



(51) International Patent Classification:
G06F 17/30 (2006.01)

(21) International Application Number:
PCT/US2012/054680

(22) International Filing Date:
11 September 2012 (11.09.2012)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/535,227 15 September 2011 (15.09.2011) US
13/420,324 14 March 2012 (14.03.2012) US

(71) Applicant (for all designated States except US): **ORACLE INTERNATIONAL CORPORATION** [US/US];
500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **BEERBOWER, Tom** [US/US]; 70 Pine Run Drive, Holland, Pennsylvania

18966 (US). **LEE, Robert, H.** [US/US]; 928 Sunset Drive, San Carlos, California 94070 (US).

(74) Agents: **MEYER, Sheldon, R.** et al.; Fliesler Meyer LLP, 650 California Street, Fourteenth Floor, San Francisco, California 94108 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,

[Continued on next page]

(54) Title: QUERY EXPLAIN PLAN IN A DISTRIBUTED DATA MANAGEMENT SYSTEM

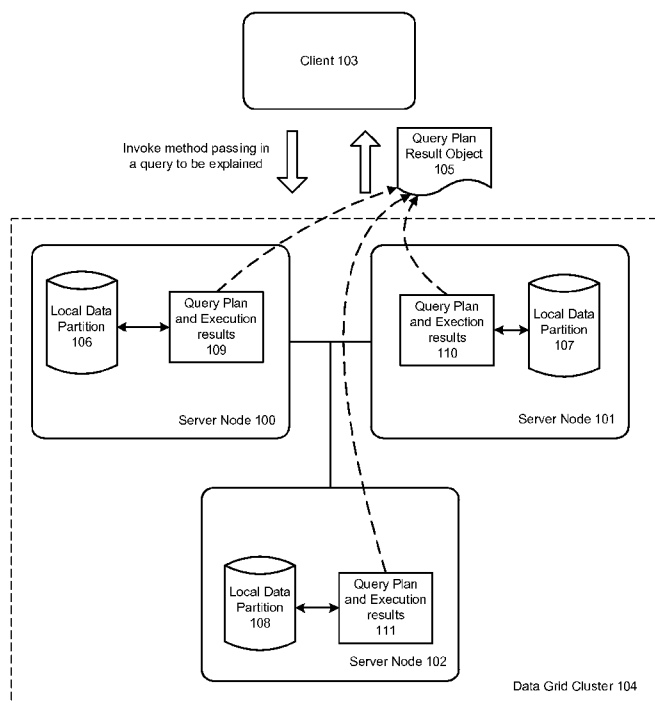


FIGURE 1

(57) Abstract: A query explain plan is described for use with a distributed data system in order to help developers and IT experts to detect bottlenecks and to determine which queries are to blame for a slow running system. In accordance with an embodiment, the distributed data grid utilizes a facility to provide a distributed explain plan. This distributed explain plan provides visibility into how a query was answered by each node in the cluster. For example, one node may have applied the filters of the query in one sequence, while another node may have used a different sequence. Additionally, the distributed query explain plan can provide execution tracing, rendering information about the execution time of each step, total execution time and the like.



TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

QUERY EXPLAIN PLAN IN A DISTRIBUTED DATA MANAGEMENT SYSTEM**COPYRIGHT NOTICE**

[0001] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

FIELD OF THE INVENTION

[0002] The current invention relates to data caching and accessing techniques in distributed computing environments.

BACKGROUND

[0003] In the context of data storage and management, query optimization attempts to determine a most optimal or efficient way to execute a given query. Because structured query language (SQL) is declarative, there are typically multiple different ways to execute a given query (query plans) and each of those ways yields a different performance while arriving at the same results. A query plan is an ordered sequence of steps used to access or modify information in SQL based data systems. A typical query optimizer considers all the possible query plans for a given query and determines which one of those query plans will be most efficient (produce the least amount of latency).

[0004] In distributed data systems where data is stored across many nodes of a cluster, it can be difficult to determine a single optimal query plan for the entire distributed system. For example, depending on the data distribution across the cluster, one query plan may be optimal for some of the nodes in the cluster, while a different query plan will be optimal for other nodes.

BRIEF SUMMARY

[0005] In accordance with various embodiments, a query explain plan is described for use in a distributed data system in order to help developers and IT experts to detect bottlenecks and to determine which queries are to blame for a slow running system. The query explain plan is an ordered set of steps used to access or modify information in SQL based data systems. The query explain plan generates and stores information about an execution plan

chosen by a query optimizer. In accordance with an embodiment, the distributed data grid utilizes a facility to provide a distributed explain plan. This distributed explain plan provides visibility into how a query was answered by each node in the cluster. For example, one node may have applied the filters of the query in one sequence, while another node may have used a different sequence. Additionally, the distributed query explain plan can provide execution tracing, rendering information about the execution time of each step, total execution time and the like.

BRIEF DESCRIPTION OF THE DRAWINGS

10 [0006] FIGURE 1 is a system level diagram of the query explain plan feature, in accordance with various embodiments of the invention.

[0007] FIGURES 2A and 2B are illustrations of an example of using the query explain plan in accordance with various embodiments of the invention.

15 [0008] FIGURE 3 is a flow chart diagram of providing the query explain plan in accordance with various embodiments of the invention.

[0009] FIGURE 4 is a block diagram that illustrates a computer system upon which an embodiment of the inventions may be implemented.

DETAILED DESCRIPTION

20 [0010] The invention is illustrated by way of example and not by way of limitation in the figures of the accompanying drawings in which like references indicate similar elements. References to embodiments in this disclosure are not necessarily to the same embodiment, and such references mean at least one. While specific implementations are discussed, it is understood that this is done for illustrative purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without departing from the scope and spirit of the invention.

25 [0011] In the following description, numerous specific details are set forth to provide a thorough description of the invention. However, it will be apparent to those skilled in the art that the invention may be practiced without these specific details. In other instances, well-known features have not been described in detail so as not to obscure the invention.

30 [0012] In accordance with various embodiments, a query explain plan is provided to a distributed data system in order to help developers and IT experts to detect bottlenecks and to determine which queries are to blame for a slow running system. In accordance with an embodiment, the distributed data grid utilizes a facility to provide a distributed explain plan.

This distributed explain plan provides visibility into how a query was answered by each node in the cluster. For example, one node may have applied the filters of the query in one sequence, while another node may have used a different sequence. Additionally, the distributed query explain plan can provide execution tracing, rendering information about the execution time of each step, total execution time and the like.

[0013] In accordance with an embodiment, to obtain the query explain plan, a method can be invoked, passing a query as a parameter. The query can be sent to all servers in the cluster and each server can return the determined optimal query plan specific to that server, as well as its execution results. These plans and results can be aggregated across all nodes in the cluster. The result of the method invocation can thus be a data structure that provides information regarding all the query plans across the cluster for that particular query. The facility can also provide means to compare and contrast the different query plans and to reconcile the results.

[0014] In accordance with an embodiment, the user can inspect the results of the query explain plan and create indexes to improve the speed of the query. An index is a data structure that can improve the speed of the data retrieval operation by creating a copy of a part of a table or by some other means as known in the art. Based on the results of the query explain plan, users can have a better idea of which tables to create indexes for in order to improve the speed of the system.

[0015] FIGURE 1 is a system level diagram of the query explain plan feature, in accordance with various embodiments of the invention. Although this diagram depicts components as logically separate, such depiction is merely for illustrative purposes. It will be apparent to those skilled in the art that the components portrayed in this figure can be combined or divided into separate software, firmware and/or hardware. Furthermore, it will also be apparent to those skilled in the art that such components, regardless of how they are combined or divided, can execute on the same computing device or can be distributed among different computing devices connected by one or more networks or other suitable communication means.

[0016] As illustrated, a client 103 can connect to a distributed data grid cluster 104 that manages a set of data across multiple servers (100, 101, 102). The client can submit queries to the data in order to obtain the necessary data in a low latency fashion. In addition, as shown in the figure, the client can request a query explain plan from the cluster by providing invoking a method against the grid and passing in a structured query language (SQL) query as a parameter to the method. The query is then distributed to each of the server nodes in the cluster. Each node evaluates the query plan (109, 110, 111) against its own partition of data

(106, 107, 108) and returns the results of that query plan as it relates to its locally stored data. The results are then aggregated into a query result object 105 that is returned to the client. The query result object 105 provides information regarding how each node in the cluster evaluated the query plan for the original SQL query locally.

5 [0017] FIGURES 2A and 2B are illustrations of an example of using the query explain plan in accordance with various embodiments of the invention. It should be noted that while this figure includes certain code and data structure, the present embodiments are not limited to the particular code or data structure illustrated herein. Many other implementations are possible within the scope of the present invention, as will be evident to one of ordinary skill in
10 the art.

[0018] More specifically, FIGURE 2A shows an example of code invoking the query plan and FIGURE 2B shows a query record object being returned by the invocation of that code. In accordance with an embodiment, the query record object can be a string that when displayed would render the data illustrated in FIGURE 2B.

15 [0019] In accordance with an embodiment, the “Name” field of the query plan results denotes the name of the query plan step. In most cases this would be the invocation of a method “toString” on the filter of the query. The index is a reference to the index lookup if applicable. The “Cost” is the estimated cost as returned from the filter effectiveness calculation.

20 [0020] The results illustrated in FIGURE 2B show the filter(s) used for the query steps along with each index (if any), the size of the key set passed to the filter and the results of calculating the effectiveness of using the filter. For example, the results above can be for a query explain plan run against a cache with a 1000 entries in a cluster with 4 storage nodes.

[0021] It should be noted that the above is just an example of what the query plan results
25 may look like. It may be possible to aggregate the partial results of the query plan in various other meaningful ways. The following describes an architecture that will allow customization of the query plan results by extending the query plan aggregator and implementing a custom query plan results object.

[0022] Calculated Effectiveness

30 [0023] In accordance with an embodiment, the calculated effectiveness is an estimated measure of how well a filter can use a collection of indexes to filter a set of keys. This value can be used to determine the placement of filters in an ArrayFilter so that the most effective filters are applied first. The value returned from IndexAwareFilter.calculateEffectiveness is really the estimated evaluation cost of applying the filter.

[0024] As an illustration, Table 1 below shows a query for used cars where make=Honda and cost < 7000, with an index on make.

Table 1

```
cache.addIndex(new ReflectionExtractor("getMake"), true, null);
...
Filter filter = new AndFilter(
    new EqualsFilter(new ReflectionExtractor("getMake"), "Honda"),
    new LessFilter(new ReflectionExtractor("getCost"), 7000));

Set keys = cache.keySet(filter);
```

5 [0025] In accordance with an embodiment, the effectiveness for each filter in the example illustrated in Table 1 would be calculated as follows:

[0026] Equals Filter - Since there is an index that could be applied, the calculated effectiveness is given as 1. The value of 1 is used since the operation of applying the index will require just a single access to the index content. This is a pure evaluation cost value since this says nothing about how effectively the index would reduce the given key set. For example, if the cache only contains Hondas, then the application of this filter may be cheap, but not effective. If there was no index that could be applied here, then the calculated effectiveness would be EVAL_COST multiplied by the number of keys. EVAL_COST is a constant value, currently set at 1000. This is intended to show the relative cost of doing a full scan to reduce the key set. As stated above, this value is purely an evaluation cost and says nothing about the actual effectiveness of the operation to reduce the key set.

[0027] Less Filter - Since there is no index that could be applied here, the calculated effectiveness would be EVAL_COST multiplied by the number of keys.

[0028] And Filter - The calculated effectiveness of this filter is based on the effectiveness of the filters that make it up. In this example, the calculated effectiveness of the Equals Filter is 1 and the calculated effectiveness of the Less Filter is some value based on the number of keys. The underlying filters are weighted based on their calculated effectiveness in an ordered list with the most effective (actually least costly) coming first. The calculated effectiveness of this And Filter is the calculated effectiveness of the first filter in its list of filters.

[0029] In this example, the given set of keys would be reduced by the Equals Filter first and the Less Filter second. If there was no index on the make attribute but there was an index on cost then it would likely be reversed with the Less Filter being applied first. The current ordering of filters is based primarily on whether or not there is an index that can be applied. Whether or not this produces the best plan depends on the actual effectiveness (the ability of the filter to reduce the key set) of applying the filter.

[0030] If the index is taken out of the equation then both filters would return the same calculated effectiveness (EVAL_COST multiplied by the number of keys). This means that one cannot tell which filter would be applied first. In order to determine which would be best, an actual effectiveness would need to be calculated. For example, if only 2% of the entries in the cache are Hondas and 99% cost less than \$7000 then it makes sense to apply the Equals Filter first since that would reduce the key set the most. Conversely, if all the entries in the cache are Hondas, then it makes sense to apply the Less Filter first since the Equals Filter would have no effect in reducing the key set.

[0031] It may be possible in some cases to make assumptions and take an educated case about effectiveness without actually applying the filter. For example, an Equals Filter with an index. Instead of only returning 1, the estimate can be enhanced by including the key set reduction factor:

```
Set setMatch = (Set) mapIndex.getIndexContents().get(getValue());  
factorReduction = setKeys.size() / setMatch.size();
```

[0032] A variation of this value could be used as the effectiveness/cost.

[0033] FIGURE 3 is a flow chart diagram of providing the query explain plan in accordance with various embodiments of the invention. Although this figure depicts functional steps in a particular sequence for purposes of illustration, the process is not necessarily limited to this particular order or steps. One skilled in the art will appreciate that the various steps portrayed in this figure can be changed, rearranged, performed in parallel or adapted in various ways. Furthermore, it is to be understood that certain steps or sequences of steps can be added to or omitted from this process, without departing from the spirit and scope of the invention.

[0034] As shown in step 300, data is stored on a plurality of nodes in a cluster, each node being capable of executing a structured query language (SQL) query on the data stored on said node. In step 301, a request is received containing the SQL query from a client connected to the cluster. In step 302, the SQL query is submitted to each node in the cluster. Each node in the cluster evaluates a query plan for the SQL query using local data on that node and returns a set of information regarding how the SQL query is executed at that node. In step 303, a query record object is created that aggregates the information from all nodes in the cluster. The query record object contains information regarding how the SQL query is executed by each node in the cluster.

[0035] In accordance with various embodiments, the query explain plan can provide a method for gathering statistics and calculating a cost for running a query in the distributed data system. It can be done on a per request basis, and does not need to be performed as part of every query. In accordance with one embodiment, the explain plan process does not actually need to run the query but can instead calculate a cost of running a query with the given filter. The implementation of the query explain plan feature does not need to impact query performance in the distributed data system. The results of the explain plan should be consumable for a non programmer such that IT can add missing indices. Existing methods for calculating effectiveness of applying an index during filter evaluation can be examined and implemented to improve query performance and provide meaningful values for the explain plan. The implementation of the query explain plan can be used with a trace feature. Trace is similar to the query explain plan except that it would actually run the query and provide additional statistics about the query performance.

[0036] In accordance with an embodiment, the query explain plan facility can be executed in at least two modes: an explain plan mode and a trace mode. In the explain plan mode, the facility does not actually execute the query, but simply returns the query plans for each cluster member. In trace mode, the facility executes the query and returns the actual performance metrics on each node in the cluster, including for example the execution time of applying each filter and the sequence in which the filters were applied.

[0037] Architecture

[0038] Below are some examples of the various objects that can be used to enable the query explain plan feature in the distributed cache. It is noted that these objects are provided purely for purposes of illustration and are not intended to be limiting to all embodiments of the invention.

[0039] QueryRecorder – A QueryRecorder aggregator can be used for obtaining a query record. This aggregator can serve as both a signal to the distributed data grid that a query explain plan is being requested as well as an actual aggregator of explain plan results. The QueryRecorder will be created with one of two types; EXPLAIN or TRACE. Creating a QueryRecorder with EXPLAIN will produce a QueryRecord object that contains an estimated cost of the query execution. Creating a QueryRecorder with TRACE will produce a QueryRecord object that contains the actual cost of the query execution. In other words, doing a trace aggregation will actually execute the query and record the effectiveness (ability to reduce the entry set) at each filter step while an explain plan aggregation will simply walk each step and record an estimated cost (as described above).

[0040] QueryRecorderFilter – A filter interface that can extend EntryFilter can be created. In accordance with an embodiment, all filters will implement this interface. This interface includes new methods for explaining and tracing the filter. For most filters the call to explain will calculate the effectiveness of applying the filter through
5 IndexAwareFilter.calculateEffectiveness and record this information as part of a query.

[0041] Filter Implementations - The standard implementation of the query explain plan for most filters will simply record information about the application of the filter into a query record step. This information may include the size of the key set prior to applying the filter, the size of the key set after applying the filter, the duration and the actual or estimated cost of
10 applying the filter. For filters that extend ArrayFilter, the implementation of the query explain plan will create a new query record step for each sub filter and call the query explain plan on each. Users can be enabled to implement or extend QueryRecorderFilter. Filters that do not extend QueryRecorderFilter will still be explainable or traceable in the default manner.

[0042] QueryContext – The explain plan and trace methods of the QueryRecorderFilter
15 include the use of a query context. The context can provide access to a backing map context which will be initialized with any information required to perform a query scan (e.g. the map of indexes).

[0043] QueryRecord – The QueryPlan interface defines the standard results object returned from a query plan operation. Since QueryPlanAggregator can be extended, and the
20 return type of aggregateResults is Object, the use of the QueryPlan interface is not required. The QueryPlan interface includes an addResults method that allows the aggregator to combine partial results. The default implementation of QueryPlan can include various methods for accessing the merged results including all of the steps of the query plan. The QueryPlan interface also includes the other interfaces that make up the components of a query plan
25 result, including PartialResults, Step and ContainmentStep.

[0044] In an embodiment, each server node 100-102 may be implemented on a computer system. FIG. 4 is a block diagram that illustrates a computer system 400 upon which an
30 embodiment of the invention may be implemented. Computer system 400 includes a bus 402 or other communication mechanism for communicating information, and a processor 404 coupled with bus 402 for processing information. Computer system 400 also includes a main memory 406, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 402 for storing information and instructions to be executed by processor 404. Main memory 406 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 404. Computer

system 400 further includes a read only memory (ROM) 408 or other static storage device coupled to bus 402 for storing static information and instructions for processor 404. A storage device 410, such as a magnetic disk or optical disk, is provided and coupled to bus 402 for storing information and instructions.

5 [0045] Computer system 400 may be coupled via bus 402 to a display 412, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 414, including alphanumeric and other keys, is coupled to bus 402 for communicating information and command selections to processor 404. Another type of user input device is cursor control 316, such as a mouse, a trackball, or cursor direction keys for communicating direction
10 information and command selections to processor 404 and for controlling cursor movement on display 412.

[0046] The invention is related to the use of computer system 400 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 400 in response to processor 404 executing one or more
15 sequences of one or more instructions contained in main memory 406. Such instructions may be read into main memory 406 from another machine-readable medium, such as storage device 410. Execution of the sequences of instructions contained in main memory 406 causes processor 404 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to
20 implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0047] Additionally, each server node 100-102 can be connected with each other via a local network, such as network 422. Furthermore, the network 422 can also connect the server nodes 100-102 to a client 103 outside the cluster of server nodes 100-102.

25 [0048] Throughout the various contexts described in this disclosure, the embodiments of the invention further encompass computer apparatus, computing systems and machine-readable media configured to carry out the foregoing systems and methods. In addition to an embodiment consisting of specifically designed integrated circuits or other electronics, the present invention may be conveniently implemented using a conventional general purpose or
30 a specialized digital computer or microprocessor programmed according to the teachings of the present disclosure, as will be apparent to those skilled in the computer art.

[0049] Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art. The invention may also be implemented by the preparation of application

specific integrated circuits or by interconnecting an appropriate network of conventional component circuits, as will be readily apparent to those skilled in the art.

[0050] The various embodiments include a computer program product which is a storage medium (media) having instructions stored thereon/in which can be used to program a general purpose or specialized computing processor(s)/device(s) to perform any of the features presented herein. The storage medium can include, but is not limited to, one or more of the following: any type of physical media including floppy disks, optical discs, DVDs, CD-ROMs, microdrives, magneto-optical disks, holographic storage, ROMs, RAMs, PRAMS, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems (including molecular memory ICs); paper or paper-based media; and any type of media or device suitable for storing instructions and/or information. The computer program product can be transmitted in whole or in parts and over one or more public and/or private networks wherein the transmission includes instructions which can be used by one or more processors to perform any of the features presented herein. The transmission may include a plurality of separate transmissions. In accordance with certain embodiments, however, the computer storage medium containing the instructions is non-transitory (i.e. not in the process of being transmitted) but rather is persisted on a physical device.

[0051] The foregoing description of the preferred embodiments of the present invention has been provided for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations can be apparent to the practitioner skilled in the art. Embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby enabling others skilled in the relevant art to understand the invention. It is intended that the scope of the invention be defined by the following claims and their equivalents.

CLAIMS

What is claimed is:

1. A system for explaining queries in a distributed data grid, said system comprising:
5 one or more microprocessors;
a cluster running on the one or more microprocessors, including a plurality of cluster nodes storing data distributed throughout the cluster, each cluster node being capable of executing a structured query language (SQL) query on the data stored on said cluster node; and
10 a client connected to the cluster that submits a request that includes the SQL query;
wherein upon receiving the request, the SQL query is submitted to each node in the cluster and wherein each node in the cluster evaluates a query plan for the SQL query using local data at said each node and returns a set of information regarding how the SQL query is executed at said node; and
15 wherein a query record object is generated that aggregates the information from all nodes in the cluster, wherein the query record object contains information regarding how the SQL query is executed by each node in the cluster.
2. The system of claim 1, wherein the query record object further provides
20 execution tracing, rendering information about the execution time of each step and total execution time.
3. The system of claim 1 or 2, wherein to obtain a query explain plan, a method is invoked passing an SQL query as a parameter to said method.
25
4. The system of any preceding claim, wherein the SQL query is sent to all cluster nodes in the cluster and wherein each cluster node returns the determined optimal query plan specific to said each cluster node.
- 30 5. The system of any preceding claim, wherein a facility is provided to compare and contrast the different query plans and to reconcile the results of the query invocation.
6. The system of any preceding claim, wherein the system is executable in at least two modes: an explain plan mode and a trace mode, wherein the explain mode produces a

query record object that contains an estimated cost of the query execution, and wherein the trace mode contains the actual cost of query execution.

7. The system of any preceding claim, wherein an index is generated by
5 employing the query record object.

8. The system of any preceding claim, wherein each node in the cluster executes a different query plan depending on a partition of data stored on the node.

10 9. A method for explaining queries in a distributed data grid, said method comprising:

storing data on a plurality of nodes in a cluster running on one or more microprocessors, each node being capable of executing a structured query language (SQL) query on the data stored on said node;

15 receiving a request containing the SQL query from a client connected to the cluster;

submitting the SQL query to each node in the cluster;

each node in the cluster evaluating a query plan for the SQL query using local data at said each node and returning a set of information regarding how the SQL query is executed at said node; and

20 generating a query record object that aggregates the information from all nodes in the cluster, wherein the query record object contains information regarding how the SQL query is executed by each node in the cluster.

10. The method of claim 9, wherein the query record object further provides
25 execution tracing, rendering information about the execution time of each step and total execution time.

11. The method of claim 9 or 10, wherein the request is for a query explain plan and wherein to obtain the query explain plan, a method is invoked passing an SQL query as a
30 parameter to said method.

12. The method of any of claims 9 to 11, wherein the SQL query is sent to all cluster nodes in the cluster and wherein each cluster node can return the determined optimal query plan specific to said each cluster node.

13. The method of any of claims 9 to 12, wherein the method further comprises comparing and contrasting the different query plans and to reconcile the results of the query invocation.

5

14. The method of any of claims 9 to 13, wherein the method is executable in at least two modes: an explain plan mode and a trace mode, wherein the explain mode produces a query record object that contains an estimated cost of the query execution, and wherein the trace mode contains the actual cost of query execution.

10

15. The method of any of claims 9 to 14, wherein an index is generated by employing the query record object.

16. The method of any of claims 9 to 15, wherein each node in the cluster executes a different query plan depending on a partition of data stored on the node.

15

17. A computer program comprising program instructions for execution by a computer system, thereby causing the computer system to perform the method of any of claims 9 to 16.

20

18. A computer readable storage medium storing the computer program of claim 17.

19. A non-transitory computer readable storage medium storing one or more sequences of instructions executed by one or more processors to perform a set of steps comprising:

25

storing data on a plurality of nodes in a cluster, each node being capable of executing a structured query language (SQL) query on the data stored on said node;

receiving a request containing the SQL query from a client connected to the cluster;

30

submitting the SQL query to each node in the cluster, wherein each node in the cluster evaluates a query plan for the SQL query using local data as said each node and returns a set of information regarding how the SQL query is executed at said node; and

generating a query record object that aggregates the information from all nodes in the cluster, wherein the query record object contains information regarding how the SQL query is executed by each node in the cluster.

5 20. The non-transitory computer readable storage medium of claim 19, wherein the query record object further provides execution tracing, rendering information about the execution time of each step and total execution time.

10 21. The non-transitory computer readable storage medium of claim 19 or 20, wherein to obtain a query explain plan, a method is invoked passing an SQL query as a parameter to said method.

15 22. The non-transitory computer readable storage medium of any of claims 19 to 21, wherein the SQL query is sent to all cluster nodes in the cluster and wherein each cluster node can return the determined optimal query plan specific to said each cluster node.

23. A system for explaining queries in a distributed data grid, said system comprising:

20 means for storing data on a plurality of nodes in a cluster running on one or more microprocessors, each node being capable of executing a structured query language (SQL) query on the data stored on said node;

 means for receiving a request containing the SQL query from a client connected to the cluster;

25 means for submitting the SQL query to each node in the cluster, wherein each node in the cluster evaluates a query plan for the SQL query using local data as said each node and returns a set of information regarding how the SQL query is executed at said node; and

 means for generating a query record object that aggregates the information from all nodes in the cluster, wherein the query record object contains information regarding how the SQL query is executed by each node in the cluster.

30

24. The system of claim 23, wherein the query record object further provides execution tracing, rendering information about the execution time of each step and total execution time.

25. The system of claim 23, wherein to obtain a query explain plan, a method is invoked passing an SQL query as a parameter to said method.

26. The system of claim 23, wherein the SQL query is sent to all cluster nodes in the cluster and wherein each cluster node can return the determined optimal query plan specific to said each cluster node.

27. The system of claim 26, wherein a facility is provided to compare and contrast the different query plans and to reconcile the results of the query invocation.

28. The system of claim 23, wherein the system is executable in at least two modes: an explain mode and a trace mode, wherein the explain mode produces a query record object that contains an estimated cost of the query execution, and wherein the trace mode contains the actual cost of query execution.

29. The system of claim 23, wherein an index is generated by employing the query record object.

30. The system of claim 23, wherein each node in the cluster executes a different query plan depending on a partition of data stored on the node.

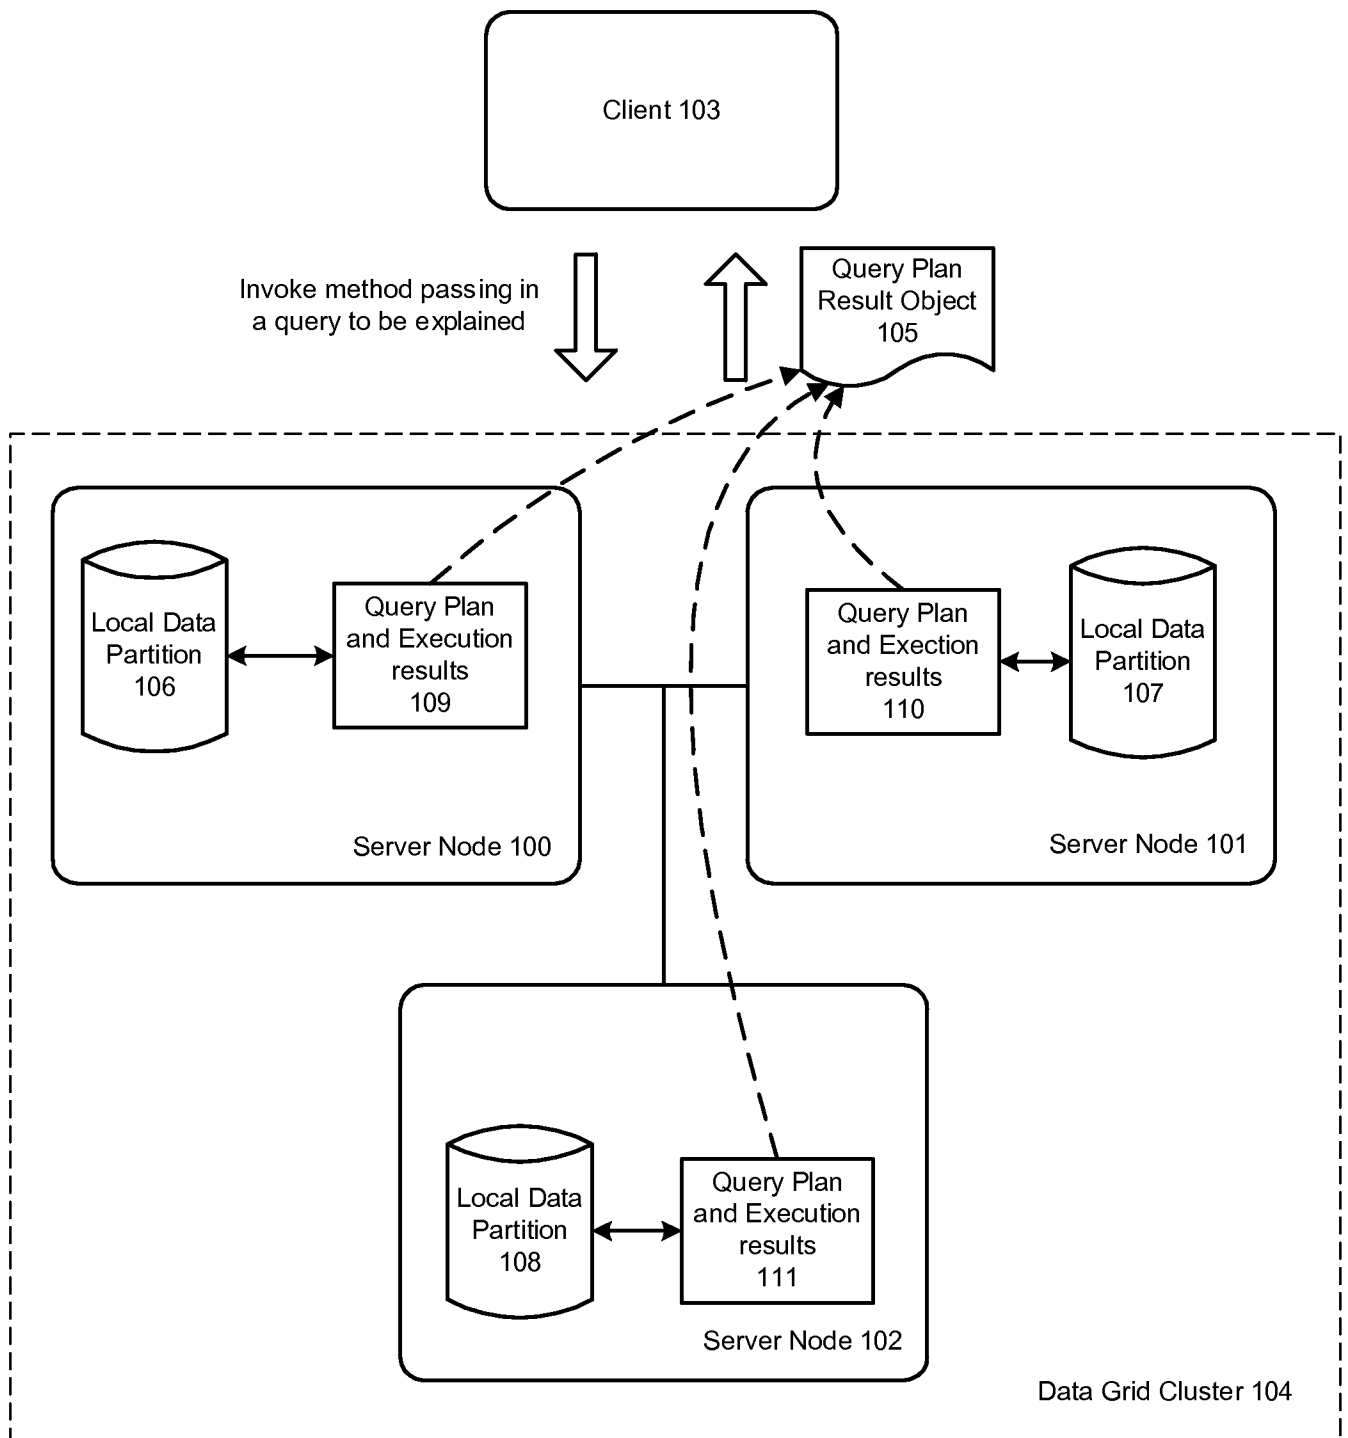
31. A system for explaining queries in a distributed data grid, said system comprising:

one or more microprocessors; and

a cluster running on the one or more microprocessors, including a plurality of cluster nodes storing data distributed throughout the cluster, each cluster node being capable of executing a structured query language (SQL) query on the data stored on said cluster node,

wherein upon receiving a request comprising the SQL query for a query explain plan from a client, the SQL query is submitted to each node in the cluster and wherein each node in the cluster evaluates a query plan for the SQL query using local data as said each node and returns a set of information regarding how the SQL query is executed at said node; and

wherein a query record object is generated that aggregates the information from all nodes in the cluster, wherein the query record object contains information regarding how the SQL query is executed by each node in the cluster.

**FIGURE 1**

```

NamedCache cache = CacheFactory.getCache("dist-person");

cache.addIndex(new ReflectionExtractor("getAge"), true, null);
...

Filter filter = new AllFilter(new Filter[]
{
    new OrFilter(
        new EqualsFilter(new ReflectionExtractor("getAge"), 16),
        new EqualsFilter(new ReflectionExtractor("getAge"), 19)),
    new EqualsFilter(new ReflectionExtractor("getLastName"), "Smith"),
    new EqualsFilter(new ReflectionExtractor("getFirstName"), "Bob"),
});

QueryRecorder agent = new QueryRecorder(RecordType.EXPLAIN);

Object resultsExplain = cache.aggregate(FILTER, agent);

```

FIGURE 2A

Explain Plan=

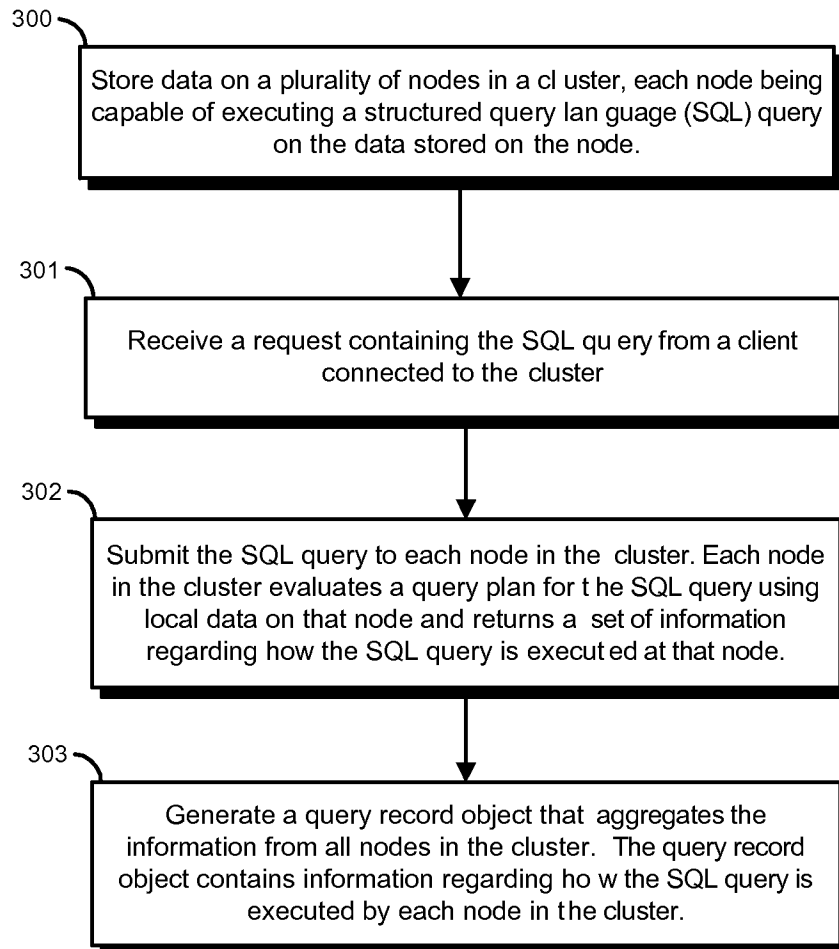
Explain Plan

Name	Index	Cost
com.tangosol.util.filter.AllFilter	----	0
com.tangosol.util.filter.OrFilter	----	0
EqualsFilter(.getAge(), 16)	0	1
EqualsFilter(.getAge(), 19)	0	1
EqualsFilter(.getLastName(), Smit	1	1500000
EqualsFilter(.getFirstName(), Bob	2	1500000

Index Lookups

Index	Description	Extractor
0	SimpleMapIndex: Extractor=.getAge(), Ord	.getAge() true
1	No index found	.getLastName() false
2	No index found	.getFirstName() false

FIGURE 2B

**FIGURE 3**

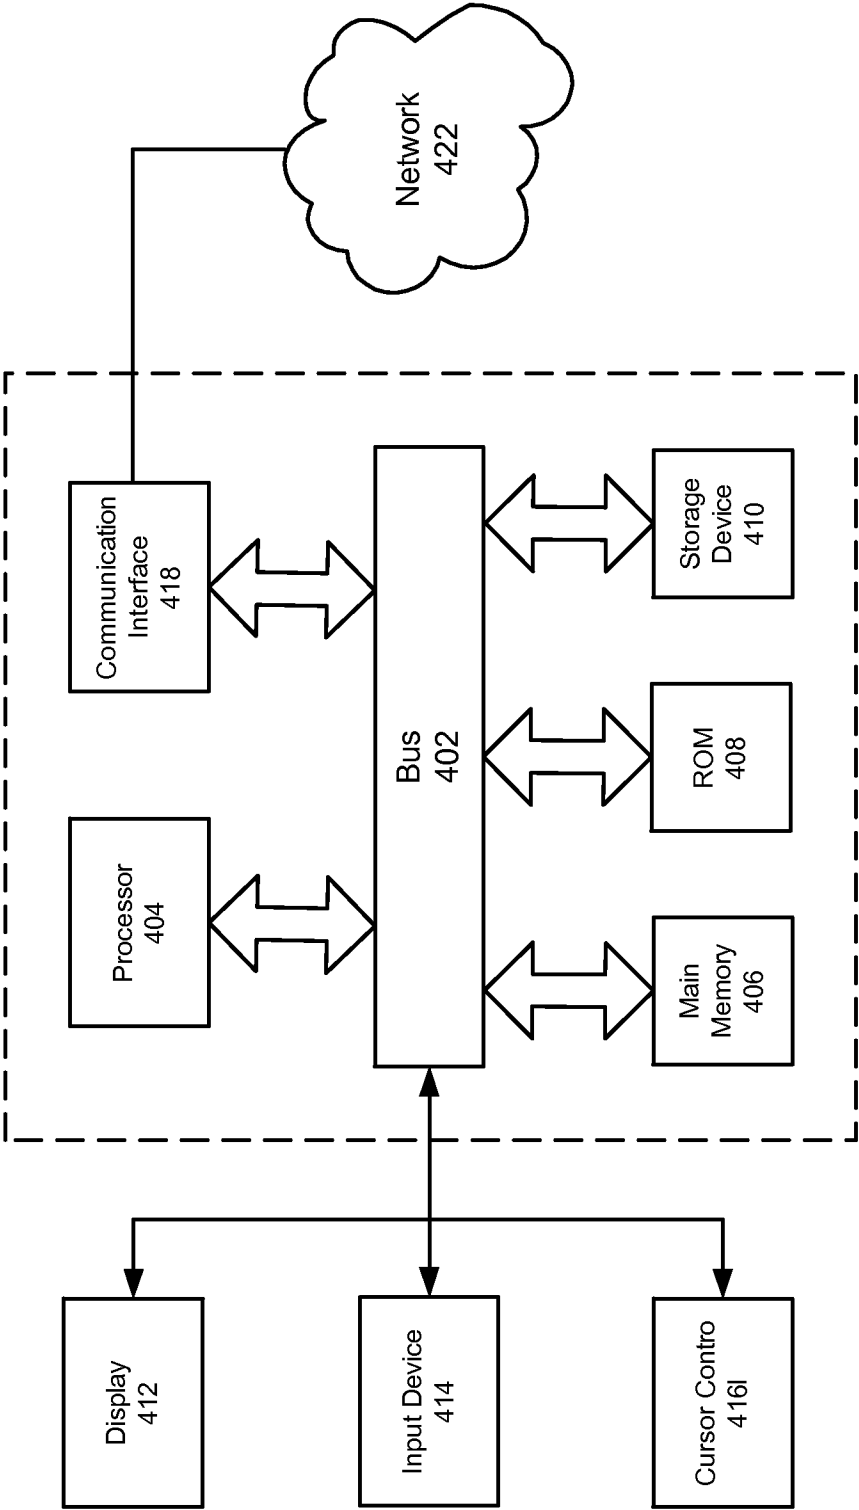


FIGURE 4

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2012/054680

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06F17/30
 ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2004/030739 A1 (YOUSEFI ZADEH HOMAYOUN [US]) 12 February 2004 (2004-02-12) paragraph [0010] paragraph [0031] -----	1-31
A	US 7 185 000 B1 (BROWN DOUGLAS P [US] ET AL) 27 February 2007 (2007-02-27) column 1, line 29 - line 45 column 2, line 46 - column 3, line 53 column 3, line 54 - column 4, line 45 column 5, line 24 - line 36 -----	1-31
A	US 6 195 653 B1 (BLEIZEFFER TERRY MICHAEL [US] ET AL) 27 February 2001 (2001-02-27) column 1, line 56 - column 2, line 57 column 2, line 60 - column 3, line 16 ----- -/--	1-31



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

8 February 2013

Date of mailing of the international search report

18/02/2013

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040,
 Fax: (+31-70) 340-3016

Authorized officer

Correia Martins, F

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2012/054680

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>KOSSMANN D: "THE STATE OF THE ART IN DISTRIBUTED QUERY PROCESSING", ACM COMPUTING SURVEYS, ACM, NEW YORK, NY, US, US, vol. 32, no. 4, 1 December 2000 (2000-12-01), pages 422-469, XP002952901, ISSN: 0360-0300, DOI: 10.1145/371578.371598 page 426, right-hand column, paragraph 2 -----</p>	1-31

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2012/054680

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004030739	A1	12-02-2004	NONE

US 7185000	B1	27-02-2007	US 7185000 B1 27-02-2007
		US 7234112	B1 19-06-2007

US 6195653	B1	27-02-2001	US 6195653 B1 27-02-2001
		US 6397207	B1 28-05-2002
