



US 20080016043A1

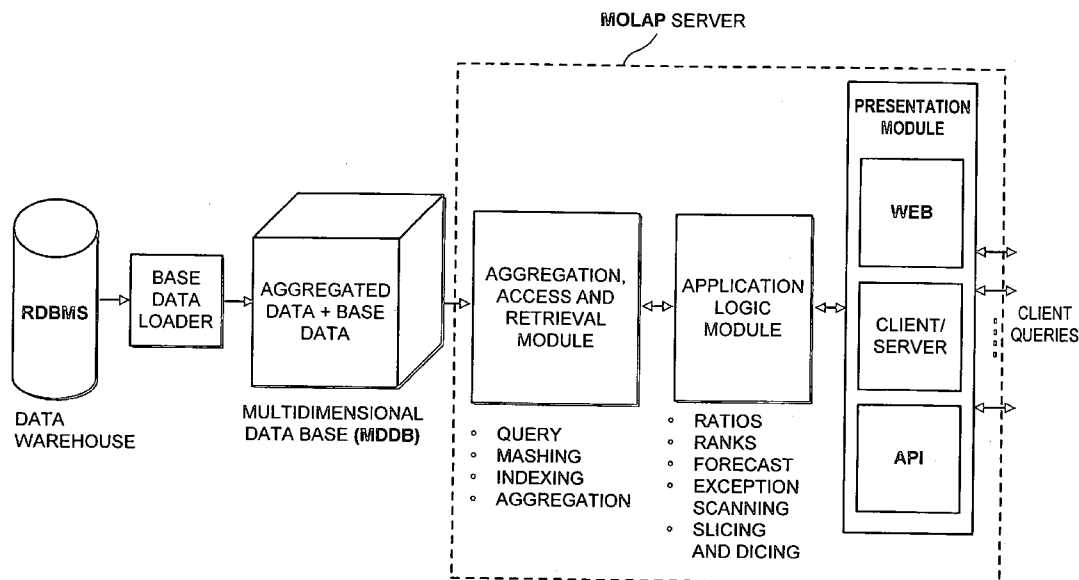
(19) **United States**(12) **Patent Application Publication**
Bakalash et al.(10) **Pub. No.: US 2008/0016043 A1**(43) **Pub. Date: Jan. 17, 2008**(54) **RELATIONAL DATABASE MANAGEMENT
SYSTEM (RDBMS) EMPLOYING
MULTI-DIMENSIONAL DATABASE (MDDDB)
FOR SERVICING QUERY STATEMENTS
THROUGH ONE OR MORE CLIENT
MACHINES**(76) Inventors: **Reuven Bakalash**, Beer Sheva (IL);
Guy Shaked, Shdema (IL); **Joseph
Caspi**, Herzlyia (IL)Correspondence Address:
Thomas J. Perkowski, Esq., P.C.
Soundview Plaza
1266 East Main Street
Stamford, CT 06902 (US)(21) Appl. No.: **11/825,644**(22) Filed: **Jul. 6, 2007****Related U.S. Application Data**

(63) Continuation of application No. 10/842,648, filed on May 10, 2004, now abandoned, which is a continuation of application No. 10/314,884, filed on Dec. 9, 2002, which is a continuation of application No. 09/796,098, filed on Feb. 28, 2001, now abandoned, which is a continuation-in-part of application No. 09/514,611, filed on Feb. 28, 2000, now Pat. No. 6,434,544, and which is a continuation-in-part of

application No. 09/634,748, filed on Aug. 9, 2000, now Pat. No. 6,385,604.

Publication Classification(51) **Int. Cl.**
G06F 7/06 (2006.01)
G06F 17/30 (2006.01)
(52) **U.S. Cl.** **707/3; 707/E17; 707/E17**(57) **ABSTRACT**

A relational database management system (RDBMS) for servicing query statements through one or more client machines. The RDBMS comprises a query interface adapted to receive query statements from the client machines. The query handling mechanism (i) receives each request from the query interface, (ii) extracts a set of dimensions associated with the request, (iii) uses the dimensions to retrieve aggregated fact data from a multi-dimensional database (MDDDB), and (iv) forwards retrieved aggregated fact data to the query processing mechanism for subsequent processing. When the query processing mechanism determines that servicing of one or more query requests require data stored in the relational tables, then the query processing mechanism automatically routes the requests to the relational data tables, so that data can be accessed from the relational tables and forwarded to the query processing mechanism for use in servicing the requests, in a manner transparent to the client machine.






Fig. 1A
(PRIOR ART)




FIG. 1B





Fig. 2A
(PRIOR ART)



G geography (e.g. cities, states, countries, continents)
T time (e.g., days, weeks, months, years)
P products (e.g. all products, by manufacturer)

Fig. 2B
(PRIOR ART)

Array structure of a
multidimensional variable

		D0					
		0	1	2	3	4	5
D2=0	D1= 0						
	D1= 1						
	D1= 2						
D2=1	D1= 0						
	D1= 1						
	D1= 2						
D2=2	D1= 0						
	D1= 1						
	D1= 2						
D2=3	D1= 0						
	D1= 1						
	D1= 2						
D2=3	D1= 0						
	D1= 1						
	D1= 2						

Fig. 2C
(PRIOR ART)

Page Allocation Table pointing on physical records of a multidimensional variable (e.g. the two first rows of a variable of FIG. 2B reside in page # 0)




Fig. 2D
(PRIOR ART)




Fig. 3A
(PRIOR ART)




Fig. 3B
(PRIOR ART)






Fig. 3C1
(PRIOR ART)



Spatial occupancy of TIME hierarchy

Fig. 3C3
(PRIOR ART)



Hierarchy of TIME dimension

Fig. 3C2
(PRIOR ART)

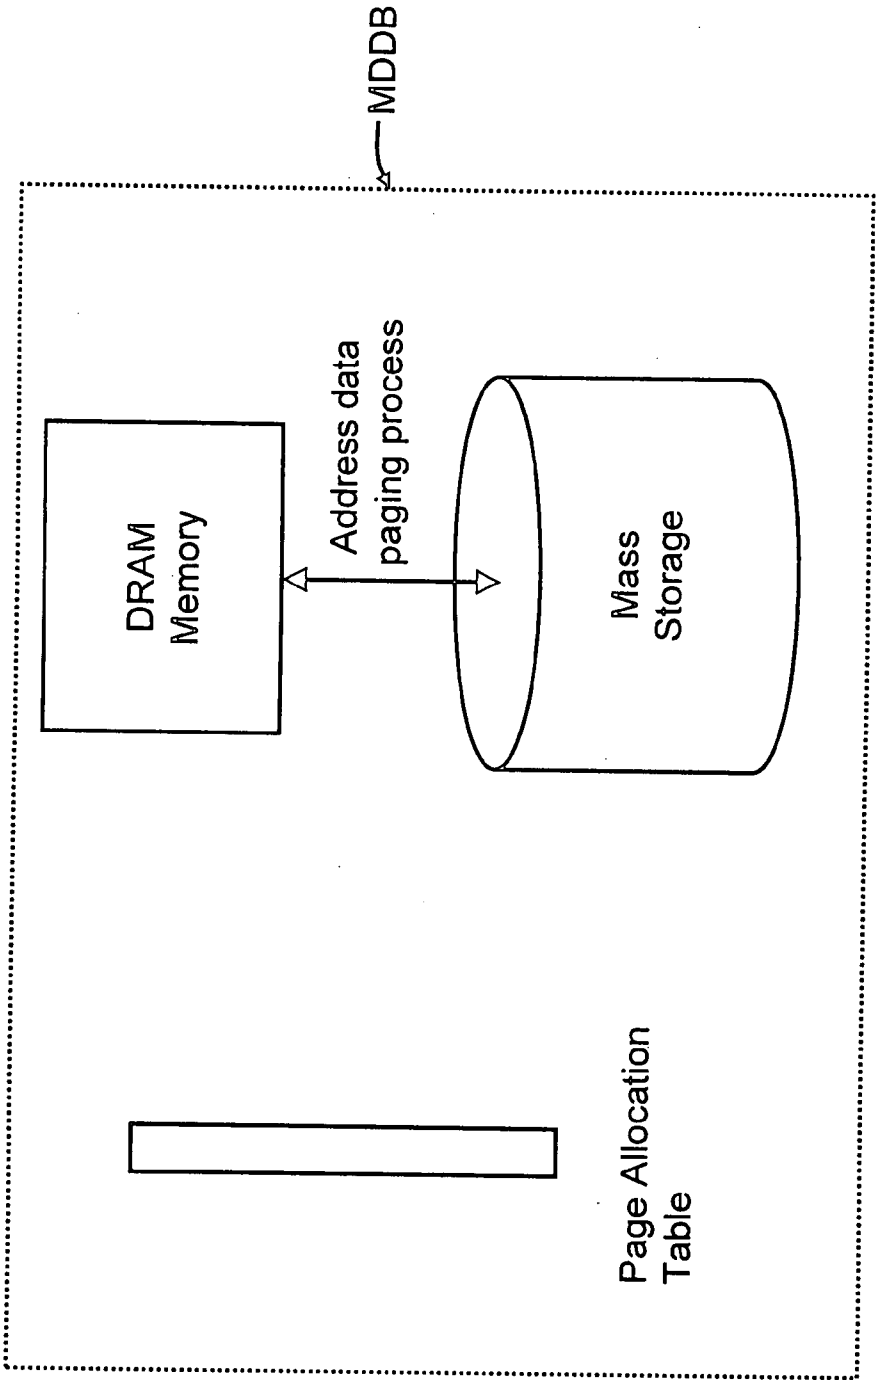


Fig. 4
(PRIOR ART)




Fig. 5
(PRIOR ART)




FIG. 6A




FIG. 6B

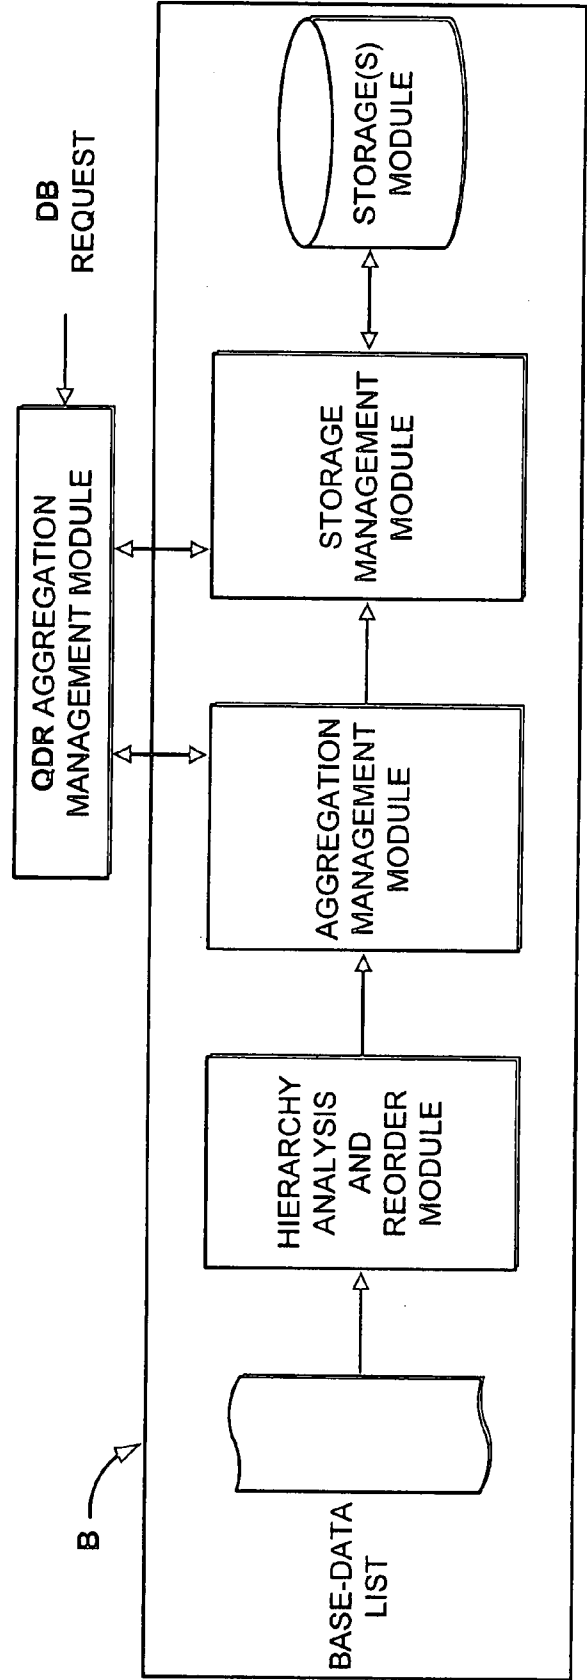


FIG. 6C

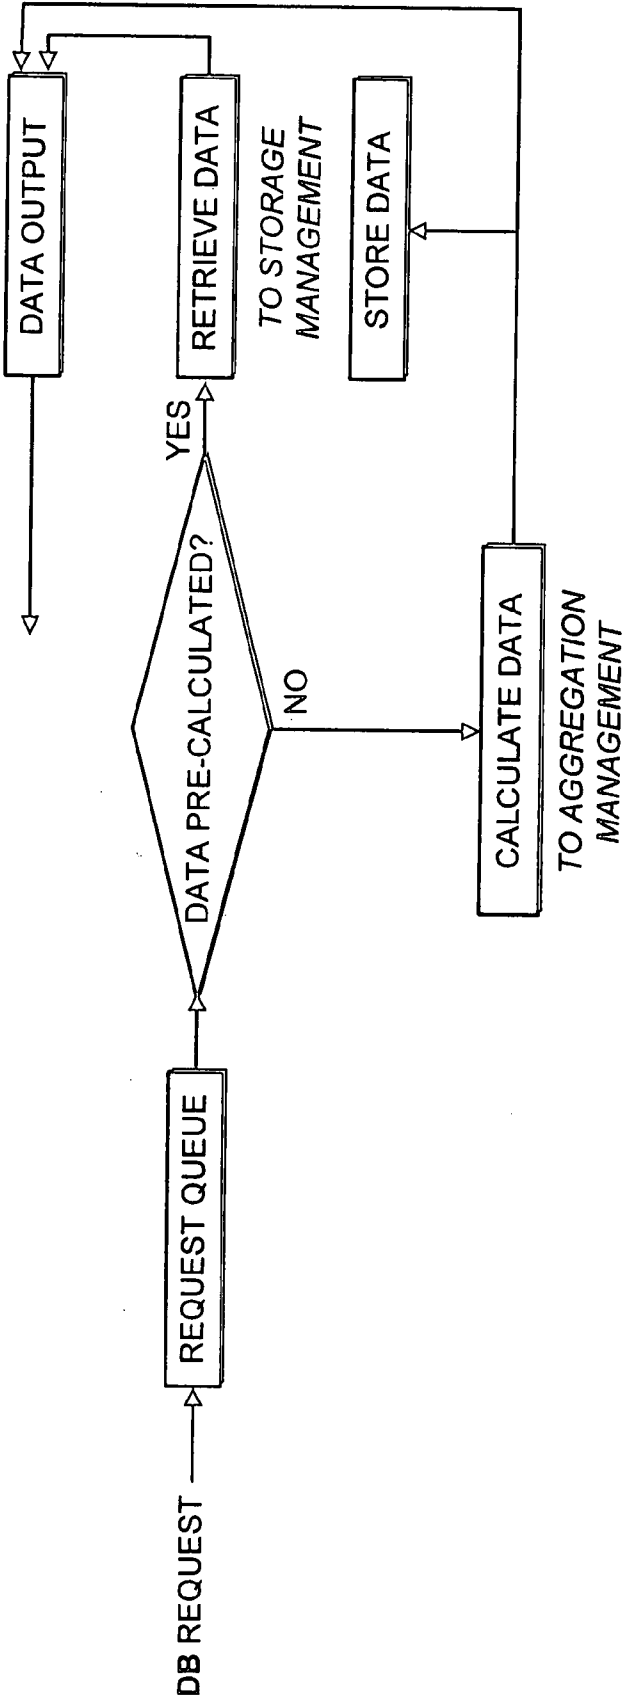


FIG. 6D




FIG. 7A




FIG. 7B

	NBR. OF DIM.	NBR. OF ATOMIC DATA DATA VALUES	LEAF NODE DENSITY %	NUMBER OF VALUES IN CUBE AFTER ROLL-UP	ORACLE EXPRESS V. 6.2	IMPLEMENTATION OF CURRENT INVENTION
D1	6	302M	9	427 M	16 h	15 m
D2	4	414M	1.27	969 M	50 m	5 m
D3	5	14,499M	0.03	63,954 M	31 h	1h 23 m
D4	6	623,494M	$8 * 10^{-4}$	7,930 G	EXCEEDS 48 h	2 h 20 m
D5	6	243,000M	10^{-8}	1,160,000 G	22 h	4 m
D6	4	7M	DEFINED AS 100	19 M	15 m	1 m

FIG. 8A




FIG. 9A






FIG. 9B



A. DIRECTED AGGREGATION IN
DIMENSION 2, CASES 1 AND 2
FIG. 9C1



A. DIRECTED AGGREGATION IN
DIMENSION 3, CASES 1 AND 2
FIG. 9C2




FIG. 10A




FIG. 10B





FIG. 11A



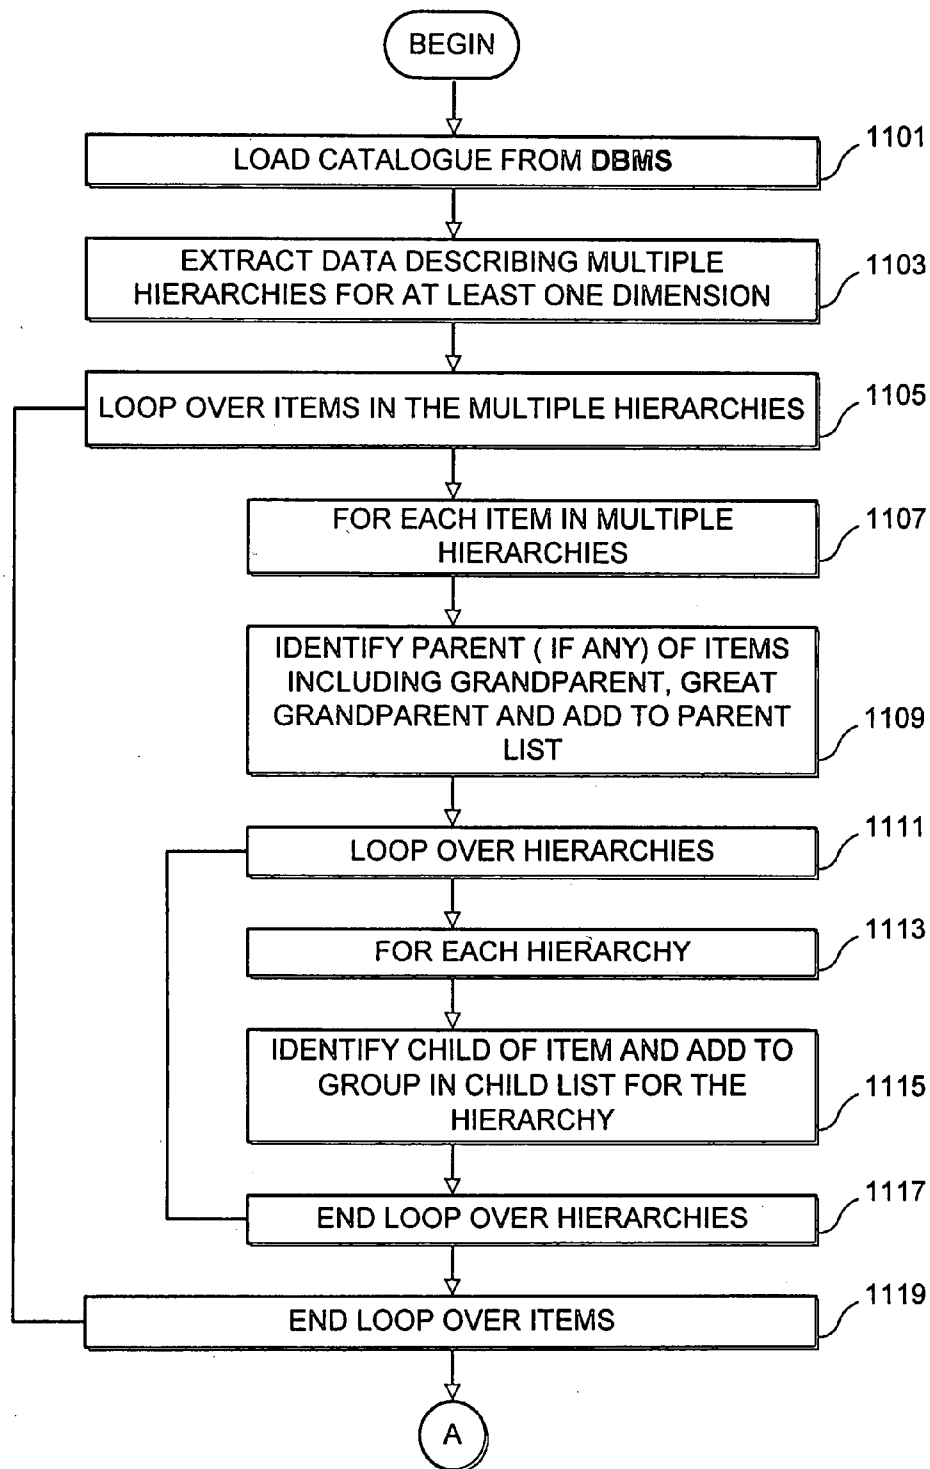


FIG. 11C(i)

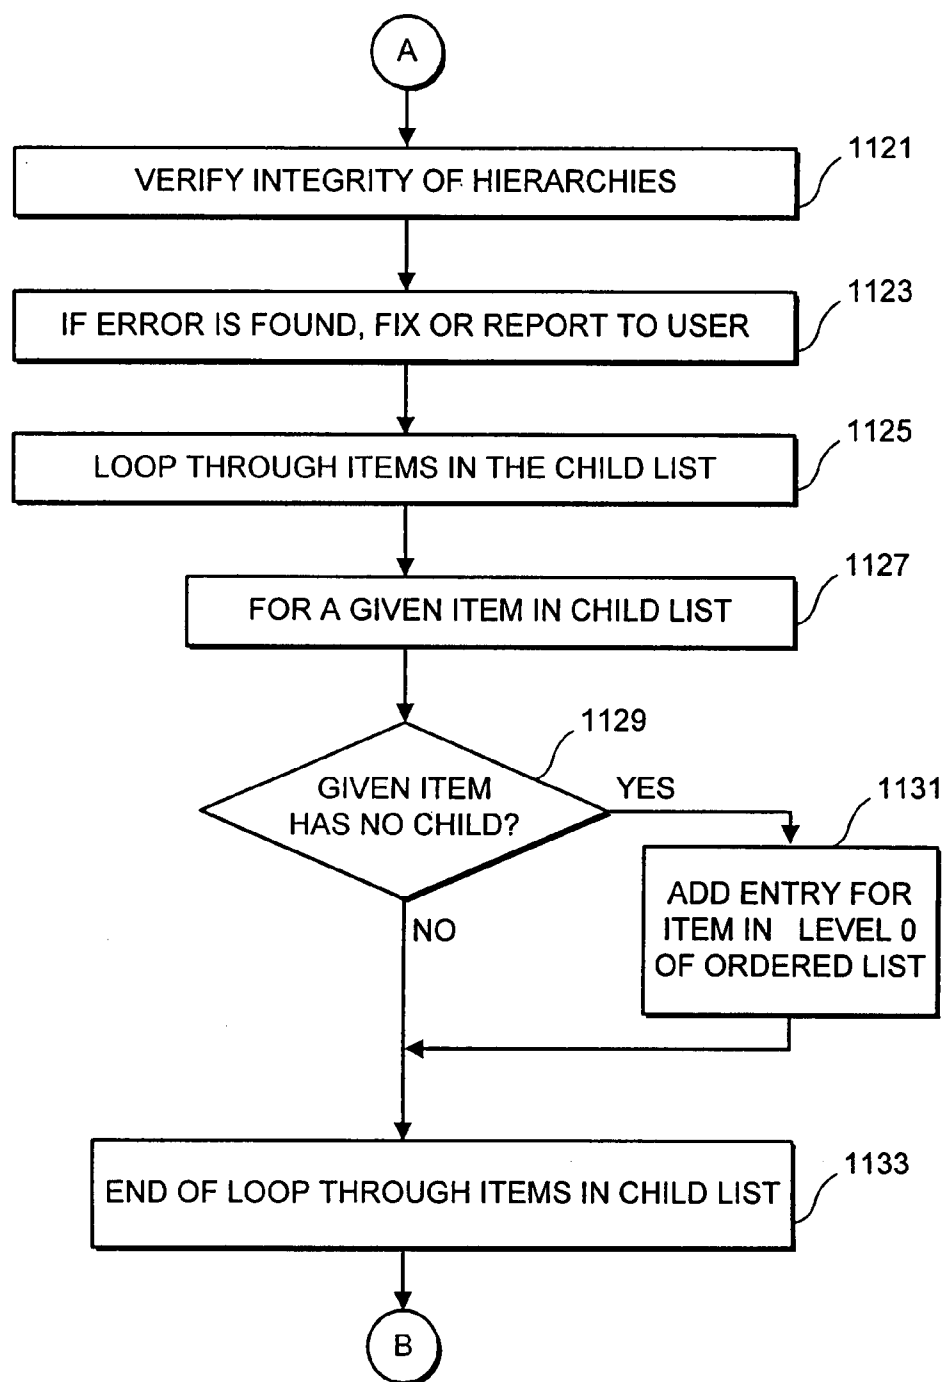


FIG. 11C(ii)

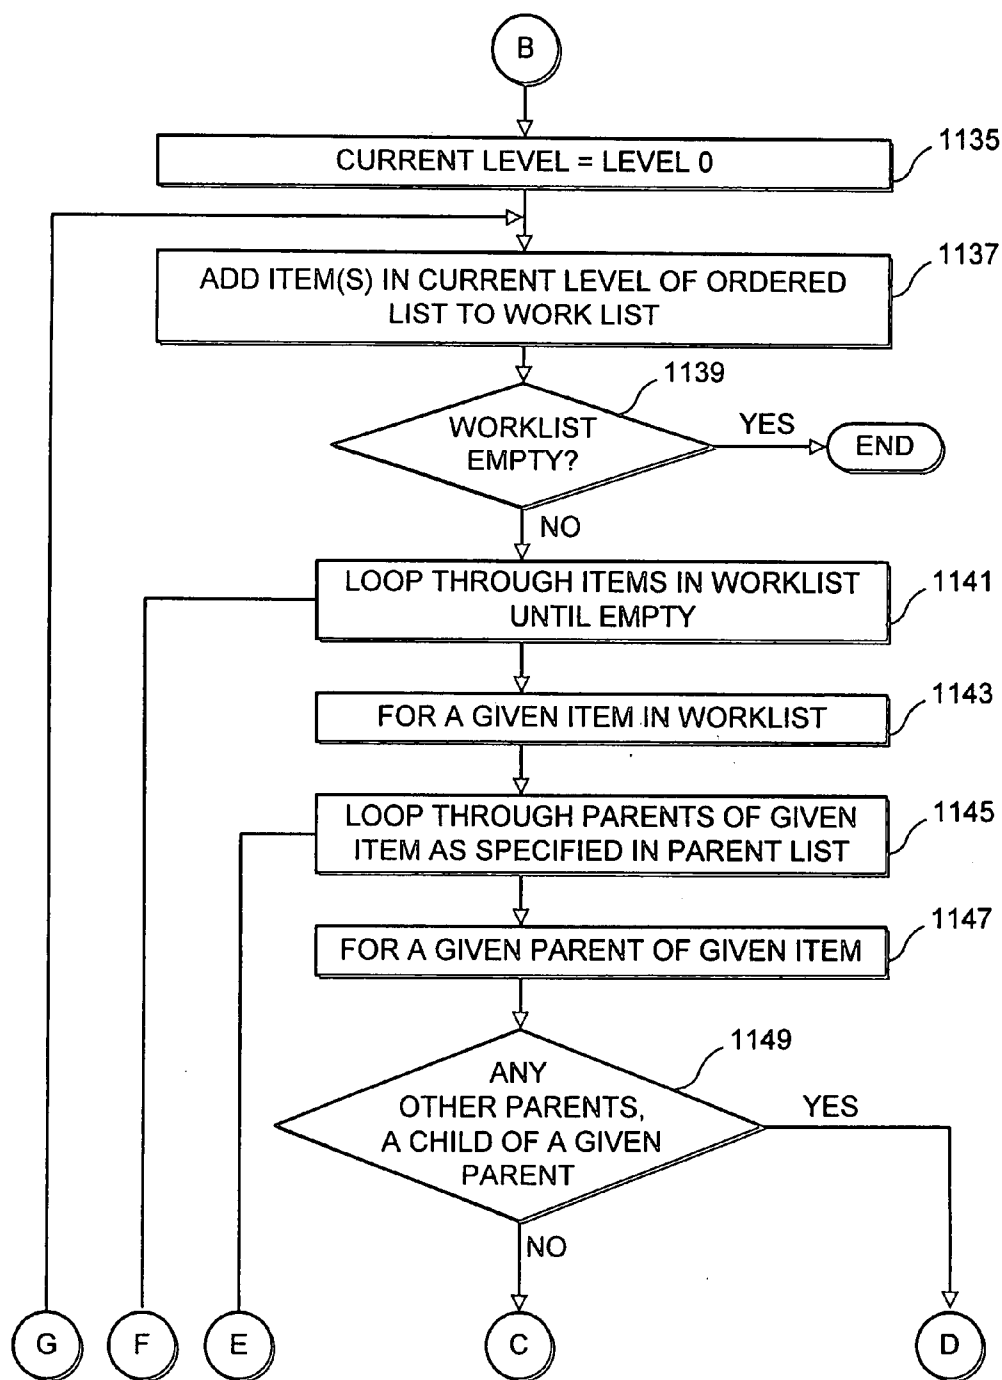


FIG. 11C(iii)




FIG. 11C(iv)

PARENT LIST	
ITEM	PARENT(S)
A	C, H, D
B	C, I, D
F	E, H, D
G	E, I, D
C	D
H	D
E	D
I	D
D	—

FIG. 11C(v)

CHILD LIST	
ITEM	CHILD(REN)
A	—
B	—
F	—
G	—
C	<A, B>
H	<F, G>
E	<A, F>
I	<B, G>
D	<A, B, F, G>, <H, I>, <C, E>

FIG. 11C(vi)

ORDERED LIST LEVEL 0	
ITEM	CHILD(REN)
A	—
B	—
F	—
G	—

FIG. 11C(vii)

ORDERED LIST LEVEL 1	
ITEM	CHILD(REN)
C	A, B
H	A, F
I	B, G
E	F, G

FIG. 11C(viii)

ORDERED LIST LEVEL 2	
ITEM	CHILD(REN)
D	C, E

FIG. 11C(ix)




FIG. 12

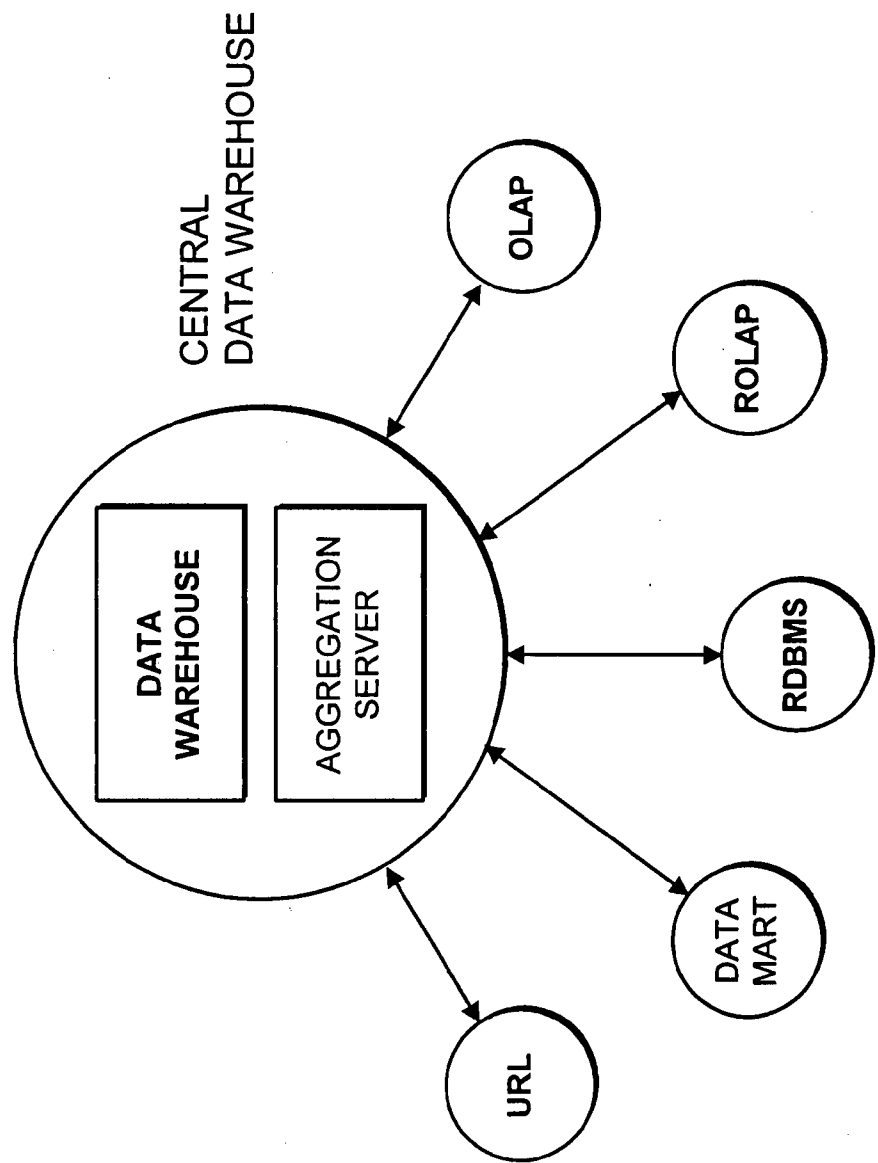


FIG. 13




FIG. 14 (PRIOR ART)




FIG. 15 (PRIOR ART)

CELLAR

WINE	YEAR	BOTTLES
CHARDONNAY	1996	4
FUME BLANK	1996	2
PINOT NOIR	1993	3
ZINFANDEL	1994	9

FIG. 16A

RESTRICT: OPERATOR:
SELECT WINE, YEAR,
BOTTLES FROM CELLAR
WHERE YEAR IS > 1995;

RESULT

WINE	YEAR	BOTTLES
CHARDONNAY	1996	4
FUME BLANK	1996	2

FIG. 16B

PROJECT: OPERATOR:
SELECT WINE, BOTTLES
FROM CELLAR;

RESULT

WINE	BOTTLES
CHARDONNAY	4
FUME BLANK	2
PINOT NOIR	3
ZINFANDEL	9

FIG. 16C




FIG. 17A




FIG. 18A




FIG. 18B




FIG. 19A




FIG. 19B

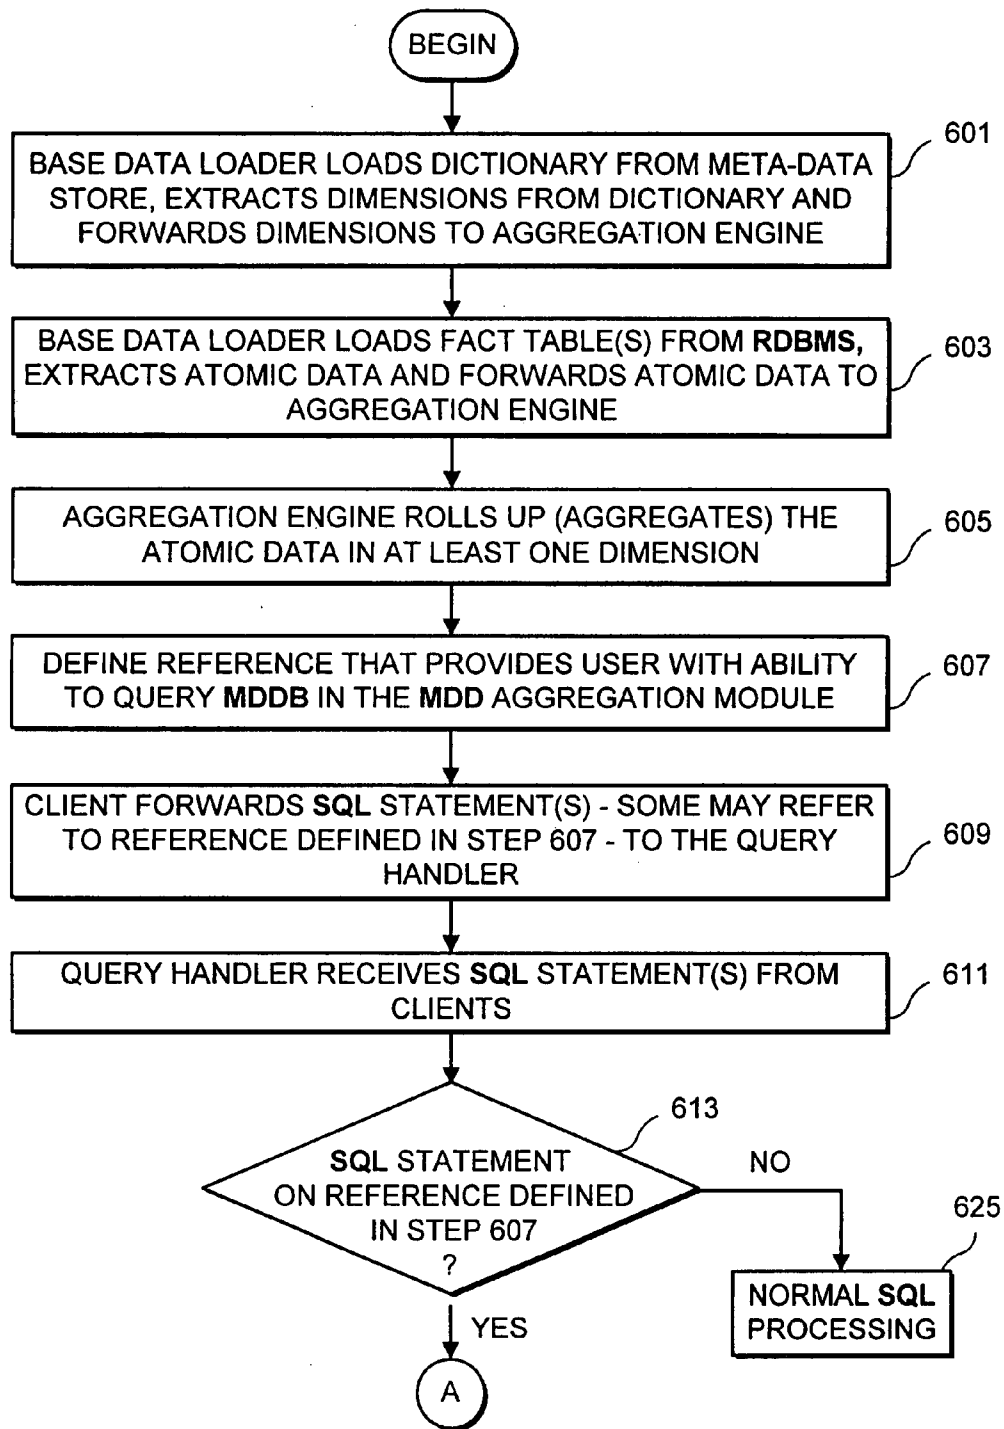


FIG. 19C(i)

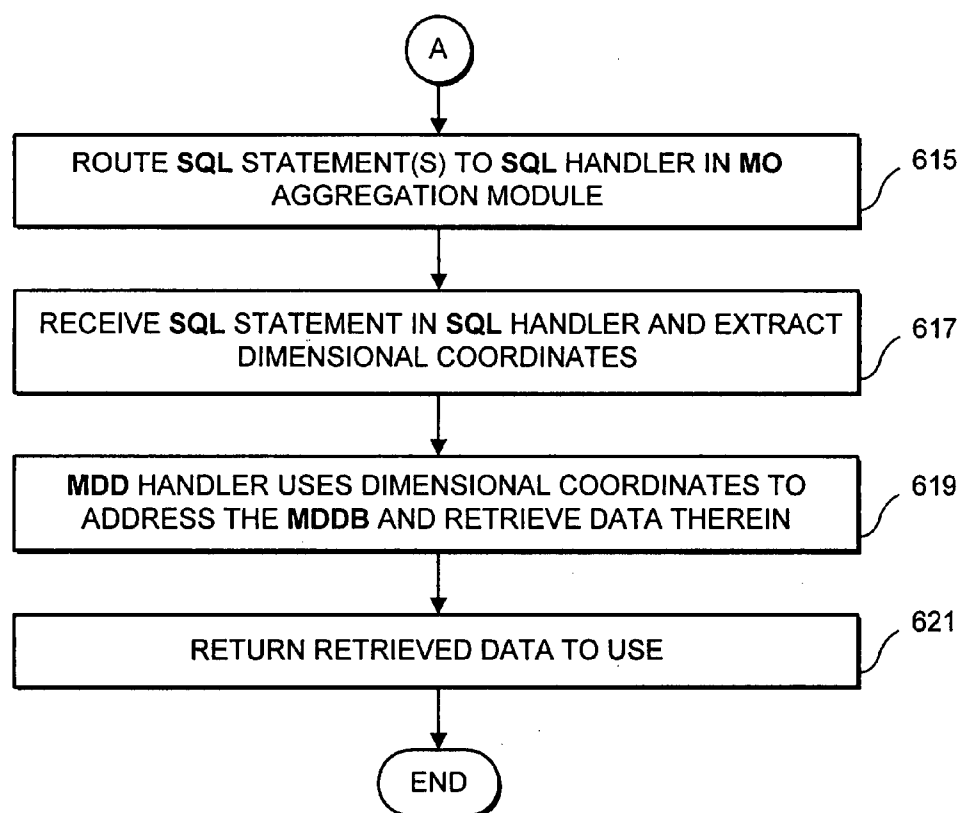


FIG. 19C(ii)

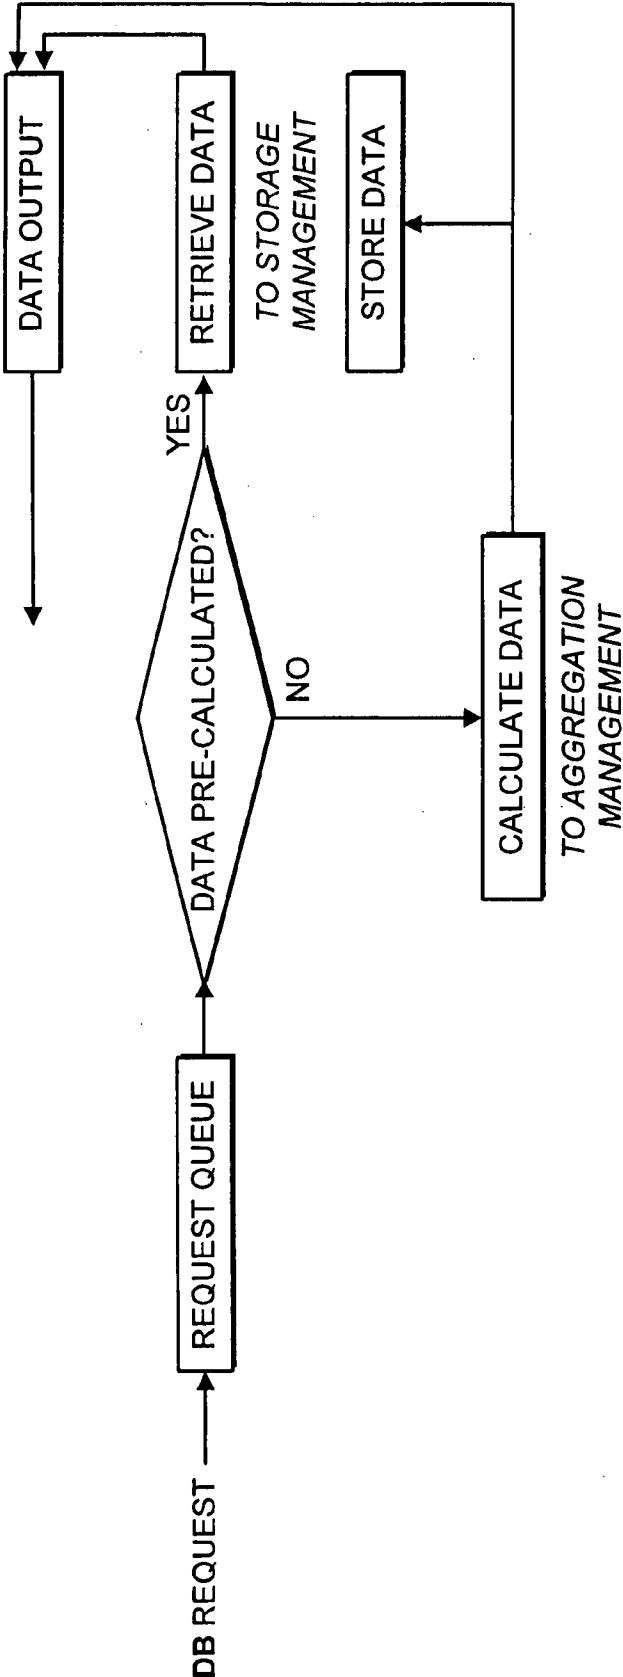


FIG. 19D




FIG. 19E




FIG. 19F




FIG. 19G

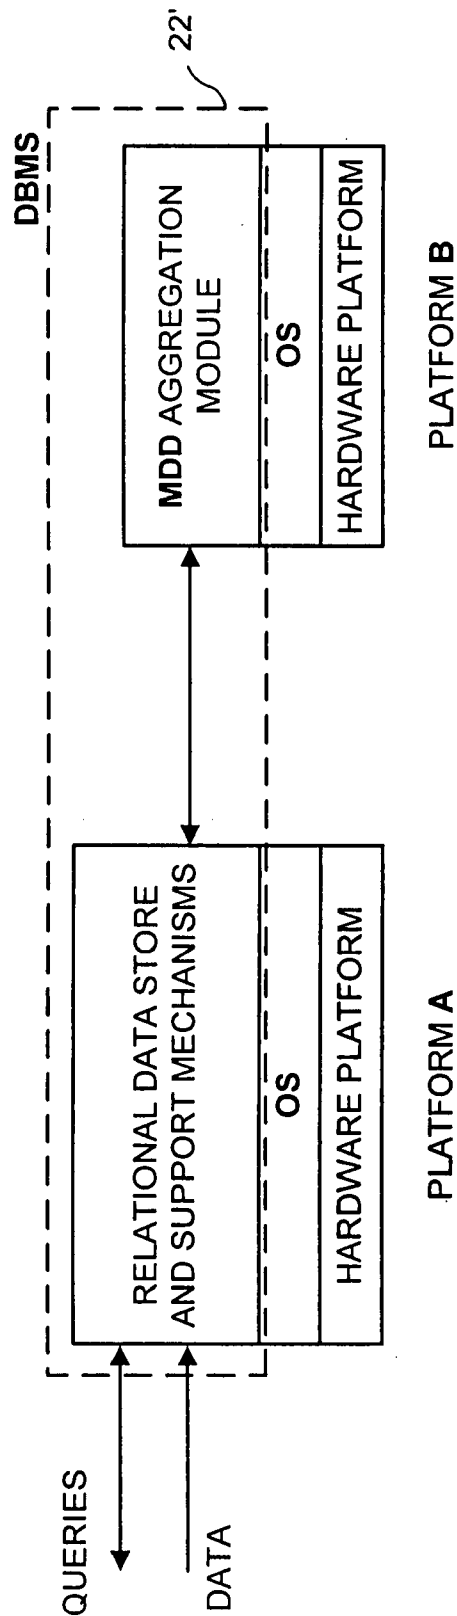


FIG. 20A




FIG. 20B




FIG. 21




FIG. 22

**RELATIONAL DATABASE MANAGEMENT
SYSTEM (RDBMS) EMPLOYING
MULTI-DIMENSIONAL DATABASE (MDDb) FOR
SERVICING QUERY STATEMENTS THROUGH
ONE OR MORE CLIENT MACHINES**

RELATED CASES

[0001] This is a Continuation of U.S. application Ser. No. 10/842,648 filed May 10, 2004, which is a Continuation of U.S. application Ser. No. 10/314,884 filed Dec. 9, 2002; which is a Continuation of U.S. application Ser. No. 09/796,098 filed Feb. 28, 2001, now abandoned; which is a Continuation-in-part of: U.S. application Ser. No. 09/514,611 filed Feb. 28, 2000, now U.S. Pat. No. 6,434,544, and U.S. application Ser. No. 09/634,748 filed Aug. 9, 2000, now U.S. Pat. No. 6,385,604; each said Application being commonly owned by HyperRoll Israel, Limited, and incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

[0002] 1. Field of Invention

[0003] The present invention relates to a method of and system for aggregating data elements in a multi-dimensional database (MDDb) supported upon a computing platform and also to provide an improved method of and system for managing data elements within a MDDb during on-line analytical processing (OLAP) operations and as an integral part of a database management system.

[0004] 2. Brief Description Of The State Of The Art

[0005] The ability to act quickly and decisively in today's increasingly competitive marketplace is critical to the success of organizations. The volume of information that is available to corporations is rapidly increasing and frequently overwhelming. Those organizations that will effectively and efficiently manage these tremendous volumes of data, and use the information to make business decisions, will realize a significant competitive advantage in the marketplace.

[0006] Data warehousing, the creation of an enterprise-wide data store, is the first step towards managing these volumes of data. The Data Warehouse is becoming an integral part of many information delivery systems because it provides a single, central location where a reconciled version of data extracted from a wide variety of operational systems is stored. Over the last few years, improvements in price, performance, scalability, and robustness of open computing systems have made data warehousing a central component of Information Technology CIT strategies. Details on methods of data integration and constructing data warehouses can be found in the white paper entitled "Data Integration: The Warehouse Foundation" by Louis Rollleigh and Joe Thomas.

[0007] Building a Data Warehouse has its own special challenges (e.g. using common data model, common business dictionary, etc.) and is a complex endeavor. However, just having a Data Warehouse does not provide organizations with the often-heralded business benefits of data warehousing. To complete the supply chain from transactional systems to decision maker, organizations need to deliver systems that allow knowledge workers to make strategic and tactical decisions based on the information stored in these warehouses. These decision support systems are referred to

as On-Line Analytical Processing (OLAP) systems. OLAP systems allow knowledge workers to intuitively, quickly, and flexibly manipulate operational data using familiar business terms, in order to provide analytical insight into a particular problem or line of inquiry. For example, by using an OLAP system, decision makers can "slice and dice" information along a customer (or business) dimension, and view business metrics by product and through time. Reports can be defined from multiple perspectives that provide a high-level or detailed view of the performance of any aspect of the business. Decision makers can navigate throughout their database by drilling down on a report to view elements at finer levels of detail, or by pivoting to view reports from different perspectives. To enable such full-functioned business analyses, OLAP systems need to (1) support sophisticated analyses, (2) scale to large numbers of dimensions, and (3) support analyses against large atomic data sets. These three key requirements are discussed further below.

[0008] Decision makers use key performance metrics to evaluate the operations within their domain, and OLAP systems need to be capable of delivering these metrics in a user-customizable format. These metrics may be obtained from the transactional databases pre-calculated and stored in the database, or generated on demand during the query process. Commonly used metrics include:

[0009] (1) Multidimensional Ratios (e.g. Percent to Total)—"Show me the contribution to weekly sales and category profit made by all items sold in the Northwest stores between July 1 and July 14."

[0010] (2) Comparisons (e.g. Actual vs. Plan, This Period vs. Last Period)—"Show me the sales to plan percentage variation for this year and compare it to that of the previous year to identify planning discrepancies."

[0011] (3) Ranking and Statistical Profiles (e.g. Top N/Bottom N, 70/30, Quartiles)—"Show me sales, profit and average call volume per day for my 20 most profitable salespeople, who are in the top 30% of the worldwide sales."

[0012] (4) Custom Consolidations—"Show me an abbreviated income statement by quarter for the last two quarters for my Western Region operations."

[0013] Knowledge workers analyze data from a number of different business perspectives or dimensions. As used hereinafter, a dimension is any element or hierarchical combination of elements in a data model that can be displayed orthogonally with respect to other combinations of elements in the data model. For example, if a report lists sales by week, promotion, store, and department, then the report would be a slice of data taken from a four-dimensional data model.

[0014] Target marketing and market segmentation applications involve extracting highly qualified result sets from large volumes of data. For example, a direct marketing organization might want to generate a targeted mailing list based on dozens of characteristics, including purchase frequency, size of the last purchase, past buying trends, customer location, age of customer, and gender of customer. These applications rapidly increase the dimensionality requirements for analysis.

[0015] The number of dimensions in OLAP systems range from a few orthogonal dimensions to hundreds of orthogonal dimensions. Orthogonal dimensions in an exemplary OLAP application might include Geography, Time, and Products.

[0016] Atomic data refers to the lowest level of data granularity required for effective decision making. In the case of a retail merchandising manager, "atomic data" may refer to information by store, by day, and by item. For a banker, atomic data may be information by account, by transaction, and by branch. Most organizations implementing OLAP systems find themselves needing systems that can scale to tens, hundreds, and even thousands of gigabytes of atomic information.

[0017] As OLAP systems become more pervasive and are used by the majority of the enterprise, more data over longer time frames will be included in the data store (i.e. data warehouse), and the size of the database will increase by at least an order of magnitude. Thus, OLAP systems need to be able to scale from present to near-future volumes of data.

[0018] In general, OLAP systems need to (1) support the complex analysis requirements of decision-makers, (2) analyze the data from a number of different perspectives (i.e. business dimensions), and (3) support complex analyses against large input (atomic-level) data sets from a Data Warehouse maintained by the organization using a relational database management system (RDBMS).

[0019] Vendors of OLAP systems classify OLAP Systems as either Relational OLAP (ROLAP) or Multidimensional OLAP (MOLAP) based on the underlying architecture thereof. Thus, there are two basic architectures for On-Line Analytical Processing systems: the ROLAP Architecture, and the MOLAP architecture.

[0020] The Relational OLAP (ROLAP) system accesses data stored in a Data Warehouse to provide OLAP analyses. The premise of ROLAP is that OLAP capabilities are best provided directly against the relational database, i.e. the Data Warehouse.

[0021] The ROLAP architecture was invented to enable direct access of data from Data Warehouses, and therefore support optimization techniques to meet batch window requirements and provide fast response times. Typically, these optimization techniques include application-level table partitioning, pre-aggregate inferencing, denormalization support, and the joining of multiple fact tables.

[0022] A typical prior art ROLAP system has a three-tier or layer client/server architecture. The "database layer" utilizes relational databases for data storage, access, and retrieval processes. The "application logic layer" is the ROLAP engine which executes the multidimensional reports from multiple users. The ROLAP engine integrates with a variety of "presentation layers," through which users perform OLAP analyses.

[0023] After the data model for the data warehouse is defined, data from on-line transaction-processing (OLTP) systems is loaded into the relational database management system (RDBMS). If required by the data model, database routines are run to pre-aggregate the data within the RDBMS. Indices are then created to optimize query access times. End users submit multidimensional analyses to the ROLAP engine, which then dynamically transforms the

requests into SQL execution plans. The SQL execution plans are submitted to the relational database for processing, the relational query results are cross-tabulated, and a multidimensional result data set is returned to the end user. ROLAP is a fully dynamic architecture capable of utilizing pre-calculated results when they are available, or dynamically generating results from atomic information when necessary.

[0024] Multidimensional OLAP (MOLAP) systems utilize a proprietary multidimensional database (MDDDB) to provide OLAP analyses. The MDDDB is logically organized as a multidimensional array (typically referred to as a multidimensional cube or hypercube or cube) whose rows/columns each represent a different dimension (i.e., relation). A data value is associated with each combination of dimensions (typically referred to as a "coordinate"). The main premise of this architecture is that data must be stored multidimensionally to be accessed and viewed multidimensionally.

[0025] As shown in FIG. 1B, prior art MOLAP systems have an Aggregation, Access and Retrieval module which is responsible for all data storage, access, and retrieval processes, including data aggregation (i.e. preaggregation) in the MDDDB. As shown in FIG. 1B, the base data loader is fed with base data, in the most detailed level, from the Data Warehouse, into the Multi-Dimensional Data Base (MDDDB). On top of the base data, layers of aggregated data are built-up by the Aggregation program, which is part of the Aggregation, Access and Retrieval module. As indicated in this figure, the application logic module is responsible for the execution of all OLAP requests/queries (e.g. ratios, ranks, forecasts, exception scanning, and slicing and dicing) of data within the MDDDB. The presentation module integrates with the application logic module and provides an interface, through which the end users view and request OLAP analyses on their client machines which may be web-enabled through the infrastructure of the Internet. The client/server architecture of a MOLAP system allows multiple users to access the same multidimensional database (MDDDB).

[0026] Information (i.e. basic data) from a variety of operational systems within an enterprise, comprising the Data Warehouse, is loaded into a prior art multidimensional database (MDDDB) through a series of batch routines. The Express™ server by the Oracle Corporation is exemplary of a popular server which can be used to carry out the data loading process in prior art MOLAP systems. As shown in FIG. 2B, an exemplary 3-D MDDDB is schematically depicted, showing geography, time and products as the "dimensions" of the database. The multidimensional data of the MDDDB is logically organized in an array structure, as shown in FIG. 2C. Physically, the Express™ server stores data in pages (or records) of an information file. Pages contain 512, or 2048, or 4096 bytes of data, depending on the platform and release of the Express™ server. In order to look up the physical record address from the database file recorded on a disk or other mass storage device, the Express™ server generates a data structure referred to as a "Page Allocation Table (PAT)". As shown in FIG. 2D, the PAT tells the Express™ server the physical record number that contains the page of data. Typically, the PAT is organized in pages. The simplest way to access a data element in

the MDDB is by calculating the “offset” using the additions and multiplications expressed by a simple formula:

$$\text{Offset} = \text{Months} + \text{Product} * (\# \text{ of_Months} * \# \text{ of_Products}) + \text{City} * (\# \text{ of_Months} * \# \text{ of_Products})$$

[0027] During an OLAP session, the response time of a multidimensional query on a prior art MDDB depends on how many cells in the MDDB have to be added “on the fly”. As the number of dimensions in the MDDB increases linearly, the number of the cells in the MDDB increases exponentially. However, it is known that the majority of multidimensional queries deal with summarized high level data. Thus, as shown in FIGS. 3A and 3B, once the atomic data (i.e. “basic data”) has been loaded into the MDDB, the general approach is to perform a series of calculations in batch in order to aggregate (i.e. pre-aggregate) the data elements along the orthogonal dimensions of the MDDB and fill the array structures thereof. For example, revenue figures for all retail stores in a particular state (i.e. New York) would be added together to fill the state level cells in the MDDB. After the array structure in the database has been filled, integer-based indices are created and hashing algorithms are used to improve query access times. Pre-aggregation of dimension D0 is always performed along the cross-section of the MDDB along the D0 dimension.

[0028] As shown in FIGS. 3C1 and 3C2, the raw data loaded into the MDDB is primarily organized at its lowest dimensional hierarchy, and the results of the pre-aggregations are stored in the neighboring parts of the MDDB.

[0029] As shown in FIG. 3C2, along the TIME dimension, weeks are the aggregation results of days, months are the aggregation results of weeks, and quarters are the aggregation results of months. While not shown in the figures, along the GEOGRAPHY dimension, states are the aggregation results of cities, countries are the aggregation results of states, and continents are the aggregation results of countries. By pre-aggregating (i.e. consolidating or compiling) all logical subtotals and totals along all dimensions of the MDDB, it is possible to carry out real-time MOLAP operations using a multidimensional database (MDDB) containing both basic (i.e. atomic) and pre-aggregated data. Once this compilation process has been completed, the MDDB is ready for use. Users request OLAP reports by submitting queries through the OLAP Application interface (e.g. using web-enabled client machines), and the application logic layer responds to the submitted queries by retrieving the stored data from the MDDB for display on the client machine.

[0030] Typically, in MDDB systems, the aggregated data is very sparse, tending to explode as the number of dimension grows and dramatically slowing down the retrieval process (as described in the report entitled “Database Explosion: The OLAP Report”, incorporated herein by reference). Quick and on line retrieval of queried data is critical in delivering on-line response for OLAP queries. Therefore, the data structure of the MDDB, and methods of its storing, indexing and handling are dictated mainly by the need of fast retrieval of massive and sparse data.

[0031] Different solutions for this problem are disclosed in the following US Patents, each of which is incorporated herein by reference in its entirety:

[0032] U.S. Pat. No. 5,822,751 “Efficient Multidimensional Data Aggregation Operator Implementation”

[0033] U.S. Pat. No. 5,805,885 “Method And System For Aggregation Objects”

[0034] U.S. Pat. No. 5,781,896 “Method And System For Efficiently Performing Database Table Aggregation Using An Aggregation Index”

[0035] U.S. Pat. No. 5,745,764 “Method And System For Aggregation Objects”

[0036] In all the prior art OLAP servers, the process of storing, indexing and handling MDDB utilize complex data structures to largely improve the retrieval speed, as part of the querying process, at the cost of slowing down the storing and aggregation. The query-bounded structure, that must support fast retrieval of queries in a restricting environment of high sparsity and multi-hierarchies, is not the optimal one for fast aggregation.

[0037] In addition to the aggregation process, the Aggregation, Access and Retrieval module is responsible for all data storage, retrieval and access processes. The Logic module is responsible for the execution of OLAP queries. The Presentation module intermediates between the user and the logic module and provides an interface through which the end users view and request OLAP analyses. The client/server architecture allows multiple users to simultaneously access the multidimensional database.

[0038] In summary, general system requirements of OLAP systems include: (1) supporting sophisticated analysis, (2) scaling to large number of dimensions, and (3) supporting analysis against large atomic data sets.

[0039] MOLAP system architecture is capable of providing analytically sophisticated reports and analysis functionality. However, requirements (2) and (3) fundamentally limit MOLAP’s capability, because to be effective and to meet end-user requirements, MOLAP databases need a high degree of aggregation.

[0040] By contrast, the ROLAP system architecture allows the construction of systems requiring a low degree of aggregation, but such systems are significantly slower than systems based on MOLAP system architecture principles. The resulting long aggregation times of ROLAP systems impose severe limitations on its volumes and dimensional capabilities.

[0041] The graphs plotted in FIG. 5 clearly indicate the computational demands that are created when searching an MDDB during an OLAP session, where answers to queries are presented to the MOLAP system, and answers thereto are solicited often under real-time constraints. However, prior art MOLAP systems have limited capabilities to dynamically create data aggregations or to calculate business metrics that have not been precalculated and stored in the MDDB.

[0042] The large volumes of data and the high dimensionality of certain market segmentation applications are orders of magnitude beyond the limits of current multidimensional databases.

[0043] ROLAP is capable of higher data volumes. However, the ROLAP architecture, despite its high volume and dimensionality superiority, suffers from several significant drawbacks as compared to MOLAP:

- [0044] Full aggregation of large data volumes are very time consuming, otherwise, partial aggregation severely degrades the query response.
- [0045] It has a slower query response
- [0046] It requires developers and end users to know SQL
- [0047] SQL is less capable of the sophisticated analytical functionality necessary for OLAP
- [0048] ROLAP provides limited application functionality
- [0049] Thus, improved techniques for data aggregation within MOLAP systems would appear to allow the number of dimensions of and the size of atomic (i.e. basic) data sets in the MDDB to be significantly increased, and thus increase the usage of the MOLAP system architecture.
- [0050] Also, improved techniques for data aggregation within ROLAP systems would appear to allow for maximized query performance on large data volumes, and reduce the time of partial aggregations that degrades query response, and thus generally benefit ROLAP system architectures.
- [0051] Thus, there is a great need in the art for an improved way of and means for aggregating data elements within a multi-dimensional database (MDDB), while avoiding the shortcomings and drawbacks of prior art systems and methodologies.
- [0052] Modern operational and informational database systems, as described above, typically use a database management system (DBMS) (such as an RDBMS system, object database system, or object/relational database system) as a repository for storing data and querying the data. FIG. 14 illustrates a data warehouse-OLAP domain that utilizes the prior art approaches described above. The data warehouse is an enterprise-wide data store. It is becoming an integral part of many information delivery systems because it provides a single, central location wherein a reconciled version of data extracted from a wide variety of operational systems is stored. Details on methods of data integration and constructing data warehouses can be found in the white paper entitled "Data Integration: The Warehouse Foundation" by Louis Rolleigh and Joe Thomas.
- [0053] Building a Data Warehouse has its own special challenges (e.g. using common data model, common business dictionary, etc.) and is a complex endeavor. However, just having a Data Warehouse does not provide organizations with the often-heralded business benefits of data warehousing. To complete the supply chain from transactional systems to decision maker, organizations need to deliver systems that allow knowledge workers to make strategic and tactical decisions based on the information stored in these warehouses. These decision support systems are referred to as On-Line Analytical Processing (OLAP) systems. Such OLAP systems are commonly classified as Relational OLAP systems or Multi-Dimensional OLAP systems as described above.
- [0054] The Relational OLAP (ROLAP) system accesses data stored in a relational database (which is part of the Data Warehouse) to provide OLAP analyses. The premise of ROLAP is that OLAP capabilities are best provided directly

against the relational database. The ROLAP architecture was invented to enable direct access of data from Data Warehouses, and therefore support optimization techniques to meet batch window requirements and provide fast response times. Typically, these optimization techniques include application-level table partitioning, pre-aggregate inferencing, denormalization support, and the joining of multiple fact tables.

[0055] As described above, a typical ROLAP system has a three-tier or layer client/server architecture. The "database layer" utilizes relational databases for data storage, access, and retrieval processes. The "application logic layer" is the ROLAP engine which executes the multidimensional reports from multiple users. The ROLAP engine integrates with a variety of "presentation layers," through which users perform OLAP analyses. After the data model for the data warehouse is defined, data from on-line transaction-processing (OLTP) systems is loaded into the relational database management system (RDBMS). If required by the data model, database routines are run to pre-aggregate the data within the RDBMS. Indices are then created to optimize query access times. End users submit multidimensional analyses to the ROLAP engine, which then dynamically transforms the requests into SQL execution plans. The SQL execution plans are submitted to the relational database for processing, the relational query results are cross-tabulated, and a multidimensional result data set is returned to the end user. ROLAP is a fully dynamic architecture capable of utilizing pre-calculated results when they are available, or dynamically generating results from the raw information when necessary.

[0056] The Multidimensional OLAP (MOLAP) systems utilize a proprietary multidimensional database (MDDB) (or "cube") to provide OLAP analyses. The main premise of this architecture is that data must be stored multidimensionally to be accessed and viewed multidimensionally. Such MOLAP systems provide an interface that enables users to query the MDDB data structure such that users can "slice and dice" the aggregated data. As shown in FIG. 15, such MOLAP systems have an aggregation engine which is responsible for all data storage, access, and retrieval processes, including data aggregation (i.e. pre-aggregation) in the MDDB, and an analytical processing and GUI module responsible for interfacing with a user to provide analytical analysis, query input, and reporting of query results to the user. In a relational database, data is stored in tables. In contrast, the MDDB is a non-relational data structure—it uses other data structures, either instead of or in addition to tables—to store data.

[0057] There are other application domains where there is a great need for improved methods of and apparatus for carrying out data aggregation operations. For example, modern operational and informational databases represent such domains. As described above, modern operational and informational databases typically utilize a relational database system (RDBMS) as a repository for storing data and querying data. FIG. 16A illustrates an exemplary table in an RDBMS; and FIGS. 16B and 16C illustrate operators (queries) on the table of FIG. 16A, and the result of such queries, respectively. The operators illustrated in FIGS. 16B and 16C are expressed as Structured Query Language (SQL) statements as is conventional in the art.

[0058] The choice of using a RDBMS as the data repository in information database systems naturally stems from the realities of SQL standardization, the wealth of RDBMS-related tools, and readily available expertise in RDBMS systems. However, the querying component of RDBMS technology suffers from performance and optimization problems stemming from the very nature of the relational data model. More specifically, during query processing, the relational data model requires a mechanism that locates the raw data elements that match the query. Moreover, to support queries that involve aggregation operations, such aggregation operations must be performed over the raw data elements that match the query. For large multi-dimensional databases, a naive implementation of these operations involves computational intensive table scans that leads to unacceptable query response times.

[0059] In order to better understand how the prior art has approached this problem, it will be helpful to briefly describe the relational database model. According to the relational database model, a relational database is represented by a logical schema and tables that implement the schema. The logical schema is represented by a set of templates that define one or more dimensions (entities) and attributes associated with a given dimension. The attributes associated with a given dimension includes one or more attributes that distinguish it from every other dimension in the database (a dimension identifier). Relationships amongst dimensions are formed by joining attributes. The data structure that represents the set of templates and relations of the logical schema is typically referred to as a catalog or dictionary. Note that the logical schema represents the relational organization of the database, but does not hold any fact data per se. This fact data is stored in tables that implement the logical schema.

[0060] Star schemas are frequently used to represent the logical structure of a relational database. The basic premise of star schemas is that information can be classified into two groups: facts and dimensions. Facts are the core data elements being analyzed. For example, units of individual item sold are facts, while dimensions are attributes about the facts. For example, dimensions are the product types purchased and the data purchase. Business questions against this schema are asked looking up specific facts (UNITS) through a set of dimensions (MARKETS, PRODUCTS, PERIOD). The central fact table is typically much larger than any of its dimension tables.

[0061] An exemplary star schema is illustrated in FIG. 17A for suppliers (the "Supplier" dimension) and parts (the "Parts" dimension) over time periods (the "Time-Period" dimension). It includes a central fact table "Supplied-Parts" that relates to multiple dimensions—the "Supplier", "Parts" and "Time-Period" dimensions. A central fact table and a dimension table for each dimension in the logical schema of FIG. 17A may be used to implement this logical schema. A given dimension table stores rows (instances) of the dimension defined in the logical schema. Each row within the central fact table includes a multi-part key associated with a set of facts (in this example, a number representing a quantity). The multi-part key of a given row (values stored in the S#,P#,TP# fields as shown) points to rows (instances) stored in the dimension tables described above. A more detailed description of star schemas and the tables used to implement star schemas may be found in C. J. Date, "An

Introduction to Database Systems," Seventh Edition, Addison-Wesley, 2000, pp. 711-715, herein incorporated by reference in its entirety.

[0062] When processing a query, the tables that implement the schema are accessed to retrieve the facts that match the query. For example, in a star schema implementation as described above, the facts are retrieved from the central fact table and/or the dimension tables. Locating the facts that match a given query involves one or more join operations. Moreover, to support queries that involve aggregation operations, such aggregation operations must be performed over the facts that match the query. For large multi-dimensional databases, a naive implementation of these operations involves computational intensive table scans that typically leads to unacceptable query response times. Moreover, since the fact tables are pre-summarized and aggregated along business dimensions, these tables tend to be very large. This point becomes an important consideration of the performance issues associated with star schemas. A more detailed discussion of the performance issues (and proposed approaches that address such issues) related to joining and aggregation of star schema is now set forth.

[0063] The first performance issue arises from computationally intensive table scans that are performed by a naive implementation of data joining. Indexing schemes may be used to bypass these scans when performing joining operations. Such schemes include B-tree indexing, inverted list indexing and aggregate indexing. A more detailed description of such indexing schemes can be found in "The Art of Indexing", Dynamic Information Systems Corporation, October 1999. All of these indexing schemes replaces table scan operations (involved in locating the data elements that match a query) with one or more index lookup operation. Inverted list indexing associates an index with a group of data elements, and stores (at a location identified by the index) a group of pointers to the associated data elements. During query processing, in the event that the query matches the index, the pointers stored in the index are used to retrieve the corresponding data elements pointed therefrom. Aggregation indexing integrates an aggregation index with an inverted list index to provide pointers to raw data elements that require aggregation, thereby providing for dynamic summarization of the raw data elements that match the user-submitted query.

[0064] These indexing schemes are intended to improve join operations by replacing table scan operations with one or more index lookup operation in order to locate the data elements that match a query. However, these indexing schemes suffer from various performance issues as follows:

[0065] Since the tables in the star schema design typically contain the entire hierarchy of attributes (e.g. in a PERIOD dimension, this hierarchy could be day>week>month>quarter>year), a multipart key of day, week, month, quarter, year has to be created; thus, multiple meta-data definitions are required (one of each key component) to define a single relationship; this adds to the design complexity, and sluggishness in performance.

[0066] Addition or deletion of levels in the hierarchy will require physical modification of the fact table, which is time consuming process that limits flexibility.

[0067] Carrying all the segments of the compound dimensional key in the fact table increases the size of the index, thus impacting both performance and scalability.

[0068] Another performance issue arises from dimension tables that contain multiple hierarchies. In such cases, the dimensional table often includes a level of hierarchy indicator for every record. Every retrieval from fact table that stores details and aggregates must use the indicator to obtain the correct result, which impacts performance. The best alternative to using the level indicator is the snowflake schema. In this schema aggregate tables are created separately from the detail tables. In addition to the main fact tables, snowflake schema contains separate fact tables for each level of aggregation. Notably, the snowflake schema is even more complicated than a star schema, and often requires multiple SQL statements to get the results that are required.

[0069] Another performance issue arises from the pairwise join problem. Traditional RDBMS engines are not design for the rich set of complex queries that are issued against a star schema. The need to retrieve related information from several tables in a single query—"join processing"—is severely limited. Many RDBMSs can join only two tables at a time. If a complex join involves more than two tables, the RDBMS needs to break the query into a series of pairwise joins. Selecting the order of these joins has a dramatic performance impact. There are optimizers that spend a lot of CPU cycles to find the best order in which to execute those joins. Unfortunately, because the number of combinations to be evaluated grows exponentially with the number of tables being joined, the problem of selecting the best order of pairwise joins rarely can be solved in a reasonable amount of time.

[0070] Moreover, because the number of combinations is often too large, optimizers limit the selection on the basis of a criterion of directly related tables. In a star schema, the fact table is the only table directly related to most other tables, meaning that the fact table is a natural candidate for the first pairwise join. Unfortunately, the fact table is the very largest table in the query, so this strategy leads to selecting a pairwise join order that generates a very large intermediate result set, severely affecting query performance.

[0071] There is an optimization strategy, typically referred to as Cartesian Joins, that lessens the performance impact of the pairwise join problem by allowing joining of unrelated tables. The join to the fact table, which is the largest one, is deferred until the very end, thus reducing the size of intermediate result sets. In a join of two unrelated tables every combination of the two tables' rows is produced, a Cartesian product. Such a Cartesian product improves query performance. However, this strategy is viable only if the Cartesian product of dimension rows selected is much smaller than the number of rows in the fact table. The multiplicative nature of the Cartesian join makes the optimization helpful only for relatively small databases.

[0072] In addition, systems that exploit hardware and software parallelism have been developed that lessens the performance issues set forth above. Parallelism can help reduce the execution time of a single query (speed-up), or handle additional work without degrading execution time (scale-up). For example, Red Brick™ has developed STAR-

join™ technology that provides high speed, parallelizable multi-table joins in a single pass, thus allowing more than two tables can be joined in a single operation. The core technology is an innovative approach to indexing that accelerates multiple joins. Unfortunately, parallelism can only reduce, not eliminate, the performance degradation issues related to the star schema.

[0073] One of the most fundamental principles of the multidimensional database is the idea of aggregation. The most common aggregation is called a roll-up aggregation. This type is relatively easy to compute: e.g. taking daily sales totals and rolling them up into a monthly sales table. The more difficult are analytical calculations, the aggregation of Boolean and comparative operators. However these are also considered as a subset of aggregation.

[0074] In a star schema, the results of aggregation are summary tables. Typically, summary tables are generated by database administrators who attempt to anticipate the data aggregations that the users will request, and then pre-build such tables. In such systems, when processing a user-generated query that involves aggregation operations, the pre-built aggregated data that matches the query is retrieved from the summary tables (if such data exists). FIGS. 18A and 18B illustrate a multi-dimensional relational database using a star schema and summary tables. In this example, the summary tables are generated over the "time" dimension storing aggregated data for "month", "quarter" and "year" time periods as shown in FIG. 18B. Summary tables are in essence additional fact tables, of higher levels. They are attached to the basic fact table creating a snowflake extension of the star schema. There are hierarchies among summary tables because users at different levels of management require different levels of summarization. Choosing the level of aggregation is accomplished via the "drill-down" feature.

[0075] Summary tables containing pre-aggregated results typically provide for improved query response time with respect to on-the-fly aggregation. However, summary tables suffer from some disadvantages:

[0076] summary tables require that database administrators anticipate the data aggregation operations that users will require; this is a difficult task in large multi-dimensional databases (for example, in data warehouses and data mining systems), where users always need to query in new ways looking for new information and patterns.

[0077] summary tables do not provide a mechanism that allows efficient drill down to view the raw data that makes up the summary table—typically a table scan of one or more large tables is required.

[0078] querying is delayed until pre-aggregation is completed.

[0079] there is a heavy time overhead because the vast majority of the generated information remains unvisited.

[0080] there is a need to synchronize the summary tables before the use.

[0081] the degree of viable parallelism is limited because the subsequent levels of summary tables must be performed in pipeline, due to their hierarchies.

[0082] for very large databases, this option is not valid because of time and storage space.

Note that it is common to utilize both pre-aggregated results and on-the-fly aggregation in support aggregation. In these system, partial pre-aggregation of the facts results in a small set of summary tables. On-the-fly aggregation is used in the case the required aggregated data does not exist in the summary tables.

[0083] Note that in the event that the aggregated data does not exist in the summary tables, table join operations and aggregation operations are performed over the raw facts in order to generate such aggregated data. This is typically referred to as on-the-fly aggregation. In such instances, aggregation indexing is used to mitigate the performance of multiple data joins associated with dynamic aggregation of the raw data. Thus, in large multi-dimensional databases, such dynamic aggregation may lead to unacceptable query response times.

[0084] In view of the problems associated with joining and aggregation within RDBMS, prior art ROLAP systems have suffered from essentially the same shortcomings and drawbacks of their underlying RDBMS.

[0085] While prior art MOLAP systems provide for improved access time to aggregated data within their underlying MDD structures, and have performance advantages when carrying out joining and aggregations operations, prior art MOLAP architectures have suffered from a number of shortcomings and drawbacks. More specifically, atomic (raw) data is moved, in a single transfer, to the MOLAP system for aggregation, analysis and querying. Importantly, the aggregation results are external to the DBMS. Thus, users of the DBMS cannot directly view these results. Such results are accessible only from the MOLAP system. Because the MDD query processing logic in prior art MOLAP systems is separate from that of the DBMS, users must procure rights to access to the MOLAP system and be instructed (and be careful to conform to such instructions) to access the MDD (or the DBMS) under certain conditions. Such requirements can present security issues, highly undesirable for system administration. Satisfying such requirements is a costly and logistically cumbersome process. As a result, the widespread applicability of MOLAP systems has been limited.

[0086] Thus, there is a great need in the art for an improved mechanism for joining and aggregating data elements within a database management system (e.g., RDBMS), and for integrating the improved database management system (e.g., RDBMS) into informational database systems (including the data warehouse and OLAP domains), while avoiding the shortcomings and drawbacks of prior art systems and methodologies.

SUMMARY AND OBJECTS OF PRESENT INVENTION

[0087] Accordingly, it is a further object of the present invention to provide an improved method of and system for managing data elements within a multidimensional database (MDDB) using a novel stand-alone (i.e. external) data aggregation server, achieving a significant increase in system performance (e.g. decreased access/search time) using a stand-alone scalable data aggregation server.

[0088] Another object of the present invention is to provide such system, wherein the stand-alone aggregation server includes an aggregation engine that is integrated with an MDDB, to provide a cartridge-style plug-in accelerator which can communicate with virtually any conventional OLAP server.

[0089] Another object of the present invention is to provide such a stand-alone data aggregation server whose computational tasks are restricted to data aggregation, leaving all other OLAP functions to the MOLAP server and therefore complementing OLAP server's functionality.

[0090] Another object of the present invention is to provide such a system, wherein the stand-alone aggregation server carries out an improved method of data aggregation within the MDDB which enables the dimensions of the MDDB to be scaled up to large numbers and large atomic (i.e. base) data sets to be handled within the MDDB.

[0091] Another object of the present invention is to provide such a stand-alone aggregation server, wherein the aggregation engine supports high-performance aggregation (i.e. data roll-up) processes to maximize query performance of large data volumes, and to reduce the time of partial aggregations that degrades the query response.

[0092] Another object of the present invention is to provide such a stand-alone, external scalable aggregation server, wherein its integrated data aggregation (i.e. roll-up) engine speeds up the aggregation process by orders of magnitude, enabling larger database analysis by lowering the aggregation times.

[0093] Another object of the present invention is to provide such a novel stand-alone scalable aggregation server for use in OLAP operations, wherein the scalability of the aggregation server enables (i) the speed of the aggregation process carried out therewithin to be substantially increased by distributing the computationally intensive tasks associated with data aggregation among multiple processors, and (ii) the large data sets contained within the MDDB of the aggregation server to be subdivided among multiple processors thus allowing the size of atomic (i.e. basic) data sets within the MDDB to be substantially increased.

[0094] Another object of the present invention is to provide such a novel stand-alone scalable aggregation server, which provides for uniform load balancing among processors for high efficiency and best performance, and linear scalability for extending the limits by adding processors.

[0095] Another object of the present invention is to provide a stand-alone, external scalable aggregation server, which is suitable for MOLAP as well as for ROLAP system architectures.

[0096] Another object of the present invention is to provide a novel stand-alone scalable aggregation server, wherein an MDDB and aggregation engine are integrated and the aggregation engine carries out a high-performance aggregation algorithm and novel storing and searching methods within the MDDB.

[0097] Another object of the present invention is to provide a novel stand-alone scalable aggregation server which can be supported on single-processor (i.e. sequential or serial) computing platforms, as well as on multi-processor (i.e. parallel) computing platforms.

[0098] Another object of the present invention is to provide a novel stand-alone scalable aggregation server which can be used as a complementary aggregation plug-in to existing MOLAP and ROLAP databases.

[0099] Another object of the present invention is to provide a novel stand-alone scalable aggregation server which carries out a novel rollup (i.e. down-up) and spread down (i.e. top-down) aggregation algorithms.

[0100] Another object of the present invention is to provide a novel stand-alone scalable aggregation server which includes an integrated MDDB and aggregation engine which carries out full pre-aggregation and/or “on-the-fly” aggregation processes within the MDDB.

[0101] Another object of the present invention is to provide such a novel stand-alone scalable aggregation server which is capable of supporting MDDB having a multi-hierarchy dimensionality.

[0102] Another object of the present invention is to provide a novel method of aggregating multidimensional data of atomic data sets originating from a RDBMS Data Warehouse.

[0103] Another object of the present invention is to provide a novel method of aggregating multidimensional data of atomic data sets originating from other sources, such as external ASCII files, MOLAP server, or other end user applications.

[0104] Another object of the present invention is to provide a novel stand-alone scalable data aggregation server which can communicate with any MOLAP server via standard ODBC, OLE DB or DLL interface, in a completely transparent manner with respect to the (client) user, without any time delays in queries, equivalent to storage in MOLAP server's cache.

[0105] Another object of the present invention is to provide a novel “cartridge-style” (stand-alone) scalable data aggregation engine which dramatically expands the boundaries of MOLAP into large-scale applications including Banking, Insurance, Retail and Promotion Analysis.

[0106] Another object of the present invention is to provide a novel “cartridge-style” (stand-alone) scalable data aggregation engine which dramatically expands the boundaries of high-volatility type ROLAP applications such as, for example, the precalculation of data to maximize query performance.

[0107] Another object of the present invention is to provide a generic plug-in cartridge-type data aggregation component, suitable for all MOLAP systems of different vendors, dramatically reducing their aggregation burdens.

[0108] Another object of the present invention is to provide a novel high performance cartridge-type data aggregation server which, having standardized interfaces, can be plugged-into the OLAP system of virtually any user or vendor.

[0109] Another object of the present invention is to provide a novel “cartridge-style” (stand-alone) scalable data aggregation engine which has the capacity to convert long batch-type data aggregations into interactive sessions.

[0110] In another aspect, it is an object of the present invention to provide an improved method of and system for joining and aggregating data elements integrated within a database management system (DBMS) using a non-relational multi-dimensional data structure (MDDB), achieving a significant increase in system performance (e.g. decreased access/search time), user flexibility and ease of use.

[0111] Another object of the present invention is to provide such an DBMS wherein its integrated data aggregation module supports high-performance aggregation (i.e. data roll-up) processes to maximize query performance of large data volumes.

[0112] Another object of the present invention is to provide such an DBMS system, wherein its integrated data aggregation (i.e. roll-up) module speeds up the aggregation process by orders of magnitude, enabling larger database analysis by lowering the aggregation times.

[0113] Another object of the present invention is to provide such a novel DBMS system for use in OLAP operations.

[0114] Another object of the present invention is to provide a novel DBMS system having an integrated aggregation module that carries out an novel rollup (i.e. down-up) and spread down (i.e. top-down) aggregation algorithms.

[0115] Another object of the present invention is to provide a novel DBMS system having an integrated aggregation module that carries out full pre-aggregation and/or “on-the-fly” aggregation processes.

[0116] Another object of the present invention is to provide a novel DBMS system having an integrated aggregation module which is capable of supporting a MDDB having a multi-hierarchy dimensionality.

[0117] These and other object of the present invention will become apparent hereinafter and in the Claims to Invention set forth herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0118] In order to more fully appreciate the objects of the present invention, the following Detailed Description of the Illustrative Embodiments should be read in conjunction with the accompanying Drawings, wherein:

[0119] FIG. 1A is a schematic representation of an exemplary prior art relational on-line analytical processing (ROLAP) system comprising a three-tier or layer client/server architecture, wherein the first tier has a database layer utilizing an RDBMS for data storage, access, and retrieval processes, the second tier has an application logic layer (i.e. the ROLAP engine) for executing the multidimensional reports from multiple users, and the third tier integrates the ROLAP engine with a variety of presentation layers, through which users perform OLAP analyses;

[0120] FIG. 1B is a schematic representation of a generalized embodiment of a prior art multidimensional on-line analytical processing (MOLAP) system comprising a base data loader for receiving atomic (i.e. base) data from a Data Warehouse realized by a RDBMS, an OLAP multidimensional database (MDDB), an aggregation, access and retrieval module, application logic module and presentation module associated with a conventional OLAP sever (e.g. Oracle's

Express Server) for supporting on-line transactional processing (OLTP) operations on the MDDDB, to service database queries and requests from a plurality of OLAP client machines typically accessing the system from an information network (e.g. the Internet);

[0121] FIG. 2A is a schematic representation of the Data Warehouse shown in the prior art system of FIG. 1B comprising numerous data tables (e.g. T1, T2 . . . Tn) and data field links, and the OLAP multidimensional database shown of FIG. 1B, comprising a conventional page allocation table (PAT) with pointers pointing to the physical storage of variables in an information storage device;

[0122] FIG. 2B is a schematic representation of an exemplary three-dimensional MDDDB and organized as a 3-dimensional Cartesian cube and used in the prior art system of FIG. 2A, wherein the first dimension of the MDDDB is representative of geography (e.g. cities, states, countries, continents), the second dimension of the MDDDB is representative of time (e.g. days, weeks, months, years), the third dimension of the MDDDB is representative of products (e.g. all products, by manufacturer), and the basic data element is a set of variables which are addressed by 3-dimensional coordinate values;

[0123] FIG. 2C is a schematic representation of a prior art array structure associated with an exemplary three-dimensional MDDDB, arranged according to a dimensional hierarchy;

[0124] FIG. 2D is a schematic representation of a prior art page allocation table for an exemplary three-dimensional MDDDB, arranged according to pages of data element addresses;

[0125] FIG. 3A is a schematic representation of a prior art MOLAP system, illustrating the process of periodically storing raw data in the RDBMS Data Warehouse thereof, serially loading of basic data from the Data Warehouse to the MDDDB, and the process of serially pre-aggregating (or pre-compiling) the data in the MDDDB along the entire dimensional hierarchy thereof;

[0126] FIG. 3B is a schematic representation illustrating that the Cartesian addresses listed in a prior art page allocation table (PAT) point to where physical storage of data elements (i.e. variables) occurs in the information recording media (e.g. storage volumes) associated with the MDDDB, during the loading of basic data into the MDDDB as well as during data preaggregation processes carried out there-within;

[0127] FIG. 3C1 is a schematic representation of an exemplary three-dimensional database used in a conventional MOLAP system of the prior art, showing that each data element contained therein is physically stored at a location in the recording media of the system which is specified by the dimensions (and subdimensions within the dimensional hierarchy) of the data variables which are assigned integer-based coordinates in the MDDDB, and also that data elements associated with the basic data loaded into the MDDDB are assigned lower integer coordinates in MDDDB Space than pre-aggregated data elements contained therewithin;

[0128] FIG. 3C2 is a schematic representation illustrating that a conventional hierarchy of the dimension of "time"

typically contains the subdimensions "days, weeks, months, quarters, etc." of the prior art;

[0129] FIG. 3C3 is a schematic representation showing how data elements having higher subdimensions of time in the MDDDB of the prior art are typically assigned increased integer addresses along the time dimension thereof;

[0130] FIG. 4 is a schematic representation illustrating that, for very large prior art MDDDBs, very large page allocation tables (PATs) are required to represent the address locations of the data elements contained therein, and thus there is a need to employ address data paging techniques between the DRAM (e.g. program memory) and mass storage devices (e.g. recording discs or RAIDs) available on the serial computing platform used to implement such prior art MOLAP systems;

[0131] FIG. 5 is a graphical representation showing how search time in a conventional (i.e. prior art) MDDDB increases in proportion to the amount of preaggregation of data therewithin;

[0132] FIG. 6A is a schematic representation of a generalized embodiment of a multidimensional on-line analytical processing (MOLAP) system of the present invention comprising a Data Warehouse realized as a relational database, a stand-alone Aggregation Server of the present invention having an integrated aggregation engine and MDDDB, and an OLAP server supporting a plurality of OLAP clients, wherein the stand-alone Aggregation Server performs aggregation functions (e.g. summation of numbers, as well as other mathematical operations, such as multiplication, subtraction, division etc.) and multi-dimensional data storage functions;

[0133] FIG. 6B is a schematic block diagram of the stand-alone Aggregation Server of the illustrative embodiment shown in FIG. 6A, showing its primary components, namely, a base data interface (e.g. OLDB, OLE-DB, ODBC, SQL, JDBC, API, etc.) for receiving RDBMS flat files lists and other files from the Data Warehouse (RDBMS), a base data loader for receiving base data from the base data interface, configuration manager for managing the operation of the base data interface and base data loader, an aggregation engine and MDDDB handler for receiving base data from the base loader, performing aggregation operations on the base data, and storing the base data and aggregated data in the MDDDB; an aggregation client interface (e.g. OLDB, OLE-DB, ODBC, SQL, JDBC, API, etc.) and input analyzer for receiving requests from OLAP client machines, cooperating with the aggregation engine and MDDDB handler to generate aggregated data and/or retrieve aggregated data from the MDDDB that pertains to the received requests, and returning this aggregated back to the requesting OLAP clients; and a configuration manager for managing the operation of the input analyzer and the aggregation client interface.

[0134] FIG. 6C is a schematic representation of the software modules comprising the aggregation engine and MDDDB handler of the stand-alone Aggregation Server of the illustrative embodiment of the present invention, showing a base data list structure being supplied to a hierarchy analysis and reorder module, the output thereof being transferred to an aggregation management module, the output thereof being transferred to a storage module via a storage manage-

ment module, and a Query Directed Roll-up (QDR) aggregation management module being provided for receiving database (DB) requests from OLAP client machines (via the aggregation client interface) and managing the operation of the aggregation and storage management modules of the present invention;

[0135] FIG. 6D is a flow chart representation of the primary operations carried out by the (DB) request serving mechanism within the QDR aggregation management module shown in FIG. 6C;

[0136] FIG. 7A is a schematic representation of a separate-platform type implementation of the stand-alone Aggregation Server of the illustrative embodiment of FIG. 6B and a conventional OLAP server supporting a plurality of client machines, wherein base data from a Data Warehouse is shown being received by the aggregation server, realized on a first hardware/software platform (i.e. Platform A) and the stand-alone Aggregation Server is shown serving the conventional OLAP server, realized on a second hardware/software platform (i.e. Platform B), as well as serving data aggregation requirements of other clients supporting diverse applications such as spreadsheet, GUI front end, and applications;

[0137] FIG. 7B is a schematic representation of a shared-platform type implementation of the stand-alone Aggregation Server of the illustrative embodiment of FIG. 6B and a conventional OLAP server supporting a plurality of client machines, wherein base data from a Data Warehouse is shown being received by the stand-alone Aggregation Server, realized on a common hardware/software platform and the aggregation server is shown serving the conventional OLAP server, realized on the same common hardware/software platform, as well as serving data aggregation requirements of other clients supporting diverse applications such as spreadsheet, GUI front end, and applications;

[0138] FIG. 8A is a data table setting forth information representative of performance benchmarks obtained by the shared-platform type implementation of the stand-alone Aggregation Server of the illustrative embodiment serving the conventional OLAP server (i.e. Oracle EXPRESS Server) shown in FIG. 7B, wherein the common hardware/software platform is realized using a Pentium II 450 Mhz, 1 GB RAM, 18 GB Disk, running the Microsoft NT operating system (OS);

[0139] FIG. 9A is a schematic representation of the first stage in the method of segmented aggregation according to the principles of the present invention, showing initial aggregation along the 1st dimension;

[0140] FIG. 9B is a schematic representation of the next stage in the method of segmented aggregation according to the principles of the present invention, showing that any segment along dimension 1, such as the shown slice, can be separately aggregated along the remaining dimensions, 2 and 3, and that in general, for an N dimensional system, the second stage involves aggregation in N-1 dimensions. The principle of segmentation can be applied on the first stage as well, however, only a large enough data will justify such a sliced procedure in the first dimension. Actually, it is possible to consider each segment as an N-1 cube, enabling recursive computation.

[0141] FIG. 9C1 is a schematic representation of the Query Directed Roll-up (QDR) aggregation method/proce-

dures of the present invention, showing data aggregation starting from existing basic data or previously aggregated data in the first dimension (D1), and such aggregated data being utilized as a basis for QDR aggregation along the second dimension (D2);

[0142] FIG. 9C2 is a schematic representation of the Query Directed Roll-up (QDR) aggregation method/procedure of the present invention, showing initial data aggregation starting from existing previously aggregated data in the second third (D3), and continuing along the third dimension (D3), and thereafter continuing aggregation along the second dimension (D2);

[0143] FIG. 10A is a schematic representation of the "slice-storage" method of storing sparse data in the disk storage devices of the MDDB of FIG. 6B in accordance with the principles of the present invention, based on an ascending-ordered index along aggregation direction, enabling fast retrieval of data;

[0144] FIG. 10B is a schematic representation of the data organization of data files and the directory file used in the storages of the MDDB of FIG. 6B, and the method of searching for a queried data point therein using a simple binary search technique due to the data files ascending order;

[0145] FIG. 11A is a schematic representation of three exemplary multi-hierarchical data structures for storage of data within the MDDB of FIG. 6B, having three levels of hierarchy, wherein the first level representative of base data is composed of items A,B,F, and G, the second level is composed of items C,E,H and I, and the third level is composed of a single item D, which is common to all three hierarchical structures;

[0146] FIG. 11B is a schematic representation of an optimized multi-hierarchical data structure merged from all three hierarchies of FIG. 11A, in accordance with the principles of the present invention;

[0147] FIG. 11C(i) through 11C(ix) represent a flow chart description (and accompanying data structures) of the operations of an exemplary hierarchy transformation mechanism of the present invention that optimally merges multiple hierarchies into a single hierarchy that is functionally equivalent to the multiple hierarchies.

[0148] FIG. 12 is a schematic representation showing the levels of operations performed by the stand-alone Aggregation Server of FIG. 6B, summarizing the different enabling components for carrying out the method of segmented aggregation in accordance with the principles of the present invention;

[0149] FIG. 13 is a schematic representation of the stand-alone Aggregation Server of the present invention shown as a component of a central data warehouse, serving the data aggregation needs of URL directory systems, Data Marts, RDBMSs, ROLAP systems and OLAP systems alike;

[0150] FIG. 14 is a schematic representation of a prior art information database system, wherein the present invention may be embodied;

[0151] FIG. 15 is a schematic representation of the prior art data warehouse and OLAP system, wherein the present invention may be embodied;

[0152] FIGS. 16A-16C are schematic representations of exemplary tables employed in a prior art Relational Database Management System (RDBMS); FIGS. 16B and 16C illustrate operators (queries) on the table of FIG. 16A, and the result of such queries, respectively;

[0153] FIG. 17A is a schematic representation of an exemplary dimensional schema (star schema) of a relational database;

[0154] FIG. 18A is a schematic representation of an exemplary multidimensional schema (star schema);

[0155] FIG. 18B is a schematic representation of tables used to implement the schema of FIG. 18A, including summary tables storing results of aggregation operations performed on the facts of the central fact table along the time-period dimension, in accordance with conventional teachings;

[0156] FIG. 19A is a schematic representation of an exemplary embodiment of a DBMS (for example, an RDBMS as shown) of the present invention comprising a relational datastore and an integrated multidimensional (MDD) aggregation module supporting queries from a plurality of clients, wherein the aggregation engine performs aggregation functions (e.g. summation of numbers, as well as other mathematical operations, such as multiplication, subtraction, division etc.) and non-relational multi-dimensional data storage functions.

[0157] FIG. 19B is a schematic block diagram of the MDD aggregation module of the illustrative embodiment of the present invention shown in FIG. 6A.

[0158] FIGS. 19C(i) and 19C(ii), taken together, set forth a flow chart representation of the primary operations carried out within the DBMS of the present invention when performing data aggregation and related support operations, including the servicing of user-submitted (e.g. natural language) queries made on such aggregated database of the present invention.

[0159] FIG. 19D is a flow chart representation of the primary operations carried out by the (DB) request serving mechanism within the MDD control module shown in FIG. 6B.

[0160] FIG. 19E is a schematic representation of the view mechanism of an DBMS that enables users to query on the aggregated data generated and/or stored in the MDD Aggregation module according to the present invention.

[0161] FIG. 19F is a schematic representation of the trigger mechanism of the DBMS that enables users to query on the aggregated data generated and/or stored in the MDD Aggregation module according to the present invention.

[0162] FIG. 19G is a schematic representation of the DBMS of the present invention, illustrating a logically partitioning into a relational part and a non-relational part. The relational part includes the relational data store (e.g., table(s) and dictionary) and support mechanisms (e.g., query handling services). The non-relational part includes the MDD Aggregation Module. Data flows bidirectionally between the relational part and the non-relational part as shown.

[0163] FIG. 20A shows a separate-platform type implementation of the DBMS system of the illustrative embodi-

ment shown in FIG. 19A, wherein the relational datastore and support mechanisms (e.g., query handling, fact table(s) and dictionary of the DBMS) reside on a separate hardware platform and/or OS system from that used to run the MDD Aggregation Module of the present invention.

[0164] FIG. 20B shows a common-platform type implementation of the DBMS system of the illustrative embodiment shown in FIG. 19A, wherein the relational datastore and support mechanisms (e.g., query handling, fact table(s) and dictionary of the DBMS) share the same hardware platform and operating system (OS) that is used to run the MDD Aggregation Module of the present invention.

[0165] FIG. 21 is a schematic representation of the DBMS of the present invention shown as a component of a central data warehouse, serving the data storage and aggregation needs of a ROLAP system (or other OLAP system).

[0166] FIG. 22 is a schematic representation of the DBMS of the present invention shown as a component of a central data warehouse, wherein the DBMS includes integrated OLAP Analysis Logic (and preferably an integrated Presentation Module) that operates cooperatively with the query handling of the DBMS system and the MDD Aggregation Module to enable users of the DBMS system to execute multidimensional reports (e.g., ratios, ranks, transforms, dynamic consolidation, complex filtering, forecasts, query governing, scheduling, flow control, pre-aggregate inferencing, denormalization support, and/or table partitioning and joins) and preferably perform traditional OLAP analyses (grids, graphs, maps, alerts, drill-down, data pivot, data surf, slice and dice, print).

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS OF THE PRESENT INVENTION

[0167] Referring now to FIGS. 6A through 13, the preferred embodiments of the method and system of the present invention will be now described in great detail hereinbelow, wherein like elements in the Drawings shall be indicated by like reference numerals.

[0168] Through this invention disclosure, the term "aggregation" and "preaggregation" shall be understood to mean the process of summation of numbers, as well as other mathematical operations, such as multiplication, subtraction, division etc.

[0169] In general, the stand-alone aggregation server and methods of and apparatus for data aggregation of the present invention can be employed in a wide range of applications, including MOLAP systems, ROLAP systems, Internet URL-directory systems, personalized on-line e-commerce shopping systems, Internet-based systems requiring real-time control of packet routing and/or switching, and the like.

[0170] For purposes of illustration, initial focus will be accorded to improvements in MOLAP systems, in which knowledge workers are enabled to intuitively, quickly, and flexibly manipulate operational data within a MDDB using familiar business terms in order to provide analytical insight into a business domain of interest.

[0171] FIG. 6A illustrates a generalized embodiment of a multidimensional on-line analytical processing (MOLAP) system of the present invention comprising: a Data Ware-

house realized as a relational database; a stand-alone cartridge-style Aggregation Server of the present invention having an integrated aggregation engine and a MDDB; and an OLAP server communicating with the Aggregation Server, and supporting a plurality of OLAP clients. In accordance with the principles of the present invention, the stand-alone Aggregation Server performs aggregation functions (e.g. summation of numbers, as well as other mathematical operations, such as multiplication, subtraction, division etc.) and multi-dimensional data storage functions.

[0172] Departing from conventional practices, the principles of the present invention teaches moving the aggregation engine and the MDDB into a separate Aggregation Server having standardized interfaces so that it can be plugged-into the OLAP server of virtually any user or vendor. This dramatic move discontinues the restricting dependency of aggregation from the analytical functions of OLAP, and by applying novel and independent algorithms. The stand-alone data aggregation server enables efficient organization and handling of data, fast aggregation processing, and fast access to and retrieval of any data element in the MDDB.

[0173] As will be described in greater detail hereinafter, the Aggregation Server of the present invention can serve the data aggregation requirements of other types of systems besides OLAP systems such as, for example, URL directory management Data Marts, RDBMS, or ROLAP systems.

[0174] The Aggregation Server of the present invention excels in performing two distinct functions, namely: the aggregation of data in the MDDB; and the handling of the resulting data base in the MDDB, for "on demand" client use. In the case of serving an OLAP server, the Aggregation Server of the present invention focuses on performing these two functions in a high performance manner (i.e. aggregating and storing base data, originated at the Data Warehouse, in a multidimensional storage (MDDB), and providing the results of this data aggregation process "on demand" to the clients, such as the OLAP server, spreadsheet applications, the end user applications. As such, the Aggregation Server of the present invention frees each conventional OLAP server, with which it interfaces, from the need of making data aggregations, and therefore allows the conventional OLAP server to concentrate on the primary functions of OLAP servers, namely: data analysis and supporting a graphical interface with the user client.

[0175] FIG. 6B shows the primary components of the stand-alone Aggregation Server of the illustrative embodiment, namely: a base data interface (e.g. OLDB, OLE-DB, ODBC, SQL, JDBC, API, etc.) for receiving RDBMS flat files lists and other files from the Data Warehouse (RDBMS), a base data loader for receiving base data from the base data interface, configuration manager for managing the operation of the base data interface and base data loader, an aggregation engine for receiving base data from the base loader, a multi-dimensional database (MDDB); a MDDB handler, an input analyzer, an aggregation client interface (e.g. OLDB, OLE-DB, ODBC, SQL, API, JDBC, etc.) and a configuration manager for managing the operation of the input analyzer and the aggregation client interface.

[0176] During operation, the base data originates at data warehouse or other sources, such as external ASCII files, MOLAP server, or others. The Configuration Manager, in

order to enable proper communication with all possible sources and data structures, configures two blocks, the Base Data Interface and Data Loader. Their configuration is matched with different standards such as OLDB, OLE-DB, ODBC, SQL, API, JDBC, etc.

[0177] As shown in FIG. 6B, the core of the data Aggregation Server of the present invention comprises: a data Aggregation Engine; a Multidimensional Data Handler (MDDB Handler); and a Multidimensional Data Storage (MDDB). The results of data aggregation are efficiently stored in the MDDB by the MDDB Handler.

[0178] As shown in FIGS. 6A and 6B, the stand-alone Aggregation Server of the present invention serves the OLAP Server (or other requesting computing system) via an aggregation client interface, which preferably conforms to standard interface protocols such as OLDB, OLE-DB, ODBC, SQL, JDBC, an API, etc. Aggregation results required by the OLAP server are supplied on demand. Typically, the OLAP Server disintegrates the query, via parsing process, into series of requests. Each such request, specifying a n-dimensional coordinate, is presented to the Aggregation Server. The Configuration Manager sets the Aggregation Client Interface and Input Analyzer for a proper communication protocol according to the client user. The Input Analyzer converts the input format to make it suitable for the MDDB Handler.

[0179] An object of the present invention is to make the transfer of data completely transparent to the OLAP user, in a manner which is equivalent to the storing of data in the MOLAP server's cache and without any query delays. This requires that the stand-alone Aggregation Server have exceptionally fast response characteristics. This object is enabled by providing the unique data structure and aggregation mechanism of the present invention.

[0180] FIG. 6C shows the software modules comprising the aggregation engine and MDDB handler components of the stand-alone Aggregation Server of the illustrative embodiment. The base data list, as it arrives from RDBMS or text files, has to be analyzed and reordered to optimize hierarchy handling, according to the unique method of the present invention, as described later with reference to FIGS. 11A and 11B.

[0181] The function of the aggregation management module is to administrate the aggregation process according to the method illustrated in FIGS. 9A and 9B.

[0182] In accordance with the principles of the present invention, data aggregation within the stand-alone Aggregation Server can be carried out either as a complete pre-aggregation process, where the base data is fully aggregated before commencing querying, or as a query directed roll-up (QDR) process, where querying is allowed at any stage of aggregation using the "on-the-fly" data aggregation process of the present invention. The QDR process will be described hereinafter in greater detail with reference to FIG. 9C. The response to a request (i.e. a basic component of a client query), by calling the Aggregation management module for "on-the-fly" data aggregation, or for accessing pre-aggregated result data via the Storage management module. The query/request serving mechanism of the present invention within the QDR aggregation management module is illustrated in the flow chart of FIG. 6D.

[0183] The function of the Storage management module is to handle multidimensional data in the storage(s) module in a very efficient way, according to the novel method of the present invention, which will be described in detail herein-after with reference to FIGS. 10A and 10B.

[0184] The request serving mechanism shown in FIG. 6D is controlled by the QDR aggregation management module. Requests are queued and served one by one. If the required data is already pre-calculated, then it is retrieved by the storage management module and returned to the client. Otherwise, the required data is calculated “on-the-fly” by the aggregation management module, and the result moved out to the client, while simultaneously stored by the storage management module, shown in FIG. 6C.

[0185] FIGS. 7A and 7B outline two different implementations of the stand-alone (cartridge-style) Aggregation Server of the present invention. In both implementations, the Aggregation Server supplies aggregated results to a client.

[0186] FIG. 7A shows a separate-platform type implementation of the MOLAP system of the illustrative embodiment shown in FIG. 6A, wherein the Aggregation Server of the present invention resides on a separate hardware platform and OS system from that used to run the OLAP server. In this type of implementation, it is even possible to run the Aggregation Server and the OLAP Server on different-type operating systems (e.g. NT, Unix, MAC OS).

[0187] FIG. 7B shows a common-platform type implementation of the MOLAP system of the illustrative embodiment shown in FIG. 6B, wherein the Aggregation Server of the present invention and OLAP Server share the same hardware platform and operating system (OS).

[0188] FIG. 8A shows a table setting forth the benchmark results of an aggregation engine, implemented on a shared/common hardware platform and OS, in accordance with the principles of the present invention. The common platform and OS is realized using a Pentium II 450 Mhz, 1 GB RAM, 18 GB Disk, running the Microsoft NT operating system. The six (6) data sets shown in the table differ in number of dimensions, number of hierarchies, measure of sparsity and data size. A comparison with ORACLE Express, a major OLAP server, is made. It is evident that the aggregation engine of the present invention outperforms currently leading aggregation technology by more than an order of magnitude.

[0189] The segmented data aggregation method of the present invention is described in FIGS. 9A through 9C2. These figures outline a simplified setting of three dimensions only; however, the following analysis applies to any number of dimensions as well.

[0190] The data is being divided into autonomic segments to minimize the amount of simultaneously handled data. The initial aggregation is practiced on a single dimension only, while later on the aggregation process involves all other dimensions.

[0191] At the first stage of the aggregation method, an aggregation is performed along dimension 1. The first stage can be performed on more than one dimension. As shown in FIG. 9A, the space of the base data is expanded by the aggregation process.

[0192] In the next stage shown in FIG. 9B, any segment along dimension 1, such as the shown slice, can be separately aggregated along the remaining dimensions, 2 and 3. In general, for an N dimensional system, the second stage involves aggregation in N-1 dimensions.

[0193] The principle of data segmentation can be applied on the first stage as well. However, only a large enough data set will justify such a sliced procedure in the first dimension. Actually, it is possible to consider each segment as an N-1 cube, enabling recursive computation.

[0194] It is imperative to get aggregation results of a specific slice before the entire aggregation is completed, or alternatively, to have the roll-up done in a particular sequence. This novel feature of the aggregation method of the present invention is that it allows the querying to begin, even before the regular aggregation process is accomplished, and still having fast response. Moreover, in relational OLAP and other systems requiring only partial aggregations, the QDR process dramatically speeds up the query response.

[0195] The QDR process is made feasible by the slice-oriented roll-up method of the present invention. After aggregating the first dimension(s), the multidimensional space is composed of independent multidimensional cubes (slices). These cubes can be processed in any arbitrary sequence.

[0196] Consequently the aggregation process of the present invention can be monitored by means of files, shared memory sockets, or queues to statically or dynamically set the roll-up order.

[0197] In order to satisfy a single query coming from a client, before the required aggregation result has been prepared, the QDR process of the present invention involves performing a fast on-the-fly aggregation (roll-up) involving only a thin slice of the multidimensional data.

[0198] FIG. 9C1 shows a slice required for building-up a roll-up result of the 2nd dimension. In case 1, as shown, the aggregation starts from an existing data, either basic or previously aggregated in the first dimension. This data is utilized as a basis for QDR aggregation along the second dimension. In case 2, due to lack of previous data, a QDR involves an initial slice aggregation along dimension 3, and thereafter aggregation along the 2nd dimension.

[0199] FIG. 9C2 shows two corresponding QDR cases for gaining results in the 3d dimension. Cases 1 and 2 differ in the amount of initial aggregation required in 2 dimension.

[0200] FIG. 10A illustrates the “Slice-Storage” method of storing sparse data on storage disks. In general, this data storage method is based on the principle that an ascending-ordered index along aggregation direction, enables fast retrieval of data. FIG. 10A illustrates a unit-wide slice of the multidimensional cube of data. Since the data is sparse, only few non-NA data points exist. These points are indexed as follows. The Data File consists of data records, in which each n-1 dimensional slice is being stored, in a separate record. These records have a varying length, according to the amount of non-NA stored points. For each registered point in the record, IND_k stands for an index in a n-dimensional cube, and Data stands for the value of a given point in the cube.

[0201] FIG. 10B illustrates a novel method for randomly searching for a queried data point in the MDDB of FIG. 6B by using a novel technique of organizing data files and the directory file used in the storages of the MDDB, so that a simple binary search technique can then be employed within the Aggregation Server of the present invention. According to this method, a metafile termed DIR File, keeps pointers to Data Files as well as additional parameters such as the start and end addresses of data record (IND_o , IND_n), its location within the Data File, record size (n), file's physical address on disk (D_Path), and auxiliary information on the record (Flags).

[0202] A search for a queried data point is then performed by an access to the DIR file. The search along the file can be made using a simple binary search due to file's ascending order. When the record is found, it is then loaded into main memory to search for the required point, characterized by its index IND_k . The attached Data field represents the queried value. In case the exact index is not found, it means that the point is a NA.

[0203] In another aspect of the present invention, a novel method is provided for optimally merging multiple hierarchies in multi-hierarchical structures. The method, illustrated in FIGS. 11A, 11B, and 11C is preferably used by the Aggregation Server of the present invention in processing the table data (base data), as it arrives from RDBMS.

[0204] According to the devised method, the inner order of hierarchies within a dimension is optimized, to achieve efficient data handling for summations and other mathematical formulas (termed in general "Aggregation"). The order of hierarchy is defined externally. It is brought from a data source to the stand-alone aggregation engine, as a descriptor of data, before the data itself. In the illustrative embodiment, the method assumes hierarchical relations of the data, as shown in FIG. 11A. The way data items are ordered in the memory space of the Aggregation Server, with regard to the hierarchy, has a significant impact on its data handling efficiency.

[0205] Notably, when using prior art techniques, multiple handling of data elements, which occurs when a data element is accessed more than once during aggregation process, has been hitherto unavoidable when the main concern is to effectively handle the sparse data. The data structures used in prior art data handling methods have been designed for fast access to a non NA data. According to prior art techniques, each access is associated with a timely search and retrieval in the data structure. For the massive amount of data typically accessed from a Data Warehouse in an OLAP application, such multiple handling of data elements has significantly degraded the efficiency of prior art data aggregation processes. When using prior art data handling techniques, the data element D shown in FIG. 11A must be accessed three times, causing poor aggregation performance.

[0206] In accordance with the data handling method of the present invention, the data is being pre-ordered for a singular handling, as opposed to multiple handling taught by prior art methods. According to the present invention, elements of base data and their aggregated results are contiguously stored in a way that each element will be accessed only once. This particular order allows a forward-only handling, never backward. Once a base data element is stored, or aggregated

result is generated and stored, it is never to be retrieved again for further aggregation. As a result the storage access is minimized. This way of singular handling greatly elevates the aggregation efficiency of large data bases. An efficient handling method as used in the present invention, is shown in FIG. 7A. The data element D, as any other element, is accessed and handled only once.

[0207] FIG. 11A shows an example of a multi-hierarchical database structure having 3 hierarchies. As shown, the base data has a dimension that includes items A,B,F, and G. The second level is composed of items C,E,H and I. The third level has a single item D, which is common to all three hierarchical structures. In accordance with the method of the present invention, a minimal computing path is always taken. For example, according to the method of the present invention, item D will be calculated as part of structure 1, requiring two mathematical operations only, rather than as in structure 3, which would need four mathematical operations. FIG. 11B depicts an optimized structure merged from all three hierarchies.

[0208] FIG. 11C(i) through 11C(ix) represent a flow chart description (and accompanying data structures) of the operations of an exemplary hierarchy transformation mechanism of the present invention that optimally merges multiple hierarchies into a single hierarchy that is functionally equivalent to the multiple hierarchies. For the sake of description, the data structures correspond to exemplary hierarchical structures described above with respect to FIGS. 11(A) and 11(B). As illustrated in FIG. 11C(i), in step 1101, a catalogue is loaded from the DBMS system. As is conventional, the catalogue includes data ("hierarchy descriptor data") describing multiple hierarchies for at least one dimension of the data stored in the DBMS. In step 1103, this hierarchy descriptor data is extracted from the catalogue. A loop (steps 1105-1119) is performed over the items in the multiple hierarchy described by the hierarchy descriptor data.

[0209] In the loop 1105-1119, a given item in the multiple hierarchy is selected (step 1107); and, in step 1109, the parent(s) (if any)—including grandparents, great-grandparents, etc.—of the given item are identified and added to an entry (for the given item) in a parent list data structure, which is illustrated in FIG. 11C(v). Each entry in the parent list corresponds to a specific item and includes zero or more identifiers for items that are parents (or grandparents, or great-grandparents) of the specific item. In addition, an inner loop (steps 1111-1117) is performed over the hierarchies of the multiple hierarchies described by the hierarchy descriptor data, wherein in step 1113 one of the multiple hierarchies is selected. In step 1115, the child of the given item in the selected hierarchy (if any) is identified and added (if need be) to a group of identifiers in an entry (for the given item) in a child list data structure, which is illustrated in FIG. 11C(vi). Each entry in the child list corresponds to a specific item and includes zero or more groups of identifiers each identifying a child of the specific item. Each group corresponds to one or more of the hierarchies described by the hierarchy descriptor data.

[0210] The operation then continues to steps 1121 and 1123 as illustrated in FIG. 11C(ii) to verify the integrity of the multiple hierarchies described by the hierarchy descriptor data (step 1121) and fix (or report to the user) any errors

discovered therein (step 1123). Preferably, the integrity of the multiple hierarchies is verified in step 1121 by iteratively expanding each group of identifiers in the child list to include the children, grandchildren, etc of any item listed in the group. If the child(ren) for each group for a specific item do not match, a verification error is encountered, and such error is fixed (or reported to the user (step 1123)). The operation then proceeds to a loop (steps 1125-1133) over the items in the child list.

[0211] In the loop (steps 1125-1133), a given item in the child list is identified in step 1127. In step 1129, the entry in the child list for the given item is examined to determine if the given item has no children (e.g., the corresponding entry is null). If so, the operation continues to step 1131 to add an entry for the item in level 0 of an ordered list data structure, which is illustrated in FIG. 11C(vii); otherwise the operation continues to process the next item of the child list in the loop. Each entry in a given level of the order list corresponds to a specific item and includes zero or more identifiers each identifying a child of the specific item. The levels of the order list described the transformed hierarchy as will readily become apparent in light of the following. Essentially, loop 1125-1133 builds the lowest level (level 0) of the transformed hierarchy.

[0212] After loop 1125-1133, operation continues to process the lowest level to derive the next higher level, and iterate over this process to build out the entire transformed hierarchy. More specifically, in step 1135, a "current level" variable is set to identify the lowest level. In step 1137, the items of the "current level" of the ordered list are copied to a work list. In step 1139, it is determined if the worklist is empty. If so, the operation ends; otherwise operation continues to step 1141 wherein a loop (steps 1141-1159) is performed over the items in the work list.

[0213] In step 1143, a given item in the work list is identified and operation continues to an inner loop (steps 1145-1155) over the parent(s) of the given item (which are specified in the parent list entry for the given item). In step 1147 of the inner loop, a given parent of the given item is identified. In step 1149, it is determined whether any other parent (e.g., a parent other than the given parent) of the given item is a child of the given parent (as specified in the child list entry for the given parent). If so, operation continues to step 1155 to process the next parent of the given item in the inner loop; otherwise, operation continues to steps 1151 and 1153. In step 1151, an entry for the given parent is added to the next level (current level+1) of the ordered list, if it does not exist there already. In step 1153, if no children of the given item (as specified in the entry for the given item in the current level of the ordered list) matches (e.g., is covered by) any child (or grandchild or great grandchild etc) of item(s) in the entry for the given parent in the next level of the ordered list, the given item is added to the entry for the given parent in the next level of the ordered list. Levels 1 and 2 of the ordered list for the example described above are shown in FIGS. 11C(viii) and 11C(ix), respectively. The children (including grandchildren and great grandchildren, etc) of an item in the entry for a given parent in the next level of the ordered list may be identified by the information encoded in the lower levels of the ordered list. After step 1153, operation continues to step 1155 to process the next parent of the given item in the inner loop (steps 1145-1155). After processing the inner loop (steps 1145-1155), operation contin-

ues to step 1157 to delete the given item from the work list, and processing continues to step 1159 to process the next item of the work list in the loop (steps 1141-1159).

[0214] After processing the loop (steps 1141-1159), the ordered list (e.g., transformed hierarchy) has been built for the next higher level. The operation continues to step 1161 to increment the current level to the next higher level, and operation returns (in step 1163) to step 1138 to build the next higher level, until the highest level is reached (determined in step 1139) and the operation ends.

[0215] FIG. 12 summarizes the components of an exemplary aggregation module that takes advantage of the hierarchy transformation technique described above. More specifically, the aggregation module includes an hierarchy transformation module that optimally merges multiple hierarchies into a single hierarchy that is functionally equivalent to the multiple hierarchies. A second module loads and indexes the base data supplied from the DBMS using the optimal hierarchy generated by the hierarchy transformation module. An aggregation engine performs aggregation operations on the base data. During the aggregation operations along the dimension specified by the optimal hierarchy, the results of the aggregation operations of the level 0 items may be used in the aggregation operations of the level 1 items, the results of the aggregation operations of the level 1 items may be used in the aggregation operations of the level 2 items, etc. Based on these operations, the loading and indexing operations of the base data, along with the aggregation become very efficient, minimizing memory and storage access, and speeding up storing and retrieval operations.

[0216] FIG. 13 shows the stand-alone Aggregation Server of the present invention as a component of a central data warehouse, serving the data aggregation needs of URL directory systems, Data Marts, RDBMSs, ROLAP systems and OLAP systems alike.

[0217] The reason for the central multidimensional database's rise to corporate necessity is that it facilitates flexible, high-performance access and analysis of large volumes of complex and interrelated data.

[0218] A stand-alone specialized aggregation server, simultaneously serving many different kinds of clients (e.g. data mart, OLAP, URL, RDBMS), has the power of delivering an enterprise-wide aggregation in a cost-effective way. This kind of server eliminates the roll-up redundancy over the group of clients, delivering scalability and flexibility.

[0219] Performance associated with central data warehouse is an important consideration in the overall approach. Performance includes aggregation times and query response.

[0220] Effective interactive query applications require near real-time performance, measured in seconds. These application performances translate directly into the aggregation requirements.

[0221] In the prior art, in case of MOLAP, a full pre-aggregation must be done before starting querying. In the present invention, in contrast to prior art, the query directed roll-up (QDR) allows instant querying, while the full pre-aggregation is done in the background. In cases a full pre-aggregation is preferred, the currently invented aggregation outperforms any prior art. For the ROLAP and RDBMS clients, partial aggregations maximize query per-

formance. In both cases fast aggregation process is imperative. The aggregation performance of the current invention is by orders of magnitude higher than that of the prior art.

[0222] The stand-alone scalable aggregation server of the present invention can be used in any MOLAP system environment for answering questions about corporate performance in a particular market, economic trends, consumer behaviors, weather conditions, population trends, or the state of any physical, social, biological or other system or phenomenon on which different types or categories of information, organizable in accordance with a predetermined dimensional hierarchy, are collected and stored within a RDBMS of one sort or another. Regardless of the particular application selected, the address data mapping processes of the present invention will provide a quick and efficient way of managing a MDDB and also enabling decision support capabilities utilizing the same in diverse application environments.

[0223] The stand-alone “cartridge-style” plug-in features of the data aggregation server of the present invention, provides freedom in designing an optimized multidimensional data structure and handling method for aggregation, provides freedom in designing a generic aggregation server matching all OLAP vendors, and enables enterprise-wide centralized aggregation.

[0224] The method of Segmented Aggregation employed in the aggregation server of the present invention provides flexibility, scalability, a condition for Query Directed Aggregation, and speed improvement.

[0225] The method of Multidimensional data organization and indexing employed in the aggregation server of the present invention provides fast storage and retrieval, a condition for Segmented Aggregation, improves the storing, handling, and retrieval of data in a fast manner, and contributes to structural flexibility to allow sliced aggregation and QDR. It also enables the forwarding and single handling of data with improvements in speed performance.

[0226] The method of Query Directed Aggregation (QDR) employed in the aggregation server of the present invention minimizes the data handling operations in multi-hierarchy data structures.

[0227] The method of Query Directed Aggregation (QDR) employed in the aggregation server of the present invention eliminates the need to wait for full aggregation to be completed, and provides build-up aggregated data required for full aggregation.

[0228] In another aspect of the present invention, an improved DBMS system (e.g., RDBMS system, object oriented database system or object/relational database system) is provided that excels in performing two distinct functions, namely: the aggregation of data; and the handling of the resulting data for “on demand” client use. Moreover, because of improved data aggregation capabilities, the DBMS of the present invention can be employed in a wide range of applications, including Data Warehouses supporting OLAP systems and the like. For purposes of illustration, initial focus will be accorded to the DBMS of the present invention. Referring now to FIGS. 19 through FIGS. 21, the preferred embodiments of the method and system of the present invention will be now described in great detail herein below.

[0229] Through this document, the term “aggregation” and “pre-aggregation” shall be understood to mean the process of summation of numbers, as well as other mathematical operations, such as multiplication, subtraction, division etc. It shall be understood that pre-aggregation operations occur asynchronously with respect to the traditional query processing operations. Moreover, the term “atomic data” shall be understood to refer to the lowest level of data granularity required for effective decision making. In the case of a retail merchandising manager, atomic data may refer to information by store, by day, and by item. For a banker, atomic data may be information by account, by transaction, and by branch.

[0230] FIG. 19A illustrates the primary components of an illustrative embodiment of an DBMS of the present invention, namely: support mechanisms including a query interface and query handler; a relational data store including one or more tables storing at least the atomic data (and possibly summary tables) and a meta-data store for storing a dictionary (sometimes referred to as a catalogue or directory); and an MDD Aggregation Module that stores atomic data and aggregated data in a MDDB. The MDDB is a non-relational data structure-it uses other data structures, either instead of or in addition to tables-to store data. For illustrative purposes, FIG. 19A illustrates an RDBMS wherein the relational data store includes fact tables and a dictionary.

[0231] It should be noted that the DBMS typically includes additional components (not shown) that are not relevant to the present invention. The query interface and query handler service user-submitted queries (in the preferred embodiment, SQL query statements) forwarded, for example, from a client machine over a network as shown. The query handler and relational data store (tables and meta-data store) are operably coupled to the MDD Aggregation Module. Importantly, the query handler and integrated MDD Aggregation Module operate to provide for dramatically improved query response times for data aggregation operations and drill-downs. Moreover, it is an object of the present invention to make user-querying of the non-relational MDDB no different than querying a relational table of the DBMS, in a manner that minimizes the delays associated with queries that involve aggregation or drill down operations. This object is enabled by providing the novel DBMS system and integrated aggregation mechanism of the present invention.

[0232] FIG. 19B shows the primary components of an illustrative embodiment of the MDD Aggregation Module of FIG. 19A, namely: a base data loader for loading the directory and table(s) of relational data store of the DBMS; an aggregation engine for receiving dimension data and atomic data from the base loader, a multi-dimensional database (MDDB); a MDDB handler and an SQL handler that operate cooperatively with the query handler of the DBMS to provide users with query access to the MDD Aggregation Module, and a control module for managing the operation of the components of the MDD aggregation module. The base data loader may load the directory and table(s) of the relational data store over a standard interface (such as OLDB, OLE-DB, ODBC, SQL, API, JDBC, etc.). In this case, the DBMS and base data loader include components that provide communication of such data over these standard interfaces. Such interface components are

well known in the art. For example, such interface components are readily available from Attunity Corporation.

[0233] During operation, base data originates from the table(s) of the DBMS. The core data aggregation operations are performed by the Aggregation Engine; a Multidimensional Data (MDDB) Handler; and a Multidimensional Data Storage (MDDB). The results of data aggregation are efficiently stored in the MDDB by the MDDB Handler. The SQL handler of the MDD Aggregation module services user-submitted queries (in the preferred embodiment, SQL query statements) forwarded from the query handler of the DBMS. The SQL handler of the MDD Aggregation module may communicate with the query handler of the DBMS over a standard interface (such as OLDB, OLE-DB, ODBC, SQL, API, JDBC, etc.). In this case, the support mechanisms of the RDBMS and SQL handler include components that provide communication of such data over these standard interfaces. Such interface components are well known in the art. Aggregation (or drill down results) are retrieved on demand and returned to the user.

[0234] Typically, a user interacts with a client machine (for example, using a web-enabled browser) to generate a natural language query, that is communicated to the query interface of the DBMS, for example over a network as shown. The query interface disintegrates the query, via parsing, into a series of requests (in the preferred embodiment, SQL statements) that are communicated to the query handler of the DBMS. It should be noted that the functions of the query interface may be implemented in a module that is not part of the DBMS (for example, in the client machine). The query handler of the DBMS forwards requests that involve data stored in the MDD of the MDD Aggregation module to the SQL handler of the MDD Aggregation module for servicing. Each request specifies a set of n-dimensions. The SQL handler of the MDD Aggregation Module extracts this set of dimensions and operates cooperatively with the MDD handler to address the MDDB using the set of dimensions, retrieve the addressed data from the MDDB, and return the results to the user via the query handler of the DBMS.

[0235] FIG. 19C(i) and 19C(ii) is a flow chart illustrating the operations of an illustrative DBMS of the present invention. In step 601, the base data loader of the MDD Aggregation Module loads the dictionary (or catalog) from the meta-data store of the DBMS. In performing this function, the base data loader may utilize an adapter (interface) that maps the data types of the dictionary of the DBMS (or that maps a standard data type used to represent the dictionary of the DBMS) into the data types used in the MDD aggregation module. In addition, the base data loader extracts the dimensions from the dictionary and forwards the dimensions to the aggregation engine of the MDD Aggregation Module.

[0236] In step 603, the base data loader loads table(s) from the DBMS. In performing this function, the base data loader may utilize an adapter (interface) that maps the data types of the table(s) of the DBMS (or that maps a standard data type used to represent the fact table(s) of the DBMS) into the data types used in the MDD Aggregation Module. In addition, the base data loader extracts the atomic data from the table(s), and forwards the atomic data to the aggregation engine.

[0237] In step 605, the aggregation engine performs aggregation operations (i.e., roll-up operation) on the atomic

data (provided by the base data loader in step 603) along at least one of the dimensions (extracted from the dictionary of the DBMS in step 601) and operates cooperatively with the MDD handler to store the resultant aggregated data in the MDDB. A more detailed description of exemplary aggregation operations according to a preferred embodiment of the present invention is set forth below with respect to the QDR process of FIGS. 9A-9C.

[0238] In step 607, a reference is defined that provides users with the ability to query the data generated by the MDD Aggregation Module and/or stored in the MDDB of the MDD Aggregation Module. This reference is preferably defined using the Create View SQL statement, which allows the user to: i) define a table name (TN) associated with the MDDB stored in the MDD Aggregation Module, and ii) define a link used to route SQL statements on the table TN to the MDD Aggregation Module. In this embodiment, the view mechanism of the DBMS enables reference and linking to the data stored in the MDDB of the MDD Aggregation Engine as illustrated in FIG. 19E. A more detailed description of the view mechanism and the Create View SQL statement may be found in C. J. Date, "An Introduction to Database Systems," Addison-Wesley, Seventh Edition, 2000, pp. 289-326, herein incorporated by reference in its entirety. Thus, the view mechanism enables the query handler of the DBMS system to forward any SQL query on table TN to the MDD aggregation module via the associated link. In an alternative embodiment, a direct mechanism (e.g., NA trigger mechanism) may be used to enable the DBMS system to reference and link to the data generated by the MDD Aggregation Module and/or stored in the MDDB of the MDD Aggregation Engine as illustrated in FIG. 19F. A more detailed description of trigger mechanisms and methods may be found in C. J. Date, "An Introduction to Database Systems," Addison-Wesley, Seventh Edition, 2000, pp. 250, 266, herein incorporated by reference in its entirety.

[0239] In step 609, a user interacts with a client machine to generate a query, and the query is communicated to the query interface. The query interface generate one or more SQL statements. These SQL statements may refer to data stored in tables of the relational datastore, or may refer to the reference defined in step 607 (this reference refers to the data stored in the MDDB of the MDD Aggregation Module). These SQL statement(s) are forwarded to the query handler of the DBMS.

[0240] In step 611, the query handler receives the SQL statement(s); and optionally transforms such SQL statement(s) to optimize the SQL statement(s) for more efficient query handling. Such transformations are well known in the art. For example, see Kimball, "Aggregation Navigation With (Almost) No MetaData", DBMS Data Warehouse Supplement, August 1996.

[0241] In step 613: the query handler determines whether the received SQL statement(s) [or transformed SQL statement(s)] is on the reference generated in step 607. If so, operation continues to step 615; otherwise normal query handling operations continue in step 625 wherein the relational datastore is accessed to extract, store, and/or manipulate the data stored therein as directed by the query, and results are returned back to the user via the client machine, if needed.

[0242] In step 615, the received SQL statement(s) [or transformed SQL statement(s)] is routed to the MDD aggregation engine for processing in step 617 using the link for the reference as described above with respect to step 607.

[0243] In step 617, the SQL statement(s) is received by the SQL handler of the MDD Aggregation Module, wherein a set of one or more N-dimensional coordinates are extracted from the SQL statement. In performing this function, SQL handler may utilize an adapter (interface) that maps the data types of the SQL statement issued by query handler of the DBMS (or that maps a standard data type used to represent the SQL statement issued by query handler of the DBMS) into the data types used in the MDD aggregation module.

[0244] In step 619, the set of N-dimensional coordinates extracted in step 617 are used by the MDD handler to address the MDDB and retrieve the corresponding data from the MDDB.

[0245] Finally, in step 621, the retrieved data is returned to the user via the DBMS (for example, by forwarding the retrieved data to the SQL handler, which returns the retrieved data to the query handler of the DBMS system, which returns the results of the user-submitted query to the user via the client machine), and the operation ends.

[0246] It should be noted that the table data (base data), as it arrives from DBMS, may be analyzed and reordered to optimize hierarchy handling, according to the unique method of the present invention, as described above with reference to FIGS. 11A, 11B, and 11C.

[0247] Moreover, the MDD control module of the MDD Aggregation Module preferably administers the aggregation process according to the method illustrated in FIGS. 9A and 9B. Thus, in accordance with the principles of the present invention, data aggregation within the DBMS can be carried out either as a complete pre-aggregation process, where the base data is fully aggregated before commencing querying, or as a query directed roll-up (QDR) process, where querying is allowed at any stage of aggregation using the “on-the-fly” data aggregation process of the present invention. The QDR process will be described hereinafter in greater detail with reference to FIG. 9C. The response to a request (i.e. a basic component of a client query) requiring “on-the-fly” data aggregation, or requiring access to pre-aggregated result data via the MDD handler is provided by a query/request serving mechanism of the present invention within the MDD control module, the primary operations of which are illustrated in the flow chart of FIG. 6D. The function of the MDD Handler is to handle multidimensional data in the storage(s) module in a very efficient way, according to the novel method of the present invention, which will be described in detail hereinafter with reference to FIGS. 10A and 10B.

[0248] The SQL handling mechanism shown in FIG. 6D is controlled by the MDD control module. Requests are queued and served one by one. If the required data is already pre-calculated, then it is retrieved by the MDD handler and returned to the client. Otherwise, the required data is calculated “on-the-fly” by the aggregation engine, and the result moved out to the client, while simultaneously stored by the MDD handler, shown in FIG. 6C.

[0249] As illustrated in FIG. 19G, the DBMS of the present invention as described above may be logically

partitioned into a relational part and a non-relational part. The relational part includes the relational datastore (e.g., table(s) and dictionary) and support mechanisms (e.g., query handling services). The non-relational part includes the MDD Aggregation Module. As described above, bidirectional data flow occurs between the relational part and the non-relational part as shown. More specifically, during data load operations, data is loaded from the relational part (i.e., the relational datastore) into the non-relational part, wherein it is aggregated and stored in the MDDB. And during query servicing operations, when a given query references data stored in the MDDB, data pertaining to the query is generated by the non-relational part (e.g., generated and/or retrieved from the MDDB) and supplied to the relational part (e.g., query servicing mechanism) for communication back to the user. Such bidirectional data flow represents an important distinguishing feature with respect to the prior art. For example, in the prior art MOLAP architecture as illustrated in FIG. 1B, unidirectional data flows occurs from the relational data base (e.g., the Data Warehouse RDBMS system) into the MDDB during data loading operations.

[0250] FIGS. 20A and 20B outline two different implementations of the DBMS of the present invention. In both implementations, the query handler of the DBMS system supplies aggregated results retrieved from the MDD to a client.

[0251] FIG. 20A shows a separate-platform implementation of the DBMS system of the illustrative embodiment shown in FIG. 19A, wherein the relational part of the DBMS reside on a separate hardware platform and/or OS system from that used to run the non-relational part (MDD Aggregation Module). In this type of implementation, it is even possible to run parts of the DBMS system and the MDD Aggregation Module on different-type operating systems (e.g. NT, Unix, MAC OS).

[0252] FIG. 20B shows a common-platform implementation of the DBMS system of the illustrative embodiment shown in FIG. 20A, wherein the relational part of the DBMS share the same hardware platform and operating system (OS) that is used to run the non-relational part (MDD Aggregation Module).

[0253] FIG. 21 shows the improved DBMS (e.g., RDBMS) of the present invention as a component of a data warehouse, serving the data storage and aggregation needs of a ROLAP system (or other OLAP systems alike). Importantly, the improved DBMS of the present invention provides flexible, high-performance access and analysis of large volumes of complex and interrelated data. Moreover, the improved Data Warehouse DBMS of the present invention can simultaneously serve many different kinds of clients (e.g. data mart, OLAP, URL) and has the power of delivering an enterprise-wide data storage and aggregation in a cost-effective way. This kind of system eliminates redundancy over the group of clients, delivering scalability and flexibility. Moreover, the improved DBMS of the present invention can be used as the data store component of in any informational database system as described above, including data analysis programs such as spread-sheet modeling programs, serving the data storage and aggregation needs of such systems.

[0254] FIG. 22 shows an embodiment of the present invention wherein the DBMS (e.g., RDBMS) of the present

invention is a component of a data warehouse-OLAP system. The DBMS operates as a traditional data warehouse, serving the data storage and aggregation needs of an enterprise. In addition, the DBMS includes integrated OLAP Analysis Logic (and preferably an integrated Presentation Module not shown) that operates cooperatively with the query handling of the DBMS system and the MDD Aggregation Module to enable users of the DBMS system to execute multidimensional reports (e.g., ratios, ranks, transforms, dynamic consolidation, complex filtering, forecasts, query governing, scheduling, flow control, pre-aggregate inferencing, denormalization support, and/or table partitioning and joins) and preferably perform traditional OLAP analyses (grids, graphs, maps, alerts, drill-down, data pivot, data surf, slice and dice, print). Importantly, the improved DBMS of the present invention provides flexible, high-performance access and analysis of large volumes of complex and interrelated data. Moreover, the improved DBMS of the present invention can simultaneously serve many different kinds of clients (e.g. data mart, other OLAP systems, URL-Directory Systems) and has the power of delivering enterprise-wide data storage and aggregation and OLAP analysis in a cost-effective way. This kind of system eliminates redundancy over the group of clients, delivering scalability and flexibility. Moreover, the improved DBMS of the present invention can be used as the data store component of in any informational database system as described above, serving the data storage and aggregation needs of such systems.

Functional Advantages Gained By The Improved DBMS Of The Present Invention

[0255] The features of the DBMS of the present invention, provides for dramatically improved response time in handling queries issued to the DBMS that involve aggregation, thus enabling enterprise-wide centralized aggregation. Moreover, in the preferred embodiment of the present invention, users can query the aggregated data in a manner no different than traditional queries on the DBMS.

[0256] The method of Segmented Aggregation employed by the novel DBMS of the present invention provides flexibility, scalability, the capability of Query Directed Aggregation, and speed improvement.

[0257] Moreover, the method of Query Directed Aggregation (QDR) employed by the novel DBMS of the present invention minimizes the data handling operations in multi-hierarchy data structures, eliminates the need to wait for full aggregation to be complete, and provides for build-up of aggregated data required for full aggregation.

[0258] It is understood that the System and Method of the illustrative embodiments described herein above may be modified in a variety of ways which will become readily apparent to those skilled in the art of having the benefit of the novel teachings disclosed herein. All such modifications and variations of the illustrative embodiments thereof shall be deemed to be within the scope and spirit of the present invention as defined by the Claims to Invention appended hereto.

What is claimed is:

1. A relational database management system (RDBMS) for servicing query statements through one or more client machines, said RDBMS comprising:

- a RDBMS query interface adapted to receive query statements from said one or more client machines;
- a query processing mechanism for processing each query statement received from said RDBMS query interface, and generating one or more query requests by disintegrating the query statement so that each said query request specifies a set of dimensions;
- a query handling mechanism operably coupled to said RDBMS query interface; and
- a relational data store having relational tables for storing fact data, and a meta-data store for storing a dictionary containing dimension data;
- a multi-dimensional database (MDDB) for storing aggregated fact data in a multi-dimensional data structure;
- a data loading mechanism for loading said meta-data and said fact data into said MDDB;
- an aggregation engine for calculating aggregated fact data from the fact data according to a multi-dimensional data aggregation process;
- a data handling mechanism for storing aggregated fact data in said MDDB and retrieving aggregated fact data from said MDDB;

wherein said query handling mechanism (i) receives each query request, (ii) extracts a set of dimensions associated with said query request, (iii) uses said dimensions to retrieve aggregated fact data from said MDDB, and (iv) forwards retrieved aggregated fact data to said query processing mechanism for subsequent processing, if needed; and

wherein when said query processing mechanism determines that servicing of one or more query requests require data stored in said relational tables, then said query processing mechanism automatically routes said one or more query requests to said relational data tables, so that data can be accessed from said relational tables and forwarded to said query processing mechanism for use in servicing said one or more query requests, in a manner transparent to said client machine; and

wherein when said query processing mechanism determines that servicing of one or more query requests require aggregated data stored in said MDDB, then said query processing mechanism automatically routes said one or more query requests to said MDDB, so that aggregated data can be accessed from said MDDB and forwarded to said query processing mechanism for use in servicing said one or more query requests, in a manner transparent to said client machine.

2. The RDBMS of claim 1, wherein said a query handling mechanism comprises:

- an RDBMS query handler operably coupled to said RDBMS query interface; and
- a MDDB query handler operably coupled to a first data communication interface, for (i) receiving each query request from said RDBMS query handler, (ii) extracting a set of dimensions associated with said query request, (iii) providing said dimensions to said data handling mechanism so that said data handling mechanism

nism can retrieve aggregated fact data from said MDDB, and (iv) forwarding retrieved aggregated fact data to said RDBMS query processing mechanism for subsequent processing, if needed; and

3. The RDBMS of claim 1, where said data loading mechanism is in data communication with said relational data store by way of a second data communication interface, for loading said meta-data and said fact data into said MDDB.

4. The RDBMS of claim 1, wherein after data loading operations, the fact data is initially aggregated along at least one of the dimensions extracted from said dictionary and aggregated multi-dimensional data results are stored in said MDDB, and

wherein during query servicing operations, the fact data is aggregated on demand along additional dimensions extracted from said query request and aggregated multi-dimensional data results are stored in said MDDB.

5. The RDBMS of claim 3, wherein said multi-dimensional data structure comprises a plurality of data storage cells, and wherein each said data storage cell is indexed with multiple dimensions and stores either base data values or aggregated multi-dimensional data values.

6. The RDBMS of claim 3, wherein during data loading operations, said data loading mechanism extracts dimension data from the dictionary in said meta-data store, and forwards the dimension data over said first data communication interface to said aggregation engine and configures said MDDB using said dimension data.

7. The RDBMS of claim 3, wherein during aggregation operations, said data loading mechanism extracts fact data from said relational tables, and forwards the fact data to said aggregation engine over said first data communication interface; and wherein said aggregation engine calculates aggregated multi-dimensional data from the fact data according to said multi-dimensional data aggregation process, and stores the aggregated multi-dimensional data in said MDDB specified by a set of multi-dimensional coordinates.

8. The RDBMS of claim 1, wherein said query statement is an SQL-type query statement and each query request is an SQL-type query request.

9. The RDBMS of claim 2, wherein said first data communication interface is a standard interface selected from the group selected from OLDB, OLE-DB, ODBC, SQL, API, and JDBC.

10. The RDBMS of claim 3, wherein said second data communication interface is a standard interface selected from the group selected from OLDB, OLE-DB, ODBC, SQL, API, and JDBC.

11. The RDBMS of claim 1, wherein said RDBMS query interface is implemented in a module within said RDBMS.

12. The RDBMS of claim 1, wherein a user interacts with said client machine using a web-enabled browser to generate a natural language type query statement, and said natural language type query statement is communicated to said RDBMS query interface.

13. The RDBMS of claim 2, wherein said MDDB query handler further includes means for transforming each query request so as to optimize the query request for efficient query handling.

14. The RDBMS of claim 2, wherein said MDDB query handler further includes an adapter for mapping the data types of the query statement issued by said RDBMS query handler, and/or for mapping a standard data type used to represent the query request issued by said RDBMS query handler into the data types used in said MDD aggregation module.

15. The RDBMS of claim 1, wherein said multi-dimensional aggregation process supported by said aggregation engine further operates as follows:

- (a) if the aggregated multi-dimensional data required to service a given query statement is already pre-calculated and stored within said MDDB, then the pre-aggregated multi-dimensional data is retrieved by said data handling mechanism and returned to said client machine via said RDBMS query handler; and (b) if the required multi-dimensional data is not already pre-aggregated and stored within said MDDB, then the required aggregated multi-dimensional data is calculated on demand by said aggregation engine, and the aggregated multi-dimensional data result is automatically forwarded to said client machine.

* * * * *