

WG 14 N1632

ISO/IEC JTC 1/SC 22 0000

Date: yyyy-mm-dd

5

Reference number of document: **ISO/IEC nnn-n**

Committee identification: ISO/IEC JTC 1/SC 22/WG 14

Secretariat: ANSI

10

Information Technology — Programming languages, their environments, and system software interfaces — Floating-point extensions for C — Part 2: Decimal floating-point arithmetic

Technologies de l'information — Langages de programmation, leurs environnements et interfaces du logiciel système — Extensions à virgule flottante pour C — Partie 2: décimal arithmétique flottante

15

Warning

This document is not an ISO International Standard. It is distributed for review and comment. It is subject to change without notice and may not be referred to as an International Standard.

Recipients of this draft are invited to submit, with their comments, notification of any relevant patent rights of which they are aware and to provide supporting documentation.

Copyright notice

5 This ISO document is a working draft or committee draft and is copyright-protected by ISO. While the reproduction of working drafts or committee drafts in any form for use by participants in the ISO standards development process is permitted without prior permission from ISO, neither this document nor any extract from it may be reproduced, stored or transmitted in any form for any other purpose without prior written permission from ISO.

Requests for permission to reproduce this document for the purpose of selling it should be addressed as shown below or to ISO's member body in the country of the requester:

10 *ISO copyright office*
Case postale 56 CH-1211 Geneva 20
Tel. +41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

15 Reproduction for sales purposes may be subject to royalty payments or a licensing agreement.

Violators may be prosecuted.

Contents

	Page
Foreword	iv
0 Introduction	iv
0.1 Background	iv
5 0.2 Purpose	v
0.3 The arithmetic model	v
0.4 The formats	vi
1 Scope	1
2 Conformance	1
10 3 Normative references	2
4 Terms and definitions	2
5 Predefined macros	2
6 Decimal floating types	2
7 Characteristics of decimal floating types <code><float.h></code>	4
15 8 Operation binding	6
9 Conversions	7
9.1 Conversions between decimal floating and integer	7
9.2 Conversions among decimal floating types, and between decimal floating types and generic floating types	7
20 9.3 Conversions between decimal floating and complex	8
9.4 Usual arithmetic conversions	8
9.5 Default argument promotion	8
10 Constants	8
11 Arithmetic operations	9
25 11.1 Operators	9
11.2 Functions	10
11.3 Conversions	11
12 Library	11
12.1 Standard headers	11
30 12.2 Floating-point environment <code><fenv.h></code>	11
12.3 Decimal mathematics <code><math.h></code>	12
12.4 New <code><math.h></code> functions	21
12.4.1 Quantum exponent functions	21
12.4.2 Decimal re-encoding functions	22
35 12.5 Formatted input/output specifiers	24
12.6 <code>strtod32</code> , <code>strtod64</code> , and <code>strtod128</code> functions <code><stdlib.h></code>	26
12.7 <code>wcstod32</code> , <code>wcstod64</code> , and <code>wcstod128</code> functions <code><wchar.h></code>	29
12.8 Type-generic macros <code><tgmath.h></code>	30
40 Bibliography	32

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of technical committees is to prepare International Standards. Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75 % of the member bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights.

ISO nnn-n was prepared by Technical Committee ISO JTC 1, *Information Technology*, Subcommittee SC 22, *Programming languages, their environments, and system software interfaces*.

ISO nnn consists of the following parts, under the general title *Floating-point extensions for C*:

- *Part 1: Binary floating-point arithmetic*
- *Part 2: Decimal floating-point arithmetic*
- *Part 3: Interchange and extended types*
- *Part 4: Supplemental functions*
- *Part 5: Supplemental attributes*

Part 1 updates ISO/IEC 9899:2011 (*Information technology — Programming languages, their environments and system software interfaces — Programming Language C*), Annex F in particular, to support all required features of ISO/IEC/IEEE 60559:2011 (*Information technology — Microprocessor Systems — Floating-point arithmetic*).

Part 2 supersedes ISO/IEC TR 24732:2008 (*Information technology – Programming languages, their environments and system software interfaces – Extension for the programming language C to support decimal floating-point arithmetic*).

Parts 3-5 specify extensions to ISO/IEC 9899:2011 for features recommended in ISO/IEC/IEEE 60559:2011.

0 Introduction

0.1 Background

Most of today's general-purpose computing architectures provide binary floating-point arithmetic in hardware. Binary floating-point is an efficient representation that minimizes memory use, and is simpler to implement than floating-point arithmetic using other bases. It has therefore become the norm for scientific computations, with almost all implementations following the IEEE 754 standard for binary floating-point arithmetic (and the equivalent international ISO/IEC 60559 standard).

5 However, human computation and communication of numeric values almost always uses decimal arithmetic and decimal notations. Laboratory notes, scientific papers, legal documents, business reports and financial statements all record numeric values in decimal form. When numeric data are given to a program or are displayed to a user, conversion between binary and decimal is required. There are inherent rounding errors involved in such conversions; decimal fractions cannot, in general, be represented exactly by binary floating-point values. These errors often cause usability and efficiency problems, depending on the application.

10 These problems are minor when the application domain accepts, or requires results to have, associated error estimates (as is the case with scientific applications). However, in business and financial applications, computations are either required to be exact (with no rounding errors) unless explicitly rounded, or be supported by detailed analyses that are auditable to be correct. Such applications therefore have to take special care in handling any rounding errors introduced by the computations.

15 The most efficient way to avoid conversion error is to use decimal arithmetic. Currently, the IBM z/Architecture (and its predecessors since System/360) is a widely used system that supports built-in decimal arithmetic. Prior to the IBM System z10 processor, however, this provided integer arithmetic only, meaning that every number and computation has to have separate scale information preserved and computed in order to maintain the required precision and value range. Such scaling is difficult to code and is error-prone; it affects execution time significantly, and the resulting program is often difficult to maintain and enhance.

20 Even though the hardware may not provide decimal arithmetic operations, the support can still be emulated by software. Programming languages used for business applications either have native decimal types (such as PL/I, COBOL, REXX, C#, or Visual Basic) or provide decimal arithmetic libraries (such as the BigDecimal class in Java). The arithmetic used in business applications, nowadays, is almost invariably decimal floating-point; the COBOL 2002 ISO standard, for example, requires that all standard decimal arithmetic calculations use 32-digit decimal floating-point.

25 The IEEE has recognized the importance of this. Decimal floating-point formats and arithmetic are major new features in the IEEE 754:2008 standard and its international equivalent IEC 60559:2011.

30 Historically there has been a close tie between IEEE 754 and C with respect to floating-point specification. IEC Technical Report 24732:2008 introduced partial C support for the decimal floating-point arithmetic in IEC 60559:2011. TR 24732, for which technical content was completed while IEEE 754-2008 was still in the later stages of development, specifies decimal types based on IEC 60559:2011 decimal formats, though it does not include all of the operations required by IEC 60559:2011.

0.2 Purpose

35 The purpose of this document is to enhance TR 24732 to cover all the requirements, plus some basic recommendations, of IEC 60559:2011 for decimal floating-point arithmetic. C implementations intending to provide an extension for decimal floating-point arithmetic supporting IEC 60559-2011 are expected to conform to Part 2 of this Technical Specification.

0.3 The arithmetic model

The model of floating-point arithmetic used in IEC 60559:2011 has three components:

- 40 • *data* - numbers and NaNs, which can be manipulated by, or be the results of, the operations it specifies
- *operations* - (addition, multiplication, conversions, etc.) which can be carried out on data
- *context* - the status of operations (namely, exceptions flags), and controls to govern the results of operations (for example, rounding modes). (IEC 60559 does not use a single term to refer to these collectively.)

45 The model defines these components in the abstract. It neither defines the way in which operations are expressed (which might vary depending on the computer language or other interface being used), nor does it define the concrete representation (specific layout in storage, or in a processor's register, for example) of data or context, except that it does define specific encodings that are to be used for data that may be exchanged between different implementations that conform to the specification.

From the perspective of the C language, *data* are represented by data types, *operations* are defined within expressions, and *context* is the floating environment specified in `<fenv.h>`. Part 2 of this Technical Specification, this document, specifies how the C language implements these components.

0.4 The formats

5 IEC 60559:2011 specifies *formats*, in terms of their radix, exponent range, and precision (significand length), to support general-purpose decimal floating-point arithmetic. It specifies operation semantics in terms of values and abstract representations of data (format members). It also specifies bit-level encodings for formats intended for data interchange.

10 C11 specifies floating-point arithmetic using a two-layer organization. The first layer provides a specification using an abstract model. The representation of a floating-point number is specified in an abstract form where the constituent components of the representation are defined (sign, exponent, significand) but not the internals of these components. In particular, the exponent range, significand size, and the base (or radix) are implementation defined. This allows flexibility for an implementation to take advantage of its underlying hardware architecture. Furthermore, certain behaviors of operations are also implementation defined, for
15 example in the area of handling of special numbers and in exceptions.

The reason for this approach is historical. At the time when C was first standardized, there were already various hardware implementations of floating-point arithmetic in common use. Specifying the exact details of a representation would make most of the existing implementations at the time not conforming.

20 C11 provides a binding to the previous 1989 version of the IEC 60559 standard by specifying an annex, Annex F, *IEC 60559 floating point arithmetic*, and adopting that standard by reference. An implementation may choose not to conform to IEC 60559 and indicates that by not defining the macro `__STDC_IEC_559__`. This means not all implementations need to support IEC 60559, and the floating-point arithmetic need not be binary.

25 Part 1 of this Technical Specification suggests changes that update C11 to provide a binding to IEC 60559:2011 for binary floating-point arithmetic.

Part 2 of this Technical Specification, this document, specifies decimal floating-point arithmetic according to IEC 60559:2011, with the constituent components of the representation defined. This is more stringent than the existing C11 approach for the floating types. Since it is expected that decimal floating-point implementations will conform to IEC 60559:2011, binding to this standard directly benefits both implementers
30 and programmers.

Information Technology — Programming languages, their environments, and system software interfaces — Floating-point extensions for C — Part 2: Decimal floating-point arithmetic

1 Scope

5 Part 2 of Technical Specification 00000 extends programming language C, as specified in IEC 9899:2011 (C11), to support decimal floating-point arithmetic conforming to ISO/IEC/IEEE 60559:2011. It covers all requirements of IEC 60559 as they pertain to C decimal floating types.

10 TS 00000-2 supersedes ISO/IEC TR 24732:2008 (Information technology – Programming languages, their environments and system software interfaces – Extension for the programming language C to support decimal floating-point arithmetic).

TS 00000-2 does not cover binary floating-point arithmetic (which is covered in TS 00000-1), nor most other optional features of IEC 60559.

2 Conformance

An implementation conforms to Part 2 of this Technical Specification if

- 15 a) It conforms to Part 1 of this Technical Specification;
- b) It meets the requirements for a conforming implementation of C11 with all the suggested changes to C11 in Part 2 of this Technical Specification; and
- 20 c) It defines `__STDC_DEC_FP__` to 201ymmL.

The library functions, macros, and types defined in Part 2 of this Technical Specification are defined by their respective headers if the macro `__STDC_WANT_DEC_FP__` is defined at the point in the source file where the appropriate header is first included.

3 Normative references

The following referenced documents are indispensable for the application of this document. Only the editions cited apply.

5 ISO/IEC 9899:2011, *Information technology — Programming languages, their environments and system software interfaces — Programming Language C*

ISO/IEC/IEEE 60559:2011, *Information technology — Microprocessor Systems — Floating-point arithmetic* (with identical content to IEEE 754-2008, *IEEE Standard for Floating-Point Arithmetic*. The Institute of Electrical and Electronic Engineers, Inc., New York, 2008)

10 ISO/IEC TR 24732:2008, *Information technology – Programming languages, their environments and system software interfaces – Extension for the programming language C to support decimal floating-point arithmetic*

ISO/IEC TS nnnnn-1:yyyy, *Information technology – Programming languages, their environments and system software interfaces – Floating-point extension for C – Part 1: Binary floating-point arithmetic*

4 Terms and definitions

15 For the purposes of this document, the terms and definitions given in ISO/IEC 9899:2011 and ISO/IEC/IEEE 60559:2011 and the following apply.

4.1 C11

standard ISO/IEC 9899:2011, *Information technology — Programming languages, their environments and system software interfaces — Programming Language C*

5 Predefined macros

The following suggested change to C11 provides the conditional feature macro for the features in Part 2 of this Technical Specification.

Suggested change to C11:

In 6.10.8.3#1, add:

25 `__STDC_DEC_FP__` The integer constant `201ymmL`, intended to indicate support of decimal floating types and operations according to IEC 60559.

6 Decimal floating types

30 Part 2 of this Technical Specification introduces three decimal floating types, designated as `_Decimal132`, `_Decimal164` and `_Decimal128`. The set of values of type `_Decimal132` is a subset of the set of values of the type `_Decimal164`; the set of values of the type `_Decimal164` is a subset of the set of values of the type `_Decimal128`.

Within the type hierarchy, decimal floating types are basic types, real types and arithmetic types.

The types `float`, `double`, and `long double` are also called generic floating types for the purpose of this Technical Report.

35 Note: C does not specify a radix for `float`, `double`, and `long double`. An implementation can choose the representation of `float`, `double`, and `long double` to be the same as the decimal floating types. In any case, the decimal floating types are distinct from `float`, `double`, and `long double` regardless of the representation.

Note: This Technical Report does not define decimal complex types or decimal imaginary types. The three complex types remain as `float _Complex`, `double _Complex`, and `long double _Complex`, and the three imaginary types remain as `float _Imaginary`, `double _Imaginary`, and `long double _Imaginary`.

5 Suggested changes to C11:

Change the first sentence of 6.2.5#10 from:

[10] There are three *real floating types*, designated as `float`, `double`, and `long double`

to:

[10] There are three *generic floating types*, designated as `float`, `double`, and `long double`.

10 Add the following paragraphs after 6.2.5#10:

[10a] There are three *decimal floating types*, designated as `_Decimal32`, `_Decimal64`, and `_Decimal128`. The set of values of the type `_Decimal32`¹ is a subset of the set of values of the type `_Decimal64`; the set of values of the type `_Decimal64` is a subset of the set of values of the type `_Decimal128`. Decimal floating types are real floating types.

15 [10b] Together, the generic floating types and the decimal floating types comprise the *real floating types*.

Add the following to 6.4.1 Keywords:

keyword:

20 `_Decimal32`
 `_Decimal64`
 `_Decimal128`

Add the following to 6.7.2 Type specifiers:

25 *type-specifier:*
 `_Decimal32`
 `_Decimal64`
 `_Decimal128`

Add the following bullets in 6.7.2#2 Constraints:

30 — `_Decimal32`

 — `_Decimal64`

 — `_Decimal128`

Add the following after 6.7.2#3:

35 [3a] The type specifiers `_Decimal32`, `_Decimal64`, and `_Decimal128` shall not be used if the implementation does not support decimal types (see 6.10.8.3).

Add the following after 6.5#8:

¹ The 32-bit format is a storage only format in IEC 60559.

[8a] Expressions involving decimal floating-point operands are evaluated according to the semantics of IEC 60559, including production of results with the preferred quantum exponent as specified in IEC 60559.

7 Characteristics of decimal floating types <float.h>

5 The characteristics of decimal floating types are defined in terms of a model specifying general decimal arithmetic (0.3). The formats are specified in IEC 60559 (0.4).

The three decimal floating types correspond to the decimal formats defined in IEC 60559 as follows:

- `__Decimal132` is a *decimal32* format, which is encoded in four consecutive octets (32 bits)
- `__Decimal164` is a *decimal64* format, which is encoded in eight consecutive octets (64 bits)
- 10 — `__Decimal128` is a *decimal128* format, which is encoded in 16 consecutive octets (128 bits)

The value of a finite number is given by $(-1)^{\text{sign}} \times \text{coefficient} \times 10^{\text{exponent}}$. Refer to IEC 60559 for details of the format.

15 These formats are characterized by the length of the coefficient and the maximum and minimum exponent. The coefficient is not normalized, so trailing zeros are significant; i.e., 1.0 is equal to but can be distinguished from 1.00. The table below shows these characteristics by format:

Format	<code>__Decimal132</code>	<code>__Decimal164</code>	<code>__Decimal128</code>
Coefficient length in digits	7	16	34
Maximum Exponent ² (E_{max})	97	385	6145
Minimum Exponent (E_{min})	-94	-382	-6142

20 If the macro `__STDC_WANT_DEC_FP__` is defined at the point in the source file where the header <float.h> is first included, the header <float.h> shall define several macros that expand to various limits and parameters of the decimal floating types. The names and meaning of these macros are similar to the corresponding macros for generic floating types.

Suggested change to C11:

Add the following after 5.2.4.2.2:

5.2.4.2.2a Characteristics of decimal floating types <float.h>

25 [1] Macros in <float.h> provide characteristics of floating types in terms of the model presented in 5.2.4.2.2. The prefixes `DEC32_`, `DEC64_`, and `DEC128_` denote the types `__Decimal132`, `__Decimal164`, and `__Decimal128` respectively.

[2] For decimal floating-point, it is often convenient to consider an alternate equivalent model where the significand is represented with integer rather than fraction digits: a floating-point number (x) is defined by the model

$$30 \quad x = sb^{(e-p)} \sum_{k=1}^p f_k b^{(p-k)}$$

where s , b , e , p , and f_k are as defined in 5.2.4.2.2, and $b = 10$.

² The maximum and minimum exponents here are for floating-point numbers expressed with significands less than 1, as in the C11 model (5.2.4.2.2). They differ (by 1) from the maximum and minimum exponents in the IEC 60559 standard, where normalized floating-point numbers are expressed with one significant digit to the left of the radix point.

[3] The term *quantum exponent* refers to $q = e - p$ and *coefficient* to $c = f_1f_2\dots f_p$, an integer between 0 and $b^p - 1$ inclusive. Thus, $x = s * c * b^q$ is represented by the triple of integers (s, c, q) .

[4] For binary floating-point following IEC 60559 (and IEC 60559), representations in the model described in 5.2.4.2.2 that have the same numerical value are indistinguishable in the arithmetic. However, for decimal floating-point, representations that have the same numerical value but different quantum exponents, e.g., $(1, 10, -1)$ representing 1.0 and $(1, 100, -2)$ representing 1.00, are distinguishable. To facilitate exact fixed-point calculation, standard decimal floating-point operations and functions have a *preferred quantum exponent*, as specified in IEC 60559, which is determined by the quantum exponents of the operands if they have decimal floating-point types (or by specific rules for conversions from other types), and they produce a result with that preferred quantum exponent, or as close to it as possible within the limitations of the type. For example, the preferred quantum exponent for addition is the minimum of the quantum exponents of the operands. Hence $(1, 123, -2) + (1, 4000, -3) = (1, 5230, -3)$ or $1.23 + 4.000 = 5.230$.

[5] Except for assignment and casts, the values of operations with decimal floating operands and values subject to the usual arithmetic conversions and of decimal floating constants are evaluated to a format whose range and precision may be greater than required by the type. The use of evaluation formats is characterized by the implementation-defined value of **DEC_EVAL_METHOD**:

- 1 indeterminate;
- 0 evaluate all operations and constants just to the range and precision of the type;
- 1 evaluate operations and constants of type **_Decimal32** and **_Decimal64** to the range and precision of the **_Decimal64** type, evaluate **_Decimal128** operations and constants to the range and precision of the **_Decimal128** type;
- 2 evaluate all operations and constants to the range and precision of the **_Decimal128** type.

[6] The integer values given in the following lists shall be replaced by constant expressions suitable for use in **#if** preprocessing directives:

- radix of exponent representation, $b(=10)$

For the generic floating-point types, this value is implementation-defined and is specified by the macro **FLT_RADIX**. For the decimal floating-point types there is no corresponding macro, since the value 10 is an inherent property of the types. Wherever **FLT_RADIX** appears in a description of a function that has versions that operate on decimal floating-point types, it is noted that for the decimal floating-point versions the value used is implicitly 10, rather than **FLT_RADIX**.

- number of digits in the coefficient

DEC32_MANT_DIG	7
DEC64_MANT_DIG	16
DEC128_MANT_DIG	34

- minimum exponent

DEC32_MIN_EXP	-94
DEC64_MIN_EXP	-382
DEC128_MIN_EXP	-6142

- maximum exponent

DEC32_MAX_EXP	97
DEC64_MAX_EXP	385
DEC128_MAX_EXP	6145

The C `encodebid` (7.12.11.10) and `decodebid` (7.12.11.11) functions provide the `encodeBinary` and `decodeBinary` operations defined (specifically for decimal floating point) in IEC 60559.

9 Conversions

9.1 Conversions between decimal floating and integer

- 5 For conversions between real floating and integer types, C11 6.3.1.4 leaves the behavior undefined if the conversion result cannot be represented (Annex F.4 tightened up the behavior.) To help writing portable code, Part 2 of this Technical Specification provides defined behavior for decimal floating types. Furthermore, it is useful to allow program execution to continue without interruption unless the program needs to check the condition.

10 Suggested changes to C11:

Change the first sentence of 6.3.1.4 paragraph 1 from:

[1] When a finite value of real floating type is converted to an integer type ...

to:

[1] When a finite value of generic floating type is converted to an integer type ...

- 15 Add the follow paragraph after 6.3.1.4 paragraph 1:

[1a] When a finite value of decimal floating type is converted to an integer type other than `_Bool`, the fractional part is discarded (i.e., the value is truncated toward zero). If the value of the integral part cannot be represented by the integer type, the “invalid” floating-point exception shall be raised and the result of the conversion is unspecified.

- 20 ***ISSUE 1: Should we capture in C11 the IEC 60559 recommendation for implicit inexact conversion to raise the “inexact” exception? Would this conflict with conventions that assignments and casts are equivalent?***

Change the first sentence of 6.3.1.4 paragraph 2 from:

[2] When a value of integer type is converted to a real floating type, ...

- 25 to:

[2] When a value of integer type is converted to a generic floating type, ...

Add the following paragraph after 6.3.1.4 paragraph 2:

- 30 [2a] When a value of integer type is converted to a decimal floating type, if the value being converted can be represented exactly in the new type, it is unchanged. If the value being converted cannot be represented exactly, the result shall be correctly rounded with exceptions raised as specified in IEC 60559.

9.2 Conversions among decimal floating types, and between decimal floating types and generic floating types

- 35 The specification is similar to the existing ones for `float`, `double`, and `long double`, except that when the result cannot be represented exactly, the behavior is tightened to become correctly rounded.

Suggested change to C11:

Add after 6.3.1.5#1.

[2] When a value of real floating type is converted to a real floating type and one or both of the types are decimal types, if the value being converted cannot be represented exactly, the result is correctly rounded with exceptions raised as specified in IEC 60559.

9.3 Conversions between decimal floating and complex

5 This is covered by C11 6.3.1.7.

9.4 Usual arithmetic conversions

In an application that is written using decimal arithmetic, mixed operations between decimal and other real types are likely to occur only when interfacing with other languages, calling existing libraries written for binary floating point arithmetic, or accessing existing data. Determining the common type for mixed operations is difficult because ranges overlap; therefore, mixed mode operations are not allowed and the programmer must use explicit casts. Implicit conversions are allowed only for simple assignment, `return` statement, and in argument passing involving prototyped functions.

Following are suggested changes to C11:

Insert the following to 6.3.1.8#1, after "This pattern is called the *usual arithmetic conversions*:"

15 If one operand is a decimal floating type, all other operands shall not be generic floating type, complex type, or imaginary type:

First if either operand is `_Decimal128`, the other operand is converted to `_Decimal128`.

Otherwise, if either operand is `_Decimal64`, the other operand is converted to `_Decimal64`.

Otherwise, if either operand is `_Decimal32`, the other operand is converted to `_Decimal32`.

20 If there are no decimal floating types in the operands:

First, if the corresponding real type of either operand is `long double`, the other operand is converted, without ... <the rest of 6.3.1.8#1 remains the same>

9.5 Default argument promotion

25 There is no default argument promotion specified for the decimal floating types. Default argument promotion covered in C11 6.5.2.2 [6] and [7] remains unchanged, and applies to generic floating types only.

10 Constants

New suffixes are added to denote decimal floating constants: `DF` for `_Decimal32`, `DD` for `_Decimal64`, and `DL` for `_Decimal128`.

30 Part 2 of this specification does not carry forward two features introduced in TR 24732: the `FLOAT_CONST_DECIMAL64` pragma and the `d` and `D` suffixes for floating constants. The pragma changed the interpretation of unsuffixed floating constants between `double` and `_Decimal64`. The suffixes provided a way to designate `double` floating constants so that the pragma would not affect them. The pragma is not included because of its potential for inadvertently reinterpreting constants. Without the pragma, the suffixes are no longer needed. Also, significant implementations use the `d` and `D` suffixes for other purposes.

35 Suggested changes to C11:

Change *floating-suffix* in 6.4.4.2 from:

floating-suffix: one of
`f l F L`

to:

floating-suffix: one of
f l F L df dd d1 DF DD DL

Add the following paragraph after 6.4.4.2#2:

5 [2a] Constraints

A *floating-suffix* df, dd, d1, DF, DD, or DL shall not be used in a *hexadecimal-floating-constant*.

Add the following paragraph after 6.4.4.2#4:

[4a] If a floating constant is suffixed by df or DF, it has type `_Decimal32`. If suffixed by dd or DD, it has type `_Decimal64`. If suffixed by d1 or DL, it has type `_Decimal128`.

10 Add the following paragraph after 6.4.4.2#5:

[5a] Decimal floating-point constants that have the same numerical value but different quantum exponents have distinguishable internal representations. The quantum exponent is specified to be the same as for the corresponding `strtod32`, `strtod64`, or `strtod128` function for the same numeric string.

15 11 Arithmetic operations

11.1 Operators

The operators *Add* (C11 6.5.6), *Subtract* (C11 6.5.6), *Multiply* (C11 6.5.5), *Divide* (C11 6.5.5), *Relational operators* (C11 6.5.8), *Equality operators* (C11 6.5.9), *Unary Arithmetic operators* (C11 6.5.3.3), and *Compound Assignment operators* (C11 6.5.16.2) when applied to decimal floating type operands shall follow the semantics as defined in IEC 60559.

Suggested changes to C11:

Add the following after 6.5.5 paragraph 2:

[2a] If either operand has decimal floating type, the other operand shall not have generic floating type, complex type, nor imaginary type.

25 Add the following after 6.5.6 paragraph 3:

[3a] If either operand has decimal floating type, the other operand shall not have generic floating type, complex type, nor imaginary type.

Add the following after 6.5.8 paragraph 2:

30 [2a] If either operand has decimal floating type, the other operand shall not have generic floating type.

Add the following after 6.5.9 paragraph 2:

[2a] If either operand has decimal floating type, the other operand shall not have generic floating type, complex type, nor imaginary type.

Add the following bullet to 6.5.15 paragraph 3:

35 — one operand has decimal floating type, and the other has arithmetic type other than generic floating type, complex type, or imaginary type;

Add the following after 6.5.16.2 paragraph 2:

[2a] If either operand has decimal floating type, the other operand shall not have generic floating type, complex type, nor imaginary type.

11.2 Functions

- 5 The headers and library supply a number of functions and macros that implement support for decimal floating-point data with the semantics specified in IEC 60559, including producing results with the preferred quantum exponent where appropriate. That support is provided by the following:

From C11 `<math.h>`, with suggested changes in Part 1 of this Technical Specification, the decimal floating-point type versions of:

10 `sqrt`, `fma`, `fabs`, `fmax`, `fmin`, `ceil`, `floor`, `trunc`, `round`, `rint`, `lround`, `llround`, `ldexp`,
`frexp`, `ilogb`, `logb`, `scalbn`, `scalbln`, `copysign`, `remainder`, `isnan`, `isinf`, `isfinite`,
`isnormal`, `signbit`, `fpclassify`, `isunordered`, `isgreater`, `isgreaterequal`, `isless`,
`islessequal` and `islessgreater`.

15 From the `<math.h>` extensions specified in Part 1 of this Technical Specification, the decimal floating-point type versions of:

`roundeven`, `nextup`, `nextdown`, `fminmag`, `fmaxmag`, `llogb`, `fadd`, `faddl`, `daddl`, `fsub`,
`fsubl`, `dsubl`, `fmul`, `fmull`, `dmull`, `fdiv`, `fdivl`, `ddivl`, `fsqrt`, `fsqrtl`, `dsqrtl`,
`ffma`, `ffmal`, `dfmal`, `fromfp`, `ufromfp`, `fromfpx`, `ufromfpx`, `canonicalize`, `iseqsig`,
20 `issignaling`, `issubnormal`, `iscanonical`, `totalorder`, `totalordermag`, `getpayload`,
`setpayload`, and `setpayloadsig`.

The `<math.h>` extensions specified below in 12.4 for the decimal-specific functions:

`quantize`, `samequantum`, `quantexp`, `encodedpd`, `decodedpd`, `encodebid`, and `decodebid`.

From C11 `<fenv.h>`, facilities dealing with decimal context:

25 `feraiseexcept`, `feclearexcept`, `fetestexcept`, `fesetexceptflag`, `fegetexceptflag`,
`fesetenv`, `fegetenv`, `feupdateenv`, and `feholdexcept`.

From `<fenv.h>` extensions specified in Part 2 of this Technical Specification, facilities dealing with decimal context:

`fe_dec_getround` and `fe_dec_setround`.

30 From the `<fenv.h>` extensions specified in Part 1 of this Technical Specification, facilities dealing with decimal context:

`fetestexceptflag`, `fesetexcept`, `fegetmode`, and `fesetmode`.

From `<stdio.h>`, decimal floating-point modified format specifiers for:

The `printf/scanf` family of functions.

35 From `<stdlib.h>` and `<wchar.h>`, with suggested changes in Part 1 of this Technical Specification, the decimal floating-point type versions of:

`strtod` and `wctod`.

From the `<stdlib.h>` extensions specified in Part 1 of this Technical Specification, the decimal floating-point type version of:

```
strfromflt.
```

From `<wchar.h>`, decimal floating-point modified format specifiers for:

- 5 The wide `wprintf/wscanf` family of functions.

11.3 Conversions

Conversions between different floating types and to/from integer types are covered in clause 9.

12 Library

12.1 Standard headers

- 10 The functions, macros, and types declared or defined in Clause 12 and its subclauses are only declared or defined by their respective headers if the macro `__STDC_WANT_DEC_FP__` is defined at the point in the source file where the appropriate header is first included.

12.2 Floating-point environment `<fenv.h>`

- 15 The floating-point environment specified in C11 7.6 applies to both generic floating types and decimal floating types. This is to implement the *context* defined in IEC 60559. The existing C11 specification gives flexibility to an implementation on which part of the environment is accessible to programs. The decimal floating-point arithmetic specifies a more stringent requirement. All the rounding directions and flags are supported.

DEC Macros	Existing C11 macros for generic floating types	IEC 60559
<code>FE_DEC_TOWARDZERO</code>	<code>FE_TOWARDZERO</code>	Toward zero
<code>FE_DEC_TONEAREST</code>	<code>FE_TONEAREST</code>	To nearest, ties even
<code>FE_DEC_UPWARD</code>	<code>FE_UPWARD</code>	Toward plus infinity
<code>FE_DEC_DOWNWARD</code>	<code>FE_DOWNWARD</code>	Toward minus infinity
<code>FE_DEC_TONEARESTFROMZERO</code>	n/a	To nearest, ties away from zero

Suggested changes to C11:

- 20 Add the following after 7.6 paragraph 6:

[6a] Decimal floating-point arithmetic and IEC 60559 binary floating-point arithmetic (Annex F) access the same floating-point exception status flags.

Add the following after 7.6 paragraph 8:

[8a] Each of the macros

- 25 `FE_DEC_DOWNWARD`
`FE_DEC_TONEAREST`
`FE_DEC_TONEARESTFROMZERO`
`FE_DEC_TOWARDZERO`
`FE_DEC_UPWARD`

- 30 is defined for use by the `fe_dec_getround` and `fe_dec_setround` functions for getting and setting the rounding direction of decimal floating-point operations. The defined macros expand to integer constant expressions whose values are distinct nonnegative values.

[8b] During translation, constant rounding direction modes for decimal floating-point arithmetic are in effect where specified. Elsewhere, during translation the decimal rounding direction mode is **FE_DEC_TONEAREST**.

[8c] At program startup the dynamic rounding direction mode for decimal floating-point arithmetic is initialized to **FE_DEC_TONEAREST**.

Add the following after 7.6.3.2:

7.6.3.3 The `fe_dec_getround` function

Synopsis

```
#define __STDC_WANT_DEC_FP__
#include <fenv.h>
int fe_dec_getround(void);
```

Description

The `fe_dec_getround` function gets the current rounding direction for decimal floating-point operations.

Returns

The `fe_dec_getround` function returns the value of the rounding direction macro representing the current rounding direction for decimal floating-point operations, or a negative value if there is no such rounding macro or the current rounding direction is not determinable.

7.6.3.4 The `fe_dec_setround` function

Synopsis

```
#define __STDC_WANT_DEC_FP__
#include <fenv.h>
int fe_dec_setround(int round);
```

Description

The `fe_dec_setround` function establishes the rounding direction for decimal floating-point operations represented by its argument `round`. If the argument is not equal to the value of a rounding direction macro, the rounding direction is not changed.

If **FLT_RADIX** is not 10, the rounding direction altered by the `fesetround` function is independent of the rounding direction altered by the `fe_dec_setround` function; otherwise if **FLT_RADIX** is 10, whether the `fesetround` and `fe_dec_setround` functions alter the rounding direction of both generic floating type and decimal floating type operations is implementation defined.

Returns

The `fe_dec_setround` function returns a zero value if and only if the argument is equal to a rounding direction macro (that is, if and only if the requested rounding direction was established).

12.3 Decimal mathematics <math.h>

The list of elementary functions specified in the mathematics library is extended to handle decimal floating-point types. These include functions specified in C11 7.12.4, 7.12.5, 7.12.6, 7.12.7, 7.12.8, 7.12.9, 7.12.10, 7.12.11, 7.12.12, and 7.12.13 and in Part 1 of this Technical Specification 14.1, 14.2, 14.3, 14.4, 14.5, 14.8, 14.9, and 14.0. The macros **HUGE_VAL_D32**, **HUGE_VAL_D64**, **HUGE_VAL_D128**, **DEC_INFINITY**, **DEC_NAN**,

`DEC_SNaN32`, `DEC_SNaN64`, and `DEC_SNaN128` are defined to help using these functions. With the exception of the decimal floating-point functions listed in 11.2, which have accuracy as specified in IEC 60559, the accuracy of decimal floating-point results is implementation-defined. The implementation may state that the accuracy is unknown. All classification macros specified in C11 7.12.3 and in Part 1 of this Technical Specification 14.7 are also extended to handle decimal floating-point types. The same applies to all comparison macros specified in C11 7.12.14 and in Part 1 of this Technical Specification 14.6.

The names of the functions are derived by adding suffixes `d32`, `d64`, and `d128` to the `double` version of the function name, except for the functions that round result to narrower type (7.12.14).

Suggested changes to C11:

10 Add after 7.12 paragraph 2.

[2a] The types

```
    _Decimal32_t
    _Decimal64_t
```

15 are decimal floating types at least as wide as `_Decimal32` and `_Decimal64`, respectively, and such that `_Decimal64_t` is at least as wide as `_Decimal32_t`. If `DEC_EVAL_METHOD` equals 0, `_Decimal32_t` and `_Decimal64_t` are `_Decimal32` and `_Decimal64`, respectively; if `DEC_EVAL_METHOD` equals 1, they are both `_Decimal64`; if `DEC_EVAL_METHOD` equals 2, they are both `_Decimal128`; and for other values of `DEC_EVAL_METHOD`, they are otherwise implementation-defined.

20

[2b] The types

```
    dpencoding32_t
    dpencoding64_t
    dpencoding128_t
    bidencoding32_t
    bidencoding64_t
    bidencoding128_t
```

25

30 represent values of decimal floating types in one of the two alternative encodings required by the IEC 60559 standard: the Densely Packed Decimal (DPD) encoding or the Binary Integer Decimal (BID) encoding. These types are used by the decimal re-encoding functions (12.4.2).

Add at the end of 7.12 paragraph 3 the following macros.

[3] The macro

```
    HUGE_VAL_D64
```

35 expands to a constant expression of type `_Decimal64` representing infinity. The macros

```
    HUGE_VAL_D32
    HUGE_VAL_D128
```

are respectively `_Decimal32` and `_Decimal128` analogs of `HUGE_VAL_D64`.

40 Add at the end of 7.12 paragraph 4 the following macro.

[4] The macro

```
    DEC_INFINITY
```

expands to a constant expression of type `_Decimal32` representing positive infinity.

Add at the end of 7.12 paragraph 5 the following macros.

[5a] The macro

```
DEC_NAN
```

expands to a constant expression of type `_Decimal32` representing a quiet NaN.

5 [5b] The signaling NaN macros

```
DEC_SNAN32
DEC_SNAN64
DEC_SNAN128
```

10 expand into a constant expression of the respective decimal floating type representing a signaling NaN. If a signaling NaN macro is used for initializing an object of the same type that has static or thread-local storage duration, the object is initialized with a signaling NaN value.

Add at the end of 7.12 paragraph 7 the following macros.

[7] The macros

15

```
FP_FAST_FMAD32
FP_FAST_FMAD64
FP_FAST_FMAD128
```

are, respectively, `_Decimal32`, `_Decimal64`, and `_Decimal128` analogs of `FP_FAST_FMA`.

20 Add the following list of function prototypes to the synopsis of the respective subclauses:

7.12.4 Trigonometric functions

25

```
_Decimal64 acosd64(_Decimal64 x);
_Decimal32 acosd32(_Decimal32 x);
_Decimal128 acosd128(_Decimal128 x);

_Decimal64 asind64(_Decimal64 x);
_Decimal32 asind32(_Decimal32 x);
_Decimal128 asind128(_Decimal128 x);

30 _Decimal64 atand64(_Decimal64 x);
_Decimal32 atand32(_Decimal32 x);
_Decimal128 atand128(_Decimal128 x);

_Decimal64 atan2d64(_Decimal64 y, _Decimal64 x);
35 _Decimal32 atan2d32(_Decimal32 y, _Decimal32 x);
_Decimal128 atan2d128(_Decimal128 y, _Decimal128 x);

_Decimal64 cosd64(_Decimal64 x);
_Decimal32 cosd32(_Decimal32 x);
40 _Decimal128 cosd128(_Decimal128 x);

_Decimal64 sind64(_Decimal64 x);
_Decimal32 sind32(_Decimal32 x);
_Decimal128 sind128(_Decimal128 x);

45 _Decimal64 tand64(_Decimal64 x);
_Decimal32 tand32(_Decimal32 x);
_Decimal128 tand128(_Decimal128 x);
```

7.12.5 Hyperbolic functions

```

    _Decimal64 acoshd64(_Decimal64 x);
    _Decimal32 acoshd32(_Decimal32 x);
    _Decimal128 acoshd128(_Decimal128 x);
5
    _Decimal64 asinhd64(_Decimal64 x);
    _Decimal32 asinhd32(_Decimal32 x);
    _Decimal128 asinhd128(_Decimal128 x);
10
    _Decimal64 atanhd64(_Decimal64 x);
    _Decimal32 atanhd32(_Decimal32 x);
    _Decimal128 atanhd128(_Decimal128 x);
15
    _Decimal64 coshd64(_Decimal64 x);
    _Decimal32 coshd32(_Decimal32 x);
    _Decimal128 coshd128(_Decimal128 x);
20
    _Decimal64 sinhd64(_Decimal64 x);
    _Decimal32 sinhd32(_Decimal32 x);
    _Decimal128 sinhd128(_Decimal128 x);
    _Decimal64 tanhd64(_Decimal64 x);
    _Decimal32 tanhd32(_Decimal32 x);
    _Decimal128 tanhd128(_Decimal128 x);

```

25 7.12.6 Exponential and logarithmic functions

```

    _Decimal64 expd64(_Decimal64 x);
    _Decimal32 expd32(_Decimal32 x);
    _Decimal128 expd128(_Decimal128 x);
30
    _Decimal64 exp2d64(_Decimal64 x);
    _Decimal32 exp2d32(_Decimal32 x);
    _Decimal128 exp2d128(_Decimal128 x);
35
    _Decimal64 expm1d64(_Decimal64 x);
    _Decimal32 expm1d32(_Decimal32 x);
    _Decimal128 expm1d128(_Decimal128 x);
40
    _Decimal64 frexp64(_Decimal64 value, int *exp);3
    _Decimal32 frexp32(_Decimal32 value, int *exp);
    _Decimal128 frexp128(_Decimal128 value, int *exp);
45
    int ilogbd64(_Decimal64 x);
    int ilogbd32(_Decimal32 x);
    int ilogbd128(_Decimal128 x);
50
    int llogbd64(_Decimal64 x);
    int llogbd32(_Decimal32 x);
    int llogbd128(_Decimal128 x);
    _Decimal64 ldexpd64(_Decimal64 x, int exp);4
    _Decimal32 ldexpd32(_Decimal32 x, int exp);
    _Decimal128 ldexpd128(_Decimal128 x, int exp);

```

³ See suggested changes to the frexp function description below.

⁴ See suggested changes to the ldexp function description below.

```

    _Decimal64 logd64(_Decimal64 x);
    _Decimal32 logd32(_Decimal32 x);
    _Decimal128 logd128(_Decimal128 x);
5
    _Decimal64 log10d64(_Decimal64 x);
    _Decimal32 log10d32(_Decimal32 x);
    _Decimal128 log10d128(_Decimal128 x);
10
    _Decimal64 log1pd64(_Decimal64 x);
    _Decimal32 log1pd32(_Decimal32 x);
    _Decimal128 log1pd128(_Decimal128 x);
15
    _Decimal64 log2d64(_Decimal64 x);
    _Decimal32 log2d32(_Decimal32 x);
    _Decimal128 log2d128(_Decimal128 x);
20
    _Decimal64 logbd64(_Decimal64 x);
    _Decimal32 logbd32(_Decimal32 x);
    _Decimal128 logbd128(_Decimal128 x);
25
    _Decimal64 modfd64(_Decimal64 value, _Decimal64 *iptr);
    _Decimal32 modfd32(_Decimal32 value, _Decimal32 *iptr);
    _Decimal128 modfd128(_Decimal128 value, _Decimal128 *iptr);
30
    _Decimal64 scalbnd64(_Decimal64 x, int n);
    _Decimal32 scalbnd32(_Decimal32 x, int n);
    _Decimal128 scalbnd128(_Decimal128 x, int n);
35
    _Decimal64 scalblnd64(_Decimal64 x, long int n);
    _Decimal32 scalblnd32(_Decimal32 x, long int n);
    _Decimal128 scalblnd128(_Decimal128 x, long int n);

```

7.12.7 Power and absolute-value functions

```

    _Decimal64 cbrtd64(_Decimal64 x);
    _Decimal32 cbrtd32(_Decimal32 x);
    _Decimal128 cbrtd128(_Decimal128 x);
35
    _Decimal64 fabsd64(_Decimal64 x);
    _Decimal32 fabsd32(_Decimal32 x);
    _Decimal128 fabsd128(_Decimal128 x);
40
    _Decimal64 hypotd64(_Decimal64 x, _Decimal64 y);
    _Decimal32 hypotd32(_Decimal32 x, _Decimal32 y);
    _Decimal128 hypotd128(_Decimal128 x, _Decimal128 y);
45
    _Decimal64 powd64(_Decimal64 x, _Decimal64 y);
    _Decimal32 powd32(_Decimal32 x, _Decimal32 y);
    _Decimal128 powd128(_Decimal128 x, _Decimal128 y);
50
    _Decimal64 sqrtd64(_Decimal64 x);
    _Decimal32 sqrtd32(_Decimal32 x);
    _Decimal128 sqrtd128(_Decimal128 x);

```

7.12.8 Error and gamma functions

```

    _Decimal64 erfd64(_Decimal64 x);
    _Decimal32 erfd32(_Decimal32 x);
    _Decimal128 erfd128(_Decimal128 x);
55

```

```

    _Decimal64 erfcd64(_Decimal64 x);
    _Decimal32 erfcd32(_Decimal32 x);
    _Decimal128 erfcd128(_Decimal128 x);
5
    _Decimal64 lgammad64(_Decimal64 x);
    _Decimal32 lgammad32(_Decimal32 x);
    _Decimal128 lgammad128(_Decimal128 x);
10
    _Decimal64 tgamma64(_Decimal64 x);
    _Decimal32 tgamma32(_Decimal32 x);
    _Decimal128 tgamma128(_Decimal128 x);

```

7.12.9 Nearest integer functions

```

    _Decimal64 ceild64(_Decimal64 x);
    _Decimal32 ceild32(_Decimal32 x);
    _Decimal128 ceild128(_Decimal128 x);
15
    _Decimal64 floord64(_Decimal64 x);
    _Decimal32 floord32(_Decimal32 x);
    _Decimal128 floord128(_Decimal128 x);
20
    _Decimal64 nearbyintd64(_Decimal64 x);
    _Decimal32 nearbyintd32(_Decimal32 x);
    _Decimal128 nearbyintd128(_Decimal128 x);
25
    _Decimal64 rintd64(_Decimal64 x);
    _Decimal32 rintd32(_Decimal32 x);
    _Decimal128 rintd128(_Decimal128 x);
30
    long int lrintd64(_Decimal64 x);
    long int lrintd32(_Decimal32 x);
    long int lrintd128(_Decimal128 x);
35
    long long int llrintd64(_Decimal64 x);
    long long int llrintd32(_Decimal32 x);
    long long int llrintd128(_Decimal128 x);
40
    _Decimal64 roundd64(_Decimal64 x);
    _Decimal32 roundd32(_Decimal32 x);
    _Decimal128 roundd128(_Decimal128 x);
45
    long int lroundd64(_Decimal64 x);
    long int lroundd32(_Decimal32 x);
    long int lroundd128(_Decimal128 x);
50
    long long int llroundd64(_Decimal64 x);
    long long int llroundd32(_Decimal32 x);
    long long int llroundd128(_Decimal128 x);
55
    _Decimal64 truncd64(_Decimal64 x);
    _Decimal32 truncd32(_Decimal32 x);
    _Decimal128 truncd128(_Decimal128 x);
    _Decimal64 roundevend64(_Decimal64 x);
    _Decimal32 roundevend32(_Decimal32 x);
    _Decimal128 roundevend128(_Decimal128 x);
    intmax_t fromfpd64(_Decimal64 x, int round, unsigned int width);

```

```

intmax_t fromfpd32(_Decimal32 x, int round, unsigned int width);
intmax_t fromfpd128(_Decimal128 x, int round, unsigned int width);
uintmax_t ufromfpd64(_Decimal64 x, int round, unsigned int width);
uintmax_t ufromfpd32(_Decimal32 x, int round, unsigned int width);
5 uintmax_t ufromfpd128(_Decimal128 x, int round, unsigned int width);

```

```

intmax_t fromfpd64(_Decimal64 x, int round, unsigned int width);
intmax_t fromfpd32(_Decimal32 x, int round, unsigned int width);
intmax_t fromfpd128(_Decimal128 x, int round, unsigned int width);
10 uintmax_t ufromfpd64(_Decimal64 x, int round, unsigned int width);
uintmax_t ufromfpd32(_Decimal32 x, int round, unsigned int width);
uintmax_t ufromfpd128(_Decimal128 x, int round, unsigned int width);

```

7.12.10 Remainder functions⁵

```

15 _Decimal64 fmodd64(_Decimal64 x, _Decimal64 y);
_Decimal32 fmodd32(_Decimal32 x, _Decimal32 y);
_Decimal128 fmodd128(_Decimal128 x, _Decimal128 y);

_Decimal64 remainderd64(_Decimal64 x, _Decimal64 y);
_Decimal32 remainderd32(_Decimal32 x, _Decimal32 y);
20 _Decimal128 remainderd128(_Decimal128 x, _Decimal128 y);

```

7.12.11 Manipulation functions

```

_Decimal64 copysignd64(_Decimal64 x, _Decimal64 y);
_Decimal32 copysignd32(_Decimal32 x, _Decimal32 y);
_Decimal128 copysignd128(_Decimal128 x, _Decimal128 y);
25 _Decimal64 nand64(const char *tagp);
_Decimal32 nand32(const char *tagp);
_Decimal128 nand128(const char *tagp);

_Decimal64 nextafterd64(_Decimal64 x, _Decimal64 y);
_Decimal32 nextafterd32(_Decimal32 x, _Decimal32 y);
30 _Decimal128 nextafterd128(_Decimal128 x, _Decimal128 y);

_Decimal64 nexttowardd64(_Decimal64 x, _Decimal128 y);
_Decimal32 nexttowardd32(_Decimal32 x, _Decimal128 y);
35 _Decimal128 nexttowardd128(_Decimal128 x, _Decimal128 y);

_Decimal64 nextupd64(_Decimal64 x);
_Decimal32 nextupd32(_Decimal32 x);
40 _Decimal128 nextupd128(_Decimal128 x);

_Decimal64 nextdownd64(_Decimal64 x);
_Decimal32 nextdownd32(_Decimal32 x);
45 _Decimal128 nextdownd128(_Decimal128 x);

_Decimal64 canonicalized64(_Decimal64 x);
_Decimal32 canonicalized32(_Decimal32 x);
_Decimal128 canonicalized128(_Decimal128 x);

```

7.12.12 Maximum, minimum, and positive difference functions

```

50 _Decimal64 fdimd64(_Decimal64 x, _Decimal64 y);
_Decimal32 fdimd32(_Decimal32 x, _Decimal32 y);

```

⁵ There is no decimal floating-point type version of the `remquo` function.


```

    _Decimal128 fdime,d128(_Decimal128 x, _Decimal128 y);

    _Decimal64 fmaxd64(_Decimal64 x, _Decimal64 y);
    _Decimal32 fmaxd32(_Decimal32 x, _Decimal32 y);
5    _Decimal128 fmaxd128(_Decimal128 x, _Decimal128 y);

    _Decimal64 fmind64(_Decimal64 x, _Decimal64 y);
    _Decimal32 fmind32(_Decimal32 x, _Decimal32 y);
10   _Decimal128 fmind128(_Decimal128 x, _Decimal128 y);

    _Decimal64 fmaxmagd64(_Decimal64 x, _Decimal64 y);
    _Decimal32 fmaxmagd32(_Decimal32 x, _Decimal32 y);
    _Decimal128 fmaxmagd128(_Decimal128 x, _Decimal128 y);

15   _Decimal64 fminmagd64(_Decimal64 x, _Decimal64 y);
    _Decimal32 fminmagd32(_Decimal32 x, _Decimal32 y);
    _Decimal128 fminmagd128(_Decimal128 x, _Decimal128 y);

```

7.12.13 Floating multiply-add

```

    _Decimal64 fmad64(_Decimal64 x, _Decimal64 y, _Decimal64 z);
20   _Decimal32 fmad32(_Decimal32 x, _Decimal32 y, _Decimal32 z);
    _Decimal128 fmad128(_Decimal128 x, _Decimal128 y, _Decimal128 z);

```

7.12.14 Functions that round result to narrower format

```

    _Decimal32 d32addd64(_Decimal64 x, _Decimal64 y);
    _Decimal32 d32addd128(_Decimal128 x, _Decimal128 y);
25   _Decimal64 d64addd128(_Decimal128 x, _Decimal128 y);

    _Decimal32 d32subd64(_Decimal64 x, _Decimal64 y);
    _Decimal32 d32subd128(_Decimal128 x, _Decimal128 y);
30   _Decimal64 d64addd128(_Decimal128 x, _Decimal128 y);

    _Decimal32 d32muld64(_Decimal64 x, _Decimal64 y);
    _Decimal32 d32muld128(_Decimal128 x, _Decimal128 y);
    _Decimal64 d64muld128(_Decimal128 x, _Decimal128 y);

35   _Decimal32 d32divd64(_Decimal64 x, _Decimal64 y);
    _Decimal32 d32divd128(_Decimal128 x, _Decimal128 y);
    _Decimal64 d64divd128(_Decimal128 x, _Decimal128 y);

    _Decimal32 d32fmad64(_Decimal64 x, _Decimal64 y, _Decimal64 z);
40   _Decimal32 d32fmad128(_Decimal128 x, _Decimal128 y, _Decimal128 z);
    _Decimal64 d64fmad128(_Decimal128 x, _Decimal128 y, _Decimal128 z);

    _Decimal32 d32sqrtd64(_Decimal64 x);
    _Decimal32 d32sqrtd128(_Decimal128 x);
45   _Decimal64 d64sqrtd128(_Decimal128 x);

```

F.10.13 Payload functions

```

    _Decimal64 getpayloadd64(const _Decimal64 *x);
    _Decimal32 getpayloadd32(const _Decimal32 *x);
50   _Decimal128 getpayloadd128(const _Decimal128 *x);

    int setpayloadd64(_Decimal64 *res, _Decimal64 pl);
    int setpayloadd32(_Decimal32 *res, _Decimal32 pl);
    int setpayloadd128(_Decimal128 *res, _Decimal128 pl);

```

```
int setpayloadsigd64(_Decimal64 *res, _Decimal64 p1);
int setpayloadsigd32(_Decimal32 *res, _Decimal32 p1);
int setpayloadsigd128(_Decimal128 *res, _Decimal128 p1);
```

Add to the end of 7.12.14 paragraph 1:

- 5 [1] ... If either argument has decimal floating type, the other argument shall have decimal floating type as well.

Replace 7.12.6.4 paragraphs 2 and 3:

[2] The `frexp` functions break a floating-point number into a normalized fraction and an integral power of 2. They store the integer in the `int` object pointed to by `exp`.

- 10 [3] If `value` is not a floating-point number or if the integral power of 2 is outside the range of `int`, the results are unspecified. Otherwise, the `frexp` functions return the value `x`, such that `x` has a magnitude in the interval $[1/2, 1)$ or zero, and `value` equals $x \times 2^{\text{exp}}$. If `value` is zero, both parts of the result are zero.

with the following:

- 15 [2] The `frexp` functions break a floating-point number into a normalized fraction and an integer exponent. They store the integer in the `int` object pointed to by `exp`. If `value` is a decimal floating-point number, the exponent is an integral power of 10; otherwise it is an integral power of 2.

- 20 [3] If `value` is not a floating-point number, the results are unspecified. Otherwise, the `frexp` functions return the value `x`, such that `x` has a magnitude in the interval $[1/10, 1)$ or zero, and `value` equals $x \times 10^{\text{exp}}$ when `value` is a decimal floating-point number, or `x` has a magnitude in the interval $[1/2, 1)$ or zero, and `value` equals $x \times 2^{\text{exp}}$ when `value` is a generic floating-point number. If `value` is zero, both parts of the result are zero.

25 **ISSUE 2: Should we change "If value is a decimal floating-point number, the exponent is an integral power of 10; otherwise it is an integral power of 2." to "If value is a generic floating-point value, the exponent is integral power of 2, otherwise if is an integral power of 10"? Can we assume generic RADIX is not 10?**

Replace 7.12.6.6 paragraphs 2 and 3:

[2] The `ldexp` functions multiply a floating-point number by an integral power of 2. A range error may occur.

- 30 [3] The `ldexp` functions return $x \times 2^{\text{exp}}$.

with the following:

[2] The `ldexp` functions multiply a decimal floating-point number by an integral power of 10, or a generic floating-point number by an integral power of 2. A range error may occur.

- 35 [3] If `x` is a decimal floating-point number, the `ldexp` functions return $x \times 10^{\text{exp}}$; otherwise they return $x \times 2^{\text{exp}}$.

ISSUE 3: For `ldexp`, and `scalbn` below, should we change "A range error may occur" to "A range error may occur for finite arguments"? Waiting for WG14 resolution of same issue for corresponding functions in C11.

Replace 7.12.6.11 paragraph 2:

[2] The `logb` functions extract the exponent of x , as a signed integer value in floating-point format. If x is subnormal it is treated as though it were normalized; thus, for positive finite x ,

$$1 \leq x \times \text{FLT_RADIX}^{-\text{logb}(x)} < \text{FLT_RADIX}$$

A domain error or pole error may occur if the argument is zero.

5 with the following:

[2] The `logb` functions extract the exponent of x , as a signed integer value in floating-point format. If x is subnormal it is treated as though it were normalized; thus, for positive finite x ,

$$1 \leq x \times b^{-\text{logb}(x)} < b$$

10 where $b = 10$ if x is a decimal floating-point number; otherwise $b = \text{FLT_RADIX}$. A domain error or range error may occur if the argument is zero.

Replace 7.12.6.13 paragraphs 2 and 3:

[2] The `scalbn` and `scalbln` functions compute $x \times \text{FLT_RADIX}^n$ efficiently, not normally by computing FLT_RADIX^n explicitly. A range error may occur.

[3] The `scalbn` and `scalbln` functions return $x \times \text{FLT_RADIX}^n$.

15 with the following:

[2] The `scalbn` and `scalbln` functions compute $x \times b^n$ (where $b = 10$ if x is a decimal floating-point number; otherwise $b = \text{FLT_RADIX}$) efficiently, not normally by computing b^n explicitly. A range error may occur.

[3] The `scalbn` and `scalbln` functions return $x \times b^n$.

20 12.4 New `<math.h>` functions

This clause suggests new functions to be added to `<math.h>`.

12.4.1 Quantum exponent functions

Suggested addition to C11:

7.12.11.5 The `quantize` functions

25 Synopsis

```
[1] #define __STDC_WANT_DEC_FP__
    #include <math.h>
    _Decimal32 quantized32 (_Decimal32 x, _Decimal32 y);
    _Decimal64 quantized64 (_Decimal64 x, _Decimal64 y);
30  _Decimal128 quantized128(_Decimal128 x, _Decimal128 y);
```

Description

35 [2] The `quantize` functions set the quantum exponent of argument x to the quantum exponent of argument y , while attempting to keep the value the same. If the quantum exponent is being increased, the value shall be correctly rounded according to the current rounding mode; if the result does not have the same value as x , the “inexact” floating-point exception shall be raised. If the quantum exponent is being decreased and the significand of the result has more digits than the type

would allow, the result is NaN and the “invalid” floating-point exception shall be raised. If one or both operands are NaN the result is NaN. Otherwise if only one operand is infinity, the result is NaN and the “invalid” floating-point exception shall be raised. If both operands are infinity, the result is **DEC_INFINITY** with the sign as **x**, converted to the type of the function. The **quantize** functions do not raise the “underflow” floating-point exception.

Returns

[3] The **quantize** functions return the number which is equal in value (except for any rounding) and sign to **x**, and which has a quantum exponent set to be equal to the quantum exponent of **y**.

7.12.11.6 The **samequantum** functions

Synopsis

```
[1] #define __STDC_WANT_DEC_FP__
    #include <math.h>
    _Bool samequantumd32 (_Decimal32 x, _Decimal32 y);
    _Bool samequantumd64 (_Decimal64 x, _Decimal64 y);
    _Bool samequantumd128 (_Decimal128 x, _Decimal128 y);
```

Description

[2] The **samequantum** functions determine if the quantum exponents of the **x** and **y** are the same. If both **x** and **y** are NaN, or infinity, they have the same quantum exponents; if exactly one operand is infinity or exactly one operand is NaN, they do not have the same quantum exponents. The **samequantum** functions raise no exception.

Returns

[3] The **samequantum** functions return nonzero (true) when **x** and **y** have the same quantum exponents, zero (false) otherwise.

7.12.11.7 The **quantexp** functions

Synopsis

```
[1] #define __STDC_WANT_DEC_FP__
    #include <math.h>
    int quantexpd32 (_Decimal32 x);
    int quantexpd64 (_Decimal64 x);
    int quantexpd128 (_Decimal128 x);
```

Description

[2] The **quantexp** functions compute the quantum exponent of a finite argument. If **x** is infinite or NaN, they compute **INT_MIN** and a domain error occurs.

Returns

[3] The **quantexp** functions return the quantum exponent of **x**.

12.4.2 Decimal re-encoding functions

Suggested addition to C11:

7.12.11.8 The **encodedpd** functions

Synopsis

```
[1] #define __STDC_WANT_DEC_FP__
#include <math.h>
dpdencodingd32_t encodedpdd32 (_Decimal32 x);
5 dpdencodingd64_t encodedpdd64 (_Decimal64 x);
dpdencodingd128_t encodedpdd128 (_Decimal128 x);
```

Description

[2] The `encodedpd` functions convert the argument into the DPD encoding. These functions preserve the value of `x` and raise no floating-point exceptions. If `x` is non-canonical, these functions may or may not produce a canonical representation.

Returns

[3] The `encodedpd` functions return the DPD encoding of `x`.

7.12.11.9 The `decodedpd` functions**Synopsis**

```
[1] #define __STDC_WANT_DEC_FP__
#include <math.h>
_Decimal32 decodedpdd32 (dpdencodingd32_t x);
_Decimal64 decodedpdd64 (dpdencodingd64_t x);
20 _Decimal128 decodedpdd128 (dpdencodingd128_t x);
```

Description

[2] The `decodedpd` functions convert the argument from the DPD encoding into a representation in the type of the function. These functions preserve the value of `x` and raise no floating-point exceptions. If `x` is non-canonical, these functions may or may not produce a canonical representation.

Returns

[3] The `decodedpd` functions return the converted representation.

7.12.11.10 The `encodebid` functions**Synopsis**

```
[1] #define __STDC_WANT_DEC_FP__
#include <math.h>
bidencodingd32_t encodebidd32 (_Decimal32 x);
bidencodingd64_t encodebidd64 (_Decimal64 x);
35 bidencodingd128_t encodebidd128 (_Decimal128 x);
```

Description

[2] The `encodebid` functions convert the argument into the BID encoding. These functions preserve the value of `x` and raise no floating-point exceptions. If `x` is non-canonical, these functions may or may not produce a canonical representation.

Returns

[3] The `encodebid` functions return the BID encoding of `x`.

7.12.11.11 The `decodebid` functions

Synopsis

```
[1] #define __STDC_WANT_DEC_FP__
    #include <math.h>
5   _Decimal32 decodebidd32 (bidencodingd32_t x);
    _Decimal64 decodebidd64 (bidencodingd64_t x);
    _Decimal128 decodebidd128 (bidencodingd128_t x);
```

Description

[2] The `decodebid` functions convert the argument from the BID encoding into a representation in the type of the function. These functions preserve the value of `x` and raise no floating-point exceptions. If `x` is non-canonical, these functions may or may not produce a canonical representation.

Returns

[3] The `decodebid` functions return the converted representation.

12.5 Formatted input/output specifiers

With the following decimal forms of the `a,A` format specifiers, the `printf` family of functions provide conversions to decimal character sequences that preserve quantum exponents, as required by IEC 60559.

Suggested changes to C11:

Add the following to 7.21.6.1 paragraph 7, to 7.21.6.2 paragraph 11, to 7.29.2.1 paragraph 7, and to 7.29.2.2 paragraph 11:

- H Specifies that a following `a`, `A`, `e`, `E`, `f`, `F`, `g`, or `G` conversion specifier applies to a `_Decimal32` argument.
- D Specifies that a following `a`, `A`, `e`, `E`, `f`, `F`, `g`, or `G` conversion specifier applies to a `_Decimal64` argument.
- DD Specifies that a following `a`, `A`, `e`, `E`, `f`, `F`, `g`, or `G` conversion specifier applies to a `_Decimal128` argument.

Change all occurrences of:

A `double` argument representing ...

in the descriptions of the `e`, `E`, `f`, `F`, `g`, and `G` conversion specifiers in 7.21.6.1 paragraph 8 and 7.29.2.1 paragraph 8 to:

A `double` or decimal floating type argument representing ...

Change the second paragraph in the description of the `a,A` conversion specifier in 7.21.6.1 paragraph 8 and 7.29.2.1. paragraph 8 from:

A `double` argument representing an infinity or NaN is converted in the style of an `f` or `F` conversion specifier.

to:

A `double` or decimal floating type argument representing an infinity or NaN is converted in the style of an `f` or `F` conversion specifier.

Add the following to 7.21.6.1 paragraph 8 and 7.29.2.1 paragraph 8, under **a,A** conversion specifiers:

If an **H**, **D**, or **DD** modifier is present and the precision is missing, then for a decimal floating type argument represented by a triple of integers (s , c , q), where n is the number of digits in the coefficient c ,

- 5 • if $0 \geq q \geq -(n+5)$, use style **f** formatting with formatting precision equal to $-q$,
- otherwise, use style **e** formatting with formatting precision equal to $n - 1$, with the exceptions that if $c = 0$ then the *digit-sequence* in the *exponent-part* shall have the value q (rather than 0), and that the exponent is always expressed with the minimum number of digits required to represent its value (the exponent never contains a leading zero).

10 If the precision is present (in the conversion specification) and is zero or at least as large as the precision p (5.2.4.2.2) of the decimal floating type, the conversion is as if the precision were missing. If the precision is present (and nonzero) and less than the precision p of the decimal floating type, the conversion first obtains an intermediate result by rounding the input in the type, according to the current rounding direction for decimal floating-point operations, to the number of digits specified by the precision, then converts the intermediate result as if the precision were missing. The length of the coefficient of the intermediate result is the smallest number, at least as large as the formatting precision, for which the quantum exponent is within the quantum exponent range of the type. The intermediate rounding may overflow.

15 EXAMPLE 1 Following are representations of `_Decimal64` arguments as triples (s , c , q) and the corresponding character sequences `printf` produces with `%Da`:

```

25           ( 1, 123, 0)        123
             (-1, 123, 0)       -123
             ( 1, 123, -2)       1.23
             ( 1, 123, 1)       1.23e+3
             (-1, 123, 1)       -1.23e+3
             ( 1, 123, -8)       0.00000123
             ( 1, 123, -9)       1.23e-7
             ( 1, 120, -8)       0.00000120
30           ( 1, 120, -9)       1.20e-7
             ( 1, 1234567890123456, 0)   1234567890123456
             ( 1, 1234567890123456, 1)   1.234567890123456e+16
             ( 1, 1234567890123456, -1)   123456789012345.6
             ( 1, 1234567890123456, -21) 0.000001234567890123456
35           ( 1, 1234567890123456, -22) 1.234567890123456e-7
             ( 1, 0, 0)         0
             (-1, 0, 0)        -0
             ( 1, 0, -6)        0.000000
             ( 1, 0, -7)        0e-7
40           ( 1, 0, 2)         0e+2
             ( 1, 5, -6)        0.000005
             ( 1, 50, -7)       0.0000050
             ( 1, 5, -7)        5e-7

```

45 EXAMPLE 2 To illustrate the effects of a precision specification, the sequence:

```

50        _Decimal32 x = 6543.00DF; // represented by the triple (1, 654300, -2)
          printf("%Ha\n", x);
          printf("%.6Ha\n", x);
          printf("%.5Ha\n", x);
          printf("%.4Ha\n", x);
          printf("%.3Ha\n", x);
          printf("%.2Ha\n", x);

```

```
printf("%.1Ha\n", x);
printf("%.0Ha\n", x);
```

assuming default rounding, results in:

```
5      6543.00
      6543.00
      6543.0
      6543
      6.54e+3
10     6.5e+3
      7e+3
      6543.00
```

EXAMPLE 3 To illustrate the effect of the exponent range, the sequence:

```
15     _Decimal32 x = 9543210e87DF; // represented by the triple (1, 9543210, 87)
      _Decimal32 y = 9500000e90DF; // represented by the triple (1, 9500000, 90)
      printf("%.6Ha\n", x);
      printf("%.5Ha\n", x);
      printf("%.4Ha\n", x);
20     printf("%.3Ha\n", x);
      printf("%.2Ha\n", x);
      printf("%.1Ha\n", x);
      printf("%.1Ha\n", y);
```

assuming default rounding, results in:

```
25     954321e88
      95432e89
      9543e90
      9540e90
30     9500e90
      10000e90
      inf
```

12.6 strtod32, strtod64, and strtod128 functions <stdlib.h>

The specifications of these functions are similar to those of `strtod`, `strtof`, and `strtold` as defined in C11 7.22.1.3. These functions are declared in `<stdlib.h>`.

Suggested addition to C11:

7.22.1.5 The `strtod32`, `strtod64`, and `strtod128` functions

Synopsis

```
40 [1] #define __STDC_WANT_DEC_FP__
      #include <stdlib.h>
      _Decimal32 strtod32 (const char * restrict nptr, char ** restrict
      endptr);
45     _Decimal64 strtod64 (const char * restrict nptr, char ** restrict
      endptr);
      _Decimal128 strtod128 (const char * restrict nptr, char ** restrict
      endptr);
```

Description

[2] The `strtod32`, `strtod64`, and `strtod128` functions convert the initial portion of the string pointed to by `nptr` to `_Decimal32`, `_Decimal64`, and `_Decimal128` representation, respectively. First, they decompose the input string into three parts: an initial, possibly empty, sequence of white-space characters (as specified by the `isspace` function), a subject sequence resembling a floating-point constant or representing an infinity or NaN; and a final string of one or more unrecognized characters, including the terminating null character of the input string. Then, they attempt to convert the subject sequence to a floating-point number, and return the result.

[3] The expected form of the subject sequence is an optional plus or minus sign, then one of the following:

- a nonempty sequence of decimal digits optionally containing a decimal-point character, then an optional exponent part as defined in 6.4.4.2;
- **INF** or **INFINITY**, ignoring case
- **NAN** or **NAN(*d-char-sequence_{opt}*)**, ignoring case in the **NAN** part, where:

d-char-sequence:
digit
d-char-sequence digit

The subject sequence is defined as the longest initial subsequence of the input string, starting with the first non-white-space character, that is of the expected form. The subject sequence contains no characters if the input string is not of the expected form.

[4] If the subject sequence has the expected form for a floating-point number, the sequence of characters starting with the first digit or the decimal-point character (whichever occurs first) is interpreted as a floating constant according to the rules of 6.4.4.2, except that it is not a hexadecimal floating number, that the decimal-point character is used in place of a period, and that if neither an exponent part nor a decimal-point character appears in a decimal floating point number, an exponent part of the appropriate type with value zero is assumed to follow the last digit in the string. If the subject sequence begins with a minus sign, the sequence is interpreted as negated (before rounding). A character sequence **INF** or **INFINITY** is interpreted as an infinity. A character sequence **NAN** or **NAN(*d-char-sequence_{opt}*)**, is interpreted as a quiet NaN; the meaning of the *d-char sequences* is implementation-defined.⁶ A pointer to the final string is stored in the object pointed to by `endptr`, provided that `endptr` is not a null pointer.

[5] If the sequence is negated, the sign *s* is set to -1 , else *s* is set to 1.

[6] If the subject sequence has the expected form for a floating-point number, then the result shall be correctly rounded as specified in IEC 60559.

[7] The coefficient *c* and the quantum exponent *q* of a finite converted floating-point number are determined from the subject sequence as follows:

- The *fractional-constant* or *digit-sequence* and the *exponent-part* (if any) are extracted from the subject sequence. If there is an *exponent-part*, then *q* is set to the value of *sign_{opt} digit-sequence* in the *exponent-part*. If there is no *exponent-part*, *q* is set to 0.
- If there is a *fractional-constant*, *q* is decreased by the number of digits to the right of the decimal point and the decimal point is removed to form a *digit-sequence*.
- *c* is set to the value of the *digit-sequence* (after any decimal point has been removed).
- Rounding required because of insufficient precision or range in the type of the result will round *c* to the full precision available in the type, and will adjust *q* accordingly within the limits of the type, provided the rounding does not yield an infinity (in which case an appropriately signed internal representation of infinity is returned). If the full precision of the type would

⁶ An implementation may use the *d-char sequence* to determine extra information to be represented in the NaN's significand.

require q to be smaller than the minimum for the type, then q is pinned at the minimum and c is adjusted through the subnormal range accordingly, perhaps to zero.

EXAMPLE Following are subject sequences of the decimal form and the resulting triples (s, c, q) produced by `strtod64`. Note that for `_Decimal64`, the precision (maximum coefficient length) is 16 and the quantum exponent range is $-398 \leq q \leq 369$.

"0"	(1,0,0)
"0.00"	(1,0,-2)
"123"	(1,123,0)
"-123"	(-1,123,0)
"1.23E3"	(1,123,1)
"1.23E+3"	(1,123,1)
"12.3E+7"	(1,123,6)
"12.0"	(1,120,-1)
"12.3"	(1,123,-1)
"0.00123"	(1,123,-5)
"-1.23E-12"	(-1,123,-14)
"1234.5E-4"	(1,12345,-5)
"-0"	(-1,0,0)
"-0.00"	(-1,0,-2)
"0E+7"	(1,0,7)
"-0E-7"	(-1,0,-7)
"12345678901234567890"	(1, 1234567890123457, 4) or (1, 1234567890123456, 4) depending on rounding mode
"1234E-400"	(1, 12, -398) or (1, 13, -398) depending on rounding mode
"1234E-402"	(1, 0, -398) or (1, 1, -398) depending on rounding mode
"1000."	(1,1000,0)
".0001"	(1,1,-4)
"1000.e0"	(1,1000,0)
".0001e0"	(1,1,-4)
"1000.0"	(1,10000,-1)
"0.0001"	(1,1,-4)
"1000.00"	(1,100000,-2)
"00.0001"	(1,1,-4)
"001000."	(1,1000,0)
"001000.0"	(1,10000,-1)
"001000.00"	(1,100000,-2)
"00.00"	(1,0,-2)
"00."	(1,0,0)
".00"	(1,0,-2)
"00.00e-5"	(1,0,-7)
"00.e-5"	(1,0,-5)
".00e-5"	(1,0,-7)

[8] In other than the "C" locale, additional locale-specific subject sequence forms may be accepted.

[9] If the subject sequence is empty or does not have the expected form, no conversion is performed; the value of `nptr` is stored in the object pointed to by `endptr`, provided that `endptr` is not a null pointer.

Returns

[10] The functions return the converted value, if any. If no conversion could be performed, `+0.E0dd` converted to type of the function is returned. If the correct value overflows and default rounding is in effect (7.12.1), plus or minus `HUGE_VAL_D64`, `HUGE_VAL_D32`, or `HUGE_VAL_D128` is returned (according to the return type and sign of the value), and the value of the macro `ERANGE` is stored in `errno`. If the result underflows (7.12.1), the functions return a value whose magnitude is no greater

than the smallest normalized positive number in the return type; whether `errno` acquires the value `ERANGE` is implementation-defined.

12.7 `wcstod32`, `wcstod64`, and `wcstod128` functions `<wchar.h>`

5 The specifications of these functions are similar to those of `wcstod`, `wcstof`, and `wcstold` as defined in C11 7.29.4.1.1. They are declared in `<wchar.h>`.

Suggested addition to C11:

7.29.4.1.3 The `wcstod32`, `wcstod64`, and `wcstod128` functions

Synopsis

```

10 [1] #define __STDC_WANT_DEC_FP__
    #include <wchar.h>
    _Decimal32 wcstod32 (const wchar_t * restrict nptr, wchar_t ** restrict
        endptr);
    _Decimal64 wcstod64 (const wchar_t * restrict nptr, wchar_t ** restrict
        endptr);
15 _Decimal128 wcstod128(const wchar_t * restrict nptr, wchar_t **
        restrict endptr);

```

Description

20 [2] The `wcstod32`, `wcstod64`, and `wcstod128` functions convert the initial portion of the wide string pointed to by `nptr` to `_Decimal32`, `_Decimal64`, and `_Decimal128` representation, respectively. First, they decompose the input string into three parts: an initial, possibly empty, sequence of white-space wide characters (as specified by the `iswspace` function), a subject sequence resembling a floating-point constant or representing an infinity or NaN; and a final wide string of one or more unrecognized wide characters, including the terminating null wide character of the input wide string.

25 Then, they attempt to convert the subject sequence to a floating-point number, and return the result.

[3] The expected form of the subject sequence is an optional plus or minus sign, then one of the following:

- a nonempty sequence of decimal digits optionally containing a decimal-point wide character, then an optional exponent part as defined in 6.4.4.2;
- 30 • `INF` or `INFINITY`, ignoring case
- `NAN` or `NAN(d-wchar-sequenceopt)`, ignoring case in the `NAN` part, where:

```

    d-wchar-sequence:
35     digit
    d-wchar-sequence digit

```

The subject sequence is defined as the longest initial subsequence of the input wide string, starting with the first non-white-space wide character, that is of the expected form. The subject sequence contains no wide characters if the input wide string is not of the expected form.

40 [4] If the subject sequence has the expected form for a floating-point number, the sequence of wide characters starting with the first digit or the decimal-point wide character (whichever occurs first) is interpreted as a floating constant according to the rules of 6.4.4.2, except that it is not a hexadecimal floating number, that the decimal-point wide character is used in place of a period, and that if neither an exponent part nor a decimal-point wide character appears in a decimal floating point number, an exponent part of the appropriate type with value zero is assumed to follow the last digit in the string.

45 If the subject sequence begins with a minus sign, the sequence is interpreted as negated (before rounding). A wide character sequence `INF` or `INFINITY` is interpreted as an infinity. A wide character sequence `NAN` or `NAN(d-wchar-sequenceopt)`, is interpreted as a quiet NaN; the meaning of

the d-wchar sequences is implementation-defined.⁷ A pointer to the final wide string is stored in the object pointed to by **endptr**, provided that **endptr** is not a null pointer.

[5] If the sequence is negated, the sign *s* is set to -1 , else *s* is set to 1 .

[6] If the subject sequence has the expected form for a floating-point number, then the result shall be correctly rounded as specified in IEC 60559.

[7] The coefficient *c* and the quantum exponent *q* of a finite converted floating-point number are determined from the subject sequence as follows:

- The *fractional-constant* or *digit-sequence* and the *exponent-part* (if any) are extracted from the subject sequence. If there is an *exponent-part*, then *q* is set to the value of *sign_{opt} digit-sequence* in the *exponent-part*. If there is no *exponent-part*, *q* is set to 0 .
- If there is a *fractional-constant*, *q* is decreased by the number of digits to the right of the decimal point and the decimal point is removed to form a *digit-sequence*.
- *c* is set to the value of the *digit-sequence* (after any decimal point has been removed).
- Rounding required because of insufficient precision or range in the type of the result will round *c* to the full precision available in the type, and will adjust *q* accordingly within the limits of the type, provided the rounding does not yield an infinity (in which case an appropriately signed internal representation of infinity is returned). If the full precision of the type would require *q* to be smaller than the minimum for the type, then *q* is pinned at the minimum and *c* is adjusted through the subnormal range accordingly, perhaps to zero.

[8] In other than the "C" locale, additional locale-specific subject sequence forms may be accepted.

[9] If the subject sequence is empty or does not have the expected form, no conversion is performed; the value of **nptr** is stored in the object pointed to by **endptr**, provided that **endptr** is not a null pointer.

Returns

[10] The functions return the converted value, if any. If no conversion could be performed, `+0.E0dd` converted to the type of the function is returned. If the correct value overflows and default rounding is in effect (7.12.1), plus or minus `HUGE_VAL_D64`, `HUGE_VAL_D32`, or `HUGE_VAL_D128` is returned (according to the return type and sign of the value), and the value of the macro `ERANGE` is stored in `errno`. If the result underflows (7.12.1), the functions return a value whose magnitude is no greater than the smallest normalized positive number in the return type; whether `errno` acquires the value `ERANGE` is implementation-defined.

12.8 Type-generic macros <tgmath.h>

All new functions added to <math.h> are subjected to the same requirement as specified in C11 7.25 to provide support for *type-generic* macro expansion. When one of the type-generic arguments has a decimal floating type, use of the type-generic macro invokes a function whose parameters have the types determined as follows:

If there are more than one real floating type arguments, usual arithmetic conversions are applied to the real floating type arguments so that they have compatible types. Then,

- If any argument has type `_Decimal128`, the type determined is `_Decimal128`.
- Otherwise, if any argument has type `_Decimal164`, or if one argument has an integer type and another argument has type `_Decimal132`, the type determined is `_Decimal164`.
- Otherwise, if any argument has type `_Decimal132`, the type determined is `_Decimal132`.
- Otherwise, the specification in C11 7.25 paragraph 3 applies.

⁷ An implementation may use the d-char sequence to determine extra information to be represented in the NaN's significand.

EXAMPLE

```
pow(2,3.0) // expands to the double version of pow:  
           // pow((double)2, (double)3.0)  
5 pow(2,3.DD) // expands to the _Decimal64 version of pow:  
           // powd64((_Decimal64)2, (_Decimal64)3.DD)
```

ISSUE 4: *Clause 12.8 needs to provide suggested changes to C11.*

ISSUE 5: *The study group is still considering other issues with tgmth for decimal.*

10 **ISSUE 6:** *We intend to add an example showing how generic selection can be used to define cbrt for tgmth to handle decimal as well as generic FP types.*

ISSUE 7: *We intend to add specification for determining the quantum exponent for decimal operation results.*

ISSUE 8: *We intend to add guidance for optimization that will preserve quantum exponents.*

Bibliography

- [1] ISO/IEC 9899:2011, *Information technology — Programming languages, their environments and system software interfaces — Programming Language C*
- 5 [2] ISO/IEC/IEEE 60559:2011, *Information technology — Microprocessor Systems — Floating-point arithmetic*
- [3] ISO/IEC TR 24732:2008, *Information technology – Programming languages, their environments and system software interfaces – Extension for the programming language C to support decimal floating-point arithmetic*
- [4] IEC 60559:1989, *Binary floating-point arithmetic for microprocessor systems, second edition*
- 10 [5] IEEE 754-2008, *IEEE Standard for Floating-Point Arithmetic*
- [6] IEEE 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*
- [7] IEEE 854-1987, *IEEE Standard for Radix-Independent Floating-Point Arithmetic*
- [8] ISO/IEC 1989:2002, *Information technology — Programming languages, their environments and system software interfaces — Programming languages COBOL.*
- 15 [9] ANSI X3.274, *Information Technology — Programming Language REXX.*