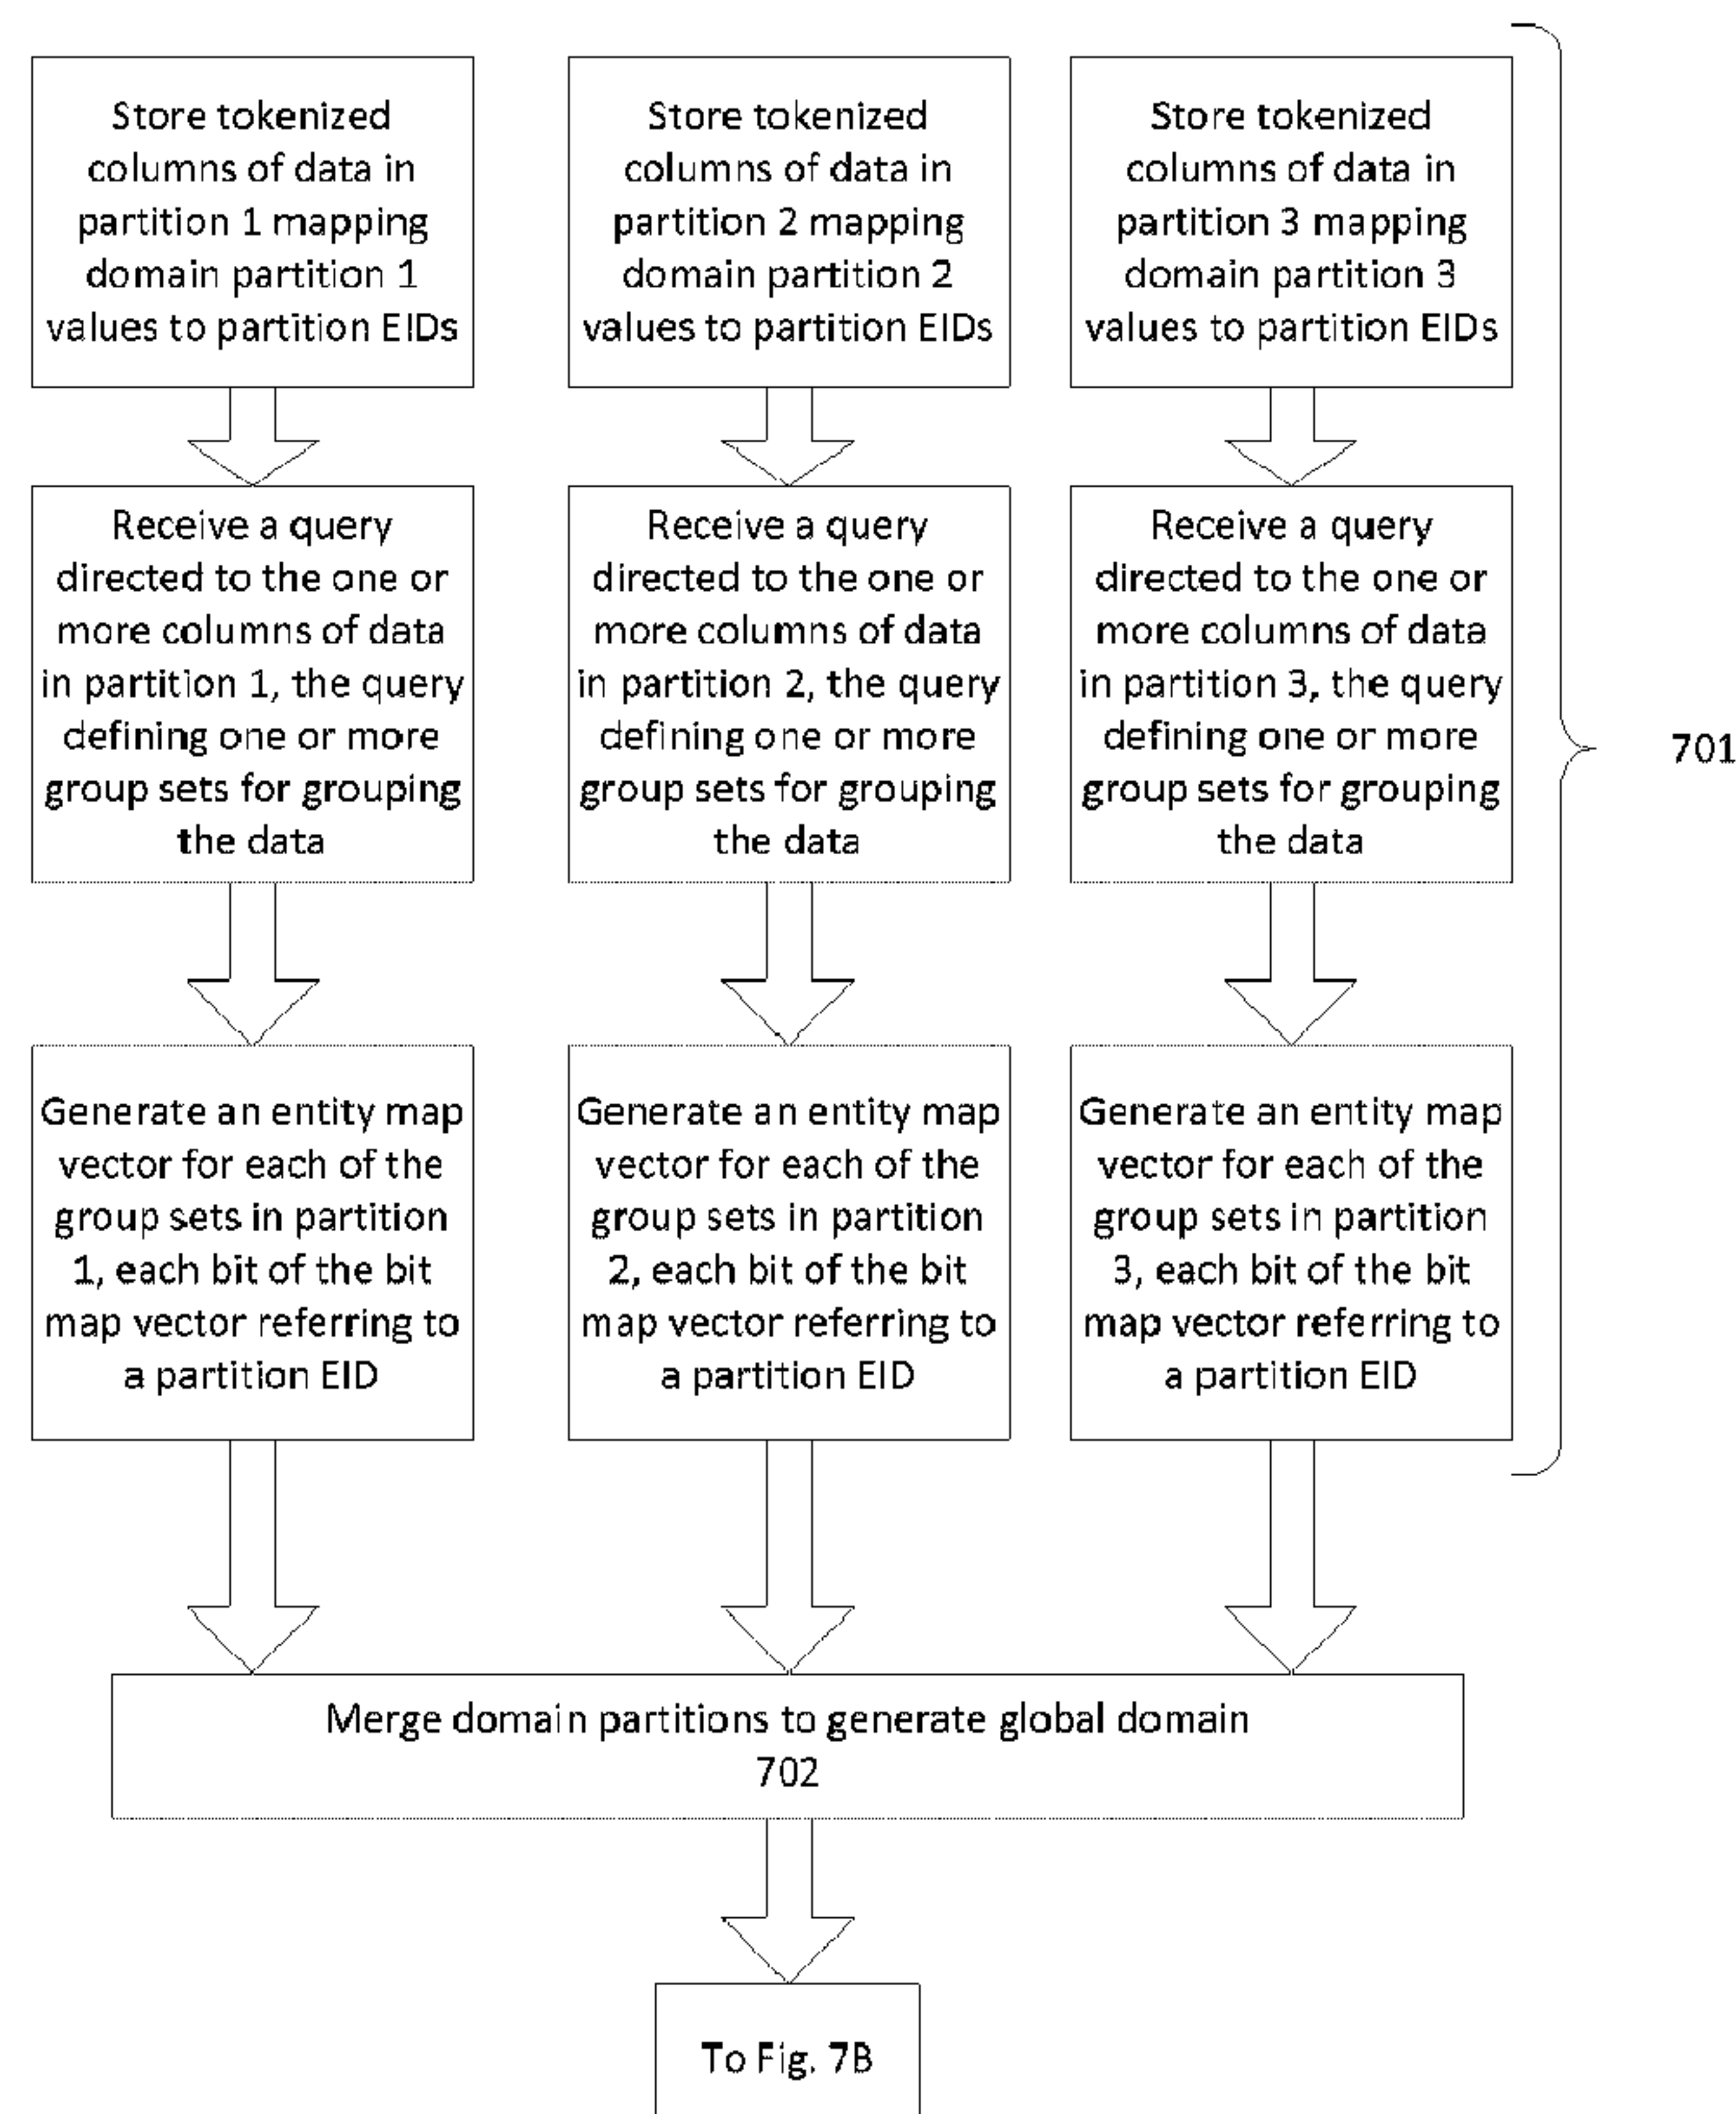




(86) Date de dépôt PCT/PCT Filing Date: 2014/03/14  
(87) Date publication PCT/PCT Publication Date: 2014/09/18  
(45) Date de délivrance/Issue Date: 2019/11/26  
(85) Entrée phase nationale/National Entry: 2015/09/14  
(86) N° demande PCT/PCT Application No.: US 2014/027907  
(87) N° publication PCT/PCT Publication No.: 2014/143791  
(30) Priorité/Priority: 2013/03/15 (US13/835,590)

(51) Cl.Int./Int.Cl. *G06F 16/22* (2019.01),  
*G06F 16/24* (2019.01), *G06F 7/00* (2006.01)  
(72) Inventeurs/Inventors:  
GRONDIN, RICHARD, CA;  
FADEITCHEV, EVGUENI, CA  
(73) Propriétaire/Owner:  
INFORMATICA LLC, US  
(74) Agent: RIDOUT & MAYBEE LLP

(54) Titre : PROCÉDE, APPAREIL ET SUPPORT LISIBLE PAR ORDINATEUR D'EXECUTION EFFICACE  
D'OPERATIONS SUR DES VALEURS DE DONNEES DISTINCTES  
(54) Title: EFFICIENTLY PERFORMING OPERATIONS ON DISTINCT DATA VALUES



(57) **Abrégé/Abstract:**


An apparatus, computer-readable medium, and computer-implemented method for efficiently performing operations on distinct data values, including storing a tokenized column of data in a table by mapping each unique data value in a corresponding domain

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

to a unique entity ID, and replacing each of the data values in the column with the corresponding entity ID to generate a column of tokenized data containing one or more entity IDs, receiving a query directed to the column of data, the query defining one or more group sets for grouping the data retrieved in response to the query, and generating an entity map vector for each group set, the length of each entity map vector equal to the number of unique entity IDs for the domain, and the value of each bit in the entity map vector indicating the presence or absence of a different unique entity ID in the group set.

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

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
18 September 2014 (18.09.2014)

WIPO | PCT

(10) International Publication Number  
**WO 2014/143791 A1**

(51) International Patent Classification:  
*G06F 17/16* (2006.01)

(21) International Application Number:  
PCT/US2014/027907

(22) International Filing Date:  
14 March 2014 (14.03.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
13/835,590 15 March 2013 (15.03.2013) US

(71) Applicant: **INFORMATICA CORPORATION**  
[US/US]; 2100 Seaport Blvd., Redwood City, CA 94063 (US).

(72) Inventors: **GRONDIN, Richard**; 3 Rue Vauquelin, Ste-julie, QC J3E 2Z1 (CA). **FADEITCHEV, Evgueni**; 5139 Prince of Wales, Montreal, QC H4V 3N2 (CA).

(74) Agent: **GREWAL, Amardeep, S.**; Reed Smith LLP, 1301 K Street, N.W., Suite 1100, East Tower, Washington, DC 20005-3317 (US).

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, 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, KM, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: EFFICIENTLY PERFORMING OPERATIONS ON DISTINCT DATA VALUES

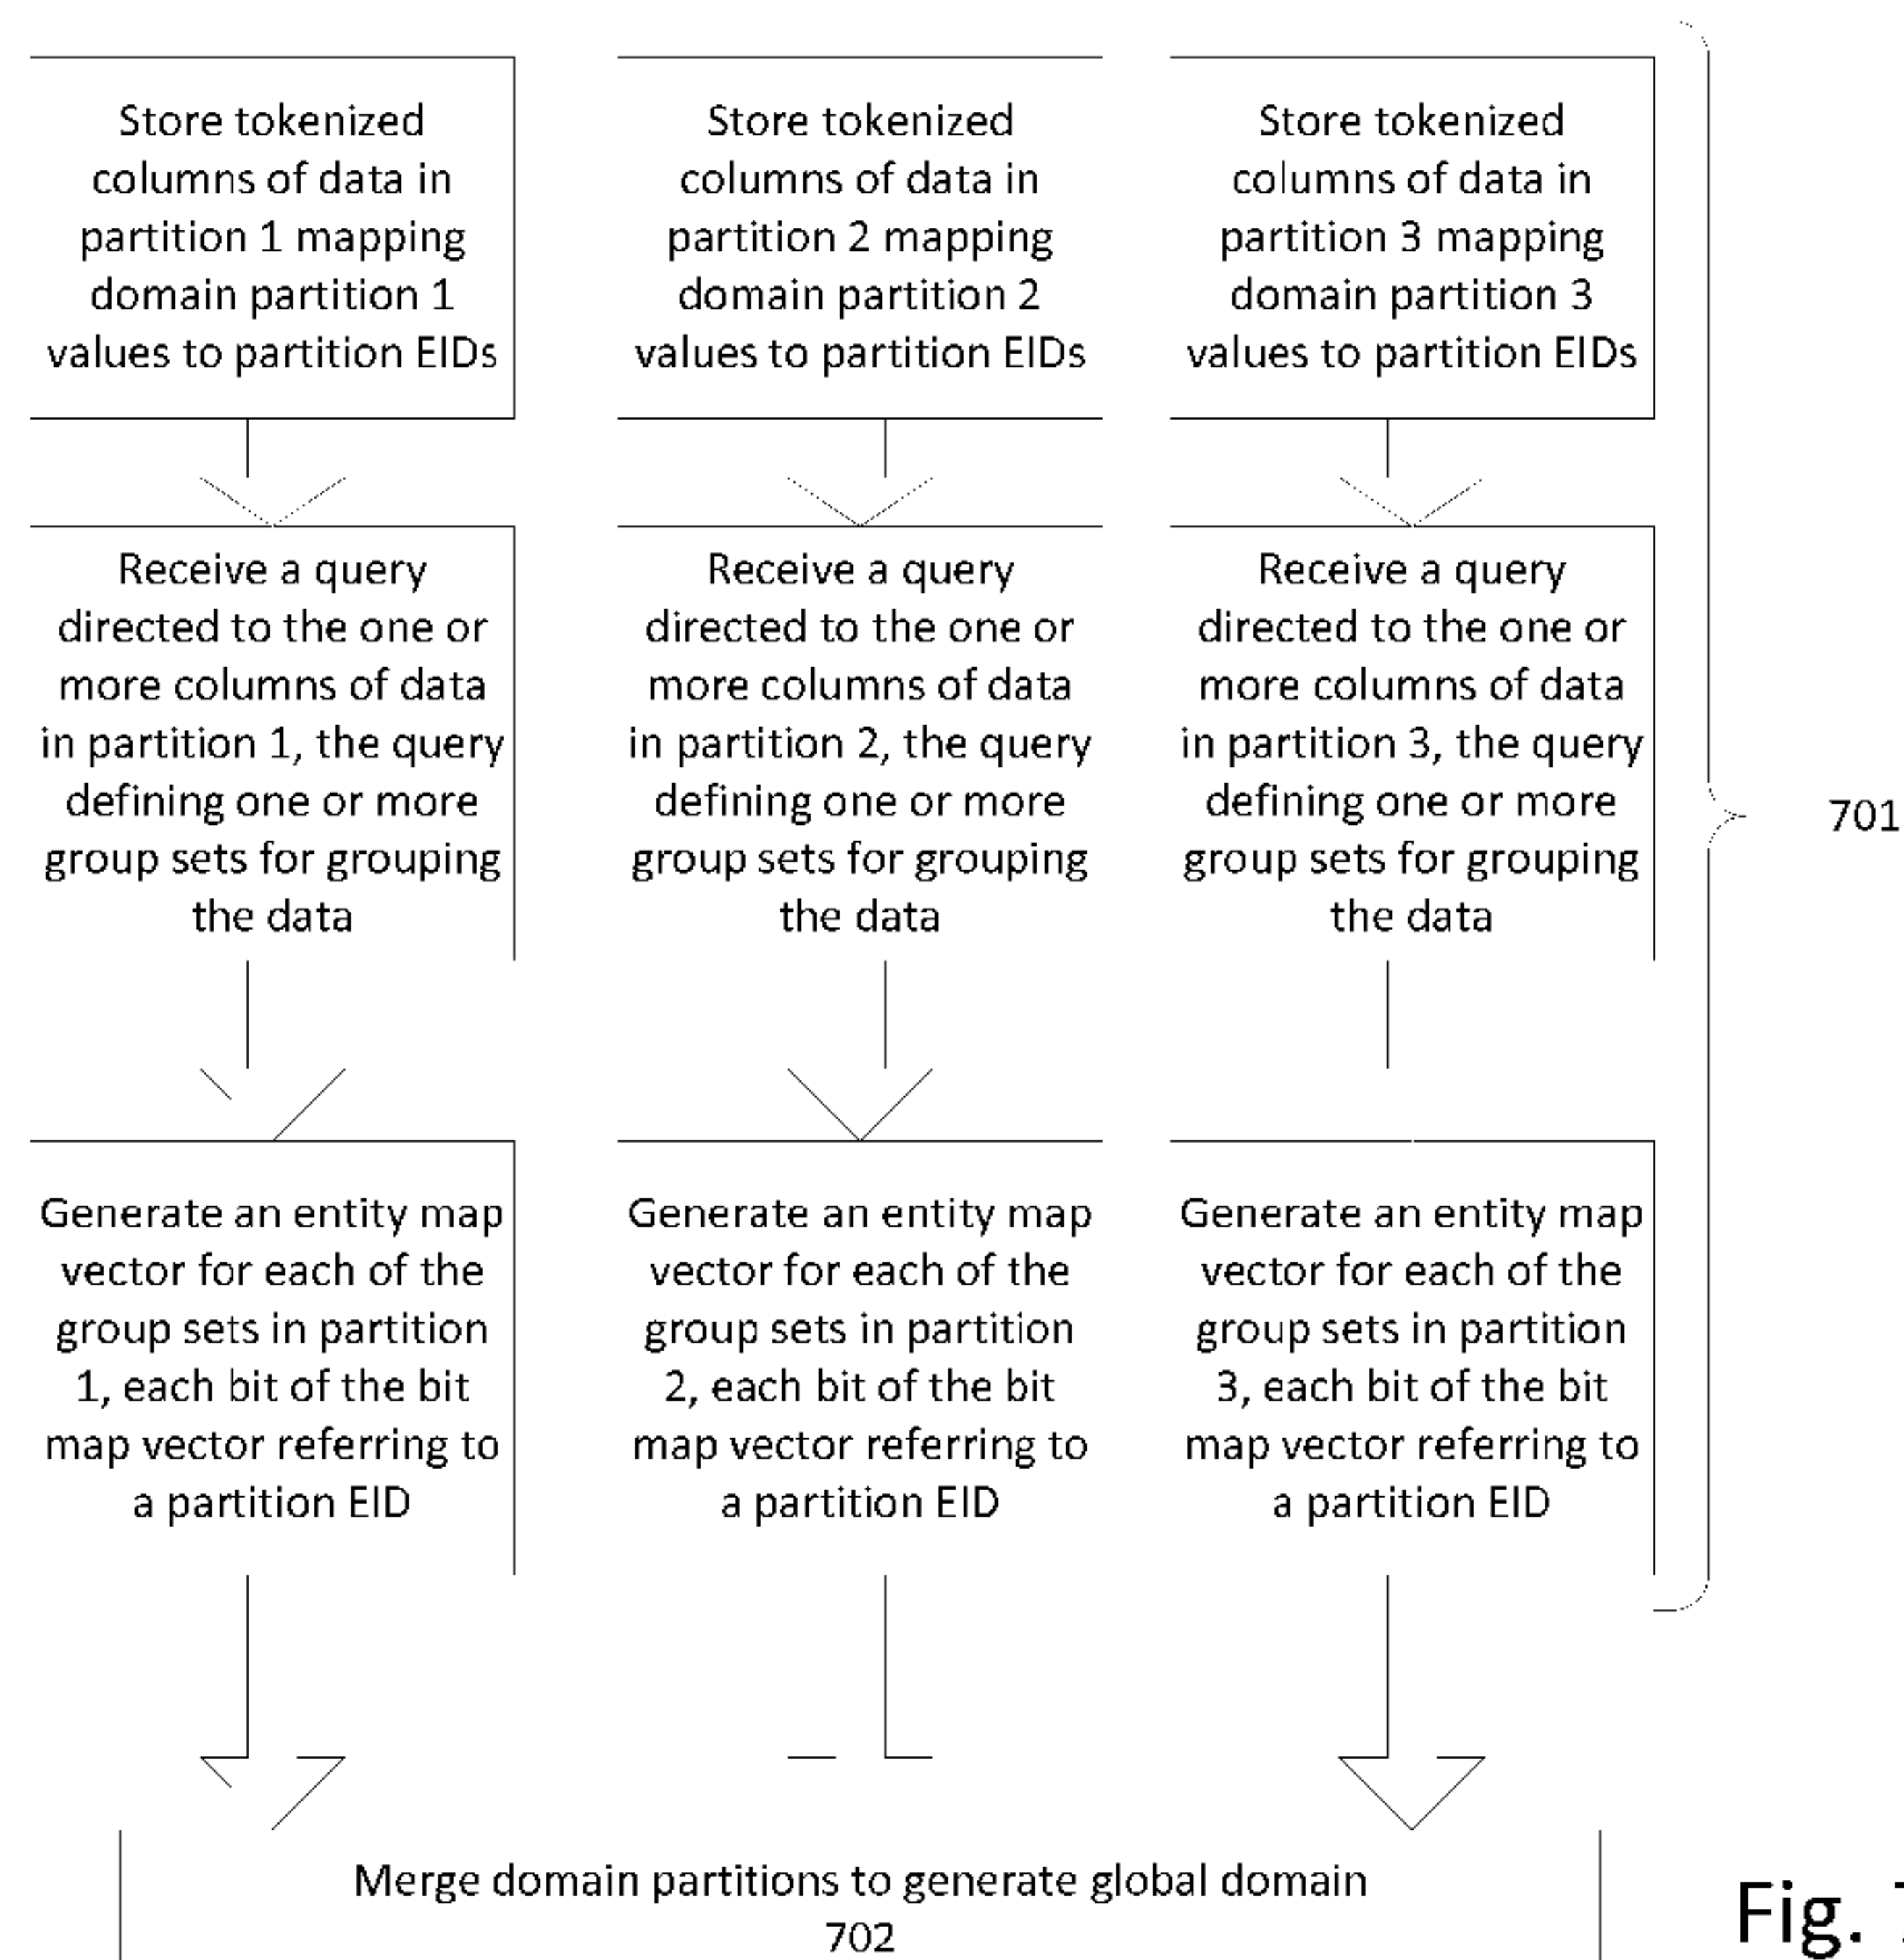


Fig. 7A

(57) Abstract: An apparatus, computer-readable medium, and computer-implemented method for efficiently performing operations on distinct data values, including storing a tokenized column of data in a table by mapping each unique data value in a corresponding domain to a unique entity ID, and replacing each of the data values in the column with the corresponding entity ID to generate a column of tokenized data containing one or more entity IDs, receiving a query directed to the column of data, the query defining one or more group sets for grouping the data retrieved in response to the query, and generating an entity map vector for each group set, the length of each entity map vector equal to the number of unique entity IDs for the domain, and the value of each bit in the entity map vector indicating the presence or absence of a different unique entity ID in the group set.

## **EFFICIENTLY PERFORMING OPERATIONS ON DISTINCT DATA VALUES**

**[0001]**

### **BACKGROUND**

**[0002]** Enterprises frequently store data in data warehouses. Such data warehouses often have large volumes of data spread out over multiple databases and database tables. Therefore, in order to compile data from the numerous tables and databases, data aggregation techniques must be utilized.

**[0003]** Low Level Aggregate (LLA) tables, which store a summary of data at a lower level, are frequently used to improve the performance of the aggregation process in a data warehouse. In the case of a distributed database stored on multiple partitions, an LLA table can be utilized for each partition and an aggregation process can combine the data from each of the LLA tables to generate an aggregated table for the entire database. Such an aggregation process is known as a roll-up process, and is useful because a higher level of summary can be generated from a lower level of summary (the LLA table), rather than from the source data.

**[0004]** Roll-up processes can be utilized for an aggregate query over multiple data partitions in the context of a Massively Parallel Processing (MPP) database architecture, or to run queries on a partitioned table.



[0005] In order to correctly aggregate data from multiple tables during a roll-up process, the underlying function that is being aggregated must be additive. For example, if a first table and a second table both have a column corresponding to spending amounts, and a user desires to compute the sum of all the spending amounts in both tables, they can additively aggregate the sum of the spending amounts in the first table and the sum of the spending amounts in the second table and arrive at the correct total sum. However, not all functions that users may wish to aggregate in a roll-up process are additive functions.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0006] Figs. 1A-1G illustrate a sample database table containing records of visits to stores by customers and some examples of low level aggregation tables generated from the sample database table.

[0007] Figs. 2A-2B are flowcharts of methods for efficiently performing operations on distinct data values according to a disclosed embodiment.

[0008] Figs. 3A-3C are tables illustrating steps for tokenizing one or more columns of data in a table according to a disclosed embodiment.

[0009] Figs. 4A-4C illustrate an LLA table generated from a tokenized table and two sample higher level aggregate tables generated from the LLA table.

[0010] Fig. 5A illustrates an LLA table generated from a tokenized table which utilizes the entity map vector data structure according to a disclosed embodiment.

[0011] Fig. 5B illustrates a method of assigning values to each of the bits in the entity map vector according to a disclosed embodiment.

- [0012] Fig. 5C illustrates a table of lexically ordered unique data values in a domain according to a disclosed embodiment.
- [0013] Fig. 5D illustrates a method of assigning values to each of the bits in the entity map vector according to a disclosed embodiment.
- [0014] Fig. 6A illustrates the steps involved in carrying out a roll-up process on the entity map vectors stored in the LLA table of Fig. 5A according to a disclosed embodiment.
- [0015] Fig. 6B illustrates the higher level table that is the result of the roll-up process shown in Fig. 6A.
- [0016] Figs. 7A-7B illustrate a flowchart for efficiently performing operations on distinct data values in a distributed database by merging domains across multiple partitions according to a disclosed embodiment.
- [0017] Figs. 8A-8D illustrate each of steps required to convert a sample set of partition specific entity IDs into global entity IDs according to a disclosed embodiment.
- [0018] Fig. 9 illustrates a data cube for storing global entity map vectors across multiple partitions according to a disclosed embodiment.
- [0019] Fig. 10 illustrates a database warehouse update operation using chronologically assigned entity IDs according to a disclosed embodiment.
- [0020] Fig. 11A illustrates domain synchronization data structures and process flow according to a disclosed embodiment.
- [0021] Fig. 11B illustrates the mapping of local entity map vectors to a global entity map vectors using a mapping table.

**[0022]** Fig. 12 illustrates an exemplary computing environment that can be used to carry out the method for efficiently performing operations on distinct data values according to a disclosed embodiment.

### **DETAILED DESCRIPTION**

**[0023]** While methods, apparatuses, and computer-readable media are described herein by way of examples and embodiments, those skilled in the art recognize that methods, apparatuses, and computer-readable media for efficiently performing operations on distinct data values are not limited to the embodiments or drawings described. It should be understood that the drawings and description are not intended to be limited to the particular form disclosed. Rather, the intention is to cover all modifications, equivalents and alternatives falling within the scope of the appended claims. Any headings used herein are for organizational purposes only and are not meant to limit the scope of the description or the claims. As used herein, the word “can” is used in a permissive sense (i.e., meaning having the potential to) rather than the mandatory sense (i.e., meaning must). Similarly, the words “include,” “including,” and “includes” mean including, but not limited to.

**[0024]** Some examples of roll-up processes for additive functions will now be discussed with reference to the table in Fig. 1A, which shows a sample database table containing records of visits to stores by customers. The table contains columns for customer name, store ID, and amount spent during the visit.

**[0025]** The query language used throughout this specification is Structured Query Language (SQL), but it is understood that any database query language can be utilized without



departing from the scope of this invention, including Contextual Query Language (CQL), XQuery, YQL, Datalog, OQL, RDQL, Multidimensional Expressions (MDX), and many others.

**[0026]** An LLA table of the number of visits to each store is shown in Fig. 1B. The “Count of Visits” column is generated by counting the number of records for each of the stores, resulting in 3 visits for store 1 and 4 visits for store 2. If the user wishes to execute a roll-up process using this table in order to determine the total number of visits to all the stores, the data can be rolled up with the function  $\text{Sum}(\text{Count of Visits})$ , which will result in a total value of 7. In this case, there is no need to go back to the original table to determine the total number of visits to stores.

**[0027]** Similarly, the table in Fig. 1C is an LLA table of the total amount spent at each of the stores. This amount is determined by summing the values in the “Amount Spent” column of the original table for each store. If a user wants to roll-up this table to determine the total amount spent at all stores, they can use the function  $\text{Sum}(\text{Amount Spent})$  on the LLA table of Fig. 1C to arrive at the correct total amount of 107.

**[0028]** The aggregation of LLA tables during a roll-up process can be accomplished for functions outside of simple sums as well. Fig. 1D shows an LLA table of the average amount spent at each store per visit. This amount is determined by averaging the values in the “Amount Spent” column of the original table for each store. Additionally, the LLA table contains a count of the number of visits to each store as discussed earlier. If a user wants to roll-up this table to determine the average amount spent at all stores, they can use the following function:

$$\text{Average Spent (total)} = \frac{(\text{Avg}_{\text{store 1}} \times \text{Visits}_{\text{store 1}}) + (\text{Avg}_{\text{store 2}} \times \text{Visits}_{\text{store 2}})}{\text{Sum}(\text{Visits})}$$



[0029] Once again, the roll-up process does not require the original table and can be accomplished using the LLA table alone.

[0030] Figs. 1E and 1F show LLA tables containing the minimum spent at each store and the maximum spent at each store, respectively. To roll-up the table in Fig. 1E to determine minimum amount spent at any store, the user can use the function Min(Minimum Spent) on the LLA table and arrive at the correct answer of 5. Similarly, to roll-up the table in Fig. 1F to determine maximum amount spent at any store, the user can use the function Max(Maximum Spent) on the LLA table and arrive at the correct answer of 31.

[0031] Referring now to Fig. 1G, an LLA table is shown which contains a count of the number distinct customers which have visited each store. This can be accomplished with a Count(distinct<Customer Name>) function that is grouped by store ID. So, in this case, 3 distinct customers (Bill, George, and Larry) visited store 1 and 4 distinct customers (Bill, George, Max, and Larry) visited store 2.

[0032] If a user wished to roll-up the LLA table in Fig. 1G, there would be no way for them to correctly determine the total number of distinct customers that frequented store 1 or store 2 without accessing the underlying main table in Fig. 1A. A count-distinct function run on the distinct customers column of the LLA table would only return a total of 2 distinct values, and a sum would return 7, when the correct number of distinct customers to visit both stores is 4. Similarly, if the LLA table contained a sum of all the distinct dollar amounts spent at each store, and the user wanted to compute the sum of all the distinct dollar amounts spent at all stores, they would have to access the original table, since otherwise the 11 dollars spent at store 1 and the 11 dollars spent at store 2 would be double counted. Additionally, if the LLA table contained an average of all the distinct amounts spent at each store, a roll-up to determine the average of all of

the distinct amounts spent at both stores would be unsuccessful and the underlying main table would have to be accessed.

**[0033]** Queries that include non-additive commands, such as count<distinct>, cannot make use of roll-up operations of LLA tables, resulting in much higher computation costs, since these commands must be run on the underlying data. In a partitioned or otherwise distributed database, this cost is magnified because the underlying data must be aggregated or exchanged between partitions in order to retrieve the correct values. For example, if column XYZ is spread out over two partitions, and the user wants to determine the number of distinct values in column XYZ, then all of the values of column XYZ on the first partition must somehow be pooled with all the values of column XYZ on the second partition to determine the number of distinct values.

**[0034]** Applicants have discovered a method and associated data structure, referred to as the Entity Map Vector (EMV), which allows users to use non-additive aggregate functions as additive, thereby allowing roll-up processes to be carried out on LLA tables for functions which would previously have required access to and processing of underlying data in tables at a lower level than the LLA tables. Applicants' technique further allows for the amount of data exchanged between nodes of a partitioned database to be greatly reduced, since LLA tables for each node can be additively aggregated even for non-additive aggregate functions.

**[0035]** Fig. 2A shows a flowchart of a method for generating an EMV according to a disclosed embodiment. At step 201, a tokenized column of data in a table of a database is stored prior to generating the EMV. This tokenization process is described with reference to Figs. 3A-3C. Additionally, as shown in Fig. 2B and described later, an EMV can be generated even without a tokenized column of data, based on a lexical ordering of the unique data values in a particular domain.



**[0036]** Fig. 3A shows a table of customer transactions before tokenization 300. The table includes tuple ID 301, and data columns for week 302, store ID 303, customer ID 304, product ID 305, quantity 306, and amount spent 307. Each of the data columns can correspond to an associated domain. For example, a domain containing customer ID's can be used to populate customer ID column of the table shown, but the same domain can be used to populate another table that lists customer ID's with customer addresses. Of course, a domain can also be used to populate multiple columns in the same table. For example, a domain of city names can be used to populate columns in a table corresponding to an origin city and a destination city. So the column for week 302 can correspond to domain D\_Week, store ID 303 can correspond to domain D\_StoreId, customer ID 304 can correspond to domain D\_CustId, product ID 305 can correspond to domain D\_ProductId, quantity 306 can correspond to domain D\_Qty, and amount spent 307 can correspond to domain D\_Spent.

**[0037]** In order to tokenize the data values in each of the data columns, the unique data values that occur in each of the domains (corresponding to the different columns) must be identified and mapped to unique token values. Mapping table 310 in Fig. 3B shows this mapping of unique data values to unique token values for each domain in the table. Each of the unique data values in each of the domains is mapped to a unique token for that domain, referred to as an Entity ID, or EID. So, for example, D\_Week 312 contains three unique data values and therefore those three unique data values are mapped to three corresponding unique EIDs. Similarly, D\_StoreID 313 also has three unique data values and three EIDs, D\_CustId 314 has four unique data values and four EIDs, D\_ProductId has four unique data values and four EIDs, D\_Qty has five unique data values and five EIDs, and D\_Spent has six unique data values and six EIDs.



[0038] The mapping tables used to map unique data values in a domain to corresponding EIDs can be stored as a look-up dictionary, so that functions and associated processes can retrieve data values using EIDs if necessary. Additionally, the mapping of unique data values to EIDs is shown as occurring in the lexical order of the unique data values, since this can simplify the assignment process, future data storage, and retrieval, but the assignment of EIDs to unique data values can occur in any order of unique data values. For example, EIDs can be assigned to be assigned to data values in chronological order, so that a first data value encountered before a second data value will have a lower EID.

[0039] Fig. 3C shows the table of customer transactions after tokenization 320. The EIDs from mapping table 310 have replaced the data values in each of the columns of data. However, even in this advanced encoded structure, any LLA tables generated via from the table 320 will not store data in the proper format to execute roll-up processes involving non-additive (distinct-type) aggregate functions.

[0040] For example, an LLA table can be generated based on the transaction table (“Trx”) 320 that lists the week, the store ID, and the distinct customers that visit each store each week using the following definition:

```
Create Table LLA_TRX as
Select Week,
       StoreId,
       Count(distinct CustId) as CD_CustId,
From Trx
Group By Week, StoreId;
```

[0041] This definition assigns the value corresponding to the distinct customers that visited each store each week to the column CD\_CustId. Additionally, the “group by” command creates group sets corresponding to week-store ID pairs. For example, group set 1(week 1, store

1), group set 2 (week 1, store 2), and so on. Any arrangement or number of group sets are possible. For example, the group sets can be based on a single column, or two or more columns.

**[0042]** Referring to Fig. 4A, the resulting LLA\_TRX table 400 is shown. The table accurately represents the number of distinct customers per group set. For example, the table correctly lists the number of distinct customers that visited store 1 in week 1 as two, or the number of distinct customers that visited store 1 in week 2 as three.

**[0043]** If a user tried to roll-up the table LLA\_TRX for each week to list the number of distinct customers that visited any of the stores each week, they would be unable to produce a correct count of the number of distinct customers for each group set, the group sets in this case being group 1(week 1), group 2(week 2), and so on. For example, if the following definition was used to create a higher level LLA table:

```
Create Table LLA_TRX_WEEK as
Select Week,
        Count(distinct CD_CustId) as CD_CD_CustId,
From LLA_TrX
Group By Week;
```

**[0044]** Then the resulting table, shown as Table 410 in Fig. 4B, would incorrectly list the number of distinct customers that visited any of the stores each week. Similarly, if a sum function was used to aggregate the table LLA\_TRX:

```
Create Table LLA_TRX_WEEK as
Select Week,
        Sum(distinct CD_CustId) as Sum_CD_CustId,
From LLA_TrX
Group By Week;
```

[0045] Then the resulting table, shown in Fig. 4C, would also incorrectly list the number of distinct customers that visited any of the stores each week. For example, the number of distinct customers in week 2, based on the table 320 in Fig. 3C, is 4 (custId's 1,2,3 and 4), but the table 410 in Fig. 4B lists 2 and the table in Fig. 4C 420 lists 5.

[0046] As discussed earlier, this is because the aggregation of distinct values is non-additive. In order to determine the correct values for the number of distinct customers per week, it is necessary to return to the original data table, making the LLA\_TRX table useless for non-additive aggregation queries.

[0047] The EMV data structure generated according to a disclosed embodiment solves this problem by storing the distinct values of each group set efficiently, and in a way that allows them to be aggregated. Returning to Fig. 2A, at step 202, a special query can be received which specifies one or more group sets and includes a new function, or other similar command, which is used to generate the EMV data structure. At step 203, the computing system can generate one or more entity map vectors. Of course, steps 201, 202, and 203 can be executed in a different order. For example, a query can first be received (step 202), then the tokenized column of data can be stored (step 201), and then one or more entity map vectors can be generated (step 203).

[0048] To facilitate steps 202-203, a new SQL function count (Additive\_distinct<Column Name>) can be defined to instruct the computing system to produce an EMV for each group set, instead of a value. An EMV can be a bit map vector, with each bit mapped to a different entity. By storing the information relating to the distinct values in a group set in an EMV, each distinct value, referred to as an entity, can be mapped to one bit of a bit map vector. So using the earlier example of the tokenized transactions table (Trx) 320 in Fig. 3C, the following definition can be used to generate the LLA table LLA\_TRX:



```

Create Table LLA_TRX as
Select  Week,
        StoreId,
        Count(Additive_distinct CustId) as CD_CustId,
From Trx
Group By Week, StoreId;

```

**[0049]** Fig. 5A shows the resulting LLA\_TRX table 500. As in the previous LLA\_TRX table, the tuple ID 501 corresponding to each group set, week 502, store ID 503, and count of distinct customers 504 are stored. However, rather than storing the count of distinct customers as a numeral, the distinct customers for each group set are stored in an EMV with each bit referring to different EID and the value of the bit indicating the presence or absence that EID in the group set.

**[0050]** Of course, an SQL extension, such as Additive\_distinct, is not required to instruct the computing system to produce an EMV for each group set. For example, the database engine can detect the creation of a table with a column that involves an aggregation function of distinct data values, such as Count (distinct<Column Name>), and automatically build an EMV to store the count of distinct data values. This allows the EMV structure to be utilized with existing front end tools.

**[0051]** The structure of an EMV will now be described with respect to the second EMV 513 corresponding to the group set (week 1, store ID 2). Referring to Fig. 5B, an empty EMV 505 is shown. Lead characters 506 serve only to label the EMV as a bit map vector and can be any suitable characters or can be omitted. The length dimension of the EMV 505 is the number of bit slots and does not include any lead characters. The length of EMV is equal to the total number of EIDs in the domain of the column that is used to generate the EMV. In other words,

the length of the EMV is equal to the number of unique values in the domain which is being used to generate the EMV. So, for example, the domain used to generate the EMVs for table 500 in

Fig. 5A is the D\_CustId domain, and from Fig. 3B we see that the D\_CustId domain has 4 EIDs. Therefore the length of the EMVs for that domain will have a length of 4, with 4 corresponding bit slots.

**[0052]** Referring back to Fig. 5B, each of the bit slots corresponds to a different EID, in the order that the EIDs are assigned. So the first slot corresponds to EID=1, the second slot corresponds to EID=2, and so on. To determine the value of the bit in each of the bit slots in an EMV, the group set corresponding to the EMV is checked to see whether it contains the EID value. So for the first bit in EMV 513, it is determined whether the group set (week 1, store 2) includes a CustId=1 at step 507. From table 320 in Fig. 3C we can see that a CustId=1 does not appear in during week 1 at store 2. Therefore the value of the first bit is 0.

**[0053]** In a similar manner, for the second bit in EMV 513, it is determined whether the group set (week 1, store 2) includes a CustId=2 at step 508. Since CustId=2 also does not appear during week 1 at store 2, the value of the second bit is also zero. For the third bit in EMV 513, it is determined whether the group set (week 1, store 2) includes a CustId=3 at step 509. From table 320 we see that CustId=3 appears twice during week 1 at store 2. Therefore the value of the third bit is equal to 1. For the fourth bit in EMV 513, it is determined whether the group set (week 1, store 2) includes a CustId=4 at step 510. Since a CustId=4 also appears during week 1 at store 2, the fourth bit is also set to 1. The resulting EMV 513 captures the distinct customers that visited store 2 during week 1 in a format that allows for additive aggregation.

**[0054]** If a user wanted to determine the number of distinct customers per group set, all they would have to do is sum each of the bits in the EMV for that group set. For example, the EMV for group set (week 1, store 1) is OB1100, so the total number of distinct customers is



$1+1+0+0 = 2$ . The EMV for group set (week 2, store 1) is OB1101 so the total number of distinct customers for that group set is  $1+1+0+1 = 3$ .

**[0055]** Referring back to Fig. 2B, a method for generating EMVs from a non-tokenized data column will now be described. Step 210 is similar to step 202 of Fig. 2A, in that a query is received which specifies one or more group sets and includes a new extension, or other similar command, which instructs the database engine to generate the EMV data structure. At step 211, an entity map vector is generated for each of the group sets, based on the lexical ordering of the unique data values in the domain corresponding to the column of data. Figs. 5C and 5D illustrate this process in greater detail with regard to an EMV generated for the group set (week 201001, store 1002) in the non-tokenized table 300 in Fig. 3A. Table 520 shows the unique data values in the domain CustId, arranged in lexical order. The EMV generation process shown in Fig. 5D uses this lexical ordering to bypass the step of generating EIDs. As shown in empty EMV 525, the length of each of the EMVs generated is equal to the number of unique data values in the domain. In this case, since there are four unique data values in the domain CustId, the length of the EMVs generated for a particular group set will be four. The position of each of the bits in each EMV corresponds to the lexical position of a corresponding unique data value in the lexically ordered list of unique data values 520 and the value of each bit in the EMV indicates the presence or absence of the corresponding unique data value in the group set. So for the first bit in EMV 533, it is determined whether the group set (week 201001, store 1002) includes the CustId which is at the lowest lexical position. This CustId is 165, and from table 300 in Fig. 3A, we see that this CustId does not occur in the group set (week 201001, store 1002), so the bit is equal to zero. Similarly, the second bit of EMV 533 refers to the unique data value with the second lowest lexical position, which is 167, and also does not occur in the group set, resulting in

a zero for the bit. The third bit of EMV 533 refers to the unique data value with the third lowest lexical position, which is 268, and we see from table 300 in Fig. 3A that this value does occur in the group set (week 201001, store 1002), so the third bit is set to 1. Similarly, the value corresponding to the fourth bit, 301, also occurs in the group set, so the fourth bit is set to 1. This results in the same EMV that was generated in Fig. 5B, but which can be generated from a non-tokenized column of data.

**[0056]** Unlike the previous examples of LLA tables which do not utilize the entity map vector, LLA tables with EMVs can be correctly aggregated during a roll-up process. For example, the following definition can be used to generate a higher level table for the number of distinct customers that visited any of the stores from the LLA table, grouped by week:

```
Create Table LLA_TRX_WEEK as
Select Week,
       Count(Distinct CD_CustId) as EMV_CD_CustId,
From LLA_Trx
Group By Week;
```

**[0057]** Referring to Fig. 6A, the process for aggregating the bit vectors across stores and grouping by week is shown. Each of the group sets which must be aggregated is determined for each week in steps 601A, 601B, and 601C. For example, the group sets that must be combined for week 1 are the CustId's corresponding to group sets week 1, store 1 and week 1, store 2. The group sets that must be combined for week 2 are the CustId's corresponding to group sets week 2, store 1, week 2, store 2, and week 2, store 3. The group sets that must be combined for week 3 are the CustId's corresponding to group sets week 3, store 1 and week 3, store 3.

**[0058]** Referring to steps 602A, 602B, and 602C, the EMVs for each of the group sets that must be combined for each week are aggregated using a Boolean OR. So, for example, if



two EMVs are being combined with a Boolean OR and only the first EMV has a “1” in the first position, then the resulting EMV will have a “1” in the first position.

**[0059]** The resulting combined EMVs are shown in steps 603A, 603B, and 603C. Each EMV corresponds to a different week group set. As discussed earlier, the count of distinct customers in each week can be determined by summing all of the bits in each of the EMVs.

**[0060]** Fig. 6B shows the resulting table 600 with the tuple ID 601 for each group set, the week number 602, and the count of distinct customers 603 based on the combined EMVs from Fig. 6A. By referring to the original table 320 in Fig. 3C, we can see the count of distinct customers 603 in table 600 correctly indicates that 4 distinct customers (1,2,3, and 4) went to any of the stores in week 1, 4 distinct customers (1,2,3, and 4) went to any of the stores in week 2, and 2 distinct customers (1 and 2) went to any of the stores in week 3.

**[0061]** Of course, aggregation of non-additive aggregate functions is not limited to counting the number of distinct elements in a group set. If a user wishes to sum a set of distinct values in a particular group set or average a set of distinct values in a group set, that can be accomplished by looking up the unique data values corresponding to each bit of the EMV in the look-up dictionary and adding it to the sum or including it in the average if the bit is “1.” Additionally, the look-up dictionary can be used to retrieve each of the unique data values corresponding to the bits in the entity map vector for any other purpose, such as customized calculations or metrics. In situations where the entity map vectors are generated from non-tokenized data values based on the lexical ordering of the unique data values, that same lexical ordering relationship can be used to retrieve the unique data values corresponding to each bit. For example, if the first bit in an EMV is a “1,” then the unique data value with the lowest lexical position can be included in a sum, average, or other calculation involving the data values in the



set. Similarly, if the second bit in an EMV is a “0,” then the unique data value corresponding to the second lowest lexical position can be left out of the sum, average, or other calculation.

**[0062]** An overview of query processing in a partitioned database and a process for EMV based representation of data in a partitioned database will now be described. Given a three partition database containing a transaction table, a logical view of that table can be presented as a View:

```
Create View V1 as
Select * from Partition1.TRX
Union All
Select * from Partition2.TRX
Union All
Select * from Partition3.TRX
```

**[0063]** Table partitioning enables efficient distributed parallel processing of requests, with the exception of non-additive aggregate functions. For example, if an original query for a transaction table named TRX similar to table 300 in Fig. 3A is “Select Week, Sum(Spent) from TRX group by Week,” but the table is distributed over three partitions, the query can be transformed into the distributed request:

```
Select Week, Sum(P_Spent) from (
    Select Week, Sum(Spent) as P_Spent from Partition1.TRX group by Week
    Union All
    Select Week, Sum(Spent) as P_Spent from Partition2.TRX group by Week
    Union All
    Select Week, Sum(Spent) as P_Spent from Partition3.TRX group by Week
) T
Group By Week
```

**[0064]** For that distributed request, each query segment, like “Select Week, Sum(Spent) as P\_Spent from Partition1.TRX group by Week” can be executed independently and their

intermediate result set could be appended together and finally reprocessed to produce the final query result.

**[0065]** In the case of a non-additive aggregate function request, such as “Select Week, Count(Distinct CustId) from TRX group by Week,” the request can be transformed into the distributed request:

```

Select Week, Count(Distinct CustId) from (
    Select Week, CustId from Partition1.TRX group by Week, CustId
    Union All
    Select Week, CustId from Partition2.TRX group by Week, CustId
    Union All
    Select Week, CustId from Partition3.Trx group by Week, CustId
) T
Group By Week

```

**[0066]** The distributed request in this case does preserve the list of unique CustIds from each partition, but can require a huge amount of data to be exchanged between each of the partitions and subsequently processed. For example, the list of CustIds in the first partition for the first week will have to be compared to the list of CustIds in the second partition for the first week to determine if there is any overlap. As a result, such a distributed request introduces a bottlenecking of the request execution process.

**[0067]** The method to transform a non-additive aggregate function to an additive one can be applied to one or more partitioned databases to reduce the amount of data exchanged between partitions and improve performance. For example, the original query “Select Week, Count(Distinct CustId) from TRX group by Week” can be transformed to:

```

Select Week, Count(Distinct CustId) from
(
    Select Week, Count(Additive_distinct CustId) CustId from Partition1.TRX group by Week
    Union All

```

```

Select Week, Count(Additive_distinct CustId) CustId from Partition2.TRX group by Week
Union All
Select Week, Count(Additive_distinct CustId) CustId from Partition3.Trx group by Week
) T
Group By Week

```

**[0068]** Fig. 7A illustrates the steps for processing this distributed request according to a disclosed embodiment. The first three steps 701 can be performed in parallel for each of the partitions and are similar to the steps for generating an EMV shown in Fig. 2A, applied to each partition. Of course, these steps do not need to be performed in parallel and can be performed in any suitable order. After the first three steps 701 are executed for each of the partitions, the result will be three sets of EMVs, one for each partition, with each digit of the EMVs referring to local partition EIDs. For example, Fig. 8A shows three hypothetical EMV sets 802A, 802B, and 802C, corresponding to the three partitions, along with three hypothetical EID mapping tables 801A, 801B, and 801C, for each of the unique data values in each of the partitions. Since the EIDs in the mapping tables 801A-801C all correspond to different unique data values, and since the number of EIDs in each partition varies, the length of the EMVs and significance of the bits in each of the EMVs in EMV sets 802A-802C varies. As a result, the EMV sets 802A-802C cannot be combined by the union operator in their present state.

**[0069]** Referring back to Fig. 7A, at step 702, a domain merge process is carried out to build a global domain from each of the domain partitions so that the EMV sets can be standardized. Continuing to Fig. 7B, at step 703 global EIDs are generated for each of the unique data values in the global domain. Since the global EIDs correspond to the same unique data values as the partition EIDs, each of the partition EIDs can be converted into global EIDs at step 704 of Fig. 7B. At step 705, the partition specific EMVs are converted into global EMVs, all of which will have the same length and will reference the same unique data values at each bit.



This is accomplished by using a partition-to-global EID mapping table to rebuild each of the EMVs created in the partition to create new ones which map to the global EIDs

**[0070]** For an illustration of this process, Fig. 8B shows a global domain 803 that results from a domain merge of the domain partitions in Fig. 8A and the corresponding global EIDs 804. Fig. 8C shows the partition-to-global EID mapping tables for partition 1 811, partition 2, 812 and partition 3, 813. As discussed earlier, the partition-to-global EID mappings are generated by cross referencing a unique data value with its partition EID and global EID. This mapping table is then used to convert the partition specific EMVs to global EMVs. The resulting global EMVs generated from conversion are shown in Fig. 8D for partition 1 821, partition 2 822, and partition 3 823.

**[0071]** Once the partition specific EMVs have been converted to global EMVs, they can be aggregated as discussed earlier with a Boolean OR. The resulting combined global EMVs will indicate the presence or absence of unique data values in a particular group set across all of the combined partitions. Additionally, the sum of the bits in a combined global EMV will equal the total number of distinct data values in the group set corresponding to that EMV across partitions.

**[0072]** This domain merge process can be utilized for database update operations. In data warehouse environments, the update process often comprises appending new data. By processing the appended data as a new partition to generate an associated partition LLA table and then combining the partition LLA table with existing LLA tables using the techniques discussed above, the new data can merged with existing LLA tables.

**[0073]** For example, the global EMVs for each of the partitions can be stored in a data cube. Fig. 9 shows an exemplary data cube 900 storing the global EIDs according to a disclosed

embodiment. As shown, the dimensions can include an EMV dimension 901, a group set dimension 902, and a partition dimension 903. When a new data update is received, it can be treated as a partition and added to the data cube using the domain merge processes discussed above.

**[0074]** As discussed earlier, EIDs may also be assigned in chronological order of appearance of the data values, rather than lexical order of the data values. By assigning EIDs in this manner, the domain merge step for an update to a data warehouse may be eliminated. Referring to Fig. 10, an existing data warehouse is shown at 1001. In this example, the chronological order in which the data values are processed is used to assign the EIDs for each data value. The table 1002 shows the EIDs for four data values processed in the following sequence: 368, 167, 165, 268. Additionally, table 1003 shows the EMVs for three group sets which contain data values. So, for example, the second EMV “OB1110” corresponds to a group set which contains the data values 368, 167, and 165.

**[0075]** Fig. 10 also shows a data update 1004 for the data warehouse. The data update is similar to the earlier discussed partitions. A local EID mapping table 1005 shows the EIDs for the data values in the update, also assigned in chronological order. Additionally, the EMVs for two group sets are shown at 1006. In this case, two of the data values in the update are new, with one of the data values already existing in the data warehouse.

**[0076]** The updated data warehouse is shown at 1007. Since the EIDs are assigned in chronological order, there is no need to perform a domain merge process. Rather, the two new data values in the data update, 392 and 163, are assigned the next two available EIDs, 5 and 6 in this case. Since there is no remapping of existing EIDs, the existing table of EMVs 1003 can be updated simply by extending the length of each EMV by adding a number of zeroes equal to the



number of new EIDs to the end of each EMV. This results in table 1009 with the updated EMVs. Of course, the zeroes do not necessarily need to be appended to the EMVs, as the absence of a character at a particular position can be interpreted as a zero.

**[0077]** The only remapping that is required for a data update when the EIDs are assigned based on chronological order of the data values is to the EMVs in the actual data update. In this case, the EMVs in update 1006 must be remapped to use the global EID table 1008 rather than the local EID table 1005 to generate the updated EMVs in table 1010. So, for example, since the first EMV in the update 1006 corresponds to a group set that includes the values 392 and 163 (1's at the first and third position of the EMV), and the global EIDs corresponding to those data values are 5 and 6, the first EMV in the updated data table 1010 corresponding to the first EMV in the data update will have 1's at the fifth and sixth positions.

**[0078]** In the scenario where one or more partition-specific EMVs are generated from a non-tokenized column of data in a first data partition and a merge is carried out with one or more other partitions, the first domain partition corresponding to the non-tokenized column of data in the first data partition can be combined with one or more other domain partitions in the set of domain partitions to generate a global domain. This global domain can be used to generate one or more global EMVs corresponding to the one or more partition-specific EMVs based on the lexical ordering of all of the unique data values in the global domain, similar to the process of generating an EMV described with regards to Fig. 2B. The length of each global EMV can be equal to the number of unique data values in the global domain, the position of each bit in a global EMV can correspond to the lexical position of a corresponding unique data value in a lexically ordered list of the unique data values in the global domain, and the value of each bit in a



global EMV can indicate the presence or absence of the corresponding unique data value in a group set.

**[0079]** A domain synchronization process will now be described which can be used to bypass the domain merge process described earlier, and which can make operation of a database in distributed system more efficient. Fig. 11A shows two partitions, 1101 and 1102, in a distributed database. Similar to the distributed database of Fig. 8A, each of the partitions includes a local EID table, with table 1104 for partition 1101 and table 1106 for partition 1102. In this example, the EIDs are shown as being assigned based upon lexical order of the data values, but the EIDs can also be assigned chronologically, as discussed earlier. Additionally, the distributed database includes a global EID table 1103. Global EID table 1103 contains mappings of all of the data values to global EIDs. Once again, these global EIDs are assigned based upon lexical order of data values, but can be assigned based upon chronological order of data values. Although global EID table 1103 is shown as being stored outside of partition 1101 and partition 1102, the global EID table can be stored on one or both of the partitions as well.

**[0080]** Each of the partitions, 1101 and 1102, also includes a local EID-to-global EID mapping table, table 1105 for partition 1101 and table 1107 for partition 1102. These mapping tables can be used to translate local EMVs to global EMVs, as will be discussed later using the examples of EMVs 1108 and 1109.

**[0081]** When an update is received at either partition 1101 or partition 1102, any new unique data values can be sent to the global EID table 1103. These new unique data values can be used to update the global EID table. For example, when the EIDs are assigned by lexical order of the unique data values, the data values in the global EID table 1103 can be re-sorted and the global EIDs can be reassigned to each of the unique data values taking into account the new

unique data values received. Note that if chronologically assigned EIDs are used, this process is greatly simplified, as new unique data values will be assigned the next available EID, and there is no need to reassign EIDs to all of the preexisting unique data values.

**[0082]** After the global EID table 1103 is updated, any new global EIDs can be propagated back to the partitions 1101 and 1102. The new global EIDs can be used to update each of the local EID-to-global EID mapping tables, 1105 and 1107. Once again, chronologically assigned EIDs are more efficient for this process, as the only new global EIDs will be for the new unique data values, whereas with lexically assigned EIDs, all of the EIDs for existing unique data values can potentially be reshuffled.

**[0083]** When a query requires EMVs from more than one partition, the local EID-to-global EID mapping tables, 1105 and 1107, can be used to convert any local EMVs to global EMVs. Referring to Fig. 11B, the conversion process is shown for sample EMV 1108 in partition 1101. As shown in Fig. 11B, each bit in the local EMV is mapped from its local position to its global position in the equivalent global EMV 1110. So, for example, since the local EID-to-global EID mapping table 1105 maps local EID “3” to global EID “5,” the value of the bit at the third position in the local EMV 1108 is assigned to the fifth position of the global EMV 1110. If there is no mapping for a position in the global EMV, such as for the fourth position in global EMV 1110, then that value can be set to zero to reflect the absence of the particular unique data value in the group set, since if that value was present it would be in the local EID-to-global EID mapping table 1105. Similarly, local EID-to-global EID mapping table 1107 can be used to generate global EMV 1111 from local EMV 1109 in partition 1102.

**[0084]** Through this domain synchronization process, the domain merge procedure described earlier can be bypassed, resulting in processing time improvement, and reducing the



number of unique data values that need to be exported at one time. Of course, a single EMV is shown in each partition for the sake of clarity only. In practice, if a local EID and value existed in a particular partition, then some group set and EMV would have to contain that value. Additionally, two partitions are shown for the sake of clarity only, and the database can contain any number of partitions.

**[0085]** One or more of the above-described techniques can be implemented in or involve one or more computer systems. Fig. 12 illustrates a generalized example of a computing environment 1200. The computing environment 1200 is not intended to suggest any limitation as to scope of use or functionality of a described embodiment.

**[0086]** With reference to Fig. 12, the computing environment 1200 includes at least one processing unit 1210 and memory 1220. The processing unit 1210 executes computer-executable instructions and can be a real or a virtual processor. In a multi-processing system, multiple processing units execute computer-executable instructions to increase processing power. The memory 1220 can be volatile memory (e.g., registers, cache, RAM), non-volatile memory (e.g., ROM, EEPROM, flash memory, etc.), or some combination of the two. The memory 1220 can store software 1280 implementing described techniques.

**[0087]** A computing environment can have additional features. For example, the computing environment 1200 includes storage 1240, one or more input devices 1250, one or more output devices 1260, and one or more communication connections 1290. An interconnection mechanism 1270, such as a bus, controller, or network interconnects the components of the computing environment 1200. Typically, operating system software or firmware (not shown) provides an operating environment for other software executing in the



computing environment 1200, and coordinates activities of the components of the computing environment 1200.

**[0088]** The storage 1240 can be removable or non-removable, and includes magnetic disks, magnetic tapes or cassettes, CD-ROMs, CD-RWs, DVDs, or any other medium which can be used to store information and which can be accessed within the computing environment 1200. The storage 1240 can store instructions for the software 1280.

**[0089]** The input device(s) 1250 can be a touch input device such as a keyboard, mouse, pen, trackball, touch screen, or game controller, a voice input device, a scanning device, a digital camera, remote control, or another device that provides input to the computing environment 1200. The output device(s) 1260 can be a display, television, monitor, printer, speaker, or another device that provides output from the computing environment 1200.

**[0090]** The communication connection(s) 1290 enable communication over a communication medium to another computing entity. The communication medium conveys information such as computer-executable instructions, audio or video information, or other data in a modulated data signal. A modulated data signal is a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media include wired or wireless techniques implemented with an electrical, optical, RF, infrared, acoustic, or other carrier.

**[0091]** Implementations can be described in the general context of computer-readable media. Computer-readable media are any available media that can be accessed within a computing environment. By way of example, and not limitation, within the computing environment 1200, computer-readable media include memory 1220, storage 1240, communication media, and combinations of any of the above.

**[0092]** Of course, Fig. 12 illustrates computing environment 1200, display device 1260, and input device 1250 as separate devices for ease of identification only. Computing environment 1200, display device 1260, and input device 1250 can be separate devices (e.g., a personal computer connected by wires to a monitor and mouse), can be integrated in a single device (e.g., a mobile device with a touch-display, such as a smartphone or a tablet), or any combination of devices (e.g., a computing device operatively coupled to a touch-screen display device, a plurality of computing devices attached to a single display device and input device, etc.). Computing environment 1200 can be a set-top box, personal computer, or one or more servers, for example a farm of networked servers, a clustered server environment, or a cloud network of computing devices.

**[0093]** Having described and illustrated the principles of our invention with reference to the described embodiment, it will be recognized that the described embodiment can be modified in arrangement and detail without departing from such principles. It should be understood that the programs, processes, or methods described herein are not related or limited to any particular type of computing environment, unless indicated otherwise. Various types of general purpose or specialized computing environments can be used with or perform operations in accordance with the teachings described herein. Elements of the described embodiment shown in software can be implemented in hardware and vice versa.

**[0094]** In view of the many possible embodiments to which the principles of our invention can be applied, we claim as our invention all such embodiments as can come within the scope of the following claims and equivalents thereto.

What is Claimed is:

1. A method for efficiently performing operations on distinct data values by one or more computing devices, the method comprising:

storing, by at least one of the one or more computing devices, a tokenized column of data in a table, the tokenized column of data created by mapping each unique data value in a domain of a database to an entity ID, and replacing each of a plurality of data values in a column of data corresponding to the domain with the corresponding entity ID to generate the column of tokenized data containing a plurality of entity IDs;

receiving, by at least one of the one or more computing devices, a query directed to the column of data, the query defining one or more group sets for grouping data retrieved in response to the query, wherein each group set in the one or more group sets corresponds to a unique group of one or more values associated with one or more other domains of the database; and

generating, by at least one of the one or more computing devices, an entity map vector for each group set in the one or more group sets by identifying any entity IDs in the tokenized column of data which are present in any rows of the table which include the unique group of one or more values corresponding to the group set, wherein the length of each entity map vector is equal to the total number of entity IDs in the domain and the value of each bit in each entity map vector indicates the presence or absence of a different entity ID in the corresponding group set.

2. The method of claim 1, wherein the mappings between each of the entity IDs and the unique data values are stored in a look-up dictionary.



3. The method of claim 1, wherein a value of “1” for a bit in an entity map vector indicates the presence of an entity ID in the group set, and a value of “0” for a bit in an entity map vector indicates the absence of the entity ID in the group set.
4. The method of claim 1, wherein the number of distinct data values in any group set is equal to the sum of the values of each of the bits of the entity map vector corresponding to that group set.
5. The method of claim 1, wherein the one or more group sets comprise at least two group sets, the method further comprising:
  - receiving, by at least one of the one or more computing devices, a command to aggregate the data from two or more of the group sets into a combined group set; and
  - performing, by at least one of the one or more computing devices, a Boolean OR operation on the entity map vectors corresponding to the two or more group sets to generate a combined entity map vector.
6. The method of claim 5, wherein the number of distinct data values in the combined group set is equal to the sum of the values of each of the bits of the combined entity map vector.
7. The method of claim 1, wherein the table is partitioned and distributed over a plurality of computing nodes, and wherein the column of data is a first column partition in a set of column partitions, the domain is a first domain partition in a set of domain partitions, the one or more

entity IDs are one or more first partition entity IDs, and the entity map vectors are one or more first partition entity map vectors, the method further comprising:

combining, by at least one of the one or more computing devices, the first domain partition with one or more other domain partitions in the set of domain partitions to generate a global domain;

tokenizing, by at least one of the one or more computing devices, one or more of the unique data values in the global domain by mapping each of the one or more unique data values to a global entity ID to generate one or more global entity IDs;

mapping, by at least one of the one or more computing devices, each of the one or more first partition entity IDs to a corresponding global entity ID in the one or more global entity IDs by cross-referencing a common unique data value; and

translating, by at least one of the one or more computing devices, each of the one or more first partition entity map vectors to one or more global entity map vectors based on the mapping of one or more first partition entity IDs to one or more global entity IDs.

8. The method of claim 7, further comprising:

performing, by at least one of the one or more computing devices, a Boolean OR operation on the one or more global entity map vectors corresponding to the first partition and one or more second global entity map vectors corresponding to a second partition to generate one or more combined global entity map vectors.

9. The method of claim 8, wherein the sum of the values of each of the bits of each of the combined global entity map vectors is the total number of distinct elements in a particular group

set corresponding to that global entity map vector across the first partition and the second partition.

10. The method of claim 7, wherein the global entity map vectors for each partition are stored in a global data cube, the global data cube including an entity map vector dimension, a group set dimension, and a partition dimension.

11. The method of claim 1, wherein the number of group sets is determined, at least in part, by the number of distinct data values in a second column of data in the table.

12. The method of claim 1, wherein the group sets are defined in reference to one or more second columns of data in the table.

13. The method of claim 1, wherein each entity ID is mapped to each unique data value in chronological order of appearance of the unique data values, such that the first entity ID is mapped to the first unique data value encountered.

14. The method of claim 13, wherein the one or more entity IDs are one or more global entity IDs and the one or more the entity map vectors are one or more preexisting global entity map vectors, the method further comprising:

receiving, by at least one of the one or more computing devices, an update comprising one or more new entity map vectors corresponding to one or more new entity IDs, wherein each of the new entity IDs corresponds to a unique data value;



assigning, by at least one of the one or more computing devices, a new global entity ID to each unique data value that is not already represented in the domain; and

generating, by at least one of the one or more computing devices, one or more global entity map vectors corresponding to the one or more new entity map vectors, wherein the length of each global entity map vector is equal to the total number of global entity IDs, and the value of each bit in each global entity map vector indicates the presence or absence of a different global entity ID in a corresponding group set.

15. The method of claim 14, further comprising:

updating, by at least one of the one or more computing devices, the pre-existing global entity map vectors by appending a number of zeroes to the end of each of the pre-existing global entity map vectors, wherein the number of zeroes is equal to the number of unique data values in the update that were not already represented in the domain.

16. The method of claim 1, wherein each entity ID is mapped to each unique data value in the lexical order of the unique data values, such that the first entity ID is mapped to the unique data value that is first in the lexical order of the unique data values.

17. The method of claim 1, wherein the numerical location of each bit in an entity map vector corresponds to the entity ID whose presence or absence in the group set the bit represents.

18. The method of claim 1, wherein storing a tokenized column of data occurs after receiving a query.

19. The method of claim 1, wherein storing a tokenized column of data occurs before receiving a query.
20. The method of claim 1, wherein the query includes a sum command pertaining to each of the one or more group sets and the method further comprises:
- for each of the one or more group sets, summing, by at least one of the one or more computing devices, the unique data values corresponding to the entity IDs which are indicated as present in the entity map vector corresponding to that group set.
21. The method of claim 1, wherein the query includes an average command pertaining to each of the one or more group sets and the method further comprises:
- for each of the one or more group sets, computing, by at least one of the one or more computing devices, the average of the unique data values corresponding to the entity IDs which are indicated as present in the entity map vector corresponding to that group set.
22. An apparatus for efficiently performing operations on distinct data values, the apparatus comprising:
- one or more processors; and
  - one or more memories operatively coupled to at least one of the one or more processors and having instructions stored thereon that, when executed by at least one of the one or more processors, cause at least one of the one or more processors to:

store a tokenized column of data in a table, the tokenized column of data created by mapping each unique data value in a domain of a database to an entity ID, and replacing each of a plurality of data values in a column of data corresponding to the domain with the corresponding entity ID to generate the column of tokenized data containing a plurality of entity IDs;

receive a query directed to the column of data, the query defining one or more group sets for grouping data retrieved in response to the query, wherein each group set in the one or more group sets corresponds to a unique group of one or more values associated with one or more other domains of the database; and

generate an entity map vector for each group set in the one or more group sets by identifying any entity IDs in the tokenized column of data which are present in any rows of the table which include the unique group of one or more values corresponding to the group set, wherein the length of each entity map vector is equal to the total number of entity IDs in the domain and the value of each bit in each entity map vector indicates the presence or absence of a different entity ID in the corresponding group set.

23. The apparatus of claim 22, wherein the mappings between each of the entity IDs and the unique data values are stored in a look-up dictionary.

24. The apparatus of claim 22, wherein a value of “1” for a bit in an entity map vector indicates the presence of an entity ID in the group set, and a value of “0” for a bit in an entity map vector indicates the absence of the entity ID in the group set.



25. The apparatus of claim 22, wherein the number of distinct data values in any group set is equal to the sum of the values of each of the bits of the entity map vector corresponding to that group set.

26. The apparatus of claim 22, wherein the one or more group sets comprise at least two group sets, and the one or more memories have further instructions stored thereon, that, when executed by at least one of the one or more processors, cause at least one of the one or more processors to:

receive a command to aggregate the data from two or more of the group sets into a combined group set; and

perform a Boolean OR operation on the entity map vectors corresponding to the two or more group sets to generate a combined entity map vector.

27. The apparatus of claim 26, wherein the number of distinct data values in the combined group set is equal to the sum of the values of each of the bits of the combined entity map vector.

28. The apparatus of claim 22, wherein the table is partitioned and distributed over a plurality of computing nodes, and wherein the column of data is a first column partition in a set of column partitions, the domain is a first domain partition in a set of domain partitions, the one or more entity IDs are one or more first partition entity IDs, and the entity map vectors are one or more first partition entity map vectors, wherein the one or more memories have further instructions stored thereon, that, when executed by at least one of the one or more processors, cause at least one of the one or more processors to:

combine the first domain partition with one or more other domain partitions in the set of domain partitions to generate a global domain;

tokenize one or more of the unique data values in the global domain by mapping each of the one or more unique data values to a global entity ID to generate one or more global entity IDs;

map each of the one or more first partition entity IDs to a corresponding global entity ID in the one or more global entity IDs by cross-referencing a common unique data value; and

translate each of the one or more first partition entity map vectors to one or more global entity map vectors based on the mapping of one or more first partition entity IDs to one or more global entity IDs.

29. The apparatus of claim 28, wherein the one or more memories have further instructions stored thereon, that, when executed by at least one of the one or more processors, cause at least one of the one or more processors to:

perform a Boolean OR operation on the one or more global entity map vectors corresponding to the first partition and one or more second global entity map vectors corresponding to a second partition to generate one or more combined global entity map vectors.

30. The apparatus of claim 29, wherein the sum of the values of each of the bits of each of the combined global entity map vectors is the total number of distinct elements in a particular group set corresponding to that global entity map vector across the first partition and the second partition.

31. The apparatus of claim 28, wherein the global entity map vectors for each partition are stored in a global data cube, the global data cube including an entity map vector dimension, a group set dimension, and a partition dimension.
32. The apparatus of claim 22, wherein the number of group sets is determined, at least in part, by the number of distinct data values in a second column of data in the table.
33. The apparatus of claim 22, wherein the group sets are defined in reference to one or more second columns of data in the table.
34. The apparatus of claim 22, wherein each entity ID is mapped to each unique data value in chronological order of appearance of the unique data values, such that the first entity ID is mapped to the first unique data value encountered.
35. The apparatus of claim 34, wherein the one or more entity IDs are one or more global entity IDs, the one or more the entity map vectors are one or more preexisting global entity map vectors, and wherein the one or more memories have further instructions stored thereon, that, when executed by at least one of the one or more processors, cause at least one of the one or more processors to:
- receive an update comprising one or more new entity map vectors corresponding to one or more new entity IDs, wherein each of the new entity IDs corresponds to a unique data value;
  - assign a new global entity ID to each unique data value that is not already represented in the domain; and



generate one or more global entity map vectors corresponding to the one or more new entity map vectors, wherein the length of each global entity map vector is equal to the total number of global entity IDs, and the value of each bit in each global entity map vector indicates the presence or absence of a different global entity ID in a corresponding group set.

36. The apparatus of claim 35, wherein the one or more memories have further instructions stored thereon, that, when executed by at least one of the one or more processors, cause at least one of the one or more processors to:

update the pre-existing global entity map vectors by appending a number of zeroes to the end of each of the pre-existing global entity map vectors, wherein the number of zeroes is equal to the number of unique data values in the update that were not already represented in the domain.

37. The apparatus of claim 22, wherein each entity ID is mapped to each unique data value in the lexical order of the unique data values, such that the first entity ID is mapped to the unique data value that is first in the lexical order of the unique data values.

38. The apparatus of claim 22, wherein the numerical location of each bit in the entity map vector corresponds to an entity ID whose presence or absence in the group set the bit represents.

39. The apparatus of claim 22, wherein the tokenized column of data is stored after receiving a query.

40. The apparatus of claim 22, wherein the tokenized column of data is stored before receiving a query.

41. The apparatus of claim 22, wherein the query includes a sum command pertaining to each of the one or more group sets and wherein the one or more memories have further instructions stored thereon, that, when executed by at least one of the one or more processors, cause at least one of the one or more processors to:

for each of the one or more group sets, sum the unique data values corresponding to the entity IDs which are indicated as present in the entity map vector corresponding to that group set.

42. The apparatus of claim 22, wherein the query includes an average command pertaining to each of the one or more group sets and wherein the one or more memories have further instructions stored thereon, that, when executed by at least one of the one or more processors, cause at least one of the one or more processors to:

for each of the one or more group sets, compute the average of the unique data values corresponding to the entity IDs which are indicated as present in the entity map vector corresponding to that group set.

43. At least one non-transitory computer-readable medium storing computer-readable instructions that, when executed by one or more computing devices, cause at least one of the one or more computing devices to:

store a tokenized column of data in a table, the tokenized column of data created by mapping each unique data value in a domain of a database to an entity ID, and replacing each of

a plurality of data values in a column of data corresponding to the domain with the corresponding entity ID to generate the column of tokenized data containing a plurality of entity IDs;

receive a query directed to the column of data, the query defining one or more group sets for grouping data retrieved in response to the query, wherein each group set in the one or more group sets corresponds to a unique group of one or more values associated with one or more other domains of the database; and

generate an entity map vector for each group set in the one or more group sets by identifying any entity IDs in the tokenized column of data which are present in any rows of the table which include the unique group of one or more values corresponding to the group set, wherein the length of each entity map vector is equal to the total number of entity IDs in the domain and the value of each bit in each entity map vector indicates the presence or absence of a different entity ID in the corresponding group set.

44. The at least one non-transitory computer-readable medium of claim 43, wherein the mappings between each of the entity IDs and the unique data values are stored in a look-up dictionary.

45. The at least one non-transitory computer-readable medium of claim 43, wherein a value of “1” for a bit in an entity map vector indicates the presence of an entity ID in the group set, and a value of “0” for a bit in an entity map vector indicates the absence of the entity ID in the group set.



46. The at least one non-transitory computer-readable medium of claim 43, wherein the number of distinct data values in any group set is equal to the sum of the values of each of the bits of the bit map vector corresponding to that group set.

47. The at least one non-transitory computer-readable medium of claim 43, wherein the one or more group sets comprise at least two group sets, the at least one non-transitory computer-readable medium further comprising additional instructions that, when executed by one or more computing devices, cause at least one of the one or more computing devices to:

receive a command to aggregate the data from two or more of the group sets into a combined group set; and

perform a Boolean OR operation on the entity map vectors corresponding to the two or more group sets to generate a combined entity map vector.

48. The at least one non-transitory computer-readable medium of claim 47, wherein the number of distinct data values in the combined group set is equal to the sum of the values of each of the bits of the combined entity map vector.

49. The at least one non-transitory computer-readable medium of claim 43, wherein the table is partitioned and distributed over a plurality of computing nodes, and wherein the column of data is a first column partition in a set of column partitions, the domain is a first domain partition in a set of domain partitions, the one or more entity IDs are one or more first partition entity IDs, and the entity map vectors are one or more first partition entity map vectors, the at least one non-transitory computer-readable medium further comprising additional instructions that, when

executed by one or more computing devices, cause at least one of the one or more computing devices to:

combine the first domain partition with one or more other domain partitions in the set of domain partitions to generate a global domain;

tokenize one or more of the unique data values in the global domain by mapping each of the one or more unique data values to a global entity ID to generate one or more global entity IDs;

map each of the one or more first partition entity IDs to a corresponding global entity ID in the one or more global entity IDs by cross-referencing a common unique data value; and

translate each of the one or more first partition entity map vectors to one or more global entity map vectors based on the mapping of one or more first partition entity IDs to one or more global entity IDs.

50. The at least one non-transitory computer-readable medium of claim 49, the at least one non-transitory computer-readable medium further comprising additional instructions that, when executed by one or more computing devices, cause at least one of the one or more computing devices to:

perform a Boolean OR operation on the one or more global entity map vectors corresponding to the first partition and one or more second global entity map vectors corresponding to a second partition to generate one or more combined global entity map vectors.

51. The at least one non-transitory computer-readable medium of claim 50, wherein the sum of the values of each of the bits of each of the combined global entity map vectors is the total

number of distinct elements in a particular group set corresponding to that global entity map vector across the first partition and the second partition.

52. The at least one non-transitory computer-readable medium of claim 49, wherein the global entity map vectors for each partition are stored in a global data cube, the global data cube including an entity map vector dimension, a group set dimension, and a partition dimension.

53. The at least one non-transitory computer-readable medium of claim 43, wherein the number of group sets is determined, at least in part, by the number of distinct data values in a second column of data in the table.

54. The at least one non-transitory computer-readable medium of claim 43, wherein the group sets are defined in reference to one or more second columns of data in the table.

55. The at least one non-transitory computer-readable medium of claim 43, wherein each entity ID is mapped to each unique data value in chronological order of appearance of the unique data values, such that the first entity ID is mapped to the first unique data value encountered.

56. The at least one non-transitory computer-readable medium of claim 55, wherein the one or more entity IDs are one or more global entity IDs, the one or more the entity map vectors are one or more preexisting global entity map vectors, and the at least one non-transitory computer-readable medium further comprises additional instructions that, when executed by one or more computing devices, cause at least one of the one or more computing devices to:



receive an update comprising one or more new entity map vectors corresponding to one or more new entity IDs, wherein each of the new entity IDs corresponds to a unique data value;

assign a new global entity ID to each unique data value that is not already represented in the domain; and

generate one or more global entity map vectors corresponding to the one or more new entity map vectors, wherein the length of each global entity map vector is equal to the total number of global entity IDs, and the value of each bit in each global entity map vector indicates the presence or absence of a different global entity ID in a corresponding group set.

57. The at least one non-transitory computer-readable medium of claim 56, the at least one non-transitory computer-readable medium further comprising additional instructions that, when executed by one or more computing devices, cause at least one of the one or more computing devices to:

update the pre-existing global entity map vectors by appending a number of zeroes to the end of each of the pre-existing global entity map vectors, wherein the number of zeroes is equal to the number of unique data values in the update that were not already represented in the domain.

58. The at least one non-transitory computer-readable medium of claim 43, wherein each entity ID is mapped to each unique data value in the lexical order of the unique data values, such that the first entity ID is mapped to the unique data value that is first in the lexical order of the unique data values.

59. The at least one non-transitory computer-readable medium of claim 43, wherein the numerical location of each bit in an entity map vector corresponds to the entity ID whose presence or absence in the group set the bit represents.

60. The at least one non-transitory computer-readable medium of claim 43, wherein the tokenized column of data is stored after receiving a query.

61. The at least one non-transitory computer-readable medium of claim 43, wherein the tokenized column of data is stored before receiving a query.

62. The at least one non-transitory computer-readable medium of claim 43, wherein the query includes a sum command pertaining to each of the one or more group sets, and the at least one non-transitory computer-readable medium further comprises additional instructions that, when executed by one or more computing devices, cause at least one of the one or more computing devices to:

for each of the one or more group sets, sum the unique data values corresponding to the entity IDs which are indicated as present in the entity map vector corresponding to that group set.

63. The at least one non-transitory computer-readable medium of claim 43, wherein the query includes an average command pertaining to each of the one or more group sets, and the at least one non-transitory computer-readable medium further comprises additional instructions that, when executed by one or more computing devices, cause at least one of the one or more computing devices to:

for each of the one or more group sets, compute the average of the unique data values corresponding to the entity IDs which are indicated as present in the entity map vector corresponding to that group set.



Customer Name	Store ID	Amount Spent
Bill	1	10
Bill	2	5
George	1	11
George	2	11
Max	2	21
Larry	1	31
Larry	2	18

Fig. 1A

Store ID	Count of Visits
1	3
2	4

Fig. 1B

Store ID	Amount Spent
1	52
2	55

Fig. 1C

Store ID	Average Amount Spent	Count of Visits
1	17 1/3	3
2	13 3/4	4

Fig. 1D

Store ID	Minimum Spent
1	10
2	5

Fig. 1E

Store ID	Maximum Spent
1	31
2	21

Fig. 1F

Store ID	Number of Distinct Customers
1	3
2	4

Fig. 1G

3/18

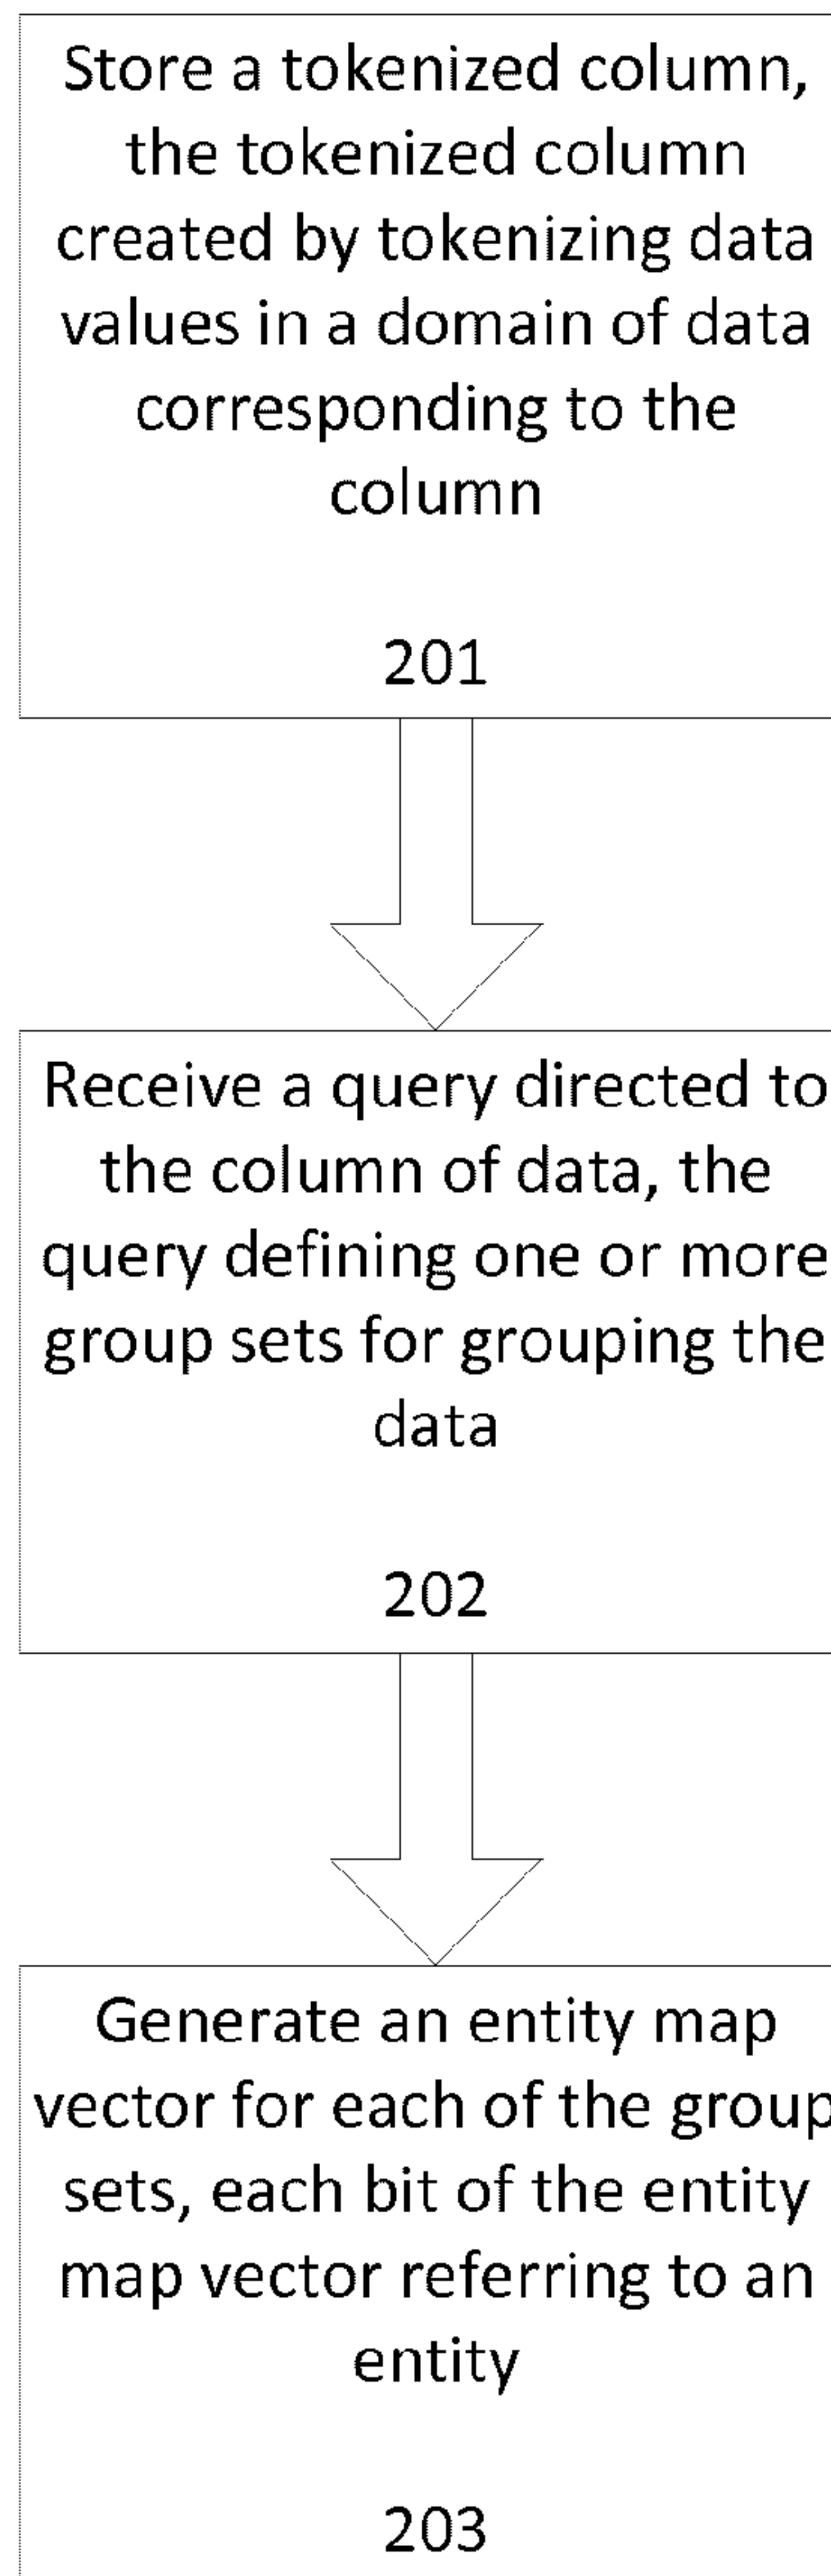


Fig. 2A



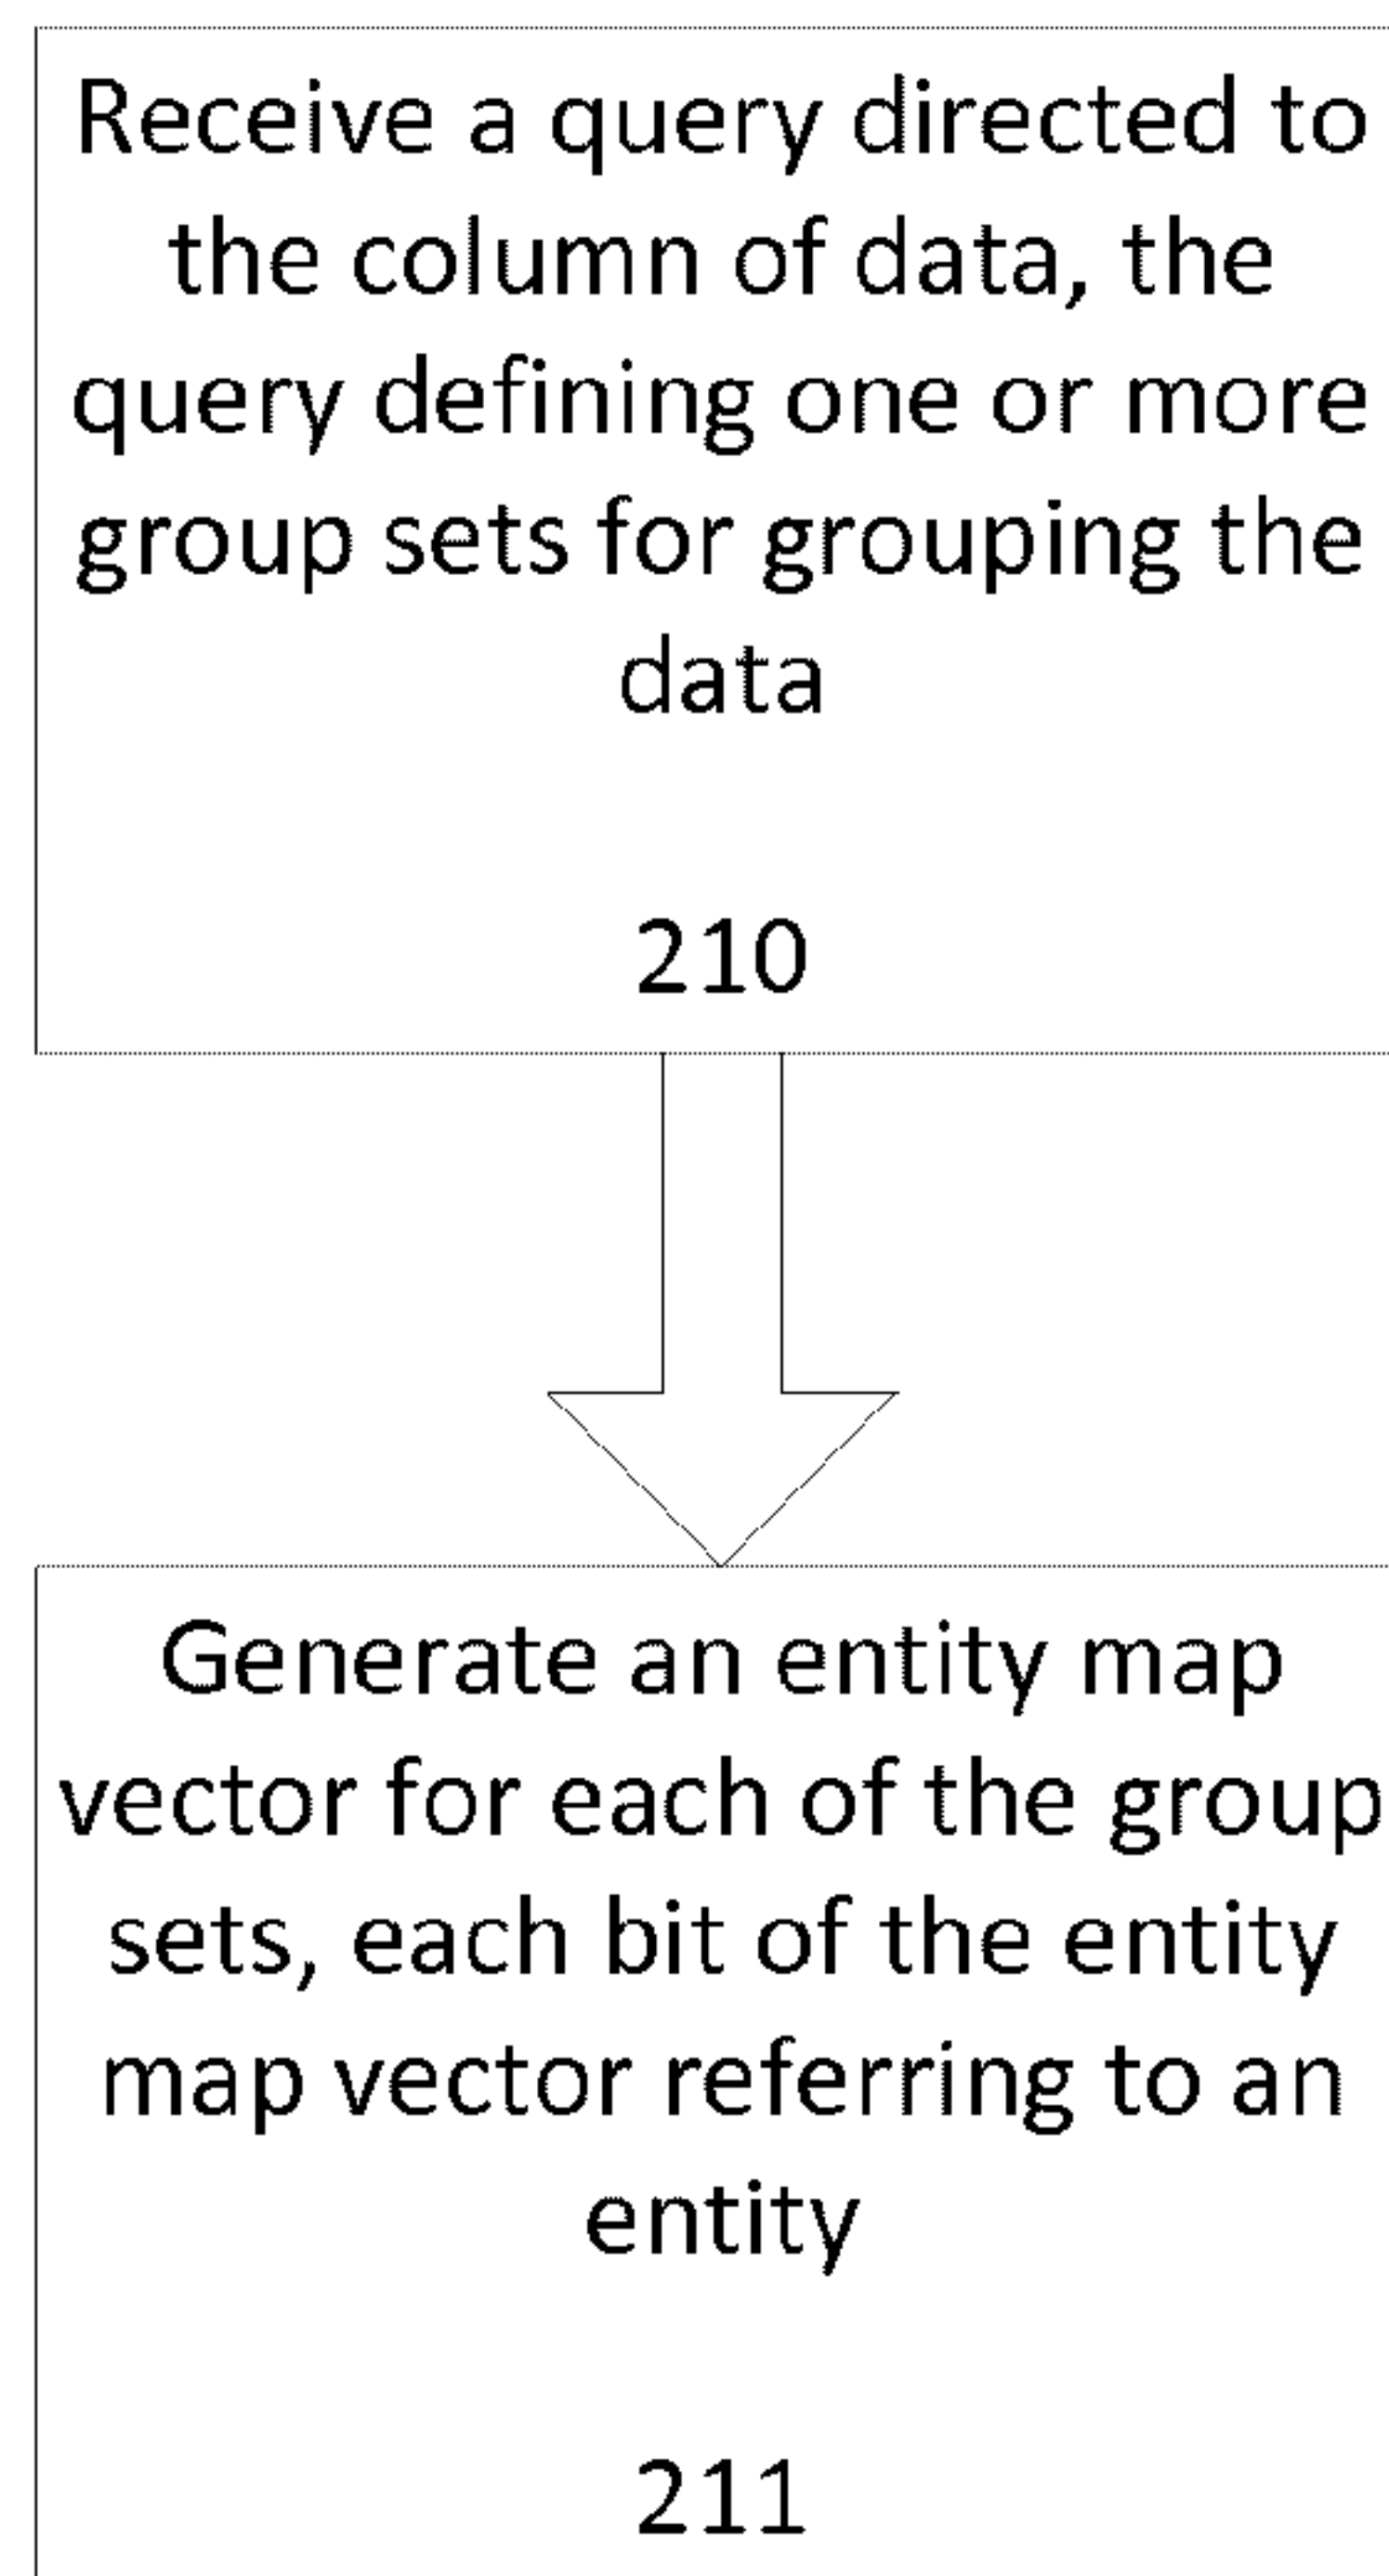


Fig. 2B

301

302

303

304

305

306

307

TupleId	Week	StoreId	CustId	ProductId	Qty	Spent
1	201001	1001	167	10	20	200
2	201001	1001	167	12	10	20
3	201001	1001	167	11	10	40
4	201001	1001	167	20	30	30
5	201001	1001	165	11	20	50
6	201001	1001	165	12	20	40
7	201001	1002	268	11	5	30
8	201001	1002	268	20	15	15
9	201001	1002	301	11	10	44
10	201002	1001	167	10	20	200
11	201002	1001	167	12	10	20
12	201002	1001	165	11	10	40
13	201002	1001	301	20	30	30
14	201002	1002	268	11	20	50
15	201002	1002	268	12	20	40
16	201002	1003	167	11	5	30
17	201002	1003	167	20	15	15
18	201003	1001	165	11	10	44
19	201003	1001	165	11	10	40
20	201003	1003	167	20	30	30

300

Fig. 3A

312

313

314

315

316

317

D_Week		D_StoreId		D_CustId		D_ProductId		D_Qty		D_Spent	
Eid	Value	Eid	Value	Eid	Value	Eid	Value	Eid	Value	Eid	Value
1	201001	1	1001	1	165	1	10	1	5	1	15
2	201002	2	1002	2	167	2	11	2	10	2	20
3	201003	3	1003	3	268	3	12	3	15	3	30
				4	301	4	20	4	20	4	40
								5	30	5	50
										6	200

310

Fig. 3B




Fig. 3C



<u>TupleId</u>	<u>Week</u>	<u>StoreId</u>	<u>CD_CustId</u>
1	1	1	2
2	1	2	2
3	2	1	3
4	2	2	3
5	2	3	3
6	3	1	4
7	3	3	4

Fig. 4A

400

<u>TupleId</u>	<u>Week</u>	<u>CD_CD_CustId</u>
1	1	1
2	2	2
3	3	1

Fig. 4B

410

<u>TupleId</u>	<u>Week</u>	<u>Sum_CD_CustId</u>
1	1	4
2	2	5
3	3	2

Fig. 4C

420

TupleId	Week	StoreId	CD_CustId
1	1	1	031100
2	1	2	030011
3	2	1	031101
4	2	2	030010
5	2	3	030100
6	3	1	031000
7	3	3	030100

Fig. 5A

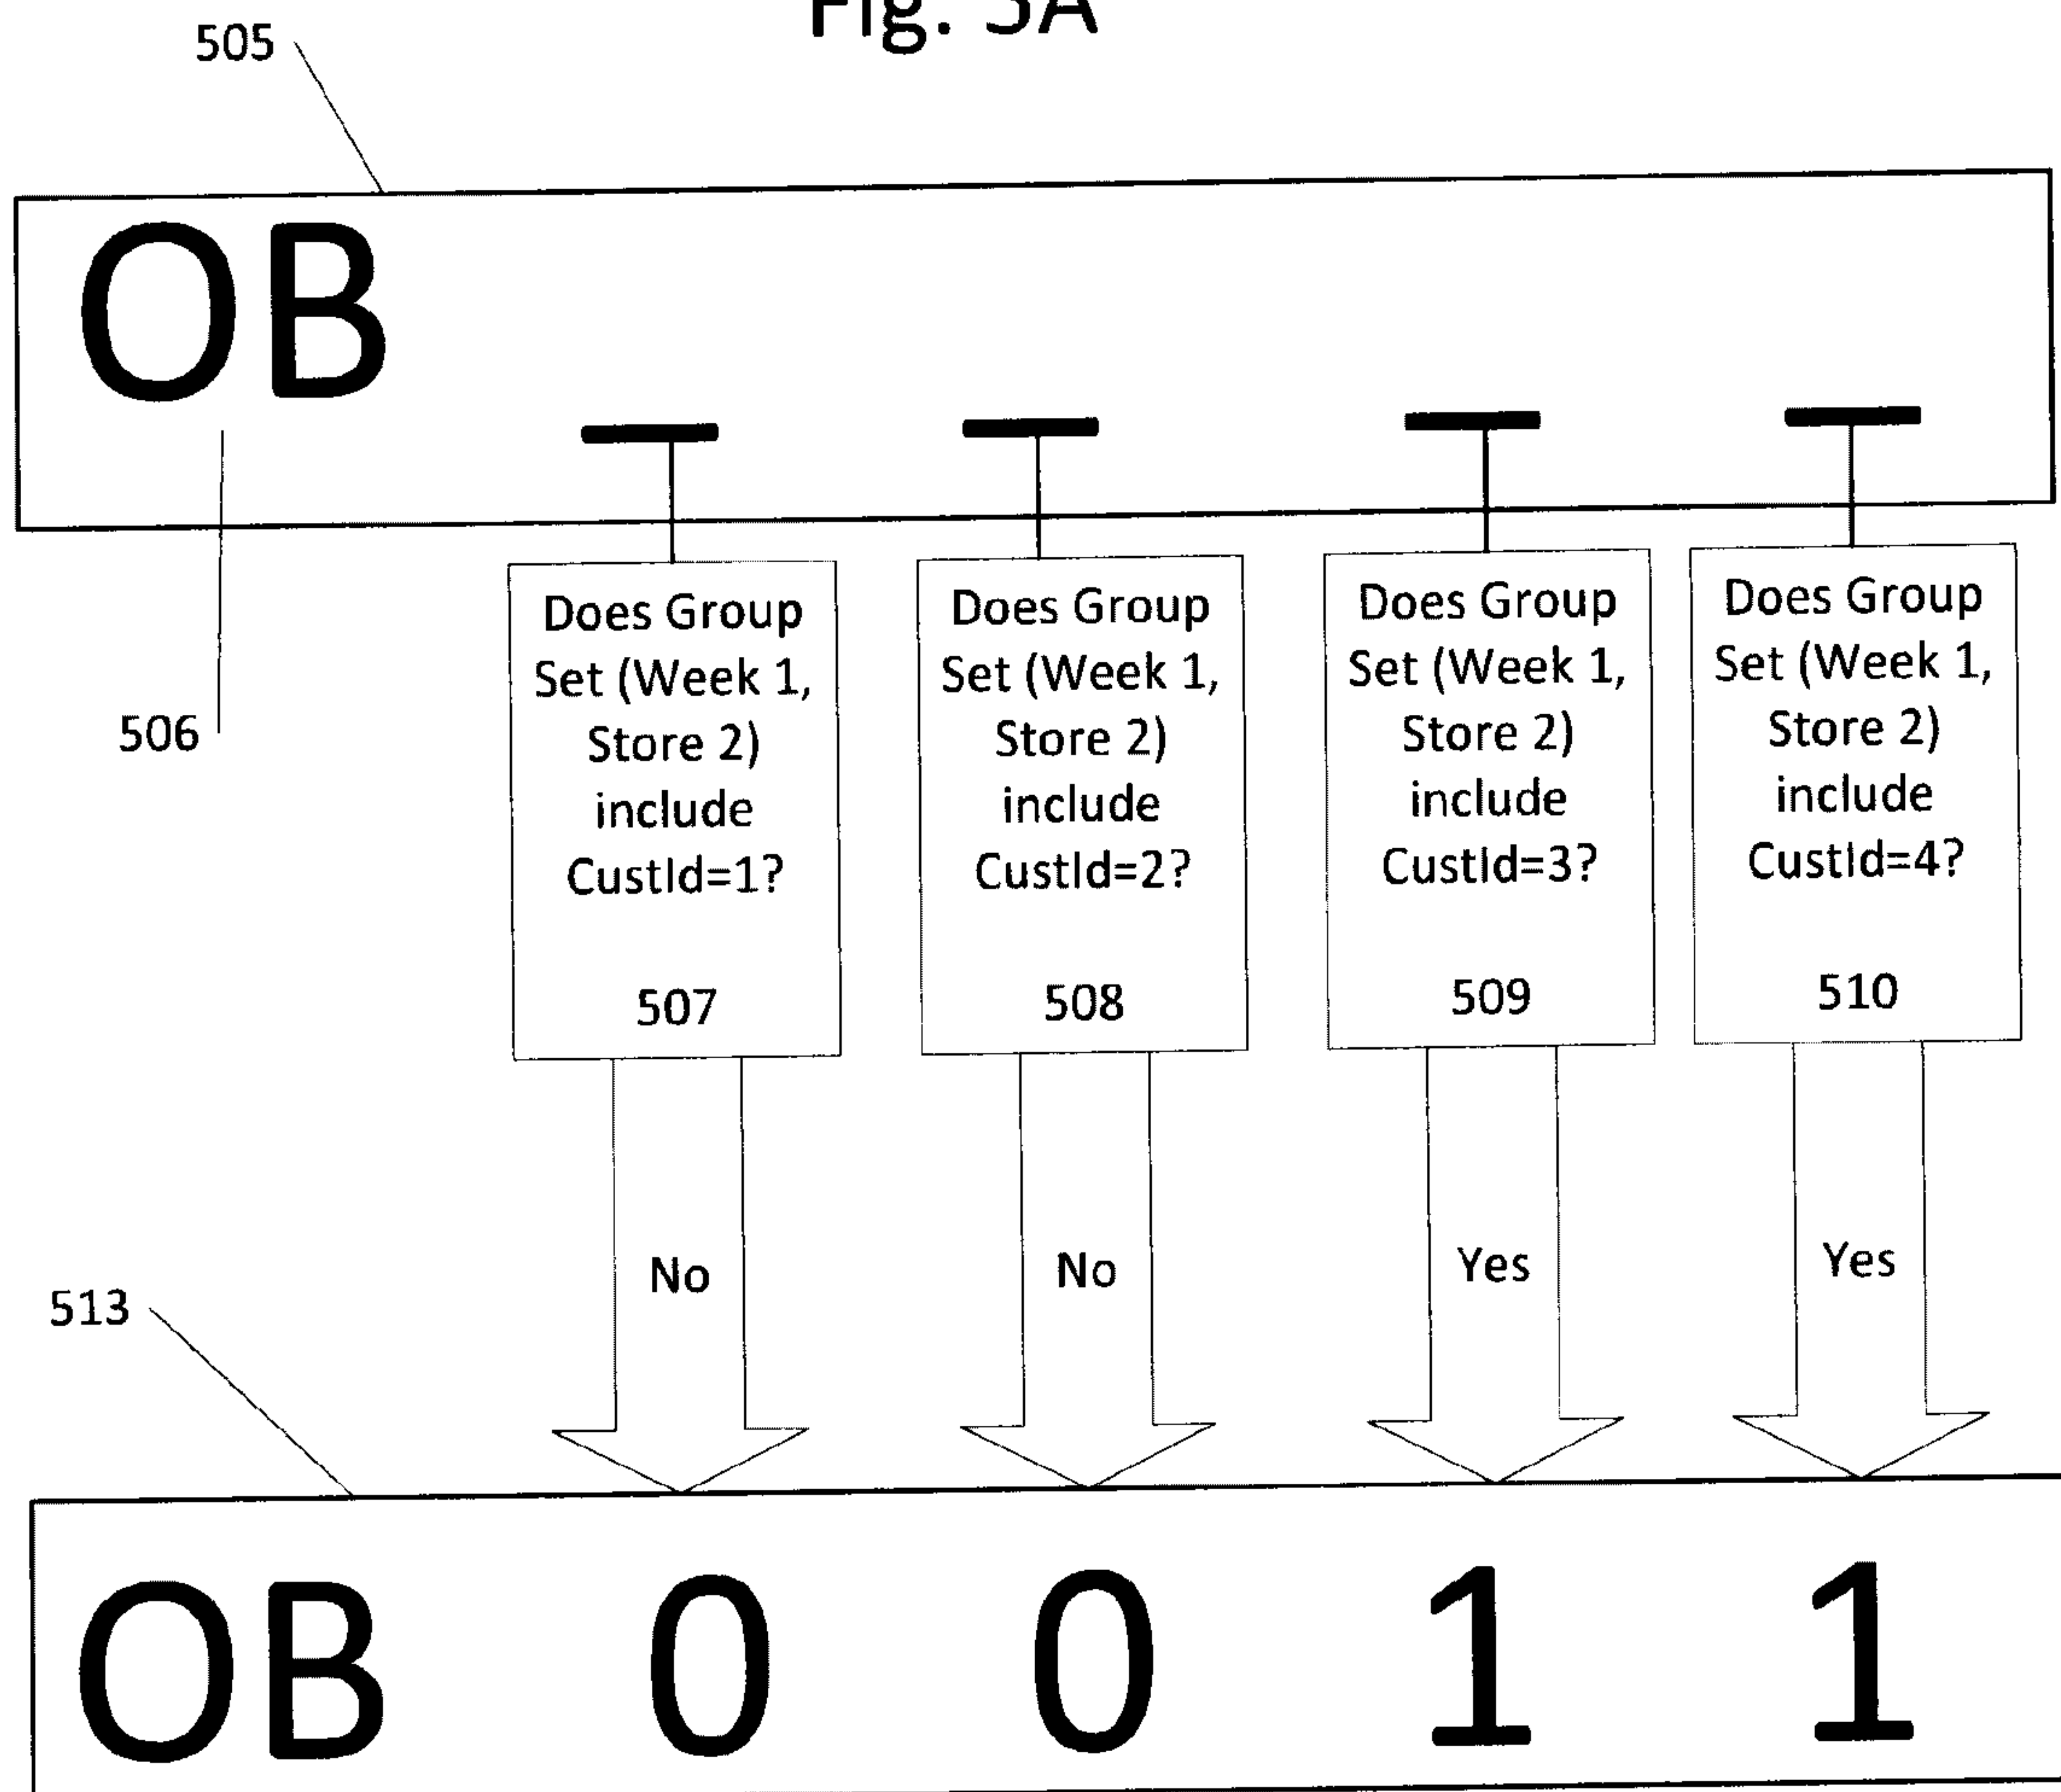


Fig. 5B

520

Lexically Ordered List of Unique Data Values in D_CustId	
165	
167	
268	
301	

Fig. 5C

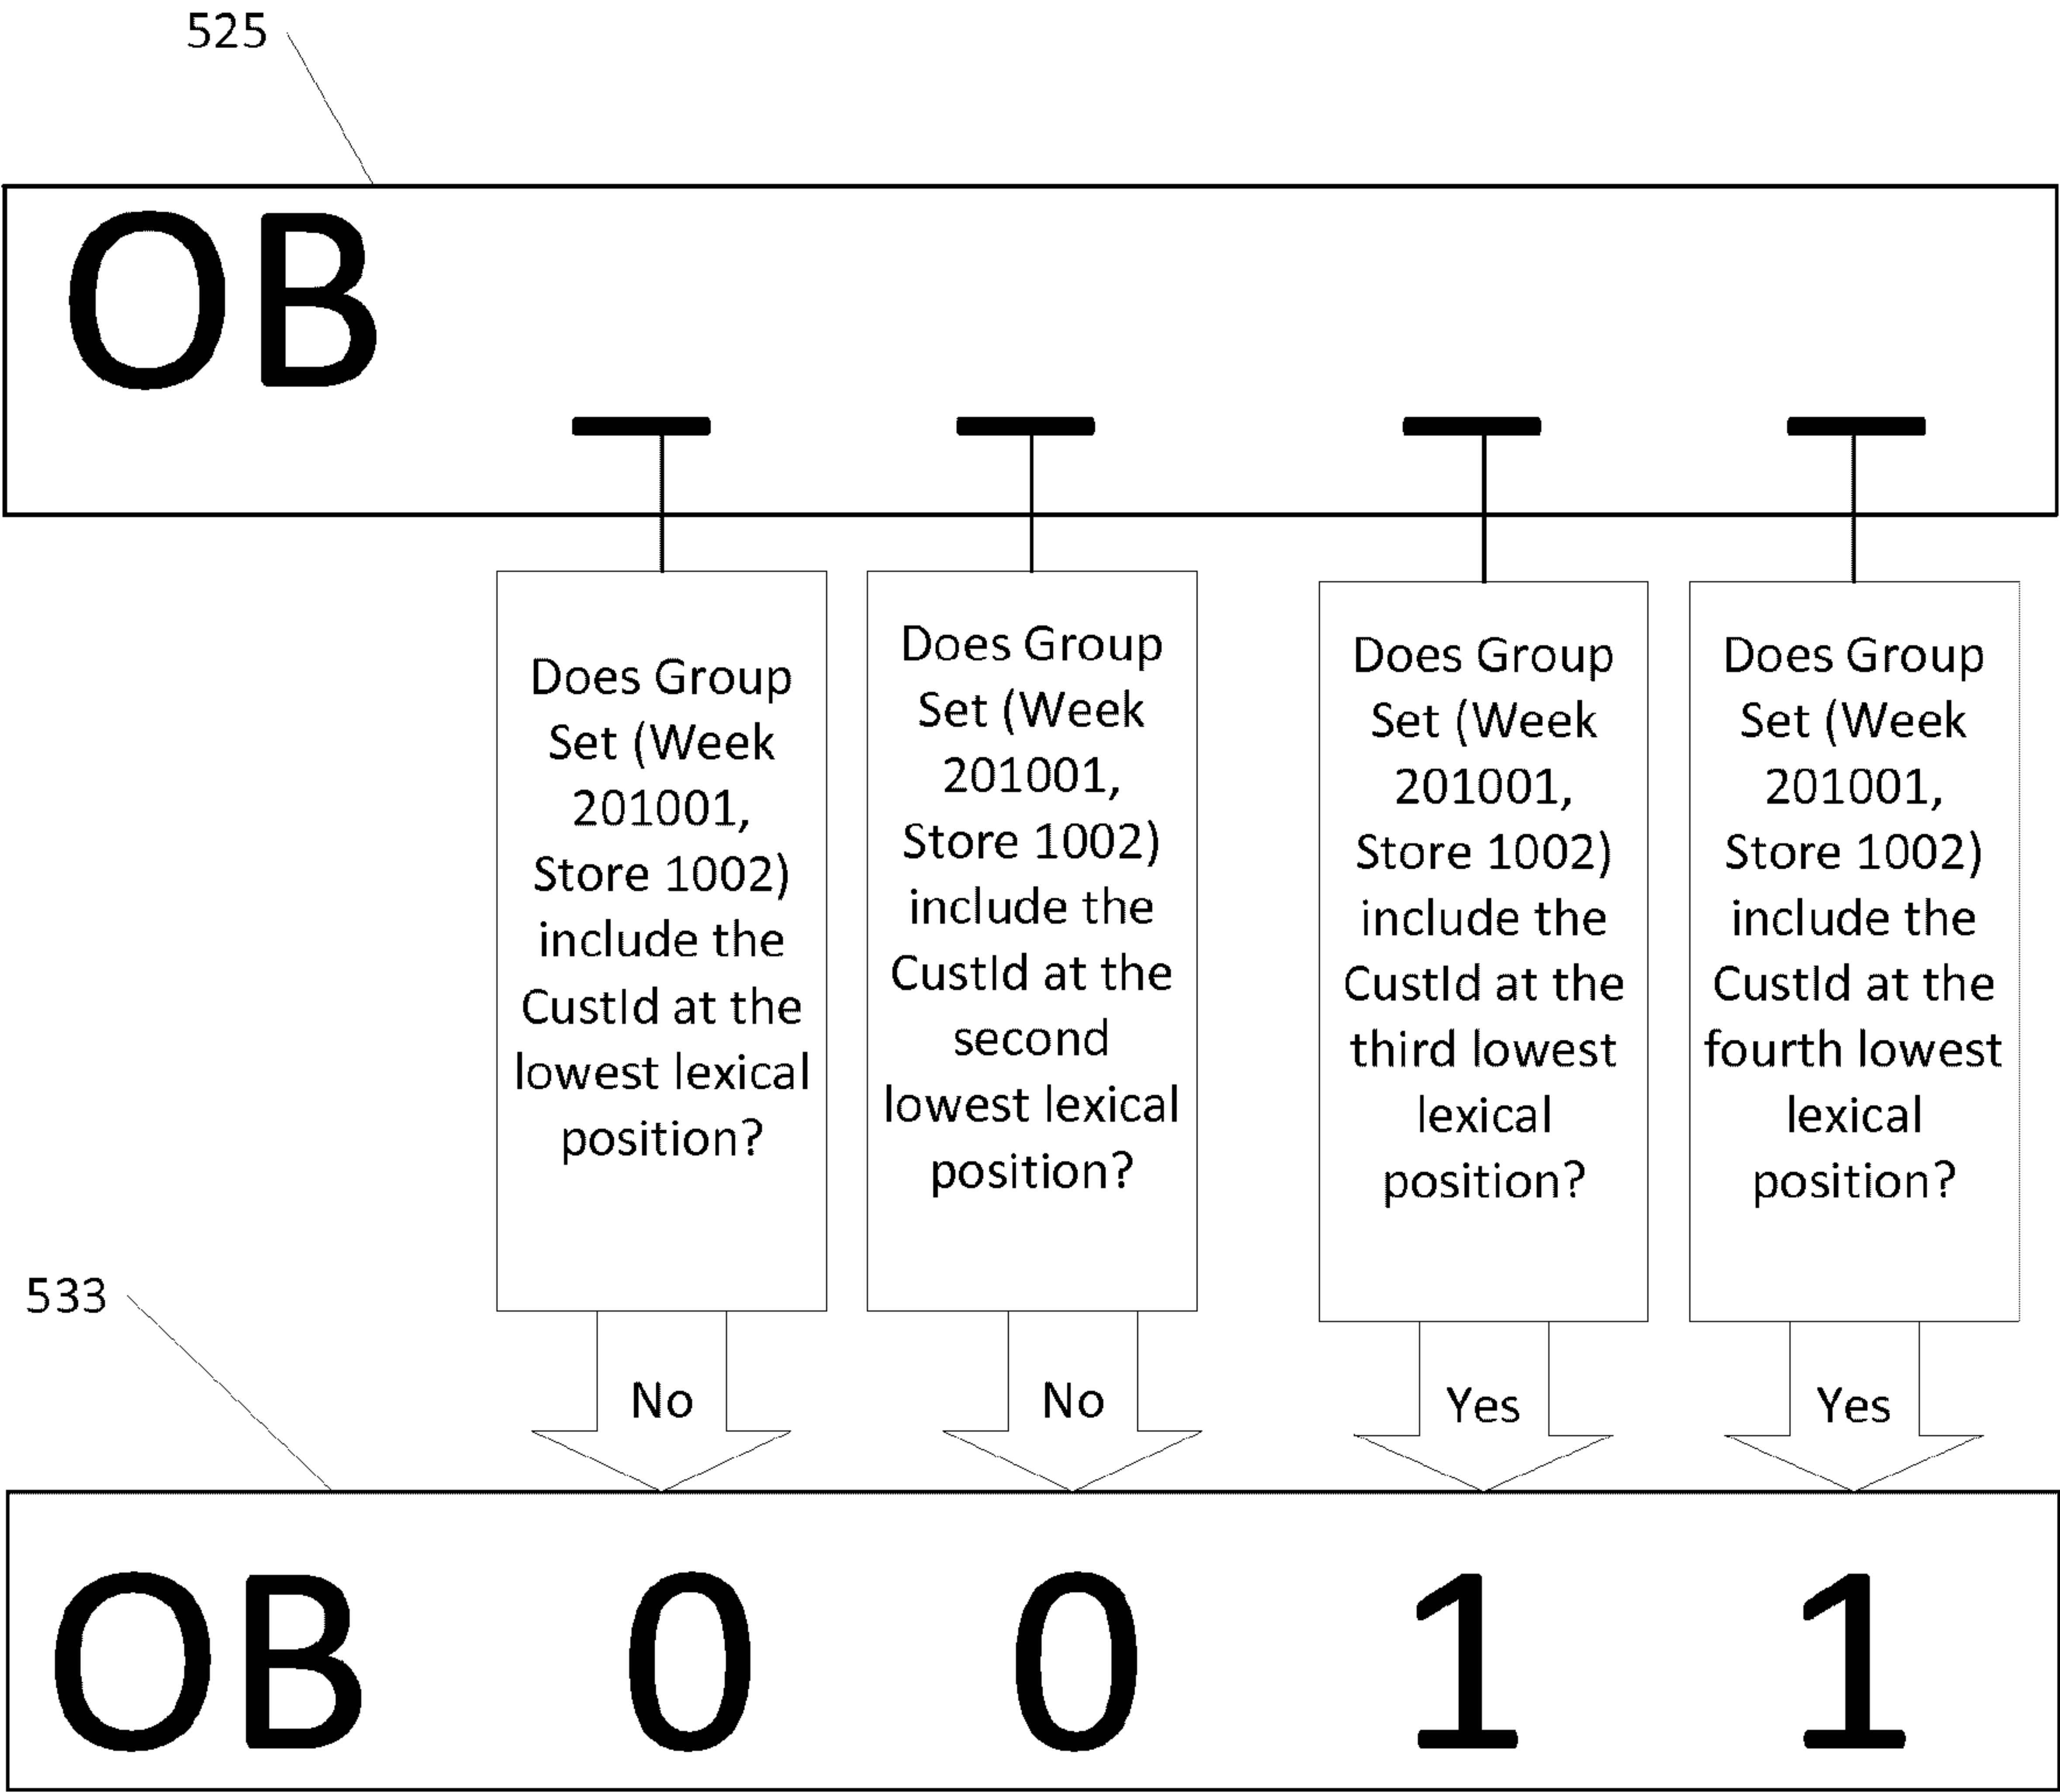


Fig. 5D



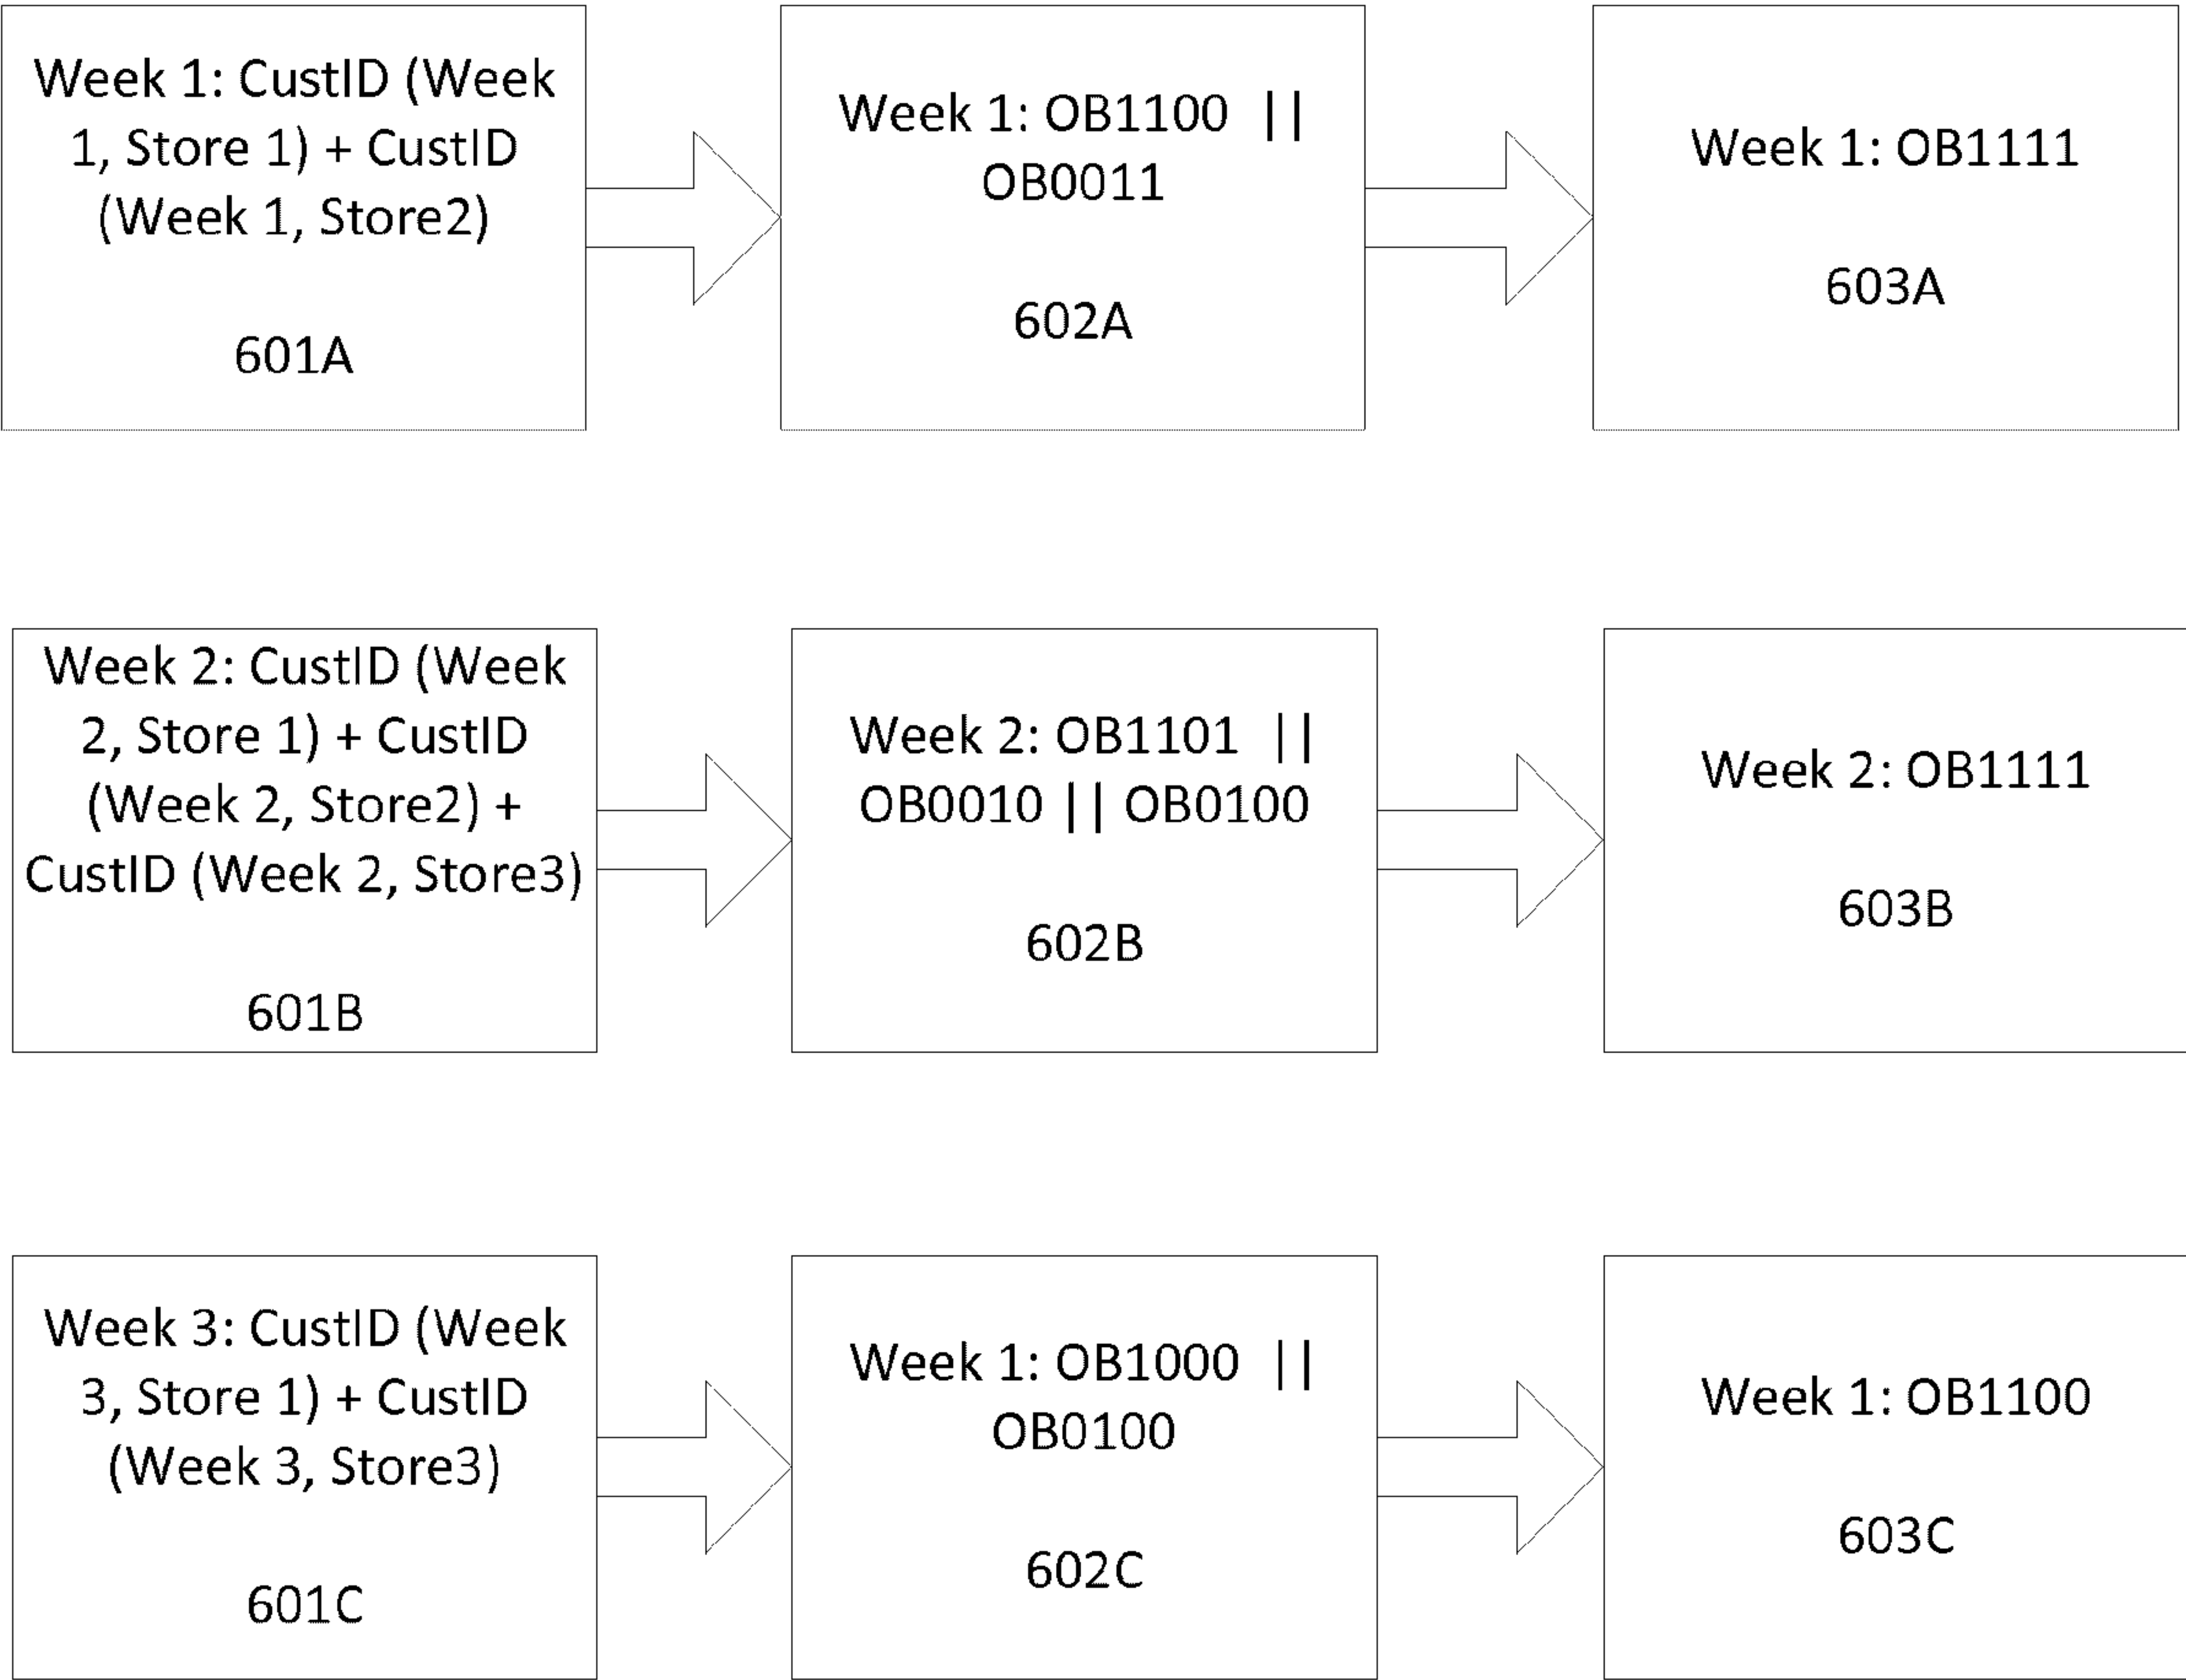


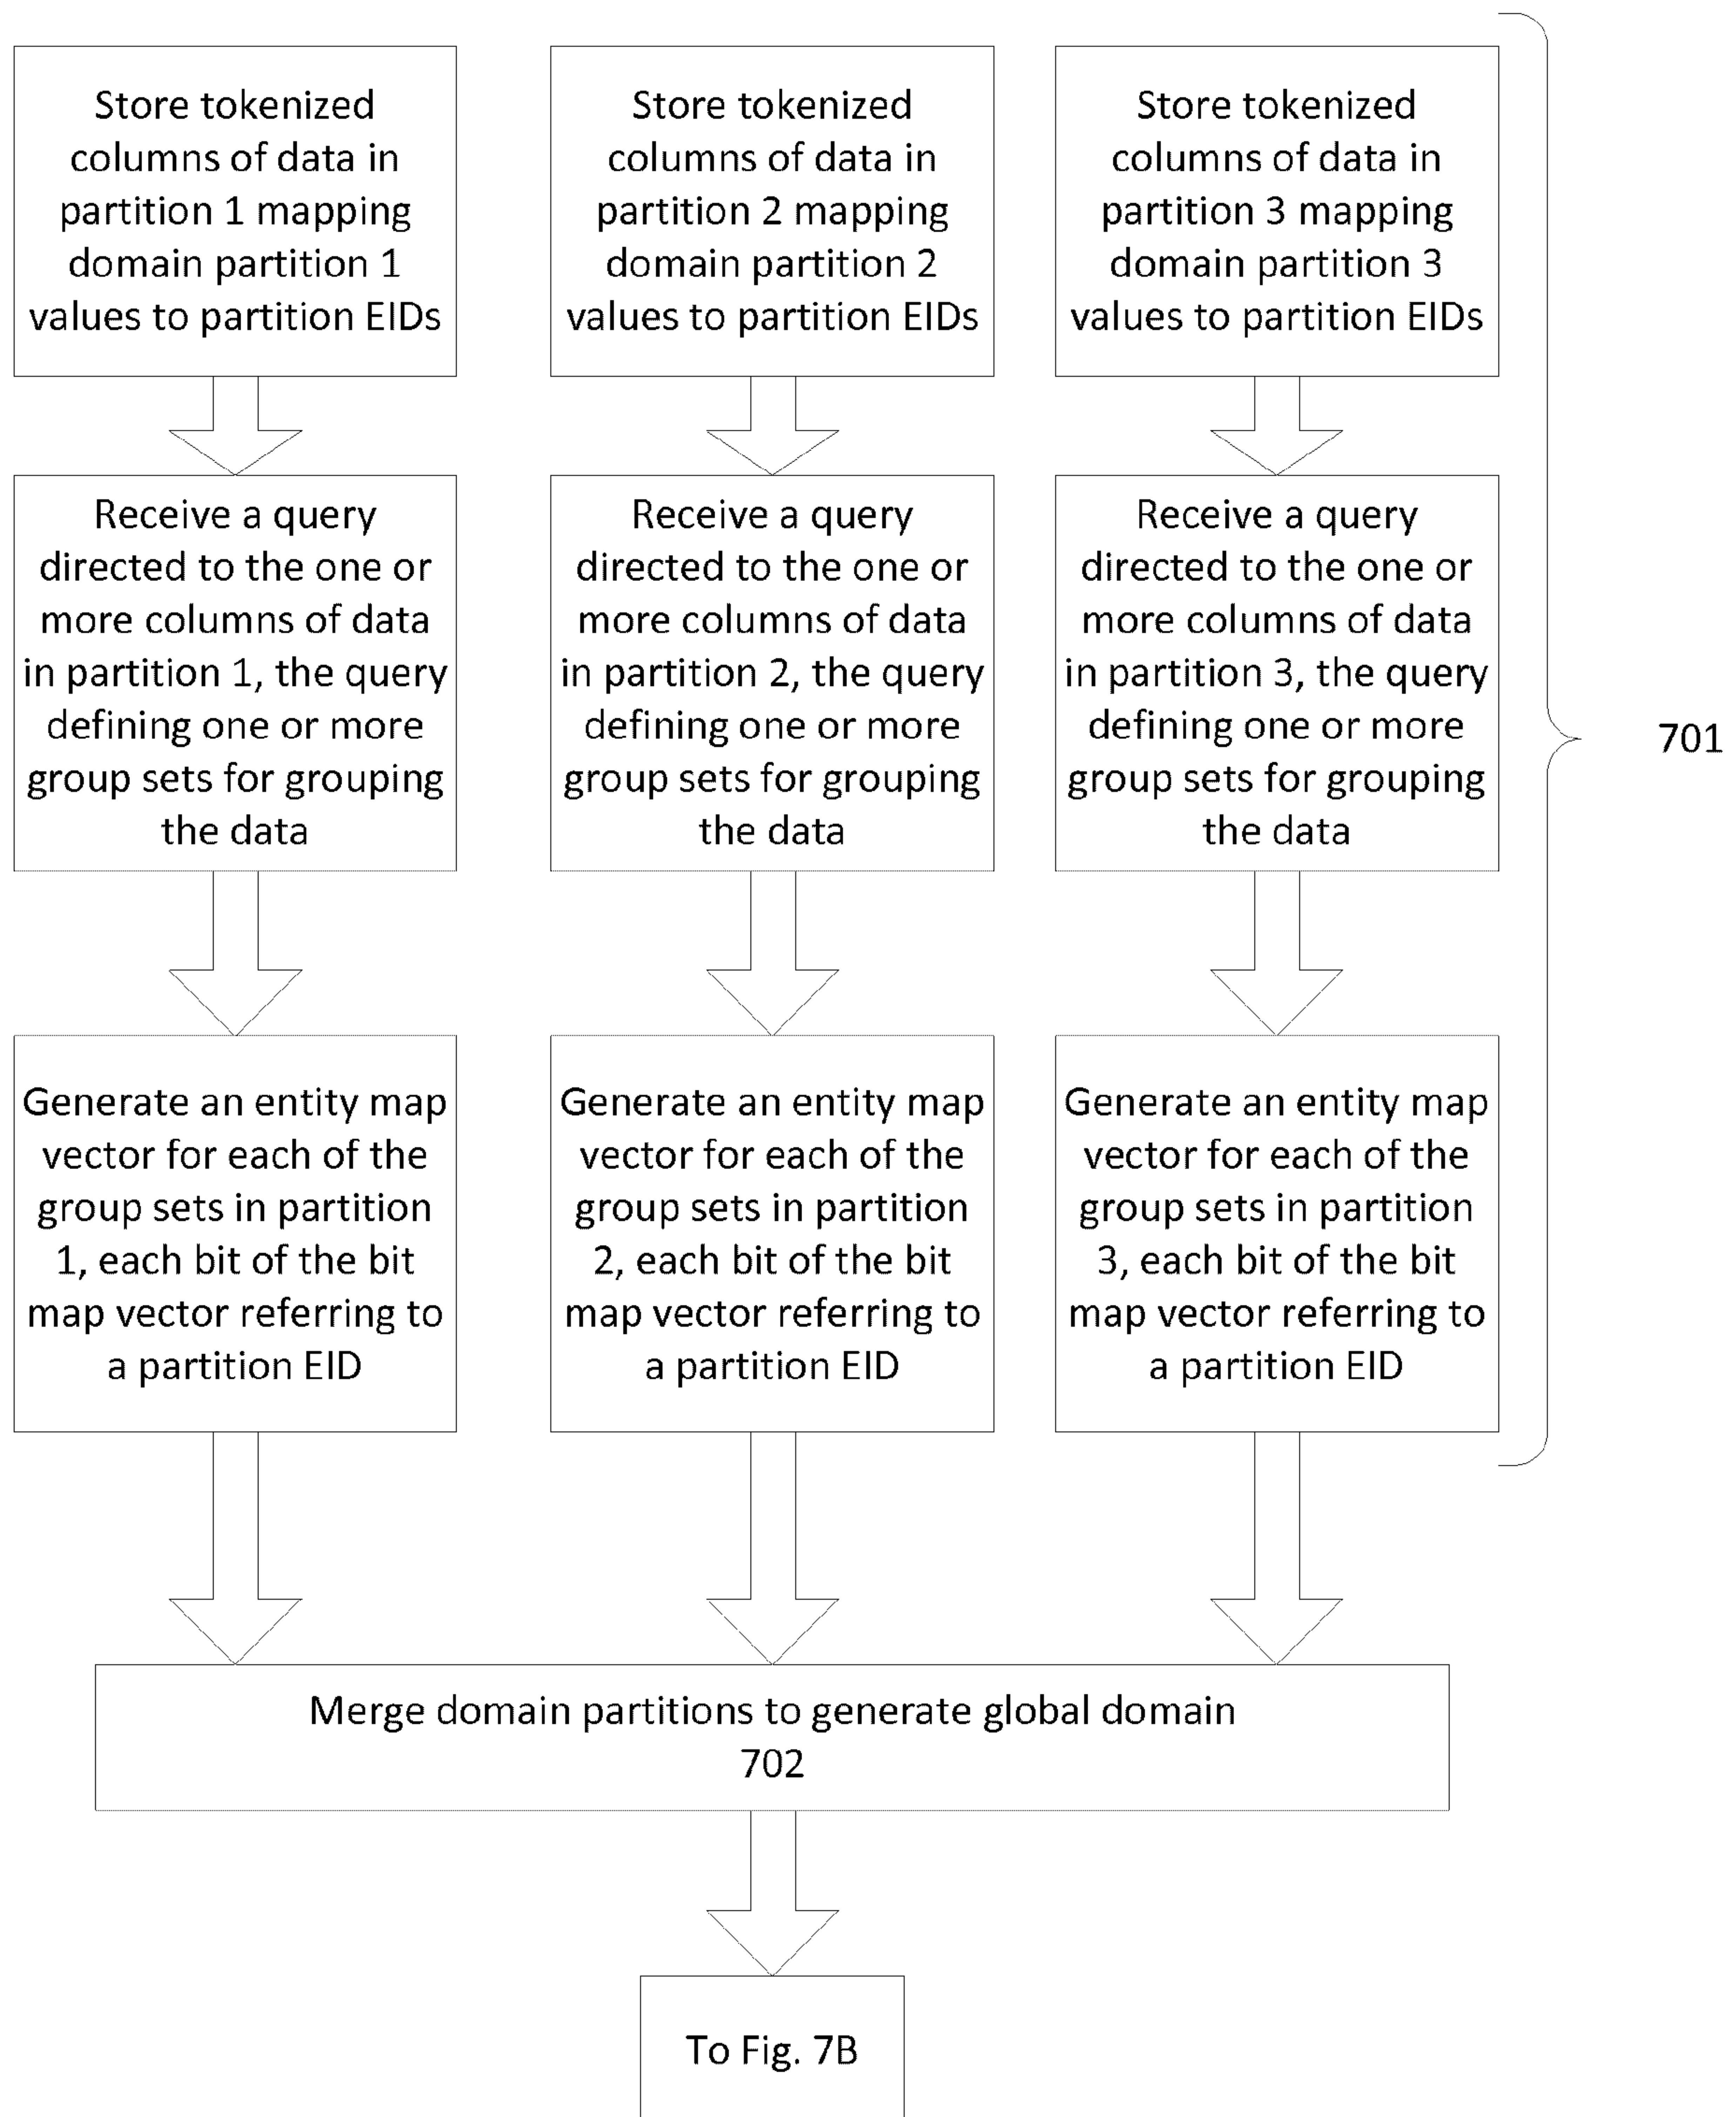
Fig. 6A

<u>611</u> TupleId	<u>612</u> Week	<u>613</u> EMV_CD CustId
1	1	4
2	2	4
3	3	2

Fig. 6B

610

Fig. 7A



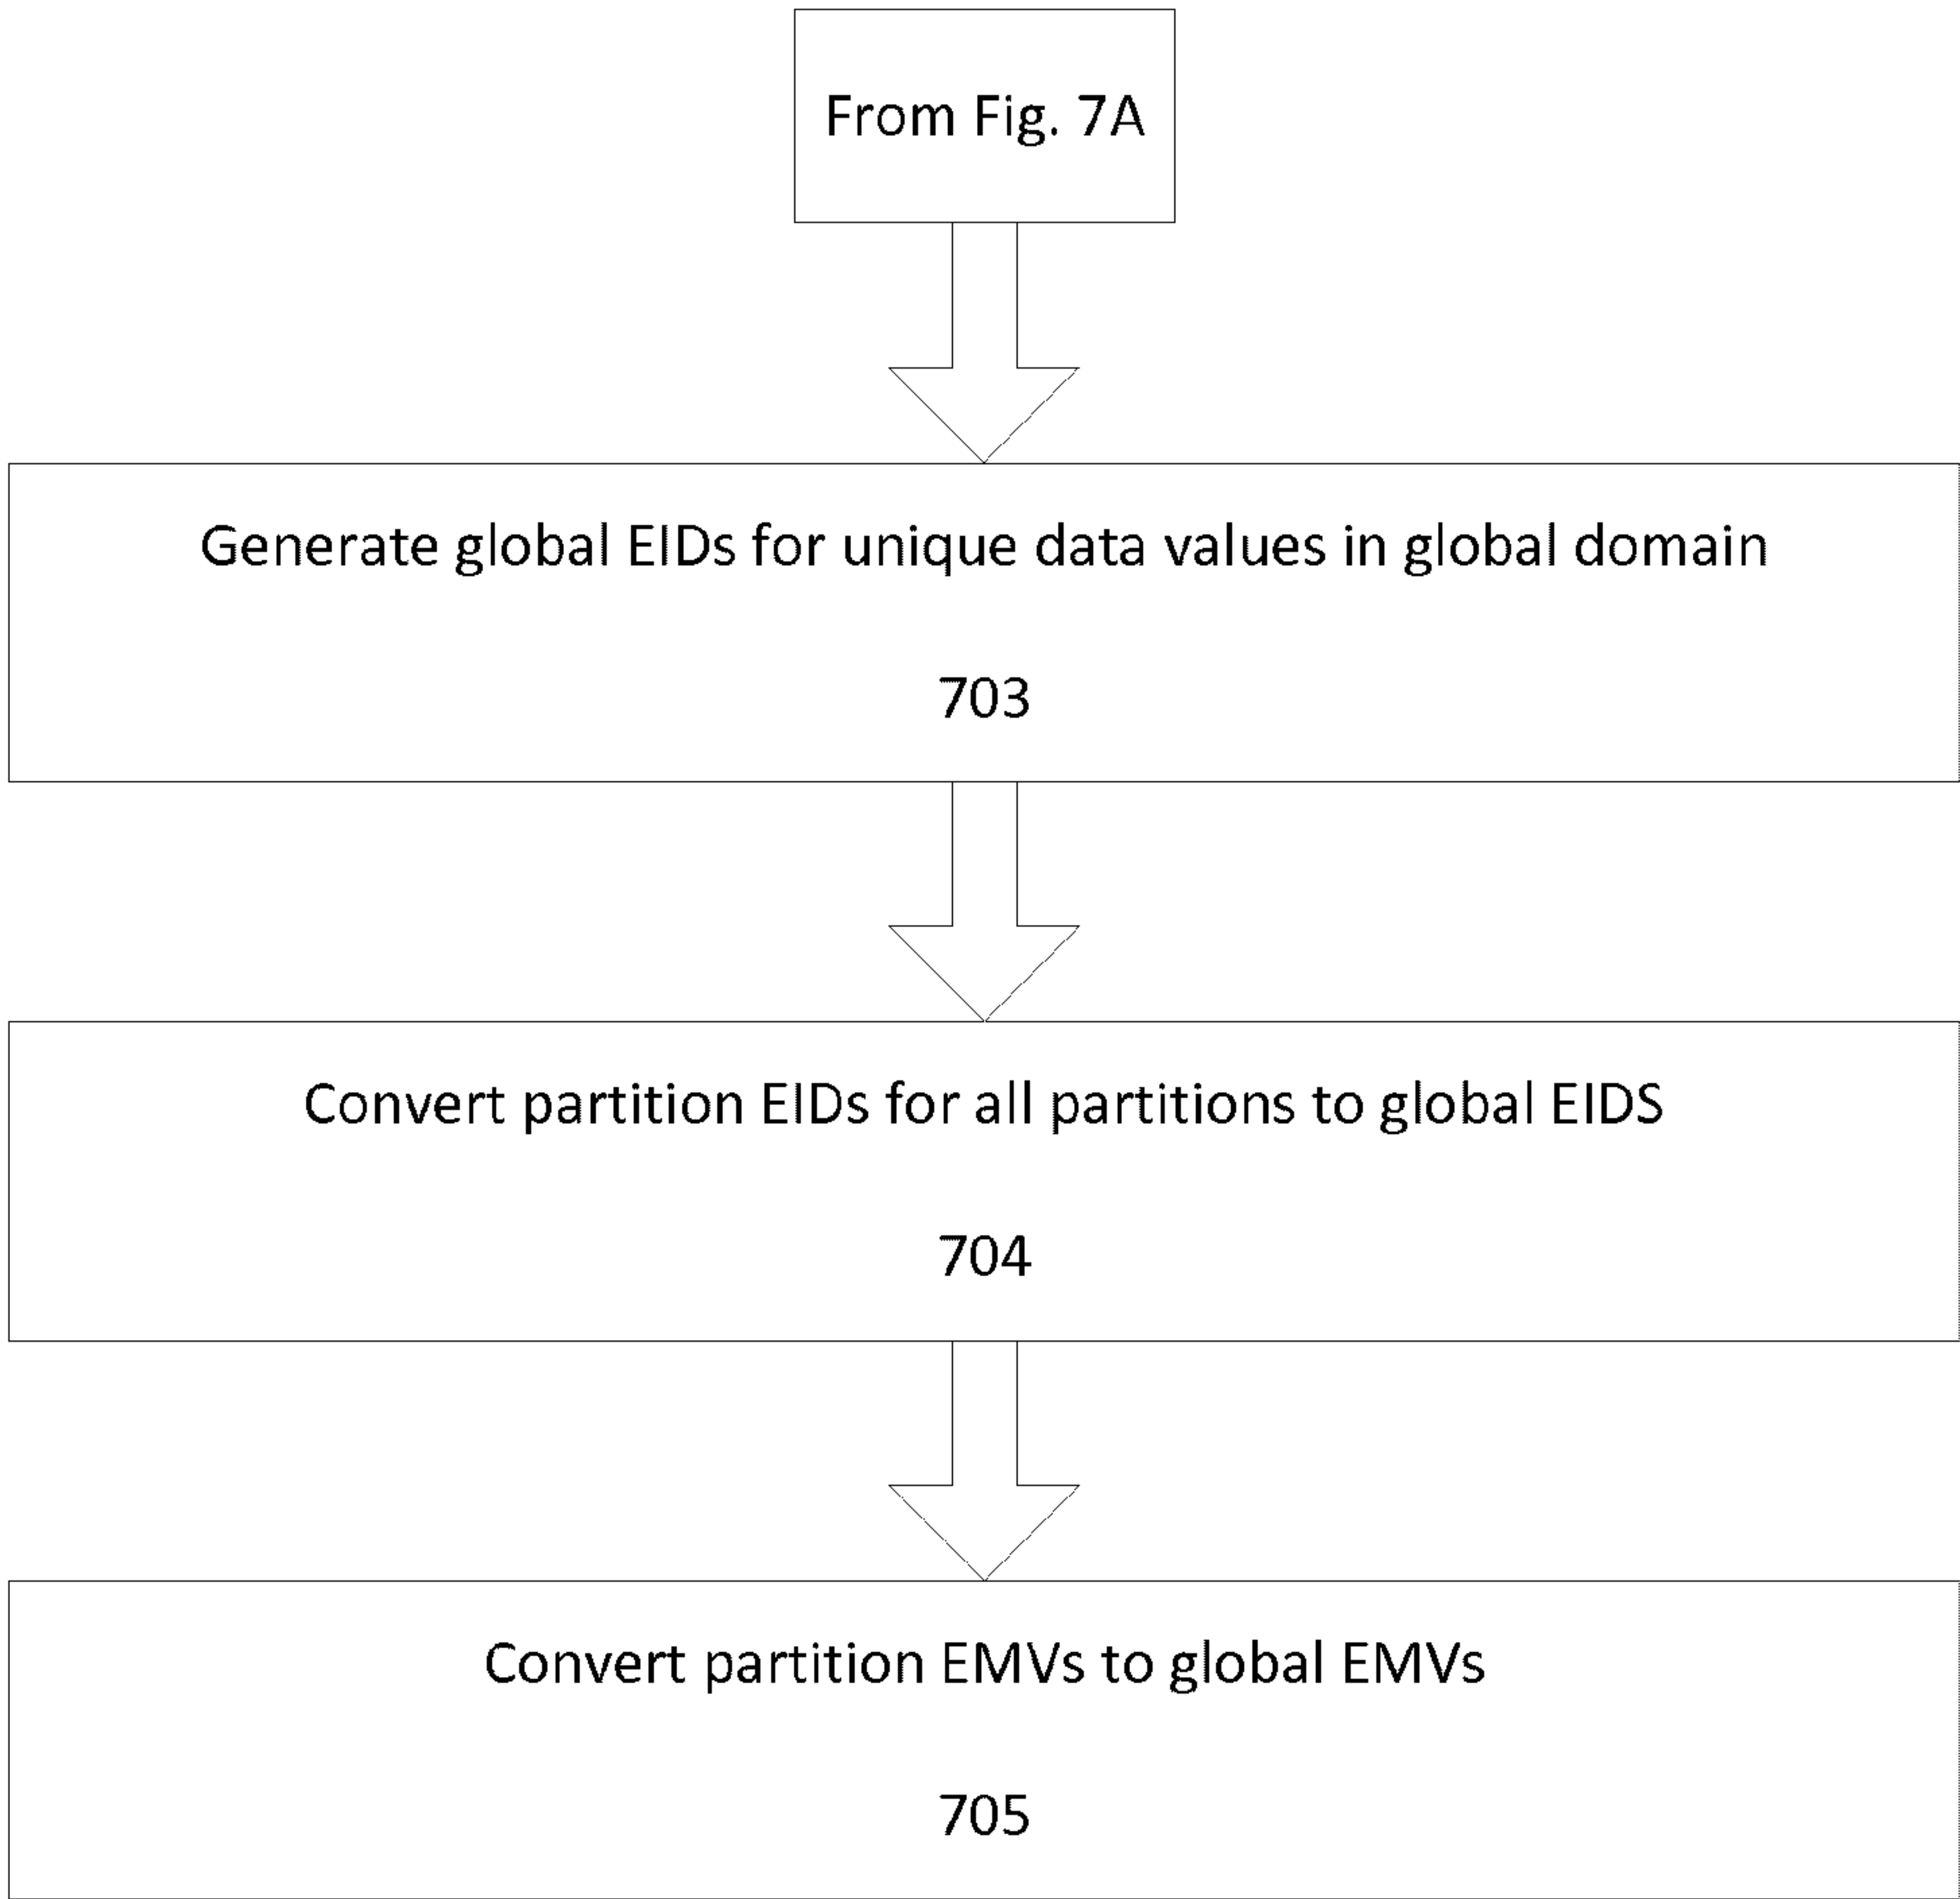


Fig. 7B



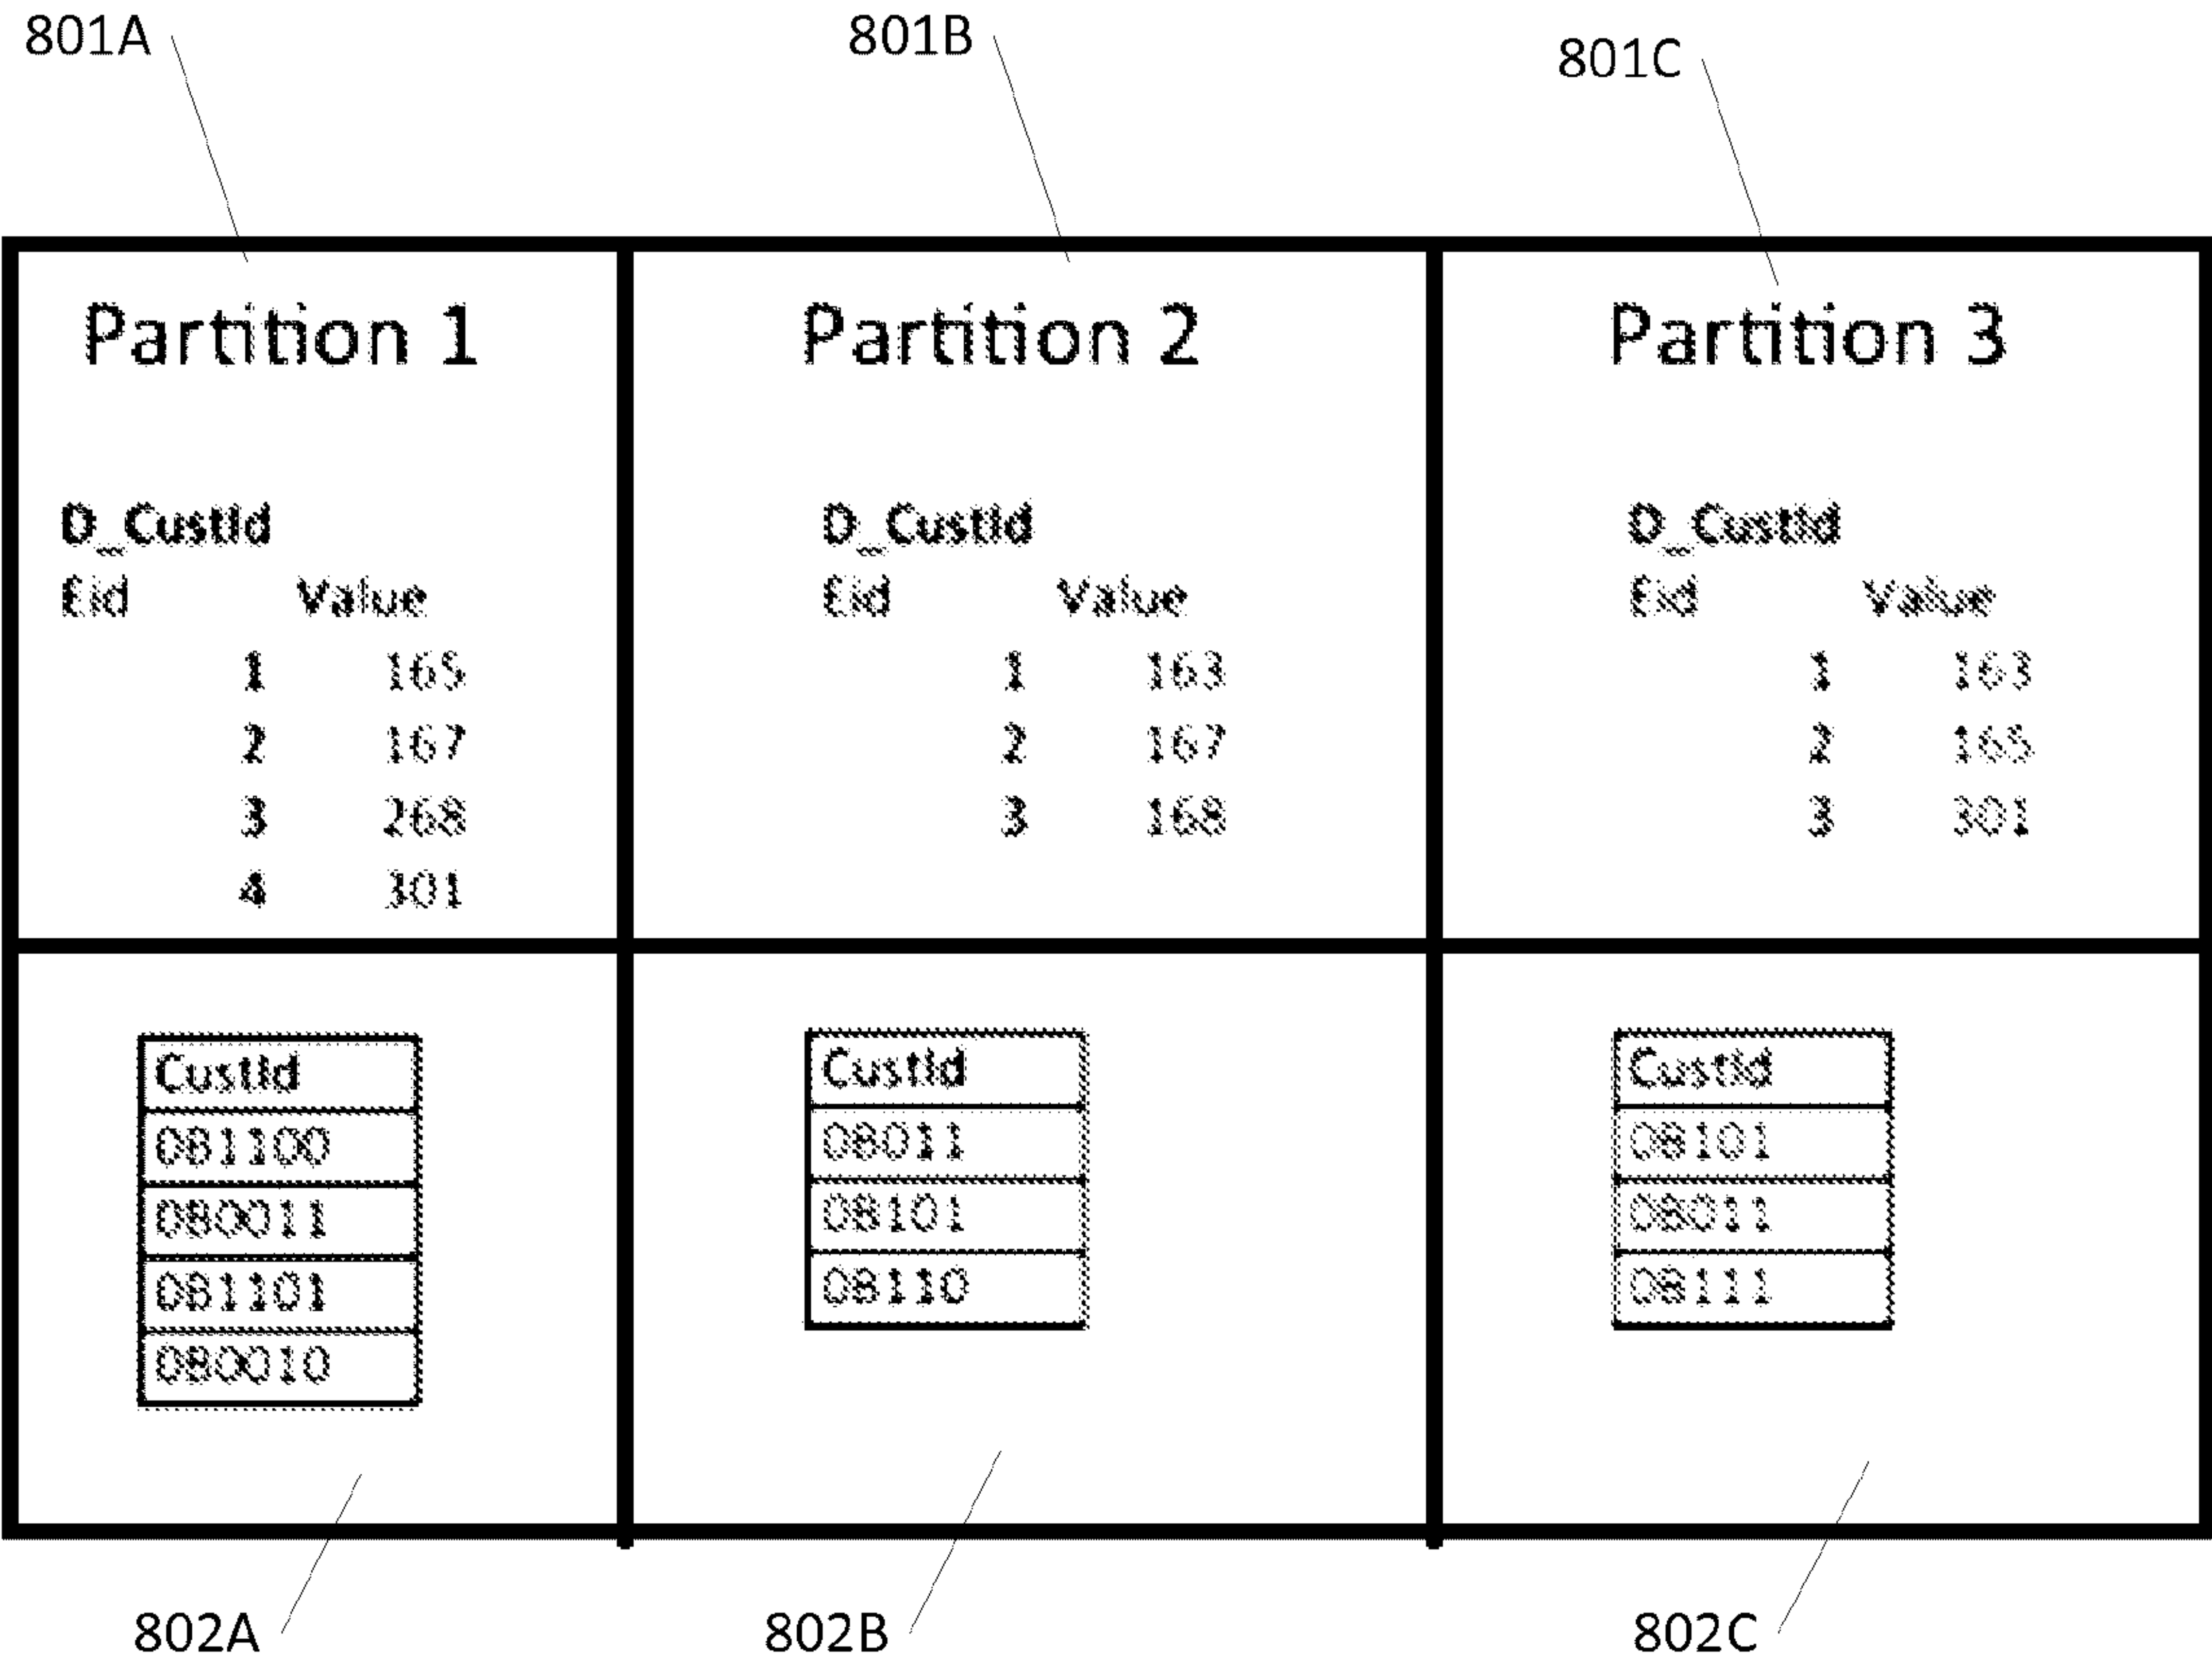


Fig. 8A




Fig. 8B

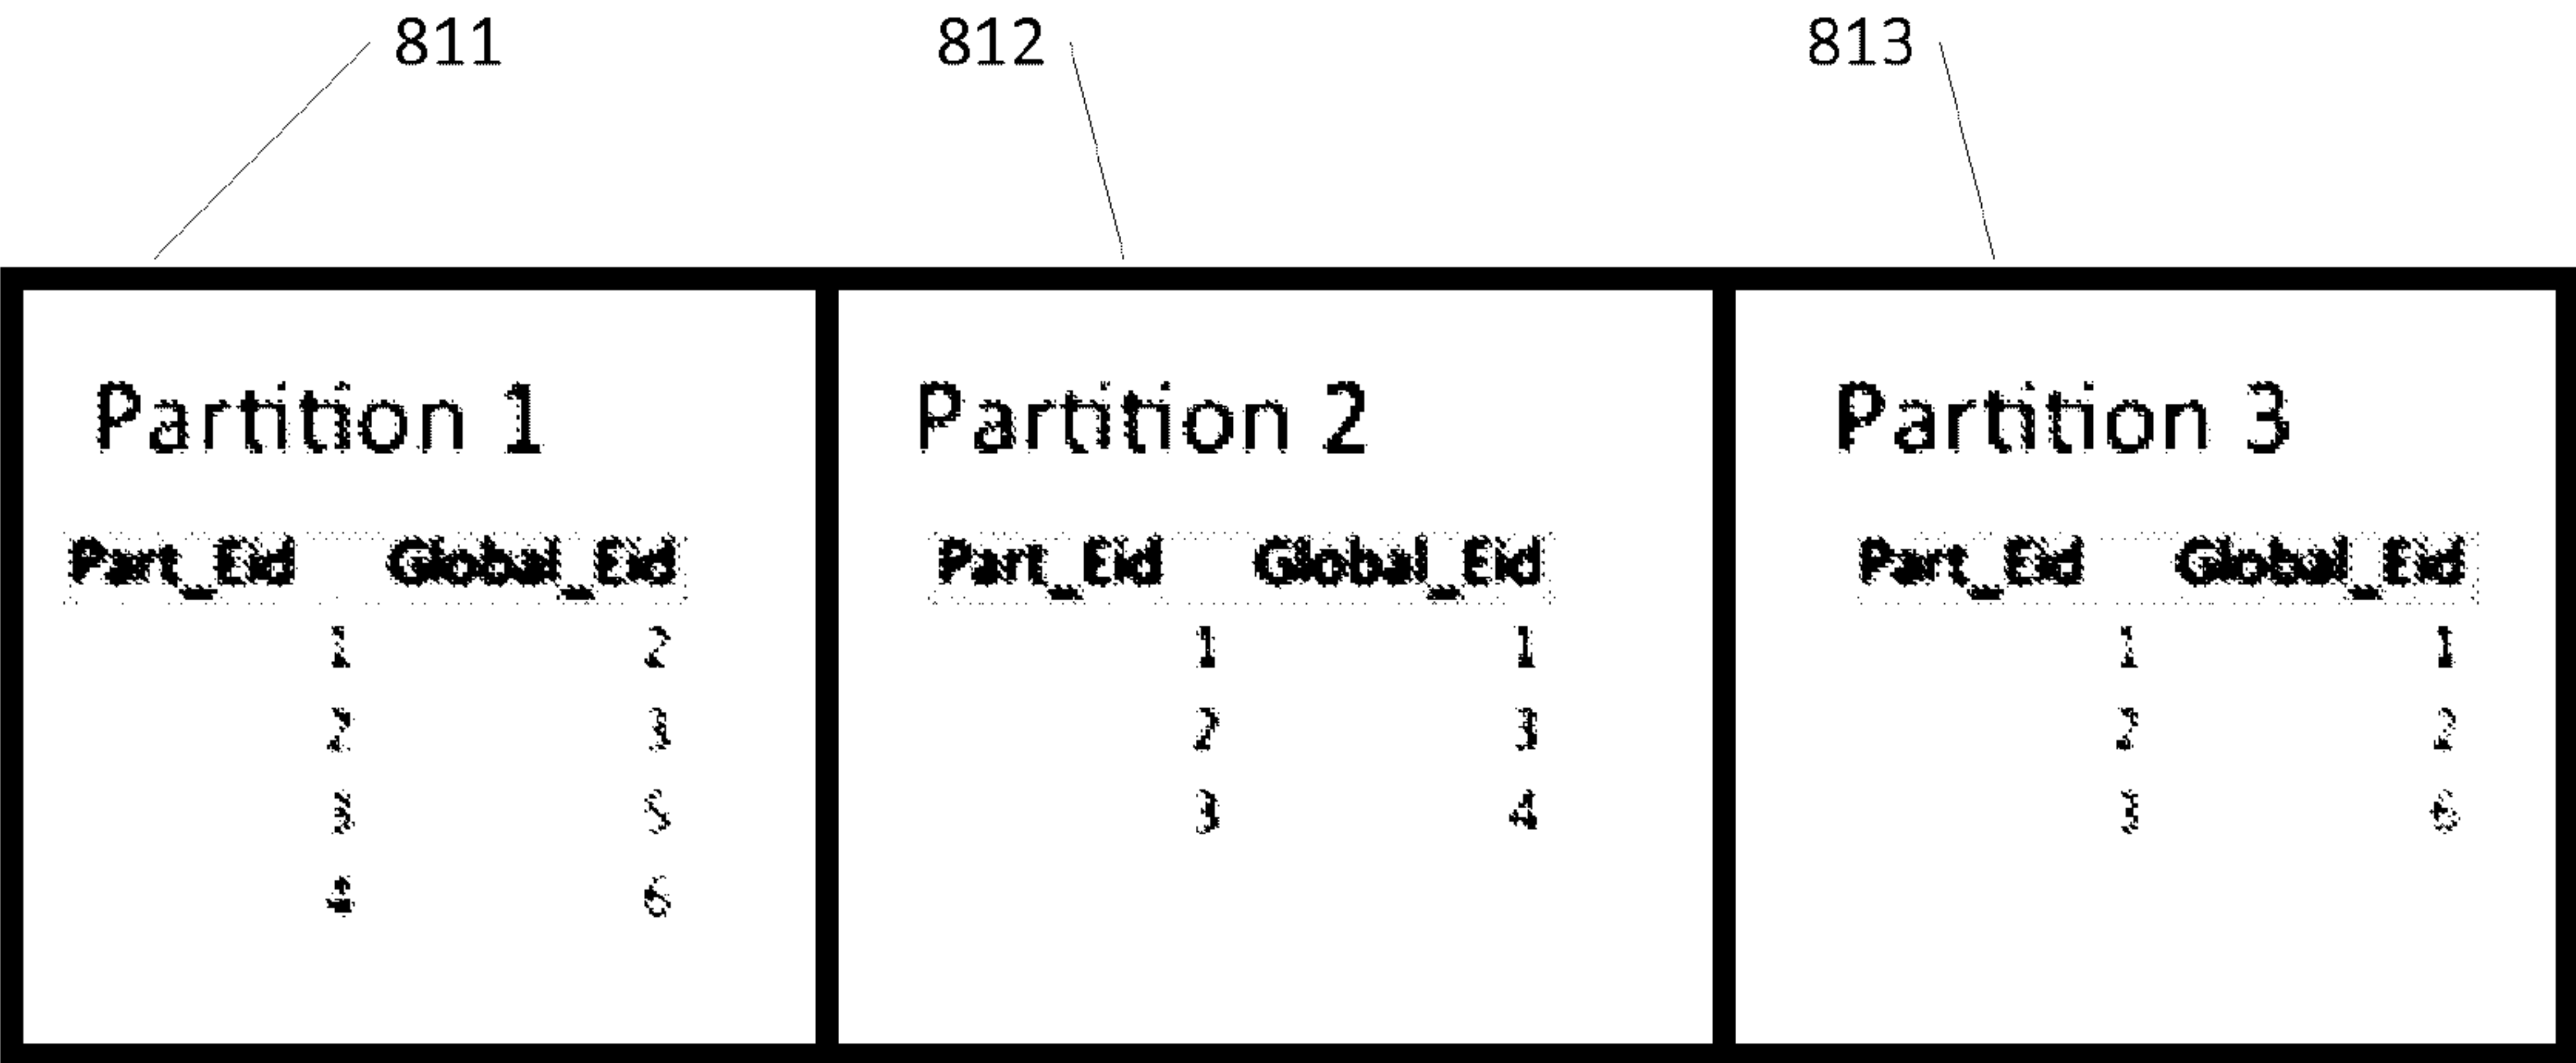


Fig. 8C

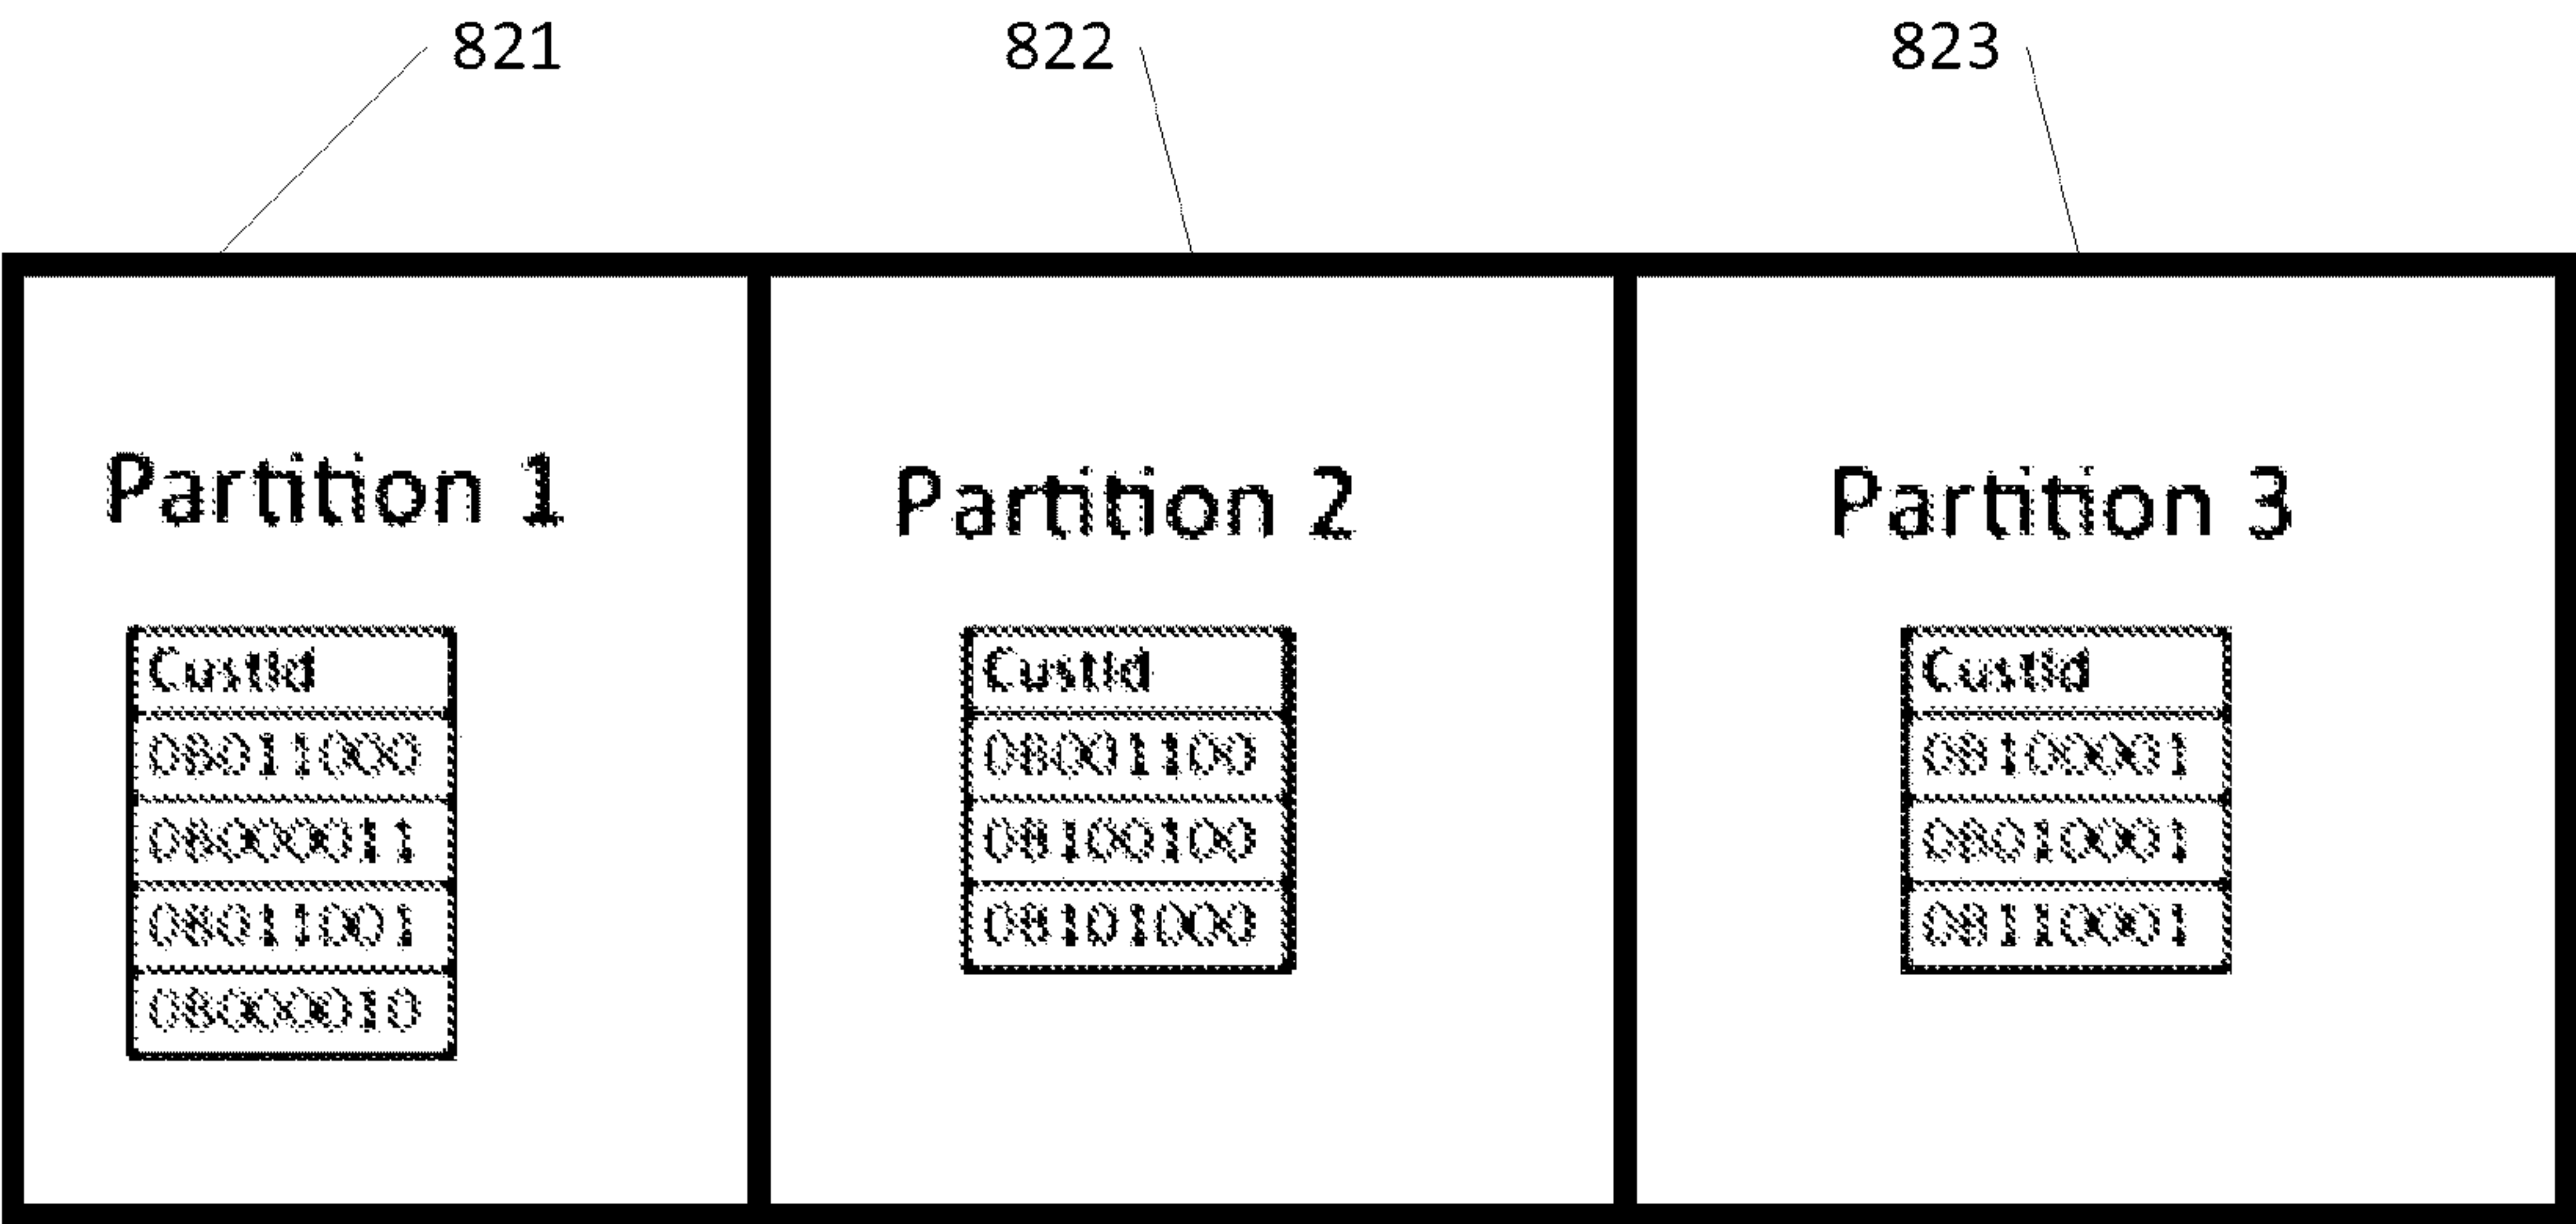


Fig. 8D




Fig. 9



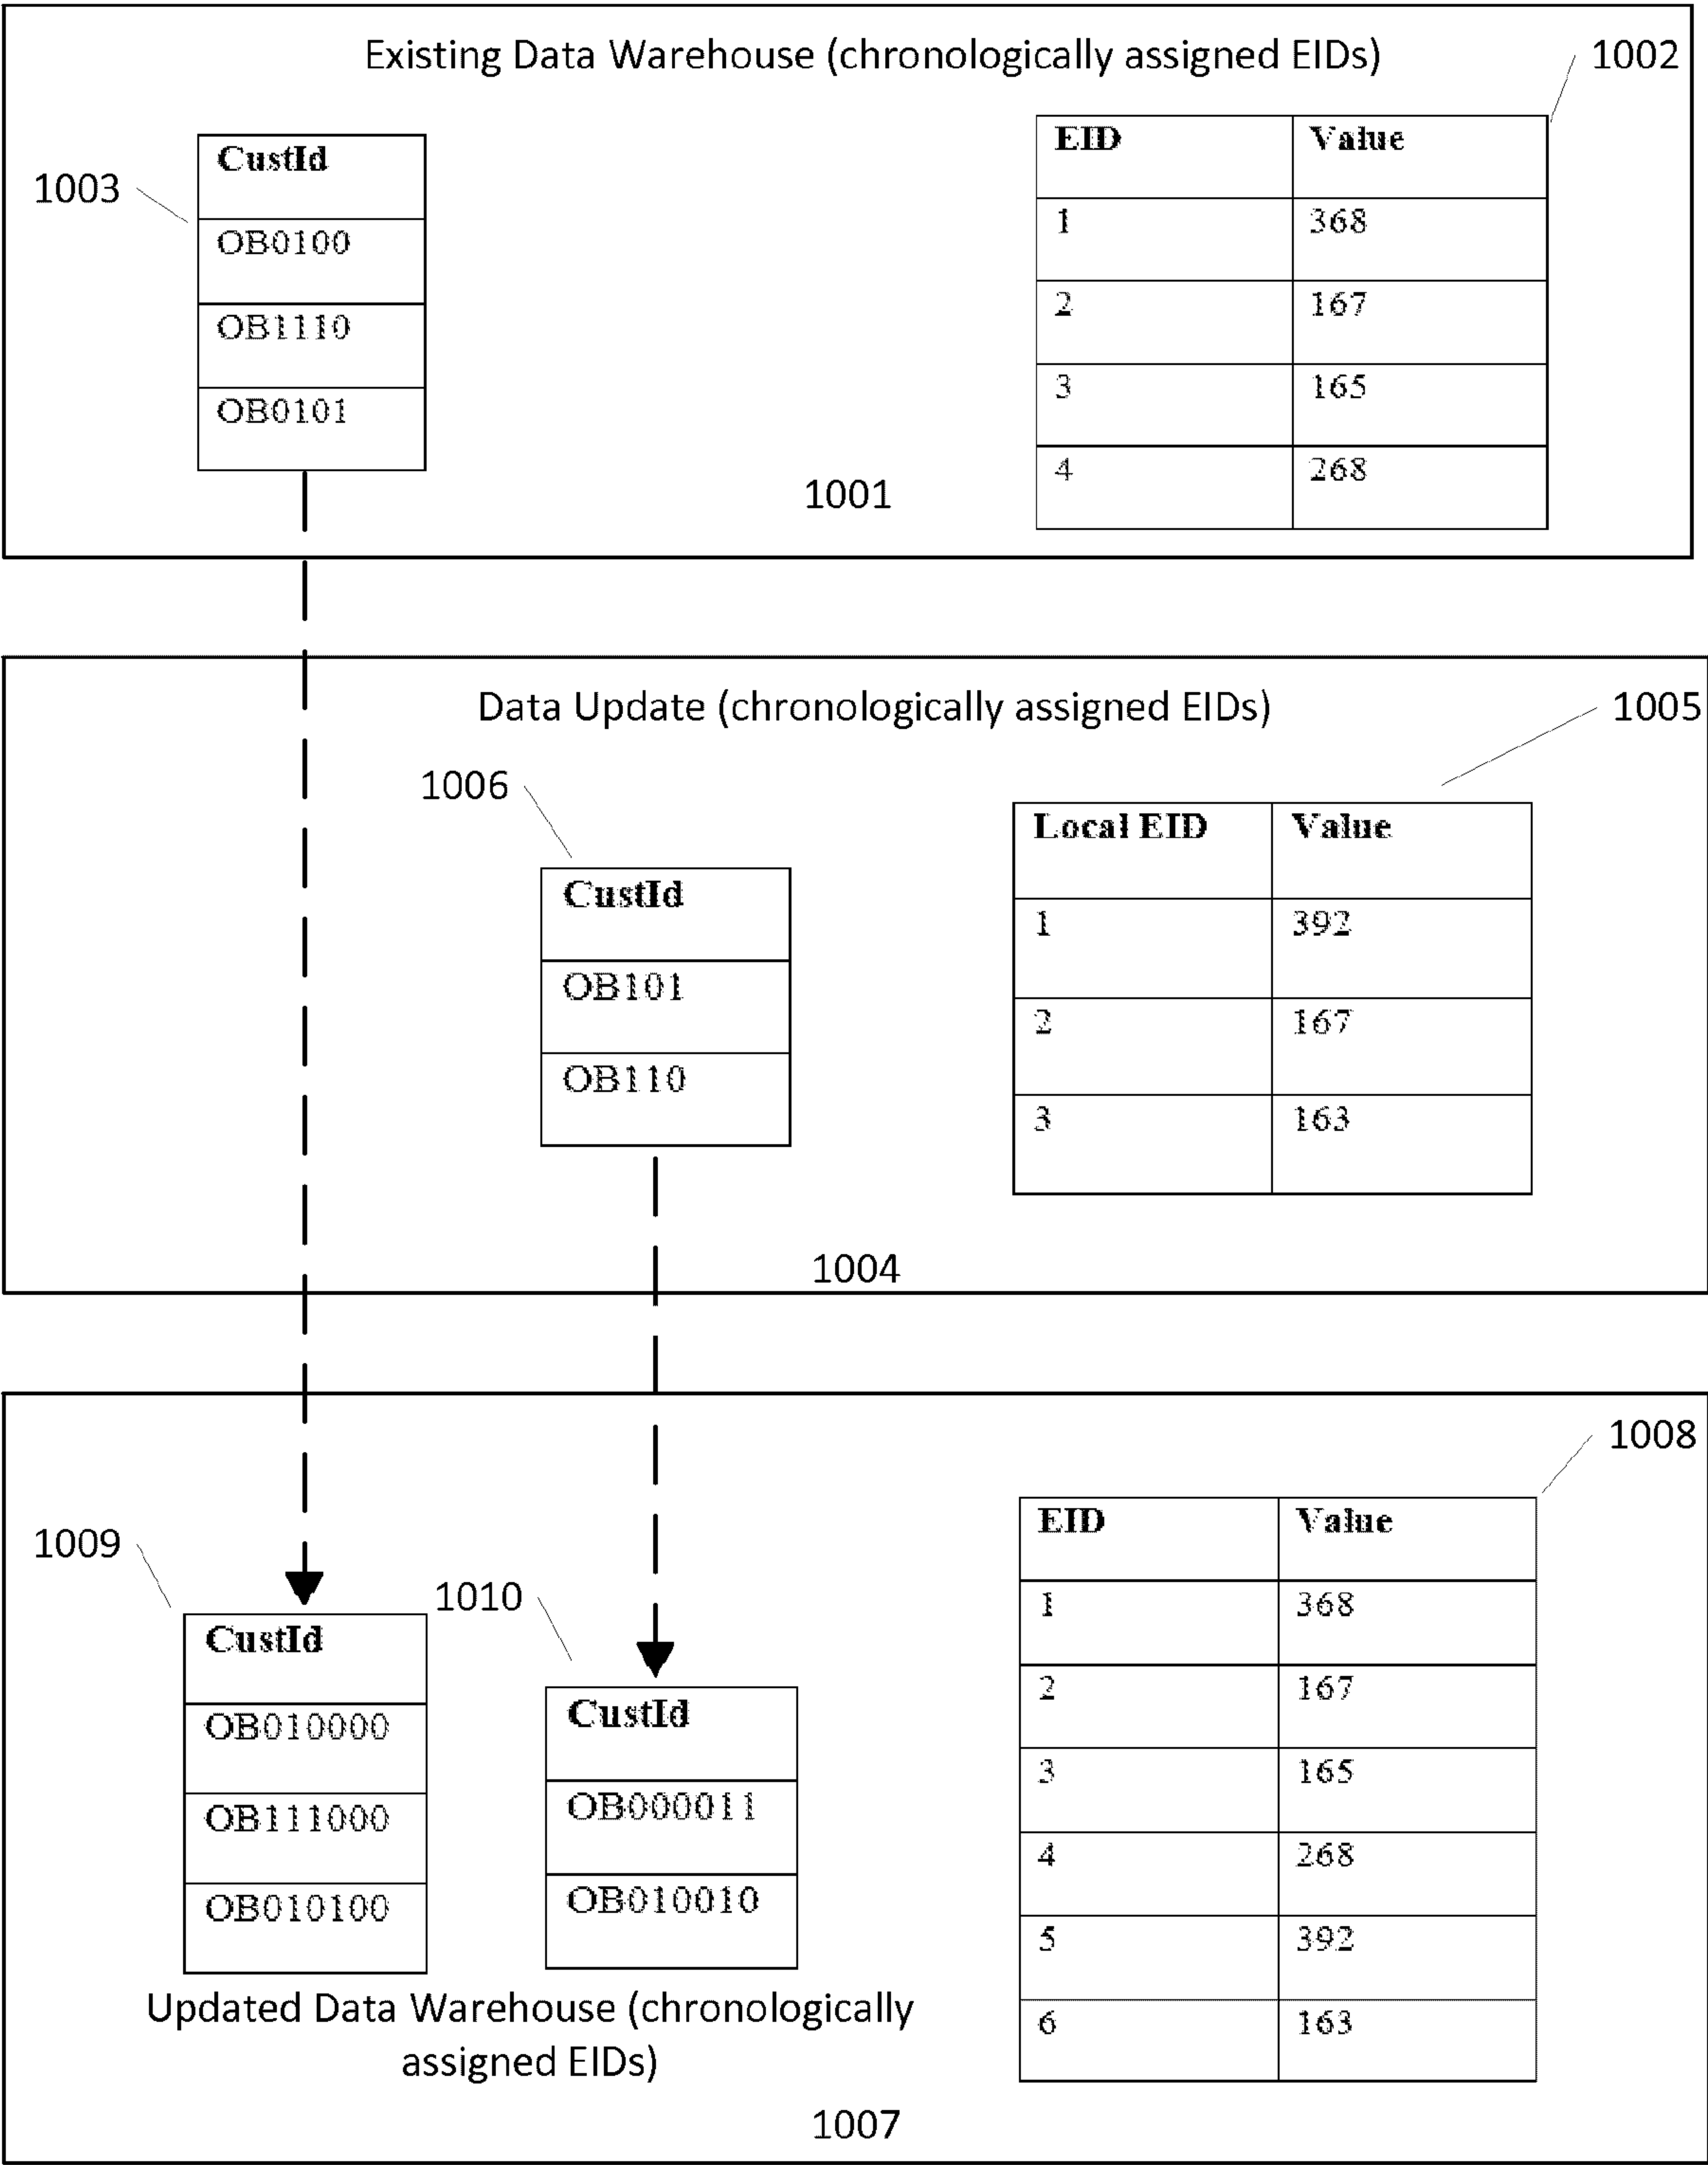




Fig. 10






Fig. 12

