



- (51) **International Patent Classification:**
G06F 17/30 (2006.01)
- (21) **International Application Number:**
PCT/US2012/040505
- (22) **International Filing Date:**
1 June 2012 (01.06.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant (for all designated States except US):** **HEW-LETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **TRAN, Hoa Binh Nga** [US/US]; 8 Federal Street, Billerica, Massachusetts 01821 (US). **LAMB, Andrew Allinson** [US/US]; 8 Federal Street, Billerica, Massachusetts 01821 (US). **FULLER, Matthew Steven** [US/US]; 8 Federal Street, Billerica, Massachusetts 01821 (US).
- (74) **Agents:** **JAKOBSEN, Kraig A.** et al.; Hewlett-Packard Company, Intellectual Property Administration, 3404 E. Harmony Road, Mail Stop 35, Fort Collins, CO 80528 (US).

- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

[Continued on next page]

- (54) **Title:** MERGING DATA FROM A SOURCE LOCATION INTO A TARGET LOCATION




FIG. 3

- (57) **Abstract:** In an example, data from a source location is merged into a target location containing existing data, in which both the source location and the target location contain tuples of data. For each tuple in the source location, during a single operation, a determination is made as to whether there is a matched tuple in the target location that satisfies a predetermined condition. For each matched tuple that satisfies the predetermined condition, the matched tuple in the target location is updated with a count value that is equal to a count of the matched tuple in the source location and the target location. In addition, for each tuple that does not have a matched tuple that satisfies the predetermined condition, the unmatched tuple is inserted into the target location.





Published:

— *with international search report (Art. 21(3))*

MERGING DATA FROM A SOURCE LOCATION INTO A TARGET LOCATION

BACKGROUND

[0001] Loading new data into existing tables is an important process in most analytic databases. New data is typically loaded into existing tables to ensure that the data contained in the tables is up to date. The new data often includes both data that is new and data that is an update to existing data. Existing methods for loading the data typically employ two separate operations, one operation to load updated data and another operation to load new data.

-2-

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Features of the present disclosure are illustrated by way of example and not limited in the following figure(s), in which like numerals indicate like elements, in which:

[0003] FIG. 1 shows a diagram of a merge operation between a source location and a target location, according to an example of the present disclosure;

[0004] FIG. 2 shows a block diagram of a machine that may implement the merge operation of FIG. 1, according to an example of the present disclosure;

[0005] FIGS. 3-6, respectively, show flow diagrams of methods for merging data from a source location into a target location containing existing data, in which both the source location and the target location contain tuples of data, according to examples of the present disclosure;

[0006] FIGS. 7A-7D, 8A-8C, 9A-9C, and 10A-10G, respectively, depict diagrams of manners in which data may be redistributed among a plurality of projections, according to examples of the present disclosure; and

[0007] FIG. 11 illustrates a schematic representation of a computing device, which may be employed to perform various functions of the machine depicted in FIG. 2, according to an example of the present disclosure.

-3-

DETAILED DESCRIPTION

[0008] For simplicity and illustrative purposes, the present disclosure is described by referring mainly to an example thereof. In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present disclosure. It will be readily apparent however, that the present disclosure may be practiced without limitation to these specific details. In other instances, some methods and structures have not been described in detail so as not to unnecessarily obscure the present disclosure.

[0009] As used throughout the present disclosure, the term "includes" means includes but not limited to, the term "including" means including but not limited to. The term "based on" means based at least in part on. In addition, the terms "a" and "an" are intended to denote at least one of a particular element.

[0010] Disclosed herein is a method for merging data from a source location into a target location containing existing data, in which both the source location and the target location contain tuples of data. Also disclosed herein are an apparatus for implementing the method and a non-transitory computer readable medium on which is stored machine readable instructions that implement the method. According to an example, the method for merging data disclosed here comprises a Merge statement, which is implemented or invoked in a VerticaTM column-stored database.

[0011] As discussed in greater detail herein below, in the method, the merge operation of the present disclosure performs both an update to existing data and an insertion of new data in the target location during a single merge operation. As such, compared with conventional data loading operations, which require separate update and loading operations, the merge operation of the present disclosure provides a relatively more efficient manner of loading data containing both updates and new data into an existing location. Particularly, in conventional data loading operations, a first operation is required to update the existing tuples with changed tuples, and a second operation is required to insert new tuples. In the second operation, an anti-join operation in the insert query

-4-

(NOT IN) is required, thus causing the second operation to be very slow. In contrast, the merge operation of the present disclosure is relatively faster because its query plan does not include anti-join operations as discussed in greater detail herein.

[0012] With reference first to FIG. 1, there is shown a diagram 100 of a merge operation between a source location and a target location, according to an example of the present disclosure. It should be understood that the data depicted in FIG. 1 is for purposes of illustration only. In addition, the data depicted in FIG. 1 is used in various examples throughout the present disclosure to provide an understanding of the merge operation of the present disclosure. The merge operation of the present disclosure may comprise, for instance, a MERGE statement in an analytic database, such as Vertica™.

[0013] As shown in FIG. 1, the data contained in a source location 110 is to be merged into the data contained in a target location 120 to generate a merged location 130. According to an example, the source location 110, the target location 120, and the merged location 130 comprise separate tables or other arrangements stored in separate nodes, e.g., databases, memories, disks, servers, client devices, etc. Alternatively, the source location 110 and the target location 120 comprise data stored in the same node. In any regard, because the target location 120 typically contains a larger amount of data than the source location 110, the data contained in the source location 110 may be merged into the target location 120. In this regard, the merged location 130 may comprise an updated version of the data contained in the target location 120.

[0014] Although the source location 110 and the target location 120 have been depicted as having the same number of columns, it should be understood that various aspects of the present disclosure may be implemented in source locations and target locations having different numbers of columns and data types. In one regard, therefore, various aspects of the present disclosure may be implemented with source and target tables that are not the same schema.

-5-

[0015] The data contained in each of the source location 110 and the target location 120 is depicted as being arranged in tables formed of respective tuples, in which, each of the tuples includes a user identification (UserID), two-dimensional geographic coordinates (X, Y), a count, and a name of a business. In the example of FIG. 1, one of the tuples 112 in the source location 110 is depicted as matching one of the tuples 122 in the target location 120. In addition, the remaining two tuples 114 and 116 in the source location 110 are depicted as not matching any of the tuples 122-126 in the target location 120. In this regard, the tuple 122 in the target location 120 is to be updated and the tuples 114 and 116 in the source location 110 are to be inserted into the target location 120 during the merge operation.

[0016] As discussed in greater detail herein, both the updating and the inserting of the data from the source location 110 into the target location 120 are performed during a single merge operation. Particularly, the single merge operation of the present disclosure requires only a single scan to be performed on the source location 110 and a single scan to be performed on the target location 120. In addition, the single merge operation of the present disclosure does not require an anti-join operation. In contrast, conventional operations that seek to load new data into an existing table require that a first scan be performed to update modified data and that a second scan, which includes an anti-join operation, be performed to insert new data. As such, the merge operation of the present disclosure may generally be more efficient as compared with conventional loading operations.

[0017] Turning now to FIG. 2, there is shown a block diagram of a machine 200 that may implement the merge operation of FIG. 1, according to an example of the present disclosure. It should be understood that the machine 200 may include additional components and that one or more of the components described herein may be removed and/or modified without departing from a scope of the machine 200.

[0018] The machine 200 is depicted as including a processor 202, a data store 204, an input/output interface 206, and a data merging manager 210. The

-6-

machine 200 comprises any of, for instance, a server, a computer, a laptop computer, a tablet computer, a personal digital assistant, a cellular telephone, or other electronic apparatus that is to perform a method for merging data from a source location into a target location disclosed herein. The machine 200 may store the target location and/or may manage the storage of data in a target location stored in a separate machine, for instance, through a network device 208, which may comprise, for instance, a router, a switch, a hub, etc.

[0019] The data merging manager 210 is depicted as including an input/output module 212, a source location scanning module 214, a target location scanning module 216, a tuple matching module 218, an operator applying module 220, a filter applying module 222, a merging module 224, and a processing module 226. The processor 202, which may comprise a microprocessor, a micro-controller, an application specific integrated circuit (ASIC), or the like, is to perform various processing functions in the machine 200. One of the processing functions includes invoking or implementing the modules 212-226 of the data merging manager 210 as discussed in greater detail herein below.

[0020] According to an example, the data merging manager 210 comprises a hardware device, such as a circuit or multiple circuits arranged on a board. In this example, the modules 212-226 comprise circuit components or individual circuits. According to another example, the data merging manager 210 comprises a volatile or non-volatile memory, such as dynamic random access memory (DRAM), electrically erasable programmable read-only memory (EEPROM), magnetoresistive random access memory (MRAM), Memristor, flash memory, floppy disk, a compact disc read only memory (CD-ROM), a digital video disc read only memory (DVD-ROM), or other optical or magnetic media, and the like. In this example, the modules 212-226 comprise software modules stored in the data merging manager 210. According to a further example, the modules 212-226 comprise a combination of hardware and software modules.

-7-

[0021] The input/output interface 206 comprises a hardware and/or a software interface. In any regard, the input/output interface 206 may be connected to a network, such as the Internet, an intranet, etc., through the network device 208, over which the data merging manager 210 may receive and communicate information, for instance, the data contained in the source location 110 and data contained in other locations. The processor 202 may store information received through the input/output interface 206 in the data store 204 and may use the information in implementing the modules 212-226. The data store 204 comprises volatile and/or non-volatile memory, such as DRAM, EEPROM, MRAM, phase change RAM (PCRAM), Memristor, flash memory, and the like. In addition, or alternatively, the data store 204 comprises a device that is to read from and write to a removable media, such as a floppy disk, a CD-ROM, a DVD-ROM, or other optical or magnetic media.

[0022] Various manners in which the modules 212-226 may be implemented are discussed in greater detail with respect to the methods 300-600 depicted in FIGS. 3-6. FIGS. 3-6, respectively depict flow diagrams of methods 300-600 for merging data from a source location into a target location containing existing data, according to examples of the present disclosure. It should be apparent to those of ordinary skill in the art that the methods 300-600 represent generalized illustrations and that other operations may be added or existing operations may be removed, modified or rearranged without departing from the scopes of the methods 300-600. Although particular reference is made to the data merging manager 210 depicted in FIG. 2 as comprising an apparatus and/or a set of machine readable instructions that may perform the operations described in the methods 300-600, it should be understood that differently configured apparatuses and/or machine readable instructions may perform the methods 300-600 without departing from the scopes of the methods 300-600. The methods 400-600 are related to the method 300 in that the methods 400-600 include operations in addition to those included in the method 300. Thus, the methods 400-600 may be construed as including all of the features discussed with respect to the method 300.

-8-

[0023] Generally speaking, the methods 300-600 may separately be implemented to merge data from a source location into a target location containing existing data, in which both the source location and the target location contain tuples of data. In addition, the data from the source location may be merged with the data contained in the target location with during a single operation. In one regard, therefore, the methods 300-600 may be implemented to merge the data in a relatively more efficient manner than is possible with current data loading operations.

[0024] With reference first to FIG. 3, at block 302, for each tuple in the source location, a determination is made as to whether there is a matched tuple in the target location that satisfies a predetermined condition, for instance, by the tuple matching module 218. The determination as to whether there is a matched tuple in the target location for each tuple in the source location may be made during a scan of each of the source location and the target location. In addition, the predetermined condition may comprise that the data contained in selected ones of the columns match. Thus, for instance, in instances where the source location contains a different number of columns than the target location, tuples may be considered as matching when less than all of the data contained in the columns of the source location and the target location match, so long as the data contained in predetermined ones of the columns match each other.

[0025] As discussed above with respect to the diagram 100 depicted in FIG. 1, the predetermined condition comprises that the UserID and the X and Y of a tuple in the source location 110 be the same as the UserID and the X and Y of a tuple in the target location 120 for the tuple in the source location 110 to be determined as matching a tuple in the target location 120. In addition, the operation at block 302 results in a determination that the tuple 112 in the source location 110 matches the tuple 122 in the target location 120. The operation at block 302 also results in a determination that the tuples 114 and 116 in the source location 110 do not match any of the tuples in the target location 120.

[0026] At block 304, for each matched tuple that satisfies the predetermined condition, the matched tuple in the target location is updated with

-9-

a count value that is equal to the count of the matched tuple in the source location and the target location, for instance, by the merging module 224. In addition, at block 306, for each tuple in the source location that does not have a matched tuple in the target location that satisfies the predetermined condition, the unmatched tuple is inserted into the target location, for instance, by the merging module 224.

[0027] As discussed above with respect to the diagram 100 depicted in FIG. 1, in the tuple 132 in the merged location 130, the count value has been increased to "2", which is an example of a matched tuple existing in the target location being updated at block 304. In addition, tuples 138 and 140 have been inserted into the merged location 130, which is an example of unmatched tuples being inserted into the target location at block 306.

[0028] Turning now to FIG. 4, at block 402, the source location is scanned, for instance, by the source location scanning module 214. In addition, at block 404, the target location is scanned, for instance, by the target location scanning module 216. The source location and the target location are scanned to identify the tuples of data contained in each of the source location and the target location. By way of example, the outcome of the scanning of the source location and the target location may be a determination that the source location 110 contains the tuples 112-116 and that the target location 120 contains the tuples 122-126.

[0029] At block 406, the source location and the target location are right outer joined, for instance, by the tuple matching module 218. The outcome of the right outer join at block 406 is a determination as to which tuples from the source location are to be inserted into the target location and which tuples in the target location are to be updated. The output tuples of the right outer join operation include both source and target data for the matched tuples, and the source data for the unmatched tuples. According to an example, during implementation of the right outer join, a matching boolean (M) is added to each output tuple, for instance, as a new column (M), to mark whether the tuple is new (M=false) or existing (M=true).

-10-

[0030] In the diagram 100 depicted in FIG. 1, there is one matched tuple 112, 122 that includes data of the join attributes: `userId (1)`, `X (10.1)`, `Y (2.7)`; other data from the source table: `count(1)`, `name ('Restaurant A')`, and other data from the target table: `count(1)`, and `name('Restaurant A')`. The remaining two unmatched tuples 114, 116 include only data from the source location 110.

[0031] At block 408, the values of the matched tuples and the unmatched tuples are applied, for instance, by the operator applying module 220. Particularly, the operator applying module 220 may apply an APPLY operator on the matched tuples and the unmatched tuples to apply appropriate expressions for the matched and unmatched tuples. For example, the count of the matched tuple in the target location may be updated to be the computed value (`count =tgt.count + src.count`). In addition, an appropriate expression for adding the unmatched tuples may be applied.

[0032] At block 410, the values of the matched tuples and the unmatched tuples are filtered, for instance, by the filter applying module 222. Particularly, the filter applying module 222 may apply a filter that outputs the correct update and/or insert tuples depending upon whether the tuples are matched or unmatched. By way of example, the filter may perform a `doUpdate` when the tuples are matched and a `doInsert` when the tuples are not matched.

[0033] At block 412, the filtered tuples are inserted into the target location, for instance, by the merging module 224. Particularly, the updated tuples and the new tuples are inserted in the target location.

[0034] At block 414, the tuples in the source location that have been determined as having a matched tuple in the target location are filtered, for instance, by the filter applying module 222. Particularly, the filter applying module 222 filters the matched tuples to output the data of the updated (matched) tuples that are to be inserted into a `DeleteVector`. In addition, at block 416, the filtered tuples are inserted into the `DeleteVector`.

[0035] As described with respect to the method 400, when an update is performed, data is neither physically updated nor removed from the target location. Instead, the existing tuples are marked removed in the `DeleteVector`

-11-

and the updated tuples are not updated but newly inserted into the target location.

[0036] According to an example, a Sideways Information Passing (SIPS) operation is performed at blocks 402 and 404, for instance, by the source location scanning module 214 and the target location scanning module 216. The target location may be much larger than the source location because only a small part of the data may need to be updated and inserted. This means that the target location may contain a relatively large number of tuples that are unmatched with the tuples in the source location, and thus joining the unmatched tuples with the source location may be unnecessary and wasteful. In one regard, therefore, the SIPS operation may be performed to eliminate unmatched tuples from the outer input before the join operation at block 406. Particularly, the join first gets data from the inner input (source location), which is usually small, and sends their join attributes (e.g., from FIG. 1, the userID, X and Y) down to the outer input (target location). In the diagram 100 of FIG. 1, there are three such tuples 112-116. The outer input operator then eliminates data whose join attributes do not match with those tuples and sends the remaining tuples to the join. In the diagram 100 of FIG. 1, the tuples 124 and 126 would not be sent to the join. In this example, at block 406, the right outer join is implemented on the source location and the target location following implementation of the SIPS operation.

[0037] Turning now to FIG. 5, there is shown a flow diagram of a method 500 for merging data from a source location into a target location, according to another example. The method 500 differs from the method 400 depicted in FIG. 4 in that the method 500 pertains to merging data from a source location into a plurality of target projections. Particularly, the method 500 pertains to the merging of data from the source location into two target projections, target projection1 and target projection2. It should, however, be understood that the method 500 may pertain to any number of target projections, including one target projection. Although not shown, the source location may also be a projection.

-12-

[0038] A projection is a set of columns that are either from a table or a join of different tables. In an analytic database, column data is usually stored redundantly in various projections with different column sort orders or data segmentation. This storage mechanism ensures that queries still work when one or more nodes are down, and improves the performance of many different queries. Hence, in the method 500, the data in the source location is merged into target projection1 and target projection2.

[0039] At block 502, the source location is scanned as discussed above with respect to block 402 in FIG. 4. In addition, at block 504, target projection2 is scanned, for instance, in a manner similar to that discussed above with respect to block 404 in FIG. 4. The selection of which of target projection1 and target projection2 that is to be scanned at block 504 may be based upon any of a variety of factors. For instance, the selection of the target projection may be made by an optimization operation that uses any of heuristics and a complex custom-built cost model, based on compression aware I/O, CPU and Network transfer costs, etc. By way of particular example, target projection2 may differ from target projection1 in that target projection1 is not sorted on the query's join attributes, e.g., UserID, X, Y, whereas target projection2 is sorted on the query's join attributes. In this regard, target projection2 may be scanned at block 504 so that a fast merge join may be used to find matched and unmatched tuples between the source location and the target projection2.

[0040] At block 506, the source location and the target projection2 are right outer merge joined, for instance, by the tuple matching module 218, to determine which tuples from the source location are to be inserted into the target projection2 and which tuples in the target projection2 are to be updated. The right outer merge join implemented at block 506 is similar to the right outer join discussed above with respect to block 406 in FIG. 4.

[0041] According to an example, a SIPS operation is performed at blocks 502 and 504, for instance, by the source location scanning module 214 and the target location scanning module 216, prior to block 506, as also discussed above.

-13-

[0042] At block 508, the values of the matched tuples and the unmatched tuples are applied, for instance, by the operator applying module 220, to apply appropriate expressions for the matched and unmatched tuples. Block 508 is similar to block 408 in FIG. 4. In addition, at block 510, the values of the matched tuples and the unmatched tuples are filtered, for instance, by the filter applying module 222, as also discussed above with respect to block 410 in FIG. 4.

[0043] At block 512, the filtered tuples are inserted into target projection1, for instance, by the merging module 224. In addition, at block 514, the filtered tuples are inserted into target projection2. Particularly, the updated tuples and the new tuples are inserted into both target projection1 and target projection2.

[0044] At block 516, the tuples in the source location that have been determined as having a matched tuple in the target projection2 are filtered, for instance, by the filter applying module 222. Particularly, the filter applying module 222 filters the matched tuples to output the data of the updated (matched) tuples that are to be inserted into a DeleteVector. In addition, at block 518, the filtered tuples are inserted into the DeleteVector for target projection1 and at block 520, the filtered tuples are inserted into the DeleteVector for target projection2.

[0045] As described with respect to the method 500, when an update is performed, data is neither physically updated nor removed from either of the target projections. Instead, the existing tuples are marked removed in the DeleteVectors of the target projections and the updated tuples are not updated but newly inserted into the target projections.

[0046] With reference now to FIG. 6, there is shown a flow diagram of a method 600 for merging data from a source location into a target location, according to another example. The method 600 differs from the method 500 depicted in FIG. 5 in that the method 600 also pertains to a pre-join projection, which is a projection used to store the results of a join between a single large fact table, containing attributes of events, with one or more dimension tables, containing descriptive attributes of the events. Although two pre-join operations

-14-

are described with respect to the method 600, it should be clearly understood that any number of pre-join operations using any number of dimension tables may be implemented without departing from a scope of the method 600.

[0047] According to an example in which the single large fact table comprises the target location 120 depicted in FIG. 1, a first dimension table1 includes the names associated with the UserIds. Thus, for purposes of illustration, UserId1 may be associated with the name "John", UserId2 may be associated with the name "Sydney", and UserId3 may be associated with the name "Emily". In addition, a second dimension table2 includes the cities associated with the X and Y coordinates. Thus, for purposes of illustration, the coordinates (4.1, 7.7) may correspond to the city "Andover", the coordinates (5.1, 7.9) may correspond to the city "Cambridge", and the coordinates (10.1, 2.7) may correspond to the city "Boston."

[0048] Blocks 602-606 are similar to blocks 502-506, respectively. In addition, and according to an example, a SIPS operation is performed at blocks 602 and 604, for instance, by the source location scanning module 214 and the target location scanning module 216, prior to block 606, as also discussed above.

[0049] At block 608, a dimension table1 is scanned, for instance, by the target scanning module 216. In addition, at block 610, a join, e.g., a hash/merge join, is performed on the output of the right outer merge join performed at block 606 and the data contained in the dimension table1, for instance, by the processing module 226. The join at block 610 generally joins the data contained in the dimension table1 with the data contained in right outer merge joined source location and the target projection2. By way of particular example, and with reference to the diagram 100 in FIG. 1, the join at block 610 may add the identification data, i.e., names, associated with the UserIds.

[0050] At block 612, a dimension table2 is scanned, for instance, by the target scanning module 216. In addition, at block 614 a join, e.g., a hash/merge join, is performed on the output of the join performed at block 610 and the data contained in the dimension table2, for instance, by the processing module 226.

-15-

The join at block 614 generally joins the data contained in the dimension table2 with the data contained in the source location, the target projection2, and the dimension table1. By way of particular example, and with reference to the diagram 100 in FIG. 1, the join at block 610 may add the city information associated with the X and Y coordinates in the target location 120.

[0051] Blocks 616-628, respectively, are similar to blocks 508-520 in FIG. 5 and are thus not described in greater detail herein.

[0052] In various instances, data is distributed across different nodes of a database cluster. According to an example, a process of data redistribution is performed before insert and join operators in any of the methods 300-600 discussed above to substantially ensure that the appropriate data reaches the appropriate operators. Particularly, a determination is made that data is to be redistributed among multiple ones of the plurality of projections (or target locations). In addition, the data is redistributed among multiple ones of the plurality of projections in response to a determination that data is to be redistributed among multiple ones of the plurality of projections.

[0053] Generally speaking, redistributing data to an INSERT operator is a unary redistribution because the source of the data is from a single operator, FILTER. To generalize this process, and for purposes of example, the INSERT operator has been named 'Operator A' and the FILTER operator has been named 'Operator B' in FIGS. 7A-7D, 8A-8C, 9A-9C, and 10A-10G, which respectively depict diagrams of manners in which data may be redistributed among a plurality of projections, according to a plurality of examples. FIG. 7A shows a general redistribution of data from Operator B 704 to Operator A 702.

[0054] According to a first example, using the heuristics that the less data transferred the faster the plan, the data may be redistributed according to any of the following manners. In a first manner, and as shown in FIG. 7B, if each node of the cluster of Operator A 702 is expecting data on the same node of Operator B 704, no redistribution is needed. Instead, data is moved up locally from Operator B 704 to Operator A 702. This occurs, for instance, when data from Operator B 704 is already segmented as required by Operator A 702.

-16-

[0055] In a second manner, if each node of the cluster of Operator A 702 is expecting all of the data, the data at each node of Operator B 704 must be broadcast (sent all) to all nodes of Operator A 702, as shown in FIG. 7C, and described further in detail below with respect to the FIGS. 8A-8C. This occurs when data from Operator B 704 is segmented while Operator A 702 requires all the data.

[0056] In a third manner, if each node of the cluster of Operator A 702 is expecting a segment/range of the data that is different from the segment/range of the data of Operator B 704 on the same node, the data at each node of Operator B 704 must be resegmented to the same segment/range with Operator A 702 and then sent to Operator A 702 as shown in FIG. 7D, and described further in detail below with respect to FIGS. 9A-9C. This occurs when data from Operator B 704 is segmented but on different segment/range required by Operator A 702.

[0057] FIGS. 8A-8C, respectively, show the broadcasting process from Operator B 704 to Operator A 702 on a three-node cluster. Particularly, FIG. 8A shows the process at node 1, where the output data of Operator B 704 is moved directly up to Operator A 702, and sent to Operator A 702 on node 2 and node 3. The Operator A 702 at node 1, besides getting data from Operator B 702 on its own node, also waits to receive data sent from Operator B 702 on node 2 and node 3. A similar process is performed on node 2 and node 3 as illustrated in FIGS. 8B and 8C, respectively.

[0058] FIGS. 9A-9C, respectively, show resegmentation processes for a three-node cluster, which include two main steps: (1) resegment the output data of the Operator B 704 on the expression specified in Operator A 702, and (2) send the right data segments to the corresponding nodes of Operator A 702. FIG. 9A demonstrates the process at node 1, where the output data of Operator B 704 is first segmented on the expression specified in Operator A 702 to three segments: seg1, seg2, and seg3. Then seg1 is moved up to the node's Operator A 702, while seg2 and seg3 are sent to Operator A 702 of node 2 and node 3, respectively. The Operator A 702 at node 1, besides getting seg1 from

-17-

Operator B 704 on its own node, also waits to receive the remaining data of seg1 sent from Operator B on node 2 and node 3. A similar process is performed on node 2 and node 3 and illustrated in FIGS. 9B and 9C, respectively.

[0059] The data segmentation expression required from the INSERT operator in the MERGE plan is the data segmentation of the target location. Depending on the data segmentation of the INSERT operators (or target location) and the data segmentation of their input operators (FILTER in this case), the corresponding data redistribution on-the-fly (none or broadcast or resegment) will be chosen.

[0060] Redistributing data to a JOIN operator 706 is said to be a binary redistribution because the source of the data is from two input operators, as shown in FIG. 10A. In general, at each node, the JOIN operator 706 will join certain segments/ranges of data from the outer input and the inner input such that its result will be a subset of the whole join result. In other words, the union of all per-node join results will be the final join results, and the intersection of per-node join results will be empty. Thus, redistributing data to a join is dependent upon the data segmentation of the two join inputs and also based on the basic heuristic, that the less data transferred the faster the plan. As such, the data distribution is classified into six categories as depicted in FIGS. 10B-10G.

[0061] Particularly, as shown in FIG. 10B, redistribution is not needed if the data of the two join inputs are co-located by either (1) both being already segmented on their join attributes, or (2) by one being segmented on the join attributes, while the other is replicated (full copy of data is stored at each node), or (3) both inputs are replicated. Resegmenting data of the outer input on the join attributes if the data of the inner input is already segmented on the join attributes, as shown in FIG. 10C. Resegmenting data of the inner input on the join attributes if the data of the outer input is already segmented on the join attributes, as shown in FIG. 10D. Resegmenting both inner and outer data on the join attributes if neither are segmented on the join attributes yet, as shown in

-18-

FIG. 10E. Broadcasting data from the outer to the inner if both are segmented but not on the join attributes and the outer's data is much smaller than the inner's, as shown in FIG. 10F. Broadcasting data from the inner to the outer if both are segmented but not on the join attributes and the inner's data is much smaller than the outer's, as shown in FIG. 10G.

[0062] According to an example, an optimizer selects one of the choices described in FIGS. 10B-10G to redistribute the data for a join based on a custom-built cost model. In addition, there is a restriction in the outer-join in that data of the outer side cannot be broadcast. This is because the outer join does not eliminate tuples of the outer side and would lead to duplicate tuples when a union of the results of all nodes is later created. As the join between the source and target tables in the MERGE plan is a right-outer-join, the data redistribution that broadcasts the inner input to outer input as shown in FIG. 10G cannot be used.

[0063] Some or all of the operations set forth in the methods 300-600 may be contained as a utility, program, or subprogram, in any desired computer accessible medium. In addition, the methods 300-600 may be embodied by computer programs, which may exist in a variety of forms both active and inactive. For example, they may exist as machine readable instructions, including source code, object code, executable code or other formats. Any of the above may be embodied on a non-transitory computer readable storage medium. Examples of non-transitory computer readable storage media include conventional computer system RAM, ROM, EPROM, EEPROM, and magnetic or optical disks or tapes. It is therefore to be understood that any electronic device capable of executing the above-described functions may perform those functions enumerated above.

[0064] Turning now to FIG. 11, there is shown a schematic representation of a computing device 800, which may be employed to perform various functions of the machine 200 depicted in FIG. 2, according to an example. The computing device 800 includes a processor 802, such as but not limited to a central processing unit; a display device 804, such as but not limited to a

-19-

monitor; a network interface 808, such as but not limited to a Local Area Network LAN, a wireless 802.11 LAN, a 3G/4G mobile WAN or a WiMax WAN; and a computer-readable medium 810. Each of these components is operatively coupled to a bus 812. For example, the bus 812 may be an EISA, a PCI, a USB, a FireWire, a NuBus, or a PDS.

[0065] The computer readable medium 810 comprises any suitable medium that participates in providing instructions to the processor 802 for execution. For example, the computer readable medium 810 may be non-volatile media, such as memory. The computer-readable medium 810 may also store an operating system 814, such as but not limited to Mac OS, MS Windows, Unix, or Linux; network applications 816; and a data merging application 818. The operating system 814 may be multi-user, multiprocessing, multitasking, multithreading, real-time and the like. The operating system 814 may also perform basic tasks, such as but not limited to recognizing input from input devices, such as but not limited to a keyboard or a keypad; sending output to the display 804; keeping track of files and directories on medium 810; controlling peripheral devices, such as but not limited to disk drives, printers, image capture device; and managing traffic on the bus 812. The network applications 816 include various components for establishing and maintaining network connections, such as but not limited to machine readable instructions for implementing communication protocols including TCP/IP, HTTP, Ethernet, USB, and FireWire.

[0066] The data merging application 818 provides various components for merging data from a source location into a target location (projection(s)) as discussed above with respect to the methods 300-600 in FIGS. 3-6. The data merging application 818 may thus comprise the input/output module 212, the source location scanning module 214, the target location scanning module 216, the tuple matching module 218, the operator applying module 220, the filter applying module 222, the merging module 224, and the processing module 226. In this regard, the data merging application 818 may include modules for performing at least one of the methods 300-600.

-20-

[0067] In certain examples, some or all of the processes performed by the application 818 may be integrated into the operating system 814. In certain examples, the processes may be at least partially implemented in digital electronic circuitry, or in computer hardware, machine readable instructions (including firmware and software), or in any combination thereof, as also discussed above.

[0068] What has been described and illustrated herein are examples of the disclosure along with some variations. The terms, descriptions and figures used herein are set forth by way of illustration only and are not meant as limitations. Many variations are possible within the scope of the disclosure, which is intended to be defined by the following claims -- and their equivalents -- in which all terms are meant in their broadest reasonable sense unless otherwise indicated.

-21-

CLAIMS

What is claimed is:

1. A method for merging data from a source location into a target location
5 containing existing data, wherein both the source location and the target
location contain tuples of data, said method comprising:
for each tuple in the source location, during a single operation,
determining, by a processor, whether there is a matched tuple in
the target location that satisfies a predetermined condition;
10 for each matched tuple that satisfies the predetermined condition,
updating the matched tuple in the target location with a count value that is equal
to a count of the matched tuple in the source location and the target location;
and
for each tuple that does not have a matched tuple that satisfies the
15 predetermined condition, inserting the unmatched tuple into the target location.
2. The method according to claim 1, further comprising:
scanning both the source location and the target location to determine
the tuples contained in each of the source location and the target location;
20 wherein determining whether there is a matched tuple in the target
location that satisfies a predetermined condition further comprises right outer
joining the source location and the target location to determine which tuples
from the source location are to be inserted into the target location and which
tuples in the target location are to be updated, inserting the matched tuples into
25 a DeleteVector, and inserting the matched tuples and the unmatched tuples into
the target location.
3. The method according to claim 2, wherein right outer joining the source
location and the target location further comprises adding a matching Boolean
30 (M) to each output tuple of the right outer join to mark whether the tuple is new
(M=false) or existing (M=true);
moving matched tuples (M=true) into the DeleteVector; and

-22-

inserting unmatched tuples (M=false) and matched tuples (M=true) into the target location.

4. The method according to claim 3, further comprising:

5 applying a first filter to correctly output the unmatched tuples into the DeleteVector;

applying appropriate expressions for the matched and unmatched tuples; and

10 applying a second filter to correctly output the unmatched tuples and the matched tuples prior to inserting the unmatched tuples and the matched tuples into the target location.

5. The method according to claim 4, further comprising:

15 implementing a sideways information passing operation that eliminates unmatching tuples from the target location before performing the right outer joining; and

wherein right outer joining the source location and the target location further comprises right outer joining the source location and the target location following implementation of the sideways information passing operation.

20

6. The method according to claim 1, wherein the target location comprises a projection, said method further comprising:

scanning both the source location and the projection to determine the tuples contained in each of the source location and the projection;

25 wherein determining whether there is a matched tuple in the target location that satisfies a predetermined condition further comprises determining whether there is a matched tuple in the projection by right outer merge joining the source location and the projection to determine which tuples from the source location are to be inserted into the projection and which tuples in the projection
30 are to be updated, inserting the matched tuples into a DeleteVector of the projection, and inserting the unmatched and the matched tuples in the projection.

-23-

7. The method according to claim 6, wherein the target location comprises a plurality of projections, said method further comprising:

5 selecting one of the plurality of projections to scan based upon at least one factor, wherein scanning both the source location and the projection further comprises scanning both the source location and the selected one of the plurality of projections;

inserting the matched tuples into DeleteVectors of each of the plurality of projections; and

10 inserting the unmatched and the matched tuples in each of the plurality of projections.

8. The method according to claim 7, further comprising:

scanning a dimension table containing descriptive attributes of events;

15 joining the descriptive attributes of events in the dimension table with the data contained in an additional projection with the right outer merge joined source location and the selected one of the plurality of projections;

inserting the unmatched tuples from the join into each of the plurality of projections; and

20 updating the matched tuples from the join in each of the plurality of projections.

9. The method according to claim 8, wherein the existing data of the target location is distributed across different nodes, said method further comprising:

25 determining that the distributed existing data is to be redistributed among the different nodes; and

redistributing the distributed existing data among the different nodes in response to a determination that the distributed existing data is to be redistributed among the different nodes.

30

10. An apparatus for merging data from a source location into a target location containing existing data, wherein both the source location and the

-24-

target location contain tuples of data, and wherein the target location comprises a column-based database, said apparatus comprising:

a module to,

scan both the source location and the target location to determine
5 the tuples contained in each of the source location and the target location;

right outer join the source location and the target location to
determine which tuples from the source location are to be inserted into the
target location and which tuples in the target location are to be updated;

for each of the tuples in the target location that are to be updated,
10 update the tuples in the target location with a count value that is equal to a
count of the tuple in the source location and the target location; and

for each of the tuples in the source location that are to be inserted
into the target location, insert the tuples into the target location; and

a processor to implement the module.

15

11. The apparatus according to claim 10, wherein the module is further to:
add a matching Boolean (M) to each output tuple of the right outer join to
mark whether the tuple is new (M=false) or existing (M=true);

insert existing tuples (M=true) into a DeleteVector; and

20 insert new tuples (M=false) and existing tuples (M=true) into the target
location.

12. The apparatus according to claim 11, wherein the module is further to:

apply a first filter to output the new tuples into the DeleteVector;

25 apply appropriate expressions for the existing and new tuples; and

apply a second filter to correctly output the new tuples and the existing
tuples prior to inserting the new tuples and the existing tuples into the target
location.

30 13. The apparatus according to claim 12, wherein the module is further to:

implement a sideways information passing operation that eliminates new
tuples from the target location before performance of the right outer join.

-25-

14. The apparatus according to claim 10, wherein the target location comprises at least one projection.

5 15. A non-transitory computer readable storage medium on which is stored machine readable instructions that when executed by a processor, implement a method for merging data from a source location into a target location containing existing data, wherein both the source location and the target location contain tuples of data, and wherein the target location comprises a column-based
10 database, said machine readable instructions comprising code to:

scan both the source location and the target location to determine the tuples contained in each of the source location and the target location;

right outer join the source location and the target location to determine which tuples from the source location are to be inserted into the target location

15 and which tuples in the target location are to be updated;

for each of the tuples in the target location that are to be updated, update the tuples in the target location with a count value that is equal to a count of the tuple in the source location and the target location; and

20 for each of the tuples in the source location that are to be inserted into the target location, insert the tuples into the target location.

100




FIG. 1






FIG. 2

3 / 12

300*FIG. 3*

4 / 12

400*FIG. 4*

5 / 12


500

FIG. 5

6 / 12








FIG. 6

7 / 12

*FIG. 7A**FIG. 7B**FIG. 7C**FIG. 7D*




FIG. 8A




FIG. 8B




FIG. 8C











FIG. 9C

FIG. 9B

FIG. 9A

10 / 12

*FIG. 10A**FIG. 10B**FIG. 10C*

*FIG. 10D**FIG. 10E**FIG. 10F**FIG. 10G*




FIG. 11

A. CLASSIFICATION OF SUBJECT MATTER**G06F 17/30(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 17/30; G06F 17/00; G06F 7/00; G06F 12/00

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: merge, tuple, predetermined, condition, match

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 7523123 B2 (HUNG-CHIH YANG et al.) 21 April 2009 See abstract; column 2, line 34 - column 3, line 37; figures 2-5.	1-15
A	US 7627613 B1 (DANIEL DULITZ et al.) 01 December 2009 See abstract; column 10, line 20 - column 11, line 28; claim 1 and figure 5.	1-15
A	US 2008-0235260 A1 (WOOK-SHIN HAN et al.) 25 September 2008 See abstract; paragraphs[0023-0028]; claim 17 and figure 1.	1-15
A	US 7089266 B2 (CHRIS STOLTE et al.) 08 August 2006 See abstract; column 23, line 33 - column 25, line 5; claim 1 and figure 6.	1-15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

20 DECEMBER 2012 (20.12.2012)

Date of mailing of the international search report

27 DECEMBER 2012 (27.12.2012)

Name and mailing address of the ISA/KR


Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan
City, 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

LEE, Myung Jin

Telephone No. 82-42-481-8474



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2012/040505

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 7523123 B2	21.04.2009	US 2008-120314 A1	22.05.2008
US 7627613 B1	01.12.2009	US 2010-0076954 A1	25.03.2010
		US 2011-276561 A1	10.11.2011
		US 7984054 B2	19.07.2011
US 2008-0235260 A1	25.09.2008	US 2008-0275907 A1	06.11.2008
US 7089266 B2	08.08.2006	AU 2004-248607 A1	23.12.2004
		CA 2526825 A1	23.12.2004
		EP 1634142 A2	15.03.2006
		EP 1634142 A4	16.01.2008
		JP 2006-526840 A	24.11.2006
		JP 2006-526840 T	24.11.2006
		US 2004-0243593 A1	02.12.2004
		US 2006-0259509 A1	16.11.2006
		US 2011-0131250 A1	02.06.2011
		US 7716173 B2	11.05.2010
		US 7882144 B1	01.02.2011
		US 8140586 B2	20.03.2012
		WO 2004-111786 A2	23.12.2004
		WO 2004-111786 A3	23.03.2006
		WO 2004-111786 A3	23.12.2004