Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2851350 A1 2013/05/16

(21) 2 851 350

(12) DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION

(13) A1

(86) Date de dépôt PCT/PCT Filing Date: 2012/07/27
(87) Date publication PCT/PCT Publication Date: 2013/05/16
(85) Entrée phase nationale/National Entry: 2014/04/07
(86) N° demande PCT/PCT Application No.: SE 2012/050852
(87) N° publication PCT/PCT Publication No.: 2013/070139
(30) Priorité/Priority: 2011/11/11 (US61/558,799)

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

(71) Demandeur/Applicant:
QLIKTECH INTERNATIONAL AB, SE

(72) Inventeurs/Inventors:
WOLGE, HAKAN, SE;
LINSEFORS, TOBIAS, SE

(74) Agent: MARKS & CLERK

(54) Titre : PROCEDE ET DISPOSITIF POUR LE CALCUL EN COLLABORATION
(54) Title: A METHOD AND A DEVICE FOR COLLABORATIVE COMPUTING

FIG. 11a

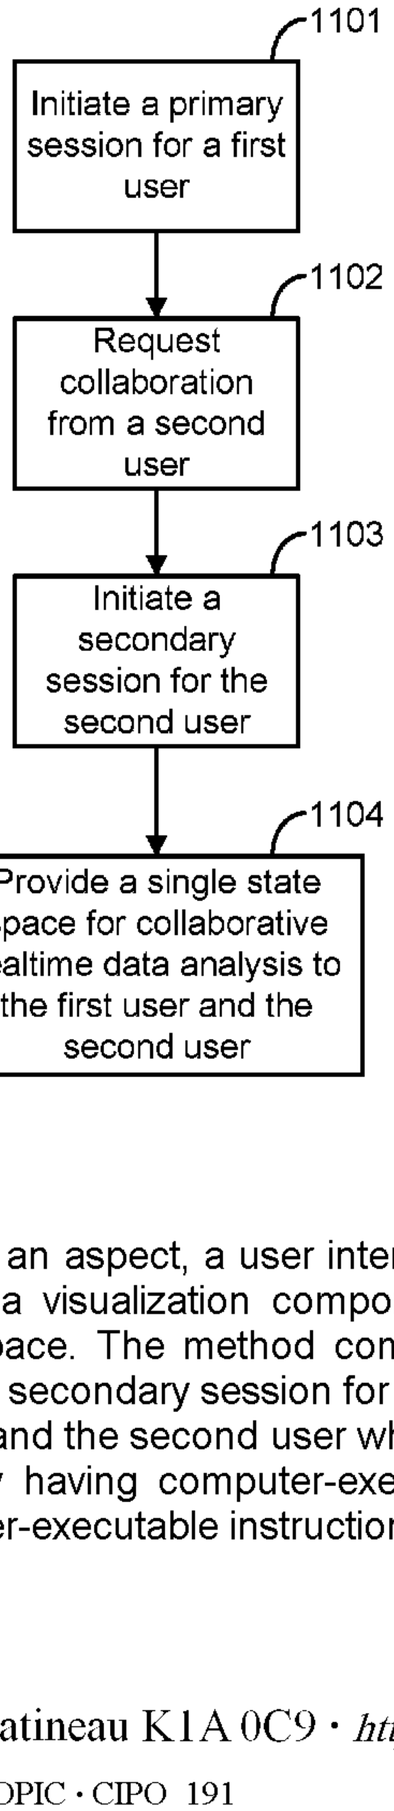

(57) Abrégé/Abstract:
A method and a device for collaborative computing. In an aspect, a user interface can be generated to facilitate dynamic display generation to view data. The system can comprise a visualization component to dynamically generate one or more visual representations of the data to present in the state space. The method comprises initiating a primary session for a first user, requesting collaboration from a second user, initiating a secondary session for the second user, and providing a single state space for collaborative real-time data analysis to the first user and the second user wherein an interaction by either user is reflected in the single state space. The device comprises a memory having computer-executable instructions; and a processor functionally coupled to the memory and configured, by the computer-executable instructions to perform the steps of the method.

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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
16 May 2013 (16.05.2013)

WIPO | PCT

(10) International Publication Number
WO 2013/070139 A3

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

(21) **International Application Number**:
PCT/SE2012/050852

(22) **International Filing Date**:
27 July 2012 (27.07.2012)

(25) **Filing Language**: English

(26) **Publication Language**: English

(30) **Priority Data**:
61/558,799 11 November 2011 (11.11.2011) US

(71) **Applicant** *(for all designated States except US)*: **QLIKTECH INTERNATIONAL AB** [SE/SE]; Scheelevägen 24-26, S-223 63 Lund (SE).

(72) **Inventors; and**
(75) **Inventors/Applicants** *(for US only)*: **WOLGÉ, Håkan** [SE/SE]; Terminsvägen 14, S-224 67 Lund (SE). **LINDSEFORS, Tobias** [SE/SE]; Elevgränden 18, S-224 67 Lund (SE).

(74) **Agent**: **HANSSON THYRESSON PATENTBYRÅ AB**; Box 73, S-201 20 Malmö (SE).

(81) **Designated States** *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, 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,

[Continued on next page]

(54) **Title**: A METHOD AND A DEVICE FOR COLLABORATIVE COMPUTING

(57) **Abstract**: A method and a device for collaborative computing. In an aspect, a user interface can be generated to facilitate dynamic display generation to view data. The system can comprise a visualization component to dynamically generate one or more visual representations of the data to present in the state space. The method comprises initiating a primary session for a first user, requesting collaboration from a second user, initiating a secondary session for the second user, and providing a single state space for collaborative real-time data analysis to the first user and the second user wherein an interaction by either user is reflected in the single state space. The device comprises a memory having computer-executable instructions; and a processor functionally coupled to the memory and configured, by the computer-executable instructions to perform the steps of the method.

FIG. 11a

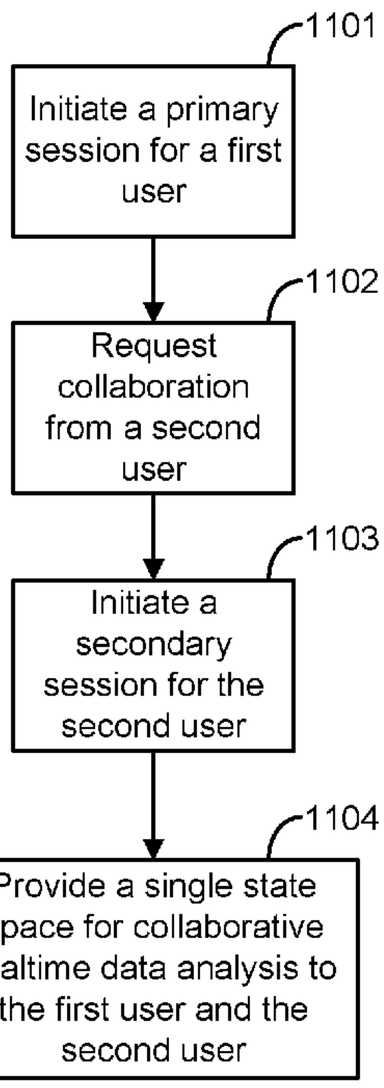

WO 2013/070139 A3

WO 2013/070139 A3

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 applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

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))*

(88) Date of publication of the international search report:
30 January 2014

A METHOD AND A DEVICE FOR COLLABORATIVE COMPUTING

TECHNICAL FIELD

[0001] The invention relates to a method for collaborative computing.

PRIOR ART

[0002] US7058621 discloses a method that operates on a database to extract and present information to a user. The database comprises data tables containing values of a number of variables. The information is to be extracted by evaluating at least one mathematical function which operates on one or more selected calculation variables. The presented information is to be partitioned on one or more selected classification variables. The method comprises the steps of identifying all boundary tables; identifying all connecting tables; electing a starting table among said boundary and connecting tables; building a conversion structure that links values of each selected variable in the boundary tables to corresponding values of one or more connecting variables in the starting table; and evaluating the mathematical function for each data record of the starting table, by using the conversion structure, such that the evaluation yields a final data structure containing a result of the mathematical function for every unique value of each classification variable.

[0003] US20100017436 discloses a method and an apparatus for extracting information. Information is extracted from a database using a computer-implemented method that involves a sequential chain of main calculations, in which a first main calculation (P1) operates a first selection item (S1) on a data set (R0) that represents the database to produce a first result (R1), and a second main calculation (P2) operates a second selection item (S2) on the first result (R1) to produce a second result (R2). The first and second results (R1, R2) are cached in computer memory (10) for re-use in subsequent iterations of the method, thereby reducing the need to execute the first and/or second main calculations (P1, P2) for extracting the information. The caching involves calculating a first selection identifier value (ID1) as a function of at least the first selection item (S1), and a second selection identifier value (ID3) as a function of at least the second selection item (S2) and the first result (R1), and storing the

first selection identifier value (ID1) and the first result (R1), and the second selection identifier value (ID3) and the second result (R2), respectively, as associated objects in a data structure (12). Each of the identifier values are generated as a statistically unique digital fingerprint by a hash function (f). The re-use involves calculating the first and second selection identifier values (ID1, ID3) during a subsequent iteration and accessing the data structure (12) to potentially retrieve the first and/or second result (R1, R2).

SUMMARY OF THE INVENTION

[0004] In an aspect, provided are methods and systems for user interaction with database methods and systems. In an aspect, a user interface can be generated to facilitate dynamic display generation to view data. The system can comprise a visualization component to dynamically generate one or more visual representations of the data to present in the state space.

[0005] In an aspect, provided are methods, systems, and computer readable media for collaborative computing comprising, initiating a primary session for a first user, requesting collaboration from a second user, initiating a secondary session for the second user, and providing a single state space for collaborative real-time data analysis to the first user and the second user wherein an interaction by either user is reflected in the single state space. In a further aspect, provided are methods, systems, and computer readable media for time shifted collaborative analysis comprising, creating a state space that reflects a selection state, creating a note, attaching the note to an object in the state space, saving the selection state, and associating the saved selection state with the note.

[0006] In another aspect, provided are methods, systems, and computer readable media for time shifted collaborative analysis comprising, creating a state space that reflects a selection state, creating a note, and attaching the note to an object in the state space. In a further aspect, provided are methods, systems, and computer readable media for time shifted collaborative analysis comprising, presenting an object in a state space having an attached note, receiving a selection of the note, and presenting the note and adjusting the state space to reflect a saved selection state associated with the note.

[0007] Additional advantages will be set forth in part in the description which follows or may be learned by practice. The advantages will be realized and attained by means of the elements and combinations particularly pointed out in the appended claims. It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive, as claimed.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate embodiments and together with the description, serve to explain the principles of the methods and systems:

Figure 1 illustrates exemplary Tables 1-5;

Figure 2 illustrates a block flow chart of an exemplary method for extracting information from a database;

Figure 3 illustrates exemplary Tables 6-12;

Figure 4 illustrates exemplary Tables 13-16;

Figure 5 illustrates exemplary Tables 17, 18, and 20-23;

Figure 6 illustrates exemplary Tables 24-29;

Figure 7 is an exemplary operating environment;

Figure 8 illustrates how a Selection operates on a Scope to generate a Data Subset;

Figure 9 illustrates an exemplary user interface;

Figure 10a illustrates another exemplary user interface;

Figure 10b illustrates another exemplary user interface;

Figure 11a is a block flow chart of an exemplary method;

Figure 11b is an exemplary operating environment;

Figure 12a illustrates an exemplary user interface;

Figure 12b illustrates another exemplary user interface;

Figure 13a is a block flow chart of an exemplary method;

Figure 13b is another block flow chart of an exemplary method;

Figure 13c is another block flow chart of an exemplary method;

Figure 14 illustrates an exemplary user interface;

Figure 15 is a block flow chart of an exemplary method;

Figure 16 illustrates an exemplary user interface;

Figures 17a-f illustrate exemplary Tables;

Figures 18a-c illustrate additional exemplary Tables; and

Figure 19 is a block flow chart of an exemplary method.

DETAILED DESCRIPTION

[0009] Before the present methods and systems are disclosed and described, it is to be understood that the methods and systems are not limited to specific methods, specific components, or to particular configurations. It is also to be understood that the terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting.

[0010] As used in the specification and the appended claims, the singular forms "a," "an" and "the" include plural referents unless the context clearly dictates otherwise. Ranges may be expressed herein as from "about" one particular value, and/or to "about" another particular value. When such a range is expressed, another embodiment includes from the one particular value and/or to the other particular value. Similarly, when values are expressed as approximations, by use of the antecedent "about," it will be understood that the particular value forms another embodiment. It will be further understood that the endpoints of each of the ranges are significant both in relation to the other endpoint, and independently of the other endpoint.

[0011] "Optional" or "optionally" means that the subsequently described event or circumstance may or may not occur, and that the description includes instances where said event or circumstance occurs and instances where it does not.

[0012] Throughout the description and claims of this specification, the word "comprise" and variations of the word, such as "comprising" and "comprises," means "including but not limited to," and is not intended to exclude, for example, other additives, components, integers or steps. "Exemplary" means "an example of" and is not intended to convey an indication of a preferred or ideal embodiment. "Such as" is not used in a restrictive sense, but for explanatory purposes.

[0013] Disclosed are components that can be used to perform the disclosed methods and systems. These and other components are disclosed herein, and it is understood that when combinations, subsets, interactions, groups, etc. of these components are disclosed that while specific reference of each various individual and collective combinations and permutation of these may not be explicitly disclosed, each is specifically contemplated and described herein, for all methods and systems. This applies to all aspects of this application including, but not limited to, steps in disclosed methods. Thus, if there are a variety of additional steps that can be performed it is understood that each of these additional steps can be performed with any specific embodiment or combination of embodiments of the disclosed methods.

[0014] The present methods and systems may be understood more readily by reference to the following detailed description of preferred embodiments and the Examples included therein and to the Figures and their previous and following description.

[0015] The methods and systems will now be described by way of examples, reference being made to **FIGS. 1-6** of the drawings, **FIG. 1** showing the content of a database after identification of relevant data tables according to the disclosed methods, **FIG. 2** showing a sequence of steps of an embodiment of the methods provided, and **FIGS. 3-6** showing exemplary data tables.

[0016] A database, as shown in **FIG. 1**, comprises a number of data tables (Tables 1-5). Each data table contains data values of a number of data variables. For example, in Table 1 each data record contains data values of the data variables “Product”, “Price” and “Part”. If there is no specific value in a field of the data record, this field is considered to hold a NULL-value. Similarly, in Table 2 each data record contains values of the variables “Date”, “Client”, “Product” and “Number”. Typically, the data values are stored in the form of ASCII-coded strings.

[0017] The methods provided can be implemented by means of a computer program. In a first step (step 101), the program reads all data records in the database, for instance using a SELECT statement which selects all the tables

of the database, e.g. Tables 1-5 in this case. Typically, the database can be read into the primary memory of the computer.

[0018] To increase the evaluation speed, it can be preferred that each unique value of each data variable in said database is assigned a different binary code and that the data records are stored in binary-coded form (step 101). This can typically be done when the program first reads the data records from the database. For each input table, the following steps are carried out. First the column names, e.g. the variables, of the table are successively read. Every time a new data variable appears, a data structure can be instantiated for it. Then, an internal table structure can be instantiated to contain all the data records in binary form, whereupon the data records are successively read and binary-coded. For each data value, the data structure of the corresponding data variable can be checked to establish if the value has previously been assigned a binary code. If so, that binary code can be inserted in the proper place in the above-mentioned table structure. If not, the data value can be added to the data structure and assigned a new binary code, preferably the next one in ascending order, before being inserted in the table structure. In other words, for each data variable, a unique binary code can be assigned to each unique data value.

[0019] Tables 6-12 of **FIG. 3** show the binary codes assigned to different data values of some data variables that are included in the database of **FIG. 1**.

[0020] After having read all data records in the database, the program analyses the database to identify all connections between the data tables (step 102). A connection between two data tables means that these data tables have one variable in common. Different algorithms for performing such an analysis are known in the art. After the analysis all data tables are virtually connected. In **FIG. 1**, such virtual connections are illustrated by double-ended arrows (a). The virtually connected data tables should form at least one so-called snowflake structure, e.g. a branching data structure in which there is one and only one connecting path between any two data tables in the database. Thus, a snowflake structure does not contain any loops. If loops do occur among the virtually connected data tables, e.g. if two tables have more than one variable in common, a snowflake structure can in some cases still be formed by means of spe-

cial algorithms known in the art for resolving such loops.

[0021] After this initial analysis, the user can start to explore the database. In doing so, the user defines a mathematical function, which could be a combination of mathematical expressions (step 103). Assume that the user wants to extract the total sales per year and client from the database in **FIG. 1**. The user defines a corresponding mathematical function “SUM(x*y)”, and selects the calculation variables to be included in this function: “Price” and “Number”. The user also selects the classification variables: “Client” and “Year”.

[0022] The computer program then identifies all relevant data tables (step 104), e.g. all data tables containing any one of the selected calculation and classification variables, such data tables being denoted boundary tables, as well as all intermediate data tables in the connecting path(s) between these boundary tables in the snowflake structure, such data tables being denoted connecting tables. For the sake of clarity, the group of relevant data tables (Tables 1-3) can be included in a first frame (A) in **FIG. 1**. Evidently, there are no connecting tables in this particular case.

[0023] In the present case, all occurrences of every value, e.g. frequency data, of the selected calculation variables must be included for evaluation of the mathematical function. In **FIG. 1**, the selected variables (“Price”, “Number”) requiring such frequency data are indicated by bold arrows (b), whereas remaining selected variables are indicated by dotted lines (b'). Now, a subset (B) can be defined that includes all boundary tables (Tables 1-2) containing such calculation variables and any connecting tables between such boundary tables in the snowflake structure. It should be noted that the frequency requirement of a particular variable can be determined by the mathematical expression in which it is included. Determination of an average or a median calls for frequency information. In general, the same is true for determination of a sum, whereas determination of a maximum or a minimum does not require frequency data of the calculation variables. It can also be noted that classification variables in general do not require frequency data.

[0024] Then, a starting table can be elected, preferably among the data tables

within subset (B), most preferably the data table with the largest number of data records in this subset (step 105). In **FIG. 1**, Table 2 can be elected as the starting table. Thus, the starting table contains selected variables ("Client", "Number"), and connecting variables ("Date", "Product"). These connecting variables link the starting table (Table 2) to the boundary tables (Tables 1 and 3).

[0025] Thereafter, a conversion structure can be built (step 106), as shown in Tables 13 and 14 of **FIG. 4**. This conversion structure can be used for translating each value of each connecting variable ("Date", "Product") in the starting table (Table 2) into a value of a corresponding selected variable ("Year", "Price") in the boundary tables (Table 3 and 1, respectively). Table 13 can be built by successively reading data records of Table 3 and creating a link between each unique value of the connecting variable ("Date") and a corresponding value of the selected variable ("Year"). It can be noted that there is no link from value 4 ("Date:1999-01-12"), since this value is not included in the boundary table. Similarly, Table 14 can be built by successively reading data records of Table 1 and creating a link between each unique value of the connecting variable ("Product") and a corresponding value of the selected variable ("Price"). In this case, value 2 ("Product: Toothpaste") can be linked to two values of the selected variable ("Price: 6.5"), since this connection occurs twice in the boundary table. Thus, frequency data can be included in the conversion structure. Also note that there is no link from value 3 ("Product: Shampoo").

[0026] When the conversion structure has been built, a virtual data record can be created. Such a virtual data record, as shown in Table 15, accommodates all selected variables ("Client", "Year", "Price", "Number") in the database. In building the virtual data record (steps 107-108), a data record can be first read from the starting table (Table 2). Then, the value of each selected variable ("Client", "Number") in the current data record of the starting table can be incorporated in the virtual data record. Also, by using the conversion structure (Tables 13-14) each value of each connecting variable ("Date", "Product") in the current data record of the starting table can be converted into a value of a corresponding selected variable ("Year", "Price"), this value also being incorporated in the virtual data record.

[0027] At this stage (step 109), the virtual data record can be used to build an intermediate data structure (Table 16). Each data record of the intermediate data structure accommodates each selected classification variable (dimension) and an aggregation field for each mathematical expression implied by the mathematical function. The intermediate data structure (Table 16) can be built based on the values of the selected variables in the virtual data record. Thus, each mathematical expression can be evaluated based on one or more values of one or more relevant calculation variables in the virtual data record, and the result can be aggregated in the appropriate aggregation field based on the combination of current values of the classification variables ("Client", "Year").

[0028] The above procedure can be repeated for all data records of the starting table (step 110). Thus, an intermediate data structure can be built by successively reading data records of the starting table, by incorporating the current values of the selected variables in a virtual data record, and by evaluating each mathematical expression based on the content of the virtual data record. If the current combination of values of classification variables in the virtual data record is new, a new data record can be created in the intermediate data structure to hold the result of the evaluation. Otherwise, the appropriate data record can be rapidly found, and the result of the evaluation can be aggregated in the aggregation field. Thus, data records are added to the intermediate data structure as the starting table can be traversed. Preferably, the intermediate data structure can be a data table associated with an efficient index system, such as an AVL or a hash structure. In most cases, the aggregation field can be implemented as a summation register, in which the result of the evaluated mathematical expression can be accumulated. In some cases, e.g. when evaluating a median, the aggregation field is instead implemented to hold all individual results for a unique combination of values of the specified classification variables. It should be noted that in some aspects only one virtual data record is needed in the procedure of building the intermediate data structure from the starting table. Thus, the content of the virtual data record can be updated for each data record of the starting table. This will minimize the memory requirement in executing the computer program.

[0029] The procedure of building the intermediate data structure will be further described with reference to Tables 15-16. In creating the first virtual data record R1, as shown in Table 15, the values of the selected variables “Client” and “Number” are directly taken from the first data record of the starting table (Table 2). Then, the value “1999-01-02” of the connecting variable “Date” can be transferred into the value “1999” of the selected variable “Year”, by means of the conversion structure (Table 13). Similarly, the value “Toothpaste” of the connecting variable “Product” can be transferred into the value “6.5” of the selected variable “Price” by means of the conversion structure (Table 14), thereby forming the virtual data record R1. Then, a data record can be created in the intermediate data structure, as shown in Table 16. In this case, the intermediate data structure has tree columns, two of which holds selected classification variables (“Client”, “Year”). The third column holds an aggregation field, in which the evaluated result of the mathematical expression (“x*y”) operating on the selected calculation variables (“Number”, “Price”) can be aggregated. In evaluating virtual data record R1, the current values (binary codes: 0,0) of the classification variables are first read and incorporated in this data record of the intermediate data structure. Then, the current values (binary codes: 2,0) of the calculation variables are read. The mathematical expression can be evaluated for these values and added to the associated aggregation field.

[0030] Next, the virtual data record can be updated based on the starting table. Since the conversion structure (Table 14) indicates a duplicate of the value “6.5” of the selected variable “Price” for the value “Toothpaste” of the connecting variable “Product”, the updated virtual data record R2 is unchanged and identical to R1. Then, the virtual data record R2 can be evaluated as described above. In this case, the intermediate data structure contains a data record corresponding to the current values (binary codes: 0,0) of the classification variables. Thus, the evaluated result of the mathematical expression can be accumulated in the associated aggregation field.

[0031] Next, the virtual data record can be updated based on the second data record of starting table. In evaluating this updated virtual data record R3, a new data record can be created in the intermediate data structure, and so on.

[0032] It should be noted that NULL values are represented by a binary code of -2 in this example. In the illustrated example, it should also be noted that any virtual data records holding a NULL value (-2) of any one of the calculation variables can be directly eliminated, since NULL values cannot be evaluated in the mathematical expression ("x*y"). It should also be noted that all NULL values (-2) of the classification variables are treated as any other valid value and are placed in the intermediate data structure.

[0033] After traversing the starting table, the intermediate data structure contains four data records, each including a unique combination of values (0,0; 1,0; 2,0; 3,-2) of the classification variables, and the corresponding accumulated result (41; 37.5; 60; 75) of the evaluated mathematical expression.

[0034] Preferably, in some aspects the intermediate data structure can also be processed to eliminate one or more classification variables (dimensions). Preferably, this can be done during the process of building the intermediate data structure, as described above. Every time a virtual data record is evaluated, additional data records can be created, or found if they already exist, in the intermediate data structure. Each of these additional data records can be destined to hold an aggregation of the evaluated result of the mathematical expression for all values of one or more classification variables. Thus, when the starting table has been traversed, the intermediate data structure will contain both the aggregated results for all unique combinations of values of the classification variables, and the aggregated results after elimination of each relevant classification variable.

[0035] This procedure of eliminating dimensions in the intermediate data structure will be further described with reference to Tables 15 and 16. When virtual data record R1 is evaluated (Table 15) and the first data record (0,0) is created in the intermediate data structure, additional data records can be created in this structure. Such additional data records are destined to hold the corresponding results when one or more dimensions are eliminated. In Table 16, a classification variable can be assigned a binary code of -1 in the intermediate data structure to denote that all values of this variable are evaluated. In this case, three additional data records are created, each holding a new combina-

tion of values (-1,0; 0,-1; -1,-1) of the classification variables. The evaluated result can be aggregated in the associated aggregation field of these additional data records. The first (-1,0) of these additional data records is destined to hold the aggregated result for all values of the classification variable "Client" when the classification variable "Year" has the value "1999". The second (0,-1) additional data record is destined to hold the aggregated result for all values of the classification variable "Year" when the classification variable "Client" is "Nisse". The third (-1,-1) additional data record is destined to hold the aggregated result for all values of both classification variables "Client" and "Year".

[0036] When virtual data record R2 is evaluated, the result can be aggregated in the aggregation field associated with the current combination of values (binary codes: 0,0) of the classification variables, as well as in the aggregation fields associated with relevant additional data records (binary codes: -1,0; 0,-1; -1,-1). When virtual data record R3 is evaluated, the result can be aggregated in the aggregation field associated with the current combination of values (binary codes: 1,0) of the classification variables. The result can also be aggregated in the aggregation field of a newly created additional data record (binary codes: 1,-1) and in the aggregation fields associated with relevant existing data records (binary codes: -1,0; -1,-1) in the intermediate data structure.

[0037] After traversing the starting table, the intermediate data structure contains eleven data records, as shown in Table 16. Preferably, if the intermediate data structure accommodates more than two classification variables, the intermediate data structure will, for each eliminated classification variable, contain the evaluated results aggregated over all values of this classification variable for each unique combination of values of remaining classification variables.

[0038] When the intermediate data structure has been built, a final data structure, e.g. a multidimensional cube, as shown in non-binary notation in Table 17 of **FIG. 5**, can be created by evaluating the mathematical function ("SUM (x*y)") based on the results of the mathematical expression ("x*y") contained in the intermediate data structure (step 111). In doing so, the results in the aggregation fields for each unique combination of values of the classification variables are combined. In the illustrated case, the creation of the final data structure is

straightforward. The content of the final data structure might then (step 112) be presented to the user in a two-dimensional table, as shown in Table 18 of **FIG. 5**. Alternatively, if the final data structure contains many dimensions, the data can be presented in a pivot table, in which the user interactively can move up and down in dimensions, as is well known in the art.

[0039] Below, a second example of the disclosed methods is described with reference to Tables 20-29 of **FIGS. 5-6**. The description will only elaborate on certain aspects of this example, namely building a conversion structure including data from connecting tables, and building an intermediate data structure for a more complicated mathematical function. In this example, the user wants to extract sales data per client from a database, which contains the data tables shown in Tables 20-23 of **FIG. 5**. For ease of interpretation, the binary coding is omitted in this example.

[0040] The user has specified the following mathematical functions, for which the result should be partitioned per Client: a) “IF(Only(Environment index)=`I`) THEN Sum(Number*Price)*2, ELSE Sum(Number*Price))”, and b) “Avg(Number*Price)”

[0041] The mathematical function (a) specifies that the sales figures should be doubled for products that belong to a product group having an environment index of `I`, while the actual sales figures should be used for other products. The mathematical function (b) has been included for reference.

[0042] In this case, the selected classification variables are “Environment index” and “Client”, and the selected calculation variables are “Number” and “Price”. Tables 20, 22 and 23 are identified as boundary tables, whereas Table 21 is identified as a connecting table. Table 20 is elected as starting table. Thus, the starting table contains selected variables (“Number”, “Client”), and a connecting variable (“Product”). The connecting variable links the starting table (Table 20) to the boundary tables (Tables 22-23), via the connecting table (Table 21).

[0043] Next, the formation of the conversion structure will be described with reference to Tables 24-26 of **FIG. 6**. A first part (Table 24) of the conversion

structure can be built by successively reading data records of a first boundary table (Table 23) and creating a link between each unique value of the connecting variable ("Product group") and a corresponding value of the selected variable ("Environment index"). Similarly, a second part (Table 25) of the conversion structure can be built by successively reading data records of a second boundary table (Table 22) and creating a link between each unique value of the connecting variable ("Price group") and a corresponding value of the selected variable ("Price"). Then, data records of the connecting table (Table 21) are read successively. Each value of the connecting variables ("Product group" and "Price group", respectively) in Tables 24 and 25 can be substituted for a corresponding value of a connecting variable ("Product") in Table 21. The result is merged in one final conversion structure, as shown in Table 26.

[0044] Then, an intermediate data structure can be built by successively reading data records of the starting table (Table 20), by using the conversion structure (Table 26) to incorporate the current values of the selected variables ("Environment index", "Client", "Number", "Price") in the virtual data record, and by evaluating each mathematical expression based on the current content of the virtual data record.

[0045] For reasons of clarity, Table 27 displays the corresponding content of the virtual data record for each data record of the starting table. As noted in connection with the first example, in some aspects only one virtual data record is needed. The content of this virtual data record can be updated, e.g. replaced, for each data record of the starting table.

[0046] Each data record of the intermediate data structure, as shown in Table 28, accommodates a value of each selected classification variable ("Client", "Environment index") and an aggregation field for each mathematical expression implied by the mathematical functions. In this case, the intermediate data structure contains two aggregation fields. One aggregation field contains the aggregated result of the mathematical expression ("x*y") operating on the selected calculation variables ("Number", "Price"), as well as a counter of the number of such operations. The layout of this aggregation field is given by the fact that an average quantity should be calculated ("Avg(x*y)"). The other ag-

gregation field can be designed to hold the lowest and highest values of the classification variable “Environment index” for each combination of values of the classification variables.

[0047] As in the first example, the intermediate data structure (Table 28) can be built by evaluating the mathematical expression for the current content of the virtual data record (each row in Table 27), and by aggregating the result in the appropriate aggregation field based on the combination of current values of the classification variables (“Client”, “Environment index”). The intermediate data structure also includes data records in which the value “<ALL>“ has been assigned to one or both of the classification variables. The corresponding aggregation fields contain the aggregated result when the one or more classification variables (dimensions) are eliminated.

[0048] When the intermediate data structure has been built, a final data structure, e.g. a multidimensional cube, can be created by evaluating the mathematical functions based on the evaluated results of the mathematical expressions contained in the intermediate data structure. Each data record of the final data structure, as shown in Table 29, accommodates a value of each selected classification variable (“Client”, “Environment index”) and an aggregation field for each mathematical function selected by the user.

[0049] The final data structure can be built based on the results in the aggregation fields of the intermediate data structure for each unique combination of values of the classification variables. When function (a) is evaluated, by sequentially reading data records of Table 28, the program first checks if both values in the last column of Table 28 is equal to `1`. If so, the relevant result contained in the first aggregation field of Table 28 is multiplied by two and stored in Table 29. If not, the relevant result contained in the first aggregation field of Table 28 is directly stored in Table 29. When function (b) is evaluated, the aggregated result of the mathematical expression (“x*y”) operating on the selected calculation variables (“Number”, “Price”) is divided by the number of such operations, both of which are stored in the first aggregation field of Table 28. The result can be stored in the second aggregation field of Table 29.

[0050] Evidently, the present methods allow the user to freely select mathematical functions and incorporate calculation variables in these functions as well as to freely select classification variables for presentation of the results.

[0051] As an alternative, albeit less memory-efficient, to the illustrated procedure of building an intermediate data structure based on sequential data records from the starting table, it is conceivable to first build a so-called join table. This join table can be built by traversing all data records of the starting table and, by use of the conversion structure, converting each value of each connecting variable in the starting table into a value of at least one corresponding selected variable in a boundary table. Thus, the data records of the join table will contain all occurring combinations of values of the selected variables. Then, the intermediate data structure can be built based on the content of the join table. For each record of the join table, each mathematical expression can be evaluated and the result can be aggregated in the appropriate aggregation field based on the current value of each selected classification variable. However, this alternative procedure requires more computer memory to extract the requested information.

[0052] It should be realized that the mathematical function could contain mathematical expressions having different, and conflicting, needs for frequency data. In this case, steps 104 110 (**FIG. 2**) are repeated for each such mathematical expression, and the results are stored in one common intermediate data structure. Alternatively, one final data structure, e.g. multidimensional cube, could be built for each mathematical expression, the contents of these cubes being fused during presentation to the user.

[0053] As will be appreciated by one skilled in the art, the methods and systems may take the form of an entirely hardware embodiment, an entirely software embodiment, or an embodiment combining software and hardware aspects. Furthermore, the methods and systems may take the form of a computer program product on a computer-readable storage medium having computer-readable program instructions (e.g., computer software) embodied in the storage medium. More particularly, the present methods and systems may take the form of web-implemented computer software. Any suitable computer-readable

storage medium may be utilized including hard disks, CD-ROMs, optical storage devices, or magnetic storage devices.

[0054] Embodiments of the methods and systems are described with reference to block diagrams and flowchart illustrations of methods, systems, apparatuses and computer program products. It will be understood that each block of the block diagrams and flowchart illustrations, and combinations of blocks in the block diagrams and flowchart illustrations, respectively, can be implemented by computer program instructions. These computer program instructions may be loaded onto a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions which execute on the computer or other programmable data processing apparatus create a means for implementing the functions specified in the flowchart block or blocks.

[0055] These computer program instructions may also be stored in a computer-readable memory that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable memory produce an article of manufacture including computer-readable instructions for implementing the function specified in the flowchart block or blocks. The computer program instructions may also be loaded onto a computer or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer or other programmable apparatus to produce a computer-implemented process such that the instructions that execute on the computer or other programmable apparatus provide steps for implementing the functions specified in the flowchart block or blocks.

[0056] Accordingly, blocks of the block diagrams and flowchart illustrations support combinations of means for performing the specified functions, combinations of steps for performing the specified functions and program instruction means for performing the specified functions. It will also be understood that each block of the block diagrams and flowchart illustrations, and combinations of blocks in the block diagrams and flowchart illustrations, can be implemented by special purpose hardware-based computer systems that perform the speci-

fied functions or steps, or combinations of special purpose hardware and computer instructions.

[0057] One skilled in the art will appreciate that provided is a functional description and that respective functions can be performed by software, hardware, or a combination of software and hardware. In an aspect, the methods and systems can comprise the Data Analysis Software **106** as illustrated in **FIG. 7** and described below. In one exemplary aspect, the methods and systems can comprise a computer **101** as illustrated in **FIG. 7** and described below.

[0058] **FIG. 7** is a block diagram illustrating an exemplary operating environment for performing the disclosed methods. This exemplary operating environment is only an example of an operating environment and is not intended to suggest any limitation as to the scope of use or functionality of operating environment architecture. Neither should the operating environment be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment.

[0059] The present methods and systems can be operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that can be suitable for use with the systems and methods comprise, but are not limited to, personal computers, server computers, laptop devices, and multiprocessor systems. Additional examples comprise set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that comprise any of the above systems or devices, and the like.

[0060] The processing of the disclosed methods and systems can be performed by software components. The disclosed systems and methods can be described in the general context of computer-executable instructions, such as program modules, being executed by one or more computers or other devices. Generally, program modules comprise computer code, routines, programs, objects, components, data structures, etc. that perform particular tasks or imple-

ment particular abstract data types. The disclosed methods can also be practiced in grid-based and distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules can be located in both local and remote computer storage media including memory storage devices.

[0061] Further, one skilled in the art will appreciate that the systems and methods disclosed herein can be implemented via a general-purpose computing device in the form of a computer **701**. The components of the computer **701** can comprise, but are not limited to, one or more processors or processing units **703**, a system memory **712**, and a system bus **713** that couples various system components including the processor **703** to the system memory **712**. In the case of multiple processing units **703**, the system can utilize parallel computing.

[0062] The system bus **713** represents one or more of several possible types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, such architectures can comprise an Industry Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA) local bus, an Accelerated Graphics Port (AGP) bus, and a Peripheral Component Interconnects (PCI), a PCI-Express bus, a Personal Computer Memory Card Industry Association (PCMCIA), Universal Serial Bus (USB) and the like. The bus **713**, and all buses specified in this description can also be implemented over a wired or wireless network connection and each of the subsystems, including the processor **703**, a mass storage device **704**, an operating system **705**, Data Analysis software **706**, data **707**, a network adapter **708**, system memory **712**, an Input/Output Interface **710**, a display adapter **709**, a display device **711**, and a human machine interface **702**, can be contained within one or more remote computing devices **714a,b,c** at physically separate locations, connected through buses of this form, in effect implementing a fully distributed system.

[0063] The computer **701** typically comprises a variety of computer readable media. Exemplary readable media can be any available media that is accessible by the computer **701** and comprises, for example and not meant to be limiting, both volatile and non-volatile media, removable and non-removable media. The system memory **712** comprises computer readable media in the form of volatile memory, such as random access memory (RAM), and/or non-volatile memory, such as read only memory (ROM). The system memory **712** typically contains data such as data **707** and/or program modules such as operating system **705** and Data Analysis software **706** that are immediately accessible to and/or are presently operated on by the processing unit **703**.

[0064] In another aspect, the computer **701** can also comprise other removable/non-removable, volatile/non-volatile computer storage media. By way of example, **FIG. 7** illustrates a mass storage device **704** which can provide non-volatile storage of computer code, computer readable instructions, data structures, program modules, and other data for the computer **701**. For example and not meant to be limiting, a mass storage device **704** can be a hard disk, a removable magnetic disk, a removable optical disk, magnetic cassettes or other magnetic storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or other optical storage, random access memories (RAM), read only memories (ROM), electrically erasable programmable read-only memory (EEPROM), and the like.

[0065] Optionally, any number of program modules can be stored on the mass storage device **704**, including by way of example, an operating system **705** and Data Analysis software **706**. In aspect, one or more of the methods described can be implemented by the Data Analysis software **706** and the processor **703**. In an aspect, the Data Analysis software **706** can comprise a visualization component for dynamically generating one or more visual representations data **707**. Each of the operating system **705** and Data Analysis software **706** (or some combination thereof) can comprise elements of the programming and the Data Analysis software **706**. **D**ata **707** can also be stored on the mass storage device **704**. Data **707** can be stored in any of one or more databases known in the art. Examples of such databases comprise, DB2®, Microsoft® Access, Mi-

crosoft® SQL Server, Oracle®, mySQL, PostgreSQL, and the like. The databases can be centralized or distributed across multiple systems.

[0066] In another aspect, the user can enter commands and information into the computer **701** via an input device (not shown). Examples of such input devices comprise, but are not limited to, a keyboard, pointing device (*e.g.*, a "mouse"), a microphone, a joystick, a scanner, tactile input devices such as gloves, and other body coverings, and the like These and other input devices can be connected to the processing unit **703** via a human machine interface **702** that is coupled to the system bus **713**, but can be connected by other interface and bus structures, such as a parallel port, game port, an IEEE 1394 Port (also known as a Firewire port), a serial port, or a universal serial bus (USB).

[0067] In yet another aspect, a display device **711** can also be connected to the system bus **713** via an interface, such as a display adapter **709**. It is contemplated that the computer **701** can have more than one display adapter **709** and the computer **701** can have more than one display device **711**. For example, a display device can be a monitor, an LCD (Liquid Crystal Display), or a projector. In addition to the display device **711**, other output peripheral devices can comprise components such as speakers (not shown) and a printer (not shown) which can be connected to the computer **701** via Input/Output Interface **710**. Any step and/or result of the methods can be output in any form to an output device. Such output can be any form of visual representation, including, but not limited to, textual, graphical, animation, audio, tactile, and the like.

[0068] The computer **701** can operate in a networked environment using logical connections to one or more remote computing devices **714a,b,c** through network **715**. By way of example, a remote computing device can be a personal computer, portable computer, a server, a router, a network computer, a peer device or other common network node, and so on. Logical connections between the computer **701** and a remote computing device **714a,b,c** can be made via a local area network (LAN) and a general wide area network (WAN). Such network connections can be through a network adapter **708**. A network adapter **708** can be implemented in both wired and wireless environments. Such networking environments are conventional and commonplace in offices,

enterprise-wide computer networks, intranets, and the Internet.

[0069] For purposes of illustration, application programs and other executable program components such as the operating system **705** are illustrated herein as discrete blocks, although it is recognized that such programs and components reside at various times in different storage components of the computing device **701**, and are executed by the data processor(s) of the computer. An implementation of Data Analysis software **706** can be stored on or transmitted across some form of computer readable media. Any of the disclosed methods can be performed by computer readable instructions embodied on computer readable media. Computer readable media can be any available media that can be accessed by a computer. By way of example and not meant to be limiting, computer readable media can comprise “computer storage media” and “communications media.” “Computer storage media” comprise volatile and non-volatile, removable and non-removable media implemented in any methods or technology for storage of information such as computer readable instructions, data structures, program modules, or other data. Exemplary computer storage media comprises, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by a computer.

[0070] The methods and systems can employ Artificial Intelligence techniques such as machine learning and iterative learning. Examples of such techniques include, but are not limited to, expert systems, case based reasoning, Bayesian networks, behavior based AI, neural networks, fuzzy systems, evolutionary computation (e.g. genetic algorithms), swarm intelligence (e.g. ant algorithms), and hybrid intelligent systems (e.g. Expert inference rules generated through a neural network or production rules from statistical learning).

[0071] In an aspect, the processor **703** can be configured for performing steps comprising, initiating a primary session for a first user, requesting collaboration from a second user, initiating a secondary session for the second user, and

providing a single state space for collaborative real-time data analysis to the first user and the second user wherein an interaction by either user is reflected in the single state space. Initiating a primary session for a first user can comprise generating a first user interface. Requesting collaboration from a second user can comprise providing the second user a URL configured to allow the second user to view the secondary session. Initiating a primary session for the second user can comprise generating a second user interface. The single state space can be a view of a dataset.

[0072] The primary and secondary sessions can result in a collaboration session that can be connected to a plurality of XML transformers connected via synchronization logic. An interaction performed by any of the plurality of XML transformers can be propagated back to the other XML transformers in the plurality of XML transformers.

[0073] In a further aspect, the processor **703** can be configured for performing steps comprising, creating a state space that reflects a selection state, creating a note, and attaching the note to an object in the state space. The state space can be a view of a dataset. The note can comprise an observation about the object. The processor can be further configured to save the selection state and associate the saved selection state with the note. The saved selection state reflects the selection state at the time the note is created. The note can comprise a note thread wherein the note thread comprises a plurality of notes. The processor can be further configured to save a plurality of selection states wherein each of the plurality of selection states is associated with one or more of the plurality of notes. The processor can be further configured to identify the object as having an attached note.

[0074] In a further aspect, the processor **703** can be configured for performing steps comprising, presenting an object in a state space having an attached note, receiving a selection of the note, and presenting the note and adjusting the state space to reflect a saved selection state associated with the note. The note can comprise an observation about the object. The state space can be a first view of a dataset. The saved selection state can be a second view of the dataset. The saved selection state can be saved contemporaneously with crea-

tion of the note. The note can comprise a note thread wherein the note thread comprises a plurality of notes wherein one or more of the plurality of notes has an associated saved selection state.

[0075] The methods and systems described above enable real-time associative data mining and visualization. In an aspect, the methods and systems can manage associations among data sets with every data point in the analytic dataset being associated with every other data point in the dataset. Datasets can be hundreds of tables with thousands of fields.

[0076] In an aspect, provided are methods and systems for user interaction with the database methods and systems disclosed. In an aspect, a user interface can be generated to facilitate dynamic display generation to view data. By way of example, a particular view of a particular dataset or data subset generated for a user can be referred to as a state space or a session. The system can comprise a visualization component to dynamically generate one or more visual representations of the data to present in the state space.

[0077] **FIG. 8** illustrates how a Selection operates on a Scope to generate a Data Subset. The Data subset can form a state space, which is based on a selection state given by the Selection. In an aspect, the selection state (or "user state") can be defined by a user clicking on list boxes and graphs in a user interface of an application. An application can be designed to host a number of graphical objects (charts, tables, etc) that evaluate one or more mathematical functions (also referred to as an "expression") on the Data subset for one or more dimensions (classification variables). The result of this evaluation creates a Chart result which is a multidimensional cube which can be visualized in one or more of the graphical objects.

[0078] The application can permit a user to explore the Scope by making different selections, by clicking on graphical objects to select variables, which causes the Chart result to change. At every time instant during the exploration, there exists a current state space, which can be associated with a current selection state that is operated on the Scope (which always remains the same).

[0079] As illustrated in **FIG. 8**, when a user makes a new selection, an infer-

ence engine calculates a data subset. Also, the identifier ID1 for the selection together with the scope can be generated based on the filters in the selection and the scope. Subsequently, the identifier ID2 for the data subset is generated based on the data subset definition, typically a bit sequence that defines the content of the data subset. Finally, ID2 can be put into a cache using ID1 as lookup identifier. Likewise, the data subset definition is put in the cache using ID2 as lookup identifier.

[0080] In **FIG. 8**, the chart calculation takes place in a similar way. Here, there are two information sets: the data subset and the relevant chart properties. The latter is typically, but not restricted to, a mathematical function together with calculation variables and classification variables (dimensions). Both of these information sets are used to calculate the chart result, and both of these information sets are also used to generate the identifier ID3 for the input to the chart calculation. ID2 was generated already in the previous step, and ID3 is generated as the first step in the chart calculation procedure.

[0081] The identifier ID3 is formed from ID2 and the relevant chart properties. ID3 can be seen as an identifier for a specific chart generation instance, which includes all information needed to calculate a specific chart result. In addition, a chart result identifier ID4 is created from the chart result definition, typically a bit sequence that defines the chart result. Finally, ID4 is put in the cache using ID3 as lookup identifier. Likewise, the chart result definition is put in the cache using ID4 as lookup identifier.

[0082] The graphical objects (or visual representations) can be substantially any display or output type including graphs, charts, trees, multi-dimensional depictions, images (computer generated or digital captures), video/audio displays describing the data, hybrid presentations where output is segmented into multiple display areas having different data analysis in each area and so forth. A user can select one or more default visual representations, however, a subsequent visual representation can be generated based off of further analysis and subsequent dynamic selection of the most suitable form for the data. As shown in **FIG. 9** several list boxes are provided on the left side of the interface and graphical objects reflecting selections (or lack of selection) in the list boxes

are displayed on the right side of the user interface. Placement of list boxes and graphical objects is a matter of design choice. In an aspect, a user can select a data point and a visualization component can instantaneously filter and re-aggregate other fields and corresponding visual representations based on the user's selection. In an aspect, the filtering and re-aggregation can be completed without querying a database. In an aspect, a visual representation can be presented to a user with color schemes applied meaningfully. For example, a user selection can be highlighted in green, datasets related to the selection can be highlighted in white, and unrelated data can be highlighted in gray. A meaningful application of a color scheme provides an intuitive navigation interface in the state space.

[0083] As shown in **FIG. 10a**, a layout including several graphical objects is provided to a user. The dataset reflects movie data. For example, movie directors, movie titles, movie actors, movie length, movie rating, movie release date, and the like. As shown in **FIG. 10b**, once the user selects a director, the graphical objects dynamically adjust in real-time. In this example, the user has selected the director "Emeric Pressburger." In response to the selection, all of the graphical objects adjust to reflect data having a relationship to "Emeric Pressburger."

[0084] Thus, the methods and systems provided enable a user to instantiate a session that enables the transformation of raw data into actionable analytics. While a single user can manipulate the interface to generate meaningful visual representations, also provided are methods and systems that facilitate collaborative sessions wherein multiple users can manipulate the interface at the same time or substantially the same time.

[0085] In an aspect, a user can share their session with one or more other users. As a result, the users can discover and develop new analyses in a real-time, collaborative environment. Each user can make selections that can be seen by all users. In some cases, restrictions can be implemented so that only some users can make selections. In a further example, transient lists (for example, searches, drop-downs, and the like) of a user can be hidden from other users.

[0086] In an aspect, two or more users can share a common session. The first time the session is generated is referred to as the primary session; while subsequent users who join are referred to as secondary sessions. In an aspect, only the primary session can invite others to join, while in another aspect, any user can invite others to join. The system can be configured such that all aspects of the secondary session mirror those of the primary session. If the primary session has section access reductions, these are mirrored in secondary sessions. Section access reductions can be a mechanism that provides data security. For example, when a user clicks on a list box, the user may be restricted to viewing a reduced amount of data versus another user with superior section access rights. For example, one user may be able to view all movie directors, whereas another user can only view one movie director. In an aspect, no checks on access rights or data security are applied to secondary sessions.

[0087] All users, primary and secondary, can share interactions with a user interface (for example, mouse clicks) that interact with the system. Any user who clicks, where that click changes a selection state, that change in state can be sent to one or more of the other clients. Any click that only affects the local client, and does not involve a message/response from the server is not shared. In the case that two or more clients click “at the same time” the server can treat each click as two or more asynchronous clicks, the same as if a single client had clicked once, and then clicked a second time canceling the first click.

[0088] In an aspect, the primary user can invite secondary users to join his/her session using a panel that drops down from the collaboration toolbar icon. Email invitations can permit the primary user to specify an email address, and some additional text that can be placed into the email body. When an “invite” button is pressed, an email can be sent to the recipient with a standard message, any additional message included by the primary user, and a URL to join the session.

[0089] An invitation to join a session can be performed using a specially formatted URL. This URL can provide a link back to the system, and the specific interface workspace. In addition, the URL can provide an additional parameter that is a one-time use key for identifying and joining the appropriate session.

Once this URL has been clicked on (e.g. sent to the server) it can be invalidated, so it can only be used once, and cannot be forwarded.

[0090] The primary user can be notified when a secondary user joins the session. This notification can be a change in state (for example, changing color) of a collaboration toolbar icon and a message connected to that toolbar icon indicating who has joined the session. Once a secondary user has joined the session, one or more other users can view a list of users currently sharing the session, and in some aspects, remove users.

[0091] In another aspect, the primary user can invite secondary users to join his/her session using a panel that drops down from the collaboration toolbar icon. An additional option for inviting secondary users is by searching user directories that are accessible to the system. A primary user can use the directory search results to invite users directly.

[0092] In an aspect, illustrated in **FIG. 11a**, provided are methods, systems, and computer readable media for collaborative computing comprising, initiating a primary session for a first user at 1101, requesting collaboration from a second user at 1102, initiating a secondary session for the second user at 1103, and providing a single state space for collaborative real-time data analysis to the first user and the second user wherein an interaction by either user is reflected in the single state space at 1104.

[0093] Initiating a primary session for a first user can comprise generating a first user interface. Requesting collaboration from a second user can comprise providing the second user a URL configured to allow the second user to view the secondary session. Initiating a primary session for the second user can comprise generating a second user interface. The first and second user interfaces can provide the users with graphical objects. The single state space can be a view of a dataset. The primary and secondary sessions can result in a collaboration session that can be connected to a plurality of XML transformers connected via synchronization logic. An interaction performed by any of the plurality of XML transformers can be propagated back to the other XML transformers in the plurality of XML transformers.

[0094] In an aspect, illustrated in **FIG. 11b**, a collaboration session can comprise a single low-level shared session that can be connected to two or more higher level XML transformers. The XML transformers can be connected via synchronization logic. Each XML transformer can be attached to an end-point of a web session and the other end-point can be connected to a web browser. Commands and selections performed by any of the XML transformers can thus affect the shared low-level session and state changes can be propagated back to both XML transformers. The XML transformer that performed the command can return the state change to the client. The other XML transformer can return the changed state through the client asynchronous mechanism.

[0095] In a further aspect, provided are methods and systems for time shifted collaboration. Within a single state space, users can create and share notes about various objects contained within the state space. These notes can be shared with one or more other users, and these other users can respond by leaving their own note comments. Each user can save a "snapshot" (bookmark) of the state space and data with each note. The notes can be searchable by users to efficient access to the note and the associated snapshot of the state space.

[0096] **FIG. 12a** illustrates a graphical object with an attached note and the note thread that can be viewed after selection of the note. **FIG12b** illustrates the change in the state space after selection of the saved selection state associated with the note.

[0097] By way of example, a user can right-click an object displayed in the state space, providing the user with a menu option to add a new note and to view existing notes, by selecting "Notes" from the context menu. Optionally, all objects in the state space with existing notes can be identified (for example, by an icon, a color change, and the like). Similarly, the number of attached notes for each object can be displayed. Thus, the resulting note can be linked to both an object and a selection state. An object can have one or more notes and one or more note threads (a series of comments based on a note). A user can create a note after the user has analyzed a dataset and accordingly arranged the state space. The user can select to attach a snapshot of the current state

space to the note. The system can then create a hidden bookmark and attaches it to the note. In an aspect, multiple snapshots of a state space can be associated with a note, reflecting for example a comparison of two different analyses.

[0098] To view a note and the associated state space, a user can select a desired note and the note text will be presented to the user. The user can then add additional information to the note thread and chose to apply the bookmark, modifying the current state space to reflect the state space associated with the note. In another example, the state space can automatically update to reflect the state space associated with the note upon note selection.

[0099] Permissions can be adjusted for notes to control access to the notes by various classes of users. For example, a class of users might be able to view notes, but not make notes whereas another class of users can make notes, edit notes, and delete notes.

[0100] The methods for time shifted collaboration can be implemented in various fashions. For example, the notes (either a single note or a note thread) can be linked to a specific selection state and stored in one single "bookmark." Hence, one bookmark can comprise several notes for each object. By applying the bookmark, the notes become visible. In a further example, the notes can be linked to several selection states: Each note can correspond to one specific selection state, and all following replies in a note thread can pertain to the same selection state. The selection state belonging to a specific note can be stored in a temporary, hidden bookmark. In a still further example, the notes can be linked to the raw data or the data in input fields. Hence, the notes can be seen as textual input fields.

[0101] In an aspect, illustrated in **FIG. 13a**, provided are methods, systems, and computer readable media for time shifted collaborative analysis comprising, creating a state space that reflects a selection state at 1301a, creating a note at 1302a, attaching the note to an object in the state space at 1303a, saving the selection state at 1304a, and associating the saved selection state with the note at 1305a. The state space can be a view of a dataset. The note can

comprise an observation about the object. The saved selection state can reflect the selection state at the time the note is created. The note can comprise a note thread wherein the note thread comprises a plurality of notes. The methods can further comprise saving a plurality of selection states wherein each of the plurality of selection states is associated with one or more of the plurality of notes. The methods can further comprise identifying the object as having an attached note.

[0102] In a further aspect, illustrated in **FIG. 13b**, provided are methods, systems, and computer readable media for time shifted collaborative analysis comprising, creating a state space that reflects a selection state at 1301b, creating a note at 1302b, and attaching the note to an object in the state space at 1303b. The state space can be a view of a dataset. The note can comprise an observation about the object. The methods can further comprise saving the selection state and associating the saved selection state with the note. The saved selection state can reflect the selection state at the time the note is created. The note can comprise a note thread wherein the note thread comprises a plurality of notes. The methods can further comprise saving a plurality of selection states wherein each of the plurality of selection states is associated with one or more of the plurality of notes. The methods can further comprise identifying the object as having an attached note.

[0103] In a further aspect, illustrated in **FIG. 13c**, provided are methods, systems, and computer readable media for time shifted collaborative analysis comprising, presenting an object in a state space having an attached note at 1301c, receiving a selection of the note at 1302c, and presenting the note and adjusting the state space to reflect a saved selection state associated with the note at 1303c. The note can comprise an observation about the object. The state space can be a first view of a dataset. The saved selection state can be a second view of the dataset. The saved selection state can be saved contemporaneously with creation of the note. The note can comprise a note thread wherein the note thread comprises a plurality of notes wherein one or more of the plurality of notes has an associated saved selection state.

[0104] In an aspect, the methods and systems provided allow a user to create

multiple states within a single space and apply these states to specific objects within the space. The user can create copies of these objects and then put those objects into different states. Objects in a given state are not affected by user selections in the other states. The methods and systems provided permit a user to generate graphical objects that represent different state spaces (and thus different selection states) in one view.

[0105] The use of alternate states permits simultaneous use of multiple selections within the space and enable comparisons of the selections in a single visual representation or in separate visual representations. A user can select data items for comparative analysis, and then make an overriding selection that impacts the comparative analysis in real-time. **FIG. 14** illustrates an exemplary implementation of alternate states.

[0106] The left-hand list boxes are logically associated with a state space X and are located in a state space X container, and the right-hand list boxes are logically associated with a state space Y and are located in a state space Y container. In this example, the result graph (chart) displays the results of evaluating a mathematical function (expression) in both the state space X and the state space Y. Thus, the user is able to define the state space X by clicking in the left-hand list boxes, causing the corresponding evaluation results to be displayed in the result graph. In the same way, the user is able to define the state space Y by clicking in the right-hand list boxes, causing the corresponding evaluation results to be displayed in the result graph.

[0107] Each state can be assigned a state identifier for system processing. In an aspect, at least two states can be made available, a default state and an inherited state. The default state can be the state where most usage occurs. Objects can inherit states from higher level objects, such as sheets and containers. This means that states are inherited as such: Document - Sheet - Sheet Objects. The sheets and sheet objects are always in the inherited state unless overridden. By way of example, a document can be an application document, a Sheet can be tab in such a document, and a container can be a region on a tab that may contain one or more Objects. An Object can be any textual or graphical object, e.g. a list box, a pie chart, a bar chart, etc. Sheets and

sheet objects (e.g. containers and graphical objects) are always in the inherited state, but it is possible for a user to override the inherited state for a sheet or a sheet object by associating the sheet or the sheet object with an explicit state space.

[0108] In an aspect, a lower level can automatically inherit the state space of a higher level. As shown in **FIG. 14**, if the sheet is assigned to the default state space X, all containers and individual objects on this sheet with also be associated with this state space, unless otherwise specified. Thus, the user only needs to associate containers/objects with the state space Y as desired.

[0109] Chart and other object expressions inherit the state of the object that contains the expression. Chart and object expressions can reference alternate states. This means that an expression, no matter where it occurs, can reference a different state than the object that contains the expression.

[0110] The methods and systems can use the default state to drive a subset of data on which to calculate charts and aggregations by taking the definition of the state in terms of Values selected per Field and determining a Set in terms of a subset of Rows per Table. This default behavior can be changed at two distinct points to enable alternate states: 1. Defining a set of data that is independent of current selections; and 2. Combining multiple sets through the use of mathematical operators such as Union, Intersection and Exception.

[0111] Alternate States plays a role in the first part; defining selection states from which sets can be generated. For processing purposes, the default state can be represented by “$”, while all the data, regardless of states and selections, can be represented by “1”. Alternate states introduces two additional syntax elements.

[0112] 1. An expression can be based on an alternate state.

Examples:

sum({[Group 1]} Sales)

calculates sales based on the selections in the state ‘Group 1’.

sum({$} Sales)

calculates sales based on the selections in the default state.

Both of these expressions can exist in a single chart. This allows users to compare multiple states within a single object. State references within expressions override the state of the object. **FIG. 14** may be seen as such an implementation. State space X may be the default state space (represented by $), and state space Y may be the state space "Group 1". Thus, the left-hand bars in the result graph may be given by the mathematical function Sum({$} Sales), whereas the right-hand bars in the result graph may be given by the mathematical function Sum({[Group 1]} Sales). This is an example of the fact that an expression, no matter where it occurs, can reference a different state than the object that contains the expression.

[0113] Instead of displaying the evaluation results for state spaces X and Y in one and the same result graph, they may be displayed in separate graphs. In such an example, one of the graphs would be associated with the expression Sum({[Group 1]} Sales) and the other graph with the expression Sum({$} Sales).

[0114] 2. Selections in a field in one state can be used as modifiers in another state.

Example:

sum({[Group 1]<Region = $::Region>} Sales)

This syntax uses the selections in the "Region" field from the default state and modifies the state 'Group 1' with them. The effect is to keep the Region field "synchronized" between the default state and 'Group 1' for this expression. Thus, selections in an object that is associated with a first state space (e.g. by the user clicking on a value in a list box associated with state space X) can be used to modify a second state space (e.g. state space Y) in addition to (or instead of) the first state space. In **FIG. 14**, this could be used to make sure that when the user makes a selection in a specific list box on the left-hand side, so as to modify the state space X, a corresponding modification (selection) is automatically made to the state space Y.

[0115] It is possible to use set operators (+, *, -, /) with states. The following

expressions are valid and will count the distinct invoice numbers that are in either the default state or State1.

Examples:

count({$ + State1} DISTINCT [Invoice Number])

counts the distinct invoice numbers in the union of the <default> state and State1.

count({1 - State1} DISTINCT [Invoice Number])

counts the distinct invoice numbers not in State1.

count({State1 * State2} DISTINCT [Invoice Number])

counts the distinct invoice numbers in that are in both the <default> state and State1.

[0116] Thus, the methods and systems provide a method of logically combining data in different state spaces by the use of logical operators known from Boolean algebra:

+ = UNION (A+B contains all elements of both A and B)

* = INTERSECT (A*B contains all elements of A that also belong to B)

- = DIFF (A-B contains all elements of A that do not belong to B)

/ = XOR (A/B contains all elements that are only found in one of A and B)

[0117] The use of Set Operators makes it possible to combine and evaluate data from two of more state spaces in one expression, e.g. for display in a graph.

[0118] In an aspect, illustrated in **FIG. 15**, provided are methods for data analysis comprising presenting a first user interface element associated with a first state space and a second user interface element associated with a second state space at 1501, receiving a selection in the first and second user interface elements at 1502, and presenting a result graph representing the a selection state of the first state space and a selection state of the second state space 1503. In an aspect, the first state space and the second state space can comprise the same dataset or different data sets.

[0119] In an aspect, provided are methods and systems for utilizing dimension

limits. Dimension limits can be set for various chart types. A user can be presented with a Dimension Limits option to control the number of dimension values displayed in a given chart. The user can select one of a plurality of values, for example: First, Largest, and Smallest. These values control the way the system sorts the values it returns to the visualization component. In an aspect, sorting only occurs for the first expression (except in pivot tables when a primary sort may override the first dimension sort). In an aspect, shown in **FIG. 16**, one or more user interface elements can be presented to apply one or more dimension limits. For example, a sliding selection tool can be presented to enable a user to apply the dimension limit "show only." The example in **FIG. 16** illustrates the application of the dimension limit show only the top 6 sales performers.

[0120] Dimension Limits may be applied for generating data to be displayed in a chart (graph, table etc). These Dimension Limits can comprise one or more of:

Show only

[0121] This option can be selected if the user wants to display the First, Largest or Smallest *x* number of values. If this option is set to 5, there will be five values displayed. If the dimension has Show Others enabled, the Others segment will take up one of the five display slots.

[0122] The First option will return the rows based on the options selected on the Sort tab of the property dialog. If the chart is a Straight Table, the rows will be returned based on the primary sort at the time. In other words, a user can change the values display by double-clicking on any column header and making that column the primary sort.

[0123] The Largest option returns the rows in descending order based on the first expression in the chart. When used in a Straight Table, the dimension values shown will remain consistent while interactively sorting the expressions. The dimensions values will (may) change when the order of the expressions is changed.

[0124] The Smallest option returns the rows in ascending order based on the

first expression in the chart. When used in a Straight Table, the dimension values shown will remain consistent while interactively sorting the expressions. The dimensions values will (may) change when the order of the expressions is changed.

Show only values that are

[0125] This option can be selected if the user wants to display all dimensions values that meet the specified condition for this option. Select to display values based on a percentage of the total, or on an exact amount. The relative to the total option enables a relative mode which is similar to the Relative option on the Expressions tab of the property dialog. The value may be entered as a calculated formula.

Show only values that accumulate to:

[0126] When this option is selected, all rows up to the current row are accumulated, and the result is compared to the value set in the option. The relative to the total option enables a relative mode which is similar to the Relative option on the Expressions tab of the property dialog, and compares the accumulated values (based on first, largest or smallest values) to the overall total. The value may be entered as a calculated formula.

[0127] Also provided are different display options comprising one or more of:

Show Others

[0128] Enabling this option will produce an Others segment in the chart. All dimension values that do not meet the comparison criteria for the display restrictions will be grouped into the Others segment. If there are dimensions after the selected dimension, Collapse Inner Dimensions will control whether individual values for the subsequent / inner dimensions display on the chart.

Global Grouping Mode

[0129] The option only applies to inner dimensions. When this option is enabled the restrictions will be calculated on the selected dimension only. All previous dimensions will be ignored. If this is disabled, the restrictions are calculated based on all preceding dimensions.

[0130] The use of Dimension Limits together with the selected option “Show

others" will now be described in relation to a simplified example, based on a data set shown in **FIG. 17a** containing variables Customer, Product and Sales, given for Customers A-F and Products X and Y:

EXAMPLE 1

[0131] Assume that the user wants to visualize the sales for each Customer. This corresponds to evaluating the mathematical function Sum(Sales) for the dimension variable Customer. This results in the following multidimensional cube (which may be visualized as a graph or a table, as shown in **FIG. 17b**):

EXAMPLE 2

[0132] Assume now that the user has applied the Dimension Limit "Show only the first 3 values" to the dimension Customer for generation of the cube, while also ticking the box "Show Others". This results in the cube shown in **FIG. 17c**. As shown, the sales are shown for Customers A and B, while the sales of the remaining Customers (C-F) are aggregated into an "Others" value.

EXAMPLE 3

[0133] Assume instead that the user has applied the Dimension Limit "Show only the largest 3 values" to the dimension Customer for generation of the cube, while also ticking the box "Show Others". This results in the cube shown in **FIG. 17d**. As shown, the sales are shown for Customers A and C, while the sales of the remaining customers (B and D-F) are aggregated into an "Others" value.

EXAMPLE 4

[0134] Assume instead that the user has applied the Dimension Limit "Show only the values that are larger or equal to 50" to the dimension Customer for generation of the cube, while also ticking the box "Show Others". This results in the cube shown in **FIG. 17e**. As shown, the sales are shown for Customers A, B and C, while the sales of the remaining customers (D-F) are aggregated into an "Others" value.

EXAMPLE 5

[0135] Assume instead that the user has applied the Dimension Limit "Show

only the largest values that accumulate to 80% of the overall total" to the dimension Customer for generation of the cube, while also ticking the box "Show Others". This results in the cube shown in **FIG. 17f**. As shown, the sales are shown for Customers A, B, C and F, while the sales of the remaining customers (D and E) are aggregated into an "Others" value.

[0136] All of the examples make use of the calculations described previously herein. It is to be understood that the above examples are simplified to facilitate the understanding of Dimension Limits. However, in a practical case, one or more complex mathematical functions may be evaluated for a large amount of data connected over a multitude of different tables.

[0137] The data may be processed in binary coded format, by using a conversion structure and based on a starting table, to sequentially evaluate a mathematical function for one or more dimensions (classification variables). This is exemplified with reference to Tables 15 and 16 in **FIG. 4**.

[0138] Here, Table 15 illustrates the use of a virtual data record which is sequentially updated for each record in the starting table, and Table 16 illustrates how an intermediate data structure is populated based on the sequentially updated content of the virtual data record. The intermediate data structure contains an aggregation field that is used for aggregating the evaluation result of a mathematical expression for each existing unique combination of values of the classification variables. In Table 16, the intermediate data structure aggregates the evaluated result for the following combinations of Client and Year: (0,0), (1,0), (2,0) (3,-2). The value -2 indicates a NULL value.

[0139] Table 16 also illustrates how dimensions are "eliminated" or "collapsed" in the intermediate data structure, which means that the mathematical expression is aggregated for all values of one or more classification variables. In this process, additional data records are added to the intermediate data structure to hold the aggregation of the evaluated result for the collapsed dimension(s). In Table 16, the intermediate data structure contains the following data records when Client is collapsed: (-1,0), (-1,-2), and the following data records when Year is collapsed: (0,-1) (1,-1), (2,-1), (3,-1), and one data record when both

Client and Year are collapsed: (-1,-1). The value -1 for a variable thus indicates that the evaluated results of all values of the variable have been aggregated.

[0140] The data in the intermediate data structure is then used for building a multidimensional cube, as shown in **FIG. 5**, Table 17. A slightly more advanced example of an intermediate data structure and a resulting multidimensional cube is illustrated in **FIG. 6**, Tables 28 and 29, respectively. Here, more complex mathematical functions are evaluated in the multidimensional cube (Table 29), and the intermediate data structure (Table 28) contains aggregation fields that aggregate the evaluation result of certain mathematical expressions that are required for correct evaluation of the mathematical functions in the multidimensional cube shown in Tables 28 and 29.

[0141] Returning to the above Examples 1-5, it should be realized that certain Dimension Limits can be applied by generating a full multidimensional cube (cfr. the Full table in Example 1 above) and simply selecting data in this cube, e.g. the 2 first Customers and their sales data (Example 2) or the 2 Customers with the largest sales and their sales data (Example 3).

[0142] A difficulty occurs when the Others value is to be evaluated, since this value cannot be defined when the multidimensional cube is generated (since its content is only known once the multidimensional cube has been generated). The Others value corresponds to an aggregation of the evaluated result for specific values of one or more classification variables (certain Customers in the above examples). In the above examples, the mathematical function is a simple summation and the evaluated result of the mathematical function for the Others value may be obtained by simply adding the sums (in the cube) for the Customers to be included in the Others value. However, if the mathematical function is more complex, e.g. if it contains an average quantity (see Tables 28-29 above), the Others value cannot be obtained by combining data in the cube.

[0143] One solution is to initiate calculation of a new multidimensional cube, which includes an aggregation field for the specific values of the classification variable(s) that define the Others value. In the context of Example 2, the new cube would be calculated to include a new Customer designated as “Others”

which includes the aggregated result for Customers C-F.

[0144] To minimize data processing, the methods and systems can make use of the intermediate data structure (e.g., the existing or previously populated intermediate data structure) to populate the multidimensional cube with the Others value. As explained in the previously, the aggregation fields of the intermediate data structure are defined to enable the dimensions to be collapsed (eliminated). In some respects, the evaluation of an Others value may be regarded as a partial elimination of a dimension in the intermediate data structure.

[0145] Thus, in Examples 2-4, the Dimension Limits identify the values of the Customer variable to be included in the cube, together with the corresponding sales. The Others value of the cube is populated by aggregating the sales for the remaining values of the variable Customer by traversing the intermediate data structure.

[0146] In Example 5, the Dimension Limit requires the total sales to be known. The total sales data is only known once the intermediate data structure has been generated (corresponding to an elimination of the dimension Customer). To populate the Others value, the intermediate data structure is traversed once more to identify the largest values (sales) in the aggregation fields for the different Customers until at least 80% of the total sales is reached, and to evaluate the content of the Others value by aggregating the sales of the remaining Customers.

[0147] There are certain situations when it may not be possible to correctly evaluate the Others value based on the intermediate data structure, e.g. if the evaluation requires special attention to frequency data (mentioned in US7058621). In one embodiment, the methods and systems comprise a component that detects a potential need for special attention to frequency data. If such a potential need is detected, the methods and systems can refuse to populate the Others value. In a variant, the methods and systems can instead initiate calculation of a new multidimensional cube that includes the Others value (e.g. using the processing intensive alternative which is generally avoided by

evaluating the Others value based on the intermediate data structure) . In one example, a potential need for special attention to frequency data may be flagged whenever the software detects, during the generation of the multidimensional cube, that more than one data record in the intermediate data structure is updated based on the content of one virtual data record.

Example of Global Grouping Mode

[0148] Assume the multidimensional cube shown in **FIG. 18a**. Here, the cube is generated to evaluate the sales for two dimensions (classification variables): Product and Customer. Assume now that the user has applied the Dimension Limit "Show only the largest 3 values" to the variable Customer, while also ticking the box "Show Others". This would result in the multidimensional cube shown in **FIG. 18b**.

[0149] As shown, the process identifies the two Customers that have the largest sales of Product X and the two Customers that have the largest sales of Product Y, and generates an Others value for Product X and an Others value for Product Y. The Others value for Product X accumulates the sales for Customers C-F, and the Others value for Product Y accumulates the sales for Customers B and D-F. The Others values are generated in the same way as described above (e.g. by traversing the intermediate data structure).

[0150] Assume instead that the user has applied the same Dimension Limit for the variable Customer, and ticked the box "Global Grouping Mode" (while also ticking the box "Show Others"). This would result in the multidimensional cube shown in **FIG. 18c**.

[0151] The Global Grouping Mode causes the process to identify the two Customers that have the largest sales of all products (e.g. Product X and Product Y combined). The cube is generated to include the sales data for Product X for these two Customers, and an Others value that accumulates the sales for the remaining Customers for Product X (e.g. Customers B and D-F), as well as the sales data for Product Y for these two Customers, and an Others value that accumulates the sales for the remaining Customers for Product Y (e.g. Customers B and D-F).

[0152] Thus, the Global Grouping Mode causes the Dimension Limits to be applied only to the selected dimension (Customer).

[0153] In an aspect, illustrated in **FIG. 19**, provided are methods for data analysis comprising performing a data processing event on a dataset resulting in a first multidimensional cube data structure at 1901 and applying one or more dimension limits to the multidimensional cube data structure resulting in a second multidimensional cube data structure at 1902. The first data processing event can comprise evaluating a mathematical function for one or more dimension variables in the data set. The one or more dimension limits can comprise show only, show only values that are, show only values that accumulate, and the like. In an aspect, the second multidimensional cube data structure can by displayed according to one or more of show others, global grouping, and the like.

[0154] A user can be presented with a Dimension Limits option to control the number of dimension values displayed in a given chart. The user can select one of a plurality of values, for example: First, Largest, and Smallest. These values control the way the system sorts the values it returns to the visualization component. In an aspect, sorting only occurs for the first expression (except in pivot tables when a primary sort may override the first dimension sort).

[0155] While the methods and systems have been described in connection with preferred embodiments and specific examples, it is not intended that the scope be limited to the particular embodiments set forth, as the embodiments herein are intended in all respects to be illustrative rather than restrictive.

[0156] Unless otherwise expressly stated, it is in no way intended that any method set forth herein be construed as requiring that its steps be performed in a specific order. Accordingly, where a method claim does not actually recite an order to be followed by its steps or it is not otherwise specifically stated in the claims or descriptions that the steps are to be limited to a specific order, it is no way intended that an order be inferred, in any respect. This holds for any possible non-express basis for interpretation, including: matters of logic with respect to arrangement of steps or operational flow; plain meaning derived from

grammatical organization or punctuation; the number or type of embodiments described in the specification.

[0157] It will be apparent to those skilled in the art that various modifications and variations can be made without departing from the scope or spirit. Other embodiments will be apparent to those skilled in the art from consideration of the specification and practice disclosed herein. It is intended that the specification and examples be considered as exemplary only, with a true scope and spirit being indicated by the following claims.

CLAIMS

1. A method for collaborative computing comprising:
initiating a primary session for a first user;
requesting collaboration from a second user;
initiating a secondary session for the second user; and
providing a single state space for collaborative real-time data analysis to the first user and the second user wherein an interaction by either user is reflected in the single state space.

2. The method of claim 1, wherein initiating a primary session for a first user comprises generating a first user interface.

3. The method of claim 1 or claim 2, wherein requesting collaboration from a second user comprises providing the second user a URL configured to allow the second user to view the secondary session.

4. The method of claim 1, wherein initiating a primary session for the second user comprises generating a second user interface.

5. The method of any one of claims 1-4, wherein the single state space is a view of a dataset.

6. The method of any one of claims 1-5, wherein the primary and secondary sessions result in a collaboration session that is connected to a plurality of XML transformers connected via synchronization logic.

7. The method of claim 6, wherein an interaction performed by any of the plurality of XML transformers is propagated back to the other XML transformers in the plurality of XML transformers.

8. A method for time shifted collaborative analysis, comprising:

creating a state space that reflects a selection state;

creating a note; and

attaching the note to an object in the state space.

9. The method of claim 8, wherein the state space is a view of a dataset.

10. The method of any one of claim 8 and claim 9, wherein the note comprises an observation about the object.

11. The method of any one of claim 8 through 10, further comprising:

saving the selection state; and

associating the saved selection state with the note.

12. The method of claim 11, wherein the saved selection state reflects the selection state at the time the note is created.

13. The method any one of claim 11 and claim 12, wherein the note comprises a note thread wherein the note thread comprises a plurality of notes.

14. The method of claim 13, further comprising saving a plurality of selection states wherein each of the plurality of selection states is associated with one or more of the plurality of notes.

15. The method of any one of claim 8 through claim 14, further comprising identifying the object as having an attached note.

16. The method of any one of claim 8 through claim 15, further comprising:

presenting an object in a state space having an attached note;

receiving a selection of the note; and

presenting the note and adjusting the state space to reflect a

saved selection state associated with the note.

17. The method of claim 16, wherein the note comprises an observation about the object.

18. The method of any one of claim 16 and claim 17, wherein the state space is a first view of a dataset.

19. The method of claim 18, wherein the saved selection state is a second view of the dataset.

20. The method of any one of claim 16 through claim 19, wherein the saved selection state was saved contemporaneously with creation of the note.

21. The method of any one of claim 16 through claim 20, wherein the note comprises a note thread wherein the note thread comprises a plurality of notes wherein one or more of the plurality of notes has an associated saved selection state.

22. A device for collaborative computing, comprising a memory having computer-executable instructions; and a processor functionally coupled to the memory and configured, by the computer-executable instructions,

to initiate a primary session for a first user;

to request collaboration from a second user;

to initiate a secondary session for the second user; and

to provide a single state space for collaborative real-time data analysis to the first user and the second user wherein an interaction by either user is reflected in the single state space.

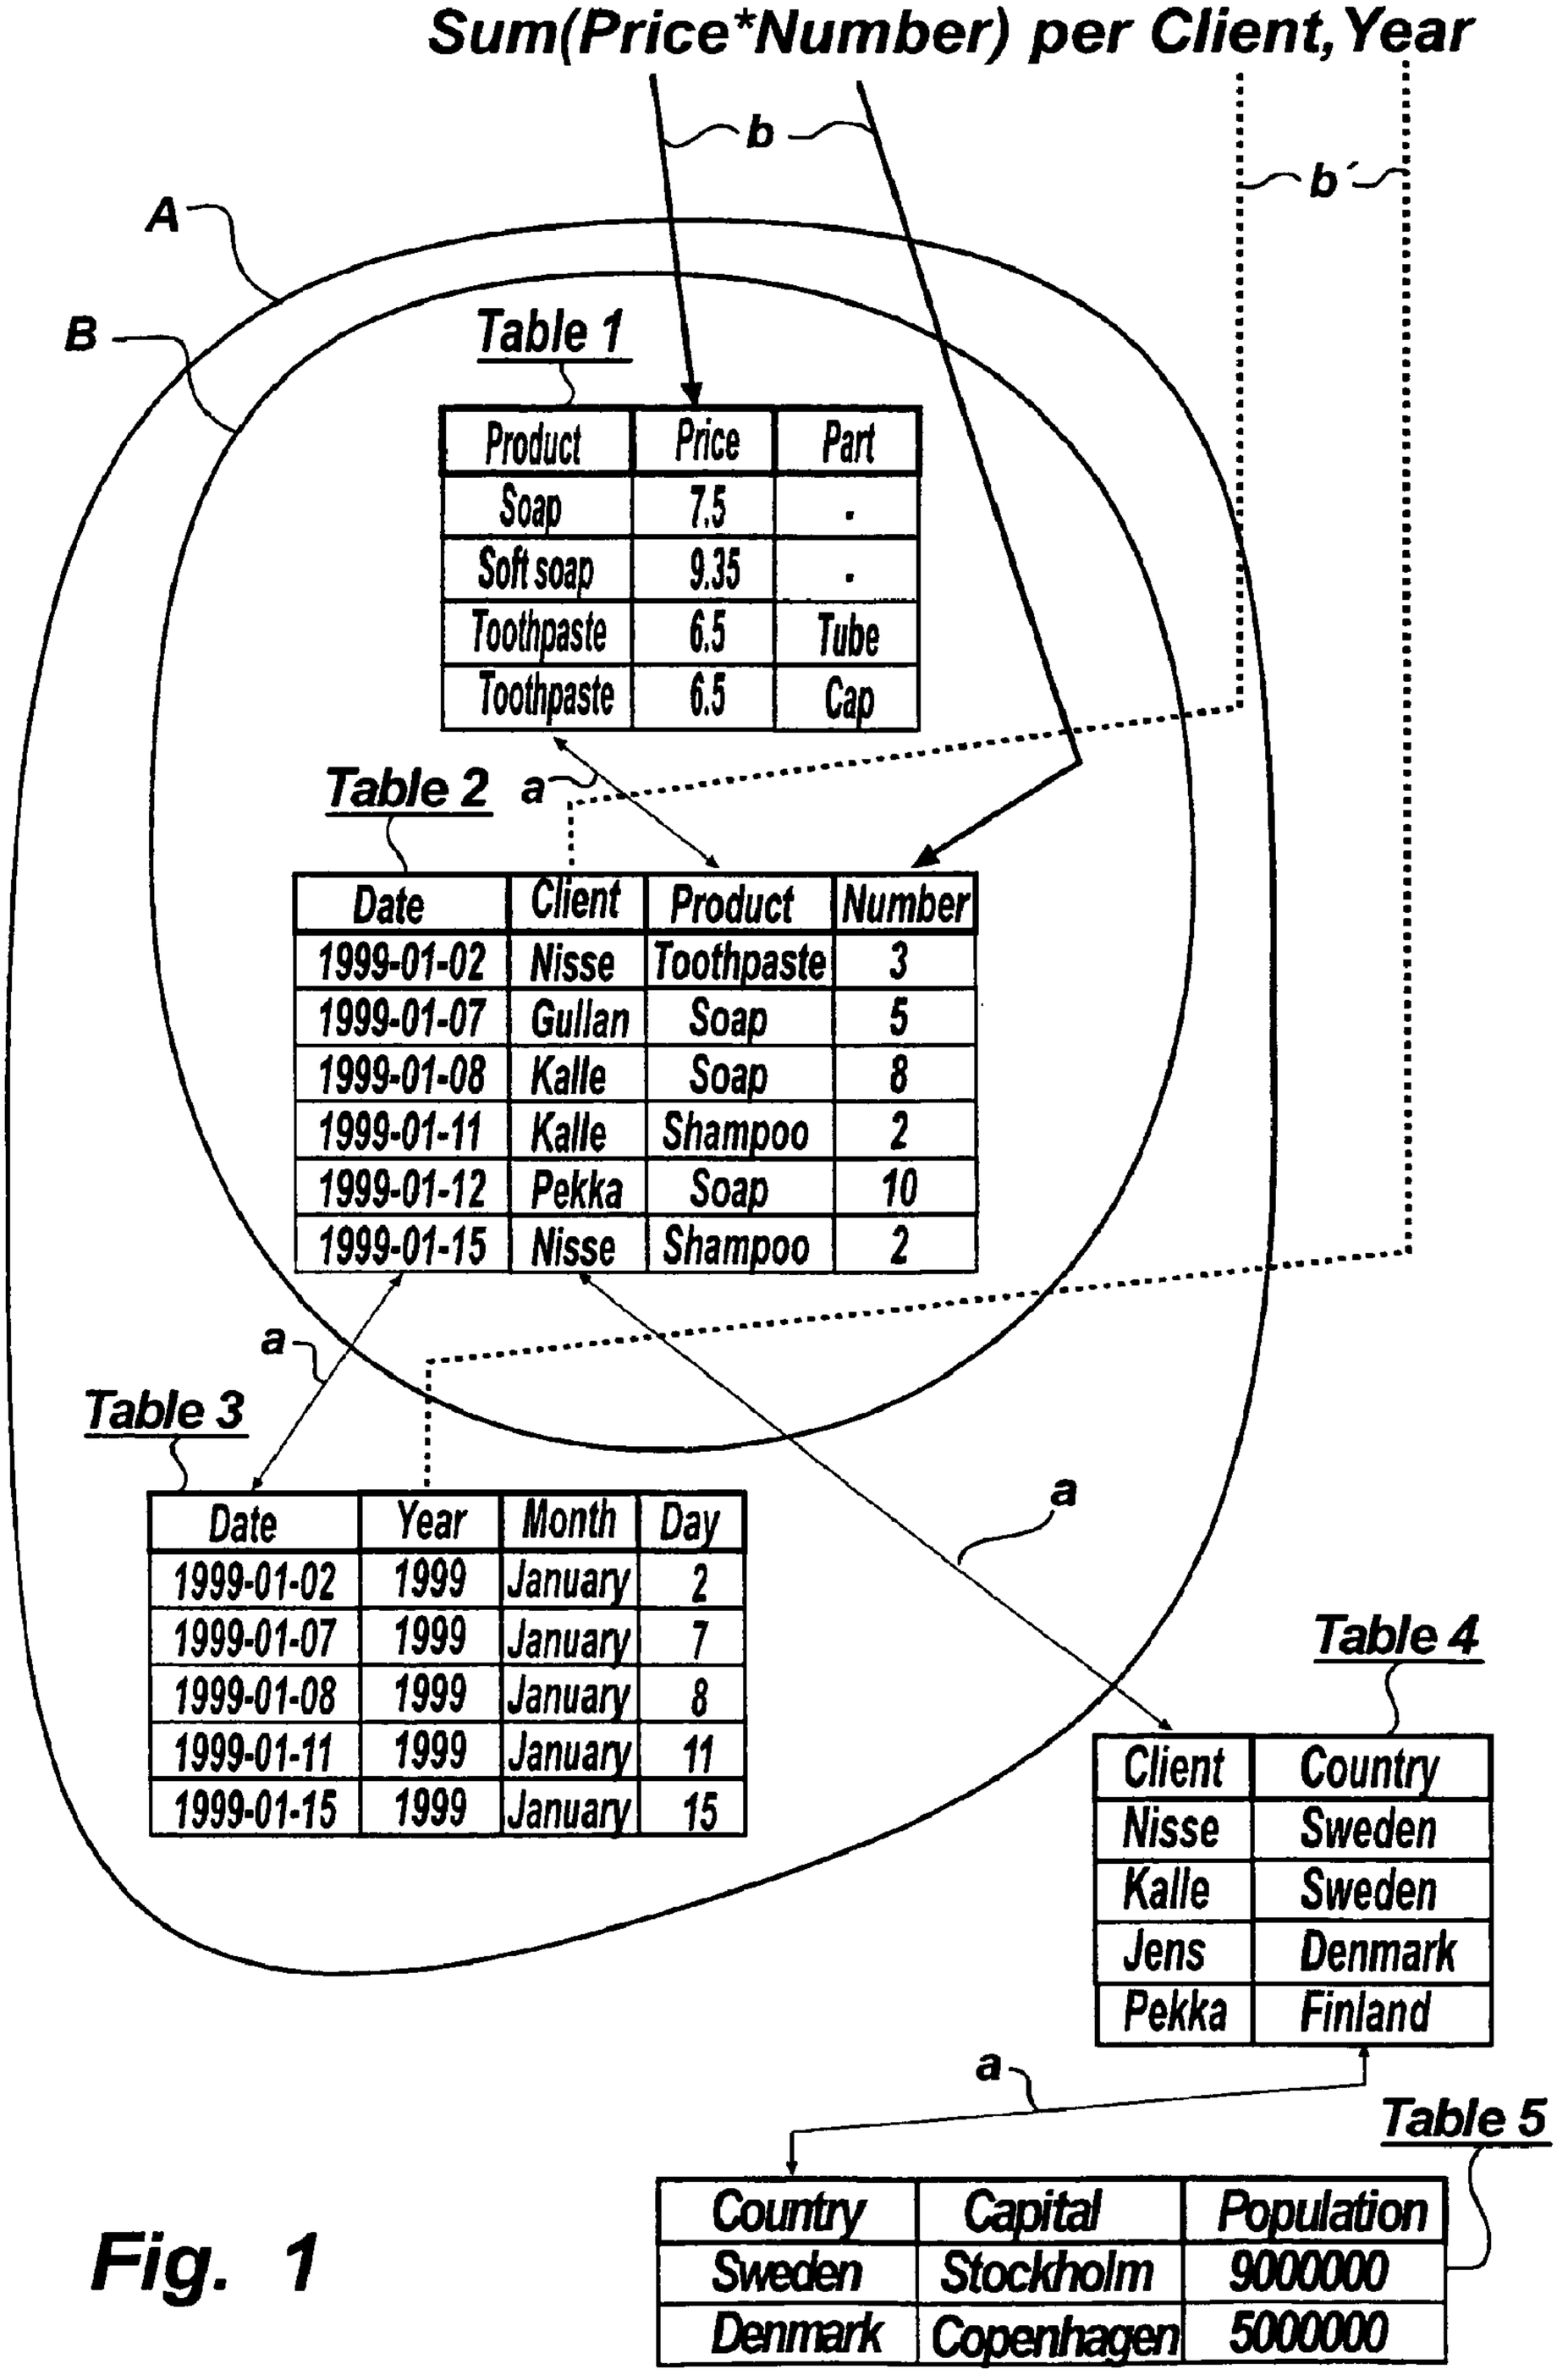
Sum(Price*Number) per Client,Year
b
b´
A
B
Table 1
Product | Price | Part
Soap | 7.5 | .
Soft soap | 9.35 | .
Toothpaste | 6.5 | Tube
Toothpaste | 6.5 | Cap
Table 2
a
Date | Client | Product | Number
1999-01-02 | Nisse | Toothpaste | 3
1999-01-07 | Gullan | Soap | 5
1999-01-08 | Kalle | Soap | 8
1999-01-11 | Kalle | Shampoo | 2
1999-01-12 | Pekka | Soap | 10
1999-01-15 | Nisse | Shampoo | 2
a
Table 3
Date | Year | Month | Day
1999-01-02 | 1999 | January | 2
1999-01-07 | 1999 | January | 7
1999-01-08 | 1999 | January | 8
1999-01-11 | 1999 | January | 11
1999-01-15 | 1999 | January | 15
a
Table 4
Client | Country
Nisse | Sweden
Kalle | Sweden
Jens | Denmark
Pekka | Finland
a
Table 5
Country | Capital | Population
Sweden | Stockholm | 9000000
Denmark | Copenhagen | 5000000

Fig. 1

Fig. 2

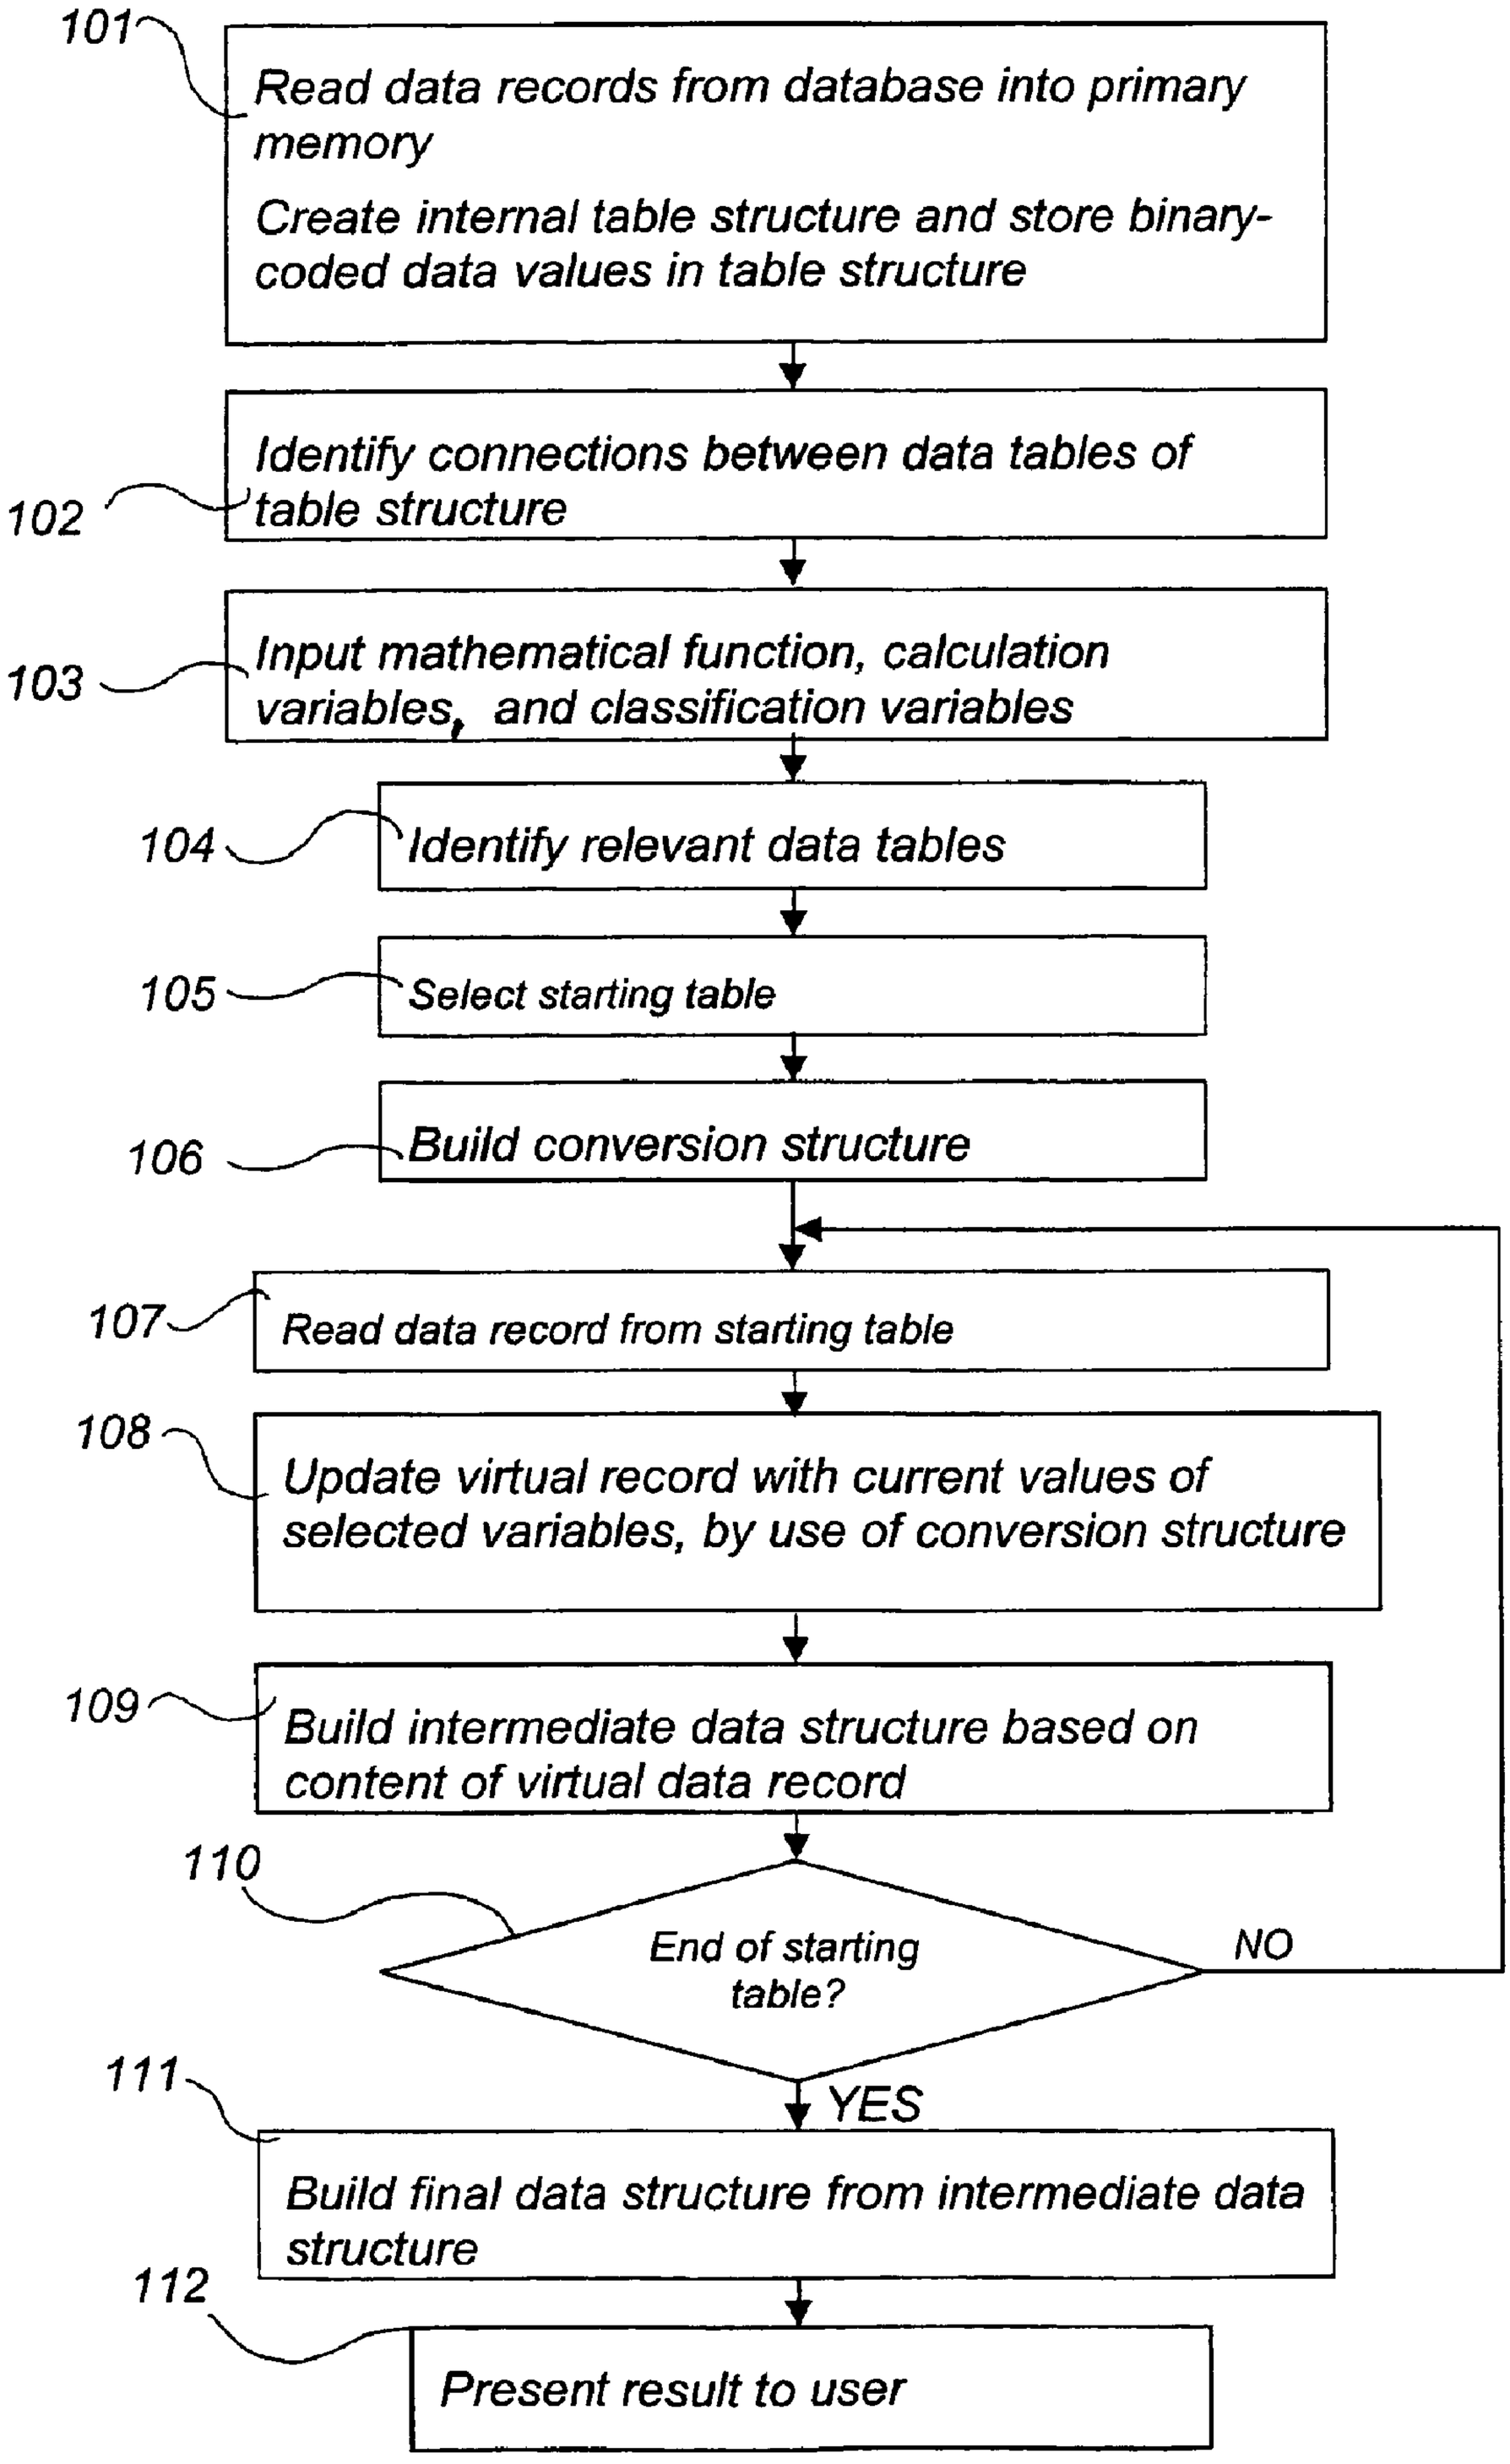
101
Read data records from database into primary memory
Create internal table structure and store binary-coded data values in table structure
102
Identify connections between data tables of table structure
103
Input mathematical function, calculation variables, and classification variables
104
Identify relevant data tables
105
Select starting table
106
Build conversion structure
107
Read data record from starting table
108
Update virtual record with current values of selected variables, by use of conversion structure
109
Build intermediate data structure based on content of virtual data record
110
End of starting table?
NO
YES
111
Build final data structure from intermediate data structure
112
Present result to user

Table 6

Product	
Soap	0
Soft soap	1
Toothpaste	2
Shampoo	3

Table 7

Price	
7.5	0
9.35	1
6.5	2

Table 8

Date	
1999-01-02	0
1999-01-07	1
1999-01-08	2
1999-01-11	3
1999-01-12	4
1999-01-15	5

Table 9

Client	
Nisse	0
Gullan	1
Kalle	2
Pekka	3
Jens	4

Table 10

Number	
3	0
5	1
8	2
2	3
10	4

Table 11

Year	
1999	0

Table 12

Country	
Sweden	0
Denmark	1
Finland	2

Fig. 3

Table 13

0	0
1	0
2	0
3	0
5	0

Table 14

Product	Price
0	0
1	1
2	2 2

Table 15

	Client	Year	Price	Number
	0			0
	0	0		0
R1	0	0	2	0
R2	0	0	2	0
R3	1	0	0	1
R4	2	0	0	2
R5	2	0	-2	3
R6	3	-2	0	4
R7	0	0	0	3

Table 16

Aggregation field

Client	Year	Σ Number*Price	
0	0	0	Sum = Sum (0) + 3 x 6.5 -> 20.5
0	0	0	Sum = Sum (20.5) + 3 x 6.5 -> 41
1	0	0	Sum = Sum (0) + 5 x 7.5 -> 37.5
2	0	0	Sum = Sum (0) + 8 x 7.5 -> 60
2	0	0	Sum = Sum (60) +2 x (NULL) -> 60
3	-2	0	Sum = Sum (0) + 10 x 7.5 -> 75
0	0	0	Sum = Sum (41) + 2 x (NULL) -> 41
		-1	0 Sum = Sum (0) +20.5 +20.5 + 37.5 + 60 + 0 + 0->138.5
		-1	-2 Sum = Sum (0) +75 -> 75
0		-1	Sum = Sum (0) + 20.5 + 20.5 + 0 -> 41
1		-1	Sum = Sum (0) + 37.5 -> 37.5
2		-1	Sum = Sum (0) + 60 + 0 -> 60
3		-1	Sum = Sum (0) + 75 -> 75
-1		-1	Sum = Sum (0) + 20.5 + 20.5 + 37.5 + 60 + 0 + 75 + 0 -> 213.5

Fig. 4

Table 17

Client	Year	Sum (Number x Price)
Nisse	1999	41
Gullan	1999	37.5
Kalle	1999	60
Pekka	<NULL>	75
<ALL>	1999	138.5
<ALL>	<NULL>	75
Nisse	<ALL>	41
Gullan	<ALL>	37.5
Kalle	<ALL>	60
Pekka	<ALL>	75
<ALL>	<ALL>	213.5

Table 18

Sum (Price*Number) Per Client, Year

	1999	<NULL>	<ALL>
Nisse	41		41
Gullan	37.5		37.5
Kalle	60		60
Pekka		75	75
<ALL>	138.5	75	75

Table 20

Date	Product	Number	Client
1998-12 -20	B	5	Nisse
1999-02-05	A	7	Kalle
1999-02-06	B	9	Kalle

Table 21

Product	Price group	Product group
A	4	Z
B	3	X

Table 22

Price group	Price
3	5.5
4	3.5

Table 23

Product group	Environment index	Legal status
X	I	YES
Y	IX	NO
Z	II	YES

Fig. 5

Table 24

Product group->Environment index

X ⟶ I
Y ⟶ IX
Z ⟶ II

Table 25

Price group ->Price

3 ⟶ 5.5
4 ⟶ 3.5

Table 26

Product->Price, Environment index

A ⟶ 3.5,II
B ⟶ 5.5,I

Table 27

Client	Environment index	Number	Price
Nisse	I	5	5.5
Kalle	II	7	3.5
Kalle	I	9	5.5

Table 28

Client	Environment index	Σ-Number x Price	Σ-Environment index
Nisse	I	Σx: 27.5, N: 1	First: I, Last: I
Kalle	II	Σx: 24.5, N: 1	First: II, Last: II
Kalle	I	Σx: 49.5, N: 1	First: I, Last: I
<ALL>	I	Σx: 77, N: 2	First: I, Last: I
<ALL>	II	Σx: 24.5, N: 1	First: II, Last: II
<ALL>	<ALL>	Σx: 101.5, N: 3	First: I, Last: II

Table 29

Client	Environment index	IF (Only (Environment index)='I', Sum(Number*Price)*2,Sum(Number*Price))	Avg(Number*Price)
Nisse	I	55.0	27.5
Kalle	II	24.5	24.5
Kalle	I	99.0	49.5
<ALL>	I	154.0	38.5
<ALL>	II	24.5	24.5
<ALL>	<ALL>	<NULL>	33.8

Fig. 6

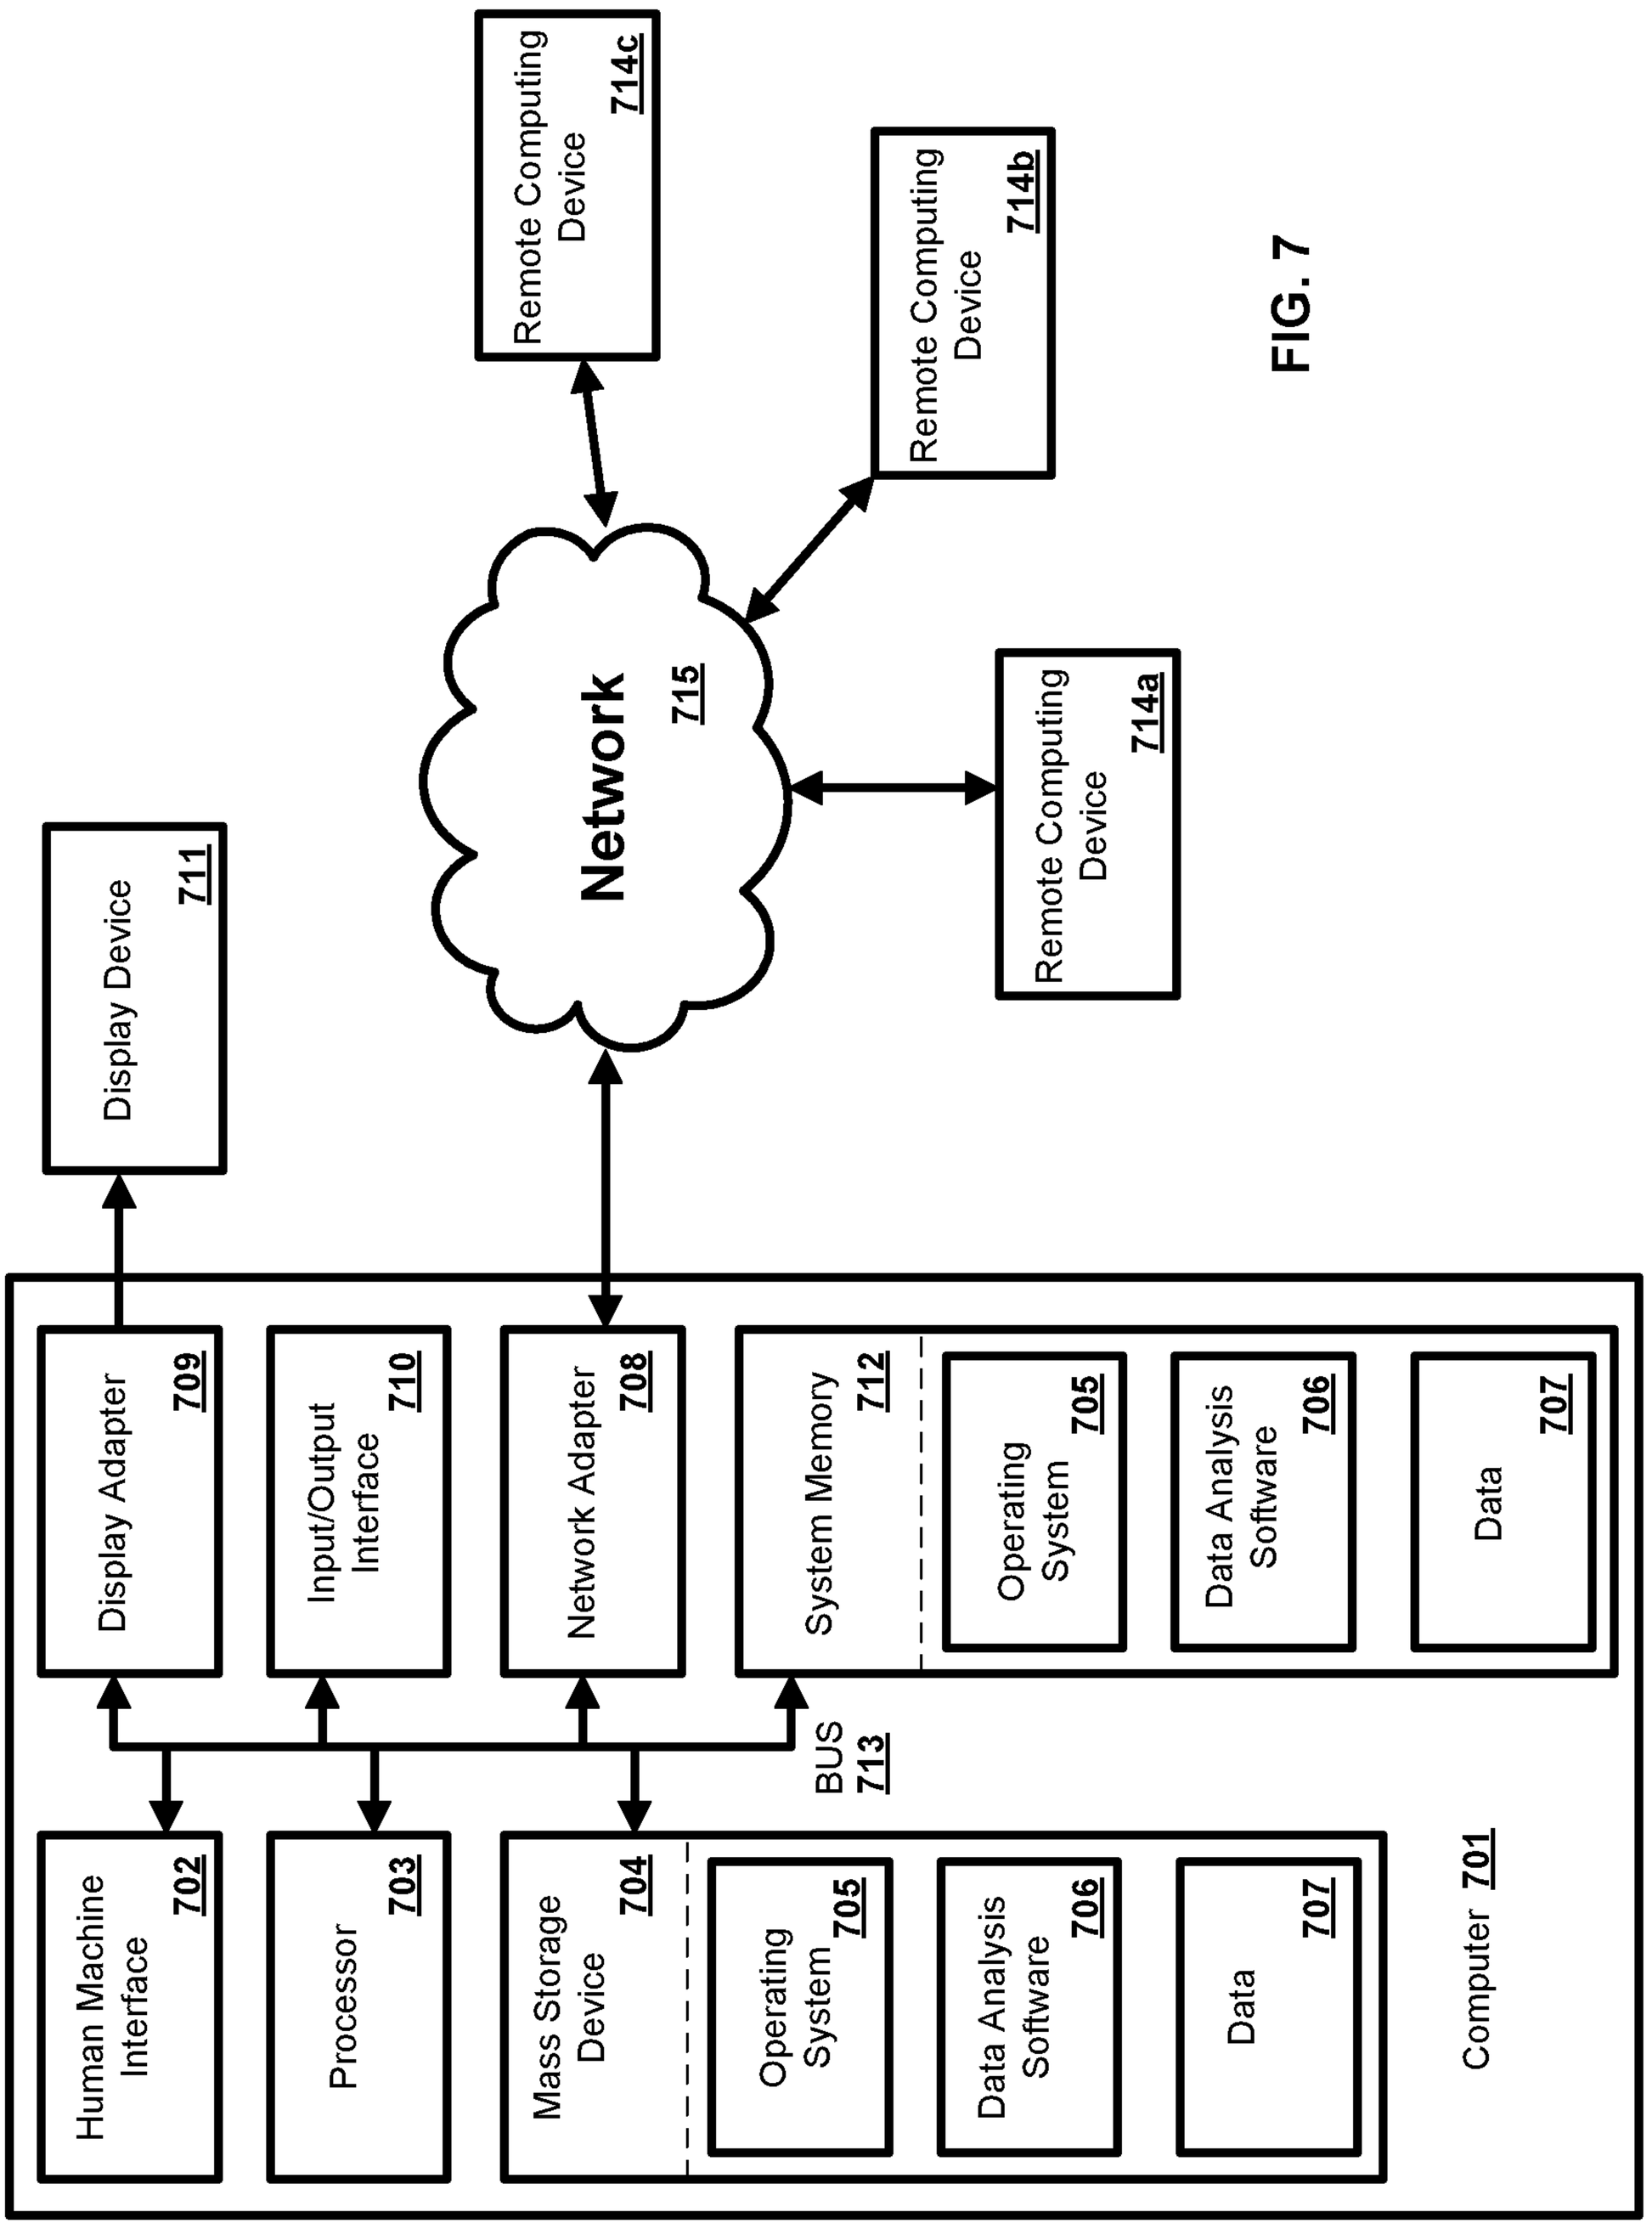
Remote Computing Device 714c
Remote Computing Device 714b
Network 715
Remote Computing Device 714a
Display Device 711
Display Adapter 709
Input/Output Interface 710
Network Adapter 708
System Memory 712
Operating System 705
Data Analysis Software 706
Data 707
BUS 713
Human Machine Interface 702
Processor 703
Mass Storage Device 704
Operating System 705
Data Analysis Software 706
Data 707
Computer 701

FIG. 7

FIG. 8

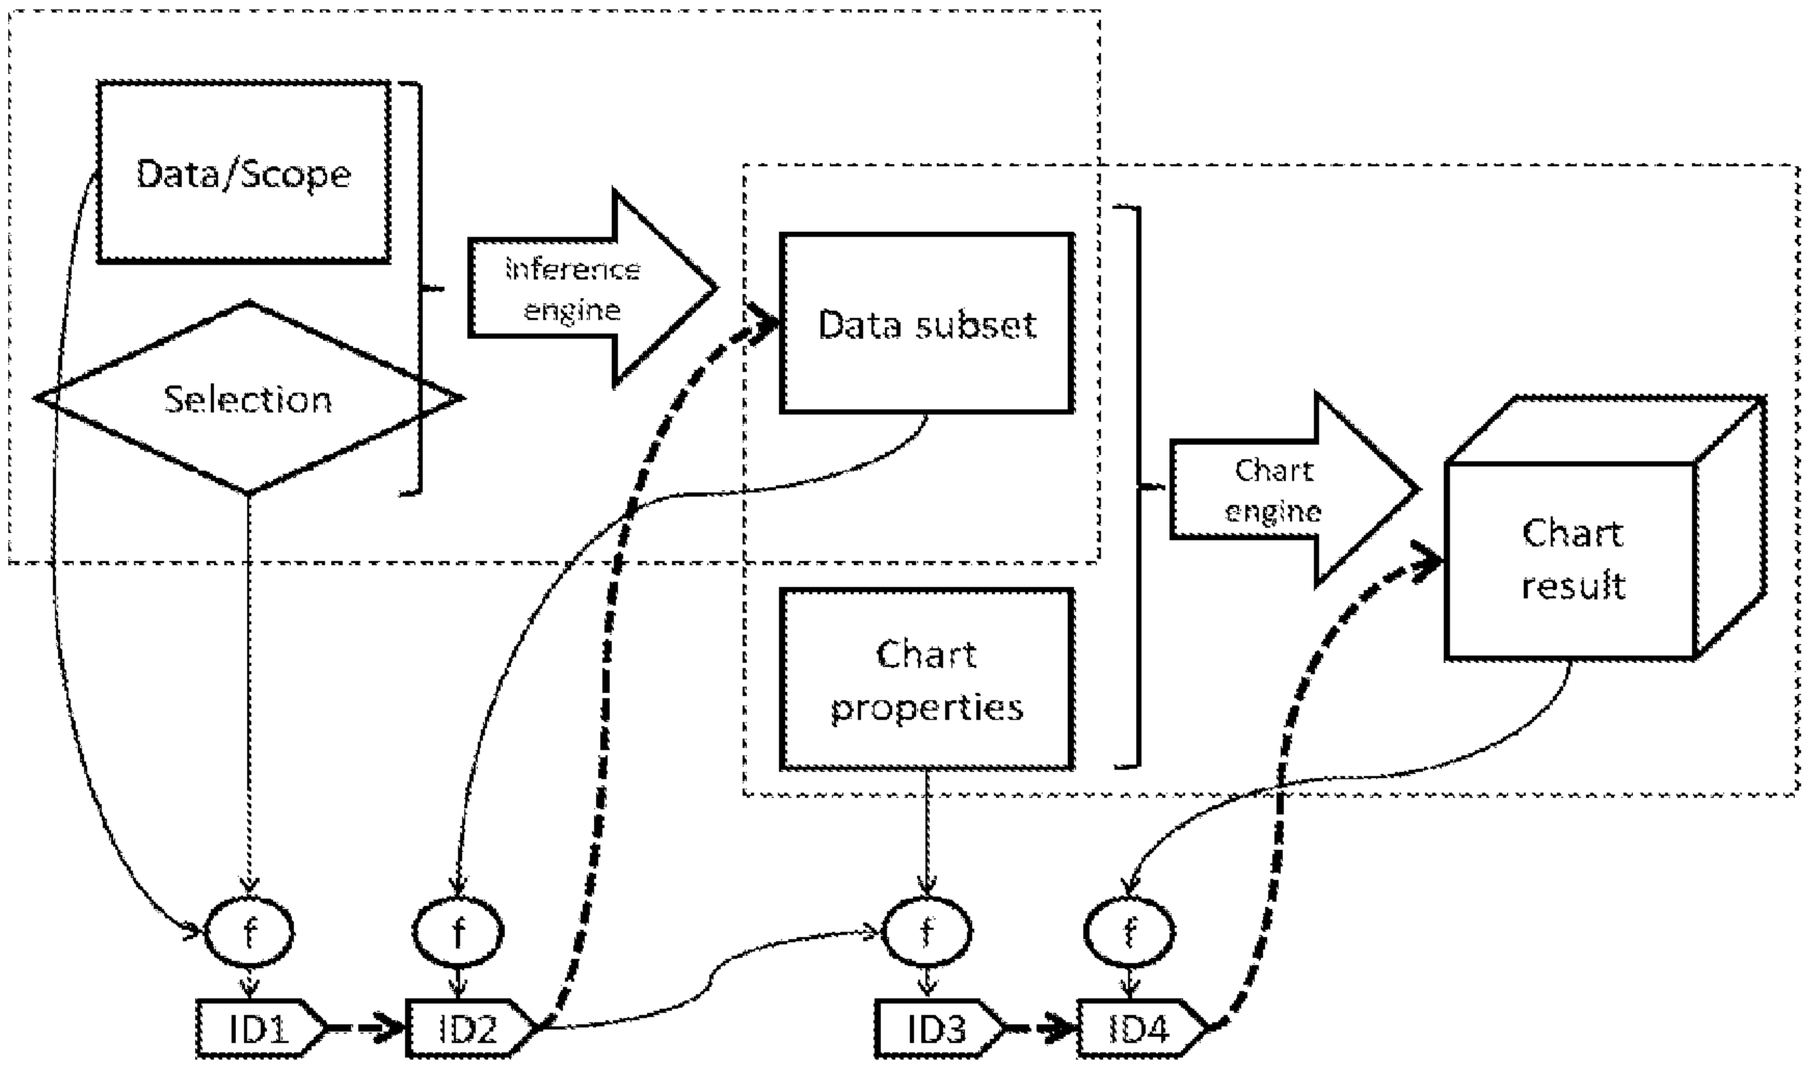
Data/Scope
Selection
inference engine
Data subset
Chart properties
Chart engine
Chart result
f
f
f
f
ID1
ID2
ID3
ID4

FIG. 9

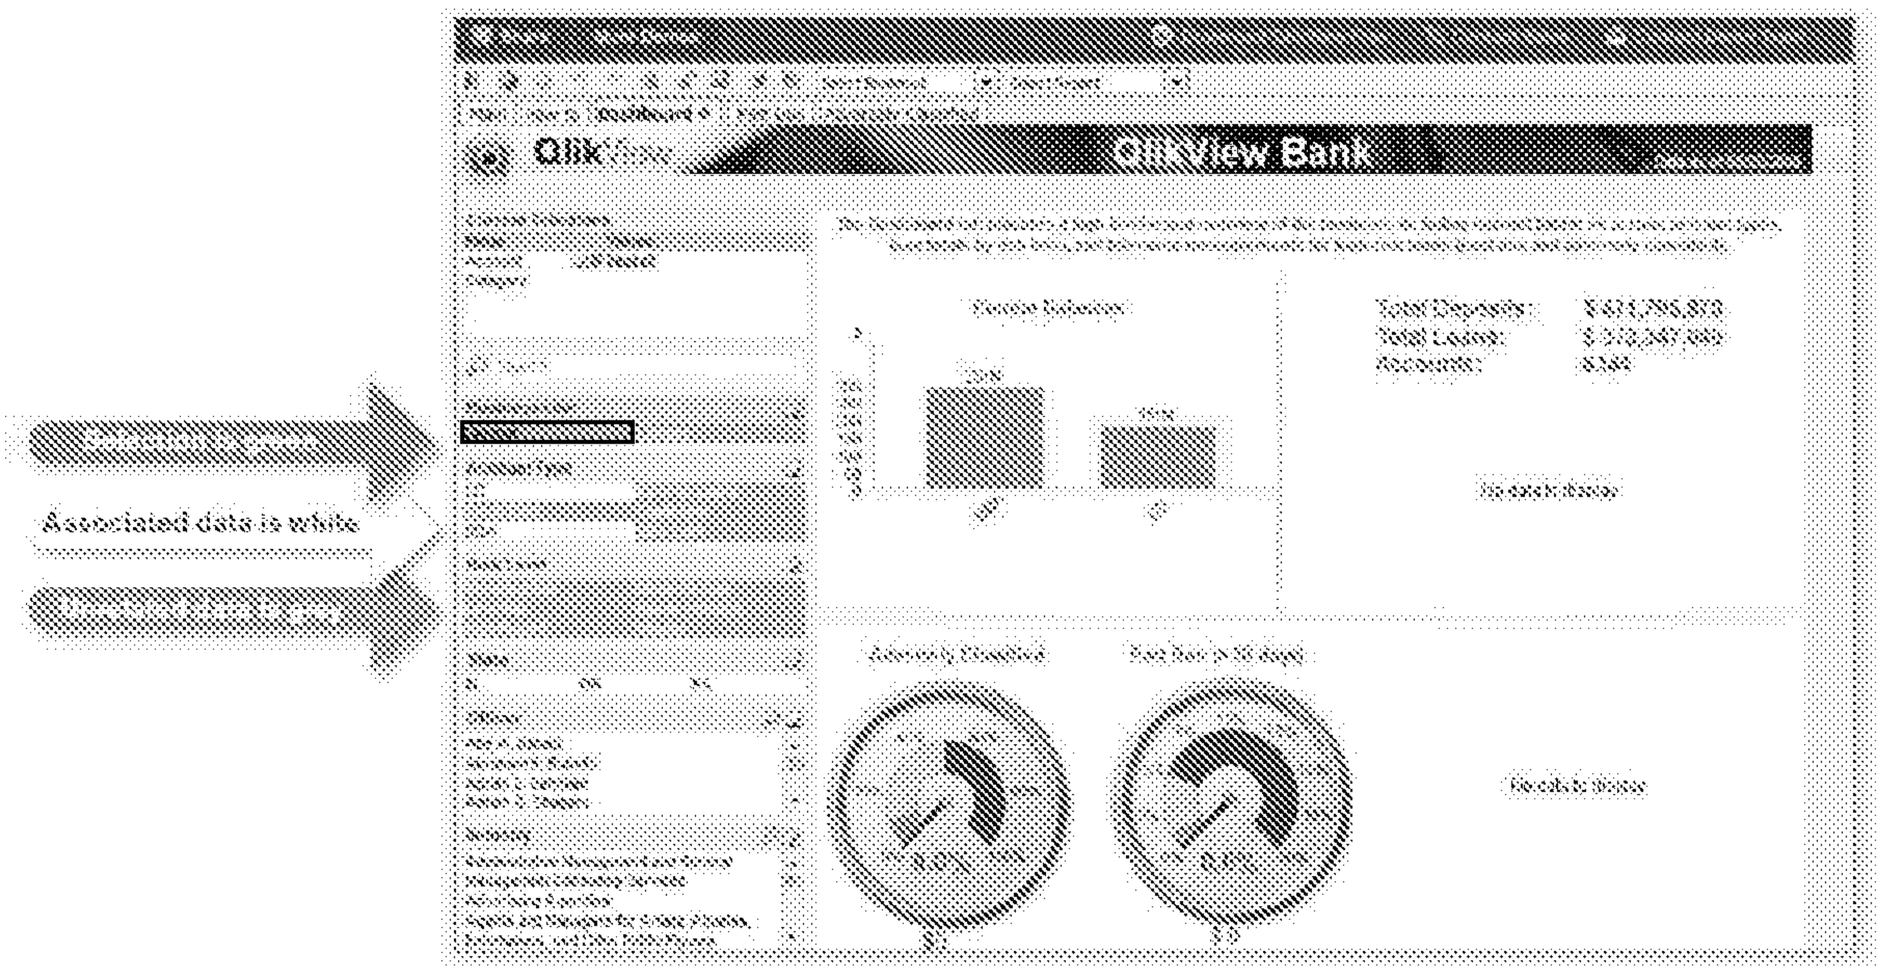

FIG. 10a

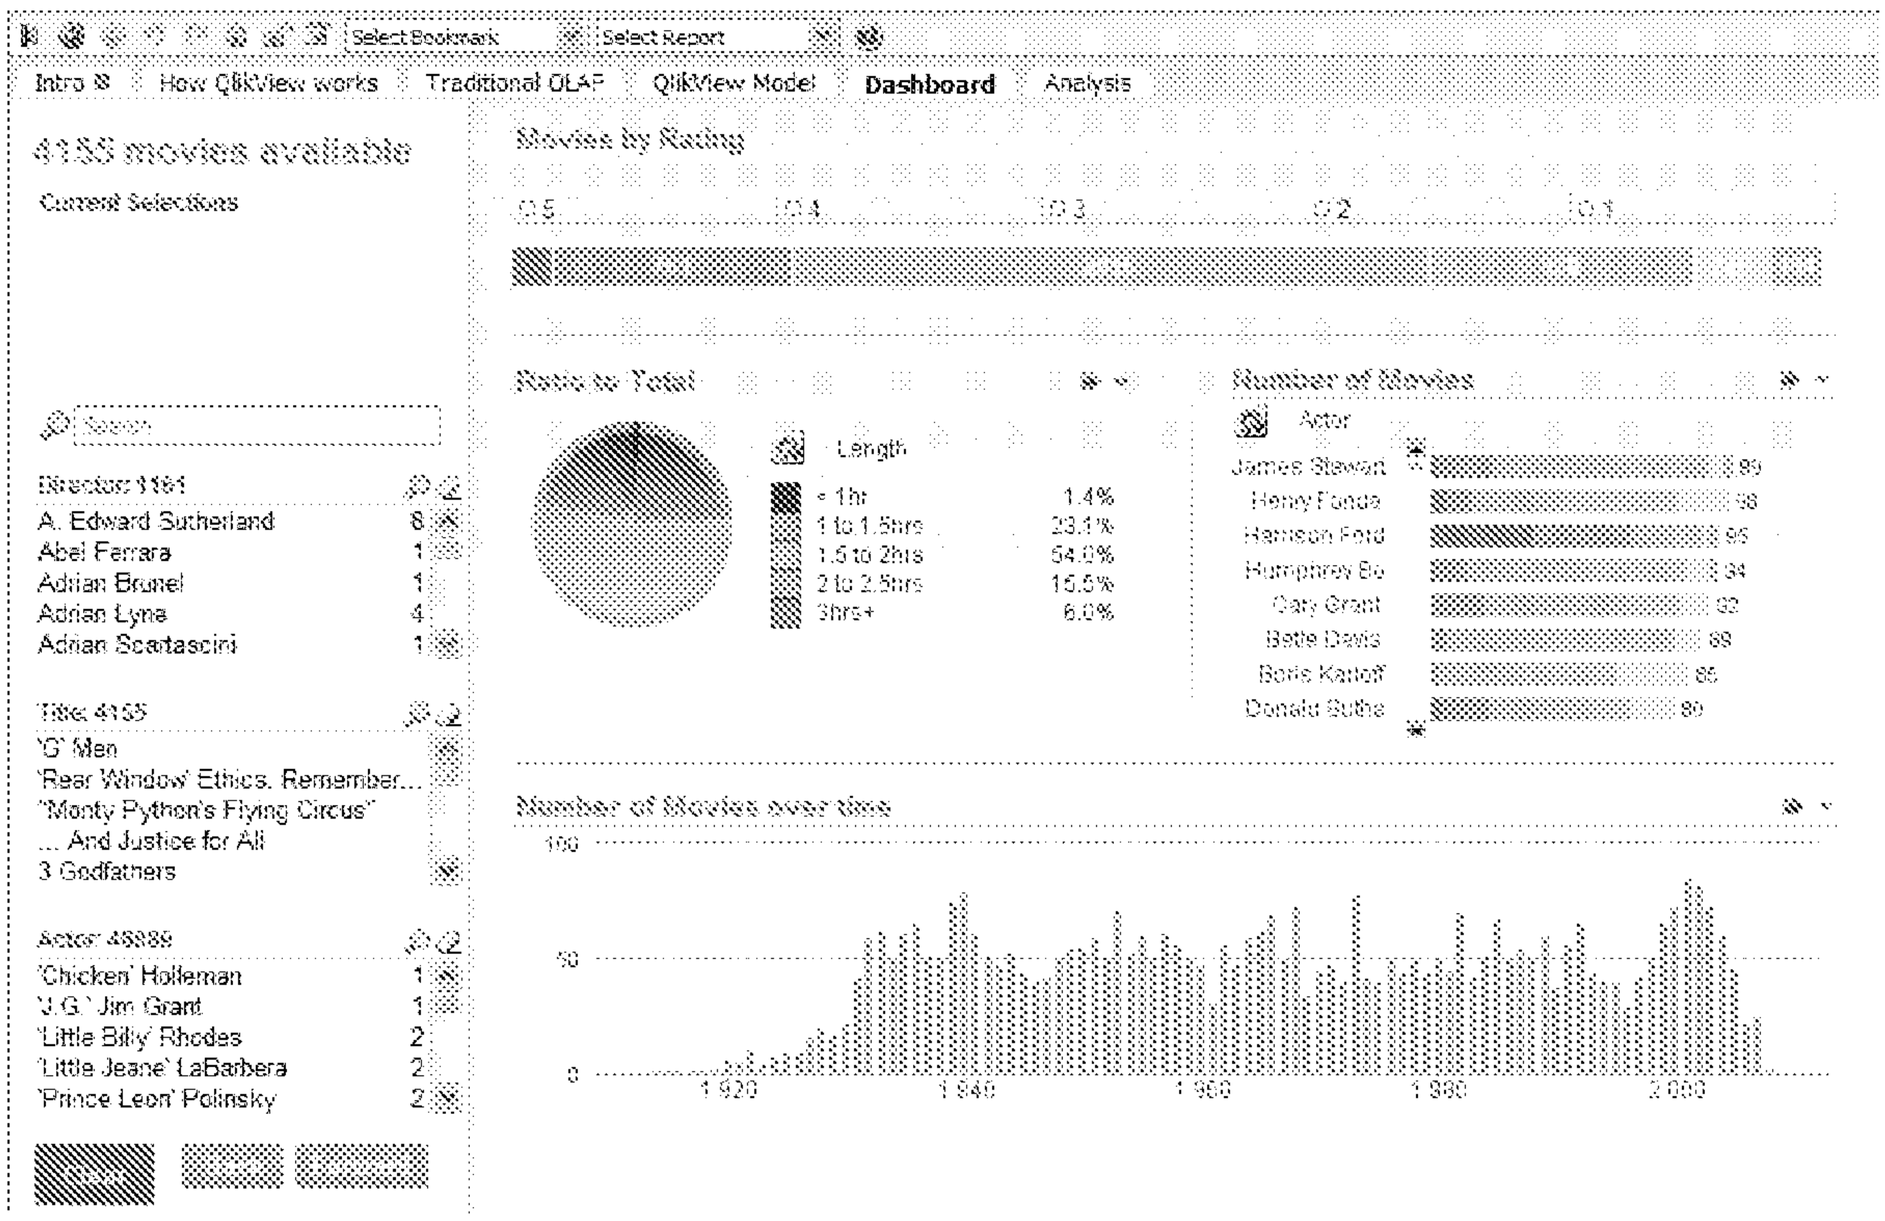

FIG. 10b

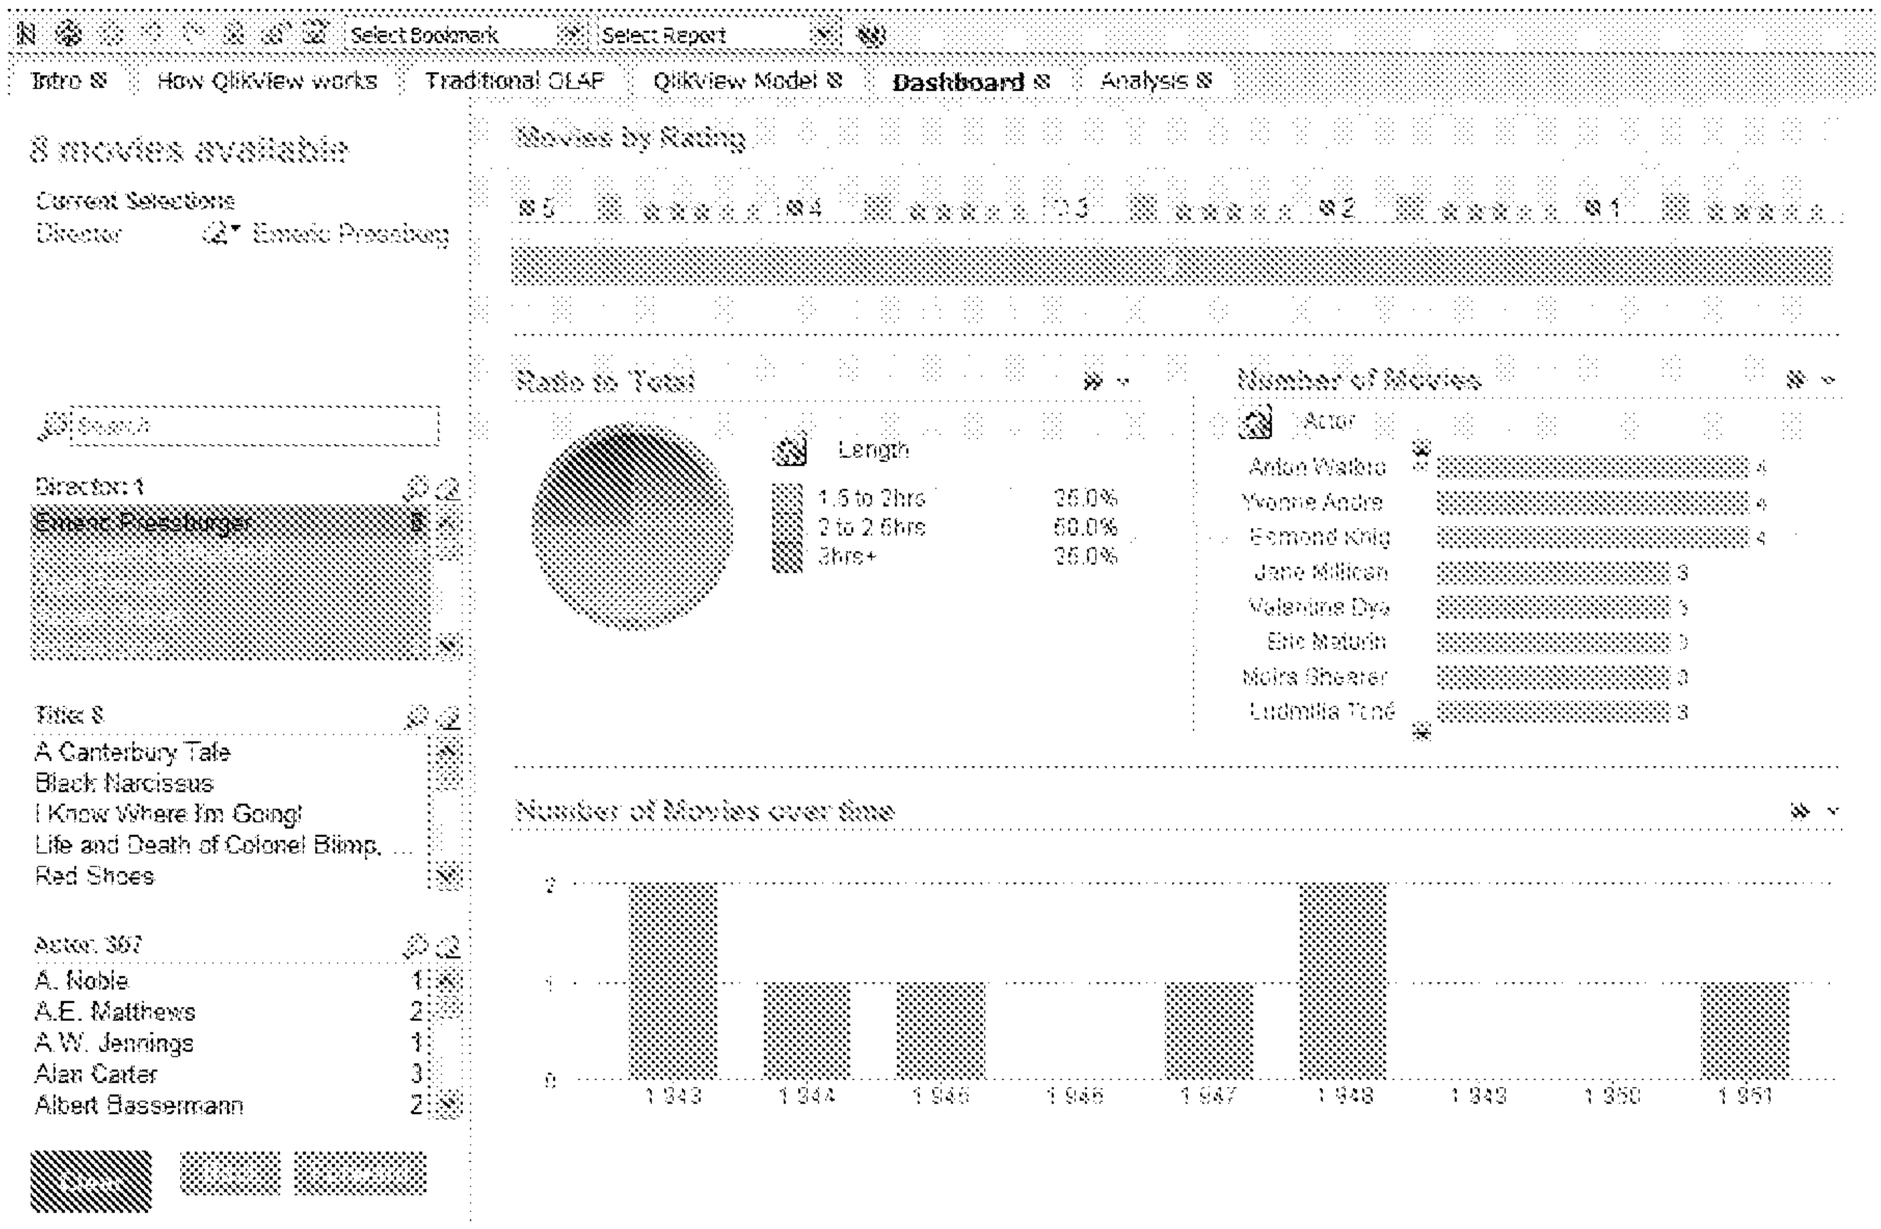

FIG. 11a

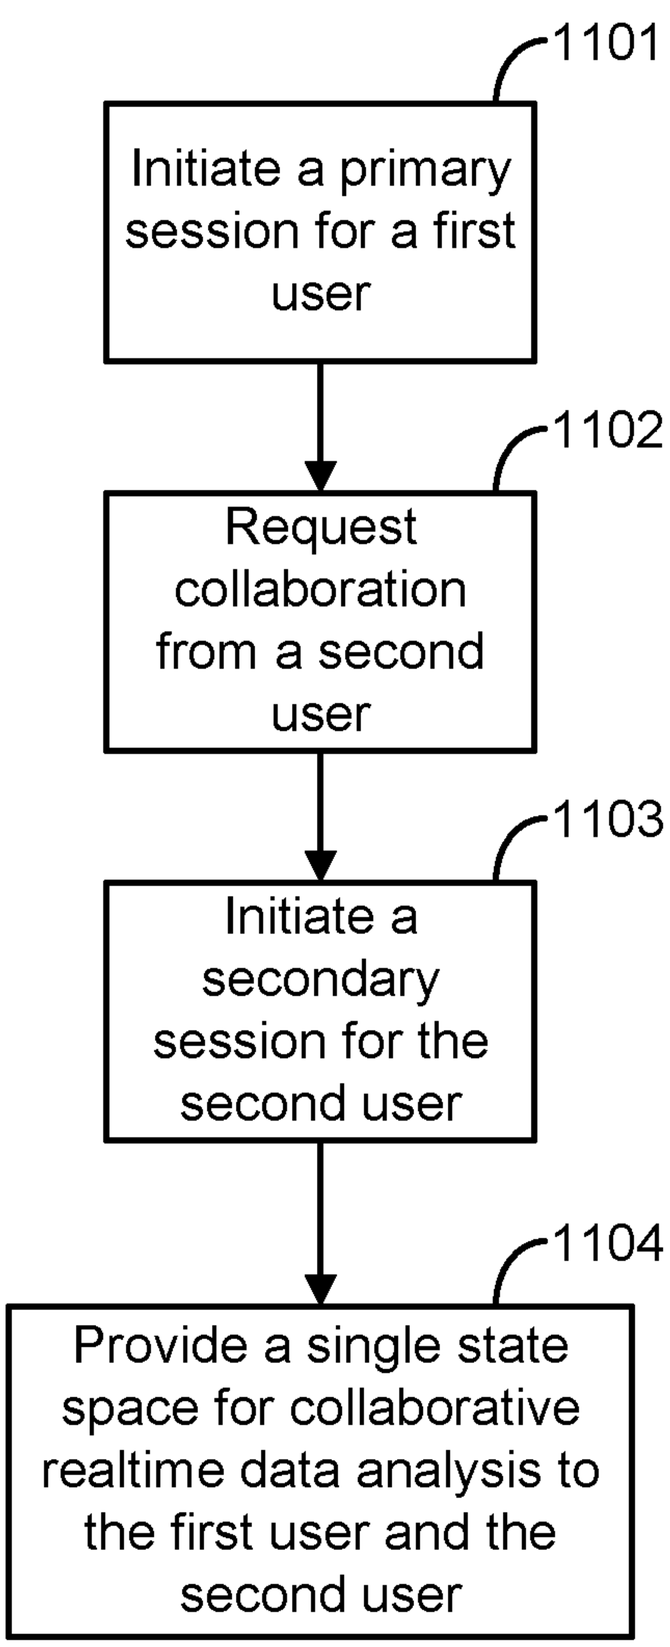

FIG. 11b

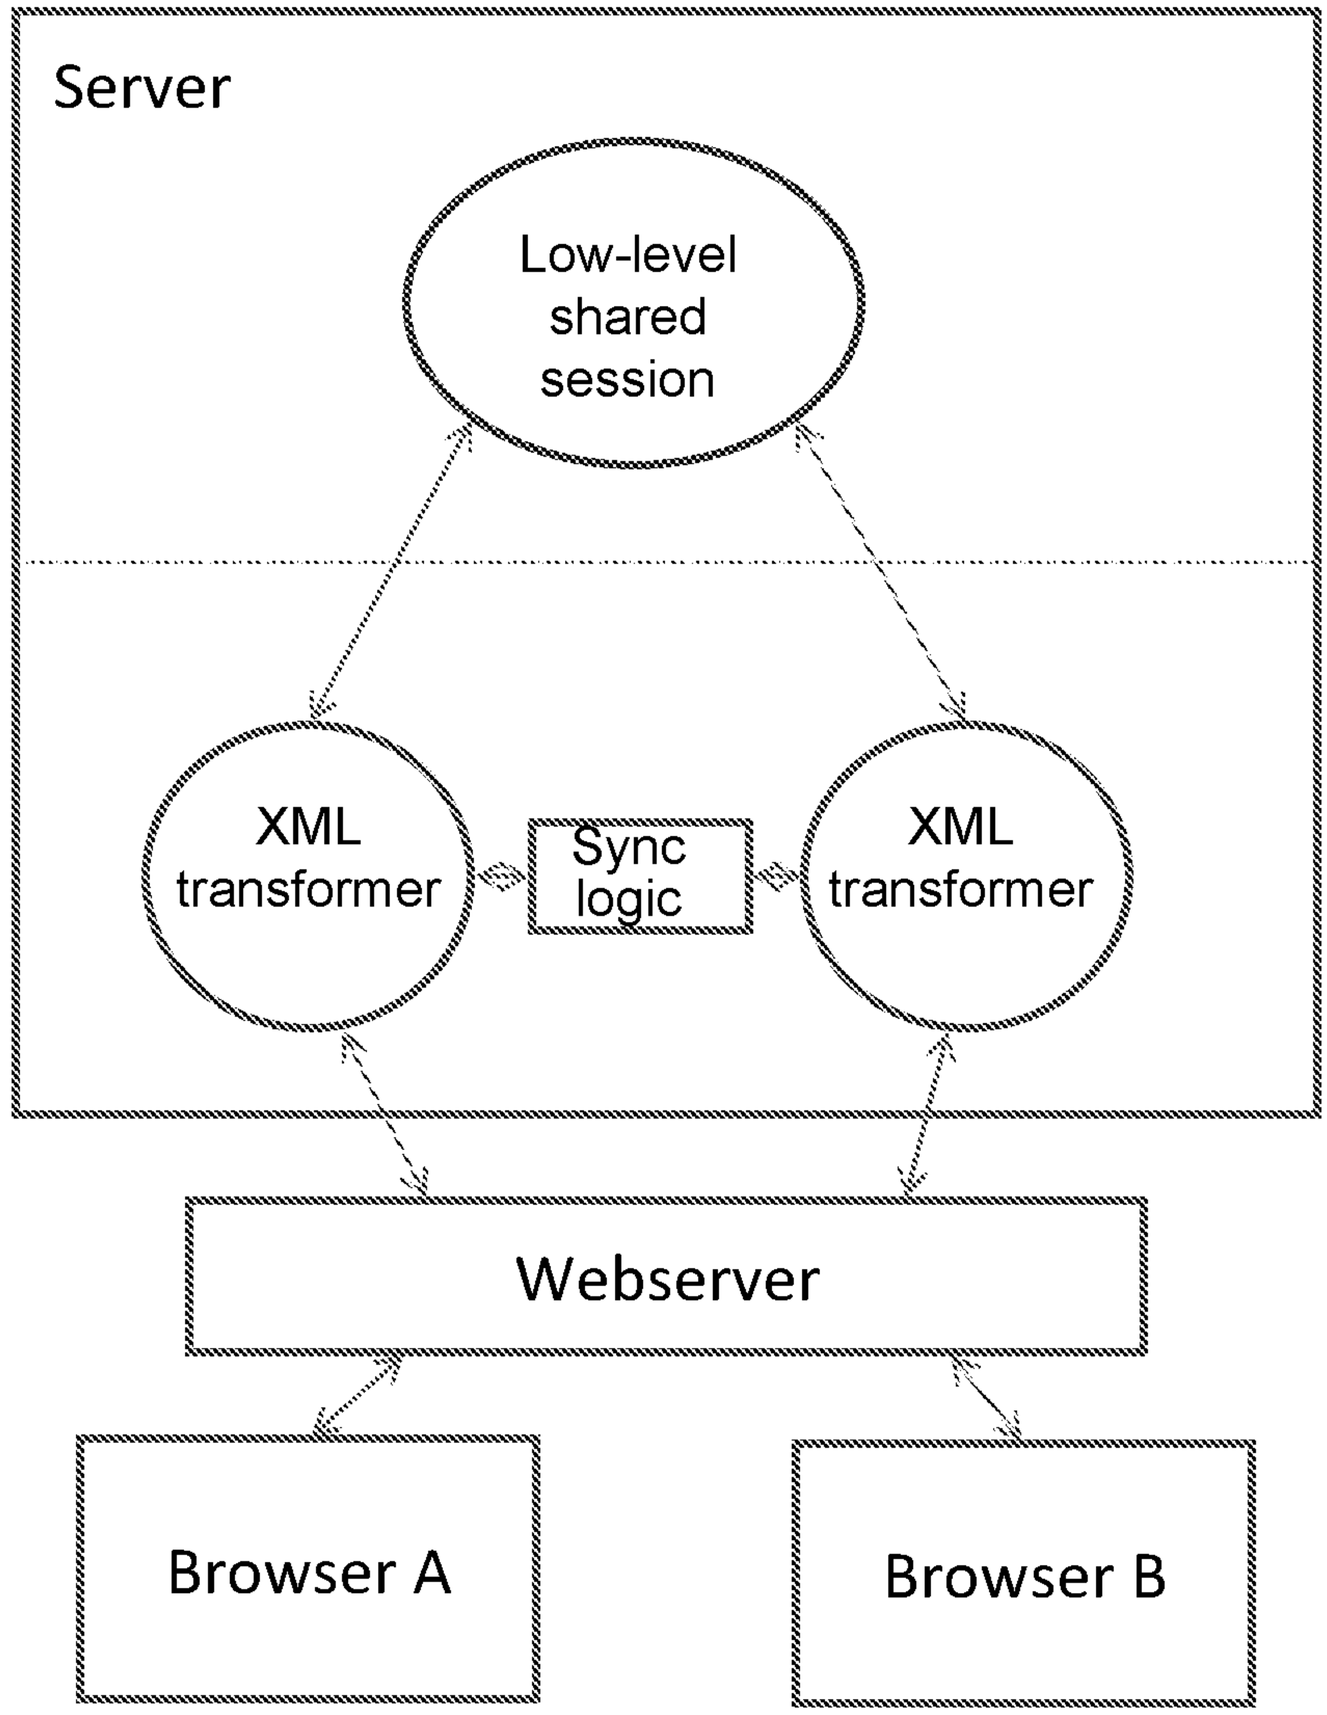

FIG. 12a

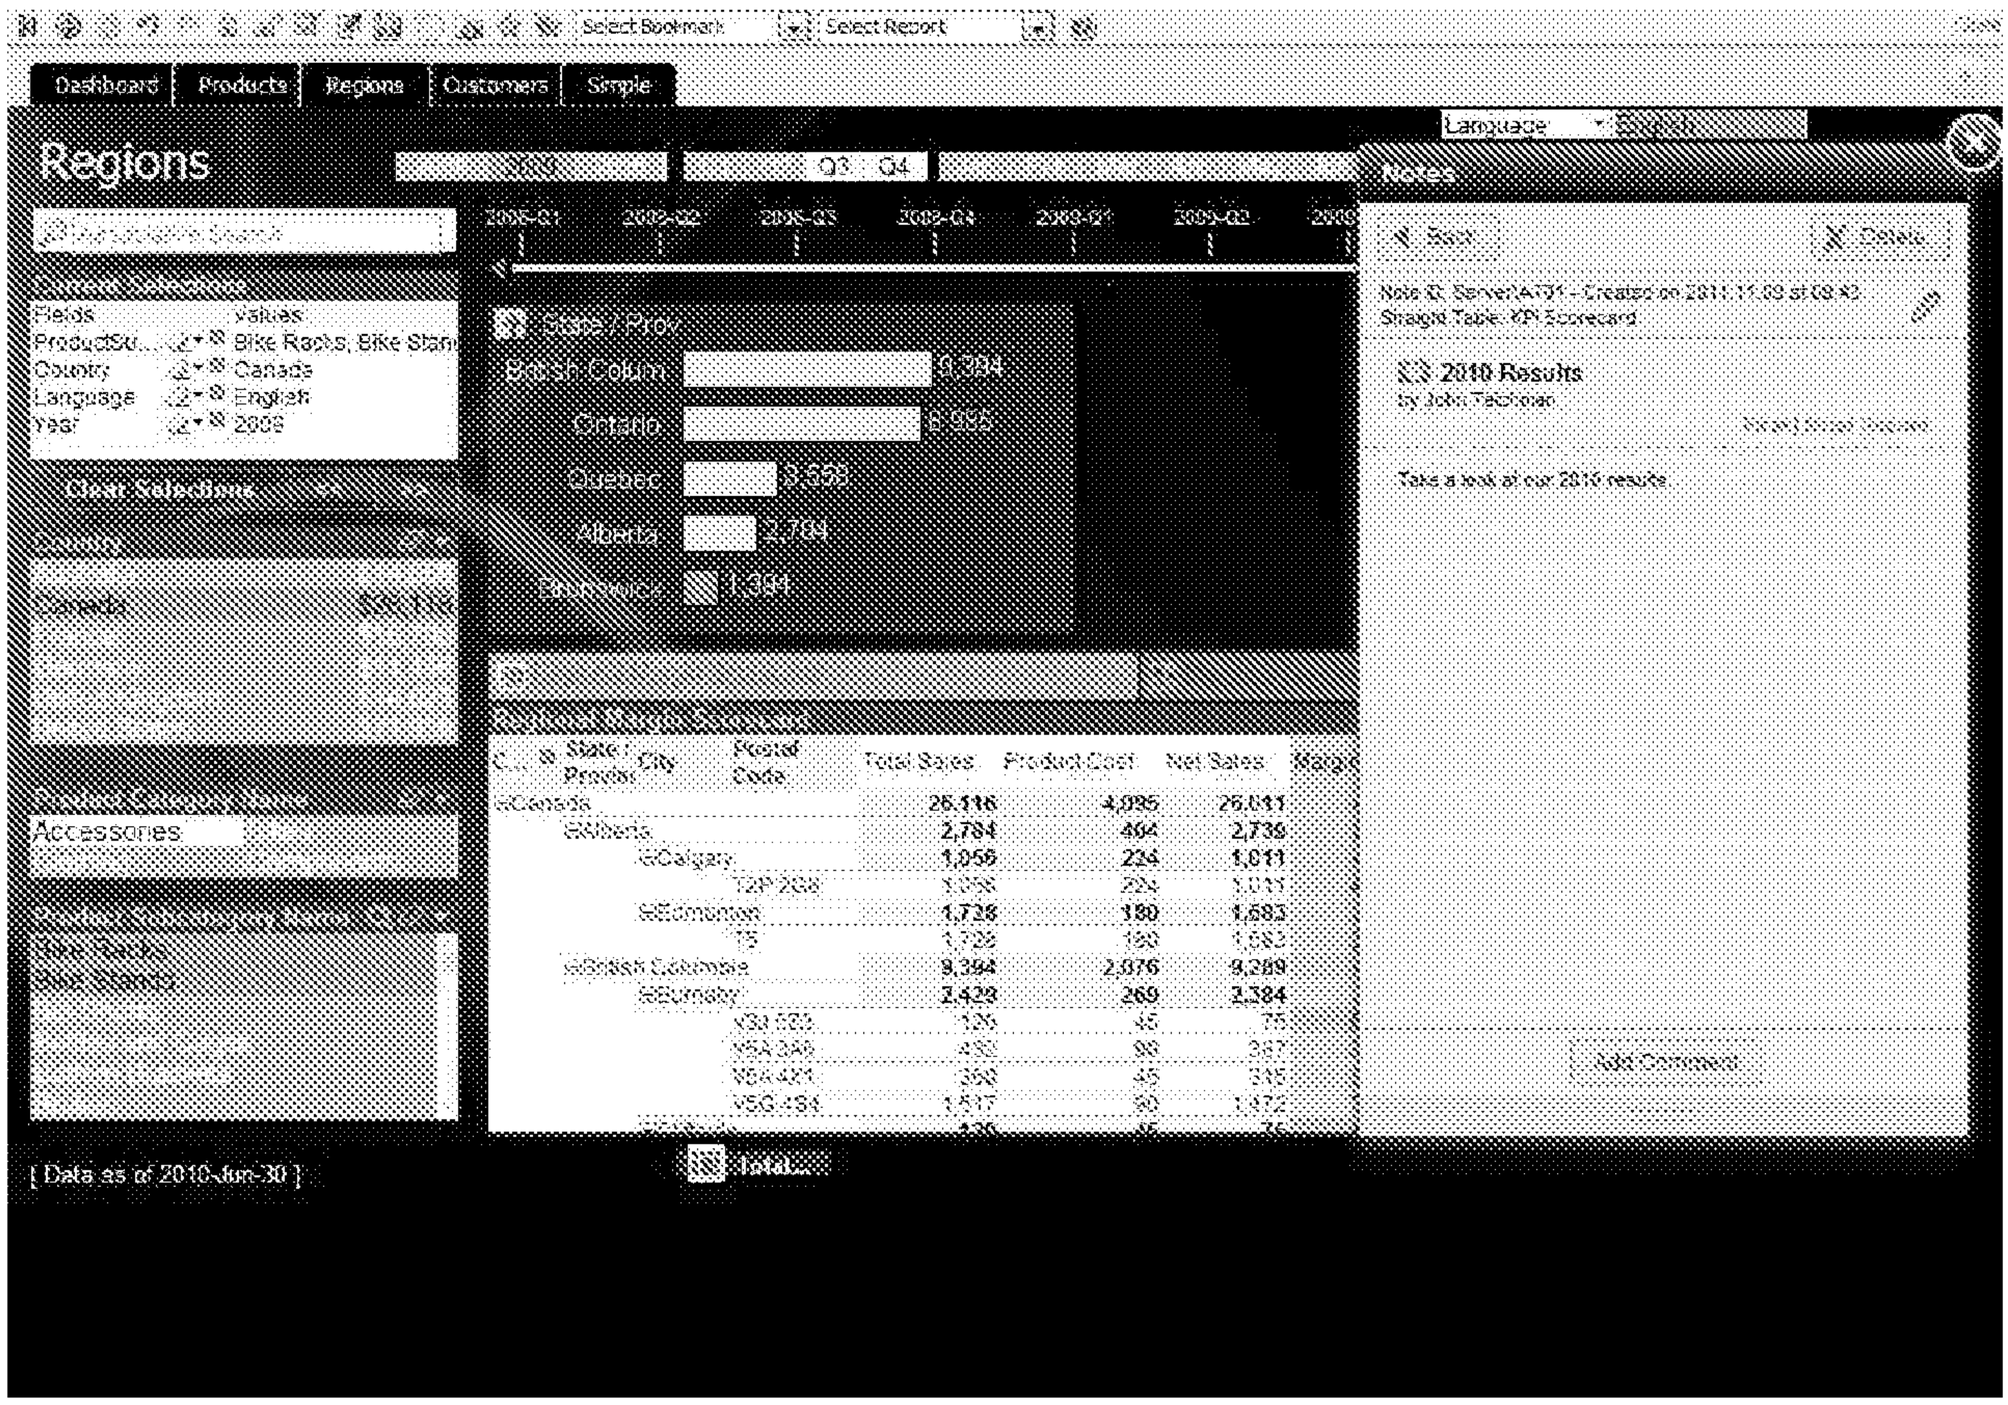

FIG. 12b

FIG. 13a

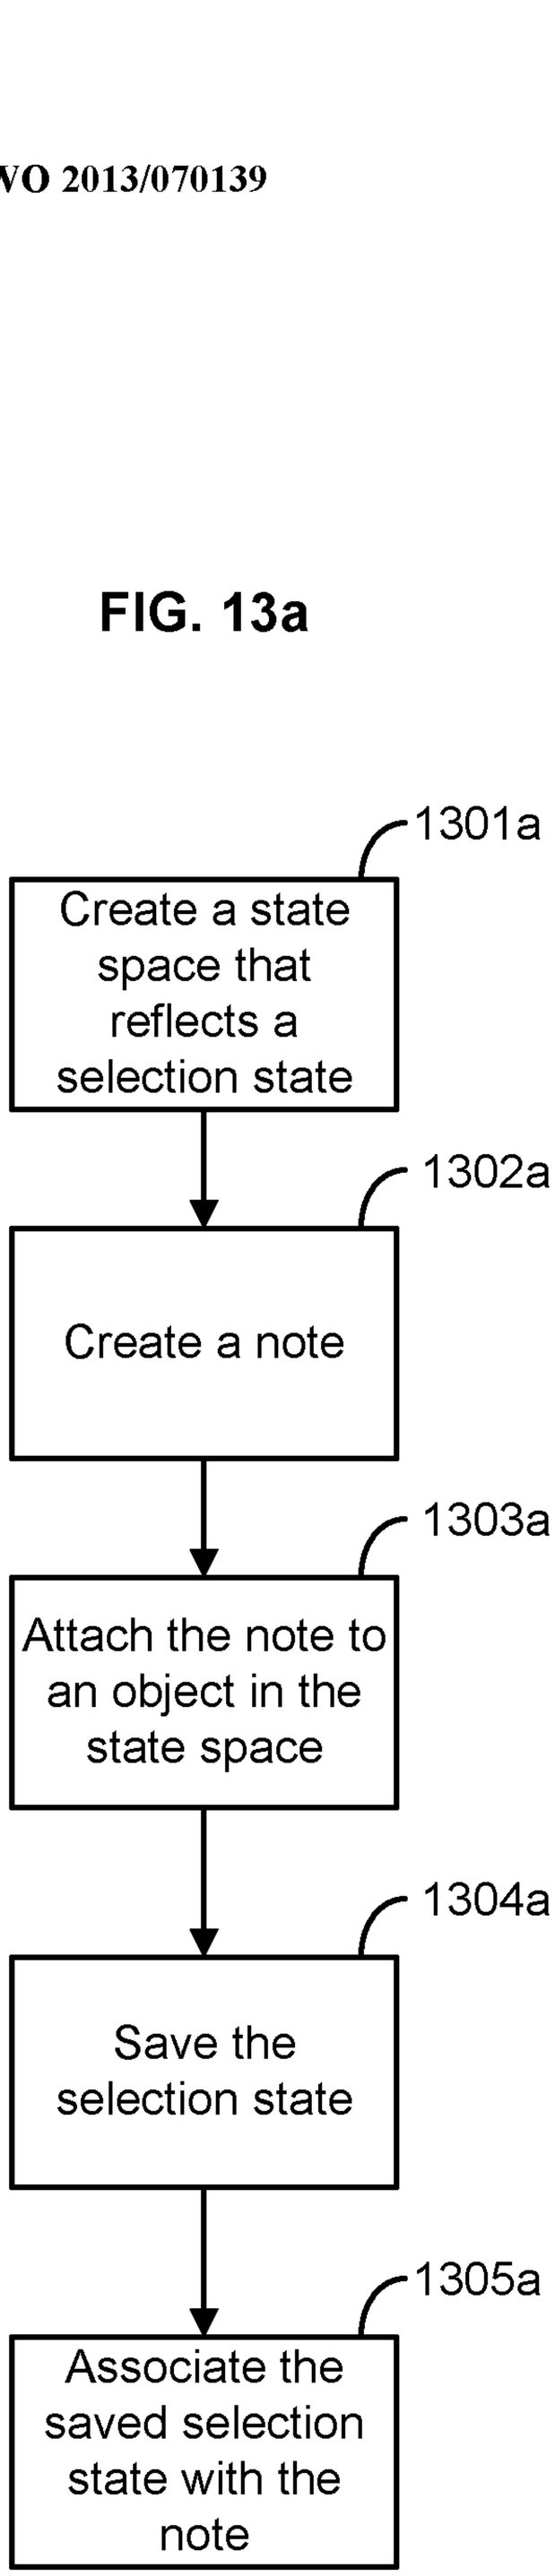

FIG. 13b

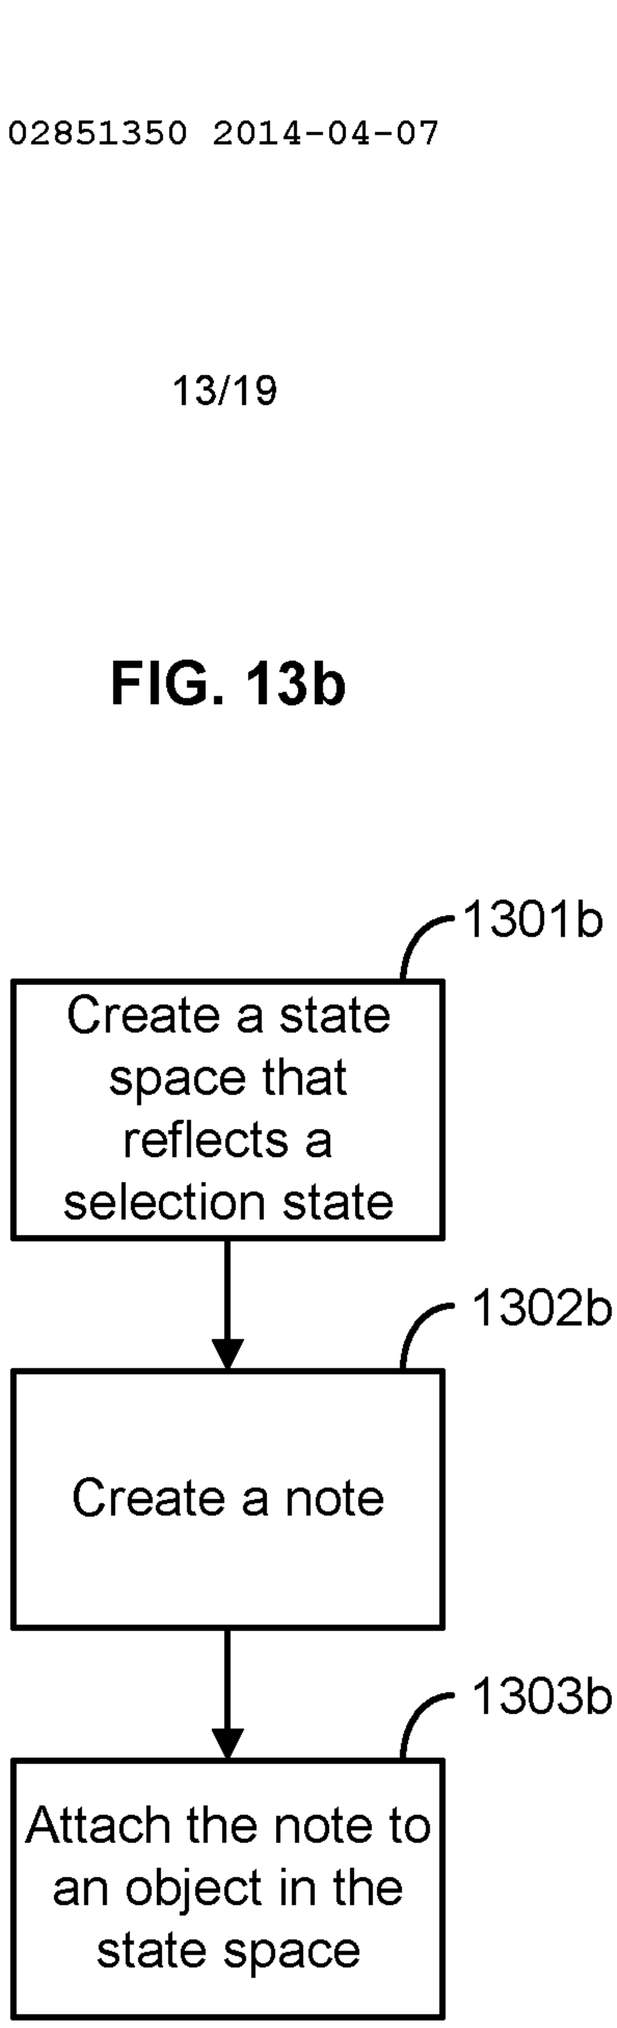

FIG. 13c

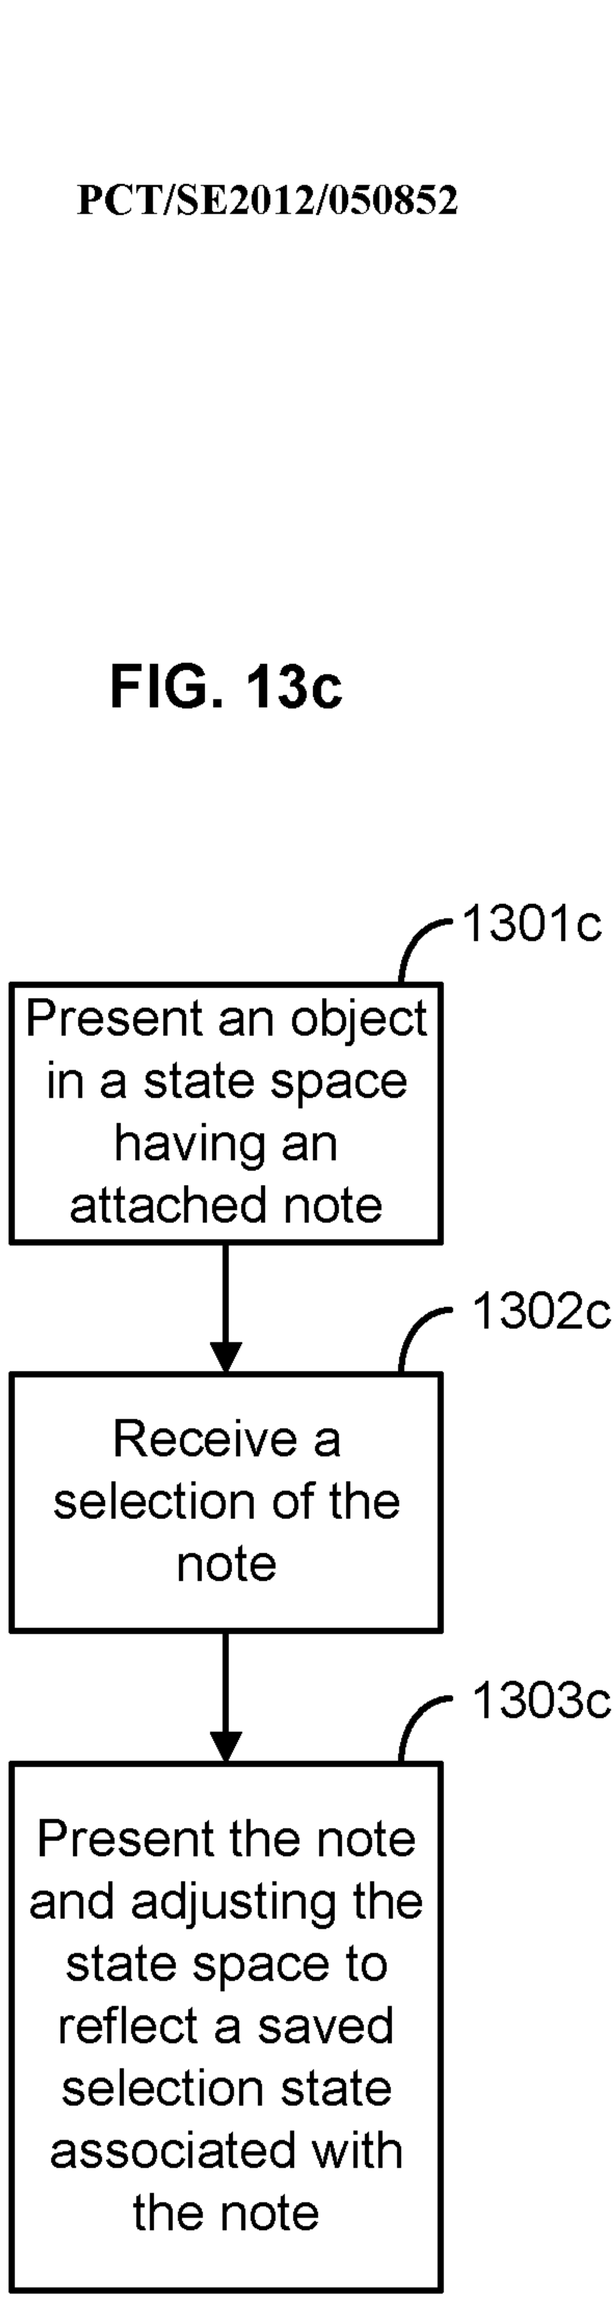

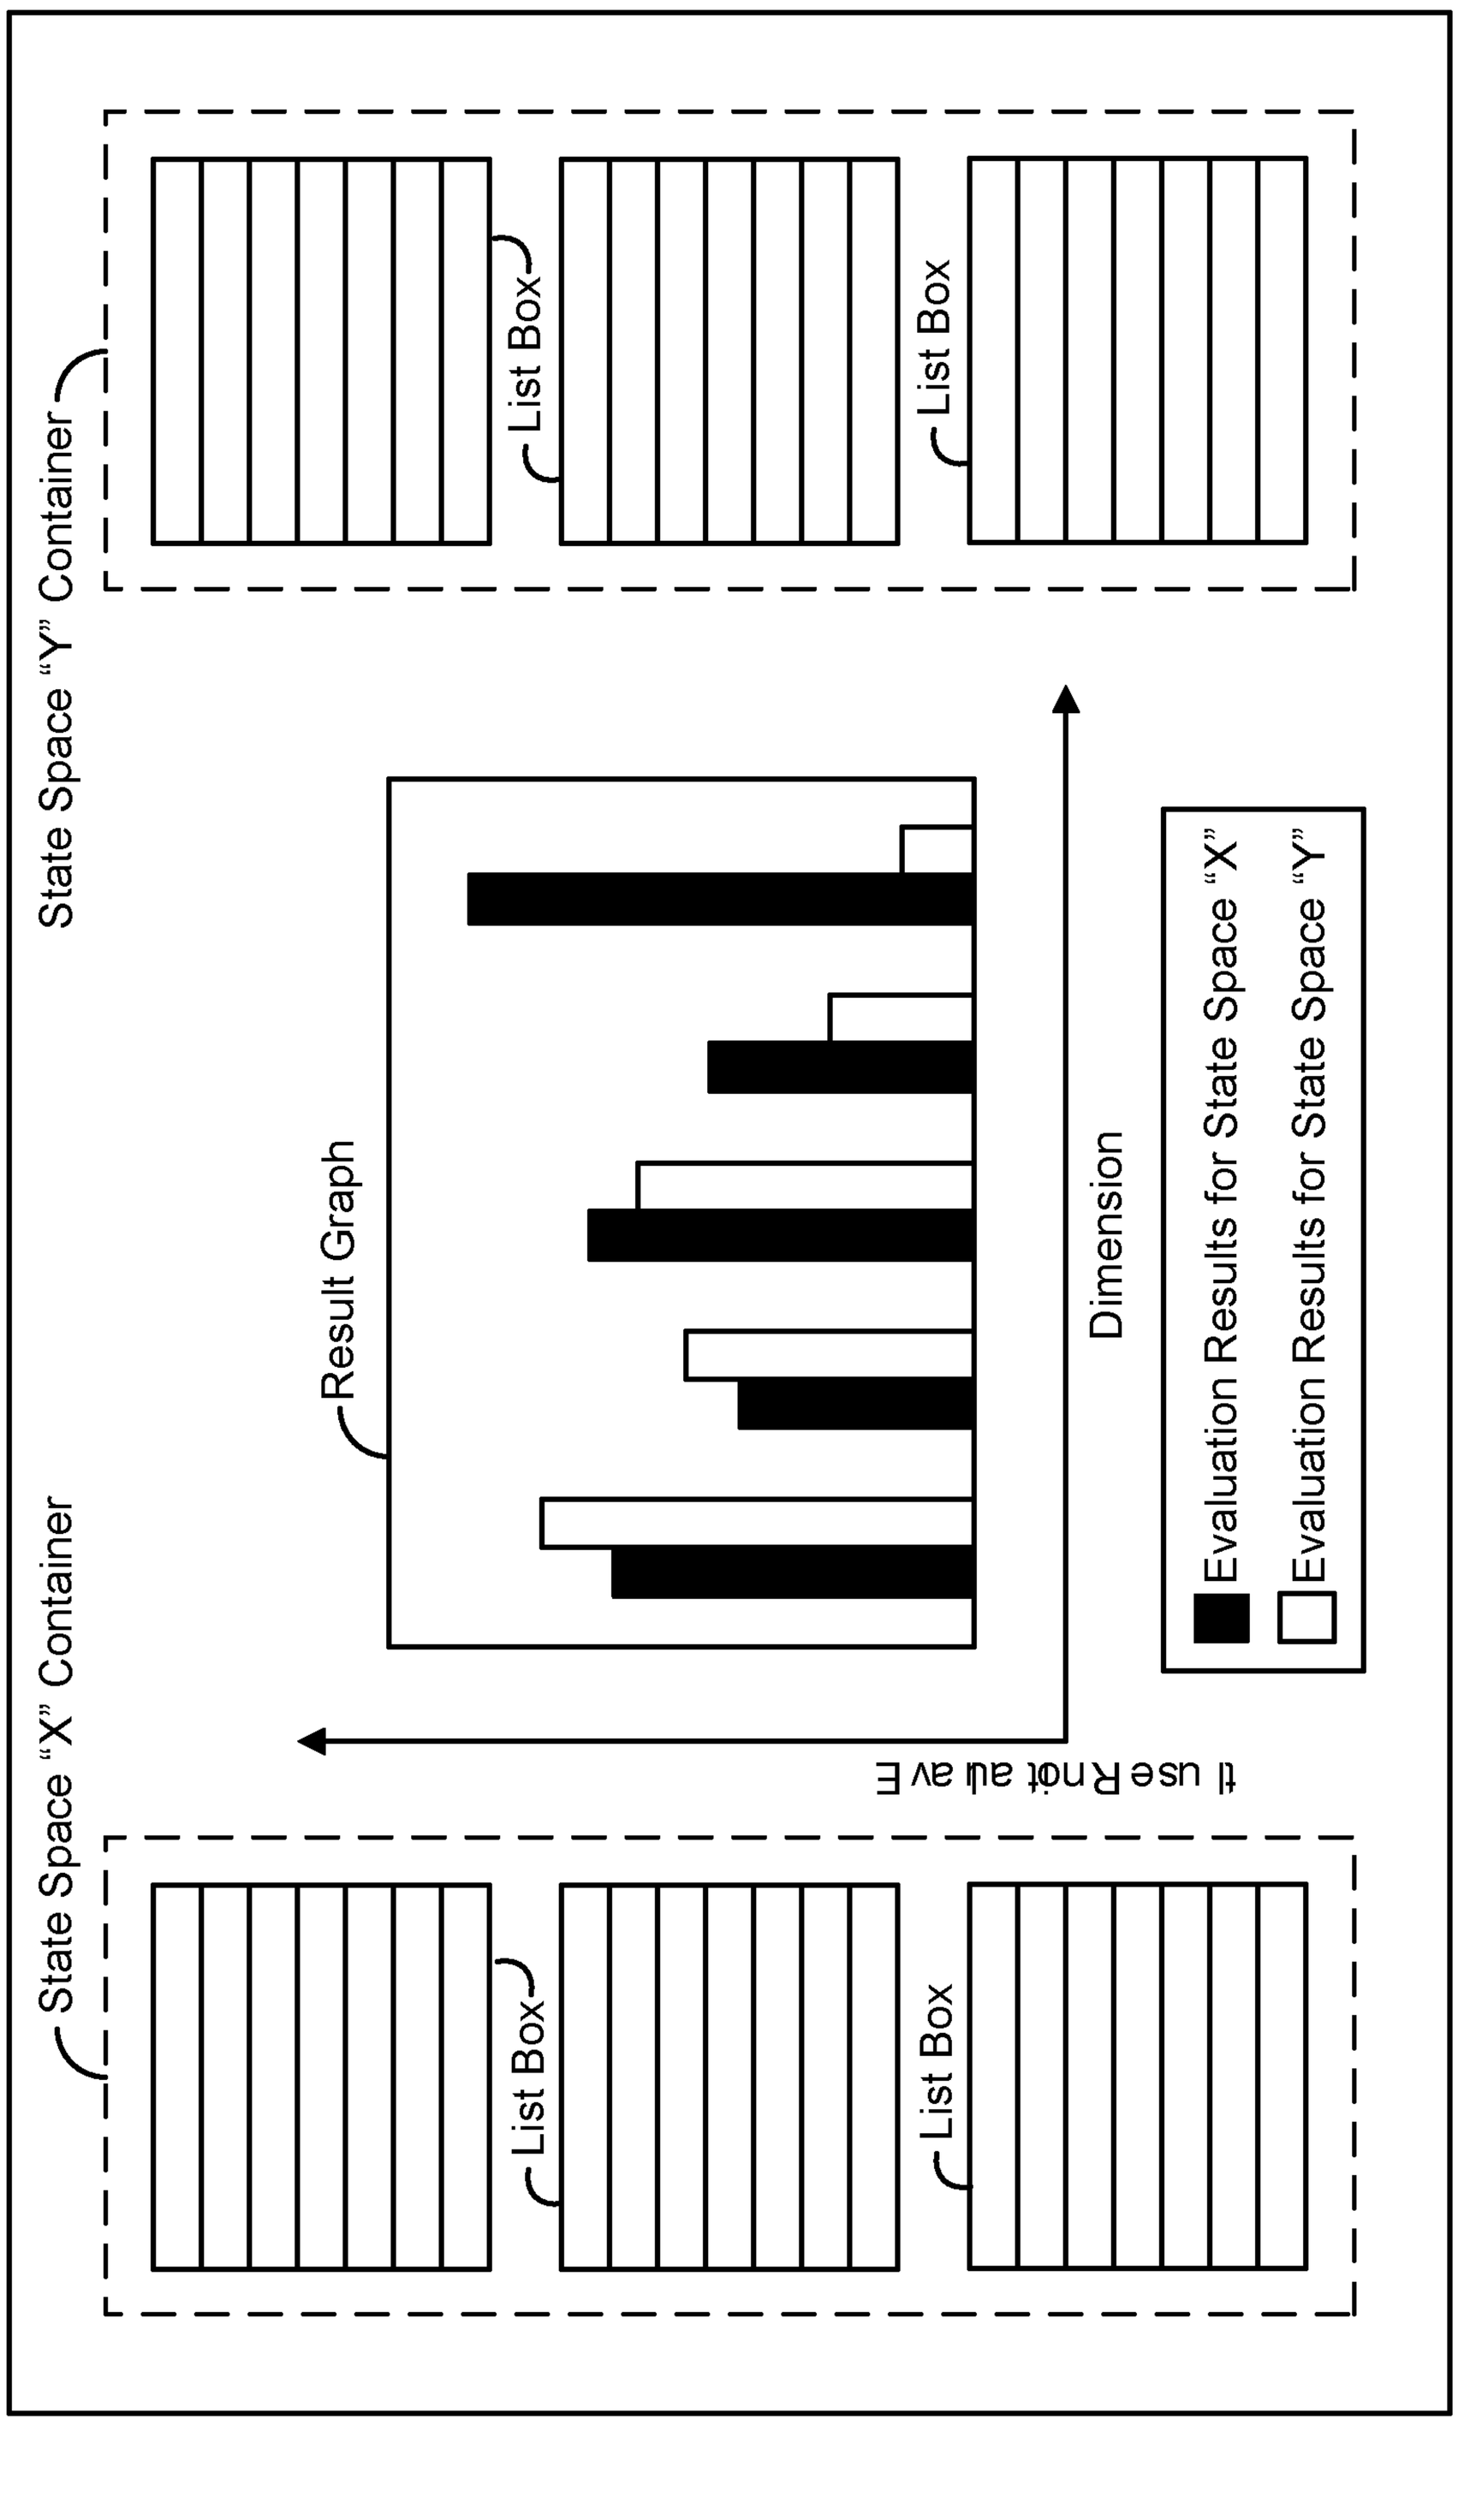
State Space "Y" Container
List Box
List Box
Result Graph
Dimension
Evaluation Results for State Space "X"
Evaluation Results for State Space "Y"
State Space "X" Container
List Box
List Box

FIG. 14

FIG. 15

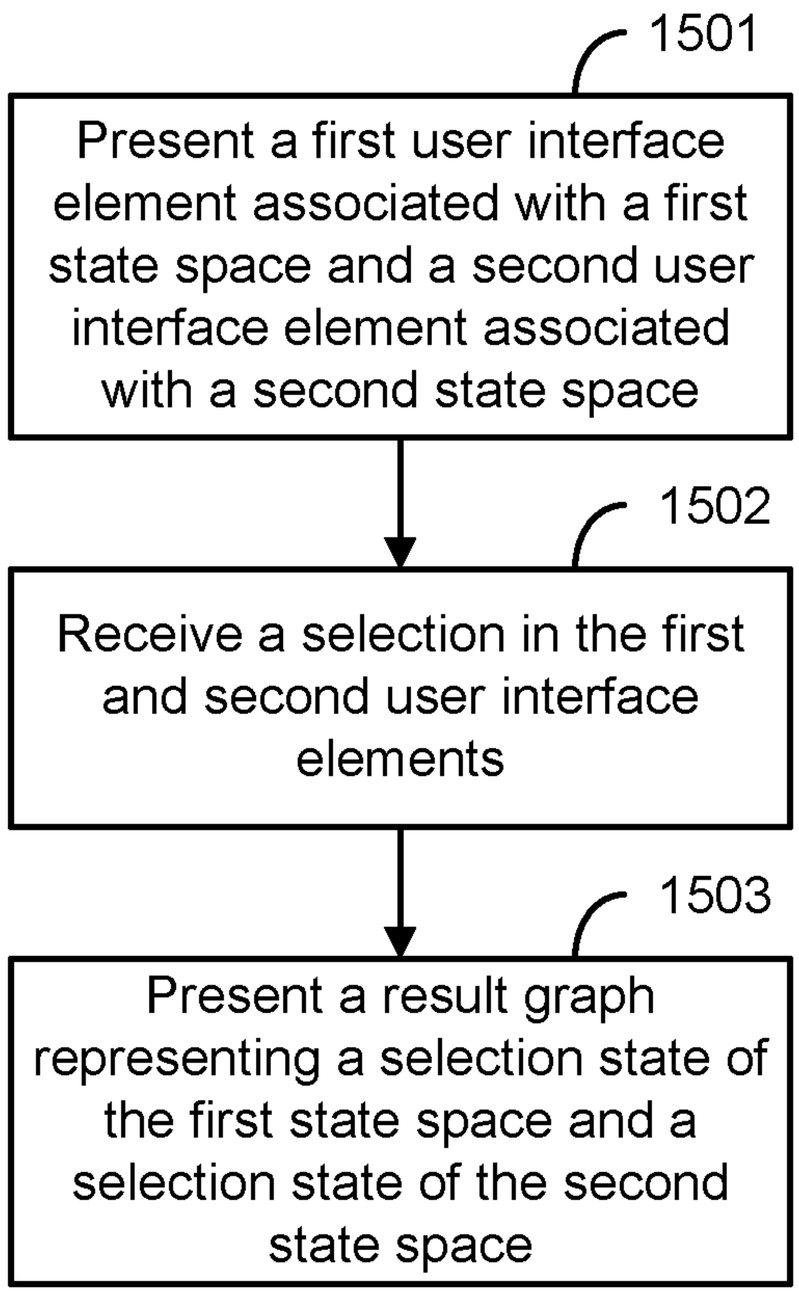
1501
Present a first user interface element associated with a first state space and a second user interface element associated with a second state space
1502
Receive a selection in the first and second user interface elements
1503
Present a result graph representing a selection state of the first state space and a selection state of the second state space

FIG. 16

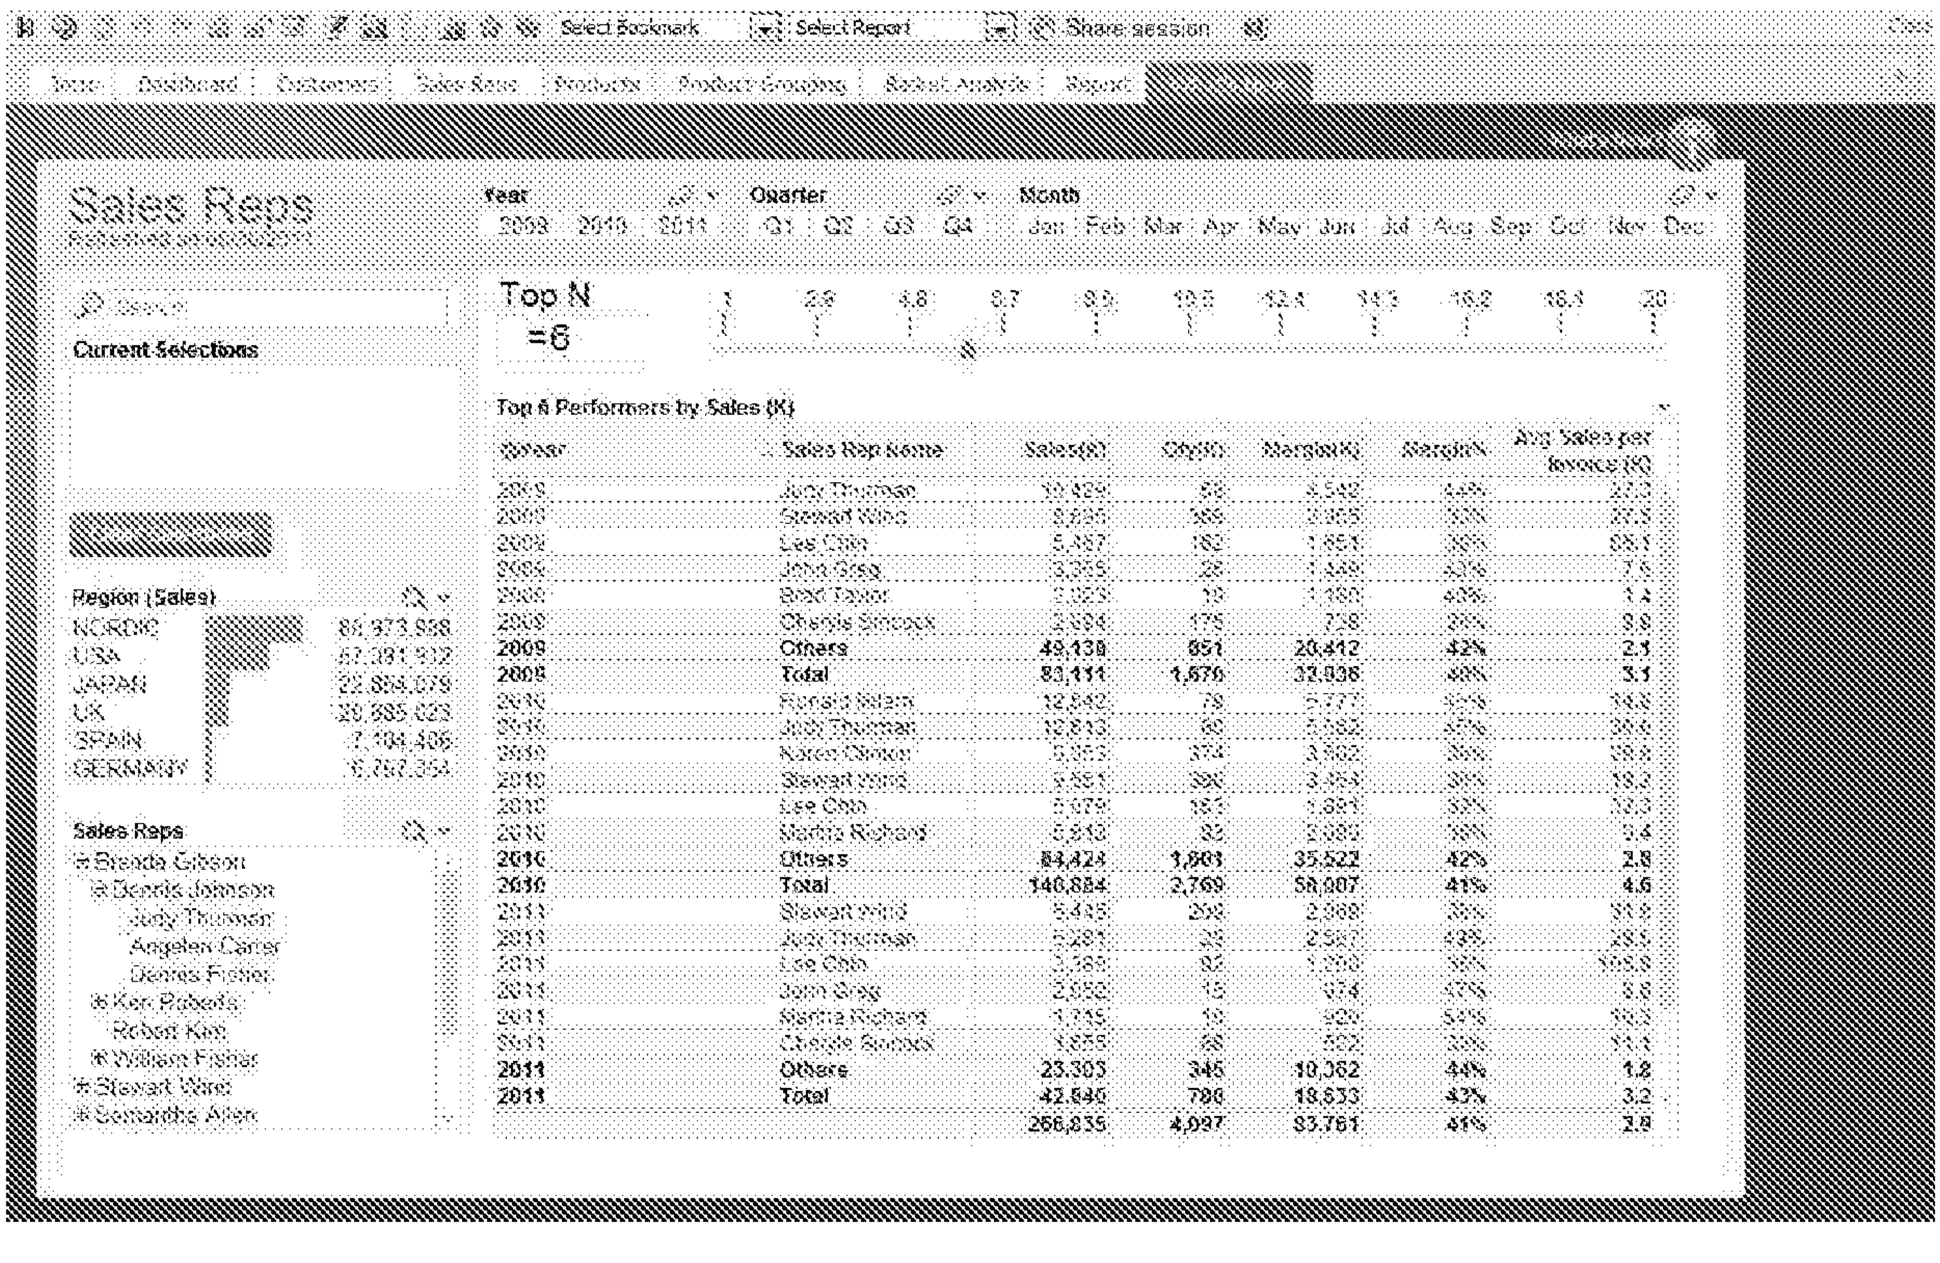
Select Bookmark
Select Report
Share session
Home
Dashboard
Customers
Sales Reps
Products
Product Grouping
Basket Analysis
Report
Sales Reps
Year
Quarter
Month
Current Selections
Region (Sales)
NORDIC 88,973,938
USA 67,381,312
JAPAN 22,886,079
UK 20,385,623
SPAIN 7,104,406
GERMANY 6,757,354
Sales Reps
Brenda Gibson
Dennis Johnson
Judy Thurman
Angelen Carter
Dennis Fisher
Ken Roberts
Robert Kerr
William Fisher
Stewart Ward
Samantha Allen
Top N
=6
Top 6 Performers by Sales (K)
Year
Sales Rep Name
Sales(K)
Qty(K)
Margin(K)
Margin%
Avg Sales per Invoice (K)
2009 Judy Thurman 10,429 53 4,542 44% 27.0
2009 Stewart Ward 9,636 569 3,955 33% 27.5
2009 Lee Cato 5,487 182 1,651 30% 26.1
2009 John Greg 3,365 22 1,448 43% 7.6
2009 Brad Taylor 2,023 10 1,160 57% 1.4
2009 Cheryle Simcock 2,004 175 738 37% 9.9
2009 Others 49,139 851 20,412 42% 2.1
2009 Total 83,111 1,676 32,936 40% 3.1
2010 Ronald Mann 12,542 79 5,777 46% 16.8
2010 Judy Thurman 12,513 60 5,582 45% 20.6
2010 Karen Clinton 5,353 374 3,502 65% 28.8
2010 Stewart Ward 5,351 330 3,454 65% 18.2
2010 Lee Cato 5,079 183 1,881 37% 10.3
2010 Martha Richard 5,013 82 2,089 42% 9.4
2010 Others 84,424 1,661 35,522 42% 2.8
2010 Total 140,884 2,769 58,807 41% 4.6
2011 Stewart Ward 5,445 268 2,368 43% 31.8
2011 Judy Thurman 5,281 23 2,587 49% 23.5
2011 Lee Cato 3,388 82 1,200 35% 105.8
2011 John Greg 2,952 15 974 33% 9.8
2011 Martha Richard 1,715 10 929 54% 19.0
2011 Cheryle Simcock 1,855 28 802 43% 11.1
2011 Others 23,305 345 10,362 44% 1.8
2011 Total 42,540 786 18,833 43% 3.2
266,535 4,097 83,761 41% 2.8

FIG. 17a

Customer

A
B
C
D
E
F

Product

X
Y

Data set

Product	Customer	Sales
X	A	35
X	B	39
X	C	33
X	D	4
X	E	7
X	F	17
Y	A	30
Y	B	15
Y	C	31
Y	E	16
Y	F	15

FIG. 17b

Full table

Customer	Sum(Sales)
A	65
B	54
C	64
D	4
E	23
F	32

FIG. 17c

First 3 (2 + others)

Customer	Sum(Sales)
A	65
B	54
Others	**123**

FIG. 17d

Largest 3 (2 + others)

Customer	Sum(Sales)
A	65
C	64
Others	**113**

FIG. 17e

Sum(Sales) >= 50

Customer	Sum(Sales)
A	65
B	54
C	64
Others	**59**

FIG. 17f

Top 80% of total sales

Customer	Sum(Sales)
A	65
C	64
B	54
E	32
Others	**27**

FIG. 18a

Full table		
Product	Customer	Sum(Sales)
X	A	35
X	B	39
X	C	33
X	D	4
X	E	7
X	F	17
Y	A	30
Y	B	15
Y	C	31
Y	E	16
Y	F	15

FIG. 18b

Largest 3 (2 + others)		
Product	Customer	Sum(Sales)
X	A	35
X	B	39
X	**Others**	**61**
Y	A	30
Y	C	31
Y	**Others**	**46**

FIG. 18c

Largest 3 (2 + others) - Global grouping		
Product	Customer	Sum(Sales)
X	A	35
X	C	33
X	**Others**	**67**
Y	A	30
Y	C	31
Y	**Others**	**46**

FIG. 19

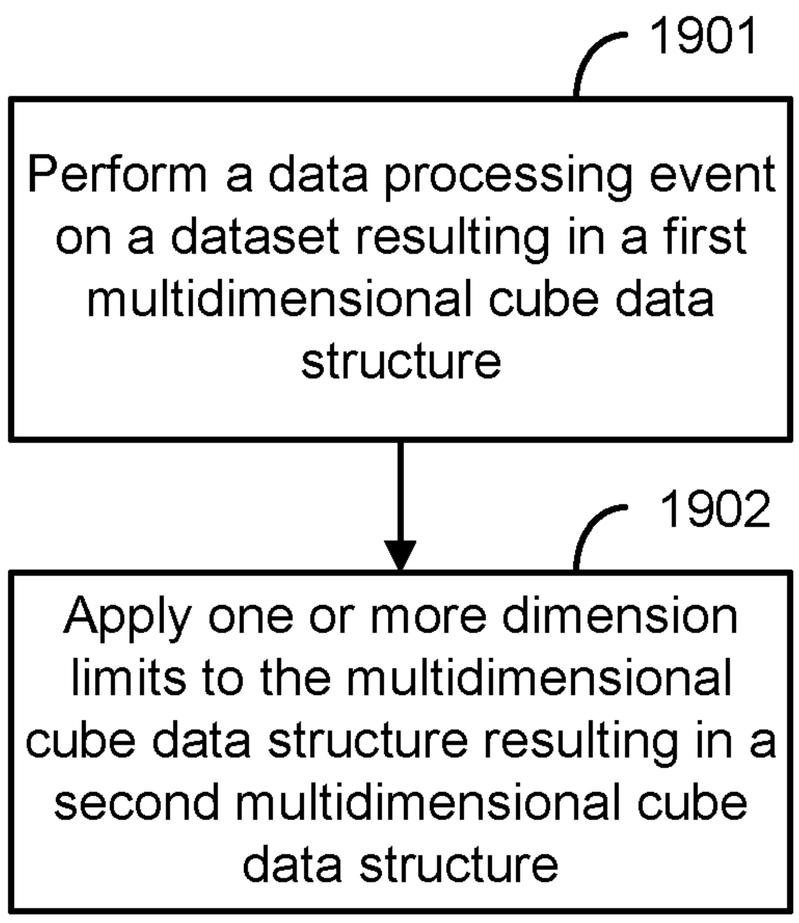
1901
Perform a data processing event on a dataset resulting in a first multidimensional cube data structure
1902
Apply one or more dimension limits to the multidimensional cube data structure resulting in a second multidimensional cube data structure

FIG. 11a

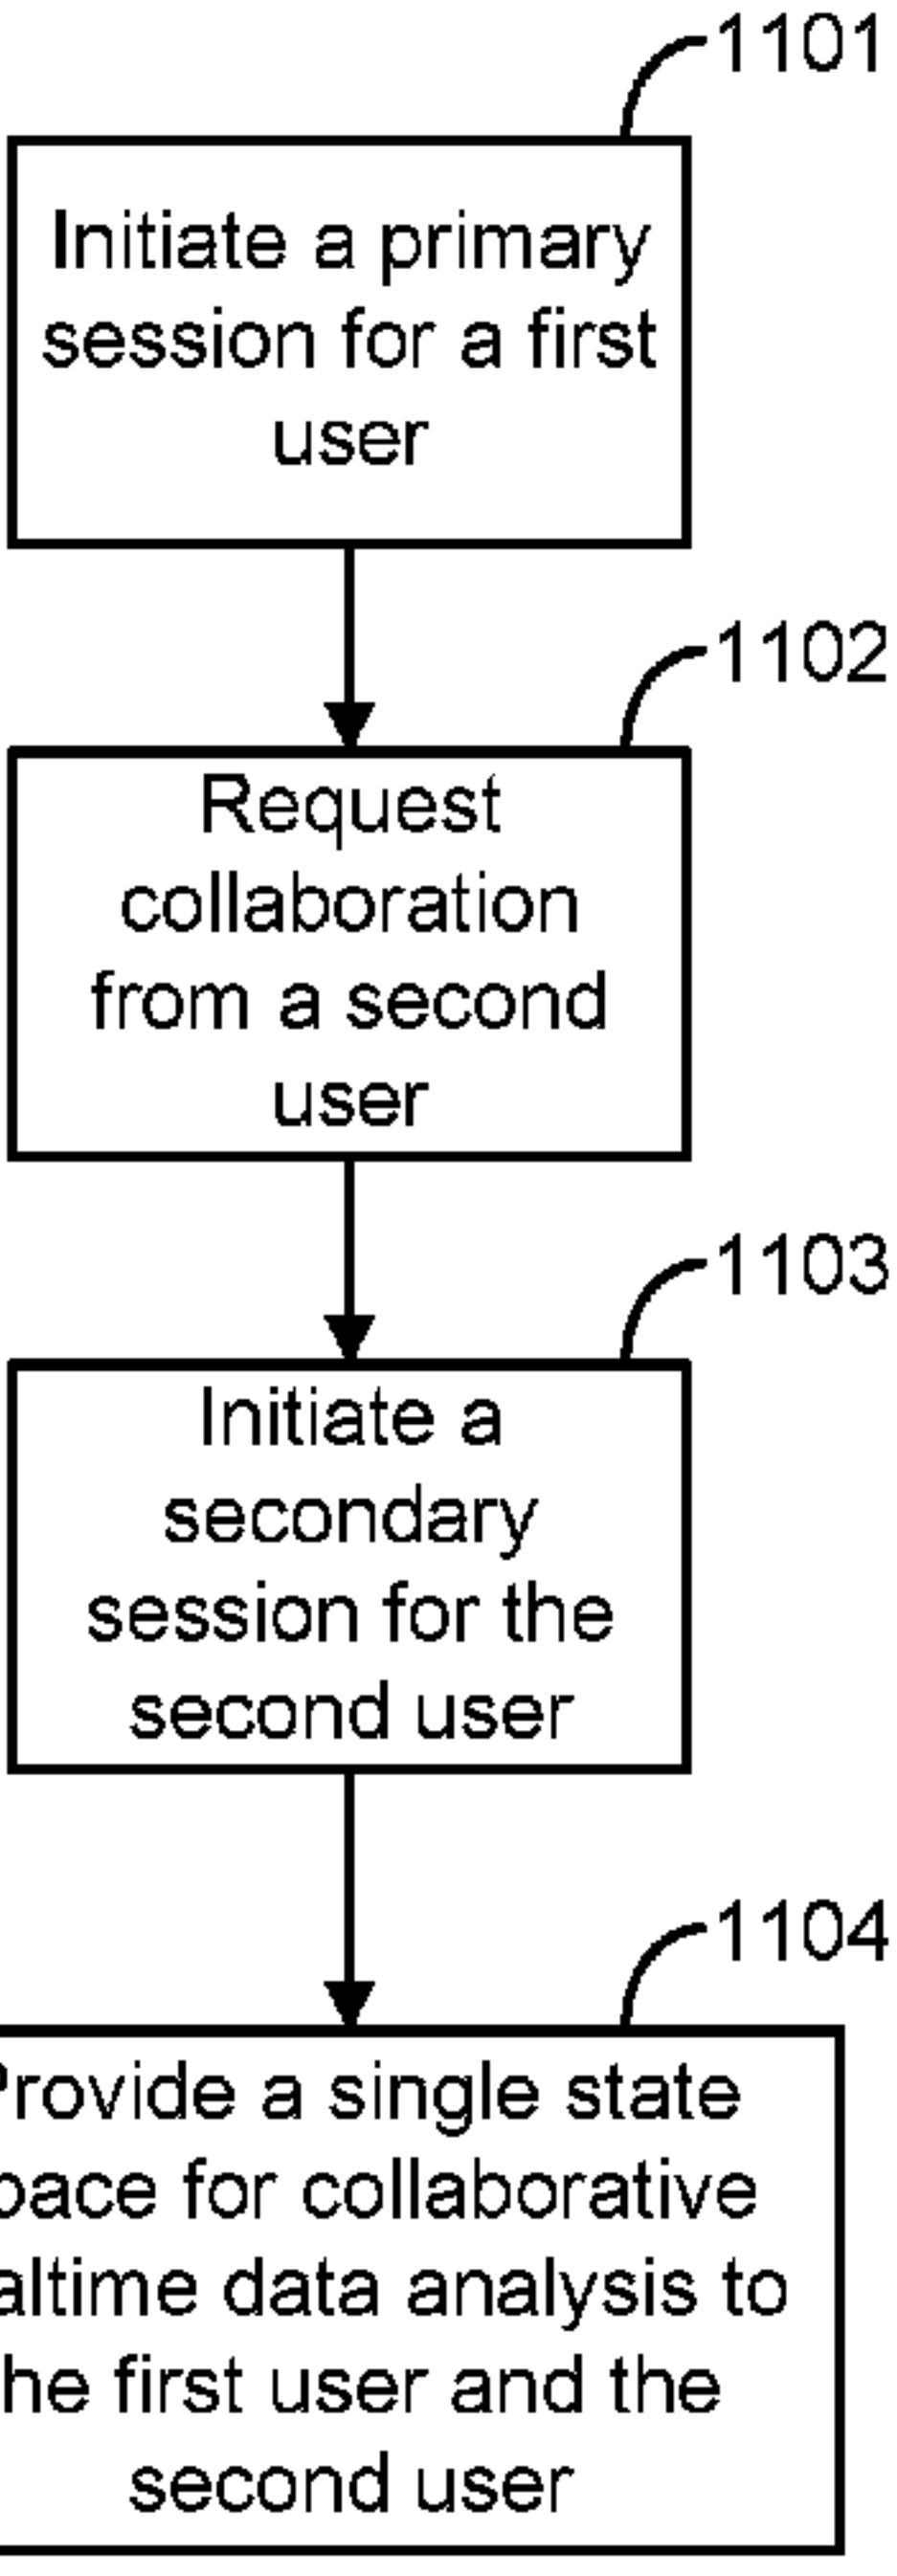
1101
Initiate a primary session for a first user
1102
Request collaboration from a second user
1103
Initiate a secondary session for the second user
1104
Provide a single state space for collaborative realtime data analysis to the first user and the second user