
$s \overset{R}{\div} x < t$	$s \neq \text{NaN} \wedge \text{ite}(\text{isNeg}(t), (s \neq \pm 0 \wedge t \neq \pm \infty), ((s \neq \pm 0 \vee t \neq \pm 0) \wedge t \neq \text{NaN}))$
$x \text{ rem } s < t$?
$s \text{ rem } x < t$?
$\sqrt[x]{x} < t$	$t \neq \pm 0 \wedge \text{isPos}(t)$
$ x < t$	$t \neq \pm 0 \wedge \text{isPos}(t)$
$-x < t$	$\text{ite}(\text{isNeg}(t), t \neq \pm \infty, t \neq \text{NaN})$
$\overset{R}{\text{rti}}(x) < t$	$\text{ite}(t \approx \pm \infty, \text{isPos}(t), t \neq \text{NaN})$

Table 6: Invertibility conditions for floating-point operators (excl. fma) with $<$.

Literal	Invertibility Condition
$x \overset{R}{+} s > t$	$(\text{isNeg}(t) \vee t \neq \pm \infty) \wedge (\text{isPos}(s) \vee (\text{isNeg}(s) \wedge s \neq \pm \infty)) \wedge t \neq \text{NaN}$
$x \overset{R}{-} s > t$	$(\text{isNeg}(t) \vee t \neq \pm \infty) \wedge (s \approx \pm \infty \Rightarrow \neg \text{isPos}(s)) \wedge s \neq \text{NaN} \wedge t \neq \text{NaN}$
$s \overset{R}{-} x > t$	$(\text{isNeg}(t) \vee t \neq \pm \infty) \wedge (\text{isPos}(s) \vee (\text{isNeg}(s) \wedge s \neq \pm \infty)) \wedge t \neq \text{NaN}$
$x \overset{R}{\cdot} s > t$	$s \neq \text{NaN} \wedge \text{ite}(\text{isNeg}(t), (t \neq \pm 0 \vee s \neq \pm 0), (s \neq \pm 0 \wedge t \neq \text{NaN} \wedge t \neq \pm \infty))$
$x \overset{R}{\div} s > t$	$s \neq \text{NaN} \wedge \text{ite}(\text{isNeg}(t), (s \neq \pm \infty \vee t \neq \pm 0), (s \neq \pm \infty \wedge t \neq \text{NaN} \wedge t \neq \pm \infty))$
$s \overset{R}{\div} x > t$	$s \neq \text{NaN} \wedge \text{ite}(\text{isNeg}(t), (t \neq \pm 0 \vee s \neq \pm 0), (s \neq \pm 0 \wedge t \neq \text{NaN} \wedge t \neq \pm \infty))$
$x \text{ rem } s > t$?
$s \text{ rem } x > t$?
$\sqrt[x]{x} > t$	$\text{ite}(t \approx \pm \infty, \text{isNeg}(t), t \neq \text{NaN})$
$ x > t$	$\text{ite}(t \approx \pm \infty, \text{isNeg}(t), t \neq \text{NaN})$
$-x > t$	$\text{ite}(t \approx \pm \infty, \text{isNeg}(t), t \neq \text{NaN})$
$\overset{R}{\text{rti}}(x) > t$	$\text{ite}(t \approx \pm \infty, \text{isNeg}(t), t \neq \text{NaN})$

Table 7: Invertibility conditions for floating-point operators (excl. fma) with $>$.

Literal	Invertibility Condition
$x \approx t$	\top
$x < t$	$\text{ite}(\text{isNeg}(t), t \not\approx \pm\infty, \text{isPos}(t))$
$x > t$	$t \not\approx \text{NaN} \wedge (t \approx \pm\infty \Rightarrow \text{isNeg}(t))$
$x \leq t$	$t \not\approx \text{NaN}$
$x \geq t$	$t \not\approx \text{NaN}$
$x \not\approx t$	\top

Table 8: Invertibility conditions for floating-point relations $\{\approx, <, >, \leq, \geq, \not\approx\}$.

Literal	Invertibility Condition
$\text{isZero}(x \overset{R}{+} s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm\infty$
$\text{isZero}(x \overset{R}{-} s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm\infty$
$\text{isZero}(s \overset{R}{-} x)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm\infty$
$\text{isZero}(x \overset{R}{\cdot} s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm\infty$
$\text{isZero}(x \overset{R}{\div} s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm 0$
$\text{isZero}(s \overset{R}{\div} x)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm\infty$
$\text{isZero}(x \text{ rem } s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm 0$
$\text{isZero}(s \text{ rem } x)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm\infty$
$\text{isZero}(\sqrt[x]{x})$	\top
$\text{isZero}(x)$	\top
$\text{isZero}(-x)$	\top
$\text{isZero}(\overset{R}{\text{rti}}(x))$	\top
$\text{isZero}(\overset{R}{\text{fma}}(x, s, t))$	$\overset{R}{\text{fma}}(-(t \overset{\text{RTP}}{\div} s), s, t) \approx \pm 0 \vee \overset{R}{\text{fma}}(-(t \overset{\text{RTN}}{\div} s), s, t) \approx \pm 0 \vee (s \approx \pm 0 \wedge t \approx \pm 0)$
$\text{isZero}(\overset{R}{\text{fma}}(s, t, x))$	$\overset{R}{\text{fma}}(s, t, -(s \overset{\text{RTP}}{\cdot} t)) \approx \pm 0 \vee \overset{R}{\text{fma}}(s, t, -(s \overset{\text{RTN}}{\cdot} t)) \approx \pm 0$

Table 9: Invertibility conditions for floating-point operators with isZero .

Literal	Invertibility Condition
$\text{isInf}(x \overset{\text{R}}{+} s)$	$s \not\approx \text{NaN}$
$\text{isInf}(x \overset{\text{R}}{-} s)$	$s \not\approx \text{NaN}$
$\text{isInf}(s \overset{\text{R}}{-} x)$	$s \not\approx \text{NaN}$
$\text{isInf}(x \overset{\text{R}}{\cdot} s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm 0$
$\text{isInf}(x \overset{\text{R}}{\div} s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm \infty$
$\text{isInf}(s \overset{\text{R}}{\div} x)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm 0$
$\text{isInf}(x \text{ rem } s)$	\perp
$\text{isInf}(s \text{ rem } x)$	\perp
$\text{isInf}(\sqrt[x]{y})$	\top
$\text{isInf}(x)$	\top
$\text{isInf}(-x)$	\top
$\text{isInf}(\overset{\text{R}}{\text{rti}}(x))$	\top
$\text{isInf}(\overset{\text{R}}{\text{fma}}(x, s, t))$	$\text{ite}(s \approx \pm 0, t \approx \pm \infty, (t \not\approx \text{NaN} \wedge s \not\approx \text{NaN}))$
$\text{isInf}(\overset{\text{R}}{\text{fma}}(s, t, x))$	$\neg(s \approx \text{NaN} \vee t \approx \text{NaN} \vee (s \approx \pm \infty \wedge t \approx \pm 0) \vee (s \approx \pm 0 \wedge t \approx \pm \infty))$

Table 10: Invertibility conditions for floating-point operators with isInf.

Literal	Invertibility Condition
$\text{isNaN}(x \overset{\text{R}}{+} s)$	\top
$\text{isNaN}(x \overset{\text{R}}{-} s)$	\top
$\text{isNaN}(s \overset{\text{R}}{-} x)$	\top
$\text{isNaN}(x \overset{\text{R}}{\cdot} s)$	\top
$\text{isNaN}(x \overset{\text{R}}{\div} s)$	\top
$\text{isNaN}(s \overset{\text{R}}{\div} x)$	\top
$\text{isNaN}(x \text{ rem } s)$	\top
$\text{isNaN}(s \text{ rem } x)$	\top
$\text{isNaN}(\sqrt[x]{y})$	\top
$\text{isNaN}(x)$	\top
$\text{isNaN}(-x)$	\top
$\text{isNaN}(\overset{\text{R}}{\text{rti}}(x))$	\top
$\text{isNaN}(\overset{\text{R}}{\text{fma}}(x, s, t))$	\top
$\text{isNaN}(\overset{\text{R}}{\text{fma}}(s, t, x))$	\top

Table 11: Invertibility conditions for floating-point operators with isNaN.

Literal	Invertibility Condition
$\text{isPos}(x \overset{R}{+} s)$	$s \neq \text{NaN} \wedge (\neg \text{isNeg}(s) \vee s \neq \pm\infty)$
$\text{isPos}(x \overset{R}{-} s)$	$(\text{isNeg}(s) \vee s \neq \pm\infty) \wedge s \neq \text{NaN}$
$\text{isPos}(s \overset{R}{-} x)$	$s \neq \text{NaN} \wedge (\neg \text{isNeg}(s) \vee s \neq \pm\infty)$
$\text{isPos}(x \overset{R}{\cdot} s)$	$s \neq \text{NaN}$
$\text{isPos}(x \overset{R}{\div} s)$	$s \neq \text{NaN}$
$\text{isPos}(s \overset{R}{\div} x)$	$s \neq \text{NaN}$
$\text{isPos}(x \text{ rem } s)$	$s \neq \text{NaN} \wedge s \neq \pm 0$
$\text{isPos}(s \text{ rem } x)$	$\text{ite}(\text{isSub}(s), (\text{isPos}(s) \vee s \neq s \cdot s^{\text{RTP}}), (\neg(s \approx \pm 0 \wedge \text{isNeg}(s)) \wedge s \neq \pm\infty \wedge s \neq \text{NaN}))$
$\text{isPos}(\sqrt[x]{y})$	\top
$\text{isPos}(x)$	\top
$\text{isPos}(-x)$	\top
$\text{isPos}(\text{rti}(x))$	\top
$\text{isPos}(\text{fma}(x, s, t))$	$\text{ite}(\text{isNeg}(t), (\text{ite}(R \approx \text{RTN}, (s \neq \text{NaN} \wedge s \neq \pm 0), \text{ite}(s \approx \pm 0, t \approx \pm 0, s \neq \text{NaN})) \wedge t \neq \pm\infty), (s \neq \text{NaN} \wedge t \neq \text{NaN}))$
$\text{isPos}(\text{fma}(s, t, x))$	$\text{ite}(t \approx \pm\infty, (\text{ite}(\text{isNeg}(t), \text{isNeg}(s), \text{isPos}(s)) \wedge s \neq \pm 0), ((\text{ite}(\text{isNeg}(t), \text{isNeg}(s), \text{isPos}(s)) \wedge t \neq \pm 0) \vee s \neq \pm\infty) \wedge s \neq \text{NaN})) \wedge t \neq \text{NaN}$

Table 12: Invertibility conditions for floating-point operators with isPos.

Literal	Invertibility Condition
$\text{isNeg}(x \overset{R}{+} s)$	$(\text{isNeg}(s) \vee s \neq \pm\infty) \wedge s \neq \text{NaN}$
$\text{isNeg}(x \overset{R}{-} s)$	$s \neq \text{NaN} \wedge (\neg \text{isNeg}(s) \vee s \neq \pm\infty)$
$\text{isNeg}(s \overset{R}{-} x)$	$(\text{isNeg}(s) \vee s \neq \pm\infty) \wedge s \neq \text{NaN}$
$\text{isNeg}(x \overset{R}{\cdot} s)$	$s \neq \text{NaN}$
$\text{isNeg}(x \overset{R}{\div} s)$	$s \neq \text{NaN}$
$\text{isNeg}(s \overset{R}{\div} x)$	$s \neq \text{NaN}$
$\text{isNeg}(x \text{ rem } s)$	$s \neq \text{NaN} \wedge s \neq \pm 0$
$\text{isNeg}(s \text{ rem } x)$	$\text{ite}(\text{isSub}(s), (\text{isNeg}(s) \vee s \neq s \cdot s^{\text{RTP}}), (\neg(s \approx \pm 0 \wedge \text{isPos}(s)) \wedge s \neq \pm\infty \wedge s \neq \text{NaN}))$
$\text{isNeg}(\sqrt[x]{y})$	\top
$\text{isNeg}(x)$	\perp
$\text{isNeg}(-x)$	\top
$\text{isNeg}(\text{rti}(x))$	\top
$\text{isNeg}(\text{fma}(x, s, t))$	$s \neq \text{NaN} \wedge t \neq \text{NaN} \wedge (\text{isPos}(t) \Rightarrow (t \neq \pm\infty \wedge \text{ite}(R \approx \text{RTN}, (s \approx \pm 0 \Rightarrow t \approx \pm 0), s \neq \pm 0)))$
$\text{isNeg}(\text{fma}(s, t, x))$	$\text{ite}(s \approx \pm\infty, \text{ite}(\text{isNeg}(s), (t \neq \pm 0 \wedge \neg \text{isNeg}(t)), (\text{isNeg}(t) \wedge t \neq \pm 0)), ((\text{ite}(\text{isNeg}(t), \text{isPos}(s), \text{isNeg}(s)) \wedge s \neq \pm 0) \vee t \neq \pm\infty)) \wedge s \neq \text{NaN} \wedge t \neq \text{NaN}$

Table 13: Invertibility conditions for floating-point operators with isNeg.

Literal Invertibility Condition	
$\text{isNorm}(x \overset{R}{+} s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm\infty$
$\text{isNorm}(x \overset{R}{-} s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm\infty$
$\text{isNorm}(s \overset{R}{-} x)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm\infty$
$\text{isNorm}(x \overset{R}{\cdot} s)$	$s \not\approx \pm 0 \wedge s \not\approx \pm\infty \wedge s \not\approx \text{NaN} \wedge \text{ite}(\text{isNorm}(s), \top,$ $\neg \text{isSub}(s \overset{\text{ite}((\text{R} \text{RTN} \vee \text{R} \text{RTP}), \text{RNE}, \text{R})}{\cdot} \max_n))$
$\text{isNorm}(x \overset{R}{\div} s)$	$s \not\approx \pm 0 \wedge s \not\approx \pm\infty \wedge s \not\approx \text{NaN}$
$\text{isNorm}(s \overset{R}{\div} x)$	$\text{isNorm}(s) \vee \text{isSub}(s)$
$\text{isNorm}(x \text{ rem } s)$	$s \approx \pm\infty \vee (\text{isNorm}(s) \wedge s \geq 2.0^{\text{RTZ}} \cdot \min_n)$
$\text{isNorm}(s \text{ rem } x)$	$\text{isNorm}(s)$
$\text{isNorm}(\sqrt[\text{R}]{x})$	\top
$\text{isNorm}(x)$	\top
$\text{isNorm}(-x)$	\top
$\text{isNorm}(\overset{R}{\text{rti}}(x))$	\top
$\text{isNorm}(\overset{R}{\text{fma}}(x, s, t))$	$t \not\approx \pm\infty \wedge s \not\approx \pm\infty \wedge t \not\approx \text{NaN} \wedge s \not\approx \text{NaN} \wedge \neg(s \approx \pm 0$ $\wedge (t \approx \pm 0 \vee \text{isSub}(t))) \wedge (\text{isNorm}(t) \vee \text{isNorm}(s) \vee$ $\text{isNorm}(\overset{\text{ite}((\text{R} \text{RTN} \vee \text{R} \text{RTP}), \text{RNE}, \text{R})}{\text{fma}}(s, \text{ite}(\text{isPos}(s) \approx \text{isPos}(t), \max_n,$ $-\max_n), t)))$
$\text{isNorm}(\overset{R}{\text{fma}}(s, t, x))$	$t \not\approx \pm\infty \wedge s \not\approx \pm\infty \wedge t \not\approx \text{NaN} \wedge s \not\approx \text{NaN}$ $\wedge ((\text{isNorm}(s) \wedge \text{isNorm}(t)) \Rightarrow$ $\overset{R}{\text{fma}}(s, t, \text{ite}(\text{isPos}(s) \approx \text{isPos}(t), -\max_n, \max_n)) \not\approx \pm\infty)$

Table 14: Invertibility conditions for floating-point operators with isNorm.

Literal Invertibility Condition	
$\text{isSub}(x \overset{R}{+} s)$	$s \not\approx s + \overset{\text{RTN}}{\max_s} \vee s \not\approx s + \overset{\text{RTP}}{-\max_s}$
$\text{isSub}(x \overset{R}{-} s)$	$s \not\approx s + \overset{\text{RTN}}{\max_s} \vee s \not\approx s + \overset{\text{RTP}}{-\max_s}$
$\text{isSub}(s \overset{R}{-} x)$	$s \not\approx s + \overset{\text{RTN}}{\max_s} \vee s \not\approx s + \overset{\text{RTP}}{-\max_s}$
$\text{isSub}(x \overset{R}{\cdot} s)$	$(\text{isNorm}(s) \vee \text{isSub}(s)) \wedge (\neg \text{isNorm}(s \overset{R}{\cdot} \min_s) \vee \neg \text{isNorm}(s \overset{R}{\cdot} -\min_s))$
$\text{isSub}(x \overset{R}{\div} s)$	$\text{isNorm}(s \overset{\text{RNE}}{\div} s) \wedge \neg \text{isNorm}(\min_s \overset{\text{RNE}}{\div} s)$
$\text{isSub}(s \overset{R}{\div} x)$	$\text{isNorm}(s \overset{\text{RNE}}{\div} \max_s) \vee (\text{isNorm}(s \overset{\text{RNE}}{\div} s) \wedge \neg \text{isNorm}(s \overset{\text{RNE}}{\div} \max_n))$
$\text{isSub}(x \text{ rem } s)$	$s \not\approx \text{NaN} \wedge s \not\approx \pm 0 \wedge s \not\approx \min_s \wedge s \not\approx -\min_s$
$\text{isSub}(s \text{ rem } x)$	$\text{isNorm}(s \overset{\text{RNE}}{\div} s)$
$\text{isSub}(\sqrt[\text{R}]{x})$	$\text{isSub}(\overset{\text{RTP}}{\sqrt{\min_s}})$
$\text{isSub}(x)$	\top
$\text{isSub}(-x)$	\top
$\text{isSub}(\overset{R}{\text{rti}}(x))$	\perp
$\text{isSub}(\overset{R}{\text{fma}}(x, s, t))$	$?$
$\text{isSub}(\overset{R}{\text{fma}}(s, t, x))$	$?$

Table 15: Invertibility conditions for floating-point operators with isSub.

B Proof of Theorem 1

Let $\exists x. P$ be unit linear with respect to x and P be quantifier-free and let \mathcal{I} be a Σ_{FP} -interpretation for the free variables of P . Then, \mathcal{I} satisfies $\neg\exists x. P$ if and only if there exists a Σ_{FP} -interpretation \mathcal{J} that extends \mathcal{I} and satisfies $\neg\text{QE}_{FP}(\exists x. P)$.

Proof. We prove this by induction on the depth of x in P .

If x has depth 1, then $\text{QE}_{FP}(\exists x. P)$ returns $\text{getIC}(x, P)$, which by Definition 1 is a formula that is equivalent to $\exists x. P$ and the statement trivially holds.

Otherwise, we have that $\text{QE}_{FP}(\exists x. P)$ is $\text{getIC}(y, Q[y]) \wedge \text{QE}_{FP}(\exists x. s_k[x] \approx y)$. Let \mathcal{I} be a Σ_{FP} -interpretation. Notice that the depth of x in $s_k[x] \approx y$ is one less than its depth in P . Hence, by the induction hypothesis, \mathcal{I} satisfies $\neg\exists x. s_k[x] \approx y$ if and only there exists an extension \mathcal{J} of \mathcal{I} that satisfies $\neg\text{QE}_{FP}(\exists x. s_k[x] \approx y)$.

Assume \mathcal{I} satisfies $\neg\exists x. P$. Assume by contraction that all extensions of \mathcal{I} do not satisfy $\neg\text{QE}_{FP}(\exists x. P)$, i.e. they all satisfy $\text{getIC}(y, Q[y])$ and $\text{QE}_{FP}(\exists x. s_k[x] \approx y)$, where notice that x is linear in $s_k[x]$. Due to the former, they all satisfy $\exists z. Q[z]$ by definition. Let \mathcal{J} be an extension where y is interpreted the same as some z such that $Q[z]$ holds. We have that \mathcal{J} satisfies $\exists x. Q[s_k[x]]$, which is $\exists x. P$. This is a contradiction since \mathcal{J} is an extension of \mathcal{I} , which satisfies $\neg\exists x. P$.

Assume there exists an extension \mathcal{J} of \mathcal{I} that satisfies $\neg\text{QE}_{FP}(\exists x. P)$, i.e. it either satisfies $\neg\text{getIC}(y, Q[y])$ or $\neg\text{QE}_{FP}(\exists x. s_k[x] \approx y)$. In the former case, we have that \mathcal{J} is a model for $\neg\exists z. Q[z]$ by definition, which implies \mathcal{J} is a model for $\neg\exists x. Q[s_k[x]]$, which is $\neg\exists x. P$. In that case we have that \mathcal{J} is an extension of some interpretation that satisfies $\neg\exists x. s_k[x] \approx y$ by the induction hypothesis, which is a contraction since y can be interpreted the same as value as $s_k[u]$ for an arbitrary u and hence that formula has no models. Thus, \mathcal{J} is a model for $\neg\exists x. P$. The restriction of \mathcal{J} to the free variables of this formula, which is \mathcal{I} , also satisfies it and the statement holds. \square