

An Automatable Formal Semantics for IEEE-754 Floating-Point Arithmetic

Last revised: Wednesday 19th November, 2014

Martin Brain
University of Oxford

Cesare Tinelli
The University of Iowa

Philipp Rümmer
Uppsala University

Thomas Wahl
Northeastern University

Abstract—Automated reasoning tools often provide little or no support to reason accurately and efficiently about floating-point arithmetic. As a consequence, software verification systems that use these tools are unable to reason reliably about programs containing floating-point calculations or may give unsound results. These deficiencies are in stark contrast to the increasing awareness that the improper use of floating-point arithmetic in programs can lead to unintuitive and harmful defects in software. To promote coordinated efforts towards building floating-point reasoning engines, this paper presents a formalisation of the IEEE-754 standard for floating-point arithmetic as a theory in many-sorted first-order logic. Benefits include a standardised syntax and unambiguous semantics, allowing tool interoperability and sharing of benchmarks, and providing a basis for automated, formal analysis of programs using floating-point data.

I. INTRODUCTION

Real values can be represented in a computer in many ways, with various level of precision: as fixed-point numbers, binary or decimal floating-point numbers, rationals, arbitrary precision reals, etc. Due to the wide availability of high-performance hardware and support in most programming languages, binary floating-point has become the dominant representation system. Aside from some industry-specific exceptions, most programs that interact with or model the real world, not least those employed in safety critical applications, rely on binary floating-point arithmetic.

This creates a significant challenge for program analysis tools: accurate reasoning about the behaviour of (numerical) programs is only possible with bit-accurate reasoning about floating-point arithmetic. Many automated verification tools, such as software model checkers, rely on Satisfiability Modulo Theories (SMT) solvers as their reasoning engines. A distinguishing feature of these solvers is their use of specialized, built-in methods to check the satisfiability of formulas in *background theories* of interest, such as for instance the theories of integer numbers, arrays, bit vectors and so on. Reasoning about floating-point accurately in SMT then requires the identification of a suitable theory of floating-point numbers. Due to significant semantic differences, reasoning in substitute theories such as the reals would generally render the solvers unsound and is thus not a satisfactory option.

In the past, designing a formal theory of floating-point arithmetic would have been prohibitively complex, as different manufactures used different floating-point formats which varied from the others in significant, structural aspects. The introduction, in 1985, and subsequent near universal adoption of the IEEE-754 floating-point standard has considerably improved the situation. However, the standard is unsatisfactory as a formal theory definition for a number of reasons: it is written in natural language; it covers various aspects that are not relevant to developing a theory of floating-point (see Section VII-B); it is lengthy (the 2008 revision has 70 pages) and, most critically for automated reasoning purposes, it does not describe a logical signature and interpretations.

This paper presents the syntax and semantics of a logical theory that formalizes floating-point arithmetic based on the IEEE-754 standard. The models of our theory correspond exactly to conforming implementations of IEEE-754. While rather general, our formalization was explicitly developed, with inputs from the SMT community, to be a reference theory for SMT solver implementors and users. Our theory was recently incorporated in the SMT-LIB standard [4], a widely used input/output language for SMT solvers.

This paper makes three specific contributions:

- 1) Discusses the challenges of reasoning about floating-point numbers and illustrates its critical uses. [Section II]
- 2) Presents mathematical structures intended to formally model binary floating-point arithmetic. [Section IV]
- 3) Provides a signature for a theory of floating-point arithmetic and an interpretation of its function and predicate symbols in terms of the mathematical structures defined earlier. [Section V]

II. FLOATING-POINT ARITHMETIC

Floating-point refers to a way of encoding subsets of the rational numbers using bit vectors of fixed width. A floating-point number consists of three such bit vectors: one for the fractional part of the number, called the *significand*, one for the *exponent* of an integer power by which the fractional part is multiplied, and a single bit for the *sign*. For example:

$$\begin{aligned} 7.28125 &= 7 + 28125 \cdot 10^{-5} = (2^2 + 2^1 + 2^0) + (2^{-2} + 2^{-5}) \\ &= 111.01001_2 = 1.1101001_2 \cdot 2^2 \end{aligned}$$

Work partially supported by the Toyota Motor Corporation, ERC project 280053, EPSRC project EP/H017585/1, DSTL under CDE Project 30713, and NSF grant CCF-1218075.

can be represented as a *binary* floating-point number with a sign bit of 0, an exponent of 2 and a significand of 1101001: the leading 1 before the binary point is omitted (for so-called *normal* numbers) and hence known as the *hidden bit*.

Arithmetic on floating-point numbers is defined as performing the operation *as if the numbers were reals*, followed by rounding the result to the nearest value representable as a floating-point number; “nearest” is defined by a set of rules called a *rounding mode*. Many floating-point systems implemented in computer processors have included special values such as infinities and not-a-number (NaN). When an underflow, overflow or other exceptional condition occurs, these special values can be returned instead of triggering an interrupt. This can simplify the control circuitry and results in faster computation but at the cost of making the floating-point number system more complex. IEEE-754 standardizes, in different floating-point formats, the sizes of the bit vectors used for number presentation, as well as the various rounding modes, and the meaning and use of special values.

Formal reasoning about algorithms and programs that use floating-point numbers has acquired a reputation of being difficult, complex and error prone [21]. There are at least two causes for this. One is the combinatorial explosion in operations due to the existence of special computational values. More precisely, IEEE-754 defines five classes of numbers: normal, subnormal (which cannot be expressed in normal form, due to limits on the exponent range), zeros (a positive and a negative one, so that approximation from above and below can be distinguished), infinities (one positive and one negative), and NaN. The behaviour of each operation depends on the classes of its arguments, potentially requiring a five-way case distinction for every operation considered.

The other cause for the intricacies of floating-point reasoning is the consequences of rounding. Rounding is vital as it permits a fixed bound on the amount of memory required to store and compute with floating-point numbers. However, it also causes floating-point arithmetic to defy some fundamental laws of traditional arithmetics, such as the associativity of addition.¹ As a result, many algebraic and automated reasoning techniques based on these laws are inapplicable to it.

While challenging, reasoning formally and possibly automatically about floating-point numbers is crucial for many purposes:

Identifying the Generation of Special-Values It is generally useful to know whether a program may generate the special floating-point values infinity and NaN. While their presence is not necessarily a bug, they fall outside classical real-number arithmetic. In some cases, the detection of subnormal numbers is valuable as they have limited precision and can thus amplify errors.

Detecting Numerical Instabilities As a consequence of rounding, there are non-zero, normal numbers x and y such that $x+y = x$. Although again this is not necessarily

a bug, it is counter-intuitive for many programmers and may be a symptom of deeper algorithmic problems. Similarly, *catastrophic cancellation* occurs in $x+y$ when the exponents of x and y are (near-)equal and their signs differ. In this case the most significant bits can cancel, effectively amplifying the errors introduced by rounding.

Exposing Undefined Behaviour With few exceptions, IEEE-754 fully defines the results of all operations on floating-point numbers but does not cover corner cases of conversions to other formats. Although this may seem like a trivial issue, it was the immediate cause of the loss of Ariane 5 Flight 501.

Generating Test-Cases Many safety standards for embedded software mandate testing coverage criteria. Automatic test case generation is vital as system complexity rises, and is particularly important for showing which paths are infeasible, as developers often miss floating-point related corner cases. These applications demand model generating and bit-precise reasoning engines.

Proving Functional Correctness Formal correctness specifications for numerical programs often include bounds on the output. Checking those bounds is a suitable task for a reasoning tool. Another common correctness measure is the difference between the result of the computation performed in floating-point and its real-arithmetic result. Care must be taken because most non-trivial numerical algorithms are not simple transcriptions of mathematical expressions into software; comparing the same algorithm over the two domains may be misleading.

All these scenarios have in common that classical mathematical reasoning, regardless of whether it is manual or automated, is ineffective when dealing with floating-point. In the rest of this paper we present our formalisation of the IEEE-754 standard as a logical theory, with the goal of facilitating the construction of automated and precise reasoning tools for it. We concentrate on arithmetic aspects, abstracting away more operational ones, such as exception handling. Also, we only consider the case of binary (as opposed to decimal) floating-point arithmetic, since it is more widely used in practice.

III. FORMAL FOUNDATIONS

IEEE-754 gives an informal definition of the semantics of floating-point operations:

Each of the computational operations that return a numeric result specified by this standard shall be performed as if it first produced an intermediate result correct to infinite precision and with unbounded range, and then rounded that intermediate result, if necessary, to fit in the destination’s format.

To formalise this it is helpful to define an extension of the real numbers so that we can treat floating-point values as if they had “infinite precision and unbounded range,” and to define a notion of rounding. Note that the extended reals are used here simply to aid the formal definition of floating-point operations. They are not the domain of interpretation of floating-point numbers. For that we will define a family of algebras over bit vector triples.

¹Since rounding is performed after every operation, the expressions $(a + b) + c$ and $a + (b + c)$ can evaluate to different values.

TABLE I
FORMALIZATION OF NaN'S BEHAVIOUR, WITH $u \in S^*$

$$\begin{aligned} u + \text{NaN} &= \text{NaN} + u = \text{NaN} & -\text{NaN} &= \text{NaN} \\ u \cdot \text{NaN} &= \text{NaN} \cdot u = \text{NaN} & \text{NaN}^{-1} &= \text{NaN} \\ \text{NaN} \leq u &\Leftrightarrow u = \text{NaN} & u \leq \text{NaN} &\Leftrightarrow u = \text{NaN} \end{aligned}$$

TABLE II
DEFINED SYMBOLS, WITH $x, y \in S^*$

$$\begin{aligned} x - y &:= x + (-y) & x \geq y &:= y \leq x \\ x/y &:= x \cdot y^{-1} & x < y &:= (x \leq y) \wedge \neg(x = y) \\ & & x > y &:= (x \geq y) \wedge \neg(x = y) \end{aligned}$$

A. Extended Reals

We extend the set \mathbb{R} of real numbers with three new elements: $+\infty$, $-\infty$ and NaN, which represent respectively positive and negative infinity and a special *not a number* value used to obtain an arithmetically closed system. For each set $S \subseteq \mathbb{R}$ we can define the following sets:

$$S^\dagger = S \cup \{+\infty, -\infty\} \quad S^* = S^\dagger \cup \{\text{NaN}\}$$

When S is the base of an ordered additive or multiplicative monoid or group, we extend \mathbb{R} 's binary operations $+$ and \cdot , unary operations $-$ and $(_)^{-1}$ and order relation \leq . Table I gives axioms defining these operations when one argument is NaN. Table III gives axioms defining these operations when one argument is $+\infty$ or $-\infty$, as well as the axiom for the inverse of zero. We remark that, although these definitions extend the operations of ordered fields and rings and can be applied to \mathbb{R} and \mathbb{Z} , the extended reals \mathbb{R}^* and extended integers \mathbb{Z}^* do not have all of the structure of \mathbb{R} and \mathbb{Z} .²

Note that (\mathbb{R}^*, \leq) is a partial order since NaN is comparable only with itself. In contrast $(\mathbb{R}^\dagger, \leq)$ is a total order. The extended reals operate as three largely algebraically independent subsets: $\{\text{NaN}\}$, $\{+\infty, -\infty\}$ and \mathbb{R} . If a sub-expression of an expression e evaluates to NaN, then the whole e evaluates to NaN — the set is closed under the basic operations. Infinities generate infinities or NaN, although the reciprocal operator maps an infinity back to a real value. Reals are of course closed under all operations except reciprocal of zero. For convenience, we will also use the usual additional symbols in Table II, which are definable in terms of the basic operations.

B. Rounding

The second concept needed to formalise the IEEE-754 definition of operations is that of rounding; a map that will take the intermediate result in \mathbb{R}^* back into the set of floating-point numbers. This will be defined as a function which selects between the two adjoints of the corresponding map into \mathbb{R}^* .

More generally, let (X, \sqsubseteq) be a partially ordered set that consists of one or more disjoint lattices, and let $v : X \rightarrow \mathbb{R}^*$ be an order-embedding function from X into the extended reals

² \mathbb{R}^* is neither an additive or multiplicative group as $(-\text{NaN}) + \text{NaN} \neq 0$ and $\text{NaN}^{-1} \cdot \text{NaN} \neq 1$, nor does it have an annihilating 0 as $\text{NaN} \cdot 0 \neq 0$. However it is an additive and multiplicative commutative monoid with the distributivity property, a structure some authors refer to as a semi-ring.

TABLE III
FORMALIZATION OF $\pm\infty$ 'S AND INVERSE OF ZERO'S BEHAVIOR, $w \in S^\dagger$

$$\begin{aligned} +\infty \leq w &\Leftrightarrow w = +\infty & w \leq -\infty &\Leftrightarrow w = -\infty \\ w \leq +\infty & & -(+\infty) &= -\infty & +\infty^{-1} &= 0 \\ -\infty \leq w & & -(-\infty) &= +\infty & -\infty^{-1} &= 0 \\ w + (+\infty) &= (+\infty) + w = \begin{cases} \text{NaN} & \text{if } w = -\infty \\ +\infty & \text{if } w \neq -\infty \end{cases} \\ w + (-\infty) &= (-\infty) + w = \begin{cases} \text{NaN} & \text{if } w = +\infty \\ -\infty & \text{if } w \neq +\infty \end{cases} \\ w \cdot +\infty &= +\infty \cdot w = \begin{cases} +\infty & \text{if } 0 < w \\ -\infty & \text{if } w < 0 \\ \text{NaN} & \text{if } w = 0 \end{cases} \\ w \cdot -\infty &= -\infty \cdot w = \begin{cases} -\infty & \text{if } 0 < w \\ +\infty & \text{if } w < 0 \\ \text{NaN} & \text{if } w = 0 \end{cases} \\ 0^{-1} &= +\infty \end{aligned}$$

such that $\{+\infty, -\infty, \text{NaN}\} \subset v(X)$. Then, the *upper adjoint* and *lower adjoint* of v are respectively the unique functions $\bar{v} : \mathbb{R}^* \rightarrow X$ and $\underline{v} : \mathbb{R}^* \rightarrow X$ such that for all $r \in \mathbb{R}^*$.

- $r \leq v(\bar{v}(r))$ and $\bar{v}(r) \sqsubseteq x$ for all $x \in X$ with $r \leq v(x)$;
- $v(\underline{v}(r)) \leq r$ and $x \sqsubseteq \underline{v}(r)$ for all $x \in X$ with $v(x) \leq r$.

The function \bar{v} maps an element r to the smallest element of X that projects above r (rounding up) while \underline{v} maps r to the largest element of X that projects below r (rounding down). Let $\mathbb{B} = \{\top, \perp\}$ be the Booleans, with \top being the true value. We define a family of (higher-order) rounding functions:

$$\text{round} : \text{RM} \times \mathbb{B} \times \mathbb{R}^* \rightarrow (X \rightarrow \mathbb{R}^*) \rightarrow (\mathbb{R}^* \rightarrow X)$$

parametrized by the partially ordered set X , which provides a systematic way of selecting between rounding up and rounding down. Given a map $v : X \rightarrow \mathbb{R}^*$, the rounding function returns one of v 's two adjoints, based on three previous inputs. The first is the *rounding mode*, chosen from the set:

$$\text{RM} = \{\text{rne}, \text{rna}, \text{rtp}, \text{rtn}, \text{rtz}\}$$

which represents the five rounding modes defined by IEEE-754, namely, round to nearest with ties picking even value (rne), round to nearest with ties away from zero (rna), round towards $+\infty$ (rtp), round towards $-\infty$ (rtn), and round towards zero (rtz). The second input of round is a Boolean value determining the sign of zero when X contains signed zeros (which is the case when X is a set of floating-point numbers). The third input is the value to be rounded, needed because the rounding direction may depend on it (for example, when rounding to the nearest element of X).

The function round is defined in Table IV. The definition relies on three auxiliary predicates lh_X , tb_X and ev_X whose own definition depend on the particular domain X . These express: when the value is in the *lower half* of the interval between two representations in X (i.e. closer to $\underline{v}(r)$ than $\bar{v}(r)$); the *tie-break* condition when it is equal distance from both; and whether a representation is even. For illustration purposes, we provide a definition of those predicates in Table V for when

TABLE IV
DEFINITION OF round

$$\begin{aligned}
\text{round}(\text{rne}, s, r)(v) &= \begin{cases} \bar{v} & r \neq 0, \neg \text{lh}_X(r, v), \neg \text{tb}_X(r, v) \\ \bar{v} & r \neq 0, \text{tb}_X(r, v), \text{ev}_X(\bar{v}(r)) \\ \underline{v} & r \neq 0, \text{tb}_X(r, v), \text{ev}_X(\underline{v}(r)) \\ \underline{v} & r \neq 0, \text{lh}_X(r, v) \\ \text{rsz}(s)(v) & r = 0 \\ \underline{v} & r = \text{NaN} \end{cases} \\
\text{round}(\text{rna}, s, r)(v) &= \begin{cases} \bar{v} & r > 0, \neg \text{lh}_X(r, v) \\ \underline{v} & r > 0, \text{lh}_X(r, v) \\ \text{rsz}(s)(v) & r = 0 \\ \bar{v} & r < 0, \neg \text{lh}_X(r, v), \neg \text{tb}_X(r, v) \\ \underline{v} & r < 0, \text{lh}_X(r, v) \vee \text{tb}_X(r, v) \\ \underline{v} & r = \text{NaN} \end{cases} \\
\text{round}(\text{rtp}, s, r)(v) &= \begin{cases} \text{rsz}(s)(v) & r = 0 \\ \bar{v} & \text{otherwise} \end{cases} \\
\text{round}(\text{rtn}, s, r)(v) &= \begin{cases} \text{rsz}(s)(v) & r = 0 \\ \underline{v} & \text{otherwise} \end{cases} \\
\text{round}(\text{rtz}, s, r)(v) &= \begin{cases} \underline{v} & r > 0 \\ \text{rsz}(s)(v) & r = 0 \\ \bar{v} & \text{otherwise} \end{cases} \\
\text{where} & \quad \text{rsz}(\top)(v) = \bar{v} \quad \text{rsz}(\perp)(v) = \underline{v}
\end{aligned}$$

TABLE V
AUXILIARY PREDICATES FOR ROUND IN THE CASE $X = \mathbb{Z}^*$

$$\begin{aligned}
\text{lh}_{\mathbb{Z}^*}(r, v) &:= r - v(\underline{v}(r)) < v(\bar{v}(r)) - r \\
\text{tb}_{\mathbb{Z}^*}(r, v) &:= r - v(\underline{v}(r)) = v(\bar{v}(r)) - r \\
\text{ev}_{\mathbb{Z}^*}(x) &:= \exists z \in \mathbb{Z}. x = 2 * z
\end{aligned}$$

$X = \mathbb{Z}$, the set of integers with the usual ordering. A definition of those predicates for sets of floating-point numbers is given later, after we formalize such sets.

The fairly elaborate definition of round is motivated by our goal to provide an accurate model of rounding as defined in IEEE-754. In particular, there is no mathematical reason for not using exclusively \bar{v} or \underline{v} in it. However, doing so would fail to reflect some properties of IEEE-754 floating-point numbers, for example “the sign of a sum [...] differs from at most one of the addends’ signs” [1].

For brevity, we will write $\text{rnd}(v, m, s, r)$ for the application $\text{round}(m, s, r)(v)(r)$ which returns the value of X that the real number r is rounded to by using v .

IV. MODELS OF FLOATING-POINT ARITHMETIC

In this section we specify a set of (many-sorted) structures in the sense of model theory. These are the intended models of a logical theory of floating-point numbers that reflects IEEE-754. In the next section we will specify a signature for such a theory and show how each sort, function and predicate symbol in the signature is interpreted over this set of structures.

A. Universe

The universe of each of our models consists of multiple sets: one for the rounding mode sort and one for each of the different floating-point precisions. Floating-point numbers other than NaN are triples of bit vectors modelling the three

components (sign, exponent and significand) of the representations in IEEE-754. We identify bit vectors of length $\nu > 0$ with elements of the function space $\mathbb{B}\mathbb{V}_\nu = \mathbb{N}_\nu \rightarrow \{0, 1\}$ where $\mathbb{N}_\nu = \{0, \dots, \nu - 1\}$. We write $\mathbf{1}_\nu$ for the unique function in $\mathbb{N}_\nu \rightarrow \{1\}$ and $\mathbf{0}_\nu$ for the unique function in $\mathbb{N}_\nu \rightarrow \{0\}$, represent respectively the bit vector of length n containing all ones and that containing all zeros. Let $\text{ub}_\nu : \mathbb{B}\mathbb{V}_\nu \rightarrow \mathbb{N}$ and $\text{sb}_\nu : \mathbb{B}\mathbb{V}_\nu \rightarrow \mathbb{Z}$ denote the usual unsigned and 2’s complement encodings of bit vectors into integers. Let $B_{\mu, \nu}$ denote the set $\mathbb{B}\mathbb{V}_1 \times \mathbb{B}\mathbb{V}_\mu \times \mathbb{B}\mathbb{V}_{\nu-1}$. For all integers $\varepsilon > 1$ and $\sigma > 1$, we define the set of floating-point numbers with ε exponent bits and σ significand bits³ as the set:

$$\mathbb{F}_{\varepsilon, \sigma} = \mathbb{F}_{\varepsilon, \sigma} \cup \{\text{NaN}\}$$

where

$$\begin{aligned}
\mathbb{F}_{\varepsilon, \sigma} &= \mathbb{F}\mathbb{Z}_{\varepsilon, \sigma} \cup \mathbb{F}\mathbb{S}_{\varepsilon, \sigma} \cup \mathbb{F}\mathbb{N}_{\varepsilon, \sigma} \cup \mathbb{F}\mathbb{I}_{\varepsilon, \sigma} \\
\mathbb{F}\mathbb{Z}_{\varepsilon, \sigma} &= \{(s, e, m) \in B_{\varepsilon, \sigma} \mid e = \mathbf{0}_\varepsilon, m = \mathbf{0}_{\sigma-1}\} \\
\mathbb{F}\mathbb{S}_{\varepsilon, \sigma} &= \{(s, e, m) \in B_{\varepsilon, \sigma} \mid e = \mathbf{0}_\varepsilon, m \neq \mathbf{0}_{\sigma-1}\} \\
\mathbb{F}\mathbb{N}_{\varepsilon, \sigma} &= \{(s, e, m) \in B_{\varepsilon, \sigma} \mid e \neq \mathbf{1}_\varepsilon, e \neq \mathbf{0}_\varepsilon\} \\
\mathbb{F}\mathbb{I}_{\varepsilon, \sigma} &= \{(s, e, m) \in B_{\varepsilon, \sigma} \mid e = \mathbf{1}_\varepsilon, m = \mathbf{0}_{\sigma-1}\}
\end{aligned}$$

The last four sets above correspond respectively to the bit vector triples used to represent zeros, subnormal numbers, normal numbers and infinities in IEEE-754, with the three components storing respectively sign, exponent and significand of the floating-point number.⁴ We will write informally -0 and $+0$ to refer to the two elements of $\mathbb{F}\mathbb{Z}_{\varepsilon, \sigma}$.

We fix a total order \sqsubseteq over $\mathbb{F}_{\varepsilon, \sigma}$ such that $(s_1, e_1, m_1) \sqsubseteq (s_2, e_2, m_2)$ if one of the following holds:

- $s_1 = 1, s_2 = 0$
- $s_1 = 0, s_2 = 0, \text{ub}_\varepsilon(e_1) < \text{ub}_\varepsilon(e_2)$
- $s_1 = 0, s_2 = 0, \text{ub}_\varepsilon(e_1) = \text{ub}_\varepsilon(e_2), \text{ub}_\sigma(m_1) \leq \text{ub}_\sigma(m_2)$
- $s_1 = 1, s_2 = 1, \text{ub}_\varepsilon(e_2) < \text{ub}_\varepsilon(e_1)$
- $s_1 = 1, s_2 = 1, \text{ub}_\varepsilon(e_1) = \text{ub}_\varepsilon(e_2), \text{ub}_\sigma(m_2) \leq \text{ub}_\sigma(m_1)$

We extend \sqsubseteq to a partial order on $\mathbb{F}_{\varepsilon, \sigma}$ by $\text{NaN} \sqsubseteq \text{NaN}$.

As discussed in Section III, we define operations over $\mathbb{F}_{\varepsilon, \sigma}$ analogously to those defined over \mathbb{R}^* by mapping floating-point values to extended reals, performing the corresponding extended reals operation and then rounding the result back to a floating-point value. To formalise this we define a function $v_{\varepsilon, \sigma} : \mathbb{F}_{\varepsilon, \sigma} \rightarrow \mathbb{R}^*$ which maps each floating-point number to the extended real it represents. Let $\text{bias}(\varepsilon) = 2^{\varepsilon-1} - 1$.

$$\begin{aligned}
v_{\varepsilon, \sigma}(f) = v_{\varepsilon, \sigma}((s, e, m)) &= \\
\begin{cases} 0 & f \in \mathbb{F}\mathbb{Z}_{\varepsilon, \sigma} \\ (-1)^{\text{ub}_1(s)} \cdot 2^{1-\text{bias}(\varepsilon)} \cdot \left(0 + \frac{\text{ub}_{\sigma-1}(m)}{2^{\sigma-1}}\right) & f \in \mathbb{F}\mathbb{S}_{\varepsilon, \sigma} \\ (-1)^{\text{ub}_1(s)} \cdot 2^{\text{ub}_\varepsilon(e) - \text{bias}(\varepsilon)} \cdot \left(1 + \frac{\text{ub}_{\sigma-1}(m)}{2^{\sigma-1}}\right) & f \in \mathbb{F}\mathbb{N}_{\varepsilon, \sigma} \\ (-1)^{\text{ub}_1(s)} \cdot (+\infty) & f \in \mathbb{F}\mathbb{I}_{\varepsilon, \sigma} \end{cases} \\
v_{\varepsilon, \sigma}(\text{NaN}) &= \text{NaN}
\end{aligned}$$

³Allowing arbitrary values for ε and σ is strictly a generalisation of IEEE-754, which only defines a handful of precisions. However, doing so supports a wider range of applications with little additional notation and effort.

⁴The significand component has length $\sigma - 1$ because the hidden bit, which is 1 for normal numbers, is not explicitly represented.

TABLE VI
AUXILIARY PREDICATES FOR ROUND IN THE CASE $X = \mathbb{F}_{\varepsilon,\sigma}$

$$\begin{aligned} \text{ev}_{\mathbb{F}_{\varepsilon,\sigma}}(f) &:= f = (s, e, m) \in \mathbb{F}_{\varepsilon,\sigma} \wedge \text{ev}_{\mathbb{Z}^*}(\text{ub}_{\sigma-1}(m)) \\ \text{lh}_{\mathbb{F}_{\varepsilon,\sigma}}(r, v) &:= \sigma' = \sigma + 1 \wedge v(\underline{v}(r)) = v_{\varepsilon,\sigma'}(\underline{v}_{\varepsilon,\sigma'}(r)) \\ \text{tb}_{\mathbb{F}_{\varepsilon,\sigma}}(r, v) &:= \sigma' = \sigma + 1 \wedge v(\underline{v}(r)) < v_{\varepsilon,\sigma'}(\underline{v}_{\varepsilon,\sigma'}(r)) = \\ &\quad v_{\varepsilon,\sigma'}(\overline{v}_{\varepsilon,\sigma'}(r)) < v_{\varepsilon,\sigma'}(\overline{v}(r)) \end{aligned}$$

For brevity we will write just v in place of $v_{\varepsilon,\sigma}$ when the values ε and σ are clear from context or not important. One can show that v is injective over $\mathbb{F}_{\varepsilon,\sigma} \setminus \text{FZ}_{\varepsilon,\sigma}$ and monotonic. Thanks to the latter we have that both \overline{v} and \underline{v} are well defined⁵ and so the function round can be used to map back from \mathbb{R}^* to $\mathbb{F}_{\varepsilon,\sigma}$. The auxiliary predicates used in the definition of rounding in the case of $X = \mathbb{F}_{\varepsilon,\sigma}$ are defined in Table VI. Both $\text{lh}_{\mathbb{F}_{\varepsilon,\sigma}}$ and $\text{tb}_{\mathbb{F}_{\varepsilon,\sigma}}$ use a set of floating-point numbers with one extra significand bit. This is equivalent to the *guard bit* used in hardware implementations, giving a point mid-way between \overline{v} and \underline{v} . The predicate $\text{tb}_{\mathbb{F}_{\varepsilon,\sigma}}$ captures the property of r being equidistant from $\underline{v}(r)$ and $\overline{v}(r)$, which means that any further significand bits would be 0. This is equivalent to the *sticky bit* used in hardware being equal to 0.

B. Relations

Having defined a universe for the models, we next define various relations which will be used as the interpretation of predicates in the theory of floating-point. Every relation is parameterised by a floating-point domain, so each definition here actually describes a whole *family* of relations.

1) *Unary Relations*: A number of unary relations (subsets) are given which allow for classification of floating-point numbers as well as determining their sign, if applicable (see Section VII-B for further discussion of sign).⁶

$$\begin{aligned} \text{isNeg}_{\varepsilon,\sigma} &= \{f \in \mathbb{F}_{\varepsilon,\sigma} \mid f = (1, e, m)\} \\ \text{isPos}_{\varepsilon,\sigma} &= \{f \in \mathbb{F}_{\varepsilon,\sigma} \mid f = (0, e, m)\} \end{aligned}$$

2) *Binary Relations*: A number of binary relations are used for comparing floating-point numbers. These are different from the equality and ordering relations on $\mathbb{F}_{\varepsilon,\sigma}$ (i.e., $=$ and \sqsubseteq) and those on \mathbb{R}^* (i.e., $=$ and \leq) and despite their names, they are not actually equality or ordering relations as they do not contain (NaN, NaN) and eq , leq and geq contain both $(+0, -0)$ and $(-0, +0)$.

$$\begin{aligned} \text{eq}_{\varepsilon,\sigma} &= \{(f, g) \in \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} \mid v(f) = v(g)\} \\ \text{leq}_{\varepsilon,\sigma} &= \{(f, g) \in \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} \mid v(f) \leq v(g)\} \\ \text{lt}_{\varepsilon,\sigma} &= \{(f, g) \in \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} \mid v(f) < v(g)\} \\ \text{geq}_{\varepsilon,\sigma} &= \{(f, g) \in \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} \mid v(f) \geq v(g)\} \\ \text{gt}_{\varepsilon,\sigma} &= \{(f, g) \in \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} \mid v(f) > v(g)\} \end{aligned}$$

⁵These are surjections for all points except $\text{FZ}_{\varepsilon,\sigma}$ which has the curious property that $-0 = \overline{v}(0) \sqsubseteq \underline{v}(0) = +0$.

⁶These definitions imply $f = \text{NaN} \Leftrightarrow \neg(\text{isNeg}_{\varepsilon,\sigma}(f) \vee \text{isPos}_{\varepsilon,\sigma}(f))$.

C. Operations

Similarly to relations, we define families (parameterised by domains) of functions which will serve as the interpretation of various operations in the theory of floating-point.

1) *Sign Operations*: The first two operations, negation and absolute value, manipulate the sign of the number. As the domains of floating-point numbers are symmetric around 0, there is no need for rounding and the operations can be defined directly on the floating-point bit vectors without using \mathbb{R}^* .

$$\begin{aligned} \text{neg}_{\varepsilon,\sigma} : \mathbb{F}_{\varepsilon,\sigma} &\rightarrow \mathbb{F}_{\varepsilon,\sigma} \\ \text{neg}_{\varepsilon,\sigma}(f) &= \begin{cases} (-s, e, m) & f = (s, e, m) \in \mathbb{F}_{\varepsilon,\sigma} \\ \text{NaN} & f = \text{NaN} \end{cases} \\ \text{abs}_{\varepsilon,\sigma} : \mathbb{F}_{\varepsilon,\sigma} &\rightarrow \mathbb{F}_{\varepsilon,\sigma} \\ \text{abs}_{\varepsilon,\sigma}(f) &= \begin{cases} (0, e, m) & f = (s, e, m) \in \mathbb{F}_{\varepsilon,\sigma} \\ \text{NaN} & f = \text{NaN} \end{cases} \end{aligned}$$

2) *Arithmetic Operations*: The main operations on floating-point numbers are those that correspond to the operations on an ordered field. These are defined using v to map the floating-point arguments to \mathbb{R}^* , then the operations are performed in \mathbb{R}^* and finally the result is mapped back with round .

$$\begin{aligned} \text{add}_{\varepsilon,\sigma} : \text{RM} \times \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} &\rightarrow \mathbb{F}_{\varepsilon,\sigma} \\ \text{add}_{\varepsilon,\sigma}(rm, f, g) &= \text{rnd}(v, rm, \text{addSign}(rm, f, g), v(f) + v(g)) \\ \text{sub}_{\varepsilon,\sigma} : \text{RM} \times \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} &\rightarrow \mathbb{F}_{\varepsilon,\sigma} \\ \text{sub}_{\varepsilon,\sigma}(rm, f, g) &= \text{rnd}(v, rm, \text{subSign}(rm, f, g), v(f) - v(g)) \\ \text{mul}_{\varepsilon,\sigma} : \text{RM} \times \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} &\rightarrow \mathbb{F}_{\varepsilon,\sigma} \\ \text{mul}_{\varepsilon,\sigma}(rm, f, g) &= \text{rnd}(v, rm, \text{mulSign}(f, g), v(f) * v(g)) \\ \text{div}_{\varepsilon,\sigma} : \text{RM} \times \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} &\rightarrow \mathbb{F}_{\varepsilon,\sigma} \\ \text{div}_{\varepsilon,\sigma}(rm, f, g) &= \begin{cases} \text{neg}_{\varepsilon,\sigma}(\text{rnd}(v, rm, \top, -(v(f)/v(g)))) & \text{xorSign}(f, g) \\ \text{rnd}(v, rm, \perp, v(f)/v(g)) & \neg\text{xorSign}(f, g) \end{cases} \\ \text{fma}_{\varepsilon,\sigma} : \text{RM} \times \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} \times \mathbb{F}_{\varepsilon,\sigma} &\rightarrow \mathbb{F}_{\varepsilon,\sigma} \\ \text{fma}_{\varepsilon,\sigma}(rm, f, g, h) &= \text{rnd}(v, rm, \text{fmaSign}(rm, f, g, h), (v(f) * v(g)) + v(h)) \end{aligned}$$

The addSign , subSign , xorSign and fmaSign predicates are defined as follows (\oplus denotes exclusive or):

$$\begin{aligned} \text{addSign}(rm, f, g) &:= \begin{cases} \text{isNeg}(f) \wedge \text{isNeg}(g) & rm \neq \text{rtn} \\ \text{isNeg}(f) \vee \text{isNeg}(g) & rm = \text{rtn} \end{cases} \\ \text{subSign}(rm, f, g) &:= \text{addSign}(rm, f, \text{neg}_{\varepsilon,\sigma}(g)) \\ \text{xorSign}(f, g) &:= \text{isNeg}(f) \oplus \text{isNeg}(g) \\ \text{fmaSign}(rm, f, g, h) &:= \text{addSign}(rm, \text{mul}_{\varepsilon,\sigma}(rm, f, g), h) \end{aligned}$$

Note that since $\text{fma}_{\varepsilon,\sigma}$ only calls round once, it is not the same as $\text{add}_{\varepsilon,\sigma}(rm, \text{mul}_{\varepsilon,\sigma}(rm, a, b), c)$. Also, $\text{div}_{\varepsilon,\sigma}$ is equal to $\text{rnd}(v, rm, \text{xorSign}(f, g), v(f)/v(g))$ at all points except positive numbers divided by -0 , where the ‘obvious’

definition gives positive infinity while the definition given above gives the correct result of minus infinity.

3) *Additional operations:* IEEE-754 defines a square root function which returns the floating-point number nearest to the square root of the real represented by the input number. Also $\text{sqrt}(rm, -0) = -0$ since -0 represents $0 \in \mathbb{R}^*$, whose square root is $0 \in \mathbb{R}^*$ and the sign is inherited when rounding is performed.

$$\begin{aligned} \text{sqrt}_{\varepsilon, \sigma} : \text{RM} \times \mathbb{F}_{\varepsilon, \sigma} &\rightarrow \mathbb{F}_{\varepsilon, \sigma} \\ \text{sqrt}_{\varepsilon, \sigma}(rm, f) &= \begin{cases} \text{rnd}(v, rm, \text{isNeg}(f), z) & 0 \leq x = v(f), \\ & z \cdot z = x, 0 \leq z \\ \text{NaN} & \text{otherwise} \end{cases} \end{aligned}$$

Although this is the natural definition of square root, it results in a few unexpected consequences. For example $\text{mul}(rm, \text{sqrt}(rm_1, x), \text{sqrt}(rm_2, x))$ is not guaranteed to be equal to x and, depending on the rounding modes rm, rm_1, rm_2 , may give an infinity, even when x is a normal number.⁷

Let $w_{\varepsilon, \sigma}$ be the restriction of $v_{\varepsilon, \sigma}$ to $\text{FI}_{\varepsilon, \sigma} \cup \{\text{NaN}\} \cup \{f \in \mathbb{F}_{\varepsilon, \sigma} \mid v_{\varepsilon, \sigma}(f) \in \mathbb{Z}\}$. The following operation rounds back to a subset of $\mathbb{F}_{\varepsilon, \sigma}$, effectively mapping the value to the nearest integer representable in the given floating-point format:

$$\begin{aligned} \text{rti}_{\varepsilon, \sigma} : \text{RM} \times \mathbb{F}_{\varepsilon, \sigma} &\rightarrow \mathbb{F}_{\varepsilon, \sigma} \\ \text{rti}_{\varepsilon, \sigma}(rm, f) &= \text{rnd}(w_{\varepsilon, \sigma}, rm, \text{isNeg}(f), v(r)) \end{aligned}$$

Let $\text{in} : \mathbb{Z}^* \rightarrow \mathbb{R}^*$ with $\text{in}(z) = z$.⁸ Similarly to $\text{rti}_{\varepsilon, \sigma}$, the remainder operation requires rounding an intermediate value to an integer:

$$\begin{aligned} \text{rem}_{\varepsilon, \sigma} : \text{RM} \times \mathbb{F}_{\varepsilon, \sigma} \times \mathbb{F}_{\varepsilon, \sigma} &\rightarrow \mathbb{F}_{\varepsilon, \sigma} \\ \text{rem}_{\varepsilon, \sigma}(rm, f, g) &= \begin{cases} f & g \in \text{FI}_{\varepsilon, \sigma} \\ \text{NaN} & f \in \text{FI}_{\varepsilon, \sigma} \cup \{\text{NaN}\} \\ \text{NaN} & g \in \text{FZ}_{\varepsilon, \sigma} \cup \{\text{NaN}\} \\ \text{rnd}(v, rm, \text{isNeg}(f), x) & x = v(f) - (v(g) * y), \\ & y = \text{rnd}(\text{in}, rm, 0, v(f)/v(g)) \end{cases} \\ \text{remrne}_{\varepsilon, \sigma} : \mathbb{F}_{\varepsilon, \sigma} \times \mathbb{F}_{\varepsilon, \sigma} &\rightarrow \mathbb{F}_{\varepsilon, \sigma} \\ \text{remrne}_{\varepsilon, \sigma}(f, g) &= \text{rem}_{\varepsilon, \sigma}(\text{rne}, f, g) \end{aligned}$$

Note that the remainder computed as above is always exact when rne is used. This remainder function is the one used by the C standard library but is not necessarily the same as the intuitive idea of remainder which can be computed via: $\text{fma}_{\varepsilon, \sigma}(rm, \text{neg}_{\varepsilon, \sigma}(\text{div}_{\varepsilon, \sigma}(rm, f, g)), g, f)$.

⁷ $x \in \mathbb{R}^*$ and thus z is uniquely defined. Given $f \in \mathbb{F}_{\varepsilon, \sigma}$ there is not necessarily a $g \in \mathbb{F}_{\varepsilon, \sigma}$ such that $f = \text{mul}_{\varepsilon, \sigma}(rm, g, g)$ and $\text{leq}_{\varepsilon, \sigma}(0, g)$. Furthermore, in the case of subnormal numbers, there are potentially multiple, distinct, positive square roots. Thus care must be taken with implicit definitions of sqrt just using mul or fma .

⁸We consider \mathbb{Z} a subset of \mathbb{R} and hence \mathbb{Z}^* a subset of \mathbb{R}^* .

The next two operations, the maximum and minimum of two floating-point numbers, are only partially specified in IEEE-754. When the value of the two arguments (in \mathbb{R}^*) is the same, then one or the other of the arguments can be returned.

$$\begin{aligned} \text{max}_{\varepsilon, \sigma} : \mathbb{F}_{\varepsilon, \sigma} \times \mathbb{F}_{\varepsilon, \sigma} &\rightarrow \mathbb{F}_{\varepsilon, \sigma} \\ \text{max}_{\varepsilon, \sigma}(f, g) &= \begin{cases} f & \text{gt}_{\varepsilon, \sigma}(f, g) \text{ or } g = \text{NaN} \\ g & \text{gt}_{\varepsilon, \sigma}(g, f) \text{ or } f = \text{NaN} \\ h & h \in \{f, g\}, \text{eq}_{\varepsilon, \sigma}(f, g) \end{cases} \\ \text{min}_{\varepsilon, \sigma} : \mathbb{F}_{\varepsilon, \sigma} \times \mathbb{F}_{\varepsilon, \sigma} &\rightarrow \mathbb{F}_{\varepsilon, \sigma} \\ \text{min}_{\varepsilon, \sigma}(f, g) &= \begin{cases} f & \text{lt}_{\varepsilon, \sigma}(f, g) \text{ or } g = \text{NaN} \\ g & \text{lt}_{\varepsilon, \sigma}(g, f) \text{ or } f = \text{NaN} \\ h & h \in \{f, g\}, \text{eq}_{\varepsilon, \sigma}(f, g) \end{cases} \end{aligned}$$

The two functions above are underspecified in case the input values f, g satisfy $\text{eq}_{\varepsilon, \sigma}$ (which is only when one is -0 and the other is $+0$). This means that we consider as acceptable models any structures with function families $\text{max}_{\varepsilon, \sigma}$ and $\text{min}_{\varepsilon, \sigma}$ that satisfy the specifications above, regardless of whether they return -0 or $+0$ for $(-0, +0)$, and for $(+0, -0)$. Although underspecification is inconvenient for implementors, it is necessary as IEEE-754 itself allows either value to be returned and compliant implementations do vary. For example, on some Intel processors the result returned by the x87 and SSE units is different.

All the preceding operations have floating-point input and outputs in the same set, $\mathbb{F}_{\varepsilon, \sigma}$. To convert between different floating-point domains the following map is needed:

$$\begin{aligned} \text{cast}_{\varepsilon, \sigma, \varepsilon', \sigma'} : \text{RM} \times \mathbb{F}_{\varepsilon', \sigma'} &\rightarrow \mathbb{F}_{\varepsilon, \sigma} \\ \text{cast}_{\varepsilon', \sigma', \varepsilon, \sigma}(rm, f) &= \text{rnd}(v_{\varepsilon, \sigma}, rm, \text{isNeg}(f), v_{\varepsilon', \sigma'}(f)) \end{aligned}$$

If $\varepsilon \geq \varepsilon'$ and $\sigma \geq \sigma'$, the rounding mode argument is irrelevant since then all values of $\mathbb{F}_{\varepsilon', \sigma'}$ are representable exactly in $\mathbb{F}_{\varepsilon, \sigma}$. However, this cannot be regarded as syntactic sugar because whether a value is a normal or a subnormal number does change depending on the floating-point domain.

D. Combinations with Other Theories

For many applications, the theory of floating-point is not sufficient to reason about the full problem; other theories such as integers, bit vectors, or reals are needed as well. This section describes the extensions required to the intended models to account for these additional domains, and possible mappings between them. IEEE-754 includes a number of functions to convert to “integer formats.” We define here such conversions as well as extensions to IEEE-754 covering conversion to real and from real and bit vectors. Many of the additional operations are underspecified in that *out of bounds* and other *error* conditions do not have specified values.

1) *Real Numbers:* For a model of the theory of floating-point to also be a model of the theory of reals, its universe has to be extended to a disjoint union with \mathbb{R} . Using the connections between $\mathbb{F}_{\varepsilon, \sigma}$ and \mathbb{R}^* , we add the following two conversion operations:

$$\begin{aligned}
\text{realToFP}_{\varepsilon,\sigma} &: \mathbb{R}M \times \mathbb{R} \rightarrow \mathbb{F}_{\varepsilon,\sigma} \\
\text{realToFP}_{\varepsilon,\sigma}(rm, r) &= \text{rnd}(v, rm, 0, r) \\
\text{FPToReal}_{\varepsilon,\sigma} &: \mathbb{F}_{\varepsilon,\sigma} \rightarrow \mathbb{R} \\
\text{FPToReal}_{\varepsilon,\sigma}(f) &= \begin{cases} v(f) & v(f) \in \mathbb{R} \\ x & x \in \mathbb{R}, \text{ otherwise} \end{cases}
\end{aligned}$$

We do not specify what the value of $\text{FPToReal}_{\varepsilon,\sigma}(f)$ is when f does not correspond to a real number. This means again that we accept as a model any structure with a function family $\text{FPToReal}_{\varepsilon,\sigma}$ that satisfies the specification above.

2) *Fixed-size Bit Vectors*: Similarly to the previous case, to form a joint model of the theories of floating-point and fixed-width bit vectors, the domain must be extended to a disjoint union with \mathbb{BV}_ν for every $\nu > 0$. Let $\bullet : \mathbb{BV}_\mu \times \mathbb{BV}_\nu \rightarrow \mathbb{BV}_{\mu+\nu}$ be the bit vector concatenation function for each $\mu, \nu > 0$. The following function converts a bit vector of length $\varepsilon + \sigma$, with $\varepsilon, \sigma > 1$, to a floating-point number in $\mathbb{F}_{\varepsilon,\sigma}$ by slicing the bit vector in three:

$$\begin{aligned}
\text{bitpatternToFP}_{\varepsilon,\sigma} &: \mathbb{BV}_{\varepsilon+\sigma} \rightarrow \mathbb{F}_{\varepsilon,\sigma} \\
\text{bitpatternToFP}_{\varepsilon,\sigma}(b) &= \begin{cases} (s, e, m) & b = s \bullet e \bullet m, \\ & (s, e, m) \in \mathbb{F}_{\varepsilon,\sigma} \\ \text{NaN} & \text{otherwise} \end{cases}
\end{aligned}$$

The function bitpatternToFP is not injective as there are multiple bit-patterns which represent NaN. This implies that it is not possible to give a reverse map without fixing the encoding of NaN to a particular value.

The next two functions first convert the bit vector to the integer value it denotes in binary, in 2's complement and unsigned format respectively, and then round that value to the corresponding floating-point. The last two functions do the inverse conversion.

$$\begin{aligned}
\text{sIntToFP}_{\nu,\varepsilon,\sigma} &: \mathbb{R}M \times \mathbb{BV}_\nu \rightarrow \mathbb{F}_{\varepsilon,\sigma} \\
\text{sIntToFP}_{\nu,\varepsilon,\sigma}(rm, b) &= \text{rnd}(v, rm, 0, \text{sb}_\nu(b)) \\
\text{uIntToFP}_{\nu,\varepsilon,\sigma} &: \mathbb{R}M \times \mathbb{BV}_\nu \rightarrow \mathbb{F}_{\varepsilon,\sigma} \\
\text{uIntToFP}_{\nu,\varepsilon,\sigma}(rm, b) &= \text{rnd}(v, rm, 0, \text{ub}_\nu(b)) \\
\text{FPToSInt}_{\nu,\varepsilon,\sigma} &: \mathbb{R}M \times \mathbb{F}_{\varepsilon,\sigma} \rightarrow \mathbb{BV}_\nu \\
\text{FPToSInt}_{\nu,\varepsilon,\sigma}(rm, f) &= \begin{cases} b & \text{sb}_\nu(b) = \text{rnd}(\text{in}, rm, 0, v(f)) \\ \text{NaN} & \text{otherwise} \end{cases} \\
\text{FPToUInt}_{\nu,\varepsilon,\sigma} &: \mathbb{R}M \times \mathbb{F}_{\varepsilon,\sigma} \rightarrow \mathbb{BV}_\nu \\
\text{FPToUInt}_{\nu,\varepsilon,\sigma}(rm, f) &= \begin{cases} b & \text{ub}_\nu(b) = \text{rnd}(\text{in}, rm, 0, v(f)) \\ \text{NaN} & \text{otherwise} \end{cases}
\end{aligned}$$

V. FROM MODELS TO THEORY

In this section we formalize a logical theory of floating-point numbers based on the structures defined in the previous section. We use the version of many-sorted logic adopted by the SMT-LIB 2 standard [4] whose main differences with

traditional many-sorted logic is that (i) it allows sorts to be denoted by terms instead of just constants and (ii) it allows sort, function and predicate symbols to be indexed by one or more natural number indices. For example, possible sorts may be not just constants like Int but also terms like $\text{Array}(\text{Int}, \text{Real})$ or BitVec_n for all $n > 0$.

A. Preliminaries

A (*logical*) *signature* is composed by a set of sort symbols (of arity ≥ 0) and a set of function symbols f with an associated *rank*, a tuple (S_1, \dots, S_n, S) of sort terms specifying the sort of f 's arguments, namely, S_1, \dots, S_n , and result, S . Constants are represented by nullary function symbols; predicate symbols by function symbols whose return sort is a distinguished sort Bool . Every signature Σ is assumed to contain Bool and constants true and false of that sort, as well as an overloaded symbol $=$ of rank (S, S, Bool) for each sort S , for the identity relation over S . Given a set of sorted variables for each of the sorts in Σ , well-sorted terms and well-sorted formulas of signature Σ are defined as usual.

For every signature Σ , a Σ -*interpretation* \mathcal{I} interprets each sort S in Σ as a non-empty set $\llbracket S \rrbracket_{\mathcal{I}}$, each variable x of sort S as an element $\llbracket x \rrbracket_{\mathcal{I}}$ of $\llbracket S \rrbracket_{\mathcal{I}}$, and each function symbol f of rank (S_1, \dots, S_n, S) as an element $\llbracket f \rrbracket_{\mathcal{I}}$ of the (total) function space $\llbracket S_1 \rrbracket_{\mathcal{I}} \times \dots \times \llbracket S_n \rrbracket_{\mathcal{I}} \rightarrow \llbracket S \rrbracket_{\mathcal{I}}$. Additionally, \mathcal{I} interprets Bool as $\mathbb{B} = \{\top, \perp\}$ and each $=$ of rank (S, S, Bool) as the function that maps $(x, y) \in S \times S$ to \top iff x is y . For each sort S , \mathcal{I} induces a mapping $\llbracket _ \rrbracket_{\mathcal{I}}$ from terms of sort S to $\llbracket S \rrbracket_{\mathcal{I}}$ as expected. A satisfaction relation \models between Σ -interpretations and Σ -formulas is also defined as expected.

In line with recent SMT literature, we define a *theory* of signature Σ as a pair $T = (\Sigma, \mathbf{I})$ where \mathbf{I} is a set of Σ -interpretations, the *models of T*, that is closed under variable reassignment.⁹ We say that a Σ -formula φ is *satisfiable* (resp., *unsatisfiable*) in T if $\mathcal{I} \models \varphi$ for some (resp., no) $\mathcal{I} \in \mathbf{I}$.

B. A Theory of Floating-Point Numbers

In the following we define a theory T_{FP} of floating-point numbers in the sense above by specifying a signature Σ_{FP} and a set of \mathbf{I}_{FP} of Σ_{FP} -interpretations.

The sorts of Σ_{FP} consist, besides Bool , of two individual sorts: RoundingMode and Real ; and two sort families: BitVec_ν , indexed by an integer $\nu > 0$, and $\text{FloatingPoint}_{\varepsilon,\sigma}$, indexed by two integers $\varepsilon, \sigma > 1$. The set of function symbols of Σ_{FP} , and their ranks, is given in Table VIII through XI.¹⁰

We define the set \mathbf{I}_{FP} as the set of *all possible* Σ_{FP} -interpretations \mathcal{I} that interpret sort and function symbol as specified in Table VII through XI in terms of the sets and functions introduced in Section IV. Many of the function symbols are *overloaded* for having different ranks; so we specify their interpretation separately for each rank.

⁹That is, every Σ -interpretation that differs from one in \mathbf{I} only in how it interprets the variables is also in \mathbf{I} .

¹⁰In those tables, we abbreviate RoundingMode , FloatingPoint , and BitVec respectively as RM , FP , and BV .

TABLE VII
SORTS AND THEIR INTERPRETATION REQUIREMENTS

$$\begin{aligned} \llbracket \text{RoundingMode} \rrbracket_{\mathcal{I}} &= \text{RM} & \llbracket \text{FloatingPoint}_{\varepsilon, \sigma} \rrbracket_{\mathcal{I}} &= \mathbb{F}_{\varepsilon, \sigma} \\ \llbracket \text{Real} \rrbracket_{\mathcal{I}} &= \mathbb{R} & \llbracket \text{BitVec}_{\nu} \rrbracket_{\mathcal{I}} &= \mathbb{B}\mathbb{V}_{\nu} \end{aligned}$$

TABLE VIII
CONSTRUCTOR SYMBOLS AND THEIR INTERPRETATION

Symbols of rank RM:

$$\begin{aligned} \llbracket \text{rne} \rrbracket_{\mathcal{I}} &= \text{rne} & \llbracket \text{rna} \rrbracket_{\mathcal{I}} &= \text{rna} & \llbracket \text{rtp} \rrbracket_{\mathcal{I}} &= \text{rtp} \\ \llbracket \text{rtn} \rrbracket_{\mathcal{I}} &= \text{rtn} & \llbracket \text{rtz} \rrbracket_{\mathcal{I}} &= \text{rtz} \end{aligned}$$

Symbols of rank $\text{FP}_{\varepsilon, \sigma}$:

$$\begin{aligned} \llbracket +\text{oo}_{\varepsilon, \sigma} \rrbracket_{\mathcal{I}} &= (0, \mathbf{1}_{\varepsilon}, \mathbf{0}_{\sigma-1}) & \llbracket +\text{zero}_{\varepsilon, \sigma} \rrbracket_{\mathcal{I}} &= (0, \mathbf{0}_{\varepsilon}, \mathbf{0}_{\sigma-1}) \\ \llbracket -\text{oo}_{\varepsilon, \sigma} \rrbracket_{\mathcal{I}} &= (1, \mathbf{1}_{\varepsilon}, \mathbf{0}_{\sigma-1}) & \llbracket -\text{zero}_{\varepsilon, \sigma} \rrbracket_{\mathcal{I}} &= (1, \mathbf{0}_{\varepsilon}, \mathbf{0}_{\sigma-1}) \\ \llbracket \text{NaN}_{\varepsilon, \sigma} \rrbracket_{\mathcal{I}} &= \text{NaN} \end{aligned}$$

Symbols of rank $(\text{BV}_1, \text{BV}_{\varepsilon}, \text{BV}_{\sigma}, \text{FP}_{\varepsilon, \sigma})$:

$$\llbracket \text{fp} \rrbracket_{\mathcal{I}} = \lambda(b_1, b_{\varepsilon}, b_{\sigma-1}). \text{bitpatternToFP}_{\varepsilon, \sigma}(b_1 \bullet b_{\varepsilon} \bullet b_{\sigma-1})$$

As shown in Table VIII, the theory has a family of symbols denoting the floating-point infinities, zeros, and NaN for each pair of exponent and significand length. It also has a ternary function symbol fp that constructs a floating point number from a triple of bit vectors respectively storing the sign, exponent and significant. This allows us to represent all non-NaN values with bit-level precision.

Table IX lists function symbols for the various arithmetic operations over floating-point numbers, and provides their semantics in terms of the operations defined in Subsection IV-C. Table X instead lists predicate symbols corresponding to the relations defined in Subsection IV-B and to the various subsets of $\mathbb{F}_{\varepsilon, \sigma}$. For simplicity and by a slight abuse of notation, in that table we identify functions h of type $D_1 \times \dots \times D_n \rightarrow \mathbb{B}$

TABLE IX
ARITHMETIC SYMBOLS AND THEIR INTERPRETATION

Symbols of rank $(\text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma})$:

$$\llbracket \text{fp.abs} \rrbracket_{\mathcal{I}} = \text{abs}_{\varepsilon, \sigma} \quad \llbracket \text{fp.neg} \rrbracket_{\mathcal{I}} = \text{neg}_{\varepsilon, \sigma}$$

Symbols of rank $(\text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma})$:

$$\begin{aligned} \llbracket \text{fp.max} \rrbracket_{\mathcal{I}} &= \text{max}_{\varepsilon, \sigma} & \llbracket \text{fp.min} \rrbracket_{\mathcal{I}} &= \text{min}_{\varepsilon, \sigma} \\ \llbracket \text{fp.rem} \rrbracket_{\mathcal{I}} &= \text{remrne}_{\varepsilon, \sigma} \end{aligned}$$

Symbols of rank $(\text{RM}, \text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma})$:

$$\llbracket \text{fp.sqrt} \rrbracket_{\mathcal{I}} = \text{sqrt}_{\varepsilon, \sigma} \quad \llbracket \text{fp.roundToIntegral} \rrbracket_{\mathcal{I}} = \text{rti}_{\varepsilon, \sigma}$$

Symbols of rank $(\text{RM}, \text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma})$:

$$\begin{aligned} \llbracket \text{fp.add} \rrbracket_{\mathcal{I}} &= \text{add}_{\varepsilon, \sigma} & \llbracket \text{fp.sub} \rrbracket_{\mathcal{I}} &= \text{sub}_{\varepsilon, \sigma} \\ \llbracket \text{fp.mul} \rrbracket_{\mathcal{I}} &= \text{mul}_{\varepsilon, \sigma} & \llbracket \text{fp.div} \rrbracket_{\mathcal{I}} &= \text{div}_{\varepsilon, \sigma} \end{aligned}$$

Symbols of rank $(\text{RM}, \text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma})$:

$$\llbracket \text{fp.fma} \rrbracket_{\mathcal{I}} = \text{fma}_{\varepsilon, \sigma}$$

TABLE X
PREDICATE SYMBOLS AND THEIR INTERPRETATION

Symbols of rank $(\text{FP}_{\varepsilon, \sigma}, \text{Bool})$:

$$\begin{aligned} \llbracket \text{fp.isNormal} \rrbracket_{\mathcal{I}} &= \text{FN}_{\varepsilon, \sigma} & \llbracket \text{fp.isNegative} \rrbracket_{\mathcal{I}} &= \text{isNeg}_{\varepsilon, \sigma} \\ \llbracket \text{fp.isSubnormal} \rrbracket_{\mathcal{I}} &= \text{FS}_{\varepsilon, \sigma} & \llbracket \text{fp.isPositive} \rrbracket_{\mathcal{I}} &= \text{isPos}_{\varepsilon, \sigma} \\ \llbracket \text{fp.isInfinite} \rrbracket_{\mathcal{I}} &= \text{FI}_{\varepsilon, \sigma} & \llbracket \text{fp.isZero} \rrbracket_{\mathcal{I}} &= \text{FZ}_{\varepsilon, \sigma} \\ \llbracket \text{fp.isNaN} \rrbracket_{\mathcal{I}} &= \{\text{NaN}\} \end{aligned}$$

Symbols of rank $(\text{FP}_{\varepsilon, \sigma}, \text{FP}_{\varepsilon, \sigma}, \text{Bool})$:

$$\begin{aligned} \llbracket \text{fp.lt} \rrbracket_{\mathcal{I}} &= \text{lt}_{\varepsilon, \sigma} & \llbracket \text{fp.leq} \rrbracket_{\mathcal{I}} &= \text{leq}_{\varepsilon, \sigma} \\ \llbracket \text{fp.gt} \rrbracket_{\mathcal{I}} &= \text{gt}_{\varepsilon, \sigma} & \llbracket \text{fp.geq} \rrbracket_{\mathcal{I}} &= \text{geq}_{\varepsilon, \sigma} \\ \llbracket \text{fp.eq} \rrbracket_{\mathcal{I}} &= \text{eq}_{\varepsilon, \sigma} \end{aligned}$$

where $\text{gt}_{\varepsilon, \sigma}$ is the converse of $\text{lt}_{\varepsilon, \sigma}$.

TABLE XI
CONVERSION SYMBOLS AND THEIR INTERPRETATION

Conversions to floating-point

$$\begin{aligned} \llbracket \text{to_fp}_{\varepsilon, \sigma} : (\text{RM}, \text{FP}_{\varepsilon', \sigma'}, \text{FP}_{\varepsilon, \sigma}) \rrbracket_{\mathcal{I}} &= \text{cast}_{\varepsilon', \sigma', \varepsilon, \sigma} \\ \llbracket \text{to_fp}_{\varepsilon, \sigma} : (\text{BV}_{\varepsilon+\sigma}, \text{FP}_{\varepsilon, \sigma}) \rrbracket_{\mathcal{I}} &= \text{bitpatternToFP}_{\varepsilon, \sigma} \\ \llbracket \text{to_fp}_{\varepsilon, \sigma} : (\text{RM}, \text{Real}, \text{FP}_{\varepsilon, \sigma}) \rrbracket_{\mathcal{I}} &= \text{realToFP}_{\varepsilon, \sigma} \\ \llbracket \text{to_fp}_{\varepsilon, \sigma} : (\text{RM}, \text{BV}_{\nu}, \text{FP}_{\varepsilon, \sigma}) \rrbracket_{\mathcal{I}} &= \text{sIntToFP}_{\nu, \varepsilon, \sigma} \\ \llbracket \text{to_fp_unsigned}_{\varepsilon, \sigma} : (\text{RM}, \text{BV}_{\nu}, \text{FP}_{\varepsilon, \sigma}) \rrbracket_{\mathcal{I}} &= \text{uIntToFP}_{\nu, \varepsilon, \sigma} \end{aligned}$$

Conversions from floating-point

$$\begin{aligned} \llbracket \text{fp.to_sbv}_{\nu} : (\text{FP}_{\varepsilon, \sigma}, \text{BV}_{\nu}) \rrbracket_{\mathcal{I}} &= \text{FPToSInt}_{\nu, \varepsilon, \sigma} \\ \llbracket \text{fp.to_ubv}_{\nu} : (\text{FP}_{\varepsilon, \sigma}, \text{BV}_{\nu}) \rrbracket_{\mathcal{I}} &= \text{FPToUInt}_{\nu, \varepsilon, \sigma} \\ \llbracket \text{fp.to_real} : (\text{FP}_{\varepsilon, \sigma}, \text{Real}) \rrbracket_{\mathcal{I}} &= \text{FPToReal}_{\nu, \varepsilon, \sigma} \end{aligned}$$

with the n -ary relations $\{(x_1, \dots, x_n) \in D_1 \times \dots \times D_n \mid h(x_1, \dots, x_n) = \top\}$.

Finally, Table XI lists function symbols corresponding to the various conversion functions introduced in Subsection IV-D as well as the casting function between floating-point sets of different precision.

VI. RELATED WORK

The earliest formalisations of floating-point [8] were limited by the diversity of floating-point formats and systems they had to cover. For example, the formalisation needed to support a range of bases, as 2, 16 and 10 were all in use. Overflow and underflow were particularly problematic as, again, systems in common usage took very different approaches to handling them. One interesting note is that the motivation behind this early work, to support writing portable numerical software, was also one of the drivers of the development of IEEE-754.

The first formalisation of IEEE-754 [5] is notable on several grounds. It was the first to use a formal language (\mathcal{Z}) and was the first used to verify algorithms for the basic operations. The verification was manual, using Hoare logic, and the algorithms were then used in the firmware of the T800 Transputer. During the formalisation, a few issues in IEEE-754 were found and the verification uncovered bugs that would have been difficult to find with testing [25]. However, foreshadowing the issue in

the Pentium 1, bugs were found in the Transputer’s handling of floating-point, introduced by the translation to machine code and manual “tidying up” [14], suggesting a need to extend the proof chain to the whole development process and a need for greater automation.

The FDIV bug in the Pentium 1 and the cost of the resultant recall spurred the use of machine-checked formal proof in the design of floating-point hardware [17], [18]. To this end, IEEE-754 was formalised in a variety of interactive theorem provers, including Isabelle [26], HOL [9], HOL Light [16] (used by Intel), ACL2 [22] (used by AMD and Centaur), PVS [23] and Coq [12], [20]. These and related approaches [2] share a number of common characteristics due to the tools used. They are axiomatic and proof theoretic, generally reducing floating-point numbers to integers and reals; are intended for use in machine checked proofs; and are generally used to verify implementations of floating-point and specific algorithms based on them. In contrast, this work (and its precursor [24]) is model theoretic; builds on computationally simple primitives; and is intended to be a formal reference for automatic theorem provers providing built-in support for reasoning about programs that use floating-point arithmetic.

A number of SMT solvers have support for early version of this theory. The most popular reasoning approach currently is based on encoding floating-point expressions as bit-vector expressions based on the circuits used to implement floating-point operations. To improve performance, over and under approximation schemes are often used. To our knowledge, the earliest implementation of this approach was given in the CBMC model checker [7]. The approach is now used in Z3 [13], MathSAT [10], SONOLAR [19] and CVC4 [3], and improving it remains an active area of research [27]. An alternative approach is based on abstraction interpretation [15], [6] and uses intervals (or other abstract domains) to over-approximate possible models, and a system of branching and learning similar to the SAT algorithm CDCL to narrow these to particular concrete models. There has also been work to integrate axiomatic approaches into SMT solvers [11], although these are not known to implement the semantics presented in this paper.

VII. DISCUSSION

This work originated in the context of the SMT-LIB initiative¹¹ with the goal of defining a standard theory of floating-point numbers for SMT solvers. The standardisation process was guided by three, somewhat conflicting, main principles:

- 1) minimise the amount of effort required to implement a solver for such a theory;
- 2) support a wide range of applications (including avoiding hardware or language specific idioms);
- 3) conform strictly with IEEE-754 in the sense that any compliant implementation of IEEE-754 would form a model of the theory and vice versa.

¹¹See www.smt-lib.org .

A. Process of Standardisation

An initial proposal for a floating-point theory for SMT was made by P. Rümmer and T. Wahl in 2010 [24]. A second draft based on that proposal was produced by C. Tinelli using initial community feedback. This draft was used by a number of system implementors to produce benchmarks and theory solvers. To draw on the experience of using the theory, a work group was formed to decide on a first official theory definition. The suggestions made and changes requested by the work group were written up by C. Tinelli and M. Brain and released to the SMT-LIB community mailing list. We used feedback and corrections from this final review to prepare the version of the theory presented here. The semantics we developed for this theory, discussed in Section III, was inspired by the initial work in [24]. However, it differs substantially from that work by providing a bit-precise semantics of IEEE-754 floating-point.

B. Limitations and Omissions

The theory presented here does not cover all of the functionality of IEEE-754. Here is a list of salient omissions and restrictions and their rationale.

- Two radices, 2 and 10, are included in IEEE-754, although compliant implementations are not required to support both. Our theory only covers radix 2 (binary) representations, as radix 10 (decimal) representations are not widely implemented, supported by programming languages, or used.
- IEEE-754 gives four “specification levels” (extended real, floating-point data, representations and bit-strings) with maps between them. The formalisation in this paper is based on the second level of specification. Thus there is only one NaN, as quiet and signalling NaN are introduced at level three. Also, NaN is not regarded to have a sign or a payload as these concepts are introduced at level four. As a consequence, functionality that uses these properties of NaN, such as *copy sign*, *sign of* and the *total order* predicate, has been omitted.
- Our formalization does not cover the IEEE-754 notion of exceptions or flags (the default handling of executions). This is simply because there is no notion of execution order in a logical formula, and so there is no meaningful way of expressing those notions directly in theory.
- The 2008 revision of IEEE-754 adds the notion of attributes, a means of attaching implicit parameters to sections of programs. Only the most commonly supported attribute, rounding direction, is modeled in our theory—by adding an explicit rounding mode parameter to operations affected by it. Some of the other attributes could be effectively modeled but we chose not to do so because they are programming language specific.
- We do not support unordered variants of comparison operators (those that return true if one or more of the arguments are NaN) since they are rarely used. It would, however, be relatively simple to model these variants.

- IEEE-754 contains recommendations for trigonometric functions and exponentials but neither are mandated. The accuracy of implementations of these functions vary significantly, making it very hard to come up with logical models that are widely applicable but also meaningfully constrained. Similarly, there are only recommendations for reduction operations, and implementations vary significantly.

VIII. CONCLUSION

This paper presents a formalization of IEEE-754 binary floating-point arithmetic as a logical theory, with the goal of enabling automated bit-precise reasoning about programs using floating-point data. The theory is defined as a set of function symbols that model several floating-point operations, and a set of mathematical structures that act as the intended models. This work represents the culmination of a community-driven process to define a standard theory of floating-point arithmetic for SMT, and is intended to provide a common formal reference for both developers and users of SMT solvers, especially in the context of verification applications.

Acknowledgments: We would like to thank F. Bobot, D. Cok, A. Griggio, F. Lapschies, L. de Moura, G. Paganelli, C. Roux, and C. Wintersteiger for their input, feedback and active involvement in the work group.

REFERENCES

- [1] I. S. Association. IEEE standard for floating-point arithmetic, 2008. <http://grouper.ieee.org/groups/754/>.
- [2] A. Ayad and C. Marché. Multi-prover verification of floating-point programs. In *Proceedings of the 5th International Conference on Automated Reasoning, IJCAR'10*, pages 127–141, Berlin, Heidelberg, 2010. Springer-Verlag.
- [3] C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanović, T. King, A. Reynolds, and C. Tinelli. Cvc4. In *Proceedings of the 23rd International Conference on Computer Aided Verification, CAV'11*, pages 171–177, Berlin, Heidelberg, 2011. Springer-Verlag.
- [4] C. Barrett, A. Stump, and C. Tinelli. The SMT-LIB Standard: Version 2.0. Technical report, Department of Computer Science, The University of Iowa, 2010.
- [5] G. Barrett. Formal methods applied to a floating-point number system. *IEEE Trans. Softw. Eng.*, 15(5):611–621, May 1989.
- [6] M. Brain, V. D'Silva, A. Griggio, L. Haller, and D. Kroening. Interpolation-based verification of floating-point programs with abstract cdbl. volume 7935, pages 412–432, June 2013.
- [7] A. Brillout, D. Kroening, and T. Wahl. Mixed abstractions for floating-point arithmetic. In *Formal Methods in Computer-Aided Design (FMCAD)*, 2009.
- [8] W. S. Brown. A simple but realistic model of floating-point computation. *ACM Trans. Math. Softw.*, 7(4):445–480, Dec. 1981.
- [9] V. A. Carreño. Interpretation of IEEE-854 floating-point standard and definition in the HOL system. Technical report, NASA Langley Research Center, 1995.
- [10] A. Cimatti, A. Griggio, B. Schaafsma, and R. Sebastiani. The MathSAT5 SMT Solver. In N. Piterman and S. Smolka, editors, *Proceedings of TACAS*, volume 7795 of *LNCS*. Springer, 2013.
- [11] S. Conchon, G. Melquiond, C. Roux, and M. Iguemelala. Built-in treatment of an axiomatic floating-point theory for SMT solvers. In P. Fontaine and A. Goel, editors, *SMT 2012*, volume 20 of *EPiC Series*, pages 12–21. EasyChair, 2013.
- [12] M. Dumas, L. Rideau, and L. Théry. A generic library for floating-point numbers and its application to exact computing. In R. J. Boulton and P. B. Jackson, editors, *Theorem Proving in Higher Order Logics*, volume 2152 of *LNCS*, pages 169–184. Springer, 2001.
- [13] L. de Moura and N. Bjørner. Z3: An efficient SMT solver. In C. Ramakrishnan and J. Rehof, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer Berlin Heidelberg, 2008.
- [14] J. Gibbons. Formal methods: Why should I care? - the development of the T800 transputer floating-point unit. In *In Proc. 13th New Zealand Computer Society Conference*, pages 207–217, 1993.
- [15] L. Haller, A. Griggio, M. Brain, and D. Kroening. Deciding floating-point logic with systematic abstraction. In G. Cabodi and S. Singh, editors, *Proceedings of the 12th Conference on Formal Methods in Computer-Aided Design (FMCAD 2012)*, pages 131–140, October 2012.
- [16] J. Harrison. Floating point verification in HOL light: The exponential function. In M. Johnson, editor, *Algebraic Methodology and Software Technology*, volume 1349 of *LNCS*, pages 246–260. Springer, 1997.
- [17] J. Harrison. Floating-point verification using theorem proving. In M. Bernardo and A. Cimatti, editors, *SFM*, volume 3965 of *Lecture Notes in Computer Science*, pages 211–242. Springer, 2006.
- [18] J. Harrison. Floating-point verification. *J. UCS*, 13(5):629–638, 2007.
- [19] E. V. Jan Peleska and F. Lapschies. Automated test case generation with SMT-solving and abstract interpretation. volume 6617 of *LNCS*, pages 298–312, SpringerAddress, April 2011. Springer.
- [20] G. Melquiond. Floating-point arithmetic in the Coq system. *Information and Computation*, 216:14–23, July 2012.
- [21] D. Monniaux. The pitfalls of verifying floating-point computations. *ACM Trans. Program. Lang. Syst.*, 30(3), 2008.
- [22] J. Moore, T. Lynch, and M. Kaufmann. A mechanically checked proof of the AMD5K86TM floating-point division program. *IEEE Transactions on Computers*, 47(9):913–926, Sep 1998.
- [23] M. Paul S. Defining the IEEE-854 floating-point standard in PVS. Technical report, NASA Langley Research Center, 1995.
- [24] P. Ruemmer and T. Wahl. An SMT-LIB theory of binary floating-point arithmetic. In *Satisfiability Modulo Theories (SMT)*, 2010.
- [25] J. Woodcock, P. G. Larsen, J. Bicarregui, and J. Fitzgerald. Formal methods: Practice and experience. *ACM Comput. Surv.*, 41(4):19:1–19:36, Oct. 2009.
- [26] L. Yu. A formal model of IEEE floating point arithmetic, July 2013. http://afp.sf.net/entries/IEEE_Floating_Point.shtml.
- [27] A. Zeljic, C. M. Wintersteiger, and P. Rümmer. Approximations for model construction. In S. Demri, D. Kapur, and C. Weidenbach, editors, *Automated Reasoning - 7th International Joint Conference, IJCAR 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 19-22, 2014. Proceedings*, volume 8562 of *Lecture Notes in Computer Science*, pages 344–359. Springer, 2014.