

Internal distribution code:

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

**Datasheet for the decision
of 11 June 2015**

Case Number: T 2151/11 - 3.5.01

Application Number: 01304846.7

Publication Number: 1178407

IPC: G06F12/10, G06F3/06

Language of the proceedings: EN

Title of invention:

Architecture for parallel distributed table driven I/O mapping

Applicant:

Hewlett-Packard Development Company, L.P.

Headword:

I/O mapping/HEWLETT-PACKARD

Relevant legal provisions:

EPC Art. 123(2), 84

Keyword:

Amendments - added subject-matter (no)

Claims - clarity (no)



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 2151/11 - 3.5.01

D E C I S I O N
of Technical Board of Appeal 3.5.01
of 11 June 2015

Appellant: Hewlett-Packard Development Company, L.P.
(Applicant) 20555 S.H. 249
Houston, TX 77070 (US)

Representative: Lawman, Matthew John Mitchell
EIP
Fairfax House
15 Fulwood Place
London, WC1V 6HU (GB)

Decision under appeal: **Decision of the Examining Division of the European Patent Office posted on 25 May 2011 refusing European patent application No. 01304846.7 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman R.R.K. Zimmermann
Members: P. Scriven
D. Prietzel-Funk

Summary of Facts and Submissions

- I. The applicant appeals against the Examining Division's decision to refuse European patent application 01304846.7 on the grounds of added subject-matter (Article 123(2) EPC), lack of clarity (Article 84 EPC), and lack of inventive step (Article 56 EPC).

- II. With the statement setting out the grounds of appeal, the appellant maintained its main request as before the Examining Division, and submitted new auxiliary requests 1 - 3. The appellant referred to two distinct objections ("Clarity1", "Clarity2") in the decision under appeal and stipulated that auxiliary request 1 was to be considered withdrawn if the Board upheld the objection under Clarity 2, and auxiliary request 2 was to be considered withdrawn if the board upheld the objection under Clarity1 or added subject matter.

Clarity1 concerned entries in the mapping table that control input and output operations.

Clarity2 concerned the issue of what actions are an appropriate response to a fault signal.

- III. The Board sent a communication pursuant to Rule 100(2) EPC, setting its provisional assessment. It noted, in particular, ambiguities and difficulties in interpretation which amounted to a lack of clarity, to added subject-matter, and possibly to a lack of disclosure.

- IV. With a letter of response dated 9 February 2015, the appellant submitted a fourth auxiliary request.

- V. The Board arranged to hold oral proceedings and issued a summons. The oral proceedings were held as scheduled, but the appellant was not represented. The appellant's requests remained as follows: that the decision under appeal be set aside and that a patent be granted on the basis of a main request, or else of one of the three auxiliary requests submitted with the statement setting out the grounds of appeal, or of the fourth auxiliary request submitted with the letter dated 9 February 2015.
- VI. Claim 1 according to the main request and to the third and fourth auxiliary requests read as follows.

MAIN REQUEST

*A storage area network system (100) comprising:
a plurality of hosts (140) linked to a plurality of storage devices over a network (130);
a plurality of agents (110) connected to the hosts (140), each agent (110) having volatile memory (11) for storing a respective first copy of a mapping table (200), the mapping table (200) having entries (210) to map virtual disk positions (220) to locations (230) on the storage devices (160), entries (210) in the mapping table (200) each including a plurality of variables (240, 250, 260, 270) each indicating, when activated, the presence of a corresponding condition, in respect of that entry, in which an input/output (I/O) operation by a said host (140) is not to be completed; and
a controller (120) separate from, and coupled to, said plurality of agents (110), the controller (120) having a non-volatile memory*

(121) for persistently storing a second copy of the mapping table (200), and the controller (120) being operative to provide centralized control of mapping between virtual disk positions (220) and locations (230) on the storage devices (160) through the administration and distribution of the mapping table (200) including by intermittently causing contents of the first copy of the mapping table (200) to be replaced by contents of the second copy of the mapping table (200), the system being so arranged that during an input and/or output (I/O) operation by a said host (140):

the host uses a said agent (110) to access one of the entries (210) in the first copy of the mapping table (200) stored by that agent (110) to determine one of the storage device locations (230),
the agent (110), where the accessed entry has an activated variable indicating that the I/O operation is not to be completed, sending a fault message to the controller (120), and the controller responds to the fault message by determining an appropriate action and commanding the agent (110) accordingly.

THIRD AUXILIARY REQUEST

A storage area network system (100) comprising:

a plurality of hosts (140) linked to a plurality of storage devices over a network (130);

a plurality of agents (110) connected to the hosts (140), each agent (110) having volatile

memory (11) for storing a respective first copy of a mapping table (200), the mapping table (200) having entries (210) to map virtual disk positions (220) to locations (230) on the storage devices (160), each said entry (210) in the mapping table (200) including a plurality of variables (240, 250, 260, 270) each indicating, when activated, the presence of a corresponding condition, in respect of the entry, in which an input/output (I/O) operation by a said host (140) between the host and the storage devices is not to be completed; and

a controller (120) separate from, and coupled to, said plurality of agents (110), the controller (120) having a non-volatile memory (121) for persistently storing a second copy of the mapping table (200), and the controller (120) being operative to provide centralized control of mapping between virtual disk positions (220) and locations (230) on the storage devices (160) by intermittently causing contents of the first copy of the mapping table (200) to be replaced by contents of the second copy of the mapping table (200), the system being so arranged that during a said I/O operation by a said host (140):

a said agent (110), in response to an I/O command issued by the host to a virtual disk, accesses one of the entries (210) in the first copy of the mapping table (200) stored by that agent (110) to determine one of the storage device locations (230),

the agent (110), in the situation in which the accessed entry has an activated variable indicating that the I/O operation is not to be

completed, not attempting such completion but sending a fault message to the controller (120), and otherwise issuing a corresponding I/O operation command to the storage devices, and the controller responds to the fault message by determining an appropriate action and commanding the agent (110) accordingly, the appropriate action being one of sending a replacement table entry (210), and instructing the agent (110) to issue an error message to the host (140).

FOURTH AUXILIARY REQUEST

A storage area network system (100) comprising:
a plurality of hosts (140) linked to a plurality of storage devices over a network (130);
a plurality of agents (110) connected to the hosts (140), such that each agent is connected to at least one host, with each host connected to a corresponding agent, each agent (110) having volatile memory (111) for storing a respective first copy of a mapping table (200), the mapping table (200) having entries (210) to map virtual disk positions (220) to locations (230) on the storage devices (160), entries (210) in the mapping table (200) each including at least one to four Boolean variables (240, 250, 260, 270) representative of respective states of the virtual disk positions, said Boolean variables each indicating, when having a value representative of an activated state, the presence of a

corresponding condition, in respect of that entry, in which an input/output (I/O) operation by a said host (140) is not to be completed; and

a controller (120) separate from, and communicatively coupled to, said plurality of agents (110), the controller (120) having a non-volatile memory (121) for persistently storing a second copy of the mapping table (200), and the controller (120) being operative to provide centralized control of mapping between virtual disk positions (220) and locations (230) on the storage devices (160) through the administration and distribution of the mapping table (200) including by intermittently causing contents of the first copy of the mapping table (200) to be replaced by contents of the second copy of the mapping table (200),

the system being so arranged that during an input and/or output (I/O) operation by a said host (140):

a said agent (110) corresponding to the said host accesses one of the entries (210) in the first copy of the mapping table (200) stored by that agent (110) to determine one of the storage device locations (230),

the agent (110), where the accessed entry has a Boolean variable value representative of an activated state indicating that the I/O operation is not to be completed, sending a fault message to the controller (120), and the controller responds to the fault message by determining an appropriate action and commanding the agent (110) accordingly.

Reasons for the Decision

Introduction

1. The invention concerns virtual storage, shared by a number of computers ("hosts"). It is not necessary to set out the details of such a system. It suffices to note that there are storage locations as seen by the storage devices themselves, storage locations as seen by the hosts, and a way of mapping between the two. The mapping is held in a table, so that when a host seeks, say, to write data to what it calls location x , the corresponding location y can be found.
2. Each host has an agent, though more than one host may share a single agent. It is part of the agents' job to convert the host's terminology (x) to the storage devices terminology (y).
3. One way of arranging things is to keep the mapping table at some location accessible to all the agents. Then, for every input or output operation to storage, an agent would consult this "central" table to find the correct conversion. This has the advantage that all agents necessarily see the same mapping, but the disadvantages that a single failure renders the whole system inoperative and that each consultation of the table involves communication over the network.
4. The invention uses a different arrangement. Each agent keeps its own mapping table. That makes access faster, and has the additional advantage that the loss of one table does not mean failure of the whole system. However, it is then necessary to ensure that all the tables in the different agents are the same. To achieve

this, the invention keeps a "second copy" of the table at a "controller", and from time to time sends this second copy to the agents.

5. The invention also foresees that the mapping tables additionally hold a number of "flags" or "state variables". An example is the "no-write" variable. This is a Boolean variable which is stored for each of the host locations. When a host seeks to write to a location *x*, its agent will consult the mapping table and find both the corresponding storage location *y* and whether "no-write" is "true" or "false". If the host wants to write and the no-write variable indicates that writing is not allowed, then the agent asks the controller what to do.

The main request, claim 1

6. The Board indicated, in its communication under Rule 100(2) EPC, that there was no basis in the application as filed for a host accessing entries in a mapping table, but that the term "the host uses said agent ... to access one of the entries", used in claim 1, allowed an interpretation under which the host, making use of the agent, did access the table. Thus the requirements of Article 123(2) EPC appeared to not be fulfilled.
7. The appellant has not argued that such an interpretation was wrong, but submitted that the fourth auxiliary request overcame this objection.
8. The appellant argued, for example in the statement setting out the grounds of appeal at point 3.1, that the disclosure of the application as filed, taken as a whole, indicated that it was the agents that had direct

access to the tables, and that the hosts' access was only indirect. The Board accepts that this is so. The claim, however, is ambiguous and allows for direct access by the hosts.

9. The Board concludes that this claim comprises subject matter that extends beyond the content of the application as filed (Article 123(2) EPC).
10. The Board also informed the appellant that it considered the term "variable" so broad that the meanings of "activated" and of "appropriate action" were unclear within the meaning of Article 84 EPC and, indeed, that there seemed to be sufficient disclosure of "appropriate action" only for the four specific Boolean variables ("invalid", "no-write", "zero", "error") set out in the original description.
11. Again, the appellant did not argue that this was incorrect, but submitted that the fourth auxiliary request overcame the objection.
12. The appellant submitted arguments regarding the term "the controller responds to the fault message by determining an appropriate action and commanding the agent accordingly". For example, at point 3.3 of the statement of grounds, the appellant expresses it thus: "On the contrary, the passage is perfectly clear: the controller is specified as responding to a fault message from an agent by commanding the agent to carry out an appropriate action," and that the controller becomes "involved in fault cases and provides a centralized control for dealing with faults, thereby ensuring consistency across the mapping agents and making it a relatively simple matter to universally change the action determined as appropriate for dealing

with a particular fault."

13. The Board sees these statements on the one hand as asserting clarity without providing any supporting argument, and on the other hand as a statement of an advantage. There is no argument regarding the meaning of "appropriate" over the broad scope of "variable."
14. The appellant, in its reply to the Board's communication under Rule 100(2) EPC, submitted that the "limitation to 'Boolean variables' clarifies any issue regarding the 'activation' of states, as Boolean variables can take only one of two possible states, meaning that these variables can have a value that represents an activated state." That is directed to claim 1 according to the fourth auxiliary request, and does not speak to the main request. The Board, therefore, sees no reason to depart from its view that it is unclear how the term "activated" could apply to a variable representing (to take one example) temperature.
15. The Board concludes that the claim is unclear (Article 84 EPC).

The first and second auxiliary requests

16. Since the main request is not allowed on the grounds of Clarity², the first auxiliary request is considered withdrawn. Since the main request is also not allowed due to added subject-matter in respect of the hosts' access to mapping tables, as in the appealed decision, the second auxiliary request is also considered withdrawn. (See point 4.1 of the statement setting out the grounds of appeal and point II., above).

The third auxiliary request, claim 1

17. This claim defines the inclusion, in the mapping table, of "a plurality of variables ... indicating, when activated, the presence of a corresponding condition ... in which an input/output (I/O) operation ... is not to be completed." As set out above (point 10. to 15.), it is unclear what "activated" means in this context. The Board concludes that the request cannot be allowed under Article 84 EPC.

The fourth auxiliary request, claim 1

18. The final clause of this claim is the same as in the main request: the controller determines an appropriate action and commands the agent accordingly. However, the lack of clarity (points 10. to 15., above) does not carry over, because this claim is restricted to Boolean variables.
19. It is perhaps arguable, that for the four Boolean variables set out in the description ("invalid", "no-write", "zero", "error") it would be clear which actions are appropriate and which are not. The Board, however, is not persuaded because appropriateness is a matter for a larger context than only the state of a variable. For example, it may sometimes be appropriate to respond to a "no-write" state with a simple error message, but at other times to provide a user with detailed information. The question of the standard by which appropriateness is to be determined is answered neither in the claim nor elsewhere in the application.

20. The questions about which actions are appropriate are multiplied when, as in this claim, the range of possible Boolean variables is left completely open, and the Board concludes that the claim lacks clarity, Article 84 EPC.

Conclusions

21. The first and second auxiliary requests are considered withdrawn. The main request, and the third and fourth auxiliary request cannot be allowed.

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar:

The Chairman:



T. Buschek

R.R.K. Zimmermann

Decision electronically authenticated