

Internal distribution code:

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

**Datasheet for the decision
of 9 May 2014**

Case Number: T 1919/10 - 3.5.06

Application Number: 97950879.3

Publication Number: 1012712

IPC: G06F9/46, G06F12/10, H04Q11/04

Language of the proceedings: EN

Title of invention:
COMPUTER INTERFACE FOR DIRECT MAPPING OF APPLICATION DATA

Applicant:
Emulex Design & Manufacturing Corporation

Headword:
Direct Memory Mapping/EMULEX

Relevant legal provisions:
EPC 1973 Art. 56

Keyword:
Inventive step - (yes)

Decisions cited:

Catchword:



**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 1919/10 - 3.5.06

D E C I S I O N
of Technical Board of Appeal 3.5.06
of 9 May 2014

Appellant: Emulex Design & Manufacturing Corporation
(Applicant) 3333 Susan Street
Costa Mesa, California 92626 (US)

Representative: Viktor, Rainer
Vossius & Partner
Postfach 86 07 67
81634 München (DE)

Decision under appeal: **Decision of the Examining Division of the
European Patent Office posted on 30 March 2010
refusing European patent application No.
97950879.3 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman: D. Rees
Members: M. Müller
W. Sekretaruk

Summary of Facts and Submissions

I. The appeal lies against the decision of the examining division, with reasons dispatched on 30 March 2010, to refuse European patent application no. 97950879.3. The decision made reference to the documents

D1: Sterbenz J *et al.*, "Axon: A High-Speed Communication Architecture for Distributed Applications", Proceedings of the Ninth Annual Joint Conference of the IEEE Computer and Communications Societies INFOCOM, pp. 415-425, IEEE Press, 1990 and

D2: Tomimitsu Y *et al.*, "An ATM Chip Set for High Performance Computer Interfaces, Affording over 100 Mbps Sustained Throughput", IEICE Transactions on Electronics, vol. E78-C, no. 12, pp. 1738-1745, 1995

and came to the conclusion that claim 1 of then main request lacked an inventive step over D1, and that the dependent claims lacked an inventive step over D1 or over the combination of D1 and D2, Article 56 EPC 1973.

II. Notice of appeal was filed on 7 June 2010, the appeal fee being paid on the same day. A statement of grounds of appeal was received on 9 August 2010. The appellant requests that the decision be set aside and that a patent be granted based on claims 1-20 according to a main request or one of first to third auxiliary requests as filed with the grounds of appeal, the other application documents being as follows:

description, pages 3-13 as published
1, 1a, 2 filed on 30 December 2008
drawings, sheets 1/5-5/5 as published

Oral proceedings were requested in the event that the main request would not be allowed.

III. Independent claims 1 and 17 of the main request read as follows:

"1. A computer interface for accomplishing a transfer from a transmitting application (60a) running on a first computer (30a), having a system bus (42a) and a first main memory (44a) connected to said system bus, to a receiving application (60b, 60c) running on a second computer (30b, 30c), said transmitting application providing an operating system (46a) with transmit virtual addresses of data to be transmitted to said receiving application, said interface comprising:

- A. a mapping memory (95a) containing a map of transmit physical addresses in said first main memory corresponding to the transmit virtual address of the data to be transmitted;
- B. a transmit work queue pointer register (100a) directly accessible to said transmitting application and said operating system;

wherein the operating system is arranged to

- 1. load into said mapping memory the transmit physical addresses corresponding to said transmit virtual addresses, and
- 2. enter into said transmit work queue pointer register a pointer that points to a transmit work queue (101a) in the first main memory and wherein each entry in the transmit work queue specifies information associated with the data to be transmitted, said entry including an identification of the transfer, an amount of the data and an identification of a location of each of said transmit virtual addresses in said mapping memory;

wherein the interface further comprises:

means for accessing said first main memory in accordance with said physical addresses contained in said mapping memory to retrieve said data therefrom; and

means for transmitting said retrieved data to said receiving application on said second computer.

17. A method for accomplishing a transfer from a transmitting application (60a) running on a first computer (30a), having a system bus (42a) and a first main memory (44a) connected to said system bus, to a receiving application (60b, 60c) running on a second computer (30b, 30c), said transmitting application providing an operating system (46a) with transmit virtual addresses of data to be transmitted to said receiving application, wherein an interface:

- A. provides a mapping memory (95a) containing a map of transmit physical addresses in said first main memory corresponding to the transmit virtual address of the data to be transmitted;
- B. provides a transmit work queue pointer register (100a) directly accessible to said transmitting application and said operating system;

and software (46a) running on said first computer:

- 1. loads into said mapping memory the transmit physical addresses corresponding to said transmit virtual addresses, and
- 2. enters into said transmit work queue pointer register a pointer that points to a transmit work queue (101a) in the first main memory and wherein each entry in the transmit work queue specifies information associated with the data to be transmitted, said entry including an identification of the transfer, an amount of the data and an identification of a location of each

of said transmit virtual addresses in said mapping memory;

wherein the interface further:

accesses said first main memory in accordance with said physical addresses contained in said mapping memory to retrieve said data therefrom; and

transmits said retrieved data to said receiving application on said second computer."

In view of the board's decision it is sufficient to say that all the independent claims of the auxiliary requests comprise all of the features of the corresponding independent claims of the main request.

- IV. With a summons to oral proceedings, the board informed the appellant of its preliminary opinion that the independent claims according to any of the pending requests were non-obvious over D1 and its intention to set aside the decision but, instead of ordering the grant of a patent, to remit the case for further prosecution to the examining division. If the appellant agreed to this course of action the board indicated that it might be possible to cancel the oral proceedings.
- V. The appellant, with a letter of response dated 23 April 2014, agreed to remittal of the present case to the first instance for further prosecution. The oral proceedings were therefore cancelled.

Reasons for the Decision

1. The application relates to an interface supporting the communication between applications on different computers based on direct memory mapping.

- 1.1 In order for two application programs to communicate with each other, a so-called "virtual circuit" VC is set up. The application initiating the VC informs the operating system that a region of local physical memory is to be used for transmission and receipt of data, and which virtual addresses will be used to refer to the locations in this region (p. 5, lines 12-19). These virtual addresses are said to be "mapped" to physical memory addresses. The other applications in the circuit will have to do the same (see e.g. p. 6, lines 20-22).
- 1.2 For each application in a VC, a pair of queues is provided, a "transmit work queue" and a "receive work queue" (see figs. 1a and 1b) which contain *inter alia* addresses of data to be transmitted or received (p. 6, lines 2-4 and 26-29).
- 1.3 The application discloses a network interface unit, claimed as an "interface", in each computer involved in a communication connection (see fig. 1b, nos. 54b and 54c, and fig. 2). This network unit comprises *inter alia* a "transmit work queue pointer register" and a "receive work queue pointer register" (see p. 5, line 28 - p. 6, line 2; p. 6, lines 29-30; fig. 2, nos. 100a and 110a) containing pointers to the local transmit and work queues, and a "mapping memory" into which the mappings from virtual to physical memory for each VC are loaded on need (see e.g., p. 6, lines 7-9; fig. 2, no 95a).
- 1.4 When a particular transmission within a given VC is to take place, the network interface retrieves the virtual addresses from the entries in the pertinent transmit work queue, uses the mapping memory to map it to the corresponding physical addresses on the sender side, and retrieves the data from the local memory at these

addresses which it eventually transmits (see esp. p. 8, lines 14-20). At the receiving end, the data is retrieved and transferred to specific physical memory locations in the host memory as identified by the local memory mapping (see p. 9, lines 6-8).

- 1.5 The board thus understands the invention as describing the communication between applications via a shared virtual address space which is, at both ends of the connection, mapped to different physical memory spaces.

2. D1 discloses an architecture for communication of distributed processes across a VHSI ("very high speed internetwork") communication substrate. This communication is enabled by so-called Network Virtual Storage NVS which establishes a common address space for all communicating processes (see sec. 3.3). This address space allows "segments" to be addressed on a non-local host (p. 419, right col., 3rd par., 1st sentence): An address comprises a segment id and a host id which are resolved by reference, in particular, to a "known segment table" KST and a "known host table" KHT (see p. 419, last par. and p. 423, par. bridging left and right cols.). The Axon architecture relies on a communications processor CMP to provide the network interface where the "NVS assists" (p. 418, left col., last full sentence). The CMP provides two "data pipes" for outgoing and incoming data, the "transmit data pipe" and the "receive data pipe" (see p. 422, right col., 3rd par.; fig. 7).

3. In the decision under appeal (see p. 4, penult. par - p. 5, 1st par.), the examining division came to the conclusion that claim 1 of the then main request was novel over D1 only in that D1 disclosed "pipelines"

while the claim required "queues" but that these were obvious alternatives for the skilled person.

- 3.1 Regarding the alleged difference, the decision states that "the length of a[n] ... item is fixed in the queue and variable in a pipeline" (p. 4, last par.).
- 3.2 The board does not in fact see this as a difference, for two reasons. First, the board disagrees that the terms "queue" and "pipeline" (or pipe) imply in general whether their entries are fixed or variable length. And second, the board is not aware of any disclosure in D1 which would imply that the two CMP pipes allow variable length entries or anything in the present application which would exclude them. The board thus considers that the transmit and receive data pipes of D1 qualify as transmit and receive work queues as claimed (with the notable difference, however, that the queues are, according to the application, not part of the network interface but are held in host memory; see fig. 1).
- 3.3 The board however agrees with the appellant's argument that the interface according to D1 discloses neither a transmit work queue pointer register nor a mapping memory as claimed (see grounds of appeal, p. 3, point 1).
 - 3.3.1 Regarding the former it is noted that the queue pointer registers as disclosed in the application contain pointers to queues (see p. 5, line 28 - p. 6, line 2) but are not queues themselves. As it appears, these queue pointer registers enable the interface according to the invention to switch between the individual work queues provided for each virtual circuit.
 - 3.3.2 Regarding the latter it is noted that the independent claims of all requests - and in particular those of the

main request as refused - specify the mapping memory to map virtual addresses to physical memory on the local machine, whereas the mapping mechanism according to D1 must map virtual addresses to memory locations on the remote hosts, as argued by the appellant (see grounds of appeal, p. 5, 1st sentence).

4. Neither of these differences was addressed in the decision under appeal, let alone acknowledged. This is sufficient reason in itself for the board to come to the conclusion that the reasons in the decision under appeal do not justify the refusal.
5. Beyond that, however, the board considers that the invention relies on a different concept of virtual memory than D1 and that these concepts are incompatible with each other.
 - 5.1 D1 enables computers to access memory at other, remote computers. This turns the totality of memory of all connected computers into a "network virtual storage" accessible to all. Any remote address in D1 identifies a unique memory segment at some remote host, and the NVS services resolve this mapping. The address mapping must be available to and will be the same for any computer accessing remote storage. Furthermore, D1 does not disclose that a computer accessing its local memory relies on the NVS. Nor would this seem to be necessary or useful.
 - 5.2 On the other hand, the application maps the same virtual addresses to two corresponding regions of local memory in the hosts connected by a VC. The mapping according to the invention is thus specific to the local computer, be it on the sending or the receiving side.

- 5.3 It might be argued to be obvious in view of D1 to incorporate the memory mapping provided by the NVS into the CMP, for instance to speed up address resolution. But this need not be decided, because even on this assumption the system of D1 would only have to map "virtual addresses" to physical memory of remote hosts rather than of the local computer.
- 5.4 The board considers that this is due to the concept of Network Virtual Storage according to D1 and the board does not see how the skilled person would, in an obvious way, modify the system of D1 so as to include the memory mapping as claimed. In other words the board considers D1 to be an inappropriate starting point for assessing the inventive step of the claimed invention.
- 5.5 The board therefore comes to the conclusion that the claims 1 and 17 of the main request show an inventive step over D1, Article 56 EPC 1973.
6. The board is not, however, in a position to order the grant a patent for two reasons:
- 6.1 The search report mentions a number of other documents which might provide a more suitable starting point than D1 to assess inventive step of the claimed invention.
- 6.2 The search division may not have realised - or may have assessed differently - the difference between the concepts of virtual memory according to D1 and according the application. The board thus cannot exclude the possibility that the concept of communication via shared virtual memory or virtual mapped memory (VMMC) as realised by the invention was not or not exhaustively searched.

6.3 The board therefore exercises its right under Article 111 (1) EPC 1973 to remit the case to the first instance for further prosecution during which the examining division should in particular determine whether the search can be considered to be complete or whether an additional search should be carried out.

Order

For these reasons it is decided that:

1. The decision under appeal is set aside.
2. The case is remitted to the department of first instance for further prosecution.

The Registrar:

The Chairman:



B. Atienza Vivancos

D. Rees

Decision electronically authenticated