

6. Level 2 description

6.1 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**⁴.

6.1.1 Interval type

The interval type, denoted by \mathbb{T} , is the **inf-sup type** derived from the IEEE 754 **binary64** format; we refer to the latter as **b64**. This interval type comprises all intervals whose endpoints are **b64** numbers, together with Empty. Since $\pm\infty$ are in **b64**, Entire is in \mathbb{T} .

An interval from \mathbb{T} is also called a **bare interval** or a **T-interval**. We use the term **T-datum** to refer to an entity that can be a T-interval or a NaI. A **T-box** is a vector with T-datum components.

6.1.2 Decorated interval type

The decorated interval type, derived from \mathbb{T} , is the set of tuples (x, d) , where $x \in \mathbb{T}$, and $d \in \mathbb{D}$. We denote this type by \mathbb{DT} .

\mathbb{DT} shall contain a “Not an Interval” datum NaI, which is identified with $(\emptyset, \text{il1})$.

6.1.3 Operations

The term **T-version** of a Level 1 operation denotes one in which any input or output that is an interval is a T-datum. For bare interval types this includes the following.

- An interval extension (see 6.4) of one of the arithmetic operations of 4.5.
- A set operation, such as intersection and convex hull of T-intervals, returning a T-interval.
- A function such as the midpoint, whose input is a T-interval and output is numeric.
- A constructor, whose input is numeric or text and output is a T-datum.

6.1.4 Exception behavior

For some operations, and some particular inputs, there might not be a valid result. At Level 1 there are several cases when no value exists. However, a Level 2 operation always returns a value. When the Level 1 result does not exist, the corresponding Level 2 operation returns either

- a special value indicating this event (e.g., NaN for most of the numeric functions in 6.7.6); or
- a value considered reasonable in practice. For example, `mid(Entire)` returns 0; a constructor given invalid input returns Empty; and one of the comparisons of 6.7.7, if any input is NaI, returns **false**.

If `intervalPart()` is called with NaI as input, the exception `IntvlPartOfNaI` is signaled (see 6.7.8).

If a bare or decorated constructor fails (see 6.7.5) the exception `UndefinedOperation` is signaled.

If a bare or decorated constructor succeeds on a accuracy-relaxed string argument (see 6.7.5) the exception `PossiblyUndefinedOperation` may be signaled.

6.2 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 `add(x, y)`.

⁴Not “data”, whose common meaning could cause confusion.

6.3 Level 2 hull operation

6.3.1 Hull in one dimension

The **interval hull** operation

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

maps an arbitrary set of reals \mathbf{s} to the tightest interval \mathbf{y} enclosing \mathbf{s} .

6.3.2 Hull in n dimensions

In n dimensions the hull, as defined mathematically in 6.3.1, is extended to act componentwise. That is, for an arbitrary subset \mathbf{s} of \mathbb{R}^n it is $\text{hull}(\mathbf{s}) = (\mathbf{y}_1, \dots, \mathbf{y}_n)$, where

$$\mathbf{y}_i = \text{hull}(\mathbf{s}_i),$$

and $\mathbf{s}_i = \{s_i \mid s \in \mathbf{s}\}$ is the projection of \mathbf{s} on the i th coordinate dimension.

6.4 Level 2 interval extensions

Let f be an n -variable scalar point function. A **T-interval extension** of f , also called a **T-version** of f , is a mapping \mathbf{f} from n -dimensional T-boxes to 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

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

for any 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.

6.5 Accuracy of operations

This subclause describes requirements and recommendations on the accuracy of operations. Here, *operation* denotes any Level 2 version, provided by the implementation, of a Level 1 operation with interval output and at least one interval input. Bare interval operations are described; the accuracy of a decorated operation is defined to be that of its interval part.

6.5.1 Measures of accuracy

Three **accuracy modes** are defined that indicate the quality of interval enclosure achieved by an operation: *tightest*, *accurate* and *valid* in order from strongest to weakest.

The term **tightness** means the strongest mode that holds uniformly for some set of evaluations. For example, for some one-argument function, an implementation might document the tightness of $\mathbf{f}(\mathbf{x})$ as being *tightest* for all \mathbf{x} contained in $[-10^{15}, 10^{15}]$ and at least *accurate* for all other \mathbf{x} .

Let $\mathbf{f}_{\text{exact}}$ denote the corresponding Level 1 operation. The weakest mode *valid* is just the property of enclosure:

$$\mathbf{f}(\mathbf{x}) \supseteq \mathbf{f}_{\text{exact}}(\mathbf{x}). \quad (11)$$

The strongest mode *tightest* is the property that $\mathbf{f}(\mathbf{x})$ equals $\mathbf{f}_{\text{tightest}}(\mathbf{x})$, the hull of the Level 1 result:

$$\mathbf{f}_{\text{tightest}}(\mathbf{x}) = \text{hull}(\mathbf{f}_{\text{exact}}(\mathbf{x})). \quad (12)$$

The intermediate mode *accurate* asserts that $\mathbf{f}(\mathbf{x})$ is *valid*, (11), and is at most slightly wider than the result of applying the *tightest* version to a slightly wider input box:

$$\mathbf{f}(\mathbf{x}) \subseteq \text{nextOut}(\mathbf{f}_{\text{tightest}}(\text{nextOut}(\text{hull}(\mathbf{x})))) . \quad (13)$$

For an interval \mathbf{x} ,

$$\text{nextOut}(\mathbf{x}) = \begin{cases} [\text{nextDown}(\underline{x}), \text{nextUp}(\bar{x})] & \text{if } \mathbf{x} = [\underline{x}, \bar{x}] \neq \emptyset, \\ \emptyset & \text{if } \mathbf{x} = \emptyset, \end{cases}$$

where **nextUp** and **nextDown** are equivalent to the corresponding functions in IEEE 754.

When x is an interval box, `nextOut` acts componentwise.

NOTE—In (13), the inner `nextOut()` aims to handle the problem of a function such as $\sin x$ 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()` relaxes the requirement for correct (rather than, say, faithful) rounding, which might be hard to achieve for some special functions at some arguments.

6.5.2 Accuracy requirements

Following the categories of functions in Table 4.1, the accuracy of the *basic operations*, the *integer functions* and the *absmax functions* shall be *tightest*. The accuracy of the *cancellative addition and subtraction* operations of 4.5.3 is specified in 6.7.3.

For all other operations in Table 4.1, the accuracy should be *accurate*.

6.5.3 Documentation requirements

An implementation shall document the tightness of each of its interval operations. This shall be done by dividing the set of possible inputs into disjoint subsets (“ranges”) and stating a tightness achieved in each range.

[Example. Sample tightness information for the \sin function might be

| Operation | Tightness | Range |
|-----------|-----------------|---|
| \sin | <i>tightest</i> | for any $x \subseteq [-10^{15}, 10^{15}]$ |
| | <i>accurate</i> | for all other x . |

]

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

6.6 Number and interval literals

6.6.1 Number literals

The following forms of number literal shall be provided.

- a) A decimal number. This comprises an optional sign, a nonempty sequence of decimal digits optionally containing a point, and an optional exponent field comprising `e` and an integer literal.⁵ The value of a decimal number is the value of the sequence of decimal digits with optional point multiplied by ten raised to the power of the value of the integer literal, negated if there is a leading `-` sign.
- b) A number in the hexadecimal-floating-constant form of the C99 standard (ISO/IEC9899, N1256 (6.4.4.2)), equivalently hexadecimal-significant form of IEEE Std 754-2008 (5.12.3). This comprises an optional sign, the string `0x`, a nonempty sequence of hexadecimal digits optionally containing a point, and an exponent field comprising `p` and an integer literal exponent. The value of a hexadecimal number is the value of the sequence of hexadecimal digits with optional point multiplied by two raised to the power of the value of the exponent, negated if there is a leading minus sign.
- c) A rational literal p/q . This comprises an integer literal p , the `/` character, and a positive-natural literal q . Its value is the value of p divided by the value of q .
- d) Either of the strings `inf` or `infinity` optionally preceded by `+`, with value $+\infty$; or preceded by `-`, with value $-\infty$.

⁵An integer literal comprises an optional sign and followed by a nonempty sequence of decimal digits.

6.6.2 Bare intervals

The following forms of bare interval literal shall be supported. Below, the number literals l and r are identified with their values. Space shown between elements of a literal denotes zero or more space characters.

- A string $[l, u]$ where l and u are optional number literals of the same radix (10 or 16) with $l \leq u$, $l < +\infty$ and $u > -\infty$, see 4.2. Its bare value is the mathematical interval $[l, u]$. Any of l and u may be omitted, with implied values $l = -\infty$ and $u = +\infty$, respectively; e.g. $[,]$ denotes Entire.
- A string $[x]$ is equivalent to $[x, x]$.

- Uncertain form: a string $m ? r v E$ where: m is a decimal number literal of form a in 6.6.1, without exponent; r is empty or is a natural-number literal *ulp-count* or is $?$; v is empty or is a *direction character*, either **u** (up) or **d** (down); and E is empty or is an *exponent field* comprising the character **e** followed by an integer literal *exponent* e . No whitespace is permitted within the string.

With **ulp** meaning 10^{-d} where d is the number of digits after the decimal point in m (or 0 if there is no decimal point), the literal $m?$ by itself denotes m with a symmetrical uncertainty of half an 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 $[m - \frac{1}{2}\text{ulp}, m]$ and $m?ru$ denotes $[m, m + r \times \text{ulp}]$. Uncertain form with radius $?$ is for unbounded intervals, e.g., $m???d$ denotes $[-\infty, m]$. The exponent field if present multiplies the whole interval by 10^e , e.g., $m?ru\text{ee}$ denotes $10^e \times [m, m + r \times \text{ulp}]$.

- Special values: the strings $[]$ and $[\text{empty}]$, whose bare value is Empty, and the string $[\text{entire}]$, whose bare value is Entire.

6.6.3 Decorated intervals

The following forms of decorated interval literal shall be supported.

- ***sx_sd***: a bare interval literal ***sx***, an underscore “_”, and a 3-character decoration string ***sd***, where ***sd*** is one of **trv**, **def**, **dac** or **com**, denoting the corresponding decoration dx .

If ***sx*** has the bare value x , and if x_{dx} is a permitted combination according to 5.4, then ***sx_sd*** has the value x_{dx} . Otherwise ***sx_sd*** has no value as a decorated interval literal.

- The string $[\text{nai}]$, with the bare value Empty and the decorated value $\text{Empty}_{\text{ill}}$.

The alphanumeric characters in the above literals are case-insensitive (e.g., $[1,1\text{e}3]_{\text{com}}$ is equivalent to $[1,1\text{E}3]_{\text{COM}}$).

6.7 Required operations

Operations in this subclause are described as functions with zero or more input arguments and one return value. It is language-defined whether they are implemented in this way.

6.7.1 Interval constants

There shall be functions **empty()** and **entire()** returning an interval with value Empty and Entire, respectively. There shall also be a decorated version of each, returning

$$\text{newDec}(\text{Empty}) = \text{Empty}_{\text{trv}} \quad \text{and} \quad \text{newDec}(\text{Entire}) = \text{Entire}_{\text{dac}},$$

respectively.

6.7.2 Elementary functions

An implementation shall provide an interval version of each arithmetic operation in Table 4.1. Its inputs and output are intervals, and it shall be a Level 2 interval extension of the corresponding point function. Recommended accuracies are given in 6.5.

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.

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 5.6—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 (9) of 5.6, 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 (9).

6.7.3 Cancellative addition and subtraction

An implementation shall provide a \mathbb{T} -version of each of the operations `cancelMinus` and `cancelPlus` in 4.5.3. Their inputs and output are \mathbb{T} -intervals.

`cancelMinus(x, y)` shall return `Empty` in the first case of (2), the hull of the result in the second, and `Entire` for each of the cases in (3).

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

These operations shall have “trivial” decorated versions, as described in 5.7.

6.7.4 Set operations

An implementation shall provide an interval version of each of the operations `intersection` and `convexHull` in 4.5.4. Its inputs and output are intervals. These operations should return the interval hull of the exact result. If either input to `intersection` is `Empty`, or both inputs to `convexHull` are `Empty`, the result shall be `Empty`.

These operations shall have “trivial” decorated versions, as described in 5.7.

6.7.5 Constructors

For the bare and decorated interval types there shall be a constructor. It returns a \mathbb{T} - or \mathbb{DT} -datum, respectively.

Bare interval constructors. A bare interval constructor call either **succeeds** or **fails**. This notion is used to determine the value returned by the corresponding decorated interval constructor.

For the constructor `numsToInterval(l, u)`, the inputs l and u are **b64** datums. If neither l nor u is NaN, and $l \leq u$, $l < +\infty$, $u > -\infty$, the result is $[l, u]$. Otherwise the call fails, and the result is `Empty`.

For the constructor `textToInterval(s)`, the input s is a string. The string of form $[l, u]$ where $l < +\infty$ and $u > -\infty$ are optional number literals is called **accuracy-relaxed** if either one of the literals is a rational number literal of form c in 6.6.1 or one of them is a decimal number literal of form a in 6.6.1 and another is a hexadecimal number literal of form b in 6.6.1.

– If s is a valid interval literal with Level 1 value x and s is not **accuracy-relaxed**, the result shall be the hull of x (the constructor succeeds).

– If s is not a valid interval interval and s is not **accuracy-relaxed**, this constructor fails, and the result is `Empty`.

– If s is **accuracy-relaxed** and $l \leq u$, the result may be any interval containing $[l, u]$ (the constructor succeeds).

– If s is **accuracy-relaxed** and $l > u$, the result may be either any interval containing $[u, l]$ (the constructor succeeds), or `Empty` (the constructor fails).

Decorated interval constructors. Let the prefix **b-** or **d-** denote the bare or decorated version of a constructor. If **b-numsToInterval(l, u)** or **b-textToInterval(s)** succeeds with result y , then **d-numsToInterval(l, u)** or **d-textToInterval(s)**, respectively, succeeds with result y and decoration **newDec(y)**, see 5.5.

1 If s is a decorated interval literal $s\mathbf{x}_{sd}$ with Level 1 value x_{dx} , see 6.6.3, and $\mathbf{b}\text{-textToInterval}(s\mathbf{x})$ succeeds
2 with result y , then $\mathbf{d}\text{-textToInterval}(s)$ succeeds with result y_{dy} , where $dy = dx$ except when $dx = \mathbf{com}$
3 and overflow has occurred, that is, x is bounded and y is unbounded. Then dy shall equal \mathbf{dac} .

4 Otherwise the call fails, and the result is \mathbf{NaI} .

5 **Exception behavior.** Exception $\mathbf{UndefinedOperation}$ is signaled by both the bare and the decorated con-
6 structor when the input is such that the bare constructor fails.

7 Exception $\mathbf{PossiblyUndefinedOperation}$ is signaled by both the bare and the decorated $\mathbf{textToInterval}(s)$
8 constructor with accuracy-relaxed s

9 – when $l > u$ and interval constructor succeeds;

10 – when $l \leq u$ and the result is wider than the hull of $[l, u]$;

11 – optionally when $l \leq u$ and the result is the hull of $[l, u]$.

12 NOTE—When signaled by the decorated constructor it will normally be ignored since returning \mathbf{NaI} gives sufficient
13 information.

14 NOTE—If the $\mathbf{textToInterval}(s)$ constructor doesn't signal exception then s is a valid interval literal and the result
15 is the hull of Level 1 value of s .

16 NOTE—The behavior of the $\mathbf{textToInterval}(s)$ constructor is implementation-dependent when s is accuracy-relaxed.
17 The least accurate implementation simply returns \mathbf{Entire} and signals $\mathbf{PossiblyUndefinedOperation}$. The most ac-
18 curate implementation fails with $\mathbf{UndefinedOperation}$ exception when $l > u$ and returns the hull of $[l, u]$ without
19 exception otherwise.

20 6.7.6 Numeric functions of intervals

21 An implementation shall provide a \mathbb{T} -version of each numeric function in Table 4.3 of 4.5.6 giving a result
22 in $\mathbf{b64}$. The mapping of a Level 1 value to a $\mathbf{b64}$ number is defined in terms of the following rounding
23 methods:

24 *Round toward positive:* x maps to the smallest $\mathbf{b64}$ number not less than x ; 0 maps to $+0$.

25 *Round toward negative:* x maps to the largest $\mathbf{b64}$ number not greater than x ; 0 maps to -0 .

26 *Round to nearest:* x maps to the $\mathbf{b64}$ number (possibly $\pm\infty$) closest to x as defined by IEEE 754-2008;
27 0 maps to $+0$.

28 NOTE—These functions help define operations of the standard but are not themselves operations of the stan-
29 dard.

30 A Level 1 value of 0 shall be returned as -0 by \mathbf{inf} , and $+0$ by all other functions in this subclause.

31 $\mathbf{inf}(x)$ returns the Level 1 value rounded toward negative.

32 $\mathbf{sup}(x)$ returns the Level 1 value rounded toward positive.

33 $\mathbf{mid}(x)$: the result is defined by the following cases, where \underline{x}, \bar{x} are the exact (Level 1) lower and upper
34 bounds of x :

| | |
|---|--|
| $x = \mathbf{Empty}$ | \mathbf{NaN} |
| $x = \mathbf{Entire}$ | 0 |
| $\underline{x} = -\infty, \bar{x}$ finite | the finite negative $\mathbf{b64}$ number of largest magnitude |
| \underline{x} finite, $\bar{x} = +\infty$ | the finite positive $\mathbf{b64}$ number of largest magnitude |
| \underline{x}, \bar{x} both finite | the Level 1 value rounded to nearest |

35 The implementation shall document how it handles the last case.

36 $\mathbf{rad}(x)$ returns \mathbf{NaN} if x is empty, and otherwise the smallest $\mathbf{b64}$ number r such that x is contained in the
37 exact interval $[m - r, m + r]$, where m is the value returned by $\mathbf{mid}(x)$.

- 1 **wid**(x) returns NaN if x is empty. Otherwise it returns the Level 1 value rounded toward positive.
- 2 **mag**(x) returns NaN if x is empty. Otherwise it returns the Level 1 value rounded toward positive.
- 3 **mig**(x) returns NaN if x is empty. Otherwise it returns the Level 1 value rounded toward negative, except
- 4 that 0 maps to +0.
- 5 Each bare interval operation in this subclause shall have a decorated version, where each input of bare interval
- 6 type is replaced by one of the corresponding decorated interval type, and the result format is that of the
- 7 bare operation. Following 5.7, if any input is NaI, the result is NaN. Otherwise the result is obtained by
- 8 discarding the decoration and applying the corresponding bare interval operation.

9 **6.7.7 Boolean functions of intervals**

- 10 An implementation shall provide the functions
- 11 – **isEmpty**(x) and **isEntire**(x) in 4.5.7,
- 12 – the functions implementing the comparison relations in Table 4.5 of 4.5.7, and
- 13 – the function **isNaI**(x) for input x of any decorated type, which returns **true** if x is NaI, else **false**.
- 14 Each bare interval operation in this subclause shall have a decorated version. Following 5.7, if any input
- 15 is NaI, the result is **false** (in particular **equal**(NaI, NaI) is **false**). Otherwise the result is obtained by
- 16 discarding the decoration and applying the corresponding bare interval operation.

17 **6.7.8 Operations on/with decorations**

- 18 An implementation shall provide the operations of 5.5. These comprise the comparison operations =, \neq , $>$,
- 19 $<$, \geq , \leq for decorations; and, for the decorated type, the operations **newDec**, **intervalPart**, **decorationPart**
- 20 and **setDec**.
- 21 A call **intervalPart**(NaI), whose value is undefined at Level 1, shall return Empty at Level 2, and shall
- 22 signal the **IntvlPartOfNaI** exception to indicate that a valid interval has been created from the ill-formed
- 23 interval.