BESCHWERDEKAMMERN   BOARDS OF APPEAL OF   CHAMBRES DE RECOURS
DES EUROPÄISCHEN   THE EUROPEAN PATENT   DE L'OFFICE EUROPEEN
PATENTAMTS   OFFICE   DES BREVETS

Internal distribution code:
(A) [ ] Publication in OJ
(B) [ ] To Chairmen and Members
(C) [X] To Chairmen
(D) [ ] No distribution

D E C I S I O N
of 18 May 2006

Case Number: T 0074/04 - 3.5.01

Application Number: 98117383.4

Publication Number: 0903664

IPC: G06F 7/48

Language of the proceedings: EN

Title of invention:
Floating-point processor with operand-format precision greater
than execution precision

Applicant:
Koninklijke Philips Electronics N.V.

Opponent:
-

Headword:
Floating-point processor/PHILIPS

Relevant legal provisions:
EPC Art. 56

Keyword:
"Inventive step (no)"

Decisions cited:
-

Catchword:
-

Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

Beschwerdekammern   Boards of Appeal   Chambres de recours

Case Number: T 0074/04 - 3.5.01

D E C I S I O N
of the Technical Board of Appeal 3.5.01
of 18 May 2006

Appellant: Koninklijke Philips Electronics N.V.
                             Groenewoudseweg 1
                             NL-5621 BA Eindhoven   (NL)

Representative: van der Veer, Johannis Leendert
                             Philips
                             Intellectual Property & Standards
                             P.O. Box 220
                             NL-5600 AE Eindhoven   (NL)

Decision under appeal: **Decision of the Examining Division of the**
                             European Patent Office posted 2 October 2003
                             refusing European application No. 98117383.4
                             pursuant to Article 97(1) EPC.

Composition of the Board:

Chairman: S. Steinbrener
Members: R. Wibergh
                  A. Pignatelli

Summary of Facts and Submissions

I.   This appeal is against the decision of the examining
        division to refuse European patent application
        No. 98117383.4.

II.   The following documents will be referred to in the
        present decision:

        D1: WO-A-95/31767
        D2: M. Darley et al., "The TMS390C602A Floating-Point
        Coprocessor for Sparc Systems", IEEE MICRO June 1990,
        p.36-47.

III.   According to the decision appealed, taken on the file
        as it was standing and with reference to the previous
        official communications, the invention according to all
        claims was obvious in view of a combination of D1,
        regarded as the closest prior art document, with D2.
        The only difference between the invention and D1 was
        that software trap means were used for performing
        floating-point operations at the highest precision.
        Such traps were however conventional, as evidenced by
        D2, and thus a matter of normal design procedure.

IV.   The decision was based on claims 1-8 as amended in
        accordance with the appellant's letter dated
        24 February 2003. Independent claim 5 read:

        " A floating-point method for performing a requested
        floating-point operation at a requested result
        precision on an operand set of operands to achieve a
        result, said requested result precision being selected

from a set of plural precisions including a highest
precision, said method comprising:

storing (S1) said operands in a highest precision
format in respective registers of a floating-point
processor;

determining (S2) an execution precision as a function
of said requested precision and the maximum apparent
precision of said operand set, said maximum apparent
precision being a function of the number of trailing
zeroes in the operand having said maximum apparent
precision;

if said execution precision is less than said highest
precision, performing (S3) said requested floating-
point operation using a hardware execution unit of said
floating-point processor; and

characterized by the steps of performing nominal
execution at a precision lower than the format
precision with which it stores operands;

if said execution precision is said highest precision,
executing (S4) a trap and performing said requested
floating-point operation in software."

Claim 1 was directed to a corresponding floating-point
processor.

V.   In the grounds of appeal the appellant argued
        essentially as follows:

        D1 disclosed a floating-point processor that performed
        operations at three execution precisions, referred to
        as single, double and extended precision. An input
        format converter converted single or double precision
        source values to extended precision format but left
        extended precision format source values in extended

1092.D

format. The resulting operands stored in respective
registers of the processor were in extended precision
format. The processor furthermore had an execution unit
that performed the operations at the precision
determined by a mode selector. Independent of whether
the mode selector selected single, double or extended
precision, it was always the execution unit of the
floating-point processor that executed the operation
(p.8, l.15—22). Therefore Dl taught that the source
values were stored in the registers of the execution
unit at a precision equal to the highest execution
precision at which the corresponding execution unit
could perform an operation. In most applications the
greater precisions were used more rarely than the
lesser precisions. Extended precision, in particular,
was encountered very infrequently in many computational
environments, and it was therefore a disadvantage of
the prior art floating-point processor according to Dl
that it comprised a relatively complex execution unit,
capable of executing extended precision operations.

D2 taught a person skilled in the art, faced with the
problem of the prior art processor, that extended
precision operations should be handled by a software
trap instead of a dedicated processor. It also taught
that the registers used by the dedicated processor
should store source values in the double-precision
format, ie the same precision as the highest execution
precision at which the corresponding execution unit can
perform an operation.

Hence, the person skilled in the art combining the
teaching of Dl and D2 would arrive at a device
comprising:

format conversion means for converting operands of an
operand not in the double-precision format to the
double-precision format;
register means for storing said operands of the operand
set in the double-precision format;
an execution unit for performing an operation on the
operands to yield an execution result, in case the
required execution precision was not the extended
precision;
trap means for executing a trap when the required
execution precision was the extended precision.

The floating-point processor so obtained had the
disadvantage that extended precision floating point
operations were executed considerably slower, because
it had no registers for storing extended precision
values, and it had to load and store these values from
and to the memory.

The floating-point processor according to claim 1 of
the present application differed from what the person
skilled in the art would obtain from documents Dl and
D2 in that the format conversion means converted
operands not in the extended-precision format to the
extended-precision format. Furthermore, the register
means stored the operands in the extended-precision
format. On the one hand the invention used trap means
for performing extended precision calculations which
allowed a simplification of the execution unit and
therefore reduced the cost of the implementation. On
the other hand, as it could use extended precision
registers there was no need to trap on high, eg
extended, precision "load" and "store" instructions.

This reduced latency and obviated the need for virtual registers in memory for storage of high precision operands and results. There was also a saving in complexity in that tag bits were not required to indicate the precision of stored operands (cf the present application, paragraph [0012]). Consequently, the floating-point processor of claim 1 was faster than one that would result from a straightforward combination of Dl and D2.

Hence, the combination of measures as claimed in claim 1 entailed advantageous technical features while this combination could not be derived in an obvious way from the documents Dl and D2.

VI.   By communication dated 18 January 2006, annexed to a summons to oral proceedings, the Board stated that the examining division's argumentation appeared convincing in spite of the appellant's counter-arguments and indicated the reasons for its view.

VII.   By letter dated 10 April 2006, the appellant withdrew its request for oral proceedings and requested a decision in accordance with the state of the file.

VIII.   Oral proceedings, at which the appellant was not represented, were held on 18 May 2006. It was verified that the appellant requested that the decision under appeal be set aside and a patent be granted on the application documents on file (cf the notice of appeal and the grounds, p.1). After deliberation the Board announced its decision.

Reasons for the Decision

1.   The appellant argues that a combination of D1 and D2
        would yield format conversion means converting operands
        to double-precision format, whereas according to the
        invention - and D1 - conversion is made to extended
        format. This argument is unusual in that a feature
        already contained in the closest prior art document,
        here D1, is said to involve a inventive step, but the
        point appears to be that the skilled person would have
        omitted this feature as a consequence of the
        combination with D2.

2.   However, the appellant's reasoning is not convincing.
        The main advantage of the circuit in D1 is that latency
        can be reduced, for example when an operand is
        specified in extended precision (cf D1 p.6, l.9-15; p.9,
        l.28-30). To this aim, the operand's "apparent
        precision" is detected (cf p.5, l.9-11; paragraph
        bridging pages 7 and 8). This requires the operand to
        be stored at maximum precision, as explained in D1. It
        is not likely that the skilled person would have left
        out any of these crucial features when incorporating
        the teaching of D2 into D1, nor is it clear how the
        operands with extended apparent precision might
        otherwise be detected. D2 does not seem to contain any
        suggestion how to replace the detection of apparent
        precision. As to the appellant's argument that the
        prior art teaches that the format precision for storing
        operands should be equal to the nominal execution
        precision of the processor, ie double precision in the
        case of the invention, it is noted that D1 explicitly
        states that the mode selector determines the maximum
        precision between the requested result precision and

the maximum apparent operand precision (p.5, l.13-15),
which is not necessarily equal to the nominal execution
precision of the processor. This quotation appears
incompatible with the appellant's suggestion that the
skilled person might opt for conversion to double-
precision format.

3.   For these reasons the subject-matter of claims 1 and 5
        does not involve an inventive step (Article 56 EPC).

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar: The Chairman:

P. Guidi   S. Steinbrener