

# **An Interval Arithmetic for Robust Error Estimation**

*Oliver Flatt*  
*University of Utah*

UUCS-22-003

School of Computing  
University of Utah  
Salt Lake City, UT 84112 USA

5 January 2022

## **Abstract**

Interval arithmetic is a simple way to compute a mathematical expression to an arbitrary accuracy, widely used for verifying floating-point computations. Yet this simplicity belies challenges. Some inputs violate preconditions or cause domain errors. Others cause the algorithm to enter an infinite loop and fail to compute a ground truth. Plus, finding valid inputs is itself a challenge when invalid and unsamplable points make up the vast majority of the input space. These issues can make interval arithmetic brittle and temperamental. This thesis introduces three extensions to interval arithmetic to address these challenges. Error intervals express rich notions of input validity and indicate whether all or some points in an interval violate implicit or explicit preconditions. Movability flags detect futile recomputations and prevent timeouts by indicating whether a higher-precision recomputation will yield a more accurate result. And input search restricts sampling to valid, samplable points, so they are easier to find. We compare these extensions to the state-of-the-art technical computing software Mathematica, and demonstrate that our extensions are able to resolve 60.3% more challenging inputs, return 10.2x fewer completely indeterminate results, and avoid 64 cases of fatal error.

## 1 INTRODUCTION

Floating-point arithmetic in scientific, financial, and engineering applications suffers from rounding error: the results of a floating-point expression can differ starkly from those of the analogous mathematical expression [25, 29, 30, 49]. Rounding error has been responsible for scientific retractions [2, 3], mispriced financial indices [20, 38], miscounted votes [51], and wartime casualties [50]. Floating-point code thus needs careful validation; sound upper bounds [19, 27, 46], semidefinite optimization [34], input generation [15, 26], and statistical methods [47] have all been proposed for such validation. Sampling-based error estimation, which can estimate typical, not worst-case, error is especially widely used [9, 42, 44, 45].

Consider the core task of computing the error for a floating-point computation at a given point. This requires computing an error-free ground-truth for a mathematical computation. Interval arithmetic [9, 31, 37, 39, 49] and its many variations [7, 28, 40] is the traditional approach to this problem. Interval arithmetic provides sound bounds on the error-free ground truth, and allows for refining those bounds by recomputing at higher precision, until the ground truth can be estimated to any accuracy. Computing the ground truth this way on a large number of sampled points gives a good estimate of the floating-point error of a computation.

In practice, however, even state of the art implementations struggle on invalid or particularly challenging input points. The bounds are meaningless for invalid inputs; the recomputation process may enter an infinite loop; and valid inputs may be too hard to find. Since error estimation requires sampling a large number of points, these challenges arise frequently. As a result, even state of the art interval arithmetic implementations enter infinite loops, give up too early, or in some exceptional cases even experience memory exhaustion and fatal errors.

We identify three particular challenges for sampling-based error estimation: invalid input points, futile recomputations, and low measure. *Invalid input points* violate explicit or implicit preconditions and cannot meaningfully be evaluated to a ground truth value; consider the square root of a negative number. Validity must thus be soundly tracked and checked, especially for inputs at the borderline between valid and invalid. *Futile recomputation* afflicts some valid inputs: interval arithmetic recomputes with ever-higher precision but without converging on a ground truth answer. These recomputations must thus be soundly cut off early, warning the user that a ground truth value cannot be computed. *Low measure* means valid inputs make up a small proportion of the input space, making it difficult for applications to find sufficiently many. Regions of valid input points must thus be efficiently and soundly identified. No general-purpose strategy exists to address these independent issues; yet all three must be addressed for interval arithmetic to be robust.

This paper proposes three improvements to interval arithmetic that address these issues. First, *error intervals* determine whether all or some of the points in an interval violate implicit or explicit preconditions. Second, *movability flags* detect most input values for which recomputation will not converge, warning the user and avoiding futile recomputation. Third, *input search* discards invalid regions of the input space, focusing sampling on valid points. Combined, our improvements soundly address the issues identified above and make interval analysis more robust.

We implement these improvements in the new interval arithmetic library called Rival and compare Rival to Mathematica's analogous N function on a benchmark suite of 481 floating-point expressions. Rival produces results in dramatically more cases than Mathematica: Rival is able to resolve 18824 inputs, while Mathematica is only able to resolve 11744 (60.3% better). In only 746 cases does Rival return a completely indeterminate result; Mathematica, however, does so in 7617 cases (10.2× worse), and among those cases sometimes enters an infinite loop, runs out of memory, or hard crashes (64 cases). Furthermore, Rival's input search saves 74.6% of invalid points from

being sampled in the first place. An ablation study shows that error intervals, movability flags, and input search effectively handle even difficult inputs across a range of domains.

This paper contributes three extensions to interval arithmetic:

- *Error intervals* to track domain errors and rich notions of input validity (Section 4);
- *Movability flags* to detect when recomputing in higher precision is futile (Section 5);
- *Input search* to uniformly sample valid inputs with high probability (Section 6).

Sections 7 and 8 demonstrate that these extensions successfully address the issues of invalid inputs, futile recomputations, and low measure. Section 9 discusses our experience deploying these extensions to users.

## 2 OVERVIEW

Imagine a simulation that evaluates  $x^y/(x^y + 2)$ . How accurate is this expression? This section describes how Rival's extensions to interval arithmetic help answer the question. For ease of exposition, this section uses two-digit decimal arithmetic as a target precision, with ordinary floating-point as a higher precision arithmetic; targetting single- or double-precision floating-point with arbitrary-precision libraries sees analogous issues.

### 2.1 Interl Arithmetic

Consider an input like  $(x, y) = (3.0, 1.1)$ . At this input point, the expression  $x^y/(x^y + 2)$  can be directly computed with two-digit decimal arithmetic:  $3.0^{1.1} = 3.34 \dots$ , which rounds to 3.3; then  $3.3 + 2 = 5.30 \dots$ , which rounds to 5.3; and finally  $3.3/5.3 = 0.622 \dots$ , which rounds to 0.62. How accurate is this result? To find out, we must compare the computed result with the true value of this expression.

One might estimate that true value using higher-precision computation. In single-precision floating-point, the input evaluates to 0.62605417, which rounds to the two-digit decimal 0.63. Two-digit decimal evaluation therefore has error: it produces 0.62 instead of the correct 0.63. Unfortunately, single-precision floating-point itself has rounding error. After all, in double-precision arithmetic the computed value is not 0.62605417 but 0.62605425  $\dots$ , and even higher precision could produce some third value. So this method of computing the true answer is suspect: how much precision is enough?

Interval arithmetic answers this question using *intervals*, written  $[a, b]$ , which represent any real value between  $a$  and  $b$  (inclusive). Ordinary mathematical operations are then extended to intervals, the two ends always rounded outward to ensure that the resulting interval always contains the mathematically exact value. Returning to the running example, the interval evaluation of  $x^y$  with  $(x, y) = (3.0, 1.1)$  is  $[3.0, 3.0]^{[1.1, 1.1]} = [3.3, 3.4]$ . Next,  $[3.3, 3.4] + [2.0, 2.0] = [5.3, 5.4]$ . Finally,  $[3.3, 3.4]/[5.3, 5.4] = [0.61, 0.65]$ , with those two endpoints computed via  $3.3/5.4 = 0.61$  and  $3.4/5.3 = 0.65$ .<sup>1</sup> Therefore, the true value is between 0.61 and 0.65. A similar interval computation can be done with any precision; with single-precision floating-point, the resulting interval is  $[0.62605417, 0.62605441]$ .

Importantly, this interval is much narrower and both endpoints round to 0.63 in two-digit decimals, meaning that the true value (which must be in the interval) does the same. These *one-value* interval therefore provide a way to evaluate the true value of the expression, to a target precision, using a higher but finite precision. If even single-precision floating-point resulted in too

<sup>1</sup>Division yields its smallest output with a small numerator and large denominator; more generally, evaluating a mathematical operation on intervals means solving an optimization problem, identifying the maximum and minimum values a function can take on over a range of inputs.

99 wide an interval, double- or higher-precision arithmetic could be tried until a one-value interval is  
100 found.

101 Interval arithmetic thus provides a fast way to compute a ground-truth value for an expression  
102 on an input point. This algorithm can be used for a range of purposes. In sampling-based error  
103 estimation, a ground truth is computed for many randomly-sampled inputs to a floating-point  
104 expression, and the actual floating-point answer is compared to the ground truth to determine  
105 how accurate the expression is. However, while this works well on most inputs, on challenging  
106 expressions it can throw errors, enter infinite loops, or fail to find valid inputs.

## 108 2.2 Invalid Input Points

109 Not all inputs to  $x^y/(x^y + 2)$  are as conceptually simple as  $(3, 1.1)$ ; consider  $(x, y) = (-1.1, 7)$ . For  
110 this input, the denominator  $x^y + 2$  evaluates to the interval  $[0.0, 0.1]$  in two-digit decimals, meaning  
111 a division by zero error is possible. But single-precision evaluation shows that the error is a mirage:  
112  $[-1.1, -1.1]^{[7,7]} + 2$  yields the interval  $[0.0512841, 0.0512843]$  for the denominator, proving that  
113 division by zero does not occur. More broadly, in interval arithmetic, an interval may contain both  
114 valid and invalid inputs for an operator, in which case an error can be possible without being  
115 guaranteed.

116 *Error intervals* formalize this reasoning. Each interval is augmented with a boolean interval  
117  $[g, p]$ , where the  $g$  tracks whether an error is guaranteed and the  $p$  tracks whether it is possible; the  
118 interval  $[\top, \perp]$  is illegal, since a guaranteed error implies a possible error. For our running example  
119 with  $(x, y) = (-1.1, 7)$ , the error interval is  $[\perp, \top]$  in two-digit decimal arithmetic, indicating that  
120 error is not guaranteed but is possible, and  $[\perp, \perp]$  in single-precision floating-point, guaranteeing  
121 that no error occurs. Importantly, error intervals can handle borderline cases with recomputation.  
122 For this point, two-digit decimal arithmetic is insufficient to determine whether the input point is  
123 valid, but single-precision is. If even single-precision weren't, double or higher precision could be  
124 tried.

125 Boolean intervals are a type of three-valued logic which can also address expressions with  
126 explicit preconditions or with conditionals that branch on floating-point comparisons. In each  
127 case a boolean interval indicates whether a condition can, or must, be true, and multiple boolean  
128 intervals can be combined with the standard boolean operators. For example, for an expression  $E(x)$   
129 with precondition  $\sin(x) > 0$ , the valid inputs  $x$  are those where  $\sin([x, x]) > 0 \wedge \neg \text{err}(E([x, x])) =$   
130  $[\top, \top]$ , where  $\text{err}(E)$  refers to the error interval of  $E$ . The key for interval arithmetic is that error and  
131 boolean intervals allow expressing rich notions of input validity that integrate with recomputation  
132 and provide sound, error-free guarantees.

## 134 2.3 Futile Recomputation

136 Recomputation guarantees soundness but not speed. To the contrary: repeated, futile recomputa-  
137 tions present a significant practical problem, most commonly due to overflow.

138 Consider our running example  $x^y/(x^y + 2)$ , but with the extreme input  $(x, y) = (10^{10}, 10^{10})$ . Now  
139  $x^y$  is huge—too large for double-precision and even for common arbitrary-precision libraries. So in  
140 interval arithmetic, it evaluates to  $[\Omega, \infty]$  for some largest finite representable value  $\Omega$ . Since  $\Omega$   
141 is already the largest representable value,  $[\Omega, \infty] + 2$  results in  $[\Omega, \infty]$  again, and since the division  
142  $[\Omega, \infty]/[\Omega, \infty]$  can produce a value as small as  $\Omega/\infty = 0$  or as large as  $\infty/\Omega = \infty$ , the final interval  
143 is  $[0, \infty]$ . This interval is too wide to be helpful, and since the problem is overflow, not rounding,  
144 recomputing at a higher precision will not help. But we need to prevent recomputation from taking  
145 futile recourse to higher and higher precisions until an error like timeout, memory exhaustion, or  
146 worse occurs.

148 *Movability flags* detect such futile recomputations. Movability flags are a pair of boolean flags  
 149 added to every interval, which track whether overflow influences each endpoint. When  $x^y$  overflows,  
 150 for example, the infinite endpoint in  $[\Omega, \infty]$  is marked “immovable”, meaning that recomputing it at  
 151 higher precision would yield the same result. These movability flags are then propagated through  
 152 interval computations:  $[\Omega, \infty] + 2$  also has an immovable infinity, and in the final result  $[0, \infty]$  both  
 153 endpoints are immovable. This guarantees that higher-precision evaluation would produce the  
 154 same uselessly-wide interval and cuts off further recomputation. Importantly, immovable endpoints  
 155 distinguish harmful overflows that produce timeouts, like these, from benign overflows that still  
 156 produce accurate results, like in  $1/e^x$ .

157 Movability flags are sound: if an expression evaluates to a wide immovable interval, recomputa-  
 158 tion cannot compute an error-free ground truth. This allows distinguishing harmful overflows from  
 159 difficult expressions that require higher precision to resolve. Movability flags are not complete—the  
 160 problem is likely undecidable [9]—but are accurate enough to dramatically cut the number of cases  
 161 of futile recomputation.

162  
 163

## 164 2.4 Low Measure

165 Many applications of interval arithmetic, such as for error estimation of floating-point expressions,  
 166 require sampling a large number of valid input points, usually via rejection sampling (repeatedly  
 167 sampling points and discarding invalid ones). Often a specific distribution, like uniform sampling,  
 168 is required as well. But for many mathematical expressions the valid, samplable points have *low*  
 169 *measure*: they are a tiny subset of the entire input space. Consider  $\sin^{-1}(x + 2007)$ : the only valid  
 170 inputs are in  $[-2008, -2006]$ , a range that contains approximately 0.0001% of double-precision  
 171 floating-point values. Rejection sampling would yield almost exclusively invalid points! Since  
 172 restrictions on multiple variables combine multiplicatively, the chance a sample is valid decreases  
 173 further for more variables.

174 *Input search* addresses this issue by identifying and ignoring regions of the input space where  
 175 all inputs are invalid or unsamplable. Input search recursively subdivides the sample space into  
 176 axis-aligned *hyperrectangles*, which assign each variable an interval. For each hyperrectangle,  
 177 input search then uses error intervals and movability flags to determine whether points in that  
 178 interval are valid and samplable. In a form of branch-and-bound search, input search then discards  
 179 hyperrectangles that return  $[\perp, \perp]$  and subdivides those that return  $[\perp, \top]$ . For example, for  
 180  $\sin^{-1}(x + 2007)$ , the interval  $[0, \infty]$  is discarded but the interval  $[-\infty, 0]$ , whose error interval is  
 181  $[\perp, \top]$ , is subdivided and searched further.

182 Input search leverages error intervals and movability flags to ensure that difficult regions of the  
 183 input space are handled soundly. This ensures that input regions that contain valid inputs are never  
 184 discarded, and that applications can choose to ignore (or include) regions where movability flags  
 185 prove unsamplability. Input search also offers control over the sampling distribution and tunes the  
 186 subdivision to the target distribution to ensure fast convergence.

187 Put together, error intervals, movability flags, and input search overcome practical issues with  
 188 interval arithmetic that plague state-of-the-art implementations and enable for more effective and  
 189 robust error estimation.

190  
 191

## 192 3 BACKGROUND

193 Arbitrary-precision interval arithmetic with recomputation is the standard method for approximatin-  
 194 g the exact result of a mathematical expression.

195  
 196

### 3.1 Arbitrary-precision Floating Point

Software libraries like MPFR [21] provide floating-point arithmetic at an arbitrary, user-configurable precision  $p$ . This arithmetic represents a finite subset  $S_p$  of the extended reals  $\mathbb{R} \cup \{\pm\infty\}$ . Any real number  $x$  can be *rounded*  $R_p(x)$  to precision  $p$  by monotonically choosing one of the two values in  $S_p$  closest to  $x$ ; we write  $R_p^\downarrow$  and  $R_p^\uparrow$  for rounding down and up specifically. Precision  $p$  implementations of functions like  $\sin$  or  $\log$  are *correctly rounded*: the implementation  $f_p$  of a function  $f$  satisfies  $f_p(x) = R_p(f(x))$ . But larger expressions can be inaccurate even in high precision:  $(1+x) - x$  evaluates to 0 for all inputs  $x > 2^p$ .

### 3.2 Interval Arithmetic

Intervals are an abstract domain over the extended reals: the interval  $[a, b] \in S_p \times S_p$  represents the set  $\{x \in \mathbb{R}^* \mid a \leq x \leq b\}$ . Interval versions  $f_p$  can be constructed for each mathematical function  $f$  to guarantee both soundness,

$$\forall x_1 \in I_1, x_2 \in I_2, \dots, f(x_1, x_2, \dots) \in f_p(I_1, I_2, \dots),$$

and weak completeness,

$$\exists x_1, y_1 \in I_1, x_2, y_2 \in I_2, \dots, f_p(I_1, \dots, I_n) = [R_p^\downarrow(f(x_1, \dots, x_n)), R_p^\uparrow(f(y_1, \dots, y_n))].$$

The points  $x_i$  and  $y_i$  are called *witness points*. For example, since  $[a_1, a_2] + [b_1, b_2]$  is at least  $a_1 + b_1$  and at most  $a_2 + b_2$ , the result must be  $[R_p^\downarrow(a_1 + b_1), R_p^\uparrow(a_2 + b_2)]$ .

### 3.3 Recomputation

Consider an expression  $e$  over the real numbers, consisting of constants, variables, and function applications. In a *target precision*  $p$ , the most accurate representation of its value at  $x$  is the *ground-truth* value  $y^* = R_p(\llbracket e \rrbracket(x))$ , where  $\llbracket e \rrbracket(x)$  is the error-free real-number result. Interval arithmetic can be used to compute  $y^*$ . Evaluate  $e$  to some interval  $[y_1, y_2]$  at precision  $q > p$ ; then

$$y_1 \leq \llbracket e \rrbracket(x) \leq y_2 \implies R_p(y_1) \leq y^* \leq R_p(y_2)$$

by soundness and monotonicity. If  $[y_1, y_2]$  is sufficiently narrow,  $R_p(y_1) = y^* = R_p(y_2)$ , which computes  $y^*$ . We call these narrow intervals *one-value* intervals; computing one just requires finding a large enough precision  $q$ .

A useful strategy for finding  $q$  is to start just a bit past  $p$  and then grow  $q$  exponentially, which guarantees that large precisions are reached quickly. Under certain assumptions, this exponential growth strategy can also be proved optimal in the sense of competitive analysis. Sometimes, though, the  $q$  required is too large to be feasible or may not exist due to limitations like a maximum exponent size. Practical implementations must thus cap the maximum value of  $q$  and return indeterminate results when the cap is reached.

## 4 BOOLEAN AND ERROR INTERVALS

Interval arithmetic is a method for establishing bounds on the output of a mathematical expression. This section introduces *boolean intervals*, which incorporate control flow into the interval arithmetic framework, and demonstrate how their application to domain errors (*error intervals*) allows integrating rich notions of input validity with the soundness and recomputation provided by interval arithmetic.

### 4.1 Modeling Control Flow

Besides real computations, floating-point expressions contain explicit and implicit control flow. However, since intervals represent a range of inputs, the results of a boolean expression can be

indeterminate (true on some inputs in the range and false on others). Boolean intervals (for explicit control flow) and error intervals (for implicit control flow) represent this possibility.

DEFINITION 1. *Boolean intervals*  $[\top, \top]$ ,  $[\perp, \perp]$ , and  $[\perp, \top]$  denote the sets  $\{\top\}$ ,  $\{\perp\}$ , and  $\{\perp, \top\}$  of booleans.

Boolean intervals form a standard three-valued logic, with  $[\perp, \top]$  representing the indeterminate truth value and with standard semantics for conjunction, negation, and disjunction. Comparison operators on intervals can return boolean intervals:  $[1, 3] < [4, 5] = [\top, \top]$ , while  $[1, 4] < [3, 5] = [\perp, \top]$ . Explicit control flow like the if operator unions its two branches when the condition is indeterminate.

Expressions also have implicit control flow due to domain errors. As with comparisons, whether a domain error occurs can be indeterminate; each expression thus returns an *error interval*,<sup>2</sup> a boolean interval describing whether a domain error occurred during the evaluation of the expression; we write  $\text{err}(e)$  for the error interval of expression  $e$ .<sup>3</sup> When an error is possible but not guaranteed, the function's output is still meaningful. For example, the expression  $\text{pow}(x, y)$  raises a domain error when  $x$  is negative and  $y$  is non-integral. Thus,  $\text{pow}([-1, 2], [1, 5])$  returns  $[-1, 32]$  with error interval  $[\perp, \top]$  because  $\text{pow}(-1, 1) = -1$  is the minimum possible output,  $\text{pow}(2, 5) = 32$  is the maximum possible output, and  $\text{pow}(-1, 2.5)$  demonstrates that domain error is possible.<sup>4</sup>

At higher precisions input intervals are narrower, so errors that were once possible may be ruled out. Error intervals make it possible to distinguish those cases (which have error interval  $[\perp, \top]$ ) from cases like  $\sqrt{-1}$  (which have error interval  $[\top, \top]$ ) where an error is guaranteed even at higher precision. Consider  $\sqrt{\cos(x)}$ . At low precisions and with large  $x$  values,  $\cos(x)$  can evaluate to  $[-1, 1]$ , so that  $\sqrt{\cos(x)}$  is possibly invalid. Without error intervals,  $\sqrt{\cos(x)}$  would indicate the possible error by returning  $[\text{NaN}, 1]$ . The error interval  $[\perp, \top]$  gives more information: it indicates that the expression could be a valid at a higher precision. But if, at a higher precision,  $\cos(x)$  is found to be strictly negative, the error interval will be  $[\top, \top]$  and cut off further recomputation.

## 4.2 Point Validity

Both explicit and implicit control flow can affect whether a point is a valid input. Combinations of boolean and error intervals allow expressing these rich notions of input validity. Consider how four validity requirements, drawn from the Herbie tool [42], can be formalized as expressions yielding boolean intervals. Each validity requirement applies to the input  $(x, y, \dots)$  of an expression  $E$ .

First, Herbie requires each input variable to have a finite value in the target precision. Given the target precision's smallest and largest finite values  $K_-$  and  $K_+$ , the comparison  $K_- \leq x \leq K_+$  expresses that requirement for the variable  $x$ . The constants  $K_-$  and  $K_+$  can be computed by ordering all values in the target precision and taking those just after/before  $\mp\infty$ . Note that the expression  $-\infty < x < +\infty$  is equivalent in the target precision, but is much weaker in higher precisions that contain values between  $K_+$  and  $+\infty$ .

Second, Herbie requires the output value, when computed without rounding error, to be finite as well. The expression  $K_- \leq E \leq K_+$  enforces this requirement. Note that boolean intervals allow incorporating arbitrary mathematical formulas into validity requirements.

Third, Herbie requires  $E$  to not raise a domain error. The expression  $-\text{err}(E)$ , which uses  $\text{err}$  to access the error interval, enforces this requirement. Since interval operations are meaningful even when an input interval contains invalid points, combining this requirement with the second

<sup>2</sup>Error intervals can be viewed as an abstraction not just of the real results of the computation, but of the Error monad the computation occurs within.

<sup>3</sup>In this way,  $\text{err}$  is analogous to a try/catch/else block.

<sup>4</sup>When an expression's error interval is  $[\top, \top]$ , of course, the interval bounds are meaningless by necessity.

one allows ignoring points that are *either* too large or invalid without using higher precision to determine which restriction applies.

Fourth, Herbie users can provide a precondition that valid inputs have to satisfy. That precondition can be interpreted as a boolean interval expression to enforce this requirement. Since interval arithmetic already accounts for rounding error, the precondition is checked without rounding error. If the chosen precision isn't high enough, the precondition returns the indeterminate boolean interval  $[\perp, \top]$ , and it is automatically recomputed at a higher precision to determine whether the input is in fact valid.

Since boolean intervals admit conjunction, all four of these requirements can be combined into a single formula that for input validity. That single formula is crucial to input search Section 6.

## 5 MOVABLE AND IMMOVABLE INTERVALS

For some inputs, computing a ground-truth value using interval arithmetic and recomputation is impossible; one common problem is an overflow that persists across precisions. When no ground-truth can be computed for certain input points, the user must be warned, and recomputation must be prevented from causing timeouts. *Movability flags* on each endpoint achieve this goal. Movability flags are set when operations overflow, and track the influence of that overflow on further computation. When overflow prevents the computation of a ground-truth value, movability flags warn the user without timing out.

### 5.1 Movable and immovable intervals

Formally, an interval is augmented with two movability flags, one associated with each of its endpoints; an endpoint is *immovable* when its movability flag is set. In examples, we write exclamation marks for immovable endpoints, so that  $[0, 1!]$  is an interval with a movable left endpoint 0 and an immovable right endpoint 1. The movability flags describe how intervals get narrower at higher precision.

**DEFINITION 2.** *One interval refines another, written  $[a', b'] < [a, b]$ , when  $b' = b$  and both are immovable or when  $b' \leq b$  and  $b$  is movable; and likewise  $a' = a$  and both are immovable or when  $a' \geq a$  and  $a$  is movable.*

The goal of movability flags is to determine whether recomputation at higher precision could possibly yield a narrower interval output. Specifically, they guarantee:

**THEOREM 3.**  $\llbracket e \rrbracket_q(x) < \llbracket e \rrbracket_p(x)$  for all  $e, x$ , and  $q > p$ .

When an expression evaluates to an interval that is not one-value and which has two immovable endpoints, any higher-precision computation will do the same, and recomputation is futile. Alternatively, an interval with a single immovable endpoint is guaranteed to produce that value (if any).

Program inputs and constants (named and numeric) form a particularly simple case of Theorem 3. Program inputs are drawn from the target precision and can be directly represented in higher precision; they are thus represented by intervals  $[!x, x!]$  with both endpoints marked immovable. Constants like  $\pi$ , however, cannot be exactly represented and have both endpoints of their interval marked movable.

For expressions larger than just constants, Theorem 3 is proved by induction. Operators in the tree must properly interpret movability flags on their inputs and set movability flags on their outputs. The key is to ensure that each operator  $f$  preserves refinement for its inputs:

**PROPERTY 4.**  $f_q(I'_1, \dots, I'_n) < f_p(I_1, \dots, I_n)$  when  $q > p$  and all  $I'_i < I_i$ .

344 Importantly, if  $f_p(I, J)$  has an immovable endpoint, then  $f_q(I', J')$  must share that immovable  
 345 endpoint. When all operators satisfy Property 4, Theorem 3 is preserved by structural induction.  
 346 The rest of this section describes how to guarantee Property 4 for various operations.

## 347 5.2 Immovability from Overflow

348  
 349 Movability flags are set when the limitations of arbitrary-precision computation make it impossible  
 350 for higher precision to make the result of a computation more precise.

351 Consider the  $\exp(x)$  function, which computes  $e^x$ . The MPFR library represents arbitrary-  
 352 precision values as  $\pm(1 + s)2^h$ , where the exponent  $h$  has a fixed upper bound  $H$ . It thus cannot  
 353 represent any values larger than  $2^{H+1}$ , so inputs  $x > \log(2^{H+1})$  to  $\exp$  overflow at any precision.<sup>5</sup>  
 354 This fact allows  $\exp$  to set movability flags on its output.

355 Consider the computation  $\exp_p([0, 10^{10!}]) = [1, \infty!]$ . The output's right endpoint is  $\exp(10^{10})$ ,  
 356 which rounds up to  $+\infty$ . Since  $10^{10}$  exceeds the overflow threshold, the rounding will occur at any  
 357 precision  $p$ ; and since  $10^{10}$  is immovable in the input interval, it will be the right endpoint of any  
 358 refinement of the input. So it is valid to mark the right output endpoint immovable. Alternatively,  
 359 for inputs like  $[10^{10}, 10^{11}]$  whose left endpoint exceeds the overflow threshold, the output's right  
 360 endpoint can be marked immovable even if both input endpoints are movable. Thus outputs can  
 361 have immovable endpoints even with entirely movable inputs. Note, however, that in this case the  
 362 left output endpoint cannot be marked immovable, since it will be the largest representable finite  
 363 value and that can typically increase with higher precision.

364 This general logic of overflow extends to the  $\exp2$  function with threshold  $H + 1$ , and to  $\text{pow}$ ,  
 365 which detects overflow via the identity  $x^n = \exp(n \log x)$ .

## 366 5.3 Propagating Immovability

367  
 368 Once one operator sets a movability flag, that movability flag must be propagated through later  
 369 computations. That requires keeping track of the exactness of intermediate computations.

370 An arbitrary-precision function can be thought of as  $R_p(f(x_1, \dots, x_n))$ , where  $f$  is the exact,  
 371 real-valued function. But when the exact result  $f(x_1, \dots, x_n)$  is representable in precision  $p$ , the  
 372 rounding function  $R_p$  is the identity. Since  $f(x_1, \dots, x_n)$  is then also representable at all higher  
 373 precisions  $q$ , we have  $f_q(x_1, \dots, x_n) = f_p(x_1, \dots, x_n)$ . This is the key to guaranteeing Property 4.

374 In general, an output interval's endpoint will be immovable if it is the exact result of a computation  
 375 on immovable inputs. Monotonic functions are a good illustration of the logic. Consider a single-  
 376 argument, monotonic function  $f$ . Its behavior on intervals is particularly simple:

$$377 f_p([a, b]) = [R_p^\downarrow(f(a)), R_p^\uparrow(f(b))].$$

379 At a higher precision, the rounding behavior could change and then so would the output interval.  
 380 However, if  $a$  is immovable, any refinement  $[a', b'] < [a, b]$  must have  $a = a'$  and  $b \leq b'$ . If  $f(a)$  is  
 381 exactly computed,  $R_p^\downarrow$  is the identity. In this case:

$$382 f_q([a', b']) = f_q([a, b']) = [R_q^\downarrow(f(a)), R_q^\uparrow(f(b'))]$$

$$383 = [f(a), R_q^\uparrow(f(b'))] < [f(a), R_p^\uparrow(f(b))] = f_p([a, b])$$

384  
 385 In other words, for monotonic functions, exact computations on immovable endpoints result in  
 386 immovable endpoints. For example,  $\text{sqrt}_p([0, 4!]) = [0, 2!]$  because the square root of 4 is exactly  
 387 2. However,  $\text{sqrt}_p([0, 2!]) = [0, 1.414 \dots]$  has a movable right endpoint because the square root  
 388  
 389

390 <sup>5</sup>The value of  $H$  for MPFR depends on the platform, so Rival calculates its overflow threshold at runtime, using 80 bits of  
 391 precision and consistent rounding up.

of 2 is not exact at any precision. Arbitrary precision libraries such as MPFR report whether a floating-point operator is exact, making this rule easy to operationalize.

For general interval functions, the task is more complex because the endpoints of the output interval need not be computed from input endpoints. General interval functions may be defined as computing  $f$  on *witness points*  $a'$  and  $b'$ :

$$f_p([a_1, b_1], \dots, [a_n, b_n]) = [R_p^\downarrow(f(a'_1, \dots, a'_n)), R_p^\uparrow(f(b'_1, \dots, b'_n))]$$

Here, the witness points  $a'$  and  $b'$  are computed as an intermediate step to minimize and maximize  $f$  over the input intervals. In such a scenario, output endpoints are immovable only if the witness points are guaranteed to be the same in all refinements of the input intervals, and when  $f$ 's output is representable:

LEMMA 5. *Suppose the left output endpoint of  $f_p$ , computed via  $f(a'_1, \dots, a'_n)$ , is representable in precision  $p$ . That endpoint may be marked immovable if, for all  $i$ , either 1)  $a'_i = a_i$  and  $a_i$  is immovable; 2)  $a'_i = b_i$  and  $b_i$  is immovable; or 3)  $a_i$  and  $b_i$  are both immovable and  $a'_i$  is computed exactly. The analogous applies to  $f_p$ 's right endpoint.*

PROOF. Since  $a'_i$  minimizes  $f$  over the intervals  $I_0 \dots I_n$ ,  $f_p(a'_1, \dots, a'_n)$  bounds  $f$  from below. If, furthermore,  $f_p(a'_1, \dots, a'_n)$  is exactly computed, then  $f_p(a'_1, \dots, a'_n) = f_q(a'_1, \dots, a'_n)$  for any higher precision  $q$ .

Consider conditions (1) and (2) first. If (1) or (2) holds,  $a'_i$  is guaranteed to be a member of any refinement  $I_i < [a_i, b_i]$  by immovability. Since it is the witness point for  $f$  at the lower precision, it is a valid witness point at any higher precision too.

Alternatively, consider condition (3). If (3) holds, the only refinement of  $[a_i, b_i]$  is itself, so  $a'_i$  is always a member of that refinement. Since  $a'_i$  is additionally computed exactly, it is computed identically and thus still a witness point at any higher precision.

Since, under (1), (2), or (3),  $a'_i$  is a valid witness point at any higher precision, the left output endpoint computed from it can be marked immovable.  $\square$

In practice, computing witness points is usually easy. Consider the absolute value function  $\text{fabs}([a, b])$ : the left witness point is  $a$ ,  $b$ , or 0, and the right witness point is  $a$  or  $b$ . In some cases, it's best to avoid computing witness points directly. For example,  $\cos$  has a minimum of  $-1$  at  $\pi$ . Instead of computing  $\pi$  inexactly,  $-1$  can be used directly, so that for example  $\cos([!3, 4!]) = [-1, 0.653 \dots]$ . In terms of Lemma 5, the witness point  $\pi$  is computed at infinite precision and then  $\cos(\pi)$  is evaluated exactly. Overflow frequently produces immovable infinite endpoints, but operations on infinities are generally exact, so Lemma 5 retains immovability in these cases too.

## 5.4 Function-specific Movability Reasoning

Function-specific reasoning sometimes provides additional cases in which movability flags may be set. The most common case is multi-argument functions where individual arguments are special values like zero or infinity. These special values allow setting movability flags even when some of the input intervals are entirely movable.

Addition is an illustrative example. Addition of anything with infinity yields infinity; thus,  $[1, +\infty!] + [1, 2] = [2, +\infty!]$ , with the right output endpoint immovable. Multiplication by an immovable 0, like addition of infinity, results in an immovable 0, even if the other argument is movable. Multiplication by infinity is similar but more complex: consider  $[-1, 1] \times [!1, +\infty!] = [-\infty, +\infty]$ , where the output interval is movable because the movable left-hand argument may refine to a strictly positive or a strictly negative interval. Thus, much like overflow detection requires knowing that the input interval is greater than a given threshold, handling infinite values in multiplication requires knowing the sign of the input interval:

```

442  $T, F, O = \emptyset, \emptyset, \{[-\infty, \infty]^n\}$ 
443 for  $i \in [1, \dots, N]$  :
444    $O' = \emptyset$ 
445   for  $x \in O$  :
446      $y = P(x)$ 
447     if stuck( $y$ ) : warn "Discarding points in  $x$ "
448     elif  $y = [\top, \top]$  :  $T = T \cup \{x\}$ 
449     elif  $y = [\perp, \perp]$  :  $F = F \cup \{x\}$ 
450     else :  $O' = O' \cup (\text{split } x \text{ along } (i \bmod n))$ 
451    $O = O'$ 
452 return  $T \cup O$ 
453
454
455

```

Fig. 1. Pseudo-code for input search. Given a precondition  $P$  over  $n$  real variables, input search finds a set  $T \cup O$  of disjoint hyperrectangles, such that  $T \cup O$  contains all valid, samplable input points. The  $T$ ,  $F$ , and  $O$  sets are initialized on the first line and then updated over the course of  $N$  search iterations. In each iteration, the hyperrectangles in  $O$  are evaluated. Stuck intervals are discarded; intervals where the precondition has undetermined truth value are split into two along a dimension picked by round robin.

LEMMA 6. Let  $c = a'_i \times b'_i$  be an endpoint of an interval returned by multiplication. The endpoint may be marked immovable if: 1) both  $a'_i$  and  $b'_i$  are immovable and  $c$  is computed exactly; or, 2)  $a'_i$  is zero and immovable (or likewise for  $b'_i$ ); or, 3)  $a'_i$  is infinite and immovable and  $[b_1, b_2]$  does not contain zero (or likewise for  $b'_i$ ).

PROOF. The first case just restates Lemma 5. In the second case,  $a'_i$  is immovable, so  $a'_i = 0$  is in any refinement. Since  $0 \times b'_i = 0$  for any  $b'_i$ , the result is immovable. Finally, in the third case, while  $b'_i$  may change when refined, its sign cannot, since  $[b_1, b_2]$  does not contain zero. Since  $a'_i$  is immovable,  $a'_i \times b'_i$  takes on a fixed sign in any refinement. Since  $a'_i$  is additionally infinity, the resulting  $+\infty$  or  $-\infty$  is immovable.  $\square$

Immovable infinite endpoints are often caused by overflow, so this multiplication-specific reasoning is essential to detecting immovability for many expressions.

## 6 INPUT SEARCH

Most applications of interval arithmetic, whether using pure sampling or input generation or exhaustive testing, start by selecting one or many valid input points. But for many expressions, valid inputs are rare, making finding one challenging. And some applications (such as sampling-based error estimation) further require uniform sampling of valid inputs. Input search addresses this issue by discarding invalid portions of the input space and focusing on valid points.

### 6.1 Intervals for Uniform Sampling

Boolean and error intervals can evaluate input validity for interval inputs, as shown in Section 4. They can thus prove a whole range of inputs to be valid or invalid at once. Consider a validity condition  $P$  expressed via boolean and error intervals, and input intervals  $I_1, I_2, \dots, I_n$ . Then  $P(I_1, I_2, \dots, I_n)$  can yield  $[\top, \top]$ ,  $[\perp, \perp]$ , or  $[\perp, \top]$ , depending on whether these intervals contain only valid points, only invalid points, or a mix.

More formally, with  $n$  free variables and a target precision  $p$ ; the total number of input points is then  $|S_p|^n$ . The hyperrectangle  $I_1 \times \dots \times I_n$ , where each free variable is bound by an interval, has weight  $w = |S_p \cap I_1|/|S_p| \cdot \dots \cdot |S_p \cap I_n|/|S_p|$ . Now consider a set of hyperrectangles that partition the total input space. Uniformly sampling from the  $[\top, \top]$  hyperrectangles and rejection sampling from the  $[\perp, \top]$  hyperrectangles, both with probability proportional to  $w$ , yields valid inputs points, uniformly sampled from the input space. If uniform sampling is unnecessary, the weights can simply be ignored; or if some non-uniform distribution is required the weight  $w$  can use the cumulative distribution function in place of the uniform size  $|S_p \cap I_k|$ .

491 This basic idea forms the basis of a branch-and-bound algorithm shown in Figure 1 that iteratively  
 492 decomposes the complete input space into a subset with a higher proportion of valid points. The  
 493 search algorithm maintains three sets of hyperrectangles,  $F$ ,  $T$ , and  $O$ , initialized with  $T = F = \{\}$   
 494 and with  $O = \{S_p^n\}$  containing the whole input space. At each step, the validity expression is  
 495 evaluated on each hyperrectangle in  $O$ . If the result is  $[\perp, \perp]$  or  $[\top, \top]$ , the interval is moved to  $F$   
 496 or  $T$ , but in the  $[\perp, \top]$  case the interval is subdivided and both halves are placed back into  $O$ . At  
 497 each step, the intervals in  $F \cup T \cup O$  cover the input space; the process terminates once no intervals  
 498 are left in  $O$ , or after a fixed number of iterations (14 in our implementation). Thus, as the search  
 499 algorithm proceeds,  $T$  and  $F$  accumulate hyperrectangles covering valid and invalid regions of  
 500 the input space, while  $O$  contains increasingly small hyperrectangles whose validity is unknown.  
 501 Once search is finished, points are sampled proportionally to weight from the hyperrectangles in  
 502  $T \cup O$ —from  $T$  directly and from  $O$  via rejection sampling.

503 Two special twists are required to make this algorithm effective and maximally general. First,  
 504 when splitting  $[\perp, \top]$  hyperrectangles, the algorithm performs best when the two resulting hy-  
 505 perrectangles contain the same number of values. Since floating-point binary representations are  
 506 sorted, the best split is midway between the *binary representations* of the endpoints in the target  
 507 precision. Second, it is important that the hyperrectangles in  $T \cup F \cup O$  strictly partition the input  
 508 space—no point is in two hyperrectangles. However, recall that intervals are inclusive on both  
 509 ends, so one point can end up in two hyperrectangles and end up oversampled. The fix for this is  
 510 to increment one endpoint to the next floating-point value when splitting a hyperrectangle. The  
 511 split value is represented in the resulting lower hyperrectangle, and the next value is represented  
 512 in the resulting higher hyperrectangle. This allows input search to guarantee uniform sampling  
 513 when that is required, such as for sampling-based error estimation.

## 514 6.2 Range Analysis

515 Hyperrectangle subdivision aids sampling because partitioning the input space can model complex  
 516 relations between variables. However, it is inefficient for the most common type of condition:  
 517 constant bounds on variables. Input search uses a static analysis of the precondition, called range  
 518 analysis, to address this issue. The main role of range analysis is to detect comparisons between  
 519 variables and constants, and propagate these comparisons through boolean operations. For each  
 520 variable, the result is a set of intervals, where inputs outside the set are provably invalid. Input  
 521 search then forms the cartesian product of these per-variable interval sets and uses the resulting  
 522 hyperrectangles as a starting point for the subdivision search.

523 Formally, range analysis returns disjoint intervals  $X_{i1} \cup X_{i2} \cup \dots \cup X_{in}$  for each variable  $x_i$ , where:

$$524 \llbracket P \rrbracket(x) \implies \forall i, x_i \in X_{i1} \cup X_{i2} \cup \dots \cup X_{in}.$$

525 Range analysis traverses the precondition from the bottom up. A comparison  $x < C$  becomes the  
 526 range table  $x \mapsto \{[-\infty, C]\}$ , and boolean operations correspond to straightforward intersections  
 527 and unions of range tables. For non-trivial comparisons like  $\exp(x) < y$ , range analysis just returns  
 528 a non-restrictive range table; these more complex cases are better handled by branch-and-bound  
 529 search.  
 530 search.

## 531 6.3 Unsamplable Inputs

532 As discussed in Section 5, some valid input points are *unsamplable*: interval arithmetic cannot  
 533 compute a ground truth value and so cannot accurately measure error. Input search must warn  
 534 the user if it finds unsamplable points, ideally providing an example to help the user debug the  
 535 issue. But unsamplable points often represent unintended inputs; users often react to the warning  
 536 by adding a precondition ruling the unsamplable points invalid. So many applications of interval  
 537 arithmetic.

540 arithmetic in fact discard unsamplable points, after issuing a warning. Since unsamplable inputs  
 541 often come in large contiguous regions, input search ought to detect them, warn the user, and  
 542 ignore them.

543 Movability flags let input search do this. Define an interval to be stuck when it has two immovable  
 544 endpoints but is not one-value.<sup>6</sup> In other words,  $x$  is unsamplable if  $\text{stuck}(\llbracket e \rrbracket_p(x))$ . Surprisingly,  
 545 stuck satisfies Property 4:

546 **THEOREM 7.** *Suppose  $\text{stuck}(\llbracket e \rrbracket_p(X))$  for a hyperrectangle  $X$ . Then for any hyperrectangle  $X' < X$   
 547 (interpreted pointwise) and precision  $q > p$ ,  $\text{stuck}(\llbracket e \rrbracket_q(X'))$  also holds.*

549 As a corollary, if all intervals in a hyperrectangle  $X$  have both endpoints movable, and  $\llbracket e \rrbracket_p(x)$  is  
 550 stuck, then all points  $x \in X$  are unsamplable, because their point intervals  $[x, x]$  refine  $X$ .

551 **PROOF.** Any refinement of a stuck interval is also stuck. Now, by Theorem 3, if  $X' < X$ ,  
 552  $\llbracket e \rrbracket_q(X') < \llbracket e \rrbracket_p(X)$ . So if  $\llbracket e \rrbracket_p(X)$  is stuck, so is  $\llbracket e \rrbracket_q(X')$ .  $\square$

554 Input search thus marks each hyperrectangle's endpoints as movable before evaluating the  
 555 validity condition. When a hyperrectangle evaluates to a stuck interval, a warning is can be raised,  
 556 informing the user to the presence of unsamplable inputs. To aid in debugging, a point is chosen  
 557 from  $X$  and provided to the user as an example unsamplable input; by Theorem 7, any point in  $X$   
 558 will do. Applications can configure input search to then discard  $X$  by moving it to the  $F$  set.

## 560 7 EVALUATION

561 We implement error intervals, movability flags, and input search in the Rival interval arithmetic  
 562 library<sup>7</sup> and evaluate it against the state-of-the-art Mathematica 12.1.1 N function. Mathematica is  
 563 a proprietary software system advertised as “the world’s definitive system for modern technical  
 564 computing” and priced at roughly \$1 500 per person per year for industrial use; our research is  
 565 supported by an educational site license. Mathematica is also widely used across all science and  
 566 engineering domains; a Google Scholar search for ‘Mathematica’ yielded 10,000 hits, while an ArXiv  
 567 search found 200 results for 2021 alone. Being a proprietary system, it’s impossible to know with  
 568 certainty how N works; however, its maximum precision flag, its configurability for target precision  
 569 or target accuracy, and its overflow and underflow warnings all suggest that N implements interval  
 570 arithmetic.

571 Our results are plotted in Figure 2. We find that among 19895 challenging input points, Rival  
 572 is able to resolve 60.3% more and has 10.2× fewer indeterminate results. Its error-handling-first  
 573 design (with error intervals and movability flags) also avoids infinite loops, memory exhaustion,  
 574 and hard crashes, which afflict Mathematica in 64 cases. Furthermore, input search can avoid 74.6%  
 575 of the invalid inputs in the first place.

### 577 7.1 Methodology

578 Mathematica’s N function can compute mathematical expressions exactly to an arbitrary precision,  
 579 precisely the task supported by Rival. We chose Mathematica as our baseline after comparing  
 580 multiple traditional interval arithmetic libraries. Of all the evaluated libraries, Mathematica was  
 581 the most robust and supported the most functions. For example, the MPFI library does not handle  
 582 errors soundly (unlike Rival’s error intervals), has no built-in support for recomputation (unlike  
 583 Rival’s movability flags), and lacks support for some mathematical functions like pow and erf.

584 <sup>6</sup>If only finite values are valid, intervals with a single immovable, infinite endpoint can also be considered stuck.

585 <sup>7</sup>Rival is written in roughly 1650 lines of Racket, including 650 lines of interval operators, 650 lines of search and recompu-  
 586 tation algorithms, 100 lines of module headers, and 250 lines of tests, and is publicly available as free software, at a URL  
 587 removed for anonymization.

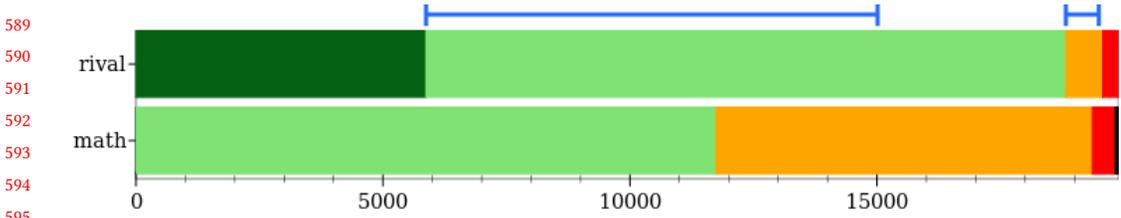


Fig. 2. Rival versus Mathematica’s N on approximately 19895 hard inputs from the 481 Herbie 1.4 benchmark suite. Rival achieves better results for each outcome threshold, including successfully sampled points (dark green), points proven invalid (light green), unsamplable points (orange), or points with unknown results (red). Additionally, some points cause Mathematica to run out of memory or crash, shown in black. The blue intervals mark invalid and unsamplable points avoided by input search.

For an interval arithmetic library, soundness and weak completeness already require optimal interval endpoints, so traditional interval arithmetic libraries generally have the same results as Rival *without* its extensions. Meanwhile, Mathematica is a widely-used commercial product with extensive support for all mathematical functions and presumably-robust implementations of all of them. Furthermore, Mathematica’s documentation suggests that N is intended to handle invalid inputs soundly and cut off recomputation in some cases of over- and under-flow, making it at least comparable to Rival.

We model our evaluation task on sampling-based error estimation; that is, on the task of randomly sampling inputs to an expression and determine the ground-truth value of the expression at those inputs. We use a collection of 481 floating-point expressions drawn from sources like numerical methods textbooks, mathematics and physics papers, and surveys of open-source code, collected in the Herbie 1.4 benchmark suite. These benchmarks consist of mathematical expressions with 0–89 operations and 0–16 variables; 4 benchmarks contain conditionals and 47 have user-defined preconditions. Each benchmark is evaluated to double precision at 8 256 randomly sampled input points (for N, to the equivalent decimal precision). Both N and Rival are limited to 10240 bits; Rival is also limited to 31-bit exponents.<sup>8</sup> In error estimation, most inputs to an expressions are “easy” points where any technique will succeed; the crux of the problem are the challenging or marginal points. To focus on these challenging points, we ignore inputs that both Mathematica and Rival can sample, and focus on the remaining more challenging points, which are 19895 out of the overall 126720. Evaluations are performed on a machine with an i7-4790K CPU (at 4.00GHz) and 32GB of DDR3 memory running Debian 10.0 (Buster), Racket 7.9 BC, and MPFR version 4.0.2-1.

The main methodological challenge is matching Mathematica’s and Rival’s semantics. Mathematica’s N function supports floating-point, fixed-precision, and exact computation; we convert the sampled input points to exact rational values to ensure that exact computation is performed. We also wrap each intermediate Mathematica operator to signal an error for Indeterminate or Complex outputs. We compare Mathematica’s and Rival’s outputs, when both tools are able to compute a value, to check that subtleties like atan2 argument order are correctly aligned between the two systems. We also capture warnings and errors, which Mathematica uses to indicate whether its evaluation is sound. These cross-checks give us confidence that Mathematica and Rival are asked to evaluate the same expressions on the same inputs. (Two points fail these cross-checks at the time of this writing. We have investigated these inputs, discovered them to trigger a bug in Mathematica, reported the bug to Wolfram support, and had it confirmed.) Rival’s input search

<sup>8</sup>Mathematica’s precision limit is set via the \$MaxExtraPrecision variable [52]. Its exponent limit is not configurable, but experiments suggest an exponent limit of approximately 37 bits.

638 has no direct analog in Mathematica, so for this experiment we randomly selected points without  
639 using input search. However, we did test whether input search would have allowed Rival to avoid  
640 sampling that point.

641 Evaluating Rival against Mathematica requires understanding Mathematica's behavior on in-  
642 valid and challenging inputs. Unfortunately, unlike Rival, Mathematica is proprietary, and the  
643 documentation is not always specific. Like Rival, Mathematica's N attempts to detect invalid inputs  
644 and futile recomputations. In general, for invalid inputs Mathematica either returns a complex  
645 number or the special Indeterminate value for some intermediate operator. In Figure 2, these cases  
646 are considered proven domain error (light green). Mathematica also raises a variety of warnings,  
647 including for overflow and underflow; we conservatively assume these warnings are sound and are  
648 reached at low precision, and mark these cases unsamplable (orange in Figure 2). For Rival this  
649 status is used only when the interval is proven stuck before the maximum precision is reached.  
650 (More broadly, we interpret ambiguous cases as generously as possible toward Mathematica to  
651 ensure a fair comparison.) When Rival or Mathematica reach their internal precision limit, we use  
652 the unknown result status (red in Figure 2).

653

## 654 7.2 Results

655 Of the 19895 hard input points across our 481 benchmarks, Mathematica is able to resolve only  
656 11744, while Rival is able to resolve 18824 points (60.3% better). Specifically, Rival samples 5872  
657 points and proves a further 12952 to be invalid; Mathematica samples only 24 and proves only 11720  
658 invalid. Rival's error intervals, which integrate the detection of invalid inputs with recomputation,  
659 account for the difference in the number of points proven invalid.

660 Not only does Rival resolve more cases, it resolves nearly a superset of the cases resolved  
661 by Mathematica. Rival resolves 5877 points not resolved by Mathematica; Mathematica resolves  
662 only 26 points that Rival cannot. (Without access to Mathematica's internals it's hard to say how  
663 Mathematica resolves those 26 points. Our best guess is that it may be simplifying the benchmark  
664 expressions before evaluating them at an input.) Furthermore, when Mathematica is unable to  
665 resolve an input, it sometimes (64 cases) enters an infinite loop or declares that insufficient memory  
666 is available.<sup>9</sup> We render these inputs in black in Figure 2; Rival never crashes in this way. Moreover,  
667 in 16 cases, the Mathematica process crashes and must be killed with SIGKILL, which may cause  
668 the user to lose work.

669 Both Rival and Mathematica are unable to resolve some inputs: 1071 points for Rival and 8151 for  
670 Mathematica. In these cases it's important to give additional information to the user: some of these  
671 points may be resolvable with more precision, time, or memory, but others cannot be resolved due  
672 to algorithmic limitations. Rival's movability flags, which prove 746 points unsamplable, soundly  
673 detect algorithmic constraints; Mathematica's various warnings affect 7617 points (10.2× worse)  
674 and are unsound, sometimes triggered even for samplable inputs. Mathematica marks unsamplable  
675 5830 points that Rival successfully samples as a result. Thus, unlike Rival's movability flags, they do  
676 not advise whether the user ought to try again with more precision or abandon the computation. At  
677 the same time, Mathematica reaches its internal precision limit 1.4× more often than Rival because  
678 Rival's movability flags allow it to identify unsamplable points at low precision.

679 For applications that need to sample valid inputs, Rival's input search is additionally helpful.  
680 Input search is able to avoid sampling 9142 invalid inputs and 668 unsamplable inputs, thereby  
681 cutting the rate of invalid and unsamplable points by 74.6% (the blue intervals in Figure 2). This  
682 means that using Rival for sampling-based error estimation would yield even better results than  
683

---

684 <sup>9</sup>We limit the computation to one second per input point; all 32GB of memory are available to Mathematica. Rival never  
685 runs out of memory, or takes longer than 250 ms to complete a computation.

686

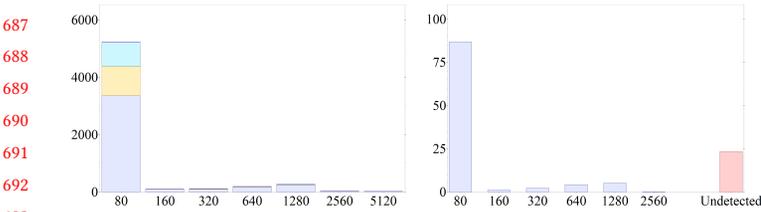


Fig. 3. Counts of valid points (bottom, blue), domain errors (middle, yellow), infinite outputs (top, light blue), and precondition failures (too rare to be visible) for points sampled without input search. 80 bits of precision suffices for most points.

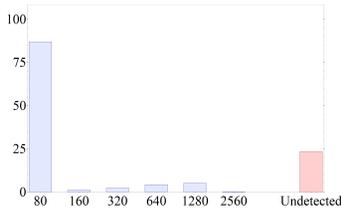


Fig. 4. Counts of unsamplable points. Rival’s movability flags detect all but 19.0% of unsamplable points, with the majority detected at just 80 bits of precision. Detecting unsamplable points avoids timeouts and user disappointment.

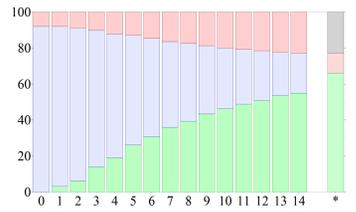


Fig. 5. Size of the  $T$  (bottom, green),  $O$  (middle, blue), and  $F$  (top, red) sets after each iteration of Rival’s input search. The “0” bar shows the results of range analysis; the “\*” bar shows the results of rejection sampling. On average, input search avoids 68.6% of invalid samples.

those suggested by Figure 2, since input search allows Rival to avoid many challenging, invalid points. The overall picture of these results is that Rival is more capable than Mathematica: it is able to sample more points, discard more invalid points, and leave fewer indeterminate cases unresolved. It also does so without timeouts, memory exhaustion, or fatal crashes. Yet it also offers sound guarantees and valuable extensions like input search.

As a result of Rival’s improvements in interval arithmetic, Rival is 9.87× faster than Mathematica. Mathematica’s timeouts, OOM, and crashes also add to its runtime, but Rival is 9.03× faster than Mathematica even after excluding these inputs. This is a consequence of Rival’s better handling of invalid inputs and futile recomputation: any interval arithmetic library spends the bulk of its time doing high-precision arithmetic. For one library to be faster than another, then, that library must identify invalid inputs, detect futile recomputation, sample fewer such points, or in some other way do fewer high-precision operations. That’s precisely what Rival’s extensions do; those extensions ultimately require just a few comparisons, so they add negligible overhead time, but they dramatically cut down on unnecessary computation. In fact, Rival is also (19.2%) faster than MPFI, another commonly-used interval arithmetic library, again due to Rival’s ability to skip unnecessary computation.

## 8 DETAILED ANALYSIS

This section expands on Section 7 to detail the performance of error intervals, movability flags, and input search individually, and to describe particularly challenging benchmarks. The data show that Rival’s extensions are effective, accomplish their task at low precision, and handle challenging inputs from a range of domains.

### 8.1 Error Intervals

Error intervals are evaluated by considering the internal precision at which they deem input points valid or invalid. Input search is disabled. Figure 3 counts valid and invalid points by the internal precision required. The vast majority of points are proven valid or invalid at 80 bits of precision, with 1280 bits the second-most common precision. This curious runner-up precision is largely due to trigonometric functions: the largest double-precision values are as large as  $2^{1024}$ , and 1280 bits is enough precision to identify their period modulo  $\pi$ . Recomputation automatically increases the precision to this level.

Compared to valid points, a larger fraction of invalid points are detected at 80 bits of precision, likely because for these points, Rival doesn't need to compute a ground truth value, just prove that the point is invalid. The majority of invalid points result from infinite outputs and domain errors, while user-specified precondition failures are rare.<sup>10</sup> Specifically among expressions with preconditions, 0.1% of sampled points are invalid, suggesting that error intervals work well even under these most challenging conditions.

The benchmark named "Harley's example"<sup>11</sup> demonstrates the importance of soundness for error and boolean intervals:

```

744     pre  $c_p > 0 \wedge c_n > 0$ 
745     let  $a = 1/(1 + e^{-\varepsilon}), b = 1/(1 + e^{-t})$  in
746      $(a^{c_p} \cdot (1 - a)^{c_p}) / (b^{c_n} \cdot (1 - b)^{c_n})$ 

```

Some inputs to this expression produce outputs just below or just above the point where the output rounds to infinity; error intervals force recomputation until it is clear whether which of the points are valid, which can require as many as 5120 bits.

## 8.2 Movability Flags

Movability flags are evaluated by considering the precision at which they can prove a point unsamplable. Figure 4 shows that movability flags detect a 81.0% of unsamplable points and thus prevent most futile recomputations. Unsamplable points are found (and warnings issued) for 21 input expressions, mostly involving ratios of exponential functions. For example, the expq2 benchmark involves the expression  $\exp(x)/(\exp(x) - 1)$ ; its unsamplable points can be detected with just 80 bits of precision. For more difficult expressions, higher precision is sometimes necessary to prove unsamplability; consider the expq3 benchmark:

$$\frac{\varepsilon \cdot (e^{(a+b) \cdot \varepsilon} - 1)}{(e^{a \cdot \varepsilon} - 1) \cdot (e^{b \cdot \varepsilon} - 1)}, \text{ with } -1 < \varepsilon < 1.$$

This benchmark is unsamplable when  $a = \varepsilon = 10^{-200}$  and  $b = 10^{200}$ : the second term in the denominator overflows, but at low precisions the first rounds down to 0, making the denominator possibly zero. Only at 2560 bits of precision<sup>12</sup> does it become clear that the denominator is very large, which is necessary to prove the expression unsamplable.

By contrast, cases where movability flags fail often seem like oversights by the benchmark developers or the users those benchmarks are derived from. For example, the regression suite's exp-w benchmark,  $e^{-w} \cdot \text{pow}(\ell, e^w)$ , fails to detect unsamplable points when  $\ell$  is negative and  $w$  large and positive,  $\text{pow}(\ell, e^w)$  raises a negative number to a possibly-infinite power. Since Rival does not track movability for error intervals, it does not detect that recomputation here is futile. But the author of this expression likely intended  $\ell$  to be close to 1, instead of negative. Similarly, movability flags fail on textbook examples of difficult-to-analyze functions like "Kahan's Monster":

```

774     pre  $y > 0$ 
775     let  $z = (|y - \sqrt{y^2 + 1}| - 1 / (y + \sqrt{y^2 + 1}))$ ,  $z_2 = z \cdot z$  in
776     if  $z_2 = 0$  then 1 else  $(\exp z_2 - 1) / z_2$ 

```

In the reals,  $z$  is exactly 0 so the program yields 1; with interval arithmetic,  $z$  is a narrow interval straddling 0 and Rival cannot rule out a division by zero in the else branch. Of course, since

<sup>10</sup>Infinite inputs are totally absent, unsurprising given the small fraction of inputs they represent.

<sup>11</sup>This benchmark appears as a case study in the original Herbie paper, suggesting its importance [42].

<sup>12</sup>Because, with a minimum double-precision exponent of -1024, the term  $e^{\varepsilon \cdot \varepsilon}$  needs at least 2048 bits to clearly differ from 1.

Benchmarks	N	Error Intervals				Movability Flags		Input Search		
		Valid	Pre.	Inf.	Error	Total	Detected	T	O	F
hamming	28	76.8%	0.0%	6.3%	16.9%	4.6%	100.0%	53.9%	20.1%	26.0%
haskell	270	67.1%	0.0%	17.3%	15.6%	0.1%	100.0%	60.5%	25.1%	14.4%
libraries	50	69.9%	0.0%	18.2%	11.9%	3.0%	32.8%	60.1%	12.2%	27.7%
mathematics	39	65.4%	2.8%	5.0%	26.9%	17.6%	88.1%	36.3%	20.8%	43.0%
numerics	38	79.5%	0.0%	15.3%	5.1%	0.7%	90.6%	41.2%	8.2%	50.6%
physics	31	51.3%	0.0%	15.9%	32.9%	6.7%	96.9%	32.9%	40.8%	26.3%
regression	13	87.3%	0.0%	5.3%	7.4%	8.3%	33.4%	53.0%	21.0%	25.9%
tutorial	3	90.1%	0.0%	9.9%	0.0%	0.0%	100.0%	84.8%	6.8%	8.4%
Herbie v1.4	481	67.3%	0.3%	14.8%	17.6%	3.0%	81.0%	54.9%	22.3%	22.8%
User Inputs	4888	63.5%	0.0%	13.5%	23.0%	5.0%	55.8%	61.1%	17.6%	21.3%
FPBench	126	79.9%	14.2%	5.9%	0.0%	2.1%	93.3%	19.8%	7.6%	72.6%

Table 1. Rival’s success on subsets and additional benchmark sets. The table shows the percentage of points that are valid or invalid for various reasons; the total fraction unsamplable points and the fraction of those detected by movability flags; and the sizes of |T|, |O|, and |F| after input search. Results for error intervals and movability flags are with input search disabled to separate the impacts of each extension.

real computation is undecidable, constructing counterexamples like this will be possible for any algorithm.

### 8.3 Input Search

Input search is evaluated by considering the fraction of the input space in the T, O, and F sets. Figure 5 shows these fractions as a percentage of the total search space, averaged over all benchmarks, at each input search iteration. Over 15 iterations, input search discards 22.8% of the total input space and proves an additional 54.9% to contain only valid points. As a result, 71.1% of sampled points are guaranteed to be valid and do not need to be rejection-sampled, and 74.6% of invalid points are avoided.

Input search helps Rival sample from benchmarks that would otherwise be infeasible to sample from. Consider  $Jmat.Real.erfi$ :<sup>13</sup>

$$\left(\frac{1}{\sqrt{\pi}}e^{x \cdot x}\right) \cdot \left(\frac{1}{|x|} + \frac{1}{2|x|^3} + \frac{3}{4|x|^5} + \frac{15}{8|x|^7}\right), \text{ with } x \geq \frac{1}{2}$$

Inputs above approximately 30 overflow, so roughly 0.15% of double-precision floating-point values are valid. But input search identifies the valid range with ease, and only 1.4% of samples are ultimately rejected.

Input search can also sometimes prove that there are no valid inputs at all. In the benchmark

$$\sqrt{\log(x) + \sin(x)} / \sin^{-1}(x + 2),$$

the term  $\sin^{-1}(x + 2)$  requires  $-3 \leq x \leq -1$ , while  $\log(x)$  requires  $x > 0$ , so the benchmark has no valid points. This benchmark was originally derived from a user submission; input search allows Rival to provide a helpful error message to the (doubtless confused) user.

## 8.4 Generality

To demonstrate the generality of error intervals, movability flags, and input search across domains, we perform a subgroup analysis. Table 1 breaks down the Herbie 1.4 benchmarks into their 8 component suites.

Rival is effective in a wide variety of circumstances. Different benchmark suites range from 51.3% (physics) to 90.1% (tutorial) valid points,<sup>14</sup> and the effectiveness of input search correspondingly varies. Movability flags are 100.0% effective in the hamming benchmark suite, which has many ratios of exponential functions, but much less effective in the regression suite, which is composed of difficult-to-analyze examples that once triggered timeouts or crashes in the Herbie project that compiled these benchmarks.

We also consider two additional benchmark sets to evaluate Rival’s generality: the 126 standard FPBench benchmarks and 4888 publically-accessible user submissions to the Herbie web demo. The FPBench suite is intended for benchmarking floating-point analysis tools and makes extensive use of preconditions. Despite this, input search samples a valid point with 93.8% probability, discarding 72.6% of the input space, and movability flags are extremely effective at detecting unsamplable points. The user-submitted benchmarks are the other extreme: unorganized, duplicative, and sometimes copied from the other two suites. Rival nonetheless preforms well—especially input search, since there are no preconditions among the user submissions so detecting valid inputs is easier. Both the subset analysis and validation sets thus show that error intervals, movability flags, and input search are effective in general across domains.

## 9 DISCUSSION

The Rival library, containing error intervals, movability flags, and input search, is publicly available as open source software.<sup>15</sup>

### 9.1 Implementing Rival

Ensuring soundness and weak completeness in Rival required care. We started by splitting each function’s domain into *monotonic regions*. For example, `fmod` has monotonic regions defined by  $ky \leq x < (k + 1)y$  for some  $k$ , monotonically increasing in  $x$  and decreasing in  $y$ .<sup>16</sup> Finding the minimum or maximum of a function then requires finding the maximum over all of the monotonic regions in the input intervals. Soundness is guaranteed by the monotonicity of the function within each region, while weak completeness is guaranteed because the output interval’s endpoints come from particular points in some monotonic region. We wrote paper proofs for each function to ensure that we accurately identified the monotonic regions, and verified that we correctly chose the maximum and minimum points using Mathematica.

To catch one-off oversights and typos, we additionally deployed millions of random tests for soundness on both wide and narrow intervals, weak completeness (by subdividing intervals), and movability (by recomputing in higher precision). These tests further improved Rival’s reliability. One random test found that `sin` internally rounded  $\pi$  to nearest instead of rounding  $\pi$  differently for the two endpoints; this is a soundness bug. Another random test found that defining `fmod(x, y)` by `trunc(x/y) · y` fails to account for the correlation between  $x/y$  and  $x$ ; this is a completeness bug fixed with a custom implementation of `fmod`. Finally, a manual audit found that movability flags weren’t

<sup>13</sup>Specifically, of the two benchmarks with that name, the one labeled “branch  $x$  greater than or equal to 5”, which we assume is mislabeled.

<sup>14</sup>The physics benchmarks also trip up input search by heavily using trigonometric functions. The tutorial suite has three very simple expressions, and its invalid points all come from a single overflow.

<sup>15</sup>URL will be available after anonymization is lifted.

<sup>16</sup>For positive  $x$ ; `fmod` is odd in its first argument.

883 being set when pow underflowed to 0, which lead to movability flags missing some unsamplable  
 884 points. We replaced the pow movability logic with the identity  $\text{pow}(x, y) = \exp(y \cdot \log(x))$  to avoid  
 885 duplicating the movability logic.

886

## 887 9.2 Deploying Rival

888 We began developing Rival three years ago to replace unsound sampling heuristics for a popular  
 889 floating point analysis and error estimation tool.<sup>17</sup> Rival was deployed to production one year ago  
 890 and has since been used by thousands of users.

891 Rival was originally a traditional interval arithmetic library aimed at soundly computing error-  
 892 free ground truths. At first, our main concern was speed: interval arithmetic generally requires  
 893 evaluating each function twice (once for each endpoint) while the existing unsound heuristics  
 894 evaluated each function once, so we expected performance to be twice as bad. These concerns were  
 895 unfounded. Instead, though “difficult” points needed higher precision in Rival (in order to soundly  
 896 prove they were correctly sampled), Rival could sample “easy” points with lower precision than the  
 897 earlier heuristics, so replacing them with Rival did not slow down sampling.

898 Instead of speed, the core problem with this initial version of Rival was its handling of invalid  
 899 inputs. Initially, Rival used NaN endpoints to indicate domain errors, much like the open source  
 900 MPFI library [43]. But those NaN endpoints would sometimes be lost in later computations, leading  
 901 to Rival erroneously computing outputs for invalid inputs. We developed error intervals to guarantee  
 902 that Rival only sampled valid points. With this problem fixed, we realized that unsamplable points  
 903 consumed an unjustifiably large share of sampling time. Movability flags were developed out of the  
 904 realization that overflow was behind most of these unsamplable points and could be easily detected  
 905 at low precision. This sped up sampling, and allowed sampling inputs for many larger expressions  
 906 that used to time out. But this only revealed that most of those sampled inputs were invalid or  
 907 unsamplable; input search was developed to address this issue.

908 All told, Rival is sounder, faster, and more user-friendly than the heuristic approach it replaced. In  
 909 rare cases, the earlier heuristics computed incorrect ground truth values, even for textbook examples  
 910 such as  $\sqrt{x+1} - \sqrt{x}$ ; Rival never does. Its input search and movability flags allow estimating error  
 911 for larger benchmarks than the earlier heuristics could support. And by rigorously handling error  
 912 cases, algorithmic limitations, and even edge cases like benchmarks with no valid inputs, Rival  
 913 allows for clearer and more actionable error messages for users. Over all, through error intervals,  
 914 movability flags, and input search, Rival successfully addressed user-visible issues and improved  
 915 the application it was embedded in.

916

## 917 10 RELATED WORK

918 *Computable Reals.* The problem of computing the exact value of a mathematical expression has  
 919 a long history. Early work [7, 8, 10, 40] explored various representations for “computable real”  
 920 values and developed algorithms for computing functions over them. One common representation  
 921 involving scaled 2-adic numbers yields particularly simple implementations [12, 41], but is slow,  
 922 especially for expressions having certain forms [33]. Interval arithmetic was later discovered to be  
 923 better suited to this task, as it tends to be faster [32].

924 Interval arithmetic has been widely available since 1967’s XSC [53]. Today, the Boost [39]  
 925 and Gaol [23] libraries are widely used, both providing interval arithmetic with double-precision  
 926 endpoints. The more modern Moore library [37] provides arbitrary-precision endpoints as well. The  
 927 IEEE 1788 standard formalizes several forms of interval arithmetic [1], including interval decorations,  
 928 which could be used to implement Rival’s error intervals and movability flags. Recent work welds  
 929

930 <sup>17</sup>Tool will be identified once anonymity is lifted.

931

computable real algorithms to decision procedures for special cases like as rational arithmetic, and provides an API for supporting the combination [9]. Despite its unique benefits for user-facing calculator applications, general-purpose error estimation generally lies outside the special cases, so Rival's methods are orthogonal. RealPaver uses interval arithmetic to characterize the solution sets of nonlinear constraints [28], to aid in pruning for global function optimization [24], similar to Rival's boolean intervals and input search. However, since RealPaver is targetted specifically at mathematical optimization problems, it neither tracks rounding error nor provides error intervals nor guarantees uniform sampling.

*Error Estimation.* There is a large literature on measuring floating-point rounding error. One approach [18, 27, 35] uses one interval to track value ranges and another interval to track rounding errors, producing a sound error bound for the whole computation. This approach can be extended to and intraprocedural, modular abstract interpretation [17, 36, 48]. Error Taylor forms can also bound floating-point error [14, 19, 46] and generally providing tighter bounds than interval-based bounds, as does semidefinite programming [34]. Both error Taylor forms and semidefinite characterize error at a fixed input point and then apply sound global optimization, such as Gelpia [11], which uses interval arithmetic to compute sound bounds.

An alternative approach to error estimation aims to generate inputs with particularly large rounding error. S3FP [15] treats the expression itself as a black box, while FPGen [26] uses symbolic execution to more efficiently target problematic input ranges. Other tools target rare errors [54–56] or attempt to maximize code coverage [4, 22]. In each case, these inputs are generated, selected, or mutated by comparing to an error-free ground-truth, usually computed with arbitrary precision or interval arithmetic. Rival's enhancements to interval arithmetic are especially attractive for input generation methodologies, since they often generate invalid inputs (which they should select against) or unsamplable inputs (which they should ignore).

*Error Detection.* Some tools aim to localize rounding error within a large program. FpDebug [6] and Herbgrind [45] use Valgrind to target x86 binaries and detect rounding error using higher-precision shadow values. PositDebug does the same but for posit, rather than floating-point, values [16]. Bao and Zhang [5] instead detect control-flow divergence and re-executing the program at higher precision. Herbie [42], an error repair tool, uses heuristics to select a sufficiently high shadow value precision; the algorithm is unsound and can compute incorrect ground truths. TAFFO replaces user-selected inputs with run-time values to dynamically substitute fixed-point for floating-point computation [13]. All of these need to compute an error-free ground truth and are susceptible to invalid or unsamplable inputs.

## 11 CONCLUSION

This paper addresses three problems common to traditional interval arithmetic implementations. Error intervals track the necessity and possibility of domain errors, and can be combined with other boolean constraints to provide a rich description of valid and invalid inputs. Movability flags add another pair of boolean flags to track the impact of overflows, identifying inputs where algorithmic limitations prevent identifying a correct ground-truth value. Finally, input search leverages both error intervals and movability flags to identify regions of valid, samplable points, even when those regions have low measure in the input space. Rival implements all three features and outperforms the state-of-the-art Mathematica 12.1.1 N function, resolving 60.3% more inputs, returning indeterminate results 10.2× less often, and avoiding all cases of fatal error.

## REFERENCES

- [1] 2015. IEEE Standard for Interval Arithmetic. *IEEE Std. 1788-2015* (2015).

- 981 [2] Micah Altman, Jeff Gill, and Michael P. McDonald. 2003. *Numerical Issues in Statistical Computing for the Social Scientist*. Springer-Verlag, 1–11 pages.
- 982 [3] Micah Altman and Michael P. McDonald. 2003. Replication with attention to numerical accuracy. *Political Analysis* 11, 3 (2003), 302–307. <http://pan.oxfordjournals.org/content/11/3/302.abstract>
- 983 [4] R. Bagnara, M. Carlier, R. Gori, and A. Gotlieb. 2013. Symbolic Path-Oriented Test Data Generation for Floating-Point Programs. In *2013 IEEE Sixth International Conference on Software Testing, Verification and Validation*. 1–10.
- 984 [5] Tao Bao and Xiangyu Zhang. 2013. On-the-fly Detection of Instability Problems in Floating-point Program Execution. *SIGPLAN Not.* 48, 10 (Oct. 2013), 817–832. <https://doi.org/10.1145/2544173.2509526>
- 985 [6] Florian Benz, Andreas Hildebrandt, and Sebastian Hack. 2012. A Dynamic Program Analysis to Find Floating-point Accuracy Problems (*PLDI '12*). ACM, New York, NY, USA, 453–462. <http://doi.acm.org/10.1145/2254064.2254118>
- 986 [7] Hans J. Boehm. 1987. Constructive Real Interpretation of Numerical Programs (*SIGPLAN '87*). ACM, 214–221. <https://doi.org/10.1145/29650.29673>
- 987 [8] Hans J. Boehm. 2004. The constructive reals as a Java Library. *Journal Logical and Algebraic Programming* 64 (2004), 3–11.
- 988 [9] Hans-J. Boehm. 2020. Towards an API for the Real Numbers. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (London, UK) (*PLDI 2020*). Association for Computing Machinery, New York, NY, USA, 562–576. <https://doi.org/10.1145/3385412.3386037>
- 989 [10] Hans J. Boehm, Robert Cartwright, Mark Riggle, and Michael J. O'Donnell. 1986. Exact Real Arithmetic: A Case Study in Higher Order Programming (*LFP '86*). ACM, 162–173. <https://doi.org/10.1145/319838.319860>
- 990 [11] Ian Briggs. 2019. Global Extrema Locator Parallelization for Interval Arithmetic. <https://github.com/soarlab/gelpia>
- 991 [12] Keith Briggs. 2006. Implementing exact real arithmetic in python, C++ and C. *Theoretical Computer Science* 351, 1 (2006), 74 – 81. <https://doi.org/10.1016/j.tcs.2005.09.058> Real Numbers and Computers.
- 992 [13] S. Cherubin, D. Cattaneo, M. Chiari, A. D. Bello, and G. Agosta. 2020. TAFFO: Tuning Assistant for Floating to Fixed Point Optimization. *IEEE Embedded Systems Letters* 12, 1 (2020), 5–8.
- 993 [14] Wei-Fan Chiang, Mark Baranowski, Ian Briggs, Alexey Solovyev, Ganesh Gopalakrishnan, and Zvonimir Rakamarić. 2017. Rigorous Floating-point Mixed-precision Tuning (*POPL*). 300–315. <https://doi.org/10.1145/3009837.3009846>
- 994 [15] Wei-Fan Chiang, Ganesh Gopalakrishnan, Zvonimir Rakamarić, and Alexey Solovyev. 2014. Efficient Search for Inputs Causing High Floating-Point Errors. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (*PPoPP '14*). Association for Computing Machinery, New York, NY, USA, 43–52. <https://doi.org/10.1145/2555243.2555265>
- 995 [16] Sangeeta Chowdhary, Jay P. Lim, and Santosh Nagarakatte. 2020. Debugging and Detecting Numerical Errors in Computation with Posits. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (London, UK) (*PLDI 2020*). Association for Computing Machinery, New York, NY, USA, 731–746. <https://doi.org/10.1145/3385412.3386004>
- 996 [17] Nasrine Damouche, Matthieu Martel, and Alexandre Chapoutot. 2015. Formal Methods for Industrial Critical Systems: 20th International Workshop, FMICS 2015 Oslo, Norway, June 22–23, 2015 Proceedings. (2015), 31–46.
- 997 [18] Eva Darulova and Viktor Kuncak. 2014. Sound Compilation of Reals (*POPL*). 14. <http://doi.acm.org/10.1145/2535838.2535874>
- 998 [19] Arnab Das, Ian Briggs, Ganesh Gopalakrishnan, Sriram Krishnamoorthy, and Pavel Panchekha. 2020. Scalable yet Rigorous Floating-Point Error Analysis. In *2020 SC20: International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. IEEE Computer Society, Los Alamitos, CA, USA, 1–14. <https://doi.org/10.1109/SC41405.2020.00055>
- 999 [20] European Commission. 1998. *The introduction of the euro and the rounding of currency amounts*. European Commission, Directorate General II Economic and Financial Affairs.
- 1000 [21] Laurent Fousse, Guillaume Hanrot, Vincent Lefèvre, Patrick Pélissier, and Paul Zimmermann. 2007. MPFR: A Multiple-Precision Binary Floating-Point Library with Correct Rounding. *ACM Trans. Math. Software* 33, 2 (June 2007), 13:1–13:15. <http://doi.acm.org/10.1145/1236463.1236468>
- 1001 [22] Zhoulai Fu and Zhendong Su. 2017. Achieving high coverage for floating-point code via unconstrained programming. *ACM SIGPLAN Notices* 52 (06 2017), 306–319. <https://doi.org/10.1145/3140587.3062383>
- 1002 [23] Frederic Goualard. 2019. Gaol: NOT Just Another Interval Library. <https://sourceforge.net/projects/gaol/>
- 1003 [24] Laurent Granvilliers and Frédéric Benhamou. 2006. Algorithm 852: RealPaver: An Interval Solver Using Constraint Satisfaction Techniques. *ACM Trans. Math. Softw.* 32, 1 (March 2006), 138–156. <https://doi.org/10.1145/1132973.1132980>
- 1004 [25] Richard Hamming. 1987. *Numerical Methods for Scientists and Engineers* (2nd ed.). Dover Publications.
- 1005 [26] Cindy Rubio-González Hui Guo. 2020. Efficient Generation of Error-Inducing Floating-Point Inputs via Symbolic Execution (*ICSE '20*).
- 1006 [27] Anastasiia Izycheva and Eva Darulova. 2017. On sound relative error bounds for floating-point arithmetic (*FMCAD*). 15–22. <https://doi.org/10.23919/FMCAD.2017.8102236>

- 1030 [28] L. Jaulin, M. Kieffer, O. Didrit, and E. Walter. 2001. *Applied Interval Analysis: With Examples in Parameter and State*  
1031 *Estimation, Robust Control and Robotics*. Springer London. [https://books.google.com/books?id=ZG0qXkYUe\\_AC](https://books.google.com/books?id=ZG0qXkYUe_AC)
- 1032 [29] William Kahan. 2000. *Miscalculating Area and Angles of a Needle-like Triangle*. Technical Report. University of  
1033 California, Berkeley. <http://www.cs.berkeley.edu/~wkahan/Triangle.pdf>
- 1034 [30] W. Kahan and Joseph D. Darcy. 1998. *How Java's Floating-Point Hurts Everyone Everywhere*. Technical Report. University  
1035 of California, Berkeley. <http://www.cs.berkeley.edu/~wkahan/JAVAhurt.pdf>
- 1036 [31] Vernon A. Lee and Hans-J. Boehm. 1990. Optimizing Programs over the Constructive Reals (*PLDI '90*). New York, NY,  
1037 USA. <https://doi.org/10.1145/93542.93558>
- 1038 [32] Vernon A. Lee, Jr. and Hans J. Boehm. 1990. Optimizing Programs over the Constructive Reals (*PLDI '90*). ACM,  
1039 102–111. <https://doi.org/10.1145/93542.93558>
- 1040 [33] Yong Li and Jun-Hai Yong. 2007. Efficient Exact Arithmetic over Constructive Reals. In *The 4th Annual Conference on*  
1041 *Theory and Applications of Models of Computation*. Shanghai, China. <https://hal.inria.fr/inria-00517598>
- 1042 [34] Victor Magron, George Constantinides, and Alastair Donaldson. 2017. Certified roundoff error bounds using semidefinite  
1043 programming. *Transactions on Mathematical Software* 43, 4 (2017), 34.
- 1044 [35] Matthieu Martel. 2009. Program Transformation for Numerical Precision (*PEPM '09*). ACM, New York, NY, USA,  
1045 101–110. <http://doi.acm.org/10.1145/1480945.1480960>
- 1046 [36] Matthieu Martel. 2018. Strongly Typed Numerical Computations. *ICFEM* (2018), 197–214.
- 1047 [37] Walter F. Mascarenhas. 2016. Moore: Interval Arithmetic in Modern C++. *CoRR* abs/1611.09567 (2016). arXiv:1611.09567  
1048 <http://arxiv.org/abs/1611.09567>
- 1049 [38] B. D. McCullough and H. D. Vinod. 1999. The Numerical Reliability of Econometric Software. *Journal of Economic*  
1050 *Literature* 37, 2 (1999), 633–665.
- 1051 [39] Guillaume Melquiond. 2006. Boost Interval Arithmetic Library. [https://www.boost.org/doc/libs/1\\_73\\_0/libs/numeric/interval/doc/interval.htm](https://www.boost.org/doc/libs/1_73_0/libs/numeric/interval/doc/interval.htm)
- 1052 [40] Marvin L. Minsky. 1967. *Computation: Finite and Infinite Machines*. Prentice-Hall, Inc.
- 1053 [41] Valérie Ménissier-Morain. 2005. Arbitrary precision real arithmetic: design and algorithms. *The Journal of Logic and*  
1054 *Algebraic Programming* 64, 1 (2005), 13 – 39. <https://doi.org/10.1016/j.jlap.2004.07.003> Practical development of exact  
1055 real number computation.
- 1056 [42] Pavel Pancheckha, Alex Sanchez-Stern, James R. Wilcox, and Zachary Tatlock. 2015. Automatically Improving Accuracy  
1057 for Floating Point Expressions (*PLDI*).
- 1058 [43] Nathalie Revol and Fabrice Rouillier. 2005. Motivations for an Arbitrary Precision Interval Arithmetic and the MPFI  
1059 Library. *Reliable Computing* 11, 4 (2005), 275–290. <https://doi.org/10.1007/s11155-005-6891-y>
- 1060 [44] Cindy Rubio-González et al. 2013. Precimonious: Tuning assistant for floating-point precision (*SC*). IEEE, 1–12.
- 1061 [45] Alex Sanchez-Stern, Pavel Pancheckha, Sorin Lerner, and Zachary Tatlock. 2018. Finding Root Causes of Floating Point  
1062 Error (*PLDI*). 256–269. <https://doi.org/10.1145/3192366.3192411>
- 1063 [46] Alexey Solovyev, Charlie Jacobsen, Zvonimir Rakamaric, and Ganesh Gopalakrishnan. 2015. Rigorous Estimation of  
1064 Floating-Point Round-off Errors with Symbolic Taylor Expansions (*FM*).
- 1065 [47] Enyi Tang, Earl Barr, Xuandong Li, and Zhendong Su. 2010. Perturbing Numerical Calculations for Statistical  
1066 Analysis of Floating-Point Program (in)Stability. In *Proceedings of the 19th International Symposium on Software*  
1067 *Testing and Analysis* (Trento, Italy) (*ISSTA '10*). Association for Computing Machinery, New York, NY, USA, 131–142.  
1068 <https://doi.org/10.1145/1831708.1831724>
- 1069 [48] Laura Titolo, Marco A Feliú, Mariano Moscato, and César A Munoz. 2018. An Abstract Interpretation Framework for  
1070 the Round-Off Error Analysis of Floating-Point Programs (*VMCAI*). 516–537.
- 1071 [49] N. Toronto and J. McCarthy. 2014. Practically Accurate Floating-Point Math. *Computing in Science Engineering* 16, 4  
1072 (July 2014), 80–95.
- 1073 [50] U.S. General Accounting Office. 1992. Patriot Missile Defense: Software Problem Led to System Failure at Dhahran,  
1074 Saudi Arabia. <http://www.gao.gov/products/IMTEC-92-26>
- 1075 [51] Debora Weber-Wulff. 1992. Rounding error changes Parliament makeup. <http://catless.ncl.ac.uk/Risks/13.37.html#subj4>
- 1076 [52] Wolfram. 2020. \$MaxExtraPrecision—Wolfram Language Documentation. <https://reference.wolfram.com/language/ref/protectT1textdollarMaxExtraPrecision.html>
- 1077 [53] WR-SWT. 2020. History of XSC-Languages and Credits. <http://www.xsc.de>
- 1078 [54] X. Yi, L. Chen, X. Mao, and T. Ji. 2017. Efficient Global Search for Inputs Triggering High Floating-Point Inaccuracies.  
In *2017 24th Asia-Pacific Software Engineering Conference (APSEC)*. 11–20.
- [55] Xin Yi, Liqian Chen, Xiaoguang Mao, and Tao Ji. 2019. Efficient Automated Repair of High Floating-Point Errors in  
Numerical Libraries. *Proc. ACM Program. Lang.* 3, POPL, Article 56 (Jan. 2019), 29 pages. <https://doi.org/10.1145/3290369>
- [56] D. Zou, R. Wang, Y. Xiong, L. Zhang, Z. Su, and H. Mei. 2015. A Genetic Algorithm for Detecting Significant Floating-  
Point Inaccuracies. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, Vol. 1. 529–539.