

Internal distribution code:

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

**Datasheet for the decision
of 22 September 2022**

Case Number: T 0873/19 - 3.5.07

Application Number: 11155687.4

Publication Number: 2369507

IPC: G06F17/30

Language of the proceedings: EN

Title of invention:

Relational database for storing business objects and method
for operating the same

Applicant:

Hasso-Plattner-Institut für
Softwaresystemtechnik GmbH

Headword:

Relational database for business objects/HASSO-PLATTNER-
INSTITUT

Relevant legal provisions:

EPC Art. 56

Keyword:

Inventive step - sole request (no)
Remittal to the department of first instance (no)

Decisions cited:

T 1965/11



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 0873/19 - 3.5.07

D E C I S I O N
of Technical Board of Appeal 3.5.07
of 22 September 2022

Appellant: Hasso-Plattner-Institut für
(Applicant) Softwaresystemtechnik GmbH
Prof.-Dr.-Helmert-Str. 2-3
14482 Potsdam (DE)

Representative: Molnia, David
Df-mp Dörries Frank-Molnia & Pohlman
Patentanwälte Rechtsanwälte PartG mbB
Theatinerstrasse 16
80333 München (DE)

Decision under appeal: **Decision of the Examining Division of the
European Patent Office posted on 23 October 2018
refusing European patent application
No. 11155687.4 pursuant to Article 97(2) EPC**

Composition of the Board:

Chair J. Geschwind
Members: C. Barel-Faucheux
P. San-Bento Furtado

Summary of Facts and Submissions

- I. The appeal lies from the decision of the examining division to refuse European patent application No. 11155687.4, which has been published as EP 2369507.
- II. The decision cited *inter alia* the following documents:
D1: EP 1 217 540 A1, published on 26 June 2002
D2: "Column-oriented DBMS", Wikipedia, 25 January 2010
- III. During oral proceedings held in the absence of the appellant, the examining division decided that neither the subject-matter of the claims of the main request nor that of the first to third auxiliary requests involved an inventive step (Article 56 EPC).
- IV. With its statement of grounds of appeal, the appellant maintained the main request and the first to third auxiliary requests considered in the contested decision. It requested that the decision under appeal be set aside and that a patent be granted on the basis of one of those requests.
- V. In a communication annexed to a summons to oral proceedings, the board stated that it appeared that neither claim 1 of the main request nor that of the first to third auxiliary requests was inventive (Article 56 EPC) with regard to the disclosure of document D1 or, alternatively, when considering as a starting point the prior art acknowledged in the background section of the description of the application.

- VI. In a letter of reply the appellant submitted further arguments in favour of inventive step.
- VII. During the oral proceedings before the board, the appellant withdrew the main request and first and third auxiliary requests and requested that the decision under appeal be set aside and that a patent be granted on the basis of the second auxiliary request considered in the decision under appeal, as the sole request.
- VIII. Claim 1 of the sole request reads as follows:

"A method for operating a relational database for storing a business object instance, wherein the business object instance has at least one attribute, wherein the relational database comprises multiple tables, wherein one table of the multiple tables is a root table for storing root records, the root table comprising a vector column, and wherein the method comprises the following steps:

(A) storing at least one attribute of the business object instance and at least one reference, e.g., [sic] a reference key, to at least one attribute of the business object instance in a subset of the multiple tables (308);

(B) determining the tables of the multiple tables that contain:

at least one attribute of the business object instance, and/or

at least one reference to at least one attribute of the business object instance (310);

(C) generating a vector with a representation suitable for identifying any possible subset of the multiple tables, the vector identifying the tables determined in step (B) (312); and

(D) storing the generated vector in the vector column of the root table, namely in a root record that, according to a business object data tree model, represents a root node of the business object instance (314);

(E) receiving a query for the stored business object instance from a client (316);

(F) identifying the root record in the root table (318);

(G) reading out the vector from the root record (320);

(H) identifying the tables determined in step (B) from the read out vector (322); and

(I) reading out all records corresponding to the business object instance in parallel (324)."

Reasons for the Decision

The application

1. The application relates to a relational database for storing a business object instance (BOI) and a method for operating the relational database (description as originally filed, page 1, lines 8 to 9).

Sole request - inventive step

2. Document D1 relates to relational database management systems (RDBMS).
 - 2.1 Document D1 discloses a relational database for storing a (business) object instance (Figure 1: "Clients" instance; Figure 2: "Policies" instance; Figure 3: "Accidents" instance; all three BOI instance tables also together forming a BOI {"Clients", "Policies",

"Accidents"}), wherein the (business) object instance has at least one attribute (Figure 1: attributes "Name", "Birth Year" and "Gender"; Figure 2: attributes "Type", "Date"; Figure 3: attributes "Date" and "Amount"), wherein the relational database comprises multiple tables (the tables of Figures 1 to 3). The "Accident Table" is a root table for storing root records and comprises columns corresponding to data attributes of the records, for example the "Date" or "Amount" columns (paragraphs [0018], [0021] to [0030], see in particular paragraph [0026] of D1: "A root table is defined as a data table for which no other data table has links pointing to its rows, such as the accident table of figure 3"; Figures 1 to 5).

- 2.2 In document D1, in the table of Figure 2, at least one attribute of the BOI "Policies", for example the attribute "Type", and the reference key (or "foreign key") "Client Link" are stored. In the table of Figure 3, at least one attribute of the BOI "Accidents", for example the attribute "Amount", and the reference key "Policy Link" are stored. Therefore, document D1 discloses step (A) of storing at least one attribute and at least one reference in the form of a reference key (paragraphs [0023] and [0024]).

Consequently, document D1 discloses the introductory part of claim 1 (except for the column being a "vector column") together with step (A) of claim 1.

Document D1 also discloses a business object data tree model having a root record (the "Accident Table") representing a root node of the BOI {"Clients", "Policies", "Accidents"} (see feature D of claim 1 and Figure 4 of document D1).

3. Therefore, the distinguishing features of claim 1 over the disclosure of document D1 are steps (B) to (I) and that the column is a "vector column".
4. Steps (B) to (D) are illustrated in Figure 2 of the application: the vector 124 comprises multiple bits, wherein each bit of the multiple bits refers to one table of the multiple tables (in the example given "eight" tables). For example, the bit vector 124 comprises eight bits. If step (B) determines that only four of the eight tables are actually used to store records for the BOI, the bit vector generated by step (C) will be " 1 0 1 1 0 0 1 0 " indicating that each of the tables 108, 136, 148, and 176 comprises at least one record corresponding to the BOI and that none of the tables 128, 164, 170, and 188 comprises at least one record corresponding to the BOI. The bit vector will be stored, in step (D), in the root node of the root table (description of the application as published, paragraph [0041]).
5. The appellant stated that the examining division had correctly identified the effect obtained by the distinguishing features as being a reduction in the number of multiple tables (in comparison to all the tables) needed to be processed when retrieving the attributes for a given BOI (statement of grounds of appeal, page 3, last paragraph).
6. The appellant also referred to decision T 1965/11 and argued that the board regarded the effect produced by a very similar subject-matter, as a technical effect. Even though the subject-matter concerned in the case T 1965/11 was not on the physical implementation level of the relational database since claim 1 did not concern aspects of where and how data were physically

stored in the database, the board had nevertheless considered the effect achieved, i.e. the cost-based optimisation of a query in a relational database system, to be technical (statement of grounds of appeal, pages 4 and 5).

- 6.1 More particularly, the board agreed that, in the technical context of query optimisation in relational database systems, the teaching of the invention related to the use of query graphs for the matching in order to substantially reduce the complexity of extracting operator trees which encoded a specific join order was based on further technical considerations and solved the problem of providing a technically feasible implementation, in particular one that achieved an acceptable time complexity for query optimisation in relational database systems (see T 1965/11, Reasons 5.3).
- 6.2 The appellant argued that claim 1 of the sole request considered in T 1965/11 only recited logical information elements and a logical location of elements. In fact, it was drafted at the very same level as the claimed invention in the case at hand. It concerned a "relational database" having "tables" that needed to be processed by "binary joins". The "operator tree" did not specify a physical location or data structure either. Rather, it was comparable to the "vector" as claimed in the present case because it was a logical data responsible for encoding what joins needed to be performed in which order (statement of grounds of appeal, page 5).
- 6.3 The examining division argued that "materialised views" of claim 1 of the sole request in case T 1965/11 were database views that were physically stored and which

had to be technically managed by the DBMS. Furthermore, cost-based query optimisation was an optimisation with respect to computer resources such as CPU, main memory or hard disk (section 12.2.2 of the decision). The board tends to agree with this statement.

7. The board is nevertheless of the opinion that reducing the number of multiple tables needed to be processed when retrieving the attributes for a given BOI is a technical effect (see section 5. above).
8. The appellant formulated the technical effect as "optimising a database query" that achieved "query execution with less computer resources". This technical effect was achieved by storing a vector in the database, the vector indicating those tables that needed to be joined and those that did not need to be joined (statement of grounds of appeal, page 6).
9. In the method of claim 1, the vector generated and stored in steps (B) to (D) is further used in steps (G) and (H). Thus, the technical effect of reducing the number of multiple tables needed to be processed when retrieving the attributes for a given BOI is achieved (see section 5. above).
10. During the oral proceedings, the appellant argued that the acknowledged (or conventional) prior art of document D1 was not limited to the description of Figures 1 to 4, but also encompassed the description of Figures 5 to 8. It also argued, in reply to the board's communication, that the board's reasoning was entirely new and based on hindsight. According to the appellant, the invention was also quite different from the disclosure of document D1 since the method of claim 1

stored the generated vector during storage time while using it during retrieval.

10.1 The board notes that Figures 1 to 3 of document D1 show an example of data structure as typically used in a conventional relational database system. Figure 4 is a diagram representing a data table tree in the example of Figures 1 to 3. Figures 5 to 7 are diagrams showing respective data graphs constructed with the tree of Figure 4 and the data of Figures 1 to 3. Figure 8 is a "flat file" representation of the data tables of Figures 1 to 3 (paragraphs [0015] to [0040]).

10.2 Figures 1 to 3 illustrate a collection of data which can be stored in a computer memory coupled with a processor arranged for running relational database management programs and show a conventional type of data organisation in a database system. In the conventional prior art of document D1, paths are defined, in the data table tree, from the root table to the leaf tables. This is because each path from a root table to a leaf table is defined by a link column of the root table pointing to the leaf table, or by a succession of link columns via one or several intermediate tables (Figure 4; paragraphs [0016], [0017] and [0028]).

10.3 The appellant has argued that figures 1 to 4 of document D1 should not be considered in isolation and that also figures 5 to 7 might be relevant. However, the board is of the opinion that all these figures refer to a conventional relational database system.

The board also notes that the description of the invention of document D1 making use of the flat file concept starts after the description of this

conventional relational database system reflected in figures 1 to 7 (see, for example, paragraphs [0015] and [0041]).

11. Therefore, starting from the conventional prior art of document D1, the objective technical problem to be solved was formulated by the appellant during the oral proceedings as "how to reduce the number of joins". The board rather considers it as "how to improve the efficiency of query execution for a stored BOI in a relational database management system".
12. Since document D1 discloses the beginning part of claim 1 together with step A of claim 1 (except for the column being a "vector column"), if one considers a (business) object instance, at least one attribute of the BOI and at least one reference to at least one attribute of the BOI has been stored in D1 in a subset of the multiple tables. When seeking to improve the efficiency of query execution for a stored BOI in a relational database management system, the skilled person would think of determining this subset of the multiple tables [for example Table 1, Table j, Table k] to derive information about the stored BOI. He would thereby directly arrive at step B.
13. It would then be obvious for the skilled person to use a bit vector of the length of the (total) number of multiple tables [having a "1" at the first, jth and kth position and "0"s at the remaining positions in the example of point 12. above] for representing the determined subset of the multiple tables. The skilled person would then arrive at step C.
14. Since, in the business object data tree model of document D1, paths are defined in the data table tree

from the root table to the leaf tables (Figure 4; paragraph [0028]) the data table tree will be read starting from this root table. It would thus also be obvious to store this information about the stored BOI in the root record in the form of a vector column. The skilled person would then arrive at step D.

15. Since a relational database is conventionally used to receive a query from a client for a stored BOI, when receiving this query in step E, the skilled person would use the vector that was stored in step D. It will thus identify it in the root record and read it out. The skilled person would then arrive at steps F and G.
16. This read-out vector would obviously be used by the skilled person to identify the subset of the multiple tables determined in step B. The skilled person would then arrive at step H.
17. After having identified the subset of the multiple tables storing at least one attribute of the BOI and at least one reference to at least one attribute of the BOI, it would also be obvious to read out all records in parallel, for example to increase the processing speed. The skilled person would then arrive at step I.
18. Thus, claim 1 of the sole request is not inventive (Article 56 EPC).

Request for remittal

19. During the oral proceedings, the appellant requested remittal "to further establish firm grounds to the common general knowledge" in case the board intended to refuse the application as a whole.

20. The board does not see any special reasons for a remittal (Article 11 RPBA). In particular, in the light of the above inventive step analysis, the board does not see a need to establish "firm grounds" as to what is to be seen as common general knowledge.

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar:

The Chair:



S. Lichtenvort

J. Geschwind

Decision electronically authenticated