

Mechanizing Olver’s Error Arithmetic

Max Fan

Cornell University
Ithaca, New York, USA
mxf@cs.cornell.edu

Ariel E. Kellison

Sandia National Laboratories
Livermore, CA, USA
aekelli@sandia.gov

Samuel D. Pollard

Sandia National Laboratories
Livermore, CA, USA
spolla@sandia.gov

We mechanize the fundamental properties of a rounding error model for floating-point arithmetic based on *relative precision*, a measure of error proposed as a substitute for relative error in rounding error analysis. A key property of relative precision is that it forms a true metric, providing a well-defined measure of distance between exact results and their floating-point approximations while offering a structured approach to propagating error bounds through sequences of computations. Our mechanization formalizes this property, establishes rules for converting between relative precision and relative error, and shows that the rounding error model based on relative precision slightly over-approximates the standard rounding error model. Finally, we demonstrate, with a simple example of the inner product of two vectors, that this alternative model facilitates a tractable approach to developing certified rounding error bounds.

1 Introduction

Many scientific and engineering domains rely on algorithms designed to solve problems in continuous mathematics. While these algorithms are typically specified using real numbers, their implementations use floating-point arithmetic. As a result, these implementations deviate from their real-number specifications due to rounding and approximation errors. These deviations make it challenging to ensure the correctness of programs that rely on floating-point arithmetic, motivating the development of numerous tools for analyzing and verifying floating-point computations [1, 6, 7, 8, 18]. In high-consequence settings, practitioners often turn to formal verification techniques, including proof assistants [16], to certify rounding error estimates and ensure rigorous guarantees of correctness.

A fundamental decision in both manual rounding error analyses and formal verification is the choice of a *rounding error model*: how does the floating-point result of a basic arithmetic operation relate to its exact real-number counterpart? This decision is crucial because it directly impacts the *soundness*, *precision*, and *tractability* of the analysis. A model that underapproximates floating-point behavior may lead to *unsound guarantees*, while an overly conservative model may introduce *pessimistic* error bounds that limit its practical usefulness. More precise models improve accuracy but often increase the complexity of verification, making proofs harder to automate. Choosing the right model is therefore a key trade-off in certified error estimation.

In this work, we provide a mechanized formalization of an alternative rounding error model based on *relative precision*, first proposed by Olver [15] as a substitute for relative error in rounding error analysis. Unlike relative error, relative precision defines a true metric, providing a well-defined measure of distance between floating-point approximations and exact results. By construction, relative precision naturally supports an alternative rounding error model that succinctly and soundly captures higher-order error terms. We mechanize the fundamental properties of this model in the Rocq¹ proof assistant [19] and

¹Rocq was renamed from Coq in 2025.

establish its relationship to the standard rounding error model for elementary floating-point operations formalized in Flocq [3].

Mechanizing previously established pen-and-paper results strengthens the foundations for automated and mechanized sound analyses, leading to more reliable tools for reasoning about floating-point error. Formalization in a proof assistant not only increases confidence in correctness but also enhances extensibility, enabling future extensions and integrations with automated verification frameworks. For example, by mechanizing Olver’s work, we discovered that two fundamental properties of relative precision from the original paper did not adequately handle signs and were thus ill-defined over negative values.

More concretely, this work contributes:

- A detailed account of Olver’s original error arithmetic (Section 3), with the results proved in Rocq.
- Corrections to make two ill-defined properties of relative precision in Olver’s original error arithmetic well-defined and true (Section 3.2).
- A demonstration of the simplicity of the proposed error analysis using an inner product example in Rocq (Section 4).

We release our code at <https://github.com/Athena-Types/mechanized-precision> under an open-source license and provide hyperlinks to the Rocq code using the \diamond symbol.

2 Overview

A rounding error analysis aims to provide an *a priori* estimate of how rounding errors impact an algorithm. The typical approach models the rounding error from a single floating-point operation using a sound overapproximation of the true error, and then combines the error bounds from each individual operation to approximate the total error accumulated over a sequence of operations. We provide the necessary definitions below, but we refer the interested reader to [14, Ch. 2] for a more complete discussion on floating point, and rounding, overflow, underflow, and exceptional values such as NaN and ∞ .

The Standard Rounding Error Model A widely used model for rounding error analysis is the *standard model* of floating-point arithmetic [9], which describes how floating-point approximations relate to exact computations under the assumption that overflow, underflow, and exceptional values do not occur. The central theorem of this model, for a given rounding function $\rho : \mathbb{R} \rightarrow \mathbb{F}$, is as follows, where \mathbb{R} and \mathbb{F} are the sets of all real and floating-point numbers, respectively.

Theorem 2.1. *Given a real number $x \in \mathbb{R}$, and assuming no underflow or overflow occurs, then there exists a $\delta \in \mathbb{R}$ such that [9, Theorem 2.2]:*

$$\rho(x) = x(1 + \delta) \quad |\delta| \leq u, \quad (1)$$

The constant u in Theorem 2.1 is the unit roundoff, which depends on the floating-point format (e.g., IEEE 754 single precision or double precision) and the rounding function $\rho : \mathbb{R} \rightarrow \mathbb{F}$; the behavior of common rounding functions and their corresponding unit roundoff values for binary formats are given in Table 1. For IEEE 754 single and double precision with round-to-nearest (with arbitrary tie breaking), the unit roundoff values are $u = 2^{-24}$ and $u = 2^{-53}$, respectively [21].

The standard model for floating-point arithmetic follows from Theorem 2.1: for floating-point numbers x and y in a format $\mathbb{F} \subset \mathbb{R}$ (again, ignoring exceptional values like NaN or ∞), for real numbers a and b , the standard model expresses the floating-point approximation $\rho(a \text{ op } b) \in \mathbb{F}$ of $(a \text{ op } b) \in \mathbb{R}$ as

$$\rho(a \text{ op } b) = (a \text{ op } b)(1 + \delta), \quad |\delta| \leq u, \quad \text{op} \in \{+, -, \times, \div\}. \quad (2)$$

Table 1: The behavior of common rounding functions and their corresponding unit roundoff values (u) for binary formats. \mathbb{F} is the set of all floating-point numbers for a given format with precision p .

Rounding Function (ρ)	Behavior	Notation	Unit Roundoff (u)
Round towards $+\infty$	$\min\{b \in \mathbb{F} \mid b \geq a\}$	$\rho_{RU}(a)$	2^{1-p}
Round towards $-\infty$	$\max\{b \in \mathbb{F} \mid b \leq a\}$	$\rho_{RD}(a)$	2^{1-p}
Round towards 0	$\rho_{RU}(a)$ if $a < 0$, otherwise $\rho_{RD}(a)$	$\rho_{RZ}(a)$	2^{1-p}
Round towards nearest	$\{b \in \mathbb{F} \mid \forall c \in \mathbb{F}, a - b \leq a - c \}$	$\rho_{RN}(a)$	2^{-p}

Absolute and Relative Error The standard rounding error model given in Equation (2) provides a structured way to approximate floating-point errors. Within this framework, two fundamental measures of error are widely used: absolute error and relative error. As functions, relative error (err_{rel}) and absolute error (err_{abs}) are defined over the real numbers \mathbb{R} , where the set of representable floating-point numbers in a specified format \mathbb{F} form a discrete subset:

$$\text{err}_{rel}(x, y) \triangleq \left| \frac{x - y}{x} \right| \quad (3)$$

$$\text{err}_{abs}(x, y) \triangleq |x - y| \quad (4)$$

Absolute error and relative error each come with distinct trade-offs. Absolute error disproportionately weights discrepancies between larger values more than smaller ones. Relative error, on the other hand, behaves poorly on small values, where it can become unbounded, and also fails to satisfy the axioms of a metric on \mathbb{R} . Specifically, relative error does not satisfy the property of symmetry, as $\text{err}_{rel}(x, y) \neq \text{err}_{rel}(y, x)$, and also violates the triangle inequality, making it unsuitable for structured reasoning about the propagation of rounding errors. A measure of error that forms a true metric provides a natural way to reason about the *composition* of errors, ensuring a tractable method for propagating error bounds through sequences of computations.

Metrics and Rounding Error Bounds For instance, consider two real numbers $a, b \in \mathbb{R}$. The exact sum of a and b is given by

$$s = a + b,$$

while the floating-point computation, using Equation (2), produces

$$\tilde{s} = (a + b)(1 + \delta) = \tilde{a} + \tilde{b}; \quad |\delta| \leq u,$$

where $\tilde{a} = a(1 + \delta)$ and $\tilde{b} = b(1 + \delta)$. If $d_{\mathbb{R}} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^{\geq 0}$ is a true metric that satisfies the triangle inequality, then the total rounding error can be bounded as follows:

$$d_{\mathbb{R}}(s, \tilde{s}) \leq d_{\mathbb{R}}(a, \tilde{a}) + d_{\mathbb{R}}(b, \tilde{b}). \quad (5)$$

Here, the first term $d_{\mathbb{R}}(a, \tilde{a})$ measures the error in approximating a by \tilde{a} , and similarly, $d_{\mathbb{R}}(b, \tilde{b})$ quantifies the error in approximating b by \tilde{b} . The triangle inequality ensures that these local errors can be systematically combined to bound the total error in the sum.

This property allows for a structured approach to error analysis: if individual rounding errors can be bounded, then their composition can be rigorously bounded using the metric structure. More generally, the same principle applies to sequences of floating-point computations beyond addition, enabling a tractable method for bounding accumulated errors in numerical algorithms.

Relative Precision A metric that closely approximates relative error, known as *relative precision*, was first introduced by Olver [15]. To simplify rounding error analysis, Olver proposed the notation

$$a \sim \tilde{a}; \quad \text{rp}(\alpha) \quad (6)$$

to indicate that

$$\tilde{a}/a = e^\delta, \quad \text{where } |\delta| \leq \alpha, \quad (7)$$

and referred to \tilde{a} as an approximation of a with *relative precision* α . Notably, Equation (7) only holds whenever a and \tilde{a} are both nonzero and have the same sign. This notation naturally leads to the following definition of a *relative precision metric* on nonzero numbers of the same sign:

$$\text{err}_{\text{rp}}(\tilde{x}, x) \triangleq |\ln(\tilde{x}/x)| \quad \text{where } x, \tilde{x} \neq 0 \text{ and } \tilde{x}/x > 0. \quad (8)$$

We next describe our Rocq formalization of the relative precision metric, a rounding error model based on Equation (7), and the key properties of both the metric and the model.

3 Formalization

This section outlines the key theorems from our Rocq formalization of an alternative rounding error model for floating-point arithmetic based on the relative precision metric.

3.1 An Alternative Model for Floating-Point Arithmetic

In the spirit of Equation (7), our mechanization establishes the following alternative to Theorem 2.1.

Theorem 3.1. *Given a real number $x \in \mathbb{R}$, and a rounding function $\rho \in \{\rho_{RU}, \rho_{RD}, \rho_{RZ}, \rho_{RN}\}$ and assuming no underflow or overflow occurs, then there exists a $\delta \in \mathbb{R}$ such that*

$$\rho(x) = xe^\delta, \quad \text{where } |\delta| \leq \frac{u}{1-u}, \quad (9, \diamond)$$

where $0 < u \ll 1$ is the unit roundoff for ρ .

Notably, the rounding error bound in Theorem 3.1 is a slight overapproximation to the bound in Theorem 2.1. This becomes clearer when considering the Taylor series expansion of e^δ :

$$\rho(x) = xe^\delta = x(1 + \delta + O(\delta^2)) \quad (10)$$

Additionally, note that $u \leq u/(1-u)$. Furthermore, while Olver's original notation in Equation (7) applies only when a and \tilde{a} are both nonzero and have the same sign, Theorem 3.1 remains valid for any real number $x \in \mathbb{R}$, provided no underflow or overflow occur.

From Theorem 3.1, we can derive an alternative model for floating-point arithmetic: for floating-point numbers $x, y \in \mathbb{F}$, the alternative model expressed the floating-point approximation $\rho(x \text{ op } y) \in \mathbb{F}$ of $(x \text{ op } y) \in \mathbb{R}$ as

$$\rho(x \text{ op } y) = (x \text{ op } y)e^\delta, \quad |\delta| \leq \frac{u}{1-u}, \quad \text{op} \in \{+, -, \times, \div\}. \quad (11)$$

The key advantage of this model is that it represents rounding errors in a way that naturally aligns with the relative precision metric, enabling a true metric-based characterization of the accuracy of a

floating-point operation relative to its exact result. In the following section, we outline the key definitions and properties of relative precision, as first introduced by Olver [15], which we have formalized in Rocq. We additionally discuss two properties that we discovered to be ill-defined in the course of mechanization, and propose fixes to make the properties well-defined. In Section 4, we demonstrate how relative precision provides a convenient and mathematically sound framework for analyzing floating-point computations.

3.2 Relative Precision

Definition 3.1 (Relative Precision (\diamond)). *Let $a, \tilde{a}, \alpha \in \mathbb{R}$. We say that \tilde{a} is an approximation to a of relative precision α iff $\alpha \geq 0$, a and \tilde{a} are nonzero and have the same sign, and*

$$\text{err}_{\text{rp}}(\tilde{a}, a) = |\ln(\tilde{a}/a)| \leq \alpha.$$

Recall from Equation (6) $a \sim \tilde{a}$; $\text{rp}(\alpha)$ denotes that \tilde{a} is an approximation to a of relative precision α . To convert from relative precision to relative error, we prove the following theorem:

Theorem 3.2. *Let $a, \tilde{a}, \alpha \in \mathbb{R}$ such that $a \sim \tilde{a}$; $\text{rp}(\alpha)$. Then,*

$$\text{err}_{\text{rel}}(a, \tilde{a}) \leq |e^\alpha - 1|. \quad (12, \diamond)$$

Given the above definition of relative precision and Theorem 3.1, we formalized the following theorem in Rocq, which characterizes the relative error of basic floating-point operations in terms of relative precision.

Theorem 3.3. *Let $a, b \in \mathbb{F}$ be two floating-point numbers such that $a + b \neq 0$. Assuming overflow and underflow do not occur, then*

$$\rho(a \text{ op } b) ; \text{rp}\left(\frac{u}{1-u}\right), \quad \text{op} \in \{+, -, \times, \div\} \quad (13, \diamond)$$

for some $0 < u \ll 1$.

In Section 4, we use this result to establish overall bounds on a floating-point inner product. Deriving these bounds relies on several desirable properties of relative precision, which we have proved in our mechanization and list below. For each, suppose $a \sim \tilde{a}$; $\text{rp}(\alpha)$.

RPI. Symmetry (\diamond): $\tilde{a} \sim a$; $\text{rp}(\alpha)$

RPII. Inclusion (\diamond): for all $\delta \geq \alpha$, $a \sim \tilde{a}$; $\text{rp}(\delta)$.

RPIII. Constant scaling (\diamond): If $k \neq 0$, then $ka \sim k\tilde{a}$; $\text{rp}(\alpha)$.

RPIV. Exponentiation, modified to take absolute values — originally ill-defined (\diamond):

If $k \in \mathbb{R}$, then $|a|^k \sim |\tilde{a}|^k$; $\text{rp}(|k|\alpha)$.

RPV. Composition (\diamond): If $b \sim \tilde{b}$; $\text{rp}(\beta)$, then $ab \sim \tilde{a}\tilde{b}$; $\text{rp}(\alpha + \beta)$.

RPVI. Triangle inequality (\diamond): If $\tilde{a} \sim \tilde{\tilde{a}}$; $\text{rp}(\delta)$, then $a \sim \tilde{\tilde{a}}$; $\text{rp}(\alpha + \delta)$.

Through formalization, we also uncovered issues in Olver's property statements. In particular, observe that property RPIV. is ill-defined over the reals for negative a and small k : $a^k \sim \tilde{a}^k$; $\text{rp}(|k|\alpha)$. To address this, we modify the property to take absolute values prior to exponentiation.

In contrast, for relative error (Equation (3)), properties RPI. and RPVI. do not hold. For RPI., consider $a = 1$ and $\tilde{a} = 1.1$. Then, $\text{err}_{rel}(a, \tilde{a}) = \frac{1}{10}$. But, $\text{err}_{rel}(\tilde{a}, a) = \frac{1}{11}$. For RPVI., consider $a = 1$, $\tilde{a} = 1.5$, and $\tilde{\tilde{a}} = 2$. Then, $\text{err}_{rel}(a, \tilde{a}) = \frac{1}{2}$ and $\text{err}_{rel}(\tilde{a}, \tilde{\tilde{a}}) = \frac{1}{3}$ but $\text{err}_{rel}(a, \tilde{\tilde{a}}) = 1 > \frac{1}{2} + \frac{1}{3}$.

As a result, relative precision satisfies the axioms of a metric on $\mathbb{R}^{>0}$, whereas relative error does not. From the above properties, the following theorems describe how relative precision propagates due to addition.

Theorem 3.4. *Let $a \sim \tilde{a}$; $\text{rp}(\alpha)$ and let $b \sim \tilde{b}$; $\text{rp}(\beta)$ and assume that a and b have the same sign and overflow and underflow do not occur. Then*

$$a + b \sim \tilde{a} + \tilde{b} ; \text{rp}\left(\ln\left(\frac{\tilde{a}e^\alpha + \tilde{b}e^\beta}{\tilde{a} + \tilde{b}}\right)\right). \quad (14, \diamond)$$

A useful corollary to Theorem 3.4 is the following, which leverages the inclusion property (Property II) for relative precision to derive a somewhat looser but more practical bound on how relative precision propagates under addition.

Corollary 3.4.1. *Let $a \sim \tilde{a}$; $\text{rp}(\alpha)$ and let $b \sim \tilde{b}$; $\text{rp}(\beta)$ and assume that a and b have the same sign. Then*

$$a + b \sim \tilde{a} + \tilde{b} ; \text{rp}(\max(\alpha, \beta)). \quad (15, \diamond)$$

We next describe *absolute precision*, another convenient formalism for describing floating-point error.

3.3 Absolute Precision

We begin with the definition of absolute precision:

Definition 3.2 (Absolute Precision (\diamond)). *Let x, \tilde{x} be real numbers. If*

$$|x - \tilde{x}| \leq \alpha \quad (16)$$

for some $\alpha \geq 0$, then \tilde{x} is an approximation to x of absolute precision α .

Following Olver, for any $a, \tilde{a} \in \mathbb{R}$, we write

$$a \sim \tilde{a} ; \text{ap}(\alpha) \quad (17)$$

to denote that \tilde{a} approximates a with absolute precision α . Clearly, if $a \sim \tilde{a} ; \text{ap}(\alpha)$ then the *absolute error* (Equation (4)) of the approximation \tilde{a} with respect to a is also bounded by α .

Absolute precision enjoys several nice properties which we have proved in our mechanization. For each, suppose $a \sim \tilde{a} ; \text{ap}(\alpha)$.

API. Symmetry (\diamond): $\tilde{a} \sim a ; \text{ap}(\alpha)$

APII. Inclusion (\diamond): for all $\delta \geq \alpha$, $a \sim \tilde{a} ; \text{ap}(\delta)$.

APIII. Addition (\diamond): If $k \in \mathbb{R}$, then $a + k \sim \tilde{a} + k ; \text{rp}(\alpha)$.

APIV. Constant Scaling (\diamond): If $k \in \mathbb{R}$, then $ka \sim k\tilde{a} ; \text{ap}(|k|\alpha)$.

APV. Composition (\diamond): If $b \sim \tilde{b} ; \text{ap}(\beta)$, then $a + b \sim \tilde{a} + \tilde{b} ; \text{ap}(\alpha + \beta)$.

APVI. Triangle inequality (\diamond): If $\tilde{a} \sim \tilde{\tilde{a}} ; \text{ap}(\delta)$, then $a \sim \tilde{\tilde{a}} ; \text{ap}(\alpha + \delta)$.

As with absolute error, absolute precision forms a metric on \mathbb{R} . From the above properties, the following theorems describe how absolute precision propagates due to multiplication and division.

Theorem 3.5. (\diamond) Let $a \sim \tilde{a}$; $\text{ap}(\alpha)$, $b \sim \tilde{b}$; $\text{ap}(\beta)$. Then

$$ab \sim \tilde{a}\tilde{b} ; \text{ap}(|\tilde{a}\beta| + |\tilde{b}\alpha| + \alpha\beta).$$

Theorem 3.6. (\diamond) Let $a \sim \tilde{a}$; $\text{ap}(\alpha)$, $b \sim \tilde{b}$; $\text{ap}(\beta)$. Suppose additionally that $|\tilde{b}| > \beta$. Then

$$\frac{a}{b} \sim \frac{\tilde{a}}{\tilde{b}} ; \text{ap}\left(\frac{|\tilde{a}|\beta + |\tilde{b}|\alpha}{|\tilde{b}|(|\tilde{b}| - \beta)}\right).$$

We uncovered another issue in the absolute precision property statements: Olver additionally claims the following equation to be equivalent to Equation (6):

$$\ln a \sim \ln \tilde{a} ; \text{ap}(\alpha) \tag{18}$$

This is not quite correct for two reasons. Firstly, the equation is ill-defined for negative a and \tilde{a} . Secondly, Equation 6 uses relative precision which restricts a and \tilde{a} to be non-zero and have the same sign. By contrast, Equation 18 uses absolute precision, which is unrestricted with respect to sign. Therefore, we modify the equation to take absolute values before logarithms and only state and prove the forwards direction, namely that Equation 6 implies Equation (19) below:

$$\ln |a| \sim \ln |\tilde{a}| ; \text{ap}(\alpha). \tag{19, \diamond}$$

We now proceed to show how this error analysis can be used in practice.

4 Error Analysis of a Simple Program

With an alternative model for floating-point arithmetic defined in Section 3.1 and the relative precision of basic arithmetic operations established in Theorem 3.3, we now leverage the metric properties of relative precision to derive a sound upper bound on the total rounding error accumulated over a sequence of floating-point operations. As an example in our mechanization, we proved the following relative precision bound for the inner product of two vectors.

Theorem 4.1. (\diamond) Let $x, y \in \mathbb{F}^n$ be vectors of floating-point numbers such that $x_i * y_i > 0$ for every $i \leq n$. Let $s = \sum_{i=1}^n (x_i * y_i)$ denote the inner product of x and y computed in infinite precision, and let \tilde{s} denote the inner product computed in floating-point arithmetic. Then

$$s \sim \tilde{s} ; \text{rp}\left(\frac{nu}{1-u}\right),$$

where u is the unit roundoff.

Proof. The proof follows naturally by induction over n , using Corollary 3.4.1 and the metric properties of relative precision. These properties simplify the analysis by ensuring that errors propagate in a structured and composable manner. In particular, they allow the error of a sum to be expressed in terms of the errors of its components, making the inductive argument both natural and straightforward.

For the base case, we show that

$$s_1 \sim \tilde{s}_1 ; \text{rp}\left(\frac{u}{1-u}\right), \quad (20)$$

which follows directly from Theorem 3.3.

For the inductive step, we assume $s_k \sim \tilde{s}_k ; \text{rp}\left(\frac{ku}{1-u}\right)$ and aim to show

$$s_{k+1} \sim \tilde{s}_{k+1} ; \text{rp}\left(\frac{(k+1)u}{1-u}\right), \quad (21)$$

where s_k is the k th partial sum $s_k = x_1y_1 + \dots + x_ky_k$. By definition, we have

$$s_{k+1} = s_k + x_{k+1}y_{k+1} \quad (22)$$

and

$$\tilde{s}_{k+1} = \rho(\tilde{s}_k + \rho(x_{k+1}y_{k+1})) \quad (23)$$

where $\rho : \mathbb{R} \rightarrow \mathbb{F}$ is a well-defined rounding function.

Introducing the intermediate value

$$\hat{s}_{k+1} = \tilde{s}_k + \rho(x_{k+1}y_{k+1}), \quad (24)$$

we apply Corollary 3.4.1, Theorem 3.3, and the induction hypothesis to obtain

$$s_{k+1} \sim \hat{s}_{k+1} ; \text{rp}\left(\frac{ku}{1-u}\right). \quad (25)$$

Furthermore, by Theorem 3.3, we also have

$$\hat{s}_{k+1} \sim \tilde{s}_{k+1} ; \text{rp}\left(\frac{u}{1-u}\right). \quad (26)$$

The result follows by the triangle inequality. \square

The restriction in Theorem 4.1 that each corresponding product x_iy_i of the vectors x and y is strictly positive for all $i \leq n$ ensures that all elements of the sums s_i and \tilde{s}_i are non-zero and of the same sign, a necessary condition for using the relative precision metric as it is defined in Definition 3.1. If the elements of the vectors do not satisfy this restriction, absolute precision provides a more appropriate measure of the error.

Since relative error bounds for computing the inner product of two vectors are well established, we can use Theorem 3.2 to compare the relative precision bound derived in Theorem 4.1 with known relative error bounds. For vectors whose components satisfy the restriction that each corresponding product x_iy_i of the vectors x and y is strictly positive for all $i \leq n$, Higham [9] shows that the relative error of the inner product can be bounded as

$$\left| \frac{s - \tilde{s}}{s} \right| \leq \frac{nu}{1 - nu}. \quad (27)$$

From equation Theorem 3.2 we can express the relative error bound in terms of relative precision. Substituting the result of Theorem 4.1 into Theorem 3.2, we obtain

$$\left| \frac{s - \tilde{s}}{s} \right| \leq \frac{nu}{1 - (n+1)u}, \quad (28)$$

which provides a slightly looser bound than Equation (27) in practice.

Thus, the cost of using relative precision instead of relative error in this setting is a slightly looser bound, but the benefit is that the metric properties of relative precision provide a structured way of composing error bounds across multiple operations.

5 Proof Engineering

Our Rocq mechanization is based on real numbers and uses the rounding modes defined in the Flocq library [3] to reason about the relationships between real numbers and their floating-point approximations; these rounding modes correspond to the rounding functions in Table 1. We represent the set of real numbers corresponding to a given floating-point format using the Flocq FLX format, which features unbounded exponents. This choice aligns with the assumptions made in the theorems and definitions from previous sections that no underflows or overflows occur.

Our definitions of absolute and relative precision are formulated over real numbers as defined in the Mathematical Components library (MathComp [13]) for real analysis. We use MathComp and SSReflect for their robust tactics in ring reasoning, number theory, and algebraic rewriting. Our mechanization of Theorem 4.1 encodes the inner product as a right fold over Rocq lists of real numbers.

6 Related Work

The standard error model is in most automated and manual floating-point error analysis tools, such as VCFLOAT2 [1], FPTaylor [18], PRECiSA, Daisy [7], Colibri2 [6], Gappa [8]. VCFLOAT2 is tool developed in Rocq for verifying that C programs implement an approximation of a real-number specification. FPTaylor, Daisy, Colibri2, and Gappa provide varying levels of automation for computing floating-point error bounds. The formalization of floating-point arithmetic in Rocq, as well as analysis of various numerical algorithms, has been described by Boldo and Melquiond [4].

Relative precision has been employed in various analyses, for example, as a loss function in modeling applications—often referred to as the accuracy ratio or quotient error, along with derived functions such as the logarithmic or squared quotient error [20]—and in accuracy statistics more generally [12], where it helps mitigate certain types of model selection bias. Botchkarev provides an overview of these metrics, particularly in the context of machine learning models [5].

Static analysis tools have also leveraged the metric properties of relative precision. For example, NumFuzz [10] incorporates floating-point error analysis into its type system using relative precision, enabling precise and scalable reasoning about the numerical accuracy of programs. Bean [11] similarly employs relative precision for backward error analysis, providing a type-based framework for analyzing numerical stability. Bean extends the definition proposed by Olver to all real numbers as follows:

$$RP(x, y) = \begin{cases} |\ln(x/y)| & \text{if } \text{sign}(x) = \text{sign}(y) \text{ and } x, y \neq 0 \\ 0 & \text{if } x = y = 0 \\ \infty & \text{otherwise} \end{cases}. \quad (29)$$

In numerical analysis, relative precision has also been used to reason about the stability of algorithms under composition. For example, Beltrán et al. [2] adopt the coordinatewise relative error metric introduced by Pryce [17] in 1984, building on Olver’s original notion of relative precision.

7 Future Work & Conclusion

In this paper we presented a formalization in Rocq of an alternative to the standard model for floating-point error, first written by Olver. By formalizing this theory in Rocq, we discovered undefined scenarios in the definitions related to sign (Property RPIV. and Equation 18) and an equivalence between Equation 6

and Equation 18 that was false due to a separate issue with signs. By making these assumptions explicit in Rocq, our analysis improves upon the soundness of relative precision.

Mechanizing this error arithmetic provides opportunities for use in other tools. We plan to investigate whether this simplifies static analysis of floating-point programs. The most straightforward future work is building relative precision for more floating-point operations, such as the IEEE 754 supported by Flocq; in this case, while the error analysis would be more complex (often, handling exceptional behaviors is a large portion of the proof engineering effort), it could allow full proofs of correctness for more robust numerical programs.

Another potential application of relative precision is backwards error analysis, which can be more convenient for detecting sensitivity of input on a program’s error; it may be that relative precision simplifies this analysis. Backwards error analysis using Equation (29) can be mechanized using MathComp’s extensive extended real number support ($\mathbb{R} \cup \{\infty\}$).

Similarly, tools such as Gappa or other (relational) abstract interpretation-based tools use heuristics to better estimate ranges of inputs to provide ranges of worst-case error analysis but can suffer from scalability issues as the number of variables and floating-point operations increases; we wish to investigate whether the relative precision-based heuristics (or a combination of standard model and Olver’s model) could simplify analysis and hence improve computational efficiency.

In conclusion, this formalization motivates further research in error analysis and provides another tool in the analyst’s toolbox to more easily write safe, correct numerical programs.

Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 2139899. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Sandia National Laboratories is a multitechnology laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525.

References

- [1] Andrew Appel and Ariel E. Kellison. VCFloat2: Floating-point error analysis in coq. In *Proceedings of the 13th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2024*, page 14–29, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704888. doi: 10.1145/3636501.3636953.
- [2] Carlos Beltrán, Vanni Noferini, and Nick Vannieuwenhoven. Condition numbers and probability: The mixed and homotopy settings. *IMA Journal of Numerical Analysis*, 44(2):886–919, 2023. doi: 10.1093/imanum/drad026. URL <https://doi.org/10.1093/imanum/drad026>.
- [3] Sylvie Boldo and Guillaume Melquiond. Flocq: A unified library for proving floating-point algorithms in coq. In *IEEE 20th Symposium on Computer Arithmetic*, pages 243–252, 2011. doi: 10.1109/ARITH.2011.40. URL <https://inria.hal.science/inria-00534854v2/document>.

- [4] Sylvie Boldo and Guillaume Melquiond. *Computer Arithmetic and Formal Proofs: Verifying Floating-Point Algorithms with the Coq System*. ISTE Press - Elsevier, United Kingdom, 1st edition, November 2017. ISBN 9781785481123. doi: 10.1016/C2015-0-01301-6.
- [5] Alexei Botchkarev. A new typology design of performance metrics to measure errors in machine learning regression algorithms. *Interdisciplinary Journal of Information, Knowledge, and Management*, 14:45–76, 2019. doi: 10.28945/4184. URL <https://arxiv.org/abs/1809.03006>.
- [6] CEA-List. Colibri: Solvers for proving programs, 2021. Available at <https://colibri.frama-c.com/index.html>.
- [7] Eva Darulova, Anastasiia Izycheva, Fariha Nasir, Fabian Ritter, Heiko Becker, and Robert Bastian. Daisy - framework for analysis and optimization of numerical programs (tool paper). In Dirk Beyer and Marieke Huisman, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 270–287, Thessaloniki, Greece, 2018. Springer International Publishing. ISBN 978-3-319-89960-2. doi: 10.1007/978-3-319-89960-2_15.
- [8] Florent de Dinechin, Christoph Lauter, and Guillaume Melquiond. Certifying the floating-point implementation of an elementary function using gappa. *IEEE Transactions on Computers*, 60(2): 242–253, 2011. doi: 10.1109/TC.2010.128.
- [9] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms, second edition*. SIAM, 2002. ISBN 978-0-898715-21-7. doi: 10.1137/1.9780898718027.
- [10] Ariel E. Kellison and Justin Hsu. Numerical fuzz: A type system for rounding error analysis. *Proc. ACM Program. Lang.*, 8(PLDI), June 2024. doi: 10.1145/3656456. URL <https://doi.org/10.1145/3656456>.
- [11] Ariel E. Kellison, Laura Zielinski, David Bindel, and Justin Hsu. Bean: A language for backward error analysis. arXiv, 2025. Available at <https://arxiv.org/abs/2501.14550v1>.
- [12] B. A. Kitchenham, L. M. Pickard, S. G. MacDonell, and M. J. Shepperd. What accuracy statistics really measure. *IEE Proceedings - Software*, 148, 2001. URL https://www.researchgate.net/publication/3422071_What_accuracy_statistics_really_measure.
- [13] Assia Mahboubi and Enrico Tassil. *Mathematical Components*. Inria-Microsoft Research Joint Center, 2022. doi: 10.5281/zenodo.3999478.
- [14] Jean-Michel Muller, Nicolas Brisebarre, Florent De Dinechin, Claude-Pierre Jeannerod, Vincent Lefèvre, Guillaume Melquiond, Nathalie Revol, and Serge Torres. *Handbook of Floating-Point Arithmetic*. Springer International Publishing AG, second edition, 2018. doi: 10.1007/978-3-319-76526-6.
- [15] F. W. J. Olver. A new approach to error arithmetic. *Journal on Numerical Analysis*, 15(2):368–393, 1978. doi: 10.1137/0715024.
- [16] Samuel D. Pollard, Jon M. Aytac, Ariel Kellison, Ignacio Laguna, Srinivas Nedunuri, Sabrina Reis, Matthew J. Sottile, and Heidi K. Thornquist. The first tri-lab workshop on formal verification: Capabilities, challenges, research opportunities, and exemplars. Technical report, Sandia National Laboratories, February 2024. URL <https://www.osti.gov/biblio/2430026>.
- [17] J. D. Pryce. A new measure of relative error for vectors. *SIAM Journal on Numerical Analysis*, 21(1):202–215, 1984. doi: 10.1137/0721015.

- [18] Alexey Solovyev, Marek S. Baranowski, Ian Briggs, Charles Jacobsen, Zvonimir Rakamarić, and Ganesh Gopalakrishnan. Rigorous estimation of floating-point round-off errors with symbolic Taylor expansions. *ACM Transactions on Programming Languages and Systems*, 41(1), December 2018. ISSN 0164-0925. doi: 10.1145/3230733.
- [19] The Coq Development Team. The Coq reference manual – release 8.20.0, 2024. URL <https://coq.inria.fr/doc/V8.20.0/refman/>.
- [20] Chris Tofallis. A better measure of relative prediction accuracy for model selection and model estimation. *Journal of the Operational Research Society*, 66:1352–1362, 2015. doi: 10.1057/jors.2014.103.
- [21] Working Group for Floating-Point Arithmetic. Ieee standard for floating-point arithmetic. *IEEE Std 754-2019 (Revision of IEEE 754-2008)*, pages 1–84, 2019. doi: 10.1109/IEEESTD.2019.8766229.