

IEEE Task P854

Minutes, 3 December 1982

The radix-free floating-point working group of the Microprocessor Standards subcommittee of the IEEE Computer Society met from 10:00 a.m. to 3:45 p.m. at Apple Computer, Cupertino. Seven people were in attendance.

Handwritten minutes from the 20 October meeting were discussed. Corrections were made and will be reflected in the typed version to be distributed through a regular mailing. The minutes were agreed to as corrected.

The next meeting of P854 will be on Thursday, January 20, 1983 at ELXSI or at Berkeley (TBD).

Ris proposed to organize a mini-symposium at the SIAM National meeting in Denver the week of June 6, 1983 on the subject of P754 and P854 status (technical and procedural) and issues. This should happen in conjunction with a P854 meeting. He was given leave to pursue this. It was agreed that Draft 1.0 could be and should be ("shall be"?) in hand prior to June 6.

Karpinski reported that the question of upward compatibility with P754 had been raised with the parent Microprocessor Standards subcommittee. Specifically, could P854 be more rigid than P754 by specification in areas where P754 had punted? (E.g., representation of NaNs on conversion to external strings.) The answer was that strict conformability is not required: "It's an engineering problem; find an engineering solution."

Tracking P754. Section 3.1 was substantially reworded to reflect changes to appear in draft 10.0 of P754. Specifically, P854 does not completely ignore issues of encodings. It makes no specification but does deal with redundant encodings (if any) and encodings intended to go beyond the scope of the standard.

It is agreed that Cody will apply all other changes made in P754 draft 10.0 which are not reflected in P854 draft 0.5e.

Projective Mode. Discussion opened with Kahan's analysis that projective mode was not worth having unless it be the default. After some further discussion, Coonen observed that in five years of P754 activity, he had yet to make a motion.

JEROME T. COONEN moved to strike projective mode from the standard. No further discussion. Passed 5-1-1.

Rounding Modes. Informal discussion about the merits of (1) not requiring directed roundings at all, and (2) not requiring a mode-based implementation of roundings. No conclusions; no action.

Relationships among Formats. Ris pointed out that intent in going from basic to extended and in going from single to double is not to adjust exponent by a factor of eight, but rather to permit eighth powers without exception. Accordingly, bounds for maximum exponent should be of the form $8e+7$ rather than $8e$. Agreed.

Kahan observed that the retained sentence relating exponent range of single extended versus double does not hang together. Agreed that essential thought is that double is an acceptable implementation of single extended.

External Interfaces to the Standard. Kahan proposed to address issues such as names of flags, traps, etc. to attempt a greater degree of inter-system transportability. Agreed. He further proposed that the com-

mittee begin its consideration of candidate protocols by examining existing or soon-to-be implementations to ensure a more likely basis in reality. Agreed. Proposals for standardization based on implementations existing or underway will be entertained starting with the 20 January meeting.

The external representation of infinities will be either an optional sign followed by "Infinity" (letters case insensitive) or an optional sign followed by "1/0". The output form is chosen by the implementation; both forms must be recognized on input. Consideration of whether to accept shortened versions of "Infinity" down to but not including "Inf" (because of existing BASIC usage) is deferred. Kahan produced a technical note on dealing with field overflow on output conversion.

"Non-Arithmetic" Operations. Those items enumerated in the appendix which the implementation may exempt from invalid operation exception in the case of signalling NaN operands should likewise be exempted from trapped underflows in the case of denormalized operands.

Inexact Exception on Input Conversions. In Section 5.6 the words "alter all" are replaced with "round off". This will better signal what is intended.

Under what conditions is "inexact" signalled on input? In decimal the signal must be accurate, unconditionally. In binary, the signal is allowed to be raised in conditions where input digits are allowed to be discarded, even when in principle the conversion could have been done exactly. Examples:

<u>Regime</u>	<u>Input String</u>	<u>Converted Value</u>	<u>Inexact Allowed</u>	<u>NOT Inexact Allowed</u>	
P754	1,073,741,824			2**30	Yes
P754	1,073,741,825			2**30	Yes
D7	1,000,000,000			10**9	
D7	1,000,000,001			10**9	

Elimination of "Denormalized". Section 2.7 will define "subnormal" in the manner in which "denormal" had been defined, and define "normal" as any non-zero number which is finite and not subnormal. The word "denormalized" will disappear except for a footnote to leave a clue for historians. Corollary changes will percolate.

Fred Ris

/lar