

Internal distribution code:

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

**Datasheet for the decision
of 14 March 2024**

Case Number: T 0928/19 - 3.5.06

Application Number: 07150065.6

Publication Number: 2073118

IPC: G06F9/50

Language of the proceedings: EN

Title of invention:

Load distribution in distributed database system

Applicant:

Nokia Solutions and Networks Oy

Headword:

Load distribution/NOKIA

Relevant legal provisions:

EPC Art. 56, 111(1)
RPBA 2020 Art. 13

Keyword:

Inventive step - non-obvious modification
Appeal decision - remittal to the department of first instance
(yes)

Decisions cited:

Catchword:



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 0928/19 - 3.5.06

D E C I S I O N
of Technical Board of Appeal 3.5.06
of 14 March 2024

Appellant:
(Applicant)

Nokia Solutions and Networks Oy
Karakaari 7
02610 Espoo (FI)

Representative:

Samson & Partner Patentanwälte mbB
Widenmayerstraße 6
80538 München (DE)

Decision under appeal:

**Decision of the Examining Division of the
European Patent Office posted on 30 October 2018
refusing European patent application No.
07150065.6 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman M. Müller
Members: T. Alecu
B. Müller

Summary of Facts and Submissions

- I. The appeal is against the decision of the Examining Division.
- II. The sole request underlying the decision under appeal was refused for
 - (a) lack of compliance with Article 123(2) EPC,
 - (b) lack of compliance with Article 84 EPC, clarity and conciseness, the latter in conjunction with Rule 43(2) (a) EPC, and
 - (c) lack of compliance with Article 56 EPC in view of D1: US 2005/114862 A1.
- III. With the statement of grounds of appeal the Appellant requested that the decision of the Examining Division be set aside and that a patent be granted on the basis of the main request or of a (sole) auxiliary request.
- IV. In a communication accompanying a summons to oral proceedings, the Board provided its provisional opinion that
 - (a) the requests were compliant with Article 123(2) EPC,
 - (b) both requests were not compliant with Article 84 EPC for lack of essential features,
 - (c) the main request was not compliant with Rule 43(2) (a) EPC, and further lacked clarity as it defined the scope of protection by reference to non-claimed entities,
 - (d) they lacked inventive step in view of D1, also with reference to the following document, which was introduced by the Board:
DA1: Kieran McCorry, "Understanding Front-End Servers", <https://www.itprotoday.com/email-and->

calendaring/understanding-front-end-servers,
19 October 2003.

- V. In reply to the summons the Appellant filed a new main request, based on the previous auxiliary request, and two new auxiliary requests, with amendments aiming to overcome both the Article 84 EPC and the Article 56 EPC objections, and argued why the new requests involved an inventive step in view of D1 and DA1.
- VI. The Board cancelled the oral proceedings and informed the Appellant that it intended to remit the case for further prosecution.
- VII. The independent claims 1 and 17 of the main request are as follows:

1. A method to distribute load after a network state change for a server among a plurality of active servers in a network, the plurality of active servers being front-end servers, wherein the front-end servers route requests for specific data items to matching back-end servers and keep a first lookup table and a partition of a second lookup table, wherein the first lookup table indicates which front-end server is responsible for which section of data stored in the back-end servers and wherein the second lookup table comprises pointers to data items in the backend servers, the method comprising:

*selecting an active server as master server;
determining (S1-S2; S21-S22), at the master server, a state change in the network;
calculating (S3; S24), at the master server, a new first lookup table;*

distributing (S4; S26), from the master server, the new first lookup table to all other active servers to update their partition of the second lookup table;

generating (S5; S26, S28, S29), at all active servers of the network, based on the new first lookup table, a modified server's partition of the second lookup table; and

routing, at all active servers of the network, incoming queries based on the new first lookup table and the modified partition of the second lookup table.

17. A system comprising a plurality of active servers (10; 201; 20i; 20n) constituting a network, the plurality of active servers being front-end servers, wherein the front-end servers route requests for specific data items to matching back-end servers, each active server comprising:

an interfacing unit (13; 23; 23n) configured to interface with the other active servers (10; 201; 20i; 20n);

a first lookup table (11; 21; 21n) indicating which front-end server is responsible for which section of data stored in the back-end servers;

a partition (12; 22; 22n) of a second lookup table, wherein the second lookup table comprises pointers to data items in the back-end servers; and

a processor (14; 24; 24n);

wherein an active server is selected as master server (10), wherein the processor (14) of the master server is configured to:

determine a state change in the network,

calculate a new first lookup table (11), distribute, via the interfacing unit, the new first lookup table (11) from the master server (10) to all other active servers to update their partition of the second lookup table;

generate a modified master server's partition of the second lookup table based on the new first lookup table; and

route incoming queries based on the new first lookup table and the modified partition of the second lookup table[,]

wherein the processor (24; 24n) of the other active servers is configured to:

receive, via said interfacing unit, the new first lookup table (21; 21n) from the master server (10);

generate a modified server's partition of the second lookup table based on the new first lookup table; and

route incoming queries based on the new first lookup table and the modified partition of the second lookup table.

VIII. The claims of the auxiliary requests are not pertinent to this decision.

Reasons for the Decision

The application

1. The application relates to a load distribution mechanism in a distributed database system using front-end servers and back-end servers, the front-end servers to route requests for specific data items to the back-end server storing the requested item. For routing, the front-end servers use a lookup table with entries of the form <key, value> pair, the key being the hash of the query search key, and the value the corresponding data address (description page 1, background section).
2. The application addresses in particular the problem of front-end server failures, or of server addition. These events are likely to result in less evenly distributed

system load and also, for the former case, losses of data (pages 2 and 3).

- 2.1 The application proposes to split the lookup table between the set of active servers according to a partitioning table - e.g. by splitting the hash values in equal ranges. One of the active servers acts as a master server which, upon a change in the number of servers, recalculates the partitioning table, i.e. the ranges of hash values, and sends it to all other active servers (pages 5, 6 and 9). The servers communicate with each other to transfer the needed lookup entries in accordance with the new partitioning table.

Main request: admittance

3. The new main request amends the previous first auxiliary request so as to overcome the Board's objections in its preliminary opinion (see below). Considering these circumstances, the Board admits this request (Article 13 RPBA 2020).

Objections under Articles 123(2) and 84 EPC

4. The Examining Division considered that the information that the "server acting as a master server" was also an "active server" was not unambiguously derivable from the application as filed. The Board agrees with the Appellant (statement of grounds of appeal, section D.I.2).c) that it is clear from the original application as a whole that the master server is selected as one of the active front-end servers (see e.g. the passages on pages 4, 2nd par., 20, 3rd par., or 36, 1st par., of the description as cited by the Appellant in the statement of grounds of appeal, section D.I.2).c)).

5. The Examining Division had raised other objections regarding clarity in the decision (see point 2.3 of the reasons, with reference to a previous communication). In response to these objections, the Appellant had amended the independent claims filed with statement of grounds of appeal. The Board agrees that these amendments overcome the clarity objections, for the reasons provided by the Appellant in its statement of grounds of appeal (D.I.4)). These reasons remain valid for the independent claims of the current main request.
6. The Examining Division also objected (decision, point 2.2 of the reasons) to the presence of multiple independent claims, one for the master server, and one for the active servers, in view of Rule 43 (2)(a) EPC. The current main request comprises only one independent claim per category.
7. The Board expressed the preliminary view that the claimed invention achieved its technical effect of redistributing workload between the active servers upon a network change only by collaboration between the master server and all other active servers. Separately claiming the partial methods to be carried out on either of them would therefore lack essential features, Article 84 EPC. An independent claim should thus comprise all the steps of re-calculating the first (partitioning) look-up table at the master server and generating the new modified partition of the second (data) lookup table at *each* (not only at *one*) of the active servers, as essential features of the invention.
- 7.1 The independent claims of the current main request comprise all these features, so also this objection is overcome.

Inventive step

8. Document D1 relates to distributed computing supporting *"adaptive workload partitioning"*, using *"an identifier key, to represent the individual components of the application [. L]oad distribution is then equivalent to modifying the partitioning of the identifier key space across a variable number of servers"* (paragraph 11).
- 8.1 If one of the servers is overloaded, it can decide to split the workload and delegate parts of it to other servers, creating child nodes. The work can be returned to the parent nodes when the child nodes have a low workload (paragraphs 37 to 40).
- 8.2 The determination of the servers that will be made responsible for part of the workload *"can be made centrally (e.g., using a table-based lookup), or [...] using a distributed lookup mechanism, such as a Distributed Hash Table (DHT)"* (paragraph 39).
- 8.3 The allocation of the identification keys to servers is maintained in a *"map function"*. *"This map function can either be centralized (every requestor must contact a specific computing node ...) or distributed (the mapping is stored in multiple nodes). One embodiment of this map function could be a static table, containing a one-to-one mapping from each unique key group to a particular server"* (paragraph 79).
- 8.4 D1 proposes further, and appears to favour, a decentralized DHT mechanism, which associates a hash value of the identifier key with a server through a DHT mapping function (paragraph 61 and further). In this approach the server responsible to handle a client request is identified on the basis of an iterative approach,

using an educated guess and the properties of the hashing algorithm (see paragraphs 67 to 78).

9. The Board understands that the Examining Division, in its analysis (decision, point 2.4.3 of the reasons), mapped the following entities of the claimed invention to those of D1:
 - (a) The claimed front-end servers to the servers of D1;
 - (b) The claimed back-end servers were not visibly taken into account in the decision;
 - (c) the first (partitioning) lookup table was read on the *"collection of all identifier keys, which defines an identifier key space"*;
 - (d) the second lookup table was read on the *"partitioning of the identifier key space across a variable number of servers"*; and
 - (e) the master server was acknowledged not to be disclosed in D1.
10. The Examining Division considered that the only difference was the presence of a master server, which recalculated and re-distributed the first look-up table, and deemed this to be part of normal design options (still point 2.4.3 of the reasons).
11. The Board is of the opinion that the Examining Division erred in its analysis.
 - 11.1 The *"collection of all identifier keys"* is not a lookup table, even less a partitioning table, so it cannot be equated with the first (or second) look-up tables claimed. At best, there is only one look-up table, i.e. the *"partitioning of the identifier key space across a variable number of servers"* which the Examining Division identified as the second lookup table.

- 11.2 The Appellant argued (statement of grounds of appeal, section D.I.5).a)) that D1 only taught a decentralized DHT approach and pointed out that in this approach there was no centralized partitioning table. But D1 also teaches a centralized approach. In this approach a centralized partitioning table is present - the table of paragraph 79 which is stored on only one node.
- 11.3 The computing nodes of D1 are all equivalent. Front-end and back-end servers are not distinguished. All computing nodes share the workload, which may be redistributed dynamically. While, formally, one may identify these nodes with either the front-end or the back-end servers as claimed, the set-up is in fact different from that claimed.
- 11.4 The redistribution is performed locally, by a server detecting a need for such distribution. In the Board's view, that server acts like a master server as claimed, even if only temporarily, as it calculates and updates the partitioning of the workload. The Appellant argued that that server was not selected, but rather selected itself. However, since no selection mechanism is defined by the claim, the Board does not see this as a difference.
12. To summarise, D1 discloses two main alternative approaches, a centralised and a decentralised approach, both using a set of active servers sharing the workload. In the decentralised approach there is no centralised look up table, whereas in the centralised one there is one, kept on one server. A master server may be said to be present, temporarily, during the redistribution of the workload.

13. An essential difference between the claimed invention and D1, whichever the approach considered, is the use of a, otherwise well-known (see DA1), front-end/back-end server architecture, where the front-end servers route the "work" to the appropriate servers. Such architecture is normally employed to ease and simplify routing to the servers doing the work, i.e. the back-end servers.
 - 13.1 The decentralised approach in D1 has its own, specifically developed, routing mechanism (see 8.4 above). The Board does not see that front-end servers fit into this in any obvious way.
 - 13.2 In the centralised approach, the skilled person might find problematic that the map function, i.e. the partitioning table as claimed, is stored only on one server, which might become overloaded. It might then appear obvious to add a layer of front-end servers, notwithstanding the fact that D1 itself has another solution to this problem, namely the aforementioned decentralised approach. But even if front-end servers were added, the skilled person would apply the teaching of D1 concerning the work redistribution, in particular that concerning table repartitioning by a temporary master server, to the active servers according to D1, which are effectively back-end servers.
 - 13.3 So the Board does not see that the skilled person would arrive at the claimed invention by solving the problem of routing work to the active servers of D1.
14. In its preliminary opinion, the Board considered instead the possibility that the "work" that needed redistribution amongst the servers of D1 was the routing work, in which case the skilled person might apply

the load redistribution approach of D1 to the front-end servers of the assumed front-end/back-end server architecture. The nature of the work would then require that the servers maintain second look-up tables.

- 14.1 Starting from the centralized approach, it might then be considered an obvious option to distribute the partitioning table to all front-end servers.
- 15. The Appellant argued that it was not apparent *"why the skilled person would change the architecture"* in this way. And that even if so, *"why the skilled person would - without knowledge of the solution claimed - consider distribution of a new first lookup table to all servers since the teaching of D1 explicitly only focuses on a local load distribution"*.
- 16. The Board acknowledges that the skilled person starting from D1, in order to arrive at the claimed invention in the manner considered would have to make a few crucial decisions.
 - 16.1 First, the skilled person would have to decide to apply the teaching of D1 to front-end servers. Notably however, the function of the active servers of D1 is to process the workload itself, the servers thus appearing rather as back-end servers.
 - 16.2 Second, if there were problems with having one table on one server only, the skilled person would have to decide not to use the local decentralised approach which is expressly disclosed in D1 and which the skilled person would be inclined to prefer.
 - 16.3 And third, even if the skilled person were to deviate from the express disclosure of D1 and not use the local

decentralised approach, the skilled person would still need to decide to distribute the partitioning table to all front-end servers.

17. Reconsidering the facts of the case, and in particular the teachings of D1, in view of the above, the Board follows the Appellant and comes to the conclusion that the skilled person could have arrived at the invention as laid out above, but that it cannot be asserted that he or she would have.

18. The Board therefore finds the claimed invention not to be obvious when starting from D1.

Remittal

19. The Board notes that document D1, relied upon by the Examining Division, is a document mainly directed to a localised load redistribution, which is conceptually the opposite of the claimed invention. Also, and perhaps more importantly, although front-end/back-end server architectures were well known in the art at the filing date of the application, the decision does not refer to any such document. More relevant prior art may therefore exist.

19.1 Under these circumstances the Board decides to remit the case to the Examining Division for further prosecution (Article 111(1) EPC), possibly including an additional search.

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.

The Registrar:

The Chairman:



L. Stridde

Martin Müller

Decision electronically authenticated