

**Internal distribution code:**

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

**Datasheet for the decision  
of 20 December 2024**

**Case Number:** T 1503/20 - 3.5.06

**Application Number:** 10187436.0

**Publication Number:** 2442228

**IPC:** G06F9/50

**Language of the proceedings:** EN

**Title of invention:**

A computer cluster arrangement for processing a computation task and method for operation thereof

**Applicant:**

ParTec AG

**Headword:**

Computer cluster arrangement/PARTEC

**Relevant legal provisions:**

EPC Art. 84, 56, 111, 123(2)

EPC R. 103(1)(a)

RPBA 2020 Art. 13(2), 12(2), 11

**Keyword:**

Inventive step - over the documents relied upon in the decision (yes)

Remittal to the examining division for further prosecution

Reimbursement of appeal fee - (no)



**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

Case Number: T 1503/20 - 3.5.06

**D E C I S I O N**  
**of Technical Board of Appeal 3.5.06**  
**of 20 December 2024**

**Appellant:** ParTec AG  
(Applicant) Possartstr. 20  
81679 München (DE)

**Representative:** Himmelsbach, Mathias  
Frohwitter  
Patent- und Rechtsanwälte  
Possartstraße 20  
81679 München (DE)

**Decision under appeal:** **Decision of the Examining Division of the  
European Patent Office posted on 17 January 2020  
refusing European patent application No.  
10187436.0 pursuant to Article 97(2) EPC.**

**Composition of the Board:**

**Chairman** M. Müller  
**Members:** M. Domingo Vecchioni  
A. Jimenez

## **Summary of Facts and Submissions**

I. The appeal is against the decision of the examining division to refuse European patent application No. 10187436.0 on the basis that the then main request and auxiliary request did not meet the requirement of inventive step, Articles 52(1) and 56 EPC, in view of the following documents:

D1: US 2009/0213127 A1,

D2: S. Pakin et al., "The reverse-acceleration model for programming petascale hybrid systems", IBM Journal of Research and Development, vol. 53, no. 5, paper 8, 2009, XP055595871.

II. In the minutes of the oral proceedings before the examining division (point 3.1), reference was also made to the following document:

D3: US 2009/217266 A1.

The board notes that is not the same document as that referred to as "D3" in the decision (point 2).

III. With the statement of grounds of appeal, the appellant requested that the decision of the examining division be set aside and a patent be granted on the basis of the main request or, alternatively, the auxiliary request underlying the contested decision.

The appellant also requested the reimbursement of the appeal fee due to a substantial procedural violation committed by the examining division, the contested

decision being not reasoned within the meaning of Rule 111(2) EPC.

- IV. Following summons to oral proceedings, the board presented its preliminary opinion in a communication pursuant to Article 15(1) RPBA dated 10 June 2024. The board was not convinced by the examining division's inventive step reasoning starting from D1 but raised objections under Articles 84 and 123(2) EPC against the main request and the auxiliary request. As regards the alleged substantial procedural violation, the board tended to consider that the contested decision was sufficiently reasoned within the meaning of Rule 111(2) EPC and that the appellant's right to heard, Article 113(1) EPC, had been respected.

The board noted that even if all outstanding objections raised in the communication were overcome, it would likely not remit the case to the examining division for grant but for further prosecution, Article 111(1) EPC. Another document,

D6: US 2010/058036 A1,

had been cited as "X" document by the search examiner for a family member of the present application, in which similar subject-matter was claimed, and it would be for the examining division to consider whether this document was relevant for the present claims.

- V. With a letter dated 18 June 2024, the appellant filed sets of claims for a new main request and new auxiliary requests 1 to 9 to replace all previous requests. The appellant commented *inter alia* on D6.

VI. In a communication dated 25 June 2024, the board indicated that only auxiliary request 9 appeared to fully overcome the objections under Articles 84 and 123(2) EPC raised by the board. The appellant was asked whether it maintained the higher-ranking. If not, the board would be minded to cancel the oral proceedings and remit the case for further prosecution, but not to reimburse the appeal fee.

VII. With a letter dated 28 June 2024, the appellant indicated that it withdrew the main request and the auxiliary requests 1 to 8, leaving the auxiliary request 9 as the only request maintained.

VIII. The oral proceedings were thereupon cancelled.

IX. The appellant thus finally requested that the decision of the examining division be set aside and a patent be granted on the basis of the auxiliary request 9 filed on 18 June 2024. Furthermore, reimbursement of the appeal fee pursuant to Rule 103(1)(a) EPC was requested.

X. Claim 1 reads as follows:

"A computer cluster arrangement for processing a computation task, comprising:

    a plurality of computation nodes (CN) for computation of the computation task, each of which interfacing a communication infrastructure (IN), at least two of said communication nodes (CN) being arranged to jointly compute at least a first part of said computation task;

    at least one booster (B) being arranged to compute at least a second part of said computation task, each of the at least one booster (B) interfacing said

communication infrastructure (IN);

wherein each of the computation nodes (CN) and each of the at least one booster (B) interfacing the communication infrastructure (IN) such that each of the computation nodes (CN) and each of the at least one booster (B) are coupled with the communication infrastructure (IN);

wherein the communication infrastructure (IN) allows each of the computation nodes (CN) to communicate with each of the at least one booster (B) over the communication infrastructure (IN) without the necessity to involve a further computation node (CN); and

a resource manager (RM) being arranged to assign at least one booster (B) to at least one of said plurality of computation nodes (CN) for computation of said second part of said computation task;

wherein the resource manager (RM) is arranged to establish an initial assignment being accomplished as a function of a predetermined assignment metric, the predetermined assignment metric being a set of rules;

wherein the resource manager (RM) is arranged to update said predetermined assignment metric during computation of said computation task; and

wherein the resource manager (RM) is arranged to establish a dynamic assignment as a function of the updated assignment metric during computation of said computation task."


## **Reasons for the Decision**

### *The application*

1. The application relates to a computer cluster arrangement comprising "computation nodes CN" and "boosters B". "Boosters" are accelerators, such as GPUs, which can assist the computation nodes in accomplishing a given computation task, such as a simulation (page 1, lines 9-15, and page 3, lines 4-7 and 31-33, of the description as filed).
2. According to the description, known in the prior art are arrangements in which accelerators are tightly coupled to the computing nodes for outsourcing computations of high resource requirements. Such a "static assignment" of accelerators to computing nodes leads "to over- or under subscription of accelerators" and does "not provide fault tolerance in case of accelerator failures". Accelerators do also not exchange information directly but via their respective computation nodes (page 1, lines 20-26; page 2, lines 4-15; figure 1).
3. The "object of the present invention" is "to provide a computer cluster arrangement which allows communication flexibility as regards data exchange between accelerator and computation nodes as well as direct access of computation nodes to any and each of the accelerators". Furthermore, an object is also "to provide dynamic coupling of accelerators to computation nodes at runtime" (page 2, lines 17-21).



4. The proposed approach is illustrated in figure 2:



The computer cluster arrangement comprises a cluster (C) of computation nodes (CN), a group (BG) of boosters (B) and a "resource manager" (RM). Each computation node and each booster "interfaces" a "communication infrastructure" (IN), such as an InfiniBand interconnect. In the proposed approach, "it is possible to directly couple boosters to the communication infrastructure without the necessity of an intermediate computation node" as in the prior art referred to in the application, and "each computation node communicates with each booster over the communication infrastructure, without the necessity to involve a further [computation] node while exchanging data from a computation node to a booster" (page 3, line 4 to page 4, line 14; page 9, lines 10-27).

5. A computation task is processed by at least one of the computation nodes CN. At least a part of the computation task may be forwarded to at least one of the boosters B. The booster computes the part and the result is delivered back to the computation node (page 9, lines 19-24).

6. The resource manager decides which booster is assigned to which computation node on the basis of an "assignment metric", which is a "set of rules". According to the description, "[f]or determining an initial assignment of boosters to computation nodes, the assignment metric is predetermined but may be altered at runtime" and "static assignment is provided at start of the processing of computation task and dynamic assignment is provided at runtime" (page 4, lines 16-28; page 9, lines 24-27).

This process is illustrated in figure 6, which is described as showing "a block diagram of control flow of a computer cluster arrangement according to an aspect of the present invention" (page 13, lines 23-30).

7. Figure 7, on the other hand, is described as showing a block diagram of control flow "implementing reserve acceleration" of "a computer cluster arrangement according to an aspect of the present invention". In this embodiment, "an acceleration of computation of computation tasks being computed by at least one booster B is performed by assigning at least one computation node CN to at least one booster B", so that "the control and information flow is reversed as regards the embodiment being shown in Fig. 6" (page 13, line 32 to page 14, line 2; see also page 6, lines 1-3).

#### *Admittance*

8. The present set of claims was filed after notification of the summons to oral proceedings. It overcomes all the objections under Articles 84 and 123(2) EPC that were raised for the first time by the board in its

preliminary opinion (see point 9 below), without introducing new issues. In view of these exceptional circumstances, the board admits this set of claims into the appeal proceedings, Article 13(2) RPBA.

*Articles 84 and 123(2) EPC*

9. The board considers that the amendments made to the claims are compliant with the requirements of Article 123(2) EPC and overcome the objections under Articles 84 and 123(2) EPC that were raised by the board in the communication pursuant to Article 15(1) RPBA dated 10 June 2024 (hereinafter: the "board's communication") in respect of the claims of the main request underlying the contested decision.

- 9.1 In claim 1, the added features

"wherein each of the computation nodes (CN) and each of the at least one booster (B) interfacing the communication infrastructure (IN) such that each of the computation nodes (CN) and each of the at least one booster (B) are coupled with the communication infrastructure (IN); wherein the communication infrastructure (IN) allows each of the computation nodes (CN) to communicate with each of the at least one booster (B) over the communication infrastructure (IN) without the necessity to involve a further computation node (CN) "

are derivable from page 4, lines 4-14, and page 6, lines 25-29, of the original description, as well as figures 2, 3 and 9 (showing a direct coupling of computation nodes and boosters via the communication infrastructure) and figure 8 (showing an indirect

coupling via "interfacing units" and the communication infrastructure).

They make clear that "interfacing" in claim 1 means that the computation nodes and boosters are coupled to the communication infrastructure and that this coupling is for the purpose of communicating data with each other via said communication infrastructure (as argued by the appellant in its letter dated 18 June 2024, points A.1 and A.5). The coupling to the communication infrastructure may be direct or indirect (e.g. be via "interfacing units" as in figure 8), but claim 1 requires that each of the computation nodes may communicate with each of the boosters via the communication infrastructure "without the necessity to involve a further computation node". Hence, a booster cannot be coupled to the communication infrastructure only via a computation node. The added features overcome thereby the objection under Article 84 EPC regarding the term "interfacing" in previous claim 1 raised in point 13 of the board's communication.

9.2 In claim 1, the amendments made to the following features:

"wherein the resource manager (RM) is arranged to establish an initial ~~the~~ assignment being accomplished as a function of the predetermined assignment metric, the predetermined assignment metric being a set of rules;  
wherein the resource manager (RM) is arranged to update said predetermined assignment metric during computation of said computation task; and  
wherein the resource manager (RM) is arranged to establish a dynamic assignment as a function of the updated assignment metric during computation of said

~~*computation task is dynamically established during processing of the computation task".*~~

are derivable from page 4, lines 16-28, of the original description and original claim 10.

These amendments clarify that the "predetermined assignment metric" is a *set of rules* that is used for an *initial* assignment of boosters to computation nodes, an *update* of said predetermined assignment metric being used for a *dynamic* assignment during computation of the computation task. This overcomes the objections under Articles 84 and 123(2) EPC regarding the features of previous claim 1 regarding the "predetermined assignment metric" raised in point 14 of the board's communication.

9.3 The objection under Article 123(2) EPC raised against the last feature of previous claim 1 relating to "reverse acceleration" in point 15 of the board's communication has become moot due to the deletion of that feature in present claim 1. The deletion itself is compliant with Article 123(2) EPC as this feature was not present in original claim 1.

9.4 The objections under Articles 84 and 123(2) EPC that were raised against independent claim 11 of the previous main request in point 16 of the board's communication have become moot due to the deletion of that independent claim in the present set of claims.

*Inventive step starting from D1*

10. The examining division considered that claim 1 of the then main request differed from the disclosure of D1 only in its last feature relating to "reverse

acceleration". This feature being obvious in view of D2, claim 1 lacked an inventive step (decision, point 12).

The appellant argued in the statement of grounds of appeal that claim 1 of the then main request differed from the disclosure of D1 in several other features. The examining division had misunderstood the meaning of "streaming attachment" as used in D1, which became clear when considering the parallel application D3. There was no motivation for a skilled person to modify the teaching of D1 in a manner which lead to the subject-matter of claim 1 (statement of grounds of appeal, pages 3 to 7)

The board notes that present claim 1 no longer comprises the sole differentiating feature identified by the examining division (see point 9.3 above). It does however consider that claim 1 is distinguished from the disclosure of D1 by several other features, as explained in the following.

11. Document D1

11.1 D1 considers, on the one hand, a "computing system" that may be "a parallel processor (i.e. one or more parallel processors) computing system, [or] a multi-processor (i.e. one or more processors connected in different configurations) computing system" (paragraph [0027]), and, on the other hand, a "group of [hardware] accelerators" (paragraphs [0018], [0029]) that may be "attached" to the computing system through a plurality of "slots" of the computing system.

The board notes that here "attached" is understood as meaning physically coupled (see also paragraphs [0003]

and [0004]). Elsewhere in D1, "attachment" seems to refer more to the allocation of an accelerator to a slot to which the accelerator should be physically attached.

The board is not convinced that D3 - a parallel application of the applicant of D1 - is relevant for the interpretation of D1. First, because D1 makes no reference to D3. Second, because D1 refers to "streaming attachment" only in paragraph [0002], and this paragraph does not appear to match the rest of the disclosure of D1 (this paragraph seems to have been erroneously kept from D3, which was filed on the same date).

- 11.2 D1 discloses a method for automatically "guiding" the attachment of all accelerators to the computing system (paragraph [0018], figure 3).
  - 11.2.1 The board considers that this method is primarily meant to be implemented as a computer program determining an assignment of accelerators of the group of accelerators to slots and hubs available in the computing system (see figure 3).
  - 11.2.2 D1 is not clear as to which entity performs then the the physical attachment of the accelerated to the determined slot. It would appear - in line with the appellant's argument - that this task is to be carried out by a human user, "guided" by the computer program (see also paragraph [0004]).
  - 11.2.3 The examining division (decision, point 12, last paragraph) considered that D1 discloses also an automated attachment by means of a computer program in view of paragraphs [0007], [0038] and [0039]. The board

is not convinced that these generic statements about a possible implementation as a computer program disclose that a physical attachment of hardware accelerated to e.g. PCI express slots (see e.g. paragraph [0033]) would be automated. The board understands them rather as explained at point 11.2.1 above.

11.2.4 Still, the board tends to consider that this computer program of D1 amounts, using the terminology of claim 1, to a "resource manager" that "assigns" "boosters" (accelerators) to *slots* of the computer system, where the computer system comprising a plurality of "computation nodes" (processors). However, in D1, the boosters/accelerators are not assigned to computation nodes/processors.

11.3 An "initial placement" may be made on the basis of the consideration that "accelerators related to tasks associated with a particular processing unit may perform more efficiently if attached near the processing unit", as "latency and other performance issues may be reduced if compared to placement of accelerators relatively far from processing units performing tasks related to the accelerators" (paragraph [0019]).

Hence, the "initial placement of accelerators may include placement near components and/or processors related to a particular accelerator task", for example "a particular accelerator may be designed for video processing and particular components/processors of the computing system may perform video processing computations" (paragraph [0031]).

The board tends to consider that the initial placement based on this rule amounts, using the terminology of claim 1, to an "initial assignment" of boosters to



slots of the computer system based on a "predetermined assignment metric".

- 11.4 D1 discloses also a "guided re-attachment of accelerators based on run-time profiling" to "further increase system performance". As it "reflects actual usage of an accelerator resource within a computer system, it can better match an accelerator to a slot of the computer system" (paragraph [0033]).

Re-attachment "based on run-time profiling" does not by itself implies re-attachment *at run-time*, i.e. during computation of a computation task, only that it is based on run-time data.

However, paragraph [0038] discloses that "the methodologies illustrated in FIG. 3 may be iterated several times (or even constantly) during computing system execution for increased performance based on run-time profiling of accelerators and/or slots".

Hence, the board considers that D1 discloses thereby, using the terminology of claim 1, a "dynamic assignment" of boosters to slots of the computer system "during computation of said computation task" and based on an "update" of the "predetermined assignment metric".

- 11.5 However, the board agrees with the appellant that claim 1 differs from the disclosure of D1 at least in that

- all the assignments performed by the resource manager of claim 1 are *assignments of at least one booster to at least one of a plurality of computation nodes for computation of a part of the computation task* (in D1

the accelerators are to slots of a computer system);

- each of the computation nodes can communicate with each of the at least one booster over the communication infrastructure *without the necessity to involve a further computation node* (D1 is silent as to the communication from a slot of a computer system to the processors comprised in it).

12. D1 is concerned with a fundamentally different purpose than the invention. D1 is not concerned with how computation tasks are offloaded from processing units to accelerators. For instance, when it is recommended in D1 (it "may be prudent") to initially attach an accelerator designed for video processing to a slot that is in close proximity to processors that perform video processing computations (paragraph [0031]), this is a general consideration prior to any concrete need to offload of a particular computation task to an accelerator.
13. The board agrees with the appellant that - without further evidence - there would have been no motivation for the skilled person, starting from D1, to change the disclosed method for attaching accelerators *to slots* into a method for assigning accelerators (boosters) *to processing units (computation nodes)* for computation of a part of a computation task, in particular as there is no one-to-one correspondence between slots and processing units in D1.
14. D2 does not alter this conclusion. It concerns a "reverse acceleration" approach in which computation nodes are assigned to boosters, which was relevant for

the differentiating feature identified by the examining division, but this feature is not present in claim 1.

15. The board is thus not convinced by the examining division's inventive step reasoning starting from D1, and concludes that the subject-matter of present claim 1 involves an inventive step over D1 and D2.

*Remittal for further prosecution*

16. The board notes that in respect of European patent application No. 19179073.2, a family member of the present application in which similar subject-matter is claimed, a different document - D6 - was cited as "X" document in the European search report.

After the board had raised the appellant's attention to D6 in the preliminary opinion, the appellant commented on inventive step with respect to D6 in its letter dated 18 June 2024. However, as indicated in Article 12(2) RPBA, the primary object of the appeal proceedings is to review the decision under appeal in a judicial manner. In the present case, D6 was not considered in the examining proceedings.

17. Hence, and in view of the complexity of D6, the board concludes that there are "special reasons" within the meaning of Article 11 RPBA for a remittal to the department of first instance for further prosecution, Article 111(1) EPC.

*Reimbursement of the appeal fee*

18. The appellant argued that the decision under appeal was not reasoned within the meaning of Rule 111(2) EPC, for the following reasons:

(a) the argument that D1 does not disclose an assignment of a booster/accelerator to a computation node - which the appellant had raised before the examining division in its letter dated 15 October 2012 and in the oral proceedings held on 3 December 2019 - had not been replied to in the decision (statement of grounds of appeal, section B.1.i, first and last paragraphs);

(b) it was not explained why D1 was considered to disclose a "resource manager" in the passages of D1 cited in the decision for that feature (statement of grounds of appeal, section B.1.ii, first paragraph);

(c) the decision contained almost exactly the same wording as contained in the summons to oral proceedings and provided no reasoning as to why the arguments of the appellant regarding non-obviousness in its letter dated 30 October 2019 and in the oral proceedings of 3 December 2019 were not considered convincing (statement of grounds of appeal, section B.2, first paragraph).

This was considered to constitute a substantial procedural violation justifying the reimbursement of the appeal fee (statement of grounds of appeal, page 1, last paragraph).

19. Rule 103(1)(a) EPC provides that the appeal shall be reimbursed in full where the board of appeal deems and appeal to be allowable, if such reimbursement is equitable by reason of a substantial procedural violation.
20. The board considers that no substantial procedural violation has occurred in the first-instance proceedings.

- 20.1 The decision comprises a self-contained and sufficiently detailed reasoning of lack of inventive step, which may be erroneous at some points but that is a matter of error of judgement, not of procedural violation.
- 20.2 Argument (a) has been taken into account by the examining division and reasons were given as to why it was not found convincing in the communication dated 31 May 2018, point 2.1, second last paragraph. This paragraph should have been taken over in the final decision, but the board considers that this procedural mistake does not amount to a substantial procedural violation.
- 20.3 Argument (b) is implicitly addressed in the decision, point 12, last paragraph, explaining why the examining division considered that D1 discloses automatic attachment through execution of a computer program.
- 20.4 As regards (c), the board considers the obviousness reasoning in view of a combination of D1 and D2 provided in the decision, point 12, to be sufficiently detailed, in particular as regards the idea of "repeating" the acceleration of the processing by invoking an "accelerator", which may be a computation node as suggested by D2 (even if this argument may turn out not to be convincing).
21. The board presented these considerations in its preliminary opinion and the appellant did not react to them. The board has no reason to deviate from these considerations.
22. The board considers thus that the decision is sufficiently reasoned within the meaning of Rule 111(2)

EPC and that the appellant's right to be heard, Article 113(1) EPC, has been respected. The request for reimbursement of the appeal fee under Rule 103(1)(a) EPC is therefore refused.

## 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 for further prosecution.
3. The request for reimbursement of the appeal fee is refused.

The Registrar:

The Chairman:



L. Stridde

Martin Müller

Decision electronically authenticated