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

Datasheet for the decision of 3 December 2014

Case Number: T 0278/10 - 3.5.01

Application Number: 02789637.2

Publication Number: 1446721

IPC: G06F12/00

Language of the proceedings: EN

Title of invention:
MEMORY ADAPTED TO PROVIDE DEDICATED AND OR SHARED MEMORY TO MULTIPLE PROCESSORS AND METHOD THEREFOR

Applicant:
Intel Corporation

Headword:
Shared memory/INTEL

Relevant legal provisions:
EPC 1973 Art. 56

Keyword:
Inventive step - (no)

Decisions cited:
T 0003/90

Catchword:

Europäisches Patentamt
European Patent Office
Office européen des brevets

Beschwerdekammern
Boards of Appeal
Chambres de recours

European Patent Office
D-80298 MUNICH
GERMANY
Tel. +49 (0) 89 2399-0
Fax +49 (0) 89 2399-4465

Case Number: T 0278/10 - 3.5.01

D E C I S I O N
of Technical Board of Appeal 3.5.01
of 3 December 2014

Appellant: (Applicant)	Intel Corporation 2200 Mission College Boulevard Santa Clara, CA 95054 (US)
Representative:	Dunlop, Hugh Christopher RGC Jenkins & Co. 26 Caxton Street London SW1H 0RJ (GB)
Decision under appeal:	**Decision of the Examining Division of the European Patent Office posted on 18 September 2009 refusing European patent application No. 02789637.2 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman	S. Wibergh
Members:	K. Bumes
	S. Fernández de Córdoba

Summary of Facts and Submissions

I. This appeal is against the decision of the examining division to refuse European patent application No. 02789637.2, entitled "Memory adapted to provide dedicated and/or shared memory to multiple processors and method therefor", for lack of inventive step (Article 56 EPC 1973).

II. The decision under appeal is based *inter alia* on the following prior art documents:

D1: US-B1-6 266 751, entitled "*Continuously sliding window method and apparatus for sharing single-ported memory banks between two agents*"; and

D2: EP-A-0 856 796, entitled "*Variable-grained memory sharing for clusters of symmetric multi-processors*".

The examining division considered document D1 to represent the closest prior art in the form of a partitioned memory arranged to give exclusive access rights to a first processor and a second processor, respectively. The problem addressed by the present application was regarded as how to share data between the processors without having to spend time reconfiguring the memory partitions. To solve that problem, the skilled person would use well-known memory sharing techniques.

III. The statement setting out the grounds of appeal includes four sets of claims as a main request and first to third auxiliary requests, respectively. The appellant requests that the refusal decision be set aside and a patent be granted based on one of the four claim sets.

(a) Apparatus claim 1 according to the Main Request is identical to the refused version of the claim:

"1. An apparatus comprising a memory device (30) including a memory array (35) having a first portion (33), a second portion (31), and a third portion (32), the first portion of the memory array being different than the second portion of the memory array, and the third portion being different than the first portion and the second portion, wherein the third portion of the memory array is accessible by both a first processor (70) and a second processor (80),

characterised in that the first portion of the memory array is accessible only by the first processor and the second portion of the memory array is accessible only by the second processor,

wherein the size of the first portion and the second portion of the memory array is dynamically alterable depending on an operational load of the first processor or the second processor."

(b) According to the First Auxiliary Request, the last clause of claim 1 reads:

"wherein the size of the first portion, the second portion, and the third portion of the memory array is dynamically alterable depending on an operational load of the first processor or the second processor."

(c) The Second Auxiliary Request appends the following clause to claim 1 of the first auxiliary request:

"; and wherein said third portion (32) is used to store data not exclusive to either of the first or second processors."

(d) The Third Auxiliary Request appends a further additional clause to claim 1:

"; and wherein the memory array is further adapted to increase the size of the first portion (33) while decreasing the size of the third portion (32) in response to an increase in the operational load of the first processor."

IV. The appellant's arguments can be summarised as follows.

The technical effect of the difference over D1 was not necessarily to avoid wasting time reconfiguring memory portions. According to the present invention, it was not necessary for data to be shared between the processors. Therefore, the problem of the present invention was how to maintain sufficient exclusive access between two processors and respective memory portions during changes in operational loads, without compromising the exclusive access of one or other of the processors.

The invention solved this problem by providing a third memory portion accessible by both first and second processors, wherein said third portion could be configured to increase exclusive portions of both processors, whilst not at the expense of either.

D1 might not be the closest prior art as it was concerned specifically with sharing data between processors at the expense of one or the other. In contrast, the present invention rather related to dedicating sufficient memory to each of the respective processors, without compromising on the efficiency of either of the processors.

V. The Board summoned the appellant to oral proceedings, as requested on an auxiliary basis, and presented its preliminary analysis of the case. On a broad construction of claim 1, none of its four versions seemed to define an inventive contribution over the memory arrangement according to D1.

VI. By a letter received on 27 November 2014, the Board was informed that neither the appellant nor the representative would attend the oral proceedings scheduled for 10 December 2014. The Board then cancelled the oral proceedings.

Reasons for the Decision

1. The appellant initially requested, as an auxiliary measure, oral proceedings before the Board. After a hearing had been appointed, the appellant informed the Board that it would not attend. This statement is regarded as a withdrawal of the request for oral proceedings (T 3/90, OJ EPO 1992, 737).

2. According to the application, which was published as

A2: WO-A2-03/042834,

dedicating different memory devices to different processors reduces the risk of memory access conflicts but increases the size, complexity and cost of the overall system. Thus, there is a need for better ways to share memory in a system having two or more processors (A2, page 1, lines 13 to 22).

In its most general aspect (original claim 1), the application proposes to divide a memory array into a first portion accessible only by a first processor and

a second portion accessible only by a second processor. The size of the first and second portions may be dynamically altered depending on an operational load of the first or second processor (A2, page 6, line 23 to page 7, line 24; original claims 5 and 6). A third memory portion may be accessible by both the first processor and the second processor (A2, page 5, lines 6 to 15; original claim 2).

Main Request

Construction of claim 1

3. According to the characterising portion of claim 1, the size of the first memory portion and the second memory portion is dynamically alterable depending on an operational load of the first processor or the second processor.

That wording encompasses an operation in which the first memory portion is resized at the expense of the second memory portion, in line with the description (page 7, lines 1 to 4) and original claim 6 (now claim 4).

Therefore, claim 1 does not solve the alleged specific problem that sufficient exclusive access between two processors and respective memory portions during changes in operational loads should be maintained "*without compromising the exclusive access of one or other of the processors*".

4. Claim 1 features a third memory portion (32) accessible by both the first processor and the second processor. According to the description (page 8, lines 8 to 11), the additional memory "*portion 32 may be used to store*

data or instructions to be used by both processor 70 and 80 [...] or be used to store information to be shared or passed between processor 70 and processor 80".

Therefore, claim 1 is not susceptible to the appellant's argument that "*[a]ccording to the present invention, it is not necessary for data to be shared between both of two processors*".

Closest prior art

5. According to the appellant, "D1 may not be the closest prior art as it is concerned specifically with sharing data between processors at the expense of one or the other".

It is true that the present application does not deal with reallocating an exclusive memory portion *for the specific purpose of handing over data from one processor to the other*. However, the wording of claim 1 does not rule out that the dynamic alteration of the first and second exclusive memory portions may be performed to share data between the processors, like in D1.

Therefore, D1 still qualifies as the closest available prior art (even though the preamble of claim 1 represents a different starting point for which no source has been cited).

Article 56 EPC 1973 - Inventive step

6. D1 (e.g. Figure 1) describes first and second memory partitions for exclusive use by two respective "processing agents" (i.e. processors), the partitioning

being flexible and alterable "on-the-fly" (i.e. dynamically, see the paragraph bridging columns 4/5 and paragraphs 1 and 2 of column 6). The partitioning can be adjusted by a processor DSP (Digital Signal Processor) depending on its operational load (column 6, line 63 to column 7, line 12).

The introductory portion of D1 (column 2, lines 24 to 44) refers to a conventional design in which each processor is provided with its own dedicated memory bank system (D1, Figure 6). To reduce the need for transferring information between memory bank systems, memory may be shared (using an arbitrator, Figure 7, or a memory management address translator, Figure 8).

D1 further states that "[i]n a multiple agent system, one shared memory block is often provided for use by all agents in the system" (column 5, lines 41/42).

7. However, those references do not unambiguously point to a parallel use of shared (non-exclusive) and dedicated (exclusive) memory portions. D1 rather focuses on a data exchange by changing the partition of the dedicated memory (310) (column 6, lines 48 to 52; column 7, lines 33 to 42).

Therefore, D1 does not disclose a combination of two dedicated, dynamically alterable memory portions and a shared third memory portion.

8. This difference of claim 1 over D1 --- an (additional) third memory portion accessible by both processors --- provides an (additional) means for exchanging data between the processors.

As a result, no memory portion has to be reconfigured to transfer data from one processor to the other. However, that effect/advantage is not achieved in the whole range claimed since claim 1 encompasses re-configuring all three memory portions (see claim 4 of the main request; description page 6, line 23 to page 7, line 4; claim 1 of auxiliary request 3).

Therefore, the only objective problem solved is how to provide an (additional) way of exchanging data between two processors.

9. When an (additional) platform for exchanging data between two processors is desired, using (also) a shared memory should be a skilled person's first thought. In this respect, the examining division relied on common general knowledge as exemplified by D2 (column 4, lines 30 to 45; column 7, lines 50 to 54). The Board reiterates that D1 itself presents memory sharing as a conventional way of reducing the need for transferring information between dedicated memory banks.

10. No synergy derives from using a non-exclusive memory portion in addition to exclusive memory portions. The parallel use just achieves what is achieved by exclusive memory portions (D1, Figure 1) and a memory shared by plural processors (D1, D2). Memory sharing goes against exclusiveness and creates access conflicts (see D1, column 2, lines 52 to 55; A2, page 1, lines 16 to 18). Those drawbacks of a shared memory are not eliminated by combining it with dedicated memory portions.

11. The idea that an exclusive memory portion (33) may be increased at the expense of a non-exclusive memory

portion (32) is mentioned in one paragraph of A2 (top of page 7). The effect of such an approach --- the processor can process a larger amount of data --- is self-evident. Basically, said idea means that more memory is dedicated to a processor while less memory is available for shared use (since the overall memory capacity normally cannot be altered during operation). This reflects a skilled person's design choice.

12. Therefore, claim 1 (main request) defines an obvious modification of the memory arrangement of D1 so that its subject-matter lacks an inventive step.

First Auxiliary Request

13. The additional feature of claim 1 --- all three memory portions are dynamically alterable --- corroborates what has been said above with respect to the main request: the claim does not safeguard the exclusiveness of the dedicated memory portions.

Therefore, the only additional effect achieved is that the size of the third memory portion is alterable during operation depending on an operational processor load.

14. In practice, the overall memory capacity of a computer system is usually not increased during operation. Therefore, dynamically increasing the size of a dedicated memory portion usually implies that the size of another memory portion has to be decreased accordingly.

When increasing the size of the first memory portion, for example, it is up to the skilled person to decrease the size of the second dedicated portion (as in D1)

and/or the size of a non-dedicated third portion of the overall memory capacity.

15. Therefore, the first auxiliary request does not provide any inventive contribution, either.

Second Auxiliary Request

16. Compared with claim 1 of the first auxiliary request the second auxiliary request additionally defines that "said third portion (32) is used to store data not exclusive to either of the first or second processors".

The definition is redundant: As the third memory portion is accessible by both processors, it is evident that data stored therein is not exclusive to either processor.

Therefore, the Board's substantive assessment of claim 1 is the same as for the previous request.

Third Auxiliary Request

17. Claim 1 of the third auxiliary request contains the additional feature that the size of the first portion (33) is increased and the size of the third portion (32) is decreased in response to an increase in the operational load of the first processor.

18. In a realistic computer architecture where not the entire memory capacity has been divided up and dedicated to individual processors, it is obviously preferable to use a non-dedicated (third) memory portion to achieve a desired increase of the dedicated first memory portion. Reducing the second memory portion (dedicated to the second processor) is clearly

undesirable as reallocating parts of a dedicated memory portion would jeopardise its exclusive operation.

Therefore, the third auxiliary request does not provide any inventive contribution, either.

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar: The Chairman:

T. Buschek S. Wibergh

Decision electronically authenticated