# SLEEF: A Portable Vectorized Library of $C$ Standard Mathematical Functions 

Naoki Shibata ${ }^{\odot}$, Member, IEEE and Francesco Petrogalli ${ }^{\odot}$


#### Abstract

In this article, we present techniques used to implement our portable vectorized library of $C$ standard mathematical functions written entirely in C language. In order to make the library portable while maintaining good performance, intrinsic functions of vector extensions are abstracted by inline functions or preprocessor macros. We implemented the functions so that they can use sub-features of vector extensions such as fused multiply-add, mask registers, and extraction of mantissa. In order to make computation with SIMD instructions efficient, the library only uses a small number of conditional branches, and all the computation paths are vectorized. We devised a variation of the Payne-Hanek argument reduction for trigonometric functions and a floating point remainder, both of which are suitable for vector computation. We compare the performance with our library to Intel SVML.


Index Terms—Parallel and vector implementations, SIMD processors, elementary functions, floating-point arithmetic

## 1 Introduction

THE instruction set architecture of most modern processors provides Single Instruction Multiple Data (SIMD) instructions that process multiple instances of data concurrently [1]. The programming model that utilizes these instructions is a key technique for many computing systems to reach their peak performance. Most software SIMD optimizations are introduced manually by programmers. However, this approach introduces a portability problem because the code needs to be re-written when targeting a new vector extension. In order to improve portability of codes with SIMD optimizations, recent compilers have introduced auto-vectorizing capability [2]. To fully exploit the SIMD capabilities of a system, the transformation for auto-vectorization of a compiler must be able to invoke a version of functions that operates on concurrent iterations, or on a vector function. This applies particularly to C mathematical functions defined in math. $h$ that are frequently called in hot-loops.

In this paper, we describe our implementation of a vectorized library of C standard math functions, called SLEEF library. SLEEF stands for SIMD Library for Evaluating Elementary Functions, and implements a vectorized version of all C99 real floating-point math functions. Our library provides 1-ULP accuracy version and 3.5-ULP accuracy version for most of the functions. We confirmed that our library satisfies such accuracy requirements on an empirical basis. Our library achieves both good performance and

- N. Shibata is with Graduate School of Information Science, Nara Institute of Science and Technology, Nara 630-0192, Japan.
E-mail: $n$-sibata@is.naist.jp.
- F. Petrogalli is with ARM110 CB1 9NJ Cambridge, United Kingdom. E-mail: Francesco.Petrogalli@arm.com.
Manuscript received 25 Apr. 2019; revised 11 Dec. 2019; accepted 13 Dec. 2019. Date of publication 18 Dec. 2019; date of current version 27 Jan. 2020. (Corresponding author: Naoki Shibata.)
Recommended for acceptance by D. S. Nikolopoulos.
Digital Object Identifier no. 10.1109/TPDS.2019.2960333
portability by abstracting intrinsic functions. This abstraction enables sub-features of vector extensions such as mask registers to be utilized while the source code of our library is shared among different vector extensions. We also implemented a version of functions that returns bit-wise consistent results across all platforms. Our library is designed to be used in conjunction with vectorizing compilers. In order to help development of vectorizing compilers, we collaborated with compiler developers in designing a Vector Function Application Binary Interface (ABI). The main difficulty in vectorizing math functions is that conditional branches are expensive. We implemented many of the functions in our library without conditional branches. We devised reduction methods and adjusted domains of polynomials so that a single polynomial covers the entire input domain. For an increased vector size, a value requiring a slow path is more likely to be contained in a vector. Therefore, we vectorized all the code paths in order to speed up the computation in such cases. We devised a variation of the Payne-Hanek range reduction and a remainder calculation method that are both suitable for vectorized implementation.

We compare the implementation of several selected functions in our library to those in other open-source libraries. We also compare the reciprocal throughput of functions in our library, Intel SVML [3], FDLIBM [4], and Vectorlibm [5]. We show that the performance of our library is comparable to that of Intel SVML.

The rest of this paper is organized as follows. Section 2 introduces related work. Section 3 discusses how portability is improved by abstracting vector extensions. Section 4 explains the development of a Vector ABI and a vectorized mathematical library. Section 5 shows an overview of the implementation of SLEEF, while comparing our library with FDLIBM and Vector-libm. Section 6 explains how our library is tested. Section 7 compares our work with prior art. In Section 8, the conclusions are presented.

## 2 Related Work

### 2.1 C Standard Math Library

The C standard library (libc) includes the standard mathematical library (libm) [6]. There have been many implementations of libm. Among them, FDLIBM [4] and the libm included in the GNU C Library [7] are the most widely used libraries. FDLIBM is a freely distributable libm developed by Sun Microsystems, Inc., and there are many derivations of this library. Gal et al. described the algorithms used in the elementary mathematical library of the IBM Israel Scientific Center [8]. Their algorithms are based on the accurate tables method developed by Gal. It achieves high performance and produces very accurate results. Crlibm is a project to build a correctly rounded mathematical library [9].

There are several existing vectorized implementations of libm. Intel Short Vector Math Library (SVML) is a highly regarded commercial library [3]. This library provides highly optimized subroutines for evaluating elementary functions which can use several kinds of vector extensions available in Intel's processors. However, this library is proprietary and only optimized for Intel's processors. There are also a few commercial and open-source implementations of vectorized libm. AMD is providing a vectorized libm called AMD Core Math Library (ACML) [10].

Some of the code from SVML is published under a free software license, and it is now published as Libmvec [11], which is a part of Glibc. This library provides functions with 4 -ULP error bound. It is coded in assembly language, and therefore it does not have good portability. C. K. Anand et al. reported their C implementation of 32 single precision libm functions tuned for the Cell BE SPU compute engine [12]. They used an environment called Coconut that enables rapid prototyping of patterns, rapid unit testing of assembly language fragments and patterns to develop their library. M. Dukhan published an open-source and portable SIMD vector libm library named Yeppp! [13], [14]. Most of vectorized implementations of libm utilizes assembly coding or intrinsic functions to specify which vector instruction is used for each operator. On the other hand, there are also other implementations of vector versions of libm which are written in a scalar fashion but rely on a vectorizing compiler to generate vector instructions and generate a vectorized binary code. Christoph Lauter published an open-source Vectorlibm library implemented with plain C [5]. VDT Mathematical Library [15], is a math library written for the compiler's auto-vectorization feature.

### 2.2 Translation of SIMD Instructions

Manilov et al. propose a C source code translator for substituting calls to platform-specific intrinsic functions in a source code with those available on the target machine [16]. This technique utilizes graph-based pattern matching to substitute intrinsics. It can translate SIMD intrinsics between extensions with different vector lengths. This rewriting is carried out through loop-unrolling.
N. Gross proposes specialized C++ templates for making the source code easily portable among different vector extensions without sacrificing performance [17]. With these templates, some part of the source code can be written in a way that resembles scalar code. In order to vectorize
algorithms that have a lot of control flow, this scheme requires the bucketing technique is applied, to compute all the paths and choose the relevant results at the end.

Clark et al. proposes a method for combining static analysis at compile time and binary translation with a JIT compiler in order to translate SIMD instructions into those that are available on the target machine [18]. In this method, SIMD instructions in the code are first converted into an equivalent scalar representation. Then, a dynamic translation phase turns the scalar representation back into architecturespecific SIMD equivalents.

Leißa et al. propose a C-like language for portable and efficient SIMD programming [19]. With their extension, writing vectorized code is almost as easy as writing traditional scalar code. There is no strict separation in host code and kernels, and scalar and vector programming can be mixed. Switching between them is triggered by the type system. The authors present a formal semantics of their extension and prove the soundness of the type system.

Most of the existing methods are aiming at translating SIMD intrinsics or instructions to those provided by a different vector extension in order to port a code. Intrinsics that are unique in a specific extension are not easy to handle, and translation works only if the source and the target architectures have equivalent SIMD instructions. Automatic vectorizers in compilers have a similar weakness. Whenever possible, we have specialized the implementation of the math functions to exploit the SIMD instructions that are specific to a target vector extension. We also want to make special handling of FMA, rounding and a few other kinds of instructions, because these are critical for both execution speed and accuracy. We want to implement a library that is statically optimized and usable with Link Time Optimization (LTO). The users of our library do not appreciate usage of a JIT compiler. In order to minimize dependency on external libraries, we want to write our library in C. In order to fulfill these requirements, we take a crosslayer approach. We have been developing our abstraction layer of intrinsics, the library implementation, and the algorithms in order to make our library run fast with any vector extensions.

## 3 Abstraction of Vector Extensions

Modern processors supporting SIMD instructions have SIMD registers that can contain multiple data [1]. For example, a 128 -bit wide SIMD register may contain four 32 -bit single-precision FP numbers. A SIMD add instruction might take two of these registers as operands, add the four pairs of numbers, and overwrite one of these registers with the resulting four numbers. We call an array of FP numbers contained in a SIMD register a vector.

SIMD registers and instruction can be exposed in a C program with intrinsic functions and types [20]. An intrinsic function is a kind of inline function that exposes the architectural features of an instruction set at $C$ level. By calling an intrinsic function, a programmer can make a compiler generate a specific instruction without hand-coded assembly. Nevertheless, the compiler can reorder instructions and allocate registers, and therefore optimize the code. When intrinsic functions corresponding to SIMD instructions are defined inside a compiler, C data types for representing vectors are also defined.

```
// Type definition
typedef svbool_t vopmask;
typedef svfloat64_t vdouble;
// Abstraction of intrinsic functions
#define ptrue svptrue_b8()
vdouble vcast_vd_d(double d) { return svdup_n_f64(d); }
vdouble vsub_vd_vd_vd(vdouble x, vdouble y) {
    return svsub_f64_x(ptrue, x, y); }
vopmask veq_vo_vd_vd(vdouble x, vdouble y) {
    return svcmpeq_f64(ptrue, x, y); }
vopmask vlt_vo_vd_vd(vdouble x, vdouble y) {
    return svcmplt_f64(ptrue, x, y); }
vopmask vor_vo_vo_vo(vopmask x, vopmask y) {
    return svorr_b_z(ptrue, x, y); }
vdouble vsel_vd_vo_vd_vd(vopmask mask, vdouble x, vdouble
        y) {
    return svsel_f64(o, x, y); }
```

Fig. 1. A part of definitions in VEAL for SVE.
In SLEEF, we use intrinsic functions to specify which assembly instruction to use for each operator. We abstract intrinsic functions for each vector extension by a set of inline functions or preprocessor macros. We implement the functions exported from the library to call abstract intrinsic functions instead of directly calling intrinsic functions. In this way, it is easy to swap the vector extension to use. We call our set of inline functions for abstracting architecture-specific intrinsics Vector Extension Abstraction Layer (VEAL).

In some of the existing vector math libraries, functions are implemented with hand-coded assembly [11]. This approach improves the absolute performance because it is possible to provide the optimal implementation for each microarchitecture. However, processors with a new microarchitecture are released every few years, and the library needs revision accordingly in order to maintain the optimal performance.

In other vector math libraries, the source code is written in a scalar fashion that is easy for compilers to auto-vectorize [5], [15]. Although such libraries have good portability, it is not easy for compilers to generate a well-optimized code. In order for each transformation rule in an optimizer to kick in, the source code must satisfy many conditions to guarantee that the optimized code runs correctly and faster. In order to control the level of optimization, a programmer must specify special attributes and compiler options.

### 3.1 Using Sub-Features of the Vector Extensions

There are differences in the features provided by different vector extensions, and we must change the function implementation according to the available features. Thanks to the level of abstraction provided by the VEALs, we implemented the functions so that all the different versions of functions can be built from the same source files with different macros enabled. For example, the availability of FMA instructions is important when implementing double-double (DD) operators [21]. We implemented DD operators both with and without FMA by manually specifying if the compiler can convert each combination of multiplication and addition instructions to an FMA instruction, utilizing VEALs.

Generally, bit masks are used in a vectorized code in order to conditionally choose elements from two vectors. In some vector extensions, a vector register with a width that matches a vector register for storing FP values, is used to store a bit mask. Some vector extensions provide narrower vector registers that are dedicated to this purpose,

```
vdouble xfdim(vdouble x, vdouble y) {
    vdouble ret = vsub_vd_vd_vd(x, y);
    ret = vsel_vd_vo_vd_vd(vor_vo_vo_vo(vlt_vo_vd_vd(ret,
        vcast_vd_d(0)), veq_vo_vd_vd(x, y)),
            vcast_vd_d(0),
            ret);
    return ret;
7 }
```

Fig. 2. Implementation of vectorized fdim (positive difference) function with VEAL.
which is SLEEF makes use of these opmask registers by providing a dedicated data type in VEALs. If a vector extension does not support an opmask, the usual bit mask is used instead of an opmask. It is also better to have an opmask as an argument of a whole math function and make that function only compute the elements specified by the opmask. By utilizing a VEAL, it is also easy to implement such a functionality.

### 3.2 Details of VEALs

Fig. 1 shows some definitions in the VEAL for SVE [22]. We abstract vector data types and intrinsic functions with typedef statements and inline functions, respectively.

The vdouble data type is for storing vectors of double precision FP numbers. vopmask is the data type for the opmask described in Section 3.1.

The function vcast_vd_d is a function that returns a vector in which the given scalar value is copied to all elements in the vector. vsub_vd_vd_vd is a function for vector subtraction between two vdouble data. veq_vo_vd_vd compares elements of two vectors of vdouble type. The results of the comparison can be used, for example, by vsel_vd_vo_vd_vd to choose a value for each element between two vector registers. Fig. 2 shows an implementation of a vectorized positive difference function using a VEAL. This function is a vectorized implementation of the fdim function in the C standard math library.

### 3.3 Making Results Bit-Wise Consistent Across All Platforms

The method of implementing math functions described so far, can deliver computation results that slightly differ depending on architectures and other conditions, although they all satisfy the accuracy requirements, and other specifications. However, in some applications, bit-wise consistent results are required.

To this extent, the SLEEF project has been working closely with Unity Technologies, ${ }^{1}$ which specializes in developing frameworks for video gaming, and we discovered that they have unique requirements for the functionalities of math libraries. Networked video games run on many gaming consoles with different architectures and they share the same virtual environment. Consistent results of simulation at each terminal and server are required to ensure fairness among all players. For this purpose, fast computation is more important than accurate computation, while the results of computation have to perfectly agree between many computing nodes, which are not guaranteed to rely on the same architecture. Usually, fixed-point arithmetic is

1. https://unity3d.com/.
used for a purpose like this, however there is a demand for modifying existing codes with FP computation to support networking.

There are also other kinds of simulation in which bit-wise identical reproducibility is important. In [23], the authors show that modeled mean climate states, variability and trends at different scales may be significantly changed or even lead to opposing results due to the round-off errors in climate system simulations. Since reproducibility is a fundamental principle of scientific research, they propose to promote bit-wise identical reproducibility as a worldwide standard.

One way to obtain bit-wise consistent values from math functions is to compute correctly rounded values. However, for applications like networked video games, this might be too expensive. SLEEF provides vectorized math functions that return bit-wise consistent results across all platforms and other settings, and this is also achieved by utilizing VEALs. The basic idea is to always apply the same sequence of operations to the arguments. The IEEE 754 standard guarantees that the basic arithmetic operators give correctly rounded results [24], and therefore the results from these operators are bit-wise consistent. Because most of the functions except trigonometric functions do not have a conditional branch in our library, producing bit-wise consistent results is fairly straightforward with VEALs. Availability of FMA instructions is another key for making results bit-wise consistent. Since FMA instructions are critical for performance, we cannot just give up using FMA instructions. In SLEEF, the bit-wise consistent versions of functions have two versions both with and without FMA instructions. We provide a non-FMA version of the functions to guarantee bit-wise consistency among extensions such as Intel SSE2 that do not have FMA instructions. Another issue is that the compiler might introduce inconsistency by FP contraction, which is the result of combining a pair of multiplication and addition operations into an FMA. By disabling FP contraction, the compiler strictly preserves the order and the type of FP operations during optimization. It is also important to make the returned values from scalar functions bit-wise consistent with the vector functions. In order to achieve this, we also made a VEAL that only uses scalar operators and data types. The bit-wise consistent and non-consistent versions of vector and scalar functions are all built from the same source files, with different VEALs and macros enabled. As described in Section 5, trigonometric functions in SLEEF chooses a reduction method according to the maximum argument of all elements in the argument vector. In order to make the returned value bit-wise consistent, the bit-wise consistent version of the functions first applies the reduction method for small arguments to the elements covered by this method. Then it applies the second method only to the elements with larger arguments which the first method does not cover.

## 4 The Development of a Vector Function ABI AND SLEEF

Recent compilers are developing new optimization techniques to automatically vectorize a code written in standard programming languages that do not support parallelization [25], [26]. Although the first SIMD and vector computing
systems [27] appeared a few decades ago, compilers with auto-vectorization capability have not been widely used until recently, because of several difficulties in implementing such functionality for modern SIMD architectures. Such difficulties include verifying whether the compiler can vectorize a loop or not, by determining data access patterns of the operations in the loop [2], [28]. For languages like C and $\mathrm{C}++$, it is also difficult to determine the data dependencies through the iteration space of the loop, because it is hard to determine aliasing conditions of the arrays processed in the loop.

### 4.1 Vector Function Application Binary Interface

Vectorizing compilers convert calls to scalar versions of math functions such as sine and exponential to the SIMD version of the math functions. The most recent versions of Intel Compiler [29], GNU Compiler [30], and Arm Compiler for HPC [31], which is based on Clang/LLVM [32], [33], are capable of this transformation, and rely on the availability of vector math libraries such as SVML [3], Libmvec [11] and SLEEF respectively to provide an implementation of the vector function calls that they generate. In order to develop this kind of transformations, a targetdependent Application Binary Interface for calling vectorized functions had to be designed.

The Vector Function ABI for AArch64 architecture [34] was designed in close relationship with the development of SLEEF. This type of ABI must standardize the mapping between scalar functions and vector functions. The existence of a standard enables interoperability across different compilers, linkers and libraries, thanks to the use of standard names defined by the specification.

The ABI includes a name mangling function, a map that converts the scalar signature to the vector one, and the calling conventions that the vector functions must obey. In particular, the name mangling function that takes the name of the scalar function to the vector function must encode all the information that is necessary to reverse the transformation back to the original scalar function. A linker can use this reverse mapping to enable more optimizations (Link Time Optimizations) that operate on object files, and does not have access to the scalar and vector function prototypes. There is a demand by users for using a different vector math library according to the usage. Reverse mapping is also handy for this purpose. A vector math library implements a function for each combination of a vector extension, a vector length and a math function to evaluate. As a result, the library exports a large number of functions. Some vector math libraries can only implement part of all the combinations. By using the reverse mapping mechanism, the compiler can check the availability of the functions by scanning the symbols exported by a library.

The Vector Function ABI is also used with OpenMP [35]. From version 4.0 onwards, OpenMP provides the directive declare simd. A user can decorate a function with this directive to inform the compiler that the function can be safely invoked concurrently on multiple instances of its arguments [36]. This means that the compiler can vectorize the function safely. This is particularly useful when the function is provided via a separate module, or an external library, for example in situations where the compiler is not able to examine the behavior of the function in the call site. The
scalar-to-vector function mapping rules stipulated in the Vector Function ABI are based on the classification of vector functions associated with the declare simd directive of OpenMP. Currently, work for implementing these OpenMP directives on LLVM is ongoing.

The Vector Function ABI specifications are provided for the Intel $x 86$ and the Armv8 (AArch64) families of vector extensions [34], [37]. The compiler generates SIMD function calls according to the compiler flags. For example, when targeting AArch64 SVE auto-vectorization, the compiler will transform a call to the standard sin function to a call to the symbol _ZGVsMxv_sin. When targeting Intel AVX-512 [38] auto-vectorization, the compiler would generate a call to the symbol_ZGVeNe8v_sin.

### 4.2 SLEEF and the Vector Function ABI

SLEEF is provided as two separate libraries. The first library exposes the functions of SLEEF to programmers for inclusion in their $C / C++$ code. The second library exposes the functions with names mangled according to the Vector Function ABI. This makes SLEEF a viable alternative to libm and its SIMD counterpart libmvec, in glibc. This also enables a user work-flow that relies on the auto-vectorization capabilities of a compiler. The compatibility with libmvec enables users to swap from libmvec to libsleef by simply changing compiler options, without changing the code that generated the vector call. The two SLEEF libraries are built from the same source code, which are configured to target the different versions via auto-generative programs that transparently rename the functions according to the rules of the target library.

## 5 Overview of Library Implementation

One of the objectives of the SLEEF project is to provide a library of vectorized math functions that can be used in conjunction with vectorizing compilers. When a non-vectorized code is automatically vectorized, the compiler converts calls to scalar math functions to calls to a SIMD version of the math functions. In order to make this conversion safe and applicable to wide variety of codes, we need functions with 1-ULP error bound that conforms to ANSI C standard. On the other hand, there are users who need better performance. Our library provides 1-ULP accuracy version and 3.5-ULP accuracy version for most of the functions. We confirmed that our library satisfies the accuracy requirements on an empirical basis. For non-finite inputs and outputs, we implemented the functions to return the same results as libm, as specified in the ANSI C standard. They do not set errno nor raise an exception.

In order to optimize a program with SIMD instructions, it is important to eliminate conditional branches as much as possible, and execute the same sequence of instructions regardless of the argument. If the algorithm requires conditional branches according to the argument, it must prepare for the case where the elements in the input vector contain both values that would make a branch happen and not happen. Recent processors have a long pipeline and therefore branch misprediction penalty can reach more than 10 cycles [39]. Making a decision for a conditional branch also requires non-negligible computation, within the scope of
our tests. A conditional move is an operator for choosing one value from two given values according to a condition. This is equivalent to a ternary operator and can be used in a vectorized code to replace a conditional branch. Some other operations are also expensive in vectorized implementation. A table-lookup is expensive. Although in-register table lookup is reported fast on Cell BE SPU [12], it is substantially slower than polynomial evaluation without any table lookup, within the scope of our tests. Most vector extensions do not provide 64-bit integer multiplication or a vector shift operator with which each element of a vector can be specified a different number of bits to shift. On the other hand, FMA and round-to-integer instructions are supported by most vector extensions. Due to the nature of the evaluation methods, dependency between operations cannot be completely eliminated. Latencies of operations become an issue when a series of dependent operations are executed. FP division and square root are not too expensive from this aspect. ${ }^{2}$

The actual structure of the pipeline in a processor is complex, and such level of details are not well-documented for most CPUs. Therefore, it is not easy to optimize the code according to such hardware implementation. In this paper, we define the latency and throughput of an instruction or a subroutine as follows [41]. The latency of an instruction or a subroutine is the delay that it generates in a dependency chain. The throughput is the maximum number of instructions or subroutines of the same kind that can be executed per unit time when the inputs are independent of the preceding instructions or subroutines. Several tools and methods are proposed for automatically constructing models of latency, throughput, and port usage of instructions [42], [43]. Within the scope of our tests, most of the instruction latency in the critical path of evaluating a vector math function tends to be dominated by FMA operations. In many processors, FMA units are implemented in a pipeline manner. Some powerful processors have multiple FMA units with out-of-order execution, and thus the throughput of FMA instruction is large, while the latency is long. In SLEEF, we try to maximize the throughput of computation in a versatile way by only taking account of dependencies among FMA operations. We regard each FMA operation as a job that can be executed in parallel and try to reduce the length of the critical path.

In order to evaluate a double-precision (DP) function to 1-ULP accuracy, the internal computation with accuracy better than 1 ULP is sometimes required. Double-double arithmetic, in which a single value is expressed by a sum of two double-precision FP values [44], [45], is used for this purpose. All the basic operators for DD arithmetic can be implemented without a conditional branch, and therefore it is suitable for vectorized implementation. Because we only need 1-ULP overall accuracy for DP functions, we use simplified DD operators with less than the full DD accuracy. In SLEEF, we omit re-normalization of DD values by default, allowing overlap between the two numbers. We carry out re-normalization only when necessary.
2. The latencies of 256 -bit DP add, divide and sqrt instructions are 4, 14 and 18 cycles, respectively on Intel Skylake processors [40].

Evaluation of an elementary function often consists of three steps: range reduction, approximation, and reconstruction [21]. An approximation step computes the elementary function using a polynomial. Since this approximation is only valid for a small domain, a number within that range is computed from the argument in a range reduction step. The reconstruction step combines the results of the first two steps to obtain the resulting number.

An argument reduction method that finds an FP remainder of dividing the argument $x$ by $\pi$ is used in evaluation of trigonometric functions. The range reduction method suggested by Cody and Waite [46], [47] is used for small arguments. The Payne and Hanek's method [48] provides an accurate range-reduction for a large argument of trigonometric function, but it is expensive in terms of operations.

There are tools available for generating the coefficients of the polynomials, such as Maple [49] and Sollya [50]. In order to fine-tune the generated coefficients, we created a tool for generating coefficients that minimizes the maximum relative error. When a SLEEF function evaluates a polynomial, it evaluates a few lowest degree terms in DD precision while other terms are computed in double-precision, in order to achieve 1-ULP overall accuracy. Accordingly, coefficients in DD precision or coefficients that can be represented by FP numbers with a few most significant bits in mantissa are used in the last few terms. We designed our tool to generate such coefficients. We use Estrin's scheme [51] to evaluate a polynomial to reduce dependency between FMA operations. This scheme reduces bubbles in the pipeline, and allows more FMA operations to be executed in parallel. Reducing latency can improve the throughput of evaluating a function because the latency and the reciprocal throughput of the entire function are close to each other.

Below, we describe and compare the implementations of selected functions in SLEEF, FDLIBM [4] and Christoph Lauter's Vector-libm [5]. We describe 1-ULP accuracy version of functions in SLEEF. The error bound specification of FDLIBM is 1 ULP.

### 5.1 Implementation of $\sin$ and cos

FDLIBM uses Cody-Waite range reduction if the argument is under $2^{18} \pi$. Otherwise, it uses the Payne-Hanek range reduction. Then, it switches between polynomial approximations of the sine and cosine functions on $[-\pi / 4, \pi / 4]$. Each polynomial has 6 non-zero terms.
sin and cos in Vector-libm have 4-ULP error bound. They use a vectorized path if all arguments are greater than 3.05e151, and less than 5.147 for sine and 2.574 for cosine. In the vectorized paths, a polynomial with 8 and 9 non-zero terms is used to approximate the sine function on $[-\pi / 2, \pi / 2]$, following Cody-Waite range reduction. In the scalar paths, Vectorlibm uses a polynomial with 10 non-zero terms.

SLEEF switches among two Cody-Waite range reduction methods with approximation with different sets of constants, and the Payne-Hanek reduction. The first version of the algorithm operates for arguments within $[-15,15]$, and the second version for arguments that are within $\left[-10^{14}, 10^{14}\right]$. Otherwise, SLEEF uses a vectorized Payne-Hanek reduction, which is described in 5.8. SLEEF only uses conditional branches for choosing a reduction method from Cody-Waite and Payne-Hanek. SLEEF uses a polynomial approximation
of the sine function on $[-\pi / 2, \pi / 2]$, which has 9 non-zero terms. The sign is set in the reconstruction step.

### 5.2 Implementation of tan

After Cody-Waite or Payne-Hanek reduction, FDLIBM reduces the argument to [0,0.67434], and uses a polynomial approximation with 13 non-zero terms. It has 10 if statements after Cody-Waite reduction.
tan in Vector-libm has 8-ULP error bound. A vectorized path is used if all arguments are less than 2.574 and greater than $3.05 \mathrm{e}-151$. After Cody-Waite range reduction, a polynomial with 9 non-zero terms for approximating sine function on $[-\pi / 2, \pi / 2]$ is used twice to approximate sine and cosine of the reduced argument. The result is obtained by dividing these values. In the scalar path, Vector-libm evaluates a polynomial with 10 non-zero terms twice.

In SLEEF, the argument is reduced in 3 levels. It first reduces the argument to $[-\pi / 2, \pi / 2]$ with Cody-Waite or Payne-Hanek range reduction. Then, it reduces the argument to $[-\pi / 4, \pi / 4]$ with $\tan a_{1}=1 / \tan \left(\pi / 2-a_{0}\right)$. At the third level, it reduces the argument to $[-\pi / 8, \pi / 8]$ with the double-angle formula. Let $a_{0}$ be the reduced argument with Cody-Waite or Payne-Hanek. $a_{1}=\pi / 2-a_{0}$ if $\left|a_{0}\right|>\pi / 4$. Otherwise, $a_{1}=a_{0}$. Then, SLEEF uses a polynomial approximation of the tangent function on $[-\pi / 8, \pi / 8]$, which has 9 non-zero terms, to approximate $\tan \left(a_{1} / 2\right)$. Let $t$ be the obtained value with this approximation. Then, $\tan a_{0} \approx$ $2 t /\left(1-t^{2}\right)$ if $\left|a_{0}\right| \leq \pi / 4$. Otherwise, $\tan a_{0} \approx\left(1-t^{2}\right) /(2 t)$. SLEEF only uses conditional branches for choosing a reduction method from Cody-Waite and Payne-Hanek. Annotated source code of tan is shown in Appendix A, which can be found on the Computer Society Digital Library at http:/ /doi.ieeecomputersociety.org/10.1109/
TPDS.2019.2960333

### 5.3 Implementation of asin and acos

FDLIBM and SLEEF first reduces the argument to [0, 0.5] using $\arcsin x=\pi / 2-2 \arcsin \sqrt{(1-x) / 2}$ and $\arccos x=$ $2 \arcsin \sqrt{(1-x) / 2}$.

Then, SLEEF uses a polynomial approximation of arcsine on $[0,0.5]$ with 12 non-zero terms.

FDLIBM uses a rational approximation with 11 terms (plus one division). For computing arcsine, FDLIBM switches the approximation method if the original argument is over 0.975 . For computing arccosine, it has three paths that are taken when $|x|<0.5, x \leq-0.5$ and $x \geq 0.5$, respectively. It has 7 and 6 if statements in asin and acos, respectively.
asin and acos in Vector-libm have 6-ULP error bound. asin and acos in Vector-libm use vectorized paths if arguments are all greater than $3.05 \mathrm{e}-151$ and $2.77 \mathrm{e}-17$, respectively. Vector-libm evaluates polynomials with $3,8,8$, and 5 terms to compute arcsine. It evaluates a polynomial with 21 terms for arccosine.

### 5.4 Implementation of atan

FDLIBM reduces the argument to $[0,7 / 16]$. It uses a polynomial approximation of the arctangent function with 11 non-zero terms. It has 9 if statements.
atan in Vector-libm have 6-ULP error bound. Vectorlibm uses vectorized paths if arguments are all greater than $1.86 \mathrm{e}-151$ and less than 2853. It evaluates four polynomials with $7,9,9$ and 4 terms in the vectorized path.

SLEEF reduces argument $a$ to $[0,1]$ using $\arctan x=$ $\pi / 2-\arctan (1 / x)$. Let $a^{\prime}=1 / a$ if $|a| \geq 1$. Otherwise, $a^{\prime}=a$. It then uses a polynomial approximation of arctangent function with 20 non-zero terms to approximate $r \approx \arctan a^{\prime}$. As a reconstruction, it computes $\arctan a \approx \pi / 2-r$ if $|a| \geq 1$. Otherwise, $\arctan a \approx r$.

### 5.5 Implementation of log

FDLIBM reduces the argument to $[\sqrt{2} / 2, \sqrt{2}]$. It then approximates the reduced argument with a polynomial that contains 7 non-zero terms in a similar way to SLEEF. It has 9 if statements.
log in Vector-libm has 4-ULP error bound. It uses a vectorized path if the input is a normalized number. It uses a polynomial with 20 non-zero terms to approximate the logarithm function on $[0.75,1.5]$. It does not use division.

SLEEF multiplies the argument $a$ by $2^{64}$, if the argument is a denormal number. Let $a^{\prime}$ be the resulting argument, $e=\left\lfloor\log _{2}\left(4 a^{\prime} / 3\right)\right\rfloor$ and $m=a^{\prime} \cdot 2^{-e}$. If $a$ is a denormal number, $e$ is subtracted 64. SLEEF uses a polynomial with 7 nonzero terms to evaluate $\log m \approx \sum_{n=0}^{6} C_{n}\left(\frac{m-1}{m+1}\right)^{2 n+1}$, where $C_{0} \ldots C_{6}$ are constants. As a reconstruction, it computes $\log a=e \log 2+\log m$.

### 5.6 Implementation of exp

All libraries reduce the argument range to $[-(\log 2) / 2$, $(\log 2) / 2]$ by finding $r$ and integer $k$ such that $x=k \log 2+$ $r,|r| \leq(\log 2) / 2$.

SLEEF then uses a polynomial approximation with 13 non-zero terms to directly approximate the exponential function of this domain. It achieves 1-ULP error bound without using a DD operation.

FDLIBM uses a polynomial with 5 non-zero terms to approximate $f(r)=r\left(e^{r}+1\right) /\left(e^{r}-1\right)$. It then computes $\exp (r)=1+2 r /(f(r)-r)$. It has 11 if statements.

The reconstruction step is to add integer $k$ to the exponent of the resulting FP number of the above computation.
exp in Vector-libm has 4-ULP error bound. A vectorized path covers almost all input domains. It uses a polynomial with 11 terms to approximate the exponential function.

### 5.7 Implementation of pow

FDLIBM computes $y \log _{2} x$ in DD precision. Then, it computes $\operatorname{pow}(x, y)=e^{\log 2 \cdot y \log _{2} x}$. It has 44 if statements.

Vector-libm does not implement pow.
SLEEF computes $e^{y \log x}$. The internal computation is carried out in DD precision. In order to compute logarithm internally, it uses a polynomial with 11 non-zero terms. The accuracy of the internal logarithm function is around 0.008 ULP. The internal exponential function in pow uses a polynomial with 13 non-zero terms.

### 5.8 The Payne-Hanek Range Reduction

Our method computes $\operatorname{rfrac}(2 x / \pi) \cdot \pi / 2$, where $\operatorname{rfrac}(a):=a$ $-\operatorname{round}(a)$. The argument $x$ is an FP number, and therefore

TABLE 1
Number of FP Operators in the Payne-Hanek Implementations

| Operator | SLEEF (1-ULP) | Vector-libm (4-ULP) |
| :--- | :---: | :---: |
| add/sub | 36 | 71 |
| mul | 5 | 18 |
| FMA | 11 | 0 |
| round | 8 | 0 |

it can be represented as $M \cdot 2^{E}$, where $M$ is an integer mantissa and $E$ is an integer exponent $E$. We now denote the integral part and the fractional part of $2^{E} \cdot 2 / \pi$ as $I(E)$ and $F(E)$, respectively. Then,

$$
\begin{aligned}
\operatorname{rfac}(2 x / \pi) & =\operatorname{rfrac}\left(M \cdot 2^{E} \cdot 2 / \pi\right) \\
& =\operatorname{rfrac}(M \cdot(I(E)+F(E))) \\
& =\operatorname{rfrac}(M \cdot F(E)) .
\end{aligned}
$$

The value $F(E)$ only depends on the exponent of the argument, and therefore, can be calculated and stored in a table, in advance. In order to compute $\operatorname{rfrac}(M \cdot F(E))$ in DD precision, $F(E)$ must be in quad-double-precision. We now denote $F(E)=F_{0}(E)+F_{1}(E)+F_{2}(E)+F_{3}(E)$, where $F_{0}(E) \ldots F_{3}(E)$ are DP numbers and $\left|F_{0}(E)\right| \geq\left|F_{1}(E)\right|$ $\geq\left|F_{2}(E)\right| \geq\left|F_{3}(E)\right|$. Then,

$$
\begin{align*}
& \operatorname{rfrac}(M \cdot F(E)) \\
= & \operatorname{rfrac}\left(M \cdot F_{0}(E)+M \cdot F_{1}(E)\right. \\
& \left.+M \cdot F_{2}(E)+M \cdot F_{3}(E)\right)  \tag{1}\\
= & \operatorname{rfrac}\left(\operatorname { r f r a c } \left(\operatorname{rfrac}\left(\operatorname{rfrac}\left(M \cdot F_{0}(E)\right)+M \cdot F_{1}(E)\right)\right.\right. \\
& \left.\left.+M \cdot F_{2}(E)\right)+M \cdot F_{3}(E)\right),
\end{align*}
$$

because $\operatorname{rfrac}(a+b)=\operatorname{rfrac}(\operatorname{rfrac}(a)+b)$. In the method, we compute (1) in DD precision in order to avoid overflow. The size of the table retaining $F_{0}(E) \ldots F_{3}(E)$ is 32 K bytes. Our method is included in the source code of tan shown in Appendix A, available in the online supplemental material.

FDLIBM seems to implement the original Payne-Hanek algorithm with more than 100 lines of C code, which includes 13 if statements, 18 for loops, 1 switch statement and 1 goto statement. The numbers of iterations of most of the for loops depend on the argument.

Vector-libm implements a non-vectorized variation of the Payne-Hanek algorithm which has some similarity with our method. In order to reduce argument $x$, it first decomposes $|x|$ into $E$ and $n$ such that $2^{E} \cdot n=|x|$. A triple-double (TD) approximation to $t(E)=2^{E} / \pi-2 \cdot\left\lfloor 2^{E-1} / \pi\right\rfloor$ is looked-up from a table. It then calculates $m=n \cdot t(E)$ in TD. The reduced argument is obtained as a product of $\pi$ and the fractional part of $m$. In Table 1, we compare the numbers of FP operators in the implementations. Note that the method in Vector-libm is used for trigonometric functions with 4-ULP error bound, while our method is used for functions with 1-ULP error bound.

### 5.9 FP Remainder

We devised an exact remainder calculation method suitable for vectorized implementation. The method is based on the long division method, where an FP number is regarded as a digit. Fig. 3 shows an example process for calculating the FP


Fig. 3. Example computation of FP remainder.
remainder of $1 \mathrm{e}+40 / 0.75$. Like a typical long division, we first find integer quotient $1.333 \mathrm{e}+40$ so that $1.333 \mathrm{e}+40 \cdot 0.75$ does not exceed $1 \mathrm{e}+40$. We multiply the found quotient with 0.75 , and then subtract it from $1 \mathrm{e}+40$ to find the dividend $4.836 \mathrm{e}+24$ for the second iteration.

Our basic algorithm is shown in Algorithm 1. If $n, d$ and $q_{k}$ are FP numbers of the same precision $p$, then $r_{k}$ is representable with an FP number of precision $2 p$. In this case, the number of iterations can be minimized by substituting $q_{k}$ with the largest FP number of precision $p$ within the range specified at line 3 . However, the algorithm still works if $q_{k}$ is any FP number of precision $p$ within the range. By utilizing this property, an implementer can use a division operator that does not return a correctly rounded result. The source code of an implementation of this algorithm is shown in Fig. 6 in Appendix B, available in the online supplemental material. A part of the proof of correctness is shown in Appendix C, available in the online supplemental material.

FDLIBM uses a method of shift and subtract. It first converts the mantissa of two given arguments into 64-bit integers, and calculates a remainder in a bit-by-bit basis. The main loop iterates $i x-i y$ times, where $i x$ and $i y$ are the exponents of the arguments of fmod. This loop includes 10 integer additions and 3 if statements. The number of iterations of the main loop can reach more than 1,000 .

Vector-libm does not implement FP remainder.

```
Algorithm 1. Exact Remainder Calculation
Input: Finite positive numbers \(n\) and \(d\)
Output: Returns \(n-d\lfloor n / d\rfloor\)
    \(r_{0}:=n, k:=0\)
    while \(d \leq r_{k}\) do
        \(q_{k}\) is an arbitrary integer satisfying \(\left(r_{k} / d\right) / 2 \leq q_{k} \leq r_{k} / d\)
        \(r_{k+1}:=r_{k}-q_{k} d\)
        \(k:=k+1\)
    end while
    return \(r_{k}\)
```


### 5.10 Handling of Special Numbers, Exception and Flags

Our implementation gives a value within the specified error bound without special handling of denormal numbers, unless otherwise noted.

When a function has to return a specific value for a specific value of an argument (such as a NaN or a negative zero) is given, such a condition is checked at the end of each function. The return value is substituted with the special value if the condition is met. This process is complicated in functions like pow, because they have many conditions for returning special values.

SLEEF functions do not give correct results if the computation mode is different from round-to-nearest. They do not
set errno nor raise an exception. This is a common behavior among vectorized math libraries including Libmvec [11] and SVML [3]. Because of SIMD processing, functions can raise spurious exceptions if they try to raise an exception.

### 5.11 Summary

FDLIBM extensively uses conditional branches in order to switch the polynomial according to the argument(sin, cos, tan, log, etc), to return a special value if the arguments are special values(pow, etc.), and to control the number of iterations (the Payne-Hanek reduction).

Vector-libm switches between a few polynomials in most of the functions. It does not provide functions with 1-ULP error bound, nevertheless, the numbers of non-zero terms in the polynomials are larger than other two libraries in some of the functions. A vectorized path is used only if the argument is smaller than 2.574 in cos and tan, although these functions are frequently evaluated with an argument up to $2 \pi$. In most of the functions, Vector-libm uses a non-vectorized path if the argument is very small or a non-finite number. For example, it processes 0 with non-vectorized paths in many functions, although 0 is a frequently evaluated argument in normal situations. If non-finite numbers are once contained in data being processed, the whole processing can become significantly slower afterward. Variation in execution time can be exploited for a side-channel attack in cryptographic applications.

SLEEF uses the fastest paths if all the arguments are under 15 for trigonometric functions, and the same vectorized path is used regardless of the argument in most of the non-trigonometric functions. SLEEF always uses the same polynomial regardless of the argument in all functions.

Although reducing the number of conditional branches has a few advantages in implementing vector math libraries, it seems to be not given a high priority in other libraries.

## 6 Testing

SLEEF includes three kinds of testers. The first two kinds of testers test the accuracy of all functions against high-precision evaluation using the MPFR library. In these tests, the computation error in ULP is calculated by comparing the values output by each SLEEF function and the values output by the corresponding function in the MPFR library, and it is checked if the error is within the specified bounds.

### 6.1 Perfunctory Test

The first kind of tester carries out a perfunctory set of tests to check if the build is correct. These tests include standards compliance tests, accuracy tests and regression tests.

In the standards compliance tests, we test if the functions return the correct values when values that require special handling are given as the argument. These argument values include $\pm \operatorname{Inf}, \mathrm{NaN}$ and $\pm 0$. Atan2 and pow are binary
functions and have many combinations of these special argument values. These are also all tested.

In the accuracy test, we test if the error of the returned values from the functions is within the specified range, when a predefined set of argument values are given. These argument values are basically chosen between a few combinations of two values at regular intervals. The trigonometric functions are also tested against argument values close to integral multiples of $\pi / 2$. Each function is tested against tens of thousands of argument values in total.

In the regression test, the functions are tested with argument values that triggered bugs in the previous library release, in order to prevent re-emergence of the same bug.

The executables are separated into a tester and IUTs (Implementation Under Test). The tests are carried out by making these two executables communicate via an input/ output pipeline, in order to enable testing of libraries for architectures which the MPFR library does not support.

### 6.2 Randomized Test

The second kind of tester is designed to run continuously. This tester generates random arguments and compare the output from each function to the output calculated with the corresponding function in the MPFR library. This tester is expected to find bugs if it is run for a sufficiently long time.

In order to randomly generate an argument, the tester generates random bits of the size of an FP value, and reinterprets the bits as an FP value. The tester executes the randomized test for all the functions in the library at several thousand arguments per second for each function on a computer with a Core i7-6700 CPU.

In the SLEEF project, we use randomized testing in order to check the correctness of functions, rather than formal verification. It is indeed true that proving correctness of implementation contributes to the reliability of implementation. However, there is a performance overhead because the way of implementation is limited in a form that is easy to prove the correctness. There would be an increased cost of maintaining the library because of the need for updating the proof each time the implementation is modified.

### 6.3 Bit-Identity Test

The third kind of tester is for testing if bit-identical results are returned from the functions that are supposed to return such results. This test is designed to compare the results among the binaries compiled with different vector extensions. For each predetermined list of arguments, we calculate an MD5 hash value of all the outputs from each function. Then, we check if the hash values match among functions for different architectures.

## 7 Performance Comparison

In this section, we present results of a performance comparison between FDLIBM Version 5.3 [4], Vector-libm [5], SLEEF 3.4, and Intel SVML [3] included in Intel C Compiler 19.

We measured the reciprocal throughput of each function by measuring the execution time of a tight loop that repeatedly calls the function in a single-threaded process. In order
to obtain useful results, we turned off optimization flags when compiling the source code of this tight loop, ${ }^{3}$ while the libraries are compiled with their default optimization options. We did not use LTO. We confirmed that the calls to the function are not compiled out or inlined by checking the assembly output from the compiler. The number of function calls by each loop is $10^{10}$, and the execution time of this loop is measured with the clock_gettime function.

We compiled SLEEF and FDLIBM using gcc-7.3.0 with "-03 -mavx2 -mfma" optimization options. We compiled Vector-libm using gcc-7.3.0 with the default "-O3 -march=native -ftree-vectorize -ftree-vector-izer-verbose=1 -fno-math-errno" options. We changed VECTOR_LENGTH in vector.h to 4 and compiled the source code on a computer with an Intel Core i76700 CPU. ${ }^{4}$ The accuracy of functions in SVML can be chosen by compiler options. We specified an "-fimf-maxerror=1.0" and an "-fimf-max-error=4.0" options for icc to obtain the 1-ULP and 4-ULP accuracy results, respectively.

We carried out all the measurements on a physical PC with Intel Core i7-6700 CPU @ 3.40 GHz without any virtual machine. In order to make sure that the CPU is always running at the same 3.4 GHz clock speed during the measurements, we turned off Turbo Boost. With this setting, 10 nano sec . corresponds to 34 clock cycles.

The following results compare the the reciprocal throughput of each function. If the implementation is vectorized and each vector has $N$ elements of FP numbers, then a single execution evaluates the corresponding mathematical function $N$ times. We generated arguments in advance and stored in arrays. Each time a function is executed, we set a randomly generated argument to each element of the argument vector (each element is set with a different value). The measurement results do not include the delay for generating random numbers.

### 7.1 Execution Time of Floating Point Remainder

We compared the reciprocal throughput of double-precision fmod functions in the libm included in Intel C Compiler 19, FDLIBM and SLEEF. All the FP remainder functions always return a correctly-rounded result. We generated a random denominator $d$ and a numerator uniformly distributed within $[1,100]$ and $[0.95 r \cdot d, 1.05 r \cdot d]$, respectively, where $r$ is varied from 1 to $10^{25}$. Fig. 4 shows the reciprocal throughput of the fmod function in each library. Please note that SVML does not contain a vectorized fmod function.

The graph of reciprocal throughput looks like a step function, because the number of iterations increases in this way.
3. If we turn on the optimizer, there is concern that the compiler optimizes away the call to a function. In order to prevent this, we have to introduce extra operations, but this also introduces overhead. After trying several configurations of the loop and optimizer settings, we decided to turn off the optimizer in favor of reproducibility, simplicity and fairness. We checked the assembly output from the compiler and confirmed that the unoptimized loop simply calls the target function and increments a counter, and therefore that the operations inside a loop are minimal.
4. Vector-libm evaluates functions with 512 bits of vector length by default. Because SLEEF and SVML are 256-bit wide, the setting is changed.


Fig. 4. Reciprocal throughput of double-precision fmod functions.

### 7.2 Comparison of Overall Execution Time

We compared the reciprocal throughput of 256 -bit wide vectorized double-precision functions in Vector-libm, SLEEF and SVML, and scalar functions in FDLIBM. We generated random arguments that were uniformly distributed within the indicated intervals for each function. In order to check execution speed of fast paths in trigonometric functions, we measured the reciprocal throughput with arguments within $[0.4,0.5]$. The result is shown in Table 2.

The reciprocal throughput of functions in SLEEF is comparable to that of SVML in all cases. This is because the latency of FP operations is generally dominant in the execution time of math functions. Because there are two levels of scheduling mechanisms, which includes the optimizer in a compiler and the out-of-order execution hardware, there is small room for making a difference to the throughput or latency.

Execution speed of FDLIBM is not very slow despite many conditional branches. This seems to be because of a smaller number of FP operations, and faster execution speed of scalar instructions compared to equivalent SIMD instructions.

Vector-libm is slow even if only the vectorized path is used. This seems to be because Vector-libm evaluates polynomials with a large number of terms. Auto-vectorizers are still developing, and the compiled binary code might not be well optimized. When a slow path has to be used, Vectorlibm is even slower since a scalar evaluation has to be carried out for each of the elements in the vector.

Vector-libm uses Horner's method to evaluate polynomials, which involves long latency of chained FP operations. In FDLIBM, this latency is reduced by splitting polynomials into even and odd terms, which can be evaluated in parallel. SLEEF uses Estrin's scheme. In our experiments, there was only a small difference between Estrin's scheme and splitting polynomials into even and odd terms with respect to execution speed.

## 8 Conclusion

In this paper, we showed that our SLEEF library shows performance comparable to commercial libraries while maintaining good portability. We have been continuously developing SLEEF since 2010. ${ }^{5}$ [52] We distribute SLEEF

TABLE 2
Reciprocal Throughput in Nano Sec

| Func, error bound, domain | Vector-libm | FDLIBM | SLEEF | SVML |
| :---: | :---: | :---: | :---: | :---: |
| $\begin{aligned} & \hline \text { sin, } 1 \text { ulp, } \\ & {[0.4,0.5]} \end{aligned}$ |  | 4.927 | 11.43 | 13.68 |
| $\begin{gathered} \sin , 4 \text { ulps, } \\ {[0.4,0.5]} \end{gathered}$ | 9.601 |  | 7.504 | 6.679 |
| $\begin{gathered} \text { sin, } 1 \text { ulp, } \\ {[0,6.28]} \end{gathered}$ |  | 18.96 | 11.41 | 13.86 |
| $\begin{gathered} \sin , 4 \text { ulps, } \\ {[0,6.28]} \end{gathered}$ | 12.48 |  | 7.507 | 6.723 |
| $\begin{gathered} \sin , 1 \text { ulp, } \\ {[0,1 e+100]} \end{gathered}$ |  | 162.3 | 48.79 | 41.72 |
| $\begin{gathered} \sin , 4 \text { ulps, } \\ {[0,1 e+100]} \end{gathered}$ | 288.6 |  | 43.82 | 34.96 |
| $\begin{gathered} \text { cos, } 1 \text { ulp, } \\ {[0.4,0.5]} \end{gathered}$ |  | 11.42 | 13.75 | 12.99 |
| $\begin{gathered} \text { cos, } 4 \text { ulps, } \\ {[0.4,0.5]} \end{gathered}$ | 9.557 |  | 7.850 | 7.917 |
| $\begin{gathered} \text { cos, } 1 \text { ulp, } \\ {[0,6.28]} \end{gathered}$ |  | 18.45 | 13.74 | 13.18 |
| $\begin{gathered} \text { cos, } 4 \text { ulps, } \\ {[0,6.28]} \end{gathered}$ | 13.97 |  | 7.850 | 7.838 |
| $\begin{gathered} \text { cos, } 1 \text { ulp, } \\ {[0,1 e+100]} \end{gathered}$ |  | 162.1 | 50.38 | 38.38 |
| $\begin{gathered} \text { cos, } 4 \text { ulps, } \\ {[0,1 e+100]} \end{gathered}$ | 360.3 |  | 46.09 | 36.57 |
| $\begin{gathered} \tan , 1 \text { ulp, } \\ {[0.4,0.5]} \end{gathered}$ |  | 7.819 | 17.30 | 15.71 |
| $\begin{gathered} \tan , 4+\text { ulps, } \\ {[0.4,0.5]} \\ \hline \end{gathered}$ | 15.58 |  | 9.367 | 7.570 |
| $\begin{gathered} \tan , 1 \text { ulp, } \\ {[0,6.28]} \end{gathered}$ |  | 22.24 | 17.28 | 15.78 |
| $\begin{gathered} \tan , 4+\mathrm{ulps}, \\ {[0,6.28]} \\ \hline \end{gathered}$ | 20.16 |  | 9.367 | 7.595 |
| $\begin{gathered} \tan , 1 \text { ulp, } \\ {[0,1 e+100]} \end{gathered}$ |  | 177.0 | 48.82 | 43.31 |
| $\begin{aligned} & \tan , 4+\mathrm{ulps} \\ & {[0,1 e+100]} \end{aligned}$ | 399.4 |  | 36.54 | 40.50 |
| $\begin{gathered} \text { asin, } 1 \text { ulp, } \\ {[-1,1]} \end{gathered}$ |  | 14.87 | 12.99 | 12.10 |
| $\begin{gathered} \text { asin, 4+ ulps, } \\ {[-1,1]} \\ \hline \end{gathered}$ | 20.75 |  | 5.552 | 9.627 |
| $\begin{gathered} \text { acos, } 1 \text { ulp, } \\ {[-1,1]} \end{gathered}$ |  | 12.07 | 16.09 | 12.11 |
| $\begin{gathered} \text { acos, } 4+\text { ulps, } \\ {[-1,1]} \end{gathered}$ | 23.62 |  | 7.572 | 10.23 |
| $\begin{aligned} & \text { atan, } 1 \text { ulp, } \\ & {[-700,700]} \end{aligned}$ |  | 10.16 | 22.12 | 19.97 |
| $\begin{gathered} \text { atan, 4+ ulps, } \\ {[-700,700]} \end{gathered}$ | 35.54 |  | 9.251 | 12.09 |
| $\begin{gathered} \text { log, } 1 \text { ulp, } \\ {[0,1 e+300]} \end{gathered}$ |  | 31.66 | 15.46 | 12.05 |
| log, 4 ulps, $[0,1 e+300]$ | 39.64 |  | 9.636 | 8.842 |
| $\begin{aligned} & \text { exp, } 1 \text { ulp, } \\ & {[-700,700]} \end{aligned}$ |  | 12.19 | 7.663 | 7.968 |
| $\begin{aligned} & \text { exp, } 4 \text { ulps, } \\ & {[-700,700]} \end{aligned}$ | 17.35 |  |  | 6.756 |
| $\begin{gathered} \text { pow, } 1 \text { ulp, } \\ {[-30,30][-30,30]} \end{gathered}$ |  | 69.40 | 55.53 | 75.18 |

under the Boost Software License [53], which is a permissive open source license. We actively communicate with developers of compilers and members of other projects in order to understand the needs of real-world users. The Vector Function ABI is important in developing vectorizing compilers. The functions that return bit-identical results are added to our library to reflect requests from our multiple partners. We thoroughly tested these functionalities, and SLEEF is already adopted in multiple commercial products.

## Acknowledgments

The authors would like to acknowledge Will Lovett and Srinath Vadlamani for their valuable input and suggestions. They are particularly grateful to Robert Werner, who reviewed the final manuscript. The authors would like to thank Prof. Leigh McDowell for his suggestions. The authors would like to thank all the developers that contributed to the SLEEF project, in particular Diana Bite and Alexandre Mutel.

## References

[1] G. Conte, S. Tommesani, and F. Zanichelli, "The long and winding road to high-performance image processing with MMX/SSE," in Proc. 5th IEEE Int. Workshop Comput. Architectures Mach. Perception, 2000, pp. 302-310.
[2] D. Naishlos, "Autovectorization in GCC," in Proc. GCC Developers Summit, 2004, pp. 105-118. [Online]. Available: http://people. redhat.com/lockhart/.gcc2004/MasterGCC-2side.pdf
[3] Intel Corporation, "Intel short vector math library," 2019. [Online]. Available: https://software.intel.com/en-us/node/523613
[4] Sun Microsystems, Inc., "Sun freely distributable libm version 5.3.," 2010. [Online]. Available: http://www.netlib.org/fdlibm/
[5] C. Lauter, "A new open-source SIMD vector libm fully implemented with high-level scalar C," in Proc. 50th Asilomar Conf. Signals Syst. Comput., 2016, pp. 407-411.
[6] ISO, ISO/IEC 9899:2011 Information technology-Programming lan-guages-C. International Organization for Standardization, 2011. [Online]. Available: https:/ /www.iso.org/standard/57853.html
[7] GNU Project, "The GNU C library (glibc)," 2018. [Online]. Available: https:/ /www.gnu.org/software/libc/
[8] S. Gal, "An accurate elementary mathematical library for the IEEE floating point standard," ACM Trans. Math. Softw., vol. 17, no. 1, pp. 26-45, Mar. 1991.
[9] C. Daramy-Loirat et al., "CR-LIBM: A correctly rounded elementary function library," in Proc. SPIE 5205, Adv. Signal Process. Algorithms Architectures Implementations XIII, 2003, vol. 5205, pp. 458-464.
[10] Advanced Micro Devices, Inc., "AMD core math library," 2013. [Online]. Available: http://developer.amd.com/tools-and-sdks/ archive/acml-product-features/
[11] GNU Project, "Libmvec in glibc," 2015. [Online]. Available: https://sourceware.org/glibc/wiki/libmvec
[12] C. K. Anand and W. Kahl, "An optimized cell be special function library generated by coconut," IEEE Trans. Comput., vol. 58, no. 8, pp. 1126-1138, Aug. 2009.
[13] M. Dukhan et al., "Yeppp! library," 2013. [Online]. Available: https://bitbucket.org/MDukhan/yeppp
[14] M. Dukhan and R. Vuduc, "Methods for high-throughput computation of elementary functions," in Proc. 10th Int. Conf. Parallel Process. Appl. Math., 2014, pp. 86-95.
[15] D. Piparo, V. Innocente, and T. Hauth, "Speeding up HEP experiment software with a library of fast and auto-vectorisable mathematical functions," J. Phys.: Conf. Ser., vol. 513, no. 5, 2014, Art. no. 052027.
[16] S. Manilov, B. Franke, A. Magrath, and C. Andrieu, "Free rider: A source-level transformation tool for retargeting platform-specific intrinsic functions," ACM Trans. Embedded Comput. Syst., vol. 16, no. 2, pp. 38:1-38:24, Dec. 2016. [Online]. Available: http:/ /doi. acm.org/10.1145/2990194
[17] M. Gross, "Neat SIMD: Elegant vectorization in C++ by using specialized templates," in Proc. Int. Conf. High Perform. Comput. Simul., 2016, pp. 848-857.
[18] N. Clark, A. Hormati, S. Yehia, S. Mahlke, and K. Flautner, "Liquid SIMD: Abstracting SIMD Hardware using Lightweight Dynamic Mapping," in Proc. IEEE 13th Int. Symp. High Perform. Comput. Architecture, 2007, pp. 216-227.
[19] R. Leißa, S. Hack, and I. Wald, "Extending a c-like language for portable SIMD programming," ACM SIGPLAN Notices, vol. 47, no. 8, pp. 65-74, Feb. 2012.
[20] Arm Limited, ARM NEON Intrinsics Reference, 2014. [Online]. Available: http://infocenter.arm.com/help/topic/com.arm.doc. ihi0073a/IHI0073A_arm_ neon_intrinsics_ref.pdf
[21] J.-M. Muller, Elementary Functions: Algorithms and Implementation. Basel, Switzerland: Birkhauser Boston, Inc., 1997.
[22] Arm Limited, "ARM C language extensions for SVE documentation," 2017. [Online]. Available: https://developer.arm. com/docs/100987/0000
[23] L. Liu et al., "Importance of bitwise identical reproducibility in earth system modeling and status report," Geoscientific Model Develop. Discuss., vol. 8, pp. 4375-4400, Jun. 2015.
[24] J.-M. Muller et al., Handbook of Floating-Point Arithmetic, 1st ed. Basel, Switzerland: Birkhauser Boston, Inc., 2009.
[25] O. Krzikalla, K. Feldhoff, R. Müller-Pfefferkorn, and W. E. Nagel, "Auto-vectorization techniques for modern SIMD architectures," in Proc. 16th Workshop Compilers Parallel Comput., 2012.
[26] D. Nuzman et al., "Vapor SIMD: Auto-vectorize once, run everywhere," in Proc. Int. Symp. Code Gener. Optim., 2011, pp. 151-160.
[27] R. M. Russell, "The CRAY-1 computer system," Commun. ACM, vol. 21, no. 1, pp. 63-72, Jan. 1978.
[28] D. Nuzman, I. Rosen, and A. Zaks, "Auto-vectorization of interleaved data for SIMD," in Proc. 27th ACM SIGPLAN Conf. Program. Lang. Design Implementation, 2006, pp. 132-143.
[29] X. Tian, A. Bik, M. Girkar, P. Grey, H. Saito, and E. Su, "Intel OpenMP C++/fortran compiler for hyper-threading technology: Implementation and performance," Intel Technol. J., vol. 6, no. 1, pp. 36-46, Feb 2002.
[30] GNU Project, "GCC, the GNU compiler collection," 1984. [Online]. Available: https://www.gnu.org/software/gcc/
[31] Arm Limited, "Arm compiler for HPC," [Online]. Available: https://developer.arm.com/products/software-development-tools/hpc/arm-c ompiler-for-hpc
[32] The LLVM compiler infrastructure project, The LLVM compiler infrastructure," 2004. [Online]. Available: https:/ /llvm.org/
[33] C. Lattner and V. Adve, "LLVM: A compilation framework for lifelong program analysis \& transformation," in Proc. Int. Symp. Code Gener. Optim., 2004, pp. 75-86.
[34] Arm Limited, "Vector function application binary interface specification for AArch64," 2018. [Online]. Available: https:/ / developer. arm.com/docs/101129/latest
[35] OpenMP Architecture Review Board, "OpenMP application program interface," 2013. [Online]. Available: https://www.openmp. org/wp-content/uploads/OpenMP4.0.0.pdf
[36] J. Lee, F. Petrogalli, G. Hunter, and M. Sato, "Extending OpenMP SIMD support for target specific code and application to ARM SVE," in Proc. Int. Workshop OpenMP, 2017, pp. 62-74.
[37] X. Tian et al., "Vector function application binary interface," Intel Corporation. 2015. [Online]. Available: https://software.intel. com/en-us/articles/vector-simd-function-abi
[38] A. Sodani et al.,"Knights landing: Second-generation intel xeon phi product," IEEE Micro, vol. 36, no. 2, pp. 34-46, Mar.-Apr. 2016.
[39] S. Eyerman, J. E. Smith, and L. Eeckhout, "Characterizing the branch misprediction penalty," in Proc. IEEE Int. Symp. Perform. Anal. Syst. Softw., 2006, pp. 48-58.
[40] Intel Corporation, "Intel 64 and IA-32 architectures optimization reference manual," 2016. [Online]. Available: https://software. intel.com/en-us/download/intel-64-and-ia-32-architectu res-optimization-reference-manual
[41] A. Fog, "Instruction tables: Lists of instruction latencies, throughputs and micro-operation breakdowns for Intel, AMD and VIA CPUs," [Online]. Available: https://www.agner.org/optimize/ instruction_tables.pdf
[42] A. Abel and J. Reineke, "Uops.info: Characterizing latency, throughput, and port usage of instructions on intel microarchitectures," in Proc. 24th Int. Conf. Architectural Support Program. Lang. Operating Syst., 2019, pp. 673-686.
[43] Google's Compiler Research Team and The LLVM compiler infrastructure project, "llvm-exegesis: Automatic measurement of instruction latency/uops," 2018. [Online]. Available: https:// github.com/google/EXEgesis
[44] T. J. Dekker, "A floating-point technique for extending the available precision," Numerische Mathematik, vol. 18, no. 3, pp. 224-242, Jun. 1971.
[45] J. Richard Shewchuk, "Adaptive precision floating-point arithmetic and fast robust geometric predicates," Discrete Comput. Geometry, vol. 18, no. 3, pp. 305-363, Oct. 1997.
[46] W. Cody and W. Waite, "Software manual for the elementary functions," Prentice-Hall, inc, Englewood cliffs, New-Jersey, 1980.
[47] W. J. Cody, "Implementation and testing of function software," in Proc. Problems Methodologies Math. Softw. Prod., Nov. 1980, pp. 24-47.
[48] M. H. Payne and R. N. Hanek, "Radian reduction for trigonometric functions," ACM SIGNUM Newslett., vol. 18, no. 1, pp. 19-24, Jan. 1983.
[49] B. W. Char, K. O. Geddes, W. M. Gentleman, and G. H. Gonnet, "The design of maple: A compact, portable, and powerful computer algebra system," in Proc. Eur. Conf. Comput. Algebra, 1983, pp. 101-115.
[50] S. Chevillard, M. Joldeş, and C. Lauter, "Sollya: An environment for the development of numerical codes," in Proc. Int. Congress Math. Softw., 2010, pp. 28-31.
[51] G. Estrin, "Organization of computer systems: The fixed plus variable structure computer," in Proc. Western Joint IRE-AIEE-ACM Comput. Conf., 1960, pp. 33-40. [Online]. Available: http://doi. acm.org/10.1145/1460361.1460365
[52] N. Shibata, "Efficient evaluation methods of elementary functions suitable for SIMD computation," Comput. Sci. - Res. Develop., vol. 25, no. 1, pp. 25-32, May 2010.
[53] D. Abrahams, D. Cabell, D. Smith, and E. Chan, "Boost Software License 1.0 (BSL-1.0)," 2003. [Online]. Available: https://www. boost.org/LICENSE_1_0.txt


Naoki Shibata received the PhD degree in computer science from Osaka University, Osaka, Japan, in 2001. He is currently an associate professor at Nara Institute of Science and Technology, Nara, Japan. He was an assistant professor at Nara Institute of Science and Technology, Nara, from 2001 to 2003 and an associate professor at Shiga University, Shiga, Japan from 2004 to 2012. His research interests include distributed and parallel systems and intelligent transportation systems. He is a member of IPSJ, ACM, and IEEE.


Francesco Petrogalli is currently a software engineer working on the development of Arm Compiler for HPC. He contributed to the implementation of the vector length agnostic (VLA) vectorizer for the scalable vector extension (SVE) of Arm, ABI specifications for the vector functions on AArch64, and open source library SLEEF. He has also worked on optimizing a variety of computational kernels that are core to machine learning algorithms.
$\triangleright$ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.

