



US 20100082603A1

(19) **United States**

(12) **Patent Application Publication**
Krompass et al.

(10) **Pub. No.: US 2010/0082603 A1**

(43) **Pub. Date: Apr. 1, 2010**

(54) **MANAGING EXECUTION OF DATABASE QUERIES**

(22) Filed: **Sep. 30, 2008**

(76) Inventors: **Stepan Krompass**, Garching (DE);
Harumi Anne Kuno, Cupertino, CA (US); **Umeshwar Dayal**,
Saratoga, CA (US); **Janet Wiener**,
Palo Alto, CA (US); **Raja Sambasivan**, Ossining, NY (US)

Publication Classification

(51) **Int. Cl.**
G06F 17/30 (2006.01)

(52) **U.S. Cl.** **707/719; 707/E17.017**

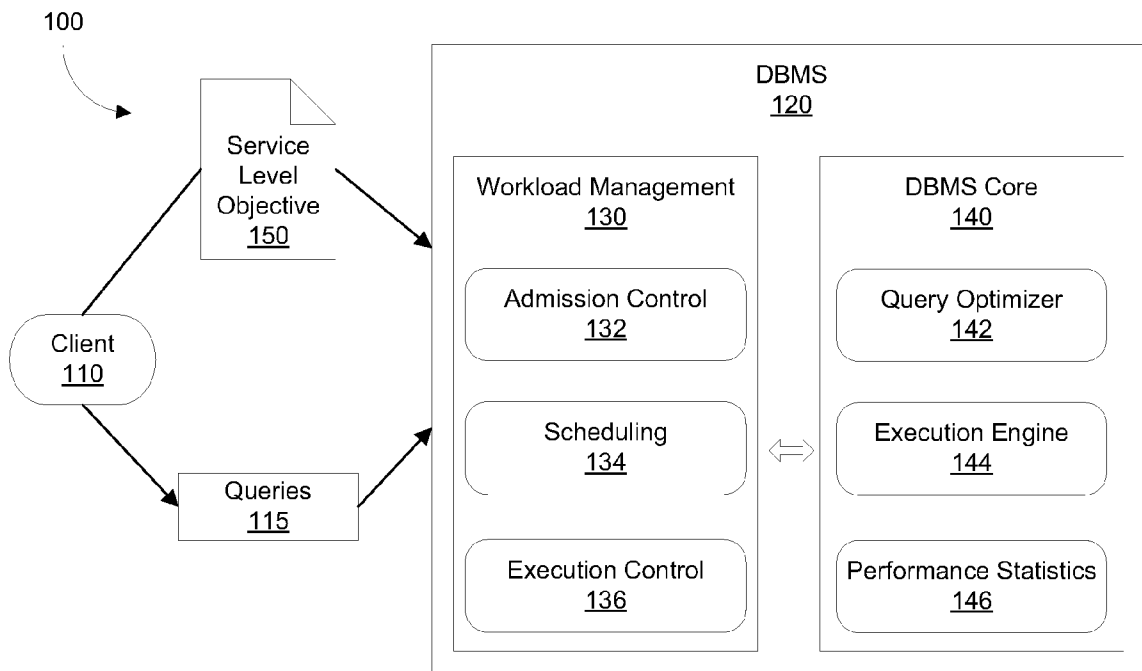
(57) **ABSTRACT**

One embodiment is a method to manage queries in a database. The method identifies a query that executes on the database for an elapsed time that is greater than a threshold and then implements a remedial action when the query executes on the database for an execution time that is greater than an estimated execution time.

Correspondence Address:

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
3404 E. Harmony Road, Mail Stop 35
FORT COLLINS, CO 80528 (US)

(21) Appl. No.: **12/242,637**



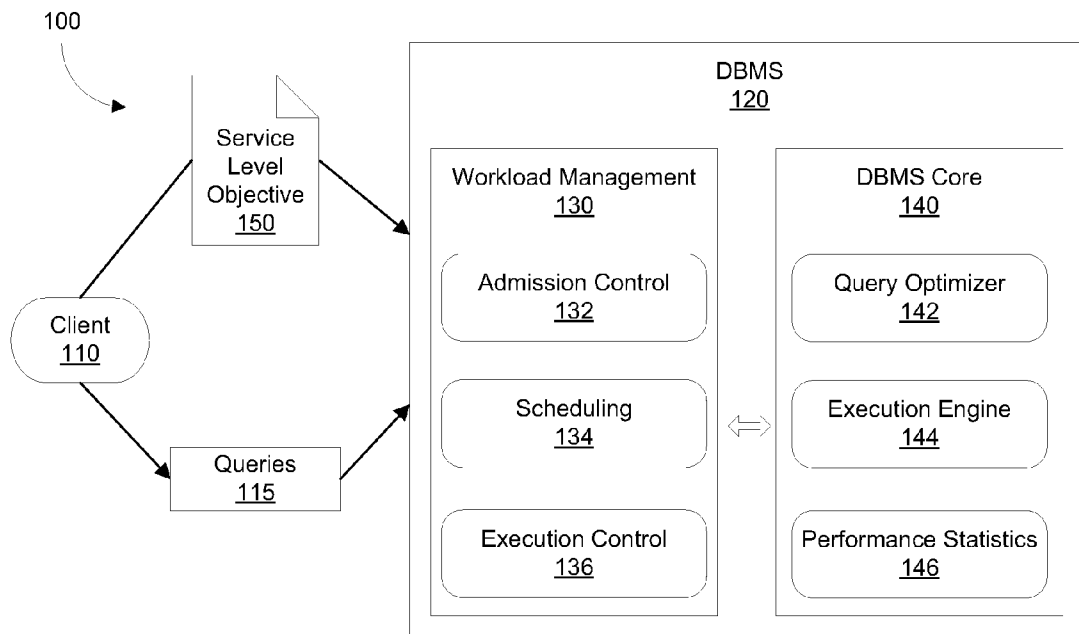


FIG. 1

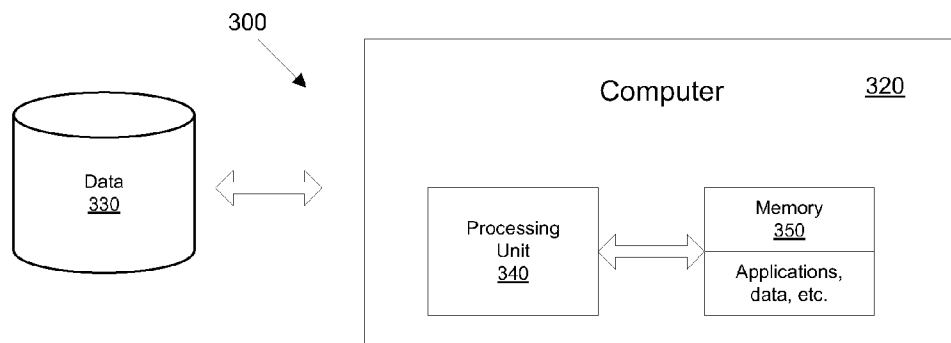


FIG. 3

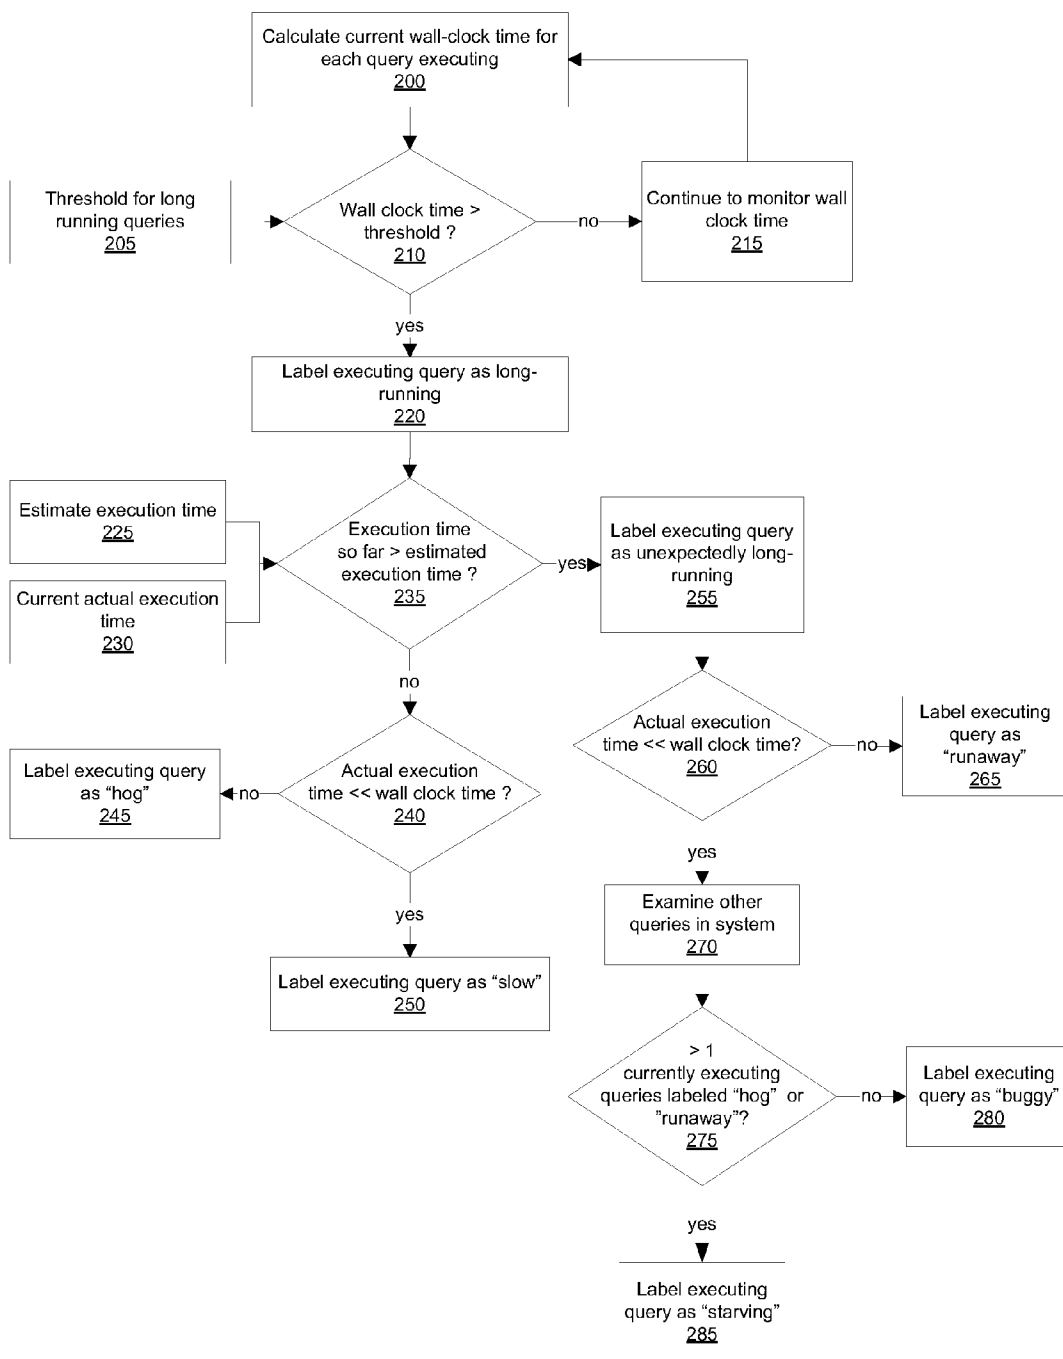


FIG. 2

MANAGING EXECUTION OF DATABASE QUERIES

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application relates to commonly assigned U.S. patent applications having attorney docket number HP 200704074-1 entitled “Reverse Mapping of Feature Space to Predict Execution in a Database” and filed on Jul. 5, 2008; attorney docket number HP 200704075-1 entitled “Predicting Performance of Executing a Query in Isolation in a Database” and filed on Jul. 5, 2008; attorney docket number HP 200704077-1 entitled “Characterizing Queries to Predict Execution in a Database” and filed on Jul. 5, 2008; attorney docket number HP 200704091-1 entitled “Predicting Performance of Multiple Queries Executing in a Database” and filed on Jul. 5, 2008, all of which are incorporated herein by reference.

BACKGROUND

[0002] Business Intelligence (BI) database systems process extremely complex queries upon massive amounts of data. This capability is important to modern enterprises, which rely upon the collection, integration, and analysis of an expanding sea of information. Analysts predict that data warehouses will grow in size and complexity.

[0003] BI queries exhibit a significant variance in response times. Most queries are known to execute in under a minute, but some small number of them require hours of execution time. According to conventional wisdom, the presence of even a few poorly written or poorly optimized queries can significantly impact the performance of a data warehouse system by taking up resources that could otherwise be used by properly functioning queries. It is not straightforward to estimate accurately how long a long-running query will take. Although customers may have service level agreements (SLAs) that spell out financial consequences, and although jobs have deadlines, the deadlines are not necessarily explicit or straightforward (or even necessarily possible) to link SLAs to deadlines.

[0004] Long running BI queries consume large resources and pose problems for administrators of databases. Administrators are regularly forced to decide which queries are hurting system performance and what action to take against such long running queries.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 is a system for managing the execution of database queries in accordance with an exemplary embodiment of the present invention.

[0006] FIG. 2 is a flow diagram for categorizing long-running database queries that are executing in accordance with an exemplary embodiment of the present invention.

[0007] FIG. 3 is a block diagram of an exemplary computer system in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0008] Exemplary embodiments in accordance with the present invention are directed to systems and methods for managing the execution of database queries, including queries that execute for unexpectedly long times. The systems and methods manage the execution of a set of queries in a

database system that is capable of distinguishing between various types of long-running queries so as to take appropriate management actions when necessary.

[0009] Exemplary embodiments distinguish between different types of long-running queries based on their performance characteristics and take different actions to address these types of long-running queries. The likely impact upon the performance of the system is characterized for each type of long-running query. This information enables database administrators to identify long-running queries but also make informed decisions about what remedial action to take (for example, cancel or kill the query, abort the query, let the query complete, suspend the query, etc.).

[0010] FIG. 1 is a database system 100 for managing the execution of database queries in accordance with an exemplary embodiment of the present invention. The system generally includes a computer or client 110 that sends queries 115 to a Database Management System (DBMS) 120 which includes a workload management component 130 and a DBMS core 140. The workload management component includes plural components or modules as admission control 132, scheduling 134, and execution control 136. The DBMS core 140 includes plural components or modules as a query optimizer 142, an execution engine 144, and performance statistics 146. Further, Service Level Objectives (SLOs) 150 are coupled between the client 110 and the DBMS 120.

[0011] The workload management architecture 130 provides fundamental workload management functionality for admission control, scheduling, and execution control. In one embodiment, each of these modules 132, 134, and 136 represents a knob that can be adjusted to select from a variety of workload management policies and algorithms.

[0012] The DBMS core 140 provides core database functionality and supply information to workload management components but does not implement workload management policies.

[0013] In one embodiment, the database system executes workloads that include one or more jobs. Each job consists of an ordered set of typed queries 115 submitted by a computer or client 110, and is associated with one or more Service Level Objectives (SLOs). Each query type maps to a tree of operators, and each operator in a tree maps in turn to its resource costs.

[0014] Policies of the admission control 132 determine the submission of queries 115 to the execution engine 144 that executes the submitted queries. The admission control 132 performs three functions in workload management. First, when a new job arrives, admission control 132 evaluates the DBMS’s multiprogramming level, and either submits or enqueues each of the job’s queries. Second, the architecture is configurable to support multiple admission queues. Policies of the admission control 132 regulate the distribution of queries among these queues, for example adding queries to queues based on estimated cost or dominant resource. Third, when the execution engine 144 has finished processing a query, admission control 132 selects the next query for execution. In one embodiment, this selection is performed by evaluating the amount of work waiting in each queue (if there are multiple queues), and selecting the first query from the queue with the most work remaining, or using another selection scheme (for example, one based on priority, job type, first-in-first-out, FIFO, etc.).

[0015] Once queries have been enqueued, the policies of the scheduler 134 determine the ordering of the queries

within a queue (for example, by estimated cost). Policies of the execution control **136** then govern the flow of the running system. As queries execute, the execution control **136** uses runtime statistics to identify problems such as overload situations or queries that are unexpectedly long-running. The execution control **136** then invokes appropriate policies to address these problems. For each executing query, the execution control **136** can monitor various metrics, such as a number of cancellations (i.e., a count of a number of times an individual query is killed), progress (i.e., a ratio of monitored to expected work terms of resource utilization), resource contention (i.e., number of queries that are competing with a query for resources), and database time (i.e., time a query is executing in the database, as opposed to waiting in a queue).

[0016] The control knobs of the workload management component **130** can be set to provide policies and corrective actions for long-running queries. The policies are the settings for the admission control **132**, scheduling **134**, and execution control **136**. The corrective actions are remedial steps taken by the execution control **136** for a given query (for example, an identified hog, runaway, slow, buggy, or starving query).

[0017] For each query executing in the database, exemplary embodiments track various parameters or characteristics with respect to time. By way of example, exemplary embodiments track estimated execution time, elapsed time, and actual consumed execution time.

[0018] The estimated execution time represents the amount of execution time that the query optimizer **142** estimates is necessary to execute this query. In one embodiment, this estimation is provided based on the query executing in isolation. Alternatively, if the query optimizer provides estimates in terms of cardinalities instead of execution time, embodiments also estimate the time needed to process a given tuple. This estimate is then used to translate from cardinality to execution time estimates.

[0019] The elapsed time represents the amount of time that the query has been running. This time can also be described as the “wall clock time” for the query. The consumed execution time represents the amount of execution time that the query has consumed thus far.

[0020] Exemplary embodiments also define a threshold that identifies long-running queries. As used herein, a query is “long-running” when it has an elapsed time that exceeds this threshold.

[0021] These different parameters (i.e., long-running query threshold, estimated elapsed time, actual elapsed time, “fair” resource usage, actual resource usage, and progress) are used to build and define a taxonomy of query performance problems. For example, exemplary embodiments identify and distinguish between different types of long-running queries. These types include queries with the following exemplary nomenclature:

[0022] (1) “heavy” queries: Queries that consume a large, but expected and fair share amount of system resources and are making good progress.

[0023] (2) “hog” queries: Queries that consume an unfairly large, but expected, amount of system resources and are making good progress.

[0024] (3) “runaway” queries: Queries that consume an unexpectedly large amount, but still a fair share, of system resources and are making good progress.

[0025] (4) “surprise-hog” queries: Queries that make good progress and consume an unexpectedly-large and

unfairly large share of system resources, potentially depriving other queries of resources.

[0026] (5) “starving” queries: Queries that take an unexpectedly long time to execute, but consume less than their fair share of resources and are making poor progress because of resource contention with other queries that are taking more than their fair share of resources.

[0027] (6) “buggy” queries: Queries that take an unexpectedly long time to execute and are not being deprived of their fair share of resources, but consume fewer resources than expected and are making poor progress because of a problem with the database system itself (e.g., a poorly implemented query operator, or an unrealistically high multi programming level (MPL)).

[0028] (7) “overload” queries: Queries that are not being deprived of their fair share of resources, but take an unexpectedly long time to execute and are making poor progress because of heavy system load.

[0029] Based on the developed taxonomy, one exemplary method first determines whether or not a query is long-running. A check or verification is then made as to whether the behavior of the long-running query is expected and/or using large amounts of system resources. Processing time (CPU time) is an example of a proxy for resource usage. Any query that performs a relatively large amount of disk Input/Output (I/O) or message transfer is also likely to use large CPU time processing the bytes read or transferred.

[0030] By way of example, heavy queries take a long time, but predictably so and allow other queries to make progress. Hog queries are also predictably long, but use more than their share of the resources. They can be scheduled when they are least likely to interfere with other queries. Runaway and surprise-hog queries behave just like heavy and hog queries—but without warning. They are the most likely to cause problems for other concurrent queries and the most important to catch. Killing (and possibly rescheduling) runaway and surprise-hog queries has the most impact on the completion time of the other queries in the workload. Starving queries are the ones impacted by hog and surprise-hog queries: they ought to be short, but are taking a long time because the hog queries do not leave them enough resources. Starving queries that are killed and restarted when there is less contention will run faster. Overload is a property of the entire system rather than a single query. It describes a system that just does not have enough resources for the current MPL, even when all of the queries are supposed to be short. There are too many queries executing concurrently and so none of them get enough resources to make reasonable progress. Running those queries at the same MPL at another time is unlikely to change their performance. The best solution is to reduce the MPL.

[0031] FIG. 2 is a flow diagram for categorizing long-running database queries that are executing in accordance with an exemplary embodiment of the present invention.

[0032] According to block **200**, a calculation is made of the current wall-clock time for each query that is executing in the database. In other words, determine the elapsed time or amount of time that each query has been running.

[0033] According to block **205**, one or more thresholds are defined for long-running queries. By way of example, when the elapsed time of the threshold is exceeded, then the executing query is deemed to be long-running.

[0034] In addition to defining long-running queries and providing numerical values for the corresponding thresholds, definitions and/or parameters from the taxonomy are also input into the method. As noted above, such parameters include, but are not limited to, estimated execution time, actual execution time, elapsed time, and classifications for long-running queries (for example, hog, runaway, starving, buggy, and slow).

[0035] According to block **210**, a determination is made as to whether the wall clock time is greater than a threshold. If the answer to this question is “no” then flow proceeds to block **215**. Here the wall clock time is continued to be monitored (for example, continuously or periodically), and flow proceeds back to block **200**. If the answer to the question is “yes” then flow proceeds to block **220**. Here, the query is labeled as being long-running since the wall clock time is greater than a predetermined threshold value (for example, a numerical value, a percentage, etc.).

[0036] Various different types of long-running queries are defined (for example, hog, runaway, starving, buggy, and slow). The remedial action (if any) depends on the type of long-running query that is identified. As such, the next portion of the method determines or identifies which type of long-running query is executing.

[0037] According to block **225**, an estimate is provided for the execution time. In one embodiment, the query optimizer estimates the time necessary to execute this query.

[0038] According to block **230**, input of the current actual execution time is provided. This time represents the amount of execution time that the query has consumed thus far.

[0039] According to block **235**, a question is asked whether the execution time so far is greater than the estimated execution time. In other words, a determination is made as to whether the actual execution time for the currently executing query is greater than the predicted or estimated execution for the query before execution commenced.

[0040] If the answer to the question in block **235** is “no” then flow proceeds to block **240**. Flow in this direction means that the actual execution time of the running query is less than the estimated execution time.

[0041] At block **240**, a question is asked whether the actual execution time is much less than the wall clock time. If the answer to this question is “yes” then flow proceeds to block **245** where the query is labeled as being a hog. Hog queries are predictable in that they take a long time to execute and utilize a relatively large amount of resources (for example, CPU resources).

[0042] If the answer to the question in block **240** is “no” then flow proceeds to block **250** where the query is labeled as being slow. Slow queries are predictable in that they take a long time to execute and utilize a relatively small amount of resources as compared to resource utilization of hog queries.

[0043] Going back to block **235**, if the answer to the question is “yes” then flow proceeds to block **255**. Flow in this direction means that the actual execution time of the running query is greater than the estimated execution time. Queries in these categories are labeled as being unexpectedly long-running queries (for example, include one of runaway, buggy, or starving).

[0044] According to block **260**, a question is asked whether the actual execution time for the unexpectedly long-running query is much less than the wall clock time. If the answer to this question is “no” then flow proceeds to block **265** and the query is labeled as being a runaway. Runaway queries behave

like hog queries (i.e., execute with a long time and consume large resources) but have no warning because they are unexpected.

[0045] If the answer to the question in block **260** is “yes” then flow proceeds to block **270**. Here, the query can be either buggy or starving. To determine how to label the query under examination, other currently executing queries in the system are examined.

[0046] Flow then proceeds to block **275** where a question is asked whether more than one currently executing query has a label of hog or runaway. In other words, another high resource query can simultaneously be executing with the query under examination.

[0047] If the answer to the question is “no” then flow proceeds to block **280**, and the query is labeled as being buggy. Buggy queries take an unexpectedly long time to execute but consume fewer resources than expected due to a problem with the database system (for example, a poorly implemented query operator).

[0048] If the answer to the question is “yes” then flow proceeds to block **285**, and the query is labeled as being starving. Starving queries take an unexpectedly long time to execute but consume fewer resources than expected because of competition or contention for another simultaneously executing query, such as a hog query or runaway query.

[0049] Once a long-running query is identified various actions can be taken to manage a running workload of queries. These actions include execution control actions such as kill (i.e., stop a running query and free resources consumed by the query), stop (i.e., stop a running query, return partial results to the requester, and free resources consumed by the query), resubmit (i.e., restart a killed query), suspend (i.e., save current state of an executing query and free resources), resume (i.e., restore state of a suspended query and resume execution), hold (i.e., do not admit new queries to the database system), and release (i.e., resume admitting new queries to the database), as well as admission control and scheduling actions such as restricting the admission or scheduling of new queries so as not to exacerbate a system overload situation.

[0050] Once a long-running query is identified, the particular action taken against the query depends on user defined parameters. As one example, a database administrator (DBA) can decide to input the following parameters into the database system: Kill starving or buggy queries if they run for longer than 120% of the expected execution time and are less than 80% complete. The killed queries can be immediately resubmitted to the database system.

[0051] As another example, runaway queries are not executed simultaneously with other queries. Since the objective is to optimize the flow time of the interactive job, one embodiment identifies and removes the batch runaway queries from the job. For identifying runaway queries, if a query consumes significantly more resources in relation to the others over a longer period of time, the query is removed from the system.

[0052] Exemplary embodiments support plural actions for moving a runaway query out of the way. For instance, the query is killed and not resubmitted. In this case an administrator or system can identify the cause of the problem before rerunning the query. Alternatively, the administrator or system kills the query and resubmits each of the killed queries sequentially after the jobs are complete. As another alterna-

tive, the database administrator or system suspends the problem queries and resumes them one after another after completing the other jobs.

[0053] Exemplary embodiments are not limited to any particular type of remedial action. By way of example, the admission control component (see FIG. 1) can determine whether newly arriving queries are admitted to the system to prevent the database system from being overloaded. Various policies can be set to provide such limits, such as providing an upper limit on a number of queries allowed to run concurrently, an upper limit on a number of concurrent users of the system, an upper limit for system resource usage, and an upper limit on a number of concurrent queries that access certain database objects (for example, views, indexes, and tables).

[0054] Optimizer costs and CPU time can also be used to prevent long-running queries from running in the system. For example, queries are rejected that have an optimizer cost exceeding a threshold specified by the DBA. Other admission controls can also be implemented with respect to long-running queries, such as no action (let the query execute to completion), warn (print or send a message to a log while the query continues), stop (stop processing the query and return results), kill (abort the query and return an error), and re-prioritize (change the priority of the query).

[0055] Embodiments in accordance with the present invention are utilized in or include a variety of systems, methods, and apparatus. FIG. 3 illustrates an exemplary embodiment as a computer system 300 for being or utilizing one or more of the computers, methods, flow diagrams and/or aspects of exemplary embodiments in accordance with the present invention.

[0056] The system 300 includes a computer system 320 (such as a host or client computer) and a repository, warehouse, or database 330. The computer system 320 comprises a processing unit 340 (such as one or more processors or central processing units, CPUs) for controlling the overall operation of memory 350 (such as random access memory (RAM) for temporary data storage and read only memory (ROM) for permanent data storage). The memory 350, for example, stores applications, data, control programs, algorithms (including diagrams and methods discussed herein), and other data associated with the computer system 320. The processing unit 340 communicates with memory 350 and data base 330 and many other components via buses, networks, etc.

[0057] Embodiments in accordance with the present invention are not limited to any particular type or number of databases and/or computer systems. The computer system, for example, includes various portable and non-portable computers and/or electronic devices. Exemplary computer systems include, but are not limited to, computers (portable and non-portable), servers, main frame computers, distributed computing devices, laptops, and other electronic devices and systems whether such devices and systems are portable or non-portable.

[0058] Exemplary embodiments do not require any expensive statistics to be maintained and are more flexible than setting absolute limits. By using the amount of query execution time in lieu of input and output cardinalities, expensive ongoing statistics are avoided. Furthermore, by considering the amount of progress done in terms of elapsed time, estimated execution time, and actual execution time, exemplary

embodiments identify problem queries based on their performance characteristics, as opposed to strictly by the amounts of resources used.

[0059] Definitions

[0060] As used herein and in the claims, the following words have the following definitions:

[0061] The terms “automated” or “automatically” (and like variations thereof) mean controlled operation of an apparatus, system, and/or process using computers and/or mechanical/electrical devices without the necessity of human intervention, observation, effort and/or decision.

[0062] A “database” is a structured collection of records or data that are stored in a computer system so that a computer program or person using a query language can consult it to retrieve records and/or answer queries. Records retrieved in response to queries provide information used to make decisions. Further, the actual collection of records is the database, whereas the DBMS is the software that manages the database.

[0063] A “database administrator” or “DBA” is a person who defines or manages a database or controls access to the database.

[0064] A “database management system” or “DBMS” is computer software designed to manage databases.

[0065] The term “elapsed time” or “wall clock time” means the actual time as measured by a clock that the query has been executing or running.

[0066] The term “execution time” means the amount of time the processor or CPU is actually executing instructions. During the execution of most programs, the CPU sits idle much of the time while the computer performs other tasks, such as fetching data from the keyboard or disk, or sending data to an output device. The execution time is, therefore, generally less than the wall-clock time (i.e., the actual time as measured by a clock that the query has been running) and includes the time a query is executing in the database and not the time waiting in a queue.

[0067] The term “long-running query” means a query executing or running for an elapsed time that exceeds a threshold.

[0068] The term “Service Level Objective” or “SLO” is a key element of a Service Level Agreement (SLA) between a Service Provider and a customer. SLOs are agreed as a means of measuring the performance of the Service Provider and are outlined as a way of avoiding disputes between the two parties based on misunderstanding. The SLA is the contract or agreement that specifies what service is to be provided, how it is supported, times, locations, costs, performance, and responsibilities of the parties involved. SLOs are specific measurable characteristics of the SLA such as availability, throughput, frequency, response time, or quality. Further, the SLOs can include one or more quality-of-Service (QoS) measurements that are combined to produce the SLO achievement value.

[0069] A “workload” is a set of queries used for the data warehouse.

[0070] In one exemplary embodiment, one or more blocks or steps discussed herein are automated. In other words, apparatus, systems, and methods occur automatically.

[0071] The methods in accordance with exemplary embodiments of the present invention are provided as examples and should not be construed to limit other embodiments within the scope of the invention. For instance, blocks in flow diagrams or numbers (such as (1), (2), etc.) should not be construed as steps that must proceed in a particular order.

Additional blocks/steps may be added, some blocks/steps removed, or the order of the blocks/steps altered and still be within the scope of the invention. Further, methods or steps discussed within different figures can be added to or exchanged with methods of steps in other figures. Further yet, specific numerical data values (such as specific quantities, numbers, categories, etc.) or other specific information should be interpreted as illustrative for discussing exemplary embodiments. Such specific information is not provided to limit the invention.

[0072] In the various embodiments in accordance with the present invention, embodiments are implemented as a method, system, and/or apparatus. As one example, exemplary embodiments and steps associated therewith are implemented as one or more computer software programs to implement the methods described herein. The software is implemented as one or more modules (also referred to as code subroutines, or “objects” in object-oriented programming). The location of the software will differ for the various alternative embodiments. The software programming code, for example, is accessed by a processor or processors of the computer or server from long-term storage media of some type, such as a CD-ROM drive or hard drive. The software programming code is embodied or stored on any of a variety of known media for use with a data processing system or in any memory device such as semiconductor, magnetic and optical devices, including a disk, hard drive, CD-ROM, ROM, etc. The code is distributed on such media, or is distributed to users from the memory or storage of one computer system over a network of some type to other computer systems for use by users of such other systems. Alternatively, the programming code is embodied in the memory and accessed by the processor using the bus. The techniques and methods for embodying software programming code in memory, on physical media, and/or distributing software code via networks are well known and will not be further discussed herein.

[0073] The above discussion is meant to be illustrative of the principles and various embodiments of the present invention. Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

What is claimed is:

1. A method for managing execution of database queries, comprising:
 - identifying a long-running query that executes on a database for an elapsed time that is greater than a threshold;
 - identifying when the long-running time is expected or unexpected with regard to optimizer cost estimations; and
 - implementing a remedial action when the long-running query executes on the database for an execution time that is greater than an estimated execution time.
2. The method of claim 1 further comprising, taking a remedial action against the long-running query when the elapsed time for the long-running query executing on the database is greater than the execution time.
3. The method of claim 1 further comprising, identifying when other queries simultaneously executing with the long-running query have execution times that are greater than estimated execution times.

4. The method of claim 1 further comprising:
 - aborting execution of the long-running query when the execution time is not << (much less than) the elapsed time;
 - allowing execution of the long-running query to continue when the execution time is << (much less than) the elapsed time.
5. The method of claim 1 further comprising:
 - identifying different types of long-running queries that are simultaneously executing the database;
 - implementing a different remedial action for each of the different types of long-running queries.
6. A tangible computer readable storage medium having instructions for causing a computer to execute a method, comprising:
 - identifying different types of long-running queries simultaneously executing on a database;
 - characterizing an impact on performance of the database for each of the different types of long-running queries; and
 - taking a remedial measure for one of the different types of long-running queries.
7. The tangible computer readable storage medium of claim 6 further comprising, identifying different types of queries that include (1) a query that consumes a large but unexpected amount of system resources, and (2) a query that consumes an unexpectedly large amount of system resources and deprives other queries of the system resources.
8. The tangible computer readable storage medium of claim 6 further comprising, identifying different types of queries that include (1) queries that execute for an unexpectedly long time but consume fewer resources than expected because of resource contention by other executing queries, and (2) queries that execute for an unexpectedly long time but consume fewer resources than expected because of a problem with a database system.
9. The tangible computer readable storage medium of claim 6 further comprising, identifying starving queries as queries that execute for an unexpectedly long time while consuming an expected or less than expected amount of system resources, and killing and then restarting the starving queries when less contention exists for resources on the database.
10. The tangible computer readable storage medium of claim 6 further comprising, identifying a query as being long-running when the query executes for an elapsed time greater than a predetermined threshold value and classifying the long-running query according to a taxonomy that includes plural different types of long-running queries.
11. A database system, comprising:
 - a database;
 - a memory for storing an algorithm; and
 - a processor for executing the algorithm to:
 - identify queries that executes in the database for an elapsed time that is greater than a threshold; and
 - identify one of the queries as being a long-running query when it executes on the database for an execution time that is greater than an estimated execution time.
12. The computer system of claim 11, wherein the processor further executes the algorithm to take a remedial action against the long-running query, the remedial action including

one of canceling the long-running query, aborting the long-running query, suspending the long-running query, or letting the long-running query complete.

13. The computer system of claim **11**, wherein the processor further executes the algorithm to distinguish between different types of long-running queries executing on the database so as to take a corrective action against the different types of long-running queries.

14. The computer system of claim **11**, wherein the processor further executes the algorithm to identify different types

of long-running queries executing on the database and determine an impact of each of the different types of long-running queries upon a performance of the database.

15. The computer system of claim **11**, wherein the processor further executes the algorithm to keep track of: an estimated execution time that a query optimizer estimates for each query, the elapsed time for which each query has been running, and the execution time that each query has consumed.

* * * * *