

P1788/D7.3
Draft Standard For Interval Arithmetic

John Pryce and Christian Keil, Technical Editors

DRAFT 7.3

DRAFT 7.3

Frontmatter

Required IEEE items to be added:

- Draft copyright statements
- Title
- Abstract and keywords
- Committee lists
- Acknowledgments (after the Introduction)

⚠ **To the P1788 working group reader.**

General. Passages in this color are my editorial comments: mostly asking for answers to or debate on a question; or giving my opinion; or noting changes made.

Introduction

This introduction explains some of the alternative interpretations, and sometimes competing objectives, that influenced the design of this standard, but is not part of the standard.

Mathematical context. Interval computation is a collaboration between human programmer and machine infrastructure which, correctly done, produces mathematically proven numerical results about continuous problems—for instance, rigorous bounds on the global minimum of a function or the solution of a differential equation. It is part of the discipline of “constructive real analysis”. In the long term, the results of such computations may become sufficiently trusted to be accepted as contributing to legal decisions. The machine infrastructure acts as a body of theorems on which the correctness of an interval algorithm relies, so it must be made as reliable as is practical. In its logical chain are many links—hardware, underlying floating-point system, etc.—over which this standard has no control. The standard aims to strengthen one specific link, by defining interval objects and operations that are theoretically well-founded and practical to implement.

This document uses the standard notation $[a, b]$ for “the interval between numbers a and b ”, with various detailed meanings depending on the underlying theory. The “classical” interval arithmetic (IA) of R.A. Moore [5] uses only bounded, closed, nonempty intervals in the real numbers \mathbb{R} —that is, $[a, b] = \{x \in \mathbb{R} \mid a \leq x \leq b\}$ where $a, b \in \mathbb{R}$ with $a \leq b$. So, for instance, division by an interval containing 0 is not defined in it. It was agreed early on that this standard should strictly extend classical IA in virtue of allowing an interval to be unbounded or empty.

Beyond this, various extensions of classical IA were considered. One choice that distinguishes between theories is: Are arithmetic operations purely algebraic, or do they involve topology? An example of the latter is containment set (cset) theory [8], which extends functions over the reals to functions over the extended reals, e.g. $\sin(+\infty)$ is the set of all possible limits of $\sin x$ as $x \rightarrow +\infty$, which is $[-1, 1]$. The complications of this were deemed to outweigh the advantages, and it was agreed that operations should be purely algebraic.

Another choice is: Is an interval a set—a subset of the number line—or is it something different? The most widely used forms of IA are *set-based* and define an interval to be a set of real numbers. They have established software to find validated solutions of linear and nonlinear algebraic equations, optimization problems, differential equations, etc.

However *Kaucher* IA and the nearly equivalent *modal* IA have significant applications. In the former an interval is formally a pair (a, b) of real numbers, which for $a \leq b$ is “proper” and identified with the normal interval $\{x \in \mathbb{R} \mid a \leq x \leq b\}$, and for $a > b$ is “improper”. In the latter, an interval is a pair (X, Q) where X is a normal interval and Q is a quantifier, either \exists or \forall . At the time of writing it finds commercial use in the graphics rendering industry. Both forms are referred to as Kaucher IA henceforth.

In view of their significance it was decided to support both set-based and Kaucher IA. Because of their different mathematical bases this led to the concept of *flavors* (see Clause 7). A flavor is a version of IA that extends classical IA in a precisely defined sense, such that when only classical intervals and restricted operations are used (avoiding, e.g., division by an interval containing zero), all flavors produce the same results, at the mathematical level and also—up to roundoff—in finite precision.

Currently the standard incorporates two flavors, set-based and Kaucher. Others could be added, though there are no current plans to do so. E.g., csets could be a flavor, since they also extend classical IA in the defined sense.

To minimize complexity, the standard for each flavor is presented as a separate sub-document within the overall standard, readable as a self-contained unit without reference to other flavors, and with common clauses that specify how a flavor extends classical IA.

The set-based flavor is presented first, on the grounds that it is relatively easy to grasp, easy to teach, and easy to interpret in the context of real-world applications. In this theory:

- Intervals are sets.
- They are subsets of the set \mathbb{R} of real numbers. At the mathematical level (Level 1 in the structure defined in Clause 5) they are precisely all topologically closed and connected subsets of \mathbb{R} . The finite-precision level (Level 2), uses the notion of an interval type, which is a finite set of Level 1 intervals.
- The interval version of an elementary function such as $\sin x$ is essentially the natural algebraic extension to sets of the corresponding pointwise function on real numbers.

Fuzzy sets, like intervals, are a way to handle uncertain knowledge, and the two topics are related. However, considering this relation was beyond the scope of this project.

Specification Levels. The 754-2008 standard describes itself as layered into four Specification Levels. To manage complexity, the present standard uses a corresponding structure. It deals mainly with Level 1, of mathematical *interval theory*, and Level 2, the finite set of *interval datums* in terms of which finite-precision interval computation is defined. It has some concern with Level 3, of *representations* of intervals as data structures; and none with Level 4, of *bit strings* and memory.

There is another important player: the programming language. It was a recognized omission of IEEE-754-1985 that it specified individual operations but not how they should be used in expressions. Optimizing compilers have, since well before that standard, used clever transformations so that it is impossible to know the precisions used and the roundings performed while evaluating an expression, or whether the compiler has even “optimized away” $(1.0 + x) - 1.0$ to become simply x . IEEE-754-2008 specifies this by placing requirements on how operations should be used in expressions, though as of this writing, few programming languages have adopted that.

The lack of any restrictions is also a problem for intervals. Thus the standard makes requirements and recommendations on language implementations, thereby defining the notion of a standard-conforming implementation of intervals within a language.

The language does not constitute a fifth level in some linear sequence; from the user’s viewpoint most current languages sit above datum level 2, alongside theory level 1, as a practical means to implement interval algorithms by manipulating Level 2 entities (though most languages have influence on Levels 3 and 4 also). This standard extends them to provide an instantiation of level 2 entities.

The Fundamental Theorem. Moore’s [5] Fundamental Theorem of Interval Arithmetic (FTIA) is central to interval computation. Roughly, it says as follows. Let f be an *explicit arithmetic expression*—that is, it is built from finitely many elementary functions (arithmetic operations) such as $+$, $-$, \times , \div , \sin , \exp , \dots , with no non-arithmetic operations such as intersection, so that it defines a real function $f(x_1, \dots, x_n)$. Then evaluating f “in interval mode” over any interval inputs $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ is guaranteed to give an enclosure of the range of f over those inputs.

A version of the FTIA holds in all variants of interval theory, but with varying hypotheses and conclusions. In the context of this standard, an expression should be evaluated entirely in one flavor, and inferences made strictly from that flavor’s FTIA; otherwise, a user may believe an FTIA holds in a case where it does not, with possibly serious effects in applications. As stated, the FTIA is about the mathematical level. Moore’s achievements were to see that “outward rounding” makes the FTIA hold also in finite precision, and to follow through the consequences. An advantage of the level structure used by the standard is that the mapping between levels 1 and 2 defines a framework where it is easily proved that

The finite-precision FTIA holds in any conforming implementation.

Generally it can only be determined *a posteriori* whether the conditions for any version of the FTIA hold; this is an important application of the standard’s *decoration system*.

For each flavor in the standard, its subdocument must state precisely the form of the FTIA it obeys, both at the mathematical level 1 and at the finite-precision level 2.

Operations.

There are several interpretations of *evaluation outside an operation's domain* and *operations as relations rather than functions*. This includes classical alternative meanings of division by an interval containing zero, or square root of an interval containing negative values. To illustrate the different interpretations, consider $\mathbf{y} = \sqrt{\mathbf{x}}$ where $\mathbf{x} = [-1, 4]$.

- (1) In *optimization*, when computing lower bounds on the objective function, it is generally appropriate to return the result $\mathbf{y} = [0, 2]$, and ignore the fact that $\sqrt{\cdot}$ has been applied to negative elements of \mathbf{x} .
- (2) In applications where one must check the hypotheses of a *fixed point theorem* are satisfied (such as solving differential equations):
 - (a) one may need to be sure that the function is defined and continuous on the input and, hence, report an illegal argument when, as in the above case, this fails; or
 - (b) one may need the result $\mathbf{y} = [0, 2]$, but must flag the fact that $\sqrt{\cdot}$ has been evaluated at points where it is undefined or not continuous.
- (3) In *constraint propagation*, the equation is often to be interpreted as: find an interval enclosing all y such that $y^2 = x$ for some $x \in [-1, 4]$. In this case the answer is $[-2, 2]$.

The standard provides means to meet these diverse needs, while aiming to preserve clarity and efficiency. A language might achieve this by binding one of the above three interpretations—usually some variant of (2)—to its built-in operations, and providing the others as library procedures.

In the context of flavors, a key idea is that of *common operation instances*: those elementary interval calculations that at the mathematical level are required to give the same result in all flavors. For example $[1, 2]/[3, 4] = [1/4, 2/3]$ is common, while division by an interval containing zero is not common.

Decorations.

Many interval algorithms are only valid if certain mathematical conditions are satisfied: for instance one may need to know that a function, defined by an expression, is everywhere continuous on a box in \mathbb{R}^n defined by n input intervals $\mathbf{x}_1, \dots, \mathbf{x}_n$. The IEEE 754 model of global flags to record events such as division by zero was considered inadequate in an era of massively parallel processing. In this standard, such events are recorded locally by *decorations*.

A *decorated interval* is an ordinary interval tagged with a few bits that encode the decoration, and record while evaluating an expression, e.g., “each elementary function was defined and continuous on its inputs”—which implies the same for the function defined by the whole expression. This makes possible a rigorous check of properties such as listed in item (2) of §. A small number of decorations is provided, designed for efficient propagation of such property information.

Care was taken to meet different user needs. *Bare* (undecorated) intervals are available for simple use without validity checks. *Decorated* intervals are recommended for serious programming, but suffer the “17-byte problem”: a typical bare interval stored as two doubles takes up 16 bytes, so a decorated one needs at least 17 bytes. With large problems on typical machine architectures this may cause inefficiencies—in data throughput if storing 17-byte data structures, or in storage if one pads the structure to, say, 32 bytes. Hence an optional *compressed* decorated interval scheme is specified, for advanced use. It aims to give the speed of 16-byte objects, at a cost in flexibility but supporting applications such as checking whether a function is defined and continuous on its inputs.

DRAFT 7.3

Contents

Frontmatter	iii
Introduction	iii
Mathematical context	iii
Specification Levels	iv
The Fundamental Theorem	iv
Operations	v
Decorations	v
Chapter 1. General Requirements	1
1. Overview	1
1.1. Scope	1
1.2. Purpose	1
1.3. Inclusions	1
1.4. Exclusions	1
1.5. Word usage	1
1.6. The meaning of conformance	2
1.7. Programming environment considerations	2
1.8. Language considerations	2
2. Normative references	3
3. Conformance Clause	4
3.1. Things to do and WIP	4
3.1.1. Sections and references here	5
3.2. The actual clause starts here	5
3.3. Set-based interval arithmetic	5
3.3.1. Documentation	5
3.4. Kaucher interval arithmetic	6
4. Notation, abbreviations, definitions	7
4.1. Frequently used notation and abbreviations	7
4.2. Definitions	7
5. Structure of the standard in levels	10
6. Expressions and the functions they define	11
7. Flavors	13
7.1. Flavors overview	13
7.2. Definition of common intervals and common evaluations	13
7.3. Loose common evaluations	14
7.4. Relation of common evaluations to flavors	14
7.5. Flavors and the Fundamental Theorem	15
8. Decoration system	15
8.1. Decorations overview	15
8.2. Decoration definition and propagation	16
8.3. Recognizing common evaluation	16
Chapter 2. Set-Based Intervals	19
9. Level 1 description	19
9.1. Non-interval Level 1 entities	19
9.2. Intervals	19
9.3. Hull	20

9.4. Functions	20
9.4.1. Function terminology	20
9.4.2. Point functions	20
9.4.3. Interval-valued functions	20
9.4.4. Constants	21
9.5. Expressions	21
9.6. Required operations	22
9.6.1. Interval literals	22
9.6.2. Interval constants	22
9.6.3. Forward-mode elementary functions	22
9.6.4. Interval case expressions and case function	22
9.6.5. Two-output division	24
9.6.6. Reverse-mode elementary functions	24
9.6.7. Cancellative addition and subtraction	25
9.6.8. Non-arithmetic (set) operations	26
9.6.9. Constructors	26
9.6.10. Numeric functions of intervals	26
9.6.11. Boolean functions of intervals	26
9.7. Recommended operations	27
9.7.1. Forward-mode elementary functions	27
9.7.2. Slope functions	28
9.7.3. Extended interval comparisons	28
10. The decoration system at Level 1	31
10.1. Decorations and decorated intervals overview	31
10.2. Definitions and basic properties	31
10.3. The ill-formed interval	32
10.4. Permitted combinations	32
10.5. Initial decoration	32
10.6. Decorations and arithmetic operations	33
10.7. Decorations and non-arithmetic operations	33
10.8. Boolean functions of decorated intervals	34
10.9. Operations on decorations	34
10.10. User-supplied functions	34
10.11. Notes on the <code>com</code> decoration	35
10.12. Compressed arithmetic with a threshold (optional)	36
10.12.1. Motivation	36
10.12.2. Compressed interval types	36
10.12.3. Operations	37
11. Level 2 description	38
11.1. Level 2 introduction	38
11.2. Naming conventions for operations	38
11.3. 754-conformance	39
11.3.1. 754-conforming mixed-type arithmetic	39
11.4. Tagging, and the meaning of equality at Level 2	39
11.5. Number formats	40
11.6. Bare and decorated interval types	40
11.6.1. Definition	40
11.6.2. Initial decoration of an interval	41
11.6.3. Inf-sup and mid-rad types	41
11.7. Multi-precision interval types	41
11.8. Explicit and implicit types, and Level 2 hull operation	41
11.8.1. Hull in one dimension	41
11.8.2. Hull in several dimensions	42
11.8.3. Type conversion for bare and decorated intervals	42
11.9. Level 2 interval extensions	42
11.10. Accuracy requirements for operations	42

11.10.1.	Measures of accuracy	42
11.10.2.	Documentation requirements	43
11.10.3.	Recommended accuracies	44
11.11.	Required operations on bare intervals	44
11.11.1.	Bare and decorated interval literals	45
11.11.2.	Interval constants	47
11.11.3.	Forward-mode elementary functions	47
11.11.4.	Interval case expressions and case function	47
11.11.5.	Reverse-mode elementary functions	47
11.11.6.	Cancellative addition and subtraction	47
11.11.7.	Set operations	48
11.11.8.	Constructors	48
11.11.9.	Numeric functions of intervals	49
11.11.10.	Boolean functions of intervals	50
11.11.11.	Reduction operations	50
11.12.	Recommended operations	50
11.12.1.	Forward-mode elementary functions	50
11.12.2.	Slope functions	51
11.12.3.	Extended interval comparisons	51
11.13.	Compressed arithmetic at Level 2	51
12.	Input and output (I/O) of intervals	52
12.1.	Overview	52
12.2.	Input	52
12.3.	Output	52
12.4.	Public representation	52
12.5.	Representations	54
12.6.	Rationale for defined hulls, text representation, and reproducibility	54
12.6.1.	The hull	54
12.6.2.	Standardised text representation	54
12.6.3.	Reproducibility	55
13.	Level 3 description	56
13.1.	Level 3 introduction	56
13.2.	Representation	56
13.3.	Operations and representation	56
13.4.	Format conversion	57
13.5.	Interchange formats	57
13.6.	Support levels for interval elementary functions	58
13.7.	Operation tables for basic interval operations	58
13.8.	Care needed with innerPlus and innerMinus	60
13.9.	Implementation of bare object arithmetic	60
13.10.	Complete arithmetic, dot product function	60
14.	Level 4 description	61
Chapter 3.	Kaucher Intervals	63
Annex A.	Details of flavor-independent requirements	65
A.1.	List of required functions	65
A.2.	List of recommended functions	65
Annex B.	Including a new flavor in the standard	67
Annex C.	Reproducibility	69
Annex D.	Set-based flavor: decoration details and examples	71
D.1.	Local decorations of arithmetic operations	71
D.1.1.	Forward-mode elementary functions	71
D.1.2.	Interval case function	71
D.2.	Examples of use of decorations	71

D.3. Implementation of compressed interval arithmetic	74
D.4. Proofs of correctness for compressed interval arithmetic	75
D.5. The fundamental theorem of decorated interval arithmetic	76
Annex E. Further material for set-based standard (informative)	79
E.1. Type conversion in mixed operations	79
E.2. The “Not an Interval” object	79
Bibliography	81

DRAFT 7.3

General Requirements

1. Overview

1.1. Scope. This standard specifies basic interval arithmetic (IA) operations selecting and following one of the commonly used mathematical interval models. This standard supports the IEEE-754-2008 floating point formats of practical use in interval computations. Exception conditions are defined and standard handling of these conditions is specified. Consistency with the model is tempered with practical considerations based on input from representatives of vendors and owners of existing systems.

The standard provides a layer between the hardware and the programming language levels. It does not mandate that any operations be implemented in hardware. It does not define any realization of the basic operations as functions in a programming language.

1.2. Purpose. The aim of the standard is to improve the availability of reliable computing in modern hardware and software environments by defining the basic building blocks needed for performing interval arithmetic. There are presently many systems for interval arithmetic in use; lack of a standard inhibits development, portability; ability to verify correctness of codes.

1.3. Inclusions. This standard specifies

- Types for interval data based on underlying numeric formats, with a special class of type derived from IEEE 754 floating point formats.
- Constructors for intervals from numeric and character sequence data.
- Addition, subtraction, multiplication, division, fused multiply add, square root; other interval-valued operations for intervals.
- Midpoint, radius and other numeric functions of intervals.
- Interval comparison relations.
- Required elementary functions.
- Conversions between different interval types.
- Conversions between interval types and external representations as text strings.
- Interval-related exceptions and their handling.

1.4. Exclusions. This standard does not specify

- Which numeric formats supported by the underlying system shall have an associated interval type.
- Details of how an implementation represents intervals at the level of programming language data types, or bit patterns.

1.5. Word usage.

In this standard three words are used to differentiate between different levels of requirements and optionality, as follows:

- **may** indicates a course of action permissible within the limits of the standard with no implied preference (“may” means “is permitted to”);
- **shall** indicates mandatory requirements strictly to be followed in order to conform to the standard and from which no deviation is permitted (“shall” means “is required to”);
- **should** indicates that among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others; or that a certain course of action is preferred but not necessarily required; or that (in the negative form) a certain course of action is deprecated but not prohibited (“should” means “is recommended to”).

Further:

- **optional** indicates features that may be omitted, but if provided shall be provided exactly as specified.
- **might** indicates the possibility of a situation that could occur, with no implication of the likelihood of that situation (“might” means “could possibly”);
- **see** followed by a number is a cross-reference to the clause or subclause of this standard identified by that number;
- **comprise** indicates members of a set are exactly those objects having some property, e.g. “the set of mathematical intervals comprises the closed, connected subsets of \mathbb{R} ”; an unqualified **consist of** merely asserts all members of a set have some property, e.g. “a binary floating point format consists of numbers with a terminating binary representation”. “Comprises” means “consists exactly of”.
- **Note** and **Example** introduce text that is informative (is not a requirement of this standard).

1.6. The meaning of conformance. §3 lists the requirements on a conforming implementation in summary form, with references to where these are stated in detail.

1.7. Programming environment considerations.

This standard does not define all aspects of a conforming programming environment. Such behavior should be defined by a programming language definition supporting this standard, if available; otherwise by a particular implementation. Some programming language specifications might permit some behaviors to be defined by the implementation.

Language-defined behavior should be defined by a programming language standard supporting this standard. Standards for languages intended to reproduce results exactly on all platforms are expected to specify behavior more tightly than do standards for languages intended to maximize performance on every platform.

Because this standard requires facilities that are not currently available in common programming languages, the standards for such languages might not be able to fully conform to this standard if they are no longer being revised. If the language can be extended by a function library or class or package to provide a conforming environment, then that extension should define all the language-defined behaviors that would normally be defined by a language standard.

Implementation-defined behavior is defined by a specific implementation of a specific programming environment conforming to this standard. Implementations define behaviors not specified by this standard nor by any relevant programming language standard or programming language extension. Conformance to this standard is a property of a specific implementation of a specific programming environment, rather than of a language specification. However a language standard could also be said to conform to this standard if it were constructed so that every conforming implementation of that language also conformed automatically to this standard.

1.8. Language considerations. All relevant languages are based on the concepts of data and transformations. In Von Neumann languages, data are held in variables, which are transformed by assignment. In functional languages, input data are supplied as arguments; the transformed form is returned as results. Dataflow languages vary considerably, but use some form of the data and transformation approach.

Similarly, all relevant languages are based on the concept of mapping the pseudo-mathematical notation that is the program code to approximate real arithmetic, nowadays almost exclusively using some form of floating-point. The unit of mapping and transformation can be individual operations and built-in functions, expressions, statements, complete procedures, or other. This standard is applicable to all of these.

The least requirement on a conforming language standard, compiler or interpreter is that it shall:

- (1) define bindings so that the programmer can specify level 2 data (in the sense of the levels defined in Clause 5) as described in this standard;
- (2) define bindings so that the programmer can specify the operations on such data as described in this standard;
- (3) define any properties of such data and operations that this standard requires to be defined;

- (4) honor the rules of interval transformations on such data and operations as described in this standard; such units of transformation that the language standard, compiler or interpreter uses.

Specifically, if the data before and after the unit of transformation are regarded as sets of mathematical intervals, the transformed form of all combinations of the elements (the real values) represented by the prior set shall be a member of the posterior set.

If a conforming language standard supports reproducible interval arithmetic it shall also:

- (5) Use the data bindings as specified in point (1) above for reproducible operations;
- (6) Define bindings to the reproducible operations as described in this standard;
- (7) Define any modes and constraints that the programmer needs to specify or obey in order to obtain reproducible results.

If a conforming language standard supports both non-reproducible and reproducible interval arithmetic it shall also:

- (8) Permit a reproducible transformation unit to be used as a component in a non-reproducible program, possibly via a suitable wrapping interface.

2. Normative references

1. IEEE Std 754, *IEEE Standard for Floating-Point Arithmetic*. References in this document are to the 2008 revision.
2. Ulrich Kulisch and Van Snyder, *The exact dot product as basic tool for long interval arithmetic*. Position paper, P1788 Working Group, version 11, July 2009.

3. Conformance Clause

3.1. Things to do and WIP. This subsection is just for keeping track of what needs to be done for this section.

To conform to OASIS Conformance Requirements for Specifications version 0.5, 1 March 2002 we have to do the following things in essence:

- ✓ provide a conformance clause—we are doing this right now
- ✓ use the proper conformance keywords—as long as we are adhering to §1.5 this is covered.
- address all topics in section 8 of the OASIS requirements
 - ✓ (8.1) Specify what conforms—an implementation of an interval arithmetic programming environment
 - ✓ (8.1.1) Modularity of the conformance—We don't have different components that individually have to conform.
 - (8.1.2) Specifying conformance claims—We don't have different levels of conformance
 - (8.2) Profiles and levels—the flavors look like profiles
 - (8.3) Extensions—probably none
 - (8.4) Discretionary items—implementation defined things
 - (8.5) Deprecated things—not yet
 - (8.6) Internationalization—english is normative, we have some characteristics for languages which might be interesting here
- examine and if appropriate address all things in section 9
 - (9.1) Implementation Conformance Statement—this might be relevant to register implementation defined points
 - (9.2) Test assertions—probable connection to the test suite
 - (9.3) Testing methodology/program

Things to consider

- a language definition may be supporting the standard, an implementation may be conforming - may we have different objects that may be the object of the standard after all, see §1.7
- an implementation might also be a library, class or package
- a reproducibility mode might look like a level? at least it is a kind of quality of the implementation
- perhaps we have to clarify (here?) the difference between conformance levels and the levels structure—if necessary add something to the definitions.
- flavors
 - shall provide at least one
 - may provide additional
 - flavors following common evaluation on common intervals
- the flavor text sounds like a flavor might be conforming to the standard? see §8.2
- flavor documentation includes documentation of the decorations
- compressed arithmetic certainly sounds like a profile—you may provide it, but if you do and want it to be conforming, you have to do it this way.
- Set-based intervals
 - Maybe reference the interval type – §9.2
 - Maybe give the type of operations: operators, named functions – §9.4
 - References should probably mostly come from level 2, not Clause 9, check if that works out
 - reference “supported type”
 - required operations – §9.6 – recommended – §9.7
 - §10.2 – §10.4 mostly background without requirements
 - I seem to remember a requirement to document if `intersectionDec` or `convexHullDec` are used as a default but cannot find it in the draft anymore.
- Set-based level 2 seems to contradict itself on the requirements of mixed-type arithmetic: §11.3 seems to require mixed-type operations for all level 2 operations—“without explicit mention”—, while §11.11 lists only a few explicitly.

- Set-based level 1 and level 2 seem to disagree on the list of required operations for an implementation: §9.6 requires all operations in the subclause, §11.11 explicitly requires all but the interval constants in the head of the subclause, while the constants seem to be required in their own subsubclause. As a side note: it might be nice to use the same headings for all subsubsections there.

3.1.1. *Sections and references here.*

- Clause 6 No items here
- Clause 8 referenced with `newDec` and `com`
- §9.1 – §9.5 probably only informational background
- §11.9 not referenced as only a definition

3.2. The actual clause starts here. An implementation of an interval arithmetic program-environment that conforms to this standard shall satisfy the following requirements

- flavor-independent requirements

A conforming implementation shall also provide at least one of the following profiles¹, defined in Clause 7, called included flavors of interval arithmetic in the context of this standard.

In addition an implementation may also provide additional flavors that

- shall provide the required operations in Clause A.1;
- should provide the recommended operations in Clause A.2;
- shall provide common evaluations on common intervals as specified in §7.4.

⚠ The following paragraph will probably turn into a requirement on the library provided by an implementation, i.e. it has to provide a plain and a decorated library. For each provided flavor an implementation shall provide a decoration system supporting the verification of the assumptions of existence, uniqueness, or nonexistence theorems as specified in Clause 8. This includes as specified in §8.2

- a `newDec` function and
- a number of decorations; if the implementation provides more than one flavor, the decorations shall include the `com` decoration.

3.3. Set-based interval arithmetic. An implementation of the set-based flavor shall satisfy the following requirements

- provide the decorations specified in §10.2
- support at least one bare interval type, see §11.6
- each supported bare interval type shall be based on a number format with an associated rounding function, see §11.5, and have an interval hull operation that maps arbitrary sets of real numbers to a corresponding interval, see §11.8
- if multi-precision interval types are supported, they shall be defined as a parameterized sequence of interval types, see §11.7
- evaluate expressions in a way that satisfies the equality principle, see §11.4,
- provide implementations of the required operations in §11.11; recommended accuracies for these operations can be found in §11.10

If an implementation wants to claim 754-conformance, see §11.3, in addition to the above it shall satisfy the following requirements

- each supported interval type has to be 754-conforming and meet the requirements for mixed-type arithmetic, see §11.3, §11.11 ²

3.3.1. *Documentation.* ⚠ This may move into a separate kind of questionnaire document that an implementation shall complete to claim conformance. Therefore it is now just a collection of items that have to be documented

A conforming implementation shall document the following things

- the interval hull operation of each supported implicit interval type, e.g. via an algorithm, see §11.8
- the achieved accuracy (i.e., tightest, accurate, or valid) for each of its interval operations

¹maybe drop this term?

²Check if we need “754-conforming part” here.

, see §11.10

3.4. Kaucher interval arithmetic. An implementation of the Kaucher flavor shall satisfy the following requirements

- Kaucher requirements

An implementation may claim its conformance to this standard in the following way

Name of implementation and version have been tested for conformance to the P1788 Standard for Interval Arithmetic, version D7.0 using the *we might put a reference to a probable test suite here* on *YYYY-MM-DD* and no nonconformities were found.

DRAFT 7.3

4. Notation, abbreviations, definitions

4.1. Frequently used notation and abbreviations.

754	IEEE-Std-754-2008 “IEEE Standard for Floating-Point Arithmetic”.
\mathbb{R}	the set of real numbers.
$\overline{\mathbb{R}}$	the set of extended real numbers, $\mathbb{R} \cup \{-\infty, +\infty\}$.
\mathbb{IR}	the set of closed real intervals, including unbounded intervals and the empty set.
$\mathbb{F}, \mathbb{G}, \dots$	generic notation for (the set of numbers, including $\pm\infty$) representable in some floating point format.
$\mathbb{IF}, \mathbb{IG}, \dots$	the members of \mathbb{IR} whose lower and upper bounds are in $\mathbb{F}, \mathbb{G}, \dots$
Empty	the empty set.
Entire	the whole real line.
NaI	Not an Interval.
NaN	Not a Number.
qNaN	quiet NaN.
sNaN	signaling NaN.
x, y, \dots [resp. f, g, \dots]	typeface/notation for a numeric value [resp. numeric function].
$\mathbf{x}, \mathbf{y}, \dots$ [resp. $\mathbf{f}, \mathbf{g}, \dots$]	typeface/notation for an interval value [resp. interval function].
f, g, \dots	typeface/notation for an expression, producing a function by evaluation.
$\text{Dom}(f)$	the domain of a point-function f .
$\text{Rge}(f \mid \mathbf{s})$	the range of a point-function f over a set \mathbf{s} ; the same as the image of \mathbf{s} under f .

[Note. Little used in this document, but used in classical interval analysis, are \mathbb{IR} , the set of bounded, nonempty closed real intervals; and \mathbb{IF} , the intervals of \mathbb{IR} whose bounds are in \mathbb{F} . The symbols \mathbb{I} and \mathbb{I} act as operators on subsets S of $\overline{\mathbb{R}}$, namely $\mathbb{I}S$ or $\mathbb{I}(S) =$ “the empty set, plus all intervals whose endpoints are in S ” and $\mathbb{I}S =$ “all nonempty intervals whose endpoints are finite and in S ”.]

4.2. Definitions.

⚠ Many of the definitions especially those for Levels 2 onward, need revision. Some should be removed. Some new ones are needed.

4.2.1. **754-conforming system.** (754 system for short) A programming environment, made up of hardware or software or both, that provides floating point arithmetic conforming to IEEE754-2008.

A **754-conforming implementation** is an implementation (see Defn 4.2.17) that is built on a 754 system and satisfies the extra requirements for 754-conformance stated at various places in the standard.

4.2.2. **accuracy mode.** A way to describe the quality of an interval version of a function. See §11.10.

4.2.3. **arithmetic operation.** A function provided by an implementation (see Defn 4.2.17). It comes in three forms: the **point** operation, which is a mathematical real function of real variables such as $\log(x)$; one or more **(bare) interval extensions** of the point operation, each of which corresponds to the finite precision interval type of its result; and one or more **decorated interval extensions**, each being the (unique) decorated version of a bare interval extension.

Together with the interval non-arithmetic operations (§9.4.1), these form the implementation’s **library**, which splits into the **point library** (a conceptual entity, being a set of mathematical functions), the **bare interval library** and the **decorated interval library**, corresponding to the above categories. The latter two may be further qualified by a result interval type, e.g., “binary64 inf-sup decorated interval library”.

The programming environment’s floating point approximations to mathematical point functions constitute the *floating point library*. The standard makes no requirements on these.

A **basic arithmetic operation** is one of the six functions $+$, $-$, \times , \div , fused multiply-add **fma** and square root **sqrt**.

Constants such as 3.456 and π are regarded as arithmetic operations whose number of arguments is zero. Details in §9.4.4.

4.2.4. **available.** To be defined, see Motion 16 ...³

4.2.5. **basic 754 format.** One of the five 754 floating point formats binary32, binary64, binary128, decimal64, decimal128.

4.2.6. **box.** See Defn 4.2.19.

4.2.7. **concrete interval type.** An interval type associated with a specific number representation. Usually the inf-sup type associated with a specific floating point format. Details in §11.5, 11.6, ??, ??.

4.2.8. **decoration.**⁴

4.2.9. **decorated interval.** A pair (interval, decoration).

4.2.10. **decorated interval library.** See Defn 4.2.3.

4.2.11. **decorated interval version.** See Defn 4.2.3.

4.2.12. **domain.** For a function with arguments taken from some set, the **domain** comprises those points in the set at which the function has a value. The domain of an arithmetic operation is part of its definition. E.g., the (point) arithmetic operation of division x/y , in this standard, has arguments (x, y) in \mathbb{R}^2 , and its domain is the set $\{(x, y) \in \mathbb{R}^2 \mid y \neq 0\}$. See also Defn 4.2.30.

4.2.13. **elementary function.** Synonymous with arithmetic operation.

4.2.14. **expression.** A symbolic object f that is either a symbolic variable or, recursively, of the form $\varphi(g_1, \dots, g_k)$, where φ is the name of a k -argument arithmetic or non-arithmetic operation and the g_i are expressions. It is an **arithmetic expression** if all its operations are arithmetic operations. Writing $f(z_1, \dots, z_n)$ makes f a **bound expression**, giving it an **argument list** comprising the variables z_i in that order, which must include all those that occur in f .

Details in Clause 6.

4.2.15. **fma.** Fused multiply-add operation, that computes $x \times y + z$. One of the basic arithmetic operations.

4.2.16. **hull.** (Full name: **interval hull.**) When not qualified by the name of a finite-precision interval type, the hull of a subset s of \mathbb{R} is the tightest (q.v.) interval containing s .

When \mathbb{T} is an explicit type, the \mathbb{T} -hull of s is the unique tightest \mathbb{T} -interval containing s .

When \mathbb{T} is an implicit type, the \mathbb{T} -hull of s is a minimal \mathbb{T} -interval containing s as specified in the definition of the type.

4.2.17. **implementation.** When used without qualification, means a realization of an interval arithmetic conforming to the specification of this standard.

4.2.18. **inf-sup.** Describes a representation of an interval based on its lower and upper bounds.

4.2.19. **interval.** A closed connected subset of \mathbb{R} ; may be empty, bounded or unbounded. May be called a 1-dimensional interval, see next paragraph. The set of all intervals is denoted $\overline{\mathbb{IR}}$.

A **box** or **interval vector** is an n -dimensional interval, i.e. a tuple $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ where the \mathbf{x}_i are intervals. Often identified with the cartesian product $\mathbf{x}_1 \times \dots \times \mathbf{x}_n \subseteq \mathbb{R}^n$, it is empty if any of the \mathbf{x}_i is empty. Details in §9.2.

4.2.20. **interval datum.** A member of the finite set of intervals representable in a specific interval type. Details in §11.6, ??, ??.

4.2.21. **interval extension.** An interval extension of a point function f is a function \mathbf{f} from intervals to intervals such that $f(x)$ belongs to $\mathbf{f}(\mathbf{x})$ whenever x belongs to \mathbf{x} and $f(x)$ is defined. Details in §9.4.3. It is the **natural** (or tightest) interval extension if $\mathbf{f}(\mathbf{x})$ is the interval hull of the range of f over \mathbf{x} , for all \mathbf{x} .

A **decorated interval extension** ...⁵

4.2.22. **i-type, interval type, radix of i-type.** The interval type associated with a number format \mathbb{F} means the set of interval datums (\mathbb{F} -intervals for short) associated with \mathbb{F} in a particular representation. When used without qualification, inf-sup representation is assumed — by contrast with, say, mid-rad representation. The radix of an i-type means the radix of its n-format. Details in §11.6, ??, ??.

4.2.23. **interval function, interval mapping.** A function from intervals to intervals is called an interval function if it is an interval extension of a point function, and an interval mapping otherwise. Details in §9.4.3.

³CK 2011-11-08 CK

⁴CK To be defined...

⁵CK Short note for completion should be added

4.2.24. **interval library.** See Defn 4.2.3.

4.2.25. **mathematical interval of constructor.** The arguments of an interval constructor, if valid, define a mathematical interval \mathbf{x} . The actual interval returned by the constructor is the tightest interval of the destination type that contains \mathbf{x} . Details in §11.11.8, ??.

4.2.26. **mid-rad.** Describes a representation of an interval based on its midpoint and radius.

4.2.27. **mixed-format.** Describes a scalar interval arithmetic operation where the formats of the operand interval(s) and the destination interval may not all be the same. X
[Note. Mixed-format does not mean that the lower and upper bounds of an individual interval object can have different number formats.]

4.2.28. **MRMF.** Mixed-radix, mixed-format. Describes a level of arithmetic support where operand(s) and destination of a scalar interval operation may be an arbitrary mix of supported formats, not necessarily of the same radix. Details in §13.6.

4.2.29. **multiple-format.** Multiple-format interval arithmetic means support for arithmetic in more than one interval format, and for conversions between these formats. X
[Note. SRSF, SRMF and MRMF interval support are all multiple-format (if more than one interval format is supported). SRMF and MRMF, but not SRSF, are mixed-format.]

4.2.30. **natural domain.** For an arithmetic expression $f(z_1, \dots, z_n)$, the natural domain is the set of $x = (x_1, \dots, x_n) \in \mathbb{R}^n$ where the expression defines a value for the associated point function $f(x)$. See Clause 6.

4.2.31. **no value vs. undefined.** Some functions do not have a defined value at the mathematical model of Level 1 (see Clause 5). This translates to a defined value at the interval datum Level 2 given at the corresponding places in the standard. Therefore the term “no value” is not to be confused with an “undefined” value, which has the common meaning of a value not determined by the standard and thus being free for the implementation to decide.

4.2.32. **non-arithmetic operation.** An operation on intervals that is not an interval extension of a point operation; includes interval intersection and union.

4.2.33. **number.** Any member of the set $\mathbb{R} \cup \{-\infty, +\infty\}$ of extended reals: a **finite number** if it belongs to \mathbb{R} , else an **infinite number**. See §9.1.

4.2.34. **n-format, number format.** A subset of the extended reals $\overline{\mathbb{R}}$, containing $\pm\infty$. In practice, it is usually the finite set of numbers representable in a specific floating point format and is then the n-format **associated** with that format. Details in §11.5.

4.2.35. **point function, point operation.** A mathematical function of real variables: that is, a map f from its domain, which is a subset of \mathbb{R}^n , to \mathbb{R}^m , where $n \geq 0, m > 0$. It is **scalar** if $m = 1$. Any arithmetic expression $f(z_1, \dots, z_n)$ defines a (usually scalar) point function, whose domain is the natural domain of f . Superseded?

4.2.36. **point library.** See Defn 4.2.3.

4.2.37. **range.** The range, $\text{Rge}(f | \mathbf{s})$, of a point function f over a subset \mathbf{s} of \mathbb{R}^n is the set of all values that f assumes at those points of \mathbf{s} where it is defined, i.e. $\{f(x) \mid x \in \mathbf{s} \text{ and } x \in \text{Dom } f\}$.

4.2.38. **SRMF.** Single-radix, mixed-format. Describes a level of arithmetic support where operand(s) and destination of an interval operation may be an arbitrary mix of supported formats, but must be of the same radix. Details in §13.6.

4.2.39. **SRSF.** Single-radix, single-format. Describes a level of arithmetic support where operand(s) and destination of an interval operation must all be of the same (supported) format. Details in §13.6.

4.2.40. **string.** A text string, or just string, is a finite sequence of characters belonging to some alphabet. See §9.1. X X

4.2.41. **support.**

- An implementation **supports** an abstract interval type $\overline{\mathbb{IF}}$ if it provides a concrete interval type that represents $\overline{\mathbb{IF}}$. Details in §11.5, 11.6, ??, ??.
- A function f is **interval-supported** in an interval type $\overline{\mathbb{IF}}$ if there is an interval version of f whose destination type is $\overline{\mathbb{IF}}$. For levels of interval support (MRMF, SRMF and SRSF) see Defn 4.2.28, 4.2.38, 4.2.39. X

4.2.42. **tightest.** Smallest in the partial order of set containment. The tightest set (unique, if it exists) with a given property is contained in every other set with that property.

4.2.43. **version.** For brevity, bare or decorated interval extensions, or floating point approximations, of a library point operation may be called versions of that operation.

Relationships between specification levels for interval arithmetic for a given flavor \mathfrak{F} and a given finite-precision interval type \mathbb{T} .		
Level 1	Number system used by flavor \mathfrak{F} . Set of allowed intervals in \mathfrak{F} . Principles of how $+$, $-$, \times , \div and other arithmetic operations are extended to intervals.	Mathematical Model level.
	\downarrow \mathbb{T} -interval hull total, many-to-one ^a	<i>identity map</i> \uparrow total, one-to-one ^b
Level 2	A finite subset \mathbb{T} of the \mathfrak{F} -intervals—the \mathbb{T} -interval datums—and operations on them.	Interval datum level.
		<i>“represents”</i> \uparrow partial, many-to-one, onto ^c
Level 3	Representations of \mathbb{T} -intervals, e.g. by two floating point numbers.	Representation level.
		<i>“encodes”</i> \uparrow partial, many-to-one, onto ^d
Level 4	Encodings 0111000...	Bit string level.

TABLE 1. Specification levels for interval arithmetic

5. Structure of the standard in levels

The standard is structured into four levels, summarized in Table 1, that match the levels defined in the 754 standard, see 754 Table 3.1.

Level 1, in Clause 9, defines the mathematical theory underlying the standard. The entities at this level are mathematical intervals and operations on them. Conforming implementations shall implement this theory.

In addition to an ordinary (bare) interval, this level defines a *decorated* interval, comprising a bare interval and a *decoration*. Decorations implement this standard’s exception handling mechanism.

Level 2, in Clause 11, is the central part of the standard. Here the mathematical theory is approximated by an implementation-defined finite set of entities and operations. A level 2 entity is called a *datum* (plural “datums” in this standard, since “data” is often misleading).

An interval datum is a mathematical interval tagged by a unique *type* \mathbb{T} . The type abstracts a *representation scheme*—a particular way of representing intervals (e.g., by storing their lower and upper bounds as IEEE `binary64` numbers). Level 2 arithmetic normally acts on intervals of a given type to produce an interval of the same type (but interval operations that act on intervals of types other than the result type are possible).

Level 3, in Clause 13, is concerned with the representation of interval datums—usually but not necessarily in terms of floating point values. A level 3 entity is an *interval object*. Representations of decorations, hence of decorated intervals, are also defined at this level.

The Level 3 requirements in this standard are few, and concern mappings from internal representations to external ones, such as interchange formats and I/O.

A level 4 entity is a *bit string*. This standard makes no Level 4 requirements.

The arrows in Table 1 denote mappings between levels. The phrases in italics name these mappings. Each phrase “total, many-to-one”, etc., labeled with a letter ^a to ^d, is descriptive of the mapping and is equivalent to the corresponding labeled fact below.

- a. Each mathematical interval (indeed each subset of \mathbb{R}) has a unique interval datum as its *\mathbb{T} -hull*—a minimal enclosing interval of the type \mathbb{T} .
- b. Each interval datum is a mathematical interval, if one ignores its type-tag.
- c. Not every interval object necessarily represents an interval datum, but when it does, that datum is unique. Each interval datum has at least one representation, and may have more than one.
- d. Not every interval encoding necessarily encodes an interval object, but when it does, that object is unique. Each interval object has at least one encoding and may have more than one.

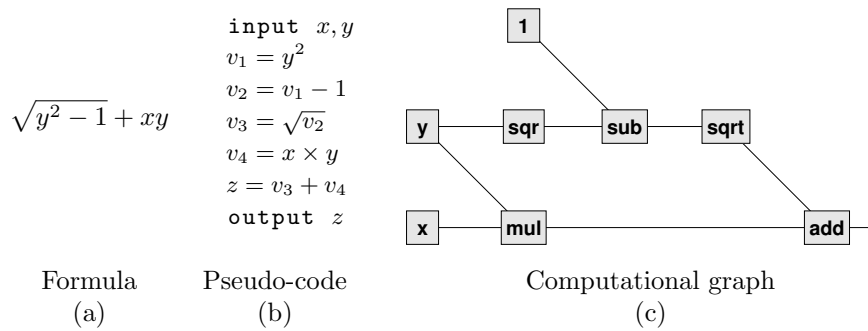


FIGURE 1. Essentially equivalent notations for an expression. The pseudo-code (b) breaks the formula into single library operations, and uses static single-assignment (SSA) form. The graph form (c) uses the names in §9.6.3, for library operations.

6. Expressions and the functions they define

An expression is some symbolic form used to define a function. Expressions are central to interval computation, because the Fundamental Theorem of Interval Arithmetic (FTIA) is about interpreting an expression in different ways: first, as defining a mathematical real point function f ; second, as defining interval mappings that give proven enclosures for the range of f over an input box \mathbf{x} . The decoration system, in suitable situations, lets one strengthen this by concluding f is everywhere defined, or everywhere continuous, on \mathbf{x} —making it possible, for example, to automatically verify the assumptions of the Brouwer Fixed Point Theorem.

The standard specifies behavior at the individual operation level, that makes such conclusions possible. It does not define a meaning of “expression” since this would raise language-dependent semantic issues outside its scope. However it is useful to list some things that an expression is and is not, to give a context for the specifications of the standard:

- (i) An expression f defines a relation between certain *inputs*, often represented by mathematical or programming variables, and certain *outputs*, via the application of named *operations* that come from a *library*. f may be represented in various ways, e.g., by a normal algebraic formula or by a segment of program code or pseudo-code or as a *computational graph*. This is illustrated for (an expression defining) the function $f(x, y) = \sqrt{y^2 - 1} + xy$ in Figure 1 (a,b,c) respectively.

Code may define a *vector* expression, with several outputs; however, all the individual library operations of the standard are *scalar*, with possibly several inputs but a single output.

- (ii) For the FTIA to apply, f must be an *arithmetic expression*, each of whose library operations is primarily a *point-operation* with real-number input(s) and output. The operations must be *generic*, each one having one or more *interval version(s)* with interval input and output, and corresponding *decorated interval version(s)*. These produce *point evaluation*, *interval evaluation* and *decorated interval evaluation* of the expression, also termed evaluation in *point mode*, *interval mode* or *decorated interval mode*.

An implementation’s library by definition comprises all its finite-precision (Level 2) versions of required or recommended operations that it provides for any of its supported interval types, as specified in §9.6, §9.7 and in Clause 11. Different Level 2 interval evaluations of f come from using library operations with different combinations of input and output types (precisions), as the implementation may provide.

The set operations `intersection` and `convexHull` are not point-operations and cannot appear in an arithmetic expression. However they are useful for efficiently *implementing* interval versions of functions defined piecewise (see Example (ii) in §10.10).

An arithmetic operation or expression may also denote a floating point function; these are not specified by this standard.

- (iii) The expression must be *explicit*. Interval analysis has ways to find range-enclosures for a function such as “ $y = g(x)$ defined implicitly by solving $\sqrt{y^2 - 1} + xy = 1$ for y ”; but the FTIA does not apply directly to such functions.

Each library point-operation has a defined domain, the set of inputs where it can be evaluated. This leads to the idea of *natural domain* $\text{Dom}(f)$ of the point function $f(x) = f(x_1, \dots, x_n)$ defined by an expression: the set of points x where f is *defined* in the sense that the whole expression can be successfully evaluated. [Example. From the domains of $/$ and $\sqrt{\cdot}$, one finds the natural domain of $\sqrt{1 + 1/x}$ is the union of the two intervals $(-\infty, -1]$ and $(0, +\infty)$.]

In the set-based flavor—see §7.5 for other flavors—Moore’s basic theorem for a scalar function is as follows, with the above notation.

Theorem 6.1 (Fundamental Theorem of Interval Arithmetic). *Let $\mathbf{y} = f(\mathbf{x})$ be the result of interval-evaluation of f over a box $\mathbf{x} = (x_1, \dots, x_n)$ using any interval versions of its component library functions. Then*

- (i) (“Weak” form of FTIA.) *In all cases, \mathbf{y} contains the range of f over \mathbf{x} , that is, the set of $f(x)$ at points of \mathbf{x} where it is defined:*

$$\mathbf{y} \supseteq \text{Rge}(f | \mathbf{x}) = \{ f(x) \mid x \in \mathbf{x} \cap \text{Dom}(f) \}. \quad (1)$$

- (ii) (“Defined” form of FTIA.) *If also each library operation in f is everywhere defined on its inputs, while evaluating \mathbf{y} , then f is everywhere defined on \mathbf{x} , that is $\text{Dom}(f) \supseteq \mathbf{x}$.*
- (iii) (“Continuous” form of FTIA.) *If in addition to (ii), each library operation in f is everywhere continuous on its inputs, while evaluating \mathbf{y} , then f is everywhere continuous on \mathbf{x} .*

It is important to note that the theorem holds in finite precision (Level 2), not just at Level 1. Each of its forms is useful; e.g., the “continuous” FTIA is needed in applications where the assumptions for applying Brouwer’s Fixed Point Theorem must be verified. Building on features in older interval systems, the standard’s decoration system, Clause 8, gives basic tools for checking the conditions for the “defined” and “continuous” forms, during evaluation of a function.

The following points are also relevant; they concern language semantics of expressions.

- (iv) The computational graph of the expression is a directed acyclic graph. When program code is involved, the FTIA applies to the expression f evaluated in a particular execution of that code. If code contains conditional statements or loops, f and its graph may change from one execution to another. Often such code is designed to define a function piecewise (e.g., over several intervals, see example in §10.10). The user is responsible for checking that the FTIA applies as intended in such cases; the standard provides no way to check automatically that a property such as continuity holds globally.
- (v) There is a mismatch between expressions and set theory when doing interval-evaluation, illustrated by the following. Define a 3-argument function by

$$f(x, y, z) = x + z, \quad (2)$$

where argument y does not occur in the expression. Interval-evaluation $\mathbf{w} = f(\mathbf{x}, \mathbf{y}, \mathbf{z})$ returns $\mathbf{x} + \mathbf{z}$, ignoring \mathbf{y} . When \mathbf{y} is empty, e.g.

$$\mathbf{w} = f([1, 2], \emptyset, [3, 4]) = [4, 6], \quad (3)$$

this is at odds with set-theory, which says \mathbf{w} encloses the range of the point-function over the product box $[1, 2] \times \emptyset \times [3, 4]$. Since this is empty, the range is empty, suggesting the computed \mathbf{w} should be \emptyset instead of $[4, 6]$.

It is language-defined which of these interpretations is adopted, but the FTIA theory in Annex D is framed so that the theoretically correct result coincides with what straightforward interval evaluation computes—here $[4, 6]$. This is done by defining the function specified by an expression to depend only on those inputs that in the evaluation’s computational graph lie on a path to the output. In (2), the function thus has two arguments x, z instead of three.

This is conveniently handled by using *keyword association* of dummy to actual arguments—for expressions, not for individual library operations for which normal positional arguments are used—so that irrelevant variables can be skipped. E.g. the argument association in (3) may be written $f(x = [1, 2], z = [3, 4])$. See Annex D for details.

7. Flavors

7.1. Flavors overview. The standard permits different interval *flavors*, which embody different foundational (Level 1) approaches to intervals. An implementation shall provide at least one flavor. For brevity, phrases such as “A flavor shall provide, or document, a feature” mean that the implementation of that flavor shall provide the feature, or its documentation describe it.

Flavor is a property of program execution context, not of an individual interval, therefore just one flavor shall be in force at any point of execution. It is recommended that at the language level, the flavor should be constant at the level of a procedure/function, or of a compilation unit.

A flavor is identified by a unique name. Certain flavors, termed **included**, are specified in this standard. The (*list to be confirmed*) flavors are the currently included flavors. The procedure for submitting a new flavor for inclusion is described in Annex B. A conforming implementation that provides one or more included flavors may also provide non-included flavors, without losing conformance for the included flavors.

The flavor concept enforces a common core of behavior that different kinds of interval arithmetic must share:

- (i) The set of required operations, identified by their names, is the same in all flavors. Similarly the set of recommended operations is the same in all flavors. See Clause A.1, Clause A.2.
- (ii) There is a set of *common intervals* whose members are—in a sense made precise below—intervals of any flavor.
- (iii) There is a set of common evaluations of library operations, with common intervals as input, that give—again in a sense made precise below—the same result in any flavor.

The result in item (iii) is a mathematically tightest (Level 1) result, ignoring any interval widening due to finite precision (Level 2).

7.2. Definition of common intervals and common evaluations. The choice of the set of common intervals, and the set of common evaluations of an operation, is a design decision that defines the flavor concept. It should aim for simplicity, and the common evaluations should be specified by a general rule that makes it easy to add a new operation to the library if needed. The choice that was made is specified in the following paragraphs.

All likely flavors extend the classical Moore arithmetic [5] on the set \mathbb{IR} of *closed bounded nonempty real intervals*, and there is no interval outside \mathbb{IR} that is supported in all these flavors. Hence, the chosen set \mathcal{C} of common intervals is \mathbb{IR} .

The common evaluations are specified in terms of graphs of interval operations. For an interval operation φ of arity k , its graph (in some flavor) is a subset of a $(k+1)$ -dimensional space of intervals, namely the set of interval $(k+1)$ -tuples $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k; \mathbf{y})$ such that $\varphi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) = \mathbf{y}$ is true in that flavor. Each such tuple is called an *operation instance*.

The general rule is that each φ has a set $\mathcal{CE}(\varphi)$ of **common evaluations**: operation instances $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k; \mathbf{y})$ such that all its components are in \mathbb{IR} and

$$\varphi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) = \mathbf{y} \quad \text{shall hold in all flavors.}$$

$\mathcal{CE}(\varphi)$ may be regarded as the *flavor-independent graph* of φ . For brevity, writing

$$\text{the evaluation } \varphi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) = \mathbf{y} \text{ is common,} \quad (4)$$

or the equivalent notation when φ is an infix operator (e.g., $\mathbf{x}_1 + \mathbf{x}_2 = \mathbf{y}$), means that $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k; \mathbf{y}) \in \mathcal{CE}(\varphi)$.

The standard defines $\mathcal{CE}(\varphi)$ as follows.

Arithmetic operation: Such an operation is an interval extension, in fact the natural interval extension (Defn 4.2.21), of the corresponding point function φ . The common operation instances are those $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k; \mathbf{y})$ such that the point function φ is *defined and continuous at each point of* the closed, bounded, nonempty box $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k)$, and \mathbf{y} equals the range of φ over this box. Then necessarily \mathbf{y} belongs to \mathbb{IR} .

Non-arithmetic operation: The common operation instances are those tuples with common inputs \mathbf{x}_i such that *the result \mathbf{y} is also common*. In particular for **convexHull**, the common operation instances are those $(\mathbf{x}_1, \mathbf{x}_2; \mathbf{y})$ with arbitrary $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{IR}$ and \mathbf{y} equal to the convex hull of $\mathbf{x}_1 \cup \mathbf{x}_2$. For **intersection**, they are those $(\mathbf{x}_1, \mathbf{x}_2; \mathbf{y})$ with

arbitrary $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{IR}$ and \mathbf{y} equal to $\mathbf{x}_1 \cap \mathbf{x}_2$, provided the latter is *nonempty* (since $\emptyset \notin \mathbb{IR}$).

[Examples.

- The evaluation $[-1, 4]/[3, 4] = [-1/3, 4/3]$ is common; but $[3, 4]/[-1, 4] = \mathbf{y}$ is not common, for any $\mathbf{y} \in \mathbb{IR}$.
- In the set-based flavor, $\sqrt{[-1, 4]} = [0, 2]$. But in the Kaucher flavor $\sqrt{[-1, 4]}$ is undefined, so $\sqrt{[-1, 4]} = \mathbf{y}$ cannot be common for any \mathbf{y} . In fact $\mathcal{CE}(\text{sqrt})$ is the set of all $([a, b]; [\sqrt{a}, \sqrt{b}])$ for which $0 \leq a \leq b < +\infty$.
- Similarly, $([-1, 4] \cap [5, 6] = \emptyset)$ in the set-based flavor, while $([-1, 4] \cap [5, 6] = [5, 4])$ in the Kaucher flavor. Thus $([-1, 4] \cap [5, 6] = \mathbf{y})$ cannot be common for any \mathbf{y} .
- The above definition for arithmetic operations requires R1 “ φ is defined and continuous at each point of \mathbf{x} ”, which is a stronger constraint on $\mathcal{CE}(\varphi)$ (results in fewer common evaluations) than R2 “the restriction of φ to \mathbf{x} is everywhere defined and continuous”. This produces a weaker constraint on what interval arithmetics can be flavors (with fewer rules, more arithmetics obey them). In particular, requirement R1 permits cset arithmetic to be a flavor, while R2 does not.

E.g., the common evaluations of the function $\text{floor}(x)$ are all $([a, b]; [k, k])$ with $k \in \mathbb{Z}$, $k < a \leq b < k + 1$. Thus $\text{floor}([1, 1.9]) = [1, 1]$ is not common, because $\text{floor}()$ is not continuous at 1, despite its restriction to $[1, 1.9]$ being everywhere continuous. If it were required to be common, cset arithmetic could not be a flavor.

]

7.3. Loose common evaluations. At Level 2, common evaluations are usually not computable because of roundoff; instead, an enclosing interval of some finite precision interval type is computed. The notion of a **loose common evaluation** takes account of this. A loose common evaluation derived from a common evaluation $\varphi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) = \mathbf{y}$ in $\mathcal{CE}(\varphi)$ is defined to be any

$$\varphi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) = \mathbf{y}' \quad (5)$$

where \mathbf{y}' is a member of \mathbb{IR} containing \mathbf{y} . An element of $\mathcal{CE}(\varphi)$ is itself loose by this definition, but may be called a **tight** common evaluation to emphasize that its \mathbf{y} is contained in each derived \mathbf{y}' , above.

Informally, for a given φ and $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k)$, the loose common evaluations describe all closed bounded intervals that might be produced by evaluating an enclosure of $\text{Rge}(\varphi | \mathbf{x})$ in finite precision.

7.4. Relation of common evaluations to flavors. The formal definition of common evaluations takes into account that the common intervals are not necessarily a subset of the intervals of a given flavor, but are identified with a subset of it by an embedding map.

[Examples.

A Kaucher interval is defined to be a pair (a, b) of real numbers—equivalently, a point in the plane \mathbb{R}^2 —which for $a \leq b$ is “proper” and identified with the normal real interval $[a, b]$, and for $a > b$ is “improper”. Thus the embedding map is $x \mapsto (\inf x, \sup x)$ for $x \in \mathbb{IR}$.

For the set-based flavor, every common interval is actually an interval of that flavor (\mathbb{IR} is a subset of $\overline{\mathbb{IR}}$), so the embedding is the identity map $x \mapsto x$ for $x \in \mathbb{IR}$.]

Formally, a flavor is identified by a pair $(\mathfrak{F}, \mathfrak{f})$ where \mathfrak{F} is a set of Level 1 entities, the *intervals* of that flavor, and \mathfrak{f} is a one-to-one *embedding map* $\mathbb{IR} \rightarrow \mathfrak{F}$. Usually, $\mathfrak{f}(\mathbf{x})$ is abbreviated to $\mathfrak{f}\mathbf{x}$.

It is then required that *operation compatibility* shall hold for each library operation φ and for each flavor $(\mathfrak{F}, \mathfrak{f})$. Namely, given $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k$ and \mathbf{y} in \mathbb{IR} ,

$$\begin{aligned} &\text{If } (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k; \mathbf{y}) \text{ is a common operation instance of } \varphi, \\ &\text{then } (\mathfrak{f}\mathbf{x}_1, \mathfrak{f}\mathbf{x}_2, \dots, \mathfrak{f}\mathbf{x}_k; \mathfrak{f}\mathbf{y}) \text{ is an operation instance of } \varphi \text{ in flavor } \mathfrak{F}. \end{aligned} \quad (6)$$

That is, if the evaluation $\varphi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) = \mathbf{y}$ is common, then $\varphi(\mathfrak{f}\mathbf{x}_1, \mathfrak{f}\mathbf{x}_2, \dots, \mathfrak{f}\mathbf{x}_k)$ must be defined in \mathfrak{F} with value $\mathfrak{f}\mathbf{y}$.

An evaluation in \mathfrak{F} of an expression, in which only (loose) common evaluations of elementary operations occur, is called a common evaluation of that expression. That is, in a flavor $(\mathfrak{F}, \mathfrak{f})$, the expression’s inputs are members of $\mathfrak{f}(\mathbb{IR})$, and each intermediate value is produced by a common evaluation of an operation so that it is also in $\mathfrak{f}(\mathbb{IR})$; hence the final result is in $\mathfrak{f}(\mathbb{IR})$.

The `com` decoration makes it possible to determine, for a specific expression and specific interval inputs, whether common evaluation has occurred, see Clause 8.

7.5. Flavors and the Fundamental Theorem. For a common evaluation of an arithmetic expression, each library operation is (i.e., can be regarded as, modulo the embedding map) defined and continuous on its inputs so that it satisfies the conditions of the strongest, “continuous” form of the FTIA, Theorem 6.1. At Level 1, using the tightest interval extension of each operation, the range enclosure obtained by a common evaluation is (again modulo the embedding map) the same, independent of flavor.

It is possible in principle for an implementation to make this true also at Level 2, by providing shared number formats and interval types that represent the same sets of reals or intervals in each flavor; and library operations on these types and formats that have identical numerical behavior in each flavor. For example, both set-based and Kaucher flavors might use intervals stored as two IEEE754 `binary64` numbers representing the lower and upper bounds, and might ensure that operations, when applied to the intervals recognized by both flavors, behave identically. Such shared behavior might be useful for testing correctness of an implementation.

Beyond common evaluations, versions of the FTIA in different flavors can be strictly incomparable. For example, the set-based FTIA handles unbounded intervals, which the Kaucher flavor does not; while Kaucher intervals have an extended FTIA, involving generalized meanings of “contains” and “interval extension” applicable to reverse-bound intervals, which has no simple interpretation in the set-based flavor.

8. Decoration system

8.1. Decorations overview. A decoration is information attached to an interval; the combination is called a decorated interval. Interval calculation has two main objectives:

- obtaining correct range enclosures for real-valued functions of real variables;
- verifying the assumptions of existence, uniqueness, or nonexistence theorems.

Traditional interval analysis targets the first objective; the decoration system, as defined in this standard, targets the second.

A decoration primarily describes a property, not of the interval it is attached to, but of the function defined by some code that produced the interval by evaluating over some input box.

For instance, if a section of code defines the expression $\sqrt{y^2 - 1} + xy$, then decorated-interval evaluation of this code with suitably initialized input intervals \mathbf{x}, \mathbf{y} gives information about the definedness, continuity, etc. of the point function $f(x, y) = \sqrt{y^2 - 1} + xy$ over the box (\mathbf{x}, \mathbf{y}) in the plane.

The decoration system is designed in a way that naive users of interval arithmetic do not notice anything about decorations, unless they inquire explicitly about their values. They only need

- call the `newDec` operation on the inputs of any function evaluation used to invoke an existence theorem,
- explicitly convert relevant floating-point constants (but not integer parameters such as the p in `pow(x, p) = xp`) to intervals,

and have the full rigor of interval calculations available. A smart implementation may even relieve users from these tasks. Expert users can inspect, set and modify decorations to improve code efficiency, but are responsible for checking that computations done in this way remain rigorously valid.

Especially in the set-based flavor, decorations are based on the desire that, from an interval evaluation of a real function f on a box \mathbf{x} , one should get not only a range enclosure $f(\mathbf{x})$ but also a guarantee that the pair (f, \mathbf{x}) has certain important properties, such as $f(x)$ being defined for all $x \in \mathbf{x}$, f restricted to \mathbf{x} being continuous, etc. This goal is achieved, in parts of a program that require it, by performing *decorated interval evaluation*, whose semantics is summarized as follows:

Each intermediate step of the original computation depends on some or all of the inputs, so it can be viewed as an intermediate function of these inputs. The result interval obtained on each intermediate step is an enclosure for the range of the corresponding intermediate function. The decoration attached to this intermediate interval reflects the available knowledge about whether this

intermediate function is guaranteed to be everywhere defined, continuous, bounded, etc., on the given inputs.

In some flavors, certain interval operations ignore decorations, i.e., give undecorated interval output. Users are responsible for the appropriate propagation of decorations by these operations.

The function f is assumed to be expressed by code, an algebraic formula, etc.—generically termed an *expression*—which can be evaluated in several modes: point evaluation, interval evaluation, or decorated interval evaluation. The standard does not specify a definition of “expression”; however, Annex D gives formal proofs in terms of a particular definition, and indicates how this relates to expressions in some programming languages.

This standard’s decoration model, in contrast with 754’s, has no status flags. A general aim, as in 754’s use of NaN and flags, is not to interrupt the flow of computation: rather, to collate information while evaluating f , that can be inspected afterwards. This enables a fully local handling of exceptional conditions in interval calculations—important in a concurrent computing environment.

An implementation may provide any of the following: (i) status flags that are raised in the event of certain decoration values being produced by an operation; (ii) means for the user to specify that such an event signals an exception, and to invoke a system- or user-defined handler as a result. [Example. *The user may be able to specify execution be terminated if an arithmetic operation is evaluated on a box that is not wholly inside its domain—an interval version of 754’s “invalid operation” exception.*] Such features are language- or implementation-defined.

8.2. Decoration definition and propagation. Each flavor shall document its set of provided decorations and their mathematical definitions. These are flavor-defined, with the exception of the decoration `com`, see §8.3.

The implementation makes the decoration system of each flavor available to the user via *decorated interval extensions* of relevant library operations. Such an operation φ , with interval inputs $\mathbf{x}_1, \dots, \mathbf{x}_k$ carrying decorations dx_1, \dots, dx_k , shall compute the same interval output \mathbf{y} as the corresponding bare interval extension of φ —hence dependent on the \mathbf{x}_i but not on the dx_i . It shall compute a *local decoration* d , dependent on the \mathbf{x}_i and possibly on \mathbf{y} , but not on the dx_i . It shall combine d with the dx_i by a flavor-defined *propagation rule* to give an output decoration dy , and return \mathbf{y} decorated by dy .

The local decoration d may convey purely Level 1 information—e.g., that φ is everywhere continuous on the box $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_k)$. It may convey Level 2 information related to the particular finite-precision interval types being used—e.g., that \mathbf{y} , though mathematically a bounded interval, became unbounded by overflow. For diagnostic use it may convey Level 3 or 4 information, e.g., how an interval is represented, or how memory is used.

If f is an expression, decorated interval evaluation of an expression means evaluation of f with decorated interval inputs and using decorated interval extensions of the expression’s library operations. Those inputs generally need to be given suitable initial decorations that lead to the most informative output-decoration. A flavor shall provide a `newDec` function for this purpose. If \mathbf{x} is a bare interval, `newDec(\mathbf{x})` equals \mathbf{x} with such an initial decoration.

It is the responsibility of each flavor to document the meaning of its decorations, and the correct use of these decorations within programs.

8.3. Recognizing common evaluation. A flavor may provide the decoration `com` with the following propagation rule for library arithmetic operations. In an implementation with more than one flavor, each flavor shall do so.

In the following, φ denotes an arbitrary interval extension of a point library arithmetic operation φ , provided by the implementation at Level 2 (typically the one associated with a particular interval type).

Let φ applied to input intervals $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k$ give the computed result \mathbf{y} , and let $\varphi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) = \mathbf{y}$ be a loose common evaluation as defined in (5); in particular \mathbf{y} is bounded. If each of the inputs \mathbf{x}_i is decorated `com`, then the output \mathbf{y} shall be decorated `com`.

Informally, `com` records that the individual operation φ took bounded nonempty input intervals and produced a bounded (necessarily nonempty) output interval. This can be interpreted as indicating “overflow did not occur”. Further, the propagation rule ensures that if the initial inputs to an

arithmetic expression f are bounded and nonempty, and are initialized with the decoration `com`, then the final result $\mathbf{y} = f(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k)$ is decorated `com` if and only if the evaluation of the whole expression was common as defined in §7.4.

Flavors should define other decoration values, but `com` is the only one that is required to have the same meaning in all flavors.

[*Examples. Reasons why an individual evaluation of φ with common inputs $\mathbf{x} = (x_1, \dots, x_k)$ may not return `com` include the following.*

Outside domain: *The implementation finds φ is not defined and continuous everywhere on \mathbf{x} .*

Examples: $\sqrt{[-4, 4]}$, $\text{sign}([0, 2])$.

Overflow: *The Level 1 result is too large to be represented. Example: Consider an interval type \mathbb{T} whose intervals are represented by their lower and upper bounds in some floating point format, let `REALMAX` be the largest finite number in that format, and \mathbf{x} be the common \mathbb{T} -interval $[0, \text{REALMAX}]$. Then $\mathbf{x} + \mathbf{x}$ cannot be enclosed in a common \mathbb{T} -interval.*

Cost: *It is too expensive to determine whether the result can be represented. A possible example is $\tan([a, b])$ where $[a, b]$ is of a high-precision interval type, and one of its endpoints happens to be very close to a singularity of $\tan(x)$.*

]

DRAFT 7.3

DRAFT 7.3

Set-Based Intervals

This Chapter contains the standard for the set-based interval flavor.

9. Level 1 description

In this clause, subclauses §9.1 to §9.5 describe the theory of mathematical intervals and interval functions that underlies this flavor. The relation between expressions and the point or interval functions that they define is specified, since it is central to the Fundamental Theorem of Interval Arithmetic. Subclauses §9.6, 9.7 list the required and recommended *arithmetic operations* (also called elementary functions) with their mathematical specifications. Clause 10 describes, at a mathematical level, the system of *decorations* that is used among other things for exception handling in this flavor of the standard.

9.1. Non-interval Level 1 entities. In addition to intervals, this flavor deals with entities of the following kinds. They may be used as inputs or outputs of operations.

- The set $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$ of **extended reals**. Following the terminology of 754 (e.g., 754§2.1.25), any member of $\overline{\mathbb{R}}$ is called a number: it is a **finite number** if it belongs to \mathbb{R} , else an **infinite number**.

An interval’s members are finite numbers, but its bounds can be infinite. Finite or infinite numbers can be inputs to interval constructors, as well as outputs from operations, e.g., the interval width operation.

- The set of **(text) strings**, namely finite sequences of **characters** chosen from some alphabet. Since Level 1 is primarily for human communication, there are no Level 1 restrictions on the alphabet used. Strings may be inputs to interval constructors, as well as inputs or outputs of read/write operations.

9.2. Intervals. The set of mathematical intervals supported by this flavor is denoted $\overline{\mathbb{IR}}$. It consists of exactly those subsets \mathbf{x} of the real line \mathbb{R} that are closed and connected in the topological sense. Thus it comprises the empty set (denoted \emptyset or Empty) together with all the nonempty intervals, denoted $[\underline{x}, \overline{x}]$, defined by

$$[\underline{x}, \overline{x}] = \{x \in \mathbb{R} \mid \underline{x} \leq x \leq \overline{x}\}, \quad (7)$$

where \underline{x} and \overline{x} , the **bounds** of the interval, are extended-real numbers satisfying $\underline{x} \leq \overline{x}$, $\underline{x} < +\infty$ and $\overline{x} > -\infty$.

[Notes.

- The above definition implies $-\infty$ and $+\infty$ can be bounds of an interval, but are never members of it. In particular, $[-\infty, +\infty]$ is the set of all real numbers satisfying $-\infty \leq x \leq +\infty$, which is the whole real line \mathbb{R} —not the whole extended real line $\overline{\mathbb{R}}$.
- Mathematical literature generally uses a round bracket, or reversed square bracket, to show that an endpoint is excluded from an interval, e.g. $(a, b]$ and $]a, b]$ denote $\{x \mid a < x \leq b\}$. Where it is convenient to change to this notation, this is pointed out, e.g., in the tables of function domains and ranges in §9.6, 9.7.
- The set of intervals $\overline{\mathbb{IR}}$ could be described more concisely as comprising all sets $\{x \in \mathbb{R} \mid \underline{x} \leq x \leq \overline{x}\}$ for arbitrary extended-real $\underline{x}, \overline{x}$. However, this obtains Empty in many ways, as $[\underline{x}, \overline{x}]$ for any bounds satisfying $\underline{x} > \overline{x}$, and also as $[-\infty, -\infty]$ or $[+\infty, +\infty]$. The description (7) was preferred as it makes a one-to-one mapping between valid pairs $\underline{x}, \overline{x}$ of endpoints and the nonempty intervals they specify.

] A **box** or **interval vector** is an n -tuple $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ whose components \mathbf{x}_i are intervals, that is a member of $\overline{\mathbb{IR}}^n$. Usually \mathbf{x} is identified with the cartesian product $\mathbf{x}_1 \times \dots \times \mathbf{x}_n$ of its components,

a subset of \mathbb{R}^n . In particular $x \in \mathbf{x}$, for $x \in \mathbb{R}^n$, means by definition $x_i \in \mathbf{x}_i$ for all $i = 1, \dots, n$; and \mathbf{x} is empty if (and only if) any of its components \mathbf{x}_i is empty.

9.3. Hull. The (interval) **hull** of an arbitrary subset \mathbf{s} of \mathbb{R}^n , written $\text{hull}(\mathbf{s})$, is the tightest member of $\overline{\mathbb{R}}^n$ that contains \mathbf{s} . (The **tightest** set with a given property is the intersection of all sets having that property, provided the intersection itself has this property.)

9.4. Functions.

9.4.1. *Function terminology.* In this flavor, operations are written as named functions; in a specific implementation they might be represented by operators (e.g., using an infix notation), or by families of type-specific functions, or by operators or functions whose names might differ from those used here.

The terms operation, function and mapping are broadly synonymous. The following summarizes usage, with references in parentheses to precise definitions of terms.

- A *point function* (§9.4.2) is a mathematical real function of real variables. Otherwise, *function* is usually used with its general mathematical meaning.
- A (point) *arithmetic operation* (§9.4.2) is a mathematical real function for which an implementation provides versions in the implementation's *library* (§9.4.2).
- A *version* of a point function f means a function derived from f ; typically a bare or decorated interval extension (§9.4.3) of f .
- An *interval arithmetic operation* is an interval extension of a point arithmetic operation (§9.4.3).
- An *interval non-arithmetic operation* is an interval-to-interval library function that is not an interval arithmetic operation (§9.4.3).
- A *constructor* is a function that creates an interval from non-interval data (§9.6.9).

9.4.2. *Point functions.* A **point function** is a (possibly partial) multivariate real function: that is, a mapping f from a subset D of \mathbb{R}^n to \mathbb{R}^m for some integers $n \geq 0, m > 0$. It is a *scalar* function if $m = 1$, otherwise a *vector* function. When not otherwise specified, scalar is assumed. The set D where f is defined is its **domain**, also written $\text{Dom } f$. To specify n , call f an n -variable point function, or denote values of f as

$$f(x_1, \dots, x_n). \quad (8)$$

The **range** of f over an arbitrary subset \mathbf{s} of \mathbb{R}^n is the set $\text{Rge}(f | \mathbf{s})$ defined by

$$\text{Rge}(f | \mathbf{s}) = \{ f(x) \mid x \in \mathbf{s} \text{ and } x \in \text{Dom } f \}. \quad (9)$$

Thus mathematically, when evaluating a function over a set, points outside the domain are ignored—e.g., $\text{Rge}(\text{sqrt} | [-1, 1]) = [0, 1]$.

Equivalently, for the case where f takes separate arguments $\mathbf{s}_1, \dots, \mathbf{s}_n$, each being a subset of \mathbb{R} , the range is written as $\text{Rge}(f | \mathbf{s}_1, \dots, \mathbf{s}_n)$. This is an alternative notation when \mathbf{s} is the cartesian product of the \mathbf{s}_i .

A (point) **arithmetic operation** is a function for which an implementation provides versions in a collection of user-available operations called its **library**. This includes functions normally written in operator form (e.g., $+$, \times) and those normally written in function form (e.g., exp , arctan). It is not specified how an implementation provides library facilities.

9.4.3. *Interval-valued functions.* A box is an interval vector $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n) \in \overline{\mathbb{R}}^n$. It is usually identified with the cartesian product $\mathbf{x}_1 \times \dots \times \mathbf{x}_n \subseteq \mathbb{R}^n$; however, the correspondence is one-to-one only when all the \mathbf{x}_j are nonempty.

Given an n -variable scalar point function f , an **interval extension** of f is a (total) mapping \mathbf{f} from n -dimensional boxes to intervals, that is $\mathbf{f} : \overline{\mathbb{R}}^n \rightarrow \overline{\mathbb{R}}$, such that $f(x) \in \mathbf{f}(\mathbf{x})$ whenever $x \in \mathbf{x}$ and $f(x)$ is defined, equivalently

$$\mathbf{f}(\mathbf{x}) \supseteq \text{Rge}(f | \mathbf{x})$$

for any box $\mathbf{x} \in \overline{\mathbb{R}}^n$, regarded as a subset of \mathbb{R}^n . The **natural interval extension** of f is the mapping \mathbf{f} defined by

$$\mathbf{f}(\mathbf{x}) = \text{hull}(\text{Rge}(f | \mathbf{x})).$$

Equivalently, using multiple-argument notation for f , an interval extension satisfies

$$\mathbf{f}(\mathbf{x}_1, \dots, \mathbf{x}_n) \supseteq \text{Rge}(f | \mathbf{x}_1, \dots, \mathbf{x}_n),$$

and the natural interval extension is defined by

$$\mathbf{f}(\mathbf{x}_1, \dots, \mathbf{x}_n) = \text{hull}(\text{Rge}(f \mid \mathbf{x}_1, \dots, \mathbf{x}_n))$$

for any intervals $\mathbf{x}_1, \dots, \mathbf{x}_n$.

In some contexts it is useful for \mathbf{x} to be a general subset of \mathbb{R}^n , or the \mathbf{x}_i to be general subsets of \mathbb{R} ; the definition is unchanged.

The natural extension is automatically defined for all interval or set arguments. The decoration system, Clause 10, gives a way of diagnosing when the underlying point function has been evaluated outside its domain.

When f is a binary operator \bullet written in infix notation, this gives the usual definition of its natural interval extension as

$$\mathbf{x} \bullet \mathbf{y} = \text{hull}(\{x \bullet y \mid x \in \mathbf{x}, y \in \mathbf{y}, \text{ and } x \bullet y \text{ is defined}\}).$$

[*Example. With these definitions, the relevant natural interval extensions satisfy $\sqrt{[-1, 4]} = [0, 2]$ and $\sqrt{[-2, -1]} = \emptyset$; also $\mathbf{x} \times [0, 0] = [0, 0]$ for any nonempty \mathbf{x} , and $\mathbf{x}/[0, 0] = \emptyset$, for any \mathbf{x} .*]

When f is a vector point function, a vector interval function with the same number of inputs and outputs as f is called an interval extension of f if each of its components is an interval extension of the corresponding component of f .

An interval-valued function in the library is called an **interval arithmetic operation** if it is an interval extension of a point arithmetic operation, and an **interval non-arithmetic operation** otherwise. Examples of the latter are interval intersection and union, $(\mathbf{x}, \mathbf{y}) \mapsto \mathbf{x} \cap \mathbf{y}$ and $(\mathbf{x}, \mathbf{y}) \mapsto \text{hull}(\mathbf{x} \cup \mathbf{y})$.

9.4.4. *Constants.* A real scalar function with no arguments—a mapping $\mathbb{R}^n \rightarrow \mathbb{R}^m$ with $n = 0$ and $m = 1$ —is a **real constant**. Languages may distinguish between a literal constant (e.g., the decimal value defined by the string 1.23e4) and a named constant (e.g., π) but the difference is not relevant on Level 1 (and easily handled by outward rounding on Level 2).

From the definition, an interval extension of a real constant is any zero-argument interval function that returns an interval containing c . The *natural extension* returns the interval $[c, c]$.

9.5. Expressions. This flavor gives the term “expression” the general meaning described in Clause 6.

9.6. Required operations.

For the functions listed in this subclause, an implementation shall provide interval versions appropriate to its supported interval types. For constants and the forward and reverse arithmetic operations in §9.6.1, 9.6.3, 9.6.4, 9.6.6, each such version shall be an interval extension (§9.4.3) of the corresponding point function—for a constant, that means any constant interval enclosing the point value. The required rounding behavior of these, and of the numeric functions of intervals in §9.6.10, is detailed in §11.9, 11.11.

The names of operations, as well as symbols used for operations (e.g., for the comparisons in §9.6.11), may not correspond to those that any particular language would use.

9.6.1. Interval literals.

An **interval literal** is a text string that denotes an interval. Level 1, which is mainly for human communication, merely assumes some agreed rules on the form and meaning of interval literals exist. A specified form and meaning shall be used in Level 2 onward: the definition is in §11.11.1. This definition is also used in Level 1 of this document for examples, where relevant. [Example. This includes the *inf-sup* form `[1.234e5,Inf]`; the *mid-rad* form `<3.1416 +- 1.2e-5>`; the *uncertain* form `3.1416?1`; and the *named interval constant* `[Empty]`.]

9.6.2. Interval constants.

The constant functions `empty()` and `entire()` have value `Empty` and `Entire` respectively.

9.6.3. Forward-mode elementary functions.

Table 1 on page 23 lists required arithmetic operations. The term *operation* includes functions normally written in function notation $f(x, y, \dots)$, as well as those normally written in unary or binary operator notation, $\bullet x$ or $x \bullet y$.

[Note. The list includes all general-computational operations in 754§5.4 except `convertFromInt`, and some recommended functions in 754§9.2.]

Proving the correctness of interval computations relies on the Fundamental Theorem of Interval Arithmetic, which in turn relies on the relation between a point-function and its interval extensions. Thus the domain of each point-function and its value at each point of the domain are specified to aid a rigorous implementation. This is mostly straightforward but needs care for functions with discontinuities, such as `pow()` and `atan2()`.

Notes to Table 1

- In describing the domain, notation such as $\{y = 0\}$ is short for $\{(x, y) \in \mathbb{R}^2 \mid y = 0\}$, etc.
- Regarded as a family of functions parameterized by the integer argument p .
- Defined as $e^{y \ln x}$ for real $x > 0$ and all real y , and 0 for $x = 0$ and $y > 0$, else there is no value at Level 1.
- $b = e, 2$ or 10 , respectively.
- The ranges shown are the mathematical range of the point function. To ensure containment, an interval result may include values just outside the mathematical range.
- `atan2(y, x)` is the principal value of the argument (polar angle) of (x, y) in the plane.
- To avoid confusion with notation for open intervals, in this table coordinates in \mathbb{R}^2 are delimited by angle brackets $\langle \rangle$.
- `sign(x)` is -1 if $x < 0$; 0 if $x = 0$; and 1 if $x > 0$.
- `ceil(x)` is the smallest integer $\geq x$. `floor(x)` is the largest integer $\leq x$. `roundTiesToEven(x)`, `roundTiesToAway(x)` are the nearest integer to x , with ties rounded to the even integer or away from zero respectively. `trunc(x)` is the nearest integer to x in the direction of zero. (As defined in the C standard §7.12.9.)
- Smallest, or largest, of its real arguments. A family of functions parameterized by the arity k .

9.6.4. Interval case expressions and case function.

Functions are often defined by conditionals: $f(x)$ equals $g(x)$ if some condition on x holds, and $h(x)$ otherwise. To handle interval extensions of such functions in a way that automatically conforms to the Fundamental Theorem of Interval Arithmetic, the ternary function `case(c, g, h)` is provided. To simplify defining its interval extension, the argument c specifying the condition is real (instead of boolean), and the condition means $c < 0$ by definition. That is,

$$\text{case}(c, g, h) = \begin{cases} \text{if } c < 0 & \text{then } g, \\ & \text{else } h. \end{cases}$$

TABLE 1. Required forward elementary functions.

Normal mathematical notation is used to include or exclude an interval endpoint, e.g., $(-\pi, \pi]$ denotes $\{x \in \mathbb{R} \mid -\pi < x \leq \pi\}$.

Name	Definition	Point function domain	Point function range	Note
$\text{neg}(x)$	$-x$	\mathbb{R}	\mathbb{R}	
$\text{add}(x, y)$	$x + y$	\mathbb{R}^2	\mathbb{R}	
$\text{sub}(x, y)$	$x - y$	\mathbb{R}^2	\mathbb{R}	
$\text{mul}(x, y)$	xy	\mathbb{R}^2	\mathbb{R}	
$\text{div}(x, y)$	x/y	$\mathbb{R}^2 \setminus \{y = 0\}$	\mathbb{R}	a
$\text{recip}(x)$	$1/x$	$\mathbb{R} \setminus \{0\}$	$\mathbb{R} \setminus \{0\}$	
$\text{sqrt}(x)$	\sqrt{x}	$[0, \infty)$	$[0, \infty)$	
$\text{fma}(x, y, z)$	$(x \times y) + z$	\mathbb{R}^3	\mathbb{R}	
$\text{case}(c, g, h)$		See §9.6.4.		
$\text{sqr}(x)$	x^2	\mathbb{R}	$[0, \infty)$	
$\text{pown}(x, p)$	$x^p, p \in \mathbb{Z}$	$\begin{cases} \mathbb{R} & \text{if } p \geq 0 \\ \mathbb{R} \setminus \{0\} & \text{if } p < 0 \end{cases}$	$\begin{cases} \mathbb{R} & \text{if } p > 0 \text{ odd} \\ [0, \infty) & \text{if } p > 0 \text{ even} \\ \{1\} & \text{if } p = 0 \\ \mathbb{R} \setminus \{0\} & \text{if } p < 0 \text{ odd} \\ (0, \infty) & \text{if } p < 0 \text{ even} \end{cases}$	b
$\text{pow}(x, y)$	x^y	$\{x > 0\} \cup \{x = 0, y > 0\}$	$[0, \infty)$	a, c
$\text{exp}, \text{exp2}, \text{exp10}(x)$	b^x	\mathbb{R}	$(0, \infty)$	d
$\text{log}, \text{log2}, \text{log10}(x)$	$\log_b x$	$(0, \infty)$	\mathbb{R}	d
$\text{sin}(x)$		\mathbb{R}	$[-1, 1]$	
$\text{cos}(x)$		\mathbb{R}	$[-1, 1]$	
$\text{tan}(x)$		$\mathbb{R} \setminus \{(k + \frac{1}{2})\pi \mid k \in \mathbb{Z}\}$	\mathbb{R}	
$\text{asin}(x)$		$[-1, 1]$	$[-\pi/2, \pi/2]$	e
$\text{acos}(x)$		$[-1, 1]$	$[0, \pi]$	e
$\text{atan}(x)$		\mathbb{R}	$(-\pi/2, \pi/2)$	e
$\text{atan2}(y, x)$		$\mathbb{R}^2 \setminus \{(0, 0)\}$	$(-\pi, \pi]$	e, f, g
$\text{sinh}(x)$		\mathbb{R}	\mathbb{R}	
$\text{cosh}(x)$		\mathbb{R}	$[1, \infty)$	
$\text{tanh}(x)$		\mathbb{R}	$(-1, 1)$	
$\text{asinh}(x)$		\mathbb{R}	\mathbb{R}	
$\text{acosh}(x)$		$[1, \infty)$	$[0, \infty)$	
$\text{atanh}(x)$		$(-1, 1)$	\mathbb{R}	
$\text{sign}(x)$		\mathbb{R}	$\{-1, 0, 1\}$	h
$\text{ceil}(x)$		\mathbb{R}	\mathbb{Z}	i
$\text{floor}(x)$		\mathbb{R}	\mathbb{Z}	i
$\text{trunc}(x)$		\mathbb{R}	\mathbb{Z}	i
$\text{roundTiesToEven}(x)$		\mathbb{R}	\mathbb{Z}	i
$\text{roundTiesToAway}(x)$		\mathbb{R}	\mathbb{Z}	i
$\text{abs}(x)$	$ x $	\mathbb{R}	$[0, \infty)$	
$\text{min}(x_1, \dots, x_k)$		\mathbb{R}^k for $k = 2, 3, \dots$	\mathbb{R}	j
$\text{max}(x_1, \dots, x_k)$		\mathbb{R}^k for $k = 2, 3, \dots$	\mathbb{R}	j

An implementation shall provide the following interval extension (see the Notes):

$$\text{case}(c, g, h) = \begin{cases} \text{if } c \text{ is empty} & \text{then } \emptyset \\ \text{elseif } c \text{ is a subset of the half-line } x < 0 & \text{then } g \\ \text{elseif } c \text{ is a subset of the half-line } x \geq 0 & \text{then } h \\ \text{else} & \text{convexHull}(g, h). \end{cases} \quad (10)$$

for any intervals c, g, h .

The function f above may be encoded as $f(x) = \text{case}(c(x), g(x), h(x))$. Then, if c, g, h are interval functions that are interval extensions of point functions c, g and h , the function

$$\mathbf{f}(\mathbf{x}) = \text{case}(\mathbf{c}(\mathbf{x}), \mathbf{g}(\mathbf{x}), \mathbf{h}(\mathbf{x})) \quad (11)$$

is automatically an interval extension of f .

[Notes.

1. Equation (10) does not define the natural interval extension, which has value Empty if any of its input arguments is empty. Its advantage is that for a function defined by a conditional expression, such as (11), it allows “short-circuiting”. That is, one can suppress evaluation of $\mathbf{h}(\mathbf{x})$ if $\bar{c} < 0$, and of $\mathbf{g}(\mathbf{x})$ if $\underline{c} \geq 0$. This is not so for the natural extension.
2. This method is less awkward than using interval comparisons as a mechanism for handling such functions. However, the resulting interval function is usually not the tightest extension of the corresponding point function. E.g., the (point) absolute value $|x|$ may be defined by

$$|x| = \text{case}(x, -x, x).$$

Then it is easy to see that formula (11), applied to a nonempty $\mathbf{x} = [\underline{x}, \bar{x}]$, gives the exact range $\{|x| \mid x \in \mathbf{x}\}$ when $\bar{x} < 0$ or $0 \leq \underline{x}$, but the poor enclosure $(-\mathbf{x}) \cup \mathbf{x}$ when $\underline{x} < 0 \leq \bar{x}$.

3. $\text{case}(c, g, h)$ is equivalent to the C expression $(c < 0 ? g : h)$.
4. Compound conditions may be expressed using the max and min operations: e.g., a real function $f(x, y)$ that equals $\sin(xy)$ in the positive quadrant of the plane, and zero elsewhere, may be written

$$f(x, y) = \text{case}(\min(x, y), 0, \sin(xy)),$$

since $\min(x, y) < 0$ is equivalent to $(x < 0 \text{ or } y < 0)$.

]

9.6.5. Two-output division.

The result of interval division \mathbf{x}/\mathbf{y} considered as a set, that is

$$\mathbf{x}/_{\text{set}} \mathbf{y} = \{x/y \mid x \in \mathbf{x}, y \in \mathbf{y} \text{ and } y \neq 0\},$$

can have zero, one or two nonempty disjoint connected components, which need not be closed. [Examples. Use the classical notation where a square or round bracket means a closed or open endpoint respectively. Then

$$\begin{aligned} [1, 2] /_{\text{set}} [0, 0] &= \emptyset, \\ [1, 2] /_{\text{set}} [1, 1] &= [1, 2], \\ [1, 1] /_{\text{set}} [1, +\infty) &= (0, 1], \\ [1, 2] /_{\text{set}} [-1, 1] &= (-\infty, -1] \cup [1, +\infty), \\ [1, 1] /_{\text{set}} \text{Entire} &= (-\infty, 0) \cup (0, +\infty), \end{aligned}$$

are cases with 0, 1, 1, 2 and 2 output components respectively. The third and fifth cases have components that are not closed.]

In applications such as the Interval Newton Method, it is useful to return enclosures of these components separately rather than the result of normal division which is the (closed) convex hull of their union, namely $\mathbf{x}/\mathbf{y} = \text{hull}(\mathbf{x}/_{\text{set}} \mathbf{y})$. The value of the operation $\text{div2pair}(\mathbf{x}, \mathbf{y})$ is an ordered pair (\mathbf{u}, \mathbf{v}) of closed intervals, namely

$$\text{div2pair}(\mathbf{x}, \mathbf{y}) = \begin{cases} (\emptyset, \emptyset) & \text{if } \mathbf{x}/_{\text{set}} \mathbf{y} \text{ is empty,} \\ (\bar{\mathbf{u}}, \emptyset) & \text{if } \mathbf{x}/_{\text{set}} \mathbf{y} \text{ has one component } \mathbf{u}, \\ (\bar{\mathbf{u}}, \bar{\mathbf{v}}) & \text{if } \mathbf{x}/_{\text{set}} \mathbf{y} \text{ has two components } \mathbf{u}, \mathbf{v}, \text{ ordered so that } \mathbf{u} < \mathbf{v}, \end{cases}$$

where $\bar{\mathbf{a}}$ denotes the topological closure of \mathbf{a} , which in this case is equivalent to $\text{hull}(\mathbf{a})$.

9.6.6. Reverse-mode elementary functions.

Constraint-satisfaction algorithms use the functions in this subclause for iteratively tightening an enclosure of a solution to a system of equations.

Given a unary arithmetic operation φ , a **reverse interval extension** of φ is a binary interval function φRev such that

$$\varphi\text{Rev}(\mathbf{c}, \mathbf{x}) \supseteq \{x \in \mathbf{x} \mid \varphi(x) \text{ is defined and in } \mathbf{c}\}, \quad (12)$$

for any intervals \mathbf{c}, \mathbf{x} .

TABLE 2. Required reverse elementary functions.

From unary functions	From binary functions
$\text{sqrRev}(c, x)$	$\text{mulRev}(b, c, x)$
$\text{recipRev}(c, x)$	$\text{divRev1}(b, c, x)$
$\text{absRev}(c, x)$	$\text{divRev2}(a, c, x)$
$\text{pownRev}(c, x, p)$	$\text{powRev1}(b, c, x)$
$\text{sinRev}(c, x)$	$\text{powRev2}(a, c, x)$
$\text{cosRev}(c, x)$	$\text{atan2Rev1}(b, c, x)$
$\text{tanRev}(c, x)$	$\text{atan2Rev2}(a, c, x)$
$\text{coshRev}(c, x)$	

Similarly, a binary arithmetic operation \bullet has two forms of reverse interval extension, which are ternary interval functions $\bullet\text{Rev}_1$ and $\bullet\text{Rev}_2$ such that

$$\bullet\text{Rev}_1(b, c, x) \supseteq \{x \in x \mid b \in b \text{ exists such that } x \bullet b \text{ is defined and in } c\}, \quad (13)$$

$$\bullet\text{Rev}_2(a, c, x) \supseteq \{x \in x \mid a \in a \text{ exists such that } a \bullet x \text{ is defined and in } c\}. \quad (14)$$

If \bullet is commutative then $\bullet\text{Rev}_1$ and $\bullet\text{Rev}_2$ agree and may be implemented simply as $\bullet\text{Rev}$.

In each of (12, 13, 14), the unique **natural reverse interval extension** is the one whose value is the interval hull of the right-hand side. Clearly, any reverse interval extension encloses this hull.

The last argument x in each of (12, 13, 14) is optional, with default $x = \mathbb{R}$ if absent.

[*Note. The argument x can be thought of as giving prior knowledge about the range of values taken by a point-variable x , which is then sharpened by applying the reverse function: see the example below.*]

Reverse operations shall be provided as in Table 2. Note $\text{pownRev}(x, p)$ is regarded as a family of unary functions parametrized by p .

[*Example.*

- Consider the function $\text{sqr}(x) = x^2$. Evaluating $\text{sqrRev}([1, 4])$ answers the question: given that $1 \leq x^2 \leq 4$, what interval can we restrict x to? Using the natural reverse extension, we have

$$\text{sqrRev}([1, 4]) = \text{hull}\{x \in \mathbb{R} \mid x^2 \in [1, 4]\} = \text{hull}([-2, -1] \cup [1, 2]) = [-2, 2].$$

- If we can add the prior knowledge that $x \in x = [0, 1.2]$, then using the optional second argument gives the tighter enclosure

$$\text{sqrRev}([1, 4], [0, 1.2]) = \text{hull}\{x \in [0, 1.2] \mid x^2 \in [1, 4]\} = \text{hull}([0, 1.2] \cap ([-2, -1] \cup [1, 2])) = [1, 1.2].$$

- One might think it suffices to apply the operation without the optional argument and intersect the result with x . This is less effective because “hull” and “intersect” do not commute. E.g., in the above, this method evaluates

$$\text{sqrRev}([1, 4]) \cap x = [-2, 2] \cap [0, 1.2] = [0, 1.2],$$

so no tightening of the enclosure x is obtained.

]

9.6.7. Cancellative addition and subtraction.

⚠ I have made this only apply to bounded intervals, since it really seems hard to frame a specification for unbounded ones that translates unambiguously to the finite precision (Level 2) situation.

Also I have deviated from the Motion 12 spec, by making the operations defined only when $\text{width}(x) \geq \text{width}(y)$. IMO it is actively harmful in applications to make it always defined. This is for the same reasons that it is harmful to replace \sqrt{x} by the everywhere defined $\sqrt{|x|}$: an application MUST test for definedness, and making it always defined leads to un-noticed wrong results from code that forgets to test. With Kaucher/modal intervals a different choice may be appropriate.

Cancellative subtraction solves the problem: Recover interval z from intervals x and y , given that one knows x was obtained as the sum $y + z$.

[*Example. In some applications one has a list of intervals a_1, \dots, a_n , and needs to form each interval s_k which is the sum of all the a_i except a_k , that is $s_k = \sum_{i=1, i \neq k}^n a_i$, for $k = 1, \dots, n$.*

Evaluating all these sums independently costs $O(n^2)$ work. However, if one forms the sum s of all the

a_i , one can obtain each s_k from s and a_k by cancellative subtraction. This method only costs $O(n)$ work.

This example illustrates that in finite precision, computing x (as a sum of terms) typically incurs at least one roundoff error, and may incur many. Thus the model underlying these cancellative operations is that x is an enclosure of an unknown true sum x_0 , whereas y is “exact”. The computed z is thus an enclosure of an unknown true z_0 such that $y + z_0 = x_0$.]

The operation `cancelPlus(x, y)` is equivalent to `cancelMinus(x, -y)` and therefore not specified separately.

For any two bounded intervals x and y , the value of the operation `cancelMinus(x, y)` is the tightest interval z such that

$$y + z \supseteq x \quad (15)$$

if such a z exists. Otherwise `cancelMinus(x, y)` has no value at Level 1.

This specification leads to the following Level 1 algorithm. If $x = \emptyset$ then $z = \emptyset$. If $x \neq \emptyset$ and $y = \emptyset$ then z has no value. If $x = [\underline{x}, \bar{x}]$ and $y = [\underline{y}, \bar{y}]$ are both nonempty and bounded, define $\underline{z} = \underline{x} - \bar{y}$ and $\bar{z} = \bar{x} - \underline{y}$. Then z is defined to be $[\underline{z}, \bar{z}]$ if $\underline{z} \leq \bar{z}$ (equivalently if $\text{width}(x) \geq \text{width}(y)$), and has no value otherwise. If either x or y is unbounded, z has no value.

[Note. Because of the cancellative nature of these operations, care is needed in finite precision to determine whether the result is defined or not. More details are given at Level 3 in §13.7.]

9.6.8. Non-arithmetic (set) operations.

The value of the operation `intersection(x, y)` is the intersection $x \cap y$ of the intervals x and y .

The value of the operation `convexHull(x, y)` is the interval hull of the union $x \cup y$ of the intervals x and y .

9.6.9. Constructors.

An interval constructor by definition is an operation that creates a bare or decorated interval from non-interval data. The following bare interval constructors shall be provided.

The operation `nums2interval(l, u)`, takes extended-real values l and u . If (see §9.2) the conditions $l \leq u$, $l < +\infty$ and $u > -\infty$ hold, its value is the nonempty interval $[l, u] = \{x \in \mathbb{R} \mid l \leq x \leq u\}$. Otherwise it has no value.

The operation `text2interval(s)` takes a text string s . If s is a valid interval literal, see §9.6.1, 11.11.1, its value is the interval denoted by s . Otherwise it has no value.

9.6.10. Numeric functions of intervals.

The operations in Table 3 shall be provided, the argument being an interval and the result a number, which for some of the operations may be infinite.

[Note. Implementations should provide an operation that returns `mid(x)` and `rad(x)` simultaneously.]

9.6.11. Boolean functions of intervals.

The following operations shall be provided, whose value is a boolean (1 = true, 0 = false) result.

There is a function `isEmpty(x)`, with value 1 if x is the empty set, 0 otherwise. There is a function `isEntire(x)`, with value 1 if x is the whole line, 0 otherwise.

There are eight boolean-valued comparison relations, which take two interval inputs. These are defined in Table 4, in which column three gives the set-theoretic definition, and column four gives an equivalent specification when both intervals are nonempty. Table 5 shows what the definitions imply when at least one interval is empty.

[Notes.

- Column two of Table 5 gives suggested symbols for use in typeset algorithms.
- All these relations, except $a \cap b$, are transitive for nonempty intervals.
- The first three are reflexive.
- `isInterior` uses the topological definition: b is a neighbourhood of each point of a . This implies, for instance, that `isInterior(Entire, Entire)` is true.
- In fact all occurrences of $<$ in column 4 of Table 4 can be replaced by $<'$.

]

TABLE 3. Required numeric functions of an interval $\mathbf{x} = [\underline{x}, \bar{x}]$.
Note `inf` can have value $-\infty$; each of `sup`, `wid`, `rad` and `mag` can have value $+\infty$.

Name	Definition
<code>inf(x)</code>	$\begin{cases} \text{lower bound of } \mathbf{x} \text{ if } \mathbf{x} \text{ is nonempty} \\ \infty \text{ if } \mathbf{x} \text{ is empty} \end{cases}$
<code>sup(x)</code>	$\begin{cases} \text{upper bound of } \mathbf{x} \text{ if } \mathbf{x} \text{ is nonempty} \\ -\infty \text{ if } \mathbf{x} \text{ is empty} \end{cases}$
<code>mid(x)</code>	$\begin{cases} \text{midpoint } (\underline{x} + \bar{x})/2 \text{ if } \mathbf{x} \text{ is nonempty bounded} \\ \text{no value if } \mathbf{x} \text{ is empty or unbounded} \end{cases}$
<code>wid(x)</code>	$\begin{cases} \text{width } \bar{x} - \underline{x} \text{ if } \mathbf{x} \text{ is nonempty} \\ \text{no value if } \mathbf{x} \text{ is empty} \end{cases}$
<code>rad(x)</code>	$\begin{cases} \text{radius } (\bar{x} - \underline{x})/2 \text{ if } \mathbf{x} \text{ is nonempty} \\ \text{no value if } \mathbf{x} \text{ is empty} \end{cases}$
<code>mag(x)</code>	$\begin{cases} \text{magnitude } \sup\{ x \mid x \in \mathbf{x}\} \text{ if } \mathbf{x} \text{ is nonempty} \\ \text{no value if } \mathbf{x} \text{ is empty} \end{cases}$
<code>mig(x)</code>	$\begin{cases} \text{mignitude } \inf\{ x \mid x \in \mathbf{x}\} \text{ if } \mathbf{x} \text{ is nonempty} \\ \text{no value if } \mathbf{x} \text{ is empty} \end{cases}$

TABLE 4. Comparisons for intervals \mathbf{a} and \mathbf{b} . Notation \forall_a means “for all a in \mathbf{a} ”, and so on. In column 4, $\mathbf{a}=[\underline{a}, \bar{a}]$ and $\mathbf{b}=[\underline{b}, \bar{b}]$, where $\underline{a}, \underline{b}$ may be $-\infty$ and \bar{a}, \bar{b} may be $+\infty$; and $<'$ is the same as $<$ except that $-\infty <' -\infty$ and $+\infty <' +\infty$ are true.

Name	Symbol	Definition	For $\mathbf{a}, \mathbf{b} \neq \emptyset$	Description
<code>equal(a, b)</code>	$\mathbf{a} = \mathbf{b}$	$\forall_a \exists_b a = b \wedge \forall_b \exists_a b = a$	$\underline{a} = \underline{b} \wedge \bar{a} = \bar{b}$	\mathbf{a} equals \mathbf{b}
<code>containedIn(a, b)</code>	$\mathbf{a} \subseteq \mathbf{b}$	$\forall_a \exists_b a = b$	$\underline{b} \leq \underline{a} \wedge \bar{a} \leq \bar{b}$	\mathbf{a} is a subset of \mathbf{b}
<code>less(a, b)</code>	$\mathbf{a} \leq \mathbf{b}$	$\forall_a \exists_b a \leq b \wedge \forall_b \exists_a a \leq b$	$\underline{a} \leq \underline{b} \wedge \bar{a} \leq \bar{b}$	\mathbf{a} is weakly less than \mathbf{b}
<code>precedes(a, b)</code>	$\mathbf{a} \prec \mathbf{b}$	$\forall_a \forall_b a \leq b$	$\bar{a} \leq \underline{b}$	\mathbf{a} is to left of but may touch \mathbf{b}
<code>isInterior(a, b)</code>	$\mathbf{a} @ \mathbf{b}$	$\forall_a \exists_b a < b \wedge \forall_a \exists_b b < a$	$\underline{b} <' \underline{a} \wedge \bar{a} <' \bar{b}$	\mathbf{a} is interior to \mathbf{b}
<code>strictlyLess(a, b)</code>	$\mathbf{a} < \mathbf{b}$	$\forall_a \exists_b a < b \wedge \forall_b \exists_a a < b$	$\underline{a} <' \underline{b} \wedge \bar{a} <' \bar{b}$	\mathbf{a} is strictly less than \mathbf{b}
<code>strictlyPrecedes(a, b)</code>	$\mathbf{a} \prec' \mathbf{b}$	$\forall_a \forall_b a < b$	$\bar{a} < \underline{b}$	\mathbf{a} is strictly to left of \mathbf{b}
<code>areDisjoint(a, b)</code>	$\mathbf{a} \not\cap \mathbf{b}$	$\forall_a \forall_b a \neq b$	$\bar{a} < \underline{b} \vee \bar{b} < \underline{a}$	\mathbf{a} and \mathbf{b} are disjoint

TABLE 5. Comparisons with empty intervals.

	$\mathbf{a} = \emptyset$ $\mathbf{b} \neq \emptyset$	$\mathbf{a} \neq \emptyset$ $\mathbf{b} = \emptyset$	$\mathbf{a} = \emptyset$ $\mathbf{b} = \emptyset$
$\mathbf{a} = \mathbf{b}$	0	0	1
$\mathbf{a} \subseteq \mathbf{b}$	1	0	1
$\mathbf{a} \leq \mathbf{b}$	0	0	1
$\mathbf{a} \prec \mathbf{b}$	1	1	1
$\mathbf{a} @ \mathbf{b}$	1	0	1
$\mathbf{a} < \mathbf{b}$	0	0	1
$\mathbf{a} \prec' \mathbf{b}$	1	1	1
$\mathbf{a} \not\cap \mathbf{b}$	1	1	1

9.7. Recommended operations.

An implementation should provide interval versions of the functions listed in this subclause. If such an interval version is provided, it shall behave as specified here.

9.7.1. Forward-mode elementary functions.

The list of recommended functions is in Table 6. Each interval version shall be an interval extension of the point function.

Notes to Table 6

TABLE 6. Recommended elementary functions.
Normal mathematical notation is used to include or exclude an interval endpoint,
e.g., $(-1, 1]$ denotes $\{x \in \mathbb{R} \mid -1 < x \leq 1\}$.

Name	Definition	Point function domain	Point function range	Note
$\text{rootn}(x, q)$	real $\sqrt[q]{x}$, $q \in \mathbb{Z} \setminus \{0\}$	$\begin{cases} \mathbb{R} & \text{if } q > 0 \text{ odd} \\ [0, \infty) & \text{if } q > 0 \text{ even} \\ \mathbb{R} \setminus \{0\} & \text{if } q < 0 \text{ odd} \\ (0, \infty) & \text{if } q < 0 \text{ even} \end{cases}$	same as domain	a
$\begin{cases} \text{expm1}(x) \\ \text{exp2m1}(x) \\ \text{exp10m1}(x) \end{cases}$	$b^x - 1$	\mathbb{R}	$(-1, \infty)$	b, c
$\begin{cases} \text{logp1}(x) \\ \text{log2p1}(x) \\ \text{log10p1}(x) \end{cases}$	$\log_b(x+1)$	$(-1, \infty)$	\mathbb{R}	b, c
$\text{compoundm1}(x, y)$	$(1+x)^y - 1$	$\{x > -1\} \cup \{x = -1, y > 0\}$	$[0, \infty)$	c, d
$\text{hypot}(x, y)$	$\sqrt{x^2 + y^2}$	\mathbb{R}^2	$[0, \infty)$	
$\text{rSqrt}(x)$	$1/\sqrt{x}$	$(0, \infty)$	$(0, \infty)$	
$\text{sinPi}(x)$	$\sin(\pi x)$	\mathbb{R}	$[-1, 1]$	e
$\text{cosPi}(x)$	$\cos(\pi x)$	\mathbb{R}	$[-1, 1]$	e
$\text{tanPi}(x)$	$\tan(\pi x)$	$\mathbb{R} \setminus \{k + \frac{1}{2} \mid k \in \mathbb{Z}\}$	\mathbb{R}	e
$\text{asinPi}(x)$	$\arcsin(x)/\pi$	$[-1, 1]$	$[-1/2, 1/2]$	e
$\text{acosPi}(x)$	$\arccos(x)/\pi$	$[-1, 1]$	$[0, 1]$	e
$\text{atanPi}(x)$	$\arctan(x)/\pi$	\mathbb{R}	$(-1/2, 1/2)$	e
$\text{atan2Pi}(y, x)$	$\text{atan2}(y, x)/\pi$	$\mathbb{R}^2 \setminus \{(0, 0)\}$	$(-1, 1]$	e, f

- Regarded as a family of functions parameterized by the integer arguments q , or r and s .
- $b = e, 2$ or 10 , respectively.
- Mathematically unnecessary, but included to let implementations give better numerical behavior for small values of the arguments.
- In describing domains, notation such as $\{y = 0\}$ is short for $\{(x, y) \in \mathbb{R}^2 \mid y = 0\}$, and so on.
- These functions avoid a loss of accuracy due to π being irrational, cf. Table 1, note e.
- To avoid confusion with notation for open intervals, in this table coordinates in \mathbb{R}^2 are delimited by angle brackets $\langle \rangle$.

9.7.2. Slope functions.

The functions in Table 7 are the commonest ones needed to efficiently implement improved range enclosures via *first- and second-order slope* algorithms. They are analytic at $x = 0$ after filling in the removable singularity there, where each has the value 1.

9.7.3. Extended interval comparisons.


The **interval overlapping operation** $\text{overlap}(\mathbf{a}, \mathbf{b})$, also written $\mathbf{a} \otimes \mathbf{b}$, arises from the work of J.F. Allen [1] on temporal logic. It may be used as an infrastructure for other interval comparisons. If implemented, it should also be available at user level; how this is done is implementation-defined or language-defined.

Allen identified 13 states of a pair (\mathbf{a}, \mathbf{b}) of nonempty intervals, which are ways in which they can be related with respect to the usual order $a < b$ of the reals. Together with three states for when either interval is empty, these define the 16 possible values of $\text{overlap}(\mathbf{a}, \mathbf{b})$.

To describe the states for nonempty intervals of positive width, it is useful to think of $\mathbf{b} = [\underline{b}, \bar{b}]$ (with $\underline{b} < \bar{b}$) as fixed, while $\mathbf{a} = [\underline{a}, \bar{a}]$ (with $\underline{a} < \bar{a}$) starts far to its left and moves to the right. Its endpoints move continuously with strictly positive velocity. Then, depending on the relative sizes of \mathbf{a} and \mathbf{b} , the value of $\mathbf{a} \otimes \mathbf{b}$ follows a path from left to right through the graph below, whose nodes represent Allen's 13 states.

TABLE 7. Recommended slope functions.

Name	Definition	Point function domain	Point function range	Note
<code>expSlope1(x)</code>	$\frac{1}{x}(e^x - 1)$	\mathbb{R}	$(0, \infty)$	
<code>expSlope2(x)</code>	$\frac{2}{x^2}(e^x - 1 - x)$	\mathbb{R}	$(0, \infty)$	
<code>logSlope1(x)</code>	$\frac{2}{x^2}(\log(1+x) - x)$	\mathbb{R}	$(0, \infty)$	
<code>logSlope2(x)</code>	$\frac{3}{x^3}(\log(1+x) - x + \frac{x^2}{2})$	\mathbb{R}	$(0, \infty)$	
<code>cosSlope2(x)</code>	$-\frac{2}{x^2}(\cos x - 1)$	\mathbb{R}	$[0, 1]$	
<code>sinSlope3(x)</code>	$-\frac{6}{x^3}(\sin x - x)$	\mathbb{R}	$(0, 1]$	
<code>asinSlope3(x)</code>	$\frac{6}{x^3}(\arcsin x - x)$	$[-1, 1]$	$[1, 3\pi - 6]$	
<code>atanSlope3(x)</code>	$-\frac{3}{x^3}(\arctan x - x)$	\mathbb{R}	$(0, 1]$	
<code>coshSlope2(x)</code>	$\frac{2}{x^2}(\cosh x - 1)$	\mathbb{R}	$[1, \infty)$	
<code>sinhSlope3(x)</code>	$\frac{3}{x^3}(\sinh x - x)$	\mathbb{R}	$[\frac{1}{2}, \infty)$	



For instance “ \mathbf{a} overlaps \mathbf{b} ”—equivalently $\mathbf{a} \odot \mathbf{b}$ has the value `overlaps`—is the case $\underline{a} < \underline{b} < \bar{a} < \bar{b}$.

The three extra values are: `bothEmpty` when $\mathbf{a} = \mathbf{b} = \emptyset$, else `firstEmpty` when $\mathbf{a} = \emptyset$, `secondEmpty` when $\mathbf{b} = \emptyset$.

Table 8 shows the 16 states, with the 13 “nonempty” states specified (a) in terms of set membership using quantifiers and (b) in terms of the endpoints $\underline{a}, \bar{a}, \underline{b}, \bar{b}$, and also (c) shown diagrammatically.

The set and endpoint specifications remove some ambiguities of the diagram view when one interval shrinks to a single point that coincides with an endpoint of the other. Such a case is allocated to `equal` when all four endpoints coincide; else to `starts`, `finishes`, `finishedBy` or `startedBy` as appropriate; never to `meets` or `metBy`.

[*Note. The 16 state values can be encoded in four bits. However, if they are then translated into patterns P in a 16-bit word, having one position equal to 1 and the rest zero, one can easily implement interval comparisons by using bit-masks.*

For instance, suppose we make the states s in Table 8’s order correspond to the 16 bits in the word, left-to-right, so $s = \text{bothEmpty}$ maps to $P(s) = 1000000000000000$, $s = \text{firstEmpty}$ maps to $P(s) = 0100000000000000$ and so on. Consider the relation `areDisjoint(a, b)`. This is true if and only if one or both of \mathbf{a} or \mathbf{b} is empty, or \mathbf{a} is “before” \mathbf{b} , or \mathbf{a} is “after” \mathbf{b} . That is, iff the logical “and” of $P(s)$ with the mask `disjointMask` = 1111000000000001 is not identically zero.

This scheme can be efficiently implemented in hardware, see for instance M. Nehmeier, S. Siegel and J. Wolff von Gudenberg [6]. All the required comparisons in this standard can be implemented in this way, as can be, e.g., the “possibly” and “certainly” comparisons of Sun’s interval Fortran. Thus the overlap operation is a primitive from which it is simple to derive all interval comparisons commonly found in the literature.]

TABLE 8. The 16 states of interval overlapping situations for intervals \mathbf{a}, \mathbf{b} .
Notation \forall_a means “for all a in \mathbf{a} ”, and so on. Phrases within a cell are joined by “and”, e.g. **starts** is specified by $(\underline{a} = \underline{b} \wedge \bar{a} < \bar{b})$.

State $\mathbf{a} \oslash \mathbf{b}$ is	Set specification	Endpoint specification	Diagram
States with either interval empty			
bothEmpty	$\mathbf{a} = \emptyset \wedge \mathbf{b} = \emptyset$		
firstEmpty	$\mathbf{a} = \emptyset \wedge \mathbf{b} \neq \emptyset$		
secondEmpty	$\mathbf{a} \neq \emptyset \wedge \mathbf{b} = \emptyset$		
States with both intervals nonempty			
before	$\forall_a \forall_b a < b$	$\bar{a} < \underline{b}$	
meets	$\forall_a \forall_b a \leq b$ $\exists_a \forall_b a < b$ $\exists_a \exists_b a = b$	$\underline{a} < \bar{a}$ $\bar{a} = \underline{b}$ $\underline{b} < \bar{b}$	
overlaps	$\exists_a \forall_b a < b$ $\exists_b \forall_a a < b$ $\exists_a \exists_b b < a$	$a < b$ $\underline{b} < \bar{a}$ $\bar{a} < \bar{b}$	
starts	$\exists_a \forall_b a \leq b$ $\exists_b \forall_a b \leq a$ $\exists_b \forall_a a < b$	$\underline{a} = \underline{b}$ $\bar{a} < \bar{b}$	
containedBy	$\exists_b \forall_a b < a$ $\exists_b \forall_a a < b$	$\underline{b} < \underline{a}$ $\bar{a} < \bar{b}$	
finishes	$\exists_b \forall_a b < a$ $\exists_a \forall_b b \leq a$ $\exists_b \forall_a a \leq b$	$\underline{b} < \underline{a}$ $\bar{a} = \bar{b}$	
equal	$\forall_a \exists_b a = b$ $\forall_b \exists_a b = a$	$\underline{a} = \underline{b}$ $\bar{a} = \bar{b}$	
finishedBy	$\exists_a \forall_b a < b$ $\exists_b \forall_a a \leq b$ $\exists_a \forall_b b \leq a$	$\underline{a} < \underline{b}$ $\bar{b} = \bar{a}$	
contains	$\exists_a \forall_b a < b$ $\exists_a \forall_b b < a$	$\underline{a} < \underline{b}$ $\bar{b} < \bar{a}$	
startedBy	$\exists_b \forall_a b \leq a$ $\exists_a \forall_b a \leq b$ $\exists_a \forall_b b < a$	$\underline{b} = \underline{a}$ $\bar{b} < \bar{a}$	
overlappedBy	$\exists_b \forall_a b < a$ $\exists_a \forall_b b < a$ $\exists_b \exists_a a < b$	$\underline{b} < \underline{a}$ $\underline{a} < \bar{b}$ $\bar{b} < \bar{a}$	
metBy	$\forall_b \forall_a b \leq a$ $\exists_b \exists_a b = a$ $\exists_b \forall_a b < a$	$\underline{b} < \bar{b}$ $\bar{b} = \underline{a}$ $\underline{a} < \bar{a}$	
after	$\forall_b \forall_a b < a$	$\bar{b} < \underline{a}$	

10. The decoration system at Level 1

10.1. Decorations and decorated intervals overview. The decoration system of the set-based flavor conforms to the principles of Clause 8.1. An implementation makes the decoration system available by providing:

- a decorated version of each interval extension of an arithmetic operation, of each interval constructor, and of some other operations;
- various auxiliary functions, e.g., to extract a decorated interval’s interval and decoration parts, and to apply a standard initial decoration to an interval.

The system is specified here at a mathematical level, with the finite-precision aspects throughout Clause 11. Subclauses §10.2, 10.3, 10.4 give the basic concepts. §10.5, 10.6 define how intervals are given an initial decoration, and how decorations are bound to library interval arithmetic operations to give correct propagation through expressions. §10.7, 10.8 are about non-arithmetic operations. §10.9 describes housekeeping operations on decorations, including comparisons, and conversion between a decorated interval and its interval and decoration parts. §10.10 discusses the decoration of user-defined arithmetic operations. The decoration `com` makes it possible to verify, under fairly restrictive conditions, whether a given computation gives the same result in different flavors; §10.11 gives explanatory notes on the `com` decoration. §10.12 defines a restricted decorated arithmetic that suffices for some important applications and is easier to implement efficiently.

In Annex D, §D.2 gives examples of the meaning and use of decorations; and §D.5 contains a rigorous theoretical foundation, including a proof of the Fundamental Theorem of Decorated Interval Arithmetic for this flavor.

10.2. Definitions and basic properties. Formally, a decoration d is a property (that is, a boolean-valued function) $p_d(f, \mathbf{x})$ of pairs (f, \mathbf{x}) , where f is a real-valued function with domain $\text{Dom}(f) \subseteq \mathbb{R}^n$ for some $n \geq 0$ and $\mathbf{x} \in \overline{\mathbb{R}}^n$ is an n -dimensional box, regarded as a subset of \mathbb{R}^n . The notation (f, \mathbf{x}) unless said otherwise denotes such a pair, for arbitrary n , f and \mathbf{x} . Equivalently, d is identified with the set of pairs for which the property holds:

$$d = \{ (f, \mathbf{x}) \mid p_d(f, \mathbf{x}) \text{ is true} \}. \quad (16)$$

The set \mathbb{D} of decorations has five members:

Value	Short description	Property	Definition
<code>com</code>	common	$p_{\text{com}}(f, \mathbf{x})$	\mathbf{x} is a bounded, nonempty subset of $\text{Dom}(f)$; f is continuous at each point of \mathbf{x} ; and the computed interval $f(\mathbf{x})$ is bounded.
<code>dac</code>	defined & continuous	$p_{\text{dac}}(f, \mathbf{x})$	\mathbf{x} is a nonempty subset of $\text{Dom}(f)$, and the restriction of f to \mathbf{x} is continuous;
<code>def</code>	defined	$p_{\text{def}}(f, \mathbf{x})$	\mathbf{x} is a nonempty subset of $\text{Dom}(f)$;
<code>trv</code>	trivial	$p_{\text{trv}}(f, \mathbf{x})$	always true (so gives no information);
<code>ill</code>	ill-formed	$p_{\text{ill}}(f, \mathbf{x})$	Not an Interval; formally $\text{Dom}(f) = \emptyset$, see §10.3.

These are listed according to the propagation order (27), which may also be thought of as a quality-order of (f, \mathbf{x}) pairs—decorations above `trv` are “good” and those below are “bad”.

A **decorated interval** is a pair, written interchangeably as (\mathbf{u}, d) or \mathbf{u}_d , where $\mathbf{u} \in \overline{\mathbb{R}}$ is a real interval and $d \in \mathbb{D}$ is a decoration. (\mathbf{u}, d) may also denote a decorated box $((\mathbf{u}_1, d_1), \dots, (\mathbf{u}_n, d_n))$, where \mathbf{u} and d are the vectors of interval parts \mathbf{u}_i and decoration parts d_i , respectively. The set of decorated intervals is denoted by $\overline{\mathbb{D}\mathbb{I}\mathbb{R}}$, and the set of decorated boxes with n components is denoted by $\overline{\mathbb{D}\mathbb{I}\mathbb{R}}^n$.

When several named intervals are involved, the decorations attached to $\mathbf{u}, \mathbf{v}, \dots$ are often named du, dv, \dots for readability, for instance (\mathbf{u}, du) or \mathbf{u}_{du} , etc.

An interval or decoration may be called a **bare** interval or decoration, to emphasize that it is not a decorated interval.

Treating the decorations as sets as in (16), `trv` is the set of all (f, \mathbf{x}) pairs, and the others are nonempty subsets of `trv`. By design they satisfy the **exclusivity rule**

$$\text{For any two decorations, either one contains the other or they are disjoint.} \quad (18)$$

Namely the definitions (17) give:

$$\text{com} \subset \text{dac} \subset \text{def} \subset \text{trv} \supset \text{ill}, \quad \text{note the change from } \subset \text{ to } \supset; \quad (19)$$

$$\text{com}, \text{dac} \text{ and } \text{def} \text{ are disjoint from } \text{ill}. \quad (20)$$

Property (18) implies that for any (f, \mathbf{x}) there is a unique tightest (in the containment order (19)), decoration such that $p_d(f, \mathbf{x})$ is true, called the **strongest decoration of (f, \mathbf{x})** , or of f over \mathbf{x} , and written $\text{dec}(f, \mathbf{x})$. That is:

$$\text{dec}(f, \mathbf{x}) = d \iff p_d(f, \mathbf{x}) \text{ holds, but } p_e(f, \mathbf{x}) \text{ fails for all } e \subset d. \quad (21)$$

[*Note. Like the exact range $\text{Rge}(f | \mathbf{x})$, the strongest decoration is theoretically well-defined, but its value for a particular f and \mathbf{x} may be impractically expensive to compute, or even undecidable.*]

10.3. The ill-formed interval. The `ill` decoration results from invalid constructions, and propagates unconditionally through arithmetic expressions. Namely, if a constructor call does not return a valid decorated interval, it returns an ill-formed one (i.e., decorated with `ill`); and the decorated interval result of a library arithmetic operation is ill-formed, if and only if one of its inputs is ill-formed. Formally, `ill` may be identified with the property $\text{Dom}(f) = \emptyset$ of (f, \mathbf{x}) pairs, see Clause E.2.

An ill-formed decorated interval is also called **NaI, Not an Interval**. Except as described in the next paragraph, an implementation shall behave as if there is only one NaI, whose interval part is indeterminate. However, the `intervalPart()` operation must return a bare interval for any decorated interval input, and for NaI this shall be `Empty`; thus NaI may be viewed as being \emptyset_{ill} .

An exception is that other information may be stored in an NaI in an implementation-defined way (like the payload of a 754 floating point NaN), and functions may be provided for a user to set and read this for diagnostic purposes. An implementation may also provide means for an exception to be raised when an NaI is produced.

[*Example. The constructor call `nums2interval(2, 1)` is invalid in this flavor, so its decorated version returns `NaI`.*]

10.4. Permitted combinations. A decorated interval \mathbf{y}_{dy} shall always be such that $\mathbf{y} \supseteq \text{Rge}(f | \mathbf{x})$ and $p_{dy}(f, \mathbf{x})$ holds, for some (f, \mathbf{x}) as in §10.2—informally, it must tell the truth about some conceivable evaluation of a function over a box. If $dy = \text{dac}$ or `def` then by definition \mathbf{x} is nonempty, and f is everywhere defined on it, so that $\text{Rge}(f | \mathbf{x})$ is nonempty, implying \mathbf{y} is nonempty. Hence the decorated intervals \emptyset_{dac} and \emptyset_{def} , and \emptyset_{com} if `com` is provided, are contradictory: implementations shall not produce them.

No other combinations are essentially forbidden.

10.5. Initial decoration. Correct use of decorations when evaluating an expression has two parts: correctly initialize the input intervals; and evaluate using decorated interval extensions of library operations. To provide correct initialization, the function `newDec()` is provided. For a single bare interval \mathbf{x} , `newDec()` decorates it with $\text{dec}(\text{Id}, \mathbf{x})$, the strongest decoration $d\mathbf{x}$ that makes $p_{d\mathbf{x}}(\text{Id}, \mathbf{x})$ true, where `Id` is the identity function $\text{Id}(x) = x$ for real x . That is,

$$\text{newDec}(\mathbf{x}) = \mathbf{x}_d \quad \text{where} \quad d = \text{dec}(\text{Id}, \mathbf{x}) = \begin{cases} \text{com} & \text{if } \mathbf{x} \text{ is nonempty and bounded,} \\ \text{dac} & \text{if } \mathbf{x} \text{ is unbounded,} \\ \text{trv} & \text{if } \mathbf{x} \text{ is empty.} \end{cases} \quad (22)$$

[*Note. The following features are used in examples in this document. It is language-defined whether an implementation provides them.*

– For a decorated interval $\mathbf{x}_{d\mathbf{x}}$, `newDec()` discards the decoration and acts on the interval part,

$$\text{newDec}(\mathbf{x}_{d\mathbf{x}}) = \text{newDec}(\mathbf{x}). \quad (23)$$

– For a vector of n bare or decorated intervals, `newDec()` acts componentwise to give a vector of n decorated intervals.

– A bare interval constant, in a context that expects a decorated interval, is implicitly promoted to a decorated interval by applying the `newDec` function. For example, `[1, 2]`, `[1, +∞]` and `∅` are promoted to $[1, 2]_{\text{com}}$, $[1, +\infty]_{\text{dac}}$ and \emptyset_{trv} respectively.

]

10.6. Decorations and arithmetic operations. Given a scalar point function φ of k variables, a **decorated interval extension** of φ —denoted here by the same name φ —adds a decoration component to a bare interval extension of φ . It has the form $\mathbf{w}_{dw} = \varphi(\mathbf{v}_{dv})$, where $\mathbf{v}_{dv} = (\mathbf{v}, dv)$ is a k -component decorated box $((\mathbf{v}_1, dv_1), \dots, (\mathbf{v}_k, dv_k))$. By the definition of a bare interval extension, the interval part \mathbf{w} depends only on the input intervals \mathbf{v} ; the decoration part dw generally depends on both \mathbf{v} and dv . In this context, NaI is regarded as being \emptyset_{i11} .

The definition of a bare interval extension implies

$$\mathbf{w} \supseteq \text{Rge}(\varphi | \mathbf{v}), \quad (\text{enclosure}). \quad (24)$$

The decorated interval extension of φ determines a dv_0 such that

$$p_{dv_0}(\varphi, \mathbf{v}) \text{ holds,} \quad (\text{a “local decoration”}). \quad (25)$$

It then evaluates the output decoration dw by

$$dw = \min\{dv_0, dv_1, \dots, dv_k\}, \quad (\text{the “min-rule”}), \quad (26)$$

where the minimum is taken with respect to the **propagation order**:

$$\text{com} > \text{dac} > \text{def} > \text{trv} > \text{i11}. \quad (27)$$

[Notes.

1. Because NaI is treated as \emptyset_{i11} , this definition implies (without treating it as a special case) that $\varphi(\mathbf{v}_{dv})$ is NaI if, and only if, some component of \mathbf{v}_{dv} is NaI.
2. Let $f(z_1, \dots, z_n)$ be an expression defining a real point function $f(x_1, \dots, x_n)$. Then decorated interval evaluation of f on a correctly initialized input decorated box \mathbf{x}_{dx} gives a decorated interval \mathbf{y}_{dy} such that not only, by the Fundamental Theorem of Interval Arithmetic, one has

$$\mathbf{y} \supseteq \text{Rge}(f | \mathbf{x}) \quad (28)$$

but also

$$p_{dy}(f, \mathbf{x}) \text{ holds.} \quad (29)$$

For instance, if the computed dy equals `def` then f is proven to be everywhere defined on the box \mathbf{x} . This is the **Fundamental Theorem of Decorated Interval Arithmetic (FTDIA)**. The rules for initializing and propagating decorations are key to its validity. They are justified, and a formal statement and proof of the FTDIA given, in Annex D.

Briefly, (22) gives the correct result for the simplest expression of all, where f is the identity $f(x) = x$, which contains no arithmetic operations. The decorations are designed so that the min-rule (26) embodies basic facts of set theory and analysis, such as “If each of a set of functions is everywhere defined [resp. continuous] on its input, their composition has the same property” and “If any of a set of functions is nowhere defined on its input, their composition has the same property”. It causes correct propagation of decorations through each arithmetic operation, and hence through a whole expression.

3. In the same way as the enclosure requirement (24) is compatible with many bare interval extensions, typically coming from different interval types at Level 2, so there may be several dv_0 satisfying the local decoration requirement (25). The ideal choice is the strongest decoration d such that $p_d(\varphi, \mathbf{v})$ holds, that is to take

$$dv_0 = \text{dec}(\varphi, \mathbf{v}). \quad (30)$$

This is easily computable in finite precision for the arithmetic operations in §9.6, 9.7—see the tables in Annex D, §D.1. However, functions may be added to the library in future for which (30) is impractical to compute for some arguments \mathbf{v} . Hence the weaker requirement (25) is made.

]

10.7. Decorations and non-arithmetic operations.

The following give interval results but are not interval extensions of point functions:

- the reverse-mode operations of §9.6.6;
- the cancellative operations `cancelPlus(x, y)` and `cancelMinus(x, y)` of §9.6.7;
- The set-oriented operations `intersection(x, y)` and `convexHull(x, y)` of §9.6.8.

The decorated interval version of each such operation distinguishes NaI inputs but otherwise returns a decoration indicating no information, as follows. If any input is NaI, the result is NaI. Otherwise the result is obtained by applying the corresponding operation to the interval parts of the inputs, and decorating its result with `trv`.

The user is responsible for applying a more informative decoration to the result via `setDec()`, where this can be deduced in a given context. Other versions of the operations may be provided in a language- or implementation-defined way, that apply such a context-adapted decoration.

The following operations give non-interval results:

- the numeric functions of §9.6.10;
- the boolean-valued functions of §9.6.11;
- the overlap function of §9.7.3.

For each such operation, if any input is NaI the result is undefined at Level 1. Otherwise, the operation acts on decorated intervals by discarding the decoration and applying the corresponding bare interval operation.

10.8. Boolean functions of decorated intervals. The equality comparison `equal`, or `=`, shall be provided for decorated intervals. NaI shall compare unequal to any decorated interval, including itself. Other input combinations shall compare equal if and only if the interval parts are equal and the decoration parts are equal.

The inequality comparison `notEqual`, or `≠`, shall be provided. It is the logical negation of `=` (so `NaI ≠ NaI` is true).

The unary function `isNaI` shall be provided. It is true if and only if its input is NaI.

10.9. Operations on decorations.

Given a decorated interval x_{dx} , the operations `intervalPart(x_{dx})` and `decorationPart(x_{dx})` shall be provided, which return x and dx respectively. NaI is deemed to equal \emptyset_{i11} , so that `intervalPart(NaI)` returns \emptyset and `decorationPart(NaI)` returns `i11`.

Given an interval x and a decoration dx , the operation `setDec(x, dx)` returns the decorated interval x_{dx} . If this would produce one of the forbidden combinations—that is, $x = \emptyset$ and dx is one of `def`, `dac` or `com`—then NaI is returned instead. `setDec($x, i11$)` returns NaI for any interval x , whether empty or not.

For decorations, comparison operations for equality `=` and its negation `≠` shall be provided, as well as comparisons `>`, `<`, `≥`, `≤` with respect to the propagation order (27).

[*Note. Careless use of the `setDec` function can negate the aims of the decoration system and lead to false conclusions that violate the FTDIA. It is provided for expert users, who may need it, e.g., to decorate the output of functions whose definition involves the intersection and convexHull operations.*]

10.10. User-supplied functions. A user may define a decorated interval extension of some point function, as defined in §10.6, to be used within expressions as if it were a library operation.

[*Examples.*

(i) *In an application, an interval extension of the function*

$$f(x) = x + 1/x$$

was required. Evaluated as written, it gives unnecessarily pessimistic enclosures: e.g., with $x = [\frac{1}{2}, 2]$, one obtains

$$f(\mathbf{x}) = [\frac{1}{2}, 2] + 1/[\frac{1}{2}, 2] = [\frac{1}{2}, 2] + [\frac{1}{2}, 2] = [1, 4],$$

much wider than $\text{Rge}(f | \mathbf{x}) = [2, 2\frac{1}{2}]$.

Thus it is useful to code a tight interval extension by special methods, e.g. monotonicity arguments, and provide this as a new library function. Suppose this has been done. To convert it to a decorated interval extension just entails adding code to provide a local decoration and combine this with the input decoration by the min-rule (26). In this case it is straightforward to compute the strongest local decoration $d = \text{dec}(f, \mathbf{x})$, as follows.


$$d = \begin{cases} \text{com} & \text{if } 0 \notin \mathbf{x} \text{ and } \mathbf{x} \text{ is nonempty and bounded,} \\ \text{dac} & \text{if } 0 \notin \mathbf{x} \text{ and } \mathbf{x} \text{ is unbounded,} \\ \text{trv} & \text{if } \mathbf{x} = \emptyset \text{ or } 0 \in \mathbf{x} \neq [0, 0]. \end{cases}$$

(ii)

The next example shows how an expert may manipulate decorations explicitly to give a function, defined piecewise by different formulas in different regions of its domain, the best possible decoration. Suppose that

$$f(x) = \begin{cases} f_1(x) := \sqrt{x^2 - 4} & \text{if } |x| > 2, \\ f_2(x) := -\sqrt{4 - x^2} & \text{otherwise,} \end{cases}$$

where $:=$ means “defined as”, see the diagram.



The function consists of three pieces, on regions $x \leq -2$, $-2 \leq x \leq 2$ and $x \geq 2$, that join continuously at region boundaries, but the standard gives no way to determine this continuity, at run time or otherwise. For instance, if f is implemented by the case function, the continuity information is lost when evaluating it on, say, $x = [1, 3]$, where both branches contribute for different values of $x \in x$.

However, a user-defined decorated interval function as defined below provides the best possible decorations.

```
function  $y_{dy} = f(x_{dx})$ 
   $u = f_1(x \cap [-\infty, -2])$ 
   $v = f_2(x \cap [-2, 2])$ 
   $w = f_1(x \cap [2, +\infty])$ 
   $y = \text{convexHull}(\text{convexHull}(u, v), w)$ 
   $dy = dx$ 
```

The user’s knowledge that f is everywhere defined and continuous is expressed by the statement $dy = dx$, propagating the input decoration unchanged. f , thus defined, can safely be used within a larger decorated interval evaluation.

]

10.11. Notes on the com decoration.

[Notes.

- The force of `com` is the Level 2 property that the computed interval $f(x)$ is bounded. Equivalently, overflow did not occur, where overflow has the generalized meaning that a finite-precision operation could not enclose a mathematically bounded result in a bounded interval of the required output type. Briefly, for a single operation, “com is dac plus bounded inputs and no overflow”.

Thus the result of interval-evaluating an arithmetic expression in finite precision is decorated `com` if and only if the evaluation is common at Level 2, meaning: each input that affects the result is nonempty and bounded, and each individual operation that affects the result is everywhere defined and continuous on its inputs and does not overflow.

- A tempting alternative is to make `com` record whether the evaluation is common at Level 1, meaning that all the relevant intervals are mathematically bounded, even if overflow occurred in finite precision. E.g., one might drop the “bounded inputs” requirement and require “mathematically bounded” instead of “actually bounded” on the output of an operation.

However, the `dac` decoration already provides such information and the suggested change gives nothing extra. Namely, if the inputs x to $f(x)$ are bounded, and the output decoration is `dac`, it follows, from the fact that a continuous function on a compact set is bounded, that the point function f is mathematically bounded on x , and all its individual operations are mathematically bounded on their inputs even if overflow may have occurred in finite precision.

For example consider $f(x) = 1/(2x)$ evaluated at $x = [1, M]$ using an inf-sup type where M is the largest representable real. This gives

$$\mathbf{y}_{dy} = f(\mathbf{x}_{com}) = 1/(2 * [1, M]_{com}) = 1/[2, +\infty]_{dac} = [0, \frac{1}{2}]_{dac}.$$

Despite the overflow, one can deduce from the final dac that the result of the multiplication was mathematically bounded.

This may be of limited use: consider $g(x) = 1/f(x) = 1/(1/(2x))$, evaluated at the same $x = [1, M]$ giving z_{dz} . The standard has no way to record that the lower bound of \mathbf{y} is mathematically positive, i.e., $1/(2M)$. Thus the Level 2 result is $z_{dz} = [2, +\infty]_{trv}$, compare $[2, 2M]_{com}$ at Level 1.

]

10.12. Compressed arithmetic with a threshold (optional).

10.12.1. *Motivation.* The **compressed decorated interval arithmetic** (compressed arithmetic for short) described here lets experienced users obtain more efficient execution in applications where the use of decorations is limited to the context described below. An implementation need not provide it; if it does so, the behavior described in this subclause is required.

Each Level 2 instance of compressed arithmetic is based on a supported Level 2 bare interval type \mathbb{T} , but is a distinct “compressed type”, with its own datums and library of operations.

The context is that of evaluating an arithmetic expression, where the use made of a decorated interval evaluation $\mathbf{y}_{dy} = f(\mathbf{x}_{dx})$ depends on a check of the result decoration dy against an application-dependent **exception threshold** τ , where $\tau \geq \text{trv}$ in the propagation order (27):

$dy \geq \tau$ represents normal computation. The decoration is not used, but one exploits the range enclosure given by the interval part and the knowledge that dy remained $\geq \tau$.

$dy < \tau$ declares an exception to have occurred. The interval part is not used, but one exploits the information given by the decoration.

10.12.2. *Compressed interval types.* For such uses, one needs to record an interval’s value, or its decoration, but never both at once. The **compressed type** of threshold τ , **associated with** \mathbb{T} , is the type each of whose datums is either a bare \mathbb{T} -interval or a bare decoration less than τ . It is denoted \mathbb{T}_τ . Two such types are the same if and only if they have the same \mathbb{T} and the same τ . A \mathbb{T}_τ datum can be any \mathbb{T} datum or any decoration except that:

- Only decorations $< \tau$ occur; in particular **com** is never used.
- The empty interval \emptyset is replaced by—equivalently, is regarded by the implementation as being—a new decoration **emp** added to the table in (17), whose defining property is

Value	Short description	Property	Definition
emp	empty	$p_{emp}(f, \mathbf{x})$	$\mathbf{x} \cap \text{Dom}(f)$ is empty;

emp lies between **trv** and **ill** in the containment order (19) and the propagation order (27):

$$\text{com} \subset \text{dac} \subset \text{def} \subset \text{trv} \supset \text{emp} \supset \text{ill}, \quad (32)$$

$$\text{com} > \text{dac} > \text{def} > \text{trv} > \text{emp} > \text{ill}.$$

Since $\tau \geq \text{trv}$ it is always true that **emp** $< \tau$, which means that as soon as an empty result is produced while evaluating an expression, the $dy < \tau$ case has occurred.

[*Note. The reason for treating \emptyset as a decoration $< \tau$ is that obtaining an empty result (e.g., by doing something like $\sqrt{[-2, -1]}$ while evaluating a function) is one of the “exceptions” that compressed interval computation should detect.*]

The only way to use compressed arithmetic with a threshold τ is to construct \mathbb{T}_τ datums. Conversion between compressed types, say from a \mathbb{T}_τ -interval to a $\mathbb{T}'_{\tau'}$ -interval, shall be equivalent to converting first to a normal decorated interval by `normalInterval()`, then between decorated interval types if $\mathbb{T} \neq \mathbb{T}'$, and finally to the output type by `τ' -compressedInterval()`.

[*Note. Since, for any practical interval type \mathbb{T} , a decoration fits into less space than an interval, one can implement arithmetic on compressed interval datums that take up the same space as a bare interval of that type. For instance if \mathbb{T} is the IEEE754 binary64 inf-sup type, a compressed interval uses 16 bytes, the same as a bare \mathbb{T} -interval; a full decorated \mathbb{T} -interval needs at least 17 bytes.*]

Because compressed intervals must behave exactly like bare intervals as long as one does not fall below the threshold, and take up the same space, there is no room to encode τ as part of the interval’s value.]

10.12.3. *Operations.* The enquiry function `isInterval(\mathbf{x})` returns true if the compressed interval \mathbf{x} is an interval, false if it is a decoration.

The constructor `τ -compressedInterval()` is provided for each threshold value τ . The result of `τ -compressedInterval(\mathbf{X})`, where $\mathbf{X} = (\mathbf{x}, dx)$ is a decorated \mathbb{T} -interval, is a \mathbb{T}_τ -interval as follows:

```
if  $dx \geq \tau$ , return the  $\mathbb{T}_\tau$ -interval with value  $\mathbf{x}$ 
else return the  $\mathbb{T}_\tau$ -interval with value  $dx$ .
```

`τ -compressedInterval(\mathbf{x})` for a bare interval \mathbf{x} is equivalent to `τ -compressedInterval(newDec(\mathbf{x}))`.

The function `normalInterval(\mathbf{x})` converts a \mathbb{T}_τ -interval to a decorated interval of the parent type, as follows:

```
if  $\mathbf{x}$  is an interval, return  $(\mathbf{x}, \tau)$ .
if  $\mathbf{x}$  is a decoration  $d$ 
  if  $d$  is ill or emp, return (Empty,  $d$ )
  else return (Entire,  $d$ ).
```

Arithmetic operations on compressed intervals shall follow *worst case semantics* rules that treat a decoration in `{trv, def, dac}` as representing a set of decorated intervals, and are necessary if the fundamental theorem is to remain valid. Namely, inputs to each operation behave as follows:

- Operations purely on bare intervals are performed as if each \mathbf{x} is the decorated interval \mathbf{x}_τ , resulting in a decorated interval \mathbf{y}_{dy} that is then converted back into a compressed interval. If $dy < \tau$, the result is the bare decoration dy , otherwise the bare interval \mathbf{y} .
- For operations with at least one bare decoration input, the result is always a bare decoration. A bare interval input is treated as in the previous item. A decoration d in `{emp, ill}` is treated as \emptyset_d . A decoration d in `{trv, def, dac}` is treated (conceptually) as \mathbf{x}_d with an arbitrary nonempty interval \mathbf{x} . The decoration com cannot occur. Performing the resulting decorated interval operation on all such possible inputs leads to a set of all possible results \mathbf{y}_{dy} . The tightest decoration (in the containment order (32)) enclosing all resulting dy is returned.

As a result each operation returns an actual or implied decoration compatible with its input, so that in an extended evaluation, the final decoration using compressed arithmetic is never stronger than that produced by full decorated interval arithmetic.

⚠ (JDP, 2013.) I conjecture that the following may be an equivalent specification, but have not checked. ■

1. Each compressed interval argument is converted to a decorated interval by `normalInterval`;
 2. the corresponding operation of the parent decorated interval type is performed;
 3. the result, if an interval, is converted back to a compressed interval by `τ -compressedInterval`.
- (33)

[*Example. Assuming $\tau > \text{def}$,*

- The division `def/[1, 2]` becomes `$\mathbf{x}_{\text{def}}/[1, 2]_\tau$` with arbitrary nonempty interval \mathbf{x} . The result is always decorated `def`, so returns `def`.
- But `[1, 2]/def` becomes `$[1, 2]_\tau/\mathbf{x}_{\text{def}}$` with arbitrary nonempty interval \mathbf{x} . The result can be decorated `def`, `trv` or `emp`, so returns the tightest decoration containing these, namely `trv`.

Since there are only a few decorations, one can prepare complete operation tables according to this rule, and only these tables need to be implemented. In Annex D, sample tables and worked examples are in Clause D.3 and a proof of correctness of the compressed arithmetic system is in Clause D.4.

If compressed arithmetic is implemented, it shall provide versions of all the required operations of §9.6, and it should provide the recommended operations of §9.7.

11. Level 2 description

11.1. Level 2 introduction. Entities and operations at Level 2 are said to have **finite precision**. From them, implementable interval algorithms may be constructed. Level 2 entities are called **datums**¹Since the standard deals with numeric functions of intervals (such as the midpoint) and interval functions of numbers (such as the construction of an interval from its lower and upper bounds), this clause involves both numeric and interval datums, as well as decoration, text string and boolean datums.

Following 754 terminology, numeric (usually but not necessarily floating point) datums are organized into **formats**. Interval datums are organized into **types**. Each format or type is a finite set of datums, with associated operations. If \mathbb{F} denotes a format, an \mathbb{F} -*number* means a member of \mathbb{F} , possibly infinite but not the “not a number” value NaN; to allow it to be NaN, it is called an \mathbb{F} -*datum*. If \mathbb{T} denotes a bare or decorated interval type, a \mathbb{T} -**interval** means a member of \mathbb{T} , possibly empty but (in the decorated case) not the “not an interval” value NaN; to allow it to be NaN, it is called a \mathbb{T} -*datum*. A \mathbb{T} -**box** means a box with \mathbb{T} -datum components.

The standard defines three kinds of interval type:

- **Bare interval types**, see §11.6, are named finite sets of (mathematical, Level 1) intervals.
- **Decorated interval types**, also see §11.6, are named finite sets of decorated intervals.
- **Compressed interval types** (optional) are named finite sets of compressed intervals. They are described in §10.12 and Level 2 makes no further requirements on them.

For each bare interval type there shall be a corresponding *derived* decorated interval type, and each decorated interval type shall be derived from a bare interval type, see §11.6.1. An implementation shall support at least one bare interval type. If 754-conforming, it shall support the inf-sup type, see §11.6.3, of at least one of the five basic formats of 754§3.3. Beyond this, which types are supported is language- or implementation-defined.

It is language-defined whether the format or type of a datum can be determined at run time.

A Level 2 operation is a finite-precision approximation to the corresponding Level 1 operation. Whereas a Level 1 operation, for some inputs, may have no value (be undefined in the mathematical sense), each Level 2 operation shall return a value for arbitrary inputs.

To describe the required functionality, the standard treats each Level 1 operation as having a number of Level 2 **versions** in which each interval input or output is given a specific interval type, and each numeric input or output is given a specific numeric format.

[*Note. An implementation might provide the functionality via differently named operations for each version, or by overloading a single operation name, or by a single operation that accepts all the required type/format combinations, or in other ways.*]

For example, let \mathbb{T}_1 and \mathbb{T}_2 be two supported interval types. The standard requires a version of addition $z = x + y$ where each of x, y, z has type \mathbb{T}_1 , and another where each of them has type \mathbb{T}_2 . An implementation might provide this functionality by having two separate operations; another might have a single operation that takes inputs of types \mathbb{T}_1 and \mathbb{T}_2 in any combination, with some rule to determine the type of the output z ; etc.]

The term \mathbb{T} -**version** of a Level 1 operation denotes one in which all interval inputs or outputs are \mathbb{T} -datums. For bare interval types this includes the following:

- (a) A \mathbb{T} -interval extension (§11.9) of one of the required or recommended arithmetic operations of §9.6, 9.7.
- (b) A set operation, such as intersection and convex hull of \mathbb{T} -intervals, returning a \mathbb{T} -interval.
- (c) A function such as the midpoint, whose input is a \mathbb{T} -interval and output is numeric.
- (d) A constructor, whose input is numeric or text and output is a \mathbb{T} -datum.
- (e) The \mathbb{T} -interval hull, regarded as a conversion operation, see §11.8.3.

Generically these comprise the **operations of the type \mathbb{T}** , for the implementation.

11.2. Naming conventions for operations. An operation is generally given a name that suits the context. For example, the addition of two interval datums x, y may be written in generic algebra notation $x + y$; or with a generic text name `addition(x, y)`; or giving full type information such as `infsup-decimal64-addition(x, y)`.

¹Not “data”, whose common meaning could cause confusion.

It may also be written as \mathbb{T} -**addition**(x, y) to show it is an operation of a particular but unspecified type \mathbb{T} . In the context of 754-conforming types it may be written as *typeOf*-**addition**(x, y) where *typeOf* has a similar meaning to 754's *formatOf*—namely, each input can be of any supported 754-conforming type of the same radix as the output type.

In a specific language or programming environment, the names used for types may differ from those used in this document.

11.3. 754-conformance. The standard defines the notion of 754-conformance, whose stronger requirements improve accuracy and programming convenience. In this context a part of an implementation means a subset of the set of supported Level 2 interval types.

A **754-conforming type** is an inf-sup type derived from a 754 floating point format—one of the five basic types or an extended precision or extendable precision format—that meets the general requirements for conformance and whose operations meet the accuracy requirements in Table 9.

A **754-conforming part** of an implementation is a subset of the set of supported 754-conforming types, that meets the requirements for mixed-type arithmetic in the next paragraphs. A **754-conforming implementation** is one where *all* supported types are 754-conforming, and the whole set of supported types meets the requirements for mixed-type arithmetic.

11.3.1. *754-conforming mixed-type arithmetic.* The 754 standard requires a conforming floating point system to provide mixed-format “*formatOf*” operations. That is, the output format is specified and the inputs may be of any format of the same radix as the output. The result is computed as if using the exact inputs and rounded to the required accuracy on output. This eliminates the problem of double rounding in mixed-format work, which otherwise can cause significant growth of errors.

A 754-conforming part of an implementation shall provide (e.g., by exploiting the *formatOf* feature at Level 3) corresponding mixed-type “*typeOf*” operations. That is, the output type is specified and the inputs may be of any type of the same radix as the output. The result shall be computed as if using the exact inputs and shall meet the accuracy requirements for each operation, specified in Table 9.

These requirements shall apply to all Level 2 operations with interval operands, without explicit mention.

11.4. Tagging, and the meaning of equality at Level 2. A Level 2 format or type is an abstraction of a particular way to represent numbers or intervals—e.g., “IEEE 64 bit binary” for numbers—focusing on the Level 1 entities represented, and hiding the Level 3 method by which it is done.

However a datum is more than just the Level 1 value: for instance the number 3.75 represented in 64 bit binary is a different datum from the same number represented in 64 bit decimal.

This is achieved by formally regarding each datum as a pair:

$$\begin{aligned} \text{number datum} &= (\text{Level 1 number, format name}), \\ \text{bare interval datum} &= (\text{Level 1 interval, type name}), \end{aligned}$$

where the name is some symbol that uniquely identifies the format or type. Since a decorated interval combines a bare interval and a decoration it thus becomes a triple at Level 2:

$$\text{decorated interval datum} = (\text{Level 1 interval, type name, decoration}).$$

The Level 1 value is said to be **tagged** by the name. It follows that distinct formats or types are disjoint sets. By convention, such names are omitted from datums except when clarity requires.

[*Example. Level 2 interval addition within a type named t is normally written $z = x + y$, though the full correct form is $(z, t) = (x, t) + (y, t)$. The full form might be used, for instance, to indicate that mixed-type addition is forbidden between types s and t but allowed between types s and u . Namely, one can say that $(x, s) + (y, t)$ is undefined, but $(x, s) + (y, u)$ is defined.*]

Two bare interval datums are defined to be equal if and only if they are the same Level 1 interval tagged by the same or equivalent name. In particular two datums of the same bare interval type are equal if and only if, dropping the tag, they are the same Level 1 interval.

[*Note. Two bare interval datums a, b of the same type are equal if and only if $\text{equal}(a, b)$ returns true. However, in the context of 754-conforming types, intervals of different types can compare equal.*]

E.g., let \mathbf{a} and \mathbf{b} be the *inf-sup* `binary32` and *inf-sup* `binary64` datums, respectively, of an interval that is represented exactly in both types, such as `[1, 2]`. Then `equal(a, b)` is true, but \mathbf{a} and \mathbf{b} are not equal in the sense of this subclause.

▲ (a) Are two equality functions needed to handle this?]

11.5. Number formats. In view of §11.4, a **number format**, or just **format**, is the set of all pairs (x, f) where x belongs to a set \mathbb{F} , and f is a name for the format.

\mathbb{F} comprises a finite subset of the extended reals \mathbb{R} , together with a value NaN. A **numeric** member of \mathbb{F} is one that is not NaN. \mathbb{F} shall contain zero, $-\infty$ and $+\infty$, and shall be symmetric: if a numeric x is in \mathbb{F} , so is $-x$.

Following the convention of omitting names, the format is normally identified with the set \mathbb{F} , and one may say a number format is a set of datums comprising NaN together with a finite, symmetric, set \mathbb{F} of extended reals that contains 0 and $\pm\infty$.

In this document the five basic formats of 754§3.3 are named `binary32`, `binary64`, `binary128`, `decimal64`, `decimal128`. Abbreviated names such as `b64` instead of `binary64` are sometimes used, and refer to the same format.

A floating-point format in the 754 sense, such as `binary64`, is identified with the number format for which \mathbb{F} is the set of extended-real numbers exactly representable in that format.

[Note. At Level 2 each format has only one zero datum and one NaN datum, but these may correspond to several values at Level 3, e.g. IEEE 754 has zeros -0 , $+0$, while NaNs are distinguished by sign bit and payload, and by being quiet or signaling. The fact that a floating point operation at Level 2 can behave differently on nominally equal datums is beyond the remit of this standard. E.g., -0 and $+0$ compare equal but $1/-0$ and $1/+0$ do not.]

A number format \mathbb{F} is said to be **compatible** with an interval type \mathbb{T} , if each non-empty \mathbb{T} -interval contains at least one finite \mathbb{F} -number.

Each format \mathbb{F} shall have an associated **rounding function** that maps any extended real number x to an element of \mathbb{F} . A constraint can be given on the rounding direction, e.g. one of the rounding direction attributes in 754§4.3. The rounded result is an element of \mathbb{F} that is closest to x subject to this constraint, with an implementation-defined rule for the distance to an infinity, and for the method of tie-breaking when more than one member of \mathbb{F} has the “closest” property.

11.6. Bare and decorated interval types.

11.6.1. *Definition.* In view of §11.4, a **bare interval type** is formally the set of all pairs (\mathbf{x}, t) where \mathbf{x} belongs to a finite subset \mathbb{T} of the mathematical intervals $\overline{\mathbb{R}}$ that contains Empty and Entire, and t is a name for the type. The **decorated interval type derived** from \mathbb{T} is formally the set of triples (\mathbf{x}, t, d) where (\mathbf{x}, t) is a \mathbb{T} -interval, and $d \in \mathbb{D}$ is a decoration that follows the rule for permitted combinations (\mathbf{x}, d) in §10.4.

Following the convention of omitting names, the type is normally regarded as being the set \mathbb{T} , and one may say a bare interval type is an arbitrary finite set \mathbb{T} of intervals that contains Empty and Entire. The derived decorated type is then regarded as a set of pairs (\mathbf{x}, d) , equivalently \mathbf{x}_d , where $\mathbf{x} \in \mathbb{T}$ and $d \in \mathbb{D}$.

Each decorated interval type shall contain a “Not an Interval” datum NaI, identified with $(\emptyset, \text{i11})$. It shall appear to be unique at Level 2, but non-Level-2 operations may be provided to set and get a payload in an NaI for diagnostic purposes, in an implementation-defined way (see §10.3).

[Example. To illustrate the flexibility allowed in defining types, let S_1 and S_2 be the sets of *inf-sup* intervals using 754 single (`binary32`) and double (`binary64`) precision respectively. That is, a member of S_1 [respectively S_2] is either empty, or an interval whose bounds are exactly representable in `binary32` [respectively `binary64`].

An implementation usually would define these as different bare interval types, by tagging members of S_1 by one type name t_1 and members of S_2 by another name t_2 —represented at Level 3 by a pair of `binary32` or of `binary64` numbers respectively. However, it might treat them as one type, with the representation by a pair of `binary32`'s being a space-saving alternative to the pair of `binary64`'s, to be used, say, for some large arrays. The resulting Level 1 intervals are exactly the same those of *inf-sup* `binary64`, so this is just another way to store the latter type; but an implementation would give it a different name, to reflect the different storage and hence different operations at the code level.]

11.6.2. *Initial decoration of an interval.* The initialization operation `newDec()` shall be provided for all supported interval types. Let \mathbb{T} be a bare type and \mathbb{DT} the derived decorated type. The input to `newDec()` is a \mathbb{T} -interval or a \mathbb{DT} -interval and its output is a \mathbb{DT} -interval, as specified in §10.5.

11.6.3. *Inf-sup and mid-rad types.*

The **inf-sup type** derived from a given number format \mathbb{F} (the type **inf-sup** \mathbb{F} , e.g., “inf-sup binary64”) is the bare interval type \mathbb{T} comprising all intervals whose endpoints are in \mathbb{F} , together with Empty. \mathbb{F} is termed the **parent format** of \mathbb{T} . When \mathbb{F} is a 754 format, the **radix** of \mathbb{T} means the radix of \mathbb{F} . Note that Entire is in \mathbb{T} because $\pm\infty \in \mathbb{F}$ by the definition of a number format, so \mathbb{T} satisfies the requirements for a bare interval type given in §11.6.1.

Mid-rad types are not specified by this standard but are useful for examples. A mid-rad bare interval type is taken to be one whose nonempty bounded intervals comprise all intervals of the form $[m - r, m + r]$, where m is in some number format \mathbb{F} , and r is in a possibly different number format \mathbb{F}' , with m, r finite and $r \geq 0$. From the definition in §11.6.1 such a type, to be conforming, must contain Empty and Entire, so at Level 3 it must have representations of these; it might also have representations of semi-bounded intervals.

11.7. Multi-precision interval types. Multi-precision floating point systems—extendable precision in 754 terminology—generally provide an (at least conceptually) infinite sequence of levels of precision, where there is a finite set \mathbb{F}_n of numbers representable at the n th level ($n = 1, 2, 3, \dots$), and $\mathbb{F}_1 \subset \mathbb{F}_2 \subset \mathbb{F}_3 \dots$. These are typically used to define a corresponding infinite sequence of interval types \mathbb{T}_n with $\mathbb{T}_1 \subset \mathbb{T}_2 \subset \mathbb{T}_3 \dots$.

[*Example. For multi-precision systems that define a nonempty \mathbb{T}_n -interval to be one whose endpoints are \mathbb{F}_n -numbers, each \mathbb{T}_n is an inf-sup type with a unique interval hull operation—explicit, in the sense of §11.8.]*

A conforming implementation must define such \mathbb{T}_n as a *parameterized sequence* of interval types. It cannot take the union over n of the sets \mathbb{T}_n as a *single* type, because this infinite set has no interval hull operation: there is generally no tightest member of it enclosing a given set of real numbers. This constrains the design of conforming multi-precision interval systems.

11.8. Explicit and implicit types, and Level 2 hull operation.

11.8.1. *Hull in one dimension.* For each bare interval type \mathbb{T} there shall be defined an **interval hull operation**

$$\mathbf{y} = \text{hull}_{\mathbb{T}}(\mathbf{s}),$$

also called the \mathbb{T} -hull, which is part of \mathbb{T} 's mathematical definition. For an implicit type, see below, the implementation's documentation shall specify the \mathbb{T} -hull, e.g., by an algorithm. For an explicit type, it is uniquely determined and need not be separately specified.

The \mathbb{T} -hull maps an arbitrary set of reals, \mathbf{s} , to a minimal \mathbb{T} -interval \mathbf{y} enclosing \mathbf{s} . Minimal is in the sense that

$$\mathbf{s} \subseteq \mathbf{y}, \text{ and for any other } \mathbb{T}\text{-interval } \mathbf{z}, \text{ if } \mathbf{s} \subseteq \mathbf{z} \subseteq \mathbf{y} \text{ then } \mathbf{z} = \mathbf{y}.$$

Since \mathbb{T} is a finite set and contains Entire, such a minimal \mathbf{y} exists for any \mathbf{s} . In general \mathbf{y} may not be unique. If it is unique for every subset \mathbf{s} of \mathbb{R} , then the type \mathbb{T} is called **explicit**, otherwise it is **implicit**.

Two types with different hull operations are different, even if they have the same set of intervals.

[*Examples. Every inf-sup type is explicit. A mid-rad type is typically implicit.*

As an example of the need for a specified hull algorithm, let \mathbb{T} be the mid-rad type (§11.6.3) where m and r use the same floating point format \mathbb{F} , say binary64, and let \mathbf{s} be the interval $[-1, 1 + \epsilon]$ where $1 + \epsilon$ is the next \mathbb{F} -number above 1. Clearly any minimal interval (m, r) enclosing \mathbf{s} has $r = 1 + \epsilon$. But m can be any of the many \mathbb{F} -numbers in the range 0 to ϵ ; each of these gives a minimal enclosure of \mathbf{s} .

A possible general algorithm, for a bounded set \mathbf{s} and a mid-rad type, is to choose $m \in \mathbb{F}$ as close as possible to the mathematical midpoint of the interval $[\underline{s}, \bar{s}] = [\inf \mathbf{s}, \sup \mathbf{s}]$ (with some way to resolve ties) and then the smallest $r \in \mathbb{F}'$ such that $r \geq \max(m - \underline{s}, \bar{s} - m)$. The cost of performing this depends on how the set \mathbf{s} is represented. If \mathbf{s} is a binary64 inf-sup interval, it is simple. If \mathbf{s} is defined as the range of a function, it might be expensive.]

For 754-conforming implementations the hull operations of the inf-sup types derived from the formats `binary32`, `binary64`, `binary128`, `decimal64` and `decimal128` are denoted respectively as

$$\text{hull}_{\text{b32}}, \text{hull}_{\text{b64}}, \text{hull}_{\text{b128}}, \text{hull}_{\text{d64}}, \text{hull}_{\text{d128}}.$$

11.8.2. *Hull in several dimensions.* In n dimensions the \mathbb{T} -hull, as defined mathematically in §11.8.1, is extended to act componentwise, namely for an arbitrary subset \mathbf{s} of \mathbb{R}^n it is $\text{hull}_{\mathbb{T}}(\mathbf{s}) = (\mathbf{y}_1, \dots, \mathbf{y}_n)$ where

$$\mathbf{y}_i = \text{hull}_{\mathbb{T}}(s_i),$$

and $s_i = \{s_i \mid s \in \mathbf{s}\}$ is the projection of \mathbf{s} on the i th coordinate dimension. It is easily seen that this is a minimal \mathbb{T} -box containing \mathbf{s} , and that if \mathbb{T} is explicit it equals the unique tightest \mathbb{T} -box containing \mathbf{s} .

11.8.3. *Type conversion for bare and decorated intervals.* An implementation shall provide conversion operations between any two supported bare interval types, and between any two supported decorated interval types. Such conversion, whether explicitly coded in a program or automatic while evaluating an expression, shall be done by applying the relevant hull operation.

For bare types, conversion of a \mathbb{T}' -interval \mathbf{x} to a \mathbb{T} -interval \mathbf{y} is therefore done by

$$\mathbf{y} = \text{hull}_{\mathbb{T}}(\mathbf{x}).$$

Let \mathbb{DT}' and \mathbb{DT} denote the corresponding decorated types. Conversion from a \mathbb{DT}' -interval \mathbf{x}_{dx} to a \mathbb{DT} -interval \mathbf{y}_{dy} is done by converting the interval parts and leaving the decoration unchanged, *except* that if the decoration is `com` and the conversion overflows (produces an unbounded interval) the decoration is changed to `dac`. That is,

$$\mathbf{y} = \text{hull}_{\mathbb{T}}(\mathbf{x});$$

$$dy = \begin{cases} \text{dac} & \text{if } dx = \text{com} \text{ and } \mathbf{y} \text{ is unbounded,} \\ dx & \text{otherwise.} \end{cases}$$

11.9. Level 2 interval extensions. Let \mathbb{T} be a bare interval type and f an n -variable scalar point function. A **\mathbb{T} -interval extension** of f , also called a **\mathbb{T} -version** of f , is a mapping \mathbf{f} from n -dimensional \mathbb{T} -boxes to \mathbb{T} -intervals, that is $\mathbf{f} : \mathbb{T}^n \rightarrow \mathbb{T}$, such that $f(x) \in \mathbf{f}(\mathbf{x})$ whenever $x \in \mathbf{x}$ and $f(x)$ is defined. Equivalently

$$\mathbf{f}(\mathbf{x}) \supseteq \text{Rge}(f \mid \mathbf{x}). \quad (34)$$

for any \mathbb{T} -box $\mathbf{x} \in \mathbb{T}^n$, regarding \mathbf{x} as a subset of \mathbb{R}^n . Generically, such mappings are called Level 2 interval extensions.

Though only defined over a finite set of boxes, a Level 2 extension of f is equivalent to a full Level 1 extension of f (§9.4.3) so that this document does not distinguish between Level 2 and Level 1 extensions. Namely define \mathbf{f}^* by

$$\mathbf{f}^*(\mathbf{s}) = \mathbf{f}(\text{hull}_{\mathbb{T}}(\mathbf{s}))$$

for any subset \mathbf{s} of \mathbb{R}^n . Then the interval $\mathbf{f}^*(\mathbf{s})$ contains $\text{Rge}(f \mid \mathbf{s})$ for any \mathbf{s} , making \mathbf{f}^* a Level 1 extension, and $\mathbf{f}^*(\mathbf{s})$ equals $\mathbf{f}(\mathbf{s})$ whenever \mathbf{s} is a \mathbb{T} -box.

11.10. Accuracy requirements for operations.

The standard makes recommendations, but no requirements, on the accuracy of operations. It makes requirements on documenting the accuracy achieved by an implementation.

In this subclause, *operation* denotes any Level 2 version, provided by the implementation, of a Level 1 operation with interval output and at least one interval input. Since for the relevant operations, the decorated versions are defined in terms of the bare versions, they need not be described separately.

11.10.1. *Measures of accuracy.* Three **accuracy modes** are defined: *tightest*, *accurate* and *valid*, that indicate the quality of interval enclosure achieved by an operation. The *tightest* and *valid* modes apply to all interval types and all operations.

Each mode is in the first instance a property of an individual evaluation, that is of a pair (\mathbf{f}, \mathbf{x}) where \mathbf{f} is an operation and \mathbf{x} is an input box. The *accurate* mode is defined only for inf-sup types (because it involves the `nextOut` function), and for interval forward and reverse arithmetic operations of §9.6.3, 9.6.6 (because *accurate* requires operations \mathbf{f} that are monotone at Level 1: $\mathbf{u} \subseteq \mathbf{v}$ implies $\mathbf{f}(\mathbf{u}) \subseteq \mathbf{f}(\mathbf{v})$).

Let $\mathbf{f}_{\text{exact}}$ denote the corresponding Level 1 operation and, for a forward or reverse arithmetic operation, let f be the underlying point function².

The weakest mode, *valid*, is just the property of enclosure:

$$\mathbf{f}_{\text{exact}}(\mathbf{x}) \subseteq \mathbf{f}(\mathbf{x}); \quad (35)$$

this is the minimal requirement for conformance. For a forward arithmetic operation, it is equivalent to (34).

The strongest mode, *tightest*, is the property that $\mathbf{f}(\mathbf{x})$ equals the value $\mathbf{f}_{\text{tightest}}(\mathbf{x})$ that gives best possible \mathbb{T} -interval enclosure of the Level 1 result:

$$\mathbf{f}_{\text{tightest}}(\mathbf{x}) = \text{hull}_{\mathbb{T}}(\mathbf{f}_{\text{exact}}(\mathbf{x})). \quad (36)$$

For a forward arithmetic operation, it is equivalent to

$$\mathbf{f}(\mathbf{x}) = \text{hull}_{\mathbb{T}}(\text{Rge}(f | \mathbf{x})). \quad (37)$$

The intermediate mode *accurate* asserts that $\mathbf{f}(\mathbf{x})$ is *valid*, (35), and is at most slightly wider than the result of applying the *tightest* version to a slightly wider input box:

$$\mathbf{f}(\mathbf{x}) \subseteq \text{nextOut}(\mathbf{f}_{\text{tightest}}(\text{nextOut}(\text{hull}_{\mathbb{T}}(\mathbf{x})))), \quad (38)$$

where the `nextOut` function is defined as follows, for an inf-sup type \mathbb{T} derived from a number format \mathbb{F} . For any \mathbb{F} -number x , define `nextUp`(x) to be $+\infty$ if $x = +\infty$, and the least member of \mathbb{F} greater than x otherwise; since $\pm\infty$ belong to \mathbb{F} by definition, this is always well-defined. Similarly `nextDown`(x) is $-\infty$ if $x = -\infty$, and the greatest member of \mathbb{F} less than x otherwise. [Note. For an IEEE 754 format, these are the `nextUp` and `nextDown` functions in 754-2008 §5.3.1.]

Then for a nonempty \mathbb{T} -interval $\mathbf{x} = [\underline{x}, \bar{x}]$, define

$$\text{nextOut}(\mathbf{x}) = [\text{nextDown}(\underline{x}), \text{nextUp}(\bar{x})]. \quad (39)$$

For a \mathbb{T} -box, `nextOut` acts componentwise.

[Notes.

- The inner `nextOut`() in (38) aims to handle the problem of a function like $\sin x$ evaluated at a very large argument, where a small relative change in the input can produce a large relative change in the result.
- The outer `nextOut`() in (38) relaxes the requirement for correct (rather than, say, faithful) rounding, which may be hard to achieve for some special functions at some arguments.
- The input box \mathbf{x} might have components of a different type from the output type \mathbb{T} . The `hullT` in (38) forces these to have type \mathbb{T} , thus ensuring each component of \mathbf{x} is widened by “at least an \mathbb{F} -ulp” (ulp = unit in the last place) at each finite endpoint.

For example, let \mathbb{T} be a 2-digit decimal inf-sup type. Then `nextOut` widens the \mathbb{T} -interval $\mathbf{x} = [2.4, 3.7]$ to $[2.3, 3.8]$ —an ulp at each end. But an operation might accept 4-digit decimal inf-sup inputs, and \mathbf{x} might be $[2.401, 3.699]$. Then `nextOut`(\mathbf{x}) is $[\text{nextDown}(2.401), \text{nextUp}(3.699)] = [2.4, 3.7]$, a widening by a small fraction of an ulp. But

$$\text{nextOut}(\text{hull}_{\mathbb{T}}([2.401, 3.699])) = \text{nextOut}([2.4, 3.7]) = [2.3, 3.8],$$

giving a widening comparable with the precision of \mathbb{T} .

]

11.10.2. *Documentation requirements.* An implementation shall document an achieved accuracy for each of its interval operations. This shall be done by dividing the operation’s domain of definition into disjoint ranges and stating an accuracy in each range. This information may be supplemented by further detail, e.g., to give accuracy data in a more appropriate way for a non-inf-sup type.

[Example. For $\sin(\mathbf{x})$ in the inf-sup `binary64` type the accuracy might be given as

Operation	Type	Mode	Range
<code>sin</code>	<code>binary64</code>	<i>tightest</i>	for any $\mathbf{x} \subseteq [-10^{15}, 10^{15}]$
		<i>accurate</i>	for all other \mathbf{x} .

²Non-interval arguments, such as the p of `pown`(\mathbf{x}, p), are ignored in the following

(a) Forward		(b) Reverse	
Name	Accuracy	Name	Accuracy
$\text{add}(x, y)$	tightest	$\text{sqrRev}(c, x)$	accurate
$\text{sub}(x, y)$	tightest	$\text{recipRev}(c, x)$	accurate
$\text{mul}(x, y)$	tightest	$\text{absRev}(c, x)$	accurate
$\text{div}(x, y)$	tightest	$\text{pownRev}c, (x, p)$	accurate
$\text{recip}(x)$	tightest	$\text{sinRev}(c, x)$	accurate
$\text{sqrt}(x)$	tightest	$\text{cosRev}(c, x)$	accurate
$\text{hypot}(x, y)$	accurate	$\text{tanRev}(c, x)$	accurate
$\text{case}(b, g, h)$	tightest	$\text{coshRev}(c, x)$	accurate
$\text{sqr}(x)$	tightest	$\text{mulRev}(b, c, x)$	accurate
$\text{pown}(x, p)$	accurate	$\text{divRev1}(b, c, x)$	accurate
$\text{pow}(x, y)$	accurate	$\text{divRev2}(a, c, x)$	accurate
$\text{exp}, \text{exp2}, \text{exp10}(x)$	accurate	$\text{powRev1}(b, c, x)$	accurate
$\text{log}, \text{log2}, \text{log10}(x)$	accurate	$\text{powRev2}(a, c, x)$	accurate
$\text{sin}(x)$	accurate	$\text{atan2Rev1}(b, c, x)$	accurate
$\text{cos}(x)$	accurate	$\text{atan2Rev2}(a, c, x)$	accurate
$\text{tan}(x)$	accurate		
$\text{asin}(x)$	accurate		
$\text{acos}(x)$	accurate		
$\text{atan}(x)$	accurate		
$\text{atan2}(y, x)$	accurate		
$\text{sinh}(x)$	accurate		
$\text{cosh}(x)$	accurate		
$\text{tanh}(x)$	accurate		
$\text{asinh}(x)$	accurate		
$\text{acosh}(x)$	accurate		
$\text{atanh}(x)$	accurate		
$\text{sign}(x)$	tightest		
$\text{ceil}(x)$	tightest		
$\text{floor}(x)$	tightest		
$\text{round}(x)$	tightest		
$\text{trunc}(x)$	tightest		
$\text{abs}(x)$	tightest		
$\text{min}(x_1, \dots, x_k)$	tightest		
$\text{max}(x_1, \dots, x_k)$	tightest		

TABLE 9. Accuracy levels for required arithmetic operations.

Such information by its nature is conservative; the actual accuracy might be better than stated, for many inputs.]

Each operation should be identified by a language- or implementation-defined name of the Level 1 operation (which may differ from that used in this standard), its output type, its input type(s) if necessary, and any other information (e.g., a library version) needed to resolve ambiguity.

11.10.3. *Recommended accuracies.* For 754-conforming types, the accuracy of operations for arbitrary interval inputs should be at least that given in Table 9, which indicates what can be achieved with acceptable efficiency at the publication date of this standard.

11.11. Required operations on bare intervals.

An implementation shall provide a \mathbb{T} -version, see §11.9, of each operation listed in subclauses §11.11.1 and §11.11.3 to §11.11.10, for each supported type \mathbb{T} . That is, those of its inputs and outputs that are intervals, are of type \mathbb{T} .

A 754-conforming implementation, or part thereof, shall provide mixed-type *typeOf* operations, as specified in §11.3.1, for the following operations, which correspond to those that 754 requires to be provided as *formatOf* operations.

add, sub, mul, div, recip, sqrt, sqr, sign, ceil, floor, round, trunc,
abs, min, max, fma.

An implementation may provide more than one version of some operations for a given type. For instance it may provide an “accurate” version in addition to a required “tightest” one, to offer a trade-off of accuracy versus speed or code size. How such a facility is provided, is language- or implementation-defined.

11.11.1. Bare and decorated interval literals.

A text string that conforms to the syntax and semantics defined in this subclause is called an **interval literal**, see §9.6.1, and denotes a bare or decorated interval, termed its *value* (as an interval literal). Any other string has no value as an interval literal. For convenience a string *s* may be termed a *valid* interval literal if it has such a value, or an *invalid* one if it does not.

A *number literal* means (within this subclause) a string that may be part of *s* and denotes an extended-real number, termed its value.

This definition of interval literals is placed in Level 2 because its use is not mandatory at Level 1. However the value of *s* is a Level 1 interval x . Within *s*, conversion of a number literal to its value shall be done as if in infinite precision, ignoring any default or explicit language-defined precision specification. Conversion of the value x to a finite-precision interval y is a separate operation. In all cases y shall contain x ; typically y is the \mathbb{T} -hull of x for some interval type \mathbb{T} .

[Examples.

- The interval denoted by the literal [1.2345] is the mathematical single-point interval $x = [1.2345, 1.2345]$. However the result of \mathbb{T} -text2interval(" [1.2345] "), where \mathbb{T} is the 754 *infsup* binary64 type, is the interval, approximately [1.2344999999999999, 1.2345000000000002], whose endpoints are the nearest binary64 numbers either side of 1.2345.
- Suppose the host language is C/C++. There, conversion of a number literal to a floating point format is considered part of its value, and a “floating-suffix” can be added to change this. E.g., 1.2345f denotes 1.2345 rounded to the nearest float (single-precision)—instead of to the nearest double, which is the default. This is ignored within an interval literal: 1.2345 and 1.2345f both denote the mathematical number 1.2345.

]

The following forms of number literal shall be supported:

- (1) A number literal³ in the syntax provided by the host language of the implementation, in decimal or any other supported radix.
- (2) Either of the strings *inf* or *infinity*, ignoring case, optionally preceded by +, with value $+\infty$; or preceded by -, with value $-\infty$.
- (3) A string in the hexadecimal-significand form of IEEE 754-2008, §5.12.3, with the value specified there.
- (4) A string p/q , that is p and q separated by the / character, where p, q are decimal integer strings, with p optionally signed and $q > 0$. Its value is the exact rational number p/q .

[Note. These categories need not be mutually exclusive, e.g., in C/C++, form (3) is included in (1).]

Bare intervals. The forms of bare interval literal that shall be supported are shown in the following list. To simplify stating the needed constraints, e.g. $l \leq u$, the number literals l, u, m, r are identified with their values. Space shown between elements of a literal, below, denotes zero or more characters that count as whitespace in the host language. For instance in the mid-rad form, whitespace may occur on either side of m and r but not between the two characters of $+-$. [Note. Some characters, e.g., *newline*, may be whitespace in some contexts but not others.]

Constants: The string [empty], ignoring case, whose value is the empty set \emptyset ; and the string [entire], ignoring case, whose value is the whole line \mathbb{R} .

³“floating constant” in C, “real literal constant” in Fortran, “real numeric literal” in Ada, etc.

Inf-sup form: A string $[l, u]$ where l and u are number literals with $l \leq u$, $l < +\infty$ and $u > -\infty$, see §9.2. Its value is the mathematical interval $[l, u]$. A string $[x]$ is equivalent to $[x, x]$.

Mid-rad form: a string $\langle m \pm r \rangle$ where m and r are number literals representing finite numbers with $r \geq 0$. Its value is the mathematical interval $[m - r, m + r]$.

Uncertain form: a string $m?rue$ where: m is a decimal number literal of form (1) above, without exponent; r is empty or is a non-negative decimal integer *ulp-count*; u is empty or is a *direction character*, either **u** (up) or **d** (down); and e is empty or is a decimal *exponent field* for the number m . No whitespace is permitted within the string.

One **ulp** equals one unit in the last place of the number m as written. The literal $m?$ by itself denotes m with a symmetrical uncertainty of $\frac{1}{2}\text{ulp}$, that is the interval $[m - \frac{1}{2}\text{ulp}, m + \frac{1}{2}\text{ulp}]$. The literal $m?r$ denotes m with a symmetrical uncertainty of r ulps, that is $[m - r \times \text{ulp}, m + r \times \text{ulp}]$. Adding **d** (down) or **u** (up) converts this to uncertainty in one direction only, e.g. $m?d$ denotes $[m - \frac{1}{2} \times \text{ulp}, m]$ and $m?ru$ denotes $[m, m + r \times \text{ulp}]$. Finally the exponent field if present multiplies the whole interval by the appropriate power of 10, e.g. $m?rue$ denotes $10^e \times [m, m + r \times \text{ulp}]$.

[Examples. Assuming the common types of decimal number literals are provided:

Form	Literal	Value
Inf-sup	[1.e-3, 1.1e-3]	[0.001, 0.0011]
	[-Inf, 2/3]	$[-\infty, 2/3]$
Mid-rad	<3.56 +- 0.002>	[3.558, 3.562]
	<5/3 +- 1/15>	[8/5, 26/15]
Uncertain	3.56?	[3.555, 3.565]
	3.56?1	[3.55, 3.57]
	3.560?2	[3.558, 3.562]
	3.560?2u	[3.560, 3.562]
	3.56?1e+1	[35.5, 35.7]

]

A formal description of bare interval literals is by the following grammar (using the notation of 754§5.12.3) which defines a `bareIntvlLiteral`, subject to the constraints on l, u, m, r stated above. `anyCase(s)` matches all upper/lower case variants of the string s , e.g. `anyCase("ab")` matches any of "ab", "Ab", "aB", "AB".

```

sign          [+ -]
digit         [0123456789]
natural       {digit} +
langNumLit    {number literal of host language}
langNumLitMant {mantissa of number literal of host language}
langNumLitExp {exponent of number literal of host language}
sp            {whitespace character of host language} *
infLit        {sign} ? ( anyCase("inf") | anyCase("infinity") )
hexNumLit     {see 754§5.12.3}
ratNumLit     {sign} ? {digit} + "/" {digit} +
numberLiteral ( {langNumLit} | {infLit} | {hexNumLit} | {ratNumLit} | )
pointIntvl    "[" {sp} {numberLiteral} {sp} "]"
infSupIntvl   "[" {sp} {numberLiteral} {sp} ", " {sp} {numberLiteral} {sp} "]"
midRadIntvl   "<" {sp} {numberLiteral} {sp} "+-" {sp} {numberLiteral} {sp} ">"
uncertIntvl   {langNumLitMant} "?" {natural} ? [du] ? {langNumLitExp} ?
constIntvl    ( anyCase("[empty]") | anyCase("[entire]") )
bareIntvlLiteral ( {pointIntvl} | {infSupIntvl} | {midRadIntvl} | {uncertIntvl} | {constIntvl} )

```

Decorated intervals. The syntax of decorated interval literals is given by the following extension to the above grammar, which defines a `decoratedIntvlLiteral`.

```

connectChar    {language-defined character}
decorationLit  ( "i11" | "trv" | "def" | "dac" | "com" )
decoratedIntvlLiteral {bareIntvlLiteral} {connectChar} {decorationLit}

```


That is, s is a bare interval literal sx , followed by a language-defined single character c followed by a 3-character decoration string sd . It is recommended that c be the underscore ‘_’, which is assumed here. The string sd is one of `ill`, `emp`, `trv`, `def`, `dac` or `com` representing the corresponding decoration dx . If sx has the value \mathbf{x} , and if \mathbf{x}_{dx} is a permitted combination according to §10.4, then \mathbf{s} has the value \mathbf{x}_{dx} . Otherwise \mathbf{s} has no value as a decorated interval literal.

11.11.2. *Interval constants.* For each supported bare interval type \mathbb{T} there shall be a \mathbb{T} -version of each constant function `empty()` and `entire()` of §9.6.2, returning a \mathbb{T} -interval with value `Empty` and `Entire` respectively. There shall also be a decorated version of each, returning `newDec(Empty) = Emptytrv` and `newDec(Entire) = Entiredac` respectively, of the derived decorated type.

11.11.3. *Forward-mode elementary functions.* Let \mathbb{T} be a supported bare interval type and \mathbb{DT} the derived decorated type. An implementation shall provide a \mathbb{T} -version of each forward arithmetic operation in §9.6.3. Its inputs and output are \mathbb{T} -intervals and it shall be a Level 2 interval extension of the corresponding point function.

[*Note. For operations, some of whose arguments are of integer type, such as integer power `pown(x, p)`, only the real arguments are replaced by intervals.*]

If \mathbb{T} is a 754-conforming type, each such operation shall have a \mathbb{T} -version with accuracy mode as in Table 9(a). For other types, the accuracy mode is language- or implementation-defined.

Each such operation shall have a decorated version with corresponding arguments of type \mathbb{DT} . It shall be a decorated interval extension as defined in §10.6—thus the interval part of its output is the same as if the bare interval operation were applied to the interval parts of its inputs.

The only freedom of choice in the decorated version is how the local decoration, denoted dv_0 in (25) of §10.6, is computed. dv_0 shall be the strongest possible (and is thus uniquely defined) if the accuracy mode of the corresponding bare interval operation is “tightest”, but otherwise is only required to obey (25).

11.11.4. *Interval case expressions and case function.* An implementation shall provide the interval `case(c, g, h)` function, see §9.6.4, for each supported type \mathbb{T} . The input \mathbf{c} is of an arbitrary supported interval type. The inputs \mathbf{g}, \mathbf{h} , and the result, are of type \mathbb{T} . The implementation shall be as if the \mathbb{T} -version of `convexHull` is used in (10) of §9.6.4.

⚠ The decorated version is TBW. Arnold Neumaier had a special recipe, which I need to look up.

11.11.5. *Reverse-mode elementary functions.* An implementation shall provide a \mathbb{T} -version of each reverse arithmetic operation in §9.6.6, for each supported bare interval type \mathbb{T} . Its inputs and output are \mathbb{T} -intervals.

For a 754-conforming type, each such operation shall have a version of that type with accuracy mode as in Table 9(b). For other types, the accuracy mode is language- or implementation-defined.

These operations shall have non-informative decorated versions, as described in §10.7.

11.11.6. *Cancellative addition and subtraction.* An implementation shall provide a \mathbb{T} -version of each of the operations `cancelMinus` and `cancelPlus` in §9.6.7, for each supported bare interval type \mathbb{T} . Its inputs and output are \mathbb{T} -intervals.

The \mathbb{T} -version shall return an enclosure of the Level 1 value if the latter is defined, and `Entire` otherwise. It shall return `Empty` if the Level 1 value is `Empty`. Thus, for the case of `cancelMinus(x, y)`, it returns `Entire` in these cases:

- either of \mathbf{x} and \mathbf{y} is unbounded;
- $\mathbf{x} \neq \emptyset$ and $\mathbf{y} = \emptyset$;
- \mathbf{x} and \mathbf{y} are nonempty bounded intervals with `width(x) < width(y)`.

`cancelPlus(x, y)` shall be equivalent to `cancelMinus(x, -y)`.

If \mathbb{T} is a 754-conforming type, the result shall be the \mathbb{T} -hull of the Level 1 result when this is defined.

[*Notes. Obtaining an accurate result may require computing in extra precision in extreme cases, namely where \mathbf{x} and \mathbf{y} are large intervals of nearly equal width. For an example let \mathbb{T} be any inf-sup type, with parent format \mathbb{F} , and let m be the smallest positive and M the largest finite \mathbb{F} -number. Consider the cases where \mathbf{x} and \mathbf{y} are $[-m, M]$ and $[m, M]$, or vice versa.*]

For any \mathbf{x}, \mathbf{y} , the Level 1 result of these operations, if defined, is always bounded but might overflow at Level 2: that is, there might be no bounded \mathbb{T} -interval containing it. An example is `cancelMinus([M, M], [-M, -M])`. The mathematical value is $[2M, 2M]$ whose \mathbb{T} -hull is $[M, +\infty]$.

An implementation should not, and for an inf-sup type it need not, return Entire in the case of overflow. By avoiding doing so, it makes the result Entire diagnostic: it occurs if, and only if, the Level 1 value is undefined.]

Decorated versions of these operations are not provided.

11.11.7. *Set operations.* An implementation shall provide a \mathbb{T} -version of each of the operations `intersection` and `convexHull` in §9.6.8, for each supported bare interval type \mathbb{T} . Its inputs and output are \mathbb{T} -intervals.

These operations shall return the \mathbb{T} -interval hull of the exact result.

[*Note. If \mathbb{T} is an inf-sup type, the above \mathbb{T} -versions always return the exact result. However, this need not be the case with (for a 754-conforming type \mathbb{T}) the mixed-type version of the operation; nor need it be the case with types that are not inf-sup.*]

These operations shall have non-informative decorated versions, as described in §10.7.

11.11.8. *Constructors.* For each supported bare or decorated interval type \mathbb{T} , there shall be a \mathbb{T} -version of each constructor in §9.6.9. It returns a \mathbb{T} -datum.

Difficulties in implementation. Both `nums2interval`, and `text2interval` when its input is a literal of inf-sup form, involve testing if a boolean value $b = (l \leq u)$ is 0 (false) or 1 (true), to determine whether the interval is empty or nonempty. In the former case, l and u are values of supported number formats within a program; in the latter, they are number literals.

Evaluating b as 0 when the true value is 1 (a “false negative”) leads to falsely returning `Empty` as an enclosure of the true nonempty interval—a containment failure. Evaluating b as 1 when the true value is 0 (a “false positive”) is undesirable, but permissible since it returns a nonempty interval as an enclosure for `Empty`. Implementations shall ensure that false negatives cannot occur, and should ensure that false positives cannot occur.

[*Note. Evaluating b correctly can be hard, if l and u have values very close in a relative sense, and are represented in different ways—e.g., if an implementation allows them to be floating point values of different radices. It could be especially challenging in an extendable-precision context.*

Language rules might cause such errors even when l and u are of the same number format. E.g., in C, if `long double` is supported and has more precision than `double`, default behavior might be to round `long double` inputs l and u to `double`, on entry to a `nums2interval` call. This would be non-conforming—the comparison $l \leq u$ requires the exact values to be used, which requires use of a version of `nums2interval` with `long double` arguments.]

Bare intervals. A bare interval constructor call either **succeeds** or **fails**. When it fails, a language- or implementation-defined exception should be signaled. This notion is used to determine the value returned by the corresponding decorated interval constructor.

For the constructor `nums2interval(l, u)`, the inputs l and u are datums of supported number formats, where the format of l may differ from that of u . For a given \mathbb{T} , each such format combination is called a kind in this subclass.

An implementation shall provide at least one kind where l and u have the same format, which is compatible with \mathbb{T} , see §11.5. For a 754-conforming type \mathbb{T} , a *formatOf* kind shall be provided, which accepts l and u having any supported 754 formats of the same radix as \mathbb{T} ’s parent format. Apart from these two requirements, what kinds are supported is language- or implementation-defined.

The constructor call succeeds in the following cases.

- The implementation determines that the Level 1 value x exists, or is unable to determine that it does not exist—see the discussion at the start of this subclass. In the first case—that is, neither l nor u is NaN, and the exact extended-real values of l and u are known to satisfy $l \leq u$, $l < +\infty$, $u > -\infty$ —the result shall be a \mathbb{T} -interval containing x . In the second case the result shall be a \mathbb{T} -interval containing l and u .

If \mathbb{T} is a 754-conforming type and l and u are of 754 formats with the same radix as \mathbb{T} , the second case cannot arise, and the result shall be the \mathbb{T} -hull of x . For other types \mathbb{T} , the tightness of the result is implementation-defined.

- The special cases $l = u = -\infty$ or $l = u = +\infty$. Here the Level 1 result is undefined. In finite precision however, l and u are likely to be finite values that overflowed.

If $l = u = +\infty$, the result shall be the tightest \mathbb{T} -interval that is unbounded above. For any type, this exists and has the form $\mathbf{x} = [\text{HUGEPOS}, +\infty]$, where **HUGEPOS** is a uniquely defined extended-real number $< +\infty$.

If $l = u = -\infty$, the result shall be $[-\infty, \text{HUGENEG}]$, defined similarly.

[*Note.* If \mathbb{T} is an inf-sup type based on a format \mathbb{F} , then **HUGEPOS** is the largest finite \mathbb{F} -number, and **HUGENEG** equals $-\text{HUGEPOS}$ by symmetry, see §11.5.]

Otherwise the call fails and the result is **Empty**.

For the constructor `text2interval(s)`, the input s is a text string. The constructor call succeeds in the following cases.

- The implementation determines that s is a valid interval literal with value \mathbf{x} , see §11.11.1. The result is a \mathbb{T} -interval containing \mathbf{x} . If \mathbb{T} is a 754-conforming type, the result shall be the \mathbb{T} -hull of \mathbf{x} . ▲ Is the last sentence always practical when s contains exotic number literals of high precision?
- s is of inf-sup form $[l, u]$ but the implementation cannot determine that $l \leq u$. The result shall be an interval containing l and u .
- s is of inf-sup form $[l, u]$ with $l = u = +\infty$ or $l = u = -\infty$. These special cases shall be treated as in `nums2interval`.

Otherwise the call fails and the result is **Empty**.

Decorated intervals. For each constructor `nums2interval(l, u)` giving a result of the bare interval type \mathbb{T} , there shall be a corresponding constructor `nums2interval(l, u, d)` giving a result of the derived decorated type.

For each constructor giving a result of the bare interval type \mathbb{T} , there shall be a corresponding decorated constructor, with an optional extra final argument d which is a decoration, and giving a result of the derived decorated type. Thus the bare constructor `nums2interval(l, u)` has the decorated forms `nums2interval(l, u)` and `nums2interval(l, u, d)`; `text2interval(s)` has the decorated forms `text2interval(s)` and `text2interval(s, d)`. The following cases arise.

- The bare constructor succeeds, returning \mathbf{x} , and d is not provided. Then the decorated form returns `newDec(x)`, see §10.5.
- The bare constructor succeeds, returning \mathbf{x} , and d is provided. Then the decorated form returns (\mathbf{x}, d) if this is a permitted combination, and **NaN** otherwise.
- Otherwise the decorated form returns **NaN**.

11.11.9. *Numeric functions of intervals.* An implementation shall provide a \mathbb{T} -version of each numeric function in Table 3 of §9.6.10, for each supported bare interval type \mathbb{T} . It shall return a result in a number format \mathbb{F} that should be compatible with \mathbb{T} , see §11.5. An implementation may provide several versions, returning results in different formats.

The implementation shall document how it breaks ties, e.g., when computing the closest \mathbb{F} -number to a value that is midway between two \mathbb{F} -numbers. If \mathbb{F} is a 754-conforming format, the tie-breaking method shall follow 754§4.3.1 and 754§4.3.3; otherwise it is language- or implementation-defined.

If \mathbb{T} is a 754-conforming type, versions shall be provided that return the result in any supported 754 format of the same radix as \mathbb{T} .

- `inf(x)` returns the Level 1 value, rounded toward $-\infty$.
- `sup(x)` returns the Level 1 value, rounded toward $+\infty$.
- `mid(x)` returns **NaN** if \mathbf{x} is **Empty**, and 0 if \mathbf{x} is **Entire**. Otherwise for nonempty $\mathbf{x} = [\underline{x}, \bar{x}]$ its value is defined by the following cases.

\underline{x}, \bar{x} both finite	the closest finite \mathbb{F} -number to the Level 1 value, breaking ties appropriately;
$\underline{x} = -\infty, \bar{x}$ finite	the finite negative \mathbb{F} -number of largest magnitude;
\underline{x} finite, $\bar{x} = +\infty$	the finite positive \mathbb{F} -number of largest magnitude;

The three tabulated cases can be obtained by computing $(\underline{x} + \bar{x})/2$ exactly in the extended reals and rounding the result to the nearest finite \mathbb{F} -number.

- `rad(x)` returns **NaN** if \mathbf{x} is empty, and otherwise the smallest \mathbb{F} -number r such that \mathbf{x} is contained in the exact interval $[m - r, m + r]$, where m is the value returned by `mid(x)`.

[*Note. $\text{rad}(x)$ may be $+\infty$ even though x is bounded, if \mathbb{F} has insufficient range. However, if \mathbb{F} is a 754 format and \mathbb{T} is the derived inf-sup type, $\text{rad}(x)$ is finite for all bounded nonempty intervals.*]

- $\text{wid}(x)$ returns the same value as $2 * \text{rad}(x)$, rounded toward $+\infty$.

[*Note. At Level 2, $\text{wid}(x)$ may be infinite though $\text{rad}(x)$ is finite.*]

- $\text{mag}(x)$ returns the Level 1 value rounded toward $+\infty$ if this value exists (if x is nonempty). Otherwise it returns NaN.
- $\text{mig}(x)$ returns the Level 1 value rounded toward zero if this value exists (if x is nonempty). Otherwise it returns NaN.

Each bare interval operation in this subclause shall have a decorated version, where each input of bare interval type is replaced by one of the corresponding decorated interval type, and the result format is that of the bare operation. Following §10.7, if any input is NaN, the result is NaN. Otherwise the result is obtained by discarding the decoration and applying the corresponding bare interval operation.

11.11.10. *Boolean functions of intervals.*

An implementation shall provide a \mathbb{T} -version of the function $\text{isEmpty}(x)$ and the function $\text{isEntire}(x)$ in §9.6.11, for each supported bare interval type \mathbb{T} .

An implementation shall provide a \mathbb{T} -version of each of the comparison relations in Table 4 of §9.6.11, for each supported bare interval type \mathbb{T} . Its inputs are \mathbb{T} -intervals.

For a 754-conforming part of an implementation, mixed-type versions of these relations shall be provided, where the inputs have arbitrary 754-conforming types of the same radix.

These comparisons shall return, in all cases, the correct value of the comparison applied to the intervals represented by the inputs as if in infinite precision. In particular equal shall return **true** if and only if its arguments are equal as defined in §11.4.

Each bare interval operation in this subclause shall have a decorated version, where each input of bare interval type is replaced by one of the corresponding decorated interval type. Following §10.7, if any input is NaN, the result is **false**, and a language- or implementation-defined exception should be signaled. Otherwise the result is obtained by discarding the decoration and applying the corresponding bare interval operation.

There shall be a function $\text{isNaI}(X)$ for input X of any decorated type, that returns **true** if X is NaN, else **false**. It shall not raise an exception, whether or not its input is NaN.

11.11.11. *Reduction operations.* An implementation that provides 754-conforming interval types shall provide the four reduction operations sum , dot , sumSquare and sumAbs of IEEE 754-2008 §9.4, correctly rounded. These shall be provided for the parent formats of each such type, and in any case for **binary64** if any 754-conforming type with binary radix is provided, and for **decimal64** if any 754-conforming type with decimal radix is provided.

Correctly rounded means that the returned result is defined as follows.

- If the exact result is defined as an extended-real number, return this after rounding to the relevant format according to the current rounding mode. An exact zero shall be returned as $+0$ in all rounding modes.
- Otherwise return NaN.

All other behavior, such as overflow, underflow, setting of IEEE 754 flags, raising of exceptions, and behavior on vectors whose length is given as non-integral, zero or negative, shall be as specified in IEEE 754-2008 §9.4. In particular evaluation is as if in exact arithmetic up to the final rounding, normally with no possibility of intermediate overflow or underflow.

Overflow shall be made impossible. It results from adding an extremely large number of large terms of the same sign, and can be prevented by giving the implementation the capacity to handle more terms than current machines can possibly sum in a reasonable time.

It is recommended that these operations be based on an implementation of Complete Arithmetic as specified in §13.10.

11.12. Recommended operations.

11.12.1. *Forward-mode elementary functions.*

The functions listed in §9.7.1 are arithmetic operations. If any of them is provided, it shall have a version for each bare and decorated interval type, in the same way as for the required operations in §9.6.3.

11.12.2. *Slope functions.*

The functions listed in §9.7.2 shall be handled in the same way as those in §11.12.1.

11.12.3. *Extended interval comparisons.*

How the operations in §9.7.3 are handled at Level 2 is implementation-defined.

11.13. Compressed arithmetic at Level 2.

How the operations in §10.12 are handled at Level 2 is implementation-defined.

DRAFT 7.3

12. Input and output (I/O) of intervals

12.1. Overview. This standard specifies conversion from a text string that holds an interval literal to an interval internal to a program (input), and the reverse (output). The method by which such strings might be read from, or written to, a character stream is language-defined.

Containment shall hold on input and output so that, when a program computes an enclosure of some quantity given an enclosure of the data, a correctly written program can ensure this holds all the way from text data to text results.

In addition to normal I/O, the standard requires each interval type \mathbb{T} to have a *public representation*. This has two parts: operations to convert any internal \mathbb{T} -interval \mathbf{x} to a string \mathbf{s} , and back again to recover \mathbf{x} exactly; and documentation of how to convert \mathbf{s} to the Level 1 interval represented by \mathbf{x} . For proprietary types in particular, this makes explicit the mathematical definition of the type, while letting its Level 3 implementation remain private.

12.2. Input. Input is provided for each supported bare or decorated interval type \mathbb{T} by the \mathbb{T} -version of `text2interval` (\mathbf{s}), where \mathbf{s} is a string, as specified in §11.11.8.

12.3. Output. Implementations shall provide a function

$$\text{interval2text}(\mathbf{x}, cs)$$

where \mathbf{x} is a bare or decorated interval of any supported type \mathbb{T} and cs is a string, the conversion specifier. It converts \mathbf{x} to a string, in a way specified by cs .

The allowed forms of cs are language-defined, and may depend on \mathbb{T} , but shall let the user specify output in any of the forms of interval literal in §11.11.1, namely:

- (i) Inf-sup form $[l, u]$, where the layouts of l and u can be specified independently.
- (ii) Mid-rad form $\langle m \mp r \rangle$, where the layouts of m and r can be specified independently.
- (iii) Uncertain form such as $m ? r$.

There should be ways to control how Empty and Entire are output, e.g., whether Entire becomes `[Entire]` or `[-Inf, Inf]`. In all cases the resulting string shall be a valid interval literal that may be read by `text2interval`. In the bare interval case, its Level 1 value \mathbf{y} shall contain \mathbf{x} . In the decorated case, the interval part of \mathbf{y} shall contain that of \mathbf{x} , and the decoration part of \mathbf{y} shall equal that of \mathbf{x} .

It shall be possible to specify output of l , u , m and r to a given number of places after the point or to a given number of significant figures, and to specify their field width. (For instance, by conversion specifiers like `f12.5` and `e12.5` in Fortran, or `%12.5f` and `%12.5e` in C.)

If \mathbb{T} is a 754-conforming type, the enclosure represented by \mathbf{s} shall be tightest possible. Namely let $\mathbf{x} = [\underline{x}, \bar{x}]$ be a \mathbb{T} -interval. For inf-sup form, l is the largest number of the specified layout that is $\leq \underline{x}$ and u is the smallest number of the specified layout that is $\geq \bar{x}$; either may be infinite in case of overflow. For mid-rad form, m is the number of the specified layout that is closest to the exact midpoint; then r is the smallest number of the specified layout such that the exact interval $[m - r, m + r]$ contains \mathbf{x} . The treatment of infinite values, overflow and tie-breaking shall follow that of the `inf`, `sup`, `mid` and `rad` functions in §11.11.9.

For other types the tightness of enclosure of \mathbf{x} by \mathbf{s} is language- or implementation-defined.

12.4. Public representation. For any supported bare interval type \mathbb{T} an implementations shall provide functions `interval2public` and `public2interval`, as follows.

- For any \mathbb{T} -interval datum \mathbf{x} the value `interval2public`(\mathbf{x}) is a string \mathbf{s} , the **public representation** of \mathbf{x} , such that `public2interval`(\mathbf{s}) = \mathbf{x} .
- The implementation's documentation shall describe how to convert \mathbf{s} to the mathematical interval $[l, u]$ represented by the datum \mathbf{x} , by means of an algorithm for obtaining l and u as decimal, hexadecimal or binary numbers, exactly or to any desired accuracy.

If \mathbb{T} is a 754-conforming type, the public representation \mathbf{s} of \mathbf{x} shall be as an interval literal (§11.11.1) that, for nonempty \mathbf{x} , is of inf-sup form. Its bounds l, u if finite shall be represented as exact decimal numbers if \mathbb{T} is a decimal type, or in the hexadecimal-significand form of 754§5.12.3 if \mathbb{T} is a binary type.

[Note. A public representation should aim for simplicity. For instance if \mathbf{x} represents an interval with small integer bounds such as $[1, 2]$, it should be straightforward to convert \mathbf{s} by hand or with the

help of a pocket calculator. A good public representation exposes the values of the parameters on which the mathematical model of the type is based.]

DRAFT 7.3

This is “EXTRA BITS” — much will be discarded from final text.

△ Q for the “Operations on decorated intervals” clause: We need to specify how the interval part of a `ill` interval behaves w.r.t. real-valued functions like `radius` (always NaN?) . Also w.r.t. “forget decoration” (gives Empty?).

12.5. Representations. [Note. Some of these ideas appeared in previous drafts of the standard or in position papers, others didn’t.]

- A *representation* of an interval type comprises a set \mathbb{IF} called an interval level-3 type, and whose members are called (level 3) *interval objects*, together with a map r from a subset of \mathbb{IF} (the “valid” objects) to \mathbb{T} . If object X maps to datum \mathbf{x} , we say X represents \mathbf{x} . (DS§3.1 Table 1 note *b*: “Not every interval object necessarily represents an interval datum, but when it does, that datum is unique. Each interval datum has at least one representation, and may have more than one.”)

[Note. A representation is not an approximation—it means just what it says. E.g., an interval object in mid-rad form with midpoint $m = 1$ and radius $r = 1e-300$ means precisely the mathematical interval $[1 - 10^{-300}, 1 + 10^{-300}]$.]

- A *text representation* of an interval type is a representation whose interval level-3 type is a set \mathbb{IT} of text strings.
- A representation is *standardized* if there is also provided a map s from the whole of \mathbb{T} into \mathbb{IF} , such that $s(\mathbf{x})$ is a representation of \mathbf{x} :

$$r(s(\mathbf{x})) = \mathbf{x} \quad \text{for each } \mathbf{x} \text{ in } \mathbb{T}.$$

The object $s(\mathbf{x})$ is called the “standard representation” of the datum \mathbf{x} .

[Example. In an inf-sup representation, r might map both the objects $(-0, 3)$ and $(+0, 3)$ to the interval datum $[0, 3]$. Suppose the standardized representation always uses $+0$; then s would map $[0, 3]$ to $(+0, 3)$.]

[Note. A standardized text representation of \mathbb{T} essentially defines a way to write any \mathbb{T} -interval out in text form and read it back exactly. Think of map s as “write” and r as “read”. See §12.6.2.]

12.6. Rationale for defined hulls, text representation, and reproducibility.

12.6.1. *The hull.* The decision whether the hull operation is made part of an interval type’s definition affects (a) inter-interval type conversion, which is done by forming the hull; (b) the definition of “tightest” standard functions; (c) hence reproducibility.

[Example. Use the notation $m \pm r$ to mean the interval $[m-r, m+r]$ written in mid-rad form. Let \mathbb{T} comprise all intervals $m \pm r$ where m and r belong to the set \mathbb{F} of 4-digit decimal floating point numbers, with some finite exponent range that is irrelevant here.

What is the conversion of the inf-sup interval $[1, 1.003]$ to mid-rad? If $\text{hull}_{\mathbb{T}}$ is not part of the definition of \mathbb{T} , one implementation can choose 1.001 ± 0.002000 , another can choose 1.002 ± 0.002000 , and both are right.]

Whether one approves of this non-uniqueness depends on one’s philosophy of the standard: should it specify minimum demands consistent with the FTIA, or should it tie things down more closely?

Personally I agree with Dan Zuras (754 chair for a number of years) that one of the main reasons why the 754 floating point standard has been so successful is because it took the hard road of specifying things down to the last bit. The parallel LIA standard, which didn’t, has sunk with barely a trace.

So let’s tie things down. What, and how far? I think requiring mid-rad implementors to define an unambiguous, platform-independent hull operation is not too much to ask. Nate Hayes agrees.

12.6.2. *Standardised text representation.* In Motion 17, persuaded by Michel Hack, I included the requirement that—for a 754-conforming inf-sup type \mathbb{T} only—there should be a way to dump any \mathbb{T} -interval \mathbf{x} to text using 754’s “hexadecimal significand” form; but we forgot to require it should be readable back again.

The current motion goes beyond inf-sup to allow *any level 3 representation whatever* of intervals, which, I believe, makes it especially important that there should be an exact textual representation \mathbf{y} of any interval datum \mathbf{x} , documented so that the user can, if desired, manually construct the mathematical interval represented. The two-way mapping (functions r and s) makes it possible to confirm that what one sees as text is what is actually being computed with.

The requirement that y depend only on the level 2 datum x , not on the possibly non-unique internal representation of x , is deliberate. It simplifies matters for the user and protects the implementer.

12.6.3. *Reproducibility.* I believe reproducibility, important for floating point, is doubly important for interval computing. With default compiler options, running the same interval code on different platforms is not expected to produce the same results down to the last bit. But a user should be able to choose a mode that (presumably at the expense of speed) ensures identical results on all platforms, for code restricting itself to some subset of language features.

There is a counter-argument that reproducibility is *less* important for interval computing: if different platforms give different results, good, because they both enclose the true result. I accept that but am unmoved by it. As Dan Zuras (13 July 2010) says, why should I trust *either* result? For instance, what if a specialised interval computing chip has something like the famous Pentium bug? The bug will probably be far easier to find if I can run the chip in “reproducibility mode”.

The key to reproducibility is reproducible behaviour of interval standard functions. The only reasonable way to specify this is to require these functions to return “tightest” results for all arguments. The remarkable work of the French experts means it will soon be practical to compute results correct to the last bit in either rounding direction for all (point) standard functions, all arguments and all sensible number formats, with little loss of speed. For inf-sup interval types this makes “tightest”, hence reproducible, standard functions entirely practicable.

For general interval types I do not know how hard it is to achieve the same level of tightness (but see Nate Hayes’ view below). In general, I believe any steps that promote reproducibility will in the long term make systems programmers and users alike more confident that our interval systems are correct.

DRAFT 7.3

13. Level 3 description

13.1. Level 3 introduction. Level 3 is where Level 2 datums are represented, and operations on them described, in terms of more primitive entities and operations. How this is done is implementation-defined. Implementation may be by hardware, software, or a combination of the two.

Level 3 entities are here called *objects*; they represent Level 2 datums and may be referred to as *concrete*, while the datums are *abstract*. An implementation shall behave as if the relation between Levels 2 and 3 is as follows.

- The set of boolean values, the set of strings, and the set \mathbb{D} of decorations are regarded as being the same at Level 3 as at Levels 1 and 2.
- The bare interval objects are organized into disjoint sets, *concrete bare interval types*, that are in one-to-one correspondence with the abstract bare interval types of Level 2. As at Level 2, a decorated interval is an ordered pair (bare interval, decoration), so this induces a one-to-one correspondence between the abstract and the concrete decorated interval types.
- The number objects are organized into disjoint sets, *concrete number formats*, that are in one-to-one correspondence with the abstract number formats of Level 2.

Thus intervals of a particular type exist in four forms: bare or decorated, and in either case abstract datums at Level 2 or concrete objects at Level 3. Similarly, numbers of a particular format exist in two forms, abstract or concrete, and the same name may be used for either.

In this document, the same name is normally used for an abstract type or format and its concrete counterpart. This convention, and the term “object”, are not intended to constrain the names that an implementation gives to types or formats, nor the data structures it uses.

[*Example. The format `binary64` may denote either the set of representations (in the sense of IEEE 754 §3.2) of `binary64` numbers, or the set of numbers thus represented. Then for instance `-0` and `+0` are different objects, but represent the same datum.*]

13.2. Representation. Individual datums of an abstract type or format are *represented* by individual objects of its concrete type or format. While the correspondence between abstract and concrete types or formats as a whole is one-to-one, that between datums and objects is not so. The property that defines a representation, for a given type or format, is:

Each datum shall be represented by at least one object. Each object shall represent at most one datum. (40)

An object that represents a datum is called *valid*; one that does not is called *invalid*.

That is, representation is a *partial function* that is *onto* but usually *not one-to-one*, from the set of objects to the set of datums of a given type or format. The set of valid objects is the domain of this function.

[*Examples. Let \mathbb{F} be a 754 format and let \mathbb{T} the derived (bare) inf-sup type. Two possible representations are:*

- **inf-sup form.** *The objects are defined to be pairs (l, u) where l, u are members of \mathbb{F} . A nonempty \mathbb{T} -interval $x = [\underline{x}, \bar{x}]$ is represented by the object (x, \bar{x}) , and (for instance—many other ways are possible) Empty is represented by (NaN, NaN) . Its valid objects are (NaN, NaN) , together with all (l, u) such that l, u are not NaN and $l \leq u$, $l < +\infty$, $u > -\infty$.*
- **neginf-sup form.** *This is as the previous, except that $x = [\underline{x}, \bar{x}]$ is represented by $(-\underline{x}, \bar{x})$. Its valid objects are (NaN, NaN) , together with all (l, u) such that l, u are not NaN and $0 \leq l + u$, $l > -\infty$, $u > -\infty$.*

If, in these descriptions l , u and NaN are viewed as Level 2 datums, then each interval has only one representative: for a nonempty \mathbb{T} -interval there are unique l and u , and there is a unique NaN to use in the fields of Empty. However l , u and NaN themselves have representations, and from this viewpoint some intervals have more than one representative: $[0, 1]$ can be either $(-0, 1)$ or $(+0, 1)$, while there are many NaNs, quiet or signaling and with different payloads, to use in $\text{Empty} = (\text{NaN}, \text{NaN})$.

]

13.3. Operations and representation. Each Level 2 (abstract) library operation is implemented by a corresponding Level 3 (concrete) operation, whose behavior shall be consistent with the abstract operation. That is, let $y = \varphi(x_1, x_2, \dots)$ be a Level 2 operation instance whose

inputs and output are any mix of number, interval, decoration, string or boolean datums, and let objects x'_1, x'_2, \dots represent x_1, x_2, \dots respectively. Then $y' = \varphi(x'_1, x'_2, \dots)$ shall be defined and be a representative of y .

Since for each Level 2 operation, the result is defined for arbitrary input datums, it follows that each Level 3 library operation has a unique result, up to representation, for arbitrary *valid* input objects. That is, if one chooses different representatives x'_i for the x_i , the result y' may be different but is still a representative of y . The result, when some inputs are invalid, is implementation-defined. An implementation shall provide means for an exception to be raised when this occurs.

To promote reproducibility, an implementation should provide a computational mode where, at least for library operations with numeric output, the representative of the output is independent of the representatives of the inputs. That is, in the notation above, y' is not changed by changing the x'_i . [Example. Let \mathbb{F} be a 754 format and \mathbb{T} the derived inf-sup type. Suppose a nonempty \mathbb{T} -interval $[l, u]$ is represented at Level 3 as the pair of \mathbb{F} -numbers (l, u) . Let f be the expression

$$1/\text{inf}(x) > 0.$$

and consider $[0, 1]$ with the two Level 3 representations $x = (-0, 1)$ and $x' = (+0, 1)$. Then $x = x'$ in the Level 2 sense, but a naive implementation gives

$$\begin{aligned} f(x) &= \left(\frac{1}{\text{inf}(x)} > 0 \right) &= \left(\frac{1}{-0} > 0 \right) &= (-\infty > 0) = \text{false}; \\ f(x') &= \left(\frac{1}{\text{inf}(x')} > 0 \right) &= \left(\frac{1}{+0} > 0 \right) &= (+\infty > 0) = \text{true}. \end{aligned}$$

The standard does not say which of these two results is “correct”. But since they differ, the equality principle is violated and such an implementation is non-XXX. One way to make it XXX is to canonicalize all operations that output an interval, to ensure that for instance all zero bounds are stored as +0.]

13.4. Format conversion. On 754 systems, level 2 interval format conversion (the hull operation) shall be implemented in terms of the floating-point operations *formatOf-convertFormat* defined in 754§5.4.2, with the appropriate outward rounding.

13.5. Interchange formats. ⚠ We need a motion on this subclause, which was my invention. ■

The purpose of interchange formats is to allow the loss-free exchange of level 2 interval data between 754-conforming implementations. This is done by imposing a standard level 3 and level 4 representation. Let \mathbb{F} be a 754 format and x a (bare) nonempty \mathbb{F} -interval datum, so that its lower bound \underline{x} and upper bound \bar{x} are \mathbb{F} -numbers. An interchange format of x is the concatenation of the bit strings of the \mathbb{F} -representations of \underline{x} and \bar{x} in that order, where:

- 0 shall be represented as +0.
- For decimal formats, any member of the number’s cohort is permitted. The choice is implementation-defined.
- When x is the empty set, \underline{x} and \bar{x} are taken as NaN. Whether qNaN or sNaN is used, and any payload, are implementation-defined.

[Note. The above rules imply an interval has a unique interchange representation if it is nonempty and in a binary format, but not generally otherwise. The reason for the rules is that the sign of a zero endpoint cannot convey any information relevant to intervals; but an implementation may potentially use cohort information, or a NaN payload.]

The interchange format for a decoration comprises ⚠ TBW. Thus a decoration occupies one byte. ■

The interchange format for a decorated interval is the concatenation of those for its interval and decoration parts, in that order.

A 754-conforming implementation shall provide an interchange format for each supported 754 interval format. Interchange formats for non-754 interval formats, and on non-754 systems, are implementation-defined. If an implementation provides other decoration attributes besides the standard ones, then how it maps them to an interchange format is implementation-defined.

13.6. Support levels for interval elementary functions. The Fundamental Theorem of Interval Arithmetic (FTIA) relies on each point elementary function e in a real expression being replaced by an interval version \mathbf{e} . Mathematically, \mathbf{e} may be an arbitrary interval extension of e , and its arguments and result are not limited by any concrete interval format.

A level 2 interval version is implemented at level 3 in terms of concrete formats such as binary64. The standard defines three levels of mixed format support, below. For each one, \mathbf{e} delivers a result of a specified ci-format \mathbb{F} , from operands of a limited number of ci-formats.

Implementations shall give at least SRSF support to all supported elementary functions. 754-conforming implementations shall give at least SRMF support.

SRSF. Single-radix, single-format. In SRSF support, e has an interval version \mathbf{e} that takes \mathbb{F} -interval operand(s) and gives an \mathbb{F} -interval result. Thus explicit format conversion is needed for any operand of a different format from \mathbb{F} .

SRMF. Single-radix, mixed-format. In SRMF support, e has an interval version \mathbf{e} that takes operand(s) of any supported interval format of the same radix as \mathbb{F} , and gives an \mathbb{F} -interval result. Thus explicit format conversion is needed for any operand whose format has a different radix from that of \mathbb{F} .

MRMF. Mixed-radix, mixed-format. In MRMF support, e has an interval version \mathbf{e} that takes operand(s) of any supported interval format, and gives an \mathbb{F} -interval result. Thus no explicit format conversion is required for any operand.

SRMF includes SRSF (SRMF support provides SRSF in particular); and MRMF includes SRMF.

MRMF support and mixed-format interval expressions of more than one operation are considered to be language issues.

[*Note. For a 754 formatOf floating-point operation, for any combination of input and output formats of the same radix, the correctly rounded result is produced, eliminating the risk of “double rounding” error in mixed-format operations.*

Most algorithms for the basic interval operations, in particular those in §13.7, can exploit the formatOf feature. That is, they can be written in terms of point operations so that arbitrary mixed formats of the same radix can be handled by essentially the same code, while remaining optimally tight at the level of a single interval operation.]

13.7. Operation tables for basic interval operations.

The tables in this subclause are an explicit realization of the general definition of interval operations given in §9.4.3. They are not normative, but are one possible basis for coding the interval versions of $+$, $-$, $*$, $/$. For fuller details see [1], [2].

Notation. In addition to the notation in §4.1 this subclause also uses, for a specified n-format \mathbb{F} :

- ∇, \triangle : the roundings downwards and upwards to the next element of \mathbb{F} ,
- $\nabla, \text{etc.}$: the operations for elements of \mathbb{F} with rounding downwards,
- $\triangle, \text{etc.}$: the operations for elements of \mathbb{F} with rounding upwards.
- $\diamond \mathbf{s}$: the same as $\text{hull}_{\mathbb{F}}(\mathbf{s})$, the \mathbb{F} -hull of a subset \mathbf{s} of \mathbb{R} .

For intervals \mathbf{a} and $\mathbf{b} \in \mathbb{IR}$ (the bounded, nonempty mathematical intervals), arithmetic operations are defined as set operations in \mathbb{R} by:

$$\mathbf{a} \circ \mathbf{b} := \{ a \circ b \mid a \in \mathbf{a} \wedge b \in \mathbf{b} \wedge a \circ b \text{ is defined} \}, \quad (41)$$

for all \mathbf{a} and $\mathbf{b} \in \mathbb{IR}$ and $\circ \in \{+, -, *, /\}$. If $0 \notin \mathbf{b}$ in case of division, then for all \mathbf{a} and $\mathbf{b} \in \mathbb{IR}$ also $\mathbf{a} \circ \mathbf{b} \in \mathbb{IR}$.

Then binary arithmetic operations in \mathbb{IF} (the bounded, nonempty level 2 interval datums) are uniquely defined by:

$$\mathbf{a} \diamond \mathbf{b} := \diamond (\mathbf{a} \circ \mathbf{b}), \quad (42)$$

for all \mathbf{a} and $\mathbf{b} \in \mathbb{IF}$ and all $\circ \in \{+, -, *, /\}$. For division we assume again that $0 \notin \mathbf{b}$.

For intervals $\mathbf{a} = [a_1, a_2]$ and $\mathbf{b} = [b_1, b_2] \in \mathbb{IF}$ these operations \diamond , for $\circ \in \{+, -, *, /\}$, have the property

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

or with the monotone roundings ∇ and \triangle ,

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

These operations and the unary operation $-\mathbf{a}$ can be expressed by more explicit formulas as shown in Tables 10–13. There the operators for intervals are simply denoted by $+$, $-$, $*$, and $/$.

These tables assume that \mathbf{a} and \mathbf{b} are nonempty and bounded. To extend them to general intervals, the first rule is that any operation with the empty set \emptyset returns the empty set. Then, the tables extend to possibly unbounded intervals of \mathbb{IF} by using the standard formulae for arithmetic operations involving $\pm\infty$, which are implemented in 754, together with one rule that goes beyond 754 arithmetic:

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

This rule is not a new mathematical law, merely a short cut to compute the bounds of the result of multiplication on unbounded intervals.

$$\begin{aligned} \text{Negation} \quad & -\mathbf{a} = [-a_2, -a_1]. \\ \text{Addition} \quad & [a_1, a_2] + [b_1, b_2] = [a_1 \nabla b_1, a_2 \triangle b_2]. \\ \text{Subtraction} \quad & [a_1, a_2] - [b_1, b_2] = [a_1 \nabla b_2, a_2 \triangle b_1]. \end{aligned}$$

TABLE 10. Negation, addition, subtraction.

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

TABLE 11. Multiplication.

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

TABLE 12. Division by interval not containing 0.

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. Whenever zero 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 following table. Here, the parentheses stress that the bounds $-\infty$ and $+\infty$ are not elements of the interval.

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 division by $[b_1, b_2]$ actually means two divisions. The results of the two divisions are already shown in Table 12, division with $0 \in \mathbf{b}$.

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	$[0, 0]$	$[0, 0]$
$a_1 < 0, a_2 \leq 0$	\emptyset	$[a_2 \nabla b_1, +\infty)$	$(-\infty, a_2 \triangle b_2]$
$[a_1, a_2], a_1 < 0 < a_2$	\emptyset	$(-\infty, +\infty)$	$(-\infty, +\infty)$
$0 \leq a_1, 0 < a_2$	\emptyset	$(-\infty, a_1 \triangle b_1]$	$[a_1 \nabla b_2, +\infty)$

TABLE 13. Division by interval containing 0.

However, in the user's program the two divisions appear as a single operation, as division by an interval $\mathbf{b} = [b_1, b_2]$ with $b_1 < 0 < b_2$ —an operation that delivers two distinct results.

▲ Prof Kulisch's motion proposed several ways to handle this situation, but listing them does not seem appropriate for the standard. Suggestions for text here, please.

13.8. Care needed with innerPlus and innerMinus. (informative)

[Example. Consider inf-sup intervals using 3 decimal digit floating point arithmetic. Let $x = [.0001, 1]$ and $y = [-1, -.0002]$. Thus x is slightly the wider, so $z_1 = \text{innerMinus}(x, y)$ is defined (its exact value is $[1.0001, 1.0002]$ whose tightest 3-digit enclosure is $[1.00, 1.01]$), while $z_2 = \text{innerMinus}(y, x)$ is not defined. However, one cannot discriminate these cases using naive 3 digit arithmetic. Comparing $\text{width}(x)$ with $\text{width}(y)$ gives the wrong result, because both are computed (rounding upward) as 1.01, suggesting z_2 is defined. Computing the bounds of z_2 , namely $[(-1.00 - .0001), (-.0002 - 1.00)]$ (with outward rounding), also gives the wrong result, namely $[-1.01, -1.00]$, again suggesting z_2 is defined.]

Only real or simulated higher precision is guaranteed to give the correct decision in all cases.

13.9. Implementation of bare object arithmetic. (informative)

▲ To be written.

13.10. Complete arithmetic, dot product function. An implementation that provides 754-conforming types should provide **complete arithmetic**, as specified in U. Kulisch and V. Snyder [4], for the parent format \mathbb{F} .

This involves providing a *complete format* datatype $C(\mathbb{F})$ associated with the relevant \mathbb{F} , and associated operations. A $C(\mathbb{F})$ datum z holds a fixed-point number of the relevant radix (2 or 10), with enough digits before and after the point to let multiply-add operations $z + x * y$ be done exactly, where x and y are arbitrary finite \mathbb{F} -numbers. It also holds one bit for sign, and 3 bits for status information (equivalent to a decoration).

[Example. For the binary64 format the recommended complete format has 4 bits for sign and status, 2134 bits before the point, and 2150 after the point, for a total of 4288 bits or 536 bytes; this allows for at least 2^{88} multiply-adds before overflow can occur.]

The following operations should be provided, see [4] for details.

- **convert** converts from a complete format to a floating point format, or vice versa, or from one complete format to another.
- **completeAdd** and **completeSub** add or subtract two complete or floating point format operands, of which at least one is complete, giving a complete format result.
- **completeMulAccum** computes $z + x * y$ where z has a complete format and x, y are of floating point format, giving a complete format result.
- **completeDotProduct**. Let a and b be vectors of length n holding floating point numbers of format \mathbb{F} . Then **completeDotProduct**(a, b) computes $a \cdot b = \sum_{k=1}^n a_k b_k$ exactly, giving a complete format result.

The result of all operations may be converted if necessary to a specified result format by application of the **convert** operation.

14. Level 4 description

⚠ Probably does not need to exist in this standard.



DRAFT 7.3

DRAFT 7.3

CHAPTER 3

Kaucher Intervals

This Chapter contains the standard for the Kaucher interval flavor.
To be included.

DRAFT 7.3

DRAFT 7.3

ANNEX A

Details of flavor-independent requirements

A.1. List of required functions

TBW

A.2. List of recommended functions

TBW

DRAFT 7.3

DRAFT 7.3

ANNEX B

Including a new flavor in the standard

To be written following discussion with the IEEE MSC working group.

DRAFT 7.3

DRAFT 7.3

Reproducibility

To be written.

⚠ Moved from earlier, hoping it is useful here.

[*Example. Suppose the set-based and Kaucher flavors co-operate by sharing a type \mathbb{T} whose intervals have lower and upper bounds that are `binary64` floating point numbers. Suppose they implement some subset of \mathbb{T} 's operations in "tightest" mode, returning the smallest \mathbb{T} -interval that encloses the Level 1 result. This specification makes such operations flavor-independent (when acting on common intervals). Then any common evaluation, that uses only this shared type and this subset of operations, gives identical results in both flavors, modulo the embedding map.*]

DRAFT 7.3

DRAFT 7.3

Set-based flavor: decoration details and examples

D.1. Local decorations of arithmetic operations

D.1.1. Forward-mode elementary functions. For each of the required functions φ of §9.6, with the decoration scheme `com > dac > def > trv > ill` of Clause 10, Tables 1 to 2 give the **strongest local decoration** for arbitrary interval inputs. That is, they give $\text{dec}(\varphi, \mathbf{x})$ for an arbitrary input box \mathbf{x} . The following facts are used to shorten the tables:

- If any input is empty, the decoration is `trv`, so the tables may assume nonempty inputs.
- Functions $\varphi(x_1, x_2, \dots)$ that are defined and continuous at all real arguments can be handled in a uniform way. This covers the required functions `neg`, `add`, `sub`, `mul`, `fma`, `sqr`, `pown(x, p)` for $p \geq 0$, `exp` and its variants, `sin`, `cos`, `atan`, `sinh`, `cosh`, `tanh`, `asinh` and `abs`, together with `min` and `max` of any number of arguments.

The functions φ in Table 2 have discontinuities at points within their domain of definition. Hence, one must note a distinction between `dac`, which requires that the restriction of φ to the input box \mathbf{x} be continuous, and `com`, which makes the stronger requirement that φ be continuous at each point of \mathbf{x} . [Example. For `floor(x)` on $[0, \frac{1}{2}]$, `dac` is true and `com` is false.] For these functions, finding the tightest interval enclosure of the range, and the local decoration, is simplified by noting that all are *increasing step functions*, that is, each one satisfies $\varphi(u) \leq \varphi(v)$ if $u \leq v$, and takes only finitely many values in any bounded interval. Further, each one is defined on the whole real line. For such an φ on an interval $[\underline{x}, \bar{x}]$ it is easy to see that

- (a) The restriction of φ to \mathbf{x} is continuous iff $\varphi(\underline{x}) = \varphi(\bar{x})$.
- (b) φ is continuous at each point of \mathbf{x} iff $\varphi(\underline{x}) = \varphi(\bar{x})$ and neither \underline{x} nor \bar{x} is a jump point of φ .

This gives a simple algorithm (given in the Table) for the range and local decoration. It relies only on φ itself and the set J of jump points of φ , so Table 2 merely displays the set J for each function.

D.1.2. Interval case function.

△ JDP March 2013. This is probably wrong now. Subclause 9.6.4 defines the function `case(c, g, h)`, and its required bare interval extension. It propagates decorations like other arithmetic operations, with local decoration d where

$$d = \begin{cases} \text{if } \mathbf{c} \text{ is empty} & \text{then } \text{trv} \\ \text{elseif } \mathbf{c} \text{ is a subset of the half- line } x < 0 \\ \text{or } \mathbf{c} \text{ is a subset of the half- line } x \geq 0 & \text{then } \text{dac} \\ \text{else} & \text{def.} \end{cases}$$

[Note. Comparisons or overlap relations, as a mechanism for handling cases, are incompatible with the decoration concept since there is no way to account for exceptions. The case function handles decorations correctly. However, in most cases, functions defined using it give very suboptimal enclosures, and it is preferable to use methods illustrated in §10.10.]

D.2. Examples of use of decorations

This subclause gives a number of examples intended to clarify decoration concepts and algorithms.

1. If $n = 1$, and f is the square root function, then the strongest decoration of $(f, [0, 1])$ is `dac`; of $(f, [-1, 1])$ is `trv`; and of $(f, [-2, -1])$ is `emp`. The expression $f(x) = \sqrt{-1 - x^2}$, as a real function, has no value for any x , so $\text{dec}(f, \mathbf{x}) = \text{ill}$ for all \mathbf{x} —though evaluation can never find this value, see examples below.

TABLE 1. Local decorations of required forward elementary functions. Normal mathematical notation is used to include or exclude an interval endpoint, e.g., $(-1, 1]$ denotes $\{x \in \mathbb{R} \mid -1 < x \leq 1\}$. The specification for each function is written as a set of mutually exclusive cases.

Function φ	Strongest local decoration, for all inputs nonempty
Everywhere continuous $\varphi(x_1, x_2, \dots)$	$\left\{ \begin{array}{l} \text{com} \text{ if inputs bounded, and result bounded at Level 2;} \\ \text{dac} \text{ otherwise.} \end{array} \right.$
$\text{div}(x, y)$	$\left\{ \begin{array}{l} \text{com} \text{ if } 0 \notin \mathbf{y}, \text{ inputs bounded, and result bounded at Level 2;} \\ \text{trv} \text{ if } 0 \in \mathbf{y}; \\ \text{dac} \text{ otherwise.} \end{array} \right.$
$\text{recip}(x)$	$\left\{ \begin{array}{l} \text{com} \text{ if } 0 \notin \mathbf{x}, \mathbf{x} \text{ bounded, and result bounded at Level 2;} \\ \text{trv} \text{ if } 0 \in \mathbf{x}; \\ \text{dac} \text{ otherwise.} \end{array} \right.$
$\text{sqrt}(x)$	$\left\{ \begin{array}{l} \text{trv} \text{ if } \mathbf{x} \not\subseteq [0, +\infty); \\ \text{com} \text{ if } \mathbf{x} \subseteq [0, +\infty), \mathbf{x} \text{ bounded, and result bounded at Level 2;} \\ \text{dac} \text{ otherwise.} \end{array} \right.$
$\text{case}(c, g, h)$	To be done.
$\text{pown}(x, p), p \geq 0$	“Everywhere continuous” case.
$\text{pown}(x, p), p < 0$	$\left\{ \begin{array}{l} \text{com} \text{ if } 0 \notin \mathbf{x}, \mathbf{x} \text{ bounded, and result bounded at Level 2;} \\ \text{trv} \text{ if } 0 \in \mathbf{x}; \\ \text{dac} \text{ otherwise.} \end{array} \right.$
$\text{pow}(x, y)$	$\left\{ \begin{array}{l} \text{trv} \text{ if } (\mathbf{x}, \mathbf{y}) \not\subseteq \mathcal{D}; \\ \text{com} \text{ if } (\mathbf{x}, \mathbf{y}) \subseteq \mathcal{D}, \text{ inputs bounded, and result bounded at Level 2;} \\ \text{dac} \text{ otherwise;} \end{array} \right.$ <p>where $\mathcal{D} = \{(x, y) \mid x > 0, \text{ or } x = 0 \text{ and } y > 0\}$.</p>
$\text{log}, \text{log2}, \text{log10}(x)$	$\left\{ \begin{array}{l} \text{trv} \text{ if } \mathbf{x} \not\subseteq (0, +\infty); \\ \text{com} \text{ if } \mathbf{x} \subseteq (0, +\infty), \mathbf{x} \text{ bounded, and result bounded at Level 2;} \\ \text{dac} \text{ otherwise.} \end{array} \right.$
$\text{tan}(x)$	$\left\{ \begin{array}{l} \text{trv} \text{ if } \mathbf{x} \not\subseteq \mathcal{D}; \\ \text{com} \text{ if } \mathbf{x} \subseteq \mathcal{D}, \mathbf{x} \text{ bounded, and result bounded at Level 2;} \\ \text{dac} \text{ otherwise.} \end{array} \right.$ <p>where $\mathcal{D} = \mathbb{R} \setminus \{\text{odd multiples of } \pi/2\}$.</p>
$\text{asin}(x), \text{acos}(x)$	$\left\{ \begin{array}{l} \text{com} \text{ if } \mathbf{x} \subseteq [-1, 1]; \\ \text{trv} \text{ otherwise.} \end{array} \right.$
$\text{atan2}(y, x)$	$\left\{ \begin{array}{l} \text{trv} \text{ if } (\mathbf{x}, \mathbf{y}) \not\subseteq \mathcal{D}; \\ \text{com} \text{ if } (\mathbf{x}, \mathbf{y}) \subseteq \mathcal{D}, \text{ inputs bounded, and result bounded at Level 2;} \\ \text{dac} \text{ otherwise;} \end{array} \right.$ <p>where $\mathcal{D} = \mathbb{R}^2 \setminus \{(x, 0) \mid x \leq 0\}$.</p> <p>Note reversal of arguments y, x compared with mathematical definition x, y.</p>
$\text{acosh}(x)$	$\left\{ \begin{array}{l} \text{trv} \text{ if } \mathbf{x} \not\subseteq [1, +\infty); \\ \text{com} \text{ if } \mathbf{x} \subseteq [1, +\infty), \mathbf{x} \text{ bounded, and result bounded at Level 2;} \\ \text{dac} \text{ otherwise.} \end{array} \right.$
$\text{atanh}(x)$	$\left\{ \begin{array}{l} \text{trv} \text{ if } \mathbf{x} \not\subseteq (-1, 1); \\ \text{com} \text{ if } \mathbf{x} \subseteq (-1, 1), \text{ and result bounded at Level 2;} \\ \text{dac} \text{ otherwise.} \end{array} \right.$

2. For a function defined by an expression, finding the strongest decoration over a box is typically hard in the same way and for the same reasons that finding the tightest interval enclosure of the exact range is hard. Straightforward interval evaluation usually does not find it. A trivial example is the expression $f(x) = \sqrt{x - x}$. As a real function it gives $f(x) = 0$, which is continuous for all x , so that $\text{dec}(f, \mathbf{x}) = \text{dac}$ for any nonempty interval \mathbf{x} . But for any \mathbf{x} of more than one point, evaluating $f(\mathbf{x})$ as in §10.6 gives trv as the decoration, because it takes the square root of an interval containing negative points.

TABLE 2. Required forward elementary functions: step functions. The set of jump points is shown. The range and local decoration are computed by the algorithm below.

Function φ	Set J of jump points of φ .
$\text{sign}(x)$	$J = \{0\}$
$\text{ceil}(x), \text{floor}(x)$	$J = \mathbb{Z}$
$\text{trunc}(x)$	$J = \mathbb{Z} \setminus \{0\}$
$\text{roundTiesToEven}(x), \text{roundTiesToAway}(x)$	$J = \{n + \frac{1}{2} \mid n \in \mathbb{Z}\}$

At an infinite endpoint, the value of φ is taken as its limiting value, e.g. $\text{sign}(-\infty) = -1$, $\text{ceil}(+\infty) = +\infty$.

Input: nonempty $\mathbf{x} = [\underline{x}, \bar{x}]$, possibly unbounded.

$\underline{y} = \varphi(\underline{x})$, $\bar{y} = \varphi(\bar{x})$

if $\underline{y} = \bar{y}$

 if $\underline{x} \notin J$ and $\bar{x} \notin J$ and \mathbf{x} is bounded

$d = \text{com}$

 else

$d = \text{dac}$

 end if

else

$d = \text{def}$

end if

Output: range enclosure $[\underline{y}, \bar{y}]$ and local decoration d .

Similarly, the computed decoration of $f(x) = \sqrt{-1 - x^2}$ in the example above will be **emp** or **trv** for any \mathbf{x} , never the correct **ill**.

Note also that though **trv** is trivial in itself, to have $\text{dec}(f, \mathbf{x}) = \text{trv}$ is not trivial: it asserts p_{emp} and p_{def} are both false. The first implies that \mathbf{x} has a point in $\text{Dom}(f)$, the second that \mathbf{x} has a point outside $\text{Dom}(f)$; together these imply \mathbf{x} is an interval of positive length.

3. Consider exact arithmetic DIE of $f(x, y) = \sqrt{x(y-x)} - 1$ with various input intervals \mathbf{x}, \mathbf{y} . Finite precision would produce valid but usually slightly different results. The natural domain $\text{Dom}(f)$ is easily seen to be the union of the regions $x > 0, y \geq x + 1/x$ and $x < 0, y \leq x + 1/x$ in the plane.

- (i) Let $\mathbf{x} = [1, 2]$, $\mathbf{y} = [3, 4]$, defining a box (\mathbf{x}, \mathbf{y}) contained in $\text{Dom } f$. Applying the **newDec** function gives initial decorated intervals $\mathbf{x}_{dx} = [1, 2]_{\text{dac}}$, $\mathbf{y}_{dy} = [3, 4]_{\text{dac}}$. The first operation is

$$\mathbf{u}_{du} = \mathbf{y}_{dy} - \mathbf{x}_{dx} = [1, 3]_{\text{dac}}.$$

Namely, subtraction is defined and continuous on all of \mathbb{R}^2 , and bounded on bounded rectangles (call this property “nice” for short), so the bare result decoration is $du' = \text{dec}(-, (\mathbf{y}, \mathbf{x})) = \text{dac}$, whence by (26) the decoration on \mathbf{u} is $du = \min\{du', dy, dx\} = \min\{\text{dac}, \text{dac}, \text{dac}\} = \text{dac}$. Multiplication is also “nice”, so the second operation similarly gives

$$\mathbf{v}_{dv} = \mathbf{x}_{dx} \times \mathbf{u}_{du} = [1, 6]_{\text{dac}}.$$

The constant 1, following §9.4.4, becomes a decorated interval function returning the constant value $[1, 1]_{\text{dac}}$. The next operation is again “nice”, and gives

$$\mathbf{w}_{dw} = \mathbf{v}_{dv} - 1 = [0, 5]_{\text{dac}}$$

Finally $\sqrt{\cdot}$ is defined, continuous and bounded on $\mathbf{w} = [0, 5]$, so, arguing similarly, one has the final result

$$\mathbf{f}_{df} = \sqrt{\mathbf{w}_{dw}} = [0, \sqrt{5}]_{\text{dac}}.$$

By the FTDIA it is thus proven that for the box $\mathbf{z} = (\mathbf{x}, \mathbf{y}) = ([1, 2], [3, 4])$,

$$[0, \sqrt{5}] \supseteq \text{Rge}(f | \mathbf{z}),$$

$$p_{\text{dac}}(f, \mathbf{z}) \text{ holds.}$$

That is, f is defined, continuous and bounded on $1 \leq x \leq 2$, $3 \leq y \leq 4$, and its range over this box is a subset of $[0, \sqrt{5}]$.

- (ii) Let $\mathbf{x} = [1, 2]$ as before, but $\mathbf{y} = [\frac{5}{2}, 4]$. The box \mathbf{z} is still contained in $\text{Dom } f$ so the true value of $\text{dec}(f, \mathbf{z})$ is still dac . However the evaluation fails to detect this because of interval widening due to the dependence problem of interval arithmetic. Namely after $\mathbf{u}_{du} = [\frac{5}{2}, 3]_{\text{dac}}$, $\mathbf{v}_{dv} = [\frac{5}{2}, 6]_{\text{dac}}$, $\mathbf{w}_{dw} = [-\frac{1}{2}, 5]_{\text{dac}}$, the final result has interval part $\mathbf{f} = \sqrt{[-\frac{1}{2}, 5]} = [0, \sqrt{5}]$ as before, but $\sqrt{\cdot}$ is not everywhere defined on \mathbf{w} , so that $d\mathbf{w}' = \text{dec}(\sqrt{\cdot}, \mathbf{w}) = \text{dec}(\sqrt{\cdot}, [-\frac{1}{2}, 5]) = \text{trv}$ giving $\text{dec}(\sqrt{\cdot}, \mathbf{w}_{dw}) = \min\{d\mathbf{w}', dv\} = \text{trv}$, so finally $\mathbf{f}_{df} = [0, \sqrt{5}]_{\text{trv}}$. This is a valid enclosure of the decorated range $[0, \sqrt{5}]_{\text{dac}}$, but we have been unable to verify the dac property.
- (iii) If $\mathbf{x} = [1, 2]$, $\mathbf{y} = [1, 1]$, the box \mathbf{z} is now wholly outside $\text{Dom } f$, and evaluation detects this, giving the exact result $\mathbf{f}_{df} = \emptyset_{\text{emp}}$. However, if $\mathbf{x} = [1, 2]$, $\mathbf{y} = [1, \frac{3}{2}]$, the box is still wholly outside $\text{Dom } f$, but owing to widening, evaluation fails to detect this, giving $\mathbf{f}_{df} = [0, 0]_{\text{trv}}$ —a valid enclosure but of little use.

D.3. Implementation of compressed interval arithmetic

Table 3 gives tables of compressed arithmetic, §10.12, for the four basic operations. Here c, d are bare decorations less than the threshold τ , and \mathbf{x}, \mathbf{y} are bare intervals. Independently of τ , if any input is the decoration ill the result is ill , else if any input is the interval \emptyset the result is \emptyset . The tables below give the remaining cases where

$$\text{trv} \leq c < \tau, \text{trv} \leq d < \tau, \text{ and } \mathbf{x}, \mathbf{y} \text{ are nonempty.} \quad (43)$$

TABLE 3. Compressed interval operations for $+, -, \times, \div$ and $\sqrt{\cdot}$ with threshold $\tau \in \{\text{trv}, \text{def}, \text{dac}, \text{com}\}$.

Binary operations, where \mathbf{x} or c is the left operand and \mathbf{y} or d is the right operand.

$+, -, \times$	\mathbf{y}	d
\mathbf{x}	Normal bare interval result	d
c	c	$\min(c, d)$

\div	$\mathbf{y} = [0, 0]$	$0 \in \mathbf{y} \neq [0, 0]$	$0 \notin \mathbf{y}$	d
\mathbf{x}	emp	If $\tau > \text{trv}$ then trv , else normal bare interval result	Normal bare interval result	trv
c	emp	trv	c	trv

Square root, where $\mathbf{x} = [\underline{x}, \bar{x}]$.

	case	
$\sqrt{\mathbf{x}}$	$\bar{x} < 0$	emp
	$\underline{x} < 0 \leq \bar{x}$	If $\tau > \text{trv}$ then trv , else normal bare interval result
	$\underline{x} \geq 0$	Normal bare interval result
\sqrt{c}		trv

Some examples of compressed arithmetic follow. In items (b) onwards, conditions (43) are assumed.

- (a) Justification for $\mathbf{emp} + \mathbf{x} = \mathbf{emp}$, independent of τ .
This promotes to $(\emptyset, \mathbf{emp}) + (\mathbf{x}, \tau) = (\emptyset, \min(\mathbf{emp}, \tau, \mathbf{emp}))$. Since $\mathbf{emp} < \tau$ this equals $(\emptyset, \mathbf{emp})$ which gives an exception (again because $\mathbf{emp} < \tau$) so is recorded as the bare decoration \mathbf{emp} . The same holds if $+$ is replaced by $-$, \times or \div .
- (b) Justification for $\mathbf{x} \times d = d$ independent of τ .
Since \mathbf{x} is nonempty and $d \geq \mathbf{trv}$, this promotes to $(\mathbf{x}, \tau) \times (\mathbf{y}, d)$ with arbitrary nonempty \mathbf{y} , giving $(\mathbf{x} \times \mathbf{y}, \min(\tau, d, e))$ where e is \mathbf{dac} if $\mathbf{x} \times \mathbf{y}$ is bounded, otherwise \mathbf{def} . Now $d < \tau$ so d cannot exceed \mathbf{def} , hence $d \leq e$, so $\min(\tau, d, e) = d$.
- (c) Justification for $c/d = \mathbf{trv}$ independent of τ .
Since $c, d \geq \mathbf{trv}$, c/d promotes to $(\mathbf{x}, c)/(\mathbf{y}, d)$ with arbitrary nonempty \mathbf{x}, \mathbf{y} , giving $(\mathbf{x}/\mathbf{y}, \min(c, d, e))$ where $e = \mathbf{emp}$ if $\mathbf{y} = [0, 0]$, else $e = \mathbf{trv}$ if $0 \in \mathbf{y}$, else $e = \mathbf{dac}$. So $\min(c, d, e) \geq \mathbf{trv}$ and can equal \mathbf{trv} , so the tightest enclosing decoration is \mathbf{trv} .
- (d) Justification for \mathbf{x}/\mathbf{y} when $0 \in \mathbf{y} \neq [0, 0]$.
 \mathbf{x}/\mathbf{y} promotes to $(\mathbf{x}, \tau)/(\mathbf{y}, \tau)$ giving $(\mathbf{x}/\mathbf{y}, \min(\tau, \tau, \mathbf{trv})) = (\mathbf{x}/\mathbf{y}, \mathbf{trv})$. If $\tau > \mathbf{trv}$ this gives an exception so the decoration \mathbf{trv} is returned; if $\tau = \mathbf{trv}$ it is not an exception, so the interval \mathbf{x}/\mathbf{y} is returned.
- (e) Justification for $\sqrt{\mathbf{x}}$ with $\mathbf{x} = [\underline{x}, \bar{x}]$ and $\underline{x} < 0 \leq \bar{x}$.
 $\sqrt{\mathbf{x}}$ promotes to $\sqrt{(\mathbf{x}, \tau)}$, giving $(\sqrt{\mathbf{x}}, \mathbf{trv})$ which in the given case equals $([0, \sqrt{\bar{x}}], \mathbf{trv})$. As with the previous item, if $\tau > \mathbf{trv}$ then \mathbf{trv} is returned; if $\tau = \mathbf{trv}$ then $\sqrt{\mathbf{x}}$ is returned.

D.4. Proofs of correctness for compressed interval arithmetic

⚠ To be completed. ■

DRAFT 7.3

D.5. The fundamental theorem of decorated interval arithmetic

We assume the seven-decoration set \mathbb{D} of decorations defined by (17). However the proof of the fundamental theorem is largely independent of the particular set of decorations chosen.

It is necessary first to clarify the case of a zero-argument arithmetic operation φ , which represents a real constant. A point argument of a general k -ary φ is a tuple $u = (u_1, \dots, u_k) \in \mathbb{R}^k$. For $k = 0$ this is the empty tuple $()$, which is the unique element of \mathbb{R}^0 .

For an interval version, an argument for general k is $\mathbf{u} = (\mathbf{u}_1, \dots, \mathbf{u}_k)$, representing the subset of \mathbb{R}^k specified by k constraints $u_1 \in \mathbf{u}_1, \dots, u_k \in \mathbf{u}_k$. For $k = 0$ there are no such constraints, so the input “box” to an interval version cannot be empty: it is always the whole of $\mathbb{R}^0 = \{()\}$.

However φ can have empty domain, in which case it is the “Not a Number” function NaN; otherwise its domain is \mathbb{R}^0 and it has a real value. Clearly, if φ is NaN then $p_{\text{in}}(\varphi, \mathbb{R}^0)$ holds, otherwise $p_{\text{bnd}}(\varphi, \mathbb{R}^0)$ holds.

We now prove:

Theorem D.5.1 (Fundamental Theorem of Decorated Interval Arithmetic, FTDIA).

Let $\mathbf{f}_{df} = f(\mathbf{x})$ be the result of evaluating an arithmetic expression $f(z_1, \dots, z_n)$ over a bare box $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n) \in \overline{\mathbb{R}}^n$ using any decorated interval version f of f . Then in addition to the enclosure

$$\mathbf{f} \supseteq \text{Rge}(f | \mathbf{x}) \quad (44)$$

given by Moore’s FTIA Theorem (page 12), we have

$$p_{df}(f, \mathbf{x}) \text{ holds.} \quad (45)$$

Proof. The case where \mathbf{x} is empty is a special case. By case (Eval1) of the definition of a decorated interval version in §10.6, $\mathbf{f}_{df} = \emptyset_{\text{ein}}$. Also $\text{Rge}(f | \mathbf{x}) = \emptyset$ and by definition, $p_{\text{ein}}(f, \mathbf{x})$ holds, so that (44, 45) hold.

Otherwise \mathbf{x} is nonempty (so each of its components is nonempty) and we proceed by induction on the number of operations in f .

The base case, where this number is zero, is that f is a variable, say z_i . Then it defines the function $f(x) = x_i$. By case (Eval2) in §10.6, $\mathbf{f} = \mathbf{x}_i = \text{Rge}(f | \mathbf{x})$ and df is such that $p_{df}(\text{Id}, \mathbf{x}_i)$ holds, where Id is the identity function $\text{Id}(x) = x$ on \mathbb{R} (specifically, $df = \text{bnd}$ if \mathbf{x}_i is bounded, and $df = \text{dac}$ otherwise). Then $p_{df}(f, \mathbf{x}) = p_{df}(\text{Id}, \mathbf{x}_i)$ holds. Thus (44, 45) hold.

Otherwise \mathbf{x} is nonempty and $f = \varphi(g_1, \dots, g_k)$ where φ is an arithmetic operation of arity $k \geq 0$, and the g_i are expressions having fewer operations than does f . When f and the g_i are regarded as point functions, this means $f(x) = \varphi(g_1(x), \dots, g_k(x))$ for $x \in \mathbb{R}^n$.

By the inductive hypothesis the theorem holds for each g_i . (The case $k = 0$, where φ is a real constant or NaN, needs no special treatment.) So (44, 45) applied to g_i give for $i = 1, \dots, k$

$$\mathbf{g}_i \supseteq \text{Rge}(g_i | \mathbf{x}), \quad (46)$$

$$p_{dg_i}(g_i, \mathbf{x}) \text{ holds.} \quad (47)$$

By the definition of a decorated interval version of f , \mathbf{f}_{df} is computed using a decorated interval extension of φ , hence by the definition in §10.6,

$$\mathbf{f} \supseteq \text{Rge}(\varphi | \mathbf{g}), \quad (48)$$

$$df = \min\{d\varphi, dg_1, \dots, dg_k\} \quad (49)$$

for some $d\varphi$ such that

$$p_{d\varphi}(\varphi, \mathbf{g}) \text{ holds.} \quad (50)$$

We show first (44) and then (45). Denote here $u_k = g_i(x)$ for some $x \in \mathbf{x}$. Then

$$u_k = g_i(x) \in \text{Rge}(g_i | \mathbf{x}) \subseteq \mathbf{g}_i. \quad (51)$$

For any $v \in \text{Rge}(f | \mathbf{x})$, there is $x \in \mathbf{x}$ such that $v = f(x)$. Then, using (51, 48),

$$\begin{aligned} v = f(x) &= \varphi(g_1(x), \dots, g_k(x)) \\ &= \varphi(u_1, \dots, u_k) \\ &= \varphi(u) \in \text{Rge}(\varphi | \mathbf{g}) \subseteq \mathbf{f}. \end{aligned}$$

Since v was arbitrary, this proves (44).

It remains to prove (45). Corresponding to the different meanings of the decorations, this is verified on a case by case basis, starting with the least decoration.

Case $df = \text{i11}$. Then either some $dg_i = \text{i11}$ or $d\varphi = \text{i11}$.

- If $dg_i = \text{i11}$, by (47) $\text{Dom } \varphi$ is empty.
- If $d\varphi = \text{i11}$, hence by (50) $\text{Dom } \varphi$ is empty.

In either case, by the definition of the point function f , $\text{Dom } f$ is empty so (45) holds.

Case $df = \text{emp}$. Then either some $dg_i = \text{emp}$ or $d\varphi = \text{emp}$.

- If $dg_i = \text{emp}$, by (47) \mathbf{x} is disjoint from $\text{Dom } g_i$.
- If $d\varphi = \text{emp}$, by (50) \mathbf{g} is disjoint from $\text{Dom } \varphi$.

In either case there is no $x \in \mathbf{x}$ for which $f(x)$ is defined. That is, \mathbf{x} is disjoint from $\text{Dom } f$, and (45) holds.

Case $df = \text{trv}$. This is always true, and nothing needs to be shown.

Case $df = \text{def}$. Then each $dg_i \geq \text{def}$, and $d\varphi \geq \text{def}$. Thus by (47, 50), each g_i is everywhere defined on \mathbf{x} , with values in \mathbf{g}_i by (46), and φ is everywhere defined on \mathbf{g} . Hence f is everywhere defined on \mathbf{x} so again (45) holds.

Case $df = \text{dac}$. This is as the **def** case with the addition that the restriction of each g_i to \mathbf{x} is everywhere defined and continuous, and the restriction of φ to \mathbf{g} is everywhere defined and continuous. Hence the restriction of f to \mathbf{x} is everywhere defined and continuous so again (45) holds.

Case $df = \text{bnd}$. Then each $dg_i \geq \text{bnd}$, and $d\varphi \geq \text{bnd}$. By similar reasoning, the restriction of f to \mathbf{x} is everywhere defined, continuous and bounded so again (45) holds.

(In fact all one needs to deduce this is that each $dg_i \geq \text{dac}$, and $d\varphi \geq \text{bnd}$.)

Case $df = \text{ein}$. This cannot occur since \mathbf{x} was assumed nonempty.

Hence all cases have been covered. This completes the induction step and the proof. \square

DRAFT

DRAFT 7.3

Further material for set-based standard (informative)

This may possibly be included in some form in the Annexes to the standard.

E.1. Type conversion in mixed operations

⚠ This needs checking by language and compiler experts and should be the subject of a separate motion. Also, does it all belong in the decorations section? Indeed should it be in Level 1 at all?

Decorated interval arithmetic is designed for maximal safety, while being simple to handle by inexperienced users. Safety requirements can be enforced only by restrictions on the kinds of type conversions permitted.

Operations between integers and decorated intervals are well-defined and hence permitted, with integers treated as constant functions.

Operations between floats and decorated intervals are error-prone and hence forbidden, since, e.g., $(2/3) * x$ in program text would generate uncovered roundoff, and $0.2 * x$ would generate uncovered conversion errors. This ensures that the user must call explicitly a conversion function **iconst** that performs the outward rounding, see §12, to convey the precise semantics of such mixed expressions. This avoids a loss of containment because of rounding errors or conversion errors.

In particular, there is no implicit type casting for real times decorated interval. Therefore, $2/3 * x$ with reals or integers 2 and 3 and a decorated interval x results in a type error when trying to evaluate the multiplication.

However, implicit type casting for text constants times interval is harmless, as text constants have no arithmetic operations defined on them, hence they can be unambiguously type cast to decorated intervals when occurring in an interval expression if the implementation language allows that. Therefore, $2/3 * x$ is allowed if the compiler translates 2 and 3 into constant functions.

Mixed operations between bare intervals and decorated intervals are also forbidden, to avoid loss of rigor through non-arithmetic operations; again, explicit conversion using the function `newDec` must be used. However, explicit, constant bare intervals in program code may be treated by the compiler as constant functions with uncertain value when the bare interval is nonempty, and as the ill-formed constant when the bare interval is empty or ill-formed.

E.2. The “Not an Interval” object

⚠ TO BE REVISED

From §9.4.4, a real scalar function with no arguments—a mapping $\mathbb{R}^n \rightarrow \mathbb{R}^m$ with $n = 0$ and $m = 1$ —is a **real constant**.

This specification of constants gives a Level 1 definition of NaN, “Not a Number”—not as a value, but as a constant function. \mathbb{R}^0 is the zero-dimensional vector space $\{0\}$ —it has one element, conventionally named 0. The real numbers c are in one-to-one correspondence with the mappings $c() : 0 \mapsto c$, so that \mathbb{R} can be identified with the *total* functions $\mathbb{R}^0 \rightarrow \mathbb{R}$. There is one *non-total* $c()$, the function `NaN()` with empty domain and, therefore, no value.

From the definition in §9.4.3, an interval extension of a real constant with value c is any zero-argument interval function that returns an interval containing c . The *natural extension* returns the interval $[c, c]$.

Its natural interval extension is the constant interval function whose value is the empty interval. the zero-argument function with empty domain is the real constant function with value NaN, “Not a Number”. It is easily seen that NaN’s natural interval extension is the interval constant function with value \emptyset , and its natural decorated interval extension is the decorated interval constant function with value `NaNI` = (\emptyset, ill) . (This was pointed out by Arnold Neumaier.)

The decorated interval NaI has behaviour that qualifies it for the role of “Not an Interval”. By definition it signals that it is the result of evaluating a null function, with empty domain.

It is returned by any invalid call to an interval constructor, such as “the interval from 3 to NaN”. It is unconditionally “sticky” within arithmetic expressions, in the sense that if any argument to an arithmetic operation is NaI, then that operation’s output is NaI.

However, it cannot be generated “new” during evaluation of any expression that uses normal operations, even if the theoretical function being defined has empty domain. For example, the expression

$$f(x) = \sqrt{-1 - x^2}$$

clearly defines, over the reals, a function with empty domain; but decorated interval evaluation can *never* notice this. With any non-NaI input, it will return $(\emptyset, \mathbf{emp})$ and not $(\emptyset, \mathbf{ill})$.

Hence, in practice, NaI behaves as one expects it to do: it records the “taint of illegitimacy” of an interval’s ancestry. A decorated interval is NaI iff it is the result of an ill-formed construction or is the computational descendant of such a result.

DRAFT 7.3

Bibliography

- [1] Allen, James F. Maintaining knowledge about temporal intervals. *Communications of the ACM* 26, 832–843, (November 1983).
- [2] Kulisch, Ulrich. Complete Interval Arithmetic and its Implementation on the Computer. Position paper, and the Dagstuhl 2008 proceedings.
- [3] Kulisch, Ulrich. *Computer Arithmetic and Validity: Theory, Implementation, and Applications*. de Gruyter, Berlin, New York, (2008).
- [4] Kulisch, Ulrich and Snyder, Van. *The exact dot product as basic tool for long interval arithmetic*. Position paper, P1788 Working Group, version 11, July 2009.
- [5] Moore, Ramon E. *Interval Analysis*. Prentice-Hall, Englewood Cliffs, N.J., (1966).
- [6] Nehmeier, Marco and Siegel, Stefan and Wolff von Gudenberg, Jürgen. Specification of hardware for interval arithmetic. *Computing* 94, 243-255, (2012).
- [7] Neumaier, Arnold. Vienna Proposal for Interval Standardization. Faculty of Mathematics, University of Vienna, (December 2008). <http://www.mat.univie.ac.at/~neum>
- [8] Pryce, John D. and Corliss, George F. Interval arithmetic with containment sets. *Computing* 78, 251–276, (2006).
- [9] Pryce, John D. P1788 Motion 6: Multi-Format Support: Text and Rationale, (2009).