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) [X] To Chairmen

D E C I S I O N
of 6 November 1996

Case Number: T 0759/95 - 3.5.1

Application Number: 89309193.4

Publication Number: 0380851

IPC: G06F15/16

Language of the proceedings: EN

Title of invention:
Modular crossbar interconnections in a digital computer

Patentee:
DIGITAL EQUIPMENT CORPORATION

Opponent:
-

Headword:
-

Relevant legal provisions:
EPC Art. 52(1), 56

Keyword:
"Inventive step (yes, after amendment)"

Decisions cited:
-

Catchword:
-

Case Number: T 0759/95 - 3.5.1

D E C I S I O N
of the Technical Board of Appeal 3.5.1
of 6 November 1996

Appellant: DIGITAL EQUIPMENT CORPORATION
111 Powdermill Road
Maynard,
Massachusetts 01754-1418 (USA)

Representative: Maggs, Michael Norman
Kilburn & Strode
30 John Street
GB - London WC1N 2DD (GB)

Decision under appeal: **Decision of the Examining Division of the European Patent Office posted 13 April 1995 refusing European patent application No. 89 309 193.4 pursuant to Article 97(1) EPC.**

Composition of the Board:

Chairman: P. K. J. Van den Berg
Members: A. S. Clelland
C. Holtz

Summary of Facts and Submissions

I. European patent application 89 309 193.4, publication No. 380 851, was refused by a decision of the Examining Division dated 13 April 1995.

II. The reason given for the refusal was that the subject-matter of each of the claims lacked an inventive step having regard to the disclosure of the following document:

D1: EP-A-0 208 319

It was argued that D1 disclosed at Figure 11B an arrangement of cascaded image transfer controllers having the same structure as the claimed system; it was obvious for the skilled person to apply the known structure, in which mass storage devices were connected to workstation computers, to a multi-processor system having the claimed features of a plurality of CPUs, at least one I/O unit and at least one main memory unit.

III. On 12 June 1995 the Applicant (Appellant) filed a notice of appeal and paid the prescribed fee. A statement of grounds, together with new claims 1 to 4, was received on 15 August 1995.

IV. In a communication from the Board the Rapporteur discussed the interpretation of claim 1 and cited a further document, representative of the common general knowledge in the art:

D2: A.S. Tanenbaum, "Computer Networks", Prentice-Hall, 1981, ISBN 0-13-164699-0, pages 315 to 317.

The Rapporteur took the preliminary view that from the teaching of D2 it could be seen that the skilled person would consider multi-processor systems to be a form of network architecture; faced with a problem in the field of multi-processor systems the skilled person would consult document D1, in which data was moved between mass storage devices and computer workstations.

V. Oral proceedings were held on 6 November 1996. At the oral proceedings the Appellant requested that the decision under appeal be set aside and as a **main request** that a patent be granted on the basis of the following documents:

Claims: 1 as filed at the oral proceedings; 2 to 4 as filed on 7 October 1996

Description: pages 5, 5A as filed at the oral proceedings; page 12 as filed on 15 April 1994; pages 1 to 4, 6 to 11 and 13 to 38 as originally filed

Drawings: sheets 1 to 7 as originally filed.

As a **first auxiliary request** the Appellant requested grant on the basis of a main claim combining the subject-matter of claims 1 and 2. As a **second auxiliary**

request the Appellant requested grant on the basis of a main claim combining the subject-matter of claims 1 to 3.

VI. Claim 1 of the **main request** reads as follows:

"A multi-processing system (10) of the kind having a system control unit (SCU) (18) for operating a plurality of system units in a parallel fashion, the system units including a plurality of central processing units (CPUs) (12), at least one input/output (I/O) unit (14), at least one main memory unit (MMU) (16), and interconnection means (20) for establishing communication paths for data transactions between system units designated by communication commands as source and destination nodes, CHARACTERIZED IN THAT

the interconnection means comprises at least two crossbar switch modules (32, 34), each module being adapted to establish a direct path or mapping between one of a fixed number of source nodes and one of the same fixed number of destination nodes defined thereupon such that a first source node is capable of data transfer to a first destination node concurrently with at least one other data transfer between a second source node and a second destination node, the source and destination nodes provided on one crossbar switch module corresponding to different ones of the system units to those provided on other crossbar switch modules, each of the modules including an expansion source node and a corresponding expansion destination node through which the crossbar switch modules are

connectable in such a way as to establish a data path between source nodes defined on one of the crossbar switch modules and destination nodes defined on another of the crossbar switch modules, and the interconnection means further including a data switch controller means (50) centrally controlling the crossbar switch modules and comprising:

means (72) for accepting control commands designating the source and destination nodes for executing a data transaction,

means (82) for determining whether the designated source and destination nodes are defined on a single crossbar switch module or are defined on separate crossbar switch modules,

means (84) for generating a control signal to the single crossbar switch module containing both the source and destination nodes in order to establish the required data path by a direct mapping between the nodes, and

means (84) for generating separate control signals to each of the crossbar switch modules on which the source and destination nodes are separately defined in order to establish the required data path by an indirect mapping between the nodes, the indirect mapping being established by (i) a first submapping between the designated source node and the expansion node of the crossbar switch module containing the source node, and (ii) a second submapping between the designated

destination node and the expansion node of the crossbar switch module containing the destination node;

and further, wherein the multi-processing system includes a plurality of input/output units, and the interconnection means can establish data paths between any two of the central processing units, between each central processing unit and each memory unit, between each central processing unit and each input/output unit, between each input/output unit and each memory unit, and between any two of the input-output units."

VII. The Appellant put forward the following arguments in favour of inventive step:

D1 did not disclose a crossbar switch in the same sense as used in the application. The document was instead concerned with a shared bus, which the skilled person would understand as preferable to the more expensive and complicated crossbar arrangement. The skilled person would therefore not replace the switches in the known arrangement by crossbar switches.

Each interface module in D1 had its own separate controller. If interface modules were cascaded as suggested by Figure 11 of D1, the resulting arrangement would provide a separate controller for each such module. The skilled person would have no reason to take the step of removing the respective controllers and replacing them by a single, central, controller. The use of a central controller implied a loss of modularity and a limit, determined by the controller

construction, as to the number of switching connections available. The skilled person would be well aware that control became considerably more complex as the number of switching connections increased and would therefore consider that central control should be avoided. As specified in claim 1 of all the requests, the single data switch controller itself determined whether the input and output addresses supplied to it were in a single module or in different modules; in the latter case commands were sent out to the respective modules for each separate message.

Although it was accepted that other interpretations of D1 were possible, that put forward by the Board was not one which the skilled person would derive on a fair reading of the document. D1 itself was of some obscurity and open to differing interpretations, but the interpretation which tied in best with the various embodiments, in particular Figures 3 and 5, was that the mass storage devices were connected with the shared bus 54 in Figure 3, which corresponded to the data bus 102 in Figure 5. Data was thus multiplexed onto a bus. This would be understood by the skilled person to give an unacceptably slow performance for use in a multi-processor arrangement.

Reasons for the Decision

1. The appeal is admissible, Articles 106 to 108 and Rule 64 EPC.

2. The revised claims do not call for observation as

regards Article 84 EPC (clarity and support) or Article 123 (2) EPC (added subject-matter).

3. *The relevant prior art*

3.1 The correct starting point for considering inventive step would appear to be a conventional multi-processor computer system incorporating a single crossbar switch for interconnecting various components of the system, as acknowledged in the introduction to the description and discussed in document D2 at pages 315 to 317. The problem to be solved is said in the application to be the difficulty of expanding the number of nodes in such a system, the complexity of connections increasing by the square of the number of nodes. A change in the system configuration to increase the number of nodes therefore necessitates a complete redesign of the network. This problem is clearly appreciated in D2 at page 317, lines 3 to 15.

3.2 One solution is known from document D1 in the context of connecting workstation computers to mass storage devices. In the preferred embodiment of D1 a total of 64 workstation computers can be connected to receive data from one of 16 mass storage devices in the form of microfiche scanners. Switching between the workstation computers and the mass storage devices is performed in interface modules 60, see Figure 3, each interface module connecting to four workstation computers and each connection comprising four lines, namely data, data clock, response and command lines. In Figure 3 the mass storage devices are apparently shown as being

connected to a bus 54, which has led the Appellant to the conclusion that data flow between the mass storage devices and interface modules is by way of bus 54. The Appellant has moreover taken the view that bus 54 in Figure 3 corresponds to data bus 102 in Figure 5.

3.3 The Board has a different understanding of the operation of D1. From a consideration of Figure 5 it appears rather that the mass storage devices are connected by way of respective channels to switching circuits 200, there being four in each interface module connected to respective workstation computers. One switching circuit, referenced 204, is shown in detail in Figure 5 and includes four multiplexers 206, 208, 210 and 211 each receiving 16 inputs from respective data, clock, response and command lines of the mass storage devices, each multiplexer thus receiving an input from each mass storage device. The signal lines from the mass storage devices are therefore connected in parallel to all the multiplexers, and bus 54 of Figure 3 is a bus in the traditional sense rather than a multiplexed bus. The outputs of the multiplexers are supplied to respective workstation computers, WC0 in the case of switching circuit 204 and, for example, WC1 for switching circuit 246 shown at the bottom of Figure 5. It can moreover be seen from the description of D1, for example at page 4 lines 11 to 14, that data transfer between a mass storage device and workstation computer is capable of being carried out concurrently with at least one other data transfer between a second mass storage device and second workstation computer. It is moreover clear that the data bus 102 of Figures 4

and 5 is connected to a central processor 100 for sending address data to the switches 200 in order to select which mass storage device is connected to a particular workstation computer, see for example page 15 lines 15 to 22; Figure 5 also makes clear, see for example the input to latch 212, that data is transferred unidirectionally from bus 102 to the individual multiplexers.

3.4 The modification shown at Figures 11A to 11C of D1 enables interface modules to be cascaded to provide interconnections between larger numbers of devices. It can be seen that particular ports and channels are reserved for connection to cascaded devices thereby increasing the number of nodes available for interconnection.

4. *Inventive step (main request)*

4.1 It will be clear from the above discussion that the skilled person, faced with the problem known from D2 of increasing the number of nodes in a multi-processor system, would find a solution in document D1 which permitted him to increase the size of the network by modular expansion. Although it was argued that the skilled person would not take D1 into account because it is concerned with the specific problem of data flow from mass storage devices to workstation computers, the Board takes the view that the problem is one of network topology, independent of the devices to be connected to the network. D2 explicitly states at page 315, point 7.3, that multi-processor systems are a

specialized form of network.

4.2 The skilled person would therefore derive from D1 the teaching that a network using a crossbar switch module can be expanded by cascading the modules and would be led by D1 to apply such cascading to a multi processor system as in the acknowledged prior art exemplified by D2. The Board considers that no inventive step would be involved in such an application.

4.3 However, in the D1 arrangement each module has its own controller. It appears from D1 that no modification to the respective controllers is necessary in the event of cascading, the system being such that the workstation computers are themselves programmed to compute whether a particular connection will be through a single or plurality of modules and to send out address signals accordingly; thus, in the event of addressing through two modules the first module is addressed and, once connection is established, is deaf to any further addressed signals, see page 34, lines 10 to 17 of D1, subsequent address signals being picked up by the second module in the path. In accordance with claim 1 of the main request on the other hand the modules are centrally controlled.

4.4 The Board takes the view that, as argued by the Appellant, it would not be obvious for the skilled person to modify the D1 arrangement to provide for central control of the modules. The resultant loss of flexibility, necessitating reconfiguration of the switch controller each time a module is added or

removed, together with the increased complexity of control required with an increasing numbers of modules, would appear to prejudice the skilled person against any such modification. The Board moreover has no prior art at its disposal which would speak against such an understanding. A modification of the multiprocessor arrangement known from D2 to employ cascading as known from D1 would therefore result in a configuration in which each switching module has its own controller and the source and destination devices are themselves responsible for specifying a signal path.

4.5 For these reasons the Board concludes that the subject-matter of claim 1 of the main request involves an inventive step.

5. Since the main request has been found allowable it is not necessary to consider the first and second auxiliary requests.

Order

For these reasons it is decided that:

1. The decision under appeal is set aside.

2. The case is remitted to the Examining Division with the order to grant a patent on the basis of the Appellant's main request.

The Registrar: The Chairman:

M. Kiehl P. K. J. Van den Berg