

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

]

12.10.2. *Accuracy requirements.* Following the categories of functions in Table 10.1, the accuracy of the *basic operations*, the *case function*, the *integer functions* and the *absmax functions* shall be *tightest*.

For all other functions in Table 10.1, for the reverse mode functions of Table 10.2, and for the recommended functions of Table 10.6 and table Table 10.7, the accuracy shall be *valid*, and, for inf-sup types, should be *accurate*.

12.10.3. *Documentation requirements.* An implementation shall document the tightness of each of its interval operations for each supported bare interval type. This shall be done by dividing the set of possible inputs into disjoint subsets (“ranges”) and stating a tightness achieved 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. Sample tightness information for the sin function might be*

Operation	Type	Tightness	Range
sin	infsup binary64	tightest accurate	for any $x \subseteq [-10^{15}, 10^{15}]$ for all other x .

]

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 needed to resolve ambiguity.

12.11. Interval and number literals.

12.11.1. Overview.

This subclause defines an **interval literal**: a (text) string that denotes a bare or decorated interval. This entails defining a **number literal**: a string within an interval literal that denotes an extended-real number; if it denotes a finite integer it is an **integer literal**.

The bare or decorated interval denoted by an interval literal is its *value* (as an interval literal). Any other string has no value in that sense. For convenience a string is called *valid* or *invalid* (as an interval literal) according as it does or does not have such a value. The usage of value, valid and invalid for number and integer literals is similar.

Interval literals are used as input to `textToInterval` §12.12.8 and in the Input/Output clause §13. In this standard, number (and integer) literals are only used within interval literals. The definitions of literals are not intended to constrain the syntax and semantics that a language might use to denote numbers and intervals in other contexts.

The definition of an interval literal s is placed in Level 2 because its use is not mandatory at Level 1, see §10.6.1. However the value of s is a bare or decorated Level 1 interval x . Within s , conversion of a number literal to its value shall be done as if in infinite precision. Conversion of x to a Level 2 datum 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} .

[*Example. The interval denoted by the literal `[1.2345]` is the Level 1 single-point interval $x = [1.2345, 1.2345]_{\text{com}}$. However the result of $\mathbb{T}\text{-textToInterval}("[1.2345"])$, where \mathbb{T} is the 754 infsup binary64 type, is the interval, approximately `[1.2344999999999999, 1.2345000000000002]_{\text{com}}`, whose endpoints are the nearest binary64 numbers either side of 1.2345.]*

The case of alphabetic characters in interval and number literals, including decorations, is ignored. It is assumed here that they have been converted to lowercase. (E.g., `inf` is equivalent to `Inf` and `[1,2]_dac` is equivalent to `[1,2]_DAC`.)

12.11.2. Number literals.

An integer literal comprises an optional sign and (i.e., followed by) a nonempty sequence of decimal digits.

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. 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-significand form of IEEE 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 optional exponent field comprising `p` and an integer literal. 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 integer literal, negated if there is a leading `-` sign.
 - (c) Either of the strings `inf` or `infinity` optionally preceded by `+`, with value $+\infty$; or preceded by `-`, with value $-\infty$.
 - (d) A rational literal p/q , that is p and q separated by the `/` character, where p, q are decimal integer literals, with q positive. Its value is the exact rational number p/q .

An implementation may support a more general form of integer and/or number literal, e.g., in the syntax of the host language of the implementation. It may restrict the support of literals, by relaxing conversion accuracy of hard cases: rational literals, long strings, etc., converting such literals to Entire, for example. It shall document such restrictions.

By default the syntax shall be that of the default locale (C locale); locale-specific variants may be provided.

12.11.3. Unit in last place. The “uncertain form” of interval literal, below, uses the notion of the *unit in the last place* of a number literal s of some radix b , possibly containing a point but without an exponent field. Ignoring the sign and any radix-specifying code (such as `0x` for hexadecimal), s is a nonempty sequence of radix- b digits optionally containing a point. Its *last place* is the integer $p = -d$ where $d = 0$ if s contains no point, otherwise d is the number of digits after the point. Then $\text{ulp}(s)$ is defined to equal b^p . When context makes clear, “ x ulps of s ” or just “ x ulps”, is used to mean $x \times \text{ulp}(s)$. [Example. For the decimal strings `123` and `123.`, as well as `0` and `0.`, the last place is `0` and one ulp is `1`. For `.123` and `0.123`, as well as `.000` and `0.000`, the last place is `-3` and one ulp is `0.001`.]

12.11.4. Bare intervals.

The following forms of bare interval literal shall be supported. 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 denotes zero or more space characters.

- (a) Special values: The strings `[]` and `[empty]`, whose bare value is Empty; the string `[entire]`, whose bare value is Entire.
- (b) Inf-sup form: A string `[l , u]` where l and u are optional number literals with $l \leq u$, $l < +\infty$ and $u > -\infty$, see §10.2 and the note on difficulties of implementation §12.12.8. 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. A string `[x]` is equivalent to `[x , x]`.
- (c) Uncertain form: a string `m ? r u E` where: m is a decimal number literal of form (a) above, without exponent; r is empty or is a non-negative decimal integer literal *ulp-count* or is the `?` character; u 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 a decimal integer literal *exponent* e . No whitespace is permitted within the string.

With ulp meaning $\text{ulp}(m)$, 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 empty or *ulp-count* is adequate for narrow (and hence bounded) intervals, but is severely restricted otherwise. Uncertain form with radius `?` is for unbounded intervals, e.g. $m???d$ denotes $[-\infty, m]$, $m???u$ denotes $[m, +\infty]$ and $m???$ denotes Entire with m being like a comment. The exponent field if present multiplies the whole interval by 10^e , e.g. $m?ru ee$ denotes $10^e \times [m, m + r \times \text{ulp}]$.

Form	Literal	Exact decorated value
Special	[]	$\text{Empty}_{\text{trv}}$
	[entire]	$[-\infty, +\infty]_{\text{dac}}$
Inf-sup	[1.e-3, 1.1e-3]	$[0.001, 0.0011]_{\text{com}}$
	[-Inf, 2/3]	$[-\infty, 2/3]_{\text{dac}}$
	[0x1.3p-1,]	$[19/32, +\infty]_{\text{dac}}$
	[,]	$[-\infty, +\infty]_{\text{dac}}$
Uncertain	3.56?1	$[3.55, 3.57]_{\text{com}}$
	3.56?1e2	$[355, 357]_{\text{com}}$
	3.560?2	$[3.558, 3.562]_{\text{com}}$
	3.56?	$[3.555, 3.565]_{\text{com}}$
	3.560?2u	$[3.560, 3.562]_{\text{com}}$
	-10?	$[-10.5, -9.5]_{\text{com}}$
	-10?u	$[-10.0, -9.5]_{\text{com}}$
	-10?12	$[-22.0, 2.0]_{\text{com}}$
	-10???u	$[-10.0, +\infty]_{\text{dac}}$
	-10??	$[-\infty, +\infty]_{\text{dac}}$
Nal	[nai]	$\text{Empty}_{\text{i11}}$
Decorated	3.56?1_def	$[3.55, 3.57]_{\text{def}}$

TABLE 12.1. Portable interval literal examples.

12.11.5. Decorated intervals.

Decorated intervals literals may denote either bare or decorated interval value depending on context. The following forms of decorated interval literal shall be supported.

- (a) The string [nai], with the bare value Empty and the decorated value $\text{Empty}_{\text{i11}}$.
- (b) A bare interval literal \mathbf{sx} .

If \mathbf{sx} has the bare value \mathbf{x} , then \mathbf{sx} has the decorated value $\text{newDec}(\mathbf{x})$ §11.5. Otherwise \mathbf{sx} has no decorated value.

- (c) A bare interval literal \mathbf{sx} , an underscore “_”, and a 3-character decoration string \mathbf{sd} ; where \mathbf{sd} is one of **trv**, **def**, **dac** or **com**, denoting the corresponding decoration \mathbf{dx} .

If \mathbf{sx} has the bare value \mathbf{x} , and if $\mathbf{x}_{\mathbf{dx}}$ is a permitted combination according to §11.4, then \mathbf{s} has the bare value \mathbf{x} and the decorated value $\mathbf{x}_{\mathbf{dx}}$. Otherwise \mathbf{s} has no value as a decorated interval literal.

[Examples. Table 12.1 illustrates valid portable interval literals. These strings are not valid portable interval literals: empty, [5?1], [1_000_000], [ganz], [entire!comment], [inf], 5???u, [nai]_i11, []_i11, []_def, [0,inf]_com.]

12.11.6. Grammar for portable literals.

Portable literals permit exchange between different implementations.

The syntax of portable integer and number literals, and of portable bare and decorated interval literals, is defined by `integerLiteral`, `numberLiteral`, `bareIntvlLiteral` and `intervalLiteral` respectively, in the grammar in Table 12.2, which uses the notation of 754§5.12.3. Lowercase is assumed, i.e., a valid string is one that after conversion to lowercase is accepted by this grammar. `\t` denotes the TAB character.

The constructor `textToInterval` §12.12.8, 13.2 of any implementation shall accept any portable interval. Implementation may restrict support of some input strings (too long strings or strings with a rational number literal). Nevertheless, the constructor shall always return a Level 2 interval (possibly Entire in this case) that contains the Level 1 interval.

An implementation may support interval literals of more general syntax (for example, with underscores in significand). In this case there shall be a value of conversion specifier \mathbf{cs} that restricts output strings of `intervalToText` §13.3 to the portable syntax.

12.12. Required operations on bare and decorated intervals.

An implementation shall provide a \mathbb{T} -version, see §12.9, of each operation listed in §12.11 to §12.12.10, for each supported type \mathbb{T} . That is, those of its inputs and outputs that are intervals, are of type \mathbb{T} .

decDigit	[0123456789]
nonzeroDecDigit	[123456789]
hexDigit	[0123456789abcdef]
spaceChar	[\t]
natural	{decDigit} +
sign	[+-]
integerLiteral	{sign} ? {natural}
decSignificand	{decDigit} * "." {decDigit} + {decDigit} + "." {decDigit} +
hexSignificand	{hexDigit} * "." {hexDigit} + {hexDigit} + "." {hexDigit} +
decNumLit	{sign} ? {decSignificand} ("e" {integerLiteral}) ?
hexNumLit	{sign} ? "0x" {hexSignificand} ("p" {integerLiteral}) ?
infNumLit	{sign} ? ("inf" "infinity")
positiveNatural	("0") * {nonzeroDecDigit} {decDigit} *
ratNumLit	{integerLiteral} "/" {positiveNatural}
numberLiteral	{decNumLit} {hexNumLit} {infNumLit} {ratNumLit}
sp	{spaceChar} *
dir	"d" "u"
pointIntvl	"[" {sp} {numberLiteral} {sp} "]"
infSupIntvl	"[" {sp} {numberLiteral} ? {sp} "," {sp} {numberLiteral} ? {sp} "]"
radius	{natural} "?"
uncertIntvl	{sign} ? {decSignificand} "?" {radius} ? {dir} ? ("e" {integerLiteral}) ?
emptyIntvl	"[" {sp} "]" "[" {sp} "empty" {sp} "]"
entireIntvl	"[" {sp} "entire" {sp} "]"
specialIntvl	{emptyIntvl} {entireIntvl}
bareIntvlLiteral	{pointIntvl} {infSupIntvl} {uncertIntvl} {specialIntvl}
Nal	"[" {sp} "nai" {sp} "]"
decorationLit	"trv" "def" "dac" "com"
intervalLiteral	{Nal} {bareIntvlLiteral} {bareIntvlLiteral} "_" {decorationLit}

TABLE 12.2. Grammar for literals: integer literal is `integerLiteral`, number literal is `numberLiteral`, bare interval literal is `bareIntvlLiteral` and decorated interval literal is `intervalLiteral`.

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: for instance two-output division, described in §12.12.4 as a function returning an ordered pair of intervals, might be implemented as a procedure `divToPair(x, y, z_1, z_2)` with input arguments x and y , and output arguments z_1 and z_2 .

An implementation, or a part thereof, that is 754-conforming shall provide mixed-type operations, as specified in §12.6.2, for the following operations, which correspond to those that 754 requires to be provided as *formatOf* operations.

`add`, `sub`, `mul`, `div`, `recip`, `sqrt`, `sqr`, `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.

12.12.1. *Interval constants.* For each supported bare interval type T there shall be a T -version of each constant function `empty()` and `entire()` of §10.6.2, returning a 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.

12.12.2. *Forward-mode elementary functions.* Let T be a supported bare interval type and $\mathbb{D}T$ the derived decorated type. An implementation shall provide a T -version of each forward arithmetic operation in §10.6.3. Its inputs and output are T -intervals and it shall be a Level 2 interval extension of the corresponding point function. Recommended accuracies are given in §12.10.