

N3163

TECHNICAL
SPECIFICATION

ISO/IEC TS
18661-4

Second edition
CFP Working Draft
2023-08-31

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

Part 4:

Supplementary functions

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

Partie 4: Fonctions supplémentaires



COPYRIGHT PROTECTED DOCUMENT

10 © ISO/IEC 2023

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
15 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
5	3 Normative references	1
	4 Terms and definitions	1
	4.1	2
	5 C standard conformance	2
	5.1 Freestanding implementations	2
10	5.2 Predefined macros	2
	5.3 Standard headers	2
	6 Reduction functions in <code><math.h></code>	2
	6.1 The <code>reduc_sum</code> functions	3
	6.2 The <code>reduc_sumabs</code> functions	4
15	6.3 The <code>reduc_sumsq</code> functions	5
	6.4 The <code>reduc_sumprod</code> functions	6
	6.5 The <code>scaled_prod</code> functions	7
	6.6 The <code>scaled_prodsum</code> functions	8
	6.7 The <code>scaled_proddiff</code> functions	9
20	7 Augmented arithmetic functions in <code><math.h></code>	11
	7.1 The <code>augadd</code> functions	12
	7.2 The <code>augsub</code> functions	13
	7.3 The <code>augmul</code> functions	14
	8 Augmented arithmetic macros in <code><tgmath.h></code>	15
25		

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

Parts 1, 2, 3, and some of Part 4 are integrated into ISO/IEC 9899: 2024 (C23).

ISO/IEC TS 18661 Part 4 Version 2, this document, supersedes ISO/IEC TS 18661-4:2015, the previous version of Part 4.

ISO/IEC TS 18661 Part 5 Version 2, a separate document, supersedes ISO/IEC TS 18661-5:2015, the previous version of Part 5.

Introduction

Background

IEC 60559 floating-point standard

5 The IEC 60559 international standard and the corresponding version of the IEEE 754 standard have equivalent content.

Floating-point standards – matching versions

IEEE 754-1985	IEC 559:1989
IEEE 754-2008	ISO/IEC/IEEE 60559:2011
IEEE 754-2019	ISO/IEC 60559:2020

10 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 standard. The stated goals of this standard were (and have remained throughout its revisions) the following:

- 1 Facilitate movement of existing programs from diverse computers to those that adhere to this standard.
- 15 2 Enhance the capabilities and safety available to programmers who, though not expert in numerical methods, may well be attempting to produce numerically sophisticated programs. However, we recognize that utility and safety are sometimes antagonists.
- 20 3 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 standard and possesses adequate capacity. When restricted to a declared subset of the standard, these programs should produce identical results on all conforming systems.
- 4 Provide direct support for
 - a. Execution-time diagnosis of anomalies
 - b. Smoother handling of exceptions
 - 25 c. Interval arithmetic at a reasonable cost
- 5 Provide for development of
 - a. Standard elementary functions such as exp and cos
 - b. Very high precision (multiword) arithmetic
 - c. Coupling of numerical and symbolic algebraic computation
- 30 6 Enable rather than preclude further refinements and extensions.

To these ends, the standard specified a floating-point model comprising 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 IEEE 754-2008 standard for floating-point arithmetic, equivalent to the ISO/IEC/IEEE 60559:2011 international standard, was a major revision. This revision:

- Specified more formats, including decimal as well as binary. It added a 128-bit binary format to its basic formats. It defined extended formats corresponding to all its basic formats. It specified 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.
- Specified more operations. It added required operations including (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. It added recommended operations including an extensive set of mathematical operations and seven reduction operations for sums and scaled products.
- Placed more emphasis on reproducible results. This is reflected in its standardization of more operations. For the most part, it completely specified behaviors. It required conversions between floating-point formats and decimal character sequences to be correctly rounded for at least three more decimal digits than is necessary to distinguish all numbers in the widest supported binary format; it completely specified such conversions involving any number of decimal digits. It specified the recommended transcendental functions to be correctly rounded.
- Added 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.
- Added other recommended features including 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 IEEE 754-2019 standard for floating-point arithmetic, equivalent to the ISO/IEC 60559:2020 international standard, was a minor revision. As such it was limited to upward-compatible editorial corrections and clarifications and minor enhancements. It added some recommended operations, including ones that might be required features in the next revision.

IEC 60559 (like IEEE 754) defines specific encodings for the exchange of floating-point data between different implementations. However, it does not define the concrete representation (specific layout in storage, or in a processor's register, for example) of data or context.

IEC 60559 (like IEEE 754) does not specify how its features are expressed in programming languages. However, its revisions have added guidance for programming language standards, in recognition of the fact that benefits of the floating-point standard, even if well supported in the hardware, are not

available to users unless the programming language provides interfaces for the features and reliable behaviors. The implementation's combination of both hardware and software determines conformance to the floating-point standard.

C support for IEC 60559

- 5 The C standard specifies floating-point arithmetic using an abstract model. The representation of a floating-point number is specified in a 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.
- 10 Furthermore, certain behaviors of most floating-point operations are also implementation-defined or unspecified, including accuracy and aspects of the way special values and exceptional conditions are handled.

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 the model would have made most of the

15 existing implementations at the time non-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

20 compatible with IEC 60559:1989.

ISO/IEC 9899:2011 (C11) and ISO/IEC 9899:2018 (C17) include refinements to the C99 floating-point specification, though are still based on IEC 60559:1989. C11 upgraded annex G from "informative" to "conditionally normative".

25 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 the operations required by ISO/IEC/IEEE 60559:2011.

30 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 the operations required by ISO/IEC/IEEE 60559:2011.

35 ISO/IEC TS 18661 provided a C language binding for ISO/IEC/IEEE 60559:2011, based on the C11 standard. ISO/IEC TS 18661 was organized into five parts:

ISO/IEC TS 18661-1:2014 – Binary floating-point arithmetic

ISO/IEC TS 18661-2:2015 – Decimal floating-point arithmetic, Second edition

ISO/IEC TS 18661-3:2015 – Interchange and extended types

40 ISO/IEC TS 18661-4:2015 – Supplementary functions

ISO/IEC TS 18661-5:2016 – Supplementary attributes

ISO/IEC 9899:2024 (C23) incorporates ISO/IEC TS 18661 Parts 1 and 2, the mathematical functions in Part 4, and, in an annex, Part 3. C23 also updates its floating-point specification to support ISO/IEC/IEEE 60559:2020.

5 This document updates ISO/IEC TS 18661-4. It retains the feature that was not incorporated into C23, namely the reduction functions. It also adds support for the augmented arithmetic operations introduced in IEC 60559:2020.

A separate document updates ISO/IEC TS 18661-5, which was not incorporated into C23.

Additional background on supplementary functions

10 This document uses the term supplementary functions to refer to functions that provide operations recommended, but not required, by IEC 60559 and that are not supported in C23.

IEC 60559 specifies and recommends reduction operations which operate on vector operands to perform widely used vector computations. These operations, which compute sums and products, may associate in any order and may evaluate in any wider format. Hence, unlike other IEC 60559 operations, they do not have unique specified results.

15 IEC 60559 also specifies and recommends augmented arithmetic operations, which are versions of operations commonly called twoSum and twoProduct. These operations can be used to implement arithmetic with extra precision, for example for double-double format. They can also be used to implement efficient reproducible dot products.

20 This document specifies functions corresponding to the IEC 60559 reduction and augmented arithmetic operations, for inclusion in the Mathematics library (`<math.h>`).

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

5 Part 4: Supplementary functions

1 Scope

This document specifies extensions to programming language C to include functions corresponding to operations specified and recommended in ISO/IEC 60559 and not supported in C23.

10 2 Conformance

An implementation may conform to one or both feature sets (reduction functions and augmented arithmetic) in this document. It so conforms if

a) it meets the requirements for a conforming implementation of C23;

and to conform to the reduction functions, the following are true:

15 b) it defines `__STDC_IEC_60559_BFP__` or `__STDC_IEC_60559_DFP__` or both, indicating support for IEC 60559 binary or decimal floating-point arithmetic, as specified in C23 Annex F;

c) it defines `__STDC_IEC_60559_FUNCS_REDUCTION__` to **20yymml** and provides the reduction functions as specified in this document ([6](#));

and to conform to the augmented arithmetic feature, the following are true:

20 d) it defines `__STDC_IEC_60559_BFP__`, indicating support for IEC 60559 binary floating-point arithmetic, as specified in C23 Annex F;

e) it defines `__STDC_IEC_60559_FUNCS_AUGMENTED_ARITHMETIC__` to **20yymml** and provides the features for augmented arithmetic specified in this document ([7](#) and [8](#)).

3 Normative references

25 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 9899:2024, *Information technology — Programming languages — C*

ISO/IEC 60559:2020, *Information technology — Microprocessor Systems — Floating-point arithmetic*

30 4 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC 9899:2024, ISO/IEC 60559:2020, and the following apply.

4.1

C23

standard ISO/IEC 9899:2024, *Information technology — Programming languages C*

5 C standard conformance

5 5.1 Freestanding implementations

The specification in C23 Clause 4 allows freestanding implementations to conform to this document.

5.2 Predefined macros

The implementation defines one or both of the following macros to indicate conformance to the specification in this document for support of the corresponding features specified and recommended in IEC 60559.

`__STDC_IEC_60559_FUNCS_REDUCTION__` The integer constant `20yymmL`.

`__STDC_IEC_60559_FUNCS_AUGMENTED_ARITHMETIC__` The integer constant `20yymmL`.

5.3 Standard headers

The identifiers specified in this document are defined or declared by the associated header if and only if the implementation defines the relevant feature macros (5.2) and

`__STDC_WANT_IEC_60559_FUNCS_EXT__`

is defined as a macro at the point in the source file where the header is first included.

This document does not specify type-generic macros in `<tgmath.h>` for its reduction functions.

6 Reduction functions in `<math.h>`

The header `<math.h>` declares the type and functions in this subclause.

The type declared is `size_t` (described in C23 7.21).

Each function in this subclause is declared in `<math.h>` if and only if the corresponding type is an IEC 60559 type supported according to C23 Annex F or an IEC 60559 arithmetic interchange or extended type supported according to C23 Annex H.

The functions in this subclause shall be implemented so that intermediate computations do not overflow or underflow. For the `reduc_sum`, `reduc_sumabs`, `reduc_sumsq`, and `reduc_sumprod` functions, the “overflow” or “underflow” floating-point exception is raised and a range error occurs if and only if the final result overflows or underflows. The `scaled_prod`, `scaled_prodsum`, and `scaled_proddiff` functions do not raise the “overflow” or “underflow” floating-point exceptions and do not cause a range error.

The reduction functions do not raise the “divide-by-zero” floating-point exception.

With IEC 60559 default exception handling, these functions raise the “inexact” floating-point exception in response to “overflow” and “underflow” exceptions; otherwise, whether they raise the “inexact” floating-point exception is unspecified.

5 Numerical results and exceptional behavior, including the “invalid” floating-point exception, might differ due to the precision of intermediates and the order of evaluation. However, only one floating-point exception is raised (other than “inexact” in response to “overflow” or “underflow”) per reduction function invocation; exceptions are not raised for each exceptional intermediate operand or result. Reduction functions may raise the “invalid” floating-point exception if a member of an array argument is a signaling NaN (see C23 F.2.1). Once an invalid floating-point exception is raised, due to signaling NaN, $\infty-\infty$, or $0\times\infty$, processing of array elements may stop.

Whether and how rounding direction modes affect functions in this subclause are implementation defined and may be indeterminate.

The preferred quantum exponent for the reduction functions for decimal floating types is unspecified.

15 For each of the following synopses, an implementation shall declare the functions suffixed with \mathbf{fN} or \mathbf{fNx} only if it defines `__STDC_WANT_IEC_60559_TYPES_EXT__` and supports the corresponding binary floating type. An implementation shall declare the functions suffixed with \mathbf{dN} or \mathbf{dNx} only if it defines `__STDC_WANT_IEC_60559_TYPES_EXT__` and supports the corresponding decimal floating type. (See C23 Annex H.)

6.1 The `reduc_sum` functions

20 Synopsis

```

#define __STDC_WANT_IEC_60559_FUNCS_EXT__
#include <math.h>

25 #ifdef __STDC_IEC_60559_FUNCS_REDUCTION__
#ifdef __STDC_IEC_60559_BFP__
double reduc_sum(size_t n, const double p[static n]);
float reduc_sumf(size_t n, const float p[static n]);
long double reduc_suml(size_t n,
    const long double p[static n]);
30 _FloatN reduc_sumfN(size_t n, const _FloatN p[static n]);
_FloatNx reduc_sumfNx(size_t n, const _FloatNx p[static n]);
#endif
#ifdef __STDC_IEC_60559_DFP__
_DecimalN reduc_sumdN(size_t n, const _DecimalN p[static n]);
35 _DecimalNx reduc_sumdNx(size_t n,
    const _DecimalNx p[static n]);
#endif
#endif

```

Description

40 The `reduc_sum` functions compute the sum of the n members of array \mathbf{p} : $\sum_{i=0}^{n-1} \mathbf{p}[i]$. If the length $n = 0$, the functions return the value $+0$. If any member of array \mathbf{p} is a NaN, the functions return a quiet NaN. If any two members of array \mathbf{p} are infinities with different signs, the functions return a quiet NaN and raise the “invalid” floating-point exception and a domain error occurs. Otherwise (if no member of \mathbf{p} is a NaN and no two members of \mathbf{p} are infinities with different signs), if any member of array \mathbf{p} is an infinity, the functions return that same infinity.

Returns

The `reduc_sum` functions return the computed sum.

6.2 The `reduc_sumabs` functions**Synopsis**

```

5  #define __STDC_WANT_IEC_60559_FUNCS_EXT__
   #include <math.h>

   #ifdef __STDC_IEC_60559_FUNCS_REDUCTION__
   #ifdef __STDC_IEC_60559_BFP__
10  double reduc_sumabs(size_t n, const double p[static n]);
   float reduc_sumabsf(size_t n, const float p[static n]);
   long double reduc_sumabsl(size_t n,
   const long double p[static n]);
   _FloatN reduc_sumabsfN(size_t n, const _FloatN p[static n]);
15  _FloatNx reduc_sumabsfNx(size_t n,
   const _FloatNx p[static n]);
   #endif
   #ifdef __STDC_IEC_60559_DFP__
20  _DecimalN reduc_sumabsdN(size_t n,
   const _DecimalN p[static n]);
   _DecimalNx reduc_sumabsdNx(size_t n,
   const _DecimalNx p[static n]);
   #endif
   #endif

```

25 Description

The `reduc_sumabs` functions compute the sum of the absolute values of the `n` members of array `p`: $\sum_{i=0}^{n-1} |p[i]|$. If the length `n` = 0, the functions return the value +0. If any member of array `p` is an infinity, the functions return $+\infty$; otherwise, if any member of array `p` is a quiet NaN, the functions return a quiet NaN.

30 Returns

The `reduc_sumabs` functions return the computed sum.

6.3 The `reduc_sumsq` functions

Synopsis

```

5  #define __STDC_WANT_IEC_60559_FUNCS_EXT__
   #include <math.h>

   #ifdef __STDC_IEC_60559_FUNCS_REDUCTION__
   #ifdef __STDC_IEC_60559_BFP__
   double reduc_sumsq(size_t n, const double p[static n]);
   float reduc_sumsqf(size_t n, const float p[static n]);
10  long double reduc_sumsql(size_t n,
   const long double p[static n]);
   _FloatN reduc_sumsqfN(size_t n, const _FloatN p[static n]);
   _FloatNx reduc_sumsqfNx(size_t n,
   const _FloatNx p[static n]);
15  #endif
   #ifdef __STDC_IEC_60559_DFP__
   _DecimalN reduc_sumsqdN(size_t n,
   const _DecimalN p[static n]);
   _DecimalNx reduc_sumsqdNx(size_t n,
20  const _DecimalNx p[static n]);
   #endif
   #endif

```

Description

25 The `reduc_sumsq` functions compute the sum of squares of the values of the `n` members of array `p`: $\sum_{i=0}^{n-1} (p[i] \times p[i])$. If the length `n` = 0, the functions return the value +0. If any member of array `p` is an infinity, the functions return $+\infty$; otherwise, if any member of array `p` is a quiet NaN, the functions return a quiet NaN.

Returns

The `reduc_sumsq` functions return the computed sum.

6.4 The `reduc_sumprod` functions

Synopsis

```

5  #define __STDC_WANT_IEC_60559_FUNCS_EXT__
   #include <math.h>

   #ifdef __STDC_IEC_60559_FUNCS_REDUCTION__
   #ifdef __STDC_IEC_60559_BFP__
   double reduc_sumprod(size_t n, const double p[static n],
10    const double q[static n]);
   float reduc_sumprodf(size_t n, const float p[static n],
    const float q[static n]);
   long double reduc_sumprodl(size_t n,
    const long double p[static n],
    const long double q[static n]);
15  _FloatN reduc_sumprodfN(size_t n, const _FloatN p[static n],
    const _FloatN q[static n]);
   _FloatNx reduc_sumprodfNx(size_t n,
    const _FloatNx p[static n],
    const _FloatNx q[static n]);
20  #endif
   #ifdef __STDC_IEC_60559_DFP__
   _DecimalN reduc_sumproddN(size_t n,
    const _DecimalN p[static n],
    const _DecimalN q[static n]);
25  _DecimalNx reduc_sumproddNx(size_t n,
    const _DecimalNx p[static n],
    const _DecimalNx q[static n]);
   #endif
   #endif

```

30 Description

The `reduc_sumprod` functions compute the dot product of the sequences of members of the arrays `p` and `q`: $\sum_{i=0}^{n-1} (p[i] \times q[i])$. If the length `n` = 0, the functions return the value +0. If any member of array `p` or `q` is a NaN, the functions return a quiet NaN. If a product is $0 \times \infty$, the functions return a quiet NaN and raise the “invalid” floating-point exception and a domain error occurs. If a sum is of infinities of different signs, the functions return a quiet NaN and raise the “invalid” floating-point exception and a domain error occurs. Otherwise (if no array member is a NaN, no product is $0 \times \infty$, and no sum is of infinities of different signs), if a term in the summation is an infinity, the functions return that same infinity.

Returns

40 The `reduc_sumprod` functions return the computed dot product.

6.5 The `scaled_prod` functions

Synopsis

```

5  #define __STDC_WANT_IEC_60559_FUNCS_EXT__
   #include <math.h>

   #ifdef __STDC_IEC_60559_FUNCS_REDUCTION__
   #ifdef __STDC_IEC_60559_BFP__
10  double scaled_prod(size_t n,
    const double p[static restrict n],
    long int * restrict sfptr);
   float scaled_prodf(size_t n,
15  const float p[static restrict n],
    long int * restrict sfptr);
   long double scaled_prodl(size_t n,
20  const long double p[static restrict n],
    long int * restrict sfptr);
   _FloatN scaled_prodfN(size_t n,
    const _FloatN p[static restrict n],
    long int * restrict sfptr);
   _FloatNx scaled_prodfNx(size_t n,
25  const _FloatNx p[static restrict n],
    long int * restrict sfptr);
   #endif
   #ifdef __STDC_IEC_60559_DFP__
   _DecimalN scaled_proddN(size_t n,
30  const _DecimalN p[static restrict n],
    long int * restrict sfptr);
   _DecimalNx scaled_proddNx(size_t n,
    const _DecimalNx p[static restrict n],
    long int * restrict sfptr);
   #endif
   #endif

```

Description

35 The `scaled_prod` functions compute a scaled product pr of the n members of the array \mathbf{p} and a scale factor sf , such that

$$pr \times b^{sf} = \prod_{i=0}^{n-1} \mathbf{p}[i]$$

where b is the radix of the type. These functions store the scale factor sf in the object pointed to by `sfptr`. If the length $n = 0$, the functions return the value +1 and store 0 in the object pointed to by `sfptr`. If any member of array \mathbf{p} is a NaN, the functions return a quiet NaN. If any two members of array \mathbf{p} are a zero and an infinity, the functions return a quiet NaN and raise the “invalid” floating-point exception and a domain error occurs. Otherwise, if any member of array \mathbf{p} is an infinity, the functions return an infinity. Otherwise, if any member of array \mathbf{p} is a zero, the functions return a zero. Otherwise, if the scale factor is outside the range of the `long int` type, the functions return a quiet NaN and raise the “invalid” floating-point exception. If a zero, infinity, or NaN is returned, the functions store 0 in the object pointed to by `sfptr`.

Returns

The `scaled_prod` functions return the computed scaled product pr .

6.6 The `scaled_prodsum` functions**Synopsis**

```

5  #define __STDC_WANT_IEC_60559_FUNCS_EXT__
   #include <math.h>

   #ifndef __STDC_IEC_60559_FUNCS_REDUCTION__
   #ifndef __STDC_IEC_60559_BFP__
10  double scaled_prodsum(size_t n,
      const double p[static restrict n],
      const double q[static restrict n],
      long int * restrict sfptr);
   float scaled_prodsumf(size_t n,
15  const float p[static restrict n],
      const float q[static restrict n],
      long int * restrict sfptr);
   long double scaled_prodsuml(size_t n,
20  const long double p[static restrict n],
      const long double q[static restrict n],
      long int * restrict sfptr);
   _FloatN scaled_prodsumfN(size_t n,
      const _FloatN p[static restrict n],
      const _FloatN q[static restrict n],
25  long int * restrict sfptr);
   _FloatNx scaled_prodsumfNx(size_t n,
      const _FloatNx p[static restrict n],
      const _FloatNx q[static restrict n],
      long int * restrict sfptr);
30  #endif
   #ifndef __STDC_IEC_60559_DFP__
   _DecimalN scaled_prodsumdN(size_t n,
      const _DecimalN p[static restrict n],
      const _DecimalN q[static restrict n],
35  long int * restrict sfptr);
   _DecimalNx scaled_prodsumdNx(size_t n,
      const _DecimalNx p[static restrict n],
      const _DecimalNx q[static restrict n],
40  long int * restrict sfptr);
   #endif
   #endif

```

Description

The `scaled_prodsum` functions compute a scaled product pr of the sums of the corresponding members of the arrays \mathbf{p} and \mathbf{q} and a scale factor sf , such that

$$pr \times b^{sf} = \prod_{i=0}^{n-1} (\mathbf{p}[i] + \mathbf{q}[i])$$

where b is the radix of the type. These functions store the scale factor sf in the object pointed to by `sfptr`. If the length $n = 0$, the functions return the value +1 and store 0 in the object pointed to by `sfptr`. If any member of array `p` or `q` is a NaN, the functions return a quiet NaN. If any sum is of infinities with different signs or if any two factors in the product are a zero and an infinity, the functions return a quiet NaN and raise the “invalid” floating-point exception and a domain error occurs. Otherwise, if any factor in the product is an infinity, the functions return an infinity. Otherwise, if any factor in the product is a zero, the functions return a zero. Otherwise, if the scale factor is outside the range of the `long int` type, the functions return a quiet NaN and raise the “invalid” floating-point exception. If a zero, infinity, or NaN is returned, the functions store 0 in the object pointed to by `sfptr`.

10 Returns

The `scaled_proddiff` functions return the computed scaled product pr .

6.7 The `scaled_proddiff` functions

Synopsis

```

15  #define __STDC_WANT_IEC_60559_FUNCS_EXT__
    #include <math.h>

    #ifdef __STDC_IEC_60559_FUNCS_REDUCTION__
    #ifdef __STDC_IEC_60559_BFP__
20  double scaled_proddiff(size_t n,
        const double p[static restrict n],
        const double q[static restrict n],
        long int * restrict sfptr);
    float scaled_proddiff(float p[static restrict n],
25  const float q[static restrict n],
        long int * restrict sfptr);
    long double scaled_proddiff(long double p[static restrict n],
30  const long double q[static restrict n],
        long int * restrict sfptr);
    _FloatN scaled_proddiff(_FloatN p[static restrict n],
        const _FloatN q[static restrict n],
        long int * restrict sfptr);
35  _FloatNx scaled_proddiff(_FloatNx p[static restrict n],
        const _FloatNx q[static restrict n],
        long int * restrict sfptr);
    #endif
40  #ifdef __STDC_IEC_60559_DFP__
    _DecimalN scaled_proddiff(_DecimalN p[static restrict n],
        const _DecimalN q[static restrict n],
        long int * restrict sfptr);

```

```

    _DecimalNx scaled_proddiffdNx(size_t n,
        const _DecimalNx p[static restrict n],
        const _DecimalNx q[static restrict n],
        long int * restrict sfptr);
5  #endif
    #endif

```

Description

The `scaled_proddiff` functions compute a scaled product pr of the differences of the corresponding members of the arrays \mathbf{p} and \mathbf{q} and a scale factor sf , such that

$$pr \times b^{sf} = \prod_{i=0}^{n-1} (\mathbf{p}[i] - \mathbf{q}[i])$$

where b is the radix of the type. These functions store the scale factor sf in the object pointed to by `sfptr`. If the length $n = 0$, the functions return the value +1 and store 0 in the object pointed to by `sfptr`. If any member of array \mathbf{p} or \mathbf{q} is a NaN, the functions return a quiet NaN. If any difference is of infinities with the same signs or if any two factors in the product are a zero and an infinity, the functions return a quiet NaN and raise the “invalid” floating-point exception and a domain error occurs. Otherwise, if any factor in the product is an infinity, the functions return an infinity. Otherwise, if any factor in the product is a zero, the functions return a zero. Otherwise, if the scale factor is outside the range of the `long int` type, the functions return a quiet NaN and raise the “invalid” floating-point exception. If a zero, infinity, or NaN is returned, the functions store 0 in the object pointed to by `sfptr`.

Returns

The `scaled_proddiff` functions return the computed scaled product pr .

EXAMPLE The scaled reduction functions support computing quantities of modest magnitudes whose intermediate results might well overflow and underflow. One example is the computation of Clebsch-Gordan coefficients or Wigner 3-j symbols for quantum physics. Expressions for these quantities involve quotients of products of factorials, and so are prone to intermediate overflow. As a simplified example, consider computing a fragment of the Clebsch-Gordan calculation.

```

#define __STDC_WANT_IEC_60559_FUNCS_EXT__
#include <math.h>

30 // compute quot = n1! * n2! / n3!

int n1 = 140, n2 = 160, n3 = 200; // factorial magnitudes
                                // 1e241, 1e284, 1e374
                                // quot magnitude 1e151

35 // products scaled to avoid overflow
double num1, num2, den;
// scale factors
long int num1e, num2e, dene;

40 // products scaled again to avoid intermediate overflow
// in final computation
double num1s, num2s, dens;
// scale factors
45 long int num1es, num2es, denes;

```

```

// result
double quot;

// arrays { 2, 3, 4, ... }
5 double num1p[n1-1], num2p[n2-1], denp[n3-1];

// n1! scaled to avoid overflow
for (int i = 2; i <= n1; i++) {
    num1p[i-2] = i;
10 }
num1 = scaled_prod(n1-1, num1p, &num1e);

// n2! scaled to avoid overflow
for (int i = 2; i <= n2; i++) {
    num2p[i-2] = i;
15 }
num2 = scaled_prod(n2-1, num2p, &num2e);

// n3! scaled to avoid overflow
for (int i = 2; i <= n3; i++) {
    denp[i-2] = i;
20 }
den = scaled_prod(n3-1, denp, &dene);

// re-scale to avoid subsequent intermediate overflow
num1es = llogb(num1);
num1s = scalbln(num1, -num1es);
num2es = llogb(num2);
num2s = scalbln(num2, -num2es);
30 denes = llogb(den);
dens = scalbln(den, -denes);

// compute result without intermediate overflow
quot = scalbln(num1s * num2s / dens,
35 num1e + num2e - dene + num1es + num2es - denes);

```

7 Augmented arithmetic functions in <math.h>

This subclause supports augmented arithmetic, as recommended by IEC 60559 for its binary formats. Each type and function in this subclause is declared in <math.h> if and only if the corresponding type is an IEC 60559 floating type supported according to C23 Annex F or an IEC 60559 arithmetic interchange or extended type supported according to C23 Annex H.

The types are structures for returning two floating-point values:

```

struct daug_t { double h; double t; };
struct faug_t { float h; float t; };
struct ldaug_t { long double h; long double t; };
45 struct _fNaug_t { _FloatN h; _FloatN t; };
struct _fNxaug_t { _FloatNx h; _FloatNx t; };

```

The *corresponding real type* of the structure refers to the type of the members.

The functions in this subclause round to nearest with ties toward zero, a rounding direction specified by IEC 60559 for use by augmented arithmetic operations. Thus, results are independent of dynamic

and constant rounding direction modes. As for other IEC 60559 operations, rounding is done with gradual underflow.

For each of the following synopses, an implementation shall declare the functions suffixed with \mathbf{fN} or \mathbf{fNx} only if it defines `__STDC_WANT_IEC_60559_TYPES_EXT__` and supports the corresponding binary floating type.

7.1 The `augadd` functions

Synopsis

```

#define __STDC_WANT_IEC_60559_FUNCS_EXT__
#include <math.h>

#ifdef __STDC_IEC_60559_FUNCS_AUGMENTED_ARITHMETIC__
#ifdef __STDC_IEC_60559_BFP__
struct daug_t augadd(double x, double y);
struct faug_t augaddf(float x, float y);
struct ldaug_t augaddl(long double x, long double y);
struct _fNaug_t augaddfN(_FloatN x, _FloatN y);
struct _fNxaug_t augaddfNx(_FloatNx x, _FloatNx y);
#endif
#endif

```

Description

The `augadd` functions compute two result values:

- h**: the sum $\mathbf{x} + \mathbf{y}$ rounded to the type using round-to-nearest with ties toward zero;
- t**: the error in **h** as a computation of $\mathbf{x} + \mathbf{y}$.

If **h** is a non-zero finite number, **t** has the value $\mathbf{x} + \mathbf{y} - \mathbf{h}$ (which is exactly representable in the type), where if **t** is zero its sign is the sign of **h**. If **h** is zero, **t** has the value of **h** (and both have the same sign). If **h** is infinite, **t** has the value of **h**. If **h** is a NaN, **t** is the same NaN.

These functions raise floating-point exceptions like the computation of **h**, except that under IEC 60559 default exception handling they raise the “inexact” floating-point exception only when the computation of **h** overflows.

A range error occurs when the computation of **h** overflows. A domain error occurs when the arguments are infinities with different signs.

Returns

These functions return the sum and error in a structure.

7.2 The augsub functions

Synopsis

```

5  #define __STDC_WANT_IEC_60559_FUNCS_EXT__
   #include <math.h>

   #ifdef __STDC_IEC_60559_FUNCS_AUGMENTED_ARITHMETIC__
   #ifdef __STDC_IEC_60559_BFP__
10  struct daug_t augsub(double x, double y);
   struct faug_t augsubf(float x, float y);
   struct ldaug_t augsubl(long double x, long double y);
   struct _fNaug_t augsubfN(_FloatN x, _FloatN y);
   struct _fNxaug_t augsubfNx(_FloatNx x, _FloatNx y);
   #endif
   #endif

```

15 Description

The **augsub** functions compute two result values:

- h**: the difference $x - y$ rounded to the type using round-to-nearest with ties toward zero;
- t**: the error in **h** as a computation of $x - y$.

20 If **h** is a non-zero finite number, **t** has the value $x - y - h$ (which is exactly representable in the type), where if **t** is zero its sign is the sign of **h**. If **h** is zero, **t** has the value of **h** (and both have the same sign). If **h** is infinite, **t** has the value of **h**. If **h** is a NaN, **t** is the same NaN.

These functions raise floating-point exceptions like the computation of **h**, except that under IEC 60559 default exception handling they raise the “inexact” floating-point exception only when the computation of **h** overflows.

25 A range error occurs when the computation of **h** overflows. A domain error occurs when the arguments are infinities with the same sign.

Returns

These functions return the difference and error in a structure.

7.3 The `augmul` functions

Synopsis

```

5  #define __STDC_WANT_IEC_60559_FUNCS_EXT__
   #include <math.h>

   #ifdef __STDC_IEC_60559_FUNCS_AUGMENTED_ARITHMETIC__
   #ifdef __STDC_IEC_60559_BFP__
10  struct daug_t augmul(double x, double y);
   struct faug_t augmulf(float x, float y);
   struct ldaug_t augmull(long double x, long double y);
   struct fNaug_t augmulfN(_FloatN x, _FloatN y);
   struct fNxaug_t augmulfNx(_FloatNx x, _FloatNx y);
   #endif
   #endif

```

Description

The `augmul` functions compute two result values:

- h**: the product $\mathbf{x} \times \mathbf{y}$ rounded to the type using round-to-nearest with ties toward zero;
- t**: the error in **h** as a computation of $\mathbf{x} \times \mathbf{y}$.

If **h** is a nonzero finite number and $\mathbf{x} \times \mathbf{y} - \mathbf{h}$ is exactly representable in the type, **t** has the value $\mathbf{x} \times \mathbf{y} - \mathbf{h}$, where if **t** is zero its sign is the sign of **h**. If **h** is a nonzero finite number and $\mathbf{x} \times \mathbf{y} - \mathbf{h}$ is not exactly representable in the type (because the magnitude of its value is too small), **t** has the value $\mathbf{x} \times \mathbf{y} - \mathbf{h}$ rounded to the type using round-to-nearest with ties toward zero. If **h** is zero, **t** has the value of **h** (and both have the same sign). If **h** is infinite, **t** has the value of **h**. If **h** is a NaN, **t** is the same NaN.

These functions raise the “overflow” and “invalid” floating-point exceptions like the computation of **h**. They raise the “underflow” floating-point exception like the computation of **t**. They raise the “inexact” floating-point exception when and only when the computation of **h** overflows or the computation of **t** is inexact.

A range error occurs when the computation of **h** overflows and may occur when the computation of **t** underflows. A domain error occurs when one argument is zero and the other an infinity.

Returns

These functions return the product and error in a structure.

EXAMPLE The augmented arithmetic operations are useful for extending precision, particularly for implementing “double-double” arithmetic, which provides a faster, though less precise and less predictable, alternative to IEC 60559 binary128 on systems which lack hardware support for binary128.

Double-double represents numbers as a pair of doubles, the second no larger in magnitude than the first, and usually much smaller. Ideally, in the pair (h, t) , h would equal the correctly rounded result of computed $(h+t)$, and t would equal the correctly rounded result of $h + t - \text{correctly-rounded-computed}(h+t)$. Performance considerations often compromise this ideal. There is no standard specification for double-double.

Let the pair (ah, at) represent a double-double value $a = ah + at$ in infinite precision, and (bh, bt) a similar double-double value b . Consider the double-double sum s represented by (sh, st) where $s = sh + st$ closely approximates $a + b = ah + at + bh + bt$. The code below uses augmented addition to compute such a sum.

```

5      #define __STDC_WANT_IEC_60559_FUNCS_EXT__
      #include <math.h>

      // components of double-double values a = 1/3, b = 2/3, and s
double ah = 0x0.AAAAAAAAAAAAA8p-1, at = 0x0.AAAAAAAAAAAAA8p-55;
10     double bh = 0x0.AAAAAAAAAAAAA8p0, bt = 0x0.AAAAAAAAAAAAA8p-54;
      double sh, st;

      struct daug_t u, v, w, y, z;

15     // compute components of s = a + b
      // exact sum is ah + at + bh + bt
      u = augadd(ah, bh); // exact sum is u.h + u.t + at + bt
      v = augadd(at, bt); // exact sum is u.h + u.t + v.h + v.t
      w = augadd(u.t, v.t); // exact sum is u.h + v.h + w.h + w.t
20     y = augadd(v.h, w.h); // exact sum is u.h + y.h + y.t + w.t
      z = augadd(u.h, y.h); // exact sum is z.h + z.t + y.t + w.t

      sh = z.h;
      st = z.t;
25

```

The code gives a good approximation to the ideal result, with absolute error $y.t + w.t$, and it is commutative and without conditional branches.

The steps for w and y could use regular addition (+) rather than `augadd`, because $w.t$ and $y.t$ are not used in the calculation. The code above gives a name to $w.t$ and $y.t$ for the didactic purpose of the error formula. It also assures a consistent result regardless of the evaluation method.

The special cases for the code might be no worse than those of any other double-double implementation fast enough to be useful. Its performance might be acceptable if `augadd` is no slower than two adds. There are faster ways to do double-double addition if hardware support for `augadd` is not available.

35 8 Augmented arithmetic macros in `<tgmath.h>`

For each family of augmented arithmetic functions in [Z](#), `<tgmath.h>` defines a macro with the name and number (two) of parameters of the corresponding unsuffixed function in the family. Thus, the macros are:

```
augadd   augsub   augmul
```

40 Both parameters are generic. The rules in C23 7.27 and H.13 apply to determine a corresponding real type based on the macro's argument types. The macro invokes the member of the function family whose `struct` return type has the determined corresponding real type.

EXAMPLE With the declarations

```

#include <tgmath.h>
int n;
float f;
5 double d;
  _Float32 f32;
  _Float64 f64;
complex float cf;
  _Decimal64 d64;

```

10 functions invoked by use of type-generic macros are shown in the following table:

<u>macro use</u>	<u>invocation</u>
augadd(f, f)	augaddf(f, f)
augadd(f, d)	augadd(f, d), the function
augsub(n, f)	augsub(n, f), the function
15 augmul(f32, f64)	augmulf64(f32, f64)
augadd(f32, d)	augadd(f32, d), the function
augsub(d, f64)	augsubf64(d, f64)
augmul(f, cf)	undefined
augadd(d, d64)	undefined
20 augsub(d64, d64)	undefined

Bibliography

- [1] IEC 559:1989, *Binary floating-point arithmetic for microprocessor systems, second edition*
- [2] IEEE 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*
- [3] IEEE 754-2008, *IEEE Standard for Floating-Point Arithmetic*
- 5 [4] IEEE 754-2019, *IEEE Standard for Floating-Point Arithmetic*
- [5] IEEE 854-1987, *IEEE Standard for Radix-Independent Floating-Point Arithmetic*
- [6] ISO/IEC 9899:2011/Cor.1:2012, *Information technology — Programming languages — C / Technical Corrigendum 1*
- [7] ISO/IEC 9899:2018, *Information technology — Programming languages — C*
- 10 [8] ISO/IEC 9899:2024, *Information technology — Programming languages — C*