Internal distribution code:

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

Datasheet for the decision of 25 May 2018

Case Number: T 2502/11 - 3.5.01

Application Number: 05732940.1

Publication Number: 1735714

IPC: G06F15/173, G06F15/80

Language of the proceedings: EN

Title of invention:
SYSTEM AND METHOD FOR CLUSTER MANAGEMENT FOR PARALLEL TASK ALLOCATION IN A MULTIPROCESSOR ARCHITECTURE

Applicant:
RAYTHEON COMPANY

Headword:
Cluster management / RAYTHEON COMPANY

Relevant legal provisions:
EPC Art. 84, 83, 56

Keyword:
Sufficiency of disclosure - (no - lack of enabling disclosure)
Claims - clarity (no)
Inventive step - (no - juxtaposition of features)

Europäisches Patentamt
European Patent Office
Office européen des brevets

Beschwerdekammern

Boards of Appeal

Chambres de recours

Boards of Appeal of the
European Patent Office
Richard-Reitzner-Allee 8
85540 Haar
GERMANY
Tel. +49 (0)89 2399-0
Fax +49 (0)89 2399-4465

Case Number: T 2502/11 - 3.5.01

D E C I S I O N
of Technical Board of Appeal 3.5.01
of 25 May 2018

Appellant: (Applicant)	RAYTHEON COMPANY 870 Winter Street Waltham, MA 02451 (US)
Representative:	Barker Brettell LLP 100 Hagley Road Edgbaston Birmingham B16 8QQ (GB)
Decision under appeal:	**Decision of the Examining Division of the European Patent Office posted on 13 July 2011 refusing European patent application No. 05732940.1 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman	W. Chandler
Members:	N. Glaser
	Y. Podbielski

Summary of Facts and Submissions

I. This appeal is against the decision of the examining division to refuse the European application No. 05732940.1 under Article 123(2) EPC and Article 84 EPC. In a section called "Further grounds" the examining division presented arguments why the application was not inventive (Article 56 EPC) over inter alia:

FEITELSON D. G.:"Job scheduling in multiprogrammed parallel systems", IBM RESEARCH REPORT, SAN JOSE, CA, US, August 1997, pages 1-4 (D1).

II. In the statement setting out the grounds of appeal, dated 23 November 2011, the appellant requested that the decision under appeal be set aside and that a patent be granted on the basis of the main or auxiliary request, filed therewith.

III. In a communication accompanying a summons to oral proceedings, the Board set out its preliminary opinion that the invention was not sufficiently disclosed (Article 83 EPC) and that claim 1 of the main and auxiliary requests was not clear (Article 84 EPC) and did not involve an inventive step (Article 56 EPC). The Board also communicated its objections under Article 123(2) EPC.

IV. In a response, the appellant did not present any new requests or arguments, but informed the Board that he would not be attending the oral proceedings. In consequence the oral proceedings were cancelled and the procedure continued in writing.

V. Claim 1 of the main request reads as follows:

A computer readable medium for cluster management of computing nodes, comprising:
code for executing a plurality of cluster agents,

each cluster agent associated with one of a plurality of nodes, each node comprising a switching fabric integrated to a card and at least two processors integrated to the card,

the cluster agent operable to determine the status of the associated node; and

code for executing a cluster management engine communicably coupled to the plurality of nodes and operable to allocate nodes in a cubic volume to minimize inter-node communication distances,

the cluster management engine operable to dynamically allocate a particular subset of the plurality of nodes to a particular job based on the determined status of each of one or more of the plurality of nodes and execute the job using the particular subset,

the cluster management engine operable to receive a job placement strategy comprising a requested allocation of nodes,

the cluster management engine operable to allocate nodes for the particular subset using a placement weight, the placement weight indicating how the cluster management engine is to allocate nodes to the particular subset between allocating nodes only if the requested job placement strategy is possible and allocating nodes immediately provided there are enough available nodes to handle the job.

VI. Claim 1 of the first auxiliary request adds the following features:

wherein, to execute the job using the particular subset, the cluster management engine is operable to:

receive a job request comprising one or more job parameters;

determine dimensions of the job based, at least in part, on the one or more job parameters;

select a policy based on the job request; and

dynamically determine the dimensions of the job further based on the selected policy;

dynamically allocate the particular subset based, at least in part, on the determined dimensions; and

execute the job using the particular subset.

VII. The appellant essentially argued as follows:

The amendment of a cluster management engine operable to receive a "job placement strategy", which comprised a requested allocation of nodes for the job, and the definition of a "placement weight" and how it was used to allocate nodes in relation to the job placement strategy, overcame the grounds for refusal.

In summary, compared to D1, the prior art did not disclose a dynamic allocation of a subset of nodes to a job based on the status of each node and based on a placement weight in combination with a job placement strategy.

Reasons for the Decision

1. Articles 84 and 83 EPC - Main and auxiliary request

1.1 The invention concerns High Performance Computing (HPC) of a cluster of one or more processors referred to as nodes. At a high level, shown in Figure 1 of the application, HPC server 102 includes a cluster management engine 130 (500 in Figure 5), a management node 105 and a grid 110 comprising a plurality of nodes 115.

1.2 The description formulates on page 2, first paragraph, various problems which the invention is said to overcome, including the problem of reducing communication overhead between nodes (increased I/O performance). This is said to achieve an efficient operation of the cluster in production-oriented environments. All of these problem formulations have in common that they are given at a general and generic level. There is a lack of a specific problem to be overcome by the invention.

1.3 The appellant claims that the use of a "placement weight" by the cluster management engine in combination with a "job placement strategy" is novel and inventive, but fails to clearly identify the specific problem to be overcome.

1.4 The placement weight is according to the description a job or policy parameter (not claimed) and represents how "aggressively" the cluster management engine should attempt to place nodes according to a placement strategy, that is, weight 0 means placement takes place only if an "optimal strategy" (not claimed) is

possible, whereas weight 1 means immediate placement (page 21, third paragraph).

1.5 This feature is disclosed within a single paragraph in the application, page 21, lines 9 to 18, which does not make explicit the technical effect (if any) on the allocation of nodes of the virtual cluster.

Furthermore, it is unclear and insufficiently disclosed how nodes are to be placed to fulfill the condition "only if the **optimum strategy** (or dimensions) is possible" (emphasis added), without providing further technical details for the person skilled in the art about how to implement such a strategy.

Furthermore, no information is given to decide under which conditions a strategy would be regarded as an "optimum strategy".

1.6 The Board therefore considers that the description at such a generic level does not enable the person skilled in the art to implement the claimed features, regarded as essential by the appellant.

1.7 Moreover, the description, page 20, fourth paragraph, tells the reader that the choice of a "Best Fit Cube" topology for the creation of a virtual cluster already minimises the distance between any two nodes. It therefore appears that the problem of reducing communication overhead is solved by choosing the right topology (feature "allocate nodes in a cubic volume"), without requiring a job placement strategy along with a placement weight.

1.8 In view of the above, the Board judges that the alleged invention is not sufficiently disclosed (Article 83 EPC).

1.9 In addition, the claims and the description employ different terms for what appear to be the same features. For instance, the terms "allocation of nodes to virtual clusters/into a job space/in a cubic volume", "placement of nodes according to task or process" or "assignment of tasks to nodes" all seem to refer to the generation of clusters of nodes, based on "policies", "placement strategies", "job parameters" or "processing and managing information".

1.10 The use of a different feature terminology obscures how the allocation of nodes is to be implemented and leads to a lack of clarity (Article 84 EPC). It is therefore not possible to clarify the unclear terms in the claims, such as "job placement strategy", based on the description.

2. Article 123(2) EPC - Main and auxiliary request

2.1 The introduction of the feature "placement weight" in claim 1 of the main and auxiliary request, contravenes Article 123(2) EPC.

The essential function of the feature "placement weight" defined in the description, page 21, lines 9 to 16, as *a job or policy parameter* has been omitted from the claims; furthermore, the feature "placement weight" is linked according to the description to an optimum job placement strategy and not to a job placement strategy in general.

The feature "placement weight" is thus an unallowable

generalisation.

2.2 Furthermore, the deletion of the feature "HPC" from the original claims contravenes Article 123(2) EPC, because this feature is mentioned throughout the application in all embodiments and implies highly scalable nodes, comprising an integrated switch, which is essential for the problem to be solved.

3. Article 56 EPC - Main and auxiliary request

3.1 D1 discloses the creation of virtual clusters based on different topologies (Table 1; 3D cubes in Figure 3; hypercubes on page 16, middle paragraph); a hypercube is said to comprise power-of-two processing entities (PE), page 19f.; section 3.2.2 on page 24 discloses that a set of PEs are connected via switches to a set of buses and that by proper settings of the switches, partitions with different numbers of PEs and buses can be created. This allows both, the computation power and the communication capacity, to be tailored to an application's needs. D1 discloses choosing the most appropriate partition size and topology, the most appropriate structure of the PEs and the most appropriate job execution mechanics for a job, pages 27 to 30.

3.2 The appellant contests that D1 discloses the features (a) "cluster management engine operable to dynamically allocate a subset of the plurality of nodes to a particular job based on the determined status of ... the plurality of nodes and execute the job ..." and (b) "... allocate nodes for the particular subject using a placement weight ...".

3.3 The Board is not convinced by the applicant's arguments and considers both features to be known from D1.

Feature (a) is about taking run-time conditions of a job/partition into account when allocating nodes to a job. This is known from D1, Sections 3.3 and 3.4, which disclose an adaptive partitioning of a cluster of nodes according to load conditions of the system. The size of the partition to be allocated to a job is based on "minimum response time", but also to maximize throughput of the system. Other criteria are job characteristics ("application characteristics").

Regarding feature (b), the Board notes that the purpose of a "placement weight" (a job parameter) is - according to the description of the application - limited to a criteria of applying an (optimum) placement strategy or not. This is known from D1, for instance, Section 3.2.1, which discloses the creation of partitions having a topology of hypercubes based on the dimensions needed by a job.

3.4 Accordingly, claim 1 of the main and auxiliary requests contain a mere juxtaposition of known features for a known purpose and do not involve inventive step (Article 56 EPC).

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar: The Chairman:

T. Buschek W. Chandler

Decision electronically authenticated