

Toward a Standard Benchmark Format and Suite for Floating-Point Analysis

Nasrine Damouche¹(✉), Matthieu Martel¹, Pavel Panchekha², Chen Qiu², Alexander Sanchez-Stern², and Zachary Tatlock²

¹ Université de Perpignan Via Domitia, Perpignan, France
`nasrine.damouche@univ-perp.fr`

² University of Washington, Seattle, USA

Abstract. We introduce FPBench, a standard benchmark format for validation and optimization of numerical accuracy in floating-point computations. FPBench is a first step toward addressing an increasing need in our community for comparisons and combinations of tools from different application domains. To this end, FPBench provides a basic floating-point benchmark format and accuracy measures for comparing different tools. The FPBench format and measures allow comparing and composing different floating-point tools. We describe the FPBench format and measures and show that FPBench expresses benchmarks from recent papers in the literature, by building an initial benchmark suite drawn from these papers. We intend for FPBench to grow into a standard benchmark suite for the members of the floating-point tools research community.

1 Introduction

The increasingly urgent demand for reliable software has led to tremendous advances in automatic program analysis and verification [4–6, 8, 18]. However, these techniques have typically focused on integer programs, and do not apply to the floating-point computations we depend on for safety-critical control in avionics or medical devices, nor the analyses carried out by scientific and computer-aided design applications. In these contexts, floating-point accuracy is critical since subtle rounding errors can lead to significant discrepancies between floating-point results and the real results developers expect. Indeed, floating-point arithmetic is notoriously unintuitive and its sensitivity to roundoff errors makes such computations fiendishly difficult to debug. Traditionally, such errors have been addressed by numerical methods experts who manually analyze and rewrite floating-point code to ensure accuracy and stability. However, these manual techniques are difficult to apply and typically do not lead to independently checkable certificates guaranteeing program accuracy.

The research community has responded to these challenges by developing a rich array of automated techniques that provide guaranteed bounds on the accuracy of floating-point computations or attempt to automatically improve

accuracy. For example, Fluctuat [12, 13], used in many companies, performs static analysis of C programs to prove a bound on the rounding errors introduced by the use of floating-point numbers instead of reals. Fluctuat also helps users debug floating-point errors by detecting the operations responsible for significant precision loss. Salsa [10] automatically improves the numerical accuracy of programs by using an abstract interpretation to guide transformations that minimize the errors arising during computations. Herbie [21] uses a heuristic search to improve the numerical accuracy of an arithmetic expression by estimating and localizing the roundoff errors of an expression using sampled points, applying a set of rules in order to improve the accuracy of the expression and combining these improvements for different input domains. Rosa [11] combines an exact SMT solver on reals with sound affine arithmetic to verify accuracy post-conditions from assertions about the accuracy of inputs. Rosa can guarantee that the desired precision can be soundly obtained in a finite-precision implementation when propagation error is included. Finally, FPTaylor [24] improves on interval arithmetic by using Taylor series to narrow the computed error bounds.

As the number of tools dedicated to analyzing and improving numerical accuracy grows, it becomes increasingly difficult to make fair comparisons between the techniques. This is because each tool is targeted to slightly different domains, uses slightly different formats for expressing benchmarks, and reports results using related but slightly different measures. Furthermore, without any standard set of floating-point benchmarks, it is difficult to identify opportunities for composing complementary tools.

To address these challenges, the floating-point research community needs a standard benchmark format and common set of measures that enables comparison and cooperation between tools. This goal is motivated by the success of standard benchmark suites like SPEC [17] and SPLASH-2 [25] in the compiler community, the DIMACS [1] format in the SAT-solving community, and the SMT-LIB [3] format in the SMT-solving community. The formats have enabled fair comparisons between tools, crisp characterizations of the tradeoffs between different approaches, and useful cooperation between tools with complementary strengths.

In this article, we propose FPBench, a general floating-point format and benchmark suite. FPBench describes a common format and a suite of accuracy measurements; we envision floating-point tools taking in FPBench formatted programs and using the FPBench accuracy measures to describe accuracy. This allows users to combine tools that perform complementary tasks and compare competing tools to choose the one best for their task. The common scientific methodology FPBench enables is crucial for demonstrating the improvements of each tool on the state of the art.

The main contributions of this article are the following:

- (i) A uniform format for expressing floating-point benchmarks, FPCore.
- (ii) A set of utilities for converting to and from FPCore programs, and working with FPCore programs.

- (iii) A set of measures on which to evaluate various floating-point tools on FPBench benchmarks.
- (iv) An initial suite of benchmarks drawing from existing floating-point literature.

The remainder of this article is organized as follows. Section 2, describes the FPBench formats. Section 3 describes the accuracy measures. Section 4 describes the utilities FPBench provides to support creating and working with benchmarks. Section 5 surveys our existing benchmark suite, highlighting representative case studies from recent tools in the literature. Finally, Sect. 6 discusses future work and concludes.

2 Benchmark Format

FPBench provides a common input format for floating-point tools. This format makes it possible to develop a collection of floating-point benchmarks common to our research community, and also allows users to compose floating-point tools. So, a common floating-point benchmark format must be easy to parse, have simple and clear semantics, support floating-point details, and be expressive enough for many domains. To satisfy these requirements, FPBench provides the FPCore format, a minimal expression-based language for floating-point computations. A common floating-point benchmark format must also be easy to translate to and from popular industrial languages like C, C++, Matlab, and Fortran. To satisfy these requirements, FPBench also provides the extended FPImp format, a basic imperative language for floating-point computations which can be automatically compiled to FPCore.

2.1 FPCore

FPCore is an S-expression format featuring mathematical functions, `if` statements, and `while` loops. All floating-point functions from C11's `math.h` and all Fortran 2003 intrinsics are supported operators, as well as standard arithmetic operators like addition and comparison; likewise, all constants defined in C11's `math.h` are available as constants. Following IEEE754 and common C and Fortran runtimes, FPCore does not prescribe the accuracy of built-in mathematical functions. However, individual benchmarks can declare the accuracy they assume for built-in operations which analyzers can take into account.

A FPCore benchmark specifies a set of inputs, a floating-point expression, and optional meta-data flags such as a name, citations, preconditions on inputs, and the floating-point precision (`binary32`, `binary64`, ...) used to evaluate that benchmark. The full FPCore syntax is as follows:¹

¹ Not shown in the grammar: FPBench uses “;” to indicate that the remainder of a line is a comment.

$$\begin{aligned}
\langle \text{FPCore} \rangle &::= (\langle \text{FPCore} \rangle (\langle \text{identifier} \rangle^*) \langle \text{property} \rangle^* \langle \text{expression} \rangle) \\
\langle \text{expression} \rangle &::= \langle \text{constant} \rangle \mid \langle \text{identifier} \rangle \mid (\langle \text{operation} \rangle \langle \text{expression} \rangle^*) \\
&\mid (\langle \text{if} \rangle \langle \text{expression} \rangle \langle \text{expression} \rangle \langle \text{expression} \rangle) \\
&\mid (\langle \text{while} \rangle \langle \text{expression} \rangle (\langle \text{loopvar} \rangle^*) \langle \text{expression} \rangle) \\
&\mid (\langle \text{let} \rangle (\langle \text{binding} \rangle^*) \langle \text{expression} \rangle) \\
\langle \text{loopvar} \rangle &::= [\langle \text{identifier} \rangle \langle \text{expression} \rangle \langle \text{expression} \rangle] \\
\langle \text{constant} \rangle &::= \langle \text{number} \rangle \mid \langle \text{E} \rangle \mid \langle \text{LOG2E} \rangle \mid \langle \text{LOG10E} \rangle \mid \langle \text{LN2} \rangle \mid \langle \text{LN10} \rangle \mid \langle \text{PI} \rangle \mid \dots \\
\langle \text{operation} \rangle &::= \langle \text{+} \rangle \mid \langle \text{*} \rangle \mid \langle \text{<} \rangle \mid \dots \mid \langle \text{abs} \rangle \mid \langle \text{acos} \rangle \mid \langle \text{and} \rangle \mid \langle \text{asin} \rangle \mid \dots \\
\langle \text{property} \rangle &::= \langle \text{propname} \rangle \langle \text{propval} \rangle \\
\langle \text{propname} \rangle &::= \langle \text{:name} \rangle \mid \langle \text{:cite} \rangle \mid \langle \text{:pre} \rangle \mid \langle \text{:type} \rangle \mid \dots \\
\langle \text{propval} \rangle &::= \langle \text{expression} \rangle \mid \langle \text{string} \rangle \mid (\langle \text{identifier} \rangle^*) \mid \langle \text{identifier} \rangle \\
\langle \text{binding} \rangle &::= [\langle \text{identifier} \rangle \langle \text{expression} \rangle]
\end{aligned}$$

Since the language is S-expression based, parentheses and braces are literals. The semantics of these programs is ordinary function evaluation, with **let** bindings evaluated simultaneously and **while** loops evaluated by simultaneously updating the loop variables until the condition is true, and then evaluating the return value:

$$\frac{H : x_i^0 \Downarrow v_i \quad H[x_i \mapsto v_i]_i : c \Downarrow \top \quad H[x_i \mapsto v_i]_i : e_i \Downarrow x'_i \quad H : (\text{while } c ([x_i \ x'_i \ e_i]_i) y) \Downarrow v}{H : (\text{while } c ([x_i \ x_i^0 \ e_i]_i) y) \Downarrow v}$$

$$\frac{H : x_i^0 \Downarrow v_i \quad H[x_i \mapsto v_i]_i : c \Downarrow \perp \quad H[x_i \mapsto v_i]_i : y \Downarrow v}{H : (\text{while } c ([x_i \ x_i^0 \ e_i]_i) y) \Downarrow v}$$

The optional properties are used to record additional information about each benchmark. Benchmarks are currently annotated with a **:name**, a **:description** of the benchmark and its inputs, the floating-point **:type** (either **binary32** or **binary64**), a precondition **:pre**, and a citation **:cite**. All FPBench tools ignore unknown attributes, so they represent an easy way to record additional benchmark information. We recommend that properties specific to a single tool be prefixed with the name of the tool.

2.2 FPImp

Where FPCore is the format consumed by tools, FPImp is a format for simplifying the translation from imperative languages to FPCore. While FPCore is a functional language with a minimal set of features for representing floating-point computations, FPImp includes common imperative features like variable assignments, multiple return values, and multi-way **if** statements. More complex features, such as arrays, pointers, records, and recursive function calls, are

left to future language extensions. In our experience, translating C, Fortran, and Matlab to FPImp is relatively easy. FPImp has expressions similar to those in FPCore, but without `if`, `while`, and `let` expressions; these are instead replaced with statements. The FPImp syntax is as follows:

$$\begin{aligned} \langle FPImp \rangle &::= (\text{'FPImp'} (\langle identifier \rangle^*) \langle property \rangle^* \langle statement \rangle^* \langle output \rangle) \\ \langle output \rangle &::= (\text{'output'} \langle expression \rangle^*) \\ \langle expression \rangle &::= \langle constant \rangle \mid \langle identifier \rangle \mid (\langle operation \rangle \langle expression \rangle^*) \\ \langle statement \rangle &::= [\text{'='} \langle identifier \rangle \langle expression \rangle] \\ &\mid (\text{'if'} \langle ifbranch \rangle^*) \\ &\mid (\text{'while'} \langle expression \rangle \langle statement \rangle^*) \\ \langle ifbranch \rangle &::= [\langle expression \rangle \langle statement \rangle^*] \\ &\mid [\text{'else'} \langle statement \rangle^*] \end{aligned}$$

As in FPCore, each FPImp benchmark includes free parameters and properties; as an imperative language, it includes a list of program statements and multiple return values instead of a single body expression. Assignments and while loops behave in the usual way. Expressions in FPImp are evaluated similarly to FPCore expressions. An `if` statement defines a many-way branch; any `else` branch must be the last branch in its `if`. The `output` statement can return several values and appears at the end of a function.

The FPImp to FPCore compiler translates FPImp functions to FPCore benchmarks while retaining all properties and keeping the same set of free parameters. It inlines assignments, converts the imperative bodies of FPImp loops to the simultaneous-assignment loops of FPCore, replaces many-way `if` statements with nested `if` s, and generates multiple FPCore benchmarks for FPImp programs with multiple outputs.

3 Accuracy Measurements

To compare floating-point tools, a common input format is not enough: a common measure of accuracy is also necessary. FPBench thus defines a collection of accuracy measures to allow tools to rigorously describe the accuracy measure they use. Given the diversity of accuracy measures in the literature, standardizing on a single accuracy measure would be difficult, and could harm the development of some classes of tools. Instead, FPBench standardizes several measures of accuracy; tools that measure accuracy should state which of the FPBench accuracy measures they use to compare tools.

3.1 Measurement Axes

Floating-point accuracy is best analyzed along several axes: scaling vs non-scaling error, forward vs. backward error, maximum vs. average error. Tools that measure error may use sound vs. statistical techniques, and tools that transform programs have several options for how to measure accuracy improvement.

Scaling vs non-scaling error (ε). There are several ways to measure the error of producing the inaccurate value \hat{x} instead of the true value x . Two common mathematical notions are the absolute and relative error:

$$\varepsilon_{\text{abs}}(x, x') = |x - \hat{x}| \quad \text{and} \quad \varepsilon_{\text{rel}}(x, x') = \left| \frac{x - \hat{x}}{x} \right|$$

Relative error scales with the quantity being measured, and thus makes sense for measuring both large and small numbers, much like the floating-point format itself. A notion of error more closely tied to the floating-point format is the Units in the Last Place (ULPs)² difference. Some tools instead use the logarithm of the ULPs, which approximately describes the number of incorrect low-order bits in \hat{x} . These two measures are defined as:³

$$\varepsilon_{\text{ulps}}(x, x') = |\mathbb{F} \cap [\min(x, \hat{x}), \max(x, \hat{x})]| \quad \text{and} \quad \varepsilon_{\text{bits}}(x, x') = \log_2 \varepsilon_{\text{ulps}}(x, x')$$

The floating-point numbers are distributed roughly exponentially, so this measure of error scales in a similar manner to relative error; however, it is better-behaved in the presence of denormal numbers and infinities.

Forward vs. backward error (ϵ). Forward error and backward error are two common measures for the error of a mathematical *computation*. For a true function f and its approximation \hat{f} , forward error measures the difference between outputs for a fixed input, while backward error measures the difference between inputs for a fixed output. Formally,⁴

$$\epsilon_{\text{fwd}}(f, \hat{f}, x) = \varepsilon(f(x), \hat{f}(x))$$

$$\epsilon_{\text{bwd}}(f, \hat{f}, x) = \min \left\{ \varepsilon(x, x') : x' \in \mathbb{F}^n \wedge \hat{f}(x') = f(x) \right\}.$$

Backward error is important for evaluating the stability of an algorithm, and in scientific applications where multiple sources of error (algorithmic error vs. sensor error) must be compared. Existing tools typically measure forward error because backward error can be tricky to compute for floating-point computations, where there may not be an input that produces the true output.

Average vs. maximum error (E). Describing the error of a floating-point computation means summarizing its behavior across multiple inputs. Existing tools use either maximum or average error for this task. Formally,⁵

$$E_{\text{max}}(f, \hat{f}) = \max \left\{ \epsilon(f, \hat{f}, x) : x \in \mathbb{F}^n \right\} \quad \text{and} \quad E_{\text{avg}}(f, \hat{f}) = \frac{1}{N} \sum_{x \in \mathbb{F}^n} \epsilon(f, \hat{f}, x).$$

² Unfortunately, this term means different things in the mathematical and programming communities. We use the definition common for programming tools [19, 21, 23].

³ We are using $|S|$ to denote the number of elements in a set S .

⁴ Where n is the number of arguments.

⁵ Where N is the number of valid inputs, and n is the number of arguments.

Worst-case error tends to be easier to measure soundly, while average error tends to be easier to measure statistically.

Sound vs. statistical techniques. Running a floating-point program on all valid inputs is intractable. Existing tools either soundly overapproximate the error using static analysis or approximate the error using statistical sampling.

Most static techniques are based on interval or affine arithmetic to overapproximate floating-point arithmetic, often using abstract interpretation. Abstract interpretation may use either non-relational [20] or relational abstract domains [2, 7, 14], and may use acceleration techniques (widening [9]) to overapproximate loops without unrolling them. While such techniques tend to provide loose over-approximations of the floating-point error of programs, they are fast and provide sound error bounds. In some embedded applications, correctness is critical and unsound techniques will not do.

In domains where correctness is not absolutely critical, sampling can provide tight approximations of error. Many sampling techniques are used, including naive random samples [21] and Markov-chain Monte Carlo [23]. These techniques involve running a program multiple times, so tend to be slower than static analysis.

Measuring improvement (ι). Tools that *transform* floating-point programs need to compare the accuracy of two floating-point programs: the original and the transformed. Several comparison measures are possible. Comparisons can use the improvement in worst-case or average error between the original \hat{f} and improved \hat{f}' implementation of the same mathematical function f :

$$\iota_{\text{imp}} = E(f, \hat{f}) - E(f, \hat{f}')$$

However, one cannot usually improve the accuracy of a computation simultaneously on all inputs. It is thus often necessary to make a computation less accurate on some points to make it more accurate overall. In this case, it may be useful to report the largest unimprovement, which measures the cost of improving accuracy:

$$\iota_{\text{wrs}}(\hat{f}, \hat{f}') = \max \left\{ \epsilon(f, \hat{f}', x) - \epsilon(f, \hat{f}, x) : x \in \mathbb{F}^n \right\}$$

Other measures, such as those describing the trade-off between accuracy and speed, are also interesting, but are less commonly used in the literature and thus not standardized in FPBench. Improvement tools could also estimate their effect on numerical stability using automatic differentiation [15] or Lyapunov exponents [22], but we do not know of any such tools.

3.2 Existing Tools

The error measures described can be applied to categorize the error measurements used by existing tools. Table 1 compares Fluctuat [12], FPTaylor [24], Herbie [21], Rosa [11], and Salsa [10].

Table 1. A comparison of how five floating-point measure error across the axes identified in this section.

Fluctuat	Absolute	Forward	Max	Sound	
FPTaylor	Absolute	Forward	Max	Sound	
Herbie	Bits	Forward	Average	Statistical	Improvement
Rosa	Absolute	Forward	Max	Sound	
Salsa	Absolute	Forward	Max	Sound	Improvement
STOKE	ULPs	Forward	Max	Statistical	Improvement

Fluctuat, FPTaylor, and Rosa all verify error bounds on the accuracy of floating-point computations. Given their need for soundness, it is natural that they use sound error analyses and estimate maximum error. Their use of absolute forward error derives from the difficulty of approximating the other forms of error statically. Herbie and Salsa are tools for improving the accuracy of programs, but differ dramatically in their approach. Salsa uses abstract interpretation to bound maximum absolute error, producing a sound overapproximation of the maximum error. Herbie, on the other hand, uses random sampling to achieve a tight statistical approximation of bits error. The tight estimates enabled by statistical techniques provide additional opportunities for Herbie to improve the accuracy of an expression, but prevent it from providing sound error bounds. Finally, STOKE uses stochastic search to optimize floating-point programs, and must compare the accuracy of floating-point programs to avoid significantly compromising their accuracy. STOKE uses a Markov-chain Monte Carlo sampling to statistically evaluate maximum ULPs error.

By exactly describing the way each tool measures accuracy, FPBench makes it possible to compare and relate tools. Unsound tools such as Herbie or STOKE can be composed with a sound verification tool to produce an accuracy guarantee, and this guarantee can be compared to the approximate error measurements those tools made statistically. Since Fluctuat, FPTaylor, Rosa, and Salsa all soundly measure maximum forward absolute error, they can be compared to determine which technique is best.

4 Tools

FPBench features a collection of compilers and measurement tools that operates in its common format, FPCore. These tools can be a community resource, increasing interoperability as well as code reuse. They also make it easier to write new floating-point analysis and transformation tools by automating what are currently common but tedious tasks.

FPImp to FPCore. The FPCore format faithfully preserves important program constructs, such as variable binding and operation ordering, while abstracting away details not relevant to floating-point semantics. However, it is syntactically very different from some of the languages from which benchmarks might originate. To make translation to FPCore from source languages like C, Fortran, and

Matlab easier, FPBench provides the FPImp format and a compiler from FPImp to FPCore. FPImp is syntactically similar to imperative languages to make it easy to translate benchmarks.

FPCore to C. Since C is a common implementation language for mathematical computations, FPBench provides a FPCore to C compiler. The FPCore to C compiler can be used to run FPCore benchmarks through the many available C analysis tools.

Average error estimation. FPBench provides a tool to statistically approximate average error using naive sampling. The statistical approach is necessary to produce accurate estimates of average error given the current state of the art. The tool can use absolute, relative, ULPs, and bits error.

We plan to continue developing community tools around the FPBench formats, especially tools for estimating the other measures of error described in Sect. 3.

5 Benchmark Suite and Examples

The FPBench suite currently includes 44 benchmarks sourced from recent papers on automatic floating-point verification and accuracy improvement. This section first summarizes these benchmarks and then details how representative examples were translated to FPBench from the input formats of various tools in the literature.

The current FPBench suite contains examples from a variety of domains, including 28 from the Herbie test suite [21], 9 from the Salsa test suite [10], 7 from the Rosa test suite [11], and one example from the FPTaylor test suite [24]. These examples range from simple test programs for early tool development up to large examples for evaluating more mature tools. The larger examples are more challenging, including loops that with up to 13 variables mutated in the loop body. As shown in Tables 2 and 3, these programs exercise the full range of functionality available in FPBench, and span a variety application areas, from control software to mathematical libraries. We intend to continue adding benchmarks to the FPBench suite.

5.1 FPTaylor

FPTaylor [24] uses series expansions and interval arithmetic to compute sound error bounds. The authors gave the following simple program as an example input for their tool:

```

1: Variables
2:   float64 x in [1.001, 2.0],
3:   float64 y in [1.001, 2.0];
4: Definitions
5:   t rnd64= x * y;
6: Expressions
7:   r rnd64= (t-1)/(t*t-1);

```

Table 2. Functions and language features used in the FPBench benchmarks. Benchmarks contain a variety of features, and many benchmarks incorporate several. Exponential functions include logarithms, the exponential function, and the power function.

Feature	Benchmarks
Basic arithmetic	44
Exponentials	13
Trigonometric	10
Comparison	12
Loops	12
Conditionals	3

Table 3. Domains which the FPBench benchmarks are drawn from. Most are general mathematical operations, useful in a variety of domains. The general expressions are the smallest, and are drawn from *Numerical Methods for Scientists and Engineers* [16] and Rosa [11].

Domain	Benchmarks
General expressions	31
Math algorithms	6
Embedded systems	4
Scientific computing	3

This program is representative of the code necessary to correct sensor data in control software, where the output of the sensor is known to be between 1.001 and 2.0. We manually translated this program to FPCore, yielding:

```
(FPCore (x y)
  :name" FPTaylor example"
  :cite (solovyev-et-al-2015)
  :type binary64
  :pre (and (and (< 1.001 x) (< x 2.0)) (and (< 1.001 y) (< y 2.0)))
  (let ([t (* x y)])
    (/ (- t 1.0) (- (* t t) 1.0))))
```

The benchmark takes inputs x and y and uses a `let` statement to represent the intermediate variables from the FPTaylor example. FPCore faithfully preserves important program constructs, such as variable binding and operation ordering, making the translation a simple matter. The benchmark additionally specifies a name and cites its source using a key to a standard BIBTEX file. Since the original program uses 64-bit floating-point numbers, the type `binary64` is specified in the benchmark. The constraints on input variables are translated to a single predicate under the `:pre` property.

5.2 Rosa

Rosa [11] soundly verifies error bounds of floating-point programs, including looping control flow through recursive function calls. Several benchmarks in its repository demonstrate this capability, including one that uses Newton’s method on a series representation of the sine function. In Rosa’s input language the original benchmark is represented as:

```

def newton(x: Real, k: LoopCounter): Real = {
  require(-1.0 < x && x < 1.0)
  if (k < 10) {
    newton(x - (x - (x**3)/6.0 + (x**5)/120.0 + (x**7)/5040.0) /
      (1.0 - (x**2)/2.0 + (x**4)/24.0 + (x**6)/720.0), k++)
  } else {
    x
  }
}
ensuring(res => -1.0 < res && res < 1.0)

```

The **require** clause denotes input preconditions, and the **ensures** clause provides the error bound to be verified. We manually translated this program to FPCore, yielding:

```

(FPCore (x0)
  :name "Rosa Example"
  :cite (darulova-kuncak-2014)
  :pre (and (< -1.0 x0) (< x0 1.0))
  :rosa-post (and (< -1.0 RES) (< RES 1.0))
  (while (< i 10)
    ([i 0 (+ i 1)]
     [x x0
      (let ([f (+ (+ (- x (/ (pow x 3) 6))
                    (/ (pow x 5) 120)) (/ (pow x 7) 5040))]
            [df (+ (+ (- 1.0 (/ (* x x) 2))
                    (/ (pow x 4) 24)) (/ (pow x 6) 720))]
              (- x (/ f df)))]])
    x))

```

Like the FPTaylor benchmark, this benchmark includes a name, citation for its source, and precondition on inputs. For completeness, we've also included the **ensuring** annotation, denoted **:rosa-post**, demonstrating the ability to add tool specific annotations by prefixing with the tool name. Note how the **while** construct from FPCore can be used to represent the tail-recursive loop from the Rosa original benchmark.

5.3 Herbie

Herbie [21] heuristically improves the accuracy of straight-line floating-point expressions. The authors demonstrate the improvements Herbie can produce using the quadratic formula for computing the roots of a second degree polynomial. It has uses from calculating trajectories to solving matrix equations. In mathematical notation, the quadratic formula is given by:⁶

$$\frac{(-b) - \sqrt{b^2 - 4ac}}{2a}$$

Herbie produces the following more-accurate variant:

$$\begin{cases} \frac{4ac}{-b + \sqrt{b^2 - 4ac}} / 2a & \text{if } b < 0 \\ (-b - \sqrt{b^2 - 4ac}) \frac{1}{2a} & \text{if } 0 \leq b \leq 10^{127} \\ -\frac{b}{a} + \frac{c}{b} & \text{if } 10^{127} < b \end{cases}$$

⁶ We use the negative variant here, as in the Herbie paper; the positive variant is analogous.

In FPCore format, the original formula is represented as:

```
(FPCore (a b c)
 :name "'NMSE p42, positive"
 :cite (hamming-1987)
 :pre (and (>= (sqr b) (* 4 (* a c))) (!= a 0))
 (/ (+ (- b) (sqrt (- (sqr b) (* 4 (* a c))))) (* 2 a)))
```

and the improved version is represented as:

```
(FPCore (a b c)
 :name "'NMSE p42, positive"
 :cite (hamming-1987)
 :pre (and (>= (sqr b) (* 4 (* a c))) (!= a 0))
 (if (< b 0)
 (/ (/ (* 4 (* a c)) (+ (- b) (sqrt (- (sqr b) (* 4 (* a c)))))
 (* 2 a))
 (if (< b 10e127)
 (* (- (- b) (- (sqr b) (* 4 (* a c)))) (/ 1 (* 2 a)))
 (+ (- (/ b a)) (/ c b)))))
```

Like the Rosa and FPTaylor examples, this benchmark gives a name, citation, and precondition. This benchmark uses FPCore's ability to write arbitrary boolean expressions as preconditions, restricting the value under the square root to be non-negative and requiring the denominator be non-zero. Unlike the FPTaylor and Rosa examples, this precondition arises from mathematical considerations, not domain knowledge. The FPCore version of Herbie's output furthermore uses `if` constructs to evaluate different expressions for different inputs, which improves accuracy. Herbie could add additional metadata to the output, such as its internal estimate of accuracy or the number of expressions considered during search, by using prefixed keys like `:herbie-accuracy-estimate`.

5.4 Salsa

Salsa [10] is a tool for soundly improving the worst-case accuracy of programs. The authors evaluate Salsa on a suite of control and numerical algorithms, including the widely used PID controller algorithm. This algorithm is used in aeronautic and avionic systems for which correctness is critical. In C, the benchmark is written as:

```
volatile double p, i, t, d, dt, invdt, m, e, eold, r;
int pid(double m0, double kp, double ki, double kd, double c){
    t = 0.0;
    invdt = 5.0;
    dt = 0.2;
    m = m0;
    eold = 0.0;
    i = 0.0;
    while (t < 100.0) {
        e = c - m;
        p = kp * e;
        i = i + ki * dt * e;
        d = kd * invdt * (e - eold);
        r = p + i + d;
        m = m + 0.01 * r; /* computing measure: the plant */
        eold = e;
        t = t + dt;
    }
    return m;
}
```

To ease the conversion of this code from C to FPCore, this program was first manually translated to the following FPImp program:

```
(FPImp (m kp ki kd)
:name ‘‘PID’’
:description ‘‘Keep a measure at its setpoint using a PID controller.’’
:cite (damouche-martel-chapoutot-nsv14)
:type binary64
:pre (and (and (< -10.0 m) (< m 10.0)) (and (< -10.0 c) (< c 10.0)))
[= t 0.0]
[= dt 0.2]
[= invdt (/ 1 dt)]
[= c 0.0]
[= eold 0.0]
[= i 0.0]
(while (< t 100.0)
[= e (- c m)]
[= p (* kp e)]
[= i (+ i (* (* ki dt) e))]
[= d (* (* kd invdt) (- e eold))]
[= r (+ (+ p i) d)]
[= m (+ m (* 0.01 r))]
[= eold e]
[= t (+ t dt)])
(output m))
```

The FPImp program was then automatically compiled, using the compiler tool in FPBench, to the following FPCore benchmark:

```
(FPCore (m0 kp ki kd c)
:name ‘‘PID’’
:description ‘‘Keep a measure at its setpoint using a PID controller.’’
:cite (damouche-martel-chapoutot-nsv14)
:type binary64
:pre (and (and (< -10.0 m0) (< m0 10.0)) (and (< -10.0 c) (< c 10.0)))
(while (< t 100.0)
([i 0.0 (+ i (* (* ki 0.2) (- c m)))]
[m m0
(let ([p (* kp (- c m))]
[d (* (* kd (/ 1 0.2)) (- (- c m) eold))])
(+ m (* 0.01 (+ (+ p (+ i (* (* ki 0.2) (- c m)))) d)))]
[eold 0.0 (- c m)]
[t 0.0 (+ t 0.2)])
m))
```

Beyond the metadata used in the FPTaylor, Rosa, and Herbie examples, this benchmark includes a `:description` tag to describe for readers what the benchmark program computes. These descriptions contain information about the distribution of inputs, the situation in which the benchmark is used, or any other information which might be useful to tool writers. The FPBench suite features descriptions for its most complex benchmarks.

6 Conclusion

The initial work on FPBench provides a foundation for evaluating and comparing floating-point analysis and optimization tools. Already the FPBench format can serve as a common language, allowing these tools to cooperatively analyze and improve floating-point code.

Though FPBench supports the composition of the floating-point tools that exist today, there is still much work to do to support the floating-point research

community as it grows. FPBench must be sufficiently expressive for the broad range of applications that represent the future of the community. In the near term, we will add additional metrics for accuracy and performance to the set of evaluators provided by the FPBench tooling and begin developing a standard set of benchmarks around the various measures. We will also expand the set of languages with direct support for compilation to and from the FPBench format. As more tools grow support for FPBench, we will provide automated comparisons of different floating-point tools. Longer term, we intend to support for mixed-precision benchmarks, fixed-point computations, and additional data structures such as vectors, records, and matrices.

We hope that FPBench encourages the already strong sense of community and collaboration around floating-point research. Toward that end, we encourage any interested readers (and tool writers) to get involved with development of FPBench by signing up for the mailing list and checking out the project website: <http://fpbench.org/>.

References

1. DIMACS challenge. Satisfiability. Suggested format (1993). <http://www.cs.ubc.ca/~hoos/SATLIB/Benchmarks/SAT/satformat.ps>
2. Adjé, A., Garoche, P.-L., Werey, A.: Quadratic zonotopes. In: Feng, X., Park, S. (eds.) APLAS 2015. LNCS, vol. 9458, pp. 127–145. Springer, Cham (2015). doi:[10.1007/978-3-319-26529-2_8](https://doi.org/10.1007/978-3-319-26529-2_8)
3. Barrett, C., Fontaine, P., Tinelli, C.: The SMT-LIB Standard: Version 2.5. Technical report, Department of Computer Science, The University of Iowa (2015). www.SMT-LIB.org
4. Bertrane, J., Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Rival, X.: Static analysis and verification of aerospace software by abstract interpretation. *Found. Trends Program. Lang.* **2**(2–3), 71–190 (2015)
5. Bessey, A., Block, K., Chelf, B., Chou, A., Fulton, B., Hallem, S., Henri-Gros, C., Kamsky, A., McPeak, S., Engler, D.: A few billion lines of code later: using static analysis to find bugs in the real world. *Commun. ACM* **53**(2), 66–75 (2010)
6. Cadar, C., Dunbar, D., Engler, D.: Klee: unassisted and automatic generation of high-coverage tests for complex systems programs. In: *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation (OSDI 2008)*, Berkeley, CA, USA, pp. 209–224. USENIX Association (2008)
7. Chapoutot, A.: Interval slopes as a numerical abstract domain for floating-point variables. In: Cousot, R., Martel, M. (eds.) SAS 2010. LNCS, vol. 6337, pp. 184–200. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15769-1_12](https://doi.org/10.1007/978-3-642-15769-1_12)
8. Clarke, E., Kroening, D., Lerda, F.: A tool for checking ANSI-C programs. In: Jensen, K., Podelski, A. (eds.) TACAS 2004. LNCS, vol. 2988, pp. 168–176. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-24730-2_15](https://doi.org/10.1007/978-3-540-24730-2_15)
9. Cousot, P., Cousot, R.: Abstract interpretation frameworks. *J. Log. Comput.* **2**(4), 511–547 (1992)
10. Damouche, N., Martel, M., Chapoutot, A.: Intra-procedural optimization of the numerical accuracy of programs. In: Núñez, M., Güdemann, M. (eds.) FMICS 2015. LNCS, vol. 9128, pp. 31–46. Springer, Cham (2015). doi:[10.1007/978-3-319-19458-5_3](https://doi.org/10.1007/978-3-319-19458-5_3)

11. Darulova, E., Kuncak, V.: Sound compilation of reals. In: Jagannathan, S., Sewell, P., (eds.), POPL 2014, pp. 235–248. ACM (2014)
12. Goubault, E.: Static analysis by abstract interpretation of numerical programs and systems, and FLUCTUAT. In: Logozzo, F., Fähndrich, M. (eds.) SAS 2013. LNCS, vol. 7935, pp. 1–3. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-38856-9_1](https://doi.org/10.1007/978-3-642-38856-9_1)
13. Goubault, E., Martel, M., Putot, S.: Some future challenges in the validation of control systems. In: European Congress on Embedded Real Time Software (ERTS 2006), Proceedings (2006)
14. Goubault, E., Putot, S.: Static analysis of finite precision computations. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 232–247. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-18275-4_17](https://doi.org/10.1007/978-3-642-18275-4_17)
15. Griewank, A., Walther, A.: Evaluating derivatives - principles and techniques of algorithmic differentiation (2 ed.). SIAM (2008)
16. Hamming, R.: Numerical Methods for Scientists and Engineers, 2nd edn. Dover Publications, New York (1987)
17. Henning, J.L.: SPEC CPU2006 benchmark descriptions. SIGARCH Comput. Archit. News **34**(4), 1–17 (2006)
18. Holzmann, G.J.: Mars code. Commun. ACM **57**(2), 64–73 (2014)
19. Lee, W., Sharma, R., Aiken, A.: Verifying bit-manipulations of floating-point. In: Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2016), New York, NY, USA. ACM (2016)
20. Martel, M.: An overview of semantics for the validation of numerical programs. In: Cousot, R. (ed.) VMCAI 2005. LNCS, vol. 3385, pp. 59–77. Springer, Heidelberg (2005). doi:[10.1007/978-3-540-30579-8_4](https://doi.org/10.1007/978-3-540-30579-8_4)
21. Wilcox, J.-R., Panckekha, P., Sanchez-Stern, A., Tatlock, Z.: Automatically improving accuracy for floating point expressions. In: Grove, D., Blackburn, S., (eds.), PLDI 2015, pp. 1–11. ACM (2015)
22. Sankaranarayanan, S.: Static analysis in the continuously changing world. In: Logozzo, F., Fähndrich, M. (eds.) SAS 2013. LNCS, vol. 7935, pp. 4–5. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-38856-9_2](https://doi.org/10.1007/978-3-642-38856-9_2)
23. Schkufza, E., Sharma, R., Aiken, A.: Stochastic optimization of floating-point programs with tunable precision. In: Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2014), New York, NY, USA. ACM (2014)
24. Solovyev, A., Jacobsen, C., Rakamarić, Z., Gopalakrishnan, G.: Rigorous estimation of floating-point round-off errors with symbolic taylor expansions. In: Bjørner, N., de Boer, F. (eds.) FM 2015: Formal Methods. Programming and Software Engineering, vol. 9109. Springer, Heidelberg (2015)
25. Woo, S.C., Ohara, M., Torrie, E., Singh, J.P., Gupta, A.: The SPLASH-2 programs: characterization and methodological considerations. In: Proceedings of the 22nd Annual International Symposium on Computer Architecture (ISCA 1995), pp. 24–36, New York, NY, USA. ACM (1995)