

Internal distribution code:

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

**Datasheet for the decision
of 24 March 2017**

Case Number: T 1965/11 - 3.5.07

Application Number: 01123065.3

Publication Number: 1193618

IPC: G06F17/30

Language of the proceedings: EN

Title of invention:

Cost-based materialised view selection for query optimisation

Applicant:

Microsoft Technology Licensing, LLC

Headword:

Cost-based materialised view selection/MICROSOFT TECHNOLOGY
LICENSING

Relevant legal provisions:

EPC Art. 56, 123(2)

Keyword:

Added subject-matter - (no)
Inventive step - (yes)

Decisions cited:

G 0003/08, T 1569/05, T 1003/09

Catchword:



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 1965/11 - 3.5.07

D E C I S I O N
of Technical Board of Appeal 3.5.07
of 24 March 2017

Appellant: Microsoft Technology Licensing, LLC
(Applicant) One Microsoft Way
Redmond, WA 98052 (US)

Representative: Grünecker Patent- und Rechtsanwälte
PartG mbB
Leopoldstraße 4
80802 München (DE)

Decision under appeal: Decision of the Examining Division of the
European Patent Office posted on 20 April 2011
refusing European patent application No.
01123065.3 pursuant to Article 97(2) EPC.

Composition of the Board:

Chairman R. Moufang
Members: M. Jaedicke
P. San-Bento Furtado

Summary of Facts and Submissions

- I. The applicant (appellant), which at the time was Microsoft Corporation, appealed against the decision of the Examining Division refusing European patent application No. 01123065.3.
- II. In the course of the appeal proceedings, the application was transferred to Microsoft Technology Licensing, LLC, which thereby obtained the status of appellant.
- III. The Examining Division decided that the subject-matter of claims 1 and 9 of the main request contained added subject-matter. Moreover, it refused to admit the first and second auxiliary requests using its discretion under Rule 137(3) EPC, stating that the first and second auxiliary requests appeared to comprise the same violations of Article 123(2) EPC as the main request.
- IV. In the written proceedings the Examining Division cited the following prior-art documents:
 - D1: WO-A-99/50762, published on 7 October 1999;
 - D2: US-A-6,026,390, published on 15 February 2000;
 - D3: US-A-5,897,632, published on 27 April 1999;
 - D4: Gupta, Ashish et al.: "Aggregate-Query Processing in Data Warehousing Environments", Proceedings of 21st International Conference on Very Large Data Bases, 11 to 15 September 1995, Zurich, Switzerland, Morgan Kaufmann, ISBN 1-55860-379-4, pages 358 to 369;

- D5: WO-A-99/50732, published on 7 October 1999;
- D6: Chang, Jae-young et al.: "An Extended Query Reformulation Technique Using Materialized Views", Proceedings of Ninth International Workshop on Database and Expert Systems Applications, 1998, Vienna, Austria, 26 to 28 August 1998, pages 931 to 936;
- D7: Zhang, Chuan et al.: "Evolving Materialized Views in Data Warehouse", Proceedings of IEEE Congress on Evolutionary Computation, Washington D.C., USA, vol. 2, 6 July 1999, pages 823 to 829.

V. In the statement of grounds of appeal, the appellant requested that the decision be set aside. It also asked the Board to decide that the claims of one of the main and the two auxiliary requests, considered in the contested decision and resubmitted with the grounds of appeal, were allowable.

VI. In a communication under Article 15(1) RPBA accompanying a summons to oral proceedings, the Board expressed its provisional opinion that the reasons for the refusal were not convincing. However, the Board raised several objections under Articles 123(2) and 84 EPC.

The Board also drew the appellant's attention to the following document:

- D8: Chaudhuri, S.: "An Overview of Query Optimization in Relational Systems", Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, pages 34 to 43,

Seattle, Washington State, USA, 1 to 4 June 1998.

VII. With its reply to the Board's communication, the appellant submitted third and fourth auxiliary requests and new description pages 3, 3a, and 3b to replace description page 3 as on file.

VIII. In the course of the oral proceedings the appellant replaced its previous requests with a sole request corresponding to the fourth auxiliary request. At the end of the oral proceedings, the chairman pronounced the Board's decision.

IX. The appellant's final request is that the decision under appeal be set aside and that a patent be granted on the basis of the following application documents:

- description pages 1, 2, and 4 to 13 as originally filed, and pages 3, 3a and 3b as filed with the letter dated 24 February 2017;
- drawings of Figures 1 to 8 as originally filed;
- claims 1 to 3 of the sole request maintained at the oral proceedings before the Board, corresponding to the fourth auxiliary request filed with letter dated 24 February 2017.

X. Claim 1 of the sole request reads as follows:

"A method for a computer-implemented query optimizer for selecting an execution plan for use in execution of a relational database query, the method comprising:

generating, by a cost-based query optimizer, a table (300) of alternatives, the table (300) of alternatives comprising several groups, one group having a root entry representing the database query and additional entries representing alternative possibilities for executing the database query and the other groups

having root entries representing sub-expressions of the database query and additional entries representing alternative possibilities for executing the respective sub-expression of the database query;

selecting candidate views for the query from a number of materialized views by using information about what database tables are referenced in the query and whether or not the query contains aggregations;

for each root entry,

extracting an operator tree for the root entry, collapsing binary joins contained in the operator tree to obtain a query graph for the root entry, the query graph listing all underlying tables along with predicates that are applied on them,

matching the query graph for the root entry and the candidate views, and

if a match is found, extending the table (300) of alternatives with the corresponding candidate view by extending the group of the root entry with the corresponding candidate view; and

using the cost based query optimizer to select an execution plan based on the extended table (300) of alternatives."

Claim 2 reads as follows:

"A computer readable medium having instructions for causing a computer to perform a method according to claim 1."

Claim 3 reads as follows:

"A computer-implemented query optimizer configured to perform a method according to claim 1."

XI. The appellant's arguments relevant to the decision are discussed in detail below.

Reasons for the Decision

1. The appeal complies with the provisions referred to in Rule 101 EPC and is therefore admissible.

The invention

2. The application relates to the optimisation of relational database queries in the presence of materialised views. According to page 1, last paragraph, of the description, the basic idea of materialised views is to store the result of a query and to use this stored result to answer similar later queries. The invention addresses the known view utilisation problem: given a user query written over base tables, as well as a collection of materialised views, which materialised views can be used to answer the query (description, page 2, first paragraph)?

According to the technical background section on page 2 of the application, prior attempts to determine which views should be used treated the problem in isolation, handled limited scenarios, and often assumed a "global" structure that covered the whole query. There was a need to deal with arbitrary queries, and to integrate view utilisation within the actual architecture of query optimisers. There was a further need to decide whether a view that could be used to answer the query, should actually be used to execute that query. Constructing a "global" structure for the user query, for the purpose of view matching, was incompatible with common optimiser architecture and was sometimes

infeasible.

According to page 7, lines 14 to 16, query optimisers are normally structured such that there is an initial simplification stage, followed by exploration of alternatives and cost-based selection of an execution plan as illustrated in Figure 2 of the application. Considering materialised views during query simplification is inadequate, because only a single solution can be generated, and there is no detailed cost information (page 7, lines 31 to 33). Hence, the invention proposes to extend the table of alternatives, which is generated by the query optimiser at the exploration stage and which compactly encodes for each sub-expression of a query the various possibilities for its execution (page 3, lines 9 to 13, and page 10, lines 1 to 21). In other words, materialised views are detected and substituted during exploration of the various query execution possibilities and added to the table of alternatives.

Admission of request

3. Since the current request was a legitimate response to the preliminary opinion of the Board, does not introduce additional complexity into the proceedings and could be dealt with without adjournment of the oral proceedings, the Board admitted it into the appeal proceedings.

Clarity and added subject-matter

4. The appealed decision found that the following features of the method of claim 1 of the then pending main request infringed the requirements of Article 123(2)

EPC:

- (a) "selecting candidate views for the query from a number of materialized views";
- (b) "extracting, for each candidate view, an operator tree from a view definition of the candidate view".

4.1 The appellant argued that present claim 1 was based on claims 1 to 3, 8 and 17 as originally filed and page 1, lines 3 to 5 and 18 to 26, page 3, lines 9 to 14 and 21 to 24, page 4, line 32, to page 5, line 1, page 8, lines 12 to 19 and 21 to 31, page 9, lines 8 to 16, page 10, lines 14 to 19, and page 11, lines 2 and 3, of the original description and original Figures 3 and 5.

4.2 The method of present claim 1 no longer contains feature (b). Feature (a) is still present, but the method of present claim 1 adds at the end of that feature the text "by using information about what database tables are referenced in the query and whether or not the query contains aggregations".

4.3 The appellant argued in the statement of grounds of appeal that feature (a) was supported by page 9, lines 8 to 14. In particular, page 9, lines 8 to 10, stated that for a particular query there was in general a number of "candidate views", as not all of the materialised views stored in conjunction with previous database queries were relevant to a particular query submitted to the database.

4.4 The Board agrees with the appellant that a database usually has a limited set of materialised views (for example due to the high costs associated with maintaining views). A particular database query will

normally refer only to a subset of all base tables and hence some materialised views will not be relevant for this query. Other views might not be relevant due to the use of aggregation. The appellant points correctly to page 9, lines 10 to 13, in support.

Hence, the Board does not share the view of the Examining Division that the cited passage on page 9 merely describes "the process of narrowing down a set of already present candidate views" and that the "origin of this initial set of candidate views ... remains undefined". The appellant suggested that the Examining Division had misinterpreted the following sentence (page 9, lines 13 to 14): "This provides the ability to narrow down the set of candidate views." The Board does not share the alleged interpretation by the Examining Division, namely that the first word ("this") of the cited sentence referred to the process of identifying views as not relevant for a particular query. The sentence simply means that the initial set of (all) materialised views is reduced to a subset of the candidate views by eliminating irrelevant views as described on page 9, lines 11 to 13.

- 4.5 Hence, the Board accepts the basis for the amendments to claim 1. Present claims 2 and 3 are directed to a computer-readable medium and to a computer-implemented query optimiser, both defined by reference to the method of claim 1.
- 4.6 Moreover, the amendments to the claims overcome the Board's objections under Article 84 EPC as detailed in its communication to the appellant.

4.7 Hence, the Board is satisfied that the claims comply with the requirements of Articles 84 and 123(2) EPC.

Novelty and inventive step

5. Inventive step was not examined in detail during the proceedings before the first instance. Nevertheless, in view of the age of the application, the Board considers it appropriate to deal with this issue itself under Article 111(1) EPC.

5.1 According to decision T 1569/05 of 26 June 2008, reasons 3.6, retrieving data in a computer database is normally considered to have technical character. While the method of claim 1 does not include the actual data retrieval, the Board considers that the cost-based optimisation of a query in a relational database system has normally technical character (see T 1003/09 of 29 April 2015, reasons 13.3 and 13.5). Such cost-based query optimisation searches for low-cost query execution plans using a cost estimate for the computer resources (such as CPU, main memory or hard disk) needed to execute a query plan (see D8, section 2, for technical background). Hence, this cost-based approach involves further technical considerations (see opinion G 3/08, "Programs for computers", OJ EPO 2011, 10, reasons 13.5) relating to the internal functioning of the computer system.

5.2 The Board summarises below the relevance of the prior art on file. Seven documents, D1 to D7, were cited in the European search report.

Document D1 describes that a query rewriter intercepts and attempts to rewrite user database queries using aggregate tables (materialised aggregate views; see for

example page 12, line 18 to page 13, line 5; page 15, line 18 to page 16, line 12) in order to improve query performance. The query rewriter uses a cost-based algorithm to choose among potential rewrites and is detailed only in document D5 (D1 refers on page 24, lines 9 to 12, to U.S. Application Serial No. 09/049,784, the priority document of D5).

Document D5 explains the query rewriting approach in Figures 6 to 9 and page 19, line 23 to page 25, line 14. A query is submitted to the rewrite system (Figure 6, reference sign 700), where it is initially screened to determine whether it should be rewritten (Figure 7). If so, then a block-wise rewrite process is performed, where a block represents part of an SQL query (Figure 8). Figure 9 describes that for each block, views from a list of precomputed (materialised) views (PCV_List) are processed. For each view, it is checked whether a rewrite of the current block using this view would lower the costs for the current block compared to the estimated cost for the best available rewrite found so far. Hence, alternatives are generated for single blocks and the cost-based selection is limited to alternative views for a single block.

Document D2 concerns the incremental maintenance of a first materialised view of data in a database, by means of an additional materialised view, if this reduces the maintenance costs.

Document D3 discloses a method of evaluating a query using materialised views, which starts by semantically analysing a first materialised view to determine whether it is usable in evaluating the query. If the view is usable, then the method rewrites the query into an equivalent query using the materialised view.

Subsequently the process is iterated for the other available materialised views. Rewriting of the query is not integrated into the plan exploration phase and is not cost-based.

Document D4 concerns methods for rewriting queries to consider materialised views. Rewriting is usually applied before the plan exploration phase.

Document D6 (see abstract and sections 1 and 3) concerns query reformulations to use materialised views, and in particular proposes rules to identify potentially useful views, but does not address the problem of generating alternative plans during the exploration phase of a query optimiser.

Document D7 addresses the following problem: based on multiple queries, select a set of views to be materialised in order that the total query and maintenance cost is minimised.

Document D8, introduced into the proceedings by the Board, provides an overview of known query optimisation techniques in relational database systems. Section 7.3 describes the optimisation problem for materialised views. It explains that the steps of enumerating and generating equivalent expressions in the presence of materialised views may overlap, but does not explain in detail how the exploration phase or the table of alternatives is implemented.

- 5.3 None of these prior-art documents addresses the problem of extending a table of alternatives generated by the query optimiser by adding further alternatives using materialised views. The invention makes it possible to find low-cost query execution plans that make use of

the available materialised views in order to improve query performance (see page 2, third paragraph). Moreover, in order to explore the search space for such low-cost query execution plans, it proposes integrating the materialised views into the table of alternatives during the plan exploration stage. For this integration, it is necessary to match query plans with materialised views in order to identify useful plan alternatives for such views. The invention teaches using query graphs for the matching in order to substantially reduce the complexity of extracting operator trees which encode a specific join order. In the technical context of query optimisation in relational database systems, this teaching is based on further technical considerations and solves the problem of providing a technically feasible implementation, in particular one that achieves an acceptable time complexity for query optimisation in relational database systems. In the oral proceedings the appellant argued along these lines in favour of inventive step, and the Board agrees.

- 5.4 The Board therefore concludes that the subject-matter of independent claims 1 to 3 involves an inventive step (Articles 52(1) and 56 EPC) over the available prior art.

Conclusion

6. Since the appellant's request complies with the provisions of the EPC, the appeal is to be allowed.

Order

For these reasons it is decided that:

1. The decision under appeal is set aside.
2. The case is remitted to the department of first instance with the order to grant a patent on the basis of the following application documents:
 - description pages 1, 2, and 4 to 13 as originally filed, and pages 3, 3a and 3b as filed with the letter dated 24 February 2017;
 - drawings of Figures 1 to 8 as originally filed;
 - claims 1 to 3 of the sole request maintained at the oral proceedings before the Board, corresponding to the fourth auxiliary request filed with the letter dated 24 February 2017.

The Registrar:

The Chairman:



I. Aperribay

R. Moufang

Decision electronically authenticated