

Mathematics and Speed for Interval Arithmetic

A Complement to IEEE 1788

Ulrich Kulisch

Institut für Angewandte und Numerische Mathematik
Karlsruher Institut für Technologie
Kaiserstrasse 12
D-76128 Karlsruhe GERMANY
Ulrich.Kulisch@kit.edu

Abstract. After a short introduction the paper begins with an axiomatic definition of rounded arithmetic. The concepts of rounding and of rounded arithmetic operations are defined in an axiomatic manner fully independent of special data formats and encodings. Basic properties of floating-point and interval arithmetic can directly be derived from this abstract mathematical model. Interval operations are defined as set operations for elements of the set $\overline{\mathbb{R}}$ of closed and connected sets of real numbers. As such they form an algebraically closed subset of the powerset of the real numbers. This property leads to explicit formulas for the arithmetic operations of floating-point intervals of $\overline{\mathbb{F}}$, which are executable on the computer. Arithmetic for intervals of $\overline{\mathbb{F}}$ forms an exception free calculus, i.e., arithmetic operations for intervals of $\overline{\mathbb{F}}$ always lead to intervals of $\overline{\mathbb{F}}$ again.

Later sections are concerned with programming support and hardware for interval arithmetic. Both are a must and absolutely necessary to move interval arithmetic more into the center of scientific computing. With some minor hardware additions interval operations can be made as fast as simple floating-point operations.

In vector and matrix spaces for real, complex, and interval data the dot product is a fundamental arithmetic operation. Computing the dot product of two vectors with floating-point components **exactly** substantially speeds up floating-point and interval arithmetic as well as the accuracy of the computed result. Hardware needed for the exact dot product is very modest. The exact dot product is essential for long real and long interval arithmetic.

Section 9 illustrates that interval arithmetic as developed in this paper has already a long tradition. Products based on these ideas have been available since 1980. Implementing what the paper advocates would have a profound effect on mathematical software. Modern processor architecture for example from Intel comes quite close to what is requested in this paper.¹

1 Introduction

A standard for floating-point arithmetic IEEE 754 was established in 1985. A revision was published in 2008. This paper was prepared in 2013 when a standard for interval arithmetic IEEE 1788 was under development. This standard was brought to a conclusion in 2015, [31]. While these two standards consider data formats, encodings, implementation, exception handling, and so on, this study focuses on the mathematics of the arithmetic for computers and on hardware support for the basic arithmetic operations. As usual mathematics is suited to focus the understanding on the essentials. High speed by hardware support is a must and vital for interval arithmetic to be more widely accepted and used by the scientific computing community.

Interval mathematics has been developed to a high standard over the last few decades. It provides methods which deliver results with guarantees. However, the arithmetic on existing processors makes these methods very slow. This paper deals with the question of how interval arithmetic can be provided effectively on computers. With more suitable processors, rigorous methods based on interval arithmetic could be comparable in speed to today's approximate methods. Interval arithmetic is a natural extension of floating-point arithmetic, not a replacement for it. As computers speed up,

¹ This article is published in: ACM Transactions on Mathematical Software, Vol. 45, No. 1, Article 5, February 2019.

from gigaflops to teraflops to petaflops, interval arithmetic becomes a principal and necessary tool for controlling the precision of a computation as well as the accuracy of the computed result.

Central for the considerations in this paper are operations for intervals of $\overline{\mathbb{R}}$. These are closed and connected sets of real numbers. For the notation of unbounded real intervals $-\infty$ and $+\infty$ are used as bounds. Since intervals are defined as sets of real numbers these bounds, however, are themselves not elements of the interval. Arithmetic for intervals of $\overline{\mathbb{R}}$ as well as for the subset of floating-point intervals of $\overline{\mathbb{F}}$ leads to a calculus that is free of exceptions. Thus arithmetic operations for intervals of $\overline{\mathbb{F}}$ always lead to intervals of $\overline{\mathbb{F}}$ again.

In vector and matrix spaces² the dot product of two vectors is a fundamental arithmetic operation in addition to the four operations $+$, $-$, $*$, $/$ for the components. Computing a verified enclosure of a system of linear equations mostly consists of dot products.

It is fascinating that this fundamental arithmetic operation also is a means to increase the speed of computing as well as the accuracy of the computed result. Actually the simplest and fastest way for computing a dot product is to compute it exactly. Here the exact products of double length are just shifted and added into a wide fixed-point register on the arithmetic unit. By pipelining, the exact dot product (EDP) can be computed in the time the processor needs to read the data, i.e., it comes with utmost speed. No other method computing a dot product can be faster. For details with respect to different data formats see References [23, 25].

The importance of the dot product for high speed verified computing has not been well understood by the scientific computing community for quite some time. So we mention it here briefly.

A guaranteed solution of a system of linear equations can be obtained in two steps. The first step I computes an approximate solution by some kind of Gaussian elimination in conventional floating-point arithmetic. The second step II, the verification step, then computes a highly accurate guaranteed enclosure of the solution. By an early estimate of Rump [36], the verification step can be done with less than 6 times the number of elementary floating-point operations needed for computing an approximation in step I. The verification step just consists of dot products. Hardware implementations of the EDP at Karlsruhe in 1993 [8, 9] and at Berkeley in 2013 [10] show that it can be computed in about one-sixth of the time needed for computing a possibly wrong result in conventional floating-point arithmetic. So, the EDP reduces the time needed for the verification step to about the time needed for computing an approximate solution by Gaussian elimination. This is a tremendous gain in computing speed. In other words, a verified solution of a system of linear equations can be computed in about twice the time needed for computing an approximate solution by some kind of Gaussian elimination.

The EDP can additionally already be applied in step I. This could reduce the computing time needed for step I and additionally increase the accuracy of the approximate solution.

The first sections of the paper derive and study basic mathematical properties of interval arithmetic. These are obtained by an abstract axiomatic approach. The representation in the paper is unique. Later sections then are concerned with speed. High speed for interval arithmetic can and should be obtained by the observation that the formulas for computing the lower and the upper bound of the five basic operations $+$, $-$, $*$, $/$, and the dot product are mutually independent of each other. So, they can and should be executed in parallel simultaneously. Thus, with some hardware investment each one of the five basic interval operations can be executed at the speed of the corresponding floating-point operation. Another great speed increase can be obtained by computing dot products exactly. High speed is essential for acceptance of interval arithmetic as a general computing tool.

A GAMM resolution required an EDP from arithmetic standards already in 1989 [1, 2]. The IFIP Working Group 2.5 on Numerical Software also required an EDP by letters to the IEEE standard committees 754 and 1788 in 2007 and 2009, respectively. In 2009, the author prepared a motion that required an EDP in IEEE 1788. The motion was accepted. Stimulated by this motion leading manufacturers now provide register memory on the arithmetic unit [41]. Access to register memory is much faster than access to main memory. About 1 K bit register memory suffice for high speed computation of the exact dot product for the double precision format. But, regrettably, IEEE 1788 withdrew the motion for an EDP in 2013. The author took this as an invitation to prepare the present paper. The subtitle of the paper, *A Complement to IEEE 1788*, indicates that the paper exceeds IEEE 1788 with respect to requirements such as speed and the EDP.

² for real, complex, interval, and complex interval data

In $\overline{\mathbb{R}}$ and $\overline{\mathbb{F}}$ division by an interval that includes zero also can be defined. With it, the Newton method reaches its ultimate elegance and power. The *extended interval Newton method* yields enclosures of all zeros of a function in a given domain. It is globally convergent and never fails, not even in rounded arithmetic [7, 25].

The reader may be well advised to begin reading this paper with section 9.

2 Axiomatic Definition of Rounded Arithmetic

Frequently mathematics is seen as the science of structures. Analysis carries three kinds of structures: an algebraic structure, an order structure, and a topological or metric structure. These are coupled by certain compatibility properties, as for instance: $a \leq b \Rightarrow a + c \leq b + c$.

It is well known that floating-point numbers and floating-point arithmetic do not obey the rules of the real numbers \mathbb{R} . However, the rounding is a monotone function. So the changes to the order structure are minimal. This is the reason why the order structure plays a key role for an axiomatic definition of computer arithmetic.

This study is restricted to the two elementary models that are covered by the two IEEE arithmetic standards, computer arithmetic on the reals and on real intervals. Abstract settings of computer arithmetic for higher dimensional spaces like complex numbers, vectors and matrices for real, complex, and interval data, can be developed following similar schemes. For more details see [25] and the literature cited there.

We begin by listing a few well-known concepts and properties of ordered sets.

Definition 1. *A relation \leq in a set M is called an order relation, and $\{M, \leq\}$ is called an ordered set³ if for all $a, b, c \in M$ the following properties hold:*

- (O1) $a \leq a$, (reflexivity)
- (O2) $a \leq b \wedge b \leq c \Rightarrow a \leq c$, (transitivity)
- (O3) $a \leq b \wedge b \leq a \Rightarrow a = b$, (antisymmetry)

An ordered set M is called linearly or totally ordered if in addition

- (O4) $a \leq b \vee b \leq a$ for all $a, b \in M$. (linearly ordered)

An ordered set M is called

- (O5) *a lattice if for any two elements $a, b \in M$, the $\inf\{a, b\}$ and the $\sup\{a, b\}$ exist.* (lattice)

- (O6) *It is called conditional completely ordered if for every bounded subset $S \subseteq M$, the $\inf S$ and the $\sup S$ exist.*

- (O7) *An ordered set M is called completely ordered or a complete lattice if for every subset $S \subseteq M$, the $\inf S$ and the $\sup S$ exist.* (complete lattice)

With these concepts, the real numbers $\{\mathbb{R}, \leq\}$ are defined as a conditional complete linearly ordered field.

In the definition of a complete lattice, the case $S = M$ is included. Therefore, $\inf M$ and $\sup M$ exist. Since they are elements of M , every complete lattice has a least and a greatest element.

If a subset $S \subseteq M$ of a complete lattice $\{M, \leq\}$ is also a complete lattice, $\{S, \leq\}$ is called a *complete sublattice* of $\{M, \leq\}$ if the two lattice operations \inf and \sup in both sets lead to the same result, i.e., if

$$\text{for all } A \subseteq S, \quad \inf_M A = \inf_S A \quad \text{and} \quad \sup_M A = \sup_S A.$$

Definition 2. *A subset S of a complete lattice $\{M, \leq\}$ is called a screen of M , if every element $a \in M$ has upper and lower bounds in S and the set of all upper and lower bounds of $a \in M$ has a least and a greatest element in S respectively. If a minus operator exists in M , a screen is called symmetric, if for all $a \in S$ also $-a \in S$.*

³ Occasionally called a partially ordered set.

As a consequence of this definition a complete lattice and a screen have the same least and greatest element. It can be shown that a screen is a complete sublattice of $\{M, \leq\}$ with the same least and greatest element, [25].

Definition 3. A mapping $\square : M \rightarrow S$ of a complete lattice $\{M, \leq\}$ onto a screen S is called a rounding if (R1) and (R2) hold:

- (R1) for all $a \in S$, $\square a := a$. (projection)
(R2) $a \leq b \Rightarrow \square a \leq \square b$. (monotone)

A rounding is called downwardly directed resp. upwardly directed if for all $a \in M$

- (R3) $\square a \leq a$ resp. $a \leq \square a$. (directed)

If a minus operator is defined in M , a rounding is called antisymmetric if

- (R4) $\square(-a) = -\square a$, for all $a \in M$. (antisymmetric)

The monotone downwardly resp. upwardly directed roundings of a complete lattice onto a screen are unique. For the proof, see [25].

Definition 4. Let $\{M, \leq\}$ be a complete lattice and $\circ : M \times M \rightarrow M$ a binary arithmetic operation in M . If S is a screen of M , then a rounding $\square : M \rightarrow S$ can be used to approximate the operation \circ in S by

- (RG) $a \boxtimes b := \square(a \circ b)$, for $a, b \in S$.

If a minus operator is defined in M and S is a symmetric screen of M , then a mapping $\square : M \rightarrow S$ with the properties (R1,2,4) and (RG) is called a semimorphism.⁴

Semimorphisms with antisymmetric roundings are particularly suited for transferring properties of the structure in M to the subset S . It can be shown [25] that semimorphisms leave a number of reasonable properties of ordered algebraic structures (ordered field, ordered vector space) invariant.

If an element $x \in M$ is bounded by $a \leq x \leq b$ with $a, b \in S$, then by (R1) and (R2) the rounded image $\square x$ is bounded by the same elements: $a \leq \square x \leq b$, i.e., $\square x$ is either the least upper (supremum) or the greatest lower (infimum) bound of x in S . Similarly, if for $x, y \in S$ the result of an operation $x \circ y$ is bounded by $a \leq x \circ y \leq b$ with $a, b \in S$, then by (R1), (R2), and (RG) also $a \leq x \boxtimes y \leq b$, i.e., $x \boxtimes y$ is either the least upper or the greatest lower bound of $x \circ y$ in S . If the rounding is upwardly or downwardly directed, the result is the least upper or the greatest lower bound, respectively.

3 Two Elementary Models

3.1 Floating-Point Arithmetic

The set $\{\overline{\mathbb{R}}, \leq\}$ with $\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$ is a complete lattice. Let \mathbb{F} denote the set of finite floating-point numbers and $\overline{\mathbb{F}} := \mathbb{F} \cup \{-\infty, +\infty\}$. Then $\overline{\mathbb{F}}$ is a screen of $\{\overline{\mathbb{R}}, \leq\}$. The least element of the set $\overline{\mathbb{R}}$ and of the subset $\overline{\mathbb{F}}$ is $-\infty$ and the greatest element is $+\infty$.

Definition 5. With a rounding $\square : \overline{\mathbb{R}} \rightarrow \overline{\mathbb{F}}$ arithmetic operations \boxtimes in $\overline{\mathbb{F}}$ are defined by

- (RG) $a \boxtimes b := \square(a \circ b)$, for $a, b \in \overline{\mathbb{F}}$ and $\circ \in \{+, -, *, /\}$,

with $b \neq 0$ in case of division.⁵

⁴ The properties (R1,2,4) and (RG) of a semimorphism can be shown to be necessary conditions for a homomorphism between ordered algebraic structures. For more details see [25]

⁵ In real analysis division by zero is not defined. It does not lead to a real number.

If a and b are adjacent floating-point numbers and $x \in \mathbb{R}$ with $a \leq x \leq b$, then because of (R1) and (R2) also $a \leq \square x \leq b$, i.e., there is never an element of \mathbb{F} between an element $x \in \mathbb{R}$ and its rounded image $\square x$. The same property holds for the operations defined by (RG): If for $x, y \in \mathbb{F}$, $a \leq x \circ y \leq b$ then by (R1), (R2), and (RG) also $a \leq x \square y \leq b$, for all $\circ \in \{+, -, *, /\}$, i.e., $x \square y$ is either the greatest lower or the least upper bound of $x \circ y$ in \mathbb{F} .

Frequently used roundings $\square : \mathbb{R} \rightarrow \overline{\mathbb{F}}$ are antisymmetric. Examples are the rounding to the nearest floating-point number, the rounding toward zero, or the rounding away from zero. A semimorphism transfers a number of useful properties of the real numbers to the floating-point numbers. The mathematical structure of \mathbb{F} can even be defined as properties of \mathbb{R} that are invariant with respect to semimorphism, [25].

For the monotone downwardly resp. upwardly directed roundings of $\overline{\mathbb{R}}$ onto $\overline{\mathbb{F}}$ often the special symbols ∇ resp. \triangle are used. These roundings are not antisymmetric. They are related by the property:

$$\nabla(-a) = -\triangle a \quad \text{and} \quad \triangle(-a) = -\nabla a. \quad (1)$$

Arithmetic operations defined by (RG) and these roundings are denoted by ∇ and \triangle , respectively, for $\circ \in \{+, -, *, /\}$. These are heavily used in interval arithmetic.

3.2 Interval Arithmetic

The set of nonempty, closed and bounded real intervals is denoted by \mathbb{IR} . An interval of \mathbb{IR} is denoted by an ordered pair $[a_1, a_2]$, with $a_1, a_2 \in \mathbb{R}$ and $a_1 \leq a_2$. The first element is the lower bound and the second is the upper bound.

Interval arithmetic over the real numbers deals with *closed*⁶ and connected sets of real numbers. Such intervals can be bounded or unbounded. For the notation of unbounded intervals $-\infty$ and $+\infty$ are used as bounds. Since $-\infty$ and $+\infty$ are not real numbers a bound $-\infty$ or $+\infty$ of an unbounded interval is not an element of the interval. Thus unbounded real intervals are frequently written as $(-\infty, a]$ or $[b, +\infty)$ or $(-\infty, +\infty)$ with $a, b \in \mathbb{R}$.

The set of bounded and unbounded real intervals is denoted by $\overline{\mathbb{IR}}$. With respect to the subset relation as an order relation the set of real intervals $\{\overline{\mathbb{IR}}, \subseteq\}$ is a complete lattice. The subset of $\overline{\mathbb{IR}}$ where all finite bounds are floating-point numbers of \mathbb{F} is denoted by $\overline{\mathbb{IF}}$. $\{\overline{\mathbb{IF}}, \subseteq\}$ is a screen of $\{\overline{\mathbb{IR}}, \subseteq\}$. In both sets $\overline{\mathbb{IR}}$ and $\overline{\mathbb{IF}}$, the infimum of a subset of $\overline{\mathbb{IF}}$ is the intersection, and the supremum is the interval hull⁷. The least element of both sets $\overline{\mathbb{IR}}$ and $\overline{\mathbb{IF}}$ is the empty set \emptyset and the greatest element is the set $\mathbb{R} = (-\infty, +\infty)$.

Definition 6. For intervals $\mathbf{a}, \mathbf{b} \in \overline{\mathbb{IR}}$ arithmetic operations $\circ \in \{+, -, *, /\}$ are defined as set operations

$$\mathbf{a} \circ \mathbf{b} := \{a \circ b \mid a \in \mathbf{a} \wedge b \in \mathbf{b}\}. \quad (2)$$

Here, for division we assume that $[0, 0] \notin \mathbf{b}$.

For $\mathbf{a} = [a_1, a_2] \in \overline{\mathbb{IR}}$ we obtain by (2) immediately⁸

$$-\mathbf{a} := (-1) * \mathbf{a} = [-a_2, -a_1] \in \overline{\mathbb{IR}}. \quad (3)$$

If in (3) $\mathbf{a} \in \overline{\mathbb{IF}}$, then also $-\mathbf{a} \in \overline{\mathbb{IF}}$, i.e., $\overline{\mathbb{IF}}$ is a symmetric screen of $\overline{\mathbb{IR}}$.

Between the complete lattice $\{\overline{\mathbb{IR}}, \subseteq\}$ and its screen $\{\overline{\mathbb{IF}}, \subseteq\}$ the monotone upwardly directed rounding $\diamond : \overline{\mathbb{IR}} \rightarrow \overline{\mathbb{IF}}$ is uniquely defined. It is characterized by the following properties:

- (R1) $\diamond \mathbf{a} = \mathbf{a}$, for all $\mathbf{a} \in \overline{\mathbb{IF}}$. (projection)
- (R2) $\mathbf{a} \subseteq \mathbf{b} \Rightarrow \diamond \mathbf{a} \subseteq \diamond \mathbf{b}$, for $\mathbf{a}, \mathbf{b} \in \overline{\mathbb{IR}}$. (monotone)
- (R3) $\mathbf{a} \subseteq \diamond \mathbf{a}$, for all $\mathbf{a} \in \overline{\mathbb{IR}}$. (upwardly directed)

⁶ A set of real numbers is called *closed*, if its complement in \mathbb{R} is open. A set $M \subset \mathbb{R}$ is called *open* if for all $a \in M$ also an ϵ -neighborhood $(a - \epsilon, a + \epsilon)$ is entirely in M .

⁷ Which is the convex hull in the one dimensional case.

⁸ An integral number a in an interval expression is interpreted as interval $[a, a]$.

For $\mathbf{a} = [a_1, a_2] \in \overline{\mathbb{IR}}$ the result of the monotone upwardly directed rounding \diamond is

$$\diamond \mathbf{a} = [\nabla a_1, \Delta a_2]. \quad (4)$$

Using (1) and (2) it is easy to see that the monotone upwardly directed rounding $\diamond : \overline{\mathbb{IR}} \rightarrow \overline{\mathbb{IF}}$ is antisymmetric, i.e.,

$$(R4) \quad \diamond(-\mathbf{a}) = -\diamond \mathbf{a}, \text{ for all } \mathbf{a} \in \overline{\mathbb{IR}}. \quad (\text{antisymmetric}).$$

An interval $\mathbf{a} = [a_1, a_2]$ is frequently interpreted as a point in \mathbb{R}^2 . This very naturally induces the order relation \leq of \mathbb{R}^2 to the set of intervals $\overline{\mathbb{IR}}$. For two intervals $\mathbf{a} = [a_1, a_2]$ and $\mathbf{b} = [b_1, b_2]$ the relation \leq is defined by $\mathbf{a} \leq \mathbf{b} :\Leftrightarrow a_1 \leq b_1 \wedge a_2 \leq b_2$.

For the \leq relation for intervals compatibility properties hold between the algebraic structure and the order structure in great similarity to the real numbers. For instance:

$$\begin{aligned} (OD1) \quad & \mathbf{a} \leq \mathbf{b} \Rightarrow \mathbf{a} + \mathbf{c} \leq \mathbf{b} + \mathbf{c}, \text{ for all } \mathbf{c}. \\ (OD2) \quad & \mathbf{a} \leq \mathbf{b} \Rightarrow -\mathbf{b} \leq -\mathbf{a}. \\ (OD3) \quad & [0, 0] \leq \mathbf{a} \leq \mathbf{b} \wedge \mathbf{c} \geq [0, 0] \Rightarrow \mathbf{a} * \mathbf{c} \leq \mathbf{b} * \mathbf{c}. \\ (OD4) \quad & [0, 0] < \mathbf{a} \leq \mathbf{b} \wedge \mathbf{c} > [0, 0] \Rightarrow [0, 0] < \mathbf{a}/\mathbf{c} \leq \mathbf{b}/\mathbf{c} \wedge \mathbf{c}/\mathbf{a} \geq \mathbf{c}/\mathbf{b} > [0, 0]. \end{aligned}$$

With respect to set inclusion as an order relation, arithmetic operations in $\{\overline{\mathbb{IR}}, \subseteq\}$ are inclusion isotone by (2), i.e., $\mathbf{a} \subseteq \mathbf{b} \Rightarrow \mathbf{a} \circ \mathbf{c} \subseteq \mathbf{b} \circ \mathbf{c}$ or equivalently

$$(OD5) \quad \mathbf{a} \subseteq \mathbf{b} \wedge \mathbf{c} \subseteq \mathbf{d} \Rightarrow \mathbf{a} \circ \mathbf{c} \subseteq \mathbf{b} \circ \mathbf{d}, \text{ for all } \circ \in \{+, -, *, /\}, 0 \notin \mathbf{b}, \mathbf{d} \text{ for } \circ = /. \text{ (inclusion isotone)}$$

Setting $\mathbf{c}, \mathbf{d} = -1$ in (OD5) delivers immediately $\mathbf{a} \subseteq \mathbf{b} \Rightarrow -\mathbf{a} \subseteq -\mathbf{b}$ which differs significantly from (OD2).

Definition 7. *With the upwardly directed rounding $\diamond : \overline{\mathbb{IR}} \rightarrow \overline{\mathbb{IF}}$ binary arithmetic operations in $\overline{\mathbb{IF}}$ are defined by semimorphism:*

$$(RG) \quad \mathbf{a} \diamond \mathbf{b} := \diamond(\mathbf{a} \circ \mathbf{b}), \text{ for all } \mathbf{a}, \mathbf{b} \in \overline{\mathbb{IF}} \text{ and all } \circ \in \{+, -, *, /\}.$$

Here, for division we assume that \mathbf{a}/\mathbf{b} is defined.

If an interval $\mathbf{a} \in \overline{\mathbb{IF}}$ is an upper bound of an interval $\mathbf{x} \in \overline{\mathbb{IR}}$, i.e., $\mathbf{x} \subseteq \mathbf{a}$, then by (R1), (R2), and (R3) also $\mathbf{x} \subseteq \diamond \mathbf{x} \subseteq \mathbf{a}$. This means $\diamond \mathbf{x}$ is the least upper bound, the supremum of \mathbf{x} in $\overline{\mathbb{IF}}$. Similarly if for $\mathbf{x}, \mathbf{y} \in \overline{\mathbb{IF}}$, $\mathbf{x} \circ \mathbf{y} \subseteq \mathbf{a}$ with $\mathbf{a} \in \overline{\mathbb{IF}}$, then by (R1), (R2), (R3), and (RG) also $\mathbf{x} \circ \mathbf{y} \subseteq \mathbf{x} \diamond \mathbf{y} \subseteq \mathbf{a}$, i.e., $\mathbf{x} \diamond \mathbf{y}$ is the least upper bound, the supremum of $\mathbf{x} \circ \mathbf{y}$ in $\overline{\mathbb{IF}}$. Occasionally the supremum $\mathbf{x} \diamond \mathbf{y}$ of the result $\mathbf{x} \circ \mathbf{y} \in \overline{\mathbb{IR}}$ is called the tightest enclosure of $\mathbf{x} \circ \mathbf{y}$.

Arithmetic operations in $\overline{\mathbb{IF}}$ are inclusion isotone, i.e.,

$$(OD5) \quad \mathbf{a} \subseteq \mathbf{b} \wedge \mathbf{c} \subseteq \mathbf{d} \Rightarrow \mathbf{a} \diamond \mathbf{c} \subseteq \mathbf{b} \diamond \mathbf{d}, \text{ for } \circ \in \{+, -, *, /\}, 0 \notin \mathbf{b}, \mathbf{d} \text{ for } \circ = /. \text{ (inclusion isotone)}$$

This is a consequence of the inclusion isotony of the arithmetic operations in $\overline{\mathbb{IR}}$, of (R2) and of (RG).

Since the arithmetic operations $\mathbf{x} \circ \mathbf{y}$ in $\overline{\mathbb{IR}}$ are defined as set operations by (2) the operations $\mathbf{x} \diamond \mathbf{y}$ for intervals of $\overline{\mathbb{IF}}$ defined by (RG) are not directly executable. The step from the definition of interval arithmetic by set operations to computer executable operations still requires some effort. This holds in particular for product sets like complex intervals and intervals of vectors and matrices of real and complex numbers. For more details see Reference[25]. We sketch it here for the most simple case of floating-point intervals.

3.3 Executable Interval Arithmetic

We first derive executable formulas for the arithmetic operations for nonempty, closed and bounded real intervals $\mathbf{a}, \mathbf{b} \in \mathbb{IR}$. These are then extended to unbounded intervals of $\overline{\mathbb{IR}}$. Arithmetic operations for intervals are defined as set operations by $\mathbf{a} \circ \mathbf{b} := \{a \circ b \mid a \in \mathbf{a} \wedge b \in \mathbf{b}\}$. For $\mathbf{a}, \mathbf{b} \in \mathbb{IR}$ with $0 \notin \mathbf{b}$ in case of division for all arithmetic operations $\circ \in \{+, -, *, /\}$, the function $a \circ b$ is a continuous function of both variables. The set $\mathbf{a} \circ \mathbf{b}$ is the range of the function $a \circ b$ over the product set $\mathbf{a} \times \mathbf{b}$.

Since \mathbf{a} and \mathbf{b} are closed and bounded intervals $\mathbf{a} \times \mathbf{b}$ is a simply connected, closed and bounded subset of \mathbb{R}^2 . In such a region the continuous function $a \circ b$ takes a minimum and a maximum as well as all values in between. Therefore

$$\mathbf{a} \circ \mathbf{b} = \left[\min_{a \in \mathbf{a}, b \in \mathbf{b}} \{a \circ b\}, \max_{a \in \mathbf{a}, b \in \mathbf{b}} \{a \circ b\} \right] = [\inf(\mathbf{a} \circ \mathbf{b}), \sup(\mathbf{a} \circ \mathbf{b})],$$

i.e., for $\mathbf{a}, \mathbf{b} \in \mathbb{IR}$, $0 \notin \mathbf{b}$ in case of division, $\mathbf{a} \circ \mathbf{b}$ is again an interval of \mathbb{IR} . In other words: **Arithmetic in \mathbb{IR} is an algebraically closed subset of the arithmetic in the powerset of the real numbers.**

Actually, the minimum and maximum is taken for operations with the bounds. For bounded intervals $\mathbf{a} = [a_1, a_2]$ and $\mathbf{b} = [b_1, b_2]$ the following formula holds for all operations with $0 \notin \mathbf{b}$ in case of division:

$$\mathbf{a} \circ \mathbf{b} = \left[\min_{i,j=1,2} (a_i \circ b_j), \max_{i,j=1,2} (a_i \circ b_j) \right] \text{ for } \circ \in \{+, -, *, /\}. \quad (5)$$

We demonstrate this in case of addition. By (OD1) we obtain $a_1 \leq \mathbf{a}$ and $b_1 \leq \mathbf{b} \Rightarrow a_1 + b_1 \leq \inf(\mathbf{a} + \mathbf{b})$. On the other hand $\inf(\mathbf{a} + \mathbf{b}) \leq a_1 + b_1$. From both inequalities we obtain by (O3): $\inf(\mathbf{a} + \mathbf{b}) = a_1 + b_1$. Analogously one obtains $\sup(\mathbf{a} + \mathbf{b}) = a_2 + b_2$. Thus

$$\mathbf{a} + \mathbf{b} = \left[\min_{a \in \mathbf{a}, b \in \mathbf{b}} \{a + b\}, \max_{a \in \mathbf{a}, b \in \mathbf{b}} \{a + b\} \right] = [\inf(\mathbf{a} + \mathbf{b}), \sup(\mathbf{a} + \mathbf{b})] = [a_1 + b_1, a_2 + b_2].$$

Similarly by making use of (OD1,2,3,4) for intervals of \mathbb{IR} and the simple sign rules $-(\mathbf{a} * \mathbf{b}) = (-\mathbf{a}) * \mathbf{b} = \mathbf{a} * (-\mathbf{b})$, $-(\mathbf{a}/\mathbf{b}) = (-\mathbf{a})/\mathbf{b} = \mathbf{a}/(-\mathbf{b})$ explicit formulas for all interval operations can be derived, [25].

Now we get by (RG) for intervals of \mathbb{IF}

$$\mathbf{a} \diamond \mathbf{b} := \diamond(\mathbf{a} \circ \mathbf{b}) = \left[\nabla \min_{i,j=1,2} (a_i \circ b_j), \Delta \max_{i,j=1,2} (a_i \circ b_j) \right]$$

and by the monotonicity of the roundings ∇ and Δ :

$$\mathbf{a} \diamond \mathbf{b} = \left[\min_{i,j=1,2} (a_i \nabla b_j), \max_{i,j=1,2} (a_i \Delta b_j) \right].$$

For bounded and nonempty intervals $\mathbf{a} = [a_1, a_2]$ and $\mathbf{b} = [b_1, b_2]$ of \mathbb{IF} the unary operation $-\mathbf{a}$ and the binary operations addition, subtraction, multiplication, and division are shown in the following tables. For details see [25]. Therein the operator symbols for intervals are simply denoted by $+$, $-$, $*$, $/$.

Minus operator $-\mathbf{a} = [-a_2, -a_1]$.

Addition $[a_1, a_2] + [b_1, b_2] = [a_1 \nabla b_1, a_2 \Delta b_2]$.

Subtraction $[a_1, a_2] - [b_1, b_2] = [a_1 \nabla b_2, a_2 \Delta b_1]$.

Multiplication	$[b_1, b_2]$	$[b_1, b_2]$	$[b_1, b_2]$
$[a_1, a_2] * [b_1, b_2]$	$b_2 \leq 0$	$b_1 < 0 < b_2$	$b_1 \geq 0$
$[a_1, a_2], a_2 \leq 0$	$[a_2 \nabla b_2, a_1 \Delta b_1]$	$[a_1 \nabla b_2, a_1 \Delta b_1]$	$[a_1 \nabla b_2, a_2 \Delta b_1]$
$a_1 < 0 < a_2$	$[a_2 \nabla b_1, a_1 \Delta b_1]$	$[\min(a_1 \nabla b_2, a_2 \nabla b_1), [a_1 \nabla b_2, a_2 \Delta b_2]$	$[\max(a_1 \Delta b_1, a_2 \Delta b_2)]$
$[a_1, a_2], a_1 \geq 0$	$[a_2 \nabla b_1, a_1 \Delta b_2]$	$[a_2 \nabla b_1, a_2 \Delta b_2]$	$[a_1 \nabla b_1, a_2 \Delta b_2]$

Division, $0 \notin \mathbf{b}$	$[b_1, b_2]$	$[b_1, b_2]$
$[a_1, a_2]/[b_1, b_2]$	$b_2 < 0$	$b_1 > 0$
$[a_1, a_2], a_2 \leq 0$	$[a_2 \nabla b_1, a_1 \triangle b_2]$	$[a_1 \nabla b_1, a_2 \triangle b_2]$
$[a_1, a_2], a_1 < 0 < a_2$	$[a_2 \nabla b_2, a_1 \triangle b_2]$	$[a_1 \nabla b_1, a_2 \triangle b_1]$
$[a_1, a_2], 0 \leq a_1$	$[a_2 \nabla b_2, a_1 \triangle b_1]$	$[a_1 \nabla b_2, a_2 \triangle b_1]$

In real analysis division by zero is not defined. In interval arithmetic, however, the interval in the denominator of a quotient may contain zero. So this case has to be considered also.

The general rule for computing the set \mathbf{a}/\mathbf{b} with $0 \in \mathbf{b}$ is to remove its zero from the interval \mathbf{b} and perform the division with the remaining set.⁹ Whenever zero in \mathbf{b} is an endpoint of \mathbf{b} , the result of the division can be obtained directly from the above table for division with $0 \notin \mathbf{b}$ by the limit process $b_1 \rightarrow 0$ or $b_2 \rightarrow 0$ respectively. The results are shown in the table for division with $0 \in \mathbf{b}$. Here, the round brackets stress that the bounds $-\infty$ and $+\infty$ are not elements of the interval.

Division, $0 \in \mathbf{b}$	$\mathbf{b} =$	$[b_1, b_2]$	$[b_1, b_2]$
$[a_1, a_2]/[b_1, b_2]$	$[0, 0]$	$b_1 < b_2 = 0$	$0 = b_1 < b_2$
$[a_1, a_2] = [0, 0]$	\emptyset	$[0, 0]$	$[0, 0]$
$[a_1, a_2], a_1 < 0, a_2 \leq 0$	\emptyset	$[a_2 \nabla b_1, +\infty)$	$(-\infty, a_2 \triangle b_2]$
$[a_1, a_2], a_1 < 0 < a_2$	\emptyset	$(-\infty, +\infty)$	$(-\infty, +\infty)$
$[a_1, a_2], 0 \leq a_1, 0 < a_2$	\emptyset	$(-\infty, a_1 \triangle b_1)$	$[a_1 \nabla b_2, +\infty)$

In addition to these operations division by an interval that contains zero as an interior point also can be defined. When zero is an interior point of the denominator, the set $[b_1, b_2]$ splits into the distinct sets $[b_1, 0)$ and $(0, b_2]$, and the division by $[b_1, b_2]$ actually means two divisions. The results of the two divisions are already shown in the table for division by $0 \in \mathbf{b}$.

However, if applied in a user's program the two divisions appear as a single operation, as division by an interval $[b_1, b_2]$ with $b_1 < 0 < b_2$, an operation that delivers two distinct results. The result could be returned as an improper interval where the left hand bound is higher than the right hand bound.

A solution to the problem would be for the computer to provide a flag for *distinct intervals*. The situation occurs if the divisor is an interval that contains zero as an interior point. In this case the flag would be raised and signaled to the user. The user may then apply a routine of his choice to deal with the situation as is appropriate for his application. This routine could be: return the entire set of real numbers $(-\infty, +\infty)$ as result and continue the computation, or continue the computation with one of the sets and ignore the other one, or put one of the sets on a list and continue the computation with the other one, or modify the operands and recompute, or stop computing, or some other action.

An alternative would be to provide a second division which in case of division by an interval that contains zero as an interior point generally delivers the result $(-\infty, +\infty)$. Then the user can decide when to use which division in his program.

Deriving explicit formulas for the result of interval operations $\mathbf{a} \circ \mathbf{b}$, $\circ \in \{+, -, *, /\}$, it was assumed so far that the intervals \mathbf{a} and \mathbf{b} are nonempty and bounded. Four kinds of *extended intervals* come from division by an interval of \mathbb{IF} that contains zero:

$$\emptyset, \quad (-\infty, a], \quad [b, +\infty), \quad \text{and} \quad (-\infty, +\infty).$$

⁹ This is in full accordance with function evaluation: When evaluating a function over a set, points outside its domain are simply ignored.

To extend the operations to these more general intervals the first rule is that any operation with the empty set \emptyset returns the empty set. Then the above tables extend to possibly unbounded intervals of $\overline{\mathbb{F}}$ by using the standard formulae for arithmetic operations involving $\pm\infty$ together with one rule that goes beyond standard rules for $\pm\infty$:

$$0 * (-\infty) = (-\infty) * 0 = 0 * (+\infty) = (+\infty) * 0 = 0.$$

This rule is not a new mathematical law, it is merely a short cut to compute the bounds of the result of multiplication on unbounded intervals. Intervals of $\overline{\mathbb{F}}$ are closed and connected sets of real numbers. $-\infty$ and $+\infty$ are not elements of these intervals. So multiplication of any such interval by 0 can only have 0 as the result:

$$(-\infty, a] * 0 = [b, +\infty) * 0 = (-\infty, +\infty) * 0 = 0.$$

In the previous section intervals of $\overline{\mathbb{R}}$ are defined as closed and connected sets of real numbers. Arithmetic for intervals of $\overline{\mathbb{F}}$ as derived here leads to an exception free calculus, i.e., arithmetic operations for intervals of $\overline{\mathbb{F}}$ always lead to intervals of $\overline{\mathbb{F}}$ again.

This is in sharp contrast to other models of interval arithmetic which consider intervals over the extended real numbers $\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$ where $-\infty$ and $+\infty$ are taken as numbers. In such models obscure arithmetic operations like $\infty - \infty, \infty/\infty, 0 * \infty$ occur which require introduction of unnatural superficial objects like *NaI* (Not an Interval).

4 Interval Arithmetic in Higher Dimensions

The axioms for computer arithmetic shown in section 2 also can be applied to define computer arithmetic in higher dimensional spaces like complex numbers, vectors and matrices for real, complex, and interval data.

If M in section 2 is the set of real matrices and $S \subseteq M$ the set of floating-point matrices, then the definition of arithmetic operations in S by (RG) ideally requires exact evaluation of scalar products of vectors with floating-point components. Indeed very effective algorithms have been developed for computing scalar products of vectors with floating-point components exactly, [23, 25]. For a brief sketch see section 7.

If M in section 2 is the set of intervals of real vectors or matrices, respectively, then the set definition of arithmetic operations in M by (2) does not lead to an interval again. The result $\mathbf{a} \circ \mathbf{b} := \{\mathbf{a} \circ \mathbf{b} \mid \mathbf{a} \in \mathbf{a} \wedge \mathbf{b} \in \mathbf{b}\}$ is a more general set. It is an element of the power set¹⁰ of vectors or matrices, respectively. To obtain an interval the upwardly directed rounding \square from the power set onto the set of intervals of M has to be applied. With it arithmetic operations for intervals $\mathbf{a}, \mathbf{b} \in M$ are defined by

$$(RG) \mathbf{a} \boxtimes \mathbf{b} := \square(\mathbf{a} \circ \mathbf{b}), \circ \in \{+, -, *\}.$$

The set S of intervals of floating-point vectors or matrices, respectively, is a screen of M . To obtain arithmetic for intervals $\mathbf{a}, \mathbf{b} \in S$ once more the monotone upwardly directed rounding, now denoted by \diamond is applied:

$$(RG) \mathbf{a} \diamond \mathbf{b} := \diamond(\mathbf{a} \boxtimes \mathbf{b}), \circ \in \{+, -, *\}.$$

This leads to the best possible operations in the interval spaces M and S . So for intervals in higher dimensional spaces generally an additional rounding step is needed. It can be shown that in all relevant cases these best possible operations can be expressed by executable formulas just using the bounds of the intervals. Showing this for more complicated cases (e.g. complex intervals or intervals of real and complex vectors and matrices) requires a more detailed study. For details see [25] and the literature cited there. The dot product of two vectors the components of which are floating-point numbers is a key operation in all these spaces. Computing it exactly can considerably speed up floating-point and interval arithmetic.

¹⁰ The power set of a set M is the set of all subsets of M .

5 Operations with Directed Roundings and Programming Languages

The IEEE standard 754 for floating-point arithmetic seems to support interval arithmetic. It requires the four basic arithmetic operations to be available with rounding to nearest, towards zero, downwards and upwards. The latter two are essential for interval arithmetic. But almost all processors that provide IEEE 754 arithmetic separate the rounding from the basic operation, which proves to be a severe drawback. In a conventional floating-point computation this does not cause any difficulties. The rounding mode is set only once. Then a large number of operations is performed with this rounding mode, one in every cycle. However, when interval arithmetic is performed the rounding mode has to be switched very frequently. The lower bound of the result of every interval operation has to be rounded downwards and the upper bound rounded upwards. Thus, the rounding mode has to be reset for every arithmetic operation. If setting the rounding mode and the arithmetic operation are equally fast this slows down the computation of each bound unnecessarily by a factor of two in comparison to conventional floating-point arithmetic. On almost all existing commercial processors, however, setting the rounding mode takes a multiple (three, five, ten) of the time that is needed for the arithmetic operation. Thus an interval operation is unnecessarily at least eight (or twenty and even more) times slower than the corresponding floating-point operation not counting the necessary case distinctions for interval multiplication and interval division. This is fatal for interval arithmetic. The rounding must be an integral part of the arithmetic operation. Every one of the rounded arithmetic operations with rounding to nearest, downwards or upwards should be equally fast.

For interval arithmetic we need to be able to call each of the operations ∇ , ∇ , ∇ , ∇ and \triangle , \triangle , \triangle , \triangle as one single operation, that is, the rounding must be inherent to each. Therefore in programming languages notations for arithmetic operations with different roundings shall be provided. They could be:

+ , - , * , / for operations with rounding to the nearest floating-point number,
+> , -> , *> , /> for operations with rounding upwards,
+< , -< , *< , /< for operations with rounding downwards, and
+| , -| , *| , /| for operations with rounding towards zero (chopping).

New operation codes are necessary! Each of these operations shall be executed as one single operation.

Frequently used programming languages regrettably do not provide four plus, minus, multiply, and divide operators for floating-point numbers. This, however, does not justify generally separating the rounding from the arithmetic operation!

Instead, a floating-point arithmetic standard should require that **every future processor** shall provide the 16 operations listed above. In languages like C++ which just provide operator overloading and do not allow user defined operators these operations could be written as functions:

addp, subp, mulp, divp, and addn, subn, muln, divn.

Here *p* stands for rounding toward positive and *n* for rounding toward negative infinity. With these operators interval operations can easily be programmed. They would be very fast and fully transferable from one processor to another.

In such languages the operations with the directed roundings are hidden within the interval operations. If operations with directed roundings are needed, interval operations are performed instead and the upper and lower bound of the result then gives the desired answer.

This is exactly the way interval arithmetic is provided in the C++ class library C-XSC which has been successfully used for more than 20 years [17, 18]. Since C++ only allows operator overloading, the operations with the directed roundings cannot and are not provided in C-XSC. These operations are rarely needed in applications.

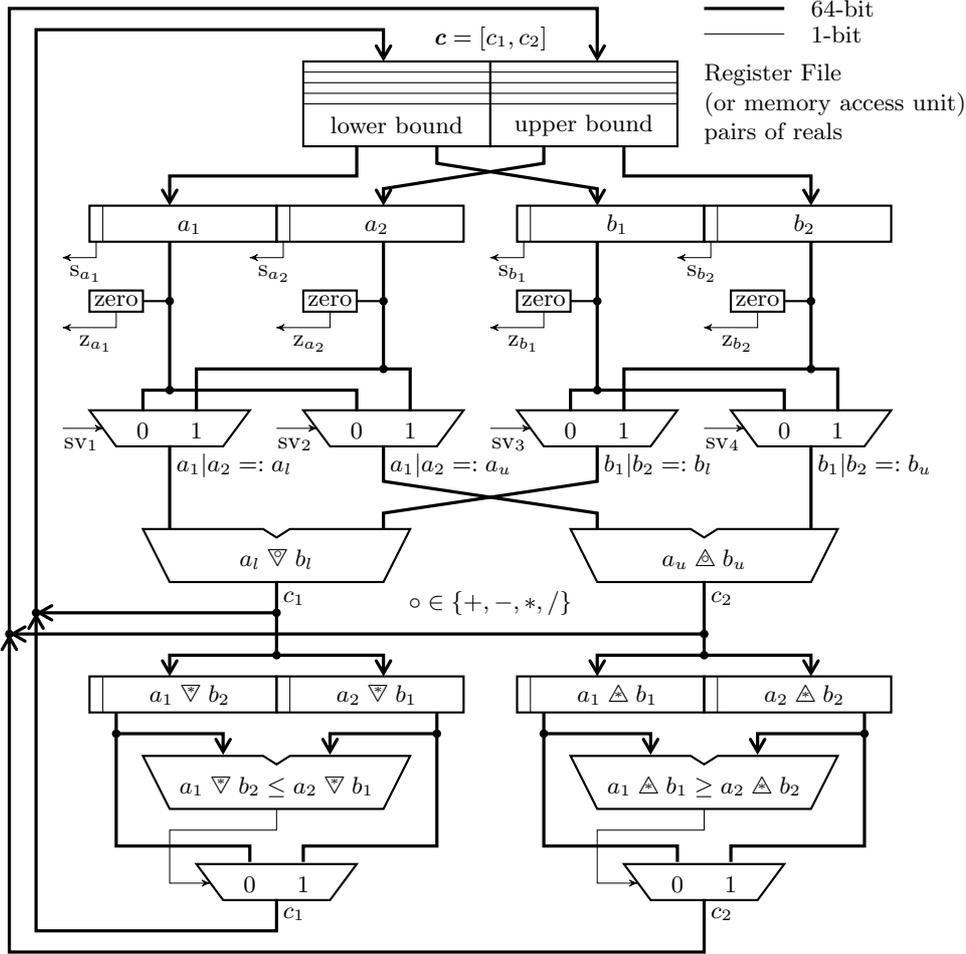
In Fortran operators like *.addp.* could be used for instance.

6 Hardware for Interval Arithmetic

Figure 1 gives a brief sketch of what hardware support for interval arithmetic may look like. It would not be hard to realize it in modern technology.

The circuitry broadly speaks for itself. The lower and the upper bound of the result of an operation are computed simultaneously. The interval operands are loaded in parallel from a register file or a

memory access unit. Then, after multiplexers have selected the appropriate operands, the lower bound of the result is computed with rounding downwards and the upper bound with rounding upwards with the selected operands. If in case of multiplication both operands contain zero as an interior point a second multiplication is necessary. In the figure it is assumed that the two multiplications are performed sequentially and that the results of both multiplications are delivered to the lower part of the circuitry. They are then forwarded to a comparison unit. Here for the lower bound of the result the lower and for the upper bound the higher of the two products is selected. This lower part of the circuitry could also be used to perform comparison relations.



operands: $\mathbf{a} = [a_1, a_2]$, $\mathbf{b} = [b_1, b_2]$, result: $\mathbf{c} = [c_1, c_2]$.
 s: sign, z: zero, o: operand select.

Fig. 1. Circuitry for Interval Operations

In the figure it is assumed that the second product is computed with the same circuitry sequentially. This slows down interval multiplication in comparison with the other interval operations. This can be avoided by doubling the circuitry for multiplication and performing the two multiplications in parallel.

The selector signals sv_1, sv_2, sv_3 , and sv_4 control the multiplexers. The circuitry shown in the figure performs all arithmetic operations.

In **addition**, the lower bounds of \mathbf{a} and \mathbf{b} are simply added with rounding downwards and the upper bounds with rounding upwards $[a_1 \nabla b_1, a_2 \triangle b_2]$. The selector signals are set to $sv_1 = 0$, $sv_2 = 1$, $sv_3 = 0$, and $sv_4 = 1$.

In **subtraction**, the bounds of \mathbf{b} are exchanged by the operand selection. Then the subtraction is performed, the lower bound being computed with rounding downwards and the upper bound with

Multiplication	$B = 0$	$b_1 < 0, b_2 \leq 0$	$b_1 < 0 < b_2$	$0 \leq b_1, 0 < b_2$
$A = 0$	$[0, 0]$	$[0, 0]$	$[0, 0]$	$[0, 0]$
$a_1 < 0, a_2 \leq 0$	$[0, 0]$	$[a_2 \nabla b_2, a_1 \triangle b_1]$ $sv = 1100$	$[a_1 \nabla b_2, a_1 \triangle b_1]$ $sv = 0100$	$[a_1 \nabla b_2, a_2 \triangle b_1]$ $sv = 0110$
$a_1 < 0 < a_2$	$[0, 0]$	$[a_2 \nabla b_1, a_1 \triangle b_1]$ $sv = 1000$	$[\min(a_1 \nabla b_2, a_2 \nabla b_1), \max(a_1 \triangle b_1, a_2 \triangle b_2)]$ $sv = 0100, sv = 1011$	$[a_1 \nabla b_2, a_2 \triangle b_2]$ $sv = 0111$
$0 \leq a_1, 0 < a_2$	$[0, 0]$	$[a_2 \nabla b_1, a_1 \triangle b_2]$ $sv = 1001$	$[a_2 \nabla b_1, a_2 \triangle b_2]$ $sv = 1011$	$[a_1 \nabla b_1, a_2 \triangle b_2]$ $sv = 0011$

Table 1. Interval multiplication.

Division	$B = 0$	$b_2 < 0$	$b_1 < b_2 = 0$	$b_1 < 0 < b_2$	$0 = b_1 < b_2$	$0 < b_1$
$A = 0$	\emptyset	$[0, 0]$	$[0, 0]$	$[0, 0]$	$[0, 0]$	$[0, 0]$
$a_1 < 0, a_2 \leq 0$	\emptyset	$[a_2 \nabla b_1, a_1 \triangle b_2]$ $sv = 1001$	$[a_2 \nabla b_1, +\infty)$ $sv = 10xx$	$(-\infty, a_2 \triangle b_2]$ $\cup [a_2 \nabla b_1, +\infty)$ $sv = 1011$	$(-\infty, a_2 \triangle b_2]$ $sv = xx11$	$[a_1 \nabla b_1, a_2 \triangle b_2]$ $sv = 0011$
$a_1 < 0 < a_2$	\emptyset	$[a_2 \nabla b_2, a_1 \triangle b_2]$ $sv = 1101$	$(-\infty, +\infty)$	$(-\infty, +\infty)$	$(-\infty, +\infty)$	$[a_1 \nabla b_1, a_2 \triangle b_1]$ $sv = 0010$
$0 \leq a_1, 0 < a_2$	\emptyset	$[a_2 \nabla b_2, a_1 \triangle b_1]$ $sv = 1100$	$(-\infty, a_1 \triangle b_1]$ $sv = xx00$	$(-\infty, a_1 \triangle b_1]$ $\cup [a_1 \nabla b_2, +\infty)$ $sv = 0100$	$[a_1 \nabla b_2, +\infty)$ $sv = 01xx$	$[a_1 \nabla b_2, a_2 \triangle b_1]$ $sv = 0110$

Table 2. Interval division.

rounding upwards $[a_1 \nabla b_2, a_2 \triangle b_1]$. The selector signals are set to $sv_1 = 0, sv_2 = 1, sv_3 = 1$, and $sv_4 = 0$.

These two operations, addition and subtraction, reveal a close relationship between the arithmetic operations shown in Tables 1 and 2 and the operand selector signals which select the lower or the upper bound of the operands \mathbf{a} and \mathbf{b} for the operation to be performed. The index of the operand component to be selected minus 1 gives the value of the selector signal $sv_i, i = 1, 2, 3, 4$ for the selection of the operand.

This simple rule also holds for the operand selections in the cases of multiplication and division. In the Tables 1 and 2 the selector signals sv_1, sv_2, sv_3, sv_4 are summarized into a selection vector $sv = (sv_1, sv_3, sv_2, sv_4)$ which is shown directly below the formulas for the arithmetic operations. In Table 2 both cases $0 \notin \mathbf{b}$ and $0 \in \mathbf{b}$ are merged into a single table.

For intervals $\mathbf{a}, \mathbf{b} \in \overline{\mathbb{R}}$ subtraction is defined by $\mathbf{a} - \mathbf{b} := \mathbf{a} + (-1) * \mathbf{b}$. Thus subtraction could be reduced to addition by negation of the operand \mathbf{b} .

7 High Speed Computing by an Exact Dot Product

The IFIP Working Group on Numerical Software and other scientists repeatedly requested that a future arithmetic standard should consider and specify an exact dot product (EDP) [5, 6], [1, 2]. In Nov. 2009 the IEEE standards committee P1788 for interval arithmetic accepted a motion for including the EDP into a future interval arithmetic standard. In Numerical Analysis the dot product is ubiquitous. It is a fundamental arithmetic operation in vector and matrix spaces for real, complex, and interval data. It is the EDP which makes residual correction effective. This has a direct and positive influence on all iterative solvers of systems of equations. The EDP is essential for fast long real and long interval arithmetic [25]. By operator overloading variable precision interval arithmetic is very easy to use. With it the result of every arithmetic expression can be guaranteed to a specified number of correct digits.

A later motion revised the motion for the EDP in 2013. IEEE 1788 now requires a correctly rounded dot product instead. Computing this via an EDP is recommended. The literature contains many excellent and intelligent methods for computing “accurate” sums and dot products or a correctly rounded dot product in particular. Several of these methods are claimed to be *fast* or even *ultimately fast*. All these methods just use conventional floating-point arithmetic. A first step computes an approximate solution. A second step then improves this to a correctly rounded dot product. The speed of the faster of these methods comes considerably close to the time T needed for computing a possibly wrong result in conventional floating-point arithmetic. The speed mildly decreases with increasing condition number.

Computing dot products exactly on the contrary considerably speeds up computing. Actually the simplest and fastest way for computing a dot product is to compute it exactly. The unrounded exact products are just shifted and added into a long fixed-point register on the arithmetic unit, see Figure 2. Fixed-point accumulation of the products is simpler than accumulation in floating-point. Many intermediate steps that are executed in a floating-point accumulation such as normalization and rounding of the products and of the intermediate sum, composition into a floating-point number and decomposition into mantissa and exponent for the next operation do not occur in the fixed-point accumulation. It simply is performed by a shift and addition of the products of full double length into a wide fixed-point register on the arithmetic unit. By pipelining the stages: loading the data, computing, shifting, and accumulation of the products, the EDP can be computed in the time the processor needs to read the data, i.e., it comes with utmost speed. No other method can be faster. Here the computing time is independent of the condition number. It is much less than T . This brings a substantial speed increase compared to a possibly wrong accumulation of the dot product using conventional floating-point arithmetic. Fixed-point addition is error free! So, the EDP is a means to considerably speed up floating-point and interval arithmetic. Hardware implementations of the EDP at Karlsruhe in 1993 [8, 9] and at Berkeley in 2013 [10] show that it can be computed in about 1/6th of the time needed for computing a possibly wrong result in conventional floating-point arithmetic. Fixed-point accumulation of the EDP into a small register memory of about 1K bit on the arithmetic unit can be faster than computing a correctly rounded dot product in software by more than one magnitude.

This section briefly illustrates how the EDP can be implemented on computers. There is indeed no simpler way of accumulating a dot product. Any method that just computes an approximation also has to consider the relative values of the summands. This results in a more complicated method.

Let $(a_i), (b_i) \in \mathbb{F}^n$ be two vectors with n components which are floating-point numbers $a_i, b_i \in \mathbb{F}(b, l, e1, e2)$, for $i = 1(1)n$. Here b is the base of the number system in use, l the number of digits of the mantissa, $e1$ is the least and $e2$ the greatest exponent respectively. We compute the sum

$$s := \sum_{i=1}^n a_i \cdot b_i = a_1 \cdot b_1 + a_2 \cdot b_2 + \dots + a_n \cdot b_n, \quad a_i \cdot b_i \in \mathbb{F}(b, 2 \cdot l, 2 \cdot e1, 2 \cdot e2), \text{ for } i = 1(1)n,$$

where all additions and multiplications are the operations for real numbers.

All summands can be taken into a fixed-point register of length $2 \cdot e2 + 2 \cdot l + 2 \cdot |e1|$ without loss of information. We call it a *complete register*, CR for short.

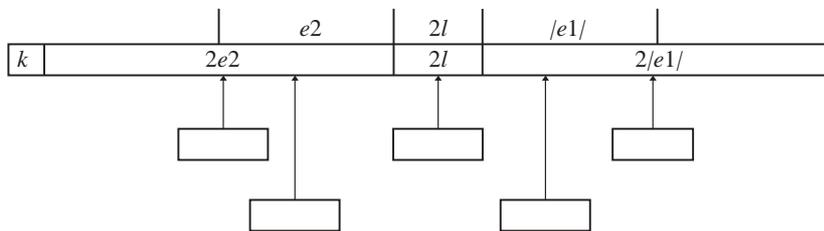


Fig. 2. Complete register for exact scalar product accumulation.

If the register is built as an accumulator with an adder, all summands could be added in without loss of information. To accommodate possible overflows, it is convenient to provide a few, say k , more

digits of base b on the left, allowing b^k accumulations to be computed without loss of information due to overflow. k can be chosen such that no overflows of the long register will occur in the lifetime of the computer.

A. For the IEEE-arithmetic *double precision* we have:

$b = 2$; 64 bits word length; 1 bit sign; 11 bits exponent; $l = 53$ bits; $e1 = -1022$, $e2 = 1023$. With $k = 92$ the entire unit consists of

$L = k + 2 \cdot e2 + 2 \cdot l + 2 \cdot |e1| = k + 4196$ bits = 4288 bits. It can be represented by 67 words of 64 bits. L is independent of n .

B. For the S/370 architecture data format *long* we have:

$b = 16$; word length: 64 bits; sign: 1 bit; mantissa: $l = 14$ hex digits; exponent range: $e1 = -64$, $e2 = 63$, hexadecimal.

$L = k + 2e2 + 2l + 2|e1| = k + 282$ digits of base 16.

With $k = 88$ bits we obtain $L = 88 + 4 \cdot 282 = 1,216$ bits. This register can be represented by 19 words of 64 bits.

In case B the size of the complete register looks much more convenient than in case A. The complete register was provided as in case B in the IBM products ACRITH and ACRITH-XSC [37–39], the Siemens product ARITHMOS [40], and in corresponding Hitachi and NAS products in the late 1980ies. Never any serious problems occurred because of the smaller exponent range.

In a floating-point computation an overflow in general means a total breakdown of the accuracy, not so in interval arithmetic. Intervals bring the continuum on the computer. An overflow of the upper bound of an interval leads to an unbounded real interval but not to a total breakdown of the accuracy, since in any case the lower bound is a finite floating-point number. For the next operation the result can already be a finite interval again.

Figure 3 informally describes the implementation of an EDP. The complete register (here represented as a chest of drawers) is organized in words of 64 bits. The exponent of the products consists of 12 bits. The leading 6 bits give the address of the three consecutive drawers to which the summand of 106 bits is added. The low end 6 bits of the exponent are used for the correct positioning of the summand within the selected drawers. A possible carry is absorbed by the next more significant word in which not all bits are 1. For fast detection of this word a flag is attached to each register word. It is set 1 if all bits of the word are 1. This means that a carry will propagate through the entire word. In the figure the flag is shown as a red (dark) point. As soon as the exponent of the summand is available the flags allow selecting and incrementing the carry word. This can be done simultaneously with adding the summand into the selected positions. Figure 4 shows a sketch for the parallel accumulation of a product into the complete register.

By pipelining, the accumulation of a product into the complete register can be done in the time the processor needs to read the data. Since every other method of computing a dot product also has to read the data this means that no such method can exceed computing the EDP in speed. For the pipelining and other solutions, see References [23] or [25]. Rounding the EDP into a correctly rounded dot product is done only once at the very end of the accumulation.

In References [23] and [25] three different solutions for pipelining the dot product computation are dealt with. They differ by the speed with which the vector components for a product can be read into the scalar product unit. They can be delivered in 32- or 64-bit portions or even at once as one 128 bit word. With increasing bus width and speed more hardware has to be invested for the multiplication and the accumulation to keep the pipeline in balance.

The hardware cost needed for the EDP is modest. It is comparable to that for a fast multiplier by an adder tree, accepted years ago and now standard technology in every modern processor. The EDP brings the same speedup for accumulations at comparable costs.

Figures 3 and 4 informally specify the implementation of the EDP on computers. For more details, see References [23, 25, 27]. Possibilities to compute the EDP by only making use of a hardware *complete register* window are also discussed there.

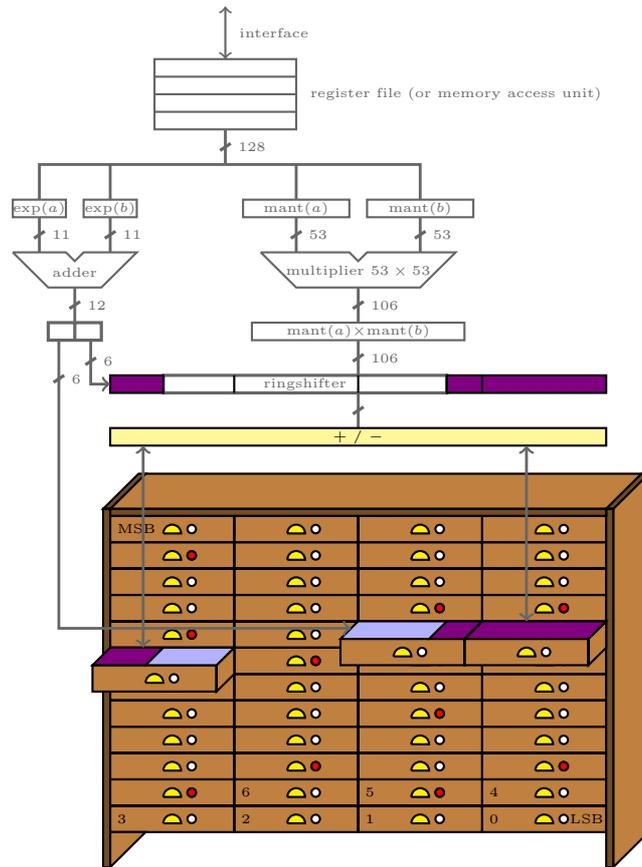


Fig. 3. Illustration for computing the exact dot product.

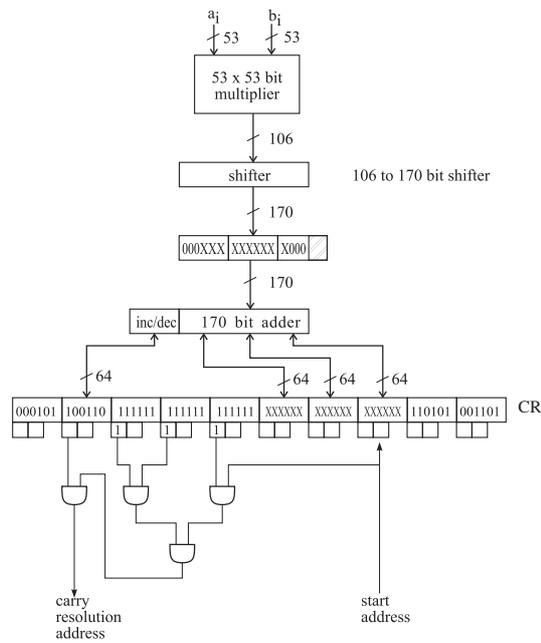


Fig. 4. Parallel accumulation of a product into the CR.

To be completely successful interval arithmetic occasionally has to be complemented by some easy way to use multiple precision real and interval arithmetic. The fast and exact dot product is the tool to provide these operations. By operator overloading multiple precision arithmetic can be used very easily. Multiple precision real and interval operations for addition, subtraction, multiplication, division, the scalar product, and the square root are defined in Reference [25]. In Reference [34] the authors show that even extremely ill conditioned matrices with a condition number of about 10^{100} can successfully be inverted using such tools.

In 1984 a hardware unit for the EDP was developed in bit-slice technology [13]. This technology was expensive and did not allow production of the unit in large quantities. Then in 1993 the EDP was implemented on a single chip using a VLSI gate-array technology [9]. It was connected with the PC via the PCI-bus and embedded into powerful PASCAL-XSC and C-XSC programming environments. Although the technology available for the chip was considerably weaker than that available for the PC, the chip computed the EDP in a fraction of the time the Intel processor needed for computing a possibly wrong result in conventional floating-point arithmetic.

Ultimately the EDP must be integrated into the arithmetic unit of the processor. A little register space on the arithmetic unit will be rewarded by a substantial increase in speed and reliability of computing.

Interval arithmetic brings guarantees and mathematics into computing, while the EDP brings higher precision, accuracy, and speed!

8 Early Super Computers

It is interesting that the technique for computing dot products exactly is not new at all. It can be traced back to the early computer by G. W. Leibniz (1675). Also old commercial mechanic calculators added numbers and products of numbers into a wide fixed-point register. It was the fastest way to use the computer. So it was applied as often as possible. No intermediate results needed to be written down and typed in again for the next operation. No intermediate roundings or normalizations had to be performed. No error analysis was necessary. As long as no underflow or overflow occurred, which would be obvious and visible, the result was always exact. It was independent of the order in which the summands were added. Rounding was only done, if required, at the very end of the accumulation.

This extremely useful and fast fifth arithmetic operation was not built into the early floating-point computers. It was too expensive for the technologies of those days. Later its superior properties had been forgotten. Thus, floating-point arithmetic is still comparatively incomplete.



Fig. 5. Mechanical computing devices equipped with the desired capability: **Burkhardt Arithmometer**, Glashütte, Germany, 1878; **Brunsviga**, Braunschweig, Germany, 1917; **MADAS**, Zürich, Switzerland, 1936; **MONROE**, New Jersey, USA, 1956.

The two lower calculators in Figure 5 are equipped with more than one long result register. In the previous section it was called a complete register CR. It is an early version of the recommended new data format *complete*.

In summary it can be said: The technique of speeding up computing by accumulation of numbers and of products of numbers exactly into a wide fixed-point register is as old as technical computing itself. It proves an excellent feeling of the old mathematicians for efficiency in computing.

Also early super computers (until ca. 2005) got their high speed by pipelining the dot product (vector processing). They provided so-called compound operations like *accumulate* or *multiply and accumulate*. The second computes the sum of products, the dot product of two vectors. Advanced programming languages offered these operations. Pipelining made them really fast. A vectorizing compiler filled them into a users program as often as possible. However, the accumulation was done in floating-point arithmetic by the so-called partial sum technique. This altered the sequence of the summands and caused errors beyond the conventional floating-point errors. So, finally this technique was abolished.

Fixed-point accumulation of the dot product, as discussed in the previous section, is simpler and faster than accumulation in floating-point arithmetic. In a very natural pipeline the accumulation is done in the time that is needed to read the data into the arithmetic unit. This means that no other method of computing a dot product can be faster, in particular not a conventional computation in double or quadruple precision floating-point arithmetic. Fixed-point accumulation is error free! It is high speed vector processing in its perfection.

9 Early Products and Motivation

Between 1976 and 1980 a powerful programming environment PASCAL-XSC for support of interval arithmetic was developed at the author's institute together with a major number of problem solving routines for all kinds of standard problems of numerical analysis. These routines delivered close bounds for the solution. The project was supported by Nixdorf Computers. The Z-80 based system was exhibited at the Hannover Exhibition in April 1980. Nixdorf Computers donated a number of such systems to the University at Karlsruhe. This made it possible to decentralize the programming education in the summer semester 1980 while batch processing and punch cards and punch tapes were standard elsewhere.

When IBM became aware of this development, a ten-year-long cooperation with the author's institute was started in 1980. The aim was to develop corresponding Fortran based products for their /370 architecture which at that time dominated the market. The corresponding IBM program products were called ACRITH (1983) and ACRITH-XSC (1990) [37–39]. Siemens and Hitachi followed soon by developing similar products for their /370 based systems. The Siemens product was called ARITHMOS [40]. The commercial success of these new products, however, was modest. During the development time other architectures appeared. The PC entered the market in 1982 and several new workstations in 1983/84. All these new systems made use of the new computer arithmetic standard IEEE 754 which was officially available since 1985. Scientific computing slowly but steadily moved from the old /370 systems to these more modern looking architectures.

In adapting the new situation the old Z-80 based PASCAL-XSC system was immediately transferred to the PC and to workstations at the author's institute. Books on PASCAL-XSC were published in German, English and Russian. See [14–16]. The book and the PASCAL-XSC system are now freely available. Further, after appearance of C and C++ a corresponding C-XSC was developed as a C++ class library again with a major number of problem solving routines [17, 18]. These books are also freely available.

High speed by support of hardware and programming languages is vital for interval arithmetic to be more widely accepted by the scientific computing community. The IEEE standard 754 seems to support interval arithmetic. It requires floating-point operations with the directed roundings. But in practice these have to be simulated by slow software. No commercial processor provides them directly by hardware. In 2008 the author published a book entitled: *Computer Arithmetic and Validity - Theory, Implementation, and Applications* [25]. It puts considerable emphasis on speeding up interval arithmetic. It extends the accuracy requirements for floating-point arithmetic as given by the IEEE standard 754 to the usual product spaces of computation like complex numbers and vectors and matrices for real, complex, and interval data. The book shows and states already in its preface that

interval arithmetic for all these spaces can efficiently be provided on the computer if two additional features are made available by fast hardware:

**I. Fast and direct hardware support for double precision interval arithmetic and
II. a fast and exact multiply and accumulate operation or, an exact dot product (EDP).**

Realization of I. and II. is discussed at detail in the book [25]. It is shown that I. and II. can be obtained at very little hardware cost. With I. interval arithmetic would be as fast as simple floating-point arithmetic. The simplest and fastest way for computing a dot product is to compute it exactly. To make II. conveniently available a new data format *complete* is used together with a few very restricted arithmetic operations. By pipelining, the EDP can be computed in the time the processor needs to read the data, i.e., it comes with utmost speed. I. and II. would boost both the speed of a computation and the accuracy of the result. Fast hardware support for I. and II. must be added to conventional floating-point arithmetic. Both are necessary extensions. Hardware implementations of the exact dot product (EDP) at Karlsruhe in 1993 [8,9] and at Berkeley [10] in 2013 show that the EDP can be computed in about 1/6th of the time needed for computing a possibly wrong result in conventional floating-point arithmetic.

In Reference [25], a real interval is defined as a closed and connected set of real numbers. Such intervals can be bounded or unbounded. In the latter case $-\infty$ or $+\infty$ occur as bounds of the interval. They are, however, not elements of the interval. Six months after publication of the book [25] the IEEE Computer Society founded a standard committee P1788 for interval arithmetic in August 2008. Work on the standard began immediately. Much emphasis is put on specifying details. The definition of a real interval as a closed and connected set of real numbers was accepted. While this paper was still in the process of being reviewed a Standard for Interval Arithmetic has been approved: IEEE 1788-2015 - Standard for Interval Arithmetic. <http://standards.ieee.org/findstds/standard/1788-2015.html>, [31] Demanding particular hardware and language support for interval arithmetic is not seen as a task of the standard.

This article is intended to complement the development of the IEEE 1788 standard by clearly stating what hardware and language support is vital and must be made available for interval arithmetic to increase its acceptance in the scientific computing community and to move interval arithmetic more into the mainstream of scientific computing.

Influenced by the requirement of an EDP modern processor architecture by leading manufacturers is coming considerably close to what is requested in this paper. See [41], and in particular pp.1-1 to 1-3 and 2-5 to 2-6. These processors provide register space of 16 K bits on the arithmetic unit. Only about 1 K bits suffice for a *complete register* and high speed computation of the EDP.

Acknowledgement: The author owes thanks to Goetz Alefeld for useful comments on the paper. He also wishes to express his thanks to the four anonymous referees. Their comments and suggestions helped to shorten and improve the paper.

References

1. IMACS and GAMM, IMACS-GAMM resolution on computer arithmetic, *Mathematics and Computers in Simulation* 31 (1989), 297–298, or in *Zeitschrift für Angewandte Mathematik und Mechanik* 70:4 (1990).
2. GAMM-IMACS proposal for accurate floating-point vector arithmetic, *GAMM Rundbrief* 2 (1993), 9–16, and *Mathematics and Computers in Simulation*, Vol. 35, IMACS, North Holland, 1993. News of IMACS, Vol. 35, No. 4, 375–382, Oct. 1993.
3. American National Standards Institute / Institute of Electrical and Electronics Engineers: *A Standard for Binary Floating-Point Arithmetic*. ANSI/IEEE Std. 754-1987, New York, 1985. (reprinted in SIG-PLAN **22**, 2, pp. 9-25, 1987). Also adopted as IEC Standard 559:1989, Revised version 2008, and in ISO/IEC/IEEE 60559:2011.
4. American National Standards Institute / Institute of Electrical and Electronics Engineers: *A Standard for Radix-Independent Floating-Point Arithmetic*. ANSI/IEEE Std. 854-1987, New York, 1987.
5. The *IFIP WG 2.5 - IEEE 754R letter*, dated September 4, 2007.
6. The *IFIP WG 2.5 - IEEE P1788 letter*, dated September 9, 2009.
7. G. Alefeld, *Intervallrechnung über den komplexen Zahlen und einige Anwendungen*, Dissertation, Universität Karlsruhe, 1968.
8. Ch. Baumhof, *A new VLSI vector arithmetic coprocessor for the PC*, in: Institute of Electrical and Electronics Engineers (IEEE), S. Knowles and W. H. McAllister (eds.), *Proceedings of the 12th Symposium on*

- Computer Arithmetic ARITH*, Bath, England, July 19–21, 1995, pp. 210–215, IEEE Computer Society Press, Piscataway, NJ, 1995.
9. Ch. Baumhof, *Ein Vektorarithmetik-Koprozessor in VLSI-Technik zur Unterstützung des Wissenschaftlichen Rechnens*, Dissertation, Universität Karlsruhe, 1996.
 10. D. Biancolin and J. Koenig, *Hardware Accelerator for Exact Dot Product*, ASPIRE Laboratory, University of California, Berkeley, 2015.
 11. G. Bohlender, *Floating-Point Computation of Functions with Maximum Accuracy*, IEEE Transactions on Computers, Vol. C-26, no. 7, July 1977.
 12. G. Bohlender, *Genaue Berechnung mehrfacher Summen, Produkte und Wurzeln von Gleitkommazahlen und allgemeine Arithmetik in höheren Programmiersprachen*, Dissertation, Universität Karlsruhe, 1978.
 13. T. Teufel, *Ein optimaler Gleitkommapprozessor*, Dissertation, Universität Karlsruhe, 1984.
 14. R. Klatte, U. Kulisch, M. Neaga, D. Ratz and Ch. Ullrich, *PASCAL-XSC – Sprachbeschreibung mit Beispielen*, Springer, Berlin Heidelberg New York, 1991.
See also <http://www2.math.uni-wuppertal.de/xsc/> or <http://www.xsc.de/>.
 15. R. Klatte, U. Kulisch, M. Neaga, D. Ratz and Ch. Ullrich, *PASCAL-XSC – Language Reference with Examples*, Springer, Berlin Heidelberg New York, 1992.
See also <http://www2.math.uni-wuppertal.de/xsc/> or <http://www.xsc.de/>.
Russian translation MIR, Moscow, 1995, third edition 2006.
See also <http://www2.math.uni-wuppertal.de/xsc/> or <http://www.xsc.de/>.
 16. R. Hammer, M. Hocks, U. Kulisch and D. Ratz, *Numerical Toolbox for Verified Computing I: Basic Numerical Problems (PASCAL-XSC)*, Springer, Berlin Heidelberg New York, 1993.
Russian translation MIR, Moskau, 2005.
 17. R. Klatte, U. Kulisch, C. Lawo, M. Rauch and A. Wiethoff, *C-XSC – A C++ Class Library for Extended Scientific Computing*, Springer, Berlin Heidelberg New York, 1993.
See also <http://www2.math.uni-wuppertal.de/xsc/> or <http://www.xsc.de/>.
 18. R. Hammer, M. Hocks, U. Kulisch and D. Ratz, *C++ Toolbox for Verified Computing: Basic Numerical Problems*. Springer, Berlin Heidelberg New York, 1995.
 19. R. Kirchner, U. Kulisch, 2006, *Hardware support for interval arithmetic*. Reliable Computing, Vol.12, 2006, 225–237.
 20. U. Kulisch, *An axiomatic approach to rounded computations*, TS Report No. 1020, Mathematics Research Center, University of Wisconsin, Madison, Wisconsin, 1969, and *Numerische Mathematik* 19 (1971), 1–17.
 21. U. Kulisch, *Implementation and Formalization of Floating-Point Arithmetics*, IBM T. J. Watson-Research Center, Report Nr. RC 4608, 1 - 50, 1973. Invited talk at the Caratheodory Symposium, Sept. 1973 in Athens, published in: The Greek Mathematical Society, C. Caratheodory Symposium, 328 - 369, 1973, and in *Computing* 14, 323–348, 1975.
 22. U. Kulisch, *Grundlagen des Numerischen Rechnens - Mathematische Begründung der Rechnerarithmetik*, Bibliographisches Institut, Mannheim Wien Zürich, 1976.
 23. U. Kulisch, *Advanced Arithmetic for the Digital Computer – Design of Arithmetic Units*, Springer, 2002.
 24. U. Kulisch, *Complete interval arithmetic and its implementation on the computer*, in: A. Cuyt, et al. (eds.), *Numerical Validation in Current Hardware Architectures*, LNCS, Vol. 5492, pp. 7–26, Springer-Verlag, Heidelberg, 2008.
 25. U. Kulisch, *Computer Arithmetic and Validity – Theory, Implementation, and Applications*, de Gruyter, Berlin, 2008, second edition 2013.
 26. U. Kulisch, *Arithmetic Operations for Floating-Point Intervals*, as Motion 5 accepted by the IEEE Standards Committee P1788 as definition of the interval operations. See [30].
 27. U. Kulisch, V. Snyder, *The Exact Dot Product as Basic Tool for Long Interval Arithmetic*. *Computing*, Vol.91, 2011, 307–313.
 28. U. Kulisch, *An Axiomatic Approach to Computer Arithmetic with an appendix on Interval Hardware*, LNCS, Springer-Verlag, Heidelberg, 484–495, 2012.
 29. U. Kulisch, *Up-to-Date Interval Arithmetic – From closed intervals to connected sets of real numbers*, Springer LNCS 9574, pp. 413 - 434, 2016.
 30. J. D. Pryce (Ed.), *P1788, IEEE Standard for Interval Arithmetic*, <http://grouper.ieee.org/groups/1788/email/pdfOWdtH2mOd9.pdf>.
 31. IEEE 1788-2015 - Standard for Interval Arithmetic. <http://standards.ieee.org/findstds/standard/1788-2015.html>.
 32. F. Blomquist, W. Hofschuster, W. Krämer, *A Modified Staggered Correction Arithmetic with Enhanced Accuracy and Very Wide Exponent Range*. In: A. Cuyt et al. (eds.): *Numerical Validation in Current Hardware Architectures*, Lecture Notes in Computer Science LNCS, vol. 5492, Springer-Verlag, Berlin Heidelberg, 41–67, 2009.
 33. M. Nehmeier, S. Siegel, J. Wolff von Gudenberg. 2012. *Specification of Hardware for Interval Arithmetic*, *Computing*, Vol. 92, 2012, 243–255.

34. S. Oishi, K. Tanabe, T. Ogita and S. M. Rump, *Convergence of Rump's method for inverting arbitrarily ill-conditioned matrices*, Journal of Computational and Applied Mathematics 205 (2007), 533–544.
35. S. M. Rump, *Kleine Fehlerschranken bei Matrixproblemen*, Dissertation, Universität Karlsruhe, 1980.
36. S. M. Rump and E. Kaucher, *Small bounds for the Solution of Systems of Linear Equations*, *Computing Supplementum 2* 157–164, 1980.
37. IBM, *IBM System/370 RPQ. High Accuracy Arithmetic*, SA 22-7093-0, IBM Deutschland GmbH (Department 3282, Schönaicher Strasse 220, D-71032 Böblingen), 1984.
38. IBM, *IBM High-Accuracy Arithmetic Subroutine Library (ACRITH)*, IBM Deutschland GmbH (Department 3282, Schönaicher Strasse 220, D-71032 Böblingen), third edition, 1986.
 1. General Information Manual, GC 33-6163-02.
 2. Program Description and User's Guide, SC 33-6164-02.
 3. Reference Summary, GX 33-9009-02.
39. IBM, *ACRITH-XSC: IBM High Accuracy Arithmetic – Extended Scientific Computation. Version 1, Release 1*, IBM Deutschland GmbH (Department 3282, Schönaicher Strasse 220, D-71032 Böblingen), 1990.
 1. General Information, GC33-6461-01.
 2. Reference, SC33-6462-00.
 3. Sample Programs, SC33-6463-00.
 4. How To Use, SC33-6464-00.
 5. Syntax Diagrams, SC33-6466-00.
40. SIEMENS, *ARITHMOS (BS 2000) Unterprogrammbibliothek für Hochpräzisionsarithmetik. Kurzbeschreibung, Tabellenheft, Benutzerhandbuch*, SIEMENS AG, Bereich Datentechnik, Postfach 83 09 51, D-8000 München 83, Bestellnummer U2900-J-Z87-1, September 1986.
41. INTEL, *Intel Architecture Instruction Set Extensions Programming Reference*, 319433-017, December 2013, <http://software.intel.com/en-us/file/319433-017pdf>.