



# One Polynomial Approximation to Produce Correctly Rounded Results of an Elementary Function for Multiple Representations and Rounding Modes

JAY P. LIM, Rutgers University, United States

SANTOSH NAGARAKATTE, Rutgers University, United States

Mainstream math libraries for floating point (FP) do not produce correctly rounded results for all inputs. In contrast, CR-LIBM and RLIBM provide correctly rounded implementations for a specific FP representation with one rounding mode. Using such libraries for a representation with a new rounding mode or with different precision will result in wrong results due to double rounding. This paper proposes a novel method to generate a single polynomial approximation that produces correctly rounded results for all inputs for multiple rounding modes and multiple precision configurations. To generate a correctly rounded library for  $n$ -bits, our key idea is to generate a polynomial approximation for a representation with  $n + 2$ -bits using the *round-to-odd* mode. We prove that the resulting polynomial approximation will produce correctly rounded results for all five rounding modes in the standard and for multiple representations with  $k$ -bits such that  $|E| + 1 < k \leq n$ , where  $|E|$  is the number of exponent bits in the representation. Similar to our prior work in the RLIBM project, we approximate the correctly rounded result when we generate the library with  $n + 2$ -bits using the round-to-odd mode. We also generate polynomial approximations by structuring it as a linear programming problem but propose enhancements to polynomial generation to handle the round-to-odd mode. Our prototype is the first 32-bit float library that produces correctly rounded results with all rounding modes in the IEEE standard for all inputs with a single polynomial approximation. It also produces correctly rounded results for any FP configuration ranging from 10-bits to 32-bits while also being faster than mainstream libraries.

CCS Concepts: • **Mathematics of computing** → **Mathematical software**; **Linear programming**; • **Theory of computation** → **Numeric approximation algorithms**.

Additional Key Words and Phrases: floating point, round-to-odd, correctly rounded math libraries

## ACM Reference Format:

Jay P. Lim and Santosh Nagarakatte. 2022. One Polynomial Approximation to Produce Correctly Rounded Results of an Elementary Function for Multiple Representations and Rounding Modes. *Proc. ACM Program. Lang.* 6, POPL, Article 3 (January 2022), 28 pages. <https://doi.org/10.1145/3498664>

## 1 INTRODUCTION

The floating point (FP) representation is widely used to approximate real numbers. The two main attributes of the FP representation are its dynamic range (*i.e.*, the range of values that can be represented) and precision (*i.e.*, the accuracy of each value represented). As some real numbers cannot be accurately represented in the FP representation, they need to be rounded to the nearest result according to the rounding mode. Further, FP performance is important in various domains ranging from scientific computing to machine learning. Hence, modern accelerators, processors,

---

Authors' addresses: Jay P. Lim, Department of Computer Science, Rutgers University, United States, [jpl169@cs.rutgers.edu](mailto:jpl169@cs.rutgers.edu); Santosh Nagarakatte, Department of Computer Science, Rutgers University, United States, [santosh.nagarakatte@cs.rutgers.edu](mailto:santosh.nagarakatte@cs.rutgers.edu).

---



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2022 Copyright held by the owner/author(s).

2475-1421/2022/1-ART3

<https://doi.org/10.1145/3498664>

and systems have explored new variants of the FP representation that vary the dynamic range and the precision. Intel’s Bfloat16 [Kalamkar et al. 2019] and FlexPoint [Köster et al. 2017], Nvidia’s TensorFloat32 [NVIDIA 2020], Microsoft’s MSFP [Rouhani et al. 2020], Facebook’s Log Number System [Johnson 2018], and Posits [Gustafson and Yonemoto 2017] are examples of such recent FP variants. All these representations need math libraries that provide approximations for various elementary functions (e.g.,  $\ln(x)$ ,  $e^x$ , ...) [Muller 2005].

An output of an elementary function is a correctly rounded result if it matches the result that is computed with infinite precision and then rounded to the target FP representation. Correctly rounded elementary functions improve the portability and reproducibility of applications. Unfortunately, mainstream FP libraries do not produce correct results for all inputs. Libraries like CR-LIBM [Daramy et al. 2003; Daramy-Loirat et al. 2006] and RLIBM [Lim et al. 2020, 2021; Lim and Nagarakatte 2021a,b] provide correctly rounded implementations for some FP representations. CR-LIBM provides correctly rounded elementary functions for the double type with a given rounding mode. As part of our RLIBM project, we have developed correctly rounded libraries for the float, bfloat16, and posit types with the round-to-nearest-ties-to-even mode.

Beyond the round-to-nearest-ties-to-even mode, there are four other rounding modes in the IEEE-754 standard. When existing correctly rounded libraries are used to generate results with other rounding modes or with other precision configurations, they can produce wrong results due to double rounding. For example, let us say we use a correctly rounded double precision library such as CR-LIBM with the round-to-nearest-ties-to-even mode to produce results for a 32-bit float type with the same rounding mode. Here, we round the result of CR-LIBM to a 32-bit float value to produce the final result. If the real value of  $f(x)$  is extremely close to the rounding boundary of two adjacent float values, then the error caused by rounding  $f(x)$  to double using the given rounding mode can be significant enough to produce the wrong result for a 32-bit float.

With existing approaches such as RLIBM and CR-LIBM, one will have to generate a new polynomial approximation for each such rounding mode and each precision configuration. Although feasible, developing efficient approximations require significant effort. Even after decades of effort, there are no efficient and correctly rounded implementations for all rounding modes in the IEEE standard even for the widely used 32-bit float type!

**This paper.** Rather than generating a correctly rounded elementary function for each individual representation and rounding mode, it would be ideal if we could generate one polynomial approximation that produces correct results for multiple rounding modes and many precision configurations. This paper proposes a novel approach to create such polynomial approximations! Our key idea is to create polynomial approximations that approximate the correctly rounded result of an elementary function  $f(x)$  with the *round-to-odd* rounding mode (i.e., the real value of  $f(x)$  is rounded with the round-to-odd mode). The round-to-odd is a non-standard rounding mode that can be described as follows. If the real value is exactly representable in the target representation, it is unchanged. Otherwise, it is rounded to the nearest value in the target representation whose bit-string when interpreted as an unsigned integer is an odd number.

The round-to-odd mode has been previously used to address double rounding issues in the context of binary FP to decimal FP conversion [Goldberg 1991] and also while performing primitive arithmetic operations with extended precision (e.g., Intel’s 80-bit floating point) and subsequently rounding the result back to a lower precision (i.e., a float or a double type) [Boldo and Melquiond 2005, 2008].

This paper makes a case for using the round-to-odd mode to generate correctly rounded results for elementary functions. To the best of our knowledge, no prior method for approximating elementary functions has used the round-to-odd mode. Further, the usage of the round-to-odd mode for approximating elementary functions is feasible because we approximate the correctly

rounded result in the RLIBM project. We discover that the round-to-odd mode has properties that enable the generation of correctly rounded elementary functions for multiple rounding modes and multiple precision configurations (Section 4.1).

If the goal is to produce correctly rounded elementary functions for multiple rounding modes of a representation with  $n$ -bits (*i.e.*,  $\mathbb{T}_n$ ), we propose to generate polynomial approximations that produce the correctly rounded results of the representation with  $n + 2$ -bits ( $\mathbb{T}_{n+2}$ ) using the round-to-odd mode. We prove that this polynomial approximation for  $\mathbb{T}_{n+2}$  produces correctly rounded results for all representations with  $k$ -bits (*i.e.*,  $\mathbb{T}_k$ ), where  $k$  is smaller than or equal to  $n$  and for all rounding modes in the IEEE standard. The only requirement in our proof is that the number of exponent bits in both the representation used to create the approximation (*i.e.*,  $\mathbb{T}_{n+2}$ ) and the target representation (*i.e.*,  $\mathbb{T}_k$ ) is identical. In summary, our approach that creates the math library with the round-to-odd mode for a configuration with two additional bits produces correctly rounded results for all target representations with any standard rounding mode.

The next task is to generate polynomial approximations using the round-to-odd mode for a representation with  $n + 2$ -bits (*i.e.*,  $\mathbb{T}_{n+2}$ ). We extend our prior work in the RLIBM project to generate polynomial approximations with the round-to-odd mode. Specifically, we approximate the correctly rounded result of an elementary function  $f(x)$  with the round-to-odd mode rather than the real value of  $f(x)$ . Subsequently, we identify an interval of real values around the correctly rounded result with the round-to-odd mode such that any value in the interval rounds to the correctly rounded result, which we call the odd interval. We show that any real value in the odd interval that rounds to the round-to-odd result in  $\mathbb{T}_{n+2}$  will subsequently round to the correctly rounded result for any representation  $\mathbb{T}_k$ .

One challenge in generating polynomial approximations using the odd intervals is the presence of singleton values with the round-to-odd mode. When the real value is exactly representable in  $\mathbb{T}_{n+2}$  and the bit-string of that value in  $\mathbb{T}_{n+2}$  is even when interpreted as an unsigned integer, the odd interval for that value will be a singleton element. We use the mathematical properties of the elementary function to identify such singletons. Concretely, we identify the inputs of an elementary function that has rational outputs (because the round-to-odd result is exactly representable in  $\mathbb{T}_{n+2}$ ). We subsequently develop efficient table-lookups to produce round-to-odd results for such inputs (Section 4.3).

Once we identify the odd interval for every input and the singleton odd intervals among them, we can subsequently use the RLIBM approach to generate polynomial approximations. Each non-singleton interval imposes a constraint on the output of the generated polynomial for a given input. Similar to our prior work in the RLIBM project [Lim 2021], we structure the problem of generating a polynomial approximation that satisfies the odd interval for each input as a linear programming problem. We use efficient range reduction and output compensation functions while accounting for numerical errors with them. To account for numerical errors, we further constrain the odd intervals given to the LP formulation. We also employ counterexample guided polynomial generation and generate piecewise polynomials for efficiency.

Our prototype, RLIBM-ALL, is open-source and publicly available [Lim and Nagarakatte 2021c]. It contains the polynomial generator for ten elementary functions and includes efficient polynomial approximations for them. The resulting polynomial generated by our prototype for the FP representation with 34-bits (*i.e.*,  $\mathbb{T}_{34}$ ) using the round-to-odd mode produces correctly rounded results for all FP representations ranging from 10-bits (*i.e.*,  $\mathbb{T}_{10}$ ) to 32-bits (*i.e.*,  $\mathbb{T}_{32}$ ) and for all rounding modes. It includes bfloat16 and tensorflow32 representations. RLIBM-ALL is the first math library that produces correctly rounded results for 32-bit floats for all rounding modes with a single polynomial approximation. Our implementations are faster than mainstream libraries for 32-bit floats while producing correctly rounded results for all inputs and for all rounding modes.






Fig. 2. When a real value  $v_{\mathbb{R}}$  is not exactly representable in  $\mathbb{T}$ , then  $v_{\mathbb{R}}$  is rounded to one of the two adjacent values  $v_{sm}, v_{lg} \in \mathbb{T}$  depending on the rounding mode. We show the range of real values ( $v_{\mathbb{R}}$ ) that round to  $v_{sm}$  (blue box) and  $v_{lg}$  (green box). (a) The *rn* mode when  $v_{sm}$  is even. (b) The *rn* mode when  $v_{lg}$  is even. (c) The *ra* mode when  $v_{\mathbb{R}} < 0$ . (d) The *ra* mode when  $v_{\mathbb{R}} > 0$ . (e) The *rz* mode when  $v_{\mathbb{R}} < 0$ . (f) The *rz* mode when  $x > 0$ . (g) The *ru* mode. (h) The *rd* mode.

**The round-to-nearest-ties-to-even (*rn*) mode.** This rounding mode rounds  $v_{\mathbb{R}}$  to a FP value that is closer to  $v_{\mathbb{R}}$  among  $v_{sm}$  and  $v_{lg}$ . If  $|v_{sm} - v_{\mathbb{R}}| < |v_{lg} - v_{\mathbb{R}}|$ , then  $v_{\mathbb{R}}$  rounds to  $v_{sm}$ . If  $|v_{sm} - v_{\mathbb{R}}| > |v_{lg} - v_{\mathbb{R}}|$ , then  $v_{\mathbb{R}}$  rounds to  $v_{lg}$ . If  $v_{\mathbb{R}}$  is exactly in the middle of  $v_{sm}$  and  $v_{lg}$  (i.e.,  $v_{\mathbb{R}} = \frac{v_{sm} + v_{lg}}{2}$ ),  $v_{\mathbb{R}}$  is rounded to a value whose bit-string is even when interpreted as an unsigned integer. The *rn* mode is the most commonly used rounding mode. Figure 2(a) and (b) illustrate rounding with the *rn* mode depending on whether  $v_{sm}$  is even or odd, respectively.

**The round-to-nearest-ties-to-away (*ra*) mode.** This rounding mode also rounds  $v_{\mathbb{R}}$  to a FP value that is closer to  $v_{\mathbb{R}}$  among  $v_{sm}$  and  $v_{lg}$ . If  $|v_{sm} - v_{\mathbb{R}}| < |v_{lg} - v_{\mathbb{R}}|$ , then  $v_{\mathbb{R}}$  rounds to  $v_{sm}$ . If  $|v_{sm} - v_{\mathbb{R}}| > |v_{lg} - v_{\mathbb{R}}|$ , then  $v_{\mathbb{R}}$  rounds to  $v_{lg}$ . When  $v_{\mathbb{R}}$  is exactly in the middle of  $v_{sm}$  and  $v_{lg}$ ,  $v_{\mathbb{R}}$  is rounded to a value that is farther away from 0. Specifically,  $v_{\mathbb{R}}$  rounds to  $v_{lg}$  if  $v_{\mathbb{R}} > 0$  because  $0 \leq v_{sm} < v_{\mathbb{R}} < v_{lg} \leq \infty$ . Similarly,  $v_{\mathbb{R}}$  rounds to  $v_{sm}$  if  $v_{\mathbb{R}} < 0$  because  $-\infty \leq v_{sm} < v_{\mathbb{R}} < v_{lg} \leq 0$ . Figure 2(c) and (d) illustrate rounding with *ra* mode depending on whether  $v_{\mathbb{R}} < 0$  or  $v_{\mathbb{R}} > 0$ , respectively.

**The round-towards-zero (*rz*) mode.** In this mode,  $v_{\mathbb{R}}$  is rounded to a value that is closer to 0. Here,  $v_{\mathbb{R}}$  is rounded to  $v_{sm}$  if  $v_{\mathbb{R}} > 0$  and  $v_{\mathbb{R}}$  is rounded to  $v_{lg}$  if  $v_{\mathbb{R}} < 0$ . The *rz* mode is equivalent to truncating the fraction bits of  $v_{\mathbb{R}}$  that cannot fit within the mantissa bits of the representation. Figure 2(e) and (f) illustrate the *rz* mode depending on whether  $v_{\mathbb{R}} < 0$  or  $v_{\mathbb{R}} > 0$ , respectively.

**The round-towards-positive-infinity (*ru*) mode.** This mode always rounds  $v_{\mathbb{R}}$  to the larger value  $v_{lg}$ , which is the value that is closer to  $+\infty$ . This mode is also known as rounding up. Figure 2(g) pictorially shows the *ru* mode.

**The round-towards-negative-infinity (*rd*) mode.** The round-towards-negative-infinity (*rd*) mode always rounds  $v_{\mathbb{R}}$  to the smaller value  $v_{sm}$  (i.e., a value that is closer to  $-\infty$ ). This mode is also known as rounding down. Figure 2(h) demonstrates the *rd* mode.

### 2.3 A Systematic Method for Rounding

We describe a systematic procedure for rounding a real number, which will be useful later for understanding our proofs. As we described above, we need to identify the two values  $v_{sm}$  and  $v_{lg}$  in the  $\mathbb{F}_{n,|E|}$  representation that are adjacent to  $v_{\mathbb{R}}$  and then decide between  $v_{sm}$  or  $v_{lg}$ . We will identify four pieces of information (*s*,  $v^-$ , *rb*, *sticky*) from the real value  $v_{\mathbb{R}}$  that will be sufficient to correctly round according to the various rounding modes. We call them rounding components. The first component, *s* represents the sign (-1 or 1) and identifies whether  $v_{\mathbb{R}}$  is positive or negative.




Fig. 3. The values of the rounding bit and the sticky bit for various real values between  $v^-$  and  $v^+$ . The gray box indicates the range of real values where both  $rb$  and  $sticky$  are 1. The purple box indicates the range of real values where  $rb = 0$  and  $sticky = 1$ .

The smaller of  $v_{sm}$  or  $v_{lg}$  in magnitude is represented by  $v^-$ . The components  $rb$  and  $sticky$  encode information about whether  $v_{\mathbb{R}}$  is in the middle, closer to  $v_{sm}$ , or closer to  $v_{lg}$ .

To identify the rounding components, we represent  $v_{\mathbb{R}}$  in the FP representation with an infinite number of mantissa bits while having the same number of exponents bits as  $\mathbb{F}_{n,|E|}$ . We call this representation extended infinite precision representation (*i.e.*,  $\mathbb{F}_{\infty,|E|}$ ). Effectively, this extended precision representation is similar to the  $\mathbb{F}_{n,|E|}$  representation but has a large number of bits for the mantissa. When  $v_{\mathbb{R}}$  is larger than the dynamic range of  $\mathbb{F}_{n,|E|}$ , we represent it with the largest representable value in  $\mathbb{F}_{\infty,|E|}$  (*i.e.*, the exponents bits correspond to the largest normal value in  $\mathbb{F}_{n,|E|}$  and all the mantissa bits are ones). We cannot use  $\infty$  to represent  $|v_{\mathbb{R}}|$  because we need to make a clear distinction between a real number and  $\infty$ . The round-towards-zero mode never rounds a real value to  $\infty$ . Similarly, the round-towards-positive-infinity mode does not round negative real values to  $-\infty$  and the round-towards-negative-infinity mode does not round positive real values to  $\infty$ .

$$B_{v_{\mathbb{R}}} = b_1 b_2 b_3 \dots b_n b_{n+1} b_{n+2} \dots \quad (1)$$

Here,  $b_1$  is the sign bit and bits  $b_2 \dots b_{|E|+1}$  represent the exponent bits. The rest of the bits starting from  $b_{|E|+2} \dots$  represent the mantissa.

**Identify rounding components.** To identify  $v_{sm}$  and  $v_{lg}$  that  $v_{\mathbb{R}}$  can round to, we identify two positive values  $v^-$  and  $v^+$  adjacent to  $|v_{\mathbb{R}}|$  (*i.e.*, the magnitude of  $v_{\mathbb{R}}$ ) in  $\mathbb{F}_{n,|E|}$ . Here,  $v^-$  represents the largest value that is smaller than or equal to  $|v_{\mathbb{R}}|$  and  $v^+$  represents the smallest value larger than  $|v_{\mathbb{R}}|$ . To identify  $v^-$ , we truncate  $B_{|v_{\mathbb{R}}|}$  to  $n$  bits,

$$B_{v^-} = 0b_2 b_3 b_4 \dots b_{n-1} b_n$$

Note that the sign bit is 0 because we are just considering the magnitude for  $v^-$ . We call  $v^-$  the truncated value, which is a rounding component. Then, the succeeding value of  $v^-$  in  $\mathbb{F}_{n,|E|}$  is  $v^+$ , which is obtained by adding 1 to  $v^-$ . We maintain the invariant:  $v^- \leq |v_{\mathbb{R}}| < v^+$ . In the context of rounding  $v_{\mathbb{R}}$  to  $\mathbb{F}_{n,|E|}$ ,  $v^-$  and  $v^+$  satisfy the following property,

$$\begin{cases} -v^+ < v_{\mathbb{R}} \leq -v^- & \text{if } v_{\mathbb{R}} < 0 \text{ (} s = -1 \text{)} \\ v^- \leq v_{\mathbb{R}} < v^+ & \text{if } v_{\mathbb{R}} \geq 0 \text{ (} s = 1 \text{)} \end{cases}$$

Once we identify  $s$ ,  $v^-$ , and  $v^+$ , we can compute  $v_{sm}$  and  $v_{lg}$  as follows. If  $v_{\mathbb{R}}$  is exactly representable in  $\mathbb{F}_{n,|E|}$ , then  $|v_{\mathbb{R}}| = v^-$ . Hence,  $v_{sm} = v_{lg} = s \times v^-$ . If  $v_{\mathbb{R}}$  is not exactly representable in  $\mathbb{F}_{n,|E|}$ , then it is guaranteed that  $v^- < |v_{\mathbb{R}}| < v^+$ . Thus,  $v_{sm} = v^-$  and  $v_{lg} = v^+$  if  $s = 1$  (*i.e.*,  $v_{\mathbb{R}} \geq 0$ ). Otherwise, when  $v_{\mathbb{R}}$  is negative (*i.e.*,  $s = -1$ ) then  $v_{sm} = -v^+$  and  $v_{lg} = -v^-$ .

**Rounding bit.** To determine the rounding decision for the *rn* and *ra* mode, we must determine whether  $v_{\mathbb{R}}$  is closer to  $s \times v^-$ , closer to  $s \times v^+$ , or exactly in the middle of the two values. We extract

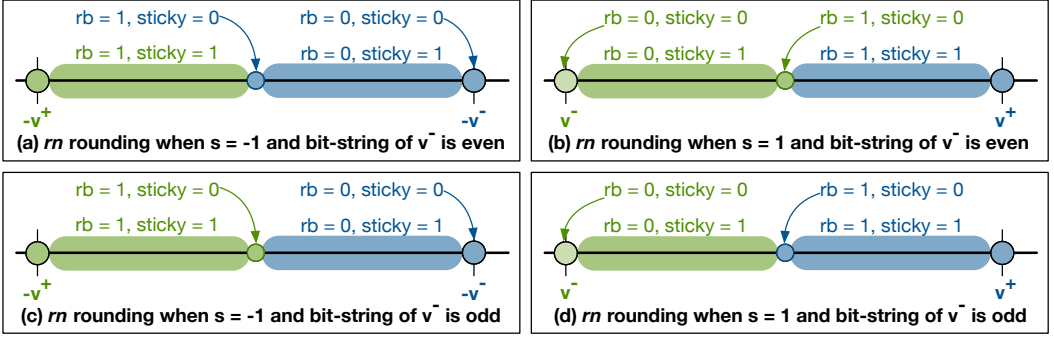


Fig. 4. The  $rn$  mode using the rounding components ( $s, v^-, rb, sticky$ ). We illustrate rounding when  $v_{\mathbb{R}}$  is positive or negative and when  $v^-$  is even or odd. When an interval is colored green, all real values in the interval round to the FP value colored green. Similarly, all real values in the interval colored blue will round to the FP value colored blue.

the  $(n + 1)^{th}$ -bit from our extended precision representation of  $B_{|v_{\mathbb{R}}|}$ , which we call as the rounding bit (*i.e.*,  $rb$ ). The rounding bit describes whether  $|v_{\mathbb{R}}|$  is closer to  $v^-$  than  $v^+$ . If the rounding bit is 0, then  $|v_{\mathbb{R}}|$  is closer to  $v^-$  (*i.e.*,  $v^- \leq |v_{\mathbb{R}}| < \frac{v^- + v^+}{2}$ ). If the rounding bit is 1, then  $|v_{\mathbb{R}}|$  is at the middle or close to  $v^+$ . Figure 3 illustrates the range of real values where the rounding bit is 0 or 1.

**Sticky bit.** While the rounding bit tells us whether  $|v_{\mathbb{R}}|$  is closer to  $v^-$ , it does not tell us whether  $|v_{\mathbb{R}}|$  is exactly equal to  $v^-$  or is exactly in the middle of  $v^-$  and  $v^+$  (*i.e.*,  $|v_{\mathbb{R}}| = \frac{v^- + v^+}{2}$ ). When we look at the bit-string  $B_{|v_{\mathbb{R}}|}$ ,  $|v_{\mathbb{R}}|$  is equal to  $v^-$  when the  $(n + 1)^{th}$ -bit (*i.e.*,  $rb$ ) is 0 and the remaining bits from the  $(n + 2)^{th}$ -bit are all zeros. If  $rb = 0$  and any bit afterwards is 1, then  $|v_{\mathbb{R}}|$  is not equal to  $v^-$ . Similarly,  $|v_{\mathbb{R}}|$  is exactly in the middle of  $v^-$  and  $v^+$  when the  $(n + 1)^{th}$ -bit is 1 and the remaining bits from the  $(n + 2)^{th}$ -bit are all 0's in  $B_{|v_{\mathbb{R}}|}$ . In both these cases, we need to determine if all the bits starting from  $b_{n+2}$  are zeros. We define the sticky bit as the bitwise OR of all bits starting from the  $(n + 2)^{th}$ -bit in the extended precision representation.

$$sticky = b_{n+2} | b_{n+3} | b_{n+4} | \dots$$

where  $|$  is the bit-wise OR operation.

Using the rounding components ( $s, v^-, rb, sticky$ ), we can identify the relationship between  $|v_{\mathbb{R}}|$  and the nearest FP values for any rounding mode in the standard.

$$\begin{cases} |v_{\mathbb{R}}| = v^- & \text{if } rb = 0 \wedge sticky = 0 \\ v^- < |v_{\mathbb{R}}| < \frac{v^- + v^+}{2} & \text{if } rb = 0 \wedge sticky = 1 \\ |v_{\mathbb{R}}| = \frac{v^- + v^+}{2} & \text{if } rb = 1 \wedge sticky = 0 \\ \frac{v^- + v^+}{2} < |v_{\mathbb{R}}| < v^+ & \text{if } rb = 1 \wedge sticky = 1 \end{cases}$$

We can compute  $v^+$  from the rounding components. Figure 3 pictorially shows the rounding bit and sticky bit for various real values between  $v^-$  and  $v^+$ .

**Rounding to various modes with the rounding components.** Figure 4 shows rounding  $v_{\mathbb{R}}$  using the rounding components for the  $rn$  mode. Similarly, Figure 5 illustrates rounding with the rounding components for the other four rounding modes.

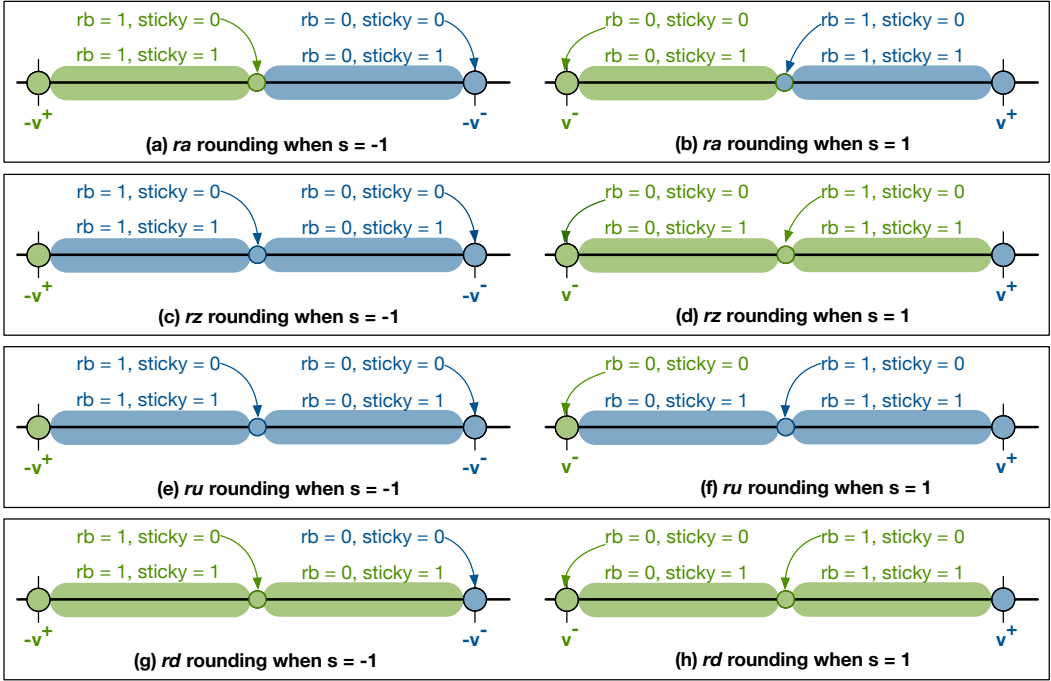


Fig. 5. Rounding decisions for various rounding modes based on the rounding components:  $(s, v^-, rb, sticky)$ . The interval of real values colored with green and blue round to the FP value colored green and blue, respectively.

## 2.4 The RLIBM Approach

We provide a brief background on our RLIBM project as we build on top of it in this paper. In the RLIBM project, we make a case for approximating the correctly rounded result rather than the real value of an elementary function [Lim et al. 2020, 2021; Lim and Nagarakatte 2021a,b]. When we approximate the correctly rounded result, there is an interval of real values around the correctly rounded result for each input such that producing any value in the interval produces the correct result. This interval is further constrained to account for numerical errors that can occur with polynomial evaluation, range reduction, and output compensation. This interval can be used to generate polynomial approximations. It represents the maximum amount of freedom available to produce the correct result. Figure 6 illustrates our RLIBM approach.

The RLIBM approach consists of four steps. The first step is to use an oracle to compute the correctly rounded result of an elementary function  $f(x)$  for each input  $x \in \mathbb{T}$ , where  $\mathbb{T}$  is the target representation. The second step is to identify an interval  $[l, h]$  around the correctly rounded result such that any value in  $[l, h]$  rounds to the correctly rounded result in  $\mathbb{T}$ , which is known as the *rounding interval*. Since polynomial evaluation, range reduction, and output compensation happen in representation with higher precision  $\mathbb{H}$ , the rounding intervals are also in  $\mathbb{H}$ . The third step is to employ range reduction to transform input  $x$  to  $x'$ . The generated polynomial will approximate the result for  $x'$ . Subsequently, we use output compensation to produce the final correctly rounded output for  $x$ . Both range reduction and output compensation happen in  $\mathbb{H}$  and can experience numerical errors. These numerical errors should not affect the generation of correctly rounded results. Hence, it is necessary to deduce intervals for the reduced domain so that the polynomial






Fig. 6. Illustration of the RLibm approach. (a) The values  $v_1$ ,  $v_2$ , and  $v_3$  are representable values in representation  $\mathbb{T}$ . The real value of  $f(x)$  for a given input  $x$  cannot be exactly represented in  $\mathbb{T}$  and it is rounded to  $v_2$ . The RLibm approach identifies the rounding interval of  $v_2$  (shown in gray box). (b) Polynomial generation using the rounding interval (*i.e.*,  $[l, h]$ ) for each input  $x$  with an LP formulation.

evaluation over the reduced input produces the correct results for the original inputs. Given  $x$  and its rounding interval  $[l, h]$ , reduced input  $x'$  is computed with range reduction. The next task before polynomial generation is to identify the reduced rounding interval for  $P(x')$  that when used with output compensation produces the correctly rounded result. We use the inverse of the output compensation function to identify the reduced interval  $[l', h']$ .

The last step is to synthesize a polynomial of a degree  $d$  using an arbitrary precision linear programming (LP) solver that satisfies the constraints (*i.e.*,  $l' \leq P(x') \leq h'$ ) when given a set of inputs  $x'$ . Approximating the correctly rounded result with the RLIBM approach provides more freedom in generating polynomials. Hence, the resulting RLIBM libraries are more efficient compared to mainstream libraries.

### 3 AN ILLUSTRATIVE EXAMPLE OF OUR APPROACH

We describe our entire approach with an end-to-end example for creating a polynomial approximation for  $\ln(x)$  that produces correctly rounded results for a 5-bit FP representation with 2 exponent bits (FP5) and a 4-bit FP with 2 exponent bits (FP4) for all standard rounding modes (*i.e.*,  $rn$ ,  $ra$ ,  $rz$ ,  $ru$ , and  $rd$ ). Figure 1(b) and Figure 1(c) show the bit-string of FP5 and FP4, respectively. Although we illustrate our approach with FP5 and FP4 for ease of exposition, it is beneficial in practice to create table-lookups for FP5 and FP4 because there are only 32 and 16 distinct bit-patterns, respectively.

The  $\ln(x)$  function is defined over the input domain  $(0, \infty)$ . The result of  $\ln(x)$  is NaN when  $x < 0$  or when  $x$  is NaN. The result is  $\infty$  when the input is  $\infty$  and  $-\infty$  when the input is 0. There are only 11 non-special case inputs, which range from 0.25 to 3.5 in FP5. Similarly, there are only 5 non-special case inputs in FP4. Now, our goal is to generate a single polynomial approximation that produces correctly rounded results for both FP5 and FP4 with all five rounding modes.

To accomplish this goal, we will generate a polynomial approximation that produces correctly rounded results for a 7-bit FP representation (FP7) with the round-to-odd mode. Here, FP7 has exactly the same number of exponent bits as FP5 and FP4 (*i.e.*, 2 exponent bits). Effectively, FP7 has 2 additional fraction bits when compared to FP5. Every value that is representable in FP5 and FP4 is also representable in FP7. While rounding with the round-to-odd mode, if the polynomial approximation produces a value that is exactly representable in FP7, then it is unchanged. Otherwise, the result of the polynomial approximation (which is implemented in double precision) is rounded to the nearest FP7 value whose bit-string is odd (*i.e.*, the last bit is a 1). When this FP7 round-to-odd result is rounded to a value in FP5 or FP4 according to any of the five standard rounding modes, it produces the correct result for them.

**Why does a correctly rounded result with the round-to-odd mode for FP7 work with FP5/FP4?** As the number of exponent bits is identical in FP7, FP5, and FP4, every value that is representable in FP5 and FP4 is also representable in FP7. Let us consider an input 1.5. We want to produce correctly rounded results of  $\ln(1.5)$  for all the rounding modes with FP5 and FP4. The

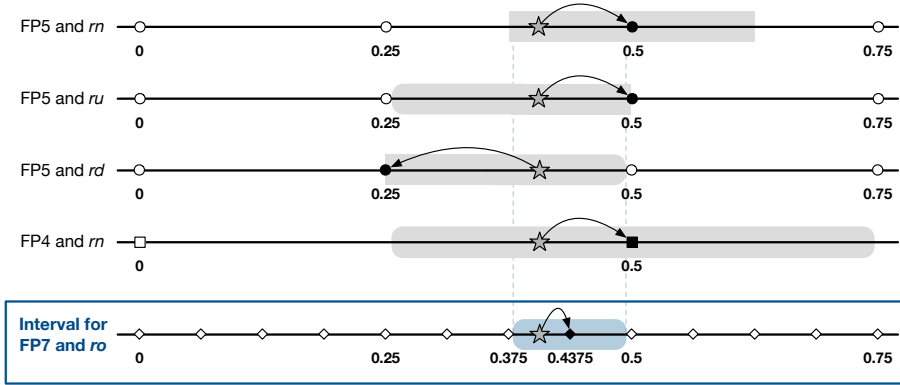


Fig. 7. The correctly rounded result of  $\ln(1.5)$  for FP5 and FP4 with some subset of the rounding modes and their rounding intervals (gray box). The gray star represents the real value of  $\ln(1.5)$ . Values that are representable in FP7, FP5, and FP4 are shown with rhombus, circle, and square, respectively. Solid shapes represents the correctly rounded result for the chosen representation and rounding mode. The last row shows the odd interval to produce the correctly rounded result of  $\ln(1.5)$  in FP7 with the round-to-odd mode. The odd interval is a subset of the intersection of the rounding intervals of these configurations.

first row of Figure 7 shows the real result (*i.e.*, a star) and the correctly rounded result of FP5 with the *rn* mode. If we want to generate the correctly rounded result for  $\ln(1.5)$  in FP5 with the *rn* mode using polynomial approximations, there is an interval of real values around the correctly rounded result such that producing any value in that interval produces the correct result (shaded in gray). The subsequent rows show the correctly rounded result and the rounding intervals for other rounding modes of FP5 and FP4 for the same input 1.5. When compared to FP5, the rounding interval for FP4 will be larger because the distance between adjacent points is larger. Intuitively, a single polynomial can produce the correctly rounded result of  $\ln(1.5)$  for both FP5 and FP4 with all rounding modes if it produces a value that lies in the common interval among all these modes and precision configurations (*i.e.*, an intersection of the rounding intervals).

We show that computing the correctly rounded result of  $\ln(1.5)$  with FP7 using the round-to-odd mode and identifying the interval around this result in FP7 is an effective way to compute the common interval described above. The last row of Figure 7 shows the correctly rounded result in FP7 with the round-to-odd mode and the interval to produce that value. The interval for the round-to-odd result in FP7 is smaller than the common interval among FP5 and FP4 with all the rounding modes because it works for many other representations beyond FP5 and FP4.

In FP7, there are three additional values between the two adjacent FP5 values. Hence, the round-to-odd result with FP7 preserves enough information to produce the correctly rounded result with FP5 and FP4 with any rounding mode. Our proofs in Section 5 show that this is a generic result for any representation with  $n$ -bits.

**Generating polynomial approximations.** The first step is to identify rounding intervals for producing the correctly rounded result of FP7 with the round-to-odd mode, which we call the odd interval. In our approach, polynomial evaluation happens with double precision. Hence, we identify an interval of values in double precision such that any value in that interval rounds to the correctly rounded result in FP7 with the round-to-odd mode. When the correctly rounded result in FP7 with the round-to-odd mode is even (*i.e.*, the bit-string is even when interpreted as an unsigned integer), the rounding interval is a singleton. For example, the odd interval for  $\ln(1.0)$  is a singleton because the correctly rounded result is 0. If the correctly rounded result with the round-to-odd mode is not




Fig. 8. (a) The set of constraints for the polynomial approximation to produce correctly rounded results for each input with the round-to-odd mode in FP7. (b) The odd intervals for each input and the resulting polynomial from our approach that produces a value in the odd interval for all inputs.

even, then we can identify the odd interval as follows. We identify the preceding value ( $l$ ) and the succeeding ( $h$ ) value corresponding to the correctly rounded result in FP7. Then, the open interval  $(l, h)$  is the odd interval. Figure 8(b) shows the odd interval for each input (shaded in blue). Next, we need to create a polynomial approximation that produces a value in the odd interval for each input.

Creating polynomial approximations with singleton odd intervals is challenging because there is no freedom for the polynomial generator. For the  $\ln(x)$  function with FP7, there is only one input (*i.e.*, 1.0) whose odd interval is a singleton. We treat it as a special case. In general, we use mathematical properties of the elementary function for larger data types to effectively handle such singleton odd intervals (Section 4.3). Figure 8(b) shows the remaining inputs and their non-singleton odd intervals.

The next step is to generate a polynomial that produces a value in the odd interval for all inputs. We show the constraints imposed by the odd interval on the output of the polynomial in Figure 8(a). As there are only 10 non-special case inputs, we encode them as a system of linear inequalities similar to our prior work in the RLIBM project and use an LP solver to solve for the coefficients of a 4<sup>th</sup> degree polynomial  $P(x)$  (see Figure 8(a)). For larger representations, we employ sophisticated range reduction, counterexample guided polynomial generation, and generate piecewise polynomials [Lim and Nagarakatte 2021a]. Figure 8(b) pictorially shows the generated polynomial, which produces a value in the odd interval for each input. This polynomial will produce the correctly rounded result of  $\ln(x)$  when the result is rounded to FP5 and FP4 with any of the five rounding modes in the standard.

#### 4 OUR APPROACH TO GENERATE A GENERIC POLYNOMIAL APPROXIMATION

Our goal is to generate a single polynomial approximation of an elementary function that produces correctly rounded results for all inputs for multiple precision and rounding configurations. Let  $\mathbb{T}_n$  be a  $n$ -bit FP representation (*i.e.*,  $\mathbb{F}_{n,|E|}$ ). Let  $\mathbb{T}_k$  be a representation where  $\mathbb{T}_k$  has no more precision bits compared to  $\mathbb{T}_n$  with the same number of exponent bits. Specifically,  $\mathbb{T}_k = \mathbb{F}_{k,|E|}$  where  $|E| + 1 < k \leq n$ . Note that all values exactly representable in  $\mathbb{T}_k$  are also exactly representable in  $\mathbb{T}_n$  (*i.e.*,  $\mathbb{T}_k \subseteq \mathbb{T}_n$ ). We define  $rm$  to be a rounding mode in the standard (*i.e.*,  $rm \in \{rn, ra, rz, ru, rd\}$ ). Our goal is to generate a polynomial approximation  $A_{\mathbb{H}}(x)$ , which is implemented in representation  $\mathbb{H}$ , of an elementary function  $f(x)$  that produces correctly rounded results for all inputs for any representation  $\mathbb{T}_k$  and any rounding mode  $rm$ . Specifically, rounding the result of  $A_{\mathbb{H}}(x)$  to any




Fig. 9. The round-to-odd (*ro*) rounding mode. We show the rounding of  $v_{\mathbb{R}}$  with the *ro* mode. Here,  $w_0$ ,  $w_1$ ,  $w_2$ ,  $w_3$ , and  $w_4$  are values representable in representation  $\mathbb{T}$ . If  $v_{\mathbb{R}}$  is exactly representable in  $\mathbb{T}$ , then  $v_{\mathbb{R}}$  rounds to that value. Otherwise,  $v_{\mathbb{R}}$  rounds to the nearest value in  $\mathbb{T}$  that is odd.

representation  $\mathbb{T}_k$  with the *rm* rounding mode must result in the same value as computing  $f(x)$  in real numbers and rounding the result to  $\mathbb{T}_k$  with the *rm* rounding mode, for all inputs in  $\mathbb{T}_k$ .

$$RN_{\mathbb{T}_k, rm}(A_{\mathbb{H}}(x)) = RN_{\mathbb{T}_k, rm}(f(x))$$

**Main insight.** To generate correct results for  $\mathbb{T}_k$ , our key insight is to create a polynomial approximation that produces the correctly rounded results for  $\mathbb{T}_{n+2}$  with the round-to-odd mode. We prove that it produces the correctly rounded result for any representation  $\mathbb{T}_k$  with all standard rounding modes when we round the round-to-odd result to the target representation (Section 5). Intuitively, our approach works because the round-to-odd result with  $\mathbb{T}_{n+2}$  maintains sufficient information about the real value of an elementary function  $f(x)$  that is required for correct rounding for all representations  $\mathbb{T}_k$  with any standard rounding mode.

To generate a polynomial approximation for  $\mathbb{T}_{n+2}$  with the round-to-odd mode, we use our RLIBM approach. We approximate the correctly rounded result rather than the real value. We extend the RLIBM approach to handle the round-to-odd mode. Specifically, we need to generate an interval of values around the correctly rounded round-to-odd result for each input, which we call the odd interval. If the generated polynomial produces a value in the odd interval for a particular input, then it produces the correct result for all representations  $\mathbb{T}_k$  and for all rounding modes. One unique challenge that we address is the presence of singleton odd intervals, which happens when the correctly rounded result in  $\mathbb{T}_{n+2}$  is even. To scale to 32-bit floats (*i.e.*,  $\mathbb{T}_{n+2} = 34$ -bit float), we employ range reduction, counterexample guided polynomial generation, and generate piecewise polynomials. Finally, we use a linear programming solver to solve for the coefficients of a polynomial given a system of linear constraints generated from the odd intervals.

#### 4.1 Generating the Correctly Rounded Result for $\mathbb{T}_{n+2}$ with the Round-to-Odd Mode

As we make a case for creating polynomial approximations for  $\mathbb{T}_{n+2}$  with the round-to-odd mode, we formally define it and describe rounding with the round-to-odd mode using the rounding components. We describe the properties of the correctly rounded result with the round-to-odd mode and provide intuition on why rounding the round-to-odd result in  $\mathbb{T}_{n+2}$  to any representation  $\mathbb{T}_k$  produces the correct result.

The round-to-odd (*ro*) is a non-standard rounding mode that has been previously used to avoid double rounding issues while converting a binary FP number to a decimal FP number [Goldberg 1991] and while performing primitive operations [Boldo and Melquiond 2005, 2008]. Given a real value  $v_{\mathbb{R}}$ , the round-to-odd mode rounds  $v_{\mathbb{R}}$  as follows. If  $v_{\mathbb{R}}$  is exactly representable as value  $v$  in the target representation, then  $v_{\mathbb{R}}$  rounds to  $v$ . Otherwise,  $v_{\mathbb{R}}$  rounds to the nearest odd value in the target representation. Figure 9 illustrates the round-to-odd rounding mode. Using the rounding components ( $s$ ,  $v^-$ ,  $rb$ , *sticky*) from Section 2.2, the round-to-odd mode can be defined as follows:




Fig. 10. An example to show that the round-to-odd result in  $\mathbb{T}_{n+2}$  maintains sufficient information to produce correctly rounded results for  $\mathbb{T}_n$  when the round-to-odd result is double rounded to  $\mathbb{T}_n$  with the  $rn$  mode. Real value is represented with a red star. Here,  $w_0, w_1, w_2, w_3,$  and  $w_4$  are representable values in  $\mathbb{T}_{n+2}$ . Values  $v_0$  and  $v_1$  are adjacent values representable in  $\mathbb{T}_n$ . As all values in  $\mathbb{T}_n$  are representable in  $\mathbb{T}_{n+2}$ ,  $w_0 = v_0$  and  $w_4 = v_1$ . Solid arrow represents directly rounding the real value to  $\mathbb{T}_n$ . The dotted arrows represent the process of double rounding from the real value to the round-to-odd result in  $\mathbb{T}_{n+2}$  and subsequently to  $\mathbb{T}_n$ . (a) When the real value is in the interval between  $w_0$  and  $w_2$ . (b) When the real value is exactly equal to  $w_2$ , which is the midpoint of  $v_0$  and  $v_1$ . (c) When the real value is between  $w_2$  and  $w_4$ .

$$v_{ro} = RN_{\mathbb{T},ro}(v_{\mathbb{R}}) = \begin{cases} s \times v^- & \text{if } IsOdd(v^-) \vee (rb = 0 \wedge sticky = 0) \\ s \times v^+ & \text{otherwise} \end{cases}$$

where  $v^+$  is the adjacent value to  $v^-$  in  $\mathbb{T}$ .

Our contribution is to use the round-to-odd mode to generate correctly rounded elementary functions for multiple representations and multiple types. Specifically, we prove the following theorem in Section 5, which forms the foundation for our approach.

**THEOREM 1.** *Let  $v_{\mathbb{R}} = f(x)$  be the real valued result of an elementary function and  $v_{ro} = RN_{\mathbb{T}_{n+2},ro}(v_{\mathbb{R}})$ . Let  $v$  be a value in the odd interval of  $v_{ro}$ . Consider a rounding mode  $rm \in \{rn, ra, rz, ru, rd\}$ . Then,*

$$RN_{\mathbb{T}_k,rm}(v) = RN_{\mathbb{T}_k,rm}(v_{\mathbb{R}})$$

We propose an efficient procedure to create polynomial approximations  $A_{\mathbb{H}}(x)$  of an elementary function  $f(x)$  that produces values in the odd interval of the correctly rounded result in  $\mathbb{T}_{n+2}$ . Using Theorem 1, rounding any value  $v$  in the odd interval (i.e.,  $A_{\mathbb{H}}(x)$ ) to  $\mathbb{T}_k$  using a rounding mode  $rm$  produces the correctly rounded result of  $f(x)$  in  $\mathbb{T}_k$  using the same rounding mode  $rm$ .

**An example to show why the round-to-odd result avoids double rounding errors.** We provide intuition on how rounding with the round-to-odd mode avoids double rounding errors in Figure 10. Any value that is representable in  $\mathbb{T}_n$  is also representable in  $\mathbb{T}_{n+2}$ . Further, there are three additional values ( $w_1, w_2, w_3$ ) in  $\mathbb{T}_{n+2}$  between  $w_0$  and  $w_4$ . Here,  $w_0$  and  $w_4$  are also representable in  $\mathbb{T}_n$ . In the round-to-odd mode, any real value between  $w_0$  and  $w_2$  rounds to  $w_1$ . Similarly, any real value between  $w_2$  and  $w_4$  rounds to  $w_3$ . If the real value is exactly equal to  $w_0$ , then the round-to-odd mode with  $\mathbb{T}_{n+2}$  also rounds to  $w_0$  (similarly  $w_2$  and  $w_4$  with  $\mathbb{T}_{n+2}$ ). Figure 10 illustrates the task of rounding the real value directly to  $\mathbb{T}_n$  with the  $rn$  mode (solid arrow) and the result produced from double rounding the  $ro$  result from  $\mathbb{T}_{n+2}$  to  $\mathbb{T}_n$  using the  $rn$  mode.

In the context of rounding a real value directly to  $\mathbb{T}_n$  with rounding components, the last bit of the round-to-odd result in  $\mathbb{T}_{n+2}$  captures the sticky bit. Similarly, the penultimate bit of the round-to-odd result in  $\mathbb{T}_{n+2}$  captures the rounding bit. In summary, the round-to-odd result in  $\mathbb{T}_{n+2}$  maintains sufficient information about the real value so that when the round-to-odd result is (double) rounded to  $\mathbb{T}_k$  with any rounding mode, it produces the correctly rounded result for  $\mathbb{T}_k$ .

#### 4.2 Polynomials for Correctly Rounded Results with the Round-to-Odd Mode in $\mathbb{T}_{n+2}$

Our strategy is to create a generic polynomial approximation that produces correctly rounded results for  $\mathbb{T}_{n+2}$  using the round-to-odd mode. Next, we describe our approach to generate such

```

1 Function GenerateGenericPolynomial( $f, \mathbb{T}_{n+2}, \mathbb{H}, X, d, RR_{\mathbb{H}}, OC_{\mathbb{H}}$ ):
2    $O \leftarrow \text{CalcResultsInRO}(f, \mathbb{T}_{n+2}, X)$ 
3    $(L, S) \leftarrow \text{CalcOddIntervals}(O, \mathbb{T}_{n+2}, \mathbb{H})$ 
4   if  $L = \emptyset$  then return ( $false, \emptyset, DNE$ )
5    $(status, P) \leftarrow \text{RLibmPolyGen}(L, \mathbb{H}, d, RR_{\mathbb{H}}, OC_{\mathbb{H}})$ 
6   return ( $status, S, P$ )

```

**Algorithm 1:** A sketch of our approach to generate piecewise polynomials of degree  $d$  for elementary function  $f(x)$  in the representation  $\mathbb{T}_{n+2}$  using the round-to-odd mode. The resulting polynomial when used with range reduction ( $RR_{\mathbb{H}}$ ) and output compensation ( $OC_{\mathbb{H}}$ ) produces correctly rounded results for all inputs  $x \in X$  with all representations  $\mathbb{T}_k$  for all standard rounding modes.  $\text{CalcResultsInRO}$  computes the round-to-odd result using an oracle (see Figure 11).  $\text{CalcOddIntervals}$  computes the set of odd intervals ( $L$ ) and set ( $S$ ) of singleton odd intervals (see Figure 11). Once we have the odd intervals and singletons, we use  $\text{RLIBM}$ 's polynomial generation procedure ( $\text{RLibmPolyGen}$ ) to obtain the generic polynomial.

```

1 Function CalcResultsInRO( $f, \mathbb{T}_{n+2}, X$ ):
2    $O \leftarrow \emptyset$ 
3   foreach  $x \in X$  do
4      $y = f(x)$ 
5      $(s, v^-, rb, sticky) \leftarrow \text{RComp}(y, \mathbb{T}_{n+2})$ 
6     if  $\text{IsOdd}(v^-) \vee (rb = 0 \wedge sticky = 0)$ 
7       then
8          $y_{ro} \leftarrow s \times v^-$ 
9       else
10         $v^+ \leftarrow \text{GetSuccVal}(v^-, \mathbb{T}_{n+2})$ 
11         $y_{ro} \leftarrow s \times v^+$ 
12      end
13       $O \leftarrow O \cup (x, y_{ro})$ 
14  end
15  return  $O$ 

```

```

1 Function CalcOddIntervals( $O, \mathbb{T}_{n+2}, \mathbb{H}$ ):
2   foreach  $(x, y_{ro}) \in O$  do
3      $L \leftarrow \emptyset$ 
4      $S \leftarrow \emptyset$ 
5     if  $\text{IsEven}(y_{ro})$  then
6        $S \leftarrow S \cup (x, y_{ro})$ 
7     end
8     else
9        $y^- \leftarrow \text{GetPrecVal}(y_{ro}, \mathbb{T}_{n+2})$ 
10       $l \leftarrow \text{GetSuccVal}(y^-, \mathbb{H})$ 
11       $y^+ \leftarrow \text{GetSuccVal}(y_{ro}, \mathbb{T}_{n+2})$ 
12       $h \leftarrow \text{GetPrecVal}(y^+, \mathbb{H})$ 
13       $L \leftarrow L \cup (x, [l, h])$ 
14    end
15    return ( $L, S$ )
16  end

```

Fig. 11.  $\text{CalcResultsInRO}$  computes the correctly rounded result of  $f(x)$  in  $\mathbb{T}_{n+2}$  using the round-to-odd rounding mode for each input  $x \in X$ .  $\text{CalcOddIntervals}$  computes the odd interval for each input  $x$  based on the correctly rounded result  $y_{ro}$  in  $\mathbb{T}_{n+2}$ . The list  $S$  is the set of inputs that have a singleton as the odd interval. The list  $L$  contains inputs and the corresponding odd intervals.  $\text{GetPrecVal}(a, \mathbb{T})$  returns the value preceding  $a$  in the representation  $\mathbb{T}$ .  $\text{GetSuccVal}(a, \mathbb{T})$  returns the value succeeding  $a$  in the representation  $\mathbb{T}$ .

a polynomial approximation. Algorithm 1 provides a high-level sketch of this process. Given an elementary function  $f(x)$  and a list of inputs  $X$  in the  $\mathbb{T}_n$  representation (i.e.,  $X \subseteq \mathbb{T}_n$ ), the first step is to compute the correctly rounded result in representation  $\mathbb{T}_{n+2}$  using the round-to-odd mode (i.e.,  $y_{ro}$  for each input  $x \in X$ ). Figure 11 shows our algorithm to compute the round-to-odd result  $y_{ro}$  for each input using the real value from the oracle.

Subsequently, we compute the odd interval of each result  $y_{ro}$  such that any real value in the odd interval rounds to  $y_{ro}$ . Figure 11 also provides our algorithm to compute the odd interval. The odd intervals of some inputs can be a singleton (i.e., only one value in the odd interval), which we handle separately (Section 4.3). Once we have a set of non-singleton odd intervals for all inputs,

we use our prior work in the RLIBM project [Lim and Nagarakatte 2021a] to generate piecewise polynomials.

At the end of this process, we will have two main components that together can produce correctly rounded results for  $f(x)$ . First, our approach produces a set  $S$  that contains inputs whose odd interval is a singleton. For the resulting math libraries to be efficient, we need a fast method to check these inputs and compute results for them either using table lookups or using function-specific mathematical properties (Section 4.3). Second, our approach produces piecewise polynomials that when used with output compensation produces correct results for all inputs when rounded to any  $\mathbb{T}_k$  with all standard rounding modes.

**Computing the round-to-odd result from a real value.** The first step in our approach is to identify the correctly rounded result  $y_{ro}$  for input  $x$ . Figure 11 provides the steps to compute the round-to-odd result in  $\mathbb{T}_{n+2}$  given a real value of  $f(x)$ . We compute the real value  $y = f(x)$  for each input  $x$  using an oracle (e.g., MPFR library). Then, we obtain the rounding components ( $s$ ,  $v^-$ ,  $rb$ ,  $sticky$ ) as described in Section 2.2. When the real value is exactly representable ( $rb = 0$  and  $sticky = 0$ ) or when  $v^-$  is odd, then the round-to-odd result is  $v^-$ . Otherwise, the round-to-odd result in  $\mathbb{T}_{n+2}$  is the value succeeding  $v^-$  in  $\mathbb{T}_{n+2}$ .

**Deducing the odd interval of an input.** Once we determine the correctly rounded result  $y_{ro}$  of  $f(x)$  in representation  $\mathbb{T}_{n+2}$  using the round-to-odd mode, the next step is compute the interval of values in representation  $\mathbb{H}$ , which is used for polynomial evaluation and range reduction, such that producing any value in the interval rounds to  $y_{ro}$ , which we call as the odd interval. The function CalcOddIntervals in Figure 11 describes the steps to compute the odd interval. If the correct rounded result  $y_{ro}$  in  $\mathbb{T}_{n+2}$  is even, then the odd interval is a singleton. In such cases, the only value that rounds to  $y_{ro}$  with the round-to-odd mode is  $y_{ro}$  itself.

Generating polynomial approximations with singletons is challenging because they limit the amount of freedom available to the polynomial generator. Hence, we identify such inputs and handle them separately. If  $y_{ro}$  is odd, then all values in  $\mathbb{H}$  that are strictly greater than the preceding value of  $y_{ro}$  in  $\mathbb{T}_{n+2}$  and strictly less than the succeeding value of  $y_{ro}$  in  $\mathbb{T}_{n+2}$  forms the odd interval. Any value in this odd interval rounds to  $y_{ro}$  in  $\mathbb{T}_{n+2}$  with the round-to-odd mode. We deduce the odd interval for each input. In Figure 11,  $L$  represents the set of non-singleton odd intervals for all inputs, which is given to the polynomial generator.

**Piecewise polynomial generation using the odd intervals.** The next step is to generate piecewise polynomials that produce a value in the odd interval for all inputs. Each input and odd interval pair (i.e.,  $(x, [l, h]) \in L$ ) specifies the constraints on the polynomial approximation  $A_{\mathbb{H}}(x)$  for each input  $x$ . Using the RLIBM methodology, we create an LP problem with these constraints to deduce the coefficients of a polynomial with degree  $d$ . Similarly, we generate piecewise polynomials and use counterexample guided polynomial generation to facilitate the entire process. We also make sure that the generated polynomial produces a value in the odd interval after range reduction and output compensation.

**Implementation of the polynomial approximation for  $\mathbb{T}_k$ .** At the end of polynomial generation, we will have a set of inputs whose odd interval is a singleton and a polynomial approximation of  $f(x)$  that produces the correct round-to-odd result in  $\mathbb{T}_{n+2}$  for all inputs. We implement the polynomial approximation as follows. Given an input  $x$ , we first check whether the input  $x$ 's odd interval is a singleton. If so, we either use the precomputed round-to-odd result with table lookups or efficiently compute the round-to-odd result using function-specific properties. Otherwise, we use perform range reduction and use Horner's method for polynomial evaluation to compute the round-to-odd result in  $\mathbb{T}_{n+2}$  for the input. Finally, we round the round-to-odd result in  $\mathbb{T}_{n+2}$  to  $\mathbb{T}_k$  using the user specified rounding mode to return the final result. We guarantee that our

implementation produces the correctly rounded result of  $f(x)$  for any representation  $\mathbb{T}_k$  with all the standard rounding modes for all inputs  $x$ .

### 4.3 Computing Round-to-Odd Results for Inputs with Singleton Odd Intervals

One of the challenging issues for polynomial generation with odd intervals is the presence of singletons, which happens when the correctly rounded result with the round-to-odd mode is even (*i.e.*, the round-to-odd result in  $\mathbb{T}_{n+2}$  matches the real value). We want to identify such inputs efficiently. As both  $\mathbb{T}_n$  and  $\mathbb{T}_{n+2}$  are finite precision representations, all values in  $\mathbb{T}_n$  and  $\mathbb{T}_{n+2}$  are rational values. If the real value matches the round-to-odd result in  $\mathbb{T}_{n+2}$  exactly, then it is a rational value. Hence, our task of identifying inputs with singleton odd intervals corresponds to the problem of identifying rational inputs that produce rational outputs for various elementary functions, which is well studied [Aigner and Ziegler 2009; Baker 1975; Cohn 1974; Niven 1956]. We first use the mathematical properties of the elementary function  $f(x)$  to identify all rational inputs such that  $f(x)$  is a rational value. Then, we check if these inputs  $x_1$  and the corresponding result  $f(x_1)$  are exactly representable in  $\mathbb{T}_n$  and  $\mathbb{T}_{n+2}$ , respectively. If so, such inputs are of interest. Then, we need to develop a quick way to identify those inputs and compute the round-to-odd results for them without using a multi-way branch.

We now describe the specific mathematical properties of elementary functions that we use to identify inputs whose odd interval is a singleton and the mechanism that we use to efficiently compute round-to-odd results for them.

**Functions  $e^x$  and  $\ln(x)$ .** From the Lindemann-Weierstrass theorem [Baker 1975], if the input  $x$  is a non-zero rational value, then  $e^x$  cannot be a rational value. Hence, the only value that will have a singleton odd interval with  $e^x$  is  $x = 0$ . Similarly,  $\ln(x)$  will produce a rational output only when  $x = 1$ , which also follows from the Lindemann-Weierstrass theorem [Baker 1975]. We have a single branch to check this input and return the pre-computed correctly rounded result.

**Functions  $2^x$  and  $10^x$ .** The function  $2^x$  can produce a rational result only when  $x$  is an integer and the value of  $2^x$  is less than the dynamic range of the  $\mathbb{T}_{n+2}$  representation. When  $\mathbb{T}_n$  is a 32-bit float,  $\mathbb{T}_{n+2}$  can represent all values of  $2^x$  for  $x$  between  $-151 \leq x \leq 127$ . Thus, any integer input between  $-151$  and  $127$  (279 inputs in total) can produce a singleton odd interval. Hence, our implementation checks whether  $x$  is an integer within a certain bound (*i.e.*,  $-151 \leq x \leq 127$ ) and directly computes the result,  $2^x$ , using bit-wise operations.

Similar to  $2^x$ ,  $10^x$  produces a rational value when  $x$  is a positive integer. In contrast to  $2^x$ ,  $10^x$  grows much faster and there are a few inputs for which  $10^x$  is exactly representable in  $\mathbb{T}_{n+2}$ . For a 32-bit float ( $\mathbb{T}_n$ ), there are only 12 inputs ranging from 0 to 11 that are exactly representable in a 34-bit float ( $\mathbb{T}_{n+2}$ ). We use a precomputed table to store the correct results for these 12 inputs and use a switch statement for it.

**Functions  $\log_2(x)$  and  $\log_{10}(x)$ .** The  $\log_2(x)$  function produces a rational result when  $x$  is a power of 2 (*i.e.*,  $x = 2^k$  and  $k$  is an integer). We use bitwise operations to check if the input is a power of two. In contrast to  $\log_2(x)$ ,  $\log_{10}(x)$  produces a rational result when  $x$  is a positive power of 10 (*i.e.*,  $x = 10^k$  and  $k$  is a positive integer). This difference between  $\log_2(x)$  and  $\log_{10}(x)$  is due to the fact that  $\mathbb{T}_n$  cannot exactly represent negative powers of 10. When we are generating a polynomial to approximate the round-to-odd result with a 34-bit float (*i.e.*,  $\mathbb{T}_{n+2}$ ), there are 11 inputs that can produce singletons. We create table-lookups for them.

**The hyperbolic functions,  $\sinh(x)$  and  $\cosh(x)$ .** If the input  $x$  is a non-zero rational value, then  $y = \sinh(x)$  or  $y = \cosh(x)$  cannot be a rational value using the Lindemann-Weierstrass theorem [Niven 1956]. Hence, the only input whose odd interval is a singleton is 0, for which we use a branch condition.



**The  $\sin\pi i(x)$  function.** The function  $\sin\pi i(x)$  is equal to  $\sin(\pi x)$ . By Niven's theorem [Niven 1956], the only rational values of  $x$  between  $0 \leq x \leq \frac{1}{2}$  where  $\sin\pi i(x)$  is also a rational value are when  $x = 0$ ,  $x = \frac{1}{6}$ , and  $x = \frac{1}{2}$ . Among these three inputs,  $\frac{1}{6}$  is not exactly representable in  $\mathbb{T}_n$ . Given that  $\sin\pi i$  is a periodic function, there are only three cases of inputs in  $x \in \mathbb{T}_n$  where the result of  $\sin\pi i(x)$  is representable in  $\mathbb{T}_{n+2}$  when we extend the domain of  $x$  to the set of all inputs:

$$\sin\pi i(x) = \begin{cases} 0 & \text{if } x \text{ is an integer} \\ 1 & \text{if } x \equiv \frac{1}{2} \pmod{2.0} \\ -1 & \text{if } x \equiv \frac{3}{2} \pmod{2.0} \end{cases}$$

We need to implement the floating point modulo operation efficiently using integer operations. Consider the case where  $\mathbb{T}_n$  is a 32-bit float. All inputs  $x \in \mathbb{T}_n$  greater than or equal to  $2^{23}$  are integers where  $\sin\pi i(x)$  is always 0. Next, if  $x < 2^{23}$ , then we need to identify whether  $x$  is either an integer, a multiple of 0.5, or a multiple of 1.5. To determine this condition, we compute  $2x$  with a 32-bit float and then cast the result (i.e.,  $2x$ ) to a 32-bit integer to obtain the value  $t$ . This operation of casting the value  $2x$  to an integer truncates the value of  $2x$  to the integral part of  $2x$ . Now if we cast  $t$  back to a 32-bit float value and the resulting float value is exactly equal to  $2x$ , then  $t$  is an integer, which implies that  $x$  is either an integer, or a multiple of 0.5, or a multiple of 1.5. Finally, we compute the result of  $\sin\pi i(x)$  based on  $t$  as shown below:

$$\sin\pi i(x) = \begin{cases} 0 & \text{if } t \equiv 0 \pmod{2} \\ 1 & \text{if } t \equiv 1 \pmod{4} \\ -1 & \text{if } t \equiv 3 \pmod{4} \end{cases}$$

**The  $\cos\pi i(x)$  function.** Similarly,  $\cos\pi i(x) = \cos(\pi x)$  produces a rational value representable in  $\mathbb{T}_{n+2}$  in the following cases.

$$\cos\pi i(x) = \begin{cases} 1 & \text{if } x \text{ is an even integer} \\ -1 & \text{if } x \text{ is an odd integer} \\ 0 & \text{if } \text{fraction}(x) \equiv 0.5 \end{cases}$$

These checks can be performed efficiently using a similar strategy illustrated for the  $\sin\pi i(x)$  function. In summary, handling singleton odd intervals efficiently is important for performance when we generate polynomials for correctly rounded results in  $\mathbb{T}_{n+2}$  with the round-to-odd mode.

## 5 PROOF THAT THE ROUND-TO-ODD RESULT WITH $\mathbb{T}_{n+2}$ PRODUCES CORRECT RESULTS FOR $\mathbb{T}_k$

We provide a proof of Theorem 1 in this section. We prove that the round-to-odd result in  $\mathbb{T}_{n+2}$  produced by our polynomial approximation when rounded to  $\mathbb{T}_k$  with any of the standard rounding modes produces the correctly rounded result for  $\mathbb{T}_k$ .

### 5.1 Unique Properties of the Round-to-Odd Result

We prove the unique properties of the round-to-odd result, which we subsequently use to prove Theorem 1. When we use  $v_{\mathbb{R}}$  to represent the real value, we refer to it in the extended infinite precision representation.

LEMMA 1. *The round-to-odd result  $v_{r,o}$  in  $\mathbb{T}_{n+2}$  preserves the sign of  $v_{\mathbb{R}}$ .*

**Proof.** The value zero is representable in  $\mathbb{T}_{n+2}$ . The only value that rounds to zero is zero itself. Hence, all positive real values will round to a positive value in the round-to-odd mode. Similarly all negative real values will round to a negative value in the round-to-odd mode.  $\square$

LEMMA 2. Let  $v_{r_o} = RN_{\mathbb{T}_{n+2}, r_o}(v_{\mathbb{R}})$ . The first  $(n + 1)$ -bits of  $v_{r_o}$  and  $v_{\mathbb{R}}$  are identical.

**Proof.** The  $v_{r_o}$  result is created using the rounding components  $(s_{v_{r_o}}, v_{v_{r_o}}^-, rb_{v_{r_o}}, sticky_{v_{r_o}})$ . The round-to-odd mode preserves the sign of  $v_{\mathbb{R}}$  in  $v_{r_o}$ . Without loss of generality, we assume  $v_{\mathbb{R}}$  is positive for the rest of the proof. Further,  $v_{v_{r_o}}^-$  is the truncated value of  $v_{\mathbb{R}}$  (see Section 2.2). Hence, all the  $(n + 2)$ -bits of  $v_{v_{r_o}}^-$  and  $v_{\mathbb{R}}$  are identical. After rounding with the round-to-odd mode,  $v_{r_o}$  is either equal to  $v_{v_{r_o}}^-$  or the succeeding value of  $v_{v_{r_o}}^-$  in  $\mathbb{T}_{n+2}$ . We prove that the  $(n + 1)$ -bits of  $v_{r_o}$  and  $v_{\mathbb{R}}$  are identical by looking at the possible values of  $v_{v_{r_o}}^-$  and its relation to  $v_{r_o}$ .

First case, when  $v_{v_{r_o}}^-$  is odd. Then,  $v_{r_o} = v_{v_{r_o}}^-$ . Hence, all the  $(n + 2)$ -bits of  $v_{r_o}$  and  $v_{\mathbb{R}}$  are identical. Second case, when  $v_{v_{r_o}}^-$  is even. Hence, the last bit of  $v_{v_{r_o}}^-$  is 0. Now, there are two sub-cases. (1) If  $rb_{v_{r_o}} = 0$  and  $sticky_{v_{r_o}} = 0$ , then  $v_{r_o} = v_{v_{r_o}}^-$ . Hence, all the  $(n + 2)$ -bits of  $v_{r_o}$  and  $v_{\mathbb{R}}$  are identical. (2) If  $rb_{v_{r_o}} \neq 0$  or  $sticky_{v_{r_o}} \neq 0$ , then  $v_{r_o}$  is equal to the succeeding value of  $v_{v_{r_o}}^-$ . The only bit that changes between  $v_{v_{r_o}}^-$  and its succeeding value is the  $(n + 2)^{th}$ -bit. Hence, the first  $(n + 1)$ -bits of  $v_{r_o}$  and  $v_{\mathbb{R}}$  are identical.  $\square$

LEMMA 3. The  $(n + 2)^{th}$ -bit of  $v_{r_o}$  is equal to the bitwise OR of all the bits of  $v_{\mathbb{R}}$  starting from the  $(n + 2)^{th}$ -bit.

**Proof.** We prove this lemma using a strategy similar to Lemma 2. Intuitively, this lemma states that the last bit of  $v_{r_o}$  is 0 if and only all bits starting from the  $(n + 2)^{th}$ -bit of  $v_{\mathbb{R}}$  is 0.

As the round-to-odd mode preserves sign, we assume  $v_{\mathbb{R}}$  is positive for the rest of the proof without loss of generality. Let us say the rounding components for  $v_{r_o}$  are  $(s_{v_{r_o}}, v_{v_{r_o}}^-, rb_{v_{r_o}}, sticky_{v_{r_o}})$ . In the round-to-odd mode with  $\mathbb{T}_{n+2}$ ,  $v_{r_o}$  will be equal to either  $v_{v_{r_o}}^-$  or a succeeding value of  $v_{v_{r_o}}^-$  in  $\mathbb{T}_{n+2}$ . Now, we look at the cases where  $v_{v_{r_o}}^-$  is odd and even to complete the proof.

In the first case,  $v_{v_{r_o}}^-$  is odd. Then,  $v_{r_o} = v_{v_{r_o}}^-$ . The  $(n + 2)^{th}$ -bit of  $v_{v_{r_o}}^-$  is 1. As  $v_{v_{r_o}}^-$  is a truncated value of  $v_{\mathbb{R}}$ , the  $(n + 2)^{th}$ -bit of  $v_{\mathbb{R}}$  is 1. Hence, the bitwise-OR of all bits of  $v_{\mathbb{R}}$  starting from the  $(n + 2)^{th}$ -bit is 1, which is equal to the  $(n + 2)^{th}$ -bit of  $v_{r_o}$  in  $\mathbb{T}_{n+2}$ .

In the second case,  $v_{v_{r_o}}^-$  is even. There are two cases depending on the values of  $rb_{v_{r_o}}$  and  $sticky_{v_{r_o}}$ . In the first sub-case,  $rb_{v_{r_o}} = 0$  and  $sticky_{v_{r_o}} = 0$ , then  $v_{r_o} = v_{v_{r_o}}^-$ . The  $(n + 2)^{th}$ -bit of  $v_{v_{r_o}}^-$  is 0. So is the  $(n + 2)^{th}$ -bit of  $v_{\mathbb{R}}$ . From the definition of rounding components for  $\mathbb{T}_{n+2}$ ,  $rb_{v_{r_o}}$  is the value of the bit at position  $(n + 3)$  in  $v_{\mathbb{R}}$  and  $sticky_{v_{r_o}}$  is the bitwise-OR of bits starting from  $(n + 4)^{th}$ -bit in  $v_{\mathbb{R}}$ . Hence, all the bits of  $v_{\mathbb{R}}$  starting from the  $(n + 2)^{th}$ -bit are 0, which matches the  $(n + 2)^{th}$ -bit of  $v_{r_o}$ .

The next sub-case is when  $rb_{v_{r_o}} \neq 0$  or  $sticky_{v_{r_o}} \neq 0$ . In this case,  $v_{r_o}$  is equal to the succeeding value of  $v_{v_{r_o}}^-$  in  $\mathbb{T}_{n+2}$ , which is odd. Hence, the  $(n + 2)^{th}$ -bit of  $v_{r_o}$  is 1. Both  $sticky_{v_{r_o}}$  and  $rb_{v_{r_o}}$  are not zeros, one of the bits starting from  $(n + 2)^{th}$ -bit in  $v_{\mathbb{R}}$  is 1. Hence, the bitwise-OR of all bits starting from the  $(n + 2)^{th}$ -bit is 1, which matches the  $(n + 2)^{th}$ -bit of  $v_{r_o}$ .  $\square$

LEMMA 4. Let  $(s_1, v_1^-, rb_1, sticky_1)$  and  $(s_2, v_2^-, rb_2, sticky_2)$  be the rounding components for two real values  $v_1$  and  $v_2$  in rounding them to a FP representation  $\mathbb{T}_n$ . If  $s_1 = s_2$ ,  $v_1^- = v_2^-$ ,  $rb_1 = rb_2$ , and  $sticky_1 = sticky_2$ , then  $RN_{\mathbb{T}_n, rm}(v_1) = RN_{\mathbb{T}_n, rm}(v_2)$  for any rounding mode  $rm$ .

**Proof.** This lemma directly follows from the definition of rounding components in Section 2.2. Intuitively, this lemma states that identifying the correctly rounded result of the real value in  $\mathbb{T}_n$  only depends on the rounding components and the rounding mode  $rm$ .  $\square$

## 5.2 Proof that Double Rounding the Round-to-Odd Result Produces Correct Results for all $\mathbb{T}_k$

We now sketch the proof of Theorem 1. We show that rounding a real value  $v_{\mathbb{R}}$  to the FP representation  $\mathbb{T}_{n+2} = \mathbb{F}_{n+2, |E|}$  using the round-to-odd mode to produce  $v_{r_o}$  and then subsequently rounding



Fig. 12. Rounding components while rounding  $v_{\mathbb{R}}$  and  $v_{r_o}$  to  $\mathbb{T}_k$ . We show the bit-string of  $v_{\mathbb{R}}$  in extended infinite precision representation. Note  $v_{r_o}$  is a value in  $\mathbb{T}_{n+2}$ .

the result ( $v_{r_o}$ ) to  $\mathbb{T}_k$  using a rounding mode  $rm$  produces the same value as rounding  $v_{\mathbb{R}}$  directly to  $\mathbb{T}_k$  using the same rounding mode  $rm$ , as long as  $|E| + 1 < k \leq n$ . More formally, we prove that

$$RN_{\mathbb{T}_k, rm}(RN_{\mathbb{T}_{n+2}, r_o}(v_{\mathbb{R}})) = RN_{\mathbb{T}_k, rm}(v_{r_o})$$

Our high-level strategy is to show that the rounding components for  $v_{\mathbb{R}}$  to  $\mathbb{T}_k$  and rounding  $v_{r_o}$  to  $\mathbb{T}_k$  are exactly the same. We prove Theorem 1 by proving the following theorem.

**THEOREM 2.** *Given a real number  $v_{\mathbb{R}}$ , representations  $\mathbb{T}_k$  and  $\mathbb{T}_{n+2}$  with same number of exponent bits that satisfy the condition  $|E| + 1 < k \leq n$ , and a rounding mode  $rm \in \{rn, ra, rz, ru, rd\}$ , then  $RN_{\mathbb{T}_k, rm}(v_{\mathbb{R}}) = RN_{\mathbb{T}_k, rm}(RN_{\mathbb{T}_{n+2}, r_o}(v_{\mathbb{R}}))$ .*

Let us say  $v_{r_o} = RN_{\mathbb{T}_{n+2}, r_o}(v_{\mathbb{R}})$ . Our goal is to prove  $RN_{\mathbb{T}_k, rm}(v_{\mathbb{R}}) = RN_{\mathbb{T}_k, rm}(v_{r_o})$ . Using Lemma 4, if the rounding components for rounding  $v_{\mathbb{R}}$  to  $\mathbb{T}_k$  is the same as the rounding components for rounding  $v_{r_o}$  to  $\mathbb{T}_k$ , then we prove that  $RN_{\mathbb{T}_k, rm}(v_{\mathbb{R}}) = RN_{\mathbb{T}_k, rm}(v_{r_o})$  for all rounding modes  $rm$ . Hence, our strategy is to show that the rounding components for rounding  $v_{\mathbb{R}}$  and  $v_{r_o}$  to  $\mathbb{T}_k$  are identical. As the the round-to-odd mode preserves the sign, we will consider  $v_{\mathbb{R}}$  to be positive in the rest of the proof.

The representation of  $v_{\mathbb{R}}$  in extended infinite precision representation ( $B_{v_{\mathbb{R}}}$ ) is as follows:

$$B_{v_{\mathbb{R}}} = b_1 b_2 \dots b_{k-1} b_k b_{k+1} b_{k+2} \dots b_n b_{n+1} b_{n+2} b_{n+3} \dots$$

Let us say ( $s_1, v_1^-, rb_1, sticky_1$ ) are the rounding components for rounding  $v_{\mathbb{R}}$  to  $\mathbb{T}_k$ . Then,  $v_1^-$  is the truncated value in  $\mathbb{T}_k$ . While rounding to  $\mathbb{T}_k$ , the rounding bit,  $rb_1$ , is the  $(k+1)^{th}$ -bit of  $v_{\mathbb{R}}$ . The sticky bit,  $sticky_1$  is the bitwise-OR of all bits starting from the  $(k+2)^{th}$ -bit of  $v_{\mathbb{R}}$ .

$$B_{v_1^-} = b_1 b_2 b_3 \dots b_{k-1} b_k, \quad rb_1 = b_{k+1}, \quad sticky_1 = b_{k+2} | b_{k+3} | \dots$$

Figure 12 pictorially shows the rounding components  $v_1^-$ ,  $rb_1$ , and  $sticky_1$  while rounding  $v_{\mathbb{R}}$  to  $\mathbb{T}_k$ .

Similarly, we next identify the rounding components ( $s_2, v_2^-, rb_2, sticky_2$ ) for rounding  $v_{r_o}$  to  $\mathbb{T}_k$ . Note that  $v_{r_o}$  is a result in  $\mathbb{T}_{n+2}$ . From Lemma 2 and Lemma 3, the bit-string of  $v_{r_o}$  is:

$$B_{v_{r_o}} = b_1 b_2 b_3 \dots b_{k-1} b_k \dots b_n b_{n+1} t, \quad t = b_{n+2} | b_{n+3} | b_{n+4} | \dots$$

Since  $k \leq n$ , there are at least one bit (i.e.,  $b_{n+1}$ ) between  $b_k$  and  $t$ , where  $t$  is the  $(n+2)^{th}$ -bit in  $v_{r_o}$ . The rounding components when we round  $v_{r_o}$  to  $\mathbb{T}_k$  are:

$$B_{v_2^-} = b_1 b_2 b_3 \dots b_{k-1} b_k, \quad rb_2 = b_{k+1} \quad sticky_2 = b_{k+2} | b_{k+3} | \dots | b_{n+1} | t$$

Figure 12 shows these components while rounding  $v_{r_o}$  to  $\mathbb{T}_k$ .

Now, we compare the rounding components when we directly round  $v_{\mathbb{R}}$  to  $\mathbb{T}_k$  with the rounding components when we round  $v_{r_o}$  to  $\mathbb{T}_k$ . The sign information ( $s_1$  and  $s_2$ ) is identical because the round-to-odd mode preserves the sign of  $v_{\mathbb{R}}$ . The truncated values,  $v_1^-$  and  $v_2^-$ , are equal because their bit-strings are identical. The rounding bit,  $rb_1$  and  $rb_2$ , is identical and is equal to  $b_{k+1}$ . Let us look at the sticky bits,  $sticky_1$  and  $sticky_2$ :

$$\begin{aligned}
sticky_2 &= b_{k+2} | b_{k+3} | \cdots | b_{n+1} | t = b_{k+2} | b_{k+3} | \cdots | b_{n+1} | b_{n+2} | b_{n+3} | b_{n+4} | \cdots \\
&= sticky_1
\end{aligned} \tag{2}$$

Hence, all the rounding components for rounding  $v_{\mathbb{R}}$  to  $\mathbb{T}_k$  directly and rounding  $v_{ro}$  to  $\mathbb{T}_k$  are identical. Hence,  $RN_{\mathbb{T}_k,rm}(v_{\mathbb{R}}) = RN_{\mathbb{T}_k,rm}(v_{ro})$  from Lemma 4.  $\square$

Theorem 1 directly follows from Theorem 2. Theorem 2 states that  $v_{ro}$ , which is produced by rounding a real value  $v_{\mathbb{R}}$  to  $\mathbb{T}_{n+2}$  using the round-to-odd mode, rounds to the same value as if  $v_{\mathbb{R}}$  is directly rounded to  $\mathbb{T}_k$  using the same rounding mode  $rm$ . If we substitute  $v_{\mathbb{R}}$  with the exact result of the elementary function  $f(x)$  for a given input  $x \in \mathbb{T}_n$ , then

$$RN_{\mathbb{T}_k,rm}(RN_{\mathbb{T}_{n+2},ro}(f(x))) = RN_{\mathbb{T}_k,rm}(f(x))$$

Further, by definition, all values in the odd interval of  $v_{ro}$  in  $\mathbb{T}_{n+2}$  round to  $v_{ro}$ . Hence, any value in the odd interval rounds to the correctly rounded result for representations  $\mathbb{T}_k$  using any rounding mode  $rm \in \{rn, ra, rz, ru, rd\}$ .

## 6 EXPERIMENTAL EVALUATION

We describe our prototype, experimental methodology, and the results of our experiments to check the correctness and performance of the generated polynomial approximations.

**Prototype.** The prototype, RLIBM-ALL, is open-source and publicly available [Lim and Nagarakatte 2021c]. RLIBM-ALL is a generator and a collection of correctly rounded implementations of polynomial approximation for multiple representations and rounding modes. RLIBM-ALL contains ten functions that produce the correctly rounded result of  $f(x)$  for the 34-bit FP representation (*i.e.*,  $\mathbb{T}_{n+2}$ ) with 8 bits of exponent (FP34) with the round-to-odd mode. As FP34 is not supported in hardware, RLIBM-ALL maintains the FP34 result in double precision. RLIBM-ALL's functions produce the correct result for all  $n$ -bit FP representations with 8-bits for the exponent with all the rounding modes in the IEEE standard where  $9 < n \leq 32$ . This includes 32-bit float, bfloat16, and tensorfloat32. RLIBM-ALL uses the MPFR library [Fousse et al. 2007] with up to 1000 precision bits to compute the oracle value of  $f(x)$ . RLIBM-ALL uses SoPlex [Gleixner et al. 2012], an exact rational LP solver, to generate the coefficients of the polynomials with a time limit of five minutes. We limit the size of the LP formulation to contain up to fifty thousand reduced input and interval constraints. We use the range reduction and output compensation functions from the RLIBM prototype. RLIBM-ALL performs range reduction, polynomial evaluation, and output compensation using the double precision. The polynomial evaluation uses the Horner's method [Borwein and Erdelyi 1995] for efficiency.

**Experimental methodology and setup.** We compare RLIBM-ALL's functions with Intel's libm, glibc's libm, CR-LIBM [Daramy-Loirat et al. 2006], and RLIBM-32. Among these, CR-LIBM provides four implementations for each elementary function that produces the correctly rounded results in double precision with the  $rn$ ,  $rz$ ,  $ru$ , and  $rd$  mode, respectively. CR-LIBM does not provide implementations for the  $ra$  mode. RLIBM-32 provides correctly rounded functions for a 32-bit float with the  $rn$  mode. To produce the result in a target representation  $\mathbb{T}$  that is not natively supported by these libraries, we first convert the input in  $\mathbb{T}$  to the representation supported by the library, use the elementary function, and round the result back to  $\mathbb{T}$ . We perform our experiments on a 2.10GHz Intel Xeon Gold 6230R machine with 192GB of RAM running Ubuntu 18.04. We disabled Intel turbo boost and hyper-threading to minimize noise. All our libraries are compiled with O3 optimizations. We use Intel's libm from the oneAPI Toolkit and glibc's libm from glibc-2.33. The test harness

Table 1. Details about the generated polynomials. For each function, we show the time taken to generate the polynomial in minutes, the size of the piecewise polynomial, the maximum degree, the number of terms, and whether the generated polynomial produces correct results in FP34 using the round-to-odd mode for all inputs.

$f(x)$	Gen. Time (Min.)	# of Polynomials	Degree	# of Terms	FP34 <i>ro</i>
$\ln(x)$	325	$2^{10}$	3	3	✓
$\log_2(x)$	420	$2^8$	3	3	✓
$\log_{10}(x)$	546	$2^8$	3	3	✓
$e^x$	241	$2^7$	4	5	✓
		$2^7$	4	5	
$10^x$	402	$2^8$	3	4	✓
		$2^8$	3	4	
$\sinh(x)$	143	$2^6$	5	3	✓
$\cosh(x)$	135	$2^5$	4	3	✓
$\sinpi(x)$	308	$2^2$	5	3	✓
$\cospi(x)$	316	$2^2$	4	3	✓
$2^x$	151	$2^7$	3	4	✓
		$2^7$	3	4	✓

for comparing glibc's libm, CR-LIBM, and RLIBM-32 is built using the gcc-10 compiler with `-O3 -static -frounding-math -fsignaling-nans` flags. Because Intel's libm is only supported in the Intel's compiler, we built a test harness that compares Intel's libm against RLIBM-ALL using the icc compiler with `-O3 -static -no-ftz -fp-model strict` flags to obtain as many accurate results as possible. To compare performance, we measure the number of cycles taken to compute the result for each input using `rtdscp`. We then measured the total time taken to compute the elementary function as the sum of the time taken by all inputs.

## 6.1 Polynomial Generation with RLIBM-ALL

Table 1 provides details on the properties of the polynomials generated by RLIBM-ALL. Our attempt was to generate piecewise polynomials with degree less than or equal to 8. We also restricted the number of sub-domains for the piecewise polynomials to  $2^{15}$ . The output compensation function for  $\sinh(x)$ ,  $\cosh(x)$ ,  $\sinpi(x)$ , and  $\cospi(x)$  uses two elementary functions and we generate two piecewise polynomials for each function. The  $e^x$ ,  $2^x$ , and  $10^x$  functions have both negative and positive reduced inputs. Hence, we create two piecewise polynomials: one for the negative reduced inputs and the other for the positive reduced inputs. As RLIBM-ALL generates piecewise polynomials for FP34 with the round-to-odd mode, the number of sub-domains used in the resulting piecewise polynomials are bigger than RLIBM-32. However, the degrees of the polynomial for each sub-domain was similar to RLIBM-32. The amount of time taken to generate the piecewise polynomials ranged from approximately 2 hours to 9 hours. About 79% of the total time on average is spent in computing the oracle result using the MPFR library. In contrast, computing the intervals and generating the polynomials using the LP solver takes 15% and 5% of the total time on average, respectively.

## 6.2 Does RLIBM-ALL Produce Correct Results?

We experimentally show that RLIBM-ALL produces correctly rounded results for all rounding modes for multiple representations. We built a harness that checks if RLIBM-ALL's functions produce correctly rounded results for all inputs with 161 different FP representations where the number of exponent bits ranged from 2 to 8 and the number of mantissa bits ranged from 1 to 23 bits (*i.e.*,  $23 * 7 = 161$ ) for all five rounding modes. RLIBM-ALL produces the correct results for all these representations with all standard rounding modes with all inputs. *RLIBM-ALL is the first efficient library that provides correctly rounded results for all rounding modes for a 32-bit float.*

**Correct results with all rounding modes for a 32-bit float.** Table 2 reports the results of our experiments to check whether existing libraries produce correct results for a 32-bit float type.

Table 2. Ability to generate correct results for a 32-bit float for all inputs with each of the five standard rounding modes with RLIBM-ALL, glibc’s double libm, Intel’s double libm, CR-LIBM, and RLIBM-32. CR-LIBM provides separate correctly rounded functions for each rounding mode except the *ra* mode, which we use to check its correctness. ✓ indicates that the library produces the correct result using a given rounding mode for all inputs. Otherwise, we use ✗.

	Using RLIBM-ALL					Using glibc double libm					Using Intel double libm				
$f(x)$	<i>rn</i>	<i>rd</i>	<i>ru</i>	<i>rz</i>	<i>ra</i>	<i>rn</i>	<i>rd</i>	<i>ru</i>	<i>rz</i>	<i>ra</i>	<i>rn</i>	<i>rd</i>	<i>ru</i>	<i>rz</i>	<i>ra</i>
$\ln(x)$	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
$\log_2(x)$	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
$\log_{10}(x)$	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
$e^x$	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗	✗	✓
$2^x$	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
$10^x$	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✗	✗	✗	✓
$\sinh(x)$	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
$\cosh(x)$	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗	✗	✓
$\sinpi(x)$	✓	✓	✓	✓	✓	N/A	N/A	N/A	N/A	N/A	✓	✓	✓	✓	✗
$\cospi(x)$	✓	✓	✓	✓	✓	N/A	N/A	N/A	N/A	N/A	✓	✗	✓	✗	✗

	Using CRLIBM					Using RLIBM-32				
$f(x)$	<i>rn</i>	<i>rd</i>	<i>ru</i>	<i>rz</i>	<i>ra</i>	<i>rn</i>	<i>rd</i>	<i>ru</i>	<i>rz</i>	<i>ra</i>
$\ln(x)$	✗	✓	✓	✓	N/A	✓	✗	✗	✗	✓
$\log_2(x)$	✓	✓	✓	✓	N/A	✓	✗	✗	✗	✓
$\log_{10}(x)$	✗	✓	✓	✓	N/A	✓	✗	✗	✗	✓
$e^x$	✓	✓	✓	✓	N/A	✓	✗	✗	✗	✓
$2^x$	N/A	N/A	N/A	N/A	N/A	✓	✗	✗	✗	✗
$10^x$	N/A	N/A	N/A	N/A	N/A	✓	✗	✗	✗	✓
$\sinh(x)$	✗	✓	✓	✓	N/A	✓	✗	✗	✗	✓
$\cosh(x)$	✓	✓	✓	✓	N/A	✓	✗	✗	✗	✓
$\sinpi(x)$	✓	✓	✓	✓	N/A	✓	✗	✗	✗	✓
$\cospi(x)$	✓	✓	✓	✓	N/A	✓	✗	✗	✗	✓

While RLIBM-ALL produces the correct float results with all five rounding modes, mainstream libraries (glibc and Intel’s libm) do not produce correctly rounded results for all rounding modes for all inputs for many of the elementary functions. When CR-LIBM’s *rn* implementations, which produces correctly rounded double results with the *rn* mode for all inputs, is used to produce the *rn* results for a 32-bit float, it does not produce correctly rounded results for all inputs with several functions due to double rounding. CR-LIBM’s *rd*, *ru*, and *rz* implementations produce correctly rounded float results for the *rd*, *ru*, and *rz* mode respectively. Double rounding with these three rounding modes using CR-LIBM’s implementations do not generate wrong results.

RLIBM-32 produces correctly rounded results for all inputs for the *rn* mode with a 32-bit float. However, it does not produce correct results for other rounding modes. In contrast, RLIBM-ALL’s produces a single polynomial approximation for an elementary function that produces correct results for all inputs and for all rounding modes.

**Correct results with all rounding modes for tensorfloat32.** Tensorfloat32 is a new 19-bit representation with the same number of exponents bits as a 32-bit float. Tensorfloat32 has the same number of exponent bits as RLIBM-ALL’s FP34. Hence, RLIBM-ALL produces correctly rounded results for all inputs and for all rounding modes with tensorfloat32. Table 3 shows that glibc’s libm, Intel’s libm, and RLibm32 do not produce correct results for all rounding modes for all ten elementary functions. Although glibc’s and Intel’s libm are designed to produce double results, which has significantly higher precision than tensorfloat32, the double rounding error still results in incorrect

Table 3. Ability to generate correct results with tensorflow32 for all inputs with various rounding modes. ✓ indicates that the library produces the correct tensorflow32 result using a given rounding mode for all inputs. Otherwise, we use ✗.

	Using RLIBM-ALL					Using glibc double libm					Using Intel double libm				
$f(x)$	$rn$	$rd$	$ru$	$rz$	$ra$	$rn$	$rd$	$ru$	$rz$	$ra$	$rn$	$rd$	$ru$	$rz$	$ra$
$\ln(x)$	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
$\log_2(x)$	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
$\log_{10}(x)$	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
$e^x$	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗	✗	✓
$2^x$	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗	✗	✓
$10^x$	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗	✗	✓
$\sinh(x)$	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗	✗	✓
$\cosh(x)$	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗	✗	✓
$\sinpi(x)$	✓	✓	✓	✓	✓	N/A	N/A	N/A	N/A	N/A	✓	✓	✓	✓	✓
$\cospi(x)$	✓	✓	✓	✓	✓	N/A	N/A	N/A	N/A	N/A	✓	✗	✓	✗	✓

	Using CRLIBM					Using RLIBM-32				
$f(x)$	$rn$	$rd$	$ru$	$rz$	$ra$	$rn$	$rd$	$ru$	$rz$	$ra$
$\ln(x)$	✓	✓	✓	✓	N/A	✗	✗	✗	✗	✗
$\log_2(x)$	✓	✓	✓	✓	N/A	✓	✓	✓	✓	✓
$\log_{10}(x)$	✓	✓	✓	✓	N/A	✗	✗	✗	✗	✗
$e^x$	✓	✓	✓	✓	N/A	✗	✗	✗	✗	✗
$2^x$	N/A	N/A	N/A	N/A	N/A	✗	✗	✗	✗	✓
$10^x$	N/A	N/A	N/A	N/A	N/A	✓	✗	✗	✗	✗
$\sinh(x)$	✓	✓	✓	✓	N/A	✗	✗	✗	✗	✗
$\cosh(x)$	✓	✓	✓	✓	N/A	✗	✗	✗	✗	✗
$\sinpi(x)$	✓	✓	✓	✓	N/A	✓	✗	✗	✗	✓
$\cospi(x)$	✓	✓	✓	✓	N/A	✓	✗	✗	✗	✓

tensorflow32 results for  $rd$ ,  $ru$ , and  $rz$  rounding modes especially with extremal values. In contrast, CR-LIBM's implementation designed for each rounding mode  $rm$  produces correctly rounded results for all inputs with the same rounding mode  $rm$ . However, using CR-LIBM's implementation for a specific rounding mode to produce the results for another rounding mode results in wrong results. RLIBM-ALL is the first collection of elementary functions for tensorflow32 that produces correct results for all inputs and for all rounding modes with a single polynomial approximation.

### 6.3 Performance Evaluation of RLIBM-ALL's Functions

Figure 13 reports the speedup of RLIBM-ALL's functions over various mainstream libraries (glibc's libm and Intel's libm) and correctly rounded libraries (CR-LIBM and RLIBM-32). Figure 13(a) presents the speedup of RLIBM-ALL's FP functions over glibc's float functions (left bar in each cluster) and double functions (right bar in each cluster). On average, RLIBM-ALL's FP functions are 1.05× and 1.1× faster than glibc's float and double functions, respectively. Figure 13(b) presents the speedup of RLIBM-ALL's FP functions over Intel's float functions (left bar in each cluster) and double functions (right bar in each cluster). On average, RLIBM-ALL has 1.34× and 1.46× speedup over Intel's float and double functions, respectively. Figure 13(c) presents the speedup of RLIBM-ALL's FP functions over CR-LIBM functions. On average, RLIBM-ALL has 1.86× speedup over CR-LIBM functions. In contrast to RLIBM-ALL, glibc's libm, Intel's libm, and CR-LIBM do not produce correct results for all inputs when used for a 32-bit float type.

Figure 13(d) presents the speedup of RLIBM-ALL's FP functions over RLIBM-32's functions in producing 32-bit float values rounded with the  $rn$  rounding mode. On average, RLIBM-ALL is almost

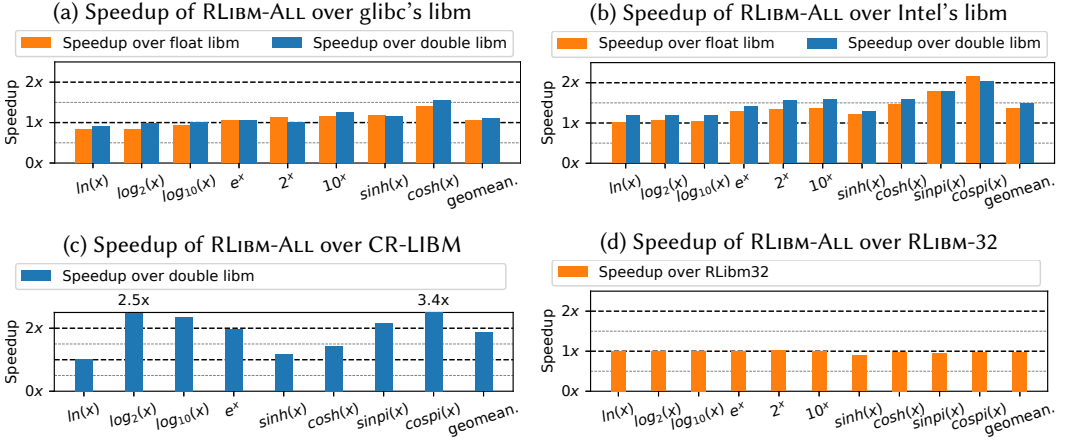


Fig. 13. (a) Speedup of RLIBM-ALL's functions compared to glibc's float functions (left) and glibc's double functions (right) when producing 32-bit float results. (b) Speedup of RLIBM-ALL's functions compared to Intel's float functions (left) and Intel's double functions (right) when producing 32-bit float results. (c) Speedup of RLIBM-ALL's functions compared to CR-LIBM functions when producing 32-bit float results. (d) Speedup of RLIBM-ALL's functions compared to RLIBM-32 functions when producing 32-bit float results.

as fast as RLIBM-32 (*i.e.*, 2% slower than RLIBM-32). RLIBM-ALL creates polynomial approximations for a 34-bit FP representation, which is the main reason for this small performance slowdown. Notably, RLIBM-ALL experiences roughly a 12% slowdown with  $\sinh(x)$  when compared to RLIBM-32. When an input  $x$  is near 0,  $\sinh(x)$  exhibits a linear behavior and  $\sinh(x) \approx x$  produces correctly rounded float values with the *rn* rounding mode. RLIBM-32 uses this property by simply returning  $x$  for inputs near 0. In contrast, RLIBM-ALL performs significantly more computation to ensure that it produces a value within the odd interval. Unlike RLIBM-32 that produces correct results for a single representation with the *rn* mode, RLIBM-ALL produces correct results for all inputs for multiple representations and all the standard rounding modes.

## 7 RELATED WORK

Seminal research over multiple decades has advanced the state-of-the-art for creating polynomial approximations [Bui and Tahar 1999; Daramy et al. 2003; Fousse et al. 2007; Jeannerod et al. 2011; Muller 2005; Remes 1934; Trefethen 2012; Ziv 1991]. Important advances in range reduction has made such approximation feasible [Boldo et al. 2009; Cody and Waite 1980; Story and Ping Tak Peter Tang 1999; Tang 1989, 1990; Tang 1991]. Simultaneously, there are verification efforts to prove bounds for math libraries [Harrison 1997a,b, 2009; Lee et al. 2017; Sawada 2002] and repair individual outputs of math libraries [Yi et al. 2019; Zou et al. 2019]. The comprehensive book on elementary functions provides detailed information on prior work [Muller 2005].

We restrict our comparison to prior work that is closely related to our work. As a correctly rounded elementary function is recommended by the IEEE standard and enables portability, a number of correctly rounded math libraries have been developed [Daramy et al. 2003; Lim et al. 2020, 2021; Lim and Nagarakatte 2021a; Ziv 1991]. They are restricted to a specific representation and a rounding mode.

CR-LIBM [Daramy et al. 2003; Lefèvre et al. 1998] is a correctly rounded collection of elementary functions for double precision. It was developed using Sollya [Chevallard et al. 2010], which is a tool



and a library for developing FP code. Sollya can generate polynomials of degree  $d$  with coefficients in a representation used for the implementation ( $\mathbb{H}$ ) that has the minimum infinity norm [Brisebarre and Chevillard 2007]. Sollya uses a modified Remez algorithm to produce polynomials. It also computes and proves the error bound on the polynomial evaluation using interval arithmetic [Chevillard et al. 2011; Chevillard and Lauter 2007]. Metalibm [Brunie et al. 2015; Kupriianova and Lauter 2014] is a customization infrastructure also built using Sollya. MetaLibm is able to automatically identify range reduction and domain splitting techniques for some transcendental functions. It has been used to trade-off correctness and performance while approximating elementary functions for float and double precision types.

A modified Remez algorithm has also been used to generate polynomials that minimizes the infinity norm compared to an ideal elementary function [Arzelier et al. 2019]. It can be useful for generating correctly rounded results for a specific precision and a rounding mode when range reduction is not necessary.

This paper builds on our prior work in the RLIBM project [Lim 2021; Lim et al. 2020, 2021; Lim and Nagarakatte 2021a,d] that creates polynomials using the correctly rounded value rather than the real value of the elementary function. Like the RLIBM project, we structure the problem of generating polynomials as an LP problem. We also use RLIBM's range reduction strategies. The RLIBM project has generated correctly rounded libraries with the commonly used  $rn$  mode for multiple types: bfloat16, posit16, 32-bit float, and posit32. However, it is necessary to create individual polynomial approximation for each representation with each rounding mode with the RLIBM project to avoid double rounding errors. In contrast, this paper shows that by generating polynomial approximations for  $\mathbb{T}_{n+2}$  with the round-to-odd mode, we can create a single polynomial approximation that works for multiple representations  $\mathbb{T}_k$  with multiple rounding modes.

## 8 CONCLUSION

This paper proposes a novel method to generate a single polynomial approximation that produces correctly rounded results for multiple representations and rounding modes. The key idea is to create a polynomial approximation that produces the correctly rounded result for  $\mathbb{T}_{n+2}$  with the round-to-odd mode when the goal is to generate correct results for  $\mathbb{T}_k$ , where  $k \leq n$ , with all rounding modes. We address the issue of singletons while generating polynomials that approximate the correctly rounded result with the round-to-odd mode. We provide the first correctly rounded implementations of elementary functions for multiple representations. We believe that our results make a strong case for mandating correctly rounded results at least with any representation that has fewer than or equal to 32-bits.

## ACKNOWLEDGMENTS

We thank John Gustafson for his suggestion that we explore a generic math library for posits and his inputs on the Minefield method and the posit representation. We thank our shepherd James Wilcox, Mridul Aanjaneya, and the POPL reviewers for their feedback on the paper. This material is based upon work supported in part by the National Science Foundation under Grant No. 1908798, Grant No. 2110861, and Grant No. 1917897. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## REFERENCES

- Martin Aigner and Gnter M. Ziegler. 2009. *Proofs from THE BOOK* (4th ed.). Springer Publishing Company, Incorporated.
- Denis Arzelier, Florent Bréhard, and Mioara Joldes. 2019. Exchange Algorithm for Evaluation and Approximation Error-Optimized Polynomials. In *2019 IEEE 26th Symposium on Computer Arithmetic (ARITH)*. 30–37. <https://doi.org/10.1109/>

ARITH.2019.00014

- Alan Baker. 1975. *Transcendental Number Theory*. Cambridge University Press.
- Sylvie Boldo, Marc Daumas, and Ren-Cang Li. 2009. Formally Verified Argument Reduction with a Fused Multiply-Add. In *IEEE Transactions on Computers*, Vol. 58. 1139–1145. <https://doi.org/10.1109/TC.2008.216>
- Sylvie Boldo and Guillaume Melquiond. 2005. When double rounding is odd. In *17th IMACS World Congress*. Paris, France, 11 pages.
- Sylvie Boldo and Guillaume Melquiond. 2008. Emulation of a FMA and Correctly Rounded Sums: Proved Algorithms Using Rounding to Odd. *IEEE Transactions on Computing* 57, 4 (April 2008), 462–471. <https://doi.org/10.1109/TC.2007.70819>
- Peter Borwein and Tamas Erdelyi. 1995. *Polynomials and Polynomial Inequalities*. Springer New York. <https://doi.org/10.1007/978-1-4612-0793-1>
- Nicolas Brisebarre and Sylvain Chevillard. 2007. Efficient polynomial L-approximations. In *18th IEEE Symposium on Computer Arithmetic (ARITH '07)*. <https://doi.org/10.1109/ARITH.2007.17>
- Nicolas Brunie, Florent de Dinechin, Olga Kupriianova, and Christoph Lauter. 2015. Code Generators for Mathematical Functions. In *2015 IEEE 22nd Symposium on Computer Arithmetic*. 66–73. <https://doi.org/10.1109/ARITH.2015.22>
- Hung Tien Bui and Sofiene Tahar. 1999. Design and synthesis of an IEEE-754 exponential function. In *Engineering Solutions for the Next Millennium. 1999 IEEE Canadian Conference on Electrical and Computer Engineering*, Vol. 1. 450–455 vol.1. <https://doi.org/10.1109/CCECE.1999.807240>
- Sylvain Chevillard, John Harrison, Mioara Joldes, and Christoph Lauter. 2011. Efficient and accurate computation of upper bounds of approximation errors. In *Theoretical Computer Science*, Vol. 412. <https://doi.org/10.1016/j.tcs.2010.11.052>
- Sylvain Chevillard, Mioara Joldes, and Christoph Lauter. 2010. Sollya: An Environment for the Development of Numerical Codes. In *Mathematical Software - ICMS 2010 (Lecture Notes in Computer Science, Vol. 6327)*. Springer, Heidelberg, Germany, 28–31. [https://doi.org/10.1007/978-3-642-15582-6\\_5](https://doi.org/10.1007/978-3-642-15582-6_5)
- Sylvain Chevillard and Christopher Lauter. 2007. A Certified Infinite Norm for the Implementation of Elementary Functions. In *Seventh International Conference on Quality Software (QSIC 2007)*. 153–160. <https://doi.org/10.1109/QSIC.2007.4385491>
- William J Cody and William M Waite. 1980. *Software manual for the elementary functions*. Prentice-Hall, Englewood Cliffs, NJ. <https://doi.org/10.1137/1024023>
- P. M. (Paul Moritz) Cohn. 1974. *Algebra [by] P. M. Cohn*. Wiley, London.
- Catherine Daramy, David Defour, Florent Dinechin, and Jean-Michel Muller. 2003. CR-LIBM: A correctly rounded elementary function library. In *Proceedings of SPIE Vol. 5205: Advanced Signal Processing Algorithms, Architectures, and Implementations XIII*, Vol. 5205. <https://doi.org/10.1117/12.505591>
- Catherine Daramy-Loirat, David Defour, Florent de Dinechin, Matthieu Gallet, Nicolas Gast, Christoph Lauter, and Jean-Michel Muller. 2006. *CR-LIBM A library of correctly rounded elementary functions in double-precision*. Research Report. Laboratoire de l'Informatique du Parallélisme. <https://hal-ens-lyon.archives-ouvertes.fr/ensl-01529804>
- 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, Article 13 (June 2007). <https://doi.org/10.1145/1236463.1236468>
- Ambros M. Gleixner, Daniel E. Steffy, and Kati Wolter. 2012. Improving the Accuracy of Linear Programming Solvers with Iterative Refinement. In *Proceedings of the 37th International Symposium on Symbolic and Algebraic Computation (Grenoble, France) (ISSAC '12)*. Association for Computing Machinery, New York, NY, USA, 187–194. <https://doi.org/10.1145/2442829.2442858>
- David Goldberg. 1991. What Every Computer Scientist Should Know About Floating-point Arithmetic. In *ACM Computing Surveys*, Vol. 23. ACM, New York, NY, USA, 5–48. <https://doi.org/10.1145/103162.103163>
- John Gustafson and Isaac Yonemoto. 2017. Beating Floating Point at Its Own Game: Posit Arithmetic. *Supercomputing Frontiers and Innovations: an International Journal* 4, 2 (June 2017), 71–86. <https://doi.org/10.14529/jsfi170206>
- John Harrison. 1997a. Floating point verification in HOL light: The exponential function. In *Algebraic Methodology and Software Technology*, Michael Johnson (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 246–260. <https://doi.org/10.1007/BFb0000475>
- John Harrison. 1997b. Verifying the Accuracy of Polynomial Approximations in HOL. In *International Conference on Theorem Proving in Higher Order Logics*. <https://doi.org/10.1007/BFb0028391>
- John Harrison. 2009. HOL Light: An Overview. In *Proceedings of the 22nd International Conference on Theorem Proving in Higher Order Logics, TPHOLS 2009 (Lecture Notes in Computer Science, Vol. 5674)*, Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel (Eds.). Springer-Verlag, Munich, Germany, 60–66. [https://doi.org/10.1007/978-3-642-03359-9\\_4](https://doi.org/10.1007/978-3-642-03359-9_4)
- Claude-Pierre Jeannerod, Hervé Knochel, Christophe Monat, and Guillaume Revy. 2011. Computing Floating-Point Square Roots via Bivariate Polynomial Evaluation. *IEEE Trans. Comput.* 60. <https://doi.org/10.1109/TC.2010.152>
- Jeff Johnson. 2018. Rethinking floating point for deep learning. arXiv:1811.01721

- Dhiraj D. Kalamkar, Dheevatsa Mudigere, Naveen Mellempudi, Dipankar Das, Kunal Banerjee, Sasikanth Avancha, Dharma Teja Vooturi, Nataraj Jammalamadaka, Jianyu Huang, Hector Yuen, Jiyan Yang, Jongsoo Park, Alexander Heinecke, Evangelos Georganas, Sudarshan Srinivasan, Abhisek Kundu, Misha Smelyanskiy, Bharat Kaul, and Pradeep Dubey. 2019. A Study of BFLOAT16 for Deep Learning Training. arXiv:1905.12322
- Urs Köster, Tristan Webb, Xin Wang, Marcel Nassar, Arjun K. Bansal, William Constable, Oguz Elibol, Stewart Hall, Luke Hornof, Amir Khosrowshahi, Carey Kloss, Ruby J. Pai, and Naveen Rao. 2017. Flexpoint: An Adaptive Numerical Format for Efficient Training of Deep Neural Networks. In *Advances in Neural Information Processing Systems*, Vol. abs/1711.02213.
- Olga Kupriianova and Christoph Lauter. 2014. Metalibm: A Mathematical Functions Code Generator. In *4th International Congress on Mathematical Software*. [https://doi.org/10.1007/978-3-662-44199-2\\_106](https://doi.org/10.1007/978-3-662-44199-2_106)
- Wonyeol Lee, Rahul Sharma, and Alex Aiken. 2017. On Automatically Proving the Correctness of Math.h Implementations. *Proceedings of the ACM on Programming Languages* 2, POPL, Article 47 (Dec. 2017), 32 pages. <https://doi.org/10.1145/3158135>
- Vincent Lefèvre, Jean-Michel Muller, and Arnaud Tisserand. 1998. Toward correctly rounded transcendentals. *IEEE Trans. Comput.* 47, 11 (1998), 1235–1243. <https://doi.org/10.1109/12.736435>
- Jay Lim. 2021. *Novel Polynomial Approximation Methods for Generating Correctly Rounded Elementary Functions*. Ph. D. Dissertation. Rutgers University.
- Jay P. Lim, Mridul Aanjaneya, John Gustafson, and Santosh Nagarakatte. 2020. A Novel Approach to Generate Correctly Rounded Math Libraries for New Floating Point Representations. arXiv:2007.05344 Rutgers Department of Computer Science Technical Report DCS-TR-753.
- Jay P. Lim, Mridul Aanjaneya, John Gustafson, and Santosh Nagarakatte. 2021. An Approach to Generate Correctly Rounded Math Libraries for New Floating Point Variants. *Proceedings of the ACM on Programming Languages* 6, POPL, Article 29 (Jan. 2021), 30 pages. <https://doi.org/10.1145/3434310>
- Jay P. Lim and Santosh Nagarakatte. 2021a. High Performance Correctly Rounded Math Libraries for 32-bit Floating Point Representations. In *42nd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'21)*. <https://doi.org/10.1145/3453483.3454049>
- Jay P Lim and Santosh Nagarakatte. 2021b. RLIBM-32: High Performance Correctly Rounded Math Libraries for 32-bit Floating Point Representations. arXiv:2104.04043 Rutgers Department of Computer Science Technical Report DCS-TR-754.
- Jay P. Lim and Santosh Nagarakatte. 2021c. *RLibm-all*. <https://github.com/rutgers-apl/rlibm-all>
- Jay P. Lim and Santosh Nagarakatte. 2021d. RLIBM-ALL: A Novel Polynomial Approximation Method to Produce Correctly Rounded Results for Multiple Representations and Rounding Modes. arXiv:2108.06756 [abs] Rutgers Department of Computer Science Technical Report DCS-TR-757.
- Jean-Michel Muller. 2005. *Elementary Functions: Algorithms and Implementation*. Birkhauser. <https://doi.org/10.1007/978-1-4899-7983-4>
- Ivan Niven. 1956. *Irrational Numbers*. Mathematical Association of America.
- NVIDIA. 2020. *TensorFloat-32 in the A100 GPU Accelerates AI Training, HPC up to 20x*. <https://blogs.nvidia.com/blog/2020/05/14/tensorfloat-32-precision-format/>
- Eugene Remes. 1934. Sur un procédé convergent d'approximations successives pour déterminer les polynômes d'approximation. *Comptes rendus de l'Académie des Sciences* 198 (1934), 2063–2065.
- Bitra Rouhani, Daniel Lo, Ritchie Zhao, Ming Liu, Jeremy Fowers, Kalin Ovtcharov, Anna Vinogradsky, Sarah Massengill, Lita Yang, Ray Bittner, Alessandro Forin, Haishan Zhu, Taesik Na, Prerak Patel, Shuai Che, Lok Chand Koppaka, Xia Song, Subhojit Som, Kaustav Das, Saurabh Tiwary, Steve Reinhardt, Sitaram Lanka, Eric Chung, and Doug Burger. 2020. Pushing the Limits of Narrow Precision Inferencing at Cloud Scale with Microsoft Floating Point. In *The Thirty-fourth Annual Conference on Neural Information Processing Systems*. ACM.
- Jun Sawada. 2002. Formal verification of divide and square root algorithms using series calculation. In *3rd International Workshop on the ACL2 Theorem Prover and its Applications*.
- Shane Story and Ping Tak Peter Tang. 1999. New algorithms for improved transcendental functions on IA-64. In *Proceedings 14th IEEE Symposium on Computer Arithmetic*. 4–11. <https://doi.org/10.1109/ARITH.1999.762822>
- Giuseppe Tagliavini, Stefan Mach, Davide Rossi, Andrea Marongiu, and Luca Benin. 2018. A transprecision floating-point platform for ultra-low power computing. In *2018 Design, Automation Test in Europe Conference Exhibition (DATE)*. 1051–1056. <https://doi.org/10.23919/DATE.2018.8342167>
- Ping-Tak Peter Tang. 1989. Table-Driven Implementation of the Exponential Function in IEEE Floating-Point Arithmetic. *ACM Trans. Math. Software* 15, 2 (June 1989), 144–157. <https://doi.org/10.1145/63522.214389>
- Ping-Tak Peter Tang. 1990. Table-Driven Implementation of the Logarithm Function in IEEE Floating-Point Arithmetic. *ACM Trans. Math. Software* 16, 4 (Dec. 1990), 378–400. <https://doi.org/10.1145/98267.98294>
- P. T. P. Tang. 1991. Table-lookup algorithms for elementary functions and their error analysis. In *[1991] Proceedings 10th IEEE Symposium on Computer Arithmetic*. 232–236. <https://doi.org/10.1109/ARITH.1991.145565>

- Lloyd N. Trefethen. 2012. *Approximation Theory and Approximation Practice (Other Titles in Applied Mathematics)*. Society for Industrial and Applied Mathematics, USA.
- Xin Yi, Liqian Chen, Xiaoguang Mao, and Tao Ji. 2019. Efficient Automated Repair of High Floating-Point Errors in Numerical Libraries. *Proceedings of the ACM on Programming Languages* 3, POPL, Article 56 (Jan. 2019), 29 pages. <https://doi.org/10.1145/3290369>
- Abraham Ziv. 1991. Fast Evaluation of Elementary Mathematical Functions with Correctly Rounded Last Bit. *ACM Trans. Math. Software* 17, 3 (Sept. 1991), 410–423. <https://doi.org/10.1145/114697.116813>
- Daming Zou, Muhan Zeng, Yingfei Xiong, Zhoulai Fu, Lu Zhang, and Zhendong Su. 2019. Detecting Floating-Point Errors via Atomic Conditions. *Proceedings of the ACM on Programming Languages* 4, POPL, Article 60 (Dec. 2019), 27 pages. <https://doi.org/10.1145/3371128>