



(22) Date de dépôt/Filing Date: 2004/05/27

(41) Mise à la disp. pub./Open to Public Insp.: 2004/11/27

(30) Priorité/Priority: 2003/05/27 (2,429,910) CA

(51) Cl.Int.⁷/Int.Cl.⁷ G06F 17/30, H04L 12/16


(71) Demandeur/Applicant:
COGNOS INCORPORATED, CA

(72) Inventeur/Inventor:
STYLES, MICHAEL E., CA

(74) Agent: GOWLING LAFLEUR HENDERSON LLP

(54) Titre : SYSTÈME ET METHODE DE TRANSFORMATION D'AGREGATS CLIENT/SERVEUR

(54) Title: SYSTEM AND METHOD OF CLIENT/SERVER AGGREGATE TRANSFORMATION



(57) Abrégé/Abstract:

A system for transforming client/server aggregates is provided. The system comprises a client/server analysis component for analyzing one or more queries that are not supported by a target database system, and a client/server transformation component for transforming the one or more queries into semantically equivalent queries that are supported by the target database system.

Abstract

A system for transforming client/server aggregates is provided. The system comprises a client/server analysis component for analyzing one or more queries that are not supported by a target database system, and a client/server transformation component for transforming the one or more queries into semantically equivalent queries that are supported by the target database system.

System and Method of Query Transformation

FIELD OF THE INVENTION

The invention relates generally to data access middleware, and in particular to a
5 system and method of query transformation.

BACKGROUND OF THE INVENTION

A typical data access environment has a multi-tier architecture. For description
purposes, it can be separated into three distinct tiers:

- 10 • Web server
- Applications
- Data

The tiers are based on business function, and are typically separated by firewalls. Client
software, such as a browser or a report-authoring tool, sits above the tiers.

15 The web server contains a firewall and one or more gateways. All web
communication is performed through a gateway. A gateway is responsible for passing on
requests to the application server, in tier 2, for execution.

 The applications tier contains one or more application servers. The application
server runs requests, such as reports and queries that are forwarded by a gateway running
20 on the web server. Typically, one of the components of the applications tier is a query
engine, which is data access middleware that provides universal data access to a variety
of heterogeneous database systems. The query engine formulates queries (typically SQL)
and passes them on to the data tier, through a native database API (such as ODBC) for
execution.

25 The data tier contains database management systems (DBMS), which manage raw
data stored in a database. Examples of such systems include Oracle, DB2, and Microsoft
SQL Server.

 Although a multi-tier architecture can be configured in several different ways, a
typical configuration places each tier on a separate computer (server). A database server
30 is typically a "high end" server, and thus can process queries at a relatively fast speed.
An application server cannot generally process queries as quickly as a database server.

 In order to solve many business questions, a query engine may generate SQL
queries that utilize the SQL/OLAP technology introduced in the SQL:1999 standard.

However, many database systems do not support this technology. Thus, the SQL queries would have to be performed by the query engine on the application server that is generally slower than the database server. It is desirable to have as much processing performed on the database server as possible.

5 There is a need to prevent or reduce the amount of local (application server) processing required to process a query. In the past, the application would be responsible for generating SQL that contained a standard GROUP BY query. Quite often, generating this type of SQL is more difficult since it is more complex.

10 One way of overcoming this problem is for the query engine to generate a GROUP BY query to compute aggregates at the lowest level of granularity, and then compute the remaining aggregates based on these values. This approach reduces the amount of processing time required on the application server.

SUMMARY OF THE INVENTION

15 It is an object of the present invention to provide a method of Client/Server Aggregate query transformation in a database system that does not support SQL:1999 standard.

20 In accordance with an embodiment of the present invention, there is provided a system for transforming client/server aggregates. The system comprises a client/server analysis component for analyzing one or more queries that are not supported by a target database system, and a client/server transformation component for transforming the one or more queries into semantically equivalent queries that are supported by the target database system.

25 In accordance with another embodiment of the present invention, there is provided a method of client/server aggregate transformation. The method comprises the steps of analyzing one or more queries that are not supported by a target database system, and transforming the one or more queries into semantically equivalent queries that are supported by the target database system.

30 In accordance with another embodiment of the present invention, there is provided a method of client/server aggregate transformation. The method comprises the steps of analyzing a client/server select list to determine if a transformation is to be performed, generating a derived table, traversing all expressions in the main select list, adding

columns or expressions that are grouping columns to the derived table select list, and converting and moving eligible aggregates into the derived table select list.

In accordance with another embodiment of the present invention, there is provided a computer data signal embodied in a carrier wave and representing sequences of instructions which, when executed by a processor, cause the processor to perform a method of client/server aggregate transformation. The method comprises the steps of analyzing one or more queries that are not supported by a target database system, and transforming the one or more queries into semantically equivalent queries that are supported by the target database system.

In accordance with another embodiment of the present invention, there is provided a computer-readable medium having computer readable code embodied therein for use in the execution in a computer of a method of client/server aggregate transformation. The method comprises the steps of analyzing one or more queries that are not supported by a target database system, and transforming the one or more queries into semantically equivalent queries that are supported by the target database system.

In accordance with another embodiment of the present invention, there is provided a computer program product for use in the execution in a computer of a query transformation system for transforming client/server aggregates. The computer program product comprises a client/server analysis component for analyzing one or more queries that are not supported by a target database system, and a client/server transformation component for transforming the one or more queries into semantically equivalent queries that are supported by the target database system.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows a typical data access environment.

Figure 2 shows a client/server aggregate transformation system, in accordance with an embodiment of the present invention.

Figure 3 shows in a flowchart an example of a method of client/server aggregate transformation, in accordance with the client/server aggregate transformation system.

Figure 4 shows in a flowchart another example of a method of client/server aggregate transformation, in accordance with the client/server aggregate transformation system.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 shows a typical data access environment 10 for processing data. Typically, data is stored in a database 11. A DBMS running on a database server 12 accesses the raw data stored in the database 11. A query engine 15, running on a report server (or application server) 13 is used to generate reports on the raw data and instruct the DBMS on the database server 12 to obtain information pertaining to the raw data in the database 11. The query engine 15 provides universal data access to a variety of heterogeneous database systems. An end user uses a client application 14, running on a client workstation, to facilitate application server 13 operations.

In order to solve many business questions, a query engine 15 generates SQL queries that utilize the SQL/online analytical programming (OLAP) technology introduced in the SQL:1999 standard. These SQL queries include SQL/OLAP functions (windowed aggregates). However, many database systems 12 do not support this technology. In order to prevent or reduce the amount of local (application server) processing required to process these types of queries, the query engine 15 attempts to generate semantically equivalent queries that can be processed in whole or in part on the database server 12 by the target database system. These semantically equivalent queries involve standard aggregate functions and the GROUP BY operator.

Figure 2 shows a Client/Server Aggregate transformation system 20, in accordance with an embodiment of the present invention. The Client/Server Aggregate transformation system 20 comprises a Client/Server analysis module 21 for analysing SQL/OLAP queries that are not supported by a target database system, and a Client/Server transformation module 22 for transforming these SQL/OLAP queries into semantically equivalent queries that are supported by the target database system.

The Client/Server Aggregate transformation system 20 is implemented as a sub-system of the query engine 15 in the data access environment 10. This transformation 20 may generate queries that can be processed in their entirety on the database server 12, or queries that require processing on both the application server 13 and the database server 12.

Advantageously, the Client/Server Aggregate transformation system 20 reduces processing that might otherwise be required on the application server by generating a semantically equivalent query, thereby improving performance in many cases. Furthermore, the Client/Server Aggregate transformation system 20 takes advantage of

functionality provided by the target database. In particular, the Client/Server Aggregate system 20 utilizes the functionality provided by standard aggregates and the GROUP BY operator.

There are two types of OLAP functions: framed functions and report functions. Framed OLAP functions contain a window frame specification (ROWS or RANGE) and an ORDER BY clause. Through window frames, capabilities such as cumulative (running) sums and moving averages can be supported. Report functions do not contain a window frame specification, and produce the same value for each row in a partition.

The Client/Server Aggregate transformation attempts to reduce the number of derived tables generated in a SQL group transformation by computing some of the aggregates locally. This transformation can be utilized when the user is willing to tolerate some (local) processing on the application server.

Figure 3 shows a flowchart of an example of a method of Client/Server Aggregate transformation (30), in accordance with an embodiment of the Client/Server Aggregate transformation system 20. The method (30) begins with analysing a query containing a client/server aggregate that is not supported by a target database system (31). Next, the query is transformed into a semantically equivalent query that is supported by the target database system (32). The method (30) is done (33).

Figure 4 shows a flowchart of an example of a method of Client/Server Aggregate transformation (40), in accordance with an embodiment of the Client/Server Aggregate transformation system 20. The method (40) begins with analyzing a select list to determine if transformations can be performed (41). If transformations cannot be performed (42), the method (40) is done (46). If a transformation is to be performed (42), a derived table is generated (43). All (grouping) columns of the longest control break are added to the select list of the derived table (44). A control break is a set of columns over which a particular OLAP function is computed. A partition is a set of rows that have the same value for columns that make up the control break. Then the select list is traversed and candidate aggregates are converted and moved into the select list of the derived table (45). The method (40) is done (46).

To determine if the transformation can be performed, all expressions and aggregates in the select list are analyzed. This analysis consists of the following:

- Assigning a *group index*. This index indicates whether the aggregate is eligible for transformation. Any aggregate assigned a group index of -1 is **not** eligible. Report

aggregates are eligible, provided they have a standard aggregate counterpart (MIN, MAX, SUM, AVG, COUNT, COUNT(*)). Framed aggregates are not eligible.

- Keeping track of all control breaks. Control break information is stored in an ordered list, with control breaks having the fewest number of columns (least amount of detail) appearing first.

The transformation can be performed if the following criteria are met:

- The keyword DISTINCT is present in the select list.
- All aggregates in the select list are eligible.
- All control breaks are compatible. Assuming that the longest control break **G** is of the form ($C_1, C_2, \dots, C_n$), all other control breaks must be of the form ($C_1, C_2, \dots, C_k$), where $0 \leq k \leq n$. Note that the control break () is compatible with all control breaks.
- Any simple column reference in the select list that is not part of an aggregate specification is contained in the **G**, or is part of an expression contained in **G**.

In performing the transformation, the following assumptions are made:

- **m** equals the total number of unique control breaks
- **n** equals the number of columns in the longest control break
- **G** represents the list of columns in the longest control break, known as the set of grouping columns
- **AGG(C_0)** represents the standard form of the aggregate **AGG(C_0) OVER(...)**

Given these assumptions, the basic steps in performing the transformation are as follows:

1. Construct a new select list (known as the inner select list) to hold columns and standard aggregates.
2. Traverse all expressions in the original select list.

3. If the expression is a simple column reference or it is contained in **G**, copy it to the inner select list. Otherwise, traverse the expression, and apply the rules outlined here.

4. Apply one of the following transformations if **m** > 1 or **k** ≠ **n**:

- 5 a. If the expression is an aggregate is of the form **AVG(C₀) OVER()**, replace it with an expression of the form **SUM(SUM(C₀)) OVER() / SUM(COUNT(C₀)) OVER()** and move the standard (nested) aggregates into the inner select list.
- 10 b. If the expression is an aggregate is of the form **AVG(C₀) OVER(PARTITION BY C₁, C₂, ..., C_k)**, replace it with an expression of the form **SUM(SUM(C₀) OVER(PARTITION BY C₁, C₂, ..., C_k)) OVER() / SUM(COUNT(C₀) OVER(PARTITION BY C₁, C₂, ..., C_k)) OVER()** and move the standard (nested) aggregates into the inner select list.
- 15 5. If the expression is an aggregate of the form **AGG(C₀) OVER(PARTITION BY C₁, C₂, ..., C_n)**, replace it with **AGG(C₀)**, and move it into the inner select.
6. If the expression is an aggregate of the form **AGG(C₀) OVER(PARTITION BY C₁, C₂, ..., C_k)**, where **k** ≠ **n**, replace it with **AGG(AGG(C₀)) OVER(PARTITION BY C₁, C₂, ..., C_k)**, and move the standard (nested) aggregate into the inner select list.
- 20 7. Construct the final query. This consists of constructing a derived table containing a GROUP BY query for computing the standard aggregates in the inner select list.

For the purpose of discussion, assume the target database for the following examples does not support the SQL/OLAP technology introduced in the SQL:1999 standard. The native SQL shown in each example is the SQL passed to the target database.

Example 1

30 In this example, multiple SUM functions with compatible control breaks are present, so the transformation is straightforward.

Original Query

```

SELECT  DISTINCT SNO, PNO,
        SUM( QTY ) OVER (),
        SUM( QTY ) OVER ( PARTITION BY SNO ),
5      SUM( QTY ) OVER ( PARTITION BY SNO, PNO )
FROM    SUPPLY

```

Transformed Query

```

SELECT  C0, C1,
10      SUM( C2 ) OVER (),
        SUM( C2 ) OVER ( PARTITION BY SNO ),
        C2
FROM (   SELECT  SNO C0, PNO C1, SUM( QTY ) C2
        FROM    SUPPLY
15      GROUP  BY SNO, PNO ) T1

```

Native SQL

```

SELECT  SNO C0, PNO C1, SUM( QTY ) C2
FROM    SUPPLY
20  GROUP  BY SNO, PNO

```

Explanation

The original query contains OLAP SUM functions computed over the partitions
 (), (SNO), and (SNO, PNO). Because of the presence of the DISTINCT keyword, and the
 25 fact that the detail columns (SNO and PNO) are part of a PARTITION BY clause, the
 query can be rewritten using a single derived table that computes the sum with the finest
 granularity (SNO, PNO). The remaining SUM functions can be computed based on this
 value, as shown in the transformed query.

30 Example 2

In this example, multiple SUM functions with compatible control breaks are
 present, but an AVG function has been introduced, making the transformation a little
 more complex.

Original Query

```

SELECT  DISTINCT SNO, PNO,
        SUM( QTY ) OVER (),
        SUM( QTY ) OVER ( PARTITION BY SNO ),
5      SUM( QTY ) OVER ( PARTITION BY SNO, PNO ),
        AVG( QTY ) OVER ()
FROM    SUPPLY

```

Transformed Query

```

10  SELECT  C0,
        C1,
        SUM( C2 ) OVER (),
        SUM( C2 ) OVER ( PARTITION BY SNO ),
        C2,
15      SUM( C2 ) OVER () / SUM( C3 ) OVER ()
FROM ( SELECT  SNO C0, PNO C1, SUM( QTY ) C2, COUNT( QTY ) C3
        FROM    SUPPLY
        GROUP   BY SNO, PNO ) T1

```

Native SQL

```

SELECT  SNO C0, PNO C1, SUM( QTY ) C2, COUNT( QTY ) C3
FROM    SUPPLY
GROUP   BY SNO, PNO

```

Explanation

The original query contains OLAP SUM functions computed over the partitions (), (SNO), and (SNO, PNO). It also contains an OLAP AVG function computed over the partition (). Because of the presence of the DISTINCT keyword, and the fact that the detail columns (SNO and PNO) are part of a PARTITION BY clause, the query can be

30 rewritten using a single derived table that computes the sum with the finest granularity (SNO, PNO). The remaining SUM functions can be computed based on this value, as shown in the transformed query. This value can be used to compute the average as well. However, we also require a count to be computed at the lowest level of granularity. The

final average is then computed as shown in the transformed query (average of an average will not work).

Example 3

- 5 In this example, the OLAP functions SUM and AVG have identical control breaks corresponding to the required detail information (SNO, PNO), so the entire query can be replaced with a simple GROUP BY query.

Original Query

```

SELECT  DISTINCT SNO, PNO,
10      SUM( QTY ) OVER ( PARTITION BY SNO, PNO ),
      AVG( QTY ) OVER ( PARTITION BY SNO, PNO ),
FROM    SUPPLY

```

Transformed Query

```

15  SELECT  SNO C0, PNO C1, SUM( QTY ) C2, AVG( QTY ) C3
FROM    SUPPLY
GROUP   BY SNO, PNO

```

Native SQL

```

20  SELECT  SNO C0, PNO C1, SUM( QTY ) C2, AVG( QTY ) C3
FROM    SUPPLY
GROUP   BY SNO, PNO

```

Explanation

- 25 The original query contains two OLAP functions computed over the partition (SNO, PNO). This partition represents the grouping columns in the transformed query. Because of the presence of the DISTINCT keyword, and the fact that the detail columns (SNO and PNO) are grouping columns, the query can be rewritten as a simple GROUP BY query.

30

Example 4

In this example, the OLAP functions SUM and AVG have identical control breaks corresponding to the required detail information (SNO, PNO), so the entire query can be replaced with a simple GROUP BY query.

5 **Original Query**

```
SELECT  DISTINCT SNO, PNO,
        SUM( QTY ) OVER (),
        SUM( QTY ) OVER ( PARTITION BY SNO )
FROM    SUPPLY
```

10

Transformed Query

```
SELECT  C0, C1, SUM( C2 ) OVER (), SUM( C2 ) OVER ( PARTITION BY C0 )
FROM (   SELECT  SNO C0, PNO C1, SUM( QTY ) C2
        FROM    SUPPLY
        GROUP   BY SNO, PNO ) T1
```

15

Native SQL

```
SELECT  SNO C0, PNO C1, SUM( QTY ) C2
FROM    SUPPLY
GROUP   BY SNO, PNO
```

20

Explanation

The original query contains two SUM functions computed over the partitions () and (SNO). The longest control break (SNO), however, does not correspond to the detail information required (SNO, PNO). Hence, the derived table generated consists of a GROUP BY query that computes a sum grouped on (SNO, PNO). Both of the outer sum values are then computed based on this value.

25

A system for query transformation may comprise an analysis component and a transformation component. The analysis component may comprise a Nested Aggregate analysis module, a SQL GROUP transformation analysis module, a summary filter transformation analysis module, and a Summary Filter analysis module. The transformation component may comprise a Nested Aggregate transformation module, a

30

SQL GROUP transformation module, a Client/Server Aggregate transformation module, and a Summary Filter transformation module.

5 The systems and methods according to the present invention may be implemented by any hardware, software or a combination of hardware and software having the functions described above. The software code, either in its entirety or a part thereof, may be stored in a computer readable memory. Further, a computer data signal representing the software code that may be embedded in a carrier wave may be transmitted via a communication network. Such a computer readable memory and a computer data signal are also within the scope of the present invention, as well as the hardware, software and
10 the combination thereof.

While particular embodiments of the present invention have been shown and described, changes and modifications may be made to such embodiments without departing from the true scope of the invention.

WHAT IS CLAIMED IS:

1. A system for transforming client/server aggregates, the system comprising:
 - a client/server analysis component for analyzing a query containing a client/server
 - 5 aggregate that are not supported by a target database system; and
 - a client/server transformation component for transforming the query into a
 - semantically equivalent query that is supported by the target database system.
2. A method of client/server aggregate transformation, the method comprising the steps
 - 10 of:
 - analyzing a query having a client/server aggregate that is not supported by a target
 - database system; and
 - transforming the query into a semantically equivalent query that is supported by
 - the target database system.
- 15 3. The method as claimed in claim 2, further comprising the steps of:
 - determining if a transformation is to be performed;
 - traversing all expressions in an original select list;
 - copying expressions contained in **G** to an inner select list, and traversing the
 - 20 expression otherwise, applying the rules outlined here;
 - applying one of the following transformations if $m > 1$ or $k \neq n$:
 - replacing an aggregate of the form **AVG(C₀) OVER()** with an expression
 - of the form **SUM(SUM(C₀)) OVER() / SUM(COUNT(C₀)) OVER()** and
 - moving the standard (nested) aggregates into the inner select list (43); and
 - 25 replacing an aggregate of the form **AVG(C₀) OVER(PARTITION BY**
 - C₁, C₂, ..., C_k)** with an expression of the form **SUM(SUM(C₀) OVER(**
 - PARTITION BY C₁, C₂, ..., C_k)) OVER() / SUM(COUNT(C₀) OVER(**
 - PARTITION BY C₁, C₂, ..., C_k)) OVER()** and moving the standard (nested)
 - aggregates into the inner select list (44); and
 - 30 constructing the final query.

4. The method as claimed in claim 3, wherein the step of determining if a transformation should be performed comprises the steps of:

assigning a group index to all aggregates;

keeping track of control breaks; and

5 determining whether certain criteria are met.

5. The method as claimed in claim 3, wherein the step of applying a client/server transformation comprises the step of replacing **AVG(C₀) OVER()** with an expression of the form **SUM(SUM(C₀)) OVER() / SUM(COUNT(C₀)) OVER()** and moving the
10 standard (nested) aggregates into the inner select list.

6. The method as claimed in claim 3, wherein the step of applying a client/server transformation comprises the step of replacing **AVG(C₀) OVER(PARTITION BY C₁, C₂, ..., C_k)** with an expression of the form **SUM(SUM(C₀)) OVER(PARTITION
15 BY C₁, C₂, ..., C_k) / SUM(COUNT(C₀)) OVER(PARTITION BY C₁, C₂, ..., C_k)** and moving the standard (nested) aggregates into the inner select list.

7. A method of client/server aggregate transformation, the method comprising the steps of:

20 analyzing a client/server select list to determine if a transformation is to be performed;

generating a derived table;

traversing all expressions in the main select list;

adding columns or expressions that are grouping columns to the derived table

25 select list; and

converting and moving eligible aggregates into the derived table select list.

8. A computer data signal embodied in a carrier wave and representing sequences of instructions which, when executed by a processor, cause the processor to perform a
30 method of client/server aggregate transformation, the method comprising the steps of:

analyzing one or more queries that are not supported by a target database system;

and

transforming the one or more queries into semantically equivalent queries that are supported by the target database system.

9. A computer-readable medium having computer readable code embodied therein for
5 use in the execution in a computer of a method of client/server aggregate transformation, the method comprising the steps of:

analyzing one or more queries that are not supported by a target database system;

and

- transforming the one or more queries into semantically equivalent queries that are
10 supported by the target database system.

10. A computer program product for use in the execution in a computer of a query transformation system for transforming client/server aggregates, the computer program product comprising:

- 15 a client/server analysis component for analyzing one or more queries that are not supported by a target database system; and

a client/server transformation component for transforming the one or more queries into semantically equivalent queries that are supported by the target database system.




Figure 1




Figure 2




Figure 3




Figure 4

20

Client Server Aggregate Transformation System

Client Server
Analysis Module

21

22

Client Server
Transformation
Module