

INTERNATIONAL
STANDARD

ISO/IEC
10967-2

Fourth committee draft
1999-09-30

Information technology —
Language independent arithmetic —
Part 2: Elementary numerical functions

Technologies de l'information —
Arithmétique indépendante de langage —
Partie 2: Fonctions numériques élémentaires

FINAL COMMITTEE DRAFT
September 30, 1999 18:52

Editor:
Kent Karlsson
IMI, Industri-Matematik International
Kungsgatan 12
SE-411 19 Göteborg
SWEDEN
Telephone: +46-31 10 22 44
Facsimile: +46-31 13 13 25
E-mail: keka@im.se

Contents

1	Scope	1
1.1	Inclusions	1
1.2	Exclusions	2
2	Conformity	2
3	Normative references	3
4	Symbols and definitions	4
4.1	Symbols	4
4.1.1	Sets and intervals	4
4.1.2	Operators and relations	4
4.1.3	Mathematical functions	4
4.1.4	Datatypes and exceptional values	5
4.2	Definitions of terms	6
5	Specifications for the numerical functions	9
5.1	Basic integer operations	9
5.1.1	The integer <i>result</i> and <i>wrap</i> helper functions	9
5.1.2	Integer maximum and minimum	10
5.1.3	Integer diminish	10
5.1.4	Integer power and arithmetic shift	10
5.1.5	Integer square root	11
5.1.6	Divisibility tests	11
5.1.7	Integer division and remainder	11
5.1.8	Greatest common divisor and least common positive multiple	12
5.1.9	Support operations for extended integer range	13
5.2	Basic floating point operations	13
5.2.1	The rounding and floating point <i>result</i> helper functions	14
5.2.2	Floating point maximum and minimum	15
5.2.3	Floating point diminish	17
5.2.4	Round, floor, and ceiling	17
5.2.5	Remainder after division with round to integer	18
5.2.6	Square root and reciprocal square root	18
5.2.7	Support operations for extended floating point precision	19
5.3	Elementary transcendental floating point operations	21
5.3.1	Maximum error requirements	21
5.3.2	Sign requirements	21
5.3.3	Monotonicity requirements	22
5.3.4	The <i>trans_result</i> helper function	22
5.3.5	Hypotenuse	22
5.3.6	Operations for exponentiations and logarithms	23

© ISO/IEC 1999

All rights reserved. No part of this publication may be reproduced or utilised in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the publisher.

ISO/IEC Copyright Office • Case Postale 56 • CH-1211 Genève 20 • Switzerland

Printed in Switzerland

5.3.6.1	Integer power of argument base	23
5.3.6.2	Natural exponentiation	24
5.3.6.3	Natural exponentiation, minus one	24
5.3.6.4	Exponentiation of 2	25
5.3.6.5	Exponentiation of 10	26
5.3.6.6	Exponentiation of argument base	26
5.3.6.7	Exponentiation of one plus the argument base, minus one	27
5.3.6.8	Natural logarithm	28
5.3.6.9	Natural logarithm of one plus the argument	28
5.3.6.10	2-logarithm	28
5.3.6.11	10-logarithm	29
5.3.6.12	Argument base logarithm	29
5.3.6.13	Argument base logarithm of one plus each argument	30
5.3.7	Operations for hyperbolic elementary functions	30
5.3.7.1	Hyperbolic sine	31
5.3.7.2	Hyperbolic cosine	31
5.3.7.3	Hyperbolic tangent	32
5.3.7.4	Hyperbolic cotangent	32
5.3.7.5	Hyperbolic secant	33
5.3.7.6	Hyperbolic cosecant	33
5.3.7.7	Inverse hyperbolic sine	33
5.3.7.8	Inverse hyperbolic cosine	34
5.3.7.9	Inverse hyperbolic tangent	34
5.3.7.10	Inverse hyperbolic cotangent	35
5.3.7.11	Inverse hyperbolic secant	35
5.3.7.12	Inverse hyperbolic cosecant	35
5.3.8	Introduction to operations for trigonometric elementary functions	36
5.3.9	Operations for radian trigonometric elementary functions	36
5.3.9.1	Radian angle normalisation	37
5.3.9.2	Radian sine	38
5.3.9.3	Radian cosine	38
5.3.9.4	Radian tangent	39
5.3.9.5	Radian cotangent	39
5.3.9.6	Radian secant	39
5.3.9.7	Radian cosecant	40
5.3.9.8	Radian cosine with sine	40
5.3.9.9	Radian arc sine	40
5.3.9.10	Radian arc cosine	41
5.3.9.11	Radian arc tangent	41
5.3.9.12	Radian arc cotangent	42
5.3.9.13	Radian arc secant	43
5.3.9.14	Radian arc cosecant	44
5.3.9.15	Radian angle from Cartesian co-ordinates	44
5.3.10	Operations for trigonometrics with given angular unit	45
5.3.10.1	Argument angular-unit angle normalisation	45
5.3.10.2	Argument angular-unit sine	46
5.3.10.3	Argument angular-unit cosine	47
5.3.10.4	Argument angular-unit tangent	47
5.3.10.5	Argument angular-unit cotangent	48
5.3.10.6	Argument angular-unit secant	48
5.3.10.7	Argument angular-unit cosecant	49

5.3.10.8	Argument angular-unit cosine with sine	49
5.3.10.9	Argument angular-unit arc sine	50
5.3.10.10	Argument angular-unit arc cosine	50
5.3.10.11	Argument angular-unit arc tangent	51
5.3.10.12	Argument angular-unit arc cotangent	51
5.3.10.13	Argument angular-unit arc secant	52
5.3.10.14	Argument angular-unit arc cosecant	53
5.3.10.15	Argument angular-unit angle from Cartesian co-ordinates	53
5.3.11	Operations for angular-unit conversions	54
5.3.11.1	Converting radian angle to argument angular-unit angle	54
5.3.11.2	Converting argument angular-unit angle to radian angle	55
5.3.11.3	Converting argument angular-unit angle to (another) argument angular-unit angle	56
5.4	Conversion operations	57
5.4.1	Integer to integer conversions	58
5.4.2	Floating point to integer conversions	58
5.4.3	Integer to floating point conversions	59
5.4.4	Floating point to floating point conversions	59
5.4.5	Floating point to fixed point conversions	60
5.4.6	Fixed point to floating point conversions	61
5.5	Numerals as operations in the programming language	62
5.5.1	Numerals for integer datatypes	62
5.5.2	Numerals for floating point datatypes	62
6	Notification	63
6.1	Continuation values	63
7	Relationship with language standards	63
8	Documentation requirements	64
 Annexes		
A	Partial conformity	67
A.1	Maximum error relaxation	67
A.2	Extra accuracy requirements relaxation	67
A.3	Relationships to other operations relaxation	68
B	Rationale	69
B.1	Scope	69
B.1.1	Inclusions	69
B.1.2	Exclusions	69
B.2	Conformity	70
B.3	Normative references	70
B.4	Symbols and definitions	70
B.4.1	Symbols	70
B.4.1.1	Sets and intervals	70
B.4.1.2	Operators and relations	70
B.4.1.3	Mathematical functions	70
B.4.1.4	Datatypes and exceptional values	71
B.4.2	Definitions of terms	71
B.5	Specifications for the numerical functions	72

B.5.1	Basic integer operations	72
B.5.1.1	The integer <i>result</i> and <i>wrap</i> helper functions	72
B.5.1.2	Integer maximum and minimum	72
B.5.1.3	Integer diminish	73
B.5.1.4	Integer power and arithmetic shift	73
B.5.1.5	Integer square root	73
B.5.1.6	Divisibility tests	73
B.5.1.7	Integer division and remainder	73
B.5.1.8	Greatest common divisor and least common positive multiple	74
B.5.1.9	Support operations for extended integer range	74
B.5.2	Basic floating point operations	74
B.5.2.1	The rounding and floating point <i>result</i> helper functions	75
B.5.2.2	Floating point maximum and minimum	76
B.5.2.3	Floating point diminish	76
B.5.2.4	Round, floor, and ceiling	76
B.5.2.5	Remainder after division and round to integer	76
B.5.2.6	Square root and reciprocal square root	76
B.5.2.7	Support operations for extended floating point precision	77
B.5.3	Elementary transcendental floating point operations	78
B.5.3.1	Maximum error requirements	78
B.5.3.2	Sign requirements	78
B.5.3.3	Monotonicity requirements	79
B.5.3.4	The <i>trans_result</i> helper function	79
B.5.3.5	Hypotenuse	79
B.5.3.6	Operations for exponentiations and logarithms	79
B.5.3.7	Operations for hyperbolic elementary functions	80
B.5.3.8	Introduction to operations for trigonometric elementary functions	81
B.5.3.9	Operations for radian trigonometric elementary functions	82
B.5.3.10	Operations for trigonometrics given angular unit	84
B.5.3.11	Operations for angular-unit conversions	84
B.5.4	Conversion operations	85
B.5.5	Numerals as operations in the programming language	85
B.5.5.1	Numerals for integer datatypes	85
B.5.5.2	Numerals for floating point datatypes	85
B.6	Notification	86
B.6.1	Continuation values	87
B.7	Relationship with language standards	87
B.8	Documentation requirements	87
C	Example bindings for specific languages	89
C.1	General comments	90
C.2	Ada	90
C.3	BASIC	95
C.4	C	98
C.5	C++	103
C.6	Fortran	108
C.7	Haskell	113
C.8	Java	118
C.9	Common Lisp	123
C.10	ISLisp	127
C.11	Modula-2	132
C.12	Pascal and Extended Pascal	136

C.13 PL/I	140
C.14 SML	145
D Bibliography	151

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialised system for world-wide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organisations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1, *Implementation of information technology*. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

International Standard ISO/IEC 10967-2 was prepared by Joint Technical Committee ISO/IEC JTC 1, Sub-Committee SC 22, *Programming languages, their environments and system software interfaces*.

ISO/IEC 10967 consists of the following parts, under the general title *Information technology — Language independent arithmetic*:

- *Part 1: Integer and floating point arithmetic*
- *Part 2: Elementary numerical functions*
- *Part 3: Complex floating point arithmetic and complex elementary numerical functions*

Additional parts will specify other arithmetic datatypes or arithmetic operations.

Introduction

Portability is a key issue for scientific and numerical software in today's heterogeneous computing environment. Such software may be required to run on systems ranging from personal computers to high performance pipelined vector processors and massively parallel systems, and the source code may be ported between several programming languages.

Part 1 of ISO/IEC 10967 specifies the basic properties of integer and floating point types that can be relied upon in writing portable software.

The aims for this Part, Part 2 of ISO/IEC 10967, are extensions of the aims for Part 1: to ensure adequate accuracy for numerical computation, predictability, notification on the production of exceptional results, and compatibility with language standards.

The content of this Part is based on Part 1, and extends Part 1's specifications to specifications for operations approximating real elementary functions, operations often required (usually without a detailed specification) by the standards for programming languages widely used for scientific software. This Part also provides specifications for conversions between the "internal" values of an arithmetic datatype, and a very close approximation in, e.g., the decimal radix. It does not cover the further transformation to decimal string format, which is usually provided by language standards. This Part also includes specifications for a number of other functions deemed useful, even though they may not be stipulated by language standards.

The numerical functions covered by this Part are computer approximations to mathematical functions of one or more real arguments. Accuracy versus performance requirements often vary with the application at hand. This is recognised by recommending that implementors support more than one library of these numerical functions. Various documentation and (program available) parameters requirements are specified to assist programmers in the selection of the library best suited to the application at hand.

Annex B is intended to be read in parallel with the standard.

Notes and annexes B to D are for information only.

Information technology — Language independent arithmetic —

Part 2: Elementary numerical functions

1 Scope

This Part of ISO/IEC 10967 defines the properties of numerical approximations for many of the real elementary numerical functions available in standard libraries for a variety of programming languages in common use for mathematical and numerical applications.

An implementor may choose any combination of hardware and software support to meet the specifications of this Part. It is the computing environment, as seen by the programmer/user, that does or does not conform to the specifications.

The term *implementation* of this Part denotes the total computing environment pertinent to this Part, including hardware, language processors, subroutine libraries, exception handling facilities, other software, and documentation.

1.1 Inclusions

The specifications of Part 1 of are included by reference in this Part.

This Part provides specifications for numerical functions for which all operand values are of integer or floating point datatypes satisfying the requirements of Part 1. Boundaries for the occurrence of exceptions and the maximum error allowed are prescribed for each specified operation. Also the result produced by giving a special value operand, such as an infinity, or a **NaN**, is prescribed for each specified floating point operation.

This Part covers most numerical functions required by the ISO/IEC standards for Ada [11], Basic [17], C [18], C++ [19], Fortran [23], ISLisp [25], Pascal [28], and PL/I [30]. In particular, specifications are provided for

- a) some additional integer operations,
- b) some additional non-transcendental floating point operations, including maximum and minimum operations,
- c) exponentiations, logarithms, hyperbolics, and
- d) trigonometrics, both in radians and for argument-given angular unit with degrees as a special case.

This Part also provides specifications for

- e) conversions between integer and floating point datatypes (possibly with different radices) conforming to the requirements of Part 1, and

- f) the conversion operations used, for example, in text input and output of integer and floating point numbers,
- g) the results produced by an included floating point operation when one or more operand values are IEC 60559 special values, and
- h) program-visible parameters that characterise certain aspects of the operations.

1.2 Exclusions

This Part provides no specifications for:

- a) Numerical functions whose operands are of more than one datatype (with one exception). This standard neither requires nor excludes the presence of such “mixed operand” operations.
- b) An interval datatype, or the operations on such data. This standard neither requires nor excludes such data or operations.
- c) A fixed point datatype, or the operations on such data. This standard neither requires nor excludes such data or operations.
- d) A rational datatype, or the operations on such data. This standard neither requires nor excludes such data or operations.
- e) Complex, matrix, statistical, or symbolic operations. This standard neither requires nor excludes such data or operations.
- f) The properties of arithmetic datatypes that are not related to the numerical process, such as the representation of values on physical media.
- g) The properties of integer and floating point datatypes that properly belong in language standards or other specification. Examples include
 - 1) the syntax of numerals and expressions in the programming language,
 - 2) the syntax used for parsed (input) or generated (output) character string forms for numerals by any specific programming language or library,
 - 3) the precedence of operators,
 - 4) the consequences of applying an operation to values of improper datatype, or to uninitialised data,
 - 5) the rules for assignment, parameter passing, and returning value,
 - 6) the presence or absence of automatic datatype coercions.

Furthermore, this Part does not provide specifications for:

- h) how numerical functions should be implemented,
- i) which algorithms are to be used for the various operations.

2 Conformity

It is expected that the provisions of this Part of ISO/IEC 10967 will be incorporated by reference and further defined in other International Standards; specifically in language standards and in language binding standards.

A binding standard specifies the correspondence between one or more operations and parameters specified in this Part and the concrete language syntax of some programming language. More generally, a binding standard specifies the correspondence between certain operations and the elements of some arbitrary computing entity. A language standard that explicitly provides such binding information can serve as a binding standard.

Conformity to this Part is always with respect to a specified set of operations. Conformity to this Part implies conformity to Part 1 for the integer and floating point datatypes used.

When a binding standard for a language exists, an implementation shall be said to conform to this Part if and only if it conforms to the binding standard. In the case of conflict between a binding standard and this Part, the specifications of the binding standard takes precedence.

When a binding standard covers only a subset of the operations defined in this Part, an implementation remains free to conform to this Part with respect to other operations independently of that binding standard.

When no binding standard for a language and some operations specified in this Part exists, an implementation conforms to this Part if and only if it provides one or more operations that together satisfy all the requirements of clauses 5 through 8 that are relevant to those operations. The implementation shall then document the binding.

An implementation is free to provide operations that do not conform to this Part, or that are beyond the scope of this Part. The implementation shall not claim or imply conformity to this Part with respect to such operations.

An implementation is permitted to have modes of operation that do not conform to this Part. A conforming implementation shall specify how to select the modes of operation that ensure conformity.

NOTES

- 1 Language bindings are essential. Clause 8 requires an implementation to supply a binding if no binding standard exists. See annex C for suggested language bindings.
- 2 A complete binding for this Part will include (explicitly or by reference) a binding for Part 1 as well, which in turn may include (explicitly or by reference) a binding for IEC 60559 as well.
- 3 It is not possible to conform to this Part without specifying to which set of operations conformity is claimed.

3 Normative references

The following standards contain provisions which, through reference in this text, constitute provisions of this Part. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this Part are encouraged to investigate the possibility of applying the most recent edition of the standards indicated below. Members of IEC and ISO maintain registers of currently valid International Standards.

IEC 60559:1989, *Binary floating-point arithmetic for microprocessor systems*.

ISO/IEC 10967-1:1994, *Information technology — Language independent arithmetic — Part 1: Integer and floating point arithmetic*.

4 Symbols and definitions

4.1 Symbols

4.1.1 Sets and intervals

In this Part, \mathcal{Z} denotes the set of mathematical integers, \mathcal{R} denotes the set of classical real numbers, and \mathcal{C} denotes the set of complex numbers over \mathcal{R} . Note that $\mathcal{Z} \subset \mathcal{R} \subset \mathcal{C}$.

$[x, z]$ designates the interval $\{y \in \mathcal{R} \mid x \leq y \leq z\}$,
 $]x, z]$ designates the interval $\{y \in \mathcal{R} \mid x < y \leq z\}$,
 $[x, z[$ designates the interval $\{y \in \mathcal{R} \mid x \leq y < z\}$, and
 $]x, z[$ designates the interval $\{y \in \mathcal{R} \mid x < y < z\}$.

NOTE – The notation using a round bracket for an open end of an interval is not used, for the risk of confusion with the notation for pairs.

4.1.2 Operators and relations

All prefix and infix operators have their conventional (exact) mathematical meaning. The conventional notation for set definition and manipulation is also used. In particular this Part uses

\Rightarrow and \Leftrightarrow for logical implication and equivalence
 $+$, $-$, $/$, $|x|$, $\lfloor x \rfloor$, $\lceil x \rceil$, and $\text{round}(x)$ on reals
 \cdot for multiplication on reals
 $<$, \leq , $=$, \neq , \geq , and $>$ between reals
 \max on non-empty upwardly closed sets of reals
 \min on non-empty downwardly closed sets of reals
 \cup , \cap , \times , \in , \notin , \subset , $\not\subset$, \neq , and $=$ with sets
 \times for the Cartesian product of sets
 \rightarrow for a mapping between sets
 $|$ for the divides relation between integers

For $x \in \mathcal{R}$, the notation $\lfloor x \rfloor$ designates the largest integer not greater than x :

$$\lfloor x \rfloor \in \mathcal{Z} \quad \text{and} \quad x - 1 < \lfloor x \rfloor \leq x$$

the notation $\lceil x \rceil$ designates the smallest integer not less than x :

$$\lceil x \rceil \in \mathcal{Z} \quad \text{and} \quad x \leq \lceil x \rceil < x + 1$$

and the notation $\text{round}(x)$ designates the integer closest to x :

$$\text{round}(x) \in \mathcal{Z} \quad \text{and} \quad x - 0.5 \leq \text{round}(x) \leq x + 0.5$$

where in case x is exactly half-way between two integers, the even integer is the result.

The *divides* relation ($|$) on integers tests whether an integer i divides an integer j exactly:

$$i|j \Leftrightarrow (i \neq 0 \text{ and } i \cdot n = j \text{ for some } n \in \mathcal{Z})$$

NOTE – $i|j$ is true exactly when j/i is defined and $j/i \in \mathcal{Z}$.

4.1.3 Mathematical functions

This Part specifies properties for a number of operations numerically approximating some of the elementary functions. The following ideal mathematical functions are defined in Chapter 4 of the *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables* [48] (e is the Napierian base):

e^x , x^y , \sqrt{x} , \ln , \log_b ,
 \sinh , \cosh , \tanh , \coth , sech , csch , $\operatorname{arcsinh}$, $\operatorname{arccosh}$, $\operatorname{arctanh}$, $\operatorname{arccoth}$, $\operatorname{arcsech}$, $\operatorname{arccsch}$,
 \sin , \cos , \tan , \cot , \sec , \csc , arcsin , arccos , arctan , arccot , arcsec , arccsc .

Many of the inverses above are multi-valued. The selection of which value to return, the principal value, so as to make the inverses into functions, is done in the conventional way. The only one over which there is some difference of conventions is the arccot function. Conventions there vary for negative arguments; either a positive return value (giving a function that is continuous over zero), or a negative value (giving a sign symmetric function). In this Part arccot refers to the continuous inverse function, and arctg refers to the sign symmetric inverse function.

$\operatorname{arccosh}(x) \geq 0$, $\operatorname{arcsech}(x) \geq 0$,
 $\operatorname{arcsin}(x) \in [-\pi/2, \pi/2]$, $\operatorname{arccos}(x) \in [0, \pi]$, $\operatorname{arctan}(x) \in]-\pi/2, \pi/2[$,
 $\operatorname{arccot}(x) \in]0, \pi[$, $\operatorname{arctg}(x) \in]-\pi/2, \pi/2[$, $\operatorname{arcsec}(x) \in [0, \pi]$, $\operatorname{arccsc}(x) \in [-\pi/2, \pi/2]$.

NOTES

- 1 *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables* [48] uses the notation arccot for what is called arctg in this Part.
- 2 $e = 2.71828\dots$ e is not in F .

4.1.4 Datatypes and exceptional values

For pairs, define:

$\operatorname{fst}((x, y)) = x$
 $\operatorname{snd}((x, y)) = y$

Square brackets are used to write finite sequences of values. $[]$ is the sequence containing no values. $[s]$, is the sequence of one value, s . $[s_1, s_2]$, is the sequence of two values, s_1 and then s_2 , etc. The colon operator is used to prepend a value to a sequence: $x : [x_1, \dots, x_n] = [x, x_1, \dots, x_n]$.

$[S]$, where S is a set, denotes the set of finite sequences, where each value in each sequence is in S .

NOTE 1 – It is always clear from context, in the text of this Part, if $[X]$ is a sequence of one element, or the set of sequences with values from X . It is also clear from context if $[x_1, x_2]$ is a sequence of two values or an interval.

The datatype **Boolean** consists of the two values **true** and **false**.

Integer datatypes and floating point datatypes are defined in Part 1.

The following symbols are defined in Part 1, and used in this Part.

Exceptional values:

underflow.

Integer parameters:

$\operatorname{bounded}_I$, maxint_I , and minint_I .

Integer helper function:

wrap_I .

Integer operations:

neg_I , add_I , sub_I , and mul_I .

Floating point parameters:

r_F , p_F , emin_F , emax_F , denorm_F , and $\operatorname{iec}_{559}_F$.

Derived floating point constants:

fmax_F , fmin_F , fminN_F , fminD_F , and $\operatorname{epsilon}_F$.

Floating point rounding constants:

$\operatorname{rnd_error}_F$.

Floating point value sets related to F :

F^* , F_D , and F_N .

Floating point helper functions:

e_F , $result_F$, and rnd_F .

Floating point operations:

neg_F , add_F , sub_F , mul_F , div_F , and ulp_F .

Floating point datatypes that conform to Part 1 shall, for use with this Part, have a value for the parameter p_F such that $p_F \geq 2 \cdot \max\{1, \log_{r_F}(2 \cdot \pi)\}$, and have a value for the parameter $emin_F$ such that $emin_F \leq -p_F - 1$.

NOTES

- 2 This implies that $fminN_F < 0.5 \cdot \epsilon_F / r_F$ in this Part, rather than just $fminN_F \leq \epsilon_F$.
- 3 These extra requirements, which do not limit the use of any existing floating point datatype, are made 1) so that angles in radians are not too degenerate within the first two cycles, plus and minus, when represented in F , and 2) in order to justly allow avoiding the underflow notification in specifications for the $expm1_F$ and $ln1p_F$ operations.
- 4 F should also be such that $p_F \geq 2 + \log_{r_F}(1000)$, to allow for a not too coarse angle resolution anywhere in the interval $[-big_angle_{r_F}, big_angle_{r_F}]$. See clause 5.3.9.

Three new exceptional values, **overflow**, **invalid**, and **pole**, are introduced in this Part replacing three other exceptional values used in Part 1. One new exceptional value, **absolute_precision_underflow**, is introduced in this Part with no correspondence in Part 1. **invalid** and **pole** are in this Part used instead of the **undefined** of Part 1. **overflow** is used instead of the **integer_overflow** and **floating_overflow** of Part 1. Bindings may still distinguish between **integer_overflow** and **floating_overflow**. The exceptional value **absolute_precision_underflow** is used when the given floating point angle value argument is so big that even a highly accurate result from a trigonometric operation is questionable, due to the fact that the density of floating point values has decreased significantly at these big angle values. For the exceptional values, a continuation value may be given in parenthesis after the exceptional value.

The following symbols represent floating point values defined in IEC 60559 and used in this Part:

-0, **+\infty**, **-\infty**, **qNaN**, and **sNaN**.

These floating point values are not part of the set F , but if iec_559_F has the value **true**, these values are included in the floating point datatype corresponding to F .

NOTE 5 – This Part uses the above five special values for compatibility with IEC 60559. In particular, the symbol **-0** (in bold) is not the application of (mathematical) unary $-$ to the value 0, and is a value logically distinct from 0.

The specifications cover the results to be returned by an operation if given one or more of the IEC 60559 special values **-0**, **+\infty**, **-\infty**, or **NaNs** as input values. These specifications apply only to systems which provide and support these special values. If an implementation is not capable of representing a **-0** result or continuation value, the actual result or continuation value shall be 0. If an implementation is not capable of representing a prescribed result or continuation value of the IEC 60559 special values **+\infty**, **-\infty**, or **qNaN**, the actual result or continuation value is binding or implementation defined.

4.2 Definitions of terms

For the purposes of this Part, the following definitions apply:

accuracy: The closeness between the true mathematical result and a computed result.

arithmetic datatype: A datatype whose non-special values are members of \mathcal{Z} , \mathcal{R} , or \mathcal{C} .

NOTE 1 – This standard specifies requirements for integer and floating point datatypes. Complex numbers are not covered here, but will be included in a subsequent Part of ISO/IEC 10967 [5].

continuation value: A computational value used as the result of an arithmetic operation when an exception occurs. Continuation values are intended to be used in subsequent arithmetic processing. A continuation value can be a value in F or an IEC 60559 special value. (Contrast with *exceptional value*. See 6.1.2 of Part 1.)

denormalisation loss: A larger than normal rounding error caused by the fact that subnormal values have less than full precision. (See 5.2.5 of Part 1 for a full definition.)

denormalised, denormal: The non-zero values of a floating point type F that provide less than the full precision allowed by that type. (See F_D in 5.2 of Part 1 for a full definition.)

error: (1) The difference between a computed value and the correct value. (Used in phrases like “rounding error” or “error bound”.)

(2) A synonym for *exception* in phrases like “error message” or “error output”. Error and exception are not synonyms in any other context.

exception: The inability of an operation to return a suitable finite numeric result from finite arguments. This might arise because no such finite result exists mathematically, or because the mathematical result cannot be represented with sufficient accuracy.

exceptional value: A non-numeric value produced by an arithmetic operation to indicate the occurrence of an exception. Exceptional values are not used in subsequent arithmetic processing. (See clause 5 of Part 1.)

NOTES

- 2 Exceptional values are used as part of the defining formalism only. With respect to this Part, they do not represent values of any of the datatypes described. There is no requirement that they be represented or stored in the computing system.
- 3 Exceptional values are not to be confused with the NaNs and infinities defined in IEC 60559. Contrast this definition with that of *continuation value* above.

helper function: A function used solely to aid in the expression of a requirement. Helper functions are not visible to the programmer, and are not required to be part of an implementation.

implementation (of this Part): The total arithmetic environment presented to a programmer, including hardware, language processors, exception handling facilities, subroutine libraries, other software, and all pertinent documentation.

literal: A syntactic entity denoting a constant value without having proper sub-entities that are expressions.

monotonic approximation: An operation $op_F : \dots \times F \times \dots \rightarrow F$, where the other arguments are kept constant, is a monotonic approximation of a predetermined mathematical function $h : \mathcal{R} \rightarrow \mathcal{R}$ if, for every $a \in F$ and $b \in F$,

- a) h is monotonic non-decreasing on $[a, b]$ implies $op_F(\dots, a, \dots) \leq op_F(\dots, b, \dots)$,
- b) h is monotonic non-increasing on $[a, b]$ implies $op_F(\dots, a, \dots) \geq op_F(\dots, b, \dots)$.

monotonic non-decreasing: A function $h : \mathcal{R} \rightarrow \mathcal{R}$ is monotonic non-decreasing on a real interval $[a, b]$ if for every x and y such that $a \leq x \leq y \leq b$, $h(x)$ and $h(y)$ are well-defined and $h(x) \leq h(y)$.

monotonic non-increasing: A function $h : \mathcal{R} \rightarrow \mathcal{R}$ is monotonic non-increasing on a real interval $[a, b]$ if for every x and y such that $a \leq x \leq y \leq b$, $h(x)$ and $h(y)$ are well-defined and $h(x) \geq h(y)$.

normalised: The non-zero values of a floating point type F that provide the full precision allowed by that type. (See F_N in 5.2 of Part 1 for a full definition.)

notification: The process by which a program (or that program's end user) is informed that an arithmetic exception has occurred. For example, dividing 2 by 0 results in a notification. (See clause 6 of Part 1 for details.)

numeral: A numeric literal. It may denote a value in \mathcal{Z} or \mathcal{R} , -0 , an infinity, or a NaN.

numerical function: A computer routine or other mechanism for the approximate evaluation of a mathematical function.

operation: A function directly available to the user/programmer, as opposed to helper functions or theoretical mathematical functions.

pole: A mathematical function f has a pole at x_0 if x_0 is finite, f is defined, finite, monotone, and continuous in at least one side of the neighbourhood of x_0 , and $\lim_{x \rightarrow x_0} f(x)$ is infinite.

precision: The number of digits in the fraction of a floating point number. (See 5.2 of Part 1.)

rounding: The act of computing a representable final result for an operation that is close to the exact (but unrepresentable) result for that operation. Note that a suitable representable result may not exist (see 5.2.6 of Part 1). (See also A.5.2.6 of Part 1 for some examples.)

rounding function: Any function $rnd : \mathcal{R} \rightarrow X$ (where X is a given discrete and unlimited subset of \mathcal{R}) that maps each element of X to itself, and is monotonic non-decreasing. Formally, if x and y are in \mathcal{R} ,

$$\begin{aligned} x \in X &\Rightarrow rnd(x) = x \\ x < y &\Rightarrow rnd(x) \leq rnd(y) \end{aligned}$$

Note that if $u \in \mathcal{R}$ is between two adjacent values in X , $rnd(u)$ selects one of those adjacent values.

round to nearest: The property of a rounding function rnd that when $u \in \mathcal{R}$ is between two adjacent values in X , $rnd(u)$ selects the one nearest u . If the adjacent values are equidistant from u , either may be chosen deterministically.

round toward minus infinity: The property of a rounding function rnd that when $u \in \mathcal{R}$ is between two adjacent values in X , $rnd(u)$ selects the one less than u .

round toward plus infinity: The property of a rounding function rnd that when $u \in \mathcal{R}$ is between two adjacent values in X , $rnd(u)$ selects the one greater than u .

shall: A verbal form used to indicate requirements strictly to be followed in order to conform to the standard and from which no deviation is permitted. (Quoted from the directives [1].)

should: A verbal form used to indicate that among several possibilities one is recommended as particularly suitable, without mentioning or excluding others; or that (in the negative form) a certain possibility is deprecated but not prohibited. (Quoted from the directives [1].)

signature (of a function or operation): A summary of information about an operation or function. A signature includes the function or operation name; a subset of allowed argument values to the operation; and a superset of results from the function or operation (including exceptional values if any), if the argument is in the subset of argument values given in the signature.

The signature

$$add_I : I \times I \rightarrow I \cup \{\mathbf{overflow}\}$$

states that the operation named add_I shall accept any pair of I values as input, and (when given such input) shall return either a single I value as its output or the exceptional value **overflow**.

A signature for an operation or function does not forbid the operation from accepting a wider range of arguments, nor does it guarantee that every value in the result range will actually be returned for some input. An operation given an argument outside the stipulated argument domain may produce a result outside the stipulated result range.

subnormal: A denormal value, the value 0, or the value $-\mathbf{0}$.

ulp: The value of one “unit in the last place” of a floating point number. This value depends on the exponent, the radix, and the precision used in representing the number. Thus, the ulp of a normalised value x (in F), with exponent t , precision p , and radix r , is r^{t-p} , and the ulp of a subnormal value is $fminD_F$. (See 5.2 of Part 1.)

5 Specifications for the numerical functions

This clause specifies a number of helper functions and operations for integer and floating point datatypes. Each operation is given a signature and is further specified by a number of cases. These cases may refer to other operations (specified in this Part or in Part 1), to mathematical functions, and to helper functions (specified in this Part or in Part 1). They also use special abstract values ($-\infty$, $+\infty$, $-\mathbf{0}$, **qNaN**, **sNaN**). For each datatype, two of these abstract values may represent several actual values each: **qNaN** and **sNaN**. Finally, the specifications may refer to exceptional values.

The signatures in the specifications in this clause specify only all non-special values as input values, and indicate as output values the superset of all non-special, special, and exceptional values that may result from these (non-special) input values. Therefore, exceptional and special values that can never result from non-special input values are not included in the signatures given. Also, signatures that, for example, include IEC 60559 special values as arguments are not given in the specifications below. This does not exclude such signatures from being valid for these operations.

5.1 Basic integer operations

Clause 5.1 of Part 1 specifies integer datatypes and a number of operations on values of an integer datatype. In this clause some additional operations on values of an integer datatype are specified.

I is the set of non-special values, $I \subseteq \mathcal{Z}$, for an integer datatype conforming to Part 1. Integer datatypes conforming to Part 1 often do not contain any **NaN** or infinity values, even though they may do so. Therefore this clause has no specifications for such values as arguments or results.

5.1.1 The integer *result* and *wrap* helper functions

The $result_I$ helper function:

$$\begin{aligned} result_I : \mathcal{Z} &\rightarrow I \cup \{\mathbf{overflow}\} \\ result_I(x) &= x && \text{if } x \in I \\ &= \mathbf{overflow} && \text{if } x \in \mathcal{Z} \text{ and } x \notin I \end{aligned}$$

The $wrap_I$ helper function:

$$\begin{aligned} wrap_I : \mathcal{Z} &\rightarrow I \\ wrap_I(x) &= x && \text{if } x \in I \\ &= x - (n \cdot (maxint_I - minint_I + 1)) && \text{if } x \in \mathcal{Z} \text{ and } x \notin I \end{aligned}$$

where $n \in \mathcal{Z}$ is chosen such that the result is in I .

NOTES

- 1 $n = \lfloor (x - minint_I) / (maxint_I - minint_I + 1) \rfloor$ if $x \in \mathcal{Z}$ and $bounded_I = \mathbf{true}$; or equivalently $n = \lceil (x - maxint_I) / (maxint_I - minint_I + 1) \rceil$ if $x \in \mathcal{Z}$ and $bounded_I = \mathbf{true}$.
- 2 For some wrapping basic arithmetic operations this n is computed by the ‘*lov*’ operations in clause 5.1.9.
- 3 The $wrap_I$ helper function is also used in Part 1.

5.1.2 Integer maximum and minimum

$$\begin{aligned} max_I : I \times I &\rightarrow I \\ max_I(x, y) &= \max\{x, y\} && \text{if } x, y \in I \end{aligned}$$

$$\begin{aligned} min_I : I \times I &\rightarrow I \\ min_I(x, y) &= \min\{x, y\} && \text{if } x, y \in I \end{aligned}$$

$$\begin{aligned} max_seq_I : [I] &\rightarrow I \cup \{\mathbf{pole}\} \\ max_seq_I([x_1, \dots, x_n]) &= \mathbf{pole}(-\infty) && \text{if } n = 0 \\ &= \max\{x_1, \dots, x_n\} && \text{if } n \geq 1 \text{ and } \{x_1, \dots, x_n\} \subseteq I \end{aligned}$$

$$\begin{aligned} min_seq_I : [I] &\rightarrow I \cup \{\mathbf{pole}\} \\ min_seq_I([x_1, \dots, x_n]) &= \mathbf{pole}(+\infty) && \text{if } n = 0 \\ &= \min\{x_1, \dots, x_n\} && \text{if } n \geq 1 \text{ and } \{x_1, \dots, x_n\} \subseteq I \end{aligned}$$

5.1.3 Integer diminish

$$\begin{aligned} dim_I : I \times I &\rightarrow I \cup \{\mathbf{overflow}\} \\ dim_I(x, y) &= result_I(\max\{0, x - y\}) && \text{if } x, y \in I \end{aligned}$$

NOTE – dim_I cannot be implemented as $max_I(0, sub_I(x, y))$ for bounded integer types, since this latter expression has other overflow properties.

5.1.4 Integer power and arithmetic shift

$$\begin{aligned} power_I : I \times I &\rightarrow I \cup \{\mathbf{overflow}, \mathbf{pole}, \mathbf{invalid}\} \\ power_I(x, y) &= result_I(x^y) && \text{if } x, y \in I \text{ and } (y > 0 \text{ or } |x| = 1) \\ &= 1 && \text{if } x \in I \text{ and } x \neq 0 \text{ and } y = 0 \\ &= \mathbf{invalid}(1) && \text{if } x = 0 \text{ and } y = 0 \\ &= \mathbf{pole}(+\infty) && \text{if } x = 0 \text{ and } y \in I \text{ and } y < 0 \\ &= \mathbf{invalid}(0) && \text{if } x, y \in I \text{ and } x \notin \{-1, 0, 1\} \text{ and } y < 0 \end{aligned}$$

$$\begin{aligned} \text{shift2}_I &: I \times I \rightarrow I \cup \{\mathbf{overflow}\} \\ \text{shift2}_I(x, y) &= \text{result}_I(\lfloor x \cdot 2^y \rfloor) && \text{if } x, y \in I \end{aligned}$$

$$\begin{aligned} \text{shift10}_I &: I \times I \rightarrow I \cup \{\mathbf{overflow}\} \\ \text{shift10}_I(x, y) &= \text{result}_I(\lfloor x \cdot 10^y \rfloor) && \text{if } x, y \in I \end{aligned}$$

5.1.5 Integer square root

$$\begin{aligned} \text{sqr}_I &: I \rightarrow I \cup \{\mathbf{invalid}\} \\ \text{sqr}_I(x) &= \lfloor \sqrt{x} \rfloor && \text{if } x \in I \text{ and } x \geq 0 \\ &= \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \in I \text{ and } x < 0 \end{aligned}$$

5.1.6 Divisibility tests

$$\begin{aligned} \text{divides}_I &: I \times I \rightarrow \mathbf{Boolean} \\ \text{divides}_I(x, y) &= \mathbf{true} && \text{if } x, y \in I \text{ and } x|y \\ &= \mathbf{false} && \text{if } x, y \in I \text{ and not } x|y \end{aligned}$$

NOTES

- 1 $\text{divides}_I(0, 0) = \mathbf{false}$, since 0 does not divide anything, not even 0.
- 2 divides_I cannot be implemented as, e.g., $\text{eq}_I(0, \text{moda}_I(y, x))$, since the remainder functions give notifications for a zero second argument.

$$\begin{aligned} \text{even}_I &: I \rightarrow \mathbf{Boolean} \\ \text{even}_I(x) &= \mathbf{true} && \text{if } x \in I \text{ and } 2|x \\ &= \mathbf{false} && \text{if } x \in I \text{ and not } 2|x \end{aligned}$$

$$\begin{aligned} \text{odd}_I &: I \rightarrow \mathbf{Boolean} \\ \text{odd}_I(x) &= \mathbf{true} && \text{if } x \in I \text{ and not } 2|x \\ &= \mathbf{false} && \text{if } x \in I \text{ and } 2|x \end{aligned}$$

5.1.7 Integer division and remainder

$$\begin{aligned} \text{divf}_I &: I \times I \rightarrow I \cup \{\mathbf{overflow}, \mathbf{pole}, \mathbf{invalid}\} \\ \text{divf}_I(x, y) &= \text{result}_I(\lfloor x/y \rfloor) && \text{if } x, y \in I \text{ and } y \neq 0 \\ &= \mathbf{pole}(+\infty) && \text{if } x \in I \text{ and } x > 0 \text{ and } y = 0 \\ &= \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x = 0 \text{ and } y = 0 \\ &= \mathbf{pole}(-\infty) && \text{if } x \in I \text{ and } x < 0 \text{ and } y = 0 \end{aligned}$$

$$\begin{aligned} \text{moda}_I &: I \times I \rightarrow I \cup \{\mathbf{invalid}\} \\ \text{moda}_I(x, y) &= x - (\lfloor x/y \rfloor \cdot y) && \text{if } x, y \in I \text{ and } y \neq 0 \\ &= \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \in I \text{ and } y = 0 \end{aligned}$$

$$\begin{aligned}
& \text{group}_I : I \times I \rightarrow I \cup \{\text{overflow}, \text{pole}, \text{invalid}\} \\
& \text{group}_I(x, y) = \text{result}_I(\lceil x/y \rceil) && \text{if } x, y \in I \text{ and } y \neq 0 \\
& = \text{pole}(+\infty) && \text{if } x \in I \text{ and } x > 0 \text{ and } y = 0 \\
& = \text{invalid}(\text{qNaN}) && \text{if } x = 0 \text{ and } y = 0 \\
& = \text{pole}(-\infty) && \text{if } x \in I \text{ and } x < 0 \text{ and } y = 0
\end{aligned}$$

$$\begin{aligned}
& \text{pad}_I : I \times I \rightarrow I \cup \{\text{invalid}\} \\
& \text{pad}_I(x, y) = (\lceil x/y \rceil \cdot y) - x && \text{if } x, y \in I \text{ and } y \neq 0 \\
& = \text{invalid}(\text{qNaN}) && \text{if } x \in I \text{ and } y = 0
\end{aligned}$$

$$\begin{aligned}
& \text{quot}_I : I \times I \rightarrow I \cup \{\text{overflow}, \text{pole}, \text{invalid}\} \\
& \text{quot}_I(x, y) = \text{result}_I(\text{round}(x/y)) && \text{if } x, y \in I \text{ and } y \neq 0 \\
& = \text{pole}(+\infty) && \text{if } x \in I \text{ and } x > 0 \text{ and } y = 0 \\
& = \text{invalid}(\text{qNaN}) && \text{if } x = 0 \text{ and } y = 0 \\
& = \text{pole}(-\infty) && \text{if } x \in I \text{ and } x < 0 \text{ and } y = 0
\end{aligned}$$

$$\begin{aligned}
& \text{remr}_I : I \times I \rightarrow I \cup \{\text{overflow}, \text{invalid}\} \\
& \text{remr}_I(x, y) = \text{result}_I(x - (\text{round}(x/y) \cdot y)) && \text{if } x, y \in I \text{ and } y \neq 0 \\
& = \text{invalid}(\text{qNaN}) && \text{if } x \in I \text{ and } y = 0
\end{aligned}$$

5.1.8 Greatest common divisor and least common positive multiple

$$\begin{aligned}
& \text{gcd}_I : I \times I \rightarrow I \cup \{\text{overflow}, \text{pole}\} \\
& \text{gcd}_I(x, y) = \text{result}_I(\max\{v \in \mathcal{Z} \mid v|x \text{ and } v|y\}) && \text{if } x, y \in I \text{ and } (x \neq 0 \text{ or } y \neq 0) \\
& = \text{pole}(+\infty) && \text{if } x = 0 \text{ and } y = 0
\end{aligned}$$

$$\begin{aligned}
& \text{lcm}_I : I \times I \rightarrow I \cup \{\text{overflow}\} \\
& \text{lcm}_I(x, y) = \text{result}_I(\min\{v \in \mathcal{Z} \mid x|v \text{ and } y|v \text{ and } v > 0\}) && \text{if } x, y \in I \text{ and } x \neq 0 \text{ and } y \neq 0 \\
& = 0 && \text{if } x, y \in I \text{ and } (x = 0 \text{ or } y = 0)
\end{aligned}$$

$$\begin{aligned}
& \text{gcd_seq}_I : [I] \rightarrow I \cup \{\text{overflow}, \text{pole}\} \\
& \text{gcd_seq}_I([x_1, \dots, x_n]) = \text{result}_I(\max\{v \in \mathcal{Z} \mid v|x_i \text{ for all } i \in \{1, \dots, n\}\}) && \text{if } \{x_1, \dots, x_n\} \subseteq I \text{ and } \{x_1, \dots, x_n\} \not\subseteq \{0\} \\
& = \text{pole}(+\infty) && \text{if } \{x_1, \dots, x_n\} \subseteq \{0\}
\end{aligned}$$

$$\begin{aligned}
& \text{lcm_seq}_I : [I] \rightarrow I \cup \{\text{overflow}\} \\
& \text{lcm_seq}_I([x_1, \dots, x_n]) = \text{result}_I(\min\{v \in \mathcal{Z} \mid x_i|v \text{ for all } i \in \{1, \dots, n\} \text{ and } v > 0\}) && \text{if } \{x_1, \dots, x_n\} \subseteq I \text{ and } 0 \notin \{x_1, \dots, x_n\} \\
& = 0 && \text{if } \{x_1, \dots, x_n\} \subseteq I \text{ and } 0 \in \{x_1, \dots, x_n\}
\end{aligned}$$

NOTE – This specification implies that $\text{lcm_seq}_I(\emptyset) = 1$.

5.1.9 Support operations for extended integer range

These operations can be used to implement extended range integer datatypes, including unbounded integer datatypes.

$$\mathit{add_wrap}_I : I \times I \rightarrow I$$

$$\mathit{add_wrap}_I(x, y) = \mathit{wrap}_I(x + y) \quad \text{if } x, y \in I$$

$$\mathit{add_ov}_I : I \times I \rightarrow \{-1, 0, 1\}$$

$$\begin{aligned} \mathit{add_ov}_I(x, y) &= ((x + y) - \mathit{add_wrap}_I(x, y)) / (\mathit{maxint}_I - \mathit{minint}_I + 1) && \text{if } x, y \in I \text{ and } \mathit{bounded}_I = \mathbf{true} \\ &= 0 && \text{if } x, y \in I \text{ and } \mathit{bounded}_I = \mathbf{false} \end{aligned}$$

$$\mathit{sub_wrap}_I : I \times I \rightarrow I$$

$$\mathit{sub_wrap}_I(x, y) = \mathit{wrap}_I(x - y) \quad \text{if } x, y \in I$$

$$\mathit{sub_ov}_I : I \times I \rightarrow \{-1, 0, 1\}$$

$$\begin{aligned} \mathit{sub_ov}_I(x, y) &= ((x - y) - \mathit{sub_wrap}_I(x, y)) / (\mathit{maxint}_I - \mathit{minint}_I + 1) && \text{if } x, y \in I \text{ and } \mathit{bounded}_I = \mathbf{true} \\ &= 0 && \text{if } x, y \in I \text{ and } \mathit{bounded}_I = \mathbf{false} \end{aligned}$$

$$\mathit{mul_wrap}_I : I \times I \rightarrow I$$

$$\mathit{mul_wrap}_I(x, y) = \mathit{wrap}_I(x \cdot y) \quad \text{if } x, y \in I$$

$$\mathit{mul_ov}_I : I \times I \rightarrow I$$

$$\begin{aligned} \mathit{mul_ov}_I(x, y) &= ((x \cdot y) - \mathit{mul_wrap}_I(x, y)) / (\mathit{maxint}_I - \mathit{minint}_I + 1) && \text{if } x, y \in I \text{ and } \mathit{bounded}_I = \mathbf{true} \\ &= 0 && \text{if } x, y \in I \text{ and } \mathit{bounded}_I = \mathbf{false} \end{aligned}$$

NOTE – The $\mathit{add_ov}_I$ and $\mathit{sub_ov}_I$ will only return -1 (for negative overflow), 0 (no overflow), and 1 (for positive overflow).

5.2 Basic floating point operations

Clause 5.2 of Part 1 specifies floating point datatypes and a number of operations on values of a floating point datatype. In this clause some additional operations on values of a floating point datatype are specified.

NOTE – Further operations on values of a floating point datatype, for elementary floating point numerical functions, are specified in clause 5.3.

F is the non-special value set, $F \subset \mathcal{R}$, for a floating point datatype conforming to Part 1. Floating point datatypes conforming to Part 1 often do contain $-\mathbf{0}$, infinity, and **NaN** values. Therefore, in this clause there are specifications for such values as arguments.

5.2.1 The rounding and floating point *result* helper functions

Floating point rounding helper functions: The floating point helper function

$$\text{down}_F : \mathcal{R} \rightarrow F^*$$

is the rounding function that rounds towards negative infinity. The floating point helper function

$$\text{up}_F : \mathcal{R} \rightarrow F^*$$

is the rounding function that rounds towards positive infinity. The floating point helper function

$$\text{nearest}_F : \mathcal{R} \rightarrow F^*$$

is the rounding function that rounds to nearest. nearest_F is partially implementation defined: the handling of ties is implementation defined, but must be sign symmetric. If $\text{iec_559}_F = \mathbf{true}$, the semantics of nearest_F is completely defined by IEC 60559: in this case ties are rounded to even last digit.

result_F is a helper function that is partially implementation defined.

$$\text{result}_F : \mathcal{R} \times (\mathcal{R} \rightarrow F^*) \rightarrow F \cup \{\mathbf{underflow}, \mathbf{overflow}\}$$

$$\begin{aligned} \text{result}_F(x, \text{nearest}_F) &= \mathbf{overflow}(+\infty) && \text{if } x \in \mathcal{R} \text{ and } \text{nearest}_F(x) > \text{fmax}_F \\ \text{result}_F(x, \text{nearest}_F) &= \mathbf{overflow}(-\infty) && \text{if } x \in \mathcal{R} \text{ and } \text{nearest}_F(x) < -\text{fmax}_F \\ \text{result}_F(x, \text{up}_F) &= \mathbf{overflow}(+\infty) && \text{if } x \in \mathcal{R} \text{ and } \text{up}_F(x) > \text{fmax}_F \\ \text{result}_F(x, \text{up}_F) &= \mathbf{overflow}(-\text{fmax}_F) && \text{if } x \in \mathcal{R} \text{ and } \text{up}_F(x) < -\text{fmax}_F \\ \text{result}_F(x, \text{down}_F) &= \mathbf{overflow}(\text{fmax}_F) && \text{if } x \in \mathcal{R} \text{ and } \text{down}_F(x) > \text{fmax}_F \\ \text{result}_F(x, \text{down}_F) &= \mathbf{overflow}(-\infty) && \text{if } x \in \mathcal{R} \text{ and } \text{down}_F(x) < -\text{fmax}_F \end{aligned}$$

otherwise:

$$\begin{aligned} \text{result}_F(x, \text{rnd}) &= x && \text{if } x = 0 \\ &= \text{rnd}(x) && \text{if } x \in \mathcal{R} \text{ and } \text{fmin}N_F \leq |x| \text{ and } |\text{rnd}(x)| \leq \text{fmax}_F \\ &= \text{rnd}(x) \text{ or } \mathbf{underflow}(c) && \text{if } x \in \mathcal{R} \text{ and } |x| < \text{fmin}N_F \text{ and } |\text{rnd}(x)| = \text{fmin}N_F \\ &&& \text{and } \text{rnd} \text{ has no denormalisation loss at } x \\ &= \text{rnd}(x) \text{ or } \mathbf{underflow}(c) && \text{if } x \in \mathcal{R} \text{ and } \text{denorm}_F = \mathbf{true} \text{ and} \\ &&& |\text{rnd}(x)| < \text{fmin}N_F \text{ and } x \neq 0 \\ &&& \text{and } \text{rnd} \text{ has no denormalisation loss at } x \\ &= \mathbf{underflow}(c) && \text{otherwise} \end{aligned}$$

where

$$\begin{aligned} c &= \text{rnd}(x) && \text{when } \text{denorm}_F = \mathbf{true} \text{ and } (\text{rnd}(x) \neq 0 \text{ or } x > 0), \\ c &= -\mathbf{0} && \text{when } \text{denorm}_F = \mathbf{true} \text{ and } \text{rnd}(x) = 0 \text{ and } x < 0, \\ c &= 0 && \text{when } \text{denorm}_F = \mathbf{false} \text{ and } x > 0, \\ c &= -\mathbf{0} && \text{when } \text{denorm}_F = \mathbf{false} \text{ and } x < 0 \end{aligned}$$

An implementation is allowed to choose between $\text{rnd}(x)$ and $\mathbf{underflow}(\text{rnd}(x))$ in the region between 0 and $\text{fmin}N_F$. However, a subnormal value without underflow notification can be chosen only if denorm_F is true and no denormalisation loss occurs at x .

NOTES

- 1 This differs from the specification of result_F as given in Part 1 in the following respects:
 - 1) the continuation values on overflow and underflow are given directly here, and 2) all instances of denormalisation loss must be accompanied with an underflow notification.
- 2 $\text{denorm}_F = \mathbf{false}$ implies $\text{iec_559}_F = \mathbf{false}$, and $\text{iec_559}_F = \mathbf{true}$ implies $\text{denorm}_F = \mathbf{true}$.

- 3 If $iec_{559_F} = \mathbf{true}$, then subnormal results that have no denormalisation loss, e.g. are exact, do not result in an underflow notification, if the notification is by recording of indicators.

Define the $result_NaN_F$, $result_NaN2_F$, and $result_NaN3_F$ helper functions:

$$result_NaN_F : F \rightarrow \{\mathbf{invalid}\}$$

$$\begin{aligned} result_NaN_F(x) &= \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN} \\ &= \mathbf{invalid}(\mathbf{qNaN}) && \text{otherwise} \end{aligned}$$

$$result_NaN2_F : F \times F \rightarrow \{\mathbf{invalid}\}$$

$$\begin{aligned} result_NaN2_F(x, y) &= \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN and } y \text{ is not a signalling NaN} \\ &= \mathbf{qNaN} && \text{if } y \text{ is a quiet NaN and } x \text{ is not a signalling NaN} \\ &= \mathbf{invalid}(\mathbf{qNaN}) && \text{otherwise} \end{aligned}$$

$$result_NaN3_F : F \times F \times F \rightarrow \{\mathbf{invalid}\}$$

$$\begin{aligned} result_NaN3_F(x, y, z) &= \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN and} \\ & && \text{not } y \text{ nor } z \text{ is a signalling NaN} \\ &= \mathbf{qNaN} && \text{if } y \text{ is a quiet NaN and} \\ & && \text{not } x \text{ nor } z \text{ is a signalling NaN} \\ &= \mathbf{qNaN} && \text{if } z \text{ is a quiet NaN and} \\ & && \text{not } x \text{ nor } y \text{ is a signalling NaN} \\ &= \mathbf{invalid}(\mathbf{qNaN}) && \text{otherwise} \end{aligned}$$

These helper functions are used to specify both NaN argument handling and to handle non-NaN-argument cases where $\mathbf{invalid}(\mathbf{qNaN})$ is the appropriate result.

5.2.2 Floating point maximum and minimum

The appropriate return value of the maximum and minimum operations given a quiet NaN (\mathbf{qNaN}) as one of the input values depends on the circumstances for each point of use. Sometimes \mathbf{qNaN} is the appropriate result, sometimes the non- \mathbf{NaN} argument is the appropriate result. Therefore, two variants each of the floating point maximum and minimum operations are specified here, and the programmer can decide which one is appropriate to use at each particular place of usage, assuming both variants are included in the binding.

$$max_F : F \times F \rightarrow F$$

$$\begin{aligned} max_F(x, y) &= \max\{x, y\} && \text{if } x, y \in F \\ &= +\infty && \text{if } x = +\infty \text{ and } y \in F \cup \{-\infty, -0\} \\ &= y && \text{if } x = -0 \text{ and } y \in F \text{ and } y \geq 0 \\ &= -0 && \text{if } x = -0 \text{ and } ((y \in F \text{ and } y < 0) \text{ or } y = -0) \\ &= y && \text{if } x = -\infty \text{ and } y \in F \cup \{+\infty, -0\} \\ &= +\infty && \text{if } y = +\infty \text{ and } x \in F \cup \{+\infty, -0\} \\ &= x && \text{if } y = -0 \text{ and } x \in F \text{ and } x \geq 0 \\ &= -0 && \text{if } y = -0 \text{ and } x \in F \text{ and } x < 0 \end{aligned}$$

$$\begin{aligned}
 &= x && \text{if } y = -\infty \text{ and } x \in F \cup \{-\infty, -\mathbf{0}\} \\
 &= \text{result_NaN}\mathcal{2}_F(x, y) && \text{otherwise}
 \end{aligned}$$

$$\min_F : F \times F \rightarrow F$$

$$\begin{aligned}
 \min_F(x, y) &= \min\{x, y\} && \text{if } x, y \in F \\
 &= y && \text{if } x = +\infty \text{ and } y \in F \cup \{-\infty, -\mathbf{0}\} \\
 &= -\mathbf{0} && \text{if } x = -\mathbf{0} \text{ and } y \in F \text{ and } y \geq 0 \\
 &= y && \text{if } x = -\mathbf{0} \text{ and } ((y \in F \text{ and } y < 0) \text{ or } y = -\mathbf{0}) \\
 &= -\infty && \text{if } x = -\infty \text{ and } y \in F \cup \{+\infty, -\mathbf{0}\} \\
 &= x && \text{if } y = +\infty \text{ and } x \in F \cup \{+\infty, -\mathbf{0}\} \\
 &= -\mathbf{0} && \text{if } y = -\mathbf{0} \text{ and } x \in F \text{ and } x \geq 0 \\
 &= x && \text{if } y = -\mathbf{0} \text{ and } x \in F \text{ and } x < 0 \\
 &= -\infty && \text{if } y = -\infty \text{ and } x \in F \cup \{-\infty, -\mathbf{0}\} \\
 &= \text{result_NaN}\mathcal{2}_F(x, y) && \text{otherwise}
 \end{aligned}$$

$$\text{mmax}_F : F \times F \rightarrow F$$

$$\begin{aligned}
 \text{mmax}_F(x, y) &= \text{max}_F(x, y) && \text{if } x, y \in F \cup \{+\infty, -\mathbf{0}, -\infty\} \\
 &= x && \text{if } x \in F \cup \{+\infty, -\mathbf{0}, -\infty\} \text{ and } y \text{ is a quiet NaN} \\
 &= y && \text{if } y \in F \cup \{+\infty, -\mathbf{0}, -\infty\} \text{ and } x \text{ is a quiet NaN} \\
 &= \text{result_NaN}\mathcal{2}_F(x, y) && \text{otherwise}
 \end{aligned}$$

$$\text{mmin}_F : F \times F \rightarrow F$$

$$\begin{aligned}
 \text{mmin}_F(x, y) &= \min_F(x, y) && \text{if } x, y \in F \cup \{+\infty, -\mathbf{0}, -\infty\} \\
 &= x && \text{if } x \in F \cup \{+\infty, -\mathbf{0}, -\infty\} \text{ and } y \text{ is a quiet NaN} \\
 &= y && \text{if } y \in F \cup \{+\infty, -\mathbf{0}, -\infty\} \text{ and } x \text{ is a quiet NaN} \\
 &= \text{result_NaN}\mathcal{2}_F(x, y) && \text{otherwise}
 \end{aligned}$$

NOTE – If one of the arguments to mmax_F or mmin_F is a quiet NaN, that argument is ignored.

$$\text{max_seq}_F : [F] \rightarrow F \cup \{-\infty, \text{pole}\}$$

$$\begin{aligned}
 \text{max_seq}_F([x_1, \dots, x_n]) &= -\infty && \text{if } n = 0 \text{ and } -\infty \text{ is available} \\
 &= \text{pole}(-\text{fmax}_F) && \text{if } n = 0 \text{ and } -\infty \text{ is not available} \\
 &= \text{max}_F(\text{max_seq}_F([x_1, \dots, x_{n-1}]), x_n) && \text{if } n \geq 2 \\
 &= x_1 && \text{if } n = 1 \text{ and } x_1 \text{ is not a NaN} \\
 &= \text{result_NaN}_F(x_1) && \text{otherwise}
 \end{aligned}$$

$$\text{min_seq}_F : [F] \rightarrow F \cup \{+\infty, \text{pole}\}$$

$$\begin{aligned}
 \text{min_seq}_F([x_1, \dots, x_n]) &= +\infty && \text{if } n = 0 \text{ and } +\infty \text{ is available} \\
 &= \text{pole}(\text{fmax}_F) && \text{if } n = 0 \text{ and } +\infty \text{ is not available} \\
 &= \min_F(\text{min_seq}_F([x_1, \dots, x_{n-1}]), x_n) && \text{if } n \geq 2 \\
 &= x_1 && \text{if } n = 1 \text{ and } x_1 \text{ is not a NaN} \\
 &= \text{result_NaN}_F(x_1) && \text{otherwise}
 \end{aligned}$$

$$\begin{aligned}
mmax_seq_F : [F] &\rightarrow F \cup \{-\infty, \mathbf{pole}\} \\
mmax_seq_F([x_1, \dots, x_n]) & \\
&= -\infty && \text{if } n = 0 \text{ and } -\infty \text{ is available} \\
&= \mathbf{pole}(-fmax_F) && \text{if } n = 0 \text{ and } -\infty \text{ is not available} \\
&= mmax_F(mmax_seq_F([x_1, \dots, x_{n-1}], x_n)) && \text{if } n \geq 2 \\
&= x_1 && \text{if } n = 1 \text{ and } x_1 \text{ is not a NaN} \\
&= result_NaN_F(x_1) && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
mmin_seq_F : [F] &\rightarrow F \cup \{+\infty, \mathbf{pole}\} \\
mmin_seq_F([x_1, \dots, x_n]) & \\
&= +\infty && \text{if } n = 0 \text{ and } +\infty \text{ is available} \\
&= \mathbf{pole}(fmax_F) && \text{if } n = 0 \text{ and } +\infty \text{ is not available} \\
&= mmin_F(mmin_seq_F([x_1, \dots, x_{n-1}], x_n)) && \text{if } n \geq 2 \\
&= x_1 && \text{if } n = 1 \text{ and } x_1 \text{ is not a NaN} \\
&= result_NaN_F(x_1) && \text{otherwise}
\end{aligned}$$

5.2.3 Floating point diminish

$$\begin{aligned}
dim_F : F \times F &\rightarrow F \cup \{\mathbf{overflow}, \mathbf{underflow}\} \\
dim_F(x, y) &= result_F(\max\{0, x - y\}, rnd_F) \\
&= dim_F(0, y) && \text{if } x, y \in F \\
&= dim_F(x, 0) && \text{if } x = -\mathbf{0} \text{ and } y \in F \cup \{-\infty, -\mathbf{0}, +\infty\} \\
&= +\infty && \text{if } y = -\mathbf{0} \text{ and } x \in F \cup \{-\infty, +\infty\} \\
&= 0 && \text{if } x = +\infty \text{ and } y \in F \cup \{-\infty\} \\
&= 0 && \text{if } x = -\infty \text{ and } y \in F \cup \{+\infty\} \\
&= +\infty && \text{if } y = +\infty \text{ and } x \in F \\
&= +\infty && \text{if } y = -\infty \text{ and } x \in F \\
&= result_NaN2_F(x, y) && \text{otherwise}
\end{aligned}$$

NOTE – dim_F cannot be implemented by $max_F(0, sub_F(x, y))$, since this latter expression has other overflow properties.

5.2.4 Round, floor, and ceiling

$$\begin{aligned}
rounding_F : F &\rightarrow F \cup \{-\mathbf{0}\} \\
rounding_F(x) &= \text{round}(x) && \text{if } x \in F \text{ and } (x \geq 0 \text{ or } \text{round}(x) \neq 0) \\
&= -\mathbf{0} && \text{if } x \in F \text{ and } x < 0 \text{ and } \text{round}(x) = 0 \\
&= x && \text{if } x \in \{-\infty, -\mathbf{0}, +\infty, \} \\
&= result_NaN_F(x) && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
floor_F : F &\rightarrow F \\
floor_F(x) &= \lfloor x \rfloor && \text{if } x \in F \\
&= x && \text{if } x \in \{-\infty, -\mathbf{0}, +\infty, \} \\
&= result_NaN_F(x) && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
\text{ceiling}_F : F &\rightarrow F \cup \{-0\} \\
\text{ceiling}_F(x) &= \lceil x \rceil && \text{if } x \in F \text{ and } (x \geq 0 \text{ or } \lceil x \rceil \neq 0) \\
&= -0 && \text{if } x \in F \text{ and } x < 0 \text{ and } \lceil x \rceil = 0 \\
&= x && \text{if } x \in \{-\infty, -0, +\infty, \} \\
&= \text{result_NaN}_F(x) && \text{otherwise}
\end{aligned}$$

NOTE 1 – Truncate to integer is specified in Part 1, by the name *intpart_F*.

$$\begin{aligned}
\text{rounding_rest}_F : F &\rightarrow F \\
\text{rounding_rest}_F(x) &= x - \text{round}(x) && \text{if } x \in F \\
&= 0 && \text{if } x = -0 \\
&= \text{result_NaN}_F(x) && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
\text{floor_rest}_F : F &\rightarrow F \\
\text{floor_rest}_F(x) &= \text{result}_F(x - \lfloor x \rfloor, \text{rnd}_F) && \text{if } x \in F \\
&= 0 && \text{if } x = -0 \\
&= \text{result_NaN}_F(x) && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
\text{ceiling_rest}_F : F &\rightarrow F \\
\text{ceiling_rest}_F(x) &= \text{result}_F(x - \lceil x \rceil, \text{rnd}_F) && \text{if } x \in F \\
&= 0 && \text{if } x = -0 \\
&= \text{result_NaN}_F(x) && \text{otherwise}
\end{aligned}$$

NOTE 2 – The rest after truncation is specified in Part 1, by the name *fractpart_F*.

5.2.5 Remainder after division with round to integer

$$\begin{aligned}
\text{remr}_F : F \times F &\rightarrow F \cup \{-0, \text{underflow}, \text{invalid}\} \\
\text{remr}_F(x, y) &= \text{result}_F(x - (\text{round}(x/y) \cdot y), \text{nearest}_F) && \text{if } x, y \in F \text{ and } y \neq 0 \text{ and} \\
&&& (x \geq 0 \text{ or } x - (\text{round}(x/y) \cdot y) \neq 0) \\
&= -0 && \text{if } x, y \in F \text{ and } y \neq 0 \text{ and} \\
&&& x < 0 \text{ and } x - (\text{round}(x/y) \cdot y) = 0 \\
&= -0 && \text{if } x = -0 \text{ and } y \in F \cup \{-\infty, +\infty\} \text{ and } y \neq 0 \\
&= x && \text{if } x \in F \text{ and } y \in \{-\infty, +\infty\} \\
&= \text{result_NaN}_F(x, y) && \text{otherwise}
\end{aligned}$$

5.2.6 Square root and reciprocal square root

$$\begin{aligned}
\text{sqr}_F : F &\rightarrow F \cup \{\text{invalid}\} \\
\text{sqr}_F(x) &= \text{nearest}_F(\sqrt{x}) && \text{if } x \in F \text{ and } x \geq 0 \\
&= x && \text{if } x \in \{-0, +\infty\} \\
&= \text{result_NaN}_F(x) && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
\text{rec_sqrt}_F &: F \rightarrow F \cup \{\mathbf{invalid}, \mathbf{pole}\} \\
\text{rec_sqrt}_F(x) &= \text{rnd}_F(1/\sqrt{x}) && \text{if } x \in F \text{ and } x > 0 \\
&= \mathbf{pole}(+\infty) && \text{if } x \in \{-0, 0\} \\
&= 0 && \text{if } x = +\infty \\
&= \text{result_NaN}_F(x) && \text{otherwise}
\end{aligned}$$

5.2.7 Support operations for extended floating point precision

These operations are useful when keeping guard digits or implementing extra precision floating point datatypes. The resulting datatypes, e.g. so-called doubled precision, do not necessarily conform to Part 1.

$$\begin{aligned}
\text{add_lo}_F &: F \times F \rightarrow F \cup \{\mathbf{underflow}\} \\
\text{add_lo}_F(x, y) &= \text{result}_F((x + y) - \text{rnd}_F(x + y), \text{rnd}_F) && \text{if } x, y \in F \\
&= -0 && \text{if } x = -0 \text{ and } y \in F \cup \{-\infty, -0, +\infty\} \\
&= -0 && \text{if } x \in F \cup \{-\infty, +\infty\} \text{ and } y = -0 \\
&= y && \text{if } x = +\infty \text{ and } y \in F \cup \{+\infty\} \\
&= y && \text{if } x = -\infty \text{ and } y \in F \cup \{-\infty\} \\
&= x && \text{if } x \in F \text{ and } y \in \{-\infty, +\infty\} \\
&= \text{result_NaN}_F(x, y) && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
\text{sub_lo}_F &: F \times F \rightarrow F \cup \{\mathbf{underflow}\} \\
\text{sub_lo}_F(x, y) &= \text{add_lo}_F(x, \text{neg}_F(y))
\end{aligned}$$

NOTE 1 – If $\text{rnd_style}_F = \text{nearest}$, then, in the absence of notifications, add_lo_F and sub_lo_F returns exact results.

$$\begin{aligned}
\text{mul_lo}_F &: F \times F \rightarrow F \cup \{\mathbf{overflow}, \mathbf{underflow}\} \\
\text{mul_lo}_F(x, y) &= \text{result}_F((x \cdot y) - \text{rnd}_F(x \cdot y), \text{rnd}_F) && \text{if } x, y \in F \\
&= \text{mul_lo}_F(0, y) && \text{if } x = -0 \text{ and } y \in F \cup \{-\infty, -0, +\infty\} \\
&= \text{mul_lo}_F(x, 0) && \text{if } x \in F \cup \{-\infty, +\infty\} \text{ and } y = -0 \\
&= \text{mul}_F(x, y) && \text{if } x \in \{-\infty, +\infty\} \text{ and } y \in F \cup \{-\infty, +\infty\} \\
&= \text{mul}_F(x, y) && \text{if } x \in F \text{ and } y \in \{-\infty, +\infty\} \\
&= \text{result_NaN}_F(x, y) && \text{otherwise}
\end{aligned}$$

NOTE 2 – In the absence of notifications, mul_lo_F returns an exact result.

$$\begin{aligned}
\text{div_rest}_F &: F \times F \rightarrow F \cup \{\mathbf{underflow}, \mathbf{invalid}\} \\
\text{div_rest}_F(x, y) &= \text{result}_F(x - (y \cdot \text{rnd}_F(x/y)), \text{rnd}_F) && \text{if } x, y \in F \\
&= \text{div_rest}_F(0, y) && \text{if } x = -0 \text{ and } y \in F \cup \{-\infty, -0, +\infty\} \\
&= x && \text{if } x \in F \text{ and } y \in \{-\infty, +\infty\} \\
&= x && \text{if } x \in \{-\infty, +\infty\} \text{ and } y \in F \\
&= \text{result_NaN}_F(x, y) && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
\mathit{sqr_rest}_F : F &\rightarrow F \cup \{\mathbf{underflow}, \mathbf{invalid}\} \\
\mathit{sqr_rest}_F(x) &= \mathit{result}_F(x - (\mathit{sqr}_F(x) \cdot \mathit{sqr}_F(x)), \mathit{rnd}_F) \\
&\quad \text{if } x \in F \text{ and } x \geq 0 \\
&= -\mathbf{0} && \text{if } x = -\mathbf{0} \\
&= +\infty && \text{if } x = +\infty \\
&= \mathit{result_NaN}_F(x) && \text{otherwise}
\end{aligned}$$

NOTE 3 – $\mathit{sqr_rest}_F(x)$ is exact when there is no **underflow**.

For the following operation F' is a floating point type conforming to Part 1.

NOTE 4 – It is expected that $p_{F'} > p_F$, i.e. F' has higher precision than F , but that is not required.

$$\begin{aligned}
\mathit{mul}_{F \rightarrow F'} : F \times F &\rightarrow F' \cup \{-\mathbf{0}, \mathbf{overflow}, \mathbf{underflow}\} \\
\mathit{mul}_{F \rightarrow F'}(x, y) &= \mathit{result}_{F'}(x \cdot y, \mathit{rnd}_{F'}) \quad \text{if } x, y \in F \text{ and } x \neq 0 \text{ and } y \neq 0 \\
&= \mathit{convert}_{F \rightarrow F'}(\mathit{mul}_F(x, y)) \\
&\quad \text{if } x \in \{-\infty, -\mathbf{0}, 0, +\infty\} \text{ and} \\
&\quad \quad y \in F \cup \{-\infty, -\mathbf{0}, +\infty\} \\
&= \mathit{convert}_{F \rightarrow F'}(\mathit{mul}_F(x, y)) \\
&\quad \text{if } y \in \{-\infty, -\mathbf{0}, 0, +\infty\} \text{ and } x \in F \text{ and } x \neq 0 \\
&= \mathit{result_NaN}_{F'}(x, y) \quad \text{otherwise}
\end{aligned}$$

5.3 Elementary transcendental floating point operations

5.3.1 Maximum error requirements

The specifications for each of the transcendental and floating point conversion operations use an approximation helper function. The approximation helper functions are ideally identical to the true mathematical functions. However, that would imply a maximum error for the corresponding operation of 0.5 ulp (i.e., the minimum value for operations that are not always exact). This Part does not require that the maximum error is only 0.5 ulp for the operations specified in clauses 5.3, 5.4, and 5.5, but allows the maximum error to be a bit bigger. To express this, the approximation helper functions need not be identical to the mathematical elementary transcendental functions, but are allowed to be approximate.

The approximation helper functions for the individual operations in these subclauses have maximum error parameters that describe the maximum *relative* error of the helper function composed with $nearest_F$, for non-subnormal results. The maximum error parameters also describe the maximum *absolute* error for subnormal results and underflow continuation values if $denorm_F = \mathbf{true}$. The relevant maximum error parameters shall be available to programs.

When the maximum error for an approximation helper function h_F , approximating f , is $max_error_op_F$, then for all arguments $x, \dots \in F^* \times \dots$ the following equation shall hold:

$$|f(x, \dots) - nearest_F(h_F(x, \dots))| \leq max_error_op_F \cdot r_F^{e_F(f(x, \dots)) - p_F}$$

NOTES

- 1 Partially conforming implementations may have greater values for maximum error parameters than stipulated below. See annex A.
- 2 For most positive (and not too small) return values t , the true result is thus claimed to be in the interval $[t - (max_error_op_F \cdot ulp_F(t)), t + (max_error_op_F \cdot ulp_F(t))]$. But if the return value is exactly r_F^n for some not too small $n \in \mathcal{Z}$, then the true result is claimed to be in the interval $[t - (max_error_op_F \cdot ulp_F(t)/r_F), t + (max_error_op_F \cdot ulp_F(t))]$. Similarly for negative return values.

The results of the approximating helper functions in this clause must be exact for certain arguments as detailed below, and may be exact for all arguments. If the approximating helper function is exact for all arguments, then the corresponding maximum error parameter should be 0.5, the minimum value.

5.3.2 Sign requirements

The approximation helper functions shall be zero exactly at the points where the approximated mathematical function is exactly zero. At points where the approximation helper functions are not zero, they shall have the same sign as the approximated mathematical function at that point.

For the radian trigonometric helper functions, these zero and sign requirements are imposed only for arguments, x , such that $|x| \leq big_angle_r_F$ (see clause 5.3.9).

NOTE – For the operations, the continuation value after an **underflow** may be zero (or negative zero) as given by $trans_result_F$, even though the approximation helper function is not zero at that point. Such zero results are required to be accompanied by an **underflow** notification. When appropriate, zero may also be returned for IEC 60559 infinities arguments. See the individual specifications.

5.3.3 Monotonicity requirements

Each approximation helper function in this clause shall be a monotonic approximation to the mathematical function it is approximating, except:

- a) For the radian trigonometric approximation helper functions, the monotonic approximation requirement is imposed only for arguments, x , such that $|x| \leq \mathit{big_angle_r}_F$ (see clause 5.3.9).
- b) The argument angular unit trigonometric and argument angular unit inverse trigonometric approximating helper functions are excepted from the monotonic approximation requirement for the angular unit argument.

5.3.4 The *trans_result* helper function

The *trans_result_F* helper function is similar to the *result_F* helper function (see 5.2.1), but is simplified compared to *result_F* concerning **underflow**: *trans_result_F* always underflows for non-zero arguments that have an absolute value less than $\mathit{fmin}N_F - (\mathit{fmin}D_F/r_F)$, whereas *result_F* does not necessarily underflow in that case. This difference from *result_F* is made since the argument to *trans_result_F* might not be exact. To return **underflow** or not, for a tiny result, based upon an inexact result would be misleading. For the operations specified using *trans_result_F* where the specification implies that there will be no denormalisation loss for certain tiny results, **underflow** is instead explicitly avoided.

$$\begin{aligned} \mathit{trans_result}_F : \mathcal{R} \times (\mathcal{R} \rightarrow F^*) &\rightarrow F \cup \{\mathbf{underflow}, \mathbf{overflow}\} \\ \mathit{trans_result}_F(x, \mathit{rnd}) &= \mathbf{underflow}(c) && \text{if } x \in \mathcal{R} \text{ and } \mathit{denorm}_F = \mathbf{true} \text{ and} \\ & && |\mathit{rnd}(x)| < \mathit{fmin}N_F \text{ and } x \neq 0 \\ &= \mathit{result}_F(x, \mathit{rnd}) && \text{otherwise} \end{aligned}$$

where

$$\begin{aligned} c &= \mathit{rnd}(x) && \text{when } \mathit{rnd}(x) \neq 0 \text{ or } x > 0, \\ c &= -\mathbf{0} && \text{when } \mathit{rnd}(x) = 0 \text{ and } x < 0 \end{aligned}$$

5.3.5 Hypotenuse

There shall be a maximum error parameter for the *hypot_F* operation:

$$\mathit{max_error_hypot}_F \in F$$

The *max_error_hypot_F* parameter shall have a value in the interval $[0.5, 1]$.

The *hypot_F^{*}* approximation helper function:

$$\mathit{hypot}_F^* : F \times F \rightarrow \mathcal{R}$$

hypot_F^{}*(x, y) returns a close approximation to $\sqrt{x^2 + y^2}$ in \mathcal{R} , with maximum error *max_error_hypot_F*.

Further requirements on the *hypot_F^{*}* approximation helper function are:

$$\begin{aligned} \mathit{hypot}_F^*(x, y) &= \mathit{hypot}_F^*(y, x) \\ \mathit{hypot}_F^*(-x, y) &= \mathit{hypot}_F^*(x, y) \\ \mathit{hypot}_F^*(x, y) &\geq \max\{|x|, |y|\} \\ \mathit{hypot}_F^*(x, y) &\leq |x| + |y| \\ \mathit{hypot}_F^*(x, y) &\geq 1 && \text{if } \sqrt{x^2 + y^2} \geq 1 \\ \mathit{hypot}_F^*(x, y) &\leq 1 && \text{if } \sqrt{x^2 + y^2} \leq 1 \end{aligned}$$

The $hypot_F$ operation:

$$\begin{aligned}
 hypot_F : F \times F &\rightarrow F \cup \{\mathbf{underflow}, \mathbf{overflow}\} \\
 hypot_F(x, y) &= trans_result_F(hypot_F^*(x, y), nearest_F) \\
 &\quad \text{if } x, y \in F \\
 &= hypot_F(0, y) && \text{if } x = -\mathbf{0} \text{ and } y \in F \cup \{-\infty, -\mathbf{0}, +\infty\} \\
 &= hypot_F(x, 0) && \text{if } y = -\mathbf{0} \text{ and } x \in F \cup \{-\infty, +\infty\} \\
 &= +\infty && \text{if } x \in \{-\infty, +\infty\} \text{ and } y \in F \cup \{-\infty, +\infty\} \\
 &= +\infty && \text{if } y \in \{-\infty, +\infty\} \text{ and } x \in F \\
 &= result_NaN2_F(x, y) && \text{otherwise}
 \end{aligned}$$

5.3.6 Operations for exponentiations and logarithms

There shall be two maximum error parameters for approximate exponentiations and logarithms:

$$\begin{aligned}
 max_error_exp_F &\in F \\
 max_error_power_F &\in F
 \end{aligned}$$

The $max_error_exp_F$ parameter shall have a value in the interval $[0.5, 1.5 \cdot rnd_error_F]$. The $max_error_power_F$ parameter shall have a value in the interval $[max_error_exp_F, 2 \cdot rnd_error_F]$.

NOTE – The “exp” operations are thus required to be at least as accurate as the “power” operations.

5.3.6.1 Integer power of argument base

The $power_{FI}^*$ approximation helper function:

$$power_{FI}^* : F \times I \rightarrow \mathcal{R}$$

$power_{FI}^*(x, y)$ returns a close approximation to x^y in \mathcal{R} , with maximum error $max_error_power_F$.

Further requirements on the $power_{FI}^*$ approximation helper function are:

$$\begin{aligned}
 power_{FI}^*(1, y) &= 1 && \text{if } y \in I \\
 power_{FI}^*(x, 0) &= 1 && \text{if } x \in F \text{ and } x \neq 0 \\
 power_{FI}^*(x, 1) &= x && \text{if } x \in F \\
 power_{FI}^*(x, y) &< fminD_F/2 && \text{if } x \in F \text{ and } x > 0 \text{ and } y \in I \text{ and } x^y < fminD_F/3 \\
 power_{FI}^*(x, y) &= power_{FI}^*(-x, y) && \text{if } x \in F \text{ and } x < 0 \text{ and } y \in I \text{ and } 2|y \\
 power_{FI}^*(x, y) &= -power_{FI}^*(-x, y) && \text{if } x \in F \text{ and } x < 0 \text{ and } y \in I \text{ and not } 2|y
 \end{aligned}$$

The $power_{FI}$ operation:

$$\begin{aligned}
 power_{FI} : F \times I &\rightarrow F \cup \{\mathbf{underflow}, \mathbf{overflow}, \mathbf{pole}\} \\
 power_{FI}(x, y) &= trans_result_F(power_{FI}^*(x, y), nearest_F) \\
 &\quad \text{if } x \in F \text{ and } x \neq 0 \text{ and } y \in I \\
 &= +\infty && \text{if } x = -\infty \text{ and } y \in I \text{ and } y > 0 \text{ and } 2|y \\
 &= -\infty && \text{if } x = -\infty \text{ and } y \in I \text{ and } y > 0 \text{ and not } 2|y \\
 &= 0 && \text{if } x = -\mathbf{0} \text{ and } y \in I \text{ and } y > 0 \text{ and } 2|y \\
 &= -\mathbf{0} && \text{if } x = -\mathbf{0} \text{ and } y \in I \text{ and } y > 0 \text{ and not } 2|y \\
 &= 0 && \text{if } x = 0 \text{ and } y \in I \text{ and } y > 0 \\
 &= +\infty && \text{if } x = +\infty \text{ and } y \in I \text{ and } y > 0 \\
 &= 1 && \text{if } x \in \{-\infty -\mathbf{0}, 0, +\infty\} \text{ and } y = 0 \\
 &= 0 && \text{if } x = -\infty \text{ and } y \in I \text{ and } y < 0 \text{ and } 2|y \\
 &= -\mathbf{0} && \text{if } x = -\infty \text{ and } y \in I \text{ and } y < 0 \text{ and not } 2|y
 \end{aligned}$$

= pole ($+\infty$)	if $x = -\mathbf{0}$ and $y \in I$ and $y < 0$ and $2 y$
= pole ($-\infty$)	if $x = -\mathbf{0}$ and $y \in I$ and $y < 0$ and not $2 y$
= pole ($+\infty$)	if $x = 0$ and $y \in I$ and $y < 0$
= 0	if $x = +\infty$ and $y \in I$ and $y < 0$
= <i>result_NaN_F</i> (x)	otherwise

NOTES

- 1 $power_{FI}(x, y)$ will overflow approximately when $x^y > fmax_F$, i.e., if $x > 1$, approximately when $y > \log_x(fmax_F)$, and if $0 < x < 1$, approximately when $y < \log_x(fmax_F)$ (which is then negative). It will not overflow when $x = 0$ or when $x = 1$.
- 2 $power_I$ (in clause 5.1.4) does not allow negative *exponents* since the exact result then is not in \mathcal{Z} . $power_F$ (in clause 5.3.6.6) does not allow any negative *bases* since the (exact) result is not in \mathcal{R} unless the exponent is integer. $power_{FI}$ takes care of this latter case, where all exponents are ensured to be integers that have not arisen from implicit floating point rounding.

5.3.6.2 Natural exponentiation

The exp_F^* approximation helper function:

$$exp_F^* : F \rightarrow \mathcal{R}$$

$exp_F^*(x)$ returns a close approximation to e^x in \mathcal{R} , with maximum error $max_error_exp_F$.

Further requirements on the exp_F^* approximation helper function are:

$$exp_F^*(1) = e$$

$$exp_F^*(x) = 1 \quad \text{if } x \in F \text{ and } exp_F^*(x) \neq e^x \text{ and } \ln(1 - (epsilon_F / (2 \cdot r_F))) < x \text{ and } x < \ln(1 + (epsilon_F / 2))$$

$$exp_F^*(x) < fminD_F / 2 \quad \text{if } x \in F \text{ and } x < \ln(fminD_F) - 3$$

The exp_F operation:

$$exp_F : F \rightarrow F \cup \{\mathbf{underflow}, \mathbf{overflow}\}$$

$$exp_F(x) = trans_result_F(exp_F^*(x), nearest_F)$$

$$= 1 \quad \text{if } x \in F$$

$$= +\infty \quad \text{if } x = -\mathbf{0}$$

$$= 0 \quad \text{if } x = +\infty$$

$$= result_NaN_F(x) \quad \text{if } x = -\infty$$

$$= result_NaN_F(x) \quad \text{otherwise}$$

NOTES

- 1 $exp_F(1) = nearest_F(e)$.
- 2 $exp_F(x)$ will overflow approximately when $x > \ln(fmax_F)$.

5.3.6.3 Natural exponentiation, minus one

The $expm1_F^*$ approximation helper function:

$$expm1_F^* : F \rightarrow \mathcal{R}$$

$expm1_F^*(x)$ returns a close approximation to $e^x - 1$ in \mathcal{R} , with maximum error $max_error_expm1_F$.

Further requirements on the $expm1_F^*$ approximation helper function are:

$$\begin{aligned}
\text{expm1}_F^*(1) &= e - 1 \\
\text{expm1}_F^*(x) &= x && \text{if } x \in F \text{ and } \text{expm1}_F^*(x) \neq e^x - 1 \text{ and} \\
&&& -\text{epsilon}_F/r_F \leq x < 0.5 \cdot \text{epsilon}_F/r_F \\
\text{expm1}_F^*(x) &= -1 && \text{if } x \in F \text{ and } \text{expm1}_F^*(x) \neq e^x - 1 \text{ and} \\
&&& x < \ln(\text{epsilon}_F/(3 \cdot r_F))
\end{aligned}$$

Relationship to the exp_F^* approximation helper function:

$$\text{expm1}_F^*(x) \leq \text{exp}_F^*(x) \quad \text{if } x \in F$$

The expm1_F operation:

$$\begin{aligned}
\text{expm1}_F : F &\rightarrow F \cup \{\mathbf{overflow}\} \\
\text{expm1}_F(x) &= \text{trans_result}_F(\text{expm1}_F^*(x), \text{nearest}_F) \\
&= x && \text{if } x \in F \text{ and } |x| \geq \text{fmin}N_F \\
&= -\mathbf{0} && \text{if } x \in F \text{ and } |x| < \text{fmin}N_F \\
&= +\infty && \text{if } x = -\mathbf{0} \\
&= -1 && \text{if } x = +\infty \\
&= \text{result_NaN}_F(x) && \text{if } x = -\infty \\
&&& \text{otherwise}
\end{aligned}$$

NOTES

- 1 **underflow** is explicitly avoided. Part 1 requires that $\text{fmin}N_F \leq \text{epsilon}_F$. This Part requires that $\text{fmin}N_F < 0.5 \cdot \text{epsilon}_F/r_F$, so that underflow can be avoided here.
- 2 $\text{expm1}_F(1) = \text{nearest}_F(e - 1)$.
- 3 $\text{expm1}_F(x)$ will overflow approximately when $x > \ln(\text{fmax}_F)$.

5.3.6.4 Exponentiation of 2

The exp2_F^* approximation helper function:

$$\text{exp2}_F^* : F \rightarrow \mathcal{R}$$

$\text{exp2}_F^*(x)$ returns a close approximation to 2^x in \mathcal{R} , with maximum error max_error_exp_F .

Further requirements on the exp2_F^* approximation helper function are:

$$\begin{aligned}
\text{exp2}_F^*(x) &= 1 && \text{if } x \in F \text{ and } \text{exp2}_F^*(x) \neq 2^x \text{ and} \\
&&& \log_2(1 - (\text{epsilon}_F/(2 \cdot r_F))) < x \text{ and} \\
&&& x < \log_2(1 + (\text{epsilon}_F/2)) \\
\text{exp2}_F^*(x) &= 2^x && \text{if } x \in F \cap \mathcal{Z} \text{ and } 2^x \in F \\
\text{exp2}_F^*(x) &< \text{fmin}D_F/2 && \text{if } x \in F \text{ and } x < \log_2(\text{fmin}D_F) - 3
\end{aligned}$$

The exp2_F operation:

$$\begin{aligned}
\text{exp2}_F : F &\rightarrow F \cup \{\mathbf{underflow}, \mathbf{overflow}\} \\
\text{exp2}_F(x) &= \text{trans_result}_F(\text{exp2}_F^*(x), \text{nearest}_F) \\
&= 1 && \text{if } x \in F \\
&= +\infty && \text{if } x = -\mathbf{0} \\
&= 0 && \text{if } x = +\infty \\
&= \text{result_NaN}_F(x) && \text{if } x = -\infty \\
&&& \text{otherwise}
\end{aligned}$$

NOTE – $\text{exp2}_F(x)$ will overflow approximately when $x > \log_2(\text{fmax}_F)$.

5.3.6.5 Exponentiation of 10

The $\text{exp}10_F^*$ approximation helper function:

$$\text{exp}10_F^* : F \rightarrow \mathcal{R}$$

$\text{exp}10_F^*(x)$ returns a close approximation to 10^x in \mathcal{R} , with maximum error max_error_exp_F .

Further requirements on the $\text{exp}10_F^*$ approximation helper function are:

$$\begin{aligned} \text{exp}10_F^*(x) &= 1 && \text{if } x \in F \text{ and } \text{exp}10_F^*(x) \neq 10^x \text{ and} \\ & && \log_{10}(1 - (\text{epsilon}_F/(2 \cdot r_F))) < x \text{ and} \\ & && x < \log_{10}(1 + (\text{epsilon}_F/2)) \\ \text{exp}10_F^*(x) &= 10^x && \text{if } x \in F \cap \mathcal{Z} \text{ and } 10^x \in F \\ \text{exp}10_F^*(x) &< \text{fmin}D_F/2 && \text{if } x \in F \text{ and } x < \log_{10}(\text{fmin}D_F) - 3 \end{aligned}$$

The $\text{exp}10_F$ operation:

$$\begin{aligned} \text{exp}10_F : F &\rightarrow F \cup \{\mathbf{underflow}, \mathbf{overflow}\} \\ \text{exp}10_F(x) &= \text{trans_result}_F(\text{exp}10_F^*(x), \text{nearest}_F) && \text{if } x \in F \\ &= 1 && \text{if } x = -\mathbf{0} \\ &= +\infty && \text{if } x = +\infty \\ &= 0 && \text{if } x = -\infty \\ &= \text{result_NaN}_F(x) && \text{otherwise} \end{aligned}$$

NOTE – $\text{exp}10_F(x)$ will overflow approximately when $x > \log_{10}(\text{max}_F)$.

5.3.6.6 Exponentiation of argument base

The power_F^* approximation helper function:

$$\text{power}_F^* : F \times F \rightarrow \mathcal{R}$$

$\text{power}_F^*(x, y)$ returns a close approximation to x^y in \mathcal{R} , with maximum error max_error_power_F .

The power_F^* helper function need be defined only for first arguments that are greater than 0.

Further requirements on the power_F^* approximation helper function are:

$$\begin{aligned} \text{power}_F^*(1, y) &= 1 && \text{if } y \in F \\ \text{power}_F^*(x, 0) &= 1 && \text{if } x \in F \text{ and } x > 0 \\ \text{power}_F^*(x, 1) &= x && \text{if } x \in F \text{ and } x > 0 \\ \text{power}_F^*(x, y) &< \text{fmin}D_F/2 && \text{if } x \in F \text{ and } x > 0 \text{ and } y \in F \text{ and } x^y < \text{fmin}D_F/3 \end{aligned}$$

Relationship to the power_{FI}^* approximation helper function:

$$\text{power}_F^*(x, y) = \text{power}_{FI}^*(x, y) \quad \text{if } x \in F \text{ and } x > 0 \text{ and } y \in I \cap F$$

The power_F operation:

$$\begin{aligned} \text{power}_F : F \times F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{underflow}, \mathbf{overflow}, \mathbf{pole}\} \\ \text{power}_F(x, y) &= \text{trans_result}_F(\text{power}_F^*(x, y), \text{nearest}_F) && \text{if } x \in F \text{ and } x > 0 \text{ and } y \in F \\ &= \text{power}_F(0, y) && \text{if } x = -\mathbf{0} \text{ and } y \in F \cup \{-\infty, -\mathbf{0}, +\infty\} \\ &= \text{power}_F(x, 0) && \text{if } y = -\mathbf{0} \text{ and } x \in F \cup \{-\infty, +\infty\} \\ &= +\infty && \text{if } x = +\infty \text{ and } ((y \in F \text{ and } y > 0) \text{ or } y = +\infty) \\ &= +\infty && \text{if } x \in F \text{ and } x > 1 \text{ and } y = +\infty \\ &= 0 && \text{if } x \in F \text{ and } 0 \leq x < 1 \text{ and } y = +\infty \\ &= 0 && \text{if } x = 0 \text{ and } y \in F \text{ and } y > 0 \end{aligned}$$

$= \mathbf{pole}(+\infty)$	if $x = 0$ and $y \in F$ and $y < 0$
$= +\infty$	if $x \in F$ and $0 \leq x < 1$ and $y = -\infty$
$= 0$	if $x \in F$ and $x > 1$ and $y = -\infty$
$= 0$	if $x = +\infty$ and $((y \in F$ and $y < 0)$ or $y = -\infty)$
$= \mathit{result_NaN}\mathcal{2}_F(x, y)$	otherwise

NOTE – $\mathit{power}_F(x, y)$ will overflow approximately when $x^y > \mathit{fmax}_F$, i.e., if $x > 1$, approximately when $y > \log_x(\mathit{fmax}_F)$, and if $0 < x < 1$, approximately when $y < \log_x(\mathit{fmax}_F)$ (which is a negative number). It will not overflow when $x = 0$ or when $x = 1$.

5.3.6.7 Exponentiation of one plus the argument base, minus one

The $\mathit{power1pm1}_F^*$ approximation helper function:

$$\mathit{power1pm1}_F^* : F \times F \rightarrow \mathcal{R}$$

$\mathit{power1pm1}_F^*(x, y)$ returns a close approximation to $(1+x)^y - 1$ in \mathcal{R} , with maximum error $\mathit{max_error_power}_F$. The $\mathit{power1pm1}_F^*$ helper function need be defined only for first arguments that are greater than -1 .

Further requirements on the $\mathit{power1pm1}_F^*$ approximation helper function are:

$\mathit{power1pm1}_F^*(-1, y) = -1$	if $y \in F$ and $y > 0$
$\mathit{power1pm1}_F^*(x, y) = -1$	if $x \in F$ and $x > -1$ and $y \in F$ and $\mathit{power1pm1}_F^*(x, y) \neq (1+x)^y - 1$ and $(1+x)^y < \mathit{epsilon}_F / (3 \cdot r_F)$
$\mathit{power1pm1}_F^*(x, 1) = 1 + x$	if $x, 1+x \in F$ and $x > -1$

Relationship to the power_F^* approximation helper function:

$$\mathit{power1pm1}_F^*(x, y) \leq \mathit{power}_F^*(1+x, y) \quad \text{if } x, 1+x \in F \text{ and } x > -1 \text{ and } y \in F$$

NOTE 1 – $\mathit{power1pm1}_F^*(x, y) \approx y \cdot \ln(1+x)$ if $x \in F$ and $x > -1$ and $y \in F$ and $|y \cdot \ln(1+x)| < \mathit{epsilon}_F / r_F$.

The $\mathit{power1pm1}_F$ operation:

$$\mathit{power1pm1}_F : F \times F \rightarrow F \cup \{-0, \mathbf{invalid}, \mathbf{underflow}, \mathbf{overflow}, \mathbf{pole}\}$$

$\mathit{power1pm1}_F(x, y)$	$= \mathit{trans_result}_F(\mathit{power1pm1}_F^*(x, y), \mathit{nearest}_F)$
	if $x \in F$ and $x > -1$ and $x \neq 0$ and $y \in F$ and $y \neq 0$
$= \mathit{mul}_F(x, y)$	if $x \in \{-0, 0\}$ and $y \in F$ and $y \neq 0$
$= \mathit{mul}_F(x, y)$	if $y \in \{-0, 0\}$ and $x \in F$ and $x > -1$
$= +\infty$	if $x = +\infty$ and $((y \in F$ and $y > 0)$ or $y = +\infty)$
$= +\infty$	if $x \in F$ and $x > 0$ and $y = +\infty$
$= -1$	if $x \in F$ and $-1 \leq x < 0$ and $y = +\infty$
$= -1$	if $x = -1$ and $y \in F$ and $y > 0$
$= \mathbf{pole}(+\infty)$	if $x = -1$ and $y \in F$ and $y < 0$
$= +\infty$	if $x \in F$ and $-1 \leq x < 0$ and $y = -\infty$
$= -1$	if $x \in F$ and $x > 0$ and $y = -\infty$
$= -1$	if $x = +\infty$ and $((y \in F$ and $y < 0)$ or $y = -\infty)$
$= \mathit{result_NaN}\mathcal{2}_F(x, y)$	otherwise

NOTE 2 – $\mathit{power1pm1}_F(x, y)$ will overflow approximately when $(1+x)^y > \mathit{fmax}_F$, i.e., if $x > 0$, approximately when $y > \log_{1+x}(\mathit{fmax}_F)$, and if $-1 < x < 0$, approximately when $y < \log_{1+x}(\mathit{fmax}_F)$. It will not overflow when $x \in \{-1, 0\}$.

5.3.6.8 Natural logarithm

The ln_F^* approximation helper function:

$$ln_F^* : \mathcal{R} \rightarrow \mathcal{R}$$

$ln_F^*(x)$ returns a close approximation to $\ln(x)$ in \mathcal{R} , with maximum error $max_error_exp_F$.

A further requirement on the ln_F^* approximation helper function is:

$$ln_F^*(e) = 1$$

The ln_F operation:

$$\begin{aligned}
 ln_F : F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{pole}\} \\
 ln_F(x) &= trans_result_F(ln_F^*(x), nearest_F) && \text{if } x \in F \text{ and } x > 0 \\
 &= \mathbf{pole}(-\infty) && \text{if } x \in \{-\mathbf{0}, 0\} \\
 &= +\infty && \text{if } x = +\infty \\
 &= result_NaN_F(x) && \text{otherwise}
 \end{aligned}$$

5.3.6.9 Natural logarithm of one plus the argument

The $ln1p_F^*$ approximation helper function:

$$ln1p_F^* : \mathcal{R} \rightarrow \mathcal{R}$$

$ln1p_F^*(x)$ returns a close approximation to $\ln(1+x)$ in \mathcal{R} , with maximum error $max_error_exp_F$.

Further requirements on the $ln1p_F^*$ approximation helper function are:

$$\begin{aligned}
 ln1p_F^*(e-1) &= 1 \\
 ln1p_F^*(x) &= x && \text{if } x \in F \text{ and } ln1p_F^*(x) \neq \ln(1+x) \text{ and} \\
 & && -0.5 \cdot epsilon_F/r_F < x \leq epsilon_F/r_F
 \end{aligned}$$

Relationship to the ln_F^* approximation helper function:

$$ln1p_F^*(x) \geq ln_F^*(x) \quad \text{if } x \in F \text{ and } x > 0$$

The $ln1p_F$ operation:

$$\begin{aligned}
 ln1p_F : F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{pole}\} \\
 ln1p_F(x) &= trans_result_F(ln1p_F^*(x), nearest_F) && \text{if } x \in F \text{ and } x > -1 \text{ and } |x| \geq fminN_F \\
 &= x && \text{if } x \in F \text{ and } |x| < fminN_F \\
 &= -\mathbf{0} && \text{if } x = -\mathbf{0} \\
 &= \mathbf{pole}(-\infty) && \text{if } x = -1 \\
 &= +\infty && \text{if } x = +\infty \\
 &= result_NaN_F(x) && \text{otherwise}
 \end{aligned}$$

NOTE – **underflow** is explicitly avoided. Part 1 requires that $fminN_F \leq epsilon_F$. This Part requires that $fminN_F < 0.5 \cdot epsilon_F/r_F$, so that underflow can be avoided here.

5.3.6.10 2-logarithm

The $log2_F^*$ approximation helper function:

$$log2_F^* : F \rightarrow \mathcal{R}$$

$log2_F^*(x)$ returns a close approximation to $\log_2(x)$ in \mathcal{R} , with maximum error $max_error_exp_F$.

A further requirement on the $log2_F^*$ approximation helper function is:

$$\log 2_F^*(x) = \log_2(x) \quad \text{if } x \in F \text{ and } \log_2(x) \in \mathcal{Z}$$

The $\log 2_F$ operation:

$$\begin{aligned} \log 2_F : F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{pole}\} \\ \log 2_F(x) &= \mathit{trans_result}_F(\log 2_F^*(x), \mathit{nearest}_F) \\ &= \mathbf{pole}(-\infty) && \text{if } x \in F \text{ and } x > 0 \\ &= +\infty && \text{if } x \in \{-\mathbf{0}, 0\} \\ &= \mathit{result_NaN}_F(x) && \text{if } x = +\infty \\ & && \text{otherwise} \end{aligned}$$

5.3.6.11 10-logarithm

The $\log 10_F^*$ approximation helper function:

$$\log 10_F^* : F \rightarrow \mathcal{R}$$

$\log 10_F^*(x)$ returns a close approximation to $\log_{10}(x)$ in \mathcal{R} , with maximum error $\mathit{max_error_exp}_F$.

A further requirement on the $\log 10_F^*$ approximation helper function is:

$$\log 10_F^*(x) = \log_{10}(x) \quad \text{if } x \in F \text{ and } \log_{10}(x) \in \mathcal{Z}$$

The $\log 10_F$ operation:

$$\begin{aligned} \log 10_F : F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{pole}\} \\ \log 10_F(x) &= \mathit{trans_result}_F(\log 10_F^*(x), \mathit{nearest}_F) \\ &= \mathbf{pole}(-\infty) && \text{if } x \in F \text{ and } x > 0 \\ &= +\infty && \text{if } x \in \{-\mathbf{0}, 0\} \\ &= \mathit{result_NaN}_F(x) && \text{if } x = +\infty \\ & && \text{otherwise} \end{aligned}$$

5.3.6.12 Argument base logarithm

The $\log base_F^*$ approximation helper function:

$$\log base_F^* : F \times F \rightarrow \mathcal{R}$$

$\log base_F^*(x, y)$ returns a close approximation to $\log_x(y)$ in \mathcal{R} , with maximum error $\mathit{max_error_power}_F$.

A further requirement on the $\log base_F^*$ approximation helper function is:

$$\log base_F^*(x, x) = 1 \quad \text{if } x \in F \text{ and } x > 0 \text{ and } x \neq 1$$

The $\log base_F$ operation:

$$\begin{aligned} \log base_F : F \times F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{pole}\} \\ \log base_F(x, y) &= \mathit{trans_result}_F(\log base_F^*(x, y), \mathit{nearest}_F) \\ &= \log base_F(0, y) && \text{if } x \in F \text{ and } x > 0 \text{ and } x \neq 1 \text{ and } y \in F \text{ and } y > 0 \\ &= \log base_F(x, 0) && \text{if } x = -\mathbf{0} \text{ and } y \in F \cup \{-\infty, -\mathbf{0}, +\infty\} \\ & && \text{if } y = -\mathbf{0} \text{ and } x \in F \cup \{-\infty, +\infty\} \\ &= \mathbf{pole}(+\infty) && \text{if } x = 1 \text{ and } y \in F \text{ and } y > 1 \\ &= \mathbf{pole}(-\infty) && \text{if } x = 1 \text{ and } y \in F \text{ and } 0 \leq y < 1 \\ &= 0 && \text{if } x = +\infty \text{ and } y \in F \text{ and } y \geq 1 \\ &= +\infty && \text{if } x \in F \text{ and } 1 \leq x \text{ and } y = +\infty \\ &= -\infty && \text{if } x \in F \text{ and } 0 < x < 1 \text{ and } y = +\infty \\ &= -\mathbf{0} && \text{if } x = 0 \text{ and } y \in F \text{ and } y \geq 1 \\ &= 0 && \text{if } x = 0 \text{ and } y \in F \text{ and } 0 < y < 1 \end{aligned}$$

$$\begin{aligned}
&= \mathbf{pole}(+\infty) && \text{if } x \in F \text{ and } 0 < x < 1 \text{ and } y = 0 \\
&= \mathbf{pole}(-\infty) && \text{if } x \in F \text{ and } 1 < x \text{ and } y = 0 \\
&= -0 && \text{if } x = +\infty \text{ and } y \in F \text{ and } 0 < y < 1 \\
&= \mathit{result_NaN2}_F(x, y) && \text{otherwise}
\end{aligned}$$

5.3.6.13 Argument base logarithm of one plus each argument

The $\mathit{logbase1p1p}_F^*$ approximation helper function:

$$\mathit{logbase1p1p}_F^* : F \times F \rightarrow \mathcal{R}$$

$\mathit{logbase1p1p}_F^*(x, y)$ returns a close approximation to $\log_{(1+x)}(1+y)$ in \mathcal{R} , with maximum error $\mathit{max_error_power}_F$.

A further requirements on $\mathit{logbase1p1p}_F^*$ approximation helper function is:

$$\mathit{logbase1p1p}_F^*(x, x) = 1 \quad \text{if } x \in F \text{ and } x > -1 \text{ and } x \neq 0$$

The $\mathit{logbase1p1p}_F$ operation:

$$\mathit{logbase1p1p}_F : F \times F \rightarrow F \cup \{-0, \mathbf{invalid}, \mathbf{underflow}, \mathbf{pole}\}$$

$$\begin{aligned}
&\mathit{logbase1p1p}_F(x, y) \\
&= \mathit{trans_result}_F(\mathit{logbase1p1p}_F^*(x, y), \mathit{nearest}_F) && \text{if } x \in F \text{ and } x > -1 \text{ and } x \neq 0 \text{ and} \\
& && y \in F \text{ and } y > -1 \text{ and } y \neq 0 \\
&= \mathit{div}_F(y, x) && \text{if } x \in \{-0, 0\} \text{ and} \\
& && ((y \in F \text{ and } y > -1 \text{ and } y \neq 0) \text{ or } y = +\infty) \\
&= \mathit{div}_F(y, x) && \text{if } y \in \{-0, 0\} \text{ and} \\
& && ((x \in F \text{ and } x > -1) \text{ or } x = +\infty) \\
&= 0 && \text{if } x = +\infty \text{ and } y \in F \text{ and } y \geq 0 \\
&= +\infty && \text{if } x \in F \text{ and } 0 < x \text{ and } y = +\infty \\
&= -\infty && \text{if } x \in F \text{ and } -1 < x < 0 \text{ and } y = +\infty \\
&= -0 && \text{if } x = -1 \text{ and } y \in F \text{ and } y \geq 0 \\
&= 0 && \text{if } x = -1 \text{ and } y \in F \text{ and } -1 < y < 0 \\
&= \mathbf{pole}(+\infty) && \text{if } x \in F \text{ and } -1 < x < 0 \text{ and } y = -1 \\
&= \mathbf{pole}(-\infty) && \text{if } x \in F \text{ and } 0 < x \text{ and } y = -1 \\
&= -0 && \text{if } x = +\infty \text{ and } y \in F \text{ and } -1 < y < 0 \\
&= \mathit{result_NaN2}_F(x, y) && \text{otherwise}
\end{aligned}$$

5.3.7 Operations for hyperbolic elementary functions

There shall be two maximum error parameters for operations corresponding to the hyperbolic and inverse hyperbolic functions:

$$\mathit{max_error_sinh}_F \in F$$

$$\mathit{max_error_tanh}_F \in F$$

The $\mathit{max_error_sinh}_F$ parameter shall have a value in the interval $[0.5, 2 \cdot \mathit{rnd_error}_F]$. The $\mathit{max_error_tanh}_F$ parameter shall have a value in the interval $[\mathit{max_error_sinh}_F, 2 \cdot \mathit{rnd_error}_F]$.

5.3.7.1 Hyperbolic sine

The \sinh_F^* approximation helper function:

$$\sinh_F^* : F \rightarrow \mathcal{R}$$

$\sinh_F^*(x)$ returns a close approximation to $\sinh(x)$ in \mathcal{R} , with maximum error $max_error_sinh_F$.

Further requirements on the \sinh_F^* approximation helper function are:

$$\begin{aligned} \sinh_F^*(x) &= x && \text{if } x \in F \text{ and } \sinh_F^*(x) \neq \sinh(x) \text{ and} \\ & && |x| < \sqrt{2 \cdot \epsilonpsilon_F / r_F} \\ \sinh_F^*(-x) &= -\sinh_F^*(x) && \text{if } x \in F \end{aligned}$$

The \sinh_F operation:

$$\begin{aligned} \sinh_F : F &\rightarrow F \cup \{\mathbf{overflow}\} \\ \sinh_F(x) &= trans_result_F(\sinh_F^*(x), nearest_F) \\ &= x && \text{if } x \in F \text{ and } |x| > fminN_F \\ &= x && \text{if } x \in F \text{ and } |x| \leq fminN_F \\ &= result_NaN_F(x) && \text{if } x \in \{-\infty, -\mathbf{0}, +\infty\} \\ & && \text{otherwise} \end{aligned}$$

NOTES

- 1 **underflow** is explicitly avoided.
- 2 $\sinh_F(x)$ will overflow approximately when $|x| > \ln(2 \cdot fmax_F)$.

5.3.7.2 Hyperbolic cosine

The \cosh_F^* approximation helper function:

$$\cosh_F^* : F \rightarrow \mathcal{R}$$

$\cosh_F^*(x)$ returns a close approximation to $\cosh(x)$ in \mathcal{R} , with maximum error $max_error_sinh_F$.

Further requirements on the \cosh_F^* approximation helper function are:

$$\begin{aligned} \cosh_F^*(x) &= 1 && \text{if } x \in F \text{ and } \cosh_F^*(x) \neq \cosh(x) \text{ and} \\ & && |x| < \sqrt{\epsilonpsilon_F} \\ \cosh_F^*(-x) &= \cosh_F^*(x) && \text{if } x \in F \end{aligned}$$

Relationship to the \sinh_F^* approximation helper function:

$$\cosh_F^*(x) \geq \sinh_F^*(x) \quad \text{if } x \in F$$

The \cosh_F operation:

$$\begin{aligned} \cosh_F : F &\rightarrow F \cup \{\mathbf{overflow}\} \\ \cosh_F(x) &= trans_result_F(\cosh_F^*(x), nearest_F) \\ &= 1 && \text{if } x \in F \\ &= \mathbf{+}\infty && \text{if } x = -\mathbf{0} \\ &= result_NaN_F(x) && \text{if } x \in \{-\infty, +\infty\} \\ & && \text{otherwise} \end{aligned}$$

NOTE – $\cosh_F(x)$ overflows approximately when $|x| > \ln(2 \cdot fmax_F)$.

5.3.7.3 Hyperbolic tangent

The \tanh_F^* approximation helper function:

$$\tanh_F^* : F \rightarrow \mathcal{R}$$

$\tanh_F^*(x)$ returns a close approximation to $\tanh(x)$ in \mathcal{R} , with maximum error max_error_tanh_F .

Further requirements on the \tanh_F^* approximation helper function are:

$$\begin{aligned} \tanh_F^*(x) &= x && \text{if } x \in F \text{ and } \tanh_F^*(x) \neq \tanh(x) \text{ and} \\ & && |x| \leq \sqrt{1.5 \cdot \text{epsilon}_F / r_F} \\ \tanh_F^*(x) &= 1 && \text{if } x \in F \text{ and } \tanh_F^*(x) \neq \tanh(x) \text{ and} \\ & && x > \text{arctanh}(1 - (\text{epsilon}_F / (3 \cdot r_F))) \\ \tanh_F^*(-x) &= -\tanh_F^*(x) && \text{if } x \in F \end{aligned}$$

The \tanh_F operation:

$$\tanh_F : F \rightarrow F$$

$$\begin{aligned} \tanh_F(x) &= \text{trans_result}_F(\tanh_F^*(x), \text{nearest}_F) && \text{if } x \in F \text{ and } |x| > \text{fmin}_F \\ &= x && \text{if } x \in F \text{ and } |x| \leq \text{fmin}_F \\ &= -0 && \text{if } x = -0 \\ &= -1 && \text{if } x = -\infty \\ &= 1 && \text{if } x = +\infty \\ &= \text{result_NaN}_F(x) && \text{otherwise} \end{aligned}$$

NOTE – **underflow** is explicitly avoided.

5.3.7.4 Hyperbolic cotangent

The coth_F^* approximation helper function:

$$\text{coth}_F^* : F \rightarrow \mathcal{R}$$

$\text{coth}_F^*(x)$ returns a close approximation to $\text{coth}(x)$ in \mathcal{R} , with maximum error max_error_tanh_F .

Further requirements on the coth_F^* approximation helper function are:

$$\begin{aligned} \text{coth}_F^*(x) &= 1 && \text{if } x \in F \text{ and } \text{coth}_F^*(x) \neq \text{coth}(x) \text{ and} \\ & && x > \text{arccoth}(1 + (\text{epsilon}_F / 4)) \\ \text{coth}_F^*(-x) &= -\text{coth}_F^*(x) && \text{if } x \in F \end{aligned}$$

The coth_F operation:

$$\begin{aligned} \text{coth}_F : F &\rightarrow F \cup \{\text{pole}, \text{overflow}\} \\ \text{coth}_F(x) &= \text{trans_result}_F(\text{coth}_F^*(x), \text{nearest}_F) && \text{if } x \in F \text{ and } x \neq 0 \\ &= \text{pole}(+\infty) && \text{if } x = 0 \\ &= \text{pole}(-\infty) && \text{if } x = -0 \\ &= -1 && \text{if } x = -\infty \\ &= 1 && \text{if } x = +\infty \\ &= \text{result_NaN}_F(x) && \text{otherwise} \end{aligned}$$

NOTE – $\text{coth}_F(x)$ overflows approximately when $|1/x| > \text{fmax}_F$.

5.3.7.5 Hyperbolic secant

The $sech_F^*$ approximation helper function:

$$sech_F^* : F \rightarrow \mathcal{R}$$

$sech_F^*(x)$ returns a close approximation to $sech(x)$ in \mathcal{R} , with maximum error $max_error_tanh_F$.

Further requirements on the $sech_F^*$ approximation helper function are:

$$\begin{aligned} sech_F^*(x) &= 1 && \text{if } x \in F \text{ and } sech_F^*(x) \neq sech(x) \text{ and} \\ & && |x| < \sqrt{\epsilonpsilon_F/r_F} \\ sech_F^*(-x) &= sech_F^*(x) && \text{if } x \in F \\ sech_F^*(x) &< fminD_F/2 && \text{if } x \in F \text{ and } x > 2 - \ln(fminD_F/4) \end{aligned}$$

The $sech_F$ operation:

$$\begin{aligned} sech_F : F &\rightarrow F \cup \{\mathbf{underflow}\} \\ sech_F(x) &= trans_result_F(sech_F^*(x), nearest_F) && \text{if } x \in F \\ &= 1 && \text{if } x = -\mathbf{0} \\ &= 0 && \text{if } x \in \{-\infty, +\infty\} \\ &= result_NaN_F(x) && \text{otherwise} \end{aligned}$$

5.3.7.6 Hyperbolic cosecant

The $csch_F^*$ approximation helper function:

$$csch_F^* : F \rightarrow \mathcal{R}$$

$csch_F^*(x)$ returns a close approximation to $csch(x)$ in \mathcal{R} , with maximum error $max_error_tanh_F$.

Further requirements on the $csch_F^*$ approximation helper function are:

$$\begin{aligned} csch_F^*(-x) &= -csch_F^*(x) && \text{if } x \in F \\ csch_F^*(x) &< fminD_F/2 && \text{if } x \in F \text{ and } x > 2 - \ln(fminD_F/4) \end{aligned}$$

Relationship to the $sech_F^*$ approximation helper function:

$$csch_F^*(x) \geq sech_F^*(x) \quad \text{if } x \in F \text{ and } x > 0$$

The $csch_F$ operation:

$$\begin{aligned} csch_F : F &\rightarrow F \cup \{\mathbf{underflow}, \mathbf{overflow}, \mathbf{pole}\} \\ csch_F(x) &= trans_result_F(csch_F^*(x), nearest_F) && \text{if } x \in F \text{ and } x \neq 0 \\ &= div_F(1, x) && \text{if } x \in \{-\infty, -\mathbf{0}, 0, +\infty\} \\ &= result_NaN_F(x) && \text{otherwise} \end{aligned}$$

NOTE – $csch_F(x)$ overflows approximately when $|1/x| > fmax_F$.

5.3.7.7 Inverse hyperbolic sine

The $arcsinh_F^*$ approximation helper function:

$$arcsinh_F^* : F \rightarrow \mathcal{R}$$

$arcsinh_F^*(x)$ returns a close approximation to $arcsinh(x)$ in \mathcal{R} , with maximum error $max_error_sinh_F$.

Further requirements on the $arcsinh_F^*$ approximation helper function are:

$$\begin{aligned} \operatorname{arcsinh}_F^*(x) &= x && \text{if } x \in F \text{ and } \operatorname{arcsinh}_F^*(x) \neq \operatorname{arcsinh}(x) \text{ and} \\ & && |x| \leq \sqrt{3 \cdot \operatorname{epsilon}_F / r_F} \\ \operatorname{arcsinh}_F^*(-x) &= -\operatorname{arcsinh}_F^*(x) && \text{if } x \in F \end{aligned}$$

The $\operatorname{arcsinh}_F$ operation:

$$\begin{aligned} \operatorname{arcsinh}_F &: F \rightarrow F \\ \operatorname{arcsinh}_F(x) &= \operatorname{trans_result}_F(\operatorname{arcsinh}_F^*(x), \operatorname{nearest}_F) \\ &= x && \text{if } x \in F \text{ and } |x| > \operatorname{fmin}N_F \\ &= x && \text{if } x \in F \text{ and } |x| \leq \operatorname{fmin}N_F \\ &= \operatorname{result_NaN}_F(x) && \text{if } x \in \{-\infty, -\mathbf{0}, +\infty\} \\ & && \text{otherwise} \end{aligned}$$

NOTE – **underflow** is explicitly avoided.

5.3.7.8 Inverse hyperbolic cosine

The $\operatorname{arccosh}_F^*$ approximation helper function:

$$\operatorname{arccosh}_F^* : F \rightarrow \mathcal{R}$$

$\operatorname{arccosh}_F^*(x)$ returns a close approximation to $\operatorname{arccosh}(x)$ in \mathcal{R} , with maximum error $\operatorname{max_error_sinh}_F$.

Relationship to the $\operatorname{arcsinh}_F^*$ approximation helper function:

$$\operatorname{arccosh}_F^*(x) \leq \operatorname{arcsinh}_F^*(x)$$

The $\operatorname{arccosh}_F$ operation:

$$\begin{aligned} \operatorname{arccosh}_F &: F \rightarrow F \cup \{\mathbf{invalid}\} \\ \operatorname{arccosh}_F(x) &= \operatorname{trans_result}_F(\operatorname{arccosh}_F^*(x), \operatorname{nearest}_F) \\ &= +\infty && \text{if } x \in F \text{ and } x \geq 1 \\ &= \operatorname{result_NaN}_F(x) && \text{if } x = +\infty \\ & && \text{otherwise} \end{aligned}$$

5.3.7.9 Inverse hyperbolic tangent

The $\operatorname{arctanh}_F^*$ approximation helper function:

$$\operatorname{arctanh}_F^* : F \rightarrow \mathcal{R}$$

$\operatorname{arctanh}_F^*(x)$ returns a close approximation to $\operatorname{arctanh}(x)$ in \mathcal{R} , with maximum error $\operatorname{max_error_tanh}_F$.

Further requirements on the $\operatorname{arctanh}_F^*$ approximation helper function are:

$$\begin{aligned} \operatorname{arctanh}_F^*(x) &= x && \text{if } x \in F \text{ and } \operatorname{arctanh}_F^*(x) \neq \operatorname{arctanh}(x) \text{ and} \\ & && |x| < \sqrt{\operatorname{epsilon}_F / r_F} \\ \operatorname{arctanh}_F^*(-x) &= -\operatorname{arctanh}_F^*(x) && \text{if } x \in F \end{aligned}$$

The $\operatorname{arctanh}_F$ operation:

$$\begin{aligned} \operatorname{arctanh}_F &: F \rightarrow F \cup \{\mathbf{invalid}, \mathbf{pole}\} \\ \operatorname{arctanh}_F(x) &= \operatorname{trans_result}_F(\operatorname{arctanh}_F^*(x), \operatorname{nearest}_F) \\ &= x && \text{if } x \in F \text{ and } \operatorname{fmin}N_F < |x| < 1 \\ &= -\mathbf{0} && \text{if } x \in F \text{ and } |x| \leq \operatorname{fmin}N_F \\ &= \mathbf{pole}(+\infty) && \text{if } x = -\mathbf{0} \\ &= \mathbf{pole}(-\infty) && \text{if } x = 1 \\ &= \operatorname{result_NaN}_F(x) && \text{if } x = -1 \\ & && \text{otherwise} \end{aligned}$$

NOTE – **underflow** is explicitly avoided.

5.3.7.10 Inverse hyperbolic cotangent

The arccoth_F^* approximation helper function:

$$\text{arccoth}_F^* : F \rightarrow \mathcal{R}$$

$\text{arccoth}_F^*(x)$ returns a close approximation to $\text{arccoth}(x)$ in \mathcal{R} , with maximum error max_error_tanh_F .

A further requirements on the arccoth_F^* approximation helper function is:

$$\text{arccoth}_F^*(-x) = -\text{arccoth}_F^*(x) \quad \text{if } x \in F$$

The arccoth_F operation:

$$\begin{aligned} \text{arccoth}_F : F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{pole}\} \\ \text{arccoth}_F(x) &= \text{trans_result}_F(\text{arccoth}_F^*(x), \text{nearest}_F) \\ &\quad \text{if } x \in F \text{ and } |x| > 1 \\ &= \mathbf{pole}(+\infty) && \text{if } x = 1 \\ &= \mathbf{pole}(-\infty) && \text{if } x = -1 \\ &= -\mathbf{0} && \text{if } x = -\infty \\ &= \mathbf{0} && \text{if } x = +\infty \\ &= \text{result_NaN}_F(x) && \text{otherwise} \end{aligned}$$

NOTE – There is no **underflow** for this operation for most kinds of floating point types, e.g. IEC 60559 ones.

5.3.7.11 Inverse hyperbolic secant

The arcsech_F^* approximation helper function:

$$\text{arcsech}_F^* : F \rightarrow \mathcal{R}$$

$\text{arcsech}_F^*(x)$ returns a close approximation to $\text{arcsech}(x)$ in \mathcal{R} , with maximum error max_error_tanh_F .

The arcsech_F operation:

$$\begin{aligned} \text{arcsech}_F : F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{pole}\} \\ \text{arcsech}_F(x) &= \text{trans_result}_F(\text{arcsech}_F^*(x), \text{nearest}_F) \\ &\quad \text{if } x \in F \text{ and } 0 < x \leq 1 \\ &= \mathbf{pole}(+\infty) && \text{if } x \in \{-\mathbf{0}, 0\} \\ &= \text{result_NaN}_F(x) && \text{otherwise} \end{aligned}$$

5.3.7.12 Inverse hyperbolic cosecant

The arccsch_F^* approximation helper function:

$$\text{arccsch}_F^* : F \rightarrow \mathcal{R}$$

$\text{arccsch}_F^*(x)$ returns a close approximation to $\text{arccsch}(x)$ in \mathcal{R} , with maximum error max_error_tanh_F .

A further requirements on the arccsch_F^* approximation helper function is:

$$\text{arccsch}_F^*(-x) = -\text{arccsch}_F^*(x) \quad \text{if } x \in F$$

Relationship to the arcsinh_F^* approximation helper function:

$$\text{arccsch}_F^*(1) = \text{arcsinh}_F^*(1)$$

The arccsch_F operation:

$$\text{arccsch}_F : F \rightarrow F \cup \{\mathbf{underflow}, \mathbf{pole}\}$$

$$\begin{aligned}
\operatorname{arccsch}_F(x) &= \operatorname{trans_result}_F(\operatorname{arccsch}_F^*(x), \operatorname{nearest}_F) \\
&\quad \text{if } x \in F \text{ and } x \neq 0 \\
&= \operatorname{div}_F(1, x) && \text{if } x \in \{-\infty, -0, 0, +\infty\} \\
&= \operatorname{result_NaN}_F(x) && \text{otherwise}
\end{aligned}$$

NOTE – There is no **underflow** for this operation for most kinds of floating point types, e.g. IEC 60559 ones.

5.3.8 Introduction to operations for trigonometric elementary functions

Two different operations for each of sin, cos, tan, cot, sec, csc, arcsin, arccos, arctan, arccot, arcctg, arcsec, and arccsc are specified. One version for radians and one version where the angular unit is given as a parameter.

For use in the specifications below, define the following mathematical functions:

$$\begin{aligned}
\operatorname{rad} &: \mathcal{R} \rightarrow \mathcal{R} \\
\operatorname{axis_rad} &: \mathcal{R} \rightarrow \{(1, 0), (0, 1), (-1, 0), (0, -1)\} \times \mathcal{R} \\
\operatorname{arc} &: \mathcal{R} \times \mathcal{R} \rightarrow \mathcal{R}
\end{aligned}$$

The *rad*, angular value normalisation, function is defined by

$$\operatorname{rad}(x) = x - \operatorname{round}(x/(2 \cdot \pi)) \cdot 2 \cdot \pi$$

The *axis_rad* function is defined by

$$\begin{aligned}
\operatorname{axis_rad}(x) &= ((1, 0), \operatorname{arcsin}(\sin(x))) && \text{if } \cos(x) \geq 1/\sqrt{2} \\
&= ((0, 1), \operatorname{arcsin}(\cos(x))) && \text{if } \sin(x) > 1/\sqrt{2} \\
&= ((-1, 0), \operatorname{arcsin}(\sin(x))) && \text{if } \cos(x) \leq -1/\sqrt{2} \\
&= ((0, -1), \operatorname{arcsin}(\cos(x))) && \text{if } \sin(x) < -1/\sqrt{2}
\end{aligned}$$

The *arc*, angle, function is defined by

$$\begin{aligned}
\operatorname{arc}(x, y) &= -\operatorname{arccos}(x/\sqrt{x^2 + y^2}) && \text{if } y < 0 \\
&= \operatorname{arccos}(x/\sqrt{x^2 + y^2}) && \text{if } y \geq 0
\end{aligned}$$

5.3.9 Operations for radian trigonometric elementary functions

There shall be one radian big-angle parameter:

$$\operatorname{big_angle_r}_F \in F$$

It should have the following default value:

$$\operatorname{big_angle_r}_F = r_F^{\lceil p_F/2 \rceil}$$

A binding or implementation can include a method to change the value the radian big-angle parameter. This method should only allow the value of this parameter to be set to a value greater than $2 \cdot \pi$ and such that $\operatorname{ulp}_F(\operatorname{big_angle_r}_F) < \pi/1000$.

NOTES

- 1 Part 1 requires that $p_F \geq 2$, but see also A.5.2.0.2 in Part 1.
- 2 This Part requires that $p_F \geq 2 \cdot \max\{1, \lceil \log_{r_F}(2 \cdot \pi) \rceil\}$, in order to allow at least the first two cycles to be in the interval $[-\operatorname{big_angle_r}_F, \operatorname{big_angle_r}_F]$.
- 3 In order to allow $\operatorname{ulp}_F(\operatorname{big_angle_r}_F) < \pi/1000$, $p_F \geq 2 + \log_{r_F}(1000)$ should hold.

There shall be two maximum error parameters for radian trigonometric operations:

$$\begin{aligned}
\operatorname{max_error_sin}_F &\in F \\
\operatorname{max_error_tan}_F &\in F
\end{aligned}$$

The $max_error_sin_F$ parameter shall have a value in the interval $[0.5, 1.5 \cdot rnd_error_F]$. The $max_error_tan_F$ parameter shall have a value in the interval $[max_error_sin_F, 2 \cdot rnd_error_F]$.

5.3.9.1 Radian angle normalisation

The rad_F^* approximation helper function:

$$rad_F^* : \mathcal{R} \rightarrow \mathcal{R}$$

$rad_F^*(x)$ returns a close approximation to $rad(x)$ in \mathcal{R} , if $|x| \leq big_angle_r_F$, with maximum error $max_error_sin_F$.

The $axis_rad_F^*$ approximation helper function:

$$axis_rad_F^* : \mathcal{R} \rightarrow \{(1, 0), (0, 1), (-1, 0), (0, -1)\} \times \mathcal{R}$$

$axis_rad_F^*(x)$ returns a close approximation to $axis_rad(x)$, if $x \leq big_angle_r_F$. The approximation consists of that the second part of the result (the offset from the indicated axis) is approximate.

Further requirements on the rad_F^* and $axis_rad_F^*$ approximation helper functions are:

$$\begin{aligned} rad_F^*(x) &= x && \text{if } |x| < \pi \\ snd(axis_rad_F^*(x)) &= rad_F^*(x) && \text{if } fst(axis_rad_F^*(x)) = (1, 0) \end{aligned}$$

The rad_F operation:

$$\begin{aligned} rad_F : F &\rightarrow F \cup \{\mathbf{underflow}, \mathbf{absolute_precision_underflow}\} \\ rad_F(x) &= trans_result_F(rad_F^*(x), nearest_F) && \text{if } x \in F \text{ and } |x| > fminN_F \text{ and } |x| \leq big_angle_r_F \\ &= x && \text{if } (x \in F \text{ and } |x| \leq fminN_F) \text{ or } x = -\mathbf{0} \\ &= \mathbf{absolute_precision_underflow}(qNaN) && \\ &= result_NaN_F(x) && \text{if } x \in F \text{ and } |x| > big_angle_r_F \\ & && \text{otherwise} \end{aligned}$$

The $axis_rad_F$ operation:

$$\begin{aligned} axis_rad_F : F &\rightarrow ((F \times F) \times F) \cup \{\mathbf{absolute_precision_underflow}\} \\ axis_rad_F(x) &= (fst(axis_rad_F^*(x)), trans_result_F(snd(axis_rad_F^*(x)), nearest_F)) && \text{if } x \in F \text{ and } |x| > fminN_F \text{ and } |x| \leq big_angle_r_F \\ &= ((1, 0), x) && \text{if } (x \in F \text{ and } |x| \leq fminN_F) \text{ or } x = -\mathbf{0} \\ &= \mathbf{absolute_precision_underflow}((qNaN, qNaN), qNaN) && \\ & && \text{if } x \in F \text{ and } |x| > big_angle_r_F \\ &= ((qNaN, qNaN), qNaN) && \\ & && \text{if } x \text{ is a quiet NaN} \\ &= \mathbf{invalid}((qNaN, qNaN), qNaN) && \\ & && \text{otherwise} \end{aligned}$$

NOTE – rad_F is simpler, easier to use, but less accurate than $axis_rad_F$. The latter may still not be sufficient for implementing the radian trigonometric operations to less than the maximum error stated by the parameters.

5.3.9.2 Radian sine

The \sin_F^* approximation helper function:

$$\sin_F^* : \mathcal{R} \rightarrow \mathcal{R}$$

$\sin_F^*(x)$ returns a close approximation to $\sin(x)$ in \mathcal{R} if $|x| \leq \text{big_angle_r}_F$, with maximum error max_error_sin_F .

Further requirements on the \sin_F^* approximation helper function are:

$$\begin{aligned} \sin_F^*(n \cdot 2 \cdot \pi + \pi/6) &= 1/2 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/6| \leq \text{big_angle_r}_F \\ \sin_F^*(n \cdot 2 \cdot \pi + \pi/2) &= 1 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/2| \leq \text{big_angle_r}_F \\ \sin_F^*(n \cdot 2 \cdot \pi + 5 \cdot \pi/6) &= 1/2 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + 5 \cdot \pi/6| \leq \text{big_angle_r}_F \\ \sin_F^*(x) &= x && \text{if } \sin_F^*(x) \neq \sin(x) \text{ and } |x| \leq \sqrt{3 \cdot \text{epsilon}_F / r_F} \\ \sin_F^*(-x) &= -\sin_F^*(x) \end{aligned}$$

The \sin_F operation:

$$\sin_F : F \rightarrow F \cup \{\text{underflow, absolute_precision_underflow}\}$$

$$\begin{aligned} \sin_F(x) &= \text{trans_result}_F(\sin_F^*(x), \text{nearest}_F) && \text{if } x \in F \text{ and } \text{fmin}N_F < |x| \text{ and } |x| \leq \text{big_angle_r}_F \\ &= \text{rad}_F(x) && \text{otherwise} \end{aligned}$$

NOTE – **underflow** is here explicitly avoided for denormal arguments, but the operation may **underflow** for other arguments.

5.3.9.3 Radian cosine

The \cos_F^* approximation helper function:

$$\cos_F^* : \mathcal{R} \rightarrow \mathcal{R}$$

$\cos_F^*(x)$ returns a close approximation to $\cos(x)$ in \mathcal{R} if $|x| \leq \text{big_angle_r}_F$, with maximum error max_error_sin_F .

Further requirements on the \cos_F^* approximation helper function are:

$$\begin{aligned} \cos_F^*(n \cdot 2 \cdot \pi) &= 1 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi| \leq \text{big_angle_r}_F \\ \cos_F^*(n \cdot 2 \cdot \pi + \pi/3) &= 1/2 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/3| \leq \text{big_angle_r}_F \\ \cos_F^*(n \cdot 2 \cdot \pi + 2 \cdot \pi/3) &= -1/2 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + 2 \cdot \pi/3| \leq \text{big_angle_r}_F \\ \cos_F^*(n \cdot 2 \cdot \pi + \pi) &= -1 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi| \leq \text{big_angle_r}_F \\ \cos_F^*(x) &= 1 && \text{if } \cos_F^*(x) \neq \cos(x) \text{ and } |x| < \sqrt{\text{epsilon}_F / r_F} \\ \cos_F^*(-x) &= \cos_F^*(x) \end{aligned}$$

The \cos_F operation:

$$\cos_F : F \rightarrow F \cup \{\text{underflow, absolute_precision_underflow}\}$$

$$\begin{aligned} \cos_F(x) &= \text{trans_result}_F(\cos_F^*(x), \text{nearest}_F) && \text{if } x \in F \text{ and } |x| \leq \text{big_angle_r}_F \\ &= 1 && \text{if } x = -\mathbf{0} \\ &= \text{rad}_F(x) && \text{otherwise} \end{aligned}$$

5.3.9.4 Radian tangent

The \tan_F^* approximation helper function:

$$\tan_F^* : \mathcal{R} \rightarrow \mathcal{R}$$

$\tan_F^*(x)$ returns a close approximation to $\tan(x)$ in \mathcal{R} if $|x| \leq \text{big_angle_r}_F$, with maximum error max_error_tan_F .

Further requirements on the \tan_F^* approximation helper function are:

$$\begin{aligned} \tan_F^*(n \cdot 2 \cdot \pi + \pi/4) &= 1 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/4| \leq \text{big_angle_r}_F \\ \tan_F^*(n \cdot 2 \cdot \pi + 3 \cdot \pi/4) &= -1 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + 3 \cdot \pi/4| \leq \text{big_angle_r}_F \\ \tan_F^*(x) &= x && \text{if } \tan_F^*(x) \neq \tan(x) \text{ and } |x| < \sqrt{\text{epsilon}_F/r_F} \\ \tan_F^*(-x) &= -\tan_F^*(x) && \end{aligned}$$

The \tan_F operation:

$$\begin{aligned} \tan_F : F &\rightarrow F \cup \{\text{underflow}, \text{overflow}, \text{absolute_precision_underflow}\} \\ \tan_F(x) &= \text{trans_result}_F(\tan_F^*(x), \text{nearest}_F) && \text{if } x \in F \text{ and } \text{fmin}N_F < |x| \text{ and } |x| \leq \text{big_angle_r}_F \\ &= \text{rad}_F(x) && \text{otherwise} \end{aligned}$$

NOTE – **underflow** is explicitly avoided for denormal arguments, but the operation may **underflow** for other arguments.

5.3.9.5 Radian cotangent

The \cot_F^* approximation helper function:

$$\cot_F^* : \mathcal{R} \rightarrow \mathcal{R}$$

$\cot_F^*(x)$ returns a close approximation to $\cot(x)$ in \mathcal{R} if $|x| \leq \text{big_angle_r}_F$, with maximum error max_error_tan_F .

Further requirements on the \cot_F^* approximation helper function are:

$$\begin{aligned} \cot_F^*(n \cdot 2 \cdot \pi + \pi/4) &= 1 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/4| \leq \text{big_angle_r}_F \\ \cot_F^*(n \cdot 2 \cdot \pi + 3 \cdot \pi/4) &= -1 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + 3 \cdot \pi/4| \leq \text{big_angle_r}_F \\ \cot_F^*(-x) &= -\cot_F^*(x) && \end{aligned}$$

The \cot_F operation:

$$\begin{aligned} \cot_F : F &\rightarrow F \cup \{\text{underflow}, \text{overflow}, \text{pole}, \text{absolute_precision_underflow}\} \\ \cot_F(x) &= \text{trans_result}_F(\cot_F^*(x), \text{nearest}_F) && \text{if } x \in F \text{ and } x \neq 0 \text{ and } |x| \leq \text{big_angle_r}_F \\ &= \text{pole}(+\infty) && \text{if } x = 0 \\ &= \text{pole}(-\infty) && \text{if } x = -0 \\ &= \text{rad}_F(x) && \text{otherwise} \end{aligned}$$

5.3.9.6 Radian secant

The \sec_F^* approximation helper function:

$$\sec_F^* : \mathcal{R} \rightarrow \mathcal{R}$$

$\sec_F^*(x)$ returns a close approximation to $\sec(x)$ in \mathcal{R} if $|x| \leq \text{big_angle_r}_F$, with maximum error max_error_tan_F .

Further requirements on the \sec_F^* approximation helper function are:

$$\begin{array}{ll}
\sec_F^*(n \cdot 2 \cdot \pi) = 1 & \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi| \leq \text{big_angle_r}_F \\
\sec_F^*(n \cdot 2 \cdot \pi + \pi/3) = 2 & \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/3| \leq \text{big_angle_r}_F \\
\sec_F^*(n \cdot 2 \cdot \pi + 2 \cdot \pi/3) = -2 & \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + 2 \cdot \pi/3| \leq \text{big_angle_r}_F \\
\sec_F^*(n \cdot 2 \cdot \pi + \pi) = -1 & \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi| \leq \text{big_angle_r}_F \\
\sec_F^*(x) = 1 & \text{if } \sec_F^*(x) \neq \sec(x) \text{ and } |x| < \sqrt{\text{epsilon}_F} \\
\sec_F^*(-x) = \sec_F^*(x) &
\end{array}$$

The \sec_F operation:

$$\begin{array}{ll}
\sec_F : F \rightarrow F \cup \{\text{overflow, absolute_precision_underflow}\} \\
\sec_F(x) & = \text{trans_result}_F(\sec_F^*(x), \text{nearest}_F) \\
& \text{if } x \in F \text{ and } |x| \leq \text{big_angle_r}_F \\
& = 1 & \text{if } x = -\mathbf{0} \\
& = \text{rad}_F(x) & \text{otherwise}
\end{array}$$

5.3.9.7 Radian cosecant

The \csc_F^* approximation helper function:

$$\csc_F^* : \mathcal{R} \rightarrow \mathcal{R}$$

$\csc_F^*(x)$ returns a close approximation to $\csc(x)$ in \mathcal{R} if $|x| \leq \text{big_angle_r}_F$, with maximum error max_error_tan_F .

Further requirements on the \csc_F^* approximation helper function are:

$$\begin{array}{ll}
\csc_F^*(n \cdot 2 \cdot \pi + \pi/6) = 2 & \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/6| \leq \text{big_angle_r}_F \\
\csc_F^*(n \cdot 2 \cdot \pi + \pi/2) = 1 & \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/2| \leq \text{big_angle_r}_F \\
\csc_F^*(n \cdot 2 \cdot \pi + 5 \cdot \pi/6) = 2 & \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + 5 \cdot \pi/6| \leq \text{big_angle_r}_F \\
\csc_F^*(-x) = -\csc_F^*(x) &
\end{array}$$

The \csc_F operation:

$$\begin{array}{ll}
\csc_F : F \rightarrow F \cup \{\text{overflow, pole, absolute_precision_underflow}\} \\
\csc_F(x) & = \text{trans_result}_F(\csc_F^*(x), \text{nearest}_F) \\
& \text{if } x \in F \text{ and } x \neq 0 \text{ and } |x| \leq \text{big_angle_r}_F \\
& = \text{pole}(+\infty) & \text{if } x = 0 \\
& = \text{pole}(-\infty) & \text{if } x = -\mathbf{0} \\
& = \text{rad}_F(x) & \text{otherwise}
\end{array}$$

5.3.9.8 Radian cosine with sine

$$\begin{array}{ll}
\text{cossin}_F : F \rightarrow (F \times F) \cup \{\text{underflow, absolute_precision_underflow}\} \\
\text{cossin}_F(x) & = (\text{cos}_F(x), \text{sin}_F(x))
\end{array}$$

5.3.9.9 Radian arc sine

The \arcsin_F^* approximation helper function:

$$\arcsin_F^* : F \rightarrow \mathcal{R}$$

$\arcsin_F^*(x)$ returns a close approximation to $\arcsin(x)$ in \mathcal{R} , with maximum error max_error_sin_F .

Further requirements on the \arcsin_F^* approximation helper function are:

$$\begin{aligned}
\arcsin_F^*(1/2) &= \pi/6 \\
\arcsin_F^*(1) &= \pi/2 \\
\arcsin_F^*(x) &= x && \text{if } \arcsin_F^*(x) \neq \arcsin(x) \text{ and} \\
&&& |x| < \sqrt{2 \cdot \text{epsilon}_F / r_F} \\
\arcsin_F^*(-x) &= -\arcsin_F^*(x)
\end{aligned}$$

Range limitation:

$$\arcsin_F^\#(x) = \max\{up_F(-\pi/2), \min\{\arcsin_F^*(x), down_F(\pi/2)\}\}$$

The \arcsin_F operation:

$$\begin{aligned}
\arcsin_F : F &\rightarrow F \cup \{\mathbf{invalid}\} \\
\arcsin_F(x) &= trans_result_F(\arcsin_F^\#(x), nearest_F) \\
&= x && \text{if } x \in F \text{ and } fminN_F < |x| \leq 1 \\
&= result_NaN_F(x) && \text{if } (x \in F \text{ and } |x| \leq fminN_F) \text{ or } x = -\mathbf{0} \\
&&& \text{otherwise}
\end{aligned}$$

NOTE – **underflow** is explicitly avoided.

5.3.9.10 Radian arc cosine

The \arccos_F^* approximation helper function:

$$\arccos_F^* : F \rightarrow \mathcal{R}$$

$\arccos_F^*(x)$ returns a close approximation to $\arccos(x)$ in \mathcal{R} , with maximum error $max_error_sin_F$.

Further requirements on the \arccos_F^* approximation helper function are:

$$\begin{aligned}
\arccos_F^*(1/2) &= \pi/3 \\
\arccos_F^*(0) &= \pi/2 \\
\arccos_F^*(-1/2) &= 2 \cdot \pi/3 \\
\arccos_F^*(-1) &= \pi
\end{aligned}$$

Range limitation:

$$\arccos_F^\#(x) = \min\{\arccos_F^*(x), down_F(\pi)\}$$

The \arccos_F operation:

$$\begin{aligned}
\arccos_F : F &\rightarrow F \cup \{\mathbf{invalid}\} \\
\arccos_F(x) &= trans_result_F(\arccos_F^\#(x), nearest_F) \\
&= \arccos_F(0) && \text{if } x \in F \text{ and } -1 \leq x \leq 1 \\
&= result_NaN_F(x) && \text{if } x = -\mathbf{0} \\
&&& \text{otherwise}
\end{aligned}$$

5.3.9.11 Radian arc tangent

The \arctan_F^* approximation helper function:

$$\arctan_F^* : F \rightarrow \mathcal{R}$$

$\arctan_F^*(x)$ returns a close approximation to $\arctan(x)$ in \mathcal{R} , with maximum error $max_error_tan_F$.

Further requirements on the \arctan_F^* approximation helper function are:

$$\begin{aligned}
\arctan_F^*(1) &= \pi/4 \\
\arctan_F^*(x) &= x && \text{if } \arctan_F^*(x) \neq \arctan(x) \text{ and} \\
&&& |x| \leq \sqrt{1.5 \cdot \text{epsilon}_F / r_F} \\
\arctan_F^*(x) &= \pi/2 && \text{if } \arctan_F^*(x) \neq \arctan(x) \text{ and } x > 3 \cdot r_F / \text{epsilon}_F \\
\arctan_F^*(-x) &= -\arctan_F^*(x)
\end{aligned}$$

Range limitation:

$$\arctan_F^\#(x) = \max\{\text{up}_F(-\pi/2), \min\{\arctan_F^*(x), \text{down}_F(\pi/2)\}\}$$

The \arctan_F operation:

$$\begin{aligned}
\arctan_F &: F \rightarrow F \\
\arctan_F(x) &= \text{trans_result}_F(\arctan_F^\#(x), \text{nearest}_F) \\
&&& \text{if } x \in F \text{ and } \text{fmin}N_F < |x| \\
&= x && \text{if } (x \in F \text{ and } |x| \leq \text{fmin}N_F) \text{ or } x = -0 \\
&= \text{up}_F(-\pi/2) && \text{if } x = -\infty \\
&= \text{down}_F(\pi/2) && \text{if } x = +\infty \\
&= \text{result_NaN}_F(x) && \text{otherwise}
\end{aligned}$$

NOTES

- 1 $\arctan_F(x) \approx \text{arc}_F(1, x)$
- 2 **underflow** is explicitly avoided.

5.3.9.12 Radian arc cotangent

This clause specifies two inverse cotangent operations. One approximating the continuous (but not sign symmetric) arccot , the other approximating the sign symmetric (but discontinuous at 0) arcctg .

The arccot_F^* approximation helper function:

$$\text{arccot}_F^* : F \rightarrow \mathcal{R}$$

$\text{arccot}_F^*(x)$ returns a close approximation to $\text{arccot}(x)$ in \mathcal{R} , with maximum error max_error_tan_F .

The arcctg_F^* approximation helper function:

$$\text{arcctg}_F^* : F \rightarrow \mathcal{R}$$

$\text{arcctg}_F^*(x)$ returns a close approximation to $\text{arcctg}(x)$ in \mathcal{R} , with maximum error max_error_tan_F .

Further requirements on the arccot_F^* and arcctg_F^* approximation helper functions are:

$$\begin{aligned}
\text{arccot}_F^*(1) &= \pi/4 \\
\text{arccot}_F^*(0) &= \pi/2 \\
\text{arccot}_F^*(-1) &= 3 \cdot \pi/4 \\
\text{arccot}_F^*(x) &= \pi && \text{if } \text{arccot}_F^*(x) \neq \text{arccot}(x) \text{ and } x < -3 \cdot r_F / \text{epsilon}_F \\
\text{arcctg}_F^*(x) &= \text{arccot}_F^*(x) && \text{if } x \geq 0 \\
\text{arcctg}_F^*(-x) &= -\text{arcctg}_F^*(x)
\end{aligned}$$

Range limitation:

$$\begin{aligned}
\text{arccot}_F^\#(x) &= \min\{\text{arccot}_F^*(x), \text{down}_F(\pi)\} \\
\text{arcctg}_F^\#(x) &= \max\{\text{up}_F(-\pi/2), \min\{\text{arcctg}_F^*(x), \text{down}_F(\pi/2)\}\}
\end{aligned}$$

The arccot_F operation:

$$\text{arccot}_F : F \rightarrow F \cup \{\mathbf{underflow}\}$$

$$\begin{aligned}
\operatorname{arccot}_F(x) &= \operatorname{trans_result}_F(\operatorname{arccot}_F^\#(x)) \\
& \quad \text{if } x \in F \\
&= \operatorname{nearest}_F(\pi/2) & \text{if } x = -\mathbf{0} \\
&= \operatorname{down}_F(\pi) & \text{if } x = -\infty \\
&= 0 & \text{if } x = +\infty \\
&= \operatorname{result_NaN}_F(x) & \text{otherwise}
\end{aligned}$$

NOTES

- 1 $\operatorname{arccot}_F(x) \approx \operatorname{arc}_F(x, 1)$.
- 2 There is no “jump” at zero for arccot_F .

The arccot_F operation:

$$\begin{aligned}
\operatorname{arccot}_F : F &\rightarrow F \cup \{\mathbf{underflow}\} \\
\operatorname{arccot}_F(x) &= \operatorname{trans_result}_F(\operatorname{arccot}_F^\#(x), \operatorname{nearest}_F) \\
& \quad \text{if } x \in F \\
&= \operatorname{up}_F(-\pi/2) & \text{if } x = -\mathbf{0} \\
&= -\mathbf{0} & \text{if } x = -\infty \\
&= 0 & \text{if } x = +\infty \\
&= \operatorname{result_NaN}_F(x) & \text{otherwise}
\end{aligned}$$

NOTE 3 – $\operatorname{arccot}_F(\operatorname{neg}_F(x)) = \operatorname{neg}_F(\operatorname{arccot}_F(x))$.

5.3.9.13 Radian arc secant

The $\operatorname{arcsec}_F^*$ approximation helper function:

$$\operatorname{arcsec}_F^* : F \rightarrow \mathcal{R}$$

$\operatorname{arcsec}_F^*(x)$ returns a close approximation to $\operatorname{arcsec}(x)$ in \mathcal{R} , with maximum error $\operatorname{max_error_tan}_F$.

Further requirements on the $\operatorname{arcsec}_F^*$ approximation helper function are:

$$\begin{aligned}
\operatorname{arcsec}_F^*(2) &= \pi/3 \\
\operatorname{arcsec}_F^*(-2) &= 2 \cdot \pi/3 \\
\operatorname{arcsec}_F^*(-1) &= \pi \\
\operatorname{arcsec}_F^*(x) &\leq \pi/2 & \text{if } x > 0 \\
\operatorname{arcsec}_F^*(x) &\geq \pi/2 & \text{if } x < 0 \\
\operatorname{arcsec}_F^*(x) &= \pi/2 & \text{if } \operatorname{arcsec}_F^*(x) \neq \operatorname{arcsec}(x) \text{ and } |x| > 3 \cdot r_F/\operatorname{epsilon}_F
\end{aligned}$$

Range limitation:

$$\begin{aligned}
\operatorname{arcsec}_F^\#(x) &= \min\{\operatorname{arcsec}_F^*(x), \operatorname{down}_F(\pi/2)\} \\
& \quad \text{if } x \geq 1 \\
&= \max\{\operatorname{up}_F(\pi/2), \min\{\operatorname{arcsec}_F^*(x), \operatorname{down}_F(\pi)\}\} \\
& \quad \text{if } x \leq -1
\end{aligned}$$

The arcsec_F operation:

$$\begin{aligned}
\operatorname{arcsec}_F : F &\rightarrow F \cup \{\mathbf{invalid}\} \\
\operatorname{arcsec}_F(x) &= \operatorname{trans_result}_F(\operatorname{arcsec}_F^\#(x), \operatorname{nearest}_F) \\
& \quad \text{if } x \in F \text{ and } |x| \geq 1 \\
&= \operatorname{up}_F(\pi/2) & \text{if } x = -\infty \\
&= \operatorname{down}_F(\pi/2) & \text{if } x = +\infty \\
&= \operatorname{result_NaN}_F(x) & \text{otherwise}
\end{aligned}$$

5.3.9.14 Radian arc cosecant

The arccsc_F^* approximation helper function:

$$\text{arccsc}_F^* : F \rightarrow \mathcal{R}$$

$\text{arccsc}_F^*(x)$ returns a close approximation to $\text{arccsc}(x)$ in \mathcal{R} , with maximum error max_error_tan_F .

Further requirements on the arccsc_F^* approximation helper function are:

$$\begin{aligned} \text{arccsc}_F^*(2) &= \pi/6 \\ \text{arccsc}_F^*(1) &= \pi/2 \\ \text{arccsc}_F^*(-x) &= -\text{arccsc}_F^*(x) \end{aligned}$$

Range limitation:

$$\text{arccsc}_F^\#(x) = \max\{\text{up}_F(-\pi/2), \min\{\text{arccsc}_F^*(x), \text{down}_F(\pi/2)\}\}$$

The arccsc_F operation:

$$\begin{aligned} \text{arccsc}_F : F &\rightarrow F \cup \{\mathbf{underflow}, \mathbf{invalid}\} \\ \text{arccsc}_F(x) &= \text{trans_result}_F(\text{arccsc}_F^\#(x), \text{nearest}_F) \\ &\quad \text{if } x \in F \text{ and } |x| \geq 1 \\ &= -\mathbf{0} && \text{if } x = -\infty \\ &= 0 && \text{if } x = +\infty \\ &= \text{result_NaN}_F(x) && \text{otherwise} \end{aligned}$$

5.3.9.15 Radian angle from Cartesian co-ordinates

The arc_F^* approximation helper function:

$$\text{arc}_F^* : F \times F \rightarrow \mathcal{R}$$

$\text{arc}_F^*(x, y)$ returns a close approximation to $\text{arc}(x, y)$ in \mathcal{R} , with maximum error max_error_tan_F .

NOTE – The arc operations are often called `arctan2` (with the co-ordinate arguments swapped), or `arccot2`.

Further requirements on the arc_F^* approximation helper function are:

$$\begin{aligned} \text{arc}_F^*(x, 0) &= 0 && \text{if } x > 0 \\ \text{arc}_F^*(x, x) &= \pi/4 && \text{if } x > 0 \\ \text{arc}_F^*(0, y) &= \pi/2 && \text{if } y > 0 \\ \text{arc}_F^*(x, -x) &= 3 \cdot \pi/4 && \text{if } x < 0 \\ \text{arc}_F^*(x, 0) &= \pi && \text{if } x < 0 \\ \text{arc}_F^*(x, -y) &= -\text{arc}_F^*(x, y) && \text{if } y \neq 0 \text{ or } x > 0 \end{aligned}$$

Range limitation:

$$\text{arc}_F^\#(x, y) = \max\{\text{up}_F(-\pi), \min\{\text{arc}_F^*(x, y), \text{down}_F(\pi)\}\}$$

The arc_F operation:

$$\begin{aligned} \text{arc}_F : F \times F &\rightarrow F \cup \{\mathbf{underflow}\} \\ \text{arc}_F(x, y) &= \text{trans_result}_F(\text{arc}_F^\#(x, y), \text{nearest}_F) \\ &\quad \text{if } x, y \in F \text{ and } (x \neq 0 \text{ or } y \neq 0) \\ &= 0 && \text{if } x = 0 \text{ and } y = 0 \\ &= \text{down}_F(\pi) && \text{if } x = -\mathbf{0} \text{ and } y = 0 \\ &= \text{arc}_F(0, y) && \text{if } x = -\mathbf{0} \text{ and } y \in F \cup \{-\infty, +\infty\} \text{ and } y \neq 0 \\ &= \text{neg}_F(\text{arc}_F(x, 0)) && \text{if } y = -\mathbf{0} \text{ and } x \in F \cup \{-\infty, -\mathbf{0}, +\infty\} \\ &= 0 && \text{if } x = +\infty \text{ and } y \in F \text{ and } y \geq 0 \\ &= -\mathbf{0} && \text{if } x = +\infty \text{ and } y \in F \text{ and } y < 0 \end{aligned}$$

$= nearest_F(\pi/4)$	if $x = +\infty$ and $y = +\infty$
$= nearest_F(\pi/2)$	if $x \in F$ and $y = +\infty$
$= nearest_F(3 \cdot \pi/4)$	if $x = -\infty$ and $y = +\infty$
$= down_F(\pi)$	if $x = -\infty$ and $y \in F$ and $y \geq 0$
$= up_F(-\pi)$	if $x = -\infty$ and $y \in F$ and $y < 0$
$= nearest_F(-3 \cdot \pi/4)$	if $x = -\infty$ and $y = -\infty$
$= nearest_F(-\pi/2)$	if $x \in F$ and $y = -\infty$
$= nearest_F(-\pi/4)$	if $x = +\infty$ and $y = -\infty$
$= result_NaN2_F(x, y)$	otherwise

5.3.10 Operations for trigonometrics with given angular unit

There shall be one big-angle parameter for argument angular-unit trigonometric operations:

$$big_angle_u_F \in F$$

It should have the following default value:

$$big_angle_u_F = \lceil r_F^{\lceil p_F/2 \rceil} / 6 \rceil$$

A binding or implementation can include a method to change the value for this parameter. This method should only allow the value of this parameter to be set to a value greater than or equal to 1 and such that $ulp_F(big_angle_F) \leq 1/2000$.

NOTE 1 – In order to allow $ulp_F(big_angle_F) \leq 1/2000$, $p_F \geq 2 + \log_{r_F}(1000)$ should hold.

There shall be a derived parameter signifying the minimum allowed angular unit:

$$min_angular_unit_F = r_F \cdot fminN_F / epsilon_F$$

NOTE 2 – That is, $min_angular_unit_F = r_F^{(emin_F - 1 + p_F)}$

To make the specifications below a bit easier to express, let

$$G_F = \{x \in F \mid min_angular_unit_F \leq |x|\}.$$

Let $T = \{1, 2, 360, 400, 6400\}$. T consists of angle values for exactly one revolution for some common non-radian angular units: cycles, half-cycles, arc degrees, grades, and mils.

There shall be two parameterised maximum error parameters for argument angular-unit trigonometric operations:

$$max_error_sinu_F : F \rightarrow F \cup \{\mathbf{invalid}\}$$

$$max_error_tanu_F : F \rightarrow F \cup \{\mathbf{invalid}\}$$

For $u \in G_F$, the $max_error_sinu_F(u)$ parameter shall have a value in the interval $[max_error_sin_F, 2]$. The $max_error_sinu_F(u)$ parameter shall have the value of $max_error_sin_F$ if $|u| \in T$. For $u \in G_F$, the $max_error_tanu_F(u)$ parameter shall have a value in the interval $[max_error_tan_F, 4]$. The $max_error_tanu_F(u)$ parameter shall have the value of $max_error_tan_F$ if $|u| \in T$. The $max_error_sinu_F(u)$ and $max_error_tanu_F(u)$ parameters return **invalid** if $u \notin G_F$.

5.3.10.1 Argument angular-unit angle normalisation

The argument angular-unit normalisation computes exactly $rad(2 \cdot \pi \cdot x/u) \cdot u / (2 \cdot \pi)$, where x is the angular value, and u is the angular unit.

The $cycle_F$ operation:

$$cycle_F : F \times F \rightarrow F \cup \{-0, \mathbf{absolute_precision_underflow}, \mathbf{invalid}\}$$

$$\begin{aligned}
\mathit{cycle}_F(u, x) &= \mathit{remr}_F(x, u) && \text{if } u \in G_F \text{ and } (x = -\mathbf{0} \text{ or} \\
&&& (x \in F \text{ and } |x/u| \leq \mathit{big_angle_u}_F)) \\
&= \mathbf{absolute_precision_underflow}(\mathbf{qNaN}) && \text{if } u \in G_F \text{ and } x \in F \text{ and } |x/u| > \mathit{big_angle_u}_F \\
&= \mathit{result_NaN}_F(u, x) && \text{otherwise}
\end{aligned}$$

The $\mathit{axis_cycle}_F$ operation:

$$\begin{aligned}
\mathit{axis_cycle}_F : F \times F &\rightarrow ((F \times F) \times (F \cup \{-\mathbf{0}\})) \cup \{\mathbf{absolute_precision_underflow}, \mathbf{invalid}\} \\
\mathit{axis_cycle}_F(u, x) &= (\mathit{axis}(u, x), \mathit{result}_F(x - (\mathit{round}(x \cdot 4/u) \cdot u/4), \mathit{rnd}_F)) && \text{if } u \in G_F \text{ and } x \in F \text{ and } |x/u| \leq \mathit{big_angle_u}_F \text{ and} \\
&&& (x/u \geq 0 \text{ or } x - (\mathit{round}(x \cdot 4/u) \cdot u/4) \neq 0) \\
&= (\mathit{axis}(u, x), -\mathbf{0}) && \text{if } u \in G_F \text{ and } x \in F \text{ and } |x/u| \leq \mathit{big_angle_u}_F \text{ and} \\
&&& x/u < 0 \text{ and } x - (\mathit{round}(x \cdot 4/u) \cdot u/4) = 0 \text{ and} \\
&= ((1, 0), -\mathbf{0}) && \text{if } u \in G_F \text{ and } x = -\mathbf{0} \\
&= \mathbf{absolute_precision_underflow}((\mathbf{qNaN}, \mathbf{qNaN}), \mathbf{qNaN}) && \text{if } u \in G_F \text{ and } x \in F \text{ and } |x/u| > \mathit{big_angle_u}_F \\
&= ((\mathbf{qNaN}, \mathbf{qNaN}), \mathbf{qNaN}) && \text{if } x \text{ is a quiet NaN and } u \text{ is not a signalling NaN} \\
&= ((\mathbf{qNaN}, \mathbf{qNaN}), \mathbf{qNaN}) && \text{if } u \text{ is a quiet NaN and } x \text{ is not a signalling NaN} \\
&= \mathbf{invalid}((\mathbf{qNaN}, \mathbf{qNaN}), \mathbf{qNaN}) && \text{otherwise}
\end{aligned}$$

where

$$\begin{aligned}
\mathit{axis}(u, x) &= (1, 0) && \text{if } \mathit{round}(x \cdot 4/u) = 4 \cdot n \\
&= (0, 1) && \text{if } \mathit{round}(x \cdot 4/u) = 4 \cdot n + 1 \\
&= (-1, 0) && \text{if } \mathit{round}(x \cdot 4/u) = 4 \cdot n + 2 \\
&= (0, -1) && \text{if } \mathit{round}(x \cdot 4/u) = 4 \cdot n + 3
\end{aligned}$$

for some $n \in \mathcal{Z}$.

NOTES

- 1 $\mathit{axis_cycle}_F(u, x)$ is exact when $\mathit{div}_F(u, 4) = u/4$.
- 2 cycle_F is an exact operation.
- 3 $\mathit{cycle}_F(u, x)$ is $-\mathbf{0}$ or has a result in the interval $[-|u/2|, |u/2|]$ if there is no notification.
- 4 A zero resulting angle is negative if the original angle value is negative.
- 5 The cycle_F operation is used also in the specifications of the unit argument trigonometric operations. This does *not* imply that the implementation has to use the cycle operation, when implementing the operations. Just that the *results* (including notifications) must be *as if* it did.

5.3.10.2 Argument angular-unit sine

The sinu_F^* approximation helper function:

$$\mathit{sinu}_F^* : F \times \mathcal{R} \rightarrow \mathcal{R}$$

$\mathit{sinu}_F^*(u, x)$ returns a close approximation to $\sin(x \cdot 2 \cdot \pi/u)$ in \mathcal{R} if $u \neq 0$, with maximum error $\mathit{max_error_sinu}_F(u)$.

Further requirements on the sinu_F^* approximation helper function:

$$\begin{aligned}
\sinu_F^*(u, n \cdot u + x) &= \sinu_F^*(u, x) && \text{if } n \in \mathcal{Z} \text{ and } u \in F \text{ and } u \neq 0 \\
\sinu_F^*(u, u/12) &= 1/2 && \text{if } u \in F \text{ and } u \neq 0 \\
\sinu_F^*(u, u/4) &= 1 && \text{if } u \in F \text{ and } u \neq 0 \\
\sinu_F^*(u, 5 \cdot u/12) &= 1/2 && \text{if } u \in F \text{ and } u \neq 0 \\
\sinu_F^*(u, -x) &= -\sinu_F^*(u, x) && \text{if } u \in F \text{ and } u \neq 0 \\
\sinu_F^*(-u, x) &= -\sinu_F^*(u, x) && \text{if } u \in F \text{ and } u \neq 0
\end{aligned}$$

NOTE – $\sinu_F^*(u, x) \approx x \cdot 2 \cdot \pi / u$ if $|x \cdot 2 \cdot \pi / u| < fminN_F$.

The \sinu_F operation:

$$\begin{aligned}
\sinu_F : F \times F &\rightarrow F \cup \{-\mathbf{0}, \text{underflow}, \text{invalid}, \text{absolute_precision_underflow}\} \\
\sinu_F(u, x) &= \text{trans_result}_F(\sinu_F^*(u, x), \text{nearest}_F) \\
&= \text{div}_F(0, u) && \text{if } \text{cycle}_F(u, x) \in F \text{ and } \text{cycle}_F(u, x) \notin \{-u/2, 0, u/2\} \\
&= \text{div}_F(-\mathbf{0}, u) && \text{if } \text{cycle}_F(u, x) \in \{0, u/2\} \\
&= \text{cycle}_F(u, x) && \text{if } \text{cycle}_F(u, x) \in \{-u/2, -\mathbf{0}\} \\
& && \text{otherwise}
\end{aligned}$$

5.3.10.3 Argument angular-unit cosine

The \cosu_F^* approximation helper function:

$$\cosu_F^* : F \times \mathcal{R} \rightarrow \mathcal{R}$$

$\cosu_F^*(u, x)$ returns a close approximation to $\cos(x \cdot 2 \cdot \pi / u)$ in \mathcal{R} if $u \neq 0$, with maximum error $\text{max_error_sinu}_F(u)$.

Further requirements on the \cosu_F^* approximation helper function:

$$\begin{aligned}
\cosu_F^*(u, n \cdot u + x) &= \cosu_F^*(u, x) && \text{if } n \in \mathcal{Z} \text{ and } u \in F \text{ and } u \neq 0 \\
\cosu_F^*(u, 0) &= 1 && \text{if } u \in F \text{ and } u \neq 0 \\
\cosu_F^*(u, u/6) &= 1/2 && \text{if } u \in F \text{ and } u \neq 0 \\
\cosu_F^*(u, u/3) &= -1/2 && \text{if } u \in F \text{ and } u \neq 0 \\
\cosu_F^*(u, u/2) &= -1 && \text{if } u \in F \text{ and } u \neq 0 \\
\cosu_F^*(u, -x) &= \cosu_F^*(u, x) && \text{if } u \in F \text{ and } u \neq 0 \\
\cosu_F^*(-u, x) &= \cosu_F^*(u, x) && \text{if } u \in F \text{ and } u \neq 0
\end{aligned}$$

NOTE – $\cosu_F^*(u, x) = 1$ should hold if $|x \cdot 2 \cdot \pi / u| < \sqrt{\text{epsilon}_F / r_F}$

The \cosu_F operation:

$$\begin{aligned}
\cosu_F : F \times F &\rightarrow F \cup \{\text{underflow}, \text{invalid}, \text{absolute_precision_underflow}\} \\
\cosu_F(u, x) &= \text{trans_result}_F(\cosu_F^*(u, x), \text{nearest}_F) \\
&= 1 && \text{if } \text{cycle}_F(u, x) \in F \\
&= \text{cycle}_F(u, x) && \text{if } \text{cycle}_F(u, x) = -\mathbf{0} \\
& && \text{otherwise}
\end{aligned}$$

5.3.10.4 Argument angular-unit tangent

The \tanu_F^* approximation helper function:

$$\tanu_F^* : F \times \mathcal{R} \rightarrow \mathcal{R}$$

$\tanu_F^*(u, x)$ returns a close approximation to $\tan(x \cdot 2 \cdot \pi / u)$ in \mathcal{R} if $u \neq 0$, with maximum error $\text{max_error_tanu}_F(u)$.

Further requirements on the \tanu_F^* approximation helper function:

$$\begin{array}{ll}
\tanu_F^*(u, n \cdot u + x) = \tanu_F^*(u, x) & \text{if } n \in \mathcal{Z} \text{ and } u \in F \text{ and } u \neq 0 \\
\tanu_F^*(u, u/8) = 1 & \text{if } u \in F \text{ and } u \neq 0 \\
\tanu_F^*(u, 3 \cdot u/8) = -1 & \text{if } u \in F \text{ and } u \neq 0 \\
\tanu_F^*(u, -x) = -\tanu_F^*(u, x) & \text{if } u \in F \text{ and } u \neq 0 \\
\tanu_F^*(-u, x) = -\tanu_F^*(u, x) & \text{if } u \in F \text{ and } u \neq 0
\end{array}$$

NOTE 1 – $\tanu_F^*(u, x) \approx x \cdot 2 \cdot \pi / u$ if $|x \cdot 2 \cdot \pi / u| < fminN_F$.

The \tanu_F operation:

$$\begin{array}{l}
\tanu_F : F \times F \rightarrow F \cup \{-\mathbf{0}, \text{pole}, \text{overflow}, \text{underflow}, \text{invalid}, \\
\quad \text{absolute_precision_underflow}\} \\
\tanu_F(u, x) = \text{trans_result}_F(\tanu_F^*(u, x), \text{nearest}_F) \\
\quad \text{if } \text{cycle}_F(u, x) \in F \text{ and} \\
\quad \quad \text{cycle}_F(u, x) \notin \{-u/2, -u/4, 0, u/4, u/2\} \\
= \text{div}_F(0, u) \quad \text{if } \text{cycle}_F(u, x) \in \{-u/2, 0\} \\
= \text{div}_F(-\mathbf{0}, u) \quad \text{if } \text{cycle}_F(u, x) \in \{-\mathbf{0}, u/2\} \\
= \text{pole}(+\infty) \quad \text{if } \text{cycle}_F(u, x) = u/4 \\
= \text{pole}(-\infty) \quad \text{if } \text{cycle}_F(u, x) = -u/4 \\
= \text{cycle}_F(u, x) \quad \text{otherwise}
\end{array}$$

NOTE 2 – The **pole** notification can arise for $\tanu_F(u, x)$ only when $u/4$ is in F .

5.3.10.5 Argument angular-unit cotangent

The \cotu_F^* approximation helper function:

$$\cotu_F^* : F \times \mathcal{R} \rightarrow \mathcal{R}$$

$\cotu_F^*(u, x)$ returns a close approximation to $\cot(x \cdot 2 \cdot \pi / u)$ in \mathcal{R} if $u \neq 0$, with maximum error $\text{max_error_tanu}_F(u)$.

Further requirements on the \cotu_F^* approximation helper function:

$$\begin{array}{ll}
\cotu_F^*(u, n \cdot u + x) = \cotu_F^*(u, x) & \text{if } n \in \mathcal{Z} \text{ and } u \in F \text{ and } u \neq 0 \\
\cotu_F^*(u, u/8) = 1 & \text{if } u \in F \text{ and } u \neq 0 \\
\cotu_F^*(u, 3 \cdot u/8) = -1 & \text{if } u \in F \text{ and } u \neq 0 \\
\cotu_F^*(u, -x) = -\cotu_F^*(u, x) & \text{if } u \in F \text{ and } u \neq 0 \\
\cotu_F^*(-u, x) = -\cotu_F^*(u, x) & \text{if } u \in F \text{ and } u \neq 0
\end{array}$$

The \cotu_F operation:

$$\begin{array}{l}
\cotu_F : F \times F \rightarrow F \cup \{-\mathbf{0}, \text{pole}, \text{overflow}, \text{underflow}, \text{invalid}, \\
\quad \text{absolute_precision_underflow}\} \\
\cotu_F(u, x) = \text{trans_result}_F(\cotu_F^*(u, x), \text{nearest}_F) \\
\quad \text{if } \text{cycle}_F(u, x) \in F \text{ and} \\
\quad \quad \text{cycle}_F(u, x) \notin \{-u/2, -u/4, 0, u/2\} \\
= -\mathbf{0} \quad \text{if } \text{cycle}_F(u, x) = -u/4 \\
= \text{div}_F(u, \tanu_F(u, x)) \quad \text{if } \text{cycle}_F(u, x) \in \{-u/2, -\mathbf{0}, 0, u/2\} \\
= \text{cycle}_F(u, x) \quad \text{otherwise}
\end{array}$$

5.3.10.6 Argument angular-unit secant

The \secu_F^* approximation helper function:

$$\secu_F^* : F \times \mathcal{R} \rightarrow \mathcal{R}$$

$secu_F^*(u, x)$ returns a close approximation to $\sec(x \cdot 2 \cdot \pi/u)$ in \mathcal{R} if $u \neq 0$, with maximum error $max_error_tanu_F(u)$.

Further requirements on the $secu_F^*$ approximation helper function:

$$\begin{array}{ll}
secu_F^*(u, n \cdot u + x) = secu_F^*(u, x) & \text{if } n \in \mathcal{Z} \text{ and } u \in F \text{ and } u \neq 0 \\
secu_F^*(u, 0) = 1 & \text{if } u \in F \text{ and } u \neq 0 \\
secu_F^*(u, u/6) = 2 & \text{if } u \in F \text{ and } u \neq 0 \\
secu_F^*(u, u/3) = -2 & \text{if } u \in F \text{ and } u \neq 0 \\
secu_F^*(u, u/2) = -1 & \text{if } u \in F \text{ and } u \neq 0 \\
secu_F^*(u, -x) = secu_F^*(u, x) & \text{if } u \in F \text{ and } u \neq 0 \\
secu_F^*(-u, x) = secu_F^*(u, x) & \text{if } u \in F \text{ and } u \neq 0 \\
secu_F^*(u, x) = 1 & \text{if } |x \cdot 2 \cdot \pi/u| < 0.5 \cdot \sqrt{\epsilonpsilon_F}
\end{array}$$

The $secu_F$ operation:

$$\begin{array}{ll}
secu_F : F \times F \rightarrow F \cup \{\mathbf{pole}, \mathbf{overflow}, \mathbf{invalid}, \mathbf{absolute_precision_underflow}\} \\
secu_F(u, x) & = trans_result_F(secu_F^*(u, x), nearest_F) \\
& \text{if } cycle_F(u, x) \in F \text{ and } cycle_F(u, x) \notin \{-u/4, u/4\} \\
& = div_F(1, cosu_F(u, x)) & \text{if } cycle_F(u, x) \in \{-u/4, -\mathbf{0}, u/4\} \\
& = cycle_F(u, x) & \text{otherwise}
\end{array}$$

5.3.10.7 Argument angular-unit cosecant

The $cscu_F^*$ approximation helper function:

$$cscu_F^* : F \times \mathcal{R} \rightarrow \mathcal{R}$$

$cscu_F^*(u, x)$ returns a close approximation to $\csc(x \cdot 2 \cdot \pi/u)$ in \mathcal{R} if $u \neq 0$, with maximum error $max_error_tanu_F(u)$.

Further requirements on the $cscu_F^*$ approximation helper function:

$$\begin{array}{ll}
cscu_F^*(u, n \cdot u + x) = cscu_F^*(u, x) & \text{if } n \in \mathcal{Z} \text{ and } u \in F \text{ and } u \neq 0 \\
cscu_F^*(u, u/12) = 2 & \text{if } u \in F \text{ and } u \neq 0 \\
cscu_F^*(u, u/4) = 1 & \text{if } u \in F \text{ and } u \neq 0 \\
cscu_F^*(u, 5 \cdot u/12) = 2 & \text{if } u \in F \text{ and } u \neq 0 \\
cscu_F^*(u, -x) = -cscu_F^*(u, x) & \text{if } u \in F \text{ and } u \neq 0 \\
cscu_F^*(-u, x) = -cscu_F^*(u, x) & \text{if } u \in F \text{ and } u \neq 0
\end{array}$$

The $cscu_F$ operation:

$$\begin{array}{ll}
cscu_F : F \times F \rightarrow F \cup \{\mathbf{pole}, \mathbf{overflow}, \mathbf{invalid}, \mathbf{absolute_precision_underflow}\} \\
cscu_F(u, x) & = trans_result_F(cscu_F^*(u, x), nearest_F) \\
& \text{if } cycle_F(u, x) \in F \text{ and } cycle_F(u, x) \notin \{-u/2, 0, u/2\} \\
& = div_F(1, sinu_F(u, x)) & \text{if } cycle_F(u, x) \in \{-u/2, -\mathbf{0}, 0, u/2\} \\
& = cycle_F(u, x) & \text{otherwise}
\end{array}$$

5.3.10.8 Argument angular-unit cosine with sine

$$\begin{array}{ll}
cossinu_F : F \times F \rightarrow (F \times (F \cup \{-\mathbf{0}\})) \cup \{\mathbf{underflow}, \mathbf{invalid}, \mathbf{absolute_precision_underflow}\} \\
cossinu_F(u, x) & = (cosu_F(u, x), sinu_F(u, x))
\end{array}$$

5.3.10.9 Argument angular-unit arc sine

The \arcsinu_F^* approximation helper function:

$$\arcsinu_F^* : F \times F \rightarrow \mathcal{R}$$

$\arcsinu_F^*(u, x)$ returns a close approximation to $\arcsin(x) \cdot u / (2 \cdot \pi)$ in \mathcal{R} , with maximum error $\max_error_sinu_F(u)$.

Further requirements on the \arcsinu_F^* approximation helper function:

$$\arcsinu_F^*(u, 1/2) = u/12$$

$$\arcsinu_F^*(u, 1) = u/4$$

$$\arcsinu_F^*(u, -x) = -\arcsinu_F^*(u, x)$$

$$\arcsinu_F^*(-u, x) = -\arcsinu_F^*(u, x)$$

NOTE – $\arcsinu_F^*(u, x) \approx u / (2 \cdot \pi)$ if $|x| < \mathit{fmin}N_F$.

Range limitation:

$$\arcsinu_F^\#(u, x) = \max\{\mathit{up}_F(-|u/4|), \min\{\arcsinu_F^*(u, x), \mathit{down}_F(|u/4|)\}\}$$

The \arcsinu_F operation:

$$\arcsinu_F : F \times F \rightarrow F \cup \{-\mathbf{0}, \mathbf{underflow}, \mathbf{invalid}\}$$

$$\begin{aligned} \arcsinu_F(u, x) &= \mathit{trans_result}_F(\arcsinu_F^\#(u, x), \mathit{nearest}_F) && \text{if } u \in G_F \text{ and } x \in F \text{ and } |x| \leq 1 \text{ and } x \neq 0 \\ &= \mathit{mul}_F(u, x) && \text{if } u \in G_F \text{ and } x \in \{-\mathbf{0}, 0\} \\ &= \mathit{result_NaN}_F(u, x) && \text{otherwise} \end{aligned}$$

5.3.10.10 Argument angular-unit arc cosine

The \arccosu_F^* approximation helper function:

$$\arccosu_F^* : F \times F \rightarrow \mathcal{R}$$

$\arccosu_F^*(u, x)$ returns a close approximation to $\arccos(x) \cdot u / (2 \cdot \pi)$ in \mathcal{R} , with maximum error $\max_error_sinu_F(u)$.

Further requirements on the \arccosu_F^* approximation helper function:

$$\arccosu_F^*(u, 1/2) = u/6$$

$$\arccosu_F^*(u, 0) = u/4$$

$$\arccosu_F^*(u, -1/2) = u/3$$

$$\arccosu_F^*(u, -1) = u/2$$

$$\arccosu_F^*(-u, x) = -\arccosu_F^*(u, x)$$

Range limitation:

$$\arccosu_F^\#(u, x) = \max\{\mathit{up}_F(-|u/2|), \min\{\arccosu_F^*(u, x), \mathit{down}_F(|u/2|)\}\}$$

The \arccosu_F operation:

$$\arccosu_F : F \times F \rightarrow F \cup \{\mathbf{underflow}, \mathbf{invalid}\}$$

$$\begin{aligned} \arccosu_F(u, x) &= \mathit{trans_result}_F(\arccosu_F^\#(u, x), \mathit{nearest}_F) && \text{if } u \in G_F \text{ and } x \in F \text{ and } |x| \leq 1 \\ &= \mathit{nearest}_F(u/4) && \text{if } u \in G_F \text{ and } x = -\mathbf{0} \\ &= \mathit{result_NaN}_F(u, x) && \text{otherwise} \end{aligned}$$

5.3.10.11 Argument angular-unit arc tangent

The $\arctan u_F^*$ approximation helper function:

$$\arctan u_F^* : F \times F \rightarrow \mathcal{R}$$

$\arctan u_F^*(u, x)$ returns a close approximation to $\arctan(x) \cdot u / (2 \cdot \pi)$ in \mathcal{R} , with maximum error $\max_error_tan u_F(u)$.

Further requirements on the $\arctan u_F^*$ approximation helper function:

$$\begin{aligned} \arctan u_F^*(u, 1) &= u/8 \\ \arctan u_F^*(u, x) &= u/4 && \text{if } \arctan u_F^*(u, x) \neq \arctan(x) \cdot u / (2 \cdot \pi) \text{ and} \\ &&& x > 3 \cdot r_F / \epsilon_F \end{aligned}$$

$$\arctan u_F^*(u, -x) = -\arctan u_F^*(u, x)$$

$$\arctan u_F^*(-u, x) = -\arctan u_F^*(u, x)$$

NOTE 1 – $\arctan u_F^*(u, x) \approx u / (2 \cdot \pi)$ if $|x| < \text{fmin} N_F$

Range limitation:

$$\arctan u_F^\#(u, x) = \max\{up_F(-|u/4|), \min\{\arctan u_F^*(u, x), down_F(|u/4|)\}\}$$

The $\arctan u_F$ operation:

$$\arctan u_F : F \times F \rightarrow F \cup \{-\mathbf{0}, \text{invalid}, \text{underflow}\}$$

$$\begin{aligned} \arctan u_F(u, x) &= \text{trans_result}_F(\arctan u_F^\#(u, x), \text{nearest}_F) \\ &&& \text{if } u \in G_F \text{ and } x \in F \text{ and } x \neq 0 \\ &= \text{mul}_F(x, u) && \text{if } u \in G_F \text{ and } x \in \{-\mathbf{0}, 0\} \\ &= up_F(-u/4) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } u > 0 \\ &= down_F(u/4) && \text{if } u \in G_F \text{ and } x = +\infty \text{ and } u > 0 \\ &= down_F(-u/4) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } u < 0 \\ &= up_F(u/4) && \text{if } u \in G_F \text{ and } x = +\infty \text{ and } u < 0 \\ &= \text{result_NaN}_F(u, x) && \text{otherwise} \end{aligned}$$

NOTE 2 – $\arctan u_F(u, x) \approx \text{arcu}_F(u, 1, x)$.

5.3.10.12 Argument angular-unit arc cotangent

This clause specifies two inverse cotangent operations. One approximating the continuous (but not sign symmetric) arccot , the other approximating the sign symmetric (but discontinuous at 0) arcctg .

The $\text{arccot} u_F^*$ approximation helper function:

$$\text{arccot} u_F^* : F \times F \rightarrow \mathcal{R}$$

$\text{arccot} u_F^*(u, x)$ returns a close approximation to $\text{arccot}(x) \cdot u / (2 \cdot \pi)$ in \mathcal{R} , with maximum error $\max_error_tan u_F(u)$.

The $\text{arcctg} u_F^*$ approximation helper function:

$$\text{arcctg} u_F^* : F \times F \rightarrow \mathcal{R}$$

$\text{arcctg} u_F^*(u, x)$ returns a close approximation to $\text{arcctg}(x) \cdot u / (2 \cdot \pi)$ in \mathcal{R} , with maximum error $\max_error_tan u_F(u)$.

Further requirements on the $\text{arccot} u_F^*$ and $\text{arcctg} u_F^*$ approximation helper functions:

$$\begin{aligned}
\operatorname{arccotu}_F^*(u, 1) &= u/8 \\
\operatorname{arccotu}_F^*(u, 0) &= u/4 \\
\operatorname{arccotu}_F^*(u, -1) &= 3 \cdot u/8 \\
\operatorname{arccotu}_F^*(u, x) &\leq u/2 && \text{if } u > 0 \\
\operatorname{arccotu}_F^*(u, x) &\geq u/2 && \text{if } u < 0 \\
\operatorname{arccotu}_F^*(u, x) &= u/2 && \text{if } \operatorname{arccotu}_F^*(u, x) \neq \operatorname{arccot}(x) \cdot u/(2 \cdot \pi) \text{ and} \\
&&& x < -3 \cdot r_F/\epsilon_F \\
\operatorname{arccotu}_F^*(-u, x) &= -\operatorname{arccotu}_F^*(u, x) \\
\operatorname{arctgu}_F^*(u, x) &= \operatorname{arccotu}_F^*(u, x) && \text{if } x \geq 0 \\
\operatorname{arctgu}_F^*(u, -x) &= -\operatorname{arctgu}_F^*(u, x)
\end{aligned}$$

Range limitation:

$$\begin{aligned}
\operatorname{arccotu}_F^\#(u, x) &= \max\{up_F(-|u/2|), \min\{\operatorname{arccotu}_F^*(u, x), \operatorname{down}_F(|u/2|)\}\} \\
\operatorname{arctgu}_F^\#(u, x) &= \max\{up_F(-|u/4|), \min\{\operatorname{arctgu}_F^*(u, x), \operatorname{down}_F(|u/4|)\}\}
\end{aligned}$$

The $\operatorname{arccotu}_F$ operation:

$$\begin{aligned}
\operatorname{arccotu}_F : F \times F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{underflow}\} \\
\operatorname{arccotu}_F(u, x) &= \operatorname{trans_result}_F(\operatorname{arccotu}_F^\#(u, x), \operatorname{nearest}_F) \\
&= \operatorname{nearest}_F(u/4) && \text{if } u \in G_F \text{ and } x \in F \\
&= \operatorname{down}_F(u/2) && \text{if } u \in G_F \text{ and } x = -\mathbf{0} \\
&= up_F(u/2) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } u > 0 \\
&= \operatorname{div}_F(u, x) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } u < 0 \\
&= \operatorname{result_NaN}_F(u, x) && \text{if } u \in G_F \text{ and } x = +\infty \\
&= \operatorname{result_NaN}_F(u, x) && \text{otherwise}
\end{aligned}$$

NOTE – $\operatorname{arccotu}_F(u, x) \approx \operatorname{arcu}_F(u, x, 1)$.

The arctgu_F operation:

$$\begin{aligned}
\operatorname{arctgu}_F : F \times F &\rightarrow F \cup \{\mathbf{invalid}, \mathbf{underflow}\} \\
\operatorname{arctgu}_F(u, x) &= \operatorname{trans_result}_F(\operatorname{arctgu}_F^\#(u, x), \operatorname{nearest}_F) \\
&= \operatorname{neg}_F(\operatorname{arctgu}_F(u, 0)) && \text{if } u \in G_F \text{ and } x \in F \\
&= \operatorname{div}_F(u, x) && \text{if } u \in G_F \text{ and } x = -\mathbf{0} \\
&= \operatorname{result_NaN}_F(u, x) && \text{if } u \in G_F \text{ and } x \in \{-\infty, +\infty\} \\
&= \operatorname{result_NaN}_F(u, x) && \text{otherwise}
\end{aligned}$$

5.3.10.13 Argument angular-unit arc secant

The $\operatorname{arcsecu}_F^*$ approximation helper function:

$$\operatorname{arcsecu}_F^* : F \times F \rightarrow \mathcal{R}$$

$\operatorname{arcsecu}_F^*(u, x)$ returns a close approximation to $\operatorname{arcsec}(x) \cdot u/(2 \cdot \pi)$ in \mathcal{R} , with maximum error $\operatorname{max_error_tanu}_F(u)$.

Further requirements on the $\operatorname{arcsecu}_F^*$ approximation helper function:

$$\begin{aligned}
\operatorname{arcsecu}_F^*(u, 2) &= u/6 \\
\operatorname{arcsecu}_F^*(u, -2) &= u/3 \\
\operatorname{arcsecu}_F^*(u, -1) &= u/2 \\
\operatorname{arcsecu}_F^*(u, x) &\leq u/4 && \text{if } x > 0 \text{ and } u > 0 \\
\operatorname{arcsecu}_F^*(u, x) &\geq u/4 && \text{if } x < 0 \text{ and } u > 0 \\
\operatorname{arcsecu}_F^*(u, x) &= u/4 && \text{if } \operatorname{arcsecu}_F^*(u, x) \neq \operatorname{arcsec}(x) \cdot u/(2 \cdot \pi) \text{ and} \\
&&& |x| > 3 \cdot r_F/\epsilon_F \\
\operatorname{arcsecu}_F^*(-u, x) &= -\operatorname{arcsecu}_F^*(u, x)
\end{aligned}$$

Range limitation:

$$\begin{aligned} \operatorname{arcsecu}_F^\#(u, x) &= \max\{up_F(-|u/4|), \min\{\operatorname{arcsecu}_F^*(u, x), \operatorname{down}_F(|u/4|)\}\} \\ &\quad \text{if } x \geq 1 \\ &= \max\{up_F(u/4), \min\{\operatorname{arcsecu}_F^*(u, x), \operatorname{down}_F(u/2)\}\} \\ &\quad \text{if } x \leq -1 \text{ and } u > 0 \\ &= \max\{up_F(u/2), \min\{\operatorname{arcsecu}_F^*(u, x), \operatorname{down}_F(u/4)\}\} \\ &\quad \text{if } x \leq -1 \text{ and } u < 0 \end{aligned}$$

The $\operatorname{arcsecu}_F$ operation:

$$\begin{aligned} \operatorname{arcsecu}_F &: F \times F \rightarrow F \cup \{\mathbf{underflow}, \mathbf{invalid}\} \\ \operatorname{arcsecu}_F(u, x) &= \operatorname{trans_result}_F(\operatorname{arcsecu}_F^\#(u, x), \operatorname{nearest}_F) \\ &\quad \text{if } u \in G_F \text{ and } x \in F \text{ and } (x \leq -1 \text{ or } x \geq 1) \\ &= \operatorname{down}_F(u/4) \quad \text{if } u \in G_F \text{ and } x = -\infty \text{ and } u > 0 \\ &= up_F(u/4) \quad \text{if } u \in G_F \text{ and } x = +\infty \text{ and } u > 0 \\ &= up_F(u/4) \quad \text{if } u \in G_F \text{ and } x = -\infty \text{ and } u < 0 \\ &= \operatorname{down}_F(u/4) \quad \text{if } u \in G_F \text{ and } x = +\infty \text{ and } u < 0 \\ &= \operatorname{result_NaN}_F(u, x) \quad \text{otherwise} \end{aligned}$$

5.3.10.14 Argument angular-unit arc cosecant

The $\operatorname{arccscu}_F^*$ approximation helper function:

$$\operatorname{arccscu}_F^* : F \times F \rightarrow \mathcal{R}$$

$\operatorname{arccscu}_F^*(u, x)$ returns a close approximation to $\operatorname{arccsc}(x) \cdot u / (2 \cdot \pi)$ in \mathcal{R} , with maximum error $\operatorname{max_error_tan}_F(u)$.

Further requirements on the $\operatorname{arccscu}_F^*$ approximation helper function:

$$\begin{aligned} \operatorname{arccscu}_F^*(u, 2) &= u/12 \\ \operatorname{arccscu}_F^*(u, 1) &= u/4 \\ \operatorname{arccscu}_F^*(u, -x) &= -\operatorname{arccscu}_F^*(u, x) \\ \operatorname{arccscu}_F^*(-u, x) &= -\operatorname{arccscu}_F^*(u, x) \end{aligned}$$

Range limitation:

$$\operatorname{arccscu}_F^\#(u, x) = \max\{up_F(-|u/4|), \min\{\operatorname{arccscu}_F^*(u, x), \operatorname{down}_F(|u/4|)\}\}$$

The $\operatorname{arccscu}_F$ operation:

$$\begin{aligned} \operatorname{arccscu}_F &: F \times F \rightarrow F \cup \{\mathbf{underflow}, \mathbf{invalid}\} \\ \operatorname{arccscu}_F(u, x) &= \operatorname{trans_result}_F(\operatorname{arccscu}_F^\#(u, x), \operatorname{nearest}_F) \\ &\quad \text{if } u \in G_F \text{ and } x \in F \text{ and } (x \geq 1 \text{ or } x \leq -1) \\ &= \operatorname{mul}_F(-u, 0) \quad \text{if } u \in G_F \text{ and } x = -\infty \\ &= \operatorname{mul}_F(u, 0) \quad \text{if } u \in G_F \text{ and } x = +\infty \\ &= \operatorname{result_NaN}_F(u, x) \quad \text{otherwise} \end{aligned}$$

5.3.10.15 Argument angular-unit angle from Cartesian co-ordinates

The arcu_F^* approximation helper function:

$$\operatorname{arcu}_F^* : F \times F \times F \rightarrow \mathcal{R}$$

$\operatorname{arcu}_F^*(u, x, y)$ returns a close approximation to $\operatorname{arc}(x, y) \cdot u / (2 \cdot \pi)$ in \mathcal{R} , with maximum error $\operatorname{max_error_tan}_F(u)$.

Further requirements on the arcu_F^* approximation helper function:

$$\begin{aligned}
arcu_F^*(u, x, x) &= u/8 && \text{if } x > 0 \\
arcu_F^*(u, 0, y) &= u/4 && \text{if } y > 0 \\
arcu_F^*(u, x, -x) &= 3 \cdot u/8 && \text{if } x < 0 \\
arcu_F^*(u, x, 0) &= u/2 && \text{if } x < 0 \\
arcu_F^*(u, x, -y) &= -arcu_F^*(u, x, y) && \text{if } y \neq 0 \text{ or } x > 0 \\
arcu_F^*(-u, x, y) &= -arcu_F^*(u, x, y)
\end{aligned}$$

Range limitation:

$$arcu_F^\#(u, x, y) = \max\{up_F(-|u/2|), \min\{arcu_F^*(u, x, y), down_F(|u/2|)\}\}$$

The $arcu_F$ operation:

$$arcu_F : F \times F \times F \rightarrow F \cup \{-\mathbf{0}, \text{underflow}, \text{invalid}\}$$

$$\begin{aligned}
arcu_F(u, x, y) &= trans_result_F(arcu_F^\#(u, x, y), nearest_F) \\
&= mul_F(u, 0) && \text{if } u \in G_F \text{ and } x, y \in F \text{ and } (x < 0 \text{ or } y \neq 0) \\
&= 0 && \text{if } u \in G_F \text{ and } x \in F \text{ and } x > 0 \text{ and } y = 0 \\
&= down_F(u/2) && \text{if } u \in G_F \text{ and } x = 0 \text{ and } y = 0 \\
&= up_F(u/2) && \text{if } u \in G_F \text{ and } x = -\mathbf{0} \text{ and } y = 0 \text{ and } u > 0 \\
&= arcu_F(u, 0, y) && \text{if } u \in G_F \text{ and } x = -\mathbf{0} \text{ and } y \in F \cup \{-\infty, +\infty\} \text{ and } \\
&&& \quad y \neq 0 \\
&= neg_F(arcu_F(u, x, 0)) && \text{if } u \in G_F \text{ and } y = -\mathbf{0} \text{ and } x \in F \cup \{-\infty, -\mathbf{0}, +\infty\} \\
&= mul_F(0, u) && \text{if } u \in G_F \text{ and } x = +\infty \text{ and } y \in F \text{ and } y \geq 0 \\
&= mul_F(0, -u) && \text{if } u \in G_F \text{ and } x = +\infty \text{ and } y \in F \text{ and } y < 0 \\
&= nearest_F(u/8) && \text{if } u \in G_F \text{ and } x = +\infty \text{ and } y = +\infty \\
&= nearest_F(u/4) && \text{if } u \in G_F \text{ and } x \in F \text{ and } y = +\infty \\
&= nearest_F(3 \cdot u/8) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } y = +\infty \\
&= down_F(u/2) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } y \in F \text{ and } \\
&&& \quad y \geq 0 \text{ and } u > 0 \\
&= up_F(-u/2) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } y \in F \text{ and } \\
&&& \quad y < 0 \text{ and } u > 0 \\
&= up_F(u/2) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } y \in F \text{ and } \\
&&& \quad y > 0 \text{ and } u < 0 \\
&= down_F(-u/2) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } y \in F \text{ and } \\
&&& \quad y \leq 0 \text{ and } u < 0 \\
&= nearest_F(-3 \cdot u/8) && \text{if } u \in G_F \text{ and } x = -\infty \text{ and } y = -\infty \\
&= nearest_F(-u/4) && \text{if } u \in G_F \text{ and } x \in F \text{ and } y = -\infty \\
&= nearest_F(-u/8) && \text{if } u \in G_F \text{ and } x = +\infty \text{ and } y = -\infty \\
&= result_NaN3_F(u, x, y) && \text{otherwise}
\end{aligned}$$

5.3.11 Operations for angular-unit conversions

5.3.11.1 Converting radian angle to argument angular-unit angle

Define the mathematical function:

$$\begin{aligned}
rad_to_cycle &: \mathcal{R} \times \mathcal{R} \rightarrow \mathcal{R} \\
rad_to_cycle(x, v) &= \arccos(\cos(x)) \cdot v / (2 \cdot \pi) \\
&&& \text{if } \sin(x) \geq 0 \text{ and } v \neq 0 \\
&= -\arccos(\cos(x)) \cdot v / (2 \cdot \pi)
\end{aligned}$$

if $\sin(x) < 0$ and $v \neq 0$

The $rad_to_cycle_F^*$ approximation helper function:

$$rad_to_cycle_F^* : \mathcal{R} \times F \rightarrow \mathcal{R}$$

$rad_to_cycle_F^*(x, v)$ returns a close approximation to $rad_to_cycle(x, v)$ in \mathcal{R} , with maximum error $max_error_sin_F$, if $|x| \leq big_angle_r_F$.

Further requirements on the $rad_to_cycle_F^*$ approximation helper function are:

$$\begin{aligned} rad_to_cycle_F^*(n \cdot 2 \cdot \pi + \pi/6, v) &= v/12 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/6| \leq big_angle_r_F \\ rad_to_cycle_F^*(n \cdot 2 \cdot \pi + \pi/4, v) &= v/8 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/4| \leq big_angle_r_F \\ rad_to_cycle_F^*(n \cdot 2 \cdot \pi + \pi/3, v) &= v/6 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/3| \leq big_angle_r_F \\ rad_to_cycle_F^*(n \cdot 2 \cdot \pi + \pi/2, v) &= v/4 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi/2| \leq big_angle_r_F \\ rad_to_cycle_F^*(n \cdot 2 \cdot \pi + 2 \cdot \pi/3, v) &= v/3 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + 2 \cdot \pi/3| \leq big_angle_r_F \\ rad_to_cycle_F^*(n \cdot 2 \cdot \pi + 3 \cdot \pi/4, v) &= 3 \cdot v/8 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + 3 \cdot \pi/4| \leq big_angle_r_F \\ rad_to_cycle_F^*(n \cdot 2 \cdot \pi + 5 \cdot \pi/6, v) &= 5 \cdot v/12 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + 5 \cdot \pi/6| \leq big_angle_r_F \\ rad_to_cycle_F^*(n \cdot 2 \cdot \pi + \pi, v) &= v/2 && \text{if } n \in \mathcal{Z} \text{ and } |n \cdot 2 \cdot \pi + \pi| \leq big_angle_r_F \\ rad_to_cycle_F^*(-x, v) &= -rad_to_cycle_F^*(x, v) && \text{if } rad_to_cycle(x, v) \neq v/2 \\ rad_to_cycle_F^*(x, -v) &= -rad_to_cycle_F^*(x, v) && \text{if } rad_to_cycle(x, v) \neq v/2 \end{aligned}$$

The $rad_to_cycle_F$ operation:

$$rad_to_cycle_F : F \times F \rightarrow F \cup \{\mathbf{underflow}, \mathbf{absolute_precision_underflow}, \mathbf{invalid}\}$$

$$\begin{aligned} rad_to_cycle_F(x, v) &= trans_result_F(rad_to_cycle_F^*(x, v), nearest_F) && \text{if } v \in G_F \text{ and } x \in F \text{ and } |x| \leq big_angle_r_F \text{ and } \\ & && x \neq 0 \\ &= mul_F(v, x) && \text{if } v \in G_F \text{ and } x \in \{-0, 0\} \\ &= \mathbf{absolute_precision_underflow}(\mathbf{qNaN}) && \\ &= result_NaN2_F(x, v) && \text{if } v \in G_F \text{ and } x \in F \text{ and } |x| > big_angle_r_F \\ & && \text{otherwise} \end{aligned}$$

5.3.11.2 Converting argument angular-unit angle to radian angle

Define the mathematical function:

$$cycle_to_rad : \mathcal{R} \times \mathcal{R} \rightarrow \mathcal{R}$$

$$\begin{aligned} cycle_to_rad(u, x) &= \arccos(\cos(x \cdot 2 \cdot \pi/u)) && \text{if } \sin(x \cdot 2 \cdot \pi/u) \geq 0 \\ &= -\arccos(\cos(x \cdot 2 \cdot \pi/u)) && \text{if } \sin(x \cdot 2 \cdot \pi/u) < 0 \end{aligned}$$

The $cycle_to_rad_F^*$ approximation helper function:

$$cycle_to_rad_F^* : F \times \mathcal{R} \rightarrow \mathcal{R}$$

$cycle_to_rad_F^*(u, x)$ returns a close approximation to $cycle_to_rad(u, x)$ in \mathcal{R} , if $u \neq 0$, with maximum error $max_error_sin_F$.

Further requirements on the $cycle_to_rad_F^*$ approximation helper function are:

$$\begin{aligned}
\text{cycle_to_rad}_F^*(u, n \cdot u + x) &= \text{cycle_to_rad}_F^*(u, x) \\
&\quad \text{if } n \in \mathcal{Z} \\
\text{cycle_to_rad}_F^*(u, u/12) &= \pi/6 \\
\text{cycle_to_rad}_F^*(u, u/8) &= \pi/4 \\
\text{cycle_to_rad}_F^*(u, u/6) &= \pi/3 \\
\text{cycle_to_rad}_F^*(u, u/4) &= \pi/2 \\
\text{cycle_to_rad}_F^*(u, u/3) &= 2 \cdot \pi/3 \\
\text{cycle_to_rad}_F^*(u, 3 \cdot u/8) &= 3 \cdot \pi/4 \\
\text{cycle_to_rad}_F^*(u, 5 \cdot u/12) &= 5 \cdot \pi/6 \\
\text{cycle_to_rad}_F^*(u, u/2) &= \pi \\
\text{cycle_to_rad}_F^*(u, -x) &= -\text{cycle_to_rad}_F^*(u, x) \\
&\quad \text{if } \text{cycle_to_rad}(u, x) \neq \pi
\end{aligned}$$

The cycle_to_rad_F operation:

$$\begin{aligned}
\text{cycle_to_rad}_F : F \times F &\rightarrow F \cup \{-0, \text{underflow}, \text{absolute_precision_underflow}, \text{invalid}\} \\
\text{cycle_to_rad}_F(u, x) &= \text{trans_result}_F(\text{cycle_to_rad}_F^*(u, x), \text{nearest}_F) \\
&\quad \text{if } \text{cycle}_F(u, x) \in F \text{ and } \text{cycle}_F(u, x) \neq 0 \\
&= \text{mul}_F(\text{cycle}_F(u, x), u) &\quad \text{if } \text{cycle}_F(u, x) \in \{-0, 0\} \\
&= \text{cycle}_F(u, x) &\quad \text{otherwise}
\end{aligned}$$

5.3.11.3 Converting argument angular-unit angle to (another) argument angular-unit angle

Define the mathematical function:

$$\begin{aligned}
\text{cycle_to_cycle} : \mathcal{R} \times \mathcal{R} \times \mathcal{R} &\rightarrow \mathcal{R} \\
\text{cycle_to_cycle}(u, x, v) &= \arccos(\cos(x \cdot 2 \cdot \pi/u)) \cdot v/(2 \cdot \pi) \\
&\quad \text{if } u \neq 0 \text{ and } v \neq 0 \text{ and } \sin(x \cdot 2 \cdot \pi/u) \geq 0 \\
&= -\arccos(\cos(x \cdot 2 \cdot \pi/u)) \cdot v/(2 \cdot \pi) \\
&\quad \text{if } u \neq 0 \text{ and } v \neq 0 \text{ and } \sin(x \cdot 2 \cdot \pi/u) < 0
\end{aligned}$$

The $\text{cycle_to_cycle}_F^*$ approximation helper function:

$$\text{cycle_to_cycle}_F^* : F \times \mathcal{R} \times F \rightarrow \mathcal{R}$$

$\text{cycle_to_cycle}_F^*(u, x, v)$ returns a close approximation to $\text{cycle_to_cycle}(u, x, v)$ in \mathcal{R} if $u \neq 0$ and $|x/u| \leq \text{big_angle_u}_F$, with maximum error max_error_sin_F .

Further requirements on the $\text{cycle_to_cycle}_F^*$ approximation helper function are:

$$\begin{aligned}
\text{cycle_to_cycle}_F^*(u, n \cdot u + x, v) &= \text{cycle_to_cycle}_F^*(u, x, v) \\
&\quad \text{if } n \in \mathcal{Z} \\
\text{cycle_to_cycle}_F^*(u, u/12, v) &= v/12 \\
\text{cycle_to_cycle}_F^*(u, u/8, v) &= v/8 \\
\text{cycle_to_cycle}_F^*(u, u/6, v) &= v/6 \\
\text{cycle_to_cycle}_F^*(u, u/4, v) &= v/4 \\
\text{cycle_to_cycle}_F^*(u, u/3, v) &= v/3 \\
\text{cycle_to_cycle}_F^*(u, 3 \cdot u/8, v) &= 3 \cdot v/8 \\
\text{cycle_to_cycle}_F^*(u, 5 \cdot u/12, v) &= 5 \cdot v/12 \\
\text{cycle_to_cycle}_F^*(u, u/2, v) &= v/2 \\
\text{cycle_to_cycle}_F^*(u, -x, v) &= -\text{cycle_to_cycle}_F^*(u, x, v) \\
&\quad \text{if } \text{cycle_to_cycle}(u, x, v) \neq v/2 \\
\text{cycle_to_cycle}_F^*(-u, x, v) &= -\text{cycle_to_cycle}_F^*(u, x, v)
\end{aligned}$$

$$\begin{aligned}
 & \text{if } \mathit{cycle_to_cycle}(u, x, v) \neq v/2 \\
 \mathit{cycle_to_cycle}_F^*(u, x, -v) &= -\mathit{cycle_to_cycle}_F^*(u, x, v) \\
 & \text{if } \mathit{cycle_to_cycle}(u, x, v) \neq v/2
 \end{aligned}$$

The $\mathit{cycle_to_cycle}_F$ operation:

$$\begin{aligned}
 \mathit{cycle_to_cycle}_F &: F \times F \times F \rightarrow F \cup \{-\mathbf{0}, \text{underflow}, \text{absolute_precision_underflow}, \text{invalid}\} \\
 \mathit{cycle_to_cycle}_F(u, x, v) & \\
 &= \mathit{trans_result}_F(\mathit{cycle_to_cycle}_F^*(u, x, v), \mathit{nearest}_F) \\
 & \quad \text{if } v \in G_F \text{ and } \mathit{cycle}_F(u, x) \in F \text{ and } \mathit{cycle}_F(u, x) \neq 0 \\
 &= \mathit{mul}_F(v, \mathit{cycle}_F(u, x)) \quad \text{if } v \in G_F \text{ and } \mathit{cycle}_F(u, x) \in \{-\mathbf{0}, 0\} \\
 &= \text{absolute_precision_underflow}(\mathbf{qNaN}) \\
 & \quad \text{if } v \in G_F \text{ and} \\
 & \quad \quad \mathit{cycle}_F(u, x) = \text{absolute_precision_underflow} \\
 &= \mathit{result_NaN}_F(u, x, v) \quad \text{otherwise}
 \end{aligned}$$

5.4 Conversion operations

Numeric conversion between different representation forms for integer and fractional values can take place under a number of different circumstances. E.g.:

- a) explicit or implicit conversion between different numeric datatypes conforming to Part 1;
- b) explicit or implicit conversion between different numeric datatypes only one of which conforms to Part 1;
- c) explicit or implicit conversion between a character string and a numeric datatype.

The latter includes outputting a numeric value as a character string, inputting a numeric value from a character string source, and converting a numeral in the source program to a value in a numeric datatype (see 5.5). This Part covers only the cases where at least one of the source and target is a numeric datatype conforming to Part 1.

When a character string is involved as either source or target of a conversion, this Part does not specify the lexical syntax for the numerals parsed or formed. A binding standard should specify the lexical syntax or syntaxes for these numerals, and, when appropriate, how the lexical syntax for the numerals can be altered. With the exception of the radix used in numerals expressing fractional values, differences in lexical syntactic details that do not affect the value in \mathcal{R} denoted by the numerals should not affect the result of the conversion.

Character string representations for integer values can include representations for $-\mathbf{0}$, $+\infty$, $-\infty$, and quiet NaNs. Character string representations for floating point and fixed point values should have formats for $-\mathbf{0}$, $+\infty$, $-\infty$, and quiet NaNs. For both integer and floating point values, character strings that are not numerals nor special values according to the lexical syntax used, shall be regarded as signalling NaNs when used as source of a numerical conversion.

For the cases where one of the datatypes involved in the conversion does not conform to Part 1, the values of some numeric datatype parameters need to be inferred. For integers, one need to infer the value for *bounded*, and if that is **true** then also values for *maxint* and *minint*. For floating point values, one need to infer the values for *r*, *p*, and *emax* or *emin*. In case a precise determination is not possible, values that are ‘safe’ for that instance should be used. ‘Safe’ values for otherwise undetermined inferred parameters are such that

- a) monotonicity of the conversion function is not affected,
- b) the error in the conversion does not exceed that specified by the maximum error parameter (see below),

- c) if the value resulting from the conversion is converted back to the source datatype by a conversion conforming to this Part, the original value should be regenerated if possible, and
- d) overflow and underflow are avoided if possible.

If, and only if, a specified infinite special value result cannot be represented in the target datatype, the infinity result shall be interpreted as **pole**. If, and only if, a specified NaN special value result cannot be represented in the target datatype, the NaN result shall be interpreted as **invalid**.

5.4.1 Integer to integer conversions

Let I and I' be non-special value sets for integer datatypes. At least one of the datatypes corresponding to I and I' conforms to Part 1.

$$\begin{aligned}
 & \text{convert}_{I \rightarrow I'} : I \rightarrow I' \cup \{\mathbf{overflow}\} \\
 & \text{convert}_{I \rightarrow I'}(x) = \text{result}_{I'}(x) && \text{if } x \in I \\
 & & = x && \text{if } x \in \{-\infty, -0, +\infty\} \\
 & & = \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN} \\
 & & = \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \text{ is a signalling NaN}
 \end{aligned}$$

NOTE – If both I and I' are conforming to Part 1, then this conversion is covered by Part 1. This operation generalises the $\text{cvt}_{I \rightarrow I'}$ of Part 1, since only one of the integer datatypes in the conversion need be conforming to Part 1.

5.4.2 Floating point to integer conversions

Let I be the non-special value set for an integer datatype conforming to Part 1. Let F be the non-special value set for a floating point datatype conforming to Part 1.

NOTE – The operations in this clause are more specific than the floating point to integer conversion in Part 1 which allows any rounding.

$$\begin{aligned}
 & \text{rounding}_{F \rightarrow I} : F \rightarrow I \cup \{-0, \mathbf{overflow}\} \\
 & \text{rounding}_{F \rightarrow I}(x) \\
 & \quad = \text{result}_I(\text{round}(x)) && \text{if } x \in F \text{ and } (x \geq 0 \text{ or } \text{round}(x) \neq 0) \\
 & \quad = -0 && \text{if } x \in F \text{ and } x < 0 \text{ and } \text{round}(x) = 0 \\
 & \quad = x && \text{if } x \in \{-\infty, -0, +\infty\} \\
 & \quad = \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN} \\
 & \quad = \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \text{ is a signalling NaN}
 \end{aligned}$$

$$\begin{aligned}
 & \text{floor}_{F \rightarrow I} : F \rightarrow I \cup \{\mathbf{overflow}\} \\
 & \text{floor}_{F \rightarrow I}(x) \\
 & \quad = \text{result}_I(\lfloor x \rfloor) && \text{if } x \in F \\
 & \quad = x && \text{if } x \in \{-\infty, -0, +\infty\} \\
 & \quad = \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN} \\
 & \quad = \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \text{ is a signalling NaN}
 \end{aligned}$$

$$\begin{aligned}
 & \text{ceiling}_{F \rightarrow I} : F \rightarrow I \cup \{-0, \mathbf{overflow}\} \\
 & \text{ceiling}_{F \rightarrow I}(x) \\
 & \quad = \text{result}_I(\lceil x \rceil) && \text{if } x \in F \text{ and } (x \geq 0 \text{ or } \lceil x \rceil \neq 0) \\
 & \quad = -0 && \text{if } x \in F \text{ and } x < 0 \text{ and } \lceil x \rceil = 0 \\
 & \quad = x && \text{if } x \in \{-\infty, -0, +\infty\} \\
 & \quad = \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN} \\
 & \quad = \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \text{ is a signalling NaN}
 \end{aligned}$$

5.4.3 Integer to floating point conversions

Let I be the non-special value set for an integer datatype. Let F be the non-special value set for a floating point datatype. At least one of the source and target datatypes is conforming to Part 1.

$$\begin{aligned}
 \text{convert}_{I \rightarrow F} &: I \rightarrow F \cup \{\mathbf{overflow}\} \\
 \text{convert}_{I \rightarrow F}(x) &= \text{result}_F(x, \text{nearest}_F) && \text{if } x \in I \\
 &= x && \text{if } x \in \{-\infty, -0, +\infty\} \\
 &= \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN} \\
 &= \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \text{ is a signalling NaN}
 \end{aligned}$$

NOTE – When both I and F conform to Part 1, integer to nearest floating point conversions are covered by Part 1. In this case the operations $\text{cvt}_{I \rightarrow F}$ and $\text{convert}_{I \rightarrow F}$ are identical.

5.4.4 Floating point to floating point conversions

Define the least radix function, lb , defined for arguments that are greater than 0:

$$\begin{aligned}
 lb &: \mathcal{Z} \rightarrow \mathcal{Z} \\
 lb(r) &= \min\{n \in \mathcal{Z} \mid n \geq 1 \text{ and } \exists m \in \mathcal{Z} : r = n^m\}
 \end{aligned}$$

Let F , F' , and F'' be non-special value sets for floating point datatypes. At least one of the source and target datatypes in the conversion conforms to Part 1.

There shall be a $\text{max_error_convert}_{F'}$ parameter that gives the maximum error when converting from F to F' and $lb(r_F) \neq lb(r_{F'})$. The $\text{max_error_convert}_{F'}$ parameter shall have a value in the interval $[0.5, 0.75]$. If $lb(r_F) = lb(r_{F'})$, the maximum error shall be 0.5 ulp when converting from F to F' , but this is not reflected in any parameter.

The $\text{convert}_{F \rightarrow F'}^*$ approximation helper functions:

$$\text{convert}_{F \rightarrow F'}^* : \mathcal{R} \rightarrow \mathcal{R}$$

$\text{convert}_{F \rightarrow F'}^*(x)$ returns a close approximation to x in \mathcal{R} , with maximum error $\text{max_error_convert}_{F'}$.

Further requirements on the $\text{convert}_{F \rightarrow F'}^*$ approximation helper functions:

$$\begin{aligned}
 \text{convert}_{F \rightarrow F'}^*(x) &= x && \text{if } x \in \mathcal{Z} \\
 \text{convert}_{F \rightarrow F'}^*(x) &> 0 && \text{if } x > 0 \\
 \text{convert}_{F \rightarrow F'}^*(-x) &= -\text{convert}_{F \rightarrow F'}^*(x) \\
 \text{convert}_{F \rightarrow F'}^*(x) &\leq \text{convert}_{F \rightarrow F'}^*(y) && \text{if } x < y
 \end{aligned}$$

Relationship to other floating point to floating point conversion approximation helper functions:

$$\text{convert}_{F \rightarrow F'}^*(x) = \text{convert}_{F'' \rightarrow F'}^*(x) \quad \text{if } lb(r_{F''}) = lb(r_F) \text{ and } x \in F \cap F''$$

The $\text{convert}_{F \rightarrow F'}$ operation:

$$\begin{aligned}
 \text{convert}_{F \rightarrow F'} &: F \rightarrow F' \cup \{\mathbf{overflow}, \mathbf{underflow}\} \\
 \text{convert}_{F \rightarrow F'}(x) &= \text{result}_{F'}(x, \text{nearest}_{F'}) && \text{if } x \in F \text{ and } lb(r_F) = lb(r_{F'}) \\
 &= \text{trans_result}_{F'}(\text{convert}_{F \rightarrow F'}^*(x), \text{nearest}_{F'}) && \text{if } x \in F \text{ and } lb(r_F) \neq lb(r_{F'}) \\
 &= x && \text{if } x \in \{-\infty, -0, +\infty\} \\
 &= \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN} \\
 &= \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \text{ is a signalling NaN}
 \end{aligned}$$

NOTE – When both datatypes conform to Part 1, and the radices for both of these floating point datatypes are the same, floating point to nearest floating point conversions are covered by Part 1. In this case the operations $cut_{F \rightarrow F'}$ and $convert_{F \rightarrow F'}$ are identical.

5.4.5 Floating point to fixed point conversions

Let F be the non-special value set for a floating point datatype conforming to Part 1. Let D be the non-special value set for a fixed point datatype.

A fixed point datatype D is a subset of \mathcal{R} , characterised by a radix, $r_D \in \mathcal{Z}$ (≥ 2), a density, $d_D \in \mathcal{Z}$ (≥ 0), and if it is bounded, a maximum positive value, $dmax_D \in D^*$ (≥ 1). Given these values, the following sets are defined:

$$D^* = \{n/(r_D^{d_D}) \mid n \in \mathcal{Z}\}$$

$$D = D^* \quad \text{if } D \text{ is not bounded}$$

$$D = D^* \cap [-dmax_D, dmax_D] \quad \text{if } D \text{ is bounded}$$

NOTE 1 – D corresponds to **scaled**(r_D , d_D) in ISO/IEC 11404 Language independent datatypes (LID) [10]. LID has no parameter corresponding to $dmax_D$ even when the datatype is bounded.

The fixed point rounding helper function:

$$nearest_D : \mathcal{R} \rightarrow D^*$$

is the rounding function that rounds to nearest, ties round to even last digit.

The fixed point result helper function, $result_D$, is like $result_F$, but for a fixed point datatype. It will return **overflow** if the rounded result is not representable:

$$result_D : \mathcal{R} \times (\mathcal{R} \rightarrow D^*) \rightarrow D \cup \{\mathbf{overflow}\}$$

$$\begin{aligned} result_D(x, rnd) &= rnd(x) && \text{if } rnd(x) \in D \text{ and } (rnd(x) \neq 0 \text{ or } x \geq 0) \\ &= -\mathbf{0} && \text{if } rnd(x) = 0 \text{ and } x < 0 \\ &= \mathbf{overflow} && \text{if } x \in \mathcal{R} \text{ and } rnd(x) \notin D \end{aligned}$$

There shall be a $max_error_convert_D$ parameter that gives the maximum error when converting from F to D and $lb(r_F) \neq lb(r_D)$. The $max_error_convert_D$ parameter shall have a value in the interval $[0.5, 0.75]$. If $lb(r_F) = lb(r_D)$, the maximum error shall be 0.5 ulp when converting from F to D , but this is not reflected in any parameter.

The $convert_{F \rightarrow D}^*$ approximation helper function:

$$convert_{F \rightarrow D}^* : \mathcal{R} \rightarrow \mathcal{R}$$

$convert_{F \rightarrow D}^*(x)$ returns a close approximation to x in \mathcal{R} , with maximum error $max_error_convert_D$.

Further requirements on the $convert_{F \rightarrow D}^*$ approximation helper functions:

$$\begin{aligned} convert_{F \rightarrow D}^*(x) &= x && \text{if } x \in \mathcal{Z} \\ convert_{F \rightarrow D}^*(x) &> 0 && \text{if } x > 0 \\ convert_{F \rightarrow D}^*(-x) &= -convert_{F \rightarrow D}^*(x) \\ convert_{F \rightarrow D}^*(x) &\leq convert_{F \rightarrow D}^*(y) && \text{if } x < y \end{aligned}$$

Relationship to other floating point to fixed point conversion approximation helper functions:

$$convert_{F \rightarrow D}^*(x) = convert_{F'' \rightarrow D}^*(x) \quad \text{if } lb(r_{F''}) = lb(r_F) \text{ and } x \in F \cap F''$$

The $convert_{F \rightarrow D}$ operation:

$$convert_{F \rightarrow D} : F \rightarrow D \cup \{-\mathbf{0}, \mathbf{overflow}\}$$

$$\begin{aligned}
\mathit{convert}_{F \rightarrow D}(x) &= \mathit{result}_D(x, \mathit{nearest}_D) && \text{if } x \in F \text{ and } lb(r_F) = lb(r_D) \\
&= \mathit{result}_D(\mathit{convert}_{F \rightarrow D}^*(x), \mathit{nearest}_D) && \text{if } x \in F \text{ and } lb(r_F) \neq lb(r_D) \\
&= x && \text{if } x \in \{-\infty, -0, +\infty\} \\
&= \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN} \\
&= \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \text{ is a signalling NaN}
\end{aligned}$$

NOTES

- 2 The datatype D need not be visible in the programming language. D may be a subtype of strings, according to some format. Even so, no datatype for strings need be present in the programming language.
- 3 This covers, among other things, “output” of floating point datatype values, to fixed point string formats. E.g. a binding may say that `float_to_fixed_string(x, m, n)` is bound to $\mathit{convert}_{F \rightarrow S_{m,n}}(x)$ where $S_{m,n}$ is strings of length m , representing fixed point values in radix 10 with n decimals. The binding should also detail how NaNs, signed zeroes and infinities are represented in $S_{m,n}$, as well as the precise format of the strings representing ordinary values. (Note that if the length of the target string is limited, the conversion may overflow.)

5.4.6 Fixed point to floating point conversions

Let F be the non-special value set for a floating point datatype conforming to Part 1. Let D and D' be the non-special value set for fixed point datatypes.

The $\mathit{convert}_{D \rightarrow F}^*$ approximation helper function:

$$\mathit{convert}_{D \rightarrow F}^* : \mathcal{R} \rightarrow \mathcal{R}$$

$\mathit{convert}_{D \rightarrow F}^*(x)$ returns a close approximation to x in \mathcal{R} , with maximum error `max_error_convert_F`.

Further requirements on the $\mathit{convert}_{D \rightarrow F}^*$ approximation helper functions:

$$\begin{aligned}
\mathit{convert}_{D \rightarrow F}^*(x) &= x && \text{if } x \in \mathcal{Z} \\
\mathit{convert}_{D \rightarrow F}^*(x) &> 0 && \text{if } x > 0 \\
\mathit{convert}_{D \rightarrow F}^*(-x) &= -\mathit{convert}_{D \rightarrow F}^*(x) \\
\mathit{convert}_{D \rightarrow F}^*(x) &\leq \mathit{convert}_{D \rightarrow F}^*(y) && \text{if } x < y
\end{aligned}$$

Relationship to other floating point and fixed point to floating point conversion approximation helper functions:

$$\begin{aligned}
\mathit{convert}_{D \rightarrow F}^*(x) &= \mathit{convert}_{D' \rightarrow F}^*(x) && \text{if } lb(r_{D'}) = lb(r_D) \text{ and } x \in D \cap D' \\
\mathit{convert}_{D \rightarrow F}^*(x) &= \mathit{convert}_{F' \rightarrow F}^*(x) && \text{if } lb(r_{F'}) = lb(r_D) \text{ and } x \in D \cap F'
\end{aligned}$$

The $\mathit{convert}_{D \rightarrow F}$ operation:

$$\begin{aligned}
\mathit{convert}_{D \rightarrow F} &: D \rightarrow F \cup \{\mathbf{overflow}, \mathbf{underflow}\} \\
\mathit{convert}_{D \rightarrow F}(x) &= \mathit{result}_F(x, \mathit{nearest}_F) && \text{if } x \in D \text{ and } lb(r_D) = lb(r_F) \\
&= \mathit{trans_result}_F(\mathit{convert}_{D \rightarrow F}^*(x), \mathit{nearest}_F) && \text{if } x \in D \text{ and } lb(r_D) \neq lb(r_F) \\
&= x && \text{if } x \in \{-\infty, -0, +\infty\} \\
&= \mathbf{qNaN} && \text{if } x \text{ is a quiet NaN} \\
&= \mathbf{invalid}(\mathbf{qNaN}) && \text{if } x \text{ is a signalling NaN}
\end{aligned}$$

5.5 Numerals as operations in the programming language

NOTE – Numerals as input, or in strings, is covered by the conversion operations above.

Each numeral is a parameterless operation. Thus, this clause introduces a very large number of operations, since the number of numerals is in principle infinite.

5.5.1 Numerals for integer datatypes

Let I' be a non-special value set for integer numerals for the datatype corresponding to I .

An integer numeral, denoting an abstract value n in $I' \cup \{-0, +\infty, -\infty, \mathbf{qNaN}, \mathbf{sNaN}\}$, for an integer datatype, I , shall result in

$$\mathit{convert}_{I' \rightarrow I}(n)$$

For each integer datatype conforming to Part 1 and made directly available, with non-special value set I , there shall be integer numerals with radix 10.

For each radix for numerals made available for a bounded integer datatype I , there shall be integer numerals for all non-negative values of I .

For each radix for numerals made available for an unbounded integer datatype I , there shall be integer numerals for all non-negative values of I smaller than 10^{20} .

For each integer datatype made directly available and that has special values:

- a) There should be a numeral for positive infinity.
- b) There should be numerals for quiet and signalling NaNs.

5.5.2 Numerals for floating point datatypes

Let D' be a non-special value set for fixed point numerals for the datatype corresponding to F . Let F' be a non-special value set for floating point numerals for the datatype corresponding to F .

A fixed point numeral, denoting an abstract value x in $D' \cup \{-0, +\infty, -\infty, \mathbf{qNaN}, \mathbf{sNaN}\}$, for a floating point datatype, F , shall result in

$$\mathit{convert}_{D' \rightarrow F}(x)$$

A floating point numeral, denoting an abstract value x in $F' \cup \{-0, +\infty, -\infty, \mathbf{qNaN}, \mathbf{sNaN}\}$, for a floating point datatype, F , shall result in

$$\mathit{convert}_{F' \rightarrow F}(x)$$

For each floating point datatype conforming to Part 1 and made directly available, with non-special value set F , there should be radix 10 floating point numerals, and there shall be radix 10 fixed point numerals.

For each radix for fixed point numerals made available for a floating point datatype F , there shall be numerals for all bounded precision and bounded range expressible non-negative values of \mathcal{R} . At least a precision ($d_{D'}$) of 20 should be available. At least a range ($dm\mathit{ax}_{D'}$) of 10^{20} should be available.

For each radix for floating point numerals made available for a floating point datatype F , there shall be numerals for all bounded precision and bounded range expressible non-negative values of \mathcal{R} . The precision and range bounds for the numerals shall be large enough to allow all non-negative values of F to be reachable.

For each floating point datatype made directly available:

- a) There shall be a numeral for positive infinity.

b) There shall be numerals for quiet and signalling NaNs.

The conversion operations used for numerals as operations should be the same as those used by default for converting strings to values in conforming integer or floating point datatypes.

6 Notification

Notification is the process by which a user or program is informed that an arithmetic operation cannot return a suitable numeric result. Specifically, a notification shall occur when any arithmetic operation returns an exceptional value. Notification shall be performed according to the requirements of clause 6 of Part 1.

An implementation shall not give notifications for operations conforming to this Part, unless the specification requires that an exceptional value results for the given arguments.

The default method of notification should be recording of indicators.

6.1 Continuation values

If notifications are handled by a recording of indicators, in the event of notification the implementation shall provide a *continuation value* to be used in subsequent arithmetic operations. Continuation values may be in *I* or *F* (as appropriate), or be special values (-0 , $-\infty$, $+\infty$, or a **qNaN**).

Floating point datatypes that satisfy the requirements of IEC 60559 have special values in addition to the values in *F*. These are: -0 , $+\infty$, $-\infty$, *signaling NaNs* (**sNaN**), and *quiet NaNs* (**qNaN**). Such values may be passed as arguments to operations, and used as results or continuation values. Floating point types that do not fully conform to IEC 60559 can also have values corresponding to -0 , $+\infty$, $-\infty$, or **NaN**.

Continuation values of -0 , $+\infty$, $-\infty$, and **NaN** are required only if the parameter *iec_559_F* has the value **true**. If the implementation can represent such special values in the result datatype, they should be used according to the specifications in this Part.

7 Relationship with language standards

A computing system often provides some of the operations specified in this Part within the context of a programming language. The requirements of the present standard shall be in addition to those imposed by the relevant programming language standards.

This Part does not define the syntax of arithmetic expressions. However, programmers need to know how to reliably access the operations specified in this Part.

NOTE 1 – Providing the information required in this clause is properly the responsibility of programming language standards. An individual implementation would only need to provide details if it could not cite an appropriate clause of the language or binding standard.

An implementation shall document the notation that should be used to invoke an operation specified in this Part and made available. An implementation should document the notation that should be used to invoke an operation specified in this Part and that could be made available.

NOTE 2 – For example, the radian arc sine operation for an argument x ($\arcsin_F(x)$) might be invoked as

<code>arcsin(x)</code>	in Pascal [28] and Ada [11]
<code>asin(x)</code>	in C [18] and Fortran [23]
<code>(asin x)</code>	in Common Lisp [43] and ISLisp [25]
<code>function asin(x)</code>	in COBOL [20]

with suitable expression of the argument (x).

An implementation shall document the semantics of arithmetic expressions in terms of compositions of the operations specified in clause 5 of this Part and in clause 5 of Part 1.

Compilers often “optimize” code as part of compilation. Thus, an arithmetic expression might not be executed as written. An implementation shall document the possible transformations of arithmetic expressions (or groups of expressions) that it permits. Typical transformations include

- a) Insertion of operations, such as datatype conversions or changes in precision.
- b) Replacing operations (or entire subexpressions) with others, such as “`cos(-x)`” \rightarrow “`cos(x)`” (exactly the same result) or “`pi - arccos(x)`” \rightarrow “`arccos(-x)`” (more accurate result) or “`exp(x)-1`” \rightarrow “`expm1(x)`” (more accurate result if $x > -1$, less accurate result if $x < -1$, different notification behaviour).
- c) Evaluating constant subexpressions.
- d) Eliminating unneeded subexpressions.

Only transformations which alter the semantics of an expression (the values produced, and the notifications generated) need be documented. Only the range of permitted transformations need be documented. It is not necessary to describe the specific choice of transformations that will be applied to a particular expression.

The textual scope of such transformations shall be documented, and any mechanisms that provide programmer control over this process should be documented as well.

NOTE 3 – It is highly desirable that programming languages intended for numerical use provide means for limiting the transformations applied to particular arithmetic expressions. Control over changes of precision is particularly useful.

8 Documentation requirements

In order to conform to this Part, an implementation shall include documentation providing the following information to programmers.

NOTE 1 – Much of the documentation required in this clause is properly the responsibility of programming language or binding standards. An individual implementation would only need to provide details if it could not cite an appropriate clause of the language or binding standard.

- a) A list of the provided operations that conform to this Part.
- b) For each maximum error parameter, the value of that parameter or definition of that parameter function. Only maximum error parameters that are relevant to the provided operations need be given.
- c) The value of the parameters *big_angle_r_F* and *big_angle_u_F*. Only big angle parameters that are relevant to the provided operations need be given.
- d) For the *nearest_F* function, the rule used for rounding halfway cases, unless *iec_559_F* is fixed to true.

- e) For each conforming operation, the continuation value provided for each notification condition. Specific continuation values that are required by this Part need not be documented. If the notification mechanism does not make use of continuation values (see clause 6), continuation values need not be documented.

NOTE 2 – Implementations that do not provide infinities or NaNs will have to document any continuation values used in place of such values.

- f) For each conforming operation, how the results depend on the rounding mode, if rounding modes are provided. Operations may be insensitive to the rounding mode, or sensitive to it, but even then need not heed the rounding mode.
- g) For each conforming operation, the notation to be used for invoking that operation.
- h) For each maximum error parameter, the notation to be used to access that parameter.
- i) The notation to be used to access the parameters *big_angle_r_F* and *big_angle_u_F*.

Since the integer and floating point datatypes used in conforming operations shall satisfy the requirements of Part 1, the following information shall also be provided by any conforming implementation.

- j) The translation of arithmetic expressions into combinations of the operations provided by any part of ISO/IEC 10967, including any use made of higher precision. (See clause 7 of Part 1.)
- k) The methods used for notification, and the information made available about the notification. (See clause 6 of Part 1.)
- l) The means for selecting among the notification methods, and the notification method used in the absence of a user selection. (See 6.3 of Part 1.)
- m) The means for selecting the modes of operation that ensure conformity.
- n) When “recording of indicators” is the method of notification, the datatype used to represent *Ind*, the method for denoting the values of *Ind* (the association of these values with the subsets of *E* must be clear), and the notation for invoking each of the “indicator” operations. (See 6.1.2 of Part 1.) In interpreting 6.1.2 of Part 1, the set of indicators *E* shall be interpreted as including all exceptional values listed in the signatures of conforming operations. In particular, *E* may need to contain **pole** and **absolute_precision_underflow**.
- o) For each of the provided operations where this Part specifies a relation to another operation specified in this Part, the binding for that other operation.
- p) For numerals conforming to this Part, which available string conversion operations, including reading from input, give exactly the same conversion results, even if the string syntaxes for ‘internal’ and ‘external’ numerals are different.

Annex A (normative)

Partial conformity

If an implementation of an operation fulfills all relevant requirements according to the normative text in this Part, except the ones relaxed in this Annex, the implementation of that operation is said to *partially conform* to this Part.

Conformity to this Part shall not be claimed for operations that only fulfill Partial conformity.

Partial conformity shall not be claimed for operations that relax other requirements than those relaxed in this Annex.

A.1 Maximum error relaxation

This Part has the following maximum error requirements for conformity.

$$\mathit{max_error_hypot}_F \in [0.5, 1]$$

$$\mathit{max_error_exp}_F \in [0.5, 1.5 * \mathit{rnd_error}_F]$$

$$\mathit{max_error_power}_F \in [\mathit{max_error_exp}_F, 2 * \mathit{rnd_error}_F]$$

$$\mathit{max_error_sinh}_F \in [0.5, 2 * \mathit{rnd_error}_F]$$

$$\mathit{max_error_tanh}_F \in [\mathit{max_error_sinh}_F, 2 * \mathit{rnd_error}_F]$$

$$\mathit{max_error_sin}_F \in [0.5, 1.5 * \mathit{rnd_error}_F]$$

$$\mathit{max_error_tan}_F \in [\mathit{max_error_sin}_F, 2 * \mathit{rnd_error}_F]$$

$$\mathit{max_error_sinu}_F : F \rightarrow F \cup \{\mathbf{invalid}\}$$

$$\mathit{max_error_tanu}_F : F \rightarrow F \cup \{\mathbf{invalid}\}$$

$$\mathit{max_error_convert}_F \in [0.5, 0.75]$$

For $u \in G_F$, the $\mathit{max_error_sinu}_F(u)$ parameter shall be in the interval $[\mathit{max_error_sin}_F, 2]$, and the $\mathit{max_error_tanu}_F(u)$ parameter shall be in the interval $[\mathit{max_error_tan}_F, 4]$. For $u \in T$, the $\mathit{max_error_sinu}_F(u)$ parameter shall be equal to $\mathit{max_error_sin}_F$, and the $\mathit{max_error_tanu}_F(u)$ parameter shall be equal to $\mathit{max_error_tan}_F$.

In a Partially conforming implementation the maximum error parameters may be greater than what is specified by this Part. The maximum error parameter values given by an implementation shall still adequately reflect the accuracy of the relevant operations, if a claim of Partial conformity is made.

A Partially conforming implementation shall document which maximum error parameters have greater values than specified by this Part, and their values.

A.2 Extra accuracy requirements relaxation

This Part has a number of extra accuracy requirements. These are detailed in the paragraphs beginning “Further requirements on the op_F^* approximation helper function are:”.

In a Partially conforming implementation these further requirements need not be fulfilled. The values returned must still be within the maximum error bounds that are given by the maximum error parameters, if a claim of Partial conformity is made.

A Partially conforming implementation shall document which extra accuracy requirements are not fulfilled by the implementation.

A.3 Relationships to other operations relaxation

This Part has a number of requirements giving relations to other operations. These are detailed in the paragraphs beginning “Relationship to the op_F^* approximation helper function:”.

In a Partially conforming implementation these relationships need not be fulfilled. The values returned must still be within the maximum error bounds that are given by the maximum error parameters, if a claim of Partial conformity is made.

A Partially conforming implementation shall document which operation relationships are not fulfilled by the implementation.

Annex B (informative)

Rationale

This annex explains and clarifies some of the ideas behind *Information technology – Language independent arithmetic – Part 2: Elementary numerical functions* (LIA-2).

B.1 Scope

B.1.1 Inclusions

LIA-2 is intended to define the meaning of some operations on Integer and floating point types as specified in LIA-1 (ISO/IEC 10967-1), in addition to the operations specified in LIA-1. LIA-2 does not specify operations for any additional arithmetic datatypes, though fixed point datatypes are used in some of the specifications for conversion operations.

The specifications for the operations covered by LIA-2 are given in sufficient detail to

- a) support detailed and accurate numerical analysis of arithmetic algorithms,
- b) enable a precise determination of conformity or non-conformity, and
- c) prevent exceptions (like overflow) from going undetected.

LIA-2 does in no way prevent language standards or implementations including further arithmetic operations, other variations of included arithmetic operations, or the inclusion of further arithmetic datatypes, like rational number or fixed point datatypes. Some of these may become the topic of standardisation in other parts of LIA.

B.1.2 Exclusions

LIA-2 is not concerned with techniques for the *implementation* of numerical functions. Even when an LIA-2 specification is made in terms of other LIA-1 or LIA-2 operations, that does *not* imply a requirement that an implementation implements the operation in terms of those other operations. It is sufficient that the result (returned value or returned continuation value, and exception behaviour) is *as if* it was implemented in terms of those other operations.

LIA-2 does not provide specifications for operations which involve no arithmetic processing, like assignment and parameter passing, though any implicit conversions done in association with such operations are in scope. The implicit conversions should be made available to the programmer as explicit conversions.

LIA-2 does not cover operations for the support of domains such as linear algebra, statistics, and symbolic processing. Such domains deserve *separate* standardisation, if standardised.

LIA-2 only covers operations that involve integer or floating point datatypes, as specified in LIA-1, and in some cases also a Boolean datatype, but then only as result. The operations covered by LIA-2 are often to some extent covered by programming language standards, like the operations `sin`, `cos`, `tan`, `arctan`, and so on.

B.2 Conformity

Conformity to this standard is dependent on the existence of language binding standards. Each programming language committee (or other organisation responsible for a programming language or other specification to which LIA-1 and LIA-2 may apply) is encouraged to produce a binding standard covering at least those operations already required by the programming language (or similar) and also specified in LIA-2.

The term “programming language” is here used in a generalised sense to include other computing entities such as calculators, spread sheets, page description languages, web-script languages, and database query languages to the extent that they provide the operations covered by LIA-2.

Suggestions for bindings are provided in Annex C. Annex C has partial binding examples for a number of existing programming languages and LIA-2. In addition to the bindings for the operations in LIA-2, it is also necessary to provide bindings for the maximum error parameters and big angle parameters specified by LIA-2. Annex C contains suggestions for these bindings. To conform to this standard, in the absence of a binding standard, an implementation should create a binding, following the suggestions in Annex C.

B.3 Normative references

The referenced IEC 60559 standard is identical to the former IEC 559 and IEEE 754 standards.

B.4 Symbols and definitions

B.4.1 Symbols

B.4.1.1 Sets and intervals

The interval notation is in common use. It has been chosen over the other commonly used interval notation because the chosen notation has no risk of confusion with the pair notation.

B.4.1.2 Operators and relations

Note that all operators, relations, and other mathematical notation used in LIA-2 is used in their conventional exact mathematical sense. They are not used to stand for operations specified by IEC 60559, LIA-1, LIA-2, or, with the exception of programme excerpts which are clearly marked, any programming language. E.g. x/u stands for the mathematically exact result of dividing x by u , whether that value is representable in any floating point datatype or not, and $x/u \neq \text{div}_F(x, u)$ is often the case. Likewise, $=$ is the mathematical equality, not the eq_F operation: $0 \neq -\mathbf{0}$, while $eq_F(0, -\mathbf{0}) = \mathbf{true}$.

B.4.1.3 Mathematical functions

The elementary functions named `sin`, `cos`, etc. used in LIA-2 are the exact mathematical functions, not any approximation. The approximations to these mathematical functions are introduced in clauses 5.3 and 5.4 and are written in a way clearly distinct from the mathematical functions. E.g., sin_F^* , cos_F^* , etc., which are unspecified mathematical functions approximating the targeted exact mathematical functions to a specified degree; sin_F , cos_F , etc., which are the operations specified by LIA-2 based on the respective approximating function; `sin`, `cos`, etc., which are programming language names bound to LIA-2 operations. `sin` is thus very different from `sin`.

B.4.1.4 Datatypes and exceptional values

The sequence types $[I]$ and $[F]$ appear as input datatypes to a few operations: max_seq_I , min_seq_I , gcd_seq_I , lcm_seq_I , max_seq_F , min_seq_F , $mmax_seq_F$, and $mmin_seq_F$.

In effect, a sequence is a finite linearly ordered collection of elements which can be indexed from 1 to the length of the sequence. Equality of two or more elements with different indices is possible. Sequences are used in LIA-2 as an abstraction of arrays, lists, other kinds of one-dimensional sequenced collections, and even variable length argument lists. As used in LIA-2 the order of the elements and number of occurrences of each element, as long as it is more than one, does not matter, so multi-sets (bags) and sets also qualify.

LIA-2 uses a modified set of exceptional values compared to LIA-1. Instead of LIA-1's **undefined**, LIA-2 uses **invalid** and **pole**. IEC 60559 distinguishes between **invalid** and **divide_by_zero** (the latter is called **pole** by LIA-2). The distinction is valid and should be recognised, since **pole** indicates that an infinite but *exact* result is (or can be, if it were available) returned, while **invalid** indicates that a result in the target datatype (extended with infinities) cannot, or should not, be returned with adequate accuracy.

LIA-1 distinguished between **integer_overflow** and **floating_overflow**. This distinction is moot, since no distinction was made between **integer_undefined** and **floating_undefined**. In addition, continuing this distinction would force LIA to start distinguishing not only **integer_overflow** and **floating_overflow**, but also **fixed_overflow**, **complex_floating_overflow**, **complex_integer_overflow**, etc. Further, there is no general consensus that maintaining this distinction is useful, and many programming languages do not require a distinction. A binding standard can still maintain this distinction, if desired.

LIA allows for three methods for handing notifications: recording of indicators, change of control flow (returnable or not), and termination of program. The preferred method is recording of indicators. This allows the computation to continue using the continuation values. For **underflow** and **pole** notifications this course of action is strongly preferred, provided that a suitable continuation value can be represented in the result datatype.

Not all occurrences of the same exceptional value need be handled the same. There may be explicit mode changes in how notifications are handled, and there may be implicit changes. E.g., **invalid** without a specified (by LIA-2 or binding) continuation value to cause change of control flow (like an Ada [11] exception), while **invalid** with a specified continuation value use recording of indicators. This should be specified by bindings or by implementations.

The operations may return any of the exceptional values **overflow**, **underflow**, **invalid**, **pole**, or **absolute_precision_underflow**. This does *not* imply that the implemented operations are to actually *return* any of these values. When these values are returned according to the LIA specification, that means that the implementation is to perform a notification handling for that exceptional value. If the notification handling is by recording of indicators, then what is actually returned by the implemented operation is the continuation value.

B.4.2 Definitions of terms

Note the LIA distinction between exceptional values, exceptions, and exception handling (handling of notification by non-returnable change of control flow; as in e.g. Ada). LIA exceptional values are not the same as Ada **exceptions**, nor are they the same as IEC 60559 special values.

B.5 Specifications for the numerical functions

The abstract values used in the specifications are independent of datatype, just like the mathematical numbers are. That they are represented differently in, say, single precision and in double precision is out of scope for LIA-2.

The specifications in LIA-2 for floating point operations give details about certain special values (they are ‘special’ in that they are not in \mathcal{R}). These special values are commonly represented in floating point datatypes, in particular all floating point datatypes conforming to IEC 60559.

B.5.1 Basic integer operations

Integer datatypes can have infinity values as well as NaN values, and also may have a -0 . A corresponding I must, however, be a subset of \mathcal{Z} . -0 is commonly available when the integer datatype is represented using radix-minus-1-complement, e.g. 1’s complement. When using, e.g., 2’s complement, the representation that would otherwise represent the most negative value can be used as a NaN. Especially for unbounded integer types, the inclusion of infinities is advisable, not for overflow, since these do not occur, but in order to have a smallest and a largest value in the type.

B.5.1.1 The integer *result* and *wrap* helper functions

The *result_I* helper function notifies overflow when the result cannot be represented in I . When an overflow occurs, and recording of indicators is the method for handling (integer) overflows, a continuation value must be given. For bounded integer datatypes, *maxint_F* and *minint_F* can be suitable continuation values. In some instances a wrapped result, see below, may be used as continuation value on overflow. Few integer datatypes offer representations for positive and negative infinity. In case such representations are offered, they can be used as continuation values on overflow, similar to their use in floating point datatypes. LIA does not specify the continuation value in this case, that is left to bindings or implementations, but LIA does require that the continuation value(s) be documented.

The *wrap_I* helper function wraps the result into a value that can be represented in I . The result is wrapped in such a way that the value returned can be used to implement extended range integer arithmetic.

B.5.1.2 Integer maximum and minimum

The operations for integer maximum and minimum are trivial, except taking the maximum or minimum of an empty sequence (empty array, empty list, zero number of parameters, or similar). The case for zero number of parameters is often syntactically excluded (as in Fortran, Common Lisp, and ISLisp), while an empty array or empty list given as a single argument must usually be possible handle at ‘runtime’. LIA specifies a **pole** notification for this case. Since no (implied mathematical) division is involved here, **pole** is here to be interpreted as “exact infinite result from finite operands”, in this case an empty list of numbers.

If infinity values are required to be available for a particular integer datatype, a binding may require the continuation values specified to be returned without any **pole** notification. When the specified continuation value, $+\infty$ or $-\infty$, is not available, other suitable continuation values may be used, and if so they must be documented. If the integer datatype is bounded, but without infinities, *maxint_F* may be used in place of $+\infty$ and *minint_F* may be used instead of $-\infty$.

Infinities as arguments are not specified for these operations, since infinities are rarely available in integer datatypes. However, compare the specification for max and min operations for floating point datatypes (clause 5.2.2).

B.5.1.3 Integer diminish

Integer diminish is sometimes called ‘monus’. This operation computes the ‘positive difference’ between two numbers.

B.5.1.4 Integer power and arithmetic shift

The integer arithmetic shift operations can be used to implement integer multiplication and integer division more quickly in special cases.

The shift operations shift either ‘right’ or ‘left’ depending on the sign of the second argument.

Any continuation value used on overflow here must be documented, either by the binding standard or by the implementation.

B.5.1.5 Integer square root

B.5.1.6 Divisibility tests

Even and odd are simple special cases offered as separately named operations in several programming languages.

B.5.1.7 Integer division and remainder

When the result of a division between integers⁴⁷ is not an integer, but the final result is required to be an integer, the quotient must be rounded. There are several ways of doing this; floor, ceiling, and unbiased round to nearest being the most important.

pad_I returns the *negative* of the remainder after division and ceiling. The reason for this is twofold: 1) for unsigned integer datatypes the remainder is ≤ 0 , and would thus often not be representable unless negated, and 2) it is intuitively easier to think of the “places left in the last unfilled group of equi-sized and packed groups” as a positive entity, a padding.

$remr_I$ can overflow only for unsigned integer datatypes ($minint_I = 0$), and does so for too many arguments, and negating it does not change this. $remr_I$ should therefore not be provided for unsigned integer datatypes. $remr_I$ rounds in the same way as $remr_F$, IEEE remainder.

When there is no exception, these operations fulfill $divf_I(x+n \cdot y, y) = divf_I(x, y) + n$, $group_I(x+n \cdot y, y) = group_I(x, y) + n$, $quot_I(x+2 \cdot n \cdot y, y) = quot_I(x, y) + 2 \cdot n$, $moda_I(x+n \cdot y, y) = moda_I(x, y)$, $pad_I(x+n \cdot y, y) = pad_I(x, y)$, $remr_I(x+2 \cdot n \cdot y, y) = remr_I(x, y)$, where $n \in \mathcal{Z}$.

Note that the div_I^t and rem_I^t from LIA-1 do not fulfill similar useful equalities, due to the disruption around 0 for this pair of operations.

When there is no exception, $divf_I(x, y) = -group_I(-x, y)$, $divf_I(x, y) = -group_I(x, -y)$, $quot_I(x, y) = -quot_I(-x, y)$, $quot_I(x, y) = -quot_I(x, -y)$, $moda_I(x, y) = -pad_I(x, -y)$, and $remr_I(x, y) = remr_I(x, -y)$.

B.5.1.8 Greatest common divisor and least common positive multiple

The greatest common divisor is useful in reducing a rational number to its lowest terms. The least common multiple is useful in converting two rational numbers to have the same denominator.

Returning 0 for $gcd_I(0, 0)$, as is sometimes suggested, would be incorrect, since the greatest common divisor for 0 and 0 should be the supremum (upper limit) of \mathcal{Z}^+ , since these all divide 0, which is infinity.

gcd_I will overflow only if $bounded_I = \mathbf{true}$, $minint_I = -maxint_I - 1$, and both arguments are $minint_I$. The greatest common divisor is then $-minint_I$, which then is not in I .

Least common positive multiple, $lcm_I(x, y)$, overflows for many “large” arguments. E.g., if x and y are relative primes, then the least common multiple is $|x \cdot y|$, which may be greater than $maxint_I$.

B.5.1.9 Support operations for extended integer range

These operations would typically be used to extend the range of the highest level integer datatype supported by the underlying hardware of an implementation.

The two parts of an integer product, $mul_{ov_I}(x, y)$ and $mul_{wrap_I}(x, y)$ together provide the complete integer product. Similarly for addition and subtraction.

The use of $wrap_I$ guarantees that **overflow** will not occur.

B.5.2 Basic floating point operations

F must be a subset of \mathcal{R} . Floating point datatypes can have infinity values as well as NaN values, and also may have a -0 . These values are not in F . The special values are, however, commonly available in floating point datatypes today, thanks to the wide adoption of IEC 60559.

Note that for some operations the exceptional value **invalid** is produced only for argument values involving -0 , $+\infty$, $-\infty$, or **sNaN**. For these operations the signature given in LIA-2 does not contain **invalid**.

A report ([57]) issued by the ANSI X3J11 committee discusses possible ways of exploiting the IEC 60559 special values. The report identifies some of its suggestions as controversial and cites [53] as justification.

In the following paragraphs summarise the specifications of IEC 60559 on the creation and propagation of signed zeros, infinities, and **NaNs**. There is also some discussion of the material in [53, 54, 51].

IEC 60559 regards 0 and -0 as almost indistinguishable. The sign is supposed to indicate the direction of approach to zero. The sign is reliable for a zero generated by underflow in a multiplication or division operation, and should be reliable also for operations that approximate elementary transcendental functions (see the LIA-2 specifications in clause 5.3). It is not reliable for a zero generated by an implied subtraction of two floating point numbers with the same value, for which case the zero is arbitrarily given a $+$ sign. The phrase “implied subtraction” indicates either the addition of two oppositely signed numbers or the subtraction of two like signed numbers.

On occurrence of floating overflow or division of a non-zero number by zero, an implementation conforming to IEC 60559 sets the appropriate status flag (if trapping is not enabled) and then continues execution with a result of $+\infty$ or $-\infty$ if rounding is to nearest. Infinities as such do *not* indicate that an overflow or division by zero has occurred; infinities can be exact values. IEC 60559 states that the arithmetic of infinities is that associated with mathematical infinities. Thus, an infinity times, plus, minus, or divided by a non-zero finite floating point number yields an

infinity for the result; no status flag is set and execution continues. These rules are not necessarily valid for infinities generated by overflow, though they are valid if the infinitary arguments are exact.

NaNs are generated by invalid operations on infinities, $0/0$, and the square root of a negative number (other than -0). Thus **NaNs** can represent unknown real or complex values, as well as totally undefined values. IEC 60559 requires that the result of any of its basic operations with one or more **NaN** arguments shall be a **NaN**. This principle is not extended to the numerical functions by [53, 57]. The controversial specifications in [57] are based on an assumption that all of these special operands represent finite non-zero real-valued numbers; see [53, 54].

The LIA-2 policy (for clauses 5.2 and 5.3) for dealing with signed zeros, infinities, and **NaNs** is as follows:

- a) The output is a quiet **NaN** for any operation for which one (or more) arguments is a quiet **NaN**, and none of the other arguments is a signalling NaN. There is then no notification.
- b) If a mathematical function $h(x)$ is such that $h(0) = 0$, the corresponding operation $op_F(x)$ returns x if $x \in \{0, -0\}$ and h has a positive derivative at 0, and $op_F(x)$ returns $neg_F(x)$ if $x \in \{0, -0\}$ and h has a negative derivative at 0.
- c) For an argument vector, \vec{x} , where that argument vector involves 0, -0 , $+\infty$, or $-\infty$, the result of the operation $op_F(\vec{x})$ is

$$\lim_{\vec{z} \rightarrow \vec{x}} h(\vec{z})$$

where an approach to zero is from the positive side if $\vec{x} = (\dots, 0, \dots)$, and the approach is from the negative side if $\vec{x} = (\dots, -0, \dots)$. There is no notification if the limit exists, is finite, and is path independent. The returned value is $+\infty$ or $-\infty$ if the limiting value is unbounded, and the approach is towards a point infinitely far from the origin. The returned value is **pole**($+\infty$) or **pole**($-\infty$) if the limiting value is unbounded, and the approach is towards a finite point. The result is -0 if the limit is zero and the approaching values are path independently negative. The result is 0 if the limit is zero and the approaching values are not path independently negative. If a path independent limit does not exist the value returned is **invalid**, and a notification occurs, with a continuation value of **qNaN** if appropriate.

An exception is made for the arc_F and $arcu_F$ operations, where it is found significantly more useful to return certain non-exceptional values for the origin and for the four double infinity argument cases, than to return an exceptional value, even with non-NaN continuation values.

B.5.2.1 The rounding and floating point *result* helper functions

The $result_F$ helper function notifies overflow when the result is too large to be approximated by a value in F . The $result_F$ helper function notifies underflow when there is (risk for) denormalisation loss for a tiny result. The $result_F$ helper function also ensures that a properly signed zero is the continuation value when a zero is appropriate. When an overflow or underflow occurs, and recording of indicators is the method for handling (floating point) overflow or underflow, a continuation value must be provided. LIA-2 specifies a continuation value, and if that can be represented in the target datatype, that value should be used as continuation value.

B.5.2.2 Floating point maximum and minimum

As for the integer case, the maximum and minimum of empty sequences need be handled, but for floating point datatypes, infinities are usually available.

For floating point one also usually have negative zero available, and returning the correct sign on a zero result for the maximum and minimum operations requires more than simple comparisons to implement. The sign of zeroes may need to be inspected using *copysign* or *isnegativezero*.

B.5.2.3 Floating point diminish

As for the integer case, this operation computes the positive difference.

B.5.2.4 Round, floor, and ceiling

Since $fmax_F$ always has an integral value according to LIA-1, no overflow can occur for these operations.

Note that the sign of a zero result is maintained in accordance with IEC 60559:

$$\begin{aligned} floor_F(x) &= neg_F(ceiling_F(neg_F(x))) \\ ceiling_F(x) &= neg_F(floor_F(neg_F(x))) \\ rounding_F(x) &= neg_F(rounding_F(neg_F(x))) \end{aligned}$$

Negative zeroes, if available, are handled in such a way as to maintain these identities.

B.5.2.5 Remainder after division and round to integer

The remainder after division and unbiased round to integer (IEC 60559 remainder, or IEEE remainder) is an exact operation, even if the floating point datatype only conforms to LIA-1, but not to the more specific IEC 60559.

Remainder after floating point division and floor to integer cannot be exact. For a small negative nominator and a positive denominator, the resulting value loses much absolute accuracy in relation to the original value. Such an operation is therefore not included in LIA-2. Similarly for floating point division and ceiling.

See also the radian normalisation and the argument angular-unit normalisation operations (5.3.9.1, 5.3.10.1).

B.5.2.6 Square root and reciprocal square root

\sqrt{x} cannot be exactly halfway between two values in F if $x \in F$. For \sqrt{x} to be exactly halfway between two values in F would require that it had exactly $(p + 1)$ digits (last digit non-zero) for its exact representation. The square of such a number would require at least $(2 \cdot p + 1)$ digits with last $p + 1$ digits not all zero, which could not equal the p -digit number x .

The extensions $sqrt_F(+\infty) = +\infty$ and $sqrt_F(-0) = -0$ are mandated by IEC 60559. LIA-2 also requires that these hold for implementations which support infinities and signed zeros. However, it should be noted that while the second is harmless, the first may lead to erroneous results for a $+\infty$ generated by an addition or subtraction with result just barely outside of $[-fmax_F, fmax_F]$ after rounding. Hence its square root would be well within the representable range. The possibility that LIA-2 should require that $sqrt_F(+\infty) = \text{invalid}(+\infty)$ was considered, but rejected because of the principle of regarding arguments as exact, even if they are not exact, when there is a non-degenerate neighbourhood around the argument point, for which the

mathematical function on \mathcal{R} is defined. In addition $sqrt_F(+\infty) = +\infty$ is already required by IEC 60559.

Note that the requirement that $sqrt_F(x) = \mathbf{invalid}(qNaN)$ for x strictly less than zero is mandated by IEC 60559. It follows that **NaNs** generated in this way represent imaginary values, which would become complex through addition and subtraction, and even imaginary infinities on multiplication by ordinary infinities.

The $rsqrt_F$ operation will increase performance for scaling a vector into a unit vector. Such an operation involves division of each component of the vector by the magnitude of the vector or, equivalently and with higher performance, multiplication by the reciprocal of the magnitude.

B.5.2.7 Support operations for extended floating point precision

These operations would typically be used to extend the precision of the highest level floating point datatype supported by the underlying hardware of an implementation. There is, however, no intent to provide a set of operations suitable for the implementation of a *complete* package for the support of calculations at an arbitrarily high level of precision.

The major motivation for including them in LIA-2 is to provide a capability for accurately evaluating residuals in an iterative algorithm. The residuals give a measure of the error in the current solution. More importantly they can be used to estimate a correction to the current solution. The accuracy of the correction depends on the accuracy of the residuals. The residuals are calculated as a difference in which the number of leading digits cancelled increases as the accuracy of the solution increases. A doubled precision calculation of the residuals is usually adequate to produce a reasonably efficient iteration.

For the basic floating point arithmetic doubled precision operations, the high parts may be calculated by the corresponding floating point operations as specified in LIA-1. Note, however, that in order to implement exact floating point addition and subtraction, rnd_F must round to nearest. If $add_F(x, y)$ rounds to *nearest* then the high and low parts represent $x + y$ exactly.

When the high parts of an addition or subtraction overflows, the low parts, as specified by LIA-2, return their results as if there was no overflow.

The product of two numbers, each with p digits of precision, is always exactly representable in at most $2 \cdot p$ digits. The high and low parts of the product will always represent the true product.

The remainder for division is more useful than a $2 \cdot p$ -digit approximation. The remainder will be exactly representable if the high part differs from the true quotient by less than one *ulp*. The true quotient can be constructed p digits at a time by division of the successive remainders by the divisor.

The remainder for square root is more useful than a low part for the same reason that the remainder is more useful for division. The remainder for the square root operation will be exactly representable only if the high part is correctly rounded to nearest, as is required by the specification for $sqrt_F$.

See *Semantics for Exact Floating Point Operations* [63] for more information on exact floating point operations.

See *Proposal for Accurate Floating-Point Vector Arithmetic* [64] for more information on exact, or high accuracy, floating point summation and dot product. These operations may be the subject of an amendment to LIA-2.

B.5.3 Elementary transcendental floating point operations

B.5.3.1 Maximum error requirements

$max_error_op_F$ measures the discrepancy between the computed value $op_F(x)$ and the true mathematical value $f(x)$ in ulps of the true value. The magnitude of the error bound is thus available to a program from the computed value $op_F(x)$. Note that for results at an exponent boundary for F , y , the error away from zero is in terms of $ulp_F(y)$, whereas the error toward zero is in terms of $ulp_F(y)/r_F$, which is the ulp of values slightly smaller in magnitude than y .

Within limits, accuracy and performance may be varied to best meet customer needs. Note also that LIA-2 does not prevent a vendor from offering two or more implementations of the various operations.

The operation specifications define the domain and range for the operations. The computational domain and range are more limited for the operations than for the corresponding mathematical functions because the arithmetic datatypes are subsets of \mathcal{R} . Thus the actual domain of $exp_F(x)$ is approximately given by $x \leq \ln(fmax_F)$. For larger values of x , $exp_F(x)$ will overflow, though for $x = +\infty$ the exact result $+\infty$ will be returned. The actual range extends over F , although there are non-negative values, $v \in F$, for which there is no $x \in F$ satisfying $exp_F(x) = v$.

The thresholds for the **overflow** and **underflow** notifications are determined by the parameters defining the arithmetic datatypes. The threshold for an **invalid** notification is determined by the domain of arguments for which the mathematical function being approximated is defined. The **pole** notification is the operation's counterpart of a mathematical pole of the mathematical function being approximated by the operation. The threshold for **absolute_precision_underflow** is determined by the parameters $big_angle_r_F$ and $big_angle_u_F$.

LIA-2 imposes a fairly tight bound on the maximum error allowed in the implementation of each operation. The tightest possible bound is given by requiring rounding to nearest, for which the accompanying performance penalty is often unacceptably high for the operations approximating elementary transcendental functions. LIA-2 does not require round to nearest for such operations, but allows for a slightly wider error bound characterised via the $max_error_op_F$ parameters. The parameters $max_error_op_F$ must be documented by the implementation for each such parameter required by LIA-2. A comparison of the values of these parameters with the values of the specified maximum value for each such parameter will give some indication of the "quality" of the routines provided. Further, a comparison of the values of this parameter for two versions of a frequently used operation will give some indication of the accuracy sacrifice made in order to gain performance.

Language bindings are free to modify the error limits provided in the specifications for the operations to meet the expected requirements of their users.

Material on the implementation of high accuracy operations is provided in for example [51, 53, 60].

B.5.3.2 Sign requirements

The requirements imply that the sign of the result or continuation value is to be reliable, except for the sign of an infinite result or continuation value, where except for a signed zero argument, it is often the case that one cannot determine the sign of the infinity. Still for sign symmetric mathematical functions, the approximating operation is also sign symmetric, including infinitary results.

B.5.3.3 Monotonicity requirements

A maximum error of 0.5 ulp implies that an approximation helper function must be a monotonic approximation to the mathematical function. When the maximum error is greater than 0.5 ulp, and the rounding is not directed, this is not automatically the case.

There is no general requirement that the approximation helper functions are *strictly* monotone on the same intervals on which the corresponding exact function is strictly monotone, however, since such a requirement cannot be made due to the fact that all floating point types are discrete, not continuous.

B.5.3.4 The *trans_result* helper function

B.5.3.5 Hypotenuse

The $hypot_F$ operation can produce an overflow only if both arguments have magnitudes very close to the overflow threshold. Care must be taken in its implementation to either avoid or properly handle overflows and underflows which might occur in squaring the arguments. The function approximated by this operation is mathematically equivalent to complex absolute value, which is needed in the calculation of the modulus and argument of a complex number. It is important for this application that an implementation satisfy the constraint on the magnitude of the result returned.

LIA-2 does not follow the recommendations in [53] and in [54] that

$$hypot_F(+\infty, \mathbf{qNaN}) = +\infty$$

$$hypot_F(-\infty, \mathbf{qNaN}) = +\infty$$

$$hypot_F(\mathbf{qNaN}, +\infty) = +\infty$$

$$hypot_F(\mathbf{qNaN}, -\infty) = +\infty$$

which are based on the claim that a \mathbf{qNaN} represents an (unknown) real valued number. This claim is not always valid, though it may sometimes be.

B.5.3.6 Operations for exponentiations and logarithms

For all of the exponentiation operations, overflow occurs for sufficiently large values of the argument(s).

There is a problem for $power_F(x, y)$ if both x and y are zero:

- Ada raises an ‘exception’ for the operation that is close in semantics to $power_F$ when both arguments are zero, in accordance with the fact that 0^0 is mathematically undefined.
- The X/OPEN Portability Guide, as well as C9x, specifies for $\mathbf{pow}(0, 0)$ a return value of 1, and no notification. This specification agrees with the recommendations in [51, 53, 54, 57].

The specification in LIA-2 follows Ada, and returns **invalid** for $power_F(0, 0)$, because of the risks inherent in returning a result which might be inappropriate for the application at hand. Note however, that $power_{FI}(0, 0)$ is 1, without any notification. The reason is that the limiting value for the corresponding mathematical function, when following either of the only two continuous paths, is 1. This also agrees with the Ada specification for a floating point value raised to a power in an integer datatype, as well as that for other programming languages which distinguish these operations.

Along any path defined by $y = k/\ln(x)$ the mathematical function x^y has the value e^k . It follows that some of the limiting values for x^y depend on the choice of k , and hence are undefined, as indicated in the specification.

The result of the $power_F$ operation is **invalid** for negative values of the base x . The reason is that the floating point exponent y might imply an implicit extraction of an even root of x , which would have a complex value for negative x . This constraint is explicit in Ada, and is widely imposed in existing numerical packages provided by vendors, as well as several other programming languages.

The arguments of $power_F$ are floating point numbers. No special treatment is provided for integer floating point values, which may be approximate. The cases for integer values of the arguments are covered by the operations $power_{FI}$ and $power_I$. In the example binding for C a specification for pow_F is supplied. pow_F combines $power_F$ and $power_{FZ}$ in a way suitable for C's **pow** operation.

For implementations of the $power_F$ operation there is an accuracy problem with an algorithm based on the following, mathematically valid, identity:

$$x^y = r_F^{y \cdot \log_{r_F}(x)}$$

The integer part of the product $y \cdot \log_{r_F}(x)$ defines the exponent of the result and the fractional part defines the reduced argument. If the exponent is large, and one calculates p_F digits of this intermediate result, there will be fewer than p_F digits for the fraction. Thus, in order to obtain a reduced argument accurately rounded to p_F digits, it may be necessary to calculate an approximation to $y \cdot \log_{r_F}(x)$ to a few more than $\log_{r_F}(emax_F) + p_F$ base r_F digits.

In Ada95 the operation most close to $power_{FI}$ is specified to be computed by successive multiplications, for which the error in the evaluation increases linearly with the size of the exponent. In a strict Ada implementation there is no way that a prescribed error limit of a few ulps can be met for large exponents.

The special exponentiation operations, corresponding to 2^x and 10^x , have specifications which are minor variations on those for $exp_F(x)$. Accuracy and performance can be increased if they are specially coded, rather than evaluated as, e.g., $exp_F(mul_F(x, ln_F(2)))$ or $power_F(2, x)$. Similar comments hold for the base 2 and base 10 logarithmic operations.

The $expm1_F$ operation has two advantages: Firstly, $expm1_F(x)$ is much more accurate than $sub_F(exp_F(x), 1)$ when the exponent argument is close to zero. Secondly, the $expm1_F$ operation does not **underflow** for “very” negative exponent arguments, something which may be advantageous if **underflow** handling is slow, and high accuracy for “very” negative arguments is not needed. Note in addition that underflow is avoided for this operation. This can be done only since LIA-2 adds requirements beyond those of LIA-1 regarding minimum precision (see clause 4). If those extra requirements were not done, underflow would not be justifiably removable for this operation. Similar argumentation applies to $ln1p_F$.

Similarly, there are two advantages with the $power1pm1_F$ operation: Firstly, $power1pm1_F(b, x)$ is much more accurate than $sub_F(power_F(add_F(1, b), x), 1)$ when the exponent argument is close to zero. Secondly, the $power1pm1_F$ operation does not **underflow** for “very” negative exponent arguments (when the base is greater than 1), something which may be advantageous if **underflow** handling is slow, and high accuracy for “very” negative arguments is not needed.

The handling of infinities and negative zero as arguments to the exponentiation and logarithm operations, like for all other LIA operations, follow the principles for dealing with these values as explained in section B.5.2.

B.5.3.7 Operations for hyperbolic elementary functions

The hyperbolic sine operation, $sinh_F(x)$, will overflow if $|x|$ is in the immediate neighbourhood of $\ln(2 * fmax)$, or greater.

The hyperbolic cosine operation, $\cosh_F(x)$, will overflow if $|x|$ is in the immediate neighbourhood of $\ln(2 * fmax)$, or greater.

The hyperbolic cotangent operation, $\coth_F(x)$, has a pole at $x = 0$.

The inverse of cosh is double valued, the two possible results having the same magnitude with opposite signs. The value returned by $\operatorname{arccosh}_F$ is always greater than or equal to 1.

The inverse hyperbolic tangent operation $\operatorname{arctanh}_F(x)$ has poles at $x = +1$ and at $x = -1$.

The inverse hyperbolic cotangent operation $\operatorname{arccoth}_F(x)$ has poles at $x = +1$ and at $x = -1$.

B.5.3.8 Introduction to operations for trigonometric elementary functions

The real trigonometric functions sin, cos, tan, cot, sec, and csc are all periodic. The period for sin, cos, sec, and csc is $2 \cdot \pi$ radians (360 degrees). The period for tan and cot is π radians (180 degrees). The mathematical trigonometric functions are perfectly periodic. Their numerical counterparts are not that perfect, for two reasons.

Firstly, the radian normalisation cannot be exact, even though it can be made very good given very many digits for the approximation(s) of π used in the angle normalisation, returning an offset from the nearest axis, and including guard digits. The unit argument normalisation, however, can be made exact regardless of the (non-zero and, in case $denorm_F = \mathbf{false}$, not too small) unit and the original angle, returning only a plain angle in F . LIA-2 requires unit argument angle normalisation to be exact.

Secondly, the length of one revolution is of course constant, but the density of floating point values gets sparser (in absolute spacing rather than relative) the larger the magnitude of the values are. This means that the number of floating point values gets sparser per revolution the larger the magnitude of the angle value is. For this reason the notification **absolute_precision_underflow** is introduced, together with two parameters (one for radians and one for other angular units). This notification is given when the magnitude of the angle value is “too big”. Exactly when the representable angle values get too sparse depends upon the application at hand, but LIA-2 gives a default value for the parameters that define the cut-off. LIA-2 also includes specifications for high accuracy angle normalisation operations. The angle normalisation operations give a result within minus half a cycle to plus half a cycle, unless the argument angular value is too big (or there is some other error).

Note that the **absolute_precision_underflow** notification is unrelated to any argument reduction problems. Argument reduction is (implicitly for radians, explicitly for other angular units) required by LIA-2 to be very accurate. But no matter how accurate the argument reduction is, floating point values are still sparser in absolute terms the larger the values are. Note also that any use of trigonometric operations for non-trigonometric purposes is out of scope for LIA-2.

LIA-2 includes angle normalisation operations, both for radians and for other angular units. The angle normalisation operations return a value within minus half a cycle and plus half a cycle. These operations should be used to keep the representation of angles at a high accuracy. The trigonometric operations return a result within about an ulp, and that high accuracy is wasted if the angular argument is not kept at a high accuracy too. LIA-2 also includes angle normalisation operations that can be used to maintain an even higher degree of accuracy, giving the offset from the nearest axis (though without any extra guard digits). To use these one needs to keep track of the nearest axis, which unfortunately complicates programs that use this latter method.

Note that $rad(x) = \arccos(\cos(x))$ if $\sin(x) > 0$ and $rad(x) = -\arccos(\cos(x))$ if $\sin(x) < 0$. The first part of $axis_rad(x)$ indicates which axis is nearest to the angle x . The second part of $axis_rad(x)$ is an angle offset from the axis that is nearest to the angle x . The second part of

$axis_rad(x)$ is equal to $rad(x)$ if $\cos(x) \geq 1/\sqrt{2}$ (i.e. if the first part of $axis_rad(x)$ is $(1, 0)$). More generally, the second part of $axis_rad(x)$ is equal to $rad(4 \cdot x)/4$.

$rad(x)$ returns the same angle as the angle value x , but the returned angle value is between $-\pi$ and π . The rad function is defined to be used as the basis for the angle normalisation operations. The $axis_rad$ function is defined to be used as the basis for a numerically more accurate radian angle normalisation operation. The arc function is defined to be used as the basis for the arcus (angle) operations, which are used for conversion from Cartesian to polar co-ordinates.

B.5.3.9 Operations for radian trigonometric elementary functions

The radian trigonometric approximation helper functions (including those for normalisation and conversion from radians) are required to have the same zero points as the approximated mathematical function only if the absolute value of the argument is less than or equal to $big_angle_r_F$. Likewise, the radian trigonometric approximation helper functions are required to have the same sign as the approximated mathematical function only if the absolute value of the argument is less than or equal to $big_angle_r_F$.

The $big_angle_r_F$ parameter may be adjusted by bindings, or even by some compiler flag, or mode setting within a program. However, this method should only allow the value of this parameter to be set to a value greater than $2 \cdot \pi$, so that at least arguments within the first two (plus and minus) cycles are allowed, and such that $ulp_F(big_angle_r_F) < \pi/1000$, so that at least 2000 evenly distributed points within the ‘last’ cycle (farthest away from 0) are distinguishable. The latter gives a rather low accuracy at the far ends of the range, especially if p_F is comparatively large, so values this large for $big_angle_r_F$ are not recommendable unless the application is such that high accuracy trigonometric operations are not needed.

For reduction of an argument given in radians, implementations use one or several approximate value(s) of π (or of a multiple of π), valid to, say, n digits. The division implied in the argument reduction cannot be valid to more than n digits, which implies a maximum absolute angle value for which the reduction yields an accurate reduced angle value.

Regarding argument reduction for radians, there is a particular problem when the result of the trigonometric operation is very small (or very big), but the angular argument is not very small. In such cases the argument reduction must be very accurate, using an extra-precise approximation to π , relative to what is normally used for arguments of similar magnitude, so that significant digits in the result are not lost. Such loss would imply non-conformance to LIA-2 by the error in the final result being greater than that specified by LIA-2. In general, extra care has to be taken when the second part of $axis_rad(x)$ is close to 0.

Note that if $big_angle_r_F$ is allowed to be increased, then, for conformity with LIA-2, the radian angle reduction may need to be more precise.

- \tan and \sec have *poles* at odd multiples of $\pi/2$ radians (90 degrees).
- \cot and \csc have *poles* at multiples of π radians (180 degrees).

All four of the corresponding operations with poles may produce **overflow** for arguments sufficiently close to the poles of the functions. The \tan_F operation produces no **pole** notification. The reason is that the poles of $\tan(x)$ are at odd multiples of $\pi/2$, which are not representable in F . The mathematical cotangent function has a pole at the origin. For a system which supports signed zeros and infinities, the continuation values are $+\infty$ and $-\infty$ for arguments of 0 and -0 respectively. Although the mathematical function \sec has poles at odd multiples of $\pi/2$, the \sec_F operation will not generate any pole notification because such arguments are not representable in F .

The **pole** notification cannot occur for any non-zero argument in radians because π is not representable in F , nor is $\pi/2$. For the angular unit argument trigonometric operations a the sign of the infinitary continuation value has been chosen arbitrarily for a **pole** which occurs for a non-zero argument. However, sign symmetry, when appropriate, is maintained.

The operations may produce **underflow** for arguments sufficiently close to their zeros. For a denormalised argument x , the \sin_F , \tan_F , \arcsin_F , and \arctan_F return x for the result, with very high accuracy. Similarly, for a denormalised argument, \cos_F and \sec_F can return a result of 1.0 with very high accuracy.

The trigonometric inverses are multiple valued. They are rendered single valued by defining a principal value range. This range is closely related to a branch cut in the complex plane for the corresponding complex function. Among the floating point numerical functions this branch cut is “visible” only for the \arcc_F operation. The arc function has a branch cut along the negative real axis. For $x < 0$ the function has a discontinuity from $-\pi$ to $+\pi$ as y passes through zero from negative to positive values. Thus for $x < 0$, systems supporting signed zeros can handle the discontinuity as follows:

$$\begin{aligned} \arcc_F(x, -\mathbf{0}) &= up_F(-\pi) \\ \arcc_F(x, \mathbf{0}) &= down_F(\pi) \end{aligned}$$

There is a problem for zero argument values for this operation. The values given for the operation $\arcc_F(x, y)$ for the four combinations of signed zeros for x and y are those given in [53]. The following table of values is given in [53] for the value of $\arcc_F(x, y)$ with both of the arguments zero:

Zero arguments		
x	y	$\arcc_F(x, y)$
0	0	0
-0	0	π
-0	-0	$-\pi$
0	-0	-0

Note that the mathematical arc function is indeterminate (undefined) for (0,0), but these result are numerically more useful than giving an **invalid** notification for such arguments. LIA-2 therefore specifies results as above.

There is also a problem for argument values of $+\infty$ or $-\infty$ for this operation. The following table of values is given in [53] for the value of $\arcc_F(x, y)$ with at least one of the arguments infinite:

Infinite arguments		
x	y	$\arcc_F(x, y)$
$+\infty$	≥ 0	0
$+\infty$	$+\infty$	$\pi/4$
finite	$+\infty$	$\pi/2$
$-\infty$	$+\infty$	$3 \cdot \pi/4$
$-\infty$	≥ 0	π
$-\infty$	-0	$-\pi$
$-\infty$	< 0	$-\pi$
$-\infty$	$-\infty$	$-3 \cdot \pi/4$
finite	$-\infty$	$-\pi/2$
$+\infty$	$-\infty$	$-\pi/4$
$+\infty$	< 0	-0
$+\infty$	-0	-0

If one of x and y is infinite and the other is finite, the result tabulated is consistent with that obtained by a conventional limiting process. However, the results of $\pi/4$, $-\pi/4$, $3 \cdot \pi/4$, and $-3 \cdot \pi/4$ corresponding to infinite values for both x and y , are of questionable validity, since only the quadrant is known, not the angle within the quadrant. However, these results are numerically more useful than giving an **invalid** notification for such arguments. LIA-2 therefore specifies results as above.

B.5.3.10 Operations for trigonometrics given angular unit

At present only Ada specifies trigonometric operations with angular unit argument. LIA-2 has adopted angular unit argument operations in order to encourage uniformity among languages which might include such operations in the future. The angular units in T appear to be particularly important and have therefore been given a tighter error bound requirement. An implementation can of course have the same (tighter) error bound for all angular units.

The *big_angle_u_F* parameter may be adjusted by bindings, or even by some compiler flag, or mode setting within a program. However, this method should only allow the value of this parameter to be set to a value greater than or equal to 1, so that at least arguments within the first two (plus and minus) cycles are allowed, and such that $ulp_F(\text{big_angle_u}_F) \leq 1/2000$, so that at least 2000 evenly distributed points within the ‘last’ cycle (farthest away from 0) are distinguishable. The latter gives a rather low accuracy at the far ends of the range, especially if p_F is comparatively large, so values this large for *big_angle_u_F* are not recommendable unless the application is such that high accuracy trigonometric operations are not needed.

The *min_angular_unit_F* parameter is specified for two reasons. Firstly, if the type F has no denormal values (*denorm_F* = **false**), some angle values in F are not representable after normalisation if the angular unit has too small magnitude (this gives the firm limit above). Secondly, even if F has denormal values (*denorm_F* = **true**), angular units with very small magnitude do not allow the representable angles to be particularly dense, not even if the angular value is within the first cycle. This does in itself not give rise to a particular limit value, but the limit value defined here is reasonable.

All of the argument angular unit trigonometric, and argument angular unit inverse trigonometric, approximation helper functions, including those for normalisation, angular unit conversion, and arc, are exempted from the monotonicity requirement for the angular unit argument.

If the angular unit argument, u , is such that $u/4 \in F$, the *tanu_F* operation has poles at odd multiples of $u/4$. This is the case for degrees ($u = 360$). As for *tanu_F*, if the angular unit argument, u , is such that $u/4 \in F$ the *secu_F* operation has poles at odd multiples of $u/4$.

The same comments hold for the *arcu_F* operation as for *arc_F* operation, except that the discontinuity in the mathematical function is from $-u/2$ to $+u/2$.

B.5.3.11 Operations for angular-unit conversions

Conversion of an angular value x from angular unit a to angular unit b appears simple: compute $x \cdot b/a$. Basing a numerical conversion of angular values directly on the above mathematical equality (e.g. *div_F(mul_F(x, b), a))* loses much absolute angular accuracy, however, especially for large angular values. Instead computing *arcu_F($b, \text{cosu}_F(a, x), \text{sinu}_F(a, x)$)* gives a more accurate result. This might still not be within the accuracy required by LIA-2 for the angular unit conversion operations specified by LIA-2.

Note that all of the angular conversion operations return an angularly normalised result. This is in order to maintain high accuracy of the angle value being represented.

Angular conversion operations are commonly found on ‘scientific’ calculators and also in Java, though then only between degrees and radians.

B.5.4 Conversion operations

Clause 5.2 of LIA-1 covers conversions from an integer type to another integer type and to a floating point type, as well as between (LIA-1 conforming) floating point types of the same radix.

LIA-2 extends these conversions to cover conversions to and from non-LIA conforming datatypes, such as conversion to and from strings, and also extends the floating point conversion specifications to also handle conversions where the radices are different.

In ordinary string formats for numerals, the string “Hello world!” is an example of a signalling NaN.

LIA-2 does not specify any string formats, not even for the special values -0 , $+\infty$, $-\infty$, and quiet NaN, but possibilities for the special values include the strings used in the text of LIA-2, as well as strings like “+infinity” or “positiva oändligheten”, etc, and the strings used may depend on preference settings, as they may also for non-special values. E.g. one may use different notation for the decimal separator character (e.g., period, comma, Arabic comma, ...), use superscript digits for exponents, or use Arabic or traditional Thai digits. String formats for numerical values, and if and how they may depend on preference settings, is also an issue for bindings or programming language specifications, not for this part of LIA.

If the value converted is greater than those representable in the target, or less than those representable in the target, even after rounding, then an overflow will result. E.g., if the target is a character string of at most 3 digits, and the target radix is 10, then an integer source value of 1000 will result in an overflow. As for other operations, if the notification handling is by recording of indicators, a suitable continuation value must be used.

Most language standards contain (partial) format specifications for conversion to and from strings, usually for a decimal representation.

B.5.5 Numerals as operations in the programming language

B.5.5.1 Numerals for integer datatypes

Negative values (except $minint_I$ if $minint_I = -maxint_I - 1$) can be obtained by using the negation operation (neg_I).

Integer numerals in radix 10 are normally available in programming languages. Other radices may also be available for integer numerals, and the radix used may be part of determining the target integer datatype. E.g., radix 10 may be for signed integer datatypes, and radix 8 or 16 may be for unsigned integer datatypes.

Syntaxes for numerals for different integer datatypes need not be different, nor need they be the same. This Part does not further specify the format for integer numerals. That is an issue for bindings.

Overflow for integer numerals can be detected at “compile time”, and warned about.

B.5.5.2 Numerals for floating point datatypes

If the numerals used as operations in a program, and numerals read from other sources use the same radix, then “internal” numerals and “external” numerals (strings) denoting the same value in \mathcal{R} and converted to the same target datatype should be converted to the same value.

Negative values (including negative 0, -0) can be obtained by using the negation operation (neg_F).

Radices other than 10 may also be available for floating point numerals.

Integer numerals may also be floating point numerals, i.e. their syntaxes need not be different. Nor need syntaxes for numerals for different floating point datatypes be different, nor need they be the same. This Part does not specify the syntax for numerals. That is an issue for bindings or programming language specifications.

Overflow/underflow for floating point numerals can be detected at “compile time”, and warned about.

B.6 Notification

An intermediate overflow on computing approximations to x^2 or y^2 during the calculation of $hypot_F(x, y) \approx \sqrt{x^2 + y^2}$ must not result in an overflow notification, unless the end result overflows. This is clear from the specification of the $hypot_F$ operation in this Part.

If a single argument operation op_F , for the corresponding mathematical function f , is such that $f(x)$ very closely approximates x , when $|x| \leq fminN_F$, then $op_F(x)$ returns x for $|x| \leq fminN_F$, and does not give a notification if there cannot be any denormalisation loss relative to $f(x)$. For details, see the individual operation specifications for $expm1_F$, $ln1p_F$, $sinh_F$, $arcsinh_F$, $tanh_F$, $arctanh_F$, sin_F , $arcsin_F$, tan_F , and $arctan_F$.

Operations specified in LIA-2 return **invalid(qNaN)** when passed a signaling NaN (sNaN) as an argument. Most operations specified in LIA-2 return **qNaN**, without any notification when passed a quiet NaN (qNaN) as an argument.

The different kinds of notifications occur under the following circumstances:

- a) **invalid**: when an argument is not valid for the operation, and no value in F^* or any special value result makes mathematical sense.
- b) **pole**: when the input operand corresponds to a pole of the mathematical function approximated by the operation.
- c) **overflow**: when the (rounded) result is outside of the range of the result datatype.
- d) **underflow**: when a sufficiently closely approximating result of the operation has a magnitude that is so small that it might not be sufficiently accurately represented in the result datatype.
- e) **absolute_precision_underflow**: when the magnitude of the angle argument to a floating point trigonometric operation exceeds the maximum value of the argument for which the density of floating point values is deemed sufficient for the operation to make sense. See clause 5.3.8 and the associated discussion in this rationale (section B.5.3.8).

In order to avoid **absolute_precision_underflow** notifications, and to maintain a high accuracy, implementors are encouraged to provide, and programmers encouraged to use, the angle normalisation operations specified in 5.3.9.1 and 5.3.10.1.

The difference between the **pole** and **overflow** notifications is that the first corresponds to a true mathematical singularity, and the second corresponds to a well-defined mathematical result that happens to lie outside the range of F .

B.6.1 Continuation values

For handling of notifications, the method that does recording of indicators (LIA-1, clause 6.1.2) is preferred.

An implementation which supports recording of indicators must supply continuation values to be used when execution is continued following the occurrence of a notification. For systems which support signed zeros, infinities and NaNs, LIA-2 specifies how these values, as well as ordinary values, are used as continuation values. Other implementations which use recording of indicators must supply other suitable continuation values and document the values selected.

B.7 Relationship with language standards

An arithmetic expression might not be executed as written.

For example, if x is declared to be single precision (SP) floating point, and calculation is done in single precision, then the expression

$$\arcsin(\mathbf{x})$$

might translate to

$$\arcsin_{SP}(\mathbf{x})$$

If the language in question did all computations in double precision (DP) floating point, the above expression might translate to

$$\arcsin_{DP}(cvt_{SP \rightarrow DP}(x))$$

Alternatively, if \mathbf{x} was declared to be an integer, and the expected result datatype is single precision float, the above expression might translate to

$$cvt_{DP \rightarrow SP}(\arcsin_{DP}(cvt_{I \rightarrow DP}(x)))$$

The datatypes involved in implicit conversions need not be accessible to the programmer. For example, trigonometric operations may be evaluated in extended double precision, even though that datatype is not made available to programmers using a particular programming language. These extra datatypes should be made available, however, and the implicit conversions should be expressible as explicit conversions. At least in order to be able to show exactly which expression is going to be evaluated without having to look at the machine code.

B.8 Documentation requirements

Annex C

(informative)

Example bindings for specific languages

This annex describes how a computing system can simultaneously conform to a language standard (or publicly available specification) and to LIA-2. It contains suggestions for binding the “abstract” operations specified in LIA-2 to concrete language syntax. The format used for these example bindings in this annex is a short form version, suitable for the purposes of this annex. An actual binding is under no obligation to follow this format. An actual binding should, however, as in the bindings examples, give the LIA-2 operation name, or parameter name, bound to an identifier by the binding.

Portability of programs can be improved if two conforming LIA-2 systems using the same language agree in the manner with which they adhere to LIA-2. For instance, LIA-2 requires that the parameter *big_angle_rF* be provided (if any conforming radian trigonometric operations are provided), but if one system provides it by means of the identifier **BigAngle** and another by the identifier **MaxAngle**, portability is impaired. Clearly, it would be best if such names were defined in the relevant language standards or binding standards, but in the meantime, suggestions are given here to aid portability.

The following clauses are suggestions rather than requirements because the areas covered are the responsibility of the various language standards committees. Until binding standards are in place, implementors can promote “de facto” portability by following these suggestions on their own.

The languages covered in this annex are

- Ada
- Basic
- C
- C++
- Fortran
- Haskell
- Java
- Common Lisp
- ISLisp
- Modula-2
- Pascal and Extended Pascal
- PL/I
- SML

This list is not exhaustive. Other languages and other computing devices (like ‘scientific’ calculators, ‘web script’ languages, and database ‘query languages’) are suitable for conformity to LIA-2.

In this annex, the parameters, operations, and exception behaviour of each language are examined to see how closely they fit the requirements of LIA-2. Where parameters, constants, or operations are not provided by the language, names and syntax are suggested. (Already provided syntax is marked with a ★.)

This annex describes only the language-level support for LIA-2. An implementation that wishes to conform must ensure that the underlying hardware and software is also configured to conform to LIA-2 requirements.

A complete binding for LIA-2 will include, or refer to, a binding for LIA-1. In turn, a complete binding for the LIA-1 may include, or refer to, a binding for IEC 60559. Such a joint LIA-2/LIA-1/IEC 60559 binding should be developed as a single binding standard. To avoid conflict with ongoing development, only the LIA-2 specific portions of such a binding are exemplified in this annex.

C.1 General comments

Most language standards permit an implementation to provide, by some means, the parameters and operations required by LIA-2 that are not already part of the language. The method for accessing these additional parameters and operations depends on the implementation and language, and is not specified in LIA-2 nor exemplified in this annex. It could include external subroutine libraries; new intrinsic functions supported by the compiler; constants and functions provided as global “macros”; and so on. The actual method of access through libraries, macros, etc. should of course be given in a real binding.

Most language standards do not constrain the accuracy of elementary numerical functions, or specify the subsequent behaviour after an arithmetic notification occurs.

In the event that there is a conflict between the requirements of the language standard and the requirements of LIA-2, the language binding standard should clearly identify the conflict and state its resolution of the conflict.

C.2 Ada

The programming language Ada is defined by ISO/IEC 8652:1995, *Information Technology – Programming Languages – Ada* [11], where the specifications for the operations for elementary functions are based on ISO/IEC 11430:1994 *Information technology – Programming languages – Generic package of elementary functions for Ada* [12].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to LIA-2 for that operation or parameter. For each of the marked items a suggested identifier is provided.

The Ada datatype **Boolean** corresponds to the LIA datatype **Boolean**.

Every implementation of Ada has at least one integer datatype, and at least one floating point datatype. The notations *INT* and *FLT* are used to stand for the names of one of these datatypes in what follows.

The LIA-2 integer operations are listed below, along with the syntax used to invoke them:

$max_I(x, y)$	$INT'Max(x, y)$	★
$min_I(x, y)$	$INT'Min(x, y)$	★
$max_seq_I(xs)$	$Max(xs)$	†
$min_seq_I(xs)$	$Min(xs)$	†
$dim_I(x, y)$	$Dim(x, y)$	†
$power_I(x, y)$	$x ** y$	★
$shift2_I(x, y)$	$Shift2(x, y)$	†
$shift10_I(x, y)$	$Shift10(x, y)$	†
$sqr_I(x)$	$Sqrt(x)$	†

$divides_I(x, y)$	$Divides(x, y)$	†
$even_I(x)$	$x \bmod 2 = 0$	*
$odd_I(x)$	$x \bmod 2 \neq 0$	*
$divf_I(x, y)$	$Div(x, y)$	†
$moda_I(x, y)$	$x \bmod y$	*
$group_I(x, y)$	$Group(x, y)$	†
$pad_I(x, y)$	$Pad(x, y)$	†
$quot_I(x, y)$	$Quot(x, y)$	†
$remr_I(x, y)$	$Rem(x, y)$	†
$gcd_I(x, y)$	$Gcd(x, y)$	†
$lcm_I(x, y)$	$Lcm(x, y)$	†
$gcd_seq_I(xs)$	$Gcd(xs)$	†
$lcm_seq_I(xs)$	$Lcm(xs)$	†
$add_wrap_I(x, y)$	$Add_Wrap(x, y)$	†
$add_ov_I(x, y)$	$Add_Over(x, y)$	†
$sub_wrap_I(x, y)$	$Sub_Wrap(x, y)$	†
$sub_ov_I(x, y)$	$Sub_Over(x, y)$	†
$mul_wrap_I(x, y)$	$Mul_Wrap(x, y)$	†
$mul_ov_I(x, y)$	$Mul_Over(x, y)$	†

where x and y are expressions of type *INT* and where xs is an expression of type array of *INT*.

The LIA-2 basic floating point operations are listed below, along with the syntax used to invoke them:

$max_F(x, y)$	$FLT'Max(x, y)$	*
$min_F(x, y)$	$FLT'Min(x, y)$	*
$mmax_F(x, y)$	$MMax(x, y)$	†
$mmin_F(x, y)$	$MMin(x, y)$	†
$max_seq_F(xs)$	$Max(xs)$	†
$min_seq_F(xs)$	$Min(xs)$	†
$mmax_seq_F(xs)$	$MMax(xs)$	†
$mmin_seq_F(xs)$	$MMin(xs)$	†
$dim_F(x, y)$	$Dim(x, y)$	†
$rounding_F(x)$	$FLT'UnbiasedRounding(x)$	*
$floor_F(x)$	$FLT'Floor(x)$	*
$ceiling_F(x)$	$FLT'Ceiling(x)$	*
$rounding_rest_F(x)$	$x - FLT'UnbiasedRounding(x)$	*
$floor_rest_F(x)$	$x - FLT'Floor(x)$	*
$ceiling_rest_F(x)$	$x - FLT'Ceiling(x)$	*
$remr_F(x, y)$	$FLT'Remainder(x, y)$	*
$sqr_t_F(x)$	$Sqrt(x)$	*
$rsqr_t_F(x)$	$RSqrt(x)$	†
$add_lo_F(x, y)$	$Add_Low(x, y)$	†
$sub_lo_F(x, y)$	$Sub_Low(x, y)$	†
$mul_lo_F(x, y)$	$Mul_Low(x, y)$	†
$div_rest_F(x, y)$	$Div_Rest(x, y)$	†
$sqr_t_rest_F(x)$	$Sqrt_Rest(x)$	†
$mul_{F \rightarrow F'}(x, y)$	$Prod(x, y)$	†

where x , y , and z are expressions of type *FLT*, and where xs is an expression of type array of

FLT.

The parameters for LIA-2 operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	$Err_Hypotenuse(x)$	†
$max_err_exp_F$	$Err_Exp(x)$	†
$max_err_power_F$	$Err_Power(x)$	†
$max_err_sinh_F$	$Err_Sinh(x)$	†
$max_err_tanh_F$	$Err_Tanh(x)$	†
$big_angle_r_F$	$Big_Radian_Angle(x)$	†
$max_err_sin_F$	$Err_Sin(x)$	†
$max_err_tan_F$	$Err_Tan(x)$	†
$min_angular_unit_F$	$Smallest_Angular_Unit(x)$	†
$big_angle_u_F$	$Big_Angle(x)$	†
$max_err_sinu_F(u)$	$Err_Sin_Cycle(u)$	†
$max_err_tanu_F(u)$	$Err_Tan_Cycle(u)$	†
$max_err_convert_F$	$Err_Convert(x)$	†
$max_err_convert_{F'}$	$Err_Convert_To_String$	†
$max_err_convert_{D'}$	$Err_Convert_To_String$	†

where x and u are expressions of type *FLT*. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them:

$hypot_F(x, y)$	$Hypotenuse(x, y)$	†
$power_{FI}(b, z)$	$b ** z$	*
$exp_F(x)$	$Exp(x)$	*
$expm1_F(x)$	$ExpM1(x)$	†
$exp2_F(x)$	$Exp2(x)$	†
$exp10_F(x)$	$Exp10(x)$	†
$power_F(b, y)$	$b ** y$	*
$power1pm1_F(b, y)$	$Power1PM1(b, y)$	†
$ln_F(x)$	$Log(x)$	*
$ln1p_F(x)$	$Log1P(x)$	†
$log2_F(x)$	$Log2(x)$	†
$log10_F(x)$	$Log10(x)$	†
$logbase_F(b, x)$	$Log(x, b)$ (note parameter order)	*
$logbase1p1p_F(b, x)$	$Log1P1P(x, b)$	†
$sinh_F(x)$	$SinH(x)$	*
$cosh_F(x)$	$CosH(x)$	*
$tanh_F(x)$	$TanH(x)$	*
$coth_F(x)$	$CotH(x)$	*
$sech_F(x)$	$SecH(x)$	†

$csch_F(x)$	$Csch(x)$	†
$arcsinh_F(x)$	$ArcSinH(x)$	*
$arccosh_F(x)$	$ArcCosH(x)$	*
$arctanh_F(x)$	$ArcTanH(x)$	*
$arcoth_F(x)$	$ArcCoth(x)$	*
$arcsech_F(x)$	$ArcSecH(x)$	†
$arccsch_F(x)$	$ArcCsch(x)$	†
$rad_F(x)$	$Rad(x)$	†
$axis_rad_F(x)$	$Rad(x, h, v)$ (note out parameters)	†
$sin_F(x)$	$Sin(x)$	*
$cos_F(x)$	$Cos(x)$	*
$tan_F(x)$	$Tan(x)$	*
$cot_F(x)$	$Cot(x)$	*
$sec_F(x)$	$Sec(x)$	†
$csc_F(x)$	$Csc(x)$	†
$coassin_F(x)$	$CosSin(x, c, s)$ (note out parameters)	†
$arcsin_F(x)$	$ArcSin(x)$	*
$arccos_F(x)$	$ArcCos(x)$	*
$arctan_F(x)$	$ArcTan(x)$	*
$arccot_F(x)$	$ArcCot(x)$	*
$arctg_F(x)$	$ArcCtg(x)$	†
$arcsec_F(x)$	$ArcSec(x)$	†
$arccsc_F(x)$	$ArcCsc(x)$	†
$arc_F(x, y)$	$ArcTan(y, x)$ or $ArcCot(x, y)$	*
$cycle_F(u, x)$	$Cycle(x, u)$ (note parameter order)	†
$axis_cycle_F(u, x)$	$Cycle(x, u, h, v)$	†
$sinu_F(u, x)$	$Sin(x, u)$ (note parameter order)	*
$cosu_F(u, x)$	$Cos(x, u)$	*
$tanu_F(u, x)$	$Tan(x, u)$	*
$cotu_F(u, x)$	$Cot(x, u)$	*
$secu_F(u, x)$	$Sec(x, u)$	†
$cscu_F(u, x)$	$Csc(x, u)$	†
$coassinu_F(u, x)$	$CosSin(x, u, c, s)$	†
$arcsinu_F(u, x)$	$ArcSin(x, u)$	*
$arccosu_F(u, x)$	$ArcCos(x, u)$	*
$arctanu_F(u, x)$	$ArcTan(x, Cycle=>u)$	*
$arccotu_F(u, x)$	$ArcCot(x, Cycle=>u)$	*
$arctgu_F(u, x)$	$ArcCtg(x, u)$	†
$arcsecu_F(u, x)$	$ArcSec(x, u)$	†
$arccscu_F(u, x)$	$ArcCsc(x, u)$	†
$arcu_F(u, x, y)$	$ArcTan(y, x, u)$ or $ArcCot(x, y, u)$	*
$rad_to_cycle_F(x, u)$	$Rad_to_Cycle(x, u)$	†
$cycle_to_rad_F(u, x)$	$Cycle_to_Rad(u, x)$	†

<i>cycle_to_cycle_F</i> (<i>u, x, v</i>)	<i>Cycle_to_Cycle</i> (<i>u, x, v</i>)	†
--	--	---

where *b, x, y, u,* and *v* are expressions of type *FLT*, and *z* is an expressions of type *INT*.

Arithmetic value conversions in Ada are always explicit and usually use the destination datatype name as the name of the conversion function, except when converting to/from strings.

<i>convert_{I→I'}</i> (<i>x</i>)	<i>INT2</i> (<i>x</i>)	★
<i>convert_{I''→I}</i> (<i>s</i>)	<i>Get</i> (<i>s, n, w</i>)	★
<i>convert_{I→I''}</i> (<i>x</i>)	<i>Put</i> (<i>s, x, base?</i>)	★
<i>convert_{I''→I}</i> (<i>f</i>)	<i>Get</i> (<i>f?, n, w?</i>)	★
<i>convert_{I→I''}</i> (<i>x</i>)	<i>Put</i> (<i>h?, x, w?, base?</i>)	★
<i>rounding_{F→I}</i> (<i>y</i>)	<i>INT</i> (<i>FLT'Unbiased_Rounding</i> (<i>y</i>))	★
<i>floor_{F→I}</i> (<i>y</i>)	<i>INT</i> (<i>FLT'Floor</i> (<i>y</i>))	★
<i>ceiling_{F→I}</i> (<i>y</i>)	<i>INT</i> (<i>FLT'Ceiling</i> (<i>y</i>))	★
<i>convert_{I→F}</i> (<i>x</i>)	<i>FLT</i> (<i>x</i>)	★
<i>convert_{F→F'}</i> (<i>y</i>)	<i>FLT2</i> (<i>y</i>)	★
<i>convert_{F''→F}</i> (<i>s</i>)	<i>Get</i> (<i>s, n, w?</i>)	★
<i>convert_{F''→F}</i> (<i>f</i>)	<i>Get</i> (<i>f?, n, w?</i>)	★
<i>convert_{F→F''}</i> (<i>y</i>)	<i>Put</i> (<i>s, y, Aft=>a?, Exp=>e?</i>)	★
<i>convert_{F→F''}</i> (<i>y</i>)	<i>Put</i> (<i>h?, y, Fore=>i?, Aft=>a?, Exp=>e?</i>)	★
<i>convert_{D→F}</i> (<i>z</i>)	<i>FLT</i> (<i>z</i>)	★
<i>convert_{D'→F}</i> (<i>s</i>)	<i>Get</i> (<i>s, n, w?</i>)	★
<i>convert_{D'→F}</i> (<i>f</i>)	<i>Get</i> (<i>f?, n, w?</i>)	★
<i>convert_{F→D}</i> (<i>y</i>)	<i>FXD</i> (<i>y</i>)	★
<i>convert_{F→D'}</i> (<i>y</i>)	<i>Put</i> (<i>s, y, Aft=>a?, Exp=>0</i>)	★
<i>convert_{F→D'}</i> (<i>y</i>)	<i>Put</i> (<i>h?, y, Fore=>i?, Aft=>a?, Exp=>0</i>)	★

where *x* is an expression of type *INT*, *y* is an expression of type *FLT*, and *z* is an expression of type *FXD*, where *FXD* is a fixed point type. *INT2* is the integer datatype that corresponds to *I'*. A ? above indicates that the parameter is optional. *e* is greater than 0.

Ada provides non-negative numerals for all its integer and floating point types. The default base is 10, but all bases from 2 to 16 can be used. There is no differentiation between the numerals for different floating point types, nor between numerals for different integer types, but integer numerals (without a point) cannot be used for floating point types, and 'real' numerals (with a point) cannot be used for integer types. Integer numerals can have an exponent part though. The details are not repeated in this example binding, see ISO/IEC 8652:1995, clause 2.4 Numeric Literals, clause 3.5.4 Integer Types, and clause 3.5.6 Real Types.

The Ada standard does not specify any numerals for infinities and NaNs. Suggestion:

+∞	<i>FLT'Infinity</i>	†
qNaN	<i>FLT'NaN</i>	†
sNaN	<i>FLT'SigNaN</i>	†

as well as string formats for reading and writing these values as character strings.

Ada has a notion of 'exception' that implies a non-returnable, but catchable, change of control flow. Ada uses its exception mechanism as its default means of notification. Ada ignores **underflow** notifications since an Ada exception is inappropriate for an **underflow** notification. On **underflow** the continuation value (specified in LIA-2) is used directly without recording the **underflow** itself. Ada uses the exception *Constraint_Error* for **pole** and **overflow** notifications,

and the exception `Numerics.Argument_Error` for **invalid** notifications. Since Ada exceptions are non-returnable changes of control flow, no continuation value is provided for these notifications.

An implementation that wishes to follow LIA-2 should provide recording of indicators as an alternative means of handling numeric notifications. Recording of indicators is the LIA-2 preferred means of handling numeric notifications.

C.3 BASIC

The programming language BASIC is defined by ANSI X3.113-1987 (R1998) [41], endorsed by ISO/IEC 10279:1991, *Information technology – Programming languages – Full BASIC* [17].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

BASIC has no user accessible datatype corresponding to the LIA datatype **Boolean**.

BASIC has one primitive computational data type, **numeric**. The model presented by the BASIC language is that of a real number with decimal radix and a specified (minimum) number of significant decimal digits. Numeric data is not declared directly, but any special characteristics are inferred from how they are used and from any **OPTIONS** that are in force.

The BASIC statement **OPTION ARITHMETIC NATIVE** ties the **numeric** type more closely to the underlying implementation. The precision and type of **NATIVE** numeric data is implementation dependent.

Since the BASIC numeric data type does not match the integer type required by the LIA-1, this binding example does not include any of the LIA-2 operations for integer data types.

The LIA-2 non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$max_F(x, y)$	<code>MAX(x, y)</code>	†
$min_F(x, y)$	<code>MIN(x, y)</code>	†
$mmax_F(x, y)$	<code>MMAX(x, y)</code>	†
$mmin_F(x, y)$	<code>MMIN(x, y)</code>	†
$max_seq_F(xs)$	<code>MAXS(xs)</code>	†
$min_seq_F(xs)$	<code>MINS(xs)</code>	†
$mmax_seq_F(xs)$	<code>MMAXS(xs)</code>	†
$mmin_seq_F(xs)$	<code>MMINS(xs)</code>	†
$dim_F(x, y)$	<code>MONUS(x, y)</code>	†
$rounding_F(x)$	<code>ROUNDING(x)</code>	†
$floor_F(x)$	<code>FLOOR(x)</code>	★
$ceiling_F(x)$	<code>CEILING(x)</code>	†
$rounding_rest_F(x)$	<code>x - ROUNDING(x)</code>	†
$floor_rest_F(x)$	<code>x - FLOOR(x)</code>	★
$ceiling_rest_F(x)$	<code>x - CEILING(x)</code>	†
$remr_F(x, y)$	<code>REMAINDER(x, y)</code>	†
$sqrt_F(x)$	<code>SQRT(x)</code>	★
$rsqrt_F(x)$	<code>RSQRT(x)</code>	†
$add_lo_F(x, y)$	<code>ADD_LOW(x, y)</code>	†

$sub_{lo}_F(x, y)$	SUB_LOW(x, y)	†
$mul_{lo}_F(x, y)$	MUL_LOW(x, y)	†
$div_{rest}_F(x, y)$	DIV_REST(x, y)	†
$sqrt_{rest}_F(x)$	SQRT_REST(x)	†

where x , y , and z are expressions of type **numeric**, and where xs is an expression of type **array** of **numeric**.

The LIA-2 parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_{err_hypot}_F$	ERR_HYPOTENUSE	†
$max_{err_exp}_F$	ERR_EXP	†
$max_{err_power}_F$	ERR_POWER	†
$max_{err_sinh}_F$	ERR_SINH	†
$max_{err_tanh}_F$	ERR_TANH	†
$big_angle_r_F$	BIG_RADIAN_ANGLE	†
$max_{err_sin}_F$	ERR_SIN	†
$max_{err_tan}_F$	ERR_TAN	†
$min_angular_unit_F$	MIN_ANGLE_UNIT	†
$big_angle_u_F$	BIG_ANGLE	†
$max_{err_sinu}_F(u)$	ERR_SIN_CYCLE(u)	†
$max_{err_tanu}_F(u)$	ERR_TAN_CYCLE(u)	†
$max_{err_convert}_F$	ERR_CONVERT	†
$max_{err_convert}'_F$	ERR_CONVERT_TO_STRING	†
$max_{err_convert}_D$	ERR_CONVERT_TO_STRING	†

where u is an expression of type **numeric**.

The LIA-2 floating point operations are listed below, along with the syntax used to invoke them. BASIC has a degree mode and a radian mode for the trigonometric operations.

$hypot_F(x, y)$	HYPOT(x, y)	†
$exp_F(x)$	EXP(x)	*
$expm1_F(x)$	EXPM1(x)	†
$exp2_F(x)$	EXP2(x)	†
$exp10_F(x)$	EXP10(x)	†
$power_F(b, y)$	POWER(b, y)	†
$power1pm1_F(b, y)$	POWER1PM1(b, y)	†
$ln_F(x)$	LOG(x)	*
$ln1p_F(x)$	LOG1P(x)	†
$log2_F(x)$	LOG2(x)	*
$log10_F(x)$	LOG10(x)	*
$logbase_F(b, x)$	LOGBASE(b, x)	†
$logbase1p1p_F(b, x)$	LOGBASE1P1P(b, x)	†
$sinh_F(x)$	SINH(x)	*
$cosh_F(x)$	COSH(x)	*
$tanh_F(x)$	TANH(x)	*
$coth_F(x)$	COTH(x)	†

$sech_F(x)$	SECH(x)	†
$csch_F(x)$	CSCH(x)	†
$arcsinh_F(x)$	ARCSINH(x)	†
$arccosh_F(x)$	ARCCOSH(x)	†
$arctanh_F(x)$	ARCTANH(x)	†
$arcoth_F(x)$	ARCCOTH(x)	†
$arcsech_F(x)$	ARCSECH(x)	†
$arccsch_F(x)$	ARCCSCH(x)	†
$rad_F(x)$	NORMANGLE(x) (when in radian mode)	†
$sin_F(x)$	SIN(x) (when in radian mode)	★
$cos_F(x)$	COS(x) (when in radian mode)	★
$tan_F(x)$	TAN(x) (when in radian mode)	★
$cot_F(x)$	COT(x) (when in radian mode)	★
$sec_F(x)$	SEC(x) (when in radian mode)	†
$csc_F(x)$	CSC(x) (when in radian mode)	†
$arcsin_F(x)$	ARCSIN(x) (when in radian mode)	★
$arccos_F(x)$	ARCCOS(x) (when in radian mode)	★
$arctan_F(x)$	ARCTAN(x) (when in radian mode)	★
$arccot_F(x)$	ARCCOT(x) (when in radian mode)	★
$arcctg_F(x)$	ARCCTG(x) (when in radian mode)	†
$arcsec_F(x)$	ARCSEC(x) (when in radian mode)	†
$arccsc_F(x)$	ARCCSC(x) (when in radian mode)	†
$arc_F(x, y)$	ANGLE(x, y) (when in radian mode)	★
$cycle_F(u, x)$	NORMANGLEU(u, x)	†
$sinu_F(u, x)$	SINU(u, x)	†
$cosu_F(u, x)$	COSU(u, x)	†
$tanu_F(u, x)$	TANU(u, x)	†
$cotu_F(u, x)$	COTU(u, x)	†
$secu_F(u, x)$	SECU(u, x)	†
$cscu_F(u, x)$	CSCU(u, x)	†
$arcsinu_F(u, x)$	ARCSINU(u, x)	†
$arccosu_F(u, x)$	ARCCOSU(u, x)	†
$arctanu_F(u, x)$	ARCTANU(u, x)	†
$arccotu_F(u, x)$	ARCCOTU(u, x)	†
$arcctgu_F(u, x)$	ARCCTGU(u, x)	†
$arcsecu_F(u, x)$	ARCSECU(u, x)	†
$arccscu_F(u, x)$	ARCCSCU(u, x)	†
$arcu_F(u, x, y)$	ANGLEU(u, x, y)	†
$cycle_F(360, x)$	NORMANGLE(x) (when in degree mode)	†
$sinu_F(360, x)$	SIN(x) (when in degree mode)	★
$cosu_F(360, x)$	COS(x) (when in degree mode)	★
$tanu_F(360, x)$	TAN(x) (when in degree mode)	★

$cotu_F(360, x)$	COT(x) (when in degree mode)	★
$secu_F(360, x)$	SEC(x) (when in degree mode)	†
$cscu_F(360, x)$	CSC(x) (when in degree mode)	†
$arcsinu_F(360, x)$	ARCSIN(x) (when in degree mode)	★
$arccosu_F(360, x)$	ARCCOS(x) (when in degree mode)	★
$arctanu_F(360, x)$	ARCTAN(x) (when in degree mode)	★
$arccotu_F(360, x)$	ARCCOT(x) (when in degree mode)	★
$arcctgu_F(360, x)$	ARCCTG(x) (when in degree mode)	†
$arcsecu_F(360, x)$	ARCSEC(x) (when in degree mode)	†
$arccscu_F(360, x)$	ARCCSC(x) (when in degree mode)	†
$arcu_F(360, x, y)$	ANGLE(x, y) (when in degree mode)	★
$rad_to_cycle_F(x, u)$	RAD_TO_CYCLE(x, u)	†
$cycle_to_rad_F(u, x)$	CYCLE_TO_RAD(u, x)	†
$cycle_to_cycle_F(u, x, v)$	CYCLE_TO_CYCLE(u, x, v)	†

where b , x , y , u , and v are expressions of type **numeric**.

Arithmetic value conversions in BASIC are always tied to reading and writing text.

$convert_{F'' \rightarrow F}(stdin)$	READ x	★
$convert_{F \rightarrow F''}(y)$	PRINT y	★
$convert_{D' \rightarrow F}(stdin)$	READ x	★

where x is a variable of type **numeric**, y is an expression of type **numeric**.

BASIC provides non-negative numerals for **numeric** in base 10.

BASIC does not specify any numerals for infinities and NaNs. Suggestion:

+∞	INFINITY	†
qNaN	NAN	†
sNaN	SIGNAN	†

as well as string formats for reading and writing these values as character strings.

BASIC has a notion of ‘exception’ that implies a non-returnable change of control flow. BASIC uses its exception mechanism as its default means of notification. BASIC ignores **underflow** notifications since a BASIC exception is inappropriate for an **underflow** notification. On **underflow** the continuation value (specified in LIA-2) is used directly without recording the **underflow** itself. BASIC uses the exception number 1003 (Numeric supplied function overflow) for **overflow**, the exception number 3001 (Division by zero) for **pole**, and the exception numbers 301x(?) for **invalid**. Since BASIC exceptions are non-returnable changes of control flow, no continuation value is provided for these notifications.

An implementation that wishes to follow LIA-2 should provide recording of indicators as an alternative means of handling numeric notifications. Recording of indicators is the LIA-2 preferred means of handling numeric notifications.

C.4 C

The programming language C is defined by ISO/IEC 9899:1990, *Information technology – Programming languages – C* [18], currently under revision (C9x FDIS).

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

The LIA-1 datatype **Boolean** is implemented by the C datatype **int** (1 = **true** and 0 = **false**). (Revised C will provide a **_Bool** datatype.)

Every implementation of C has integral datatypes **int**, **long int**, **unsigned int**, and **unsigned long int**. *INT* is used below to designate one of the integer datatypes.

C has three floating point datatypes: **float**, **double**, and **long double**. *FLT* is used below to designate one of the floating point datatypes.

The LIA-2 integer operations are listed below, along with the syntax used to invoke them:

<i>max_I</i> (<i>x</i> , <i>y</i>)	imax (<i>x</i> , <i>y</i>)	†
<i>min_I</i> (<i>x</i> , <i>y</i>)	imin (<i>x</i> , <i>y</i>)	†
<i>max_seq_I</i> (<i>xs</i>)	imax_arr (<i>xs</i> , <i>nr_of_items</i>)	†*
<i>min_seq_I</i> (<i>xs</i>)	imin_arr (<i>xs</i> , <i>nr_of_items</i>)	†*
<i>dim_I</i> (<i>x</i> , <i>y</i>)	idim (<i>x</i> , <i>y</i>)	†
<i>power_I</i> (<i>x</i> , <i>y</i>)	ipower (<i>x</i> , <i>y</i>)	†
<i>shift2_I</i> (<i>x</i> , <i>y</i>)	shift2 (<i>x</i> , <i>y</i>)	†
<i>shift10_I</i> (<i>x</i> , <i>y</i>)	shift10 (<i>x</i> , <i>y</i>)	†
<i>sqr_I</i> (<i>x</i>)	isqrt (<i>x</i>)	†
<i>divides_I</i> (<i>x</i> , <i>y</i>)	does_divide (<i>x</i> , <i>y</i>)	†
<i>divides_I</i> (<i>x</i> , <i>y</i>)	<i>x</i> != 0 && <i>y</i> % <i>x</i> == 0	*
<i>even_I</i> (<i>x</i>)	<i>x</i> % 2 == 0	*
<i>odd_I</i> (<i>x</i>)	<i>x</i> % 2 != 0	*
<i>div_I</i> (<i>x</i> , <i>y</i>)	div (<i>x</i> , <i>y</i>)	†
<i>mod_I</i> (<i>x</i> , <i>y</i>)	mod (<i>x</i> , <i>y</i>)	†
<i>group_I</i> (<i>x</i> , <i>y</i>)	group (<i>x</i> , <i>y</i>)	†
<i>pad_I</i> (<i>x</i> , <i>y</i>)	pad (<i>x</i> , <i>y</i>)	†
<i>quot_I</i> (<i>x</i> , <i>y</i>)	quot (<i>x</i> , <i>y</i>)	†
<i>remr_I</i> (<i>x</i> , <i>y</i>)	iremainder (<i>x</i> , <i>y</i>)	†
<i>gcd_I</i> (<i>x</i> , <i>y</i>)	gcd (<i>x</i> , <i>y</i>)	†
<i>lcm_I</i> (<i>x</i> , <i>y</i>)	lcm (<i>x</i> , <i>y</i>)	†
<i>gcd_seq_I</i> (<i>xs</i>)	gcd_arr (<i>xs</i> , <i>nr_of_items</i>)	†*
<i>lcm_seq_I</i> (<i>xs</i>)	lcm_arr (<i>xs</i> , <i>nr_of_items</i>)	†*
<i>add_wrap_I</i> (<i>x</i> , <i>y</i>)	add_wrap (<i>x</i> , <i>y</i>)	†
<i>add_ov_I</i> (<i>x</i> , <i>y</i>)	add_over (<i>x</i> , <i>y</i>)	†
<i>sub_wrap_I</i> (<i>x</i> , <i>y</i>)	sub_wrap (<i>x</i> , <i>y</i>)	†
<i>sub_ov_I</i> (<i>x</i> , <i>y</i>)	sub_over (<i>x</i> , <i>y</i>)	†
<i>mul_wrap_I</i> (<i>x</i> , <i>y</i>)	mul_wrap (<i>x</i> , <i>y</i>)	†
<i>mul_ov_I</i> (<i>x</i> , <i>y</i>)	mul_over (<i>x</i> , <i>y</i>)	†

where *x* and *y* are expressions of the same integer type and where *xs* is an expression of type array of an integer type. (The operations marked with * needs one name per integer datatype.)

The LIA-2 non-transcendental floating point operations are listed below, along with the syntax (type generic macros) used to invoke them:

<i>min_F</i> (<i>x</i> , <i>y</i>)	nmin (<i>x</i> , <i>y</i>)	†
<i>max_F</i> (<i>x</i> , <i>y</i>)	nmax (<i>x</i> , <i>y</i>)	†
<i>mmin_F</i> (<i>x</i> , <i>y</i>)	fmin (<i>x</i> , <i>y</i>)	*(C9x)

$mmax_F(x, y)$	$fmax(x, y)$	★(C9x)
$min_seq_F(xs)$	$nmin_arr(xs, nr_of_items)$	†*
$max_seq_F(xs)$	$nmax_arr(xs, nr_of_items)$	†*
$mmin_seq_F(xs)$	$fmin_arr(xs, nr_of_items)$	†*
$mmax_seq_F(xs)$	$fmax_arr(xs, nr_of_items)$	†*
$dim_F(x, y)$	$fdim(x, y)$	★(C9x)
$rounding_F(x)$	$nearbyint(x)$ (when in round to nearest mode)	★(C9x)
$floor_F(x)$	$floor(x)$	★
$ceiling_F(x)$	$ceil(x)$	★
$rounding_rest_F(x)$	$x - nearbyint(x)$ (when in round to nearest mode)	★(C9x)
$floor_rest_F(x)$	$x - floor(x)$	★
$ceiling_rest_F(x)$	$x - ceil(x)$	★
$remr_F(x, y)$	$remainder(x, y)$	★(C9x)
$sqr_F(x)$	$sqr(x)$	★
$rsqr_F(x)$	$rsqr(x)$	†
$add_lo_F(x, y)$	$add_low(x, y)$	†
$sub_lo_F(x, y)$	$sub_low(x, y)$	†
$mul_lo_F(x, y)$	$mul_low(x, y)$	†
$div_rest_F(x, y)$	$div_rest(x, y)$	†
$sqr_rest_F(x)$	$sqr_rest(x)$	†
$mul_{F \rightarrow F'}(x, y)$	$dprod(x, y)$	†

where x , y and z are expressions of the same floating point type, and where xs is an expression of type array of a floating point type.

The LIA-2 parameters for operations approximating real valued transcendental functions can be accessed by the following syntax (type generic macros):

$max_err_hypot_F$	$err_hypot(x)$	†
$max_err_exp_F$	$err_exp(x)$	†
$max_err_power_F$	$err_power(x)$	†
$max_err_sinh_F$	$err_sinh(x)$	†
$max_err_tanh_F$	$err_tanh(x)$	†
$big_angle_r_F$	$big_radian_angle(x)$	†
$max_err_sin_F$	$err_sin(x)$	†
$max_err_tan_F$	$err_tan(x)$	†
$min_angular_unit_F$	$smallest_angle_unit(x)$	†
$big_angle_u_F$	$big_angle(x)$	†
$max_err_sinu_F(u)$	$err_sin_cycle(u)$	†
$max_err_tanu_F(u)$	$err_tan_cycle(u)$	†
$max_err_convert_F$	$err_convert(x)$	†
$max_err_convert_{F'}$	$err_convert_to_string$	†
$max_err_convert_{D'}$	$err_convert_to_string$	†

where x and u are expressions of a floating point type. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

C has a `pow` operation that does not conform to LIA-2, but may be specified in LIA-2 terms as:

$$\begin{aligned}
 \mathit{pow}_F(x, y) &= 1 && \text{if } x \text{ is a quiet NaN and } y = 0 \\
 &= \mathit{pow}_F(x, 0) && \text{if } y = -0 \\
 &= \mathit{power}_{FZ}(x, y) && \text{if } y \in F \cap Z \\
 &= \mathit{power}_F(x, y) && \text{otherwise}
 \end{aligned}$$

C has a `hypot` operation that does not conform to LIA-2, but may be specified in LIA-2 terms as:

$$\begin{aligned}
 \mathit{hhypot}_F(x, y) &= +\infty && \text{if } x \text{ is a quiet NaN and } y \in \{-\infty, +\infty\} \\
 &= +\infty && \text{if } x \in \{-\infty, +\infty\} \text{ and } y \text{ is a quiet NaN} \\
 &= \mathit{hypot}_F(x, y) && \text{otherwise}
 \end{aligned}$$

The LIA-2 elementary floating point operations are listed below, together with the non-LIA-2 pow_F and hhypot_F , along with the syntax (type generic macros) used to invoke them:

$\mathit{hypot}_F(x, y)$	<code>hypotenuse(x, y)</code>	†
$\mathit{hhypot}_F(x, y)$	<code>hypot(x, y)</code>	* Not LIA-2!
$\mathit{power}_{FI}(b, z)$	<code>poweri(b, z)</code>	†
$\mathit{exp}_F(x)$	<code>exp(x)</code>	*
$\mathit{expm1}_F(x)$	<code>expm1(x)</code>	*(C9x)
$\mathit{exp2}_F(x)$	<code>exp2(x)</code>	*
$\mathit{exp10}_F(x)$	<code>exp10(x)</code>	†
$\mathit{power}_F(b, y)$	<code>power(b, y)</code>	†
$\mathit{pow}_F(b, y)$	<code>pow(b, y)</code>	* Not LIA-2!
$\mathit{power1pm1}_F(b, y)$	<code>power1pm1(b, y)</code>	†
$\mathit{ln}_F(x)$	<code>log(x)</code>	*
$\mathit{ln1p}_F(x)$	<code>log1p(x)</code>	*(C9x)
$\mathit{log2}_F(x)$	<code>log2(x)</code>	*
$\mathit{log10}_F(x)$	<code>log10(x)</code>	*
$\mathit{logbase}_F(b, x)$	<code>logbase(b, x)</code>	†
$\mathit{logbase1p1p}_F(b, x)$	<code>logbase1p1p(b, x)</code>	†
$\mathit{sinh}_F(x)$	<code>sinh(x)</code>	*
$\mathit{cosh}_F(x)$	<code>cosh(x)</code>	*
$\mathit{tanh}_F(x)$	<code>tanh(x)</code>	*
$\mathit{coth}_F(x)$	<code>coth(x)</code>	†
$\mathit{sech}_F(x)$	<code>sech(x)</code>	†
$\mathit{csch}_F(x)$	<code>csch(x)</code>	†
$\mathit{arcsinh}_F(x)$	<code>asinh(x)</code>	*
$\mathit{arccosh}_F(x)$	<code>acosh(x)</code>	*
$\mathit{arctanh}_F(x)$	<code>atanh(x)</code>	*
$\mathit{arcoth}_F(x)$	<code>acoth(x)</code>	†
$\mathit{arcsech}_F(x)$	<code>asech(x)</code>	†
$\mathit{arcsch}_F(x)$	<code>acsch(x)</code>	†
$\mathit{rad}_F(x)$	<code>radian(x)</code>	†
$\mathit{axis_rad}_F(x)$	<code>axis_rad(x, &h, &v)</code> (note out parameters)	†
$\mathit{sin}_F(x)$	<code>sin(x)</code>	*

$\cos_F(x)$	<code>cos(x)</code>	★
$\tan_F(x)$	<code>tan(x)</code>	★
$\cot_F(x)$	<code>cot(x)</code>	†
$\sec_F(x)$	<code>sec(x)</code>	†
$\csc_F(x)$	<code>csc(x)</code>	†
$\text{cossin}_F(x)$	<code>cossin(x, &c, &s)</code>	†
$\arcsin_F(x)$	<code>asin(x)</code>	★
$\arccos_F(x)$	<code>acos(x)</code>	★
$\arctan_F(x)$	<code>atan(x)</code>	★
$\text{arccot}_F(x)$	<code>acot(x)</code>	†
$\text{arctg}_F(x)$	<code>actg(x)</code>	†
$\text{arcsec}_F(x)$	<code>asec(x)</code>	†
$\text{arccsc}_F(x)$	<code>acsc(x)</code>	†
$\text{arc}_F(x, y)$	<code>atan2(y, x)</code>	★
$\text{cycle}_F(u, x)$	<code>cycle(u, x)</code>	†
$\text{axis_cycle}_F(u, x)$	<code>axis_cycle(u, x, &h, &v)</code>	†
$\sinu_F(u, x)$	<code>sinu(u, x)</code>	†
$\cosu_F(u, x)$	<code>cosu(u, x)</code>	†
$\tanu_F(u, x)$	<code>tanu(u, x)</code>	†
$\cotu_F(u, x)$	<code>cotu(u, x)</code>	†
$\secu_F(u, x)$	<code>secu(u, x)</code>	†
$\cscu_F(u, x)$	<code>cscu(u, x)</code>	†
$\text{cossinu}_F(u, x)$	<code>cossinu(u, x, &c, &s)</code>	†
$\arcsinu_F(u, x)$	<code>asinu(u, x)</code>	†
$\arccosu_F(u, x)$	<code>acosu(u, x)</code>	†
$\arctanu_F(u, x)$	<code>atanu(u, x)</code>	†
$\text{arccotu}_F(u, x)$	<code>acotu(u, x)</code>	†
$\text{arctgu}_F(u, x)$	<code>actgu(u, x)</code>	†
$\text{arcsecu}_F(u, x)$	<code>asecu(u, x)</code>	†
$\text{arccscu}_F(u, x)$	<code>acscu(u, x)</code>	†
$\text{arcu}_F(u, x, y)$	<code>atan2u(u, y, x)</code>	†
$\text{rad_to_cycle}_F(x, u)$	<code>radian_to_cycle(x, u)</code>	†
$\text{cycle_to_rad}_F(u, x)$	<code>cycle_to_radian(u, x)</code>	†
$\text{cycle_to_cycle}_F(u, x, v)$	<code>cycle_to_cycle(u, x, v)</code>	†

where b , x , y , u , and v are expressions of the same floating point type.

Arithmetic value conversions in C can be explicit or implicit. The explicit arithmetic value conversions are usually expressed as ‘casts’, except when converting to/from strings. The rules for when implicit conversions are applied is not repeated here, but work as if a cast had been applied.

$\text{convert}_{I \rightarrow I'}(x)$	<code>(INT2)x</code>	★
$\text{convert}_{I'' \rightarrow I}(s)$	<code>sscanf(s, "%no", &r)</code>	★
$\text{convert}_{I'' \rightarrow I}(s)$	<code>sscanf(s, "%nd", &r)</code>	★
$\text{convert}_{I'' \rightarrow I}(s)$	<code>sscanf(s, "%nx", &r)</code>	★
$\text{convert}_{I'' \rightarrow I}(f)$	<code>fscanf(f, "%no", &r)</code>	★
$\text{convert}_{I'' \rightarrow I}(f)$	<code>fscanf(f, "%nd", &r)</code>	★

$convert_{I'' \rightarrow I}(f)$	<code>fscanf(f, "%nx", &r)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>sprintf(s, "%no", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>sprintf(s, "%nd", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>sprintf(s, "%nx", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>fprintf(h, "%no", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>fprintf(h, "%nd", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>fprintf(h, "%nx", x)</code>	★
$rounding_{F \rightarrow I}(y)$	<code>(INT)nearbyint(y)</code> (when in round to nearest mode)	★
$floor_{F \rightarrow I}(y)$	<code>(INT)floor(y)</code>	★
$ceiling_{F \rightarrow I}(y)$	<code>(INT)ceil(y)</code>	★
$convert_{I \rightarrow F}(x)$	<code>(FLT)x</code>	★
$convert_{F \rightarrow F'}(y)$	<code>(FLT2)y</code>	★
$convert_{F'' \rightarrow F}(s)$	<code>sscanf(s, "%e", &r)</code>	★
$convert_{F'' \rightarrow F}(f)$	<code>fscanf(f, "%e", &r)</code>	★
$convert_{F \rightarrow F''}(y)$	<code>sprintf(s, "%.de", x)</code>	★
$convert_{F \rightarrow F''}(y)$	<code>fprintf(h, "%.de", x)</code>	★
$convert_{D' \rightarrow F}(s)$	<code>sscanf(s, "%f", &r)</code>	★
$convert_{D' \rightarrow F}(f)$	<code>fscanf(f, "%f", &r)</code>	★
$convert_{F \rightarrow D'}(y)$	<code>sprintf(s, "%.df", x)</code>	★
$convert_{F \rightarrow D'}(y)$	<code>fprintf(h, "%.df", x)</code>	★

where x is an expression of type *INT*, y is an expression of type *FLT*, and z is an expression of type *FXD*, where *FXD* is a fixed point type. *INT2* is the integer datatype that corresponds to I' . A ? above indicates that the parameter is optional. e is greater than 0.

C9x provides non-negative numerals for all its integer and floating point types. The default base is 10, but base 8 (for integers) and 16 (both integer and float) can be used too. Numerals for different integer types are distinguished by suffixes. Numerals for different floating point types are distinguished by suffix: **f** for **float**, no suffix for **double**, **l** for **long double**. Numerals for floating point types must have a '.' in them. The details are not repeated in this example binding, see ISO/IEC FDIS 9899, clause 6.4.4.1 Integer constants, and clause 6.4.4.2 Floating constants.

C9x specifies numerals (as macros) for infinities and NaNs for **float**:

+∞	INFINITY	★
qNaN	NAN	★
sNaN	SIGNAN	†

as well as string formats for reading and writing these values as character strings.

C9x has two ways of handling arithmetic errors. One, for backwards compatibility, is by assigning to **errno**. The other is by recording of indicators, the method preferred by LIA-2, which can be used for floating point errors. For C9x, the **absolute_precision_underflow** notification is ignored. The behaviour for notification upon integer operations initiating a notification is, however, not defined by C9x.

C.5 C++

The programming language C++ is defined by ISO/IEC 14882:1998, *Programming languages – C++* [19].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

This example binding recommends that all identifiers suggested here be defined in the namespace `std::math`.

The LIA-1 datatype **Boolean** is implemented by the C++ datatype `bool`.

Every implementation of C++ has integral datatypes `int`, `long int`, `unsigned int`, and `unsigned long int`. *INT* is used below to designate one of the integer datatypes.

C++ has three floating point datatypes: `float`, `double`, and `long double`. *FLT* is used below to designate one of the floating point datatypes.

The LIA-2 integer operations are listed below, along with the syntax used to invoke them:

$max_I(x, y)$	<code>max(x, y)</code>	★
$min_I(x, y)$	<code>min(x, y)</code>	★
$max_seq_I(xs)$	<code>xs.max()</code>	★
$min_seq_I(xs)$	<code>xs.min()</code>	★
$dim_I(x, y)$	<code>dim(x, y)</code>	†
$power_I(x, y)$	<code>power(x, y)</code>	†
$sqr_I(x)$	<code>sqr(x)</code>	†
$shift2_I(x, y)$	<code>shift2(x, y)</code>	†
$shift10_I(x, y)$	<code>shift10(x, y)</code>	†
$divides_I(x, y)$	<code>does_divide(x, y)</code>	†
$divides_I(x, y)$	<code>y != 0 && y % x == 0</code>	★
$even_I(x)$	<code>x % 2 == 0</code>	★
$odd_I(x)$	<code>x % 2 != 0</code>	★
$div_I(x, y)$	<code>div(x, y)</code>	†
$moda_I(x, y)$	<code>mod(x, y)</code>	†
$group_I(x, y)$	<code>group(x, y)</code>	†
$pad_I(x, y)$	<code>pad(x, y)</code>	†
$quot_I(x, y)$	<code>quot(x, y)</code>	†
$remr_I(x, y)$	<code>iremainder(x, y)</code>	†
$gcd_I(x, y)$	<code>gcd(x, y)</code>	†
$lcm_I(x, y)$	<code>lcm(x, y)</code>	†
$gcd_seq_I(xs)$	<code>xs.gcd()</code>	†
$lcm_seq_I(xs)$	<code>xs.lcm()</code>	†
$add_wrap_I(x, y)$	<code>add_wrap(x, y)</code>	†
$add_ov_I(x, y)$	<code>add_over(x, y)</code>	†
$sub_wrap_I(x, y)$	<code>sub_wrap(x, y)</code>	†
$sub_ov_I(x, y)$	<code>sub_over(x, y)</code>	†
$mul_wrap_I(x, y)$	<code>mul_wrap(x, y)</code>	†
$mul_ov_I(x, y)$	<code>mul_over(x, y)</code>	†

where x and y are expressions of the same integer type and where xs is an expression of type valarray of an integer type.

The LIA-2 non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$max_F(x, y)$	<code>nmax(x, y)</code>	†
$min_F(x, y)$	<code>nmin(x, y)</code>	†
$mmax_F(x, y)$	<code>max(x, y)</code>	*(unclear)
$mmin_F(x, y)$	<code>min(x, y)</code>	*(unclear)
$max_seq_F(xs)$	<code>xs.nmax()</code>	†
$min_seq_F(xs)$	<code>xs.nmin()</code>	†
$mmax_seq_F(xs)$	<code>xs.max()</code>	*(unclear)
$mmin_seq_F(xs)$	<code>xs.min()</code>	*(unclear)
$dim_F(x, y)$	<code>dim(x, y)</code>	†
$rounding_F(x)$	<code>round(x)</code>	†
$floor_F(x)$	<code>floor(x)</code>	*
$ceiling_F(x)$	<code>ceil(x)</code>	*
$rounding_rest_F(x)$	<code>x - round(x)</code>	†
$floor_rest_F(x)$	<code>x - floor(x)</code>	*
$ceiling_rest_F(x)$	<code>x - ceil(x)</code>	*
$mul_{F \rightarrow F'}(x, y)$	<code>dprod(x, y)</code>	†
$remr_F(x, y)$	<code>remainder(x, y)</code>	†
$sqr_F(x)$	<code>sqr(x)</code>	*
$rsqr_F(x)$	<code>reciprocal_sqr(x)</code>	†
$add_lo_F(x, y)$	<code>add_low(x, y)</code>	†
$sub_lo_F(x, y)$	<code>sub_low(x, y)</code>	†
$mul_lo_F(x, y)$	<code>mul_low(x, y)</code>	†
$div_rest_F(x, y)$	<code>div_rest(x, y)</code>	†
$sqr_rest_F(x)$	<code>sqr_rest(x)</code>	†

where x , y and z are expressions of the same floating point type, and where xs is an expression of type valarray of a floating point type.

The parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	<code>err_hypotenuse<FLT>()</code>	†
$max_err_exp_F$	<code>err_exp<FLT>()</code>	†
$max_err_power_F$	<code>err_power<FLT>()</code>	†
$max_err_sinh_F$	<code>err_sinh<FLT>()</code>	†
$max_err_tanh_F$	<code>err_tanh<FLT>()</code>	†
$big_angle_r_F$	<code>big_radian_angle<FLT>()</code>	†
$max_err_sin_F$	<code>err_sin<FLT>()</code>	†
$max_err_tan_F$	<code>err_tan<FLT>()</code>	†
$min_angular_unit_F$	<code>smallest_angle_unit<FLT>()</code>	†
$big_angle_u_F$	<code>big_angle<FLT>()</code>	†
$max_err_sinu_F(u)$	<code>err_sin_cycle(u)</code>	†
$max_err_tanu_F(u)$	<code>err_tan_cycle(u)</code>	†
$max_err_convert_F$	<code>err_convert<FLT>()</code>	†
$max_err_convert_{F'}$	<code>err_convert_to_string()</code>	†

max_err_convert_D `err_convert_to_string()` †

where *u* is an expression of a floating point type. Several of the parameter functions are constant for each type (and library).

The LIA-2 elementary floating point operations are listed below, along with the syntax (type generic macros) used to invoke them:

<i>hypot_F</i> (<i>x</i> , <i>y</i>)	<code>hypotenusu</code> (<i>x</i> , <i>y</i>)	†
<i>power_{FI}</i> (<i>b</i> , <i>z</i>)	<code>poweri</code> (<i>b</i> , <i>z</i>)	†
<i>exp_F</i> (<i>x</i>)	<code>exp</code> (<i>x</i>)	★
<i>expm1_F</i> (<i>x</i>)	<code>expm1</code> (<i>x</i>)	†
<i>exp2_F</i> (<i>x</i>)	<code>exp2</code> (<i>x</i>)	†
<i>exp10_F</i> (<i>x</i>)	<code>exp10</code> (<i>x</i>)	†
<i>power_F</i> (<i>b</i> , <i>y</i>)	<code>power</code> (<i>b</i> , <i>y</i>)	†
<i>pow_F</i> (<i>b</i> , <i>y</i>)	<code>pow</code> (<i>b</i> , <i>y</i>)	★ Not LIA-2! (See C.)
<i>power1pm1_F</i> (<i>b</i> , <i>y</i>)	<code>power1pm1</code> (<i>b</i> , <i>y</i>)	†
<i>ln_F</i> (<i>x</i>)	<code>log</code> (<i>x</i>)	★
<i>ln1p_F</i> (<i>x</i>)	<code>log1p</code> (<i>x</i>)	†
<i>log2_F</i> (<i>x</i>)	<code>log2</code> (<i>x</i>)	†
<i>log10_F</i> (<i>x</i>)	<code>log10</code> (<i>x</i>)	★
<i>logbase_F</i> (<i>b</i> , <i>x</i>)	<code>logbase</code> (<i>b</i> , <i>x</i>)	†
<i>logbase1p1p_F</i> (<i>b</i> , <i>x</i>)	<code>logbase1p1p</code> (<i>b</i> , <i>x</i>)	†
<i>sinh_F</i> (<i>x</i>)	<code>sinh</code> (<i>x</i>)	★
<i>cosh_F</i> (<i>x</i>)	<code>cosh</code> (<i>x</i>)	★
<i>tanh_F</i> (<i>x</i>)	<code>tanh</code> (<i>x</i>)	★
<i>coth_F</i> (<i>x</i>)	<code>coth</code> (<i>x</i>)	†
<i>sech_F</i> (<i>x</i>)	<code>sech</code> (<i>x</i>)	†
<i>csch_F</i> (<i>x</i>)	<code>csch</code> (<i>x</i>)	†
<i>arcsinh_F</i> (<i>x</i>)	<code>asinh</code> (<i>x</i>)	★
<i>arccosh_F</i> (<i>x</i>)	<code>acosh</code> (<i>x</i>)	★
<i>arctanh_F</i> (<i>x</i>)	<code>atanh</code> (<i>x</i>)	★
<i>arcoth_F</i> (<i>x</i>)	<code>acoth</code> (<i>x</i>)	†
<i>arcsech_F</i> (<i>x</i>)	<code>asech</code> (<i>x</i>)	†
<i>arccsch_F</i> (<i>x</i>)	<code>acsch</code> (<i>x</i>)	†
<i>rad_F</i> (<i>x</i>)	<code>rad</code> (<i>x</i>)	†
<i>axis_rad_F</i> (<i>x</i>)	<code>axis_rad</code> (<i>x</i> , & <i>h</i> , & <i>v</i>) (note out parameters)	†
<i>sin_F</i> (<i>x</i>)	<code>sin</code> (<i>x</i>)	★
<i>cos_F</i> (<i>x</i>)	<code>cos</code> (<i>x</i>)	★
<i>tan_F</i> (<i>x</i>)	<code>tan</code> (<i>x</i>)	★
<i>cot_F</i> (<i>x</i>)	<code>cot</code> (<i>x</i>)	†
<i>sec_F</i> (<i>x</i>)	<code>sec</code> (<i>x</i>)	†
<i>csc_F</i> (<i>x</i>)	<code>csc</code> (<i>x</i>)	†
<i>coassin_F</i> (<i>x</i>)	<code>coassin</code> (<i>x</i> , & <i>c</i> , & <i>s</i>)	†
<i>arcsin_F</i> (<i>x</i>)	<code>asin</code> (<i>x</i>)	★
<i>arccos_F</i> (<i>x</i>)	<code>acos</code> (<i>x</i>)	★

$arctan_F(x)$	<code>atan(x)</code>	★
$arccot_F(x)$	<code>acot(x)</code>	†
$arcctg_F(x)$	<code>actg(x)</code>	†
$arcsec_F(x)$	<code>asec(x)</code>	†
$arccsc_F(x)$	<code>acsc(x)</code>	†
$arc_F(x, y)$	<code>atan2(y, x)</code>	★
$cycle_F(u, x)$	<code>cycle(u, x)</code>	†
$axis_cycle_F(u, x)$	<code>axis_cycle(u, x, &h, &v)</code>	†
$sinu_F(u, x)$	<code>sinu(u, x)</code>	†
$cosu_F(u, x)$	<code>cosu(u, x)</code>	†
$tanu_F(u, x)$	<code>tanu(u, x)</code>	†
$cotu_F(u, x)$	<code>cotu(u, x)</code>	†
$secu_F(u, x)$	<code>secu(u, x)</code>	†
$cscu_F(u, x)$	<code>cscu(u, x)</code>	†
$cozzsinu_F(x)$	<code>cozzsinu(u, x, &c, &s)</code>	†
$arcsinu_F(u, x)$	<code>asinu(u, x)</code>	†
$arccosu_F(u, x)$	<code>acosu(u, x)</code>	†
$arctanu_F(u, x)$	<code>atanu(u, x)</code>	†
$arccotu_F(u, x)$	<code>acotu(u, x)</code>	†
$arctgu_F(u, x)$	<code>actgu(u, x)</code>	†
$arcsecu_F(u, x)$	<code>asecu(u, x)</code>	†
$arccscu_F(u, x)$	<code>acscu(u, x)</code>	†
$arcu_F(u, x, y)$	<code>atan2u(u, y, x)</code>	†
$rad_to_cycle_F(x, u)$	<code>radian_to_cycle(x, u)</code>	†
$cycle_to_rad_F(u, x)$	<code>cycle_to_radian(u, x)</code>	†
$cycle_to_cycle_F(u, x, v)$	<code>cycle_to_cycle(u, x, v)</code>	†

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*

Arithmetic value conversions in C++ are can be explicit or implicit. The rules for when implicit conversions are applied is not repeated here. The explicit arithmetic value conversions are usually expressed as ‘casts’, except when converting to/from strings.

$convert_{I \rightarrow I'}(x)$	<code>(INT2)x</code>	★
$convert_{I'' \rightarrow I}(s)$	<code>sscanf(s, "%no", &r)</code>	★
$convert_{I'' \rightarrow I}(s)$	<code>sscanf(s, "%nd", &r)</code>	★
$convert_{I'' \rightarrow I}(s)$	<code>sscanf(s, "%nx", &r)</code>	★
$convert_{I'' \rightarrow I}(f)$	<code>fscanf(f, "%no", &r)</code>	★
$convert_{I'' \rightarrow I}(f)$	<code>fscanf(f, "%nd", &r)</code>	★
$convert_{I'' \rightarrow I}(f)$	<code>fscanf(f, "%nx", &r)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>sprintf(s, "%no", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>sprintf(s, "%nd", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>sprintf(s, "%nx", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>fprintf(h, "%no", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>fprintf(h, "%nd", x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>fprintf(h, "%nx", x)</code>	★
$rounding_{F \rightarrow I}(y)$	<code>(INT)nearbyint(y)</code> (when in round to nearest mode)	★
$floor_{F \rightarrow I}(y)$	<code>(INT)floor(y)</code>	★

$ceiling_{F \rightarrow I}(y)$	$(INT)ceil(y)$	★
$convert_{I \rightarrow F}(x)$	$(FLT)x$	★
$convert_{F \rightarrow F'}(y)$	$(FLT2)y$	★
$convert_{F'' \rightarrow F}(s)$	$sscanf(s, "%e", &r)$	★
$convert_{F'' \rightarrow F}(f)$	$fscanf(f, "%e", &r)$	★
$convert_{F \rightarrow F''}(y)$	$sprintf(s, "%.de", x)$	★
$convert_{F \rightarrow F''}(y)$	$fprintf(h, "%.de", x)$	★
$convert_{D' \rightarrow F}(s)$	$sscanf(s, "%f", &r)$	★
$convert_{D' \rightarrow F}(f)$	$fscanf(f, "%f", &r)$	★
$convert_{F \rightarrow D'}(y)$	$sprintf(s, "%.df", x)$	★
$convert_{F \rightarrow D'}(y)$	$fprintf(h, "%.df", x)$	★

where x is an expression of type INT , y is an expression of type FLT , and z is an expression of type FXD , where FXD is a fixed point type. $INT2$ is the integer datatype that corresponds to I' . A ? above indicates that the parameter is optional. e is greater than 0.

C++ provides non-negative numerals for all its integer and floating point types in base 10. Numerals for different integer types are distinguished by suffixes. Numerals for different floating point types are distinguished by suffix: **f** for **float**, no suffix for **double**, **l** for **long double**. Numerals for floating point types must have a '.' in them. The details are not repeated in this example binding, see ISO/IEC 14882, clause 2.9.1 Integer literals, and clause 2.9.4 Floating literals.

C++ does not specify numerals for infinities and NaNs. Suggestion:

+∞	INFINITY	†
qNaN	NAN	†
sNaN	SIGNAN	†

as well as string formats for reading and writing these values as character strings.

C++ has completely undefined behaviour on arithmetic notification. An implementation wishing to conform to LIA-2 should provide a means for recording of indicators, similar to C9x.

C.6 Fortran

The programming language Fortran is defined by ISO/IEC 1539-1:1997, *Information technology – Programming languages – Fortran – Part 1: Base language* [23].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

The Fortran datatype **LOGICAL** corresponds to the LIA datatype **Boolean**.

Every implementation of Fortran has one integer datatype, denoted as **INTEGER**, and two floating point data type denoted as **REAL** (single precision) and **DOUBLE PRECISION**.

An implementation is permitted to offer additional **INTEGER** types with a different range and additional **REAL** types with different precision or range, parameterised with the **KIND** parameter.

The LIA-2 integer operations are listed below, along with the syntax used to invoke them:

$max_I(x, y)$	MAX(x, y)	★
$min_I(x, y)$	MIN(x, y)	★
$max_seq_I(xs)$	MAX($xs[1], xs[2], \dots, xs[n]$) or MAXVAL(xs)	★
$min_seq_I(xs)$	MIN($xs[1], xs[2], \dots, xs[n]$) or MINVAL(xs)	★
$dim_I(x, y)$	DIM(x, y)	★
$power_I(x, y)$	$x ** y$	★
$shift2_I(x, y)$	SHIFT2(x, y)	†
$shift10_I(x, y)$	SHIFT10(x, y)	†
$sqrI(x)$	ISQRT(x)	†
$divides_I(x, y)$	DIVIDES(x, y)	†
$even_I(x)$	MODULO($x, 2$) == 0	★
$odd_I(x)$	MODULO($x, 2$) != 0	★
$divf_I(x, y)$	DIV(x, y)	†
$moda_I(x, y)$	MODULO(x, y)	★
$group_I(x, y)$	GROUP(x, y)	†
$pad_I(x, y)$	PAD(x, y)	†
$quot_I(x, y)$	QUOTIENT(x, y)	†
$remr_I(x, y)$	REMAINDER(x, y)	†
$gcd_I(x, y)$	GCD(x, y)	†
$lcm_I(x, y)$	LCM(x, y)	†
$gcd_seq_I(xs)$	GCDVAL(xs)	†
$lcm_seq_I(xs)$	LCMVAL(xs)	†
$add_wrap_I(x, y)$	ADD_WRAP(x, y)	†
$add_ov_I(x, y)$	ADD_OVER(x, y)	†
$sub_wrap_I(x, y)$	SUB_WRAP(x, y)	†
$sub_ov_I(x, y)$	SUB_OVER(x, y)	†
$mul_wrap_I(x, y)$	MUL_WRAP(x, y)	†
$mul_ov_I(x, y)$	MUL_OVER(x, y)	†

where x and y are expressions of type INTEGER and where xs is an expression of type array of INTEGER.

The additional non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$max_F(x, y)$	MAX(x, y)	★
$min_F(x, y)$	MIN(x, y)	★
$mmax_F(x, y)$	MMAX(x, y)	†
$mmin_F(x, y)$	MMIN(x, y)	†
$max_seq_F(xs)$	MAX($xs[1], xs[2], \dots, xs[n]$) or MAXVAL(xs)	★
$min_seq_F(xs)$	MIN($xs[1], xs[2], \dots, xs[n]$) or MINVAL(xs)	★
$mmax_seq_F(xs)$	MMAX($xs[1], xs[2], \dots, xs[n]$) or MMAXVAL(xs)	†
$mmin_seq_F(xs)$	MMIN($xs[1], xs[2], \dots, xs[n]$) or MMINVAL(xs)	†
$dim_F(x, y)$	DIM(x, y)	★
$rounding_F(x)$	IEEE_RINT(x) (if in round to nearest mode) (★)	
$floor_F(x)$	IEEE_RINT(x) (if in round towards $-\infty$ mode) (★)	
$ceiling_F(x)$	IEEE_RINT(x) (if in round towards $+\infty$ mode) (★)	
$rounding_rest_F(x)$	$x - \text{IEEE_RINT}(x)$ (if in round to nearest mode) (★)	

$\mathit{floor_rest}_F(x)$	$x - \text{IEEE_RINT}(x)$ (if in round towards $-\infty$ mode) (*)	
$\mathit{ceiling_rest}_F(x)$	$x - \text{IEEE_RINT}(x)$ (if in round towards $+\infty$ mode) (*)	
$\mathit{remr}_F(x, y)$	$\text{IEEE_REM}(x, y)$	(*)
$\mathit{sqr}_F(x)$	$\text{SQRT}(x)$	*
$\mathit{rsqr}_F(x)$	$\text{RSQRT}(x)$	†
$\mathit{mul}_{F \rightarrow F'}(x, y)$	$\text{DPROD}(x, y)$	*
$\mathit{add_lo}_F(x, y)$	$\text{ADD_LOW}(x, y)$	†
$\mathit{sub_lo}_F(x, y)$	$\text{SUB_LOW}(x, y)$	†
$\mathit{mul_lo}_F(x, y)$	$\text{MUL_LOW}(x, y)$	†
$\mathit{div_rest}_F(x, y)$	$\text{DIV_REST}(x, y)$	†
$\mathit{sqr_rest}_F(x)$	$\text{SQRT_REST}(x)$	†

where x , y and z are expressions of type *FLT*, and where xs is an expression of type array of *FLT*.

The LIA-2 parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$\mathit{max_err_hypot}_F$	$\text{ERR_HYPOTENUSE}(x)$	†
$\mathit{max_err_exp}_F$	$\text{ERR_EXP}(x)$	†
$\mathit{max_err_power}_F$	$\text{ERR_POWER}(x)$	†
$\mathit{max_err_sinh}_F$	$\text{ERR_SINH}(x)$	†
$\mathit{max_err_tanh}_F$	$\text{ERR_TANH}(x)$	†
$\mathit{big_angle_r}_F$	$\text{BIG_RADIAN_ANGLE}(x)$	†
$\mathit{max_err_sin}_F$	$\text{ERR_SIN}(x)$	†
$\mathit{max_err_tan}_F$	$\text{ERR_TAN}(x)$	†
$\mathit{min_angular_unit}_F$	$\text{MIN_ANGLE_UNIT}(x)$	†
$\mathit{big_angle_u}_F$	$\text{BIG_ANGLE}(x)$	†
$\mathit{max_err_sinu}_F(u)$	$\text{ERR_SIN_CYCLE}(u)$	†
$\mathit{max_err_tanu}_F(u)$	$\text{ERR_TAN_CYCLE}(u)$	†
$\mathit{max_err_convert}_F$	$\text{ERR_CONVERT}(x)$	†
$\mathit{max_err_convert}_{F'}$	$\text{ERR_CONVERT_TO_STRING}$	†
$\mathit{max_err_convert}_{D'}$	$\text{ERR_CONVERT_TO_STRING}$	†

where b , x and u are expressions of type *FLT*. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them:

$\mathit{hypot}_F(x, y)$	$\text{HYPOT}(x, y)$	†
$\mathit{power}_{FI}(b, z)$	$b ** z$	*
$\mathit{exp}_F(x)$	$\text{EXP}(x)$	*
$\mathit{expm1}_F(x)$	$\text{EXPM1}(x)$	†
$\mathit{exp2}_F(x)$	$\text{EXP2}(x)$	†
$\mathit{exp10}_F(x)$	$\text{EXP10}(x)$	†
$\mathit{power}_F(b, y)$	$b ** y$	*
$\mathit{power1pm1}_F(b, y)$	$\text{POWER1PM1}(b, y)$	†

$\ln_F(x)$	LOG(x)	*
$\ln1p_F(x)$	LOG1P(x)	†
$\log^2_F(x)$	LOG2(x)	†
$\log10_F(x)$	LOG10(x)	*
$\logbase_F(b, x)$	LOGBASE(b, x)	†
$\logbase1p1p_F(b, x)$	LOGBASE1P1P(b, x)	†
$\sinh_F(x)$	SINH(x)	*
$\cosh_F(x)$	COSH(x)	*
$\tanh_F(x)$	TANH(x)	*
$\coth_F(x)$	COTH(x)	†
$\operatorname{sech}_F(x)$	SECH(x)	†
$\operatorname{csch}_F(x)$	CSCH(x)	†
$\operatorname{arcsinh}_F(x)$	ASINH(x)	†
$\operatorname{arccosh}_F(x)$	ACOSH(x)	†
$\operatorname{arctanh}_F(x)$	ATANH(x)	†
$\operatorname{arccoth}_F(x)$	ACOTH(x)	†
$\operatorname{arcsech}_F(x)$	ASECH(x)	†
$\operatorname{arccsch}_F(x)$	ACSCH(x)	†
$\operatorname{rad}_F(x)$	RAD(x)	†
$\sin_F(x)$	SIN(x)	*
$\cos_F(x)$	COS(x)	*
$\tan_F(x)$	TAN(x)	*
$\cot_F(x)$	COT(x)	†
$\sec_F(x)$	SEC(x)	†
$\csc_F(x)$	CSC(x)	†
$\operatorname{arcsin}_F(x)$	ASIN(x)	*
$\operatorname{arccos}_F(x)$	ACOS(x)	*
$\operatorname{arctan}_F(x)$	ATAN(x)	*
$\operatorname{arccot}_F(x)$	ACOT(x)	†
$\operatorname{arctg}_F(x)$	ACTG(x)	†
$\operatorname{arcsec}_F(x)$	ASEC(x)	†
$\operatorname{arccsc}_F(x)$	ACSC(x)	†
$\operatorname{arc}_F(x, y)$	ATAN2(y, x)	*
$\operatorname{cycle}_F(u, x)$	CYCLE(u, x)	†
$\operatorname{sinu}_F(u, x)$	SINU(u, x)	†
$\operatorname{cosu}_F(u, x)$	COSU(u, x)	†
$\operatorname{tanu}_F(u, x)$	TANU(u, x)	†
$\operatorname{cotu}_F(u, x)$	COTU(u, x)	†
$\operatorname{secu}_F(u, x)$	SECU(u, x)	†
$\operatorname{cscu}_F(u, x)$	CSCU(u, x)	†
$\operatorname{arcsinu}_F(u, x)$	ASINU(u, x)	†
$\operatorname{arccosu}_F(u, x)$	ACOSU(u, x)	†

$arctanu_F(u, x)$	ATANU(u, x)	†
$arccotu_F(u, x)$	ACOTU(u, x)	†
$arcctgu_F(u, x)$	ACTGU(u, x)	†
$arcsecu_F(u, x)$	ASECU(u, x)	†
$arccscu_F(u, x)$	ACSCU(u, x)	†
$arcu_F(u, x, y)$	ATAN2U(u, y, x)	†
$cycle_F(360, x)$	DEGREES(x)	†
$sinu_F(360, x)$	SIND(x)	†
$cosu_F(360, x)$	COSD(x)	†
$tanu_F(360, x)$	TAND(x)	†
$cotu_F(360, x)$	COTD(x)	†
$secu_F(360, x)$	SECD(x)	†
$cscu_F(360, x)$	CSCD(x)	†
$arcsinu_F(360, x)$	ASIND(x)	†
$arccosu_F(360, x)$	ACOSD(x)	†
$arctanu_F(360, x)$	ATAND(x)	†
$arccotu_F(360, x)$	ACOTD(x)	†
$arcctgu_F(360, x)$	ACTGD(x)	†
$arcsecu_F(360, x)$	ASECD(x)	†
$arccscu_F(360, x)$	ACSCD(x)	†
$arcu_F(360, x, y)$	ATAN2D(y, x)	†
$rad_to_cycle_F(x, u)$	RAD_TO_CYCLE(x, u)	†
$cycle_to_rad_F(u, x)$	CYCLE_TO_RAD(u, x)	†
$cycle_to_cycle_F(u, x, v)$	CYCLE_TO_CYCLE(u, x, v)	†

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*

Arithmetic value conversions in Fortran are always explicit, and the conversion function is named like the target type, except when converting to/from strings.

$convert_{I \rightarrow I'}(x)$	INT($x, kind$)	*
	<i>lbla</i> FORMAT (Bn)	*(binary)
$convert_{I'' \rightarrow I}(f)$	READ (UNIT= f , FMT= $lbla$) r	*
$convert_{I \rightarrow I''}(x)$	WRITE (UNIT= h , FMT= $lbla$) x	*
	<i>lblb</i> FORMAT (On)	*(octal)
$convert_{I'' \rightarrow I}(f)$	READ (UNIT= f , FMT= $lblb$) r	*
$convert_{I \rightarrow I''}(x)$	WRITE (UNIT= h , FMT= $lblb$) x	*
	<i>lblc</i> FORMAT (In)	*(decimal)
$convert_{I'' \rightarrow I}(f)$	READ (UNIT= f , FMT= $lblc$) r	*
$convert_{I \rightarrow I''}(x)$	WRITE (UNIT= h , FMT= $lblc$) x	*
	<i>lbld</i> FORMAT (Zn)	*(hexadecimal)
$convert_{I'' \rightarrow I}(f)$	READ (UNIT= f , FMT= $lbld$) r	*
$convert_{I \rightarrow I''}(x)$	WRITE (UNIT= h , FMT= $lbld$) x	*
$rounding_{F \rightarrow I}(y)$	ROUND($y, kind?$)	†
$floor_{F \rightarrow I}(y)$	FLOOR($y, kind?$)	*
$ceiling_{F \rightarrow I}(y)$	CEILING($y, kind?$)	*

$convert_{I \rightarrow F}(x)$	REAL($x, kind$) or sometimes DBLE(x)	*
$convert_{F \rightarrow F'}(y)$	REAL($y, kind$) or sometimes DBLE(y)	*
	<i>blb</i> FORMAT ($Fw.d$)	*
	<i>blf</i> FORMAT ($Dw.d$)	*
	<i>blg</i> FORMAT ($Ew.d$)	*
	<i>blh</i> FORMAT ($Ew.dEe$)	*
	<i>bli</i> FORMAT ($ENw.d$)	*
	<i>blj</i> FORMAT ($ENw.dEe$)	*
	<i>blk</i> FORMAT ($ESw.d$)	*
	<i>bll</i> FORMAT ($ESw.dEe$)	*
$convert_{F'' \rightarrow F}(f)$	READ (UNIT= f , FMT= $lblx$) t	*
$convert_{F \rightarrow F''}(y)$	WRITE (UNIT= h , FMT= $lblx$) y	*
$convert_{D' \rightarrow F}(f)$	READ (UNIT= f , FMT= $lblx$) t	*

where x is an expression of type *INT*, y is an expression of type *FLT*, and z is an expression of type *FXD*, where *FXD* is a fixed point type. *INT2* is the integer datatype that corresponds to I' .

Fortran provides base 10 non-negative numerals for all of its integer and floating point types. Numerals for floating point types must have a '.' in them. The KIND of the a numeral is indicated by a suffix. The details are not repeated in this example binding, see ISO/IEC 1539-1, clause 4.3.1.1 Integer type, and clause 4.3.1.2 Real type.

Fortran does not specify numerals for infinities and NaNs. Suggestion:

$+\infty$	INFINITY	†
qNaN	NAN	†
sNaN	SIGNAN	†

as well as string formats for reading and writing these values as character strings.

Fortran provides recording of indicators for floating point arithmetic notifications, the LIA-2 preferred method. See ISO/IEC TR 15580:1998, *Information technology – Programming languages – Fortran – Floating-point exception handling* [24]. **absolute_precision_underflow** notifications are however ignored.

C.7 Haskell

The programming language Haskell is defined by *Report on the programming language Haskell 98* [66], together with *Standard libraries for the Haskell 98 programming panguage* [67].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

The Haskell datatype `Bool` corresponds to the LIA datatype **Boolean**.

Every implementation of Haskell has at least two integer datatypes **Integer**, which is unlimited, and **Int**, and at least two floating point datatypes, **Float**, and **Double**. The notation *INT* is used to stand for the name of one of the integer datatypes, and *FLT* is used to stand for the name of one of the floating point datatypes in what follows.

The LIA-2 integer operations are listed below, along with the syntax used to invoke them:

$max_I(x, y)$	<code>max x y</code> or <code>x 'max' y</code>	*
$min_I(x, y)$	<code>min x y</code> or <code>x 'min' y</code>	*
$max_seq_I(xs)$	<code>maximum xs</code>	*
$min_seq_I(xs)$	<code>minimum xs</code>	*
$dim_I(x, y)$	<code>dim x y</code> or <code>x 'dim' y</code>	†
$power_I(x, y)$	<code>x ^ y</code> or <code>(^) x y</code>	*
$shift2_I(x, y)$	<code>shift2 x y</code> or <code>x 'shift2' y</code>	†
$shift10_I(x, y)$	<code>shift10 x y</code> or <code>x 'shift10' y</code>	†
$sqr_I(x)$	<code>isqrt x</code>	†
$divides_I(x, y)$	<code>divides x y</code> or <code>x 'divides' y</code>	†
$even_I(x)$	<code>even x</code>	*
$odd_I(x)$	<code>odd x</code>	*
$div_I(x, y)$	<code>div x y</code> or <code>x 'div' y</code>	*
$mod_I(x, y)$	<code>mod x y</code> or <code>x 'mod' y</code>	*
$group_I(x, y)$	<code>grp x y</code> or <code>x 'grp' y</code>	†
$pad_I(x, y)$	<code>pad x y</code> or <code>x 'pad' y</code>	†
$quot_I(x, y)$	<code>ratio x y</code> or <code>x 'ratio' y</code>	†
$remr_I(x, y)$	<code>remainder x y</code> or <code>x 'remainder' y</code>	†
$gcd_I(x, y)$	<code>gcd x y</code> or <code>x 'gcd' y</code>	*
$lcm_I(x, y)$	<code>lcm x y</code> or <code>x 'lcm' y</code>	*
$gcd_seq_I(xs)$	<code>gcd_seq xs</code>	†
$lcm_seq_I(xs)$	<code>lcm_seq xs</code>	†
$add_wrap_I(x, y)$	<code>x +: y</code>	†
$add_ov_I(x, y)$	<code>x ++ y</code>	†
$sub_wrap_I(x, y)$	<code>x -: y</code>	†
$sub_ov_I(x, y)$	<code>x -:+ y</code>	†
$mul_wrap_I(x, y)$	<code>x *: y</code>	†
$mul_ov_I(x, y)$	<code>x *:+ y</code>	†

where x and y are expressions of type *INT* and where xs is an expression of type [*INT*].

The LIA-2 non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$max_F(x, y)$	<code>max x y</code> or <code>x 'max' y</code>	*
$min_F(x, y)$	<code>min x y</code> or <code>x 'min' y</code>	*
$mmax_F(x, y)$	<code>mmax x y</code> or <code>x 'mmax' y</code>	†
$mmin_F(x, y)$	<code>mmin x y</code> or <code>x 'mmin' y</code>	†
$max_seq_F(xs)$	<code>maximum xs</code>	*
$min_seq_F(xs)$	<code>minimum xs</code>	*
$mmax_seq_F(xs)$	<code>mmaximum xs</code>	†
$mmin_seq_F(xs)$	<code>mminimum xs</code>	†
$dim_F(x, y)$	<code>dim x y</code> or <code>x 'dim' y</code>	†
$rounding_F(x)$	<code>fromInteger (round x)</code>	*
$floor_F(x)$	<code>fromInteger (floor x)</code>	*
$ceiling_F(x)$	<code>fromInteger (ceiling x)</code>	*
$rounding_rest_F(x)$	<code>x - fromInteger (round x)</code>	*
$floor_rest_F(x)$	<code>x - fromInteger (floor x)</code>	*

$ceiling_rest_F(x)$	$x - \text{fromInteger } (\text{ceiling } x)$	*
$remr_F(x, y)$	$\text{remainder } x \ y \text{ or } x \text{ 'remainder' } y$	†
$sqr_t_F(x)$	$\text{sqrt } x$	*
$rsqr_t_F(x)$	$\text{rsqrt } x$	†
$add_lo_F(x, y)$	$x +:- y$	†
$sub_lo_F(x, y)$	$x -:- y$	†
$mul_lo_F(x, y)$	$x *:- y$	†
$div_rest_F(x, y)$	$x /:* y$	†
$sqr_t_rest_F(x)$	$\text{sqrt_rest } x$	†
$mul_{F \rightarrow F'}(x, y)$	$\text{prod } x \ y$	†

where x , y and z are expressions of type FLT , and where xs is an expression of type $[FLT]$.

The binding for the floor, round, and ceiling operations here take advantage of the unbounded `Integer` type in Haskell, and that `Integer` is the default integer type.

The LIA-2 parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	$\text{err_hypotenuse } x$	†
$max_err_exp_F$	$\text{err_exp } x$	†
$max_err_power_F$	$\text{err_power } x$	†
$max_err_sinh_F$	$\text{err_sinh } x$	†
$max_err_tanh_F$	$\text{err_tanh } x$	†
$big_angle_r_F$	$\text{big_radian_angle } x$	†
$max_err_sin_F$	$\text{err_sin } x$	†
$max_err_tan_F$	$\text{err_tan } x$	†
$min_angular_unit_F$	$\text{min_angle_unit } x$	†
$big_angle_u_F$	$\text{big_angle } x$	†
$max_err_sinu_F(u)$	$\text{err_sin_cycle } u$	†
$max_err_tanu_F(u)$	$\text{err_tan_cycle } u$	†
$max_err_convert_F$	$\text{err_convert } x$	†
$max_err_convert_{F'}$	$\text{err_convert } ""$	†
$max_err_convert_{D'}$	$\text{err_convert } ""$	†

where b , x and u are expressions of type FLT . Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them:

$hypot_F(x, y)$	$\text{hypotenuse } x \ y$	†
$power_{FI}(b, z)$	$b \ \wedge\wedge \ z \text{ or } (\wedge\wedge) \ b \ z$	*
$exp_F(x)$	$\text{exp } x$	*
$expm1_F(x)$	$\text{expM1 } x$	†
$exp2_F(x)$	$\text{exp2 } x$	†
$exp10_F(x)$	$\text{exp10 } x$	†
$power_F(b, y)$	$b \ ** \ y \text{ or } (**) \ b \ y$	*

$power1p_{m1_F}(b, y)$	power1PM1 <i>b y</i> or <i>b</i> 'power1PM1' <i>y</i>	†
$ln_F(x)$	log <i>x</i>	*
$ln1p_F(x)$	log1P <i>x</i>	†
$log2_F(x)$	log2 <i>x</i>	†
$log10_F(x)$	log10 <i>x</i>	†
$logbase_F(b, x)$	logBase <i>b x</i> or <i>b</i> 'logBase' <i>x</i>	*
$logbase1p1p_F(b, x)$	logBase1P1P <i>b x</i>	†
$sinh_F(x)$	sinh <i>x</i>	*
$cosh_F(x)$	cosh <i>x</i>	*
$tanh_F(x)$	tanh <i>x</i>	*
$coth_F(x)$	coth <i>x</i>	†
$sech_F(x)$	sech <i>x</i>	†
$csch_F(x)$	csch <i>x</i>	†
$arcsinh_F(x)$	asinh <i>x</i>	*
$arccosh_F(x)$	acosh <i>x</i>	*
$arctanh_F(x)$	atanh <i>x</i>	*
$arccoth_F(x)$	acoth <i>x</i>	†
$arcsech_F(x)$	asech <i>x</i>	†
$arccsch_F(x)$	acsch <i>x</i>	†
$rad_F(x)$	radians <i>x</i>	†
$axis_rad_F(x)$	axis_radians <i>x</i>	†
$sin_F(x)$	sin <i>x</i>	*
$cos_F(x)$	cos <i>x</i>	*
$tan_F(x)$	tan <i>x</i>	*
$cot_F(x)$	cot <i>x</i>	†
$sec_F(x)$	sec <i>x</i>	†
$csc_F(x)$	csc <i>x</i>	†
$cosSin_F(x)$	cosSin <i>x</i>	†
$arcsin_F(x)$	asin <i>x</i>	*
$arccos_F(x)$	acos <i>x</i>	*
$arctan_F(x)$	atan <i>x</i>	*
$arccot_F(x)$	acot <i>x</i>	†
$arcctg_F(x)$	actg <i>x</i>	†
$arcsec_F(x)$	asec <i>x</i>	†
$arccsc_F(x)$	acsc <i>x</i>	†
$arc_F(x, y)$	atan2 <i>y x</i>	*
$cycle_F(u, x)$	cycle <i>u x</i>	†
$axis_cycle_F(u, x)$	axis_cycle <i>u x</i>	†
$sinu_F(u, x)$	sinU <i>u x</i>	†
$cosu_F(u, x)$	cosU <i>u x</i>	†
$tanu_F(u, x)$	tanU <i>u x</i>	†
$cotu_F(u, x)$	cotU <i>u x</i>	†
$secu_F(u, x)$	secU <i>u x</i>	†

$cscu_F(u, x)$	<code>cscU u x</code>	†
$cos\sinu_F(x)$	<code>cosSinU u x</code>	†
$arcsinu_F(u, x)$	<code>asinU u x</code>	†
$arccosu_F(u, x)$	<code>acosU u x</code>	†
$arctanu_F(u, x)$	<code>atanU u x</code>	†
$arccotu_F(u, x)$	<code>acotU u x</code>	†
$arctgu_F(u, x)$	<code>acotU u x</code>	†
$arcsecu_F(u, x)$	<code>asecU u x</code>	†
$arccscu_F(u, x)$	<code>acscU u x</code>	†
$arcu_F(u, x, y)$	<code>atan2U u y x</code>	†
$rad_to_cycle_F(x, u)$	<code>rad_to_cycle x u</code>	†
$cycle_to_rad_F(u, x)$	<code>cycle_to_rad u x</code>	†
$cycle_to_cycle_F(u, x, v)$	<code>cycle_to_cycle u x v</code>	†

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*.

Arithmetic value conversions in Haskell are always explicit. They are done with the overloaded `fromIntegral` and `fromFractional` operations.

$convert_{I \rightarrow I'}(x)$	<code>fromIntegral x</code>	★
$convert_{I'' \rightarrow I}(x)$	<code>read s</code>	★
$convert_{I \rightarrow I''}(x)$	<code>show x</code>	★
$rounding_{F \rightarrow I}(y)$	<code>round(y)</code>	★
$floor_{F \rightarrow I}(y)$	<code>floor(y)</code>	★
$ceiling_{F \rightarrow I}(y)$	<code>ceiling(y)</code>	★
$convert_{I \rightarrow F}(x)$	<code>fromIntegral x</code>	★
$convert_{F \rightarrow F'}(y)$	<code>fromFractional y</code>	★
$convert_{F'' \rightarrow F}(s)$	<code>read s</code>	.★
$convert_{F \rightarrow F''}(y)$	<code>show y</code>	...★
$convert_{D' \rightarrow F}(s)$	<code>read s</code>	.★
$convert_{F \rightarrow D'}(y)$	<code>show y</code>	...★

where x is an expression of type *INT*, y is an expression of type *FLT*, and z is an expression of type *FXD*, where *FXD* is a fixed point type.

Haskell provides non-negative numerals for all its integer and floating point types in base 10. There is no differentiation between the numerals for different floating point types, nor between numerals for different integer types, and integer numerals can be used for floating point values. Integer numerals stand for a value in `Integer` (the unbounded integer type) and an implicit `fromIntegral` operation is applied to it. Fractional numerals stand for a value in `Rational` (the unbounded type of rational numbers) and an implicit `fromRational` operation is applied to it.

Haskell does not specify any numerals for infinities and NaNs. Suggestion:

<code>+\infty</code>	<code>infinity</code>	†
<code>qNaN</code>	<code>quietNaN</code>	†
<code>sNaN</code>	<code>sigallingNaN</code>	†

as well as string formats for reading and writing these values as character strings.

Haskell has the notion of **error**, which results in a change of ‘control flow’, which cannot be returned from, nor caught. An **error** results in the termination of the program. (There are suggestions to improve this.) **pole** for integer types and **invalid** (in general) are considered to be **error**. No notification results for **underflow**, and the continuation value (specified by LIA-2) is used directly, since recording of indicators is not available and **error** is inappropriate for **underflow**. The handling of integer **overflow** is implementation dependent. The handling of floating point **overflow** and **pole** should be to return a suitable infinity (specified by LIA-2), if possible, without any notification, since recording of indicators is not available.

C.8 Java

The programming language Java is defined by *The Java Language Specification* [65].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided. The LIA-2 operations that are provided in Java 2 (marked “★” below) are in the final class `java.lang.Math`.

The Java datatype `boolean` corresponds to the LIA datatype **Boolean**.

Every implementation of Java has the integral datatypes `int`, and `long`.

Java has two floating point datatypes, `float` and `double`, which must conform to IEC 60559.

The LIA-2 integer operations are listed below, along with the syntax used to invoke them:

$min_I(x, y)$	<code>min(x, y)</code>	★
$max_I(x, y)$	<code>max(x, y)</code>	★
$min_seq_I(xs)$	<code>min_arr(xs)</code>	†
$max_seq_I(xs)$	<code>max_arr(xs)</code>	†
$dim_I(x, y)$	<code>dim(x, y)</code>	†
$sqr_I(x)$	<code>sqr(x)</code>	†
$power_I(x, y)$	<code>power(x, y)</code>	†
$divides_I(x, y)$	<code>divides(x, y)</code>	†
$even_I(x)$	<code>x % 2 == 0</code>	★
$odd_I(x)$	<code>x % 2 != 0</code>	★
$div_I(x, y)$	<code>div(x, y)</code>	†
$moda_I(x, y)$	<code>mod(x, y)</code>	†
$group_I(x, y)$	<code>group(x, y)</code>	†
$pad_I(x, y)$	<code>pad(x, y)</code>	†
$quot_I(x, y)$	<code>quot(x, y)</code>	†
$remr_I(x, y)$	<code>rem(x, y)</code>	†
$gcd_I(x, y)$	<code>gcd(x, y)</code>	†
$lcm_I(x, y)$	<code>lcm(x, y)</code>	†
$gcd_seq_I(xs)$	<code>gcd_arr(xs)</code>	†
$lcm_seq_I(xs)$	<code>lcm_arr(xs)</code>	†
$add_wrap_I(x, y)$	<code>add_wrap(x, y)</code>	†
$add_ov_I(x, y)$	<code>add_over(x, y)</code>	†
$sub_wrap_I(x, y)$	<code>sub_wrap(x, y)</code>	†
$sub_ov_I(x, y)$	<code>sub_over(x, y)</code>	†

$mul_wrap_I(x, y)$	<code>mul_wrap(x, y)</code>	†
$mul_ov_I(x, y)$	<code>mul_over(x, y)</code>	†

where x and y are expressions of type *INT* and where xs is an expression of type `array of INT`.

The LIA-2 non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$min_F(x, y)$	<code>min(x, y)</code>	*
$max_F(x, y)$	<code>max(x, y)</code>	*
$mmax_F(x, y)$	<code>mmax(x, y)</code>	†
$mmin_F(x, y)$	<code>mmin(x, y)</code>	†
$min_seq_F(xs)$	<code>min_arr(xs)</code>	†
$max_seq_F(xs)$	<code>max_arr(xs)</code>	†
$mmax_seq_F(xs)$	<code>mmax(xs)</code>	†
$mmin_seq_F(xs)$	<code>mmin(xs)</code>	†
$rounding_F(x)$	<code>rint(x)</code>	*
$floor_F(x)$	<code>floor(x)</code>	*
$ceiling_F(x)$	<code>ceil(x)</code>	*
$dim_F(x, y)$	<code>dim(x, y)</code>	†
$dprod_{F \rightarrow F'}(x, y)$	<code>dprod(x, y)</code>	†
$remr_F(x, y)$	<code>IEEEremainder(x, y)</code>	*
$sqr_F(x)$	<code>sqr(x)</code>	*
$rsqr_F(x)$	<code>rsqr(x)</code>	†
$add_lo_F(x, y)$	<code>add_low(x, y)</code>	†
$sub_lo_F(x, y)$	<code>sub_low(x, y)</code>	†
$mul_lo_F(x, y)$	<code>mul_low(x, y)</code>	†
$div_rest_F(x, y)$	<code>div_rest(x, y)</code>	†
$sqr_rest_F(x)$	<code>sqr_rest(x)</code>	†

where x , y and z are expressions of type *FLT*, and where xs is an expression of type `array of FLT`.

The LIA-2 parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	<code>err_hypotenuse(x)</code>	†
$max_err_exp_F$	<code>err_exp(x)</code>	†
$max_err_power_F(b, x)$	<code>err_power(b, x)</code>	†
$max_err_sinh_F$	<code>err_sinh(x)</code>	†
$max_err_tanh_F$	<code>err_tanh(x)</code>	†
$big_radian_angle_F$	<code>big_radian_angle(x)</code>	†
$max_err_sin_F$	<code>err_sin(x)</code>	†
$max_err_tan_F$	<code>err_tan(x)</code>	†
$min_angular_unit_F$	<code>smallest_angular_unit(x)</code>	†
big_angle_F	<code>big_angle(x)</code>	†
$max_err_sinu_F(u)$	<code>err_sin_cycle(u)</code>	†
$max_err_tanu_F(u)$	<code>err_tan_cycle(u)</code>	†

$max_err_convert_F$	<code>err_convert(x)</code>	†
$max_err_convert_{F'}$	<code>err_convert_to_string</code>	†
$max_err_convert_{D'}$	<code>err_convert_to_string</code>	†

where b , x and u are expressions of type *FLT*. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them. These are defined only for double not for float.

$hypot_F(x, y)$	<code>hypotenuse(x, y)</code>	†
$power_{FI}(b, z)$	<code>poweri(b, z)</code>	†
$exp_F(x)$	<code>exp(x)</code>	*
$expm1_F(x)$	<code>expm1(x)</code>	†
$exp2_F(x)$	<code>exp2(x)</code>	†
$exp10_F(x)$	<code>exp10(x)</code>	†
$power_F(b, y)$	<code>power(b, y)</code>	†
$pow_F(b, y)$	<code>pow(b, y)</code>	* Not LIA-2!
$power1pm1_F(b, y)$	<code>power1pm1(b, y)</code>	†
$ln_F(x)$	<code>log(x)</code>	*
$ln1p_F(x)$	<code>log1p(x)</code>	†
$log2_F(x)$	<code>log2(x)</code>	†
$log10_F(x)$	<code>log10(x)</code>	†
$logbase_F(b, x)$	<code>log(b, x)</code>	†
$logbase1p1p_F(b, x)$	<code>log1p1p(b, x)</code>	†
$sinh_F(x)$	<code>sinh(x)</code>	†
$cosh_F(x)$	<code>cosh(x)</code>	†
$tanh_F(x)$	<code>tanh(x)</code>	†
$coth_F(x)$	<code>coth(x)</code>	†
$sech_F(x)$	<code>sech(x)</code>	†
$csch_F(x)$	<code>csch(x)</code>	†
$arcsinh_F(x)$	<code>asinh(x)</code>	†
$arccosh_F(x)$	<code>acosh(x)</code>	†
$arctanh_F(x)$	<code>atanh(x)</code>	†
$arccoth_F(x)$	<code>acoth(x)</code>	†
$arcsech_F(x)$	<code>asech(x)</code>	†
$arccsch_F(x)$	<code>acsch(x)</code>	†
$rad_F(x)$	<code>radian(x)</code>	†
$axis_rad_F(x)$	<code>axis_rad(x)</code>	†
$sin_F(x)$	<code>sin(x)</code>	*
$cos_F(x)$	<code>cos(x)</code>	*
$tan_F(x)$	<code>tan(x)</code>	*
$cot_F(x)$	<code>cot(x)</code>	†
$sec_F(x)$	<code>sec(x)</code>	†
$csc_F(x)$	<code>csc(x)</code>	†

$\text{arcsin}_F(x)$	$\text{asin}(x)$	*
$\text{arccos}_F(x)$	$\text{acos}(x)$	*
$\text{arctan}_F(x)$	$\text{atan}(x)$	*
$\text{arccot}_F(x)$	$\text{acot}(x)$	†
$\text{arcctg}_F(x)$	$\text{actg}(x)$	†
$\text{arcsec}_F(x)$	$\text{asec}(x)$	†
$\text{arccsc}_F(x)$	$\text{acsc}(x)$	†
$\text{arc}_F(x, y)$	$\text{atan2}(y, x)$	*
$\text{cycle}_F(u, x)$	$\text{cycle}(u, x)$	†
$\text{axis_cycle}_F(u, x)$	$\text{axis_cycle}(u, x)$	†
$\text{sinu}_F(u, x)$	$\text{sinu}(u, x)$	†
$\text{cosu}_F(u, x)$	$\text{cosu}(u, x)$	†
$\text{tanu}_F(u, x)$	$\text{tanu}(u, x)$	†
$\text{cotu}_F(u, x)$	$\text{cotu}(u, x)$	†
$\text{secu}_F(u, x)$	$\text{secu}(u, x)$	†
$\text{cscu}_F(u, x)$	$\text{cscu}(u, x)$	†
$\text{arcsinu}_F(u, x)$	$\text{asinu}(u, x)$	†
$\text{arccosu}_F(u, x)$	$\text{acosu}(u, x)$	†
$\text{arctanu}_F(u, x)$	$\text{atanu}(u, x)$	†
$\text{arccotu}_F(u, x)$	$\text{acotu}(u, x)$	†
$\text{arcctgu}_F(u, x)$	$\text{actgu}(u, x)$	†
$\text{arcsecu}_F(u, x)$	$\text{asecu}(u, x)$	†
$\text{arccscu}_F(u, x)$	$\text{acscu}(u, x)$	†
$\text{arcu}_F(u, x, y)$	$\text{atan2u}(u, y, x)$	†
$\text{rad_to_cycle}_F(x, u)$	$\text{radian_to_cycle}(x, u)$	†
$\text{cycle_to_rad}_F(u, x)$	$\text{cycle_to_radian}(u, x)$	†
$\text{cycle_to_cycle}_F(u, x, v)$	$\text{cycle_to_cycle}(u, x, v)$	†
$\text{rad_to_cycle}_F(x, 360)$	$\text{toDegrees}(x)$	*
$\text{cycle_to_rad}_F(360, x)$	$\text{toRadians}(x)$	*

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*

Arithmetic value conversions in Java can be explicit or implicit. The rules for when implicit conversions are applied is not repeated here. The explicit arithmetic value conversions are usually expressed as ‘casts’, except when converting to/from strings.

$\text{convert}_{I \rightarrow I'}(x)$	$(INT?)x$	*
$\text{convert}_{I'' \rightarrow I}(s)$	$\text{Integer.parseInt}(s)$	*
$\text{convert}_{I'' \rightarrow I}(s)$	$\text{Integer.parseInt}(s, radix)$	*
$\text{convert}_{I'' \rightarrow I}(s)$	$\text{Long.parseLong}(s)$	*
$\text{convert}_{I'' \rightarrow I}(s)$	$\text{Long.parseLong}(s, radix)$	*
$\text{convert}_{I \rightarrow I''}(x)$	$\text{Integer.toString}(x)$	*
$\text{convert}_{I \rightarrow I''}(x)$	$\text{Integer.toString}(x, radix)$	*
$\text{convert}_{I \rightarrow I''}(x)$	$\text{Integer.toBinaryString}(x)$	*
$\text{convert}_{I \rightarrow I''}(x)$	$\text{Integer.toOctalString}(x)$	*
$\text{convert}_{I \rightarrow I''}(x)$	$\text{Integer.toHexString}(x)$	*
$\text{convert}_{I \rightarrow I''}(x)$	$\text{Long.toString}(x)$	*
$\text{convert}_{I \rightarrow I''}(x)$	$\text{Long.toString}(x, radix)$	*
$\text{convert}_{I \rightarrow I''}(x)$	$\text{Long.toBinaryString}(x)$	*

$convert_{I \rightarrow I''}(x)$	<code>Long.toOctalString(x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>Long.toHexString(x)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>""+x</code>	★
$rounding_{F \rightarrow I}(y)$	<code>(INT)rint(y)</code>	★
$floor_{F \rightarrow I}(y)$	<code>(INT)floor(y)</code>	★
$ceiling_{F \rightarrow I}(y)$	<code>(INT)ceil(y)</code>	★
$convert_{I \rightarrow F}(x)$	<code>(FLT)x</code>	★
$convert_{F \rightarrow F'}(y)$	<code>(FLT2)y</code>	★
$convert_{F'' \rightarrow F}(s)$	<code>Float.parseFloat(s)</code>	★
$convert_{F'' \rightarrow F}(s)$	<code>Double.parseDouble(s)</code>	★
$convert_{F \rightarrow F''}(y)$	<code>Float.toString(x)</code>	★
$convert_{F \rightarrow F''}(y)$	<code>Double.toString(x)</code>	★
$convert_{D' \rightarrow F}(s)$	<code>Float.parseFloat(s)</code>	★
$convert_{D' \rightarrow F}(s)$	<code>Double.parseDouble(s)</code>	★

where x is an expression of type INT , y is an expression of type FLT , and z is an expression of type FXD , where FXD is a fixed point type. $INT2$ is the integer datatype that corresponds to I' . A ? above indicates that the parameter is optional. e is greater than 0.

Java provides non-negative numerals for all its integer and floating point types. The default base is 10, but for integers base 8 and 16 can be used too. Numerals for different integer types are distinguished by suffixes. Numerals for different floating point types are distinguished by suffix: **f** for **float**, no suffix for **double**, **l** for **long double**. Numerals for floating point types must have a '.' in them. The details are not repeated in this example binding, see *The Java Language Specification*, clause 3.10.1 Integer literals, and clause 3.10.2 Floating-point literals.

Java specifies numerals for infinities and NaNs:

+∞	<code>Float.POSITIVE_INFINITY</code>	★
+∞	<code>Double.POSITIVE_INFINITY</code>	★
-∞	<code>Float.NEGATIVE_INFINITY</code>	★
-∞	<code>Double.NEGATIVE_INFINITY</code>	★
qNaN	<code>Float.NaN</code>	★
qNaN	<code>Double.NaN</code>	★
sNaN	<code>Float.SigNaN</code>	†
sNaN	<code>Double.SigNaN</code>	†

as well as string formats for writing these values as character strings. However, infinities and NaNs cannot be converted *from* string.

Java has a notion of 'exception' that implies a non-returnable, but catchable, change of control flow. Java uses its exception mechanism as its default means of notification. Java ignores **underflow** notifications since a Java exception is inappropriate for an **underflow** notification. On **underflow** the continuation value (specified in LIA-2) is used directly without recording the **underflow** itself. Java also ignores **pole** and **overflow** notifications for floating point operations, and the continuation value (specified in LIA-2) is used directly without recording the **pole** or **overflow** itself. Java uses the exception `java.lang.ArithmeticException` for **invalid** notifications and for **pole** notifications for integer operations. Java, however, ignores **pole** and **invalid** for `log` and `sqrt`. Java uses `java.lang.NumberFormatException` for **invalid** (and **pole**) notifications for operations that convert from string. Since Java exceptions are non-returnable changes of control flow, no continuation value is provided for these notifications.

An implementation that wishes to follow LIA-2 should provide recording of indicators as an alternative means of handling numeric notifications, including those that Java ignores when the numeric notification handling mechanism is by Java exceptions. Recording of indicators is the LIA-2 preferred means of handling numeric notifications.

C.9 Common Lisp

The programming language Common Lisp is defined by ANSI X3.226-1994, *Information Technology – Programming Language – Common Lisp* [43].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

Common Lisp does not have a single datatype that corresponds to the LIA-1 datatype **Boolean**. Rather, **NIL** corresponds to **false** and **T** corresponds to **true**.

Every implementation of Common Lisp has one unbounded integer datatype. Any mathematical integer is assumed to have a representation as a Common Lisp data object, subject only to total memory limitations.

Common Lisp has four floating point types: **short-float**, **single-float**, **double-float**, and **long-float**. Not all of these floating point types must be distinct.

The additional integer operations are listed below, along with the syntax used to invoke them:

$min_I(x, y)$	(min x y)	★
$max_I(x, y)$	(max x y)	★
$min_seq_I(xs)$	(min . xs) or (min x_1 x_2 ... x_n)	★
$max_seq_I(xs)$	(max . xs) or (max x_1 x_2 ... x_n)	★
$dim_I(x, y)$	(dim x y)	†
$sqr_I(x)$	(isqrt x)	†
$power_I(x, y)$	(expt x y) (returns a rational on negative power)	★
$shift2_I(x, y)$	(shift2 x y)	†
$shift10_I(x, y)$	(shift10 x y)	†
$divides_I(x, y)$	(dividesp x y)	†
$even_I(x)$	(evenp x)	★
$odd_I(x)$	(oddp x)	★

(the **floor**, **ceiling**, and **round** can also accept floating point arguments)

	(multiple-value-bind (flr md) (floor x y))	
$divf_I(x, y)$	flr or (floor x y)	★
$moda_I(x, y)$	md or (mod x y)	★
	(multiple-value-bind (ceil pd) (ceiling x y))	
$group_I(x, y)$	ceil or (ceiling x y)	★
$pad_I(x, y)$	(- pd)	★
	(multiple-value-bind (rnd rm) (round x y))	
$quot_I(x, y)$	rnd or (round x y)	★
$remr_I(x, y)$	rm	★

$gcd_I(x, y)$	(gcd x y)	(deviation: (gcd 0 0) is 0)	*
$lcm_I(x, y)$	(lcm x y)		*
$gcd_seq_I(xs)$	(gcd . xs)	or (gcd x_1 x_2 ... x_n)	*
$lcm_seq_I(xs)$	(lcm . xs)	or (lcm x_1 x_2 ... x_n)	*
$add_wrap_I(x, y)$	(add-wrap x y)		†
$add_ov_I(x, y)$	(add-over x y)		†
$sub_wrap_I(x, y)$	(sub-wrap x y)		†
$sub_ov_I(x, y)$	(sub-over x y)		†
$mul_wrap_I(x, y)$	(mul-wrap x y)		†
$mul_ov_I(x, y)$	(mul-over x y)		†

where x and y are expressions of type *INT* and where xs is an expression of type list of *INT*.

The LIA-2 non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$min_F(x, y)$	(min x y)		*
$max_F(x, y)$	(max x y)		*
$min_seq_F(xs)$	(min . xs)	or (min x_1 x_2 ... x_n)	*
$max_seq_F(xs)$	(max . xs)	or (max x_1 x_2 ... x_n)	*
	(multiple-value-bind (flr frem) (ffloor x))		
$floor_F(x)$	(ffloor x)	or flr	*
$floor_rest_F(x)$	frem		*
	(multiple-value-bind (rnd rrem) (fround x))		
$rounding_F(x)$	(fround x)	or rnd	*
$rounding_rest_F(x)$	rrem		*
	(multiple-value-bind (cln crem) (fceiling x))		
$ceiling_F(x)$	(fceiling x)	or cln	*
$ceiling_rest_F(x)$	crem		*
$dim_F(x, y)$	(dim x y)		†
	(multiple-value-bind (rqt remainder) (fround x y))		
$remr_F(x, y)$	remainder		*
$sqr_F(x)$	(sqrt x)	(returns a complex on negative arg.)	*
$rsqr_F(x)$	(rsqrt x)		†
$dprod_{F \rightarrow F'}(x, y)$	(prod x y)		†
$add_lo_F(x, y)$	(add-low x y)		†
$sub_lo_F(x, y)$	(sub-low x y)		†
$mul_lo_F(x, y)$	(mul-low x y)		†
$div_rest_F(x, y)$	(div-rest x y)		†
$sqr_rest_F(x)$	(sqr-rest x)		†

where x , y and z are data objects of the same floating point type, and where xs is a data objects that is a list of data objects of (the same, in this binding) floating point type. Note that Common Lisp allows mixed number types in many of its operations. This example binding does not explain that in detail.

The LIA-2 parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	(err-hypotenuse x)		†
---------------------	-----------------------	--	---

<i>max_err_exp_F</i>	(err-exp <i>x</i>)	†
<i>max_err_power_F</i>	(err-power <i>x</i>)	†
<i>max_err_sinh_F</i>	(err-sinh <i>x</i>)	†
<i>max_err_tanh_F</i>	(err-tanh <i>x</i>)	†
<i>big_radian_angle_F</i>	(big-radian-angle <i>x</i>)	†
<i>max_err_sin_F</i>	(err-sin <i>x</i>)	†
<i>max_err_tan_F</i>	(err-tan <i>x</i>)	†
<i>min_angular_unit_F</i>	(minimum-angular-unit <i>x</i>)	†
<i>big_angle_u_F</i>	(big-angle <i>x</i>)	†
<i>max_err_sinu_F(u)</i>	(err-sin-cycle <i>u</i>)	†
<i>max_err_tanu_F(u)</i>	(err-tan-cycle <i>u</i>)	†
<i>max_err_convert_F</i>	(err_convert <i>x</i>)	†
<i>max_err_convert_F</i>	err-convert-to-string	†
<i>max_err_convert_D</i>	err-convert-to-string	†

where *b*, *x* and *u* are expressions of type *FLT*. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them:

<i>hypot_F(x, y)</i>	(hypotenuse <i>x y</i>)	†
<i>power_{FI}(b, z)</i>	(expt <i>b z</i>)	★
<i>exp_F(x)</i>	(exp <i>x</i>)	★
<i>exp2_F(x)</i>	(exp2 <i>x</i>)	†
<i>exp10_F(x)</i>	(exp10 <i>x</i>)	†
<i>expm1_F(x)</i>	(expm1 <i>x</i>)	†
<i>power_F(b, y)</i>	(expt <i>b y</i>) (deviation: (expt 0.0 0.0) is 1)	★
<i>power1pm1_F(b, y)</i>	(expt1pm1 <i>b y</i>)	†
<i>ln_F(x)</i>	(log <i>x</i>) (returns a complex on negative arg.)	★
<i>ln1p_F(x)</i>	(log1p <i>x</i>)	†
<i>log2_F(x)</i>	(log2 <i>x</i>)	†
<i>log10_F(x)</i>	(log10 <i>x</i>)	†
<i>logbase_F(b, x)</i>	(log <i>x b</i>) (note parameter order)	★
<i>logbase1p1p_F(b, x)</i>	(log1p <i>x b</i>)	†
<i>sinh_F(x)</i>	(sinh <i>x</i>)	★
<i>cosh_F(x)</i>	(cosh <i>x</i>)	★
<i>tanh_F(x)</i>	(tanh <i>x</i>)	★
<i>coth_F(x)</i>	(coth <i>x</i>)	†
<i>sech_F(x)</i>	(sech <i>x</i>)	†
<i>csch_F(x)</i>	(csch <i>x</i>)	†
<i>arcsinh_F(x)</i>	(asinh <i>x</i>)	★
<i>arccosh_F(x)</i>	(acosh <i>x</i>) (returns a complex when <i>x</i> < 1)	★

$\text{arctanh}_F(x)$	(<code>atanh</code> x)	(returns a complex when $ x > 1$) *
$\text{arccoth}_F(x)$	(<code>acoth</code> x)	†
$\text{arcsech}_F(x)$	(<code>asech</code> x)	†
$\text{arccsch}_F(x)$	(<code>acsch</code> x)	†
$\text{rad}_F(x)$	(<code>radians</code> x)	†
$\text{axis_rad}_F(x)$	(<code>axis_rad</code> x)	†
$\text{sin}_F(x)$	(<code>sin</code> x)	*
$\text{cos}_F(x)$	(<code>cos</code> x)	*
$\text{tan}_F(x)$	(<code>tan</code> x)	*
$\text{cot}_F(x)$	(<code>cot</code> x)	†
$\text{sec}_F(x)$	(<code>sec</code> x)	†
$\text{csc}_F(x)$	(<code>csc</code> x)	†
$\text{arcsin}_F(x)$	(<code>asin</code> x)	(returns a complex when $ x > 1$) *
$\text{arccos}_F(x)$	(<code>acos</code> x)	(returns a complex when $ x > 1$) *
$\text{arctan}_F(x)$	(<code>atan</code> x)	*
$\text{arccot}_F(x)$	(<code>acot</code> x)	†
$\text{arcctg}_F(x)$	(<code>actg</code> x)	†
$\text{arcsec}_F(x)$	(<code>asec</code> x)	†
$\text{arccsc}_F(x)$	(<code>acsc</code> x)	†
$\text{arc}_F(x, y)$	(<code>atan</code> y x)	*
$\text{cycle}_F(u, x)$	(<code>cycle</code> u x)	†
$\text{axis_cycle}_F(u, x)$	(<code>axis_cycle</code> u x)	†
$\text{sinu}_F(u, x)$	(<code>sinU</code> u x)	†
$\text{cosu}_F(u, x)$	(<code>cosU</code> u x)	†
$\text{tanu}_F(u, x)$	(<code>tanU</code> u x)	†
$\text{cotu}_F(u, x)$	(<code>cotU</code> u x)	†
$\text{secu}_F(u, x)$	(<code>secU</code> u x)	†
$\text{cscu}_F(u, x)$	(<code>cscU</code> u x)	†
$\text{arcsinu}_F(u, x)$	(<code>asinU</code> u x)	†
$\text{arccosu}_F(u, x)$	(<code>acosU</code> u x)	†
$\text{arctanu}_F(u, x)$	(<code>atanU</code> u x)	†
$\text{arccotu}_F(u, x)$	(<code>acotU</code> u x)	†
$\text{arcctgu}_F(u, x)$	(<code>actgU</code> u x)	†
$\text{arcsecu}_F(u, x)$	(<code>asecU</code> u x)	†
$\text{arccscu}_F(u, x)$	(<code>acscU</code> u x)	†
$\text{arcu}_F(u, x, y)$	(<code>atanU</code> u y x)	†
$\text{rad_to_cycle}_F(x, u)$	(<code>rad_to_cycle</code> x u)	†
$\text{cycle_to_rad}_F(u, x)$	(<code>cycle_to_rad</code> u x)	†
$\text{cycle_to_cycle}_F(u, x, v)$	(<code>cycle_to_cycle</code> u x v)	†

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*.

Arithmetic value conversions in Common Lisp are can be explicit or implicit. The rules for when implicit conversions are done is implementation defined.

$\text{convert}_{I \rightarrow I''}(x)$	(<code>format</code> <code>nil</code> " <code>~wB</code> " x)	*(binary)
$\text{convert}_{I \rightarrow I''}(x)$	(<code>format</code> <code>nil</code> " <code>~wO</code> " x)	*(octal)
$\text{convert}_{I \rightarrow I''}(x)$	(<code>format</code> <code>nil</code> " <code>~wD</code> " x)	*(decimal)

$convert_{I \rightarrow I'}(x)$	(format nil "~wX" x)	*(hexadecimal)
$convert_{I \rightarrow I'}(x)$	(format nil "~r,wR" x)	*(radix r)
$convert_{I \rightarrow I'}(x)$	(format nil "~@R" x)	*(roman numeral)
$rounding_{F \rightarrow I}(y)$	(round y)	*
$floor_{F \rightarrow I}(y)$	(floor y)	*
$ceiling_{F \rightarrow I}(y)$	(ceiling y)	*
$convert_{I \rightarrow F}(x)$	(float x kind)	*
$convert_{F \rightarrow F'}(y)$	(float y kind)	*
$convert_{F \rightarrow F''}(y)$	(format nil "~wF" y)	*
$convert_{F \rightarrow F''}(y)$	(format nil "~w, e, k, cE" y)	*
$convert_{F \rightarrow F''}(y)$	(format nil "~w, e, k, cG" y)	*
$convert_{F \rightarrow D'}(y)$	(format nil "~r, w, 0, #F" y)	*

where x is an expression of type *INT*, y is an expression of type *FLT*, and z is an expression of type *FXD*, where *FXD* is a fixed point type. Conversion from string to numeric value is in Common Lisp done via a general read procedure, which reads Common Lisp ‘S-expressions’.

Common Lisp provides non-negative numerals for all its integer and floating point types in base 10. There is no differentiation between the numerals for different floating point datatypes, nor between numerals for different integer types, and integer numerals can be used for floating point values.

Common Lisp does not specify numerals for infinities and NaNs. Suggestion:

$+\infty$	infinity- <i>FLT</i>	†
qNaN	nan- <i>FLT</i>	†
sNaN	signan- <i>FLT</i>	†

as well as string formats for reading and writing these values as character strings.

Common Lisp has a notion of ‘exception’, but it is unclear if it is used for any of the arithmetic operations for overflow or pole. However, Common Lisp has no notion of compile time type checking, and an operation can return differently typed values for different arguments. When justifiable, Common Lisp arithmetic operations returns a complex floating point value rather than giving a notification, even if the argument(s) to the operation were not complex. For instance, (`sqrt -1`) (quietly) returns a representation of $0 + i$.

C.10 ISLisp

The programming language ISLisp is defined by ISO/IEC 13816:1997, *Information technology – Programming languages, their environments and system software interfaces – Programming language ISLISP* [25].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

ISLisp does not have a single datatype that corresponds to the LIA datatype **Boolean**. Rather, NIL corresponds to **false** and T corresponds to **true**.

Every implementation of ISLisp has one unbounded integer datatype. Any mathematical integer is assumed to have a representation as a ISLisp data object, subject only to total memory limitations.

ISLisp has one floating point type required to conform to IEC 60559.

The additional integer operations are listed below, along with the syntax used to invoke them:

$min_I(x, y)$	(min x y)	*
$max_I(x, y)$	(max x y)	*
$min_seq_I(xs)$	(min . xs) or (min x_1 x_2 ... x_n)	*
$max_seq_I(xs)$	(max . xs) or (max x_1 x_2 ... x_n)	*
$dim_I(x, y)$	(dim x y)	†
$sqr_I(x)$	(isqrt x)	*
$power_I(x, y)$	(expt x y) (deviation: (expt 0 0) is 1)	*
$shift2_I(x, y)$	(shift2 x y)	†
$shift10_I(x, y)$	(shift10 x y)	†
$divides_I(x, y)$	(dividesp x y)	†
$even_I(x)$	(evenp x)	†
$odd_I(x)$	(oddp x)	†
$div_I(x, y)$	(div x y)	*
$moda_I(x, y)$	(mod x y)	*
$group_I(x, y)$	(group x y)	†
$pad_I(x, y)$	(pad x y)	†
$quot_I(x, y)$	(quot x y)	†
$remr_I(x, y)$	(remainder x y)	†
$gcd_I(x, y)$	(gcd x y) (deviation: (gcd 0 0) is 0)	*
$lcm_I(x, y)$	(lcm x y)	*
$gcd_seq_I(xs)$	(gcds xs)	†
$lcm_seq_I(xs)$	(lcms xs)	†
$add_wrap_I(x, y)$	(add_wrap x y)	†
$add_ov_I(x, y)$	(add_over x y)	†
$sub_wrap_I(x, y)$	(sub_wrap x y)	†
$sub_ov_I(x, y)$	(sub_over x y)	†
$mul_wrap_I(x, y)$	(mul_wrap x y)	†
$mul_ov_I(x, y)$	(mul_over x y)	†

where x and y are expressions of type *INT* and where xs is an expression of type list of *INT*.

The LIA-2 non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$min_F(x, y)$	(min x y)	*
$max_F(x, y)$	(max x y)	*
$mmin_F(x, y)$	(mmin x y)	†
$mmax_F(x, y)$	(mmax x y)	†
$min_seq_F(xs)$	(min . xs) or (min x_1 x_2 ... x_n)	*
$max_seq_F(xs)$	(max . xs) or (max x_1 x_2 ... x_n)	*
$mmin_seq_F(xs)$	(mmin . xs) or (mmin x_1 x_2 ... x_n)	†
$mmax_seq_F(xs)$	(mmax . xs) or (mmax x_1 x_2 ... x_n)	†
$floor_F(x)$	(float (floor x))	*

$rounding_F(x)$	(float (round x))	★
$ceiling_F(x)$	(float (ceiling x))	★
$dim_F(x, y)$	(dim $x y$)	†
$dprod_{F \rightarrow F'}(x, y)$	(prod $x y$)	†
$remr_F(x, y)$	(remainder $x y$)	†
$sqr_F(x)$	(sqrt x)	★
$rsqr_F(x)$	(rsqrt x)	†
$add_lo_F(x, y)$	(add_low $x y$)	†
$sub_lo_F(x, y)$	(sub_low $x y$)	†
$mul_lo_F(x, y)$	(mul_low $x y$)	†
$div_rest_F(x, y)$	(div_rest $x y$)	†
$sqr_rest_F(x)$	(sqr_rest x)	†

where x , y and z are data objects of the same floating point type, and where xs is an data objects that are lists of data objects of the same floating point type.

The LIA-2 parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	(err-hypotenuse x)	†
$max_err_exp_F$	(err-exp x)	†
$max_err_power_F$	(err-power x)	†
$max_err_sinh_F$	(err-sinh x)	†
$max_err_tanh_F$	(err-tanh x)	†
$big_radian_angle_F$	(big-radian-angle x)	†
$max_err_sin_F$	(err-sin x)	†
$max_err_tan_F$	(err-tan x)	†
$min_angular_unit_F$	(minimum-angular-unit x)	†
big_angle_F	(big-angle x)	†
$max_err_sinu_F(u)$	(err-sin-cycle u)	†
$max_err_tanu_F(u)$	(err-tan-cycle u)	†
$max_err_convert_F$	err-convert-to-string	†
$max_err_convert_D$	err-convert-to-string	†

where b , x and u are expressions of type *FLT*. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them:

$hypot_F(x, y)$	(hypotenuse $x y$)	†
$power_{FI}(b, z)$	(expt $b z$)	★
$exp_F(x)$	(exp x)	★
$expm1_F(x)$	(expm1 x)	†
$exp2_F(x)$	(exp2 x)	†
$exp10_F(x)$	(exp10 x)	†
$power_F(b, y)$	(expt $b y$)	★

$power1p m_1 F(b, y)$	(<code>expm1 b y</code>)	†
$ln_F(x)$	(<code>log x</code>)	*
$ln1p_F(x)$	(<code>log1p x</code>)	†
$log2_F(x)$	(<code>log2 x</code>)	†
$log10_F(x)$	(<code>log10 x</code>)	†
$logbase_F(b, x)$	(<code>logbase b x</code>)	†
$logbase1p1p_F(b, x)$	(<code>logbase1p b x</code>)	†
$sinh_F(x)$	(<code>sinh x</code>)	*
$cosh_F(x)$	(<code>cosh x</code>)	*
$tanh_F(x)$	(<code>tanh x</code>)	*
$coth_F(x)$	(<code>coth x</code>)	†
$sech_F(x)$	(<code>sech x</code>)	†
$csch_F(x)$	(<code>csch x</code>)	†
$arcsinh_F(x)$	(<code>asinh x</code>)	†
$arccosh_F(x)$	(<code>acosh x</code>)	†
$arctanh_F(x)$	(<code>atanh x</code>)	*
$arccoth_F(x)$	(<code>acoth x</code>)	†
$arcsech_F(x)$	(<code>asech x</code>)	†
$arccsch_F(x)$	(<code>acsch x</code>)	†
$axis_rad_F(x)$	(<code>axis_rad x</code>)	†
$rad_F(x)$	(<code>radians x</code>)	†
$sin_F(x)$	(<code>sin x</code>)	*
$cos_F(x)$	(<code>cos x</code>)	*
$tan_F(x)$	(<code>tan x</code>)	*
$cot_F(x)$	(<code>cot x</code>)	†
$sec_F(x)$	(<code>sec x</code>)	†
$csc_F(x)$	(<code>csc x</code>)	†
$arcsin_F(x)$	(<code>asin x</code>)	*
$arccos_F(x)$	(<code>acos x</code>)	*
$arctan_F(x)$	(<code>atan x</code>)	*
$arccot_F(x)$	(<code>acot x</code>)	†
$arcctg_F(x)$	(<code>actg x</code>)	†
$arcsec_F(x)$	(<code>asec x</code>)	†
$arccsc_F(x)$	(<code>acsc x</code>)	†
$arc_F(x, y)$	(<code>atan2 y x</code>)	*
$axis_cycle_F(u, x)$	(<code>axis_cycle u x</code>)	†
$cycle_F(u, x)$	(<code>cycle u x</code>)	†
$sinu_F(u, x)$	(<code>sinU u x</code>)	†
$cosu_F(u, x)$	(<code>cosU u x</code>)	†
$tanu_F(u, x)$	(<code>tanU u x</code>)	†
$cotu_F(u, x)$	(<code>cotU u x</code>)	†
$secu_F(u, x)$	(<code>secU u x</code>)	†
$cscu_F(u, x)$	(<code>cscU u x</code>)	†

$\text{arcsinu}_F(u, x)$	<code>(asinU u x)</code>	†
$\text{arccosu}_F(u, x)$	<code>(acosU u x)</code>	†
$\text{arctanu}_F(u, x)$	<code>(atanU u x)</code>	†
$\text{arccotu}_F(u, x)$	<code>(acotU u x)</code>	†
$\text{arccctgu}_F(u, x)$	<code>(actgU u x)</code>	†
$\text{arcsecu}_F(u, x)$	<code>(asecU u x)</code>	†
$\text{arccscu}_F(u, x)$	<code>(acscU u x)</code>	†
$\text{arcu}_F(u, x, y)$	<code>(atan2U u y x)</code>	†
$\text{rad_to_cycle}_F(x, u)$	<code>(rad_to_cycle x u)</code>	†
$\text{cycle_to_rad}_F(u, x)$	<code>(cycle_to_rad u x)</code>	†
$\text{cycle_to_cycle}_F(u, x, v)$	<code>(cycle_to_cycle u x v)</code>	†

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*.

Arithmetic value conversions in ISLisp are can be explicit or implicit. The rules for when implicit conversions are done is implementation defined.

$\text{convert}_{I \rightarrow I''}(x)$	<code>(format g "~B" x)</code>	*(binary)
$\text{convert}_{I \rightarrow I''}(x)$	<code>(format g "~O" x)</code>	*(octal)
$\text{convert}_{I \rightarrow I''}(x)$	<code>(format g "~D" x)</code>	*(decimal)
$\text{convert}_{I \rightarrow I''}(x)$	<code>(format g "~X" x)</code>	*(hexadecimal)
$\text{convert}_{I \rightarrow I''}(x)$	<code>(format g "~rR" x)</code>	*(radix r)
$\text{convert}_{I \rightarrow I''}(x)$	<code>(format-integer g x r)</code>	*(radix r)
$\text{rounding}_{F \rightarrow I}(y)$	<code>(round y)</code>	*
$\text{floor}_{F \rightarrow I}(y)$	<code>(floor y)</code>	*
$\text{ceiling}_{F \rightarrow I}(y)$	<code>(ceiling y)</code>	*
$\text{convert}_{I \rightarrow F}(x)$	<code>(float x kind)</code>	*
$\text{convert}_{F \rightarrow F'}(y)$	<code>(float y kind)</code>	*
$\text{convert}_{F \rightarrow F''}(y)$	<code>(format g "~G" y)</code>	*
$\text{convert}_{F \rightarrow F''}(y)$	<code>(format-float g y)</code>	*

where x is an expression of type *INT*, y is an expression of type *FLT*, and z is an expression of type *FXD*, where *FXD* is a fixed point type. Conversion from string to numeric value is in ISLisp done via a general read procedure, which reads ISLisp ‘S-expressions’.

ISLisp provides non-negative numerals for its integer and floating point types in base is 10.

ISLisp does not specify numerals for infinities and NaNs. Suggestion:

<code>+∞</code>	<code>infinity</code>	†
<code>qNaN</code>	<code>nan</code>	†
<code>sNaN</code>	<code>signan</code>	†

as well as string formats for reading and writing these values as character strings.

ISLisp has a notion of ‘error’ that implies a catchable, possibly returnable, change of control flow. ISLisp uses its exception mechanism as its default means of notification. ISLisp ignores **underflow** notifications. On **underflow** the continuation value (specified in LIA-2) is used directly without recording the **underflow** itself. ISLisp uses the error `domain-error` for **invalid** and some **pole** notifications, the error `arithmetic-error` for **overflow** notifications, and the error `division-by-zero` for other **pole** notifications.

An implementation that wishes to follow LIA-2 should provide recording of indicators as an alternative means of handling numeric notifications. Recording of indicators is the LIA-2 preferred means of handling numeric notifications.

C.11 Modula-2

The programming language Modula-2 is defined by ISO/IEC 10514-1:1996, *Information technology – Programming languages – Part 1: Modula-2, Base Language* [26]. An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

The Modula-2 datatype **Boolean** corresponds to the LIA datatype **Boolean**.

The additional integer operations are listed below, along with the syntax used to invoke them:

$min_I(x, y)$	<code>imin(x, y)</code>	†
$max_I(x, y)$	<code>imax(x, y)</code>	†
$min_seq_I(xs)$	<code>iminArr(xs)</code>	†
$max_seq_I(xs)$	<code>imaxArr(xs)</code>	†
$dim_I(x, y)$	<code>idim(x, y)</code>	†
$sqr_I(x)$	<code>isqrt(x)</code>	†
$power_I(x, y)$	<code>ipower(x, y)</code>	†
$divides_I(x, y)$	<code>divides(x, y)</code>	†
$even_I(x)$	<code>not odd(x)</code>	*
$odd_I(x)$	<code>odd(x)</code>	*
$div_I(x, y)$	<code>div(x, y)</code>	†
$mod_I(x, y)$	<code>x mod y</code>	*
$group_I(x, y)$	<code>group(x, y)</code>	†
$pad_I(x, y)$	<code>pad(x, y)</code>	†
$quot_I(x, y)$	<code>ratio(x, y)</code>	†
$remr_I(x, y)$	<code>residue(x, y)</code>	†
$gcd_I(x, y)$	<code>gcd(x, y)</code>	†
$lcm_I(x, y)$	<code>lcm(x, y)</code>	†
$gcd_seq_I(xs)$	<code>gcdarr(xs)</code>	†
$lcm_seq_I(xs)$	<code>lcmarr(xs)</code>	†
$add_wrap_I(x, y)$	<code>addwrap(x, y)</code>	†
$add_ov_I(x, y)$	<code>addover(x, y)</code>	†
$sub_wrap_I(x, y)$	<code>subwrap(x, y)</code>	†
$sub_ov_I(x, y)$	<code>subover(x, y)</code>	†
$mul_wrap_I(x, y)$	<code>mulwrap(x, y)</code>	†
$mul_ov_I(x, y)$	<code>mulover(x, y)</code>	†

where x and y are expressions of type *INT* and where xs is an expression of type `array [] of INT`.

The additional non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$min_F(x, y)$	<code>min(x, y)</code>	†
$max_F(x, y)$	<code>max(x, y)</code>	†

$mmin_F(x, y)$	<code>mmin(x, y)</code>	†
$mmax_F(x, y)$	<code>mmax(x, y)</code>	†
$min_seq_F(xs)$	<code>minarr(xs)</code>	†
$max_seq_F(xs)$	<code>maxarr(xs)</code>	†
$mmin_seq_F(xs)$	<code>mminarr(xs)</code>	†
$mmax_seq_F(xs)$	<code>mmaxarr(xs)</code>	†
$dim_F(x, y)$	<code>dim(x, y)</code>	†
$rounding_F(x)$	<code>rounding(x)</code>	†
$floor_F(x)$	<code>floor(x)</code>	†
$ceiling_F(x)$	<code>ceiling(x)</code>	†
$rounding_rest_F(x)$	<code>x - rounding(x)</code>	†
$floor_rest_F(x)$	<code>x - floor(x)</code>	†
$ceiling_rest_F(x)$	<code>x - ceiling(x)</code>	†
$dprod_{F \rightarrow F'}(x, y)$	<code>prod(x, y)</code>	†
$remr_F(x, y)$	<code>remainder(x, y)</code>	†
$sqr_t_F(x)$	<code>sqr_t(x)</code>	★
$rsqr_t_F(x)$	<code>rsqr_t(x)</code>	†
$add_lo_F(x, y)$	<code>addlow(x, y)</code>	†
$sub_lo_F(x, y)$	<code>sublow(x, y)</code>	†
$mul_lo_F(x, y)$	<code>mullow(x, y)</code>	†
$div_rest_F(x, y)$	<code>divrest(x, y)</code>	†
$sqr_t_rest_F(x)$	<code>sqr_trest(x)</code>	†

where x , y and z are expressions of type *FLT*, and where xs is an expression of type array [] of *FLT*.

The LIA-2 parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	<code>err_hypotenuse(x)</code>	†
$max_err_exp_F$	<code>err_exp(x)</code>	†
$max_err_power_F$	<code>err_power(x)</code>	†
$max_err_sinh_F$	<code>err_sinh(x)</code>	†
$max_err_tanh_F$	<code>err_tanh(x)</code>	†
$big_radian_angle_F$	<code>big_radian_angle(x)</code>	†
$max_err_sin_F$	<code>err_sin(x)</code>	†
$max_err_tan_F$	<code>err_tan(x)</code>	†
$min_angular_unit_F$	<code>min_angle_unit(x)</code>	†
$big_angle_u_F$	<code>big_angle(x)</code>	†
$max_err_sinu_F(u)$	<code>err_sin_cycle(u)</code>	†
$max_err_tanu_F(u)$	<code>err_tan_cycle(u)</code>	†
$max_err_convert_F$	<code>err_convert(x)</code>	†
$max_err_convert_F$	<code>err_convert_to_string</code>	†
$max_err_convert_D$	<code>err_convert_to_string</code>	†

where x and u are expressions of type *FLT*. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point

types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them:

$hypot_F(x, y)$	<code>hypotenuse(x, y)</code>	†
$power_{FI}(b, z)$	<code>powerI(b, z)</code>	†
$exp_F(x)$	<code>exp(x)</code>	*
$expm1_F(x)$	<code>expm1(x)</code>	†
$exp2_F(x)$	<code>exp2(x)</code>	†
$exp10_F(x)$	<code>exp10(x)</code>	†
$power_F(b, y)$	<code>power(b, y)</code>	*
$power1pm1_F(b, y)$	<code>power1PM1(b, y)</code>	†
$ln_F(x)$	<code>ln(x)</code>	*
$ln1p_F(x)$	<code>ln1P(x)</code>	†
$log2_F(x)$	<code>log2(x)</code>	†
$log10_F(x)$	<code>log10(x)</code>	†
$logbase_F(b, x)$	<code>log(x, b)</code>	†
$logbase1p1p_F(b, x)$	<code>log1P1P(x, b)</code>	†
$sinh_F(x)$	<code>sinh(x)</code>	†
$cosh_F(x)$	<code>cosh(x)</code>	†
$tanh_F(x)$	<code>tanh(x)</code>	†
$coth_F(x)$	<code>coth(x)</code>	†
$sech_F(x)$	<code>sech(x)</code>	†
$csch_F(x)$	<code>csch(x)</code>	†
$arcsinh_F(x)$	<code>arcsinh(x)</code>	†
$arccosh_F(x)$	<code>arccosh(x)</code>	†
$arctanh_F(x)$	<code>arctanh(x)</code>	†
$arccoth_F(x)$	<code>arccoth(x)</code>	†
$arcsech_F(x)$	<code>arcsech(x)</code>	†
$arccsch_F(x)$	<code>arccsch(x)</code>	†
$rad_F(x)$	<code>radian(x)</code>	†
$axis_rad_F(x)$	<code>axis_rad(x)</code>	†
$sin_F(x)$	<code>sin(x)</code>	*
$cos_F(x)$	<code>cos(x)</code>	*
$tan_F(x)$	<code>tan(x)</code>	*
$cot_F(x)$	<code>cot(x)</code>	†
$sec_F(x)$	<code>sec(x)</code>	†
$csc_F(x)$	<code>csc(x)</code>	†
$arcsin_F(x)$	<code>arcsin(x)</code>	*
$arccos_F(x)$	<code>arccos(x)</code>	*
$arctan_F(x)$	<code>arctan(x)</code>	*
$arccot_F(x)$	<code>arccot(x)</code>	†
$arcctg_F(x)$	<code>arcctg(x)</code>	†
$arcsec_F(x)$	<code>arcsec(x)</code>	†
$arccsc_F(x)$	<code>arccsc(x)</code>	†

$arc_F(x, y)$	$angle(x, y)$	†
$cycle_F(u, x)$	$cycle(u, x)$	†
$axis_cycle_F(u, x)$	$axis_cycle(u, x)$	†
$sinu_F(u, x)$	$sinu(u, x)$	†
$cosu_F(u, x)$	$cosu(u, x)$	†
$tanu_F(u, x)$	$tanu(u, x)$	†
$cotu_F(u, x)$	$cotu(u, x)$	†
$secu_F(u, x)$	$secu(u, x)$	†
$cscu_F(u, x)$	$cscu(u, x)$	†
$arcsinu_F(u, x)$	$arcsinu(u, x)$	†
$arccosu_F(u, x)$	$arccosu(u, x)$	†
$arctanu_F(u, x)$	$arctanu(u, x)$	†
$arccotu_F(u, x)$	$arccotu(u, x)$	†
$arcctgu_F(u, x)$	$arcctgu(u, x)$	†
$arcsecu_F(u, x)$	$arcsecu(u, x)$	†
$arccscu_F(u, x)$	$arccscu(u, x)$	†
$arcu_F(u, x, y)$	$angleu(u, x, y)$	†
$rad_to_cycle_F(x, u)$	$Radian_to_cycle(x, u)$	†
$cycle_to_rad_F(u, x)$	$Cycle_to_radian(u, x)$	†
$cycle_to_cycle_F(u, x, v)$	$Cycle_to_cycle(u, x, v)$	†

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*

Arithmetic value conversions in C are can be explicit or implicit. The rules for when implicit conversions are applied is not repeated here. The explicit arithmetic value conversions are usually expressed as ‘casts’, except when converting to/from strings.

$convert_{I \rightarrow I'}(x)$	$INT(x)$	*
$convert_{I'' \rightarrow I'}(f)$	$ReadCard(f, r)$	*
$convert_{I'' \rightarrow I}(f)$	$ReadInt(f, r)$	*
$convert_{I' \rightarrow I''}(x)$	$WriteCard(h, x)$	*
$convert_{I \rightarrow I''}(x)$	$WriteInt(h, x)$	*
$rounding_{F \rightarrow I}(y)$	$round(y)$	*
$floor_{F \rightarrow I}(y)$	$floor(y)$	*
$ceiling_{F \rightarrow I}(y)$	$ceiling(y)$	*
$convert_{I \rightarrow F}(x)$	$FLT(x)$	*
$convert_{F \rightarrow F'}(y)$	$FLT2(y)$	*
$convert_{F'' \rightarrow F}(f)$	$ReadReal(f, z)$	*
$convert_{F \rightarrow F''}(y)$	$WriteFloat(f, y, a, w)$	*
$convert_{F \rightarrow F''}(y)$	$WriteEng(h, y, a, w)$	*
$convert_{F \rightarrow F''}(y)$	$WriteReal(h, y, a, w)$	*
$convert_{D' \rightarrow F}(f)$	$ReadReal(f, z)$	*
$convert_{F \rightarrow D'}(y)$	$WriteFixed(h, y, a, w)$	*

where x is an expression of type *INT*, y is an expression of type *FLT*, and z is an expression of

type FXD , where FXD is a fixed point type. $INT2$ is the integer datatype that corresponds to I' . A ? above indicates that the parameter is optional. e is greater than 0.

Modula-2 provides base 8, 10, and 16 non-negative numerals for all its integer types, and base 10 non-negative numerals for all its floating point types. Numerals for floating point types must have a '.' in them. The details are not repeated in this example binding, see ISO/IEC 10514-1, clause 6.8.7.1 Whole Number Literals, and clause 6.8.7.2 Real Literals.

Modula-2 does not specify numerals for infinities and NaNs. Suggestion:

$+\infty$	INFINITY	†
qNaN	NAN	†
sNaN	SIGNAN	†

as well as string formats for reading and writing these values as character strings.

Modula-2 has a notion of 'exception' that implies a non-returnable, but catchable, change of control flow. Modula-2 uses its exception mechanism as its default means of notification. Modula-2 ignores **underflow** notifications since an Modula-2 exception is inappropriate for an **underflow** notification. On **underflow** the continuation value (specified in LIA-2) is used directly without recording the **underflow** itself. Modula-2 uses the exceptions **WHOLE-ZERO-DIVISION**, **WHOLE-ZERO-REMAINDER**, **NEGATIVE-SQRT-ARG**, **NONPOSITIVE-LN-ARG**, **NONPOSITIVE-POWER-ARG**, **TAN-OVERFLOW** (for pole, not overflow?), **ARCSIN-ARG-MAGNITUDE**, and **ARCCOS-ARG-MAGNITUDE** for **pole** and **invalid** notifications. The exceptions **WHOLE-OVERFLOW** and **REAL-OVERFLOW** are used for **overflow** notifications. Since Modula-2 exceptions are non-returnable changes of control flow, no continuation value is provided for these notifications.

An implementation that wishes to follow LIA-2 should provide recording of indicators as an alternative means of handling numeric notifications. Recording of indicators is the LIA-2 preferred means of handling numeric notifications.

C.12 Pascal and Extended Pascal

The programming language Pascal is defined by ISO/IEC 7185:1990, *Information technology - Programming languages - Pascal* [28]. The programming language Extended Pascal is defined in ISO/IEC 10206:1991 *Information technology - Programming languages - Extended Pascal* [29].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

The Pascal datatype **Boolean** corresponds to the LIA datatype **Boolean**.

The additional integer operations are listed below, along with the syntax used to invoke them:

$min_I(x, y)$	Imin (x, y)	†
$max_I(x, y)$	Imax (x, y)	†
$min_seq_I(xs)$	IminArr (xs)	†
$max_seq_I(xs)$	ImaxArr (xs)	†
$dim_I(x, y)$	Idim (x, y)	†
$sqr_I(x)$	Isqrt (x)	†
$power_I(x, y)$	x pow y	*(Extended Pascal)
$divides_I(x, y)$	Divides (x, y)	†
$even_I(x)$	(not Odd) (x)	*

$odd_I(x)$	$Odd(x)$	★
$divf_I(x, y)$	$Divi(x, y)$	†
$moda_I(x, y)$	$Modulo(x, y)$	†
$group_I(x, y)$	$Group(x, y)$	†
$pad_I(x, y)$	$Pad(x, y)$	†
$quot_I(x, y)$	$Ratio(x, y)$	†
$remr_I(x, y)$	$Residue(x, y)$	†
$gcd_I(x, y)$	$Gcd(x, y)$	†
$lcm_I(x, y)$	$Lcm(x, y)$	†
$gcd_seq_I(xs)$	$GcdArr(xs)$	†
$lcm_seq_I(xs)$	$LcmArr(xs)$	†
$add_wrap_I(x, y)$	$AddWrap(x, y)$	†
$add_ov_I(x, y)$	$AddOver(x, y)$	†
$sub_wrap_I(x, y)$	$SubWrap(x, y)$	†
$sub_ov_I(x, y)$	$SubOver(x, y)$	†
$mul_wrap_I(x, y)$	$MulWrap(x, y)$	†
$mul_ov_I(x, y)$	$MulOver(x, y)$	†

where x and y are expressions of type *INT* and where xs is an expression of type **array of INT**.

The LIA-2 non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$min_F(x, y)$	$Min(x, y)$	†
$max_F(x, y)$	$Max(x, y)$	†
$mmin_F(x, y)$	$MMin(x, y)$	†
$mmax_F(x, y)$	$MMax(x, y)$	†
$min_seq_F(xs)$	$MinArr(xs)$	†
$max_seq_F(xs)$	$MaxArr(xs)$	†
$mmin_seq_F(xs)$	$MMinarr(xs)$	†
$mmax_seq_F(xs)$	$MMaxarr(xs)$	†
$rounding_F(x)$	$Rounding(x)$	†
$floor_F(x)$	$Floor(x)$	†
$ceiling_F(x)$	$Ceiling(x)$	†
$dim_F(x, y)$	$Dim(x, y)$	†
$dprod_{F \rightarrow F'}(x, y)$	$Prod(x, y)$	†
$remr_F(x, y)$	$Remainder(x, y)$	†
$sqrt_F(x)$	$Sqrt(x)$	★
$rsqrt_F(x)$	$Rsqrt(x)$	†
$add_lo_F(x, y)$	$AddLow(x, y)$	†
$sub_lo_F(x, y)$	$SubLow(x, y)$	†
$mul_lo_F(x, y)$	$MulLow(x, y)$	†
$div_rest_F(x, y)$	$DivRest(x, y)$	†
$sqrt_rest_F(x)$	$SqrtRest(x)$	†

where x , y and z are expressions of type *FLT*, and where xs is an expression of type **array of FLT**.

The LIA-2 parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	$Err_hypotenuse(x)$	†
---------------------	----------------------	---

$max_err_exp_F$	$Err_exp(x)$	†
$max_err_power_F$	$Err_power(x)$	†
$max_err_sinh_F$	$Err_sinh(x)$	†
$max_err_tanh_F$	$Err_tanh(x)$	†
$big_radian_angle_F$	$Big_radian_angle(x)$	†
$max_err_sin_F$	$Err_sin(x)$	†
$max_err_tan_F$	$Err_tan(x)$	†
$min_angular_unit_F$	$Min_angle_unit(x)$	†
big_angle_F	$Big_angle(x)$	†
$max_err_sinu_F(u)$	$Err_sin_cycle(u)$	†
$max_err_tanu_F(u)$	$Err_tan_cycle(u)$	†
$max_err_convert_F$	$Err_convert(x)$	†
$max_err_convert_{F'}$	$Err_convert_to_string$	†
$max_err_convert_{D'}$	$Err_convert_to_string$	†

where x and u are expressions of type *FLT*. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them:

$hypot_F(x, y)$	$Hypotenuse(x, y)$	†
$power_{FI}(b, z)$	$b \text{ pow } z$	★(Extended Pascal)
$exp_F(x)$	$Exp(x)$	★
$expm1_F(x)$	$ExpM1(x)$	†
$exp2_F(x)$	$Exp2(x)$	†
$exp10_F(x)$	$Exp10(x)$	†
$power_F(b, y)$	$b ** y$	★(Extended Pascal)
$power1pm1_F(b, y)$	$Power1PM1(b, y)$	†
$ln_F(x)$	$Ln(x)$	★
$ln1p_F(x)$	$Ln1P(x)$	†
$log2_F(x)$	$Log2(x)$	†
$log10_F(x)$	$Log10(x)$	†
$logbase_F(b, x)$	$Log(x, b)$	†
$logbase1p1p_F(b, x)$	$Log1P1P(x, b)$	†
$sinh_F(x)$	$Sinh(x)$	†
$cosh_F(x)$	$Cosh(x)$	†
$tanh_F(x)$	$Tanh(x)$	†
$coth_F(x)$	$Coth(x)$	†
$sech_F(x)$	$Sech(x)$	†
$csch_F(x)$	$Csch(x)$	†
$arcsinh_F(x)$	$Arcsinh(x)$	†
$arccosh_F(x)$	$Arccosh(x)$	†

$\operatorname{arctanh}_F(x)$	$\operatorname{Arctanh}(x)$	†
$\operatorname{arccoth}_F(x)$	$\operatorname{Arccoth}(x)$	†
$\operatorname{arcsech}_F(x)$	$\operatorname{Arcsech}(x)$	†
$\operatorname{arccsch}_F(x)$	$\operatorname{Arccsch}(x)$	†
$\operatorname{rad}_F(x)$	$\operatorname{Radian}(x)$	†
$\operatorname{axis_rad}_F(x)$	$\operatorname{Axis_Radian}(x, h, v)$	†
$\operatorname{sin}_F(x)$	$\operatorname{Sin}(x)$	★
$\operatorname{cos}_F(x)$	$\operatorname{Cos}(x)$	★
$\operatorname{tan}_F(x)$	$\operatorname{Tan}(x)$	†
$\operatorname{cot}_F(x)$	$\operatorname{Cot}(x)$	†
$\operatorname{sec}_F(x)$	$\operatorname{Sec}(x)$	†
$\operatorname{csc}_F(x)$	$\operatorname{Csc}(x)$	†
$\operatorname{arcsin}_F(x)$	$\operatorname{Arcsin}(x)$	†
$\operatorname{arccos}_F(x)$	$\operatorname{Arccos}(x)$	†
$\operatorname{arctan}_F(x)$	$\operatorname{Arctan}(x)$	★
$\operatorname{arccot}_F(x)$	$\operatorname{Arccot}(x)$	†
$\operatorname{arctg}_F(x)$	$\operatorname{Arccot}(x)$	†
$\operatorname{arcsec}_F(x)$	$\operatorname{Arcsec}(x)$	†
$\operatorname{arccsc}_F(x)$	$\operatorname{Arccsc}(x)$	†
$\operatorname{arc}_F(x, y)$	$\operatorname{Angle}(x, y)$	†
$\operatorname{cycle}_F(u, x)$	$\operatorname{Cycle}(u, x)$	†
$\operatorname{axis_cycle}_F(u, x)$	$\operatorname{Axis_Cycle}(u, x, h, v)$	†
$\operatorname{sinu}_F(u, x)$	$\operatorname{SinU}(u, x)$	†
$\operatorname{cosu}_F(u, x)$	$\operatorname{CosU}(u, x)$	†
$\operatorname{tanu}_F(u, x)$	$\operatorname{TanU}(u, x)$	†
$\operatorname{cotu}_F(u, x)$	$\operatorname{CotU}(u, x)$	†
$\operatorname{secu}_F(u, x)$	$\operatorname{SecU}(u, x)$	†
$\operatorname{cscu}_F(u, x)$	$\operatorname{CscU}(u, x)$	†
$\operatorname{arcsinu}_F(u, x)$	$\operatorname{ArcsinU}(u, x)$	†
$\operatorname{arccosu}_F(u, x)$	$\operatorname{ArccosU}(u, x)$	†
$\operatorname{arctanu}_F(u, x)$	$\operatorname{ArctanU}(u, x)$	†
$\operatorname{arccotu}_F(u, x)$	$\operatorname{ArccotU}(u, x)$	†
$\operatorname{arctgu}_F(u, x)$	$\operatorname{ArccotU}(u, x)$	†
$\operatorname{arcsecu}_F(u, x)$	$\operatorname{ArcsecU}(u, x)$	†
$\operatorname{arccscu}_F(u, x)$	$\operatorname{ArccscU}(u, x)$	†
$\operatorname{arcu}_F(u, x, y)$	$\operatorname{AngleU}(u, x, y)$	†
$\operatorname{rad_to_cycle}_F(x, u)$	$\operatorname{RadianToCycle}(x, u)$	†
$\operatorname{cycle_to_rad}_F(u, x)$	$\operatorname{CycleToRadian}(u, x)$	†
$\operatorname{cycle_to_cycle}_F(u, x, v)$	$\operatorname{CycleToCycle}(u, x, v)$	†

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*

Arithmetic value conversions in C are can be explicit or implicit. The rules for when implicit conversions are applied is not repeated here. The explicit arithmetic value conversions are usually expressed as ‘casts’, except when converting to/from strings.

$convert_{I'' \rightarrow I}(f)$	<code>read(f?, r)</code>	★
$convert_{I \rightarrow I''}(x)$	<code>write(h?, x:n?)</code>	★
$rounding_{F \rightarrow I}(y)$	<code>round(y)</code>	★
$floor_{F \rightarrow I}(y)$	<code>floor(y)</code>	†
$ceiling_{F \rightarrow I}(y)$	<code>ceiling(y)</code>	†
$convert_{F'' \rightarrow F}(f)$	<code>read(f?, m)</code>	★
$convert_{F \rightarrow F''}(y)$	<code>write(h?, y:i)</code>	★
$convert_{D' \rightarrow F}(f)$	<code>read(f?, m)</code>	★
$convert_{F \rightarrow D'}(y)$	<code>write(h?, y:i:a)</code>	★

where x is an expression of type *INT*, y is an expression of type *FLT*, and z is an expression of type *FXD*, where *FXD* is a fixed point type. *INT2* is the integer datatype that corresponds to I' . A ? above indicates that the parameter is optional. e is greater than 0.

Pascal provides base 10 non-negative numerals for its only integer type and only floating point type. Numerals for floating point types must have a ‘.’ in them. The details are not repeated in this example binding, see ISO/IEC FDIS 9899, clause xxxxxx, and clause yyyy.

Pascal does not specify numerals for infinities and NaNs. Suggestion:

+∞	INFINITY	†
qNaN	NAN	†
sNaN	SIGNAN	†

as well as string formats for reading and writing these values as character strings.

Pascal has the notion of ‘error’, which results in a change of ‘control flow’, which cannot be returned from, nor caught. An ‘error’ results in the termination of the program. **pole** for integer types and **invalid** (in general) are considered to be **error**. No notification results for **underflow**, and the continuation value (specified by LIA-2) is used directly, since recording of indicators is not available and ‘error’ is inappropriate for **underflow**. The handling of integer **overflow** is implementation dependent. The handling of floating point **overflow** and **pole** should be to return a suitable infinity (specified by LIA-2), if possible, without any notification, since recording of indicators is not available.

An implementation that wishes to follow LIA-2 should provide recording of indicators as an alternative means of handling numeric notifications. Recording of indicators is the LIA-2 preferred means of handling numeric notifications.

C.13 PL/I

The programming language PL/I is defined by ANSI X3.53-1976 (R1998), *Programming languages – PL/I* [44], and endorsed by ISO 6160:1979, *Programming languages – PL/I* [30]. The programming language General Purpose PL/I is defined by ISO/IEC 6522:1992, *Information technology – Programming languages – PL/I general-purpose subset* [31], also: ANSI X3.74-1987 (R1998).

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

The LIA datatype **Boolean** is implemented in the PL/I datatype BIT(1) (1 = **true** and 0 = **false**).

An implementation of PL/I provides at least one integer data type, and at least one floating point data type. The attribute **FIXED**($n,0$) selects a signed integer datatype with at least n (decimal or binary) digits of storage. The attribute **FLOAT**(k) selects a floating point datatype with at least n (decimal or binary) digits of precision.

The LIA-2 integer operations are listed below, along with the syntax used to invoke them:

$min_I(x, y)$	<code>min(x, y)</code>	*
$max_I(x, y)$	<code>max(x, y)</code>	*
$min_seq_I(xs)$	<code>min(xs[1], xs[2], ..., xs[n])</code>	*
$max_seq_I(xs)$	<code>max(xs[1], xs[2], ..., xs[n])</code>	*
$dim_I(x, y)$	<code>dim(x, y)</code>	†
$sqr_I(x)$	<code>sqr(x)</code>	†
$power_I(x, y)$	<code>x ** y</code>	*
$divides_I(x, y)$	<code>divides(x, y)</code>	†
$even_I(x)$	<code>mod(x) = 0</code>	*
$odd_I(x)$	<code>mod(x) /= 0</code>	*
$div_I(x, y)$	<code>divi(x, y)</code>	†
$moda_I(x, y)$	<code>mod(x, y)</code>	*
$group_I(x, y)$	<code>group(x, y)</code>	†
$pad_I(x, y)$	<code>pad(x, y)</code>	†
$quot_I(x, y)$	<code>ratio(x, y)</code>	†
$remr_I(x, y)$	<code>residue(x, y)</code>	†
$gcd_I(x, y)$	<code>gcd(x, y)</code>	†
$lcm_I(x, y)$	<code>lcm(x, y)</code>	†
$gcd_seq_I(xs)$	<code>gcd(xs)</code>	†
$lcm_seq_I(xs)$	<code>lcm(xs)</code>	†
$add_wrap_I(x, y)$	<code>add_wrap(x, y)</code>	†
$add_ov_I(x, y)$	<code>add_over(x, y)</code>	†
$sub_wrap_I(x, y)$	<code>sub_wrap(x, y)</code>	†
$sub_ov_I(x, y)$	<code>sub_over(x, y)</code>	†
$mul_wrap_I(x, y)$	<code>mul_wrap(x, y)</code>	†
$mul_ov_I(x, y)$	<code>mul_over(x, y)</code>	†

where x and y are expressions of type *INT* and where xs is an expression of type **array** of *INT*.

The LIA-2 non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$min_F(x, y)$	<code>min(x, y)</code>	*
$max_F(x, y)$	<code>max(x, y)</code>	*
$min_seq_F(xs)$	<code>min(xs[1], xs[2], ..., xs[n])</code>	*
$max_seq_F(xs)$	<code>max(xs[1], xs[2], ..., xs[n])</code>	*
$rounding_F(x)$	<code>round(x)</code>	*
$floor_F(x)$	<code>floor(x)</code>	†
$ceiling_F(x)$	<code>ceil(x)</code>	†
$dim_F(x, y)$	<code>dim(x, y)</code>	†
$dprod_{F \rightarrow F'}(x, y)$	<code>prod(x, y)</code>	†

$remr_F(x, y)$	<code>remainder(x, y)</code>	†
$sqr_F(x)$	<code>sqrt(x)</code>	★
$rsqr_F(x)$	<code>rsqrt(x)</code>	†
$add_{lo_F}(x, y)$	<code>add_low(x, y)</code>	†
$sub_{lo_F}(x, y)$	<code>sub_low(x, y)</code>	†
$mul_{lo_F}(x, y)$	<code>mul_low(x, y)</code>	†
$div_{rest_F}(x, y)$	<code>div_rest(x, y)</code>	†
$sqr_{rest_F}(x)$	<code>sqrt_rest(x)</code>	†

where x , y and z are expressions of type *FLT*, and where xs is an expression of type `array` of *FLT*.

The parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	<code>err_hypotenuse(x)</code>	†
$max_err_exp_F$	<code>err_exp(x)</code>	†
$max_err_power_F$	<code>err_power(x)</code>	†
$max_err_sinh_F$	<code>err_sinh(x)</code>	†
$max_err_tanh_F$	<code>err_tanh(x)</code>	†
$big_radian_angle_F$	<code>big_radian_angle(x)</code>	†
$max_err_sin_F$	<code>err_sin(x)</code>	†
$max_err_tan_F$	<code>err_tan(x)</code>	†
$min_angular_unit_F$	<code>min_angle_unit(x)</code>	†
big_angle_F	<code>big_angle(x)</code>	†
$max_err_sinu_F(u)$	<code>err_sin_cycle(u)</code>	†
$max_err_tanu_F(u)$	<code>err_tan_cycle(u)</code>	†
$max_err_convert_F$	<code>err_convert_to_string</code>	†
$max_err_convert_D$	<code>err_convert_to_string</code>	†

where x and u are expressions of type *FLT*. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them:

$hypot_F(x, y)$	<code>hypotenuse(x, y)</code>	†
$power_{FI}(b, z)$	<code>poweri(b, z)</code>	†
$exp_F(x)$	<code>exp(x)</code>	★
$expm1_F(x)$	<code>expm1(x)</code>	†
$exp2_F(x)$	<code>exp2(x)</code>	†
$exp10_F(x)$	<code>exp10(x)</code>	†
$power_F(b, y)$	<code>power(b, y)</code>	†
$power1pm1_F(b, y)$	<code>power1pm1(b, y)</code>	†
$ln_F(x)$	<code>log(x)</code>	★
$ln1p_F(x)$	<code>log1p(x)</code>	†
$log2_F(x)$	<code>log2(x)</code>	★

$\log_{10_F}(x)$	$\log_{10}(x)$	*
$\log_{base_F}(b, x)$	$\log(b, x)$	†
$\log_{base1p1p_F}(b, x)$	$\log_{1p1p}(b, x)$	†
$\sinh_F(x)$	$\sinh(x)$	*
$\cosh_F(x)$	$\cosh(x)$	*
$\tanh_F(x)$	$\tanh(x)$	*
$\coth_F(x)$	$\coth(x)$	†
$\operatorname{sech}_F(x)$	$\operatorname{sech}(x)$	†
$\operatorname{csch}_F(x)$	$\operatorname{csch}(x)$	†
$\operatorname{arcsinh}_F(x)$	$\operatorname{arcsinh}(x)$	*
$\operatorname{arccosh}_F(x)$	$\operatorname{arccosh}(x)$	*
$\operatorname{arctanh}_F(x)$	$\operatorname{arctanh}(x)$	*
$\operatorname{arcoth}_F(x)$	$\operatorname{arcoth}(x)$	†
$\operatorname{arcsech}_F(x)$	$\operatorname{arcsech}(x)$	†
$\operatorname{arccsch}_F(x)$	$\operatorname{arccsch}(x)$	†
$\operatorname{rad}_F(x)$	$\operatorname{rad}(x)$	†
$\operatorname{axis_rad}_F(x)$	$\operatorname{axis_rad}(x)$	†
$\sin_F(x)$	$\sin(x)$	*
$\cos_F(x)$	$\cos(x)$	*
$\tan_F(x)$	$\tan(x)$	*
$\cot_F(x)$	$\cot(x)$	*
$\sec_F(x)$	$\sec(x)$	†
$\csc_F(x)$	$\csc(x)$	†
$\operatorname{arcsin}_F(x)$	$\operatorname{arcsin}(x)$	*
$\operatorname{arccos}_F(x)$	$\operatorname{arccos}(x)$	*
$\operatorname{arctan}_F(x)$	$\operatorname{arctan}(x)$	*
$\operatorname{arccot}_F(x)$	$\operatorname{arccot}(x)$	†
$\operatorname{arcctg}_F(x)$	$\operatorname{arcctg}(x)$	†
$\operatorname{arcsec}_F(x)$	$\operatorname{arcsec}(x)$	†
$\operatorname{arccsc}_F(x)$	$\operatorname{arccsc}(x)$	†
$\operatorname{arc}_F(x, y)$	$\operatorname{arc}(x, y)$	*
$\operatorname{cycle}_F(u, x)$	$\operatorname{cycle}(u, x)$	†
$\operatorname{axis_cycle}_F(u, x)$	$\operatorname{axis_cycle}(u, x)$	†
$\operatorname{sinu}_F(u, x)$	$\operatorname{sin}(u, x)$	†
$\operatorname{cosu}_F(u, x)$	$\operatorname{cos}(u, x)$	†
$\operatorname{tanu}_F(u, x)$	$\operatorname{tan}(u, x)$	†
$\operatorname{cotu}_F(u, x)$	$\operatorname{cot}(u, x)$	†
$\operatorname{secu}_F(u, x)$	$\operatorname{sec}(u, x)$	†
$\operatorname{cscu}_F(u, x)$	$\operatorname{csc}(u, x)$	†
$\operatorname{arcsinu}_F(u, x)$	$\operatorname{arcsin}(u, x)$	†
$\operatorname{arccosu}_F(u, x)$	$\operatorname{arccos}(u, x)$	†
$\operatorname{arctanu}_F(u, x)$	$\operatorname{arctan}(u, x)$	*
$\operatorname{arccotu}_F(u, x)$	$\operatorname{arccot}(u, x)$	†

$arcctgu_F(u, x)$	$arcctg(u, x)$	★
$arcsecu_F(u, x)$	$arcsec(u, x)$	†
$arccscu_F(u, x)$	$arccsc(u, x)$	†
$arcu_F(u, x, y)$	$arc(u, x, y)$	†
$sinu_F(360, x)$	$sind(x)$	★
$cosu_F(360, x)$	$cosd(x)$	★
$tanu_F(360, x)$	$tand(x)$	★
$cotu_F(360, x)$	$cotd(x)$	★
$secu_F(360, x)$	$secd(x)$	†
$cscu_F(360, x)$	$cscd(x)$	†
$arcsinu_F(360, x)$	$arcsind(x)$	★
$arccosu_F(360, x)$	$arccosd(x)$	★
$arctanu_F(360, x)$	$arctand(x)$	★
$arccotu_F(360, x)$	$arccotd(x)$	★
$arcctgu_F(360, x)$	$arcctgd(x)$	★
$arcsecu_F(360, x)$	$arcsecd(x)$	†
$arccscu_F(360, x)$	$arccscd(x)$	†
$arcu_F(360, x, y)$	$arcd(y, x)$	†
$rad_to_cycle_F(x, u)$	$rad_to_cycle(x, u)$	†
$cycle_to_rad_F(u, x)$	$cycle_to_rad(u, x)$	†
$cycle_to_cycle_F(u, x, v)$	$cycle_to_cycle(u, x, v)$	†

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*.

Arithmetic value conversions in PL/I are can be explicit or implicit. The rules for when implicit conversions are applied is not repeated here. The explicit arithmetic value conversions are usually expressed as ‘casts’, except when converting to/from strings.

$convert_{I \rightarrow I'}(x)$	$FIXED(x, p)$	★
$convert_{I'' \rightarrow I}(f)$	$GET FILE (f)? EDIT (r) (F(w));$	★
$convert_{I \rightarrow I''}(x)$	$PUT FILE (h)? EDIT (x) (F(w));$	★
$rounding_{F \rightarrow I}(y)$	$FIXED(ROUND(y, 0), p)$	★
$floor_{F \rightarrow I}(y)$	$FIXED(FLOOR(y), p)$	★
$ceiling_{F \rightarrow I}(y)$	$FIXED(CEIL(y), p)$	★
$convert_{I \rightarrow F}(x)$	$FLOAT(x, p)$	★
$convert_{I \rightarrow F}(x)$	$DECIMAL(x, p)$	★
$convert_{I \rightarrow F}(x)$	$BINARY(x, p)$	★
$convert_{F \rightarrow F'}(y)$	$FLOAT(y, p)$	★
$convert_{F \rightarrow F'}(y)$	$DECIMAL(y, p)$	★
$convert_{F \rightarrow F'}(y)$	$BINARY(y, p)$	★
$convert_{F'' \rightarrow F}(f)$	$GET FILE (f)? EDIT (t) (E(w, a));$	★
$convert_{F \rightarrow F''}(y)$	$PUT FILE (h)? EDIT (y) (E(w, a));$	★
$convert_{D' \rightarrow F}(f)$	$GET FILE (f)? EDIT (t) (F(w, a));$	★
$convert_{F \rightarrow D'}(y)$	$FIXED(y, p, a)$	★
$convert_{F \rightarrow D'}(y)$	$PUT FILE (h)? EDIT (y) (F(w, a));$	★

where x is an expression of type INT , y is an expression of type FLT , and z is an expression of type FXD , where FXD is a fixed point type. $INT2$ is the integer datatype that corresponds to I' . A ? above indicates that the parameter is optional. a is greater than 0.

PL/I provides base 10 non-negative numerals for all its integer and floating point types.

PL/I does not specify numerals for infinities and NaNs. Suggestion:

$+\infty$	INFINITY	†
qNaN	NAN	†
sNaN	SIGNAN	†

as well as string formats for reading and writing these values as character strings.

PL/I has a notion of ‘condition’ that implies a non-returnable, but catchable (in an ON-unit), change of control flow. PL/I uses its condition mechanism as its default means of notification. PL/I uses the condition UNDERFLOW for **underflow** notifications. PL/I uses the condition ZERODIVIDE for **pole** notifications, and the conditions FIXEDOVERFLOW, SIZE, and OVERFLOW for **overflow** notifications, and the exception UNDEFINED (†) for **invalid** notifications. Since PL/I exceptions are non-returnable changes of control flow, no continuation value is provided for these notifications. This is inappropriate, especially for underflow, so UNDERFLOW notifications are ignored if there is no ON-clause for UNDERFLOW in the program.

An implementation that wishes to follow LIA-2 should provide recording of indicators as an alternative means of handling numeric notifications. Recording of indicators is the LIA-2 preferred means of handling numeric notifications.

C.14 SML

The programming language SML is defined by *The Definition of Standard ML (Revised)* [68].

An implementation should follow all the requirements of LIA-2 unless otherwise specified by this language binding.

The operations or parameters marked “†” are not part of the language and should be provided by an implementation that wishes to conform to the LIA-2 for that operation. For each of the marked items a suggested identifier is provided.

The SML datatype **Boolean** corresponds to the LIA datatype **Boolean**.

Every implementation of SML has at least one integer datatype, **int**, and at least one floating point datatype, **real**. The notation INT is used to stand for the name of one of the integer datatypes, and FLT is used to stand for the name of one of the floating point datatypes in what follows.

The LIA-2 integer operations are listed below, along with the syntax used to invoke them:

$min_I(x, y)$	$x \text{ min } y$ or op min (x, y)	*
$max_I(x, y)$	$x \text{ max } y$ or op max (x, y)	*
$min_seq_I(xs)$	minimum xs	†
$max_seq_I(xs)$	maximum xs	†
$dim_I(x, y)$	$x \text{ dim } y$ or op dim (x, y)	†
$sqr_I(x)$	isqrt x	†
$power_I(x, y)$	$x \text{ pow } y$ or op pow (x, y)	†
$divides_I(x, y)$	divides (x, y)	†
$even_I(x)$	even x	†
$odd_I(x)$	odd x	†
$divf_I(x, y)$	$x \text{ div } y$ or op div (x, y)	*

$mod_I(x, y)$	$x \text{ mod } y$ or $op \text{ mod } (x, y)$	*
$group_I(x, y)$	$group (x, y)$	†
$pad_I(x, y)$	$pad (x, y)$	†
$quot_I(x, y)$	$ratio (x, y)$	†
$remr_I(x, y)$	$residue (x, y)$	†
$gcd_I(x, y)$	$gcd (x, y)$	*
$lcm_I(x, y)$	$lcm (x, y)$	*
$gcd_seq_I(xs)$	$gcd_seq xs$	†
$lcm_seq_I(xs)$	$lcm_seq xs$	†
$add_wrap_I(x, y)$	$x +: y$	†
$add_ov_I(x, y)$	$x ++ y$	†
$sub_wrap_I(x, y)$	$x -: y$	†
$sub_ov_I(x, y)$	$x -:+ y$	†
$mul_wrap_I(x, y)$	$x *: y$	†
$mul_ov_I(x, y)$	$x *:+ y$	†

where x and y are expressions of type *INT* and where xs is an expression of type *INT list*.

The additional non-transcendental floating point operations are listed below, along with the syntax used to invoke them:

$min_F(x, y)$	$x \text{ min } y$ or $op \text{ min } (x, y)$	*
$max_F(x, y)$	$x \text{ max } y$ or $op \text{ max } (x, y)$	*
$mmin_F(x, y)$	$x \text{ mmin } y$ or $op \text{ mmin } (x, y)$	†
$mmax_F(x, y)$	$x \text{ mmax } y$ or $op \text{ mmax } (x, y)$	†
$min_seq_F(xs)$	$minimum xs$	†
$max_seq_F(xs)$	$maximum xs$	†
$mmin_seq_F(xs)$	$mminimum xs$	†
$mmax_seq_F(xs)$	$mmaximum xs$	†
$rounding_F(x)$	$realRound x$	†
$floor_F(x)$	$realFloor x$	*
$ceiling_F(x)$	$realCeil x$	*
$dim_F(x, y)$	$dim (x, y)$	†
$dprod_{F \rightarrow F'}(x, y)$	$prod (x, y)$	†
$remr_F(x, y)$	$remainder (x, y)$	†
$sqr_F(x)$	$sqr x$	*
$rsqr_F(x)$	$rsqr x$	†
$add_lo_F(x, y)$	$x +:- y$	†
$sub_lo_F(x, y)$	$x -:- y$	†
$mul_lo_F(x, y)$	$x *:- y$	†
$div_rest_F(x, y)$	$x /:* y$	†
$sqr_rest_F(x)$	$sqr_rest x$	†

where x , y and z are expressions of type *FLT*, and where xs is an expression of type *FLT list*.

The binding for the floor, round, and ceiling operations here take advantage of the unlimited *Integer* type in SML, and that *Integer* is the default integer type.

The parameters for operations approximating real valued transcendental functions can be accessed by the following syntax:

$max_err_hypot_F$	$err_hypotenuse x$	†
---------------------	---------------------	---

<i>max_err_exp_F</i>	<code>err_exp x</code>	†
<i>max_err_power_F</i>	<code>err_power x</code>	†
<i>max_err_sinh_F</i>	<code>err_sinh x</code>	†
<i>max_err_tanh_F</i>	<code>err_tanh x</code>	†
<i>big_angle_r_F</i>	<code>big_radian_angle x</code>	†
<i>max_err_sin_F</i>	<code>err_sin x</code>	†
<i>max_err_tan_F</i>	<code>err_tan x</code>	†
<i>min_angular_unit_F</i>	<code>min_angular_unit x</code>	†
<i>big_angle_u_F</i>	<code>big_angle x</code>	†
<i>max_err_sinu_F(u)</i>	<code>err_sin_cycle u</code>	†
<i>max_err_tanu_F(u)</i>	<code>err_tan_cycle u</code>	†
<i>max_err_convert_F</i>	<code>err_convert(x)</code>	†
<i>max_err_convert_{F'}</i>	<code>err_convert_to_string</code>	†
<i>max_err_convert_{D'}</i>	<code>err_convert_to_string</code>	†

where x and u are expressions of type *FLT*. Several of the parameter functions are constant for each type (and library), the argument is then used only to differentiate among the floating point types.

The LIA-2 elementary floating point operations are listed below, along with the syntax used to invoke them:

<i>hypot_F(x, y)</i>	<code>hypotenuse (x, y)</code>	†
<i>power_{FI}(b, z)</i>	<code>b ^^ z</code> or <code>op ^^ (x, y)</code>	†
<i>exp_F(x)</i>	<code>exp x</code>	*
<i>expm1_F(x)</i>	<code>expM1 x</code>	†
<i>exp2_F(x)</i>	<code>exp2 x</code>	†
<i>exp10_F(x)</i>	<code>exp10 x</code>	†
<i>power_F(b, y)</i>	<code>b ** y</code>	†
<i>pow_F(b, y)</i>	<code>b pow y</code> or <code>op pow (x, y)</code>	* Not LIA-2! (See C.)
<i>power1pm1_F(b, y)</i>	<code>power1PM1 (b, y)</code>	†
<i>ln_F(x)</i>	<code>ln x</code>	*
<i>ln1p_F(x)</i>	<code>ln1P x</code>	†
<i>log2_F(x)</i>	<code>log2 x</code>	†
<i>log10_F(x)</i>	<code>log10 x</code>	*
<i>logbase_F(b, x)</i>	<code>log_base (b, x)</code>	†
<i>logbase1p1p_F(b, x)</i>	<code>log_base1P1P (b, x)</code>	†
<i>sinh_F(x)</i>	<code>sinh x</code>	*
<i>cosh_F(x)</i>	<code>cosh x</code>	*
<i>tanh_F(x)</i>	<code>tanh x</code>	*
<i>coth_F(x)</i>	<code>coth x</code>	†
<i>sech_F(x)</i>	<code>sech x</code>	†
<i>csch_F(x)</i>	<code>csch x</code>	†
<i>arcsinh_F(x)</i>	<code>arcsinh x</code>	†

$\text{arccosh}_F(x)$	<code>arccosh x</code>	†
$\text{arctanh}_F(x)$	<code>arctanh x</code>	†
$\text{arccoth}_F(x)$	<code>arccoth x</code>	†
$\text{arcsech}_F(x)$	<code>arcsech x</code>	†
$\text{arccsch}_F(x)$	<code>arccsch x</code>	†
$\text{rad}_F(x)$	<code>radians x</code>	†
$\text{axis_rad}_F(x)$	<code>axis_rad x</code>	†
$\text{sin}_F(x)$	<code>sin x</code>	*
$\text{cos}_F(x)$	<code>cos x</code>	*
$\text{tan}_F(x)$	<code>tan x</code>	*
$\text{cot}_F(x)$	<code>cot x</code>	†
$\text{sec}_F(x)$	<code>sec x</code>	†
$\text{csc}_F(x)$	<code>csc x</code>	†
$\text{arcsin}_F(x)$	<code>arcsin x</code>	*
$\text{arccos}_F(x)$	<code>arccos x</code>	*
$\text{arctan}_F(x)$	<code>arctan x</code>	*
$\text{arccot}_F(x)$	<code>arccot x</code>	†
$\text{arcctg}_F(x)$	<code>arcctg x</code>	†
$\text{arcsec}_F(x)$	<code>arcsec x</code>	†
$\text{arccsc}_F(x)$	<code>arccsc x</code>	†
$\text{arc}_F(x, y)$	<code>arctan2 (y, x)</code>	*
$\text{cycle}_F(u, x)$	<code>cycle (u, x)</code>	†
$\text{axis_cycle}_F(u, x)$	<code>axis_cycle (u, x)</code>	†
$\text{sinu}_F(u, x)$	<code>sinU (u, x)</code>	†
$\text{cosu}_F(u, x)$	<code>cosU (u, x)</code>	†
$\text{tanu}_F(u, x)$	<code>tanU (u, x)</code>	†
$\text{cotu}_F(u, x)$	<code>cotU (u, x)</code>	†
$\text{secu}_F(u, x)$	<code>secU (u, x)</code>	†
$\text{cscu}_F(u, x)$	<code>cscU (u, x)</code>	†
$\text{arcsinu}_F(u, x)$	<code>arcsinU (u, x)</code>	†
$\text{arccosu}_F(u, x)$	<code>arccosU (u, x)</code>	†
$\text{arctanu}_F(u, x)$	<code>arctanU (u, x)</code>	†
$\text{arccotu}_F(u, x)$	<code>arccotU (u, x)</code>	†
$\text{arcctgu}_F(u, x)$	<code>arcctgU (u, x)</code>	†
$\text{arcsecu}_F(u, x)$	<code>arcsecU (u, x)</code>	†
$\text{arccscu}_F(u, x)$	<code>arccscU (u, x)</code>	†
$\text{arcu}_F(u, x, y)$	<code>arctan2U (u, y, x)</code>	†
$\text{rad_to_cycle}_F(x, u)$	<code>rad_to_cycle (x, u)</code>	†
$\text{cycle_to_rad}_F(u, x)$	<code>cycle_to_rad (u, x)</code>	†
$\text{cycle_to_cycle}_F(u, x, v)$	<code>cycle_to_cycle (u, x, v)</code>	†

where b , x , y , u , and v are expressions of type *FLT*, and z is an expressions of type *INT*

Type conversions in SML are always explicit.

$\text{convert}_{I \rightarrow I'}(x)$	<code>fromLarge x or toLarge x</code>	*
--	---	---

$convert_{I'' \rightarrow I}(s)$	<code>fromString s</code>	★
$convert_{I'' \rightarrow I}(s)$	<code>scan radix getc s</code>	★
$convert_{I \rightarrow I''}(x)$	<code>toString x</code>	★
$rounding_{F \rightarrow I}(y)$	<code>round y</code>	★
$rounding_{F \rightarrow I}(y)$	<code>toInt IEEEReal.TO_NEAREST y</code>	★
$rounding_{F \rightarrow I}(y)$	<code>toLargeInt IEEEReal.TO_NEAREST y</code>	★
$floor_{F \rightarrow I}(y)$	<code>floor y</code>	★
$floor_{F \rightarrow I}(y)$	<code>toInt IEEEReal.TO_NEGINF y</code>	★
$floor_{F \rightarrow I}(y)$	<code>toLargeInt IEEEReal.TO_NEGINF y</code>	★
$ceiling_{F \rightarrow I}(y)$	<code>ceiling y</code>	★
$ceiling_{F \rightarrow I}(y)$	<code>toInt IEEEReal.TO_POSINF y</code>	★
$ceiling_{F \rightarrow I}(y)$	<code>toLargeInt IEEEReal.TO_POSINF y</code>	★
$convert_{I \rightarrow F}(x)$	<code>fromInt x</code>	★
$convert_{I \rightarrow F}(x)$	<code>fromLargeInt x</code>	★
$convert_{F \rightarrow F'}(y)$	<code>toLarge y</code>	★
$convert_{F \rightarrow F'}(y)$	<code>fromLarge IEEEReal.TO_NEAREST y</code>	★
$convert_{F'' \rightarrow F}(s)$	<code>fromString s</code>	★
$convert_{F'' \rightarrow F}(s)$	<code>fromDecimal s</code>	★
$convert_{F'' \rightarrow F}(s)$	<code>scan getc s</code>	★
$convert_{F \rightarrow F''}(y)$	<code>fmt (SCI a) y</code>	★
$convert_{F \rightarrow F''}(y)$	<code>toDecimal y</code>	★
$convert_{D' \rightarrow F}(s)$	<code>fromString s</code>	★
$convert_{D' \rightarrow F}(s)$	<code>scan getc s</code>	★
$convert_{F \rightarrow D'}(y)$	<code>fmt (FIX a) y</code>	★

where x is an expression of type INT , y is an expression of type FLT , and z is an expression of type FXD , where FXD is a fixed point type. $INT2$ is the integer datatype that corresponds to I' .

SML provides non-negative base 10 numerals for all its integer and floating point types. There is no differentiation between the numerals for different floating point types, nor between numerals for different integer types, but integer numerals cannot be used for floating point values. The details are not repeated in this example binding, see *The Definition of Standard ML (Revised)* [68].

SML specifies numerals for infinities, but not NaNs:

$+\infty$	<code>posInf</code>	★
$-\infty$	<code>negInf</code>	★
qNaN	<code>NaN</code>	†
sNaN	<code>sigNaN</code>	†

An implementation wishing to conform to LIA-2 should also provide string formats for reading and writing these values as character strings.

SML has a notion of ‘exception’ that implies a non-returnable, but catchable, change of control flow. SML uses its exception mechanism as its default means of notification. SML ignores **underflow** notifications since an SML exception is inappropriate for an **underflow** notification. On **underflow** the continuation value (specified in LIA-2) is used directly without recording the **underflow** itself. SML uses the exception `Div` for **pole** notifications, the exception `Overflow` for **overflow** notifications, and the exception `Domain` for **invalid** notifications (except for `sin`, `cos`, or `tan` given an infinitary argument, where the **invalid** notification is ignored). Since SML

exceptions are non-returnable changes of control flow, no continuation value is provided for these notifications.

An implementation that wishes to follow LIA-2 should provide recording of indicators as an alternative means of handling numeric notifications. Recording of indicators is the LIA-2 preferred means of handling numeric notifications.

Annex D (informative)

Bibliography

This annex gives references to publications relevant to LIA-2.

EDITOR'S NOTE – The naming of the standards appears unsystematic, but is as taken from ISO's, IEC's and ANSI's websites.

International standards documents

- [1] ISO/IEC JTC1 Directives, Part 3: *Drafting and presentation of International Standards*, 1989.
- [2] IEC 60559:1989, *Binary floating-point arithmetic for microprocessor systems*. (Also: ANSI/IEEE Standard 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*.)
- [3] ISO/IEC 10967-1:1994, *Information technology – Language independent arithmetic – Part 1: Integer and floating point arithmetic*, (LIA-1).
- [4] ISO/IEC FCD4 10967-2, *Information technology – Language independent arithmetic – Part 2: Elementary numerical functions*, 1999, (LIA-2). (This document.)
- [5] ISO/IEC 10967-3, *Information technology – Language independent arithmetic – Part 3: Complex floating point arithmetic and complex elementary numerical functions*, (LIA-3). (To be published.)
- [6] ISO 6093:1985, *Information processing – Representation of numerical values in character strings for information interchange*.
- [7] ISO/IEC TR 10176:1998, *Information technology – Guidelines for the preparation of programming language standards*.
- [8] ISO/IEC TR 10182:1993, *Information technology – Programming languages, their environments and system software interfaces – Guidelines for language bindings*.
- [9] ISO/IEC 13886:1996, *Information technology – Language-Independent Procedure Calling*, (LIPC).
- [10] ISO/IEC 11404:1996, *Information technology – Programming languages, their environments and system software interfaces – Language-independent datatypes*, (LID).
- [11] ISO/IEC 8652:1995, *Information technology – Programming languages – Ada*.
- [12] ISO/IEC 11430:1994, *Information technology – Programming languages – Generic package of elementary functions for Ada*.
- [13] ISO/IEC 13813:1998, *Information technology – Programming languages – Generic packages of real and complex type declarations and basic operations for Ada (including vector and matrix types)*.
- [14] ISO/IEC 13814:1998, *Information technology – Programming languages – Generic package of complex elementary functions for Ada*.
- [15] ISO 8485:1989, *Programming languages – APL*.
- [16] ISO/IEC DIS 13751, *Information technology – Programming languages, their environments and system software interfaces – Programming language APL, extended*, 1999.

- [17] ISO/IEC 10279:1991, *Information technology – Programming languages – Full BASIC*. (Essentially an endorsement of ANSI X3.113-1987 (R1998) [41].)
- [18] ISO/IEC 9899:1990, *Programming languages – C*. Currently under revision: ISO/IEC FDIS 9899, 1999.
- [19] ISO/IEC 14882:1998, *Programming languages – C++*.
- [20] ISO 1989:1985, *Programming languages – COBOL*. (Endorsement of ANSI X3.23-1985 (R1991) [42].) Currently under revision (1998).
- [21] ISO/IEC 16262:1998, *Information technology – ECMAScript language specification*.
- [22] ISO/IEC 15145:1997, *Information technology – Programming languages – FORTH*. (Also: ANSI X3.215-1994.)
- [23] ISO/IEC 1539-1:1997, *Information technology – Programming languages – Fortran – Part 1: Base language*.
- [24] ISO/IEC TR 15580:1998, *Information technology – Programming languages – Fortran – Floating-point exception handling*.
- [25] ISO/IEC 13816:1997, *Information technology – Programming languages, their environments and system software interfaces – Programming language ISLISP*.
- [26] ISO/IEC 10514-1:1996, *Information technology – Programming languages – Part 1: Modula-2, Base Language*.
- [27] ISO/IEC 10514-2:1998, *Information technology – Programming languages – Part 2: Generics Modula-2*.
- [28] ISO 7185:1990, *Information technology – Programming languages – Pascal*.
- [29] ISO/IEC 10206:1991, *Information technology – Programming languages – Extended Pascal*.
- [30] ISO 6160:1979, *Programming languages – PL/I*. (Endorsement of ANSI X3.53-1976 (R1998) [44].)
- [31] ISO/IEC 6522:1992, *Information technology – Programming languages – PL/I general-purpose subset*. (Also: ANSI X3.74-1987 (R1998).)
- [32] ISO/IEC 13211-1:1995, *Information technology – Programming languages – Prolog – Part 1: General core*.
- [33] ISO/IEC 9075:1992, *Information technology – Database languages – SQL*.
- [34] ISO/IEC 8824-1:1995, *Information technology – Abstract Syntax Notation One (ASN.1) – Part 1: Specification of basic notation*.
- [35] ISO/IEC 9001:1994, *Quality systems – Model for quality assurance in design, development, production, installation and servicing*.
- [36] ISO/IEC 9126:1991, *Information technology – Software product evaluation – Quality characteristics and guidelines for their use*.
- [37] ISO/IEC 12119:1994, *Information technology – Software packages – Quality requirements and testing*.
- [38] ISO/IEC 14598-1:1999, *Information technology – Software product evaluation – Part 1: General overview*.

National standards documents

- [39] ANSI/IEEE Standard 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*.
- [40] ANSI/IEEE Standard 854-1987, *IEEE Standard for Radix-Independent Floating-Point Arithmetic*.
- [41] ANSI X3.113-1987 (R1998), *Information technology – Programming languages – Full BASIC*.
- [42] ANSI X3.23-1985 (R1991), *Programming languages – COBOL*.
- [43] ANSI X3.226-1994, *Information Technology – Programming Language – Common Lisp*.
- [44] ANSI X3.53-1976 (R1998), *Programming languages – PL/I*.
- [45] ANSI/IEEE 1178-1990, *IEEE Standard for the Scheme Programming Language*.
- [46] ANSI/NCITS 319-1998, *Information Technology – Programming Languages – Smalltalk*.

Books, articles, and other documents

- [47] J. S. Squire (ed), *Ada Letters*, vol. XI, No. 7, ACM Press (1991).
- [48] M. Abramowitz and I. Stegun (eds), *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, Tenth Printing, 1972, Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402.
- [49] J. Du Croz and M. Pont, *The Development of a Floating-Point Validation Package*, NAG Newsletter, No. 3, 1984.
- [50] J. W. Demmel and X. Li, *Faster Numerical Algorithms via Exception Handling*, 11th International Symposium on Computer Arithmetic, Winsor, Ontario, June 29 - July 2, 1993.
- [51] D. Goldberg, *What Every Computer Scientist Should Know about Floating-Point Arithmetic*. ACM Computing Surveys, Vol. 23, No. 1, March 1991.
- [52] J. R. Hauser, *Handling Floating-Point Exceptions in Numeric Programs*. ACM Transactions on Programming Languages and Systems, Vol. 18, No. 2, March 1986, Pages 139-174.
- [53] W. Kahan, *Branch Cuts for Complex Elementary Functions, or Much Ado about Nothing's Sign Bit*, Chapter 7 in *The State of the Art in Numerical Analysis* ed. by M Powell and A Iserles (1987) Oxford.
- [54] W. Kahan, *Lecture Notes on the Status of IEEE Standard 754 for Binary Floating-Point Arithmetic*, Panel Discussion of *Floating-Point Past, Present and Future*, May 23, 1995, in a series of San Francisco Bay Area Computer Historical Perspectives, sponsored by SUN Microsystems Inc.
- [55] U. Kulisch and W. L. Miranker, *Computer Arithmetic in Theory and Practice*, Academic Press, 1981.
- [56] U. Kulisch and W. L. Miranker (eds), *A New Approach to Scientific Computation*, Academic Press, 1983.
- [57] D. C. Sorenson and P. T. P. Tang, *On the Orthogonality of Eigenvectors Computed by Divide-and-Conquer Techniques*, SIAM Journal of Numerical Analysis, Vol. 28, No. 6, p. 1760, algorithm 5.3.
- [58] *Floating-Point C Extensions* in Technical Report Numerical C Extensions Committee X3J11, April 1995, SC22/WG14 N403, X3J11/95-004.

- [59] David M. Gay, *Correctly Rounded Binary-Decimal and Decimal-Binary Conversions*, AT&T Bell Laboratories, Numerical Analysis Manuscript 90-10, November 1990.
- [60] M. Payne and R. Hanek, *Radian Reduction for Trigonometric Functions*, SIGNUM Newsletter, Vol. 18, January 1983.
- [61] M. Payne and R. Hanek, *Degree Reduction for Trigonometric Functions*, SIGNUM Newsletter, Vol. 18, April 1983.
- [62] N. L. Schryer, *A Test of a Computer's Floating-Point Unit*, Computer Science Technical Report No. 89, AT&T Bell Laboratories, Murray Hill, NJ, 1981.
- [63] G. Bohlender, W. Walter, P Kornerup, D. W. Matula, *Semantics for Exact Floating Point Operations*, IEEE Arithmetic 10, 1992.
- [64] W. Walter et al., *Proposal for Accurate Floating-Point Vector Arithmetic*, Mathematics and Computers in Simulation, vol. 35, no. 4, pp. 375-382, IMACS, 1993.
- [65] James Gosling, Bill Joy, Guy Steele, *The Java Language Specification*.
- [66] Simon Peyton Jones et al., *Report on the programming language Haskell 98*, February 1999.
- [67] Simon Peyton Jones et al., *Standard libraries for the Haskell 98 programming language*, February 1999.
- [68] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen, *The Definition of Standard ML (Revised)*, The MIT Press, 1997, ISBN: 0-262-63181-4.