

WG14 N2146
Meeting notes

C Floating Point Study Group Teleconference

2017-03-21
9 AM PDT / 12 PM EDT

Attendees: Rajan, Jim, Fred, Ian, David H., Mike, David C.

New agenda items:
None.

Last meeting action items:

Rajan: Submit DRS2:DDR9 example as a paper for a supplement to the STC in the DR and change all the %A to %a (ensuring it conforms to the specification for %a). - Done (N2126)

Jim: Create a DR against part 2 to specify the capitalization effects of %a. - Done (N2125)

Jim: DRs2:DDR2: DECIMAL_DIG: Have DECIMAL_DIG to be the largest supported floating type as a new replacement Suggested TC for DRs2:DDR2. - Not done due to further discussion. Remove from list.

Jim: Add in the encoding *_DIG macros (as per http://wiki.edg.com/pub/CFP/WebHome/characteristics_for_non-arith_formats.pdf) and obsolescing DECIMAL_DIG as a proposal to WG14 as a new enhancement. - Not done due to further discussion. Remove from list.

Jim: Include Type-generic macros for functions that round result to narrower type as part of the third set of DR's. - Done (N2125).

Fred: Propose adding in #pragma STDC FENV_ROUND DEFAULT as a means to set the static rounding mode to the default rounding mode. - Done (N2128) via a different method.

New action items:

Rajan: Bring up "Understanding" points 4 and 5 when discussing DR501 (from Jim's email on 2017/03/17).

David H: Ask about the IEEE dependencies (rounding mode, infinities) for the augmented functions and how they could be used for non-IEEE formats.

Jim: Ask David Keaton what we should be doing to close off the group or extend it (IEEE-754:2018 binding for example).

Next Meeting:

Tuesday April 25th, 2017, 12:00 EDT, 9:00 PDT
Same teleconference number.

Discussion:

IEEE 754 revision:

Near the end.

Couple of major items without text: Max/min, augmented.

What to do the old min/max? Remove? Deprecate? Currently deprecated is not a concept. Instead we say you can do x and y, but you should do y.

Currently everything added is a recommendation. Next revision may consider making changes that are required and may cause changes.

C++ liaison:
No update.

Proposals for the C standard (C2X):
Links in email sent by Jim today.

DRs:

N2127 (<http://wiki.edg.com/pub/CFP/WebHome/n2127.pdf>):

Proposed change to Part 3 for DR 501.

This paper changes part 2 in a simpler way which means part 3 does not need to change anymore.

N2126

(http://wiki.edg.com/pub/CFP/WebHome/N2126_Example_of_the_effect_of_the_change_for_CF_P_DR11.htm):

No capital A to avoid the %A specification issue (which is handled via another paper).

Discuss DECIMAL_DIG issues (Jim's 3/17/17 email):

"Understanding":

Point 2:

Ian: Regarding extension floating-point types vs extended integer types: We should see about changing the C standard to support this.

Jim: This is essentially what we're doing in part 3.

Point 3:

Common extensions are not necessarily conforming.

Point 4:

Since it doesn't say long double, this follows.

"Alternatives":

Point 2: Seems to be more acceptable to be a C2X proposal since WG14 did not seem to like making changes to C due to a TS.

Point 4: Will have to be a proposal for a new version of part 3 (or part of C if the Part 3 proposal for C2X gets in).

Alternative 2 seems to be the consensus.

Can bring up "Understanding" points 4 and 5 when discussing DR501.

Binding for IEEE 754-2018 (Jim's email on 2017/03/20):

Augmented operations: Return two results, head and tail.

WG14 said a struct is better to return two results.

Non-IEEE format needs to be looked at. Rounding mode, infinities for example do not always exist.

David H: Perhaps say the two parts need to add up to the correct result to leave slack in terms of rounding.

Jim: There are algorithms for reproducible summations (the purpose of these functions). Those algorithms likely need this rounding property to work correctly. It does not need to be a rounding mode, just the result has to follow that property.

David H: The property was an issue of efficiency. May not be possible for non-IEEE due to implicit dependencies.

Ian: For AIX double-double (vs Linux double-double), having the high part being close to infinity may have issues as rounding away from zero may result in an infinity.

Ian: No rounding mode that works in general for IEEE.

David H: It does for Decimal.

It was not intended that they add another mode, just that the semantics have this.

Ian: Most implementations do not have the rounding mode needed.

David H: Expected that we can implement this in hardware.

Ian: If we are expecting hardware to add this rounding for these operations, would it work

for a general rounding mode?

David H: No, since most hardware only has 2 bits for rounding mode and adding this would result in another bit.

Non-IEEE implementations: double-double to get quad-quad, short/half (hard to imagine), IBM Hex format (only has round to zero), packed decimal?

Jim: Annex F allows double-double for 'long double'.

*David H: Ask about the IEEE dependencies (rounding mode, infinities) for the augmented functions and how they could be used for non-IEEE formats.

Min/Max:

NaN preference or number preference.

Can do fminISO, fminIEC, fminSTD or something like that.

Mike: fminimum can be used. No reason to abbreviate.

Fred: In Annex F we don't say anything about zeros or SNaNs.

Jim: In the TS we do say SNaNs always signal. The signed zero will likely break a lot of implementations of the existing fmin (which is more like minimumNumber).

Rajan: There are long names already for the atomic functions. Can just do fminimum_magnitude_number.

We can discuss this next meeting after thinking about it.

Other:

Marius (pre-meeting question): When will things be finalized?

Fred: After C2X?

Jim: We need to stay with it until it is in the standard. This will take a while.

If parts go into the standard but not all of them, we need to revise the TS parts that didn't go in (Ex. rebase on C2X).

Rajan: Current WG14 mandate is for IEEE-754:2008 binding. The 2018 binding may be something we add on or a new group created.

*Jim: Ask David Keaton what we should be doing to close off the group or extend it (IEEE-754:2018 binding for example).

Jim: WG14 will likely be happy to have us help integrate any parts into C2X.
