IEEE 754 Revision

Paul Zimmermann

Dagstuhl Seminar 06021, 8-13 January, 2006

History

1985: IEEE 754 adopted by IEEE/ANSI

2001: revision started

regular meetings every month (SF bay area)

attended the Dec 2005 meeting

Committee

Chair is Dan Zuras

Scientists (Prof. Kahan)

Machine vendors (IBM, HP, Intel, AMD, Sun, Apple)

Applications: Mori (SAP), Oberman (NVIDIA)

Other interests: Godard (Out-of-the-box computing)

Main changes

Reference: October 20, 2005 draft

decimal formats (see below)

transcendentals (section T)

varying-width formats (section 3.5)

base conversion: should be correct in the whole range (section 5.6)

tail operations (section 5.13)

fused multiply add

new round-to-nearest-away mode

Formats

binary32, binary64

new: binary16 (storage), binary128 (quad), decimal32 (storage), decimal64, decimal128

no “single extended” and “double extended” formats any more

binary16: 11 bits, $-14 \leqslant e \leqslant +15$

binary128: 113 bits, $-16382 \leqslant e \leqslant +16383$

decimal32: 7 digits, $-95 \leqslant e \leqslant +96$

decimal64: 16 digits, $-383 \leqslant e \leqslant +384$

decimal128: 34 digits, $-6143 \leqslant e \leqslant +6144$

Which hardware decimal format (IBM or Intel)?

Transcendentals

- correctly-rounded versions (0.5 and 0/1 ulp)
- **and** well-rounded versions if more efficient (twice the error bound: 1.0 and 0/2.0 ulps)
- for exp, expm1, log, log1p, $\sin(\pi x)$, $\cos(\pi x)$, atan$(x)/\pi$ on their whole input domain
- $\sin$, $\cos$ on $(-2\pi, 2\pi)$, atan on $[-\tan(\pi/2), \tan(\pi/2)]$, well-rounded outside

Give your opinion before it is too late

All documents are public.

web site: `http://grouper.ieee.org/groups/754/`

mailing list: `stds-754@ieee.org`

drafts: `http://www.validlab.com/754R/`

Remote participation possible.