BESCHWERDEKAMMERN DES EUROPÄISCHEN PATENTAMTS

BOARDS OF APPEAL OF THE EUROPEAN PATENT OFFICE

CHAMBRES DE RECOURS DE L'OFFICE EUROPEEN DES BREVETS

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

Datasheet for the decision of 15 March 2011

Case Number: T 0374/07 - 3.5.06

Application Number: 97304617.0

Publication Number: 0817072

IPC: G06F 12/08

Language of the proceedings: EN

Title of invention:
A multiprocessing system configured to store coherency state within multiple subnodes of a processing node

Applicant:
Oracle America, Inc.

Headword:
Multiprocessing system/ORACLE

Relevant legal provisions:
RPBA Art. 15

Relevant legal provisions (EPC 1973):
EPC Art. 56

Keyword:
"Inventive step - no"

Catchword:

Europäisches Patentamt | European Patent Office | Office européen des brevets

Beschwerdekammern | Boards of Appeal | Chambres de recours

Case Number: T 0374/07 - 3.5.06

D E C I S I O N
of the Technical Board of Appeal 3.5.06
of 15 March 2011

Appellant: Oracle America, Inc.
500 Oracle Parkway
Redwood City, CA 94065 (US)

Representative: Harris, Ian Richard
D Young & Co LLP
120 Holborn
London EC1N 2DY (GB)

Decision under appeal: **Decision of the Examining Division of the European Patent Office posted 3 August 2006 refusing European patent application No. 97304617.0 pursuant to Article 97(1) EPC.**

Composition of the Board:

Chairman: D. H. Rees
Members: G. Zucka
M-B. Tardo-Dino

Summary of Facts and Submissions

I. The appeal is against the decision by the examining division dispatched on 3 August 2006 to refuse European patent application 97304617.0 on the basis that the subject-matter of claims 1, 2, 9, 10, 21 and 22 (received on 18 April 2006) is not novel, Article 54 EPC 1973, in view of the following document:

D1: Lenoski D et al.: "The Stanford DASH multiprocessor", COMPUTER, IEEE Computer Society, Long Beach, CA, US, vol. 25, no. 3, 1 March 1992, pages 63-79.

II. A notice of appeal was received on 25 September 2006, the appeal fee being paid on 26 September 2006. A statement of the grounds of the appeal was received on 01 December 2006.

III. The appellant requested oral proceedings in the event that the board was minded to refuse his request that "the [appealed] decision is overturned and the patent application allowed in the form as subject to the decision under appeal (*i.e.* as currently on file)".

IV. The board issued a summons to oral proceedings. In an annex to the summons, the board set out its preliminary opinion on the appeal. In particular, it pointed out that, from the applicant's arguments, it was apparent that a key feature of the invention is the provision of a modular structure, which minimises the necessary directory space and makes it easy to increase or decrease the amount of memory in a given node by changing the number of sub-nodes. It would be essential,

within this modular structure, to make a distinction between different kinds of sub-nodes. In particular, the modularity depends on the existence of, firstly, a "controller sub-node" and, secondly, "snooper sub-nodes". It is the addition or deletion of the latter kind of sub-nodes that would allow an easy change in the amount of shared memory in a processing node. This is, however, not reflected in the independent claims. The board was of the preliminary opinion that the subject-matter of the claims, with their present wording, was not inventive.

V. The appellant announced in a letter dated 3 February 2011 that he would not attend the oral proceedings. The oral proceedings were held on 15 March 2011, in the absence of the appellant.

VI. The board understood the appellant's substantive request to be as follows: that the decision under appeal be set aside and that a patent be granted on the basis of the documents already submitted to the examining division, *viz.* claims 1-22 as filed on 18 April 2006; description pages 1-5, 7-72 as originally filed, and pages 6, 6a as filed on 25 May 2005 (the minutes of the oral proceedings are incorrect on this point); drawings, sheets 1/14-14/14 filed with letter of 10 September 1997.

VII. Claim 1 reads as follows:

"A method of operating a computer system (10) comprising a plurality of processing nodes (12A, 12B) linked by a network (14), wherein said computer system

has a distributed shared memory architecture, said method comprising:

providing a memory (22) and multiple subnodes within a processing node, wherein said memory is logically divided into memory portions and access rights for a particular memory portion are stored in a corresponding subnode;

communicating between a local bus (20) of the processing node and the network using a first subnode (300);

storing a first plurality of coherency states corresponding to a first plurality of coherency units stored within a first memory portion within said processing node in a second subnode (302) coupled to the local bus; and

storing a second plurality of coherency states corresponding to a second plurality of coherency units stored within a second memory portion within said processing node in the first subnode;

wherein the coherency state maintained for a corresponding coherency unit indicates the access rights to that coherency unit, and wherein a coherency unit is a number of contiguous bytes of memory which are treated as a unit for coherency purposes".

Claim 9 relates to a "system interface" and contains apparatus features corresponding to the method features of claim 1.

VIII. At the end of the oral proceedings, the board announced its decision.

Reasons for the decision

1. *The admissibility of the appeal*

In view of the facts set out at points I and II above, the appeal is admissible, since it complies with the EPC formal admissibility requirements.

2. *The appellant's non-attendance at the oral proceedings*

2.1 As announced in advance, the duly summoned appellant did not attend the oral proceedings.

2.2 In accordance with Article 15(3) RPBA, the board relied for its decision only on the appellant's written submissions. The board was in a position to decide at the conclusion of the oral proceedings, since the case was ready for decision (Article 15(6) RPBA), and the voluntary absence of the appellant was not a reason for delaying a decision (Article 15(3) RPBA).

3. *Novelty, Article 54(1, 2) EPC 1973*

3.1 In terms of claim 1, D1 discloses (see figures 5-7 and the corresponding text passages) a method of operating a computer system comprising a plurality of processing nodes linked by a network, wherein said computer system has a distributed shared memory architecture, said method comprising:

providing a memory within a processing node, wherein the access rights for a particular memory portion are stored in a sub-node;

communicating between a local bus of the processing node and the network using a sub-node.

Moreover, there are coherence units, being a number of contiguous bytes of memory which are treated as a unit for coherence purposes, and there is a coherence state maintained for a corresponding coherence unit indicating the access rights to that coherence unit. This is the function of the so-called "directory memory".

These common features of D1 and the claimed invention were acknowledged by the appellant in the statement of grounds of the appeal, Section 4.

3.2 The subject-matter of claim 1 differs from D1 essentially by the following features:

storing a first plurality of coherency states corresponding to a first plurality of coherency units stored within a first memory portion within said processing node in a second sub-node coupled to the local bus; and

storing a second plurality of coherency states corresponding to a second plurality of coherency units stored within a second memory portion within said processing node in the first sub-node.

3.3 The subject-matter of claim 9 differs from D1 by the apparatus features corresponding to the above method features.

3.4 Conclusion on novelty

It follows from the above analysis that the subject-matter of claims 1 and 9 is novel.

4. *Inventive step, Article 56 EPC 1973*

4.1 The *alleged* problem which the application intends to solve by the features that distinguish the claimed subject-matter from D1 is how to deal with the drawbacks of the type of memory (typically SRAM modules) that are necessary to store access rights. These drawbacks include a smaller bit density and a higher price for such memory. A system of the type commonly described in D1 and the present application using the obvious directly-mapped directory memory may be prohibitively expensive because it needs a great deal of this expensive memory to accommodate access rights for the entire address space.

4.2 In the application, the problem is ameliorated without abandoning the directly-mapped nature of the directory by adopting a scheme which allows the directory memory to be limited to the size appropriate to the actual amount of main memory installed in the system, thus in turn allowing the size of the directory memory to be increased as the size of the main memory is increased. As the system is expanded, modules containing both main and directory memory are added (see description page 14, lines 1 to 15, page 17, lines 20 to 26, and figures 1A and 1B).

4.3 D1 also addresses the above problem (see D1 page 67, column 1 line 56 to column 2 line 21) but deals with it in a different manner, *viz.* by using the directory memory more efficiently, by abandoning the one-to-one direct mapping between bits in the directory memory and areas of main memory. It thereby achieves linear or near-linear performance growth as the number of

processors increases from a few to a few thousand, by distributing the memory among processing nodes and using a network with scalable bandwidth to connect the nodes. It would be clear to the skilled person that abandoning the direct mapping makes the system more complicated.

The *objective* problem which is solved by the present application is, therefore, to find an alternative to the solution provided by D1 without abandoning the one-to-one direct mapping between bits in the directory memory and areas of main memory.

4.4 The skilled person knows that, if the total memory increases, the size of the directory necessary to contain the access rights for the different memory portions will also increase. In computer systems, modularisation is a very common measure to deal with such a situation. And therefore, the skilled person who wants to tackle the above objective problem, will naturally consider a modular structure.

From the description (see page 55 *sqq.*), it is apparent that this is not the complete solution to the problem and, for the invention actually to work, a distinction has to be made between different kinds of sub-nodes, in particular a "controller sub-node" and "snooper sub-nodes". It is the addition or deletion of the latter kind of sub-nodes that allows an easy change in the amount of shared memory in a processing node.

However, this is not reflected in the independent claims 1 and 9. Both claims merely include the feature that the sub-nodes store coherence states and specify

that two functions are combined in the first sub-node, without any apparent reason. The different roles of the different kinds of sub-nodes, as specified in the description, are not reflected in the claims.

What remains in the claims as they are worded is nothing more than a standard modular structure, which, as set out above, would be a common measure for the skilled man and can, therefore, not be considered inventive.

5. *Conclusion*

The subject-matter of claims 1 and 9 does not involve an inventive step, Article 56 EPC 1973. For this reason, the applicant's request is not allowable.

Order

For this reason, it is decided that:

The appeal is dismissed.

The Registrar: The Chairman:

B. Atienza Vivancos D. H. Rees