

PD ISO/IEC TS 18661-2:2015



BSI Standards Publication

Information Technology — Programming languages, their environments, and system software interfaces — Floating- point extensions for C

Part 2: Decimal floating-point arithmetic

bsi.

...making excellence a habit.
AVA Documents™
Document Sharing Hub

National foreword

This Published Document is the UK implementation of ISO/IEC TS 18661-2:2015.

The UK participation in its preparation was entrusted to Technical Committee IST/5, Programming languages, their environments and system software interfaces.

A list of organizations represented on this committee can be obtained on request to its secretary.

This publication does not purport to include all the necessary provisions of a contract. Users are responsible for its correct application.

© The British Standards Institution 2015. Published by BSI Standards Limited 2015

ISBN 978 0 580 85487 3

ICS 35.060

Compliance with a British Standard cannot confer immunity from legal obligations.

This Published Document was published under the authority of the Standards Policy and Strategy Committee on 28 February 2015.

Amendments issued since publication

Date	Text affected
------	---------------

TECHNICAL SPECIFICATION

ISO/IEC TS 18661-2

First edition
2015-02-15

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: Arithmétique décimal en virgule flottante



Reference number
ISO/IEC TS 18661-2:2015(E)

AVA Documents
Document Sharing Hub
© ISO/IEC 2015



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2015

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

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

Published in Switzerland

Foreword	iv
Introduction	v
1 Scope	1
2 Conformance	1
3 Normative references	1
4 Terms and definitions	1
5 C standard conformance	2
5.1 Freestanding implementations	2
5.2 Predefined macros	2
5.3 Standard headers	3
6 Decimal floating types	10
7 Characteristics of decimal floating types <float.h>	11
8 Operation binding	15
9 Conversions	16
9.1 Conversions between decimal floating and integer types	16
9.2 Conversions among decimal floating types, and between decimal floating and standard floating types	17
9.3 Conversions between decimal floating and complex types	18
9.4 Usual arithmetic conversions	18
9.5 Default argument promotion	18
10 Constants	18
11 Arithmetic operations	19
11.1 Operators	19
11.2 Functions	20
11.3 Conversions	21
11.4 Expression transformations	21
12 Library	21
12.1 Standard headers	21
12.2 Decimal floating-point environment in <fenv.h>	21
12.3 Decimal mathematics in <math.h>	25
12.4 Decimal-only functions in <math.h>	34
12.4.1 Quantum and quantum exponent functions	34
12.4.2 Decimal re-encoding functions	36
12.5 Formatted input/output specifiers	38
12.6 strtodN functions in <stdlib.h>	40
12.7 wcstodN functions in <wchar.h>	43
12.8 strfromdN functions in <stdlib.h>	44
12.9 Type-generic math for decimal in <tgmath.h>	45
Bibliography	50

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide 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 organizations, 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.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the WTO principles in the Technical Barriers to Trade (TBT) see the following URL: [Foreword - Supplementary information](#)

The committee responsible for this document is ISO/IEC JTC 1, *Information technology*, Subcommittee, SC 22, *Programming languages, their environments, and system software interfaces*.

ISO/IEC/TS 18661 consists of the following parts, under the general title *Information technology—Programming languages, their environments, and system software interfaces — Floating-point extensions for C*:

- *Part 1: Binary floating-point arithmetic*
- *Part 2: Decimal floating-point arithmetic*

The following parts are under preparation:

- *Part 3: Interchange and extended types*
- *Part 4: Supplementary functions*
- *Part 5: Supplementary attributes*

ISO/IEC/TS 18661-1 updates ISO/IEC 9899:2011, *Information technology — 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*.

ISO/IEC/TS 18661-2 supersedes ISO/IEC/TR 24732:2009, *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 18661-3, ISO/IEC TS 18661-4, and ISO/IEC TS 18661-5 specify extensions to ISO/IEC 9899:2011 for features recommended in ISO/IEC/IEEE 60559:2011.

Introduction

Background

IEC 60559 floating-point standard

The IEEE 754:1985 standard for binary floating-point arithmetic was motivated by an expanding diversity in floating-point data representation and arithmetic which made writing robust programs, debugging, and moving programs between systems exceedingly difficult. Now, the great majority of systems provide data formats and arithmetic operations according to this International Standard. The IEC 60559:1989 international standard is equivalent to IEEE 754-1985 standard. Its stated goals are the following:

- a) facilitate movement of existing programs from diverse computers to those that adhere to this International Standard;
- b) enhance the capabilities and safety available to programmers who, though not expert in numerical methods, can well be attempting to produce numerically sophisticated programs. However, we recognize that utility and safety are sometimes antagonists;
- c) encourage experts to develop and distribute robust and efficient numerical programs that are portable, by way of minor editing and recompilation, onto any computer that conforms to this International Standard and possesses adequate capacity. When restricted to a declared subset of the standard, these programs should produce identical results on all conforming systems;
- d) provide direct support for
 - 1) execution-time diagnosis of anomalies,
 - 2) smoother handling of exceptions, and
 - 3) interval arithmetic at a reasonable cost;
- e) provide for the development of
 - 1) standard elementary functions such as exp and cos,
 - 2) very high precision (multiword) arithmetic, and
 - 3) coupling of numerical and symbolic algebraic computation;
- f) enable rather than preclude further refinements and extensions.

To these ends, the standard specified a floating-point model comprised of the following:

- *formats* – for binary floating-point data, including representations for Not-a-Number (NaN) and signed infinities and zeros;
- *operations* – basic arithmetic operations (addition, multiplication, etc.) on the format data to compose a well-defined, closed arithmetic system; also specified conversions between floating-point formats and decimal character sequences, and a few auxiliary operations;
- *context* – status flags for detecting exceptional conditions (invalid operation, division by zero, overflow, underflow, and inexact) and controls for choosing different rounding methods.

The ISO/IEC/IEEE 60559:2011 international standard is equivalent to the IEEE 754-2008 standard for floating-point arithmetic which is a major revision to IEEE 754-1985.

The revised standard specifies more formats including decimal as well as binary. It adds a 128-bit binary format to its basic formats. It also defines extended formats for all of its basic formats. It then specifies data interchange formats (which may or may not be arithmetic), including a 16-bit binary format and an unbounded tower of wider formats. To conform to the floating-point standard, an implementation must provide at least one of the basic formats, along with the required operations.

The revised standard specifies more operations. New requirements include, among others, arithmetic operations that round their result to a narrower format than the operands (with just one rounding), more conversions with integer types, more classifications and comparisons, and more operations for managing flags and modes. New recommendations include an extensive set of mathematical functions and seven reduction functions for sums and scaled products.

The revised standard places more emphasis on the reproducible results which is reflected in its standardization of more operations. For most parts, behaviors are completely specified. The standard requires conversions between floating-point formats and decimal character sequences to be correctly rounded for at least three more decimal digits than what is required to distinguish all numbers in the widest supported binary format. It also fully specifies conversions involving any number of decimal digits. It then recommends that transcendental functions be correctly rounded.

The revised standard requires a way to specify a constant rounding direction for a static portion of code with details left to programming language standards. This feature potentially allows rounding control without incurring the overhead of runtime access to a global (or thread) rounding mode.

Other features recommended by the revised standard include alternate methods for exception handling, controls for expression evaluation (allowing or disallowing various optimizations), support for fully reproducible results, and support for program debugging.

The revised standard, like its predecessor, defines its model of floating-point arithmetic 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 can be exchanged between different implementations that conform to the specification.

IEC 60559 does not include bindings of its floating-point model for particular programming languages. However, the revised standard does include guidance for programming language standards in recognition of the fact that features of the floating-point standard, even if well supported in the hardware, are not available to users unless the programming language provides a commensurate level of support. The implementation's combination of both hardware and software determines conformance to the floating-point standard.

C support for IEC 60559

The C standard specifies floating-point arithmetic using an abstract model. The representation of a floating-point number is specified in an abstract form where the constituent components (sign, exponent, significand) of the representation are defined, 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 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, before the floating-point standard was established, there were various hardware implementations of floating-point arithmetic in common use. Specifying the exact details of a representation would have made most of the existing implementations at the time not conforming.

Beginning with ISO/IEC 9899:1999, (C99), C has included an optional second level of specification for implementations supporting the floating-point standard. C99, in conditionally normative Annex F, introduced nearly complete support for the IEC 60559:1989 standard for binary floating-point arithmetic. Also, C99's informative Annex G offered a specification of complex arithmetic that is compatible with IEC 60559:1989.

ISO/IEC 9899:2011, (C11) includes refinements to the C99 floating-point specification, though it is still based on IEC 60559. C11:1989 upgraded Annex G from "informative" to "conditionally normative".

ISO/IEC/TR 24732:2009 introduced partial C support for the decimal floating-point arithmetic in ISO/IEC/IEEE 60559:2011. ISO/IEC/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 ISO/IEC/IEEE 60559:2011 decimal formats, though it does not include all of the operations required by ISO/IEC/IEEE 60559:2011.

Purpose

The purpose of this International Standard is to provide a C language binding for ISO/IEC/IEEE 60559:2011 based on the C11 standard that delivers the goals of ISO/IEC/IEEE 60559 to users and is feasible to be implemented. It is then organized into five parts.

ISO/IEC/TS 18661-1 provides changes to C11 that cover all the requirements plus some basic recommendations of ISO/IEC/IEEE 60559:2011 for binary floating-point arithmetic. C implementations intending to support ISO/IEC/IEEE 60559:2011 are expected to conform to conditionally normative Annex F as enhanced by the changes in ISO/IEC TS 18661-1.

ISO/IEC/TS 18661-2 enhances ISO/IEC/TR 24732 to cover all the requirements plus some basic recommendations of ISO/IEC/IEEE 60559:2011 for decimal floating-point arithmetic. C implementations intending to provide an extension for decimal floating-point arithmetic supporting ISO/IEC/IEEE 60559:2011 are expected to conform to ISO/IEC TS 18661-2.

ISO/IEC/TS 18661-3 (Interchange and extended types), ISO/IEC/TS 18661-4 (Supplementary functions), and ISO/IEC/TS 18661-5 (Supplementary attributes) cover recommended features of ISO/IEC/IEEE 60559:2011. C implementations intending to provide extensions for these features are expected to conform to the corresponding parts.

Additional background on decimal floating-point arithmetic

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/IEEE 60559 standard).

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.

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 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.

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 also affects execution time significantly and the resulting program is often difficult to maintain and enhance.

Even though the hardware might 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.

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

Information Technology — Programming languages, their environments, and system software interfaces — Floating-point extensions for C —

Part 2: Decimal floating-point arithmetic

1 Scope

This part of ISO/IEC/TS 18661 extends programming language C as specified in ISO/IEC 9899:2011, (C11) with changes specified in ISO/IEC/TS 18661-1, 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.

This part of ISO/IEC/TS 18661 supersedes ISO/IEC/TR 24732:2009.

This part of ISO/IEC/TS 18661 does not cover binary floating-point arithmetic (which is covered in ISO/IEC/TS 18661-1), nor does it cover most optional features of IEC 60559.

2 Conformance

An implementation conforms to this part of ISO/IEC/TS 18661 if

- a) it meets the requirements for a conforming implementation of C11 with all the changes to C11 specified in ISO/IEC/TS 18661-1 and in this part of ISO/IEC/TS 18661, and
- b) it defines `__STDC_IEC_60559_DFP__` to 201yymmL.

NOTE Conformance to this part of ISO/IEC/TS 18661 does not include all the requirements of ISO/IEC/TS 18661-1. An implementation can conform to either or both of ISO/IEC/TS 18661-1 and this part of ISO/IEC/TS 18661.

3 Normative references

The following documents, in whole or in part, are normatively referenced in this document and are indispensable for its application. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 9899:2011, *Information technology — Programming languages — C*

ISO/IEC/IEEE 60559:2011, *Information technology — Microprocessor Systems — Floating-Point arithmetic*

ISO/IEC/TS 18661-1, *Information technology — Programming languages, their environments, and system software interfaces — Floating-point extensions for C — Part 1: Binary floating-point arithmetic*

4 Terms and definitions

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

4.1 C11

standard ISO/IEC 9899:2011, *Information technology — Programming languages C*, including *Technical Corrigendum 1* (ISO/IEC 9899-1:2011/Cor:2012)

5 C standard conformance

5.1 Freestanding implementations

The following change to C11 + TS18661-1 expands the conformance requirements for freestanding implementations so that they might conform to this part of ISO/IEC/TS 18661.

Change to C11 + TS18661-1:

Replace the fourth sentence of 4#6:

The strictly conforming programs that shall be accepted by a conforming freestanding implementation that defines `__STDC_IEC_60559_BFP__` may also use features in the contents of the standard headers `<fenv.h>` and `<math.h>` and the numeric conversion functions (7.22.1) of the standard header `<stdlib.h>`.

with:

The strictly conforming programs that shall be accepted by a conforming freestanding implementation that defines `__STDC_IEC_60559_BFP__` or `__STDC_IEC_60559_DFP__` may also use features in the contents of the standard headers `<fenv.h>` and `<math.h>` and the numeric conversion functions (7.22.1) of the standard header `<stdlib.h>`.

5.2 Predefined macros

The following change to C11 + TS18661-1 replaces `__STDC_DEC_FP__`, the conformance macro for decimal floating-point arithmetic specified in TR 24732, with `__STDC_IEC_60559_DFP__`, for consistency with the conformance macro for ISO/IEC/TS 18661-1. Note that an implementation may continue to define `__STDC_DEC_FP__`, so that programs that use `__STDC_DEC_FP__` may remain valid under the changes in Part 2 of Technical Specification 18661.

Change to C11 + TS18661-1:

In 6.10.8.3#1, add:

`__STDC_IEC_60559_DFP__` The integer constant 201ymmL, intended to indicate support of decimal floating types and conformance with Annex F for IEC 60559 decimal floating-point arithmetic.

The following change to C11 + TS18661-1 specifies the applications of Annex F to binary and decimal floating-point arithmetic.

Change to C11 + TS18661-1:

Replace F.1#3:

[3] An implementation that defines `__STDC_IEC_60559_BFP__` to 201ymmL shall conform to the specifications in this Annex.356) Where a binding between the C language and IEC 60559 is indicated, the IEC 60559-specified behavior is adopted by reference, unless stated otherwise.

with:

[3] An implementation that defines `__STDC_IEC_60559_BFP__` to 201ymmL shall conform to the specifications in this annex for binary floating-point arithmetic.356)

[4] An implementation that defines `__STDC_IEC_60559_DFP__` to 201ymmL shall conform to the specifications for decimal floating-point arithmetic in the following subclauses of this annex:

- F.2.1 Infinities and NaNs
- F.3 Operations
- F.4 Floating to integer conversions
- F.6 The return statement
- F.7 Contracted expressions
- F.8 Floating-point environment
- F.9 Optimization
- F.10 Mathematics `<math.h>`

For the purpose of specifying these conformance requirements, the macros, functions, and values mentioned in the subclauses listed above are understood to refer to the corresponding macros, functions, and values for decimal floating types. Likewise, the “rounding direction mode” is understood to refer to the rounding direction mode for decimal floating-point arithmetic.

[5] Where a binding between the C language and IEC 60559 is indicated, the IEC 60559-specified behavior is adopted by reference, unless stated otherwise.

5.3 Standard headers

The new identifiers added to C11 library headers by this part of ISO/IEC/TS 18661 are defined or declared by their respective headers only if `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where the appropriate header is first included. The macro `__STDC_WANT_IEC_60559_DFP_EXT__` replaces the macro `__STDC_WANT_DEC_FP__` specified in TR 24732 for the same purpose. The following changes to C11 + TS18661-1 list these identifiers in each applicable library subclause.

Changes to C11 + TS18661-1:

In 5.2.4.2.1#1a, change:

[1a] The following identifiers are defined only if `__STDC_WANT_IEC_60559_BFP_EXT__` is defined as a macro at the point in the source file where `<limits.h>` is first included:

to:

[1a] The following identifiers are defined only if `__STDC_WANT_IEC_60559_BFP_EXT__` or `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where `<limits.h>` is first included:

After 5.2.4.2.2#6a, insert the paragraph:

[6b] The following identifiers are defined only if `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where `<float.h>` is first included:

for $N = 32, 64$, and 128 :

<code>DECN_MANT_DIG</code>	<code>DECN_MAX</code>	<code>DECN_TRUE_MIN</code>
<code>DECN_MIN_EXP</code>	<code>DECN_EPSILON</code>	
<code>DECN_MAX_EXP</code>	<code>DECN_MIN</code>	

After 7.6#3a, insert the paragraph:

[3b] The following identifiers are declared only if `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where `<fenv.h>` is first included:

`fe_dec_getround`

`fe_dec_setround`

Change 7.12#1a from:

[1a] The following identifiers are defined or declared only if `__STDC_WANT_IEC_60559_BFP_EXT__` is defined as a macro at the point in the source file where `<math.h>` is first included:

<code>FP_INT_UPWARD</code>	<code>FP_FAST_FSUB</code>
<code>FP_INT_DOWNWARD</code>	<code>FP_FAST_FSUBL</code>
<code>FP_INT_TOWARDZERO</code>	<code>FP_FAST_DSUBL</code>
<code>FP_INT_TONEAREST- FROMZERO</code>	<code>FP_FAST_FMUL</code>
<code>FP_INT_TONEAREST</code>	<code>FP_FAST_FMULL</code>
<code>FP_LLOGB0</code>	<code>FP_FAST_DMULL</code>
<code>FP_LLOGBNAN</code>	<code>FP_FAST_FDIV</code>
<code>SNANF</code>	<code>FP_FAST_FDIVL</code>
<code>SNAN</code>	<code>FP_FAST_DDIVL</code>
<code>SNANL</code>	<code>FP_FAST_FSQRT</code>
<code>FP_FAST_FADD</code>	<code>FP_FAST_FSQRTL</code>
<code>FP_FAST_FADDL</code>	<code>FP_FAST_DSQRTL</code>
<code>FP_FAST_DADDL</code>	

<code>iseqsig</code>	<code>fmaxmagf</code>	<code>ffmal</code>
<code>iscanonical</code>	<code>fmaxmagl</code>	<code>dfmal</code>
<code>issignaling</code>	<code>fminmag</code>	<code>fsqrt</code>
<code>issubnormal</code>	<code>fminmagf</code>	<code>fsqrtl</code>
<code>iszero</code>	<code>fminmagl</code>	<code>dsqrtl</code>
<code>fromfp</code>	<code>nextup</code>	<code>totalorder</code>
<code>fromfpf</code>	<code>nextupf</code>	<code>totalorderf</code>
<code>fromfpl</code>	<code>nextupl</code>	<code>totalorderl</code>
<code>ufromfp</code>	<code>nextdown</code>	<code>totalordermag</code>
<code>ufromfpf</code>	<code>nextdownf</code>	<code>totalordermagf</code>
<code>ufromfpl</code>	<code>nextdownl</code>	<code>totalordermagl</code>
<code>fromfpx</code>	<code>fadd</code>	<code>canonicalize</code>
<code>fromfpxf</code>	<code>faddl</code>	<code>canonicalizef</code>
<code>fromfpxl</code>	<code>daddl</code>	<code>canonicalizel</code>

ufromfpx	fsub	getpayload
ufromfpxf	fsubl	getpayloadf
ufromfpxl	dsubl	getpayloadl
roundeven	fmul	setpayload
roundevenf	fmull	setpayloadf
roundevenl	dmull	setpayloadl
llogb	fdiv	setpayloadsig
llogbf	fdivl	setpayloadsigf
llogbl	ddivl	setpayloadsigl
fmaxmag	ffma	

to:

[1a] The following identifiers are defined only if `__STDC_WANT_IEC_60559_BFP_EXT__` or `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where `<math.h>` is first included:

FP_INT_UPWARD	FP_LLOGBNAN
FP_INT_DOWNWARD	iseqsig
FP_INT_TOWARDZERO	iscanonical
FP_INT_TONEARESTFROMZERO	issignaling
FP_INT_TONEAREST	issubnormal
FP_LLOGB0	iszero

[1b] The following identifiers are defined or declared only if `__STDC_WANT_IEC_60559_BFP_EXT__` is defined as a macro at the point in the source file where `<math.h>` is first included:

SNANF	ufromfpxf	dmull
SNAN	ufromfpxl	fdiv
SNANL	roundeven	fdivl
FP_FAST_FADD	roundevenf	ddivl
FP_FAST_FADDL	roundevenl	ffma
FP_FAST_DADDL	llogb	ffmal
FP_FAST_FSUB	llogbf	dfmal
FP_FAST_FSUBL	llogbl	fsqrt
FP_FAST_DSUBL	fmaxmag	fsqrtl
FP_FAST_FMUL	fmaxmagf	dsqrtl
FP_FAST_FMULL	fmaxmagl	totalorder

FP_FAST_DMULL	fminmag	totalorderf
FP_FAST_FDIV	fminmagf	totalorderl
FP_FAST_FDIVL	fminmagl	totalordermag
FP_FAST_DDIVL	nextup	totalordermagf
FP_FAST_FSQRT	nextupf	totalordermagl
FP_FAST_FSQRTL	nextupl	canonicalize
FP_FAST_DSQRTL	nextdown	canonicalizef
fromfp	nextdownf	canonicalizel
fromfpf	nextdownl	getpayload
fromfpl	fadd	getpayloadf
ufromfp	faddl	getpayloadl
ufromfpf	daddl	setpayload
ufromfpl	fsub	setpayloadf
fromfpx	fsubl	setpayloadl
fromfpxf	dsubl	setpayloadsig
fromfpxl	fmul	setpayloadsigf
ufromfpx	fmull	setpayloadsigl

[1c] The following identifiers are defined or declared only if `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where `<math.h>` is first included:

<code>_Decimal32_t</code>	<code>DEC_INFINITY</code>
<code>_Decimal64_t</code>	<code>DEC_NAN</code>

and for $N = 32, 64, 128$:

<code>HUGE_VAL_DN</code>	<code>modfdN</code>	<code>remainderdN</code>
<code>SNANDN</code>	<code>scalbndN</code>	<code>copysigndN</code>
<code>FP_FAST_FMADN</code>	<code>scalblndN</code>	<code>nandN</code>
<code>acosdN</code>	<code>cbrtdN</code>	<code>nextafterdN</code>
<code>asindN</code>	<code>fabsdN</code>	<code>nexttowarddN</code>
<code>atandN</code>	<code>hypotdN</code>	<code>nextupdN</code>
<code>atan2dN</code>	<code>powdN</code>	<code>nextdowndN</code>
<code>cosdN</code>	<code>sqrtdN</code>	<code>canonicalizedN</code>
<code>sindN</code>	<code>erfdN</code>	<code>fdimdN</code>
<code>tandN</code>	<code>erfcdN</code>	<code>fmaxdN</code>

acoshdN	lgammadN	fmindN
asinhdN	tgammadN	fmaxmagdN
atanhdN	ceildN	fminmagdN
coshdN	floordN	fmadN
sinhdN	nearbyintdN	totalorderdN
tanhdN	rintdN	totalordermagdN
expdN	lrintdN	getpayloaddN
exp2dN	llrintdN	setpayloaddN
expm1dN	rounddN	setpayloadsigdN
frexpdpN	lrounddN	quantizedN
ilogbdN	llrounddN	samequantumdN
llogbdN	truncdN	quantumdN
ldexpdpN	roundevendN	llquantexpdpN
logdN	fromfpdpN	encodedecdpN
log10dpN	ufromfpdpN	decodedecdpN
log1pdpN	fromfpxdpN	encodebindN
log2dpN	ufromfpxdpN	decodebindN
logbdN	fmoddN	

and for $(M,N) = (32,64), (32,128), (64,128)$:

FP_FAST_DMADDDN	FP_FAST_DMFMADN	dMmuldpN
FP_FAST_DMSUBDN	FP_FAST_DMSQRTDN	dMdivdpN
FP_FAST_DMMULDN	dMadddpN	dMfmadpN
FP_FAST_DMDIVDN	dMsubdpN	dMsqrtdpN

In 7.20#4a, change:

[4a] The following identifiers are defined only if `__STDC_WANT_IEC_60559_BFP_EXT__` is defined as a macro at the point in the source file where `<stdint.h>` is first included:

to:

[4a] The following identifiers are defined only if `__STDC_WANT_IEC_60559_BFP_EXT__` or `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where `<stdint.h>` is first included:

After 7.22#1a, insert the paragraph:

[1b] The following identifiers are declared only if `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where `<stdlib.h>` is first included:

strfromd32 strfromd128 strtod64

strfromd64

strtod32

strtod128

Change 7.25#1a from:

[1a] The following identifiers are defined as type-generic macros only if `__STDC_WANT_IEC_60559_BFP_EXT__` is defined as a macro at the point in the source file where `<tgmath.h>` is first included:

<code>roundeven</code>	<code>fromfpx</code>	<code>fmul</code>
<code>llogb</code>	<code>ufromfpx</code>	<code>dmul</code>
<code>fmaxmag</code>	<code>totalorder</code>	<code>fdiv</code>
<code>fminmag</code>	<code>totalordermag</code>	<code>ddiv</code>
<code>nextup</code>	<code>fadd</code>	<code>ffma</code>
<code>nextdown</code>	<code>dadd</code>	<code>dfma</code>
<code>fromfp</code>	<code>fsub</code>	<code>fsqrt</code>
<code>ufromfp</code>	<code>dsub</code>	<code>dsqrt</code>

to:

[1a] The following identifiers are defined as type-generic macros only if `__STDC_WANT_IEC_60559_BFP_EXT__` or `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where `<tgmath.h>` is first included:

<code>roundeven</code>	<code>nextup</code>	<code>fromfpx</code>
<code>llogb</code>	<code>nextdown</code>	<code>ufromfpx</code>
<code>fmaxmag</code>	<code>fromfp</code>	<code>totalorder</code>
<code>fminmag</code>	<code>ufromfp</code>	<code>totalordermag</code>

[1b] The following identifiers are defined as type-generic macros only if `__STDC_WANT_IEC_60559_BFP_EXT__` is defined as a macro at the point in the source file where `<tgmath.h>` is first included:

<code>fadd</code>	<code>fmul</code>	<code>ffma</code>
<code>dadd</code>	<code>dmul</code>	<code>dfma</code>
<code>fsub</code>	<code>fdiv</code>	<code>fsqrt</code>
<code>dsub</code>	<code>ddiv</code>	<code>dsqrt</code>

[1c] The following identifiers are defined as type-generic macros only if `__STDC_WANT_IEC_60559_DFP_EXT__` is defined as a macro at the point in the source file where `<tgmath.h>` is first included:

<code>d32add</code>	<code>d64add</code>	<code>quantize</code>
<code>d32sub</code>	<code>d64sub</code>	<code>samequantum</code>
<code>d32mul</code>	<code>d64mul</code>	<code>quantum</code>
<code>d32div</code>	<code>d64div</code>	<code>llquantexp</code>
<code>d32fma</code>	<code>d64fma</code>	
<code>d32sqrt</code>	<code>d64sqrt</code>	

6 Decimal floating types

This part of ISO/IEC/TS 18661 introduces three decimal floating types, designated as `_Decimal32`, `_Decimal64` and `_Decimal128`. These types support the IEC 60559 decimal formats: decimal32, decimal64, and decimal128.

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

This part of ISO/IEC/TS 18661 introduces the term *standard floating types* to refer to the types `float`, `double`, and `long double`, which are the floating types the C Standard requires unconditionally.

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. Regardless of the representation, the decimal floating types are distinct from the types `float`, `double`, and `long double`.

NOTE This part of ISO/IEC/TS 18661 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`.

Changes to C11 + TS18661-1:

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 *standard floating types*, designated as `float`, `double`, and `long double`.

Add the following paragraphs after 6.2.5#10:

[10a] There are three *decimal floating types*, designated as `_Decimal32`, `_Decimal64`, and `_Decimal128`. Respectively, they have the IEC 60559 formats: decimal32, decimal64, and decimal128. Decimal floating types are real floating types.

[10b] The standard floating types and the decimal floating types are collectively called the *real floating types*.

In 6.2.5#10a, attach a footnote to the wording:

they have the IEC 60559 formats: decimal32

where the footnote is:

*) IEC 60559 specifies decimal32 as a data-interchange format that does not require arithmetic support; however, `_Decimal32` is a fully supported arithmetic type.

Add the following to 6.4.1 Keywords:

keyword:

`_Decimal32`
`_Decimal64`
`_Decimal128`

Add the following to 6.7.2 Type specifiers:

type-specifier:

`_Decimal32`
`_Decimal64`
`_Decimal128`