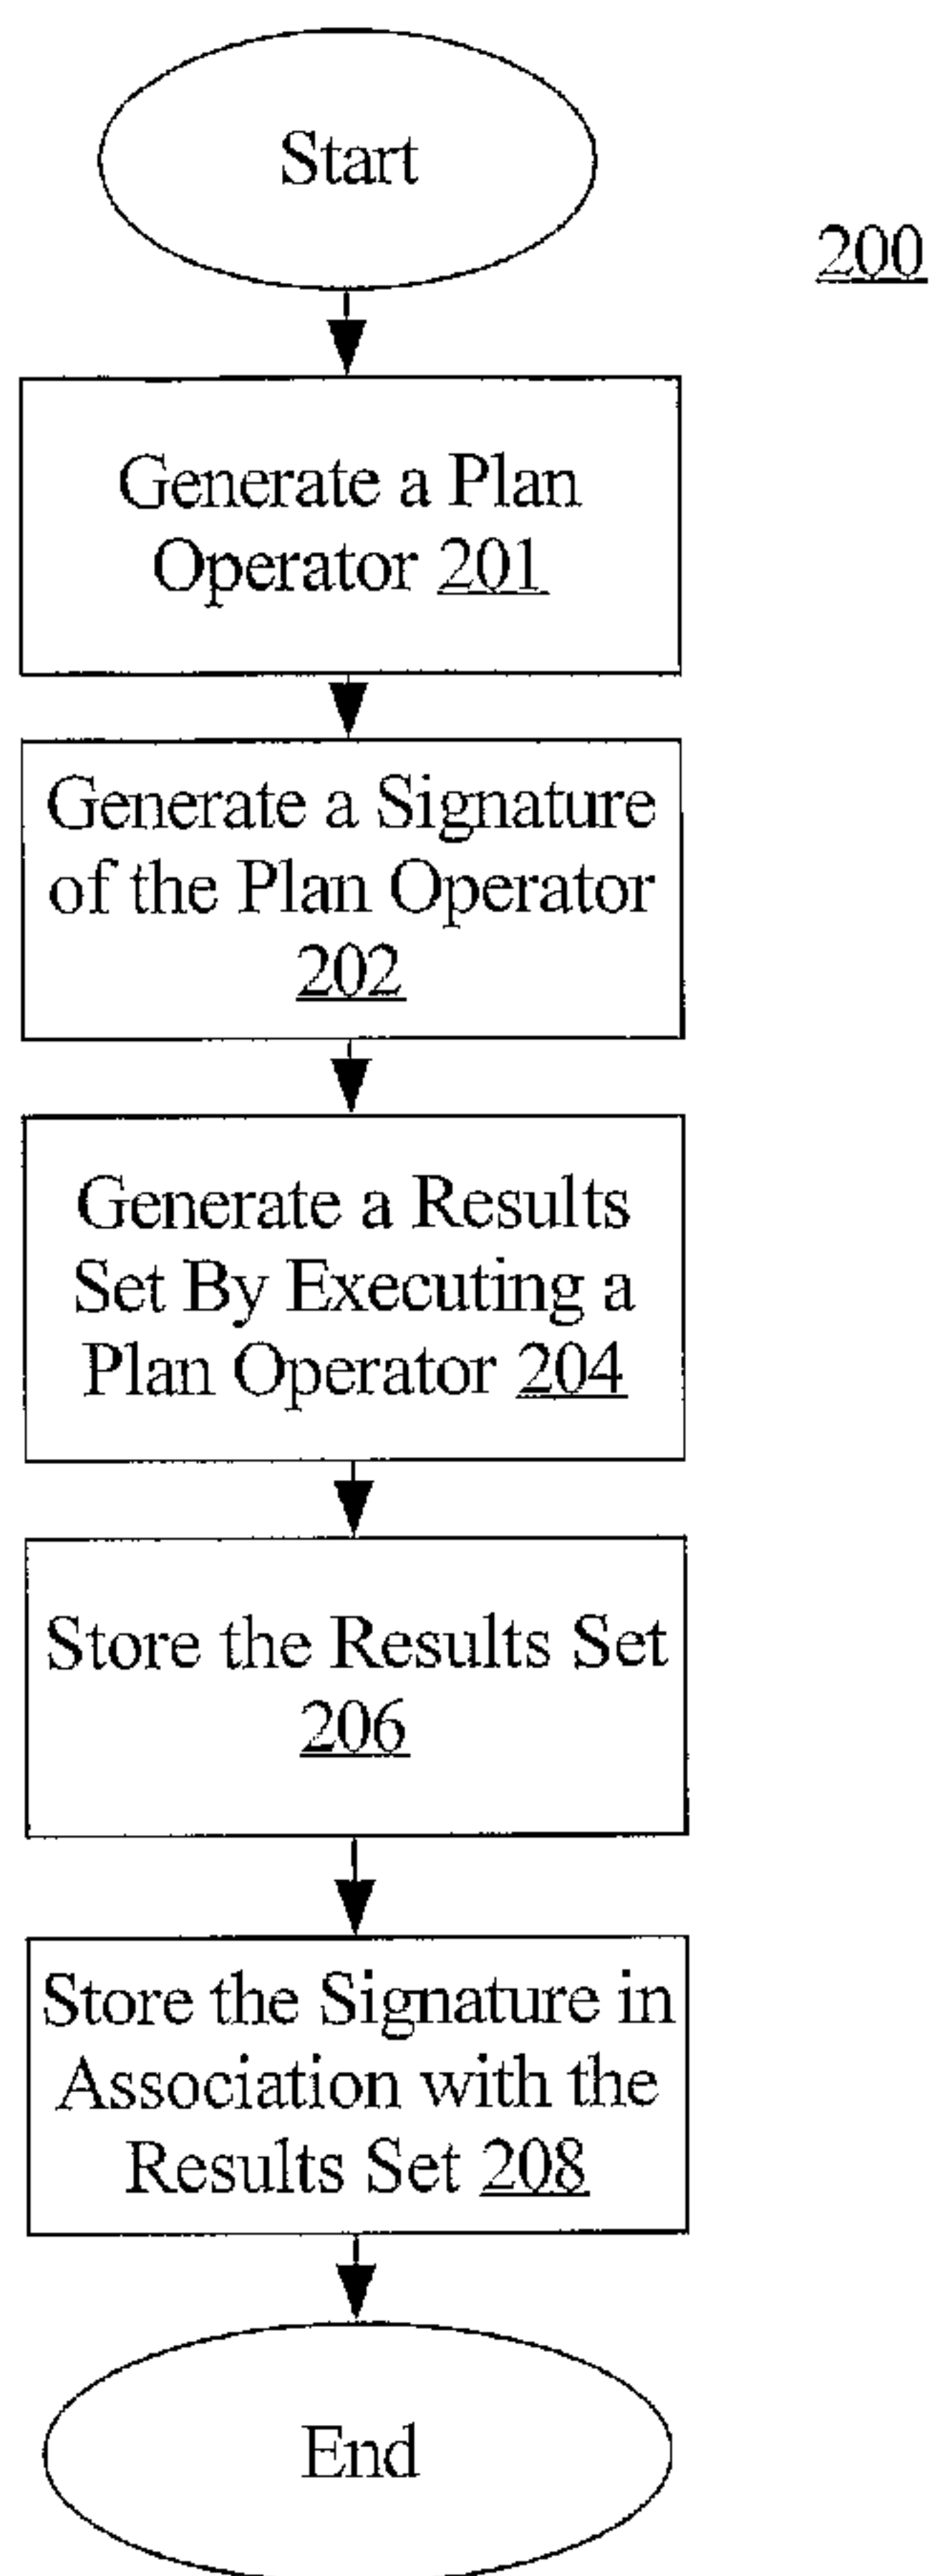




(86) Date de dépôt PCT/PCT Filing Date: 2005/04/20
 (87) Date publication PCT/PCT Publication Date: 2005/11/10
 (85) Entrée phase nationale/National Entry: 2006/10/06
 (86) N° demande PCT/PCT Application No.: US 2005/013648
 (87) N° publication PCT/PCT Publication No.: 2005/106717
 (30) Priorité/Priority: 2004/04/22 (US10/830,510)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01)
 (71) Demandeur/Applicant:
 ORACLE INTERNATIONAL CORPORATION, US
 (72) Inventeur/Inventor:
 OZBUTUN, CETIN, US
 (74) Agent: PAUL SMITH INTELLECTUAL PROPERTY LAW

(54) Titre : MISE EN ANTEMEMOIRE PARTIELLE DE REQUETE
 (54) Title: PARTIAL QUERY CACHING



(57) Abrégé/Abstract:

When a query is performed an execution plan is formed, which is made up of a plurality of plan operators. Before executing any given plan operator, a storage area is checked. If a results set has been stored corresponding to the plan operator, then the

(57) Abrégé(suite)/Abstract(continued):

cached results set is used instead of the executing the plan operator. If no results set exists that corresponds to the particular plan operator, the plan operator is executed, and the results set that is output may be stored for use by another plan operator of another execution plan.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
10 November 2005 (10.11.2005)

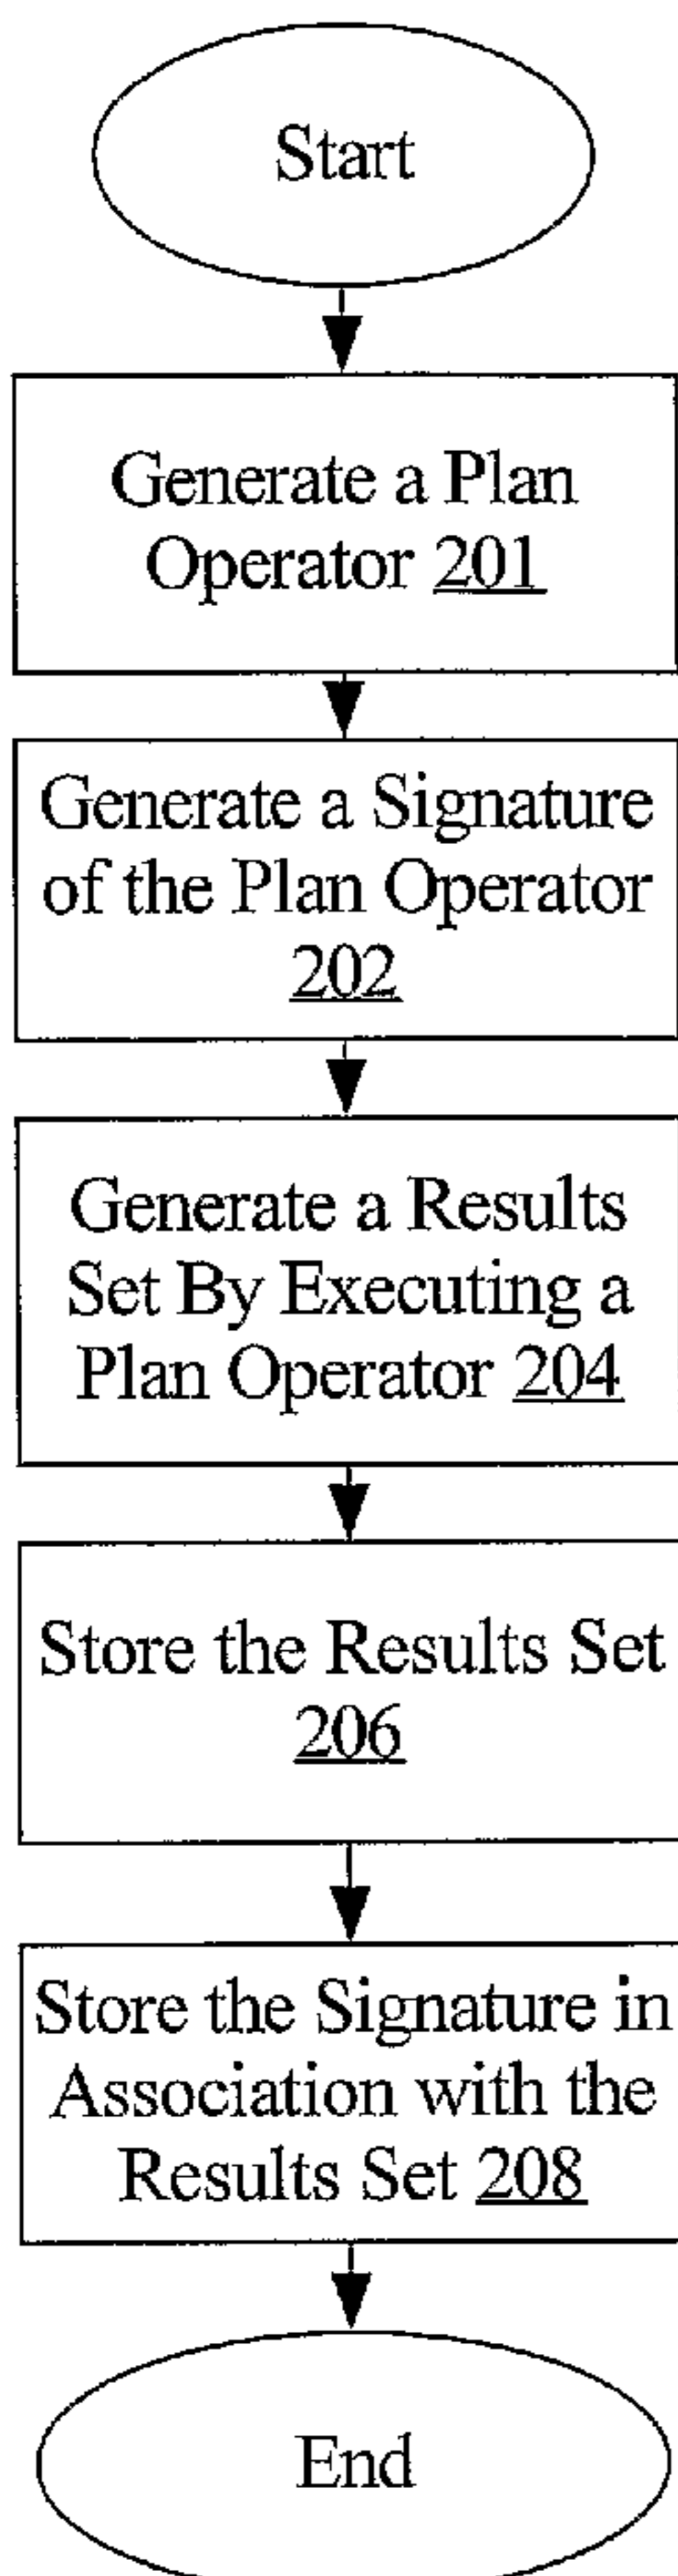
PCT

(10) International Publication Number
WO 2005/106717 A1

- (51) International Patent Classification⁷: **G06F 17/30**
- (21) International Application Number: PCT/US2005/013648
- (22) International Filing Date: 20 April 2005 (20.04.2005)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
10/830,510 22 April 2004 (22.04.2004) US
- (71) Applicant (for all designated States except US): **ORACLE INTERNATIONAL CORPORATION** [US/US]; 500 Oracle Parkway, Redwood City, CA 94065 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **OZBUTUN, Cetin** [US/US]; 360 Elm Street, San Carlos, California 94070 (US).
- (74) Agents: **HICKMAN, Brian, D.** et al.; HICKMAN PALERMO TRUONG & BECKER LLP, 2055 Gateway Place, Suite 550, San Jose, CA 95110-1089 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

[Continued on next page]

(54) Title: PARTIAL QUERY CACHING



200

(57) Abstract: When a query is performed an execution plan is formed, which is made up of a plurality of plan operators. Before executing any given plan operator, a storage area is checked. If a results set has been stored corresponding to the plan operator, then the cached results set is used instead of the executing the plan operator. If no results set exists that corresponds to the particular plan operator, the plan operator is executed, and the results set that is output may be stored for use by another plan operator of another execution plan.

WO 2005/106717 A1

WO 2005/106717 A1



FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO,
SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN,
GQ, GW, ML, MR, NE, SN, TD, TG).

— *before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments*

Published:

— *with international search report*

*For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.*

PARTIAL QUERY CACHING

FIELD OF THE INVENTION

[0001] The present invention relates to processing database commands, and in particular to caching results sets of queries.

BACKGROUND OF THE INVENTION

[0002] The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section. Similarly, the problems identified with an approach should not be assumed to have been recognized in the prior art, unless otherwise indicated.

[0003] When a database server receives a statement representing a request to perform a database operation, the database server may generate an execution plan. An execution plan defines the steps and operations performed by a database server to carry out a request. The statement used to request database operations may be in the form of a database command that conforms to a database language, such as Structured Query Language (SQL).

[0004] It is desirable to return results sets of queries quickly. As more and more people use databases to get information they are looking for, there is increased likelihood that some of the queries issued against a database are the same. In some systems, results sets of at least certain queries are stored. When a query is received a text based search is performed to find the previously executed query. If it is found that the same identical query (having the same text) is being presented for a second time, and if the results set for that query was stored, the system returns the cached results set rather than performing the query a second time. The more people that use the database, the greater the performance enhancement gained by saving the results sets of queries.

[0005] However, many queries are not identical, but similar. Many queries may process subsets of identical data within their results sets. Nonetheless, because the queries are not identical, a cached results set that is likely to have identical subsets of data as the expected results set cannot be reused, despite the similarity of the results sets, and each of the similar queries are executed.

[0006] Based on the forgoing, there is still a need to store and reuse results sets of queries.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0008] FIG. 1 is an execution plan diagram.

[0009] FIG. 2 is a flowchart showing a method of executing an plan operator, according to an embodiment of the invention.

[0010] FIG. 3 is a flow chart showing a method of storing a results set, according to an embodiment of the invention.

[0011] FIG. 4 is a block diagram that illustrates a computer system upon which an embodiment of the present invention may be implemented.

DETAILED DESCRIPTION OF THE INVENTION

[0012] A method and apparatus for partial query caching is described. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

[0013] Several features are described hereafter that can each be used independently of one another or with any combination of the other features. However, any individual feature may not address any of the problems discussed above or may only address one of the problems discussed above. Some of the problems discussed above may not be fully addressed by any of the features described herein. Although headings are provided, information related to a particular heading, but not found in the section having that heading, may also be found elsewhere in the specification.

[0014] Described herein are approaches that use at least a portion of the results previously generated for a query for a subsequent computation of another query. A results set generated for an operation defined by the execution plan of a query is used as the results set for an identical or similar operation of an execution of another query.

EXECUTION PLANS

[0015] Database servers generate execution plans that are optimized for efficiency. When determining what steps to include in an execution plan, and the order in which the steps are performed, a database server accounts for many factors that affect efficiency. For example, a query may request rows that contain a first value in a first column and a second value in a second column. If the first column is indexed, then a database server may generate an execution plan that uses the index to access data more efficiently. Such an execution plan may include the following steps, executed in the order indicated below.

1. Access the index and generate row IDs of rows that contain the first value in the first column.
2. Obtain the set of rows identified by the row IDs.
3. Filter the set of rows for the second value in the second column.

[0016] A query execution plan comprises steps referred to as “plan operators”. A plan operator defines a set of operations that are performed in an execution plan and that may be executed by a database server to access one or more row sources. A row source may be data stored in a database (e.g., one or more rows in a table), or one or more output rows of another plan operator. Execution plans are explained in chapter 22 – 24 of *Oracle 8i Concepts*, by Oracle Corporation, February, 1999. The contents of *Oracle 8i Concepts* are incorporated herein by reference.

ILLUSTRATIVE EXECUTION PLAN DIAGRAM

[0017] An execution plan can be represented by an execution plan diagram. FIG. 1 is an execution plan diagram provided to illustrate various features of an example of an execution plan that includes partial query caching. In FIG. 1, execution plan 102 is generated for query 104 (by the database server). Query 104 is the query

```
SELECT t1.col4, t2.col2
FROM t1, t2
WHERE t1.col1=t2.col2 AND t1.col3=192
GROUP BY t1.col4;
```

[0018] Query 104 contains predicate 105, which is $t1.col1=t2.col2$ and $t1.col3=192$. A predicate is a logical expression that defines criteria that must be satisfied by the rows in the tables against which a query is applied.

[0019] In execution plan 102, plan operators are depicted using blocks. Specifically, group plan operator 106, join plan operator 108, scan table and filter plan operator 112, and scan table plan operator 116 are the plan operators that make up execution plan 102. Branches between a parent plan operator and a child plan operator represent a flow of rows from the child plan operator to the parent plan operator. Thus, the output rows of a child plan operator are input rows to a parent plan operator. Plan operators are described herein as performing the operations that the plan operators themselves define. However, this is simply a convenient way of expressing that a database server is performing the operations defined by the plan operators when the database server executes the plan operators.

[0020] In execution plan 102, scan table and filter plan operator 112 and scan table plan operator 116 perform a scan of column col1 of table t1 and of column col2 of table t2, respectively. However, scan table and filter plan operator 112 additionally filters the scanned rows of table t1 so that only rows having col3=192 are selected and thereby included in the output of scan table and filter plan operator 112. Join plan operator 108 joins the rows that are output from scan table and filter plan operator 112 and scan table plan operator 116 to form the rows that are output by join plan operator 108. Group plan operator 106 groups the rows created by join plan operator 108 into groups. As a result of group plan operator 106, each row in the output represents a different group (in other words, each t1.col4 in the output is unique).

FUNCTIONAL OVERVIEW

[0021] In partial query caching, the results sets of individual plan operators of query are computed and stored in a storage area for potential reuse for other queries. The storage area is an operator output storage area, which will be referred to as output storage. The results sets of the plan operator are referred to herein as cached results sets. Each cached results set will have a signature generated based on the definition of the results set's respective plan operator. Plan operators with identical or similar definitions, and hence potentially identical or similar results sets, will have the same or similar signatures. To compute (i.e., generate results for) a plan operator, a signature is generated for the plan operator, and is used to search a storage area for a cached results set with the same signature. If a cached results set is found, the cached results set's plan operator may be further examined to determine whether the cached results set may be used as the results set for the plan operator being computed. In an another embodiment, if a cached results set for a given signature is

found, the plan operator is not further examined, and the cached results set is used as the results set of the plan operator. If a cached results set is not found, then the plan operator is executed.

Additionally, the results set that was generated by executing the plan operator may be stored in the output storage for later use by another corresponding plan operator.

[0022] In an embodiment, first a search is performed for the entire query by searching for a results set of a plan operator that encompasses all of the plan operators of the execution plan of the query. If a results set was not stored for a plan operator corresponding to the entire query, then for one or more of the plan operators that make up the execution plan, the output storage is searched to see if cached results sets were stored for the corresponding plan operators.

[0023] In an embodiment, the plan operators are searched for, starting with those plan operators that encompass the largest portions of the search (e.g., a plan operator that encompasses the entire query), and proceeding to searching for the plan operators that encompass next largest portions of the search, which are encompassed within the plan operators of the largest portions of the query. In other words, if cached results set are not found for a parent plan operator, a search is performed for cached results sets for the plan operators that are the immediate children of the parent plan operator, if any exist.

EXECUTING A QUERY

[0024] FIG. 2 is a flowchart showing a method 200 of executing a plan operator, according to an embodiment of the invention. Method 200 shows steps performed when executing a plan operator. In step 201, a plan operator is generated for a query. In step 202, a signature of the plan operator is generated. In step 204, a results set is generated by executing the plan operator. In step 206, the results set, generated in step 204, is stored in the output storage (e.g., cached). The output storage is discussed further below. In step 208, the signature and optionally the plan operator are stored in a manner in which the signature and optionally the plan operator are associated with the cached results set.

[0025] FIG. 3 is a flow chart showing a method 300 of utilizing cached results sets, according to an embodiment of the invention. In step 302 a plan operator is generated for a query. In step 303, a signature is generated for plan operator that is being evaluated. In step 304, the output storage is searched for cached results set. In performing the search, a comparison may be performed between the signature of the plan operator that is being evaluated and the signatures

stored in association with the cached results sets. In an embodiment, an exact match may not be required. In optional step 306, if a match is found, a further evaluation is performed in which the stored plan operator is compared to the plan operator being evaluated to see if the cached results set may be used. In step 308, if step 306 is included in the method 300, then if in step 306 it is determined that the cached results set is usable, the cached results set is returned (either to the user as a results set or for use as input by another plan operator) in lieu of executing the plan operator. In step 308, if step 306 is not included in method 300, then if in step 304 a match is found (e.g., if the two signatures are found to be sufficiently similar), the cached results set is returned in lieu of executing the plan operator.

[0026] Referring to FIG. 1, in carrying out method 300, upon receiving query 104, the database server formulates query plan 102. Then the database server checks the plan operator output storage area to see if cached results sets were stored for the plan operators that make up execution plan 102. Plan operators that constitute larger portions of the query are searched for first, because if results sets corresponding to the plan operators that make up the larger portions are found, there is no need to search for the plan operators that make up the smaller portions encompassed by the larger portions.

[0027] For example, in execution plan 102, first a search is done for a cached results set corresponding to group plan operator 106. If a results set is found, the results set found is returned to the user as the results set of query 104 rather than executing group plan operator 106 a second time. Additionally, there is no need to search for results sets for the remaining plan operators, because the remaining plan operators of execution plan 102 are encompassed within group plan operator 106.

[0028] However, if a cached results set corresponding to group plan operator 106 is not found, then the database server searches for a cached results set corresponding to join plan operator 108. If a cached results set is found that corresponds to join plan operator 108, then that cached results set is returned for use as input to group plan operator 106, and there is no need to execute the remaining operators of execution plan 102. After executing group plan operator 106 on the cached results set corresponding to join plan operator 108, the results set of group plan operator 106 (based on the cached results set corresponding to join plan operator 108) is returned to the user as the results set of query 104. The results set of executing group plan operator 106 may be stored

in the output storage in case another query includes group plan operator 106 or a similar plan operator.

[0029] If a cached results set is not found that corresponds to join plan operator 108, then the database server searches for cached results sets corresponding to scan table and filter plan operator 112 and scan table plan operator 116. If a cached results set is found that corresponds to scan table and filter plan operator 112, the cached results set is returned for use by join plan operator 108. If a cached results set is not found for scan table and filter plan operator 112, then the database server executes scan table and filter plan operator 112 and returns the results set of executing scan table and filter plan operator 112 for use by join plan operator 108. Additionally, the results sets of scan table and filter plan operator 112 may be stored for later use. Similarly, if a cached results set are found for scan table plan operator 116, the cached results set is returned for use by join plan operator 108. If a cached results set is not found for scan table plan operator 116, then the database server executes scan table plan operator 116 and returns the results set of executing scan table plan operator 116 for use by join plan operator 108. Additionally, the results set of scan table plan operator 116 may be stored for later use. Join plan operator 108 is then executed on the results sets returned, whether the results sets returned originated from executing scan table and filter plan operator 112 and scan table plan operator 116, originated from finding the cached results sets corresponding to the scan table and filter plan operator 112 and scan table plan operator 116, or originated from finding cached results sets stored for one of, and performing the other of, scan table and filter plan operator 112 and scan table plan operator 116. Next group plan operator 106 is executed using the results set of join plan operator 108, and the results set of group plan operator 106 are returned to the user as the results set of query 104. As mentioned above, the results set of executing group plan operator 106 and/or join plan operator 108 may be stored in the plan operator output storage in case another query includes one of group plan operator 106 or join plan operator 108.

[0030] In an embodiment, the results set of a given operator may not necessarily be stored. Instead, a cost benefit analysis may be performed to determine whether to store the results set of the given operator or whether to not store the results set and have the operator executed a second time if another query includes the same operator. Some of the factors that may be used in determining whether to store the results set of the given operator are the time and resources necessary for executing the operator and the amount of storage area available. For example, if the time required

to execute the plan operator is less than the expected time required to search the output storage for the cached results set, the results set may not be stored. Similarly, if the storage area required to store the results set is greater than the available storage area in the output storage (either because the results set is so large or because the storage area is full or nearly, for example), the results set may also not be stored. In contrast, if the results set does not require much storage area or there is a large amount of storage area available, and the plan operator requires a great deal of time to execute, the results set may be stored. Other factors that may be used in determining whether to store results set of a newly executed plan operator, as will be apparent to those skilled in the art.

[0031] Two different queries may each be formulated by a different set of database statements, but may nonetheless have the same execution plan. For example, a first query may include a NATURAL join, while a second query may be the same as the first query except instead of using a NATURAL join the join key is explicitly named in a USING clause. By formulating an execution plan and searching for results sets of the execution plan, the results sets of one query may be returned as the results sets of another query even though the two queries are formed by different sets of database statements.

[0032] The output storage may be any sort of memory, such as those described below in the HARDWARE OVERVIEW. In an embodiment, the output storage is one or more storage areas in one or more temporary or short term storage devices (e.g., volatile storage devices). In an embodiment, the short term output storage may have relatively fast access time (e.g., flash memory or RAM) when compared to long term memory (e.g., a hard drive). In an embodiment, the output storage is one or more storage areas in one or more long term storage devices (e.g., non-volatile storage devices). In an embodiment, the output storage is a combination of one or more storage areas in one or more long term storage devices and short term storage devices.

SIGNATURES OF PLAN OPERATORS

[0033] In an embodiment, a signature is established for each plan operator based on the plan operator's definition. The signatures are generated such that plan operators with the same or similar definitions have the same signature. In an embodiment, the signature may be based on a data structure used to define or describe a plan operator. For example, the data structure associated with the plan operator may be used to generate a hash value, which may be mapped to numerical values that are used for forming a sequence numbers used to compose a signature.

[0034] The results set of a plan operator may be stored with the signature of the plan operator. When a new query is received, the results sets associated with previously executed plan operators are searched for by searching for a results set that is associated with a plan operator having the same signature. Using signatures, two queries may be written differently, but their respective execution plans may share plan operators with the same signature.

[0035] Use of the signatures allows the database server to quickly find results sets that may be used for a query. In an embodiment, each plan operator is mapped to signature that represents an exact description of that plan operator. In an alternative embodiment, each plan operator is mapped to a signature that is also used for a set and/or range of other plan operators, and the results set associated with this signature is used for each of the plan operators within the set and/or range of the signature. In this embodiment, the signature does not necessarily represent an exact description of the plan operator. In an embodiment, after finding a match of signatures the plan operators are compared in order to evaluate whether the results set from one plan operator are expected to be usable as the results set of the other plan operator. In another embodiment, if the two signatures are match or are sufficiently close the cached results set found is used without a further check of the plan operator that produced the cached results set.

PLAN OPERATORS THAT ARE SIMILAR BUT NOT IDENTICAL

[0036] Many times queries are similar but not identical. By searching for plan operators corresponding to queries that are not identical, the speed at which results sets may be returned for the query may be increased. Accordingly, in an embodiment, when searching for a results set of a plan operator, if a cached results set for a different plan operator is found that is similar, the results set of the similar operation may be used instead of executing the actual operation. For example, if a plan operator generated for a query involves finding a set of data associated with the dates November 1, 2000-January 1, 2004, and a cached results set is found for a plan operator that finds data associated with the dates November 1, 2000-February 23, 2004, the cached results set found may be used instead of executing the actual plan operator generated for the query. In one embodiment, the cached results set found is used "as is", without further modification. In another embodiment, the cached results set found is further filtered to remove some or all of the extra data in the cached results set. In another embodiment, the cached results set or the signature of the cached results set is evaluated to determine whether the user is likely to mind receiving the

difference between the results requested and cached results set found, which in this example is the extra data associated with the dates Jan 2, 2004 through February 23, 2004. If it is expected that the user will not mind the difference, the cached results set found is returned rather than performing the plan operator. If it is expected that the user will mind the difference, the plan operator is executed or the cached results set is filtered to remove the extra data, Jan 2, 2004 through February 23, 2004.

[0037] In an embodiment, two plan operators may be considered similar if it is expected that a typical user is unlikely to care which results set is returned. For example, in an embodiment, if the content of a cached results set is expected to be nearly the same differing by only a small percentage of rows as that of plan operator's, then the cached result may be used as the results set of plan operator. Whether similar but different, or dissimilar, cached results sets may be used in this way may be controlled by user input. For example, a user may include a command or a parameter in a query that specifies whether similar, but different, cached results sets may be returned, or a variable in user's database session may be set to specify whether similar, but different, cached results sets may be returned, via an Application Programmer Interface ("API"). Parameters or variables may also specify the extent to which two plan operators or signatures may be different, while still allowing the cached results set of one to be used for the results set of the other.

[0038] Also, in an embodiment, if a first plan operator is expected to return all of a results set of a second plan operator in addition to some extra data, and the extra data is expected not to be so voluminous as to obscure the content corresponding to the results set of the second plan operator, then the results set of the first plan operator may be accepted as the results set of the second plan operator as is or after being modified. Similarly, in an embodiment, if a first plan operator is expected to return a smaller results set than a second plan operator, but the missing data is expected not to be important or a small amount, then the results set of the first plan operator may be accepted as the results set of the second plan operator as is or after being modified.

STALE CACHED RESULTS SETS

[0039] Even though plan operators for given queries may be identical, the results set for the plan operators may be different as a result of the queries being executed at different times. In an embodiment, when a cached results set corresponding to a given plan operator is found, a check is performed to see if the cached results set found is stale. For example, in an embodiment, the time a

cached results set was computed is stored in association with cached results set in the plan operator output storage. If the age of a the cached results set is greater than a threshold age, the cached results set is no longer preserved in cache or used as a results set for the plan operator of another query. The threshold age may be based on user input as described above.

HARDWARE OVERVIEW

[0040] FIG. 4 is a block diagram that illustrates a computer system 400, is one example of a machine upon which an embodiment of the invention may be implemented. Specifically computer system 400 may be used as the database server and/or may be programmed to implement the partial query caching method described above. 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.

[0041] 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 416, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 404 and for controlling cursor movement on display 412. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0042] Computer system 400 may be used for storing results sets of plan operators. According to one embodiment of the invention, storage of results sets of plan operators and searching for cached results sets of previously executed plan operators in memory 112 is provided by computer system 400 in response to processor 404 executing one or more sequences of one or

more instructions contained in main memory 406. Such instructions may be read into main memory 406 from another computer-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. One or more processors in a multi-processing arrangement may also be employed to execute the sequences of instructions contained in main memory 406. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0043] The term “computer-readable medium” as used herein is one example of “machine-readable medium”, and refers to any medium that participates in providing instructions to processor 404 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 410. Volatile media includes dynamic memory, such as main memory 406. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 402. Transmission media can also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communications.

[0044] Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punch cards, paper tape, any other physical medium with patterns of holes, a RAM, a PROM, an EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0045] Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 404 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 400 can receive the data on the telephone line and use an infrared transmitter to convert the data to an infrared signal. An infrared detector coupled to bus 402 can receive the data carried in the infrared signal and place the data on bus 402. Bus 402 carries the data to main memory 406, from which processor 404 retrieves and executes the instructions. The instructions received by main memory 406 may optionally be stored on storage device 410 either before or after execution by processor 404.

[0046] Computer system 400 also includes a communication interface 418 coupled to bus 402. Communication interface 418 provides a two-way data communication coupling to a network link 420 that is connected to a local network 422. For example, communication interface 418 may be an Integrated Services Digital Network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 418 may be a Local Area Network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 418 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0047] Network link 420 typically provides data communication through one or more networks to other data devices. For example, network link 420 may provide a connection through local network 422 to a host computer 424 or to data equipment operated by an Internet Service Provider (ISP) 426. ISP 426 in turn provides data communication services through the worldwide packet data communication network now commonly referred to as the "Internet" 428. Local network 422 and Internet 428 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 420 and through communication interface 418, which carry the digital data to and from computer system 400, are exemplary forms of carrier waves transporting the information.

[0048] Computer system 400 can send messages and receive data, including program code, through the network(s), network link 420 and communication interface 418. In the Internet example, a server 420 might transmit a requested code for an application program through Internet 428, ISP 426, local network 422 and communication interface 418. In accordance with the invention, one such downloaded application provides for storage of results sets associated with previously executed plan operators in output storage, as described herein.

[0049] The received code may be executed by processor 404 as it is received, and/or stored in storage device 410, or other non-volatile storage for later execution. In this manner, computer system 400 may obtain application code in the form of a carrier wave.

[0050] In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this application, in the specific form in which

such claims issue, including any subsequent correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

CLAIMS

What is claimed is:

5

1. A method for computing a query at a database server comprising:

receiving a first query at the database server from a user;

generating a first plan for executing the first query through the database server,

wherein the plan includes a first plurality of plan operators that include at least

10

a first plan operator that outputs a first results set;

generating a first signature for the first plan operator;

generating the first results set by executing the first plan operator;

storing the first results set in association with the first signature;

after storing the first results set in association with the first signature:

15

receiving a second query from the user;

generating a second plan for executing the second query through the database

server, wherein the second execution plan includes a second plurality

of plan operators that include at least a second plan operator;

determining, based on the first signature, that at least a subset of the first

20

results set may be used as a second results set of said second operator;

and

computing a second results set for the second query based on the subset of the

first results set that was stored.

25

2. The method of claim 1, wherein the signature is generated based on operations performed by the first plan operator.

3. The method of claim 1, wherein the step of determining includes at least

comparing a signature of the second plan operator with stored signatures of a plurality

30

of plan operators to determine if each of the plurality of plan operators is the second plan operator.

4. The method of claim 1, wherein the first plan operator and the second plan operator are identical.
5. The method of claim 1, where the first plan operator and the second plan operator are not identical.
6. The method of claim 1, where the first plan operator and the second plan operator are similar.
- 10 7. The method of claim 1, wherein the first plan operator and the second plan operator are identical, but the first results set and the second results set are different.
- 15 8. The method of claim 1, wherein the first execution plan includes a third plan operator, the method further comprises: outputting a third results set in response to the third plan operator, wherein the third results set is not stored.
9. The method of claim 1, wherein the step of determining includes at least determining when the first results set was computed.
- 20 10. The method of claim 1, wherein a difference exist between the first plan operator and the second plan operator, and the step of computing includes at least modifying the first results set based on the difference.
- 25 11. The method of claim 1, wherein receiving a second query from the user includes receiving the second query from a different user.
12. The method of claim 1, further comprising computing the subset of the first results set by applying a filter to the first results set that was stored, wherein the subset is generated based on a condition contained in the second query.
- 30 13. A machine-readable medium carrying one or more sequences of instructions, which when executed by one or more processors, causes the one or more processors to perform a method comprising the steps of:

- receiving a first query at the database server from a user;
generating a first plan for executing the first query through the database server,
wherein the plan includes a first plurality of plan operators that include at least
a first plan operator that outputs a first results set;
- 5 generating a first signature for the first plan operator;
generating the first results set by executing the first plan operator;
storing the first results set in association with the first signature;
after storing the first results set in association with the first signature:
receiving a second query from the user;
- 10 generating a second plan for executing the second query through the database
server, wherein the second execution plan includes a second plurality
of plan operators that include at least a second plan operator;
determining, based on the first signature, that at least a subset of the first
results set may be used as a second results set of said second operator;
- 15 and
computing a second results set for the second query based on the subset of the
first results set that was stored.
14. The machine-readable medium of claim 13, wherein the signature is generated
20 based on operations performed by the first plan operator.
15. The machine-readable medium of claim 13, wherein the step of determining
includes at least comparing a signature of the second plan operator with stored
signatures of a plurality of plan operators to determine if each of the plurality of plan
25 operators is the second plan operator.
16. The machine-readable medium of claim 13, wherein the first plan operator and
the second plan operator are identical.
- 30 17. The machine-readable medium of claim 13, where the first plan operator and the
second plan operator are not identical.

18. The machine-readable medium of claim 13, where the first plan operator and the second plan operator are similar.

19. The machine-readable medium of claim 13, wherein the first plan operator and
5 the second plan operator are identical, but the first results set and the second results set are different.

20. The machine-readable medium of claim 13, wherein the first execution plan
10 includes a third plan operator, the method further comprises: outputting a third results set in response to the third plan operator, wherein the third results set is not stored.

21. The machine-readable medium of claim 13, wherein the step of determining includes at least determining when the first results set was computed.

15 22. The machine-readable medium of claim 13, wherein a difference exist between the first plan operator and the second plan operator, and the step of computing includes at least modifying the first results set based on the difference.

20 23. The machine-readable medium of claim 13, wherein receiving a second query from the user includes receiving the second query from a different user.

24. The machine-readable medium of claim 13, further comprising computing the subset of the first results set by applying a filter to the first results set that was stored, wherein the subset is generated based on a condition contained in the second query.

25

26-SEP-06 17:40 VON -WÄCHTERSCHÄUSER & HARTZ

+49-89-223759

T-927 P.11/14 F-526

Oracle International Corporation
50277-2742 / CX-13655PCT/US 2005/013648
June 26, 2006

(Clean Copy)

CLAIMS

1. A method for computing a query at a database server comprising:
receiving a first query at the database server from a user;
generating a first plan for executing the first query through the database server, wherein
the plan includes a first plurality of plan operators that include at least a first plan
operator that outputs a first results set;
generating a first signature for the first plan operator;
generating the first results set by executing the first plan operator;
storing the first results set in association with the first signature;
after storing the first results set in association with the first signature:
receiving a second query from the user;
generating a second plan for executing the second query through the database
server, wherein the second execution plan includes a second plurality of
plan operators that include at least a second plan operator;
determining, based on the first signature, that at least a subset of the first results
set may be used as a second results set of said second operator; and
computing a second results set for the second query based on the subset of the first
results set that was stored.
2. The method of claim 1, wherein the signature is generated based on operations
performed by the first plan operator.
3. The method of claim 1, wherein the step of determining includes at least comparing a
signature of the second plan operator with stored signatures of a plurality of plan
operators to determine if each of the plurality of plan operators is the second plan
operator.
4. The method of claim 1, wherein the first plan operator and the second plan operator
are identical.

26-SEP-06 17:40 VON -WÄCHTERSCHÄUSER & HARTZ
Oracle International Corporation
50277-2742 / CX-13655

+49-89-223759

T-927 P.12/14 F-526
PCT/JP 2005/013548
June 26, 2006

(Clean Copy)

5. The method of claim 1, where the first plan operator and the second plan operator are not identical.
6. The method of claim 1, where the first plan operator and the second plan operator are similar.
7. The method of claim 1, wherein the first plan operator and the second plan operator are identical, but the first results set and the second results set are different.
8. The method of claim 1, wherein the first execution plan includes a third plan operator, the method further comprises: outputting a third results set in response to the third plan operator, wherein the third results set is not stored.
9. The method of claim 1, wherein the step of determining includes at least determining when the first results set was computed.
10. The method of claim 1, wherein a difference exist between the first plan operator and the second plan operator, and the step of computing includes at least modifying the first results set based on the difference.
11. The method of claim 1, wherein receiving a second query from the user includes receiving the second query from a different user.
12. The method of claim 1, further comprising computing the subset of the first results set by applying a filter to the first results set that was stored, wherein the subset is generated based on a condition contained in the second query.
13. A machine-readable medium carrying one or more sequences of instructions, which when executed by one or more processors, causes the one or more processors to perform a method comprising the steps of:

26-SEP-06 17:41 VON -WÄCHTERSÄUSER & HARTZ

+49-89-223759

T-927 P.13/14 F-526

Oracle International Corporation
50277-2742 / CX-13655PCT/US 2005/013648
June 26, 2006

(Clean Copy)

receiving a first query at the database server from a user;
 generating a first plan for executing the first query through the database server, wherein
 the plan includes a first plurality of plan operators that include at least a first plan
 operator that outputs a first results set;
 generating a first signature for the first plan operator;
 generating the first results set by executing the first plan operator;
 storing the first results set in association with the first signature;
 after storing the first results set in association with the first signature:
 receiving a second query from the user;
 generating a second plan for executing the second query through the database
 server, wherein the second execution plan includes a second plurality of
 plan operators that include at least a second plan operator;
 determining, based on the first signature, that at least a subset of the first results
 set may be used as a second results set of said second operator; and
 computing a second results set for the second query based on the subset of the first
 results set that was stored.

14. The machine-readable medium of claim 13, wherein the signature is generated based
 on operations performed by the first plan operator.

15. The machine-readable medium of claim 13, wherein the step of determining includes
 at least comparing a signature of the second plan operator with stored signatures of a
 plurality of plan operators to determine if each of the plurality of plan operators is the
 second plan operator.

16. The machine-readable medium of claim 13, wherein the first plan operator and the
 second plan operator are identical.

17. The machine-readable medium of claim 13, where the first plan operator and the
 second plan operator are not identical.

26-SEP-06 17:41

VON WÄCHTERSHÄUSER & HARTZ

+49-89-223759

T-927 P.14/14 F-526

Oracle International Corporation
50277-2742 / CX-13655PCT/US 2005/013648
June 26, 2006

(Clean Copy)

18. The machine-readable medium of claim 13, where the first plan operator and the second plan operator are similar.
19. The machine-readable medium of claim 13, wherein the first plan operator and the second plan operator are identical, but the first results set and the second results set are different.
20. The machine-readable medium of claim 13, wherein the first execution plan includes a third plan operator, the method further comprises: outputting a third results set in response to the third plan operator, wherein the third results set is not stored.
21. The machine-readable medium of claim 13, wherein the step of determining includes at least determining when the first results set was computed.
22. The machine-readable medium of claim 13, wherein a difference exist between the first plan operator and the second plan operator, and the step of computing includes at least modifying the first results set based on the difference.
23. The machine-readable medium of claim 13, wherein receiving a second query from the user includes receiving the second query from a different user.
24. The machine-readable medium of claim 13, further comprising computing the subset of the first results set by applying a filter to the first results set that was stored, wherein the subset is generated based on a condition contained in the second query.

(Version with Correction Marks)

CLAIMS

1. (Amended) A method for ~~processing requests to~~ **computing a query at** a database server comprising:

receiving a first query at the database server from a user;

generating a first plan for executing the first query through the database server, wherein the plan includes a first plurality of plan operators that include at least a first plan operator that outputs a first results set;

generating a first signature for the first plan operator;

generating the first results set by executing the first plan operator;

storing the first results set **in association with the first signature;**

after storing the first results set in association with the first signature:

receiving a second query from the user;

generating a second plan for executing the second query through the database server, wherein the second execution plan includes a second plurality of plan operators that include at least a second plan operator ~~that corresponds to the first plan operator;~~

determining, based on the first signature, that at least a subset of the first results set may be used as a second results set of said second operator;

~~computing the subset of the first results set by applying a filter to the first results set that was stored, wherein the subset is generated based on a condition contained in the second query; and~~

computing a second results set for the second query based on the subset of the first results set that was stored.

~~2.~~ (Canceled).

~~2.~~ ~~3.~~ The method of claim ~~1~~ ~~2~~, wherein the signature is generated based on operations performed by the first plan operator.

(Version with Correction Marks)

3. 4. (Amended) The method of claim **1 ~~2~~**, wherein the step of determining includes at least comparing a signature of the second plan operator with stored signatures of a plurality of plan operators to determine if each of the plurality of plan operators is the second plan operator.

4. 5. The method of claim 1, wherein the first plan operator and the second plan operator are identical.

5. 6. The method of claim 1, where the first plan operator and the second plan operator are not identical.

6. 7. The method of claim 1, where the first plan operator and the second plan operator are similar.

7. 8. The method of claim 1, wherein the first plan operator and the second plan operator are identical, but the first results set and the second results set are different.

8. 9. The method of claim 1, wherein the first execution plan includes a third plan operator, the method further comprises: outputting a third results set in response to the third plan operator, wherein the third results set is not stored.

9. 10. The method of claim 1, wherein the step of determining includes at least determining when the first results set was computed.

10. 11. The method of claim 1, wherein a difference exist between the first plan operator and the second plan operator, and the step of computing includes at least modifying the first results set based on the difference.

(Version with Correction Marks)

~~11. 12.~~ The method of claim 1, wherein receiving a second query from the user includes receiving the second query from a different user.

~~13.~~ (Canceled)

12. (New) **The method of claim 1, further comprising computing the subset of the first results set by applying a filter to the first results set that was stored, wherein the subset is generated based on a condition contained in the second query.**

13. 14. (Amended) A machine-readable medium carrying one or more sequences of instructions, which when executed by one or more processors, causes the one or more processors to perform a method comprising the steps of:
receiving a first query at the database server from a user;
generating a first plan for executing the first query through the database server, wherein the plan includes a first plurality of plan operators that include at least a first plan operator that outputs a first results set;
generating a first signature for the first plan operator;
generating the first results set by executing the first plan operator;
storing the first results set **in association with the first signature;**
after storing the first results set in association with the first signature:
receiving a second query from the user;
generating a second plan for executing the second query through the database server, wherein the second execution plan includes a second plurality of plan operators that include at least a second plan operator ~~that corresponds to the first plan operator;~~
determining, based on the first signature, that at least a subset of the first results set may be used as a second results set of said second operator;

(Version with Correction Marks)

~~computing the subset of the first results set by applying a filter to the first results set that was stored, wherein the subset is generated based on a condition contained in the second query; and~~
computing a second results set for the second query based on the subset of the first results set that was stored.

~~15.~~ (Canceled)

14. ~~16.~~ (Amended) The machine-readable medium of claim 13 ~~15~~, wherein the signature is generated based on operations performed by the first plan operator.

15. ~~17.~~ The machine-readable medium of claim 13 ~~15~~, wherein the step of determining includes at least comparing a signature of the second plan operator with stored signatures of a plurality of plan operators to determine if each of the plurality of plan operators is the second plan operator.

16. ~~18.~~ The machine-readable medium of claim 13 ~~14~~, wherein the first plan operator and the second plan operator are identical.

17. ~~19.~~ The machine-readable medium of claim 13 ~~14~~, where the first plan operator and the second plan operator are not identical.

18. ~~20.~~ The machine-readable medium of claim 13 ~~14~~, where the first plan operator and the second plan operator are similar.

19. ~~21.~~ The machine-readable medium of claim 13 ~~14~~, wherein the first plan operator and the second plan operator are identical, but the first results set and the second results set are different.

(Version with Correction Marks)

~~20, 22.~~ The machine-readable medium of claim ~~13~~ 14, wherein the first execution plan includes a third plan operator, the method further comprises: outputting a third results set in response to the third plan operator, wherein the third results set is not stored.

~~21, 23.~~ The machine-readable medium of claim ~~13~~ 14, wherein the step of determining includes at least determining when the first results set was computed.

~~22, 24.~~ The machine-readable medium of claim ~~13~~ 14, wherein a difference exist between the first plan operator and the second plan operator, and the step of computing includes at least modifying the first results set based on the difference.

~~23, 25.~~ The machine-readable medium of claim ~~13~~ 14, wherein receiving a second query from the user includes receiving the second query from a different user.

~~26.~~ (Canceled)

24. (New) The machine-readable medium of claim 13, further comprising computing the subset of the first results set by applying a filter to the first results set that was stored, wherein the subset is generated based on a condition contained in the second query.

SELECT t1.col4, t2.col2
FROM t1, t2
WHERE t1.col1=t2.col2 AND t1.col3=192
GROUP BY t1.col4;

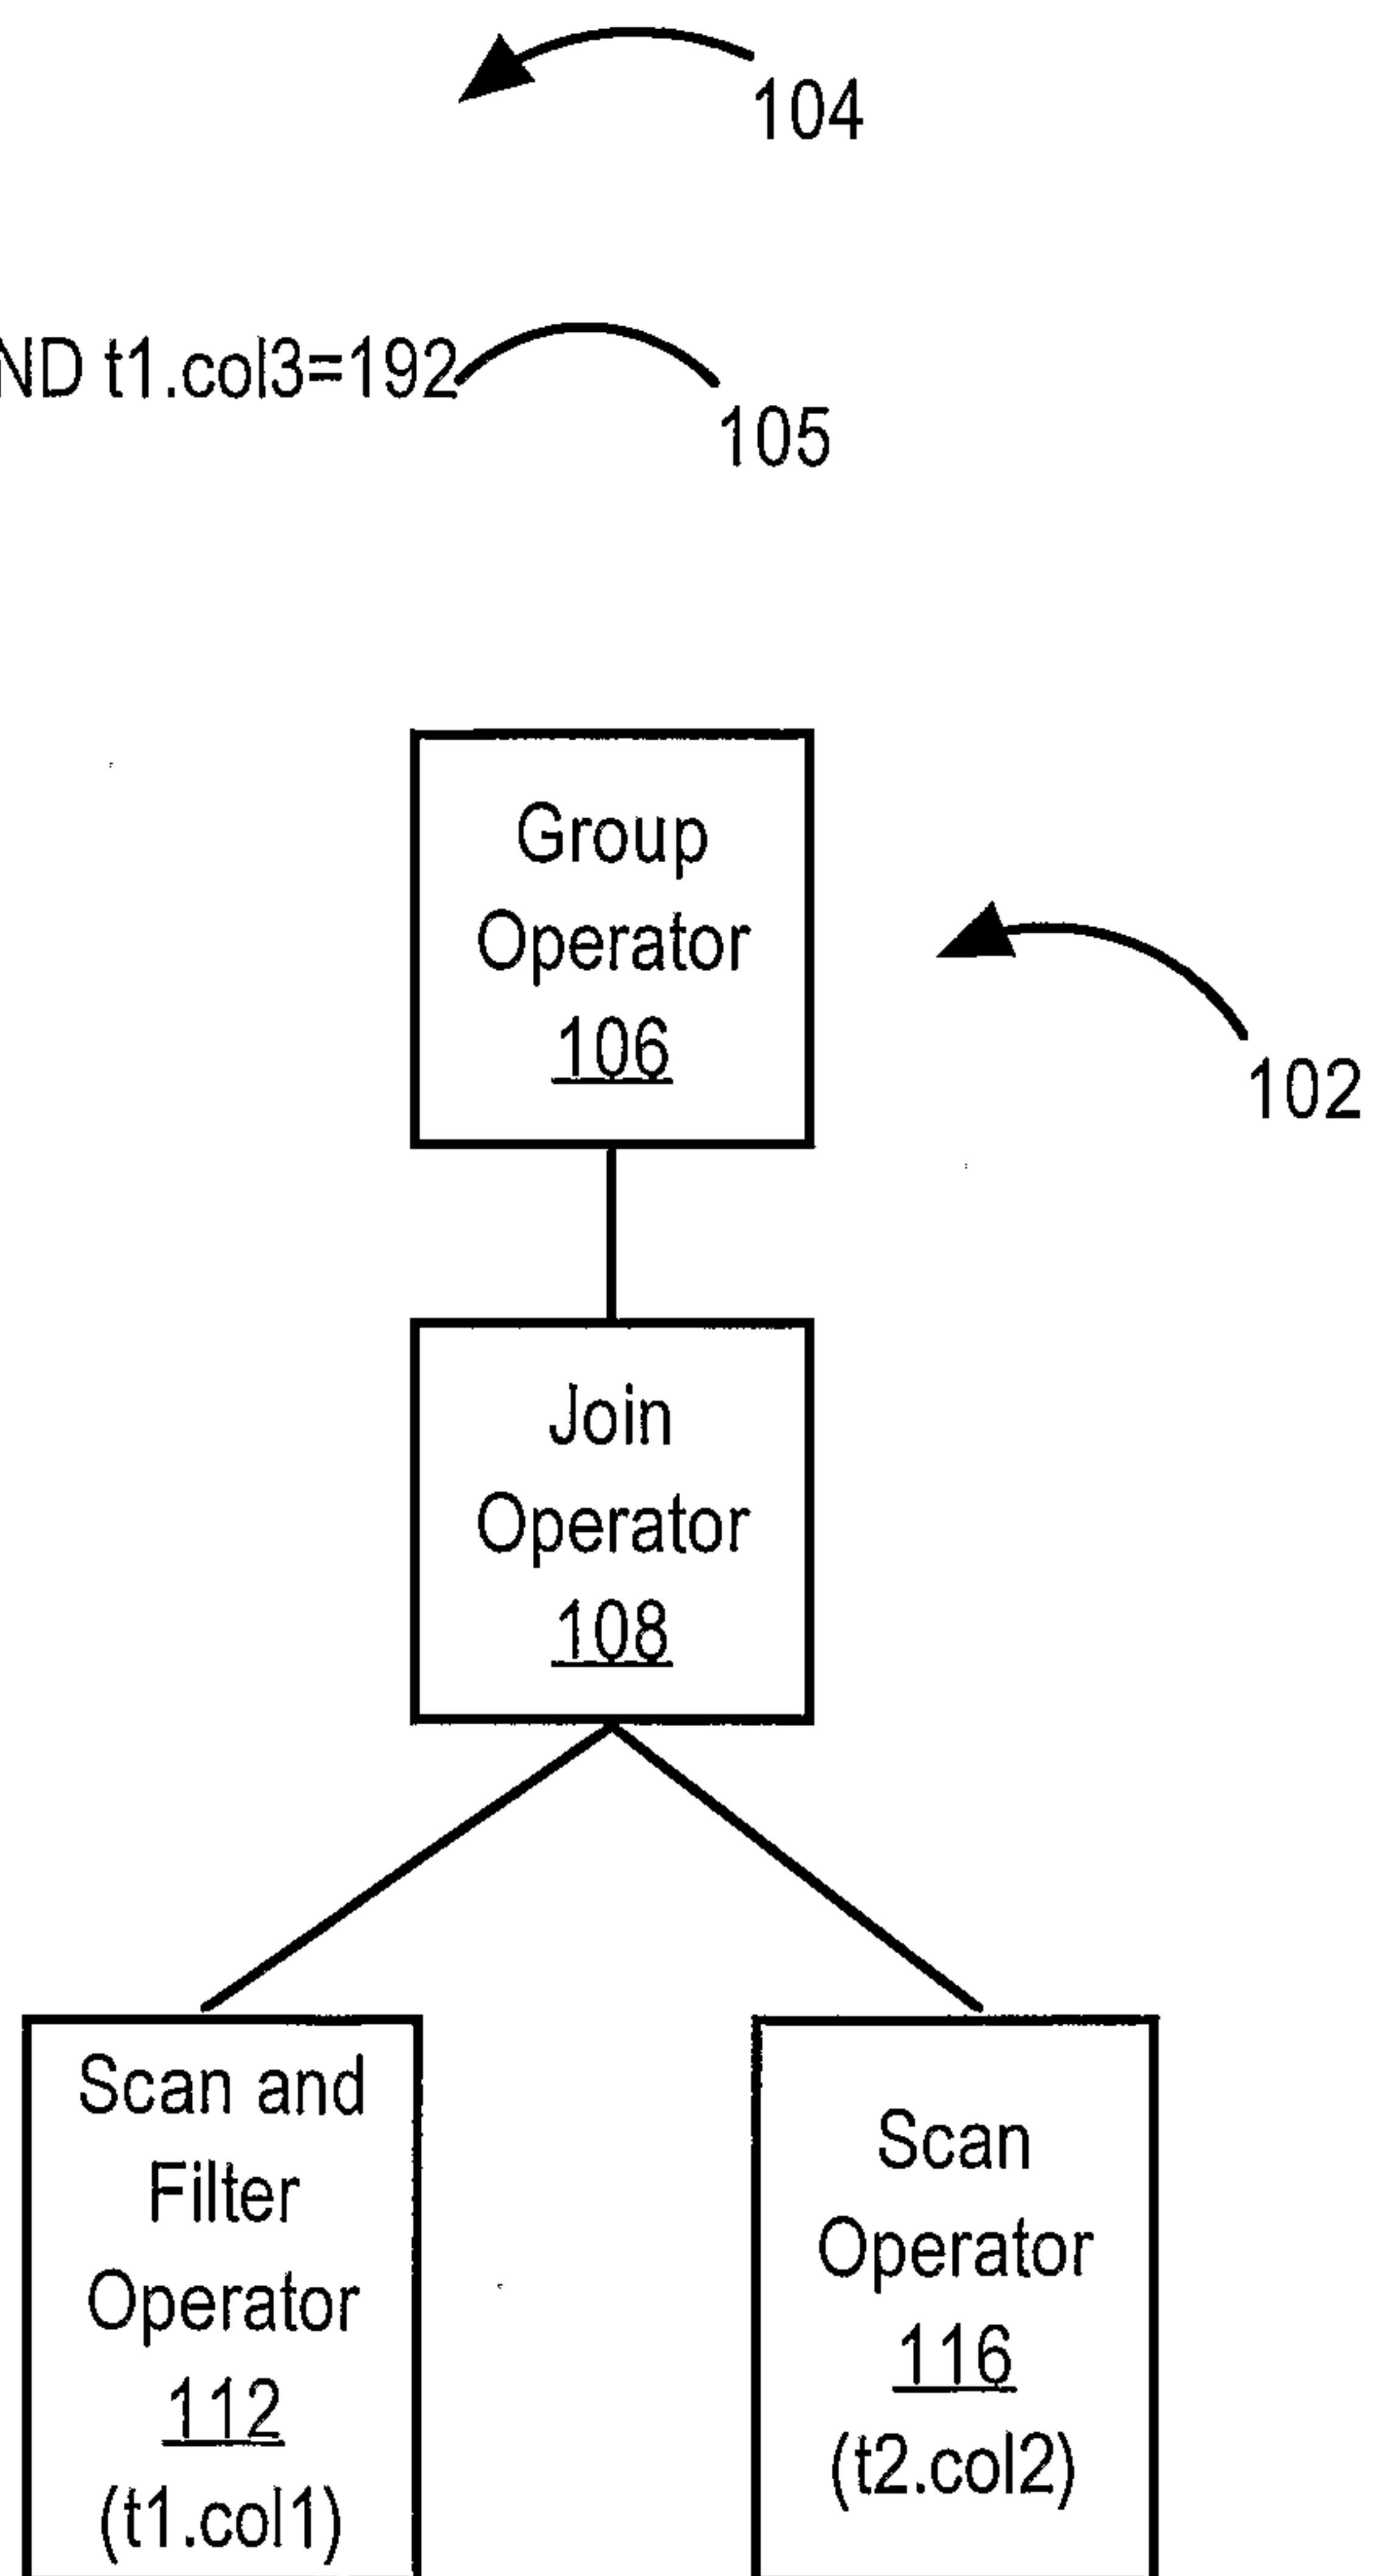
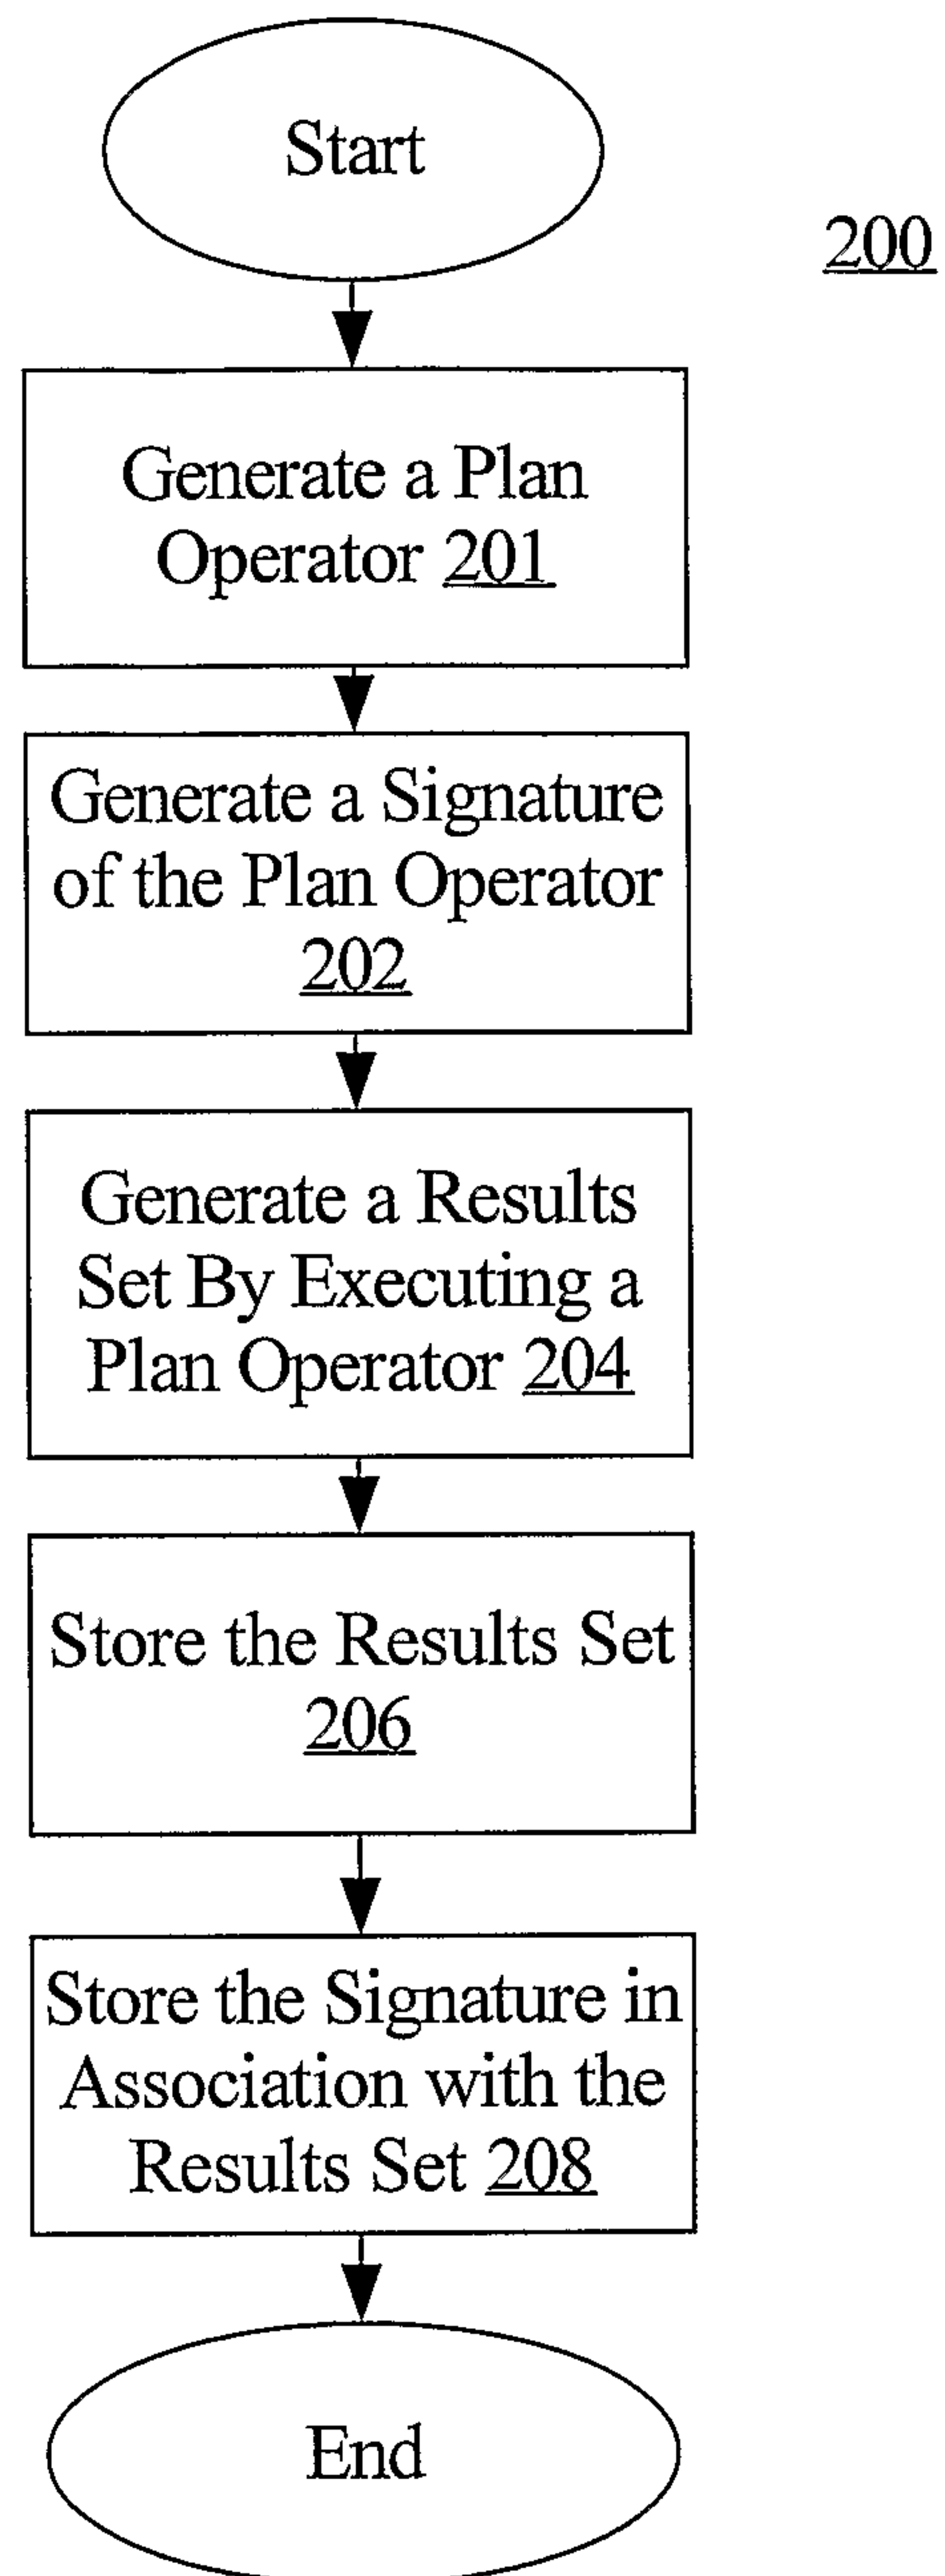
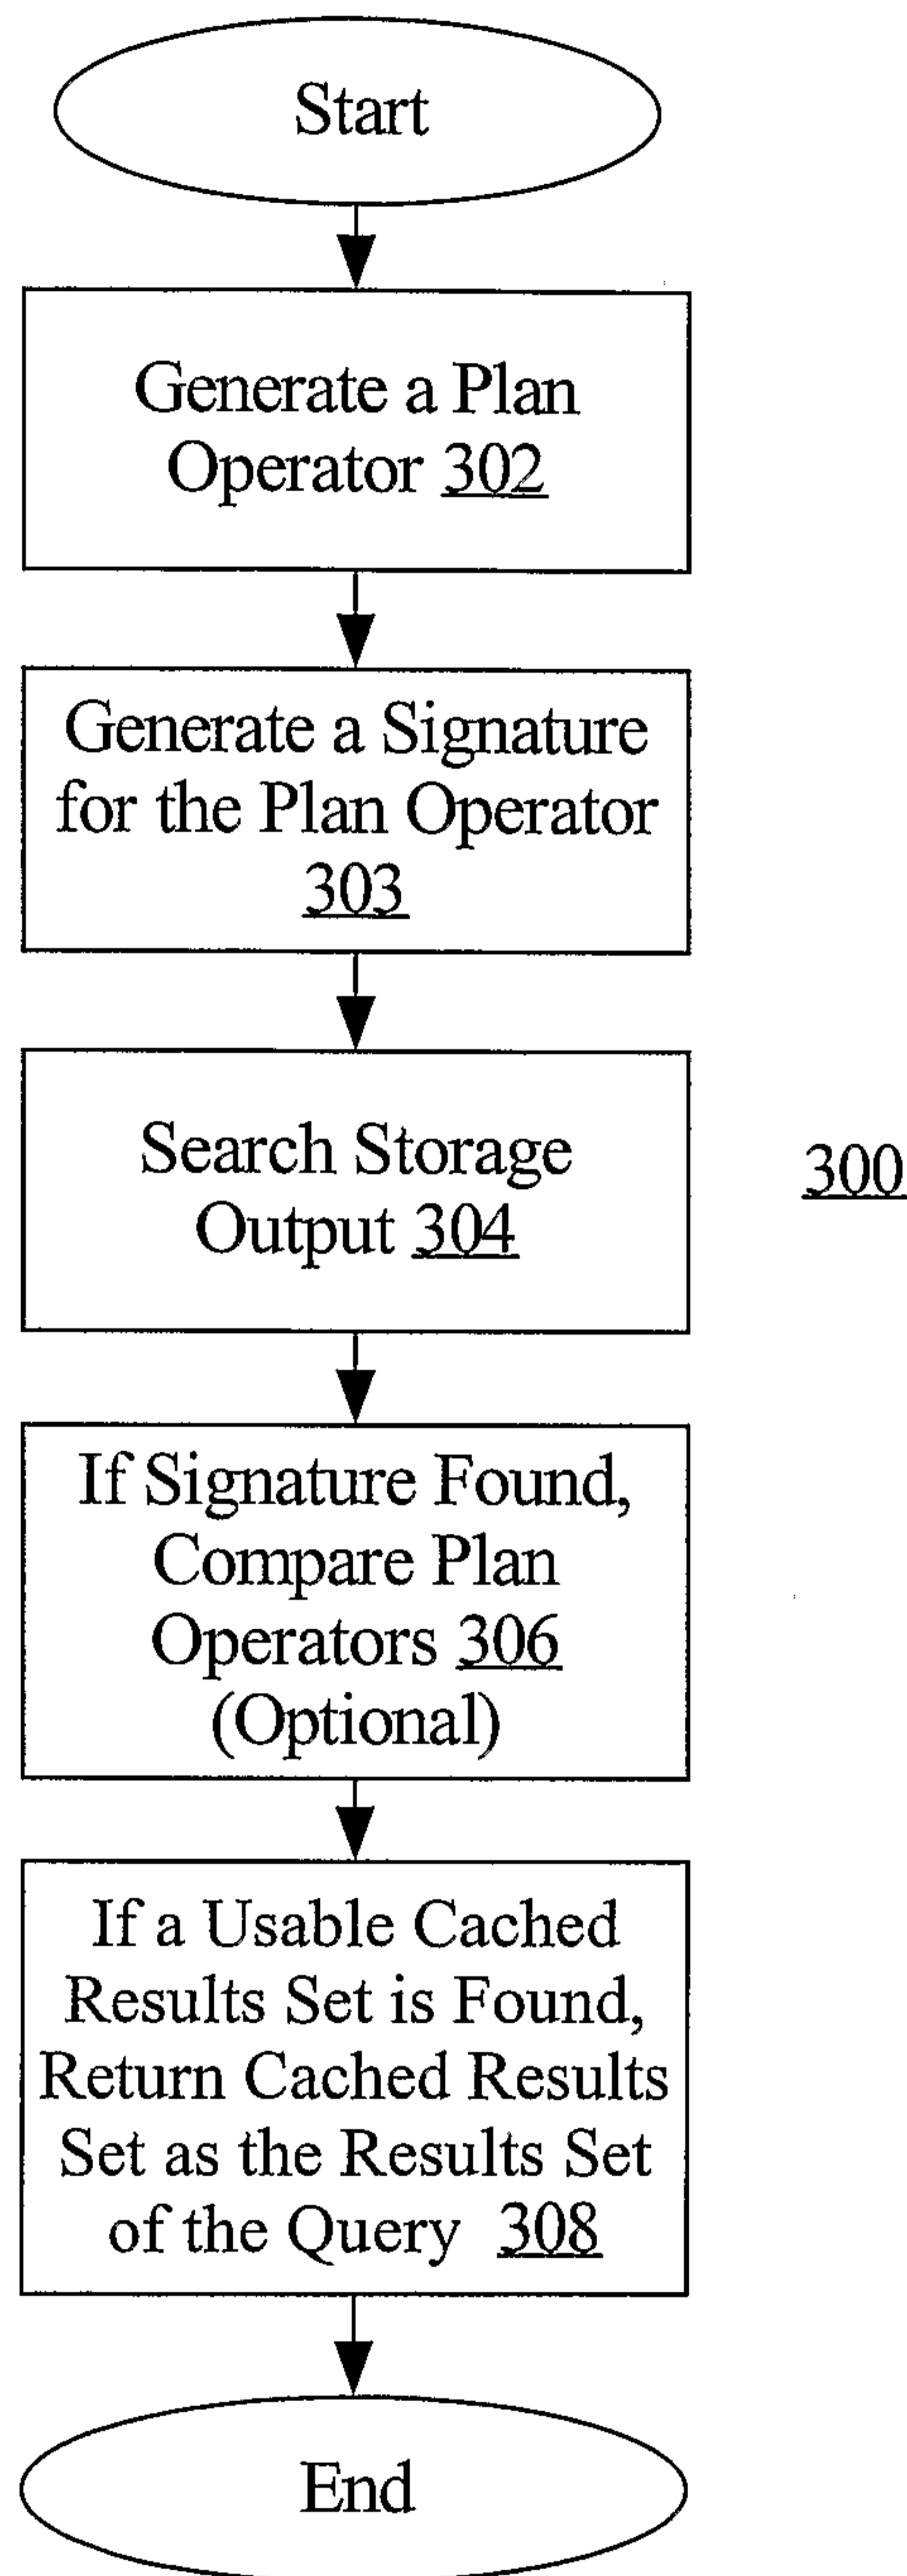


FIG. 1

2/4

**FIG. 2**

3/4

**FIG. 3**

4/4

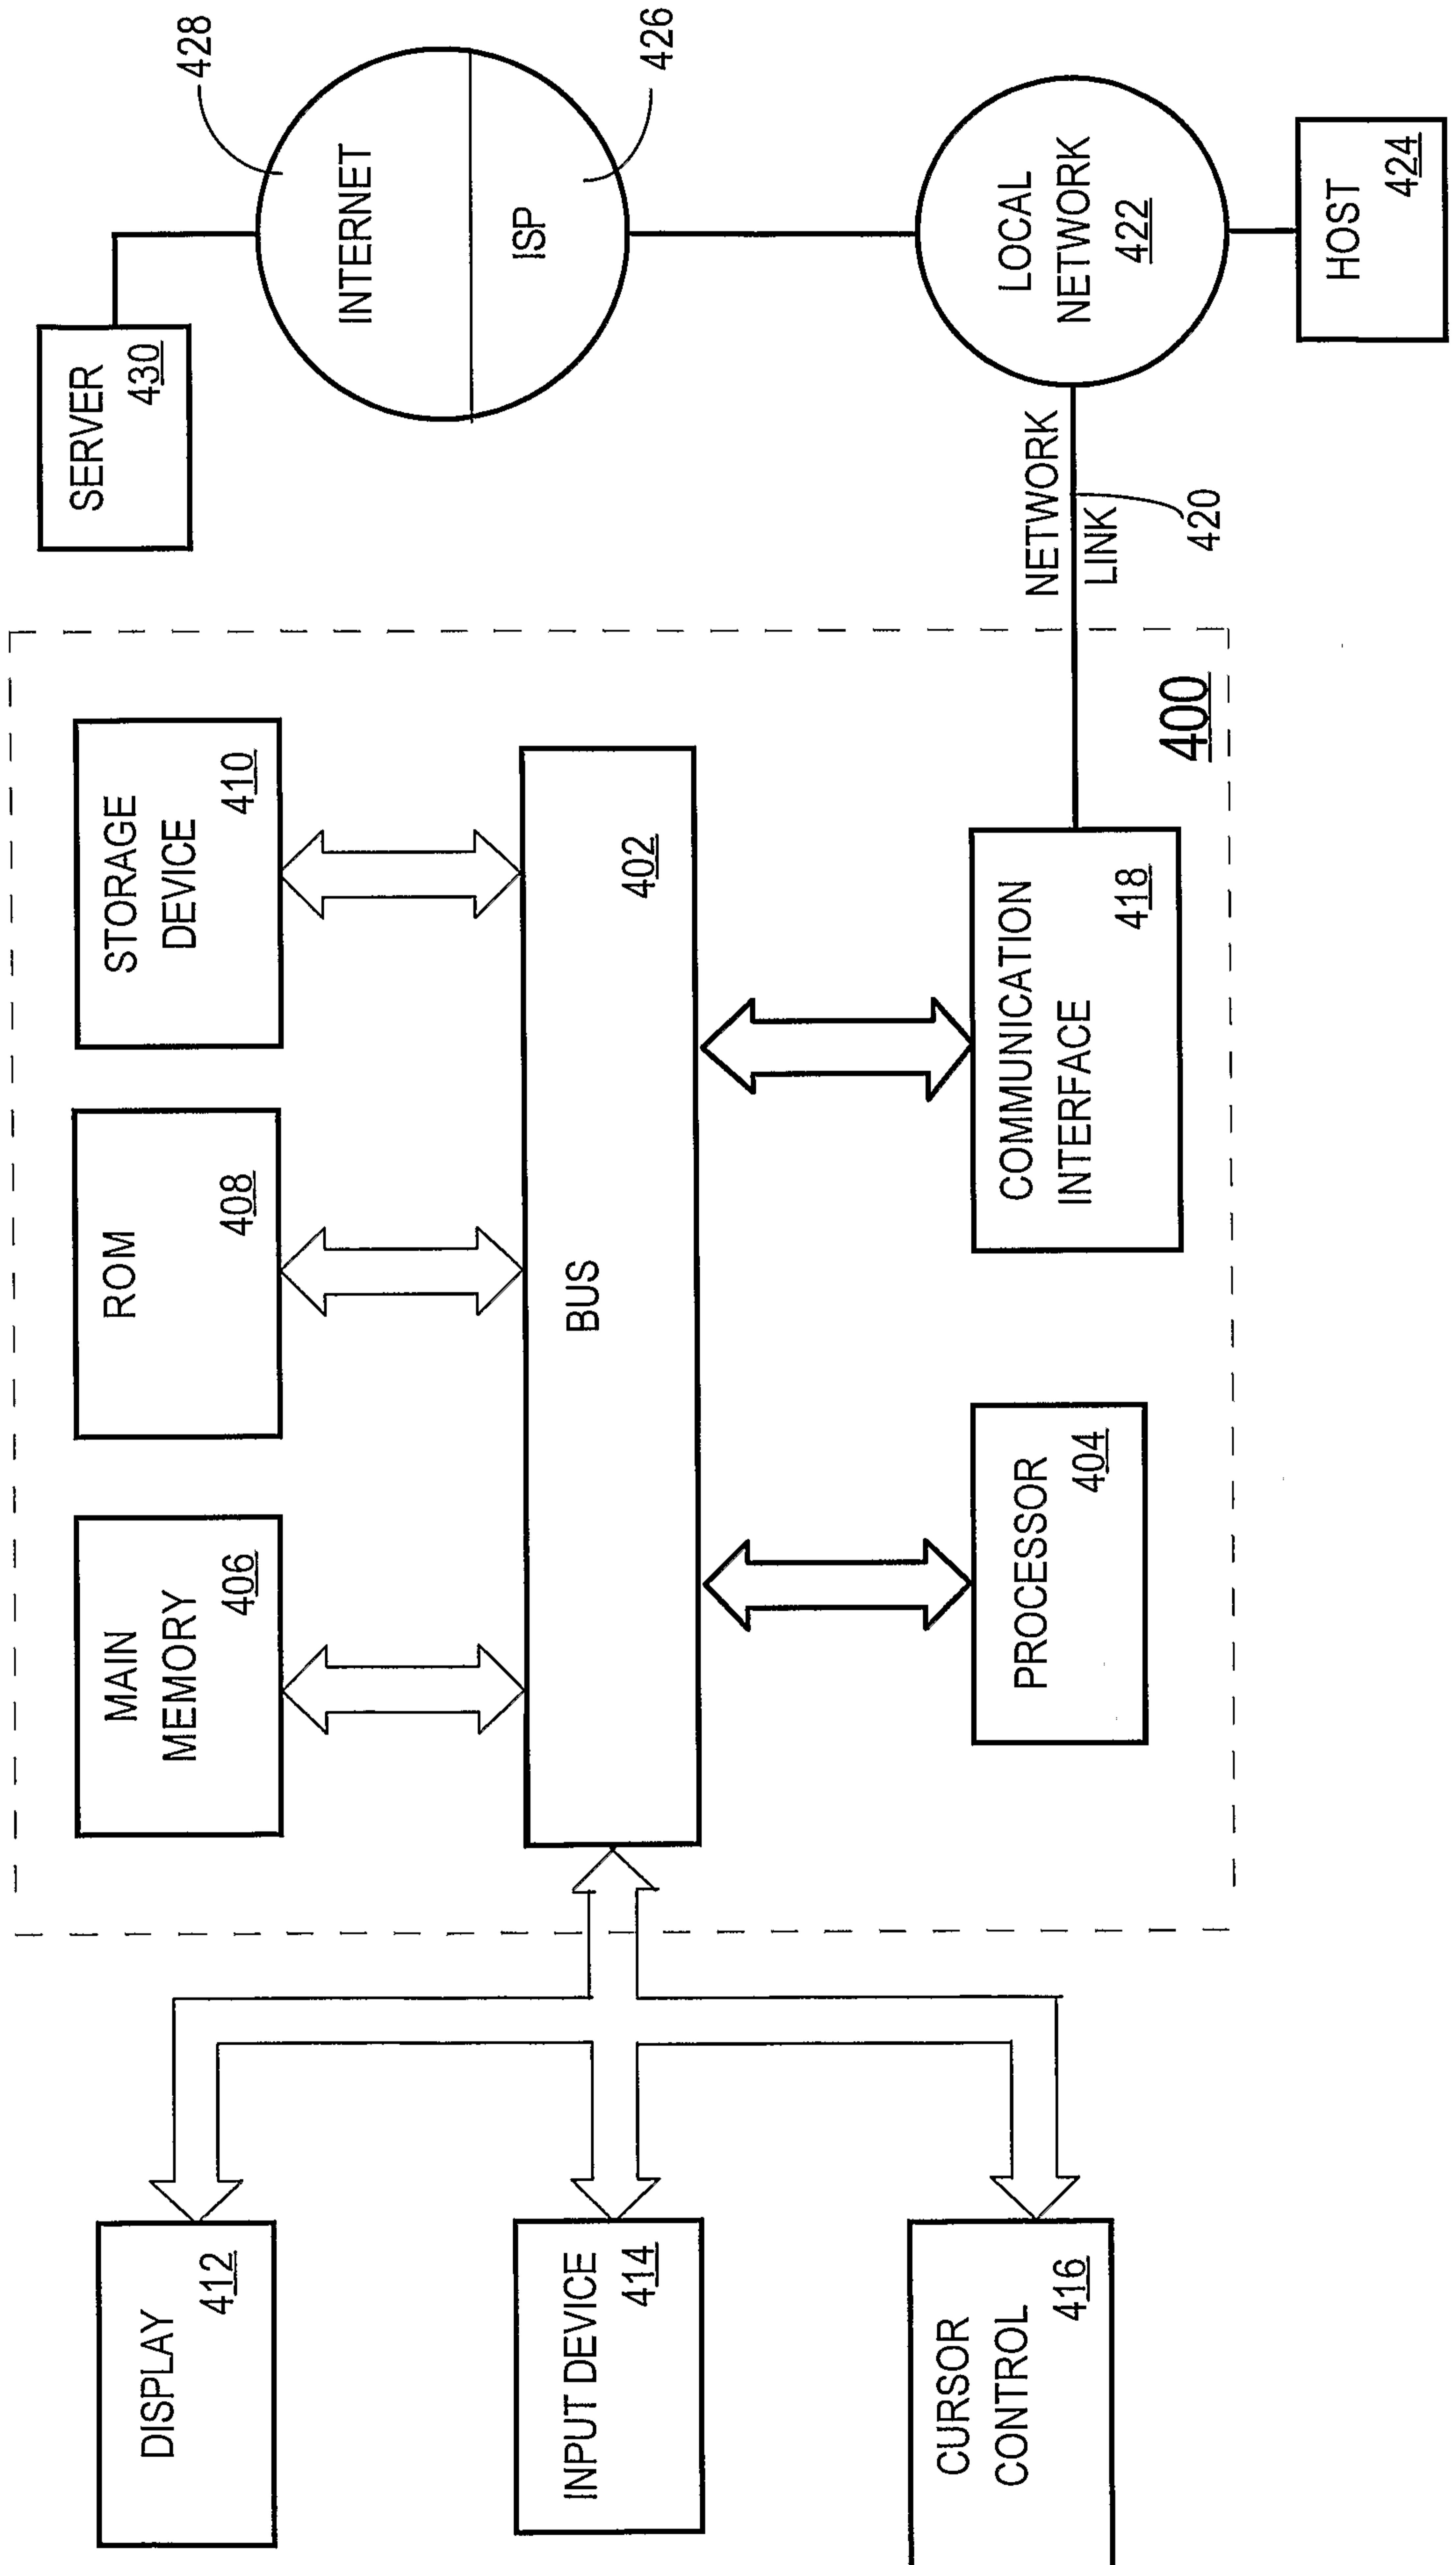


FIG. 4

