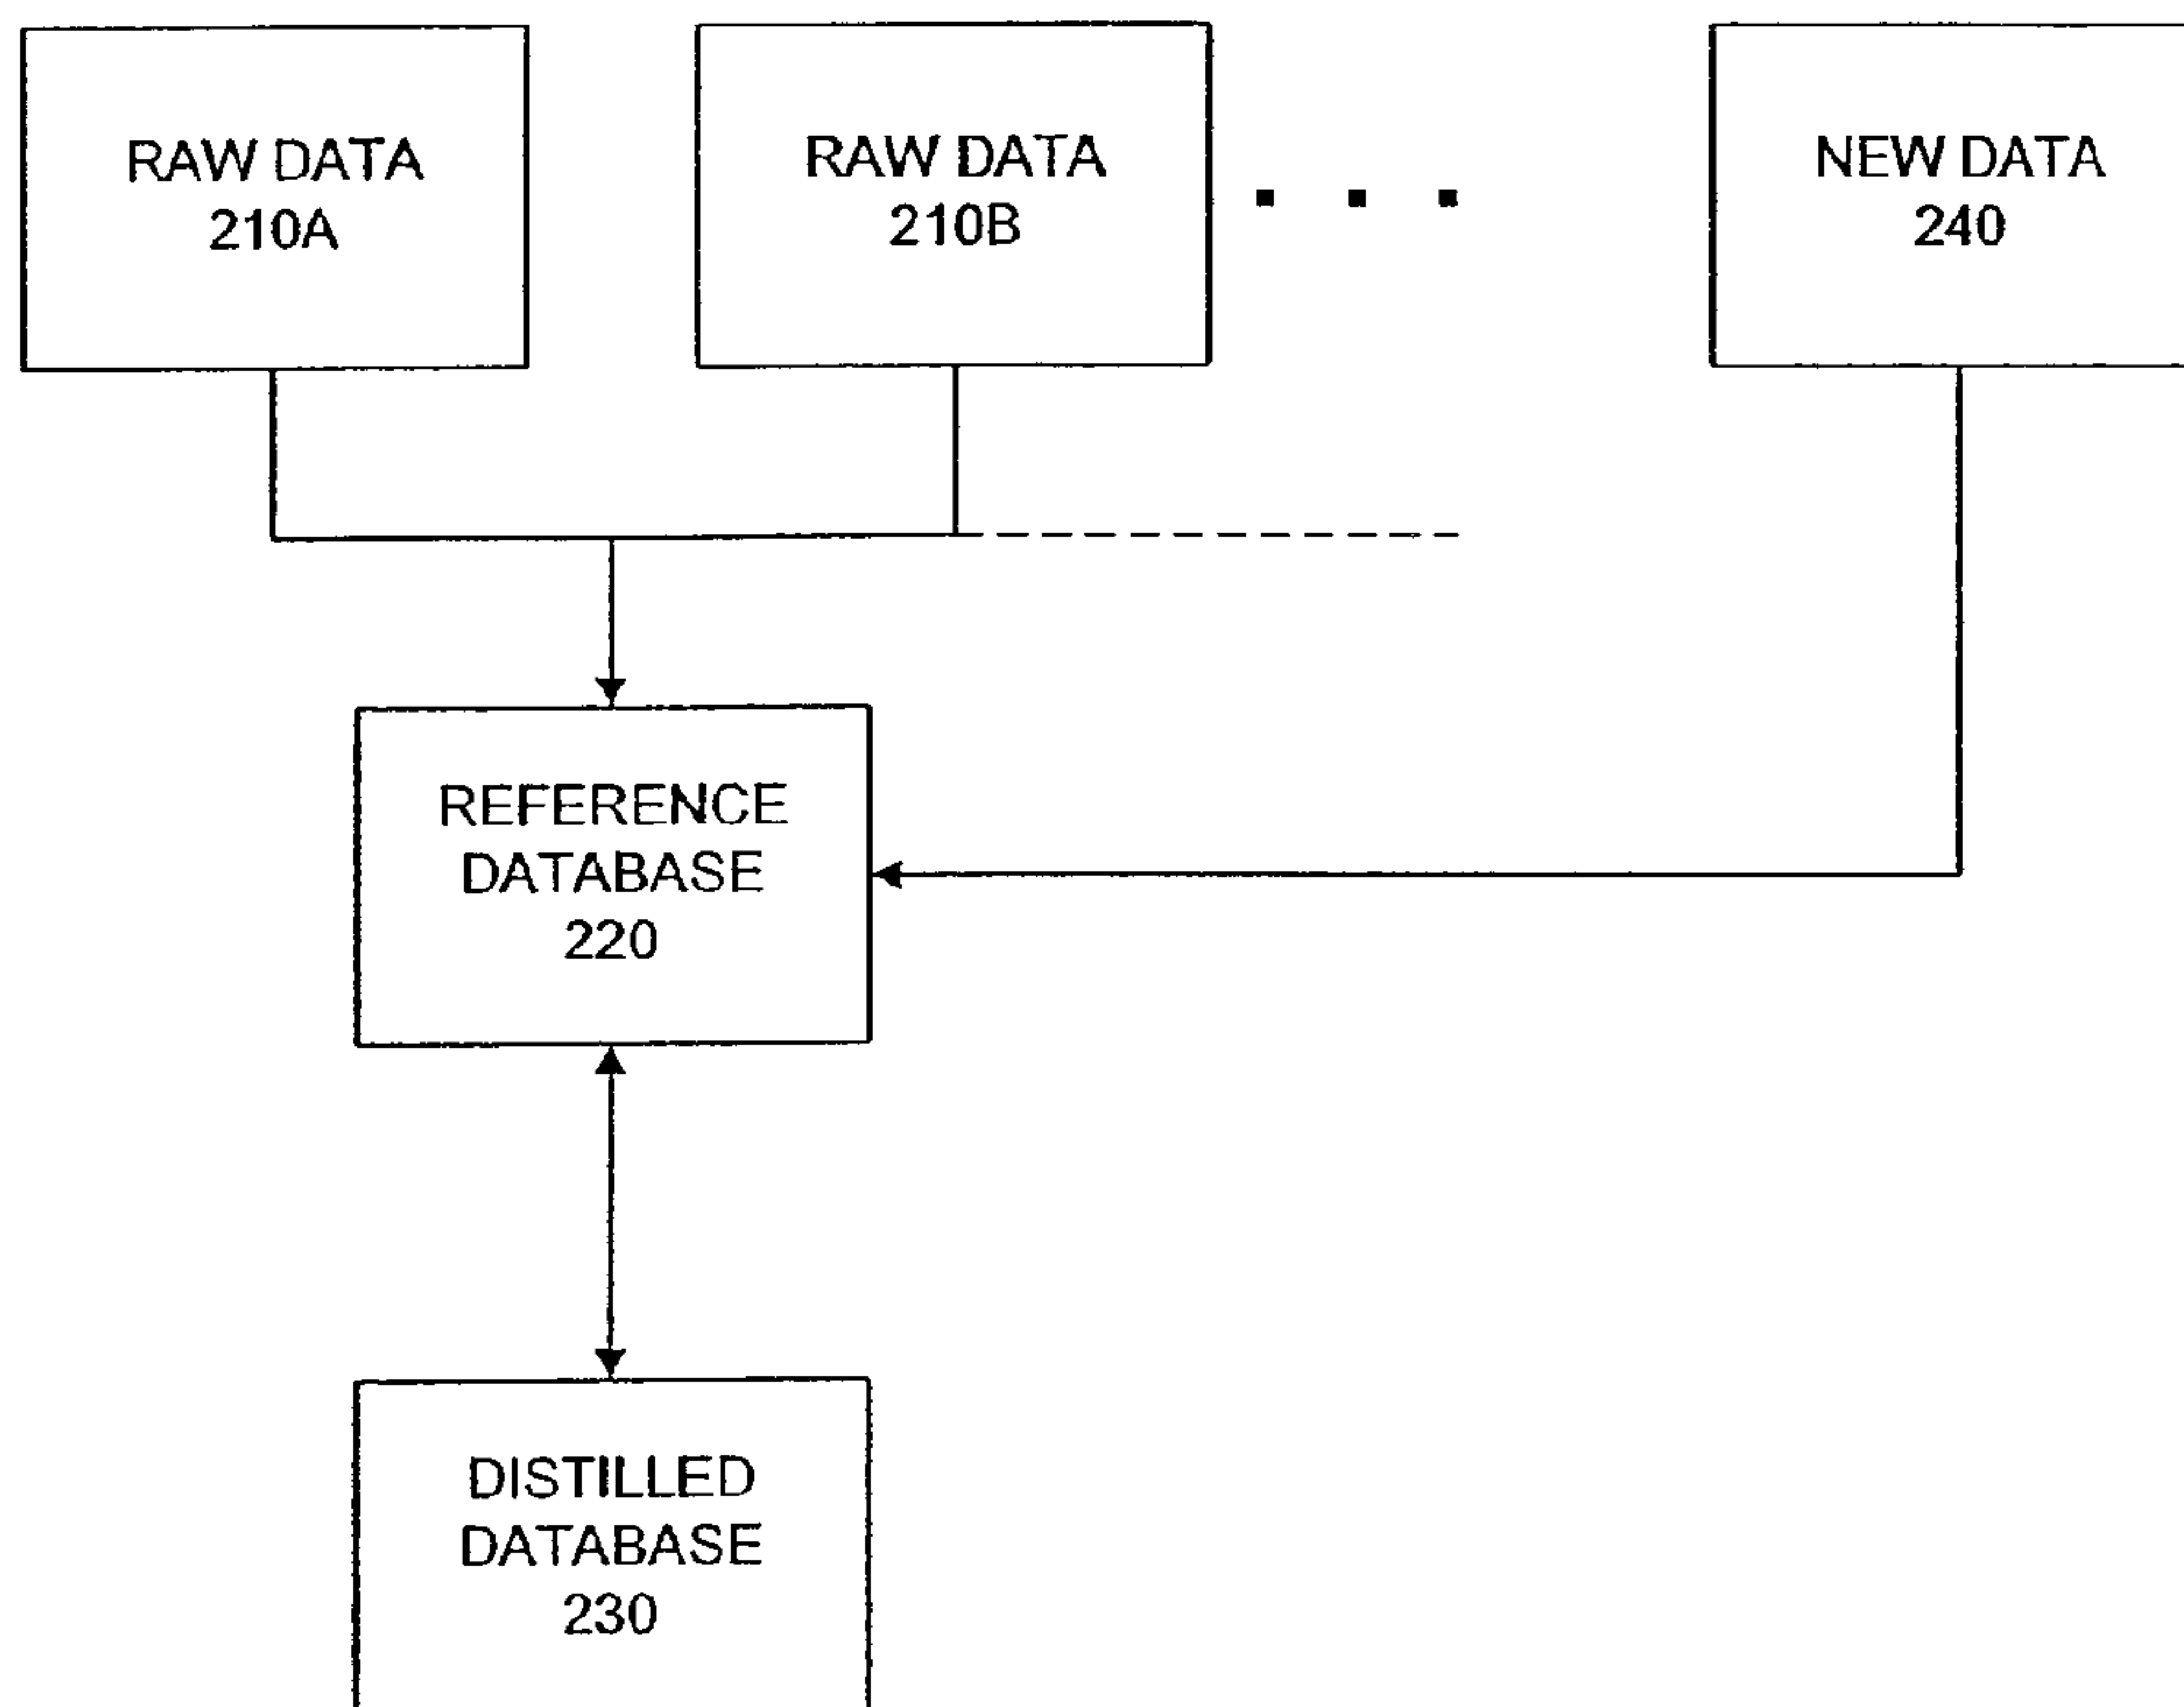




(86) Date de dépôt PCT/PCT Filing Date: 2000/07/14
(87) Date publication PCT/PCT Publication Date: 2001/01/25
(45) Date de délivrance/Issue Date: 2007/12/11
(85) Entrée phase nationale/National Entry: 2002/01/16
(86) N° demande PCT/PCT Application No.: US 2000/019195
(87) N° publication PCT/PCT Publication No.: 2001/006414
(30) Priorités/Priorities: 1999/07/20 (US09/357,301);
1999/10/06 (US09/412,970); 2000/07/14 (US09/617,047)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01),
G06F 17/16 (2006.01)
(72) Inventeur/Inventor:
GRUENWALD, BJORN J., US
(73) Propriétaire/Owner:
PRIMENTIA, INC., US
(74) Agent: BORDEN LADNER GERVAIS LLP

(54) Titre : PROCEDE ET SYSTEME D'ORGANISATION DE DONNEES
(54) Title: METHOD AND SYSTEM FOR ORGANIZING DATA



(57) Abrégé/Abstract:

A system and method for organizing raw data from one or more sources uses an improved mechanism for identifying duplicate data between fields e.g., columns) in the databases. The fields may be similar fields within a single database or similar or identical field within a pair of databases and as organized as arrays or field vectors. The present invention sorts each of the field vectors and if necessary, partitions them by common value. A number of comparisons required to identify the duplicate data between the field vectors is reduced by feeding back a difference between the compared values. This difference is used to adjust indices into the field vectors for subsequent comparison.



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

CORRECTED VERSION

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 January 2001 (25.01.2001)

PCT

(10) International Publication Number
WO 01/06414 A2

(51) International Patent Classification⁷: G06F 17/30

Bjorn, J. [US/US]; 526 Pineville Road, Newtown, PA 18940-0330 (US).

(21) International Application Number: PCT/US00/19195

(22) International Filing Date: 14 July 2000 (14.07.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/357,301 20 July 1999 (20.07.1999) US
09/412,970 6 October 1999 (06.10.1999) US
09/617,047 14 July 2000 (14.07.2000) US

(71) Applicant (for all designated States except US): INMENTIA, INC. [US/US]; 526 Pineville Road, Newtown, PA 18940-0330 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): GRUENWALD,

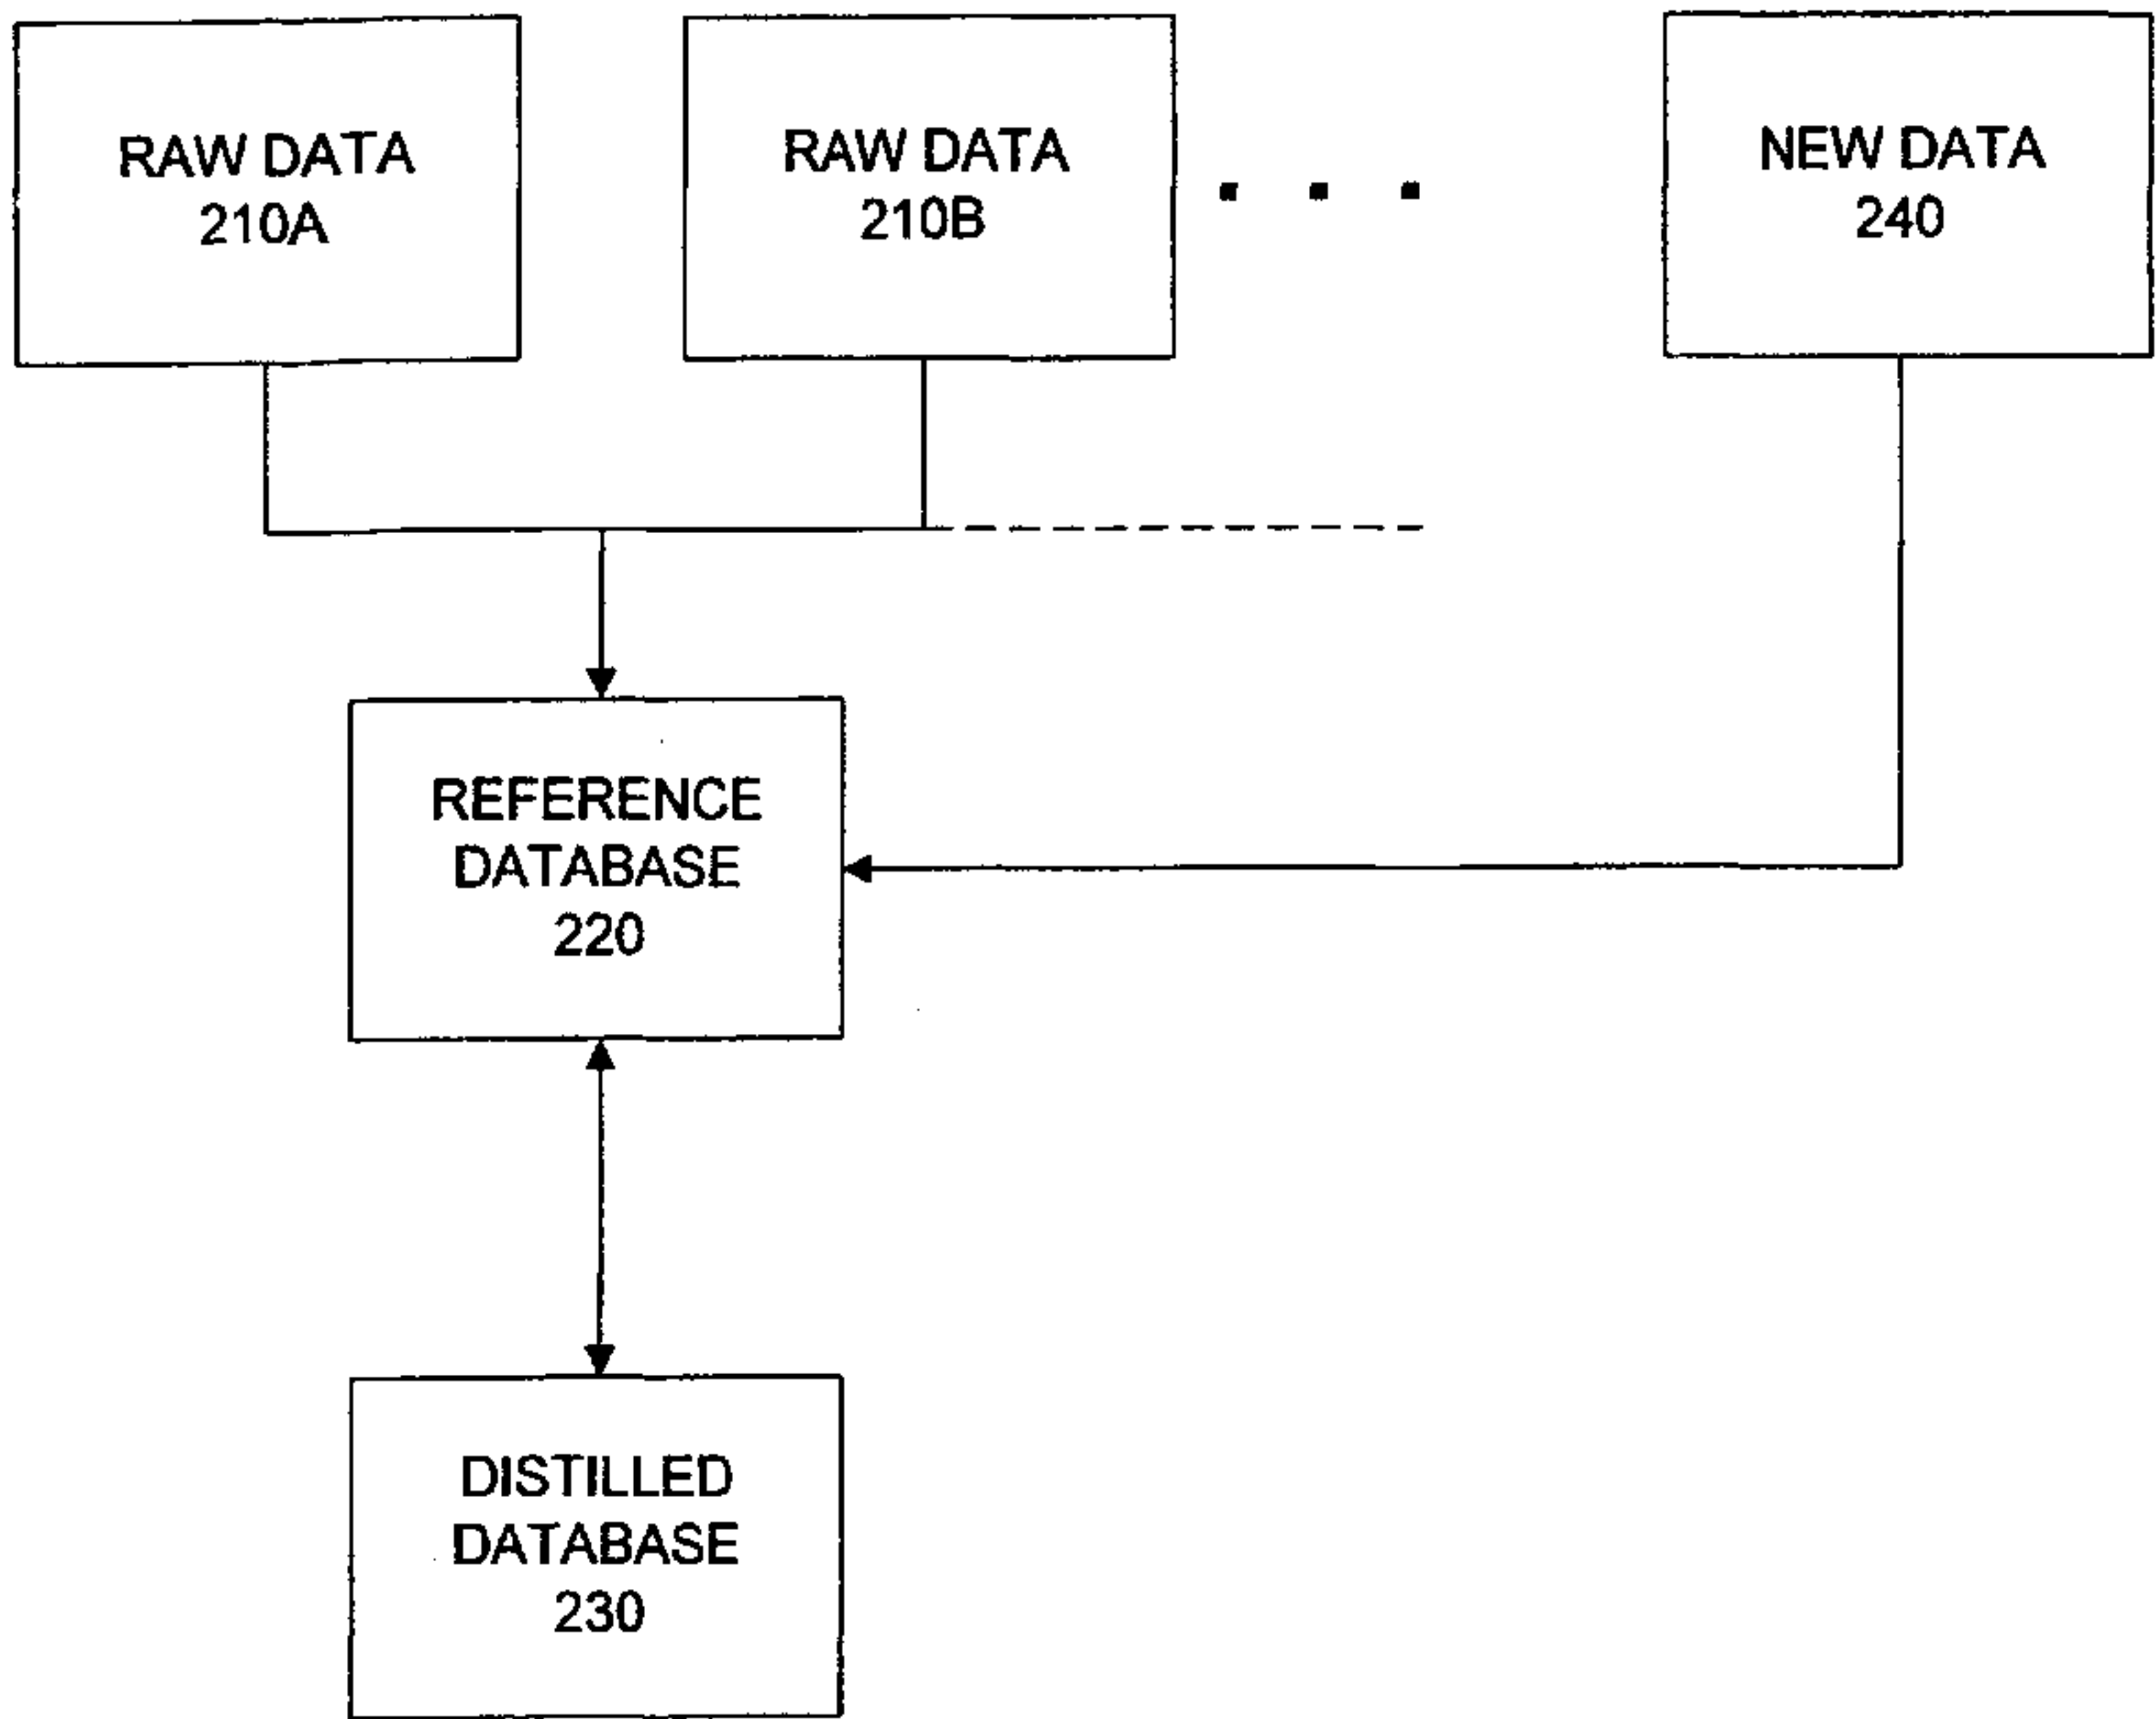
(74) Agents: BARAN, Alexandra, J. et al.; Cooley Godward LLP, One Freedom Square-Reston Town Center, Suite 1700, 11951 Freedom Drive, Reston, VA 20190-5601 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR ORGANIZING DATA



(57) Abstract: A system and method for organizing raw data from one or more sources uses an improved mechanism for identifying duplicate data between fields e.g., columns) in the databases. The fields may be similar fields within a single database or similar or identical field within a pair of databases and as organized as arrays or field vectors. The present invention sorts each of the field vectors and if necessary, partitions them by common value. A number of comparisons required to identify the duplicate data between the field vectors is reduced by feeding back a difference between the compared values. This difference is used to adjust indices into the field vectors for subsequent comparison.

WO 01/06414 A2

WO 01/06414 A2



IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

(15) Information about Correction:

see PCT Gazette No. 14/2001 of 5 April 2001, Section II

Published:

— *Without international search report and to be republished upon receipt of that report.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(48) Date of publication of this corrected version:

5 April 2001

Method and System for Organizing Data

The present invention relates to database systems and more particularly, to a system and method for organizing and/or finding data in a database system.

5 Computerized database systems have long been used and their basic concepts are well known. A good introduction to database systems may be found in C. J. DATE, INTRODUCTION TO DATABASE SYSTEMS (Addison Wesley, 6th ed. 1994).

In general, database systems are designed to organize, store and retrieve data in such a way that the data in the database is useful. For example, the data, or partitioned
10 sets of the data, may be searched, sorted, organized and/or combined with other data. To a large extent, the usefulness of a particular database system, is dependent on the integrity (i.e., the accuracy and/or correctness) of the data in the database system. Data integrity is affected by the degree of "disorder" in the data stored. Disorder may occur in the form of erroneous or incomplete data such as duplicate data, fragmented data,
15 false data, etc. In many database systems, from time to time, existing data may be edited and processed, and as a result, additional errors may be introduced. In some database systems, new data may be introduced. Additionally, as database systems are upgraded with new hardware and/or software, data conversion may be required or additional fields may become necessary. Furthermore, in some applications, the data in the
20 database may simply become outdated over time.

Regardless of the preventative steps taken, some degree of disorder is eventually introduced in conventional database systems. This degree of disorder increases exponentially over time until eventually, the data in a conventional database becomes entirely useless. As a result, even a small degree of disorder eventually affects the integrity of the database system.
25

Unfortunately, identifying and correcting disorder in the data are often difficult, if not impossible, tasks particularly in large database systems. Traditionally, such tasks are performed manually, making these tasks time-consuming, expensive, and subject to human error. Furthermore, due to the very nature of the task, much of the disorder may
30 go largely undetected. What is needed is a system and method for organizing data in a database system to overcome these and other associated problems.

The present invention provides a system and method for organizing data in a database system. The present invention derives a distilled database of accurate data from raw data extracted from one or more raw data sources. The raw data is converted from its original format(s) to a numeric format.

5 According to one embodiment of the present invention, the raw data is represented as a vector having numeric elements. Once the raw data is represented numerically, various mathematical operations such as correlation functions, pattern recognition methods, or other similar numeric methods, may be performed on these vectors to determine how content in a particular vector corresponds to others vectors in a “dis-
10 tilled” or reference database. The distilled database is formed from sets of one or more related vectors that are believed to be unique (e.g., orthogonal) with respect to the other sets. These sets represent the best information available from the raw data. After all the raw data has been incorporated into the distilled database, new data may be screened to ensure that new errors are not introduced into the distilled database. The new data may
15 be also evaluated to determine whether it is unique or whether it includes better information than that already present in the distilled database. The new data is added to the distilled database accordingly.

 According to one embodiment of the present invention, raw data are converted into a numeric format based on a number system having an appropriate radix. An appropriate radix is determined according to the type of information included in the raw
20 data. For example, for raw data generally comprised of alpha-numeric characters, an appropriate radix may be greater than or equal to the number of different alpha-numeric characters present in the raw data. Using such a number system allows raw data to be represented numerically, allowing for manipulation through various well-known
25 mathematical operations.

 According to one embodiment of the present invention, the number system may be selected so that the numbers themselves retain semantic significance to the raw data they represent. In other words, the numerals in the number system are selected so that they correspond to the raw data. For example, in the case of raw data comprised of al-
30 phanumeric characters, the numerals are selected to correspond to the alphanumeric characters they represent. When the numerals in the number system are subsequently displayed, they appear as the alphanumeric characters they represent.

According to one embodiment of the present invention, once the raw data is represented as vectors in an appropriate number system, the represented data may be efficiently manipulated in the database (e.g., sorted, etc.) using various well-known techniques. Furthermore, various well-known mathematical operations may be performed on the vectors to analyze the data content. These mathematical operations may include correlation functions, eigenvector analyses, pattern recognition methods, and others as would be apparent.

According to one embodiment of the present invention, the raw data is incorporated into a distilled database. The distilled database represents the best information extracted from the raw data without having any data disorder.

According to one embodiment of the present invention, new data may be compared to the distilled database to determine whether the new data actually includes any new information or content not already present in the distilled database. Any new information not already in the distilled database is added to the distilled database without adding any disorder. In this manner, the integrity of the distilled database may be maintained.

According to the invention a method for processing information comprises the steps of selecting an appropriate number system based on a range of possible values of a data element included in the information, representing said data element as a digit in the number system; and operating on said data element represented in the number system to process the information.

According to one embodiment of the present invention, the step of selecting an appropriate number system comprises the step of selecting a number system with a radix at least equal to and approximately the same as to an order of the alphanumeric characters “0”-“9” and “A”-“Z”.

According to one embodiment of the present invention, the step of selecting an appropriate number system comprises the step of selecting a number system with a radix greater than an order of the alphanumeric characters “0”-“9” and “A”-“Z”.

According to one embodiment of the present invention, the step of selecting an appropriate number system comprises the step of selecting a number system with a radix at least equal to an order of the alphanumeric characters “0”-“9”, “A”-“Z”, and “a”-“z”.

According to one embodiment of the present invention, the step of selecting an appropriate number system comprises the step of selecting a base 40 number system.

According to one embodiment of the present invention, the information includes financial information, scientific information, industrial information or chemical information.

The method of claim 16, wherein the step of assigning the digits further comprises assigning the digits A-Z in the number system to the alphanumeric characters "a"- "z", respectively.

According to one embodiment of the present invention, said step of comparing said vector with a distilled matrix comprises performing an eigenvector analysis or performing a pattern recognition analysis or determining a dot product between said vector and a vector in said distilled matrix or determining a cross product between said vector and a vector in said distilled matrix or determining a difference between said vector and a vector in said distilled matrix or determining a sum of said vector and a vector in said distilled matrix or determining a determinant of said distilled matrix or determining a magnitude of said vector or determining a direction of said vector.

The invention is, in general, characterized as stated in the independent claims, whereas the dependent claims contain preferred embodiments of the invention.

Preferred embodiments of the invention are described with reference following drawings. In the drawings, like reference numbers indicate identical or functionally similar elements. Additionally, the left-most digit(s) of a reference number identifies the drawing in which the reference number first appears.

FIG. 1 illustrates a processing system in which the present invention may be implemented.

FIG. 2 illustrates stages of data processed by one embodiment of the present invention.

FIG. 3 is a flow diagram for converting raw data from its original format into a numeric format in accordance with one embodiment of the present invention.

FIG. 4 illustrates a data record suitable for use with the present invention.

FIG. 5 illustrates raw data tables suitable for use with the present invention.

FIG. 6 illustrates reference data tables, representing data formatted in accordance with an embodiment of the present invention.

FIG. 7 is a flow diagram for analyzing reference data in accordance with an embodiment of the present invention.

FIG. 8 illustrates distilled data table, representing related data correlated in accordance with an embodiment of the present invention.

5 FIG. 9 illustrates an example of data clustering in a two-dimensional space.

FIG. 10 is a flow diagram for identifying duplicate data among a pair of field vectors.

FIG. 11 is a flow diagram for identifying duplicate data among a pair of field vectors in further detail.

10 FIG. 12 illustrates an example of identifying duplicate data among a pair of field vectors.

The present invention is directed to a system and method for organizing data in a database system. The present invention is described below with respect to various exemplary embodiments, particularly with respect to various database applications. However, various features of the present invention may be extended to other areas as would
15 be apparent. In general, the present invention may be applicable to many database applications where large amounts of potentially unrelated data must be compiled, stored, manipulated, and/or analyzed to determine the various relationships present in the content represented by the data. More particularly, the present invention provides a method for achieving and maintaining the integrity (i.e., accuracy and correctness) of data in a
20 database system, even when that data initially possesses a high degree of disorder. As used herein, disorder refers to data that is duplicative, erroneous, incomplete, imprecise, false or otherwise incorrect or redundant. Disorder may present itself in the database system in many ways as would be apparent.

One embodiment of the present invention is used to maintain a database associated with accounts receivable. In this embodiment, a company may collect data relating
25 to various persons, businesses and/or accounts from one or more sources. These sources may include, for example, credit card companies, financial institutions, banks, retail, and wholesale businesses and other such sources. While each of these sources may provide data relating to various accounts, each source may provide data representing different information based on its own needs. Furthermore, this data may be organized in
30 entirely different ways. For example, a wholesale distributor may have data corresponding to accounts receivable corresponding to business accounts. Such data may be

organized by account numbers, with each data record having data fields identifying an account number, a business associated with that account number, an address of that business, and an amount owed on the account. A retail company may have data records representing similar information but based on accounts corresponding to individuals as well as businesses.

In other embodiments of the present invention, other types of sources may provide different types of data. For example, the scientific institutions may provide scientific data with respect to various areas of research. Industrial companies may provide industrial data with respect to raw materials, manufacturing, production, and/or supply. Courts or other types of legal institutions may provide legal data with respect to legal status, judgments, bankruptcy, and/or liens. As would be apparent, the present invention may use data from a wide variety of sources.

In another embodiment of the present invention, a database may be maintained to implement an integrated billing and order control system. In addition to billing-type information from sources similar to those described above, this embodiment may include data records corresponding to inventory, data records corresponding to suppliers of the inventory, and data records corresponding to purchasers of the inventory. Inventory data may be organized by part numbers, with each data record having data fields identifying an internal part number, an external part number (i.e., supplier part number), a quantity on hand, a quantity expected to ship, a quantity expected to be received, a wholesale price, and a retail price. Supplier data may be organized by a supplier number; and customer data may be organized by a customer number. Data records corresponding to each of these records may include data fields identifying a part number, a part price, a quantity ordered, a ship data, and other such information.

Another embodiment of the present invention may include an enterprise storage system that consolidates corporate information from multiple, dissimilar sources and makes that information available to users on the corporate network regardless of the type of the data, the type of computer that generated the data, or the type of computer that requested the data. Still another embodiment of the present invention includes a business intelligence system that warehouses and markets information and allows that information to be processed and analyzed on-line.

The present invention enables raw data collected from different sources to be analyzed and distilled into a collection of accurate data, organized in a way that is useful for a particular application. Using the above example of an integrated billing and order control system, explained more fully below, the present invention may produce a
5 distilled database in which related data, such as data relating to a particular supplier or customer, may be identified as such. In this example, duplicate data corresponding to the same supplier or customer may be identified and/or discarded, and erroneous data associated with the supplier or customer may be identified, analyzed, and possibly corrected.

10 In general, the present invention may be implemented in hardware or software, or a combination of both. Preferably, the present invention is implemented as a software program executing in a programmable processing system including a processor, a data storage system, and input and output devices. An example of such a system 100 is illustrated in FIG. 1. System 100 may include a processor 110, a memory 120, a storage
15 device 130, and an I/O controller 140, coupled to one another by a processor bus 150. I/O controller 140 is also coupled via an I/O bus 160 to various input and output devices, such as a keyboard 170, a mouse 180, and a display 190. Other components may be included in the system 100 as would be apparent.

FIG. 2 illustrates various forms of data processed by the present invention. Raw
20 data 210 may be collected from one or more sources, such as raw data 210A and raw data 210B. As used herein, "raw data" simply refers to data as it is received from a particular source. Additional sources of raw data 210 may be included as would be apparent. As explained below, raw data 210 from various sources is preferably converted into a numeric format and stored in a reference database 220. Using a process referred
25 to herein as "data dialysis," the present invention "purifies" raw data 210 to form reference data in reference database 220. Reference database 220 includes all the information found in raw data 210 including duplicate, incomplete, inconsistent, and erroneous data.

Distilled data stored in a distilled database 230 is derived from the reference
30 data of reference database 220. Distilled data represents the "accurate" data available from raw data 210. Distilled database 230 includes the unique information found in raw

data 210. Distilled data thus represents the best information available from raw data 210.

As also explained below, the present invention further provides for using distilled database 230 to analyze and verify new data 240, which may also be used to update the reference database 220 and distilled database 230 as appropriate.

While the present invention has numerous embodiments, to clarify its description, a preferred embodiment is explained with reference to FIGS. 3-8 in a context of an integrated billing and order control system. In this embodiment, raw data 210 is a collection of data collected from various sources, such as order processing, shipping, receiving, accounts payable and accounts receivable, etc. This raw data 210 may include data records that are related but have different data fields, duplicate data records, data records having one or more erroneous data fields, etc. To address such errors, the present invention converts raw data 210 from their original formats and data structures (which may vary based on the source) into a numeric format and stores this reference data in reference database 220.

According to the present invention, the reference data is then compared and analyzed to distill the best information available. In one embodiment of the present invention, this best information may be stored as distilled data in distilled database 230. This process is now described.

20

Collecting Raw Data

FIG. 3 illustrates the process by which raw data 210 is converted into reference data in reference database 220 according to one embodiment of the present invention. In a step 310, raw data 210 is collected from a raw data source. As illustrated in FIG. 2, raw data 210 may include data from one or more sources such as raw data 210A and raw 210B. As used herein, "data" refers to the physical digital representation of information, and data "content" refers to the meaning of, or information included in or represented by that data. The different records in raw data 210 may include similar types of data content. For example, in a billing context, different records in raw data 210 may all include data content relating to a particular account.

30

Raw data 210 will typically be received in the form of data records 400, as illustrated in FIG. 4. Each data record 400 generally includes related information, such as information for a specific individual, company, or account. Each data record 400 stores this information in one or more data fields 410. Examples of possible data fields 410 include, for example, an account number, a last name, a first name, a company name, an account balance, etc. Each data field 410, in turn, may include one or more data elements 420 for representing information for that specific record and specific field. Data elements 420 may exist in various formats, such as alphanumeric, numeric, ASCII, and EBCDIC, or other representation as would be apparent. Raw data 210 collected from different sources may be formatted differently. Data records 400 may include different data fields 410, and the information included in data fields 410 may be represented using data elements 420 in different formats, as would also be apparent.

Examples of raw data 210 are illustrated in raw data tables 510, 520, and 530 of FIG. 5. Data records, such as data record 510-1 and data record 510-2, are illustrated as rows of raw data tables 510, 520, and 530, whereas data fields, such as data field 510-A and data field 510-B, are illustrated as columns of raw data tables 510, 520, and 530. Either data fields or data records can be thought of as ordinary mathematical vectors or tensors and manipulated accordingly. The tables illustrated in FIG. 5 are examples of data that might be found in various embodiments of the present invention. In other embodiments, data may come from many sources and may be formatted as databases having a much larger number of data records and/or data fields, as would be apparent.

Conversion to Numeric Format

Referring to FIG. 3, in a step 320, the present invention converts raw data 210 from its original representation (which may be in alphanumeric, numeric, ASCII, EBCDIC, or other similar formats) to a numeric representation. This ensures that reference data is represented in the same manner. Thus, the reference data, including that data from different sources, may be similarly processed.

According to the present invention, raw data 210 is converted from its original representation into an appropriate numeric representation. An appropriate numeric representation uses a number system in which each possible value of data element 420

may be represented by a unique digit or value in the number system. In other words, a radix for the number system is selected such that the radix is at least as great as the number of possible values for a particular data element. For example, in a biotechnology application for detecting nucleotide sequences of Adenine (A), Guanine (G), Cytosine (C), and Thymine (T) in nucleic acids, each data element may be one of only four values: A, G, C, and T. In such an application, a radix of four for the number system may be sufficient to represent each data element as a unique number. One such number system may include the numbers A, G, C, and T. In some embodiments of the present invention, it may be desirable to use a radix at least one greater than the number of different possible value of data element 420 in order to provide a number representative of an empty field. In this case, such a number system may include the numbers A, G, C, T, and ^, where ^ is the empty field value.

According to a preferred embodiment of the present invention, data elements 420 in raw data 210 are comprised of characters such as alphanumeric characters. In this preferred embodiment, a radix of 40 is selected to represent the alphanumeric characters as illustrated in the table below. (Note that a minimum radix of 36 is required.) This radix is selected to accommodate the ten numeric characters "0"- "9" and the twenty-six alphabetic characters "A" to "Z" as well as to allow for several additional characters. In this embodiment, uppercase and lowercase characters are not distinguished from one another.

As illustrated in Table 1, the base-40 number system includes the numbers 0-9, followed by A-Z, further followed by four additional numbers. One of these numbers may be used to represent an empty field. This number is used to represent a data field 410 that is empty or has no value (in contrast to a zero value). Other numbers may be used, for example, to represent other types of information such as spaces or used as control information.

Alpha-Numeric	Base-10 Number	Base-40 Number	Alpha-Numeric	Base 10 Number	Base-40 Number
0	0	0	K or k	20	K
1	1	1	L or l	21	L

2	2	2	M or m	22	M
3	3	3	N or n	23	N
4	4	4	O or o	24	O
5	5	5	P or p	25	P
6	6	6	Q or q	26	Q
7	7	7	R or r	27	R
8	8	8	S or s	28	S
9	9	9	T or t	29	T
A or a	10	A	U or u	30	U
B or b	11	B	V or v	31	V
C or c	12	C	W or w	32	W
D or d	13	D	X or x	33	X
E or e	14	E	Y or y	34	Y
F or f	15	F	Z or z	35	Z
G or g	16	G	--	36	[
H or h	17	H	--	37	\
I or i	18	I	--	38]
J or j	19	J	--	39	^

Table 1

Representation of raw data 210 in a base-40 format has numerous benefits. One benefit is that raw data 210 may be represented in a numeric fashion, facilitating straightforward mathematical manipulation. Another benefit is that proper selection of both the radix and the numerals in the number system allows the represented content to maintain semantic significance, facilitating recognition the content of raw data 210 in its representation in the numeric format. For example, the word "JOHN" represented by the four alphanumeric characters "J" "O" "H" "N" may be represented in various number systems. One such number system is a base-40 number system. Using Table 1, representing the alphanumeric characters "JOHN" as a base-40 number would result in the "tetradecimal" value 'JOHN', which is equivalent to the decimal value 1,255,103

($19 \cdot 40^3 + 24 \cdot 40^2 + 17 \cdot 40^1 + 23 \cdot 40^0$, where base-40 'J' equals decimal 19, etc.). Note that the base-10 number loses semantic significance from the content of raw data 210 whereas the base-40 number retains semantic significance, as the number 'JOHN' is recognizable as the content "JOHN." Semantic significance provides the benefits of a
5 numeric representation while maintaining the ability to convey semantic content.

In some embodiments of the present invention, the selection of a radix and its corresponding number system may depend upon the number of bits used by processor 110. The number of bits used by processor 110 and the radix chosen for the number system define the number characters that can be represented by a data word in proces-
10 sor 110. This relationship is governed according to the following equation:

$$N = B * \ln(2) / \ln(R),$$

where N is the number of whole characters represented by a data word of processor
15 110, B is the number of bits per data word, and R is the selected radix. This relationship limits the number of data elements 420 of raw data 210 that may fit in a data word. For example, in a 32-bit machine, the maximum number of characters that may fit in a data word using a base-40 number system is six ($32 \cdot \ln(2) / \ln(40) = 6.013$). The maximum number of characters that may fit in a data word using a base-41 number system is only
20 five ($32 \cdot \ln(2) / \ln(41) = 5.973$). Thus, in some embodiments of the present invention, in addition to having a radix sufficiently large to maintain semantic significance, the radix may also be selected to maximize the number of characters represented by a single data word and/or to facilitate rapid mathematical operations based on advantages or specific designs of various processors. In the embodiment with raw data comprised of alphanu-
25 meric characters, an appropriate radix may range from 36 to 40. This range maintains semantic significance while maximizing the number of characters represented by the 32-bit data word. Other types of raw data and other sizes of data word may dictate other appropriate radix ranges in other embodiments of the present invention.

The embodiment of the present invention described above does not distinguish
30 between uppercase and lowercase characters. However, other embodiments of the present invention may distinguish between these types of characters. Accordingly, a base-

64 representation (“0”-“9”, “A”-“Z”, “a”-“z”, and two other values) may be appropriate to distinguish between these characters as would be apparent.

The number of data elements 420 in each data field 410 also dictates the precision required by the number as represented in processor 110. As described above, each data field 410 may only be six characters or data elements 420 wide for single precision operations in a 32-bit machine. In some embodiments of the present invention, this may be insufficient. In these embodiments, double, triple, or even quadruple precision may be required to represent the entire data field 410 as a single value. Double precision numbers are sufficient for up to twelve character data fields 410; triple precision numbers are sufficient for up to eighteen characters; and quadruple precision numbers are sufficient for up to twenty-four characters.

Alternate embodiments of the present invention may accommodate large data fields by breaking a large data field into one or more smaller data fields. The large data fields may be broken at natural boundaries such as those defined by spaces. For example, a data field representing an address such as “123 West Main Street” may be broken into four smaller data fields: ‘123’, ‘West’, ‘Main’, and ‘Street’. The large data fields may also be broken at data word boundaries. In the address example above, the smaller data fields might be: ‘123We’, ‘st\Mai’, ‘n\Stre’, and ‘et’, where the number ‘\’ is used to represent a space. Other embodiments of the present invention may accommodate large data fields in other manners as would be apparent.

Data Structure Conversion

As illustrated in FIG. 3, in a step 330, raw data 210 represented as a number is stored in a predefined data structure. In one embodiment of the present invention, this data structure is a single-field table as illustrated by Tables 610-670 of FIG. 6. This data structure may vary. For example, in other embodiments of the present invention, the data structure may be a multiple-field table instead of a single-field table. In these embodiments, the data structures may be implemented with standard features such as table headers and indices, and as explained in greater detail below, may also include probability values for each record. These probability values represent the likelihood that the data in that record is complete. Higher probability values may indicate a higher

probability of completeness, and lower probability values similarly may indicate a lower probability of completeness. This is described in further detail below. Initially, the probability values are set to 0. Other embodiments may also include key numbers or identification numbers to aid in sorting and in maintaining relationships among the data records.

In a preferred embodiment of the present invention, raw data 210 illustrated in FIG. 5 includes three tables 510, 520, and 530. Table 510 may represent raw data 210 from, for example, a company's accounts receivable system. Columns of table 510 represent data fields for an account number, a last name, a first initial, and additional fields for listing various orders processed for a particular individual. Rows of table 510 (such as 510-1 and 510-2) represent data records for different individuals. Tables 520 and 530 may represent raw data 210 maintained by credit card companies. Columns of tables 520 and 530 represent data fields for an account number, a last name, a first name, and an address. Rows of tables 520 and 530 represent data records for specific accounts.

In the preferred embodiment, step 330 converts raw data 210 from the format illustrated in FIG. 5 into a format illustrated in FIG. 6. FIG. 6 illustrates raw data 210, combined from the various raw data tables 510, 520, 530 of FIG. 5, represented as numbers in a base-40 number system, and formatted as new tables (tables 610-670), which together may comprise reference database 220.

Each reference database table 610-670 corresponds to an individual field from raw data tables 510, 520, and 530 of FIG. 5. More specifically, data records of reference data tables 610-670 correspond to the data records of raw data table 510, followed by the data records of raw data table 520, followed by the data records of raw data table 530. In one embodiment of the present invention, where a raw data table record has no information for a particular data field 410 represented in a reference table 610-670, an empty field value is entered in that field in the reference table. For example, the first data record 510-1 of Table 510 has no information about an address, and thus an empty field value is placed in the first position of table 670.

Data is preferably stored in reference database 220 in such a way that all data corresponding to a single data record in a raw data table is readily identified. In the embodiment represented in FIGS. 5 and 6, for example, data corresponding to any specific data record of the raw data tables (tables 510, 520, 530) is preferably represented in

reference tables 610-670 as a “vector” of numeric data stored at an index *i* across reference tables 610-670. For example, data corresponding to the sixth record 520-6 of raw data table 520 (illustrated as account number “A60” belonging to “Jennifer Brown,” residing at “51 Fourth Street”) is represented in reference database tables 610-670 as a vector having coefficients formed from the tenth records 610-10, 620-10, 630-10, 640-10, 650-10, 660-10, and 670-10 of the tables 610-670.

As illustrated in FIG. 6, reference database 220 includes a new table 610 that does not correspond to any data field 410 in raw data 210 illustrated in FIG. 5. This table is a “key table” that identifies the related data in these data vectors. As described below, reference database 220 comprised of the tables illustrated in FIG. 6 may include additional key tables for data fields. These may include a personal identification number (“PIDN”), an account identification number (“AIDN”), or other types of identification numbers. These key tables or identification numbers may be used to identify sets of related data vectors in reference database 220.

In this example, key table 610 has a single field “PIDN,” which stands for personal identification number. Key table 610 provides a unique identifier such that a specific PIDN number never refers to more than one person represented in raw data 210. In other words, the PIDN number reflects the fact that many multiple records in raw data 210 may refer to the same person.

Preferably, each data record in the key table 610 initially corresponds to a different data record represented in the raw data tables 510, 520, and 530. For example, in FIG. 6, data record 610-10 in the key table 610 is implemented such that it includes identifiers (such as pointers or indices) for corresponding data in reference tables 620-670, which together corresponds to a single record 520-6 in raw data table 520.

Initially, while a single PIDN does not refer to multiple individuals, a single individual may correspond to multiple PIDNs. For example, in FIG. 6, vector 4 (defined by PIDN 4) and vector 9 (defined by PIDN 9) appear to refer to the same person, but as illustrated, this person is initially assigned to two PIDN numbers -- PIDN 4 and PIDN 9. As described below, the present invention enables a determination whether PIDN 4 and PIDN 9 do, in fact, refer to the same individual, and if so, assigns a single PIDN to this individual. Alternatively, some embodiments may assign a new PIDN number to individuals so determined and a reference to the old PIDN number may be retained.

As discussed above, in this embodiment, records are represented in the reference database tables 610-670 as vectors having coefficients of base-40 numbers across eight one-field tables. This numeric representation allows the data to be analyzed using straightforward mathematical operations that may be used to, for example, produce correlations, calculate eigenvectors, perform various coordinate transformations, and utilize various pattern recognition analyses. These operations may, in turn, be used to provide or derive information about the records and their relationships to one another. By using small, one-field tables, these operations may be performed quickly. In addition, as will be illustrated, representation in base-40 numbers with raw data 210 including alphanumeric characters allows content of raw data 210 to retain its semantic significance.

Data Dialysis

Referring back to FIG. 2, once reference database 220 is created as illustrated in FIG. 6, a data dialysis process 700 is applied to distill the most accurate data for inclusion in distilled database 230. Data dialysis 700 is now described with reference to FIG. 7.

Partitioning the Reference Data

In a step 710, reference database 220 is preferably partitioned or sorted into sets based on some criteria. These sorting criteria may vary. For example, as illustrated in table 810 of FIG. 8, in this embodiment, data records may be sorted into sets based on last name, with the values arranged in increasing numeric order (recall that content of raw data is now represented as base-40 numbers in reference database 220). Table 810 is derived from reference database table 620 illustrated in FIG. 6, with each entry of table 810 defined by a unique last name and having a corresponding set of table 620 records matching that last name. In the representation illustrated, table 810 includes a field for defining the set (in this case, a last name), as well as identifiers for members of the set (such as indices, pointers or other appropriated references -- in this case PIDNs).

In some embodiments of the present invention, not all vectors in reference database 220 will have data for the field on which the sets are based. Such vectors may be handled in various manners. For example, all vectors in reference database 220 having no data for that data field may be regarded as members of a single, additional set. Alternatively, each vector in reference database 220 having no data for that data field may be regarded as the single member of its own set.

Identifying Duplicate Data

Returning to FIG. 7, in a step 720, those data records within the partitioned sets identified as duplicates are marked. In some embodiments of the present invention, duplicate data may be unnecessary and may be discarded. In other embodiments, all information remains in reference database 220 as all information, even erroneous, incomplete, or duplicate information may be better than no information and may be useful for some purpose, such as identifying fraud or theft.

In some embodiments of the present invention, comparing a pair of vectors may identify duplicates. Various operations may be used, as would be apparent. In a simple example, a straightforward vector subtraction may be performed to measure the degree of similarity between two records. Other techniques may be used to identify duplicate vectors such as using “look-up” tables to identify common names, nicknames, abbreviations, etc.

Table 810 of FIG. 8 illustrates that the last name “Smith” corresponds to PIDNs 2, 4, 8, 9, and 11, representing vectors formed from entries 2, 4, 8, 9, and 11 of the reference database tables 610-670 illustrated in FIG. 6:

For PIDN 2: [SMITH, J , 98-002, A40, A60, ^]
 For PIDN 4: [SMITH, J , 98-004, A50, B10, ^]
 For PIDN 8: [SMITH, Jennifer, ^ , A40, ^ , 300 Pine St.]
 For PIDN 9: [SMITH, John, ^ , A50, ^ , 37 Hunt Dr.]
 For PIDN 11: [SMITH, Jhon, ^ , B10, ^ , 85 Belmont Ave.]

Vector (or matrix) operations comparing the vectors and thresholds for determining when two entries are similar enough to be regarded as duplicates may be de-

defined as appropriate for various embodiments. In a simple example, the sum of the absolute differences between corresponding coefficients of a pair of vectors may indicate a similarity between the corresponding pair of records. This pair of vectors may be considered duplicates if a first vector is not inconsistent with any field of a second vector, and does not provide any additional data. In this embodiment, additional rules would also be defined, for example, for comparing entries of different lengths (e.g., right aligning character strings corresponding to numbers, and left aligning character strings corresponding to letters), for recognizing commonly misspelled or spelling variations of words, and for recognizing transposed letters in words. This processing may be performed by various mechanisms, as would be apparent. In the example of Table 810 of FIG. 8, none of the data records are exact duplicates, and so none are marked in step 720.

Correlating Data

15

Referring back to FIG. 7, in a step 730, the preferred embodiment of the present invention correlates data records remaining within each set and in a step 740, further partitions the data records into independent subsets of data records. In general, the “correlation” between two vectors is a measurement of how closely one is related to the other, and specific methods of correlation may vary depending on the intended application. A general discussion and examples of correlation functions may be found in references such as NUMERICAL RECIPES IN C: THE ART OF SCIENTIFIC COMPUTING (Cambridge University Press, 2nd ed. 1992) by William H. Press, et al. Other techniques and examples may be found in THE ART OF COMPUTER PROGRAMMING (Addison-Wesley Pub., 1998) by Donald E. Knuth.

As an example, a simple measurement of the correlation between vectors is their dot product, which may be weighted as appropriate. Depending on the application, the dot product may be calculated on only a subset of the vector coefficients, or may be defined to compare not only corresponding coefficients, but also other pairs of coefficients determined to be in related fields (i.e., comparing a “first name” coefficient of a first vector with a “middle name” coefficient of a second vector). As with the operations for identifying duplicate data, the correlation function may be appropriately tai-

30

lored for its intended application. For example, a correlation function may be defined to appropriately compare entries of different lengths and to appropriately distinguish between significant and insignificant differences, as would be apparent.

5 In the embodiment explained with reference to the tables of FIGs. 5, 6, and 8, an example of a correlation function compares vectors corresponding to the members of a set sharing the same last name to identify independent subsets of vectors. Again, this determination may be based on application-specific criteria. In this example, independent vectors may be defined to be those vectors representing different individuals.

10 As a result of applying the correlation function, a correlation parameter reflecting the degree of independence of a pair of vectors is assigned. For example, a high value may be assigned to indicate a high degree of similarity, and a low value may be assigned to indicate a limited degree of similarity. The correlation value is then compared to a predetermined threshold value -- which again, may vary in different applications -- to determine whether the two records corresponding to those vectors are con-
15 sidered to be independent.

Based on the correlation values, in a step 740, the preferred embodiment partitions the data records into subsets of independent data records within each set. In the examples of FIG. 5, 6, and Table 810 of FIG. 8, members of an independent subset may be identified as those members having: the same last name (taking into consideration
20 misspellings and spelling variations); relatively similar first names (taking into consideration misspellings, spelling variations, nicknames, and combinations of first and middle names and initials); having one or more matching account numbers; and having no more than three addresses (to allow for work and home addresses, and one change of address).

25 Results of applying such a function are illustrated in Table 820 of FIG. 8. The individuals identified are:

Jennifer Brown, PIDN 10;
Howard Lee, PIDNs 3 and 6;
Carole Lee, PIDN 7;
30 Jennifer Smith, PIDNs 2 and 8;
John Smith, PIDNs 4 and 11;
John Smith, PIDN 9;

Ann Zane, PIDNs 1, 5, and 12; and
Molly Zane, PIDN 13.

Other operations for correlating the vectors are available. These may include computing dot products, cross products, lengths, direction vectors, and a plethora of
5 other functions and algorithms used for evaluation according to well-known techniques.

FIG. 9 illustrates a two-dimensional example of a concept referred to as clustering which is used conceptually to describe some general aspects of the present invention. In FIG. 9, four clusters exist as a collection of two-dimensional points. These clusters are identified as: (a,b), (c,d), (e,f), and (g,h). As illustrated, each cluster is
10 formed from one or more points in the two-dimensional space. Each point corresponds to a data record that represents (with more or less accuracy) the “true” value of the cluster in the space. As illustrated, clusters (a,b,) and (c,d) are fairly easy to distinguish from one another and from clusters (e,f) and (g,h). However, in this simple example, clusters (e,f) and (g,h) are not easily distinguished from one another. Extending the
15 space (i.e., adding additional data fields to the vectors), may increase the separation between clusters such as (e,f) and (g,h) so that they become more readily distinguished from one another. Alternately, extending the space may indicate that (g,h) is a point that belongs to cluster (e,f) or even cluster (c,d). In the abstract, the space may be extended infinitely, resulting in a Hilbert space, which has various well-known characteristics. These characteristics may be exploited by the present invention for large, albeit
20 not infinite, vectors as would be apparent.

Furthermore, while adding additional data fields to the vectors (i.e., extending the space) may separate clusters from one another to aid in their correlation, deleting data fields from the vectors (i.e., reducing the space) may also identify some correlations. In some embodiments of the present invention, reducing the space may identify
25 certain clusters that are in fact representing the same individual or other unique entity. For example, one record in a database may have ten data fields exactly identical to the same ten data fields in a second record in the database. These data fields may correspond to a first name, a birth date, an address, a mother’s maiden name, etc. However,
30 these two records may have two fields that are different. These two fields may correspond to a last name and a social security number. In some cases, these records may correspond to the same individual. The present invention simplifies the process for

identifying these types of records that would be difficult, if not impossible, to detect using conventional methods.

Thus, removing one or more particular data fields from a vector and reducing the corresponding space may reveal clusters that otherwise would not be apparent. Doing this for data fields traditionally used for identification purposes (e.g., last name, social security number, etc.) may reveal duplicate records in databases. This may be particularly useful for identifying fraud. Removing data fields where a vector includes an empty field value for that data field may also reveal clusters that would not otherwise be apparent.

Furthermore, once the clusters are identified as representing the same individual or entity, the best information for the individual or entity may be extracted from the information provided by each record or "black dot."

The principles of the present invention may be extended beyond simple vectors and data fields. For example, the present invention may be extended through the use of tensors representing objects in a multi-dimensional space. In this manner, the present invention may be used to represent the parameters of various physical phenomenon to gain additional insight into their operation and effect. Such application may be particularly useful for deciphering the human gene and aid in the efforts of programs such as the Human Genome Project.

20

Handling Stranded Data

Referring again to FIG. 7, in a step 750, the preferred embodiment of the present invention evaluates "stranded" data records. Stranded data records are those records from reference database 220 that were not partitioned into any set in step 710. In some embodiments, reference database 220 may include a large number of tables corresponding to data fields and a large number of vectors having data for various combinations of fields. For example, in an embodiment having a reference database 220 including 20 tables for different data fields and 1000 vectors defined by related data records for each table, suppose only 800 of those 1000 vectors have data for the field "last name," by which the sets were created in step 710. Step 710 may not partition those 200 vectors with no "last name" data into any set, or to partition each of those 200

vectors into its own set. In either case, the result is that those 200 vectors are not correlated with any others in steps 720, 730, and 740. Step 750 may evaluate those vectors.

Methods of evaluation may vary. For example, one embodiment may correlate each stranded entry with one member of each subset identified in step 740. Depending on the resulting correlation values, that vector may be added to the subset with which it is most highly correlated, or may define a new subset. Alternatively, in some embodiments, it may be determined that such evaluation is too time-consuming and/or costly and step 750 may be completely skipped.

10 *Repeating the Correlation Process*

Steps 710-750 may be repeated as needed for specific embodiments. As noted above, some embodiments will have reference data 220 having a large number of fields and a large number of entries, with many entries having data for only a subset of fields. In such a case, performing steps 710-750 on a single field is unlikely to derive all relevant information. Even in the simple example explained with reference to FIGS. 5, 6, and 8, correlating on the single field "last name" may provide only partial information about the correlation between those entries. For example, Jennifer Smith, corresponding to PIDNs 2 and 8 in FIG. 6, may be the same individual as Jennifer Brown, corresponding to PIDN 10, because PIDNs 2 and 10 may share a common account number. Performing the correlation on the last name field may not identify these PIDNs as corresponding to the same individual because they were evaluated only against other PIDNs sharing the same last name. Performing a correlation on the account number field may provide additional information about whether these PIDNs are related.

25 Thus, correlation across various data fields may be necessary to fully evaluate the degree of relatedness of the data in reference database 220.

Using Correlation Results to Update Reference Data

30 Once steps 710-760 are completed, reference database 220 has been distilled into a distilled database 230, as illustrated in FIG. 2. In some embodiments of the present invention, these two databases are handled separately and coexist with one another.

In other embodiments of the present invention, a single database exists with records marked or otherwise identified as belonging to reference database 220 or distilled database 230. This may be accomplished by assigning by using different ranges of PIDNs for the records in the two databases. Furthermore, relationships between records in the two databases may be maintained by adding a constant value to the PIDN for the record in reference database 220 to generate a PIDN for the record in distilled database 230. For example, a record with a PIDN of 12345 in reference database 220 may have a PIDN of 9012345 in distilled database 230. In this manner, the two databases may be treated as distinct portions of a single database.

Using the Distilled Data

Once data dialysis process 700 is complete, distilled database 230 identifies subsets of data records from the reference database 220 as related records, and as noted above, probabilities may be determined for fields in the reference database 220 to provide a qualitative measure of their completeness. This may be accomplished by assigning a probability of completeness to each of the individual data fields and then using them to compute an overall probability of completeness for the data record. For example, for a data field representing a first name, a value of 'J' may be assigned a low probability (e.g., 0 or 0.1), a value of 'JOHN' may be assigned a higher probability (e.g., 0.7 or 0.8), and a value of 'JONATHAN' may be assigned the highest probability (e.g., 0.9 or 1.0). These values may be assigned somewhat arbitrarily or according to some hypothesis of structure. However, these values help identify which data fields in the set are most likely to include the most complete information or in other words, the most probable data.

Use of the present invention may determine a significant amount of information about the records and their relationship to each other, and may be specifically tailored for particular applications. Furthermore, using standard database operations, distilled database 230 (which references records of the reference database 220) may be manipulated to provide formatted reports as needed. For example, an embodiment may be tailored to generate a report listing subsets of related records, with records of a subset providing information about a specific individual or entity. The records within such a

subset may provide information, for example about different fields of information; aliases and/or variations of names, addresses, social security numbers, etc., used by the individual; and fields -- such as occupation, address, and account numbers -- for which that individual may have more than one entry.

5 Recalling that all data is represented in numerical base-40 format, the subsets may be ordered numerically in the report. The base-40 format provides the additional advantage of representing alphabetical characters as their respective letters (as illustrated in the conversion table above). Thus, while the report will show entries in numerical representation, that representation retains the semantic significance of the data
10 it represents, allowing the data to be manually read and analyzed. For example, if the report shows records for an individual having entries for names including J SMITH, JOHN SMITH, JOHN G SMITH, G SMITH, and GERALD SMITH, a person reading that report would understand that this individual uses various first names, including his first name or initial, his middle name or initial, or some combination thereof.

15

Adding New Data

As with conventional database applications, new data may be added from time to time. As illustrated in FIG. 2, the present invention accounts for adding new (or
20 changed) data 240, which will affect reference database 220 and distilled database 230.

Generally, new data records 240 may be formatted as described with reference to FIG. 3, and entered into the existing reference database 220. Additionally, new data records 240 may be measured against distilled database 230 to determine if new information or content is available in new data record 240. For example, a new data record
25 240 may be correlated with data records from distilled database 230 to determine whether that new data record 240 is related to any data records already present in distilled database 230. If so, and new data record 240 contains information or content not already present in distilled database 230, new data record 240 may be used to update distilled database 230. For example, if new data record 240 included information for an
30 individual named John Smith that corresponds to data records already present in distilled database 230 but provided the additional information that Mr. Smith's middle

name was Greg, that additional information may be appropriately added to distilled database 230.

Changes to data records in reference database 220 and distilled database 230 may be handled using standard database protection operations, as described in references such as C. J. DATE, INTRODUCTION TO DATABASE SYSTEMS (Addison Wesley, 6th ed. 1994) (see specifically, Part IV), referenced above. For example, in the case that changes are made to reference database 220 by an authorized database administrator, related data records in reference database 220 are updated as determined by standard relational definitions and where appropriate, in accordance with relations defined in distilled database 230.

Identifying Duplicate Data Between Field Vectors

One problem associated with conventional databases is a difficulty in merging records from a first database, such as raw data 210A, with those from a second database, such as raw data 210B. Records in these databases having shared or duplicate data need to be identified so that the content included therein may be merged as a single record in a database such as reference database 220 or distilled database 230. For example, both databases 210 may include one or more entries for JOHN SMITH. If the respective records in the databases 210 represent the same individual John Smith, then the content of each of the records should be merged as a single record in, for example, distilled database 230.

Conventional brute force methods for identifying such duplicate data in these databases involve comparing a data record from the first database with every data record in the second database, and repeating this process for each record in the first database. This process is time consuming, computationally intensive, and accordingly, costly. In fact, the number of computations is geometrically related to the number of records in each of the two databases.

One process for reducing the time and number of computations required to identify the duplicate data in the databases 210 is described below with reference to FIGS. 10-12. In the process described below, a particular field common or similar among the databases is selected, for example a name field or an address field. This field

is arranged as a table or an array for each of the databases that includes the value of the selected field for each of the records. For example, as discussed above, each table 610-670 represents a particular field of each of the data records in a database. For purposes of this discussion, these tables are referred to as field vectors.

5 According to the present invention, each of the field vectors are sorted in numerical order, and if necessary, partitioned into sets of identical data as described above with respect to FIGS. 7 and 8. For example, multiple records associated with JOHN SMITH would be partitioned together within the field vector. Preferably, information regarding the location of the partitions between the sets is stored.

10 Once the field vectors are sorted and partitioned, a value of the first element of a first field vector is compared with a value of the first element of a second field vector. Essentially, if the value in the first field vector is greater than the value in the second field vector, an index into the second field vector is advanced or otherwise adjusted to a position within the next partitioned set to obtain a next value in the second field vector.
15 This next value in the second field vector is then compared to the value in the first field vector. This continues as long as the value in the first field vector is greater than the value in the second field vector.

 On the other hand, if the value of the first field vector is less than the value of the second field vector, an index into the first field vector is advanced or otherwise adjusted to a position within the next partitioned set to obtain a next value in the first field
20 vector. This next value in the first field vector is then compared to the value in the second field vector. This continues as long as the value in the first field vector is less than the value in the second field vector.

 When the value of the first field vector equals the value in the second field
25 vector, the process has identified duplicate data that is then preferably stored in a common field vector. After storing the identified duplicate data, the index into the first field vector and the index into the second field vector are both advanced or otherwise adjusted to a position within the next partitioned set of their respective field vectors.

 The process thus described may be viewed as feedback control mechanism that
30 adjusts the index into either of the arrays based on the difference between the values in the field vectors. In the embodiment described above, a positive difference generates an adjustment to the index of the second field vector whereas a negative difference gener-

ates an adjustment to the index of the first field vector. This process results in a linear relationship between the number of values in the field vectors and the number of computations (*i.e.*, comparisons) required as opposed to the geometric relationship associated with conventional methods.

5 The present invention may be extended to sorting mechanisms as well. In cases where a particular value must be inserted into a field vector (*i.e.*, a record must be inserted into a database) based on an ordering of the values in the vector (*e.g.*, alphabetically, numerically, etc.), a difference between the particular value and a value of one of the elements in the vector is computed. This difference is “fed back” to adjust the index
10 into the vector to generate the next value from the vector. Using well-established methods of control theory, the index adjustments may be integrated to determine the proper location of the value to be inserted. In addition to the integrator, a proportional gain may be applied to the difference to establish a desired system performance as would be apparent.

15 The present invention is now described with reference to FIGS. 10-12. FIG. 10 is a flow diagram for identifying duplicate data within a pair of field vectors. The field vectors may be from a single source such as raw data 210A (*e.g.*, when comparing a Residential Address Field with a Mailing Address in a single database) or from multiple sources such as raw data 210A and raw data 210B (*e.g.*, when comparing a Name
20 Field between two databases).

 For purposes of this description, the pair of field vectors are referred to as a first field vector (“FV1”) and a second field vector (“FV2”), respectively. Preferably, the data in these field vectors are base-40 numbers that represent alphanumeric data as described above. However, in some embodiments of the present invention, the data may
25 exist in other forms as well.

 In a step 1010, the first field vector is sorted in numerical order. In a step 1020, the second field vector is also sorted in numerical order. In one embodiment of the present invention, the vectors are sorted in increasing numerical order, although other embodiments of the present invention may sort the vectors in decreasing order as would be
30 apparent.

 In a step 1030, partitioned sets within the first field vector having common values are identified. Likewise, in a step 1040, partitioned sets within the second field

vector having common values are also identified. Steps 1010-1040 perform a similar function to the step of partitioning reference database 220 described above with reference to FIGS. 7 and 8. In some embodiments of the present invention, the field vectors may not include any partitioned sets as the common values within each field vector may have been eliminated. However, in a preferred embodiment of the present invention, the common values within a particular field vector are maintained.

In a step 1050, a common value vector that identifies the common values between the first and second field vectors is determined, preferably using the partitioned sets. Step 1050 is described in further detail with reference to FIG. 11.

FIG. 11 is a flow diagram for identifying common values between a pair of field vectors. In a step 1110, three vector indices are initialized. A first vector index, I, is an index into the first field vector FV1; a second vector index, J, is an index into the second field vector FV2; and a third vector index, K, is an index into the common value vector ("CV"). As mentioned above, the common value vector includes the values shared by both first and second field vectors. Indices I and J are initialized to locate a first position in each of the first and second field vectors, respectively. Index K is initialized to locate a position for a next common value to be included in the common value vector.

In a decision step 1120, the present invention determines whether the value in the I-th position of the first field vector is greater than or equal to the value of the J-th position of the second field vector. If so, processing continues at a decision step 1130; otherwise, processing continues at a step 1170. Step 1170 is performed, effectively, when the value in the I-th position of the first field vector is less than the value of the J-th position of the second field vector. In step 1170, the first index I is adjusted to locate the beginning of the next partitioned set in the first field vector. After step 1170, processing continues at a decision step 1160.

In decision step 1130, the present invention determines whether the value in the I-th position of the first field vector is equal to the value of the J-th position of the second field vector. If so, processing continues at a decision step 1140; otherwise processing continues at a step 1180. Step 1180 is performed, effectively, when the value in the I-th position of the first field vector is greater than value of the J-th position of the second field vector. In step 1180, the second index J is adjusted to locate the beginning

of the next partitioned set in the second field vector. After step 1180, processing continues at decision step 1160.

Step 1140 is performed, effectively, when the value in the I-th position of the first field vector is equal to the value of the J-th position of the second field vector. In
5 step 1140, the value included in both the first and second field vectors is placed in the common value vector.

In a step 1150, the third index K is incremented to locate the position in the common value vector of the next common value to be identified. The first index I is adjusted to locate the beginning of the next partitioned set in the first field vector. The
10 second index J is adjusted to locate the beginning of the next partitioned set in the second field vector.

In decision step 1160, the present invention determines whether additional partitioned sets exist in both the first field vector and the second field vector. If so, processing continues at step 1120. If no partitioned sets remain in either the first field vector
15 or the second field vector, processing ends. When processing ends, the common value vector includes all the duplicate data identified between the first and second field vectors.

FIG. 12 illustrates an example of identifying duplicate data between field vectors according to the present invention. Steps 1010 and 1030 sort and partition field
20 vector 1 ("FV1") and steps 1020 and 1040 sort and partition a field vector 2 ("FV2"). The operation of step 1050 is now described with reference to steps 1110-1180 where traversal through steps 1120 to step 1160 and back to step 1120 is referred to as a "loop."

In a first loop, the first element (*i.e.*, 0-th position) of FV1 is compared with the
25 first element of FV2. (This is illustrated in FIG. 12 as a line between FV1 and FV2 having arrows on both ends and annotated with 1). In this example, a value '8' of FV1 is compared with a value '8' of FV2. Decision steps 1120 and 1130 determine that these values are equal and, in step 1140, the value '8' is placed in the common value vector. (This is illustrated in FIG. 12 as a line between FV2 and the COMMON
30 VALUE VECTOR having arrows on both ends and annotated with 1'.) Step 1150 adjusts the indices of both field vectors to point at the next partitioned set. Decision step

1160 determines that more partitioned sets exist in both field vectors and a second loop is started.

In the second loop, the next element of FV1 is compared with the next element of FV2. In this example, a value '9' of FV1 is compared with a value '9' of FV2. These values are again determined to be equal and the value '9' is placed in the common value vector. As before, step 1150 adjusts both indices to point at the next partitioned sets in their respective field vectors. Decision step 1160 determines that more partitioned sets exist in both field vectors and a third loop is started.

In the third loop, the next element of FV1 is compared with the next element of FV2. In this example a value '10' of FV1 is compared with a value '12' of FV2. Decision step 1120 determines that the value in FV1 is not greater than or equal to the value in FV2 and, in step 1170, the index to FV1 is adjusted to point at the next partitioned set therein. Decision step 1160 determines that more partitioned sets exist in both field vectors and a fourth loop is started.

In the fourth loop, the next element of FV1 is compared with the previous value of FV2. In this example, a value '12' of FV1 is compared with the previously compared value of '12' of FV2. Decision steps 1120 and 1130 determine that the values are equal, and in step 1140, the value '12' is placed in the common value vector. Step 1150 adjusts both indices to point at the next partitioned sets in their respective field vectors. Decision step 1160 determines that more partitioned sets exist in both field vectors and a fifth loop is started.

In the fifth loop, the next element of FV1 is compared with the next value of FV2. In this example, a value '15' of FV1 is compared with a value '18' of FV2. Decision step 1120 determines that the value in FV1 is not greater than or equal to the value in FV2 and, in step 1170, the index to FV1 is adjusted to point at the next partitioned set therein. Because no more partitioned sets exist in FV1, processing ends.

In this example, five loops with a maximum of two comparisons per loop are required to identify three common values between the two field vectors. In a brute force method, 132 comparisons ($12 * 11$) are required.

Pre-Encoding Information

In various embodiments of the present invention, prior to, or in some embodiments, contemporaneously therewith, converting data from its original format into a numeric format, the data is pre-encoded into an intermediate encoded format. This pre-
5 encoding further reduces or compresses the information in the original format to the encoded format. Once in the encoded format, the data can be subsequently represented in an appropriate numeric format as described above. These embodiments of the present invention are best described by way of examples.

In one embodiment of the present invention, phonemes are used to represent the
10 data in its original format as the encoded format. In this embodiment, phonemes may be used to encode words, portions of words (*e.g.*, syllables), or phrases of words. Thus, identical or similar sounding words or syllables are represented using the same phonemes. For example, the names "John" or "Jon" would be represented using the same phonemes. In some embodiments, the name "Joan" may also be represented using the
15 same phonemes as those used for the names "John" and "Jon". According to the present invention, each phoneme is subsequently represented as a digit in an appropriate number system based in part on the phonemes utilized.

For example, a particular language may be broken down into its finite number of "sounds" or phonemes and represented as digits within an appropriate number system. In this manner, text may be encoded based on phonetics rather than particular
20 spellings thereby minimizing the effect of spelling errors, for example, with the use of search engines.

These embodiments of the present invention may be extended for speech, speech recognition, and artificial speech rendering mechanisms. In particular, aural
25 speech phonemes (as opposed to corresponding text phonemes) may also be represented as described above in an appropriate number system and used to simplify speech recognition and speech rendering as described above.

In other embodiments of the present invention, words, phrases, idioms, sentences, and/or ideas may be pre-encoded and then subsequently be represented as num-

bers in an appropriate number system as described above. Such embodiments may be used, for example, to improve automated language translation systems. These embodiments may also be used to improve search engines. Large portions of text that refer to one or more ideas or concepts may be pre-encoded based on each of the ideas or
5 concepts conveyed. These embodiments provide for conceptual searching as opposed to identifying and/or locating specific words or phrases that may or not appear with the passage.

In another embodiment of the present invention, raw address information is pre-encoded into coordinates expressed, for example, as longitude and latitude and subsequently represented in an appropriate number system, for example, a base 60 number
10 system. Such a system may be particularly useful for mapping operations, navigation systems or tracking systems.

In another embodiment of the present invention, raw fingerprint data is pre-encoded into various parameters, registration points, or other identifying indicia appropriate for classifying fingerprints, each of which are subsequently represented as a corresponding digit in an appropriate number system. Each fingerprint may thus be represented by a value in a field, or alternatively, each fingerprint may be represented as a vector of fields. This resulting data may be organized and maintained in a database of such information based on fingerprints collected from individuals for a variety of purposes (*i.e.*, both criminal and non-criminal). These may include fingerprints collected
15 by forensic scientists, security officers, background investigators, etc. The present invention is ideally suited for cleaning existing fingerprint databases, merging those databases into a reference database, adding new fingerprint information as it becomes available, and matching fingerprint information with that in the reference database.
20

It should be understood, that in embodiments employing pre-encoding, in many cases, the underlying original data must be pre-processed into the intermediate format. Thus, in order for the present invention to be employed in a search context, the information to be searched must be pre-encoded or "pre-processed". In some cases, this pre-processing may result in the loss of semantic significance as described above with respect to other embodiments of the present invention.
25
30

Exemplary Embodiments

Various embodiments of the present invention may be used for many different applications, some of which have been described and/or alluded to above. For example, in the application described above, the invention may be used to combine billing information collected from multiple sources to derive a distilled database in which related data records are recognized and duplicate and erroneous data records are eliminated. As suggested, this may be particularly useful in cases, for example, involving fraud. Typically, persons using credit card or other forms of retail fraud make minor changes to certain pieces of their personal information while leaving the majority of it the same. For example, oftentimes, digits in a social security number may be transposed or an alias may be used. Often, however, other information such as the person's address, date of birth, mother's maiden name, etc., is used identically. These types of fraud are readily identified by the present invention, even though they are difficult to identify by human analyses.

Other possible applications include uses in telemarketing, to compile a list of targeted individuals or addresses; in mail-order catalogs, to reduce a number of catalogs sent to the same individual or family; or to merge records from various vendors selling similar databases. Still another potential application is in the medical research or diagnostics fields, in which nucleotide sequences of Adenine (A), Guanine (G), Cytosine (C), and Thymine (T) in nucleic acids may be identified. Another application for use by taxing organizations such as the Internal Revenue Service, state and local governments, etc., organizes and maintains accurate rolls and tax basis information.

In other embodiments, the present invention may be used as a gatekeeper for a particular database at the outset to maintain integrity of the database from the very beginning, rather than achieving integrity in the database at a later date. In these embodiments, no raw data 210 is present and only new data 240 exists. Before new data 240 is added to the database, it is measured against distilled database 230 to determine whether new data 240 includes additional information or content. If so, only that new information or content is added to distilled database 230 by updating an existing record in distilled database 230 to reflect the new information or content as would be apparent.

In another embodiment of the present invention, a mailing service, such as the United States Postal Service, or a courier delivery service, such as Airborne Express,

Federal Express, United Parcel Service, etc, uses the present invention to maintain a list of valid delivery addresses. An address associated with an item to be delivered is checked against a reference database of addresses to identify any inaccuracies in the address. Inaccurate addresses may either be corrected (*e.g.*, for transposed numbers,
5 etc.) or the sender may be contacted to verify the address. New addresses may be added to the reference database as they become available, for example, as items are successfully delivered. In addition, certain senders may be identified as prone to misaddressing items or providing incorrect addresses. If appropriate, these senders may be notified accordingly.

10 In addition to using the present invention to matching fragments of DNA sequences as discussed above, genetic researchers (*e.g.*, drug companies, seed companies, animal breeders, etc.) may also use the present invention to represent palpable, tangible, and/or objective characteristics of individuals in a set and use this information to identify the individual genes or gene sequences responsible for these characteristics.

15 In another embodiment, the present invention is used for signal (packet) switching and routing data on a network, such as the Internet. Incoming packets are examined for a destination address and sequence information and sorted into an appropriate output queue in the proper order. In this embodiment, the present invention's ability to sort numbers provides a distinct advantage over conventional systems. This
20 coupled with an expanded address space as a result of using an alternate number system (as opposed to a conventional number system presently employed) provides an improved method of network addressing and communication protocols.

In another embodiment, the present invention is used for rendering and displaying objects in a three-dimensional environment. These activities require tremendous amounts of sorting in order to determine which objects to display in the foreground and which objects are correspondingly obscured in the background as well as to
25 determine lighting characteristics for each of the objects (*i.e.* shadowing, etc.).

While this invention has been described in a preferred embodiment, other embodiments and variations are within the scope of the following claims. For example,
30 formatting process 300 may format data using different radices or other character sets, and may use various data structures. The data structures may represent multiple fields, and depending on the application, will represent a variety of fields. For example, in a

credit application, fields may include an account status, an account number, and a legal status, in addition to personal information about the account holder. In a medical diagnostic application, fields may include various alleles or other genetic characteristics detected in tissue samples.

**THE EMBODIMENTS OF THE INVENTION IN WHICH AN EXCLUSIVE
PROPERTY OR PRIVILEGE IS CLAIMED ARE DEFINED AS FOLLOWS:**

1. A method operative on a computer for converting information from at least
5 one raw database accessible from the computer into a distilled database accessible
from the computer, the raw database including a plurality of raw records, each of
the plurality of raw records including at least one non-numeric raw data field, each
non-numeric raw data field including a plurality of data elements, the method
comprising:
- 10 converting the at least one non-numeric raw data field in a raw record into a
numeric value, the numeric value being generated by representing each of the
plurality of data elements in the at least one non-numeric raw data field as a digit
in a number system, the number system selected based on a range of possible
values for any of the plurality of data elements in the at least one non-numeric raw
15 data field, the number system having a unique digit for representing each possible
value of a data element in the at least one non-numeric raw data field;
- forming a vector including at least the numeric value, the vector being
representative of the raw record;
- determining whether or not to include said vector representative of the raw
20 record in the distilled database by comparing said vector with at least one other
vector from the distilled database, and
- upon determining to include said vector in the distilled database, storing
said vector in the distilled database, wherein the numeric value stored in the
distilled database retains its value in the number system.
- 25
2. The method of claim 1, wherein said determining whether or not to include
said vector representative of the raw record in the distilled database comprises
numerically comparing the numeric value from said vector with a corresponding
numeric value from at least one other vector from the distilled database.

3. The method of claim 1, wherein said determining whether or not to include said vector representative of the raw record in the distilled database comprises determining a dot product between said vector and the at least one other vector from the distilled database.

5

4. The method of claim 1, wherein said determining whether or not to include said vector representative of the raw record in the distilled database comprises performing a pattern recognition analysis between said vector and the at least one other vector from the distilled database.

10

5. The method of claim 1, wherein said determining whether or not to include said vector representative of the raw record in the distilled database comprises determining a cross product between at least said vector and the at least one other vector from the distilled database.

15

6. The method of claim 1, wherein said determining whether or not to include said vector representative of the raw record in the distilled database comprises determining a vector magnitude of said vector.

20 7. The method of claim 1, wherein said determining whether or not to include said vector representative of the raw record in the distilled database comprises determining a vector direction of said vector.

8. A method operative on a computer for converting information from at least one raw database accessible from the computer into a distilled database accessible from the computer, the raw database including a plurality of raw records, each of the plurality of raw records including at least one raw data field, each raw data field including a plurality of data elements, the method comprising:

25

converting each of the at least one raw data fields in a raw record into a numeric value, each numeric value being generated by representing each of the plurality of data elements in the raw data field as a digit in a number system, the number system selected based on a range of possible values for any of the plurality
 5 of data elements in the raw data field;

forming a vector representative of the raw record, said vector including the numeric values generated for each of the at least one raw data field in the raw record;

determining whether or not to include said vector representative of the raw
 10 record in the distilled database by comparing said vector with at least one other vector from the distilled database, and

upon determining to include said vector in the distilled database, storing said vector in the distilled database, wherein the numeric value stored in the distilled database retains its value in the number system.

15

9. A method operative on a computer for converting information from at least one raw database accessible from the computer into a distilled database accessible from the computer, the raw database including a plurality of raw records, the plurality of raw records including at least one raw data field, each raw data field
 20 including a plurality of data elements, the distilled database including at least one field vector, the at least one field vector including a plurality of numeric values, the method comprising:

converting at least one raw data field into a numeric value by representing each of the plurality of data elements in the raw data field as a digit in the numeric
 25 value, the numeric value having a representation in a number system selected based on a range of possible values for any of the plurality of data elements in the at least one raw data field;

determining whether or not to include the numeric value corresponding to the at least one raw data field in the at least one field vector by numerically

comparing the numeric value with at least one of the plurality of numeric values in the at least one field vector, and

upon determining to include the numeric value, storing the numeric value as a value in the number system.

5

10. The method of claim 9, wherein said determining whether or not to include the numeric value corresponding to the at least one raw data field in the at least one field vector comprises sorting the plurality of numeric values in the at least one field vector into a numeric order.

10

11. The method of claim 9, wherein the number system has a radix at least equal to a number of the possible values of a data element in the raw data field.

15

12. The method of claim 9, wherein the plurality of data elements include at least one alphanumeric character.

13. The method of claim 9, wherein the raw data field includes chemical information, and wherein the number system has a radix at least equal to a number of chemical structures included in the chemical information.

20

14. The method of claim 9, wherein said converting at least one raw data field into a numeric value comprises converting at least one raw data field into a numeric value so that the numeric value represented in the number system is recognizable as the raw data field.

25

15. The method of claim 10, wherein the radix also maximizes a number of data elements that can be represented by a data word associated with the computer.

16. A method operative on a computer for converting information from at least one raw database accessible from the computer into a distilled database accessible from the computer, the raw database including a plurality of raw records, each of the plurality of raw records including at least one non-numeric raw data field, each
5 non-numeric raw data field including a plurality of data elements, the method comprising:

converting the at least one non-numeric of the raw data field in a raw record into a numeric value, the numeric value being generated by representing each of the plurality of data elements in the at least on non-numeric raw data field as a
10 digit in a number system, the number system selected based on a range of possible values for any of the plurality of data elements in the at least one non-numeric raw data field, the number system having a unique digit for representing each possible value of a data element in the at least one non-numeric raw data field, each unique digit selected so as to be recognizable as the corresponding data element;

15 forming a vector including at least the numeric value, the vector being representative of the raw record, and

determining whether or not to include said vector representative of the raw record in the distilled database by comparing said vector with at least one other vector from the distilled database,

20 wherein the numeric value represented in the number system retains semantic significance of the at least one non-numeric raw data field.

1/12

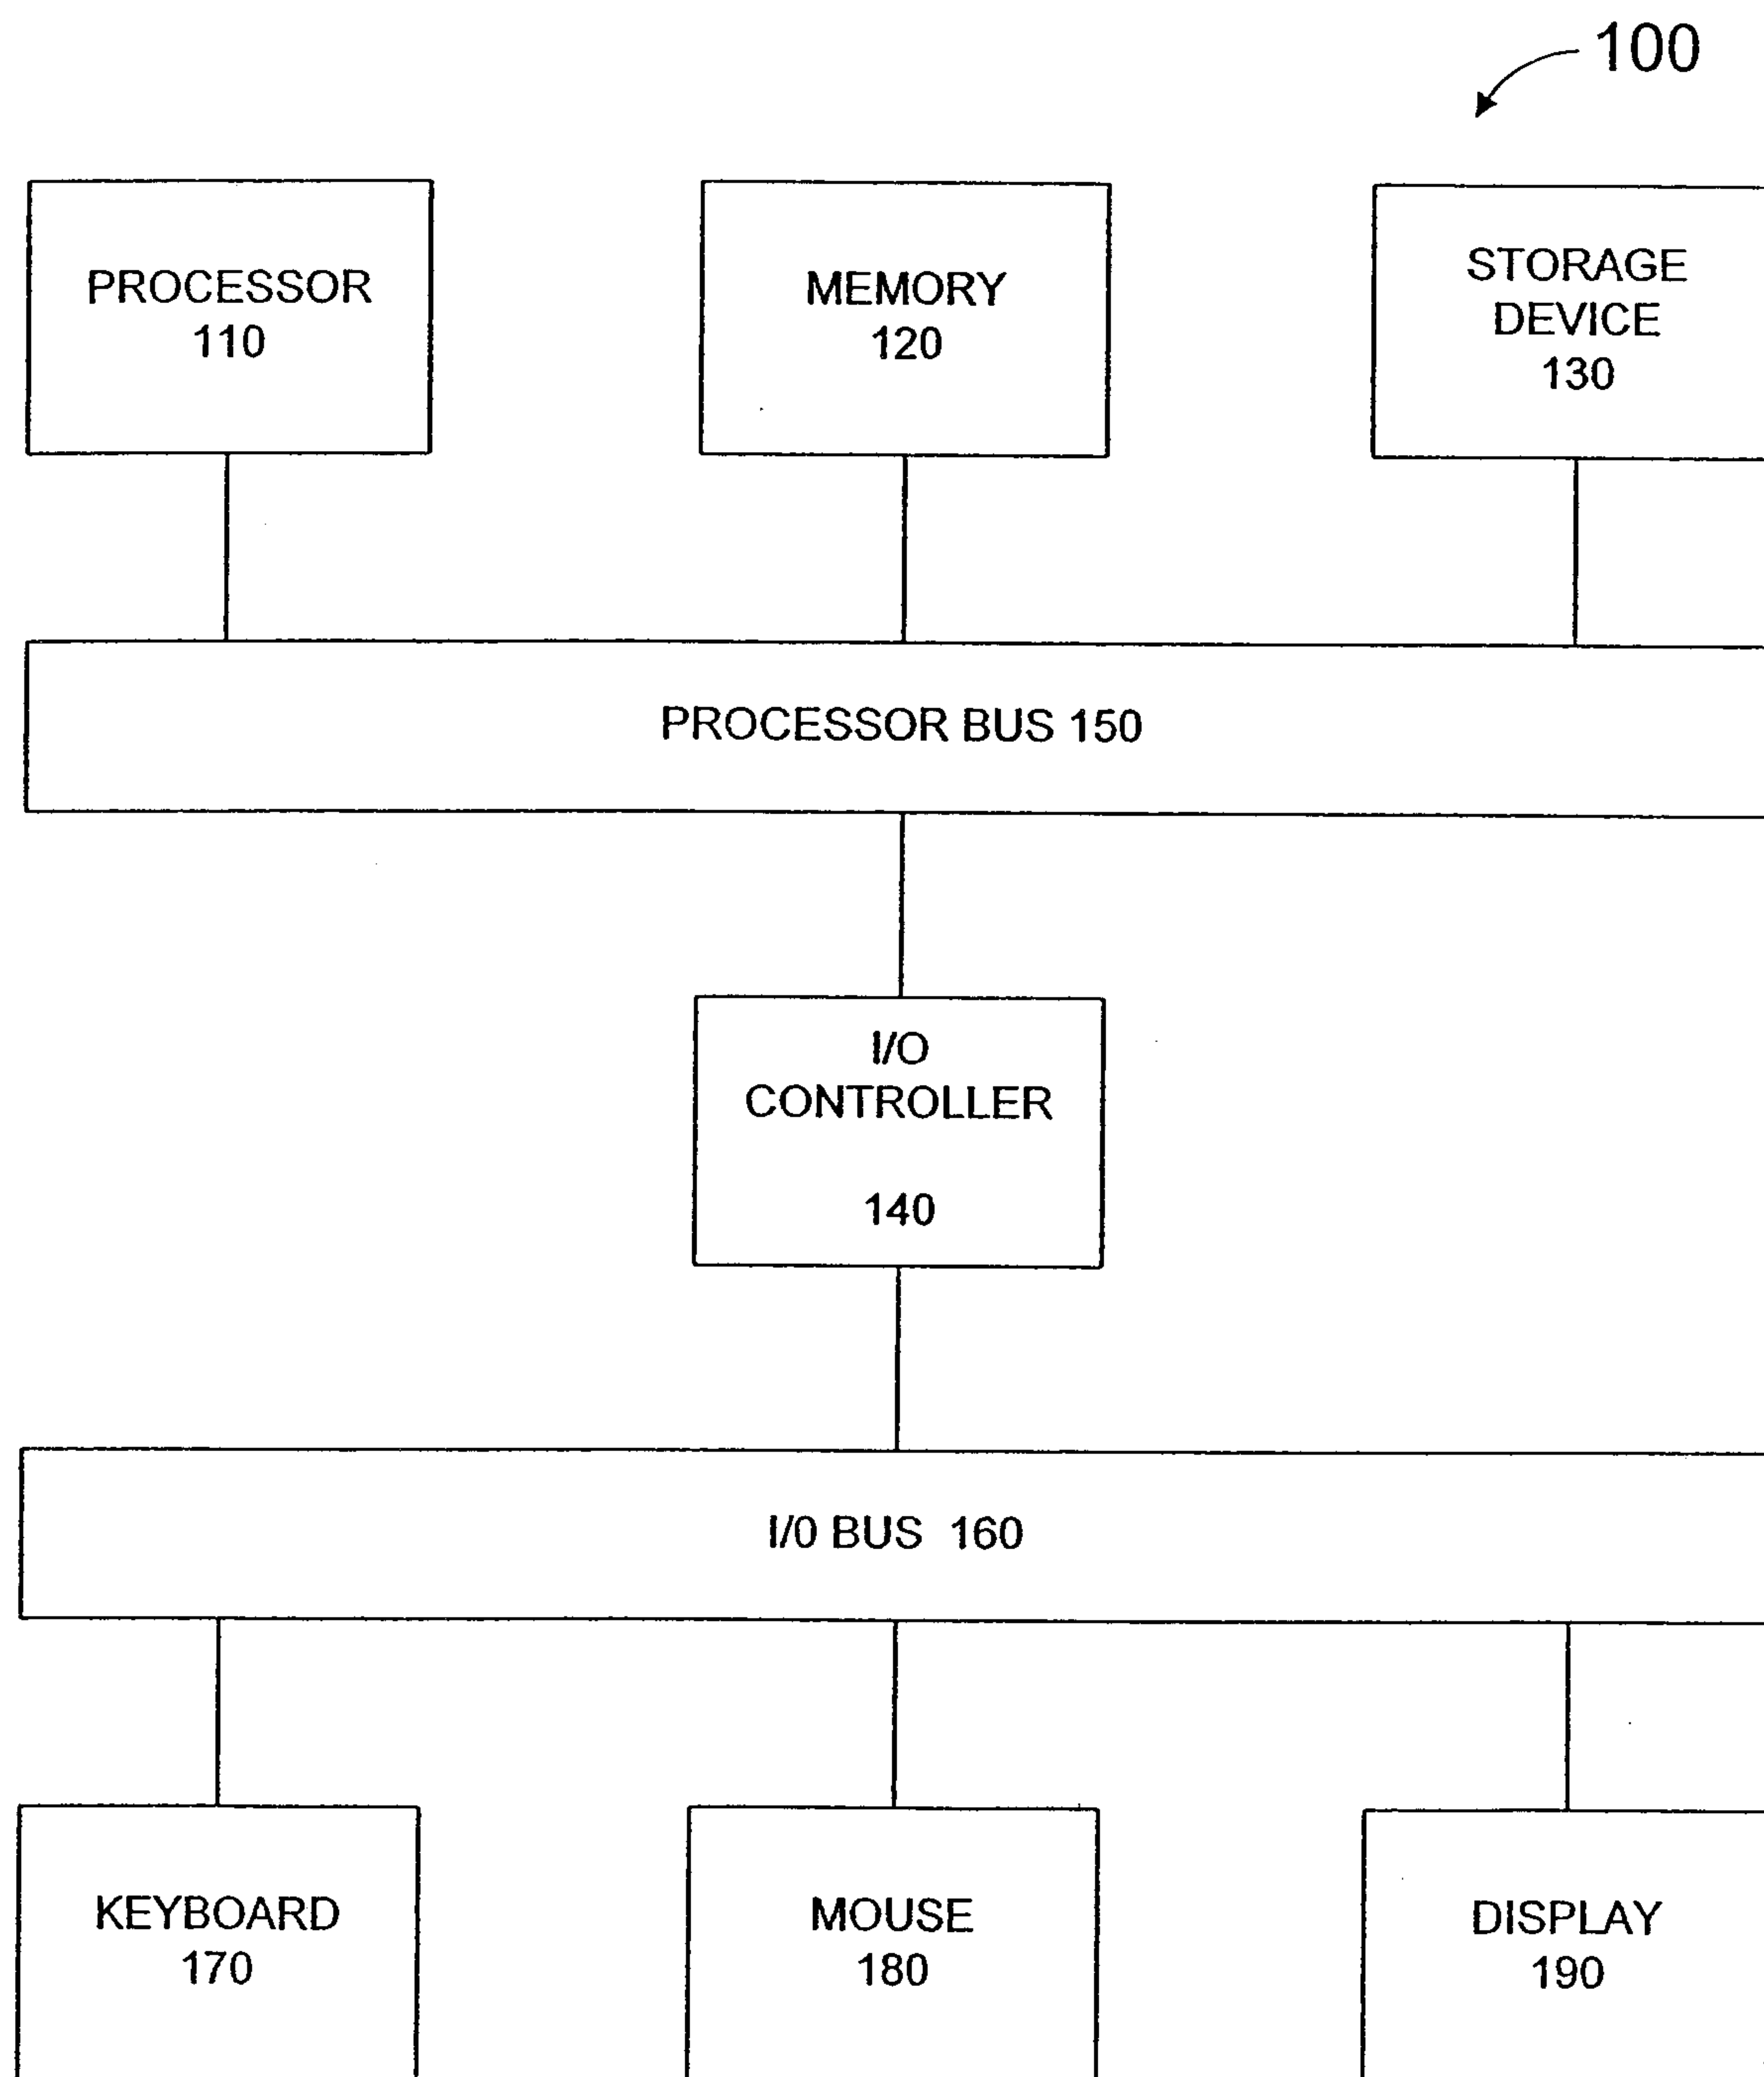


FIGURE 1

2/12

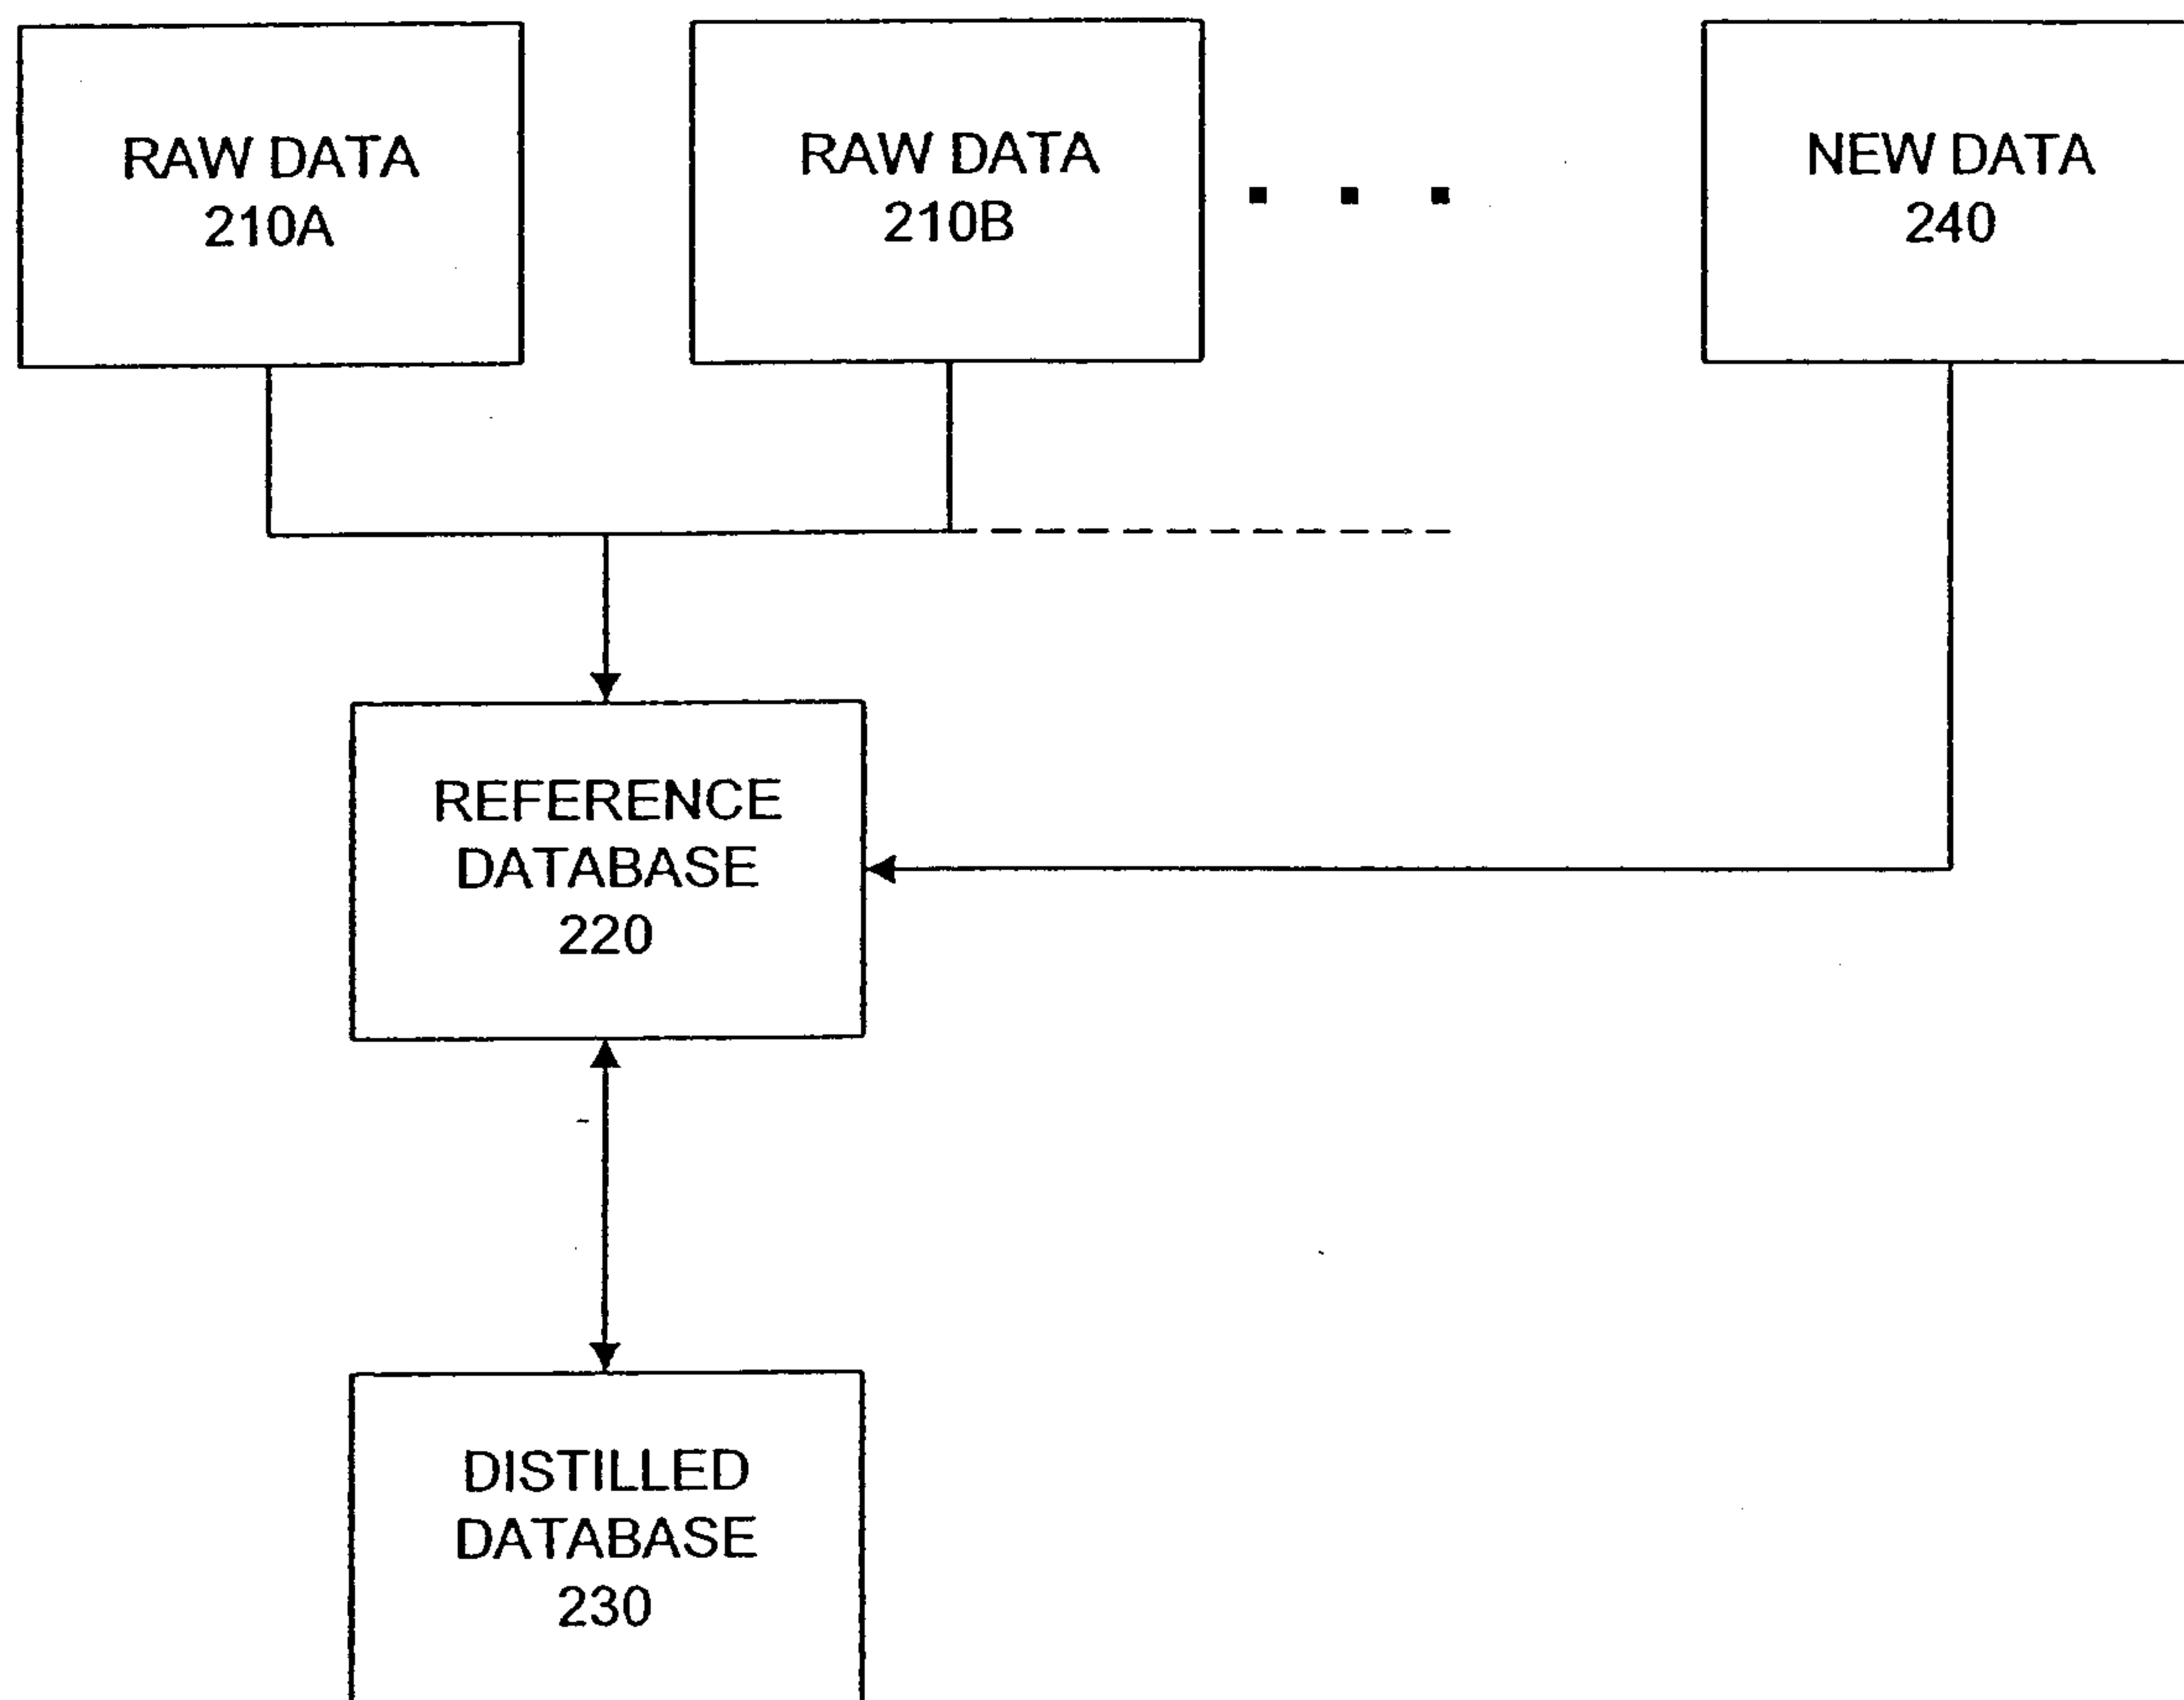


FIGURE 2

3/12

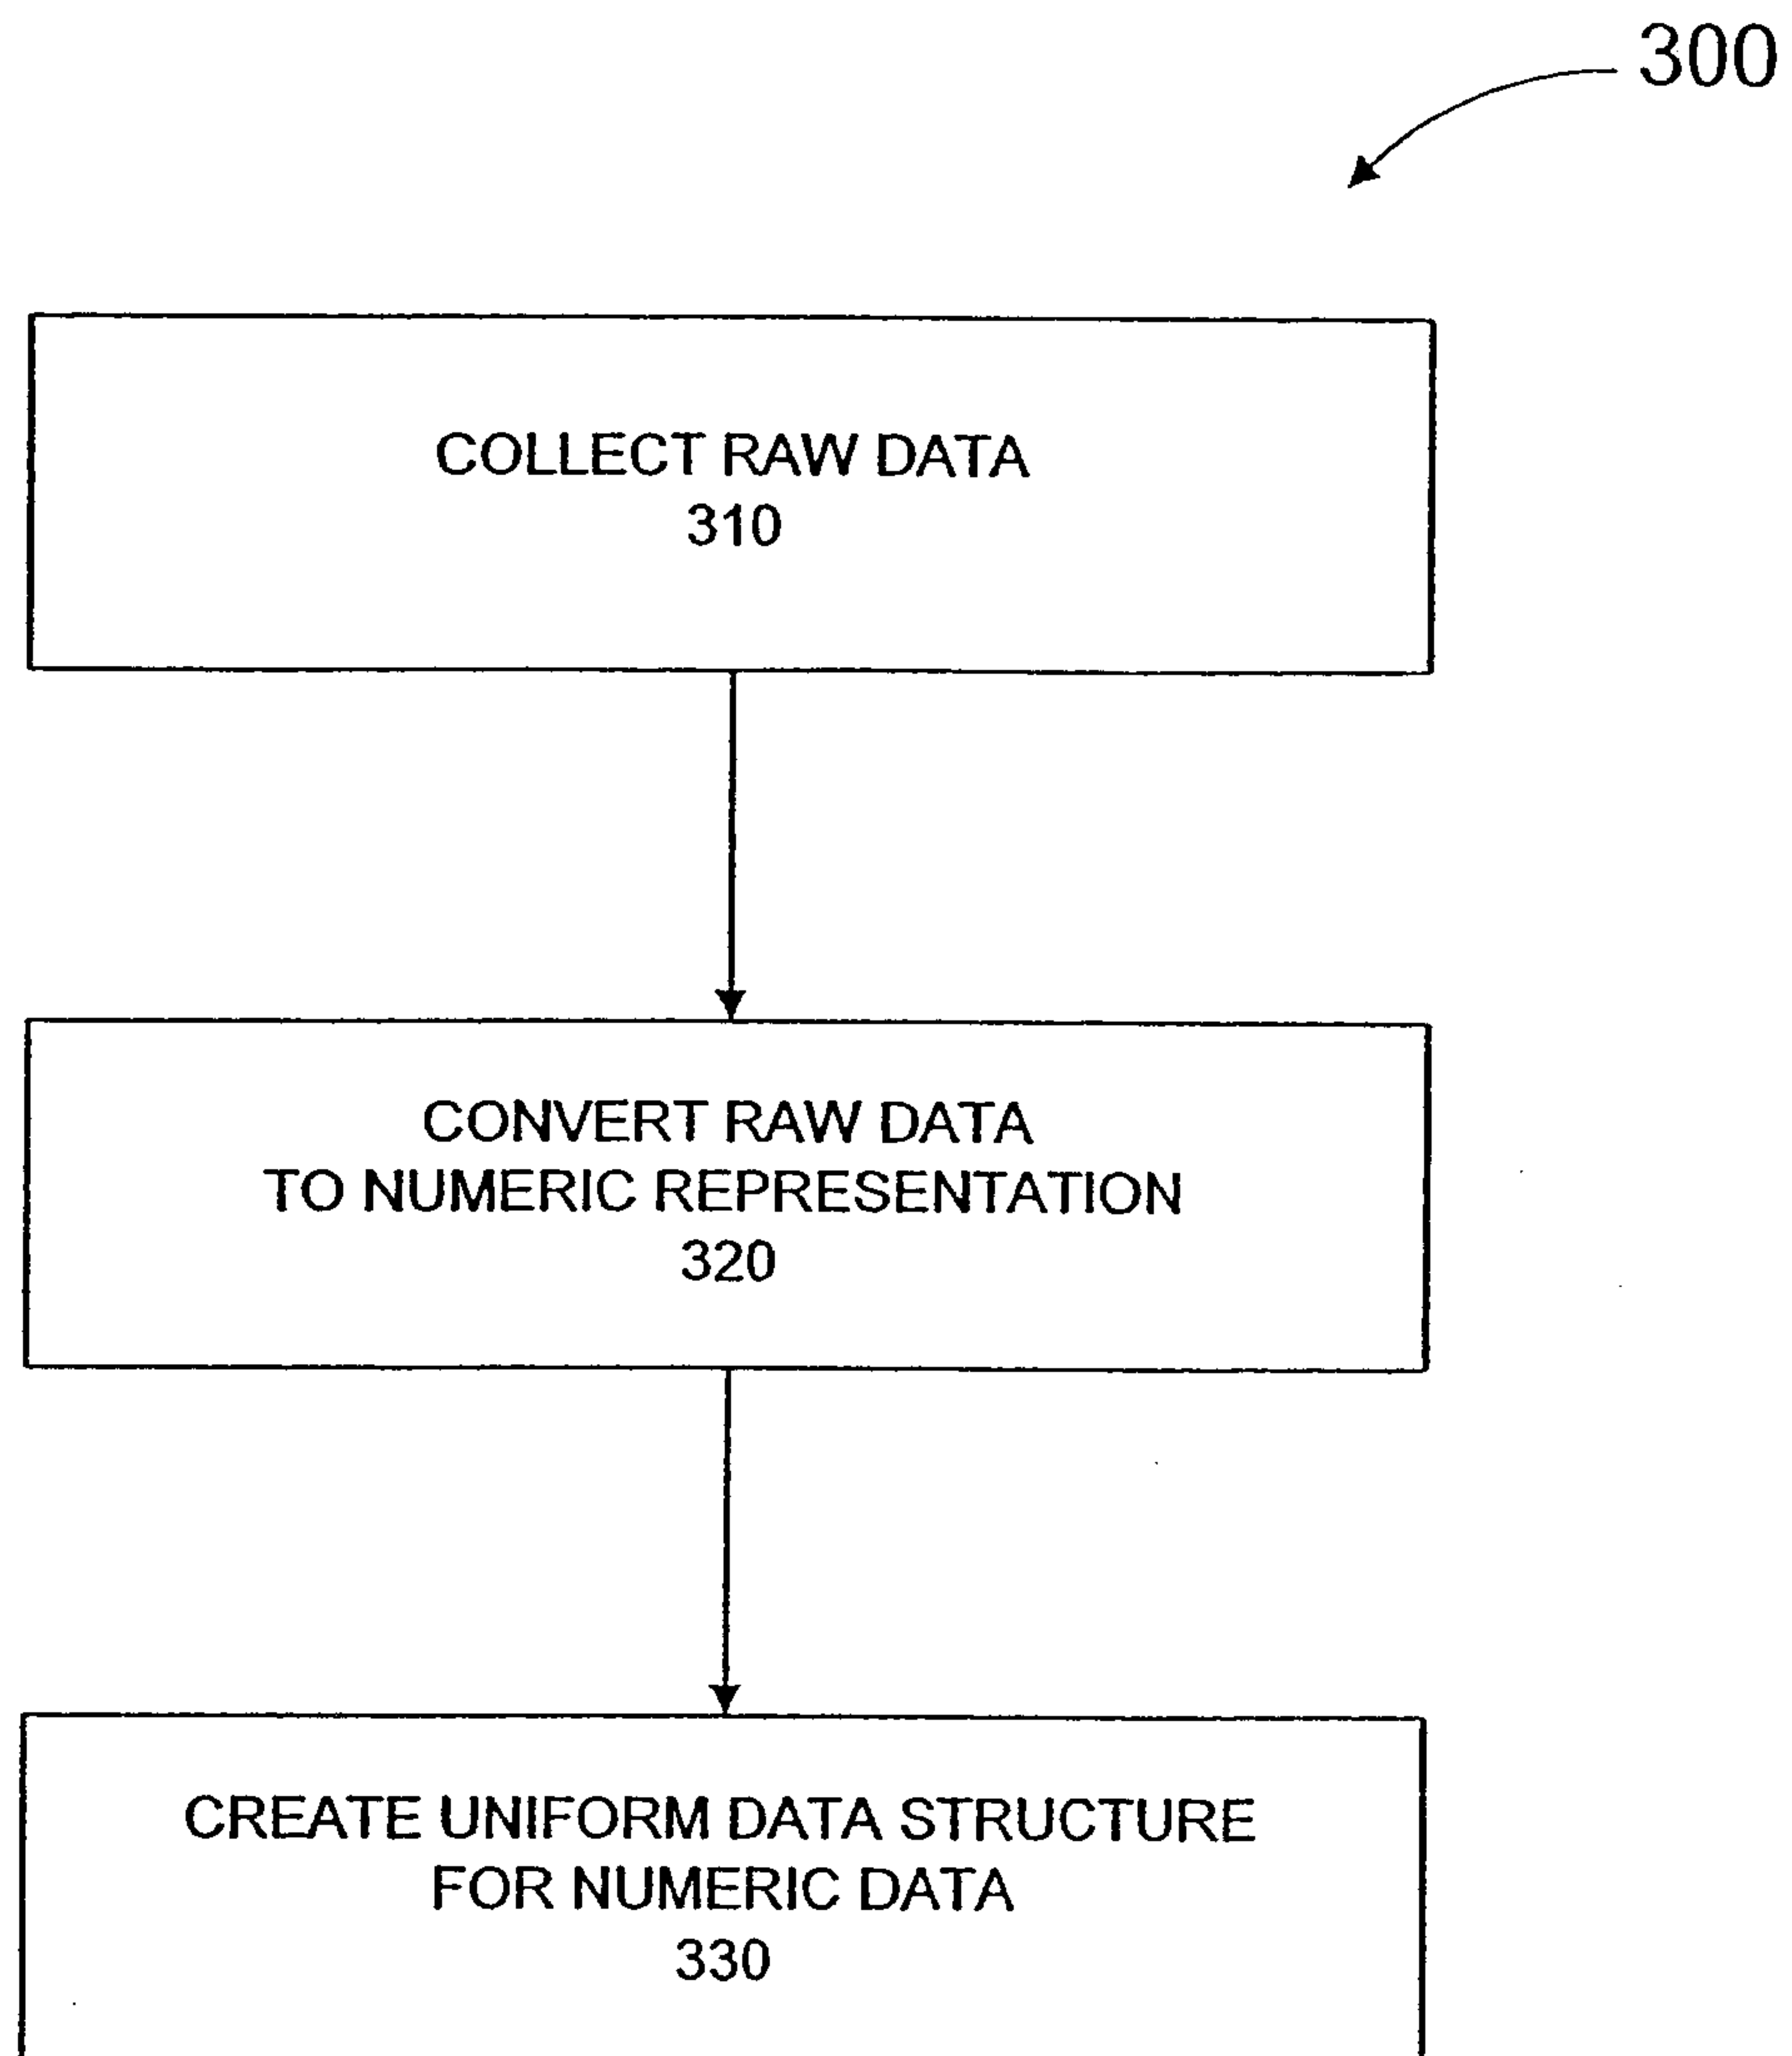


FIGURE 3




FIGURE 4

5/12

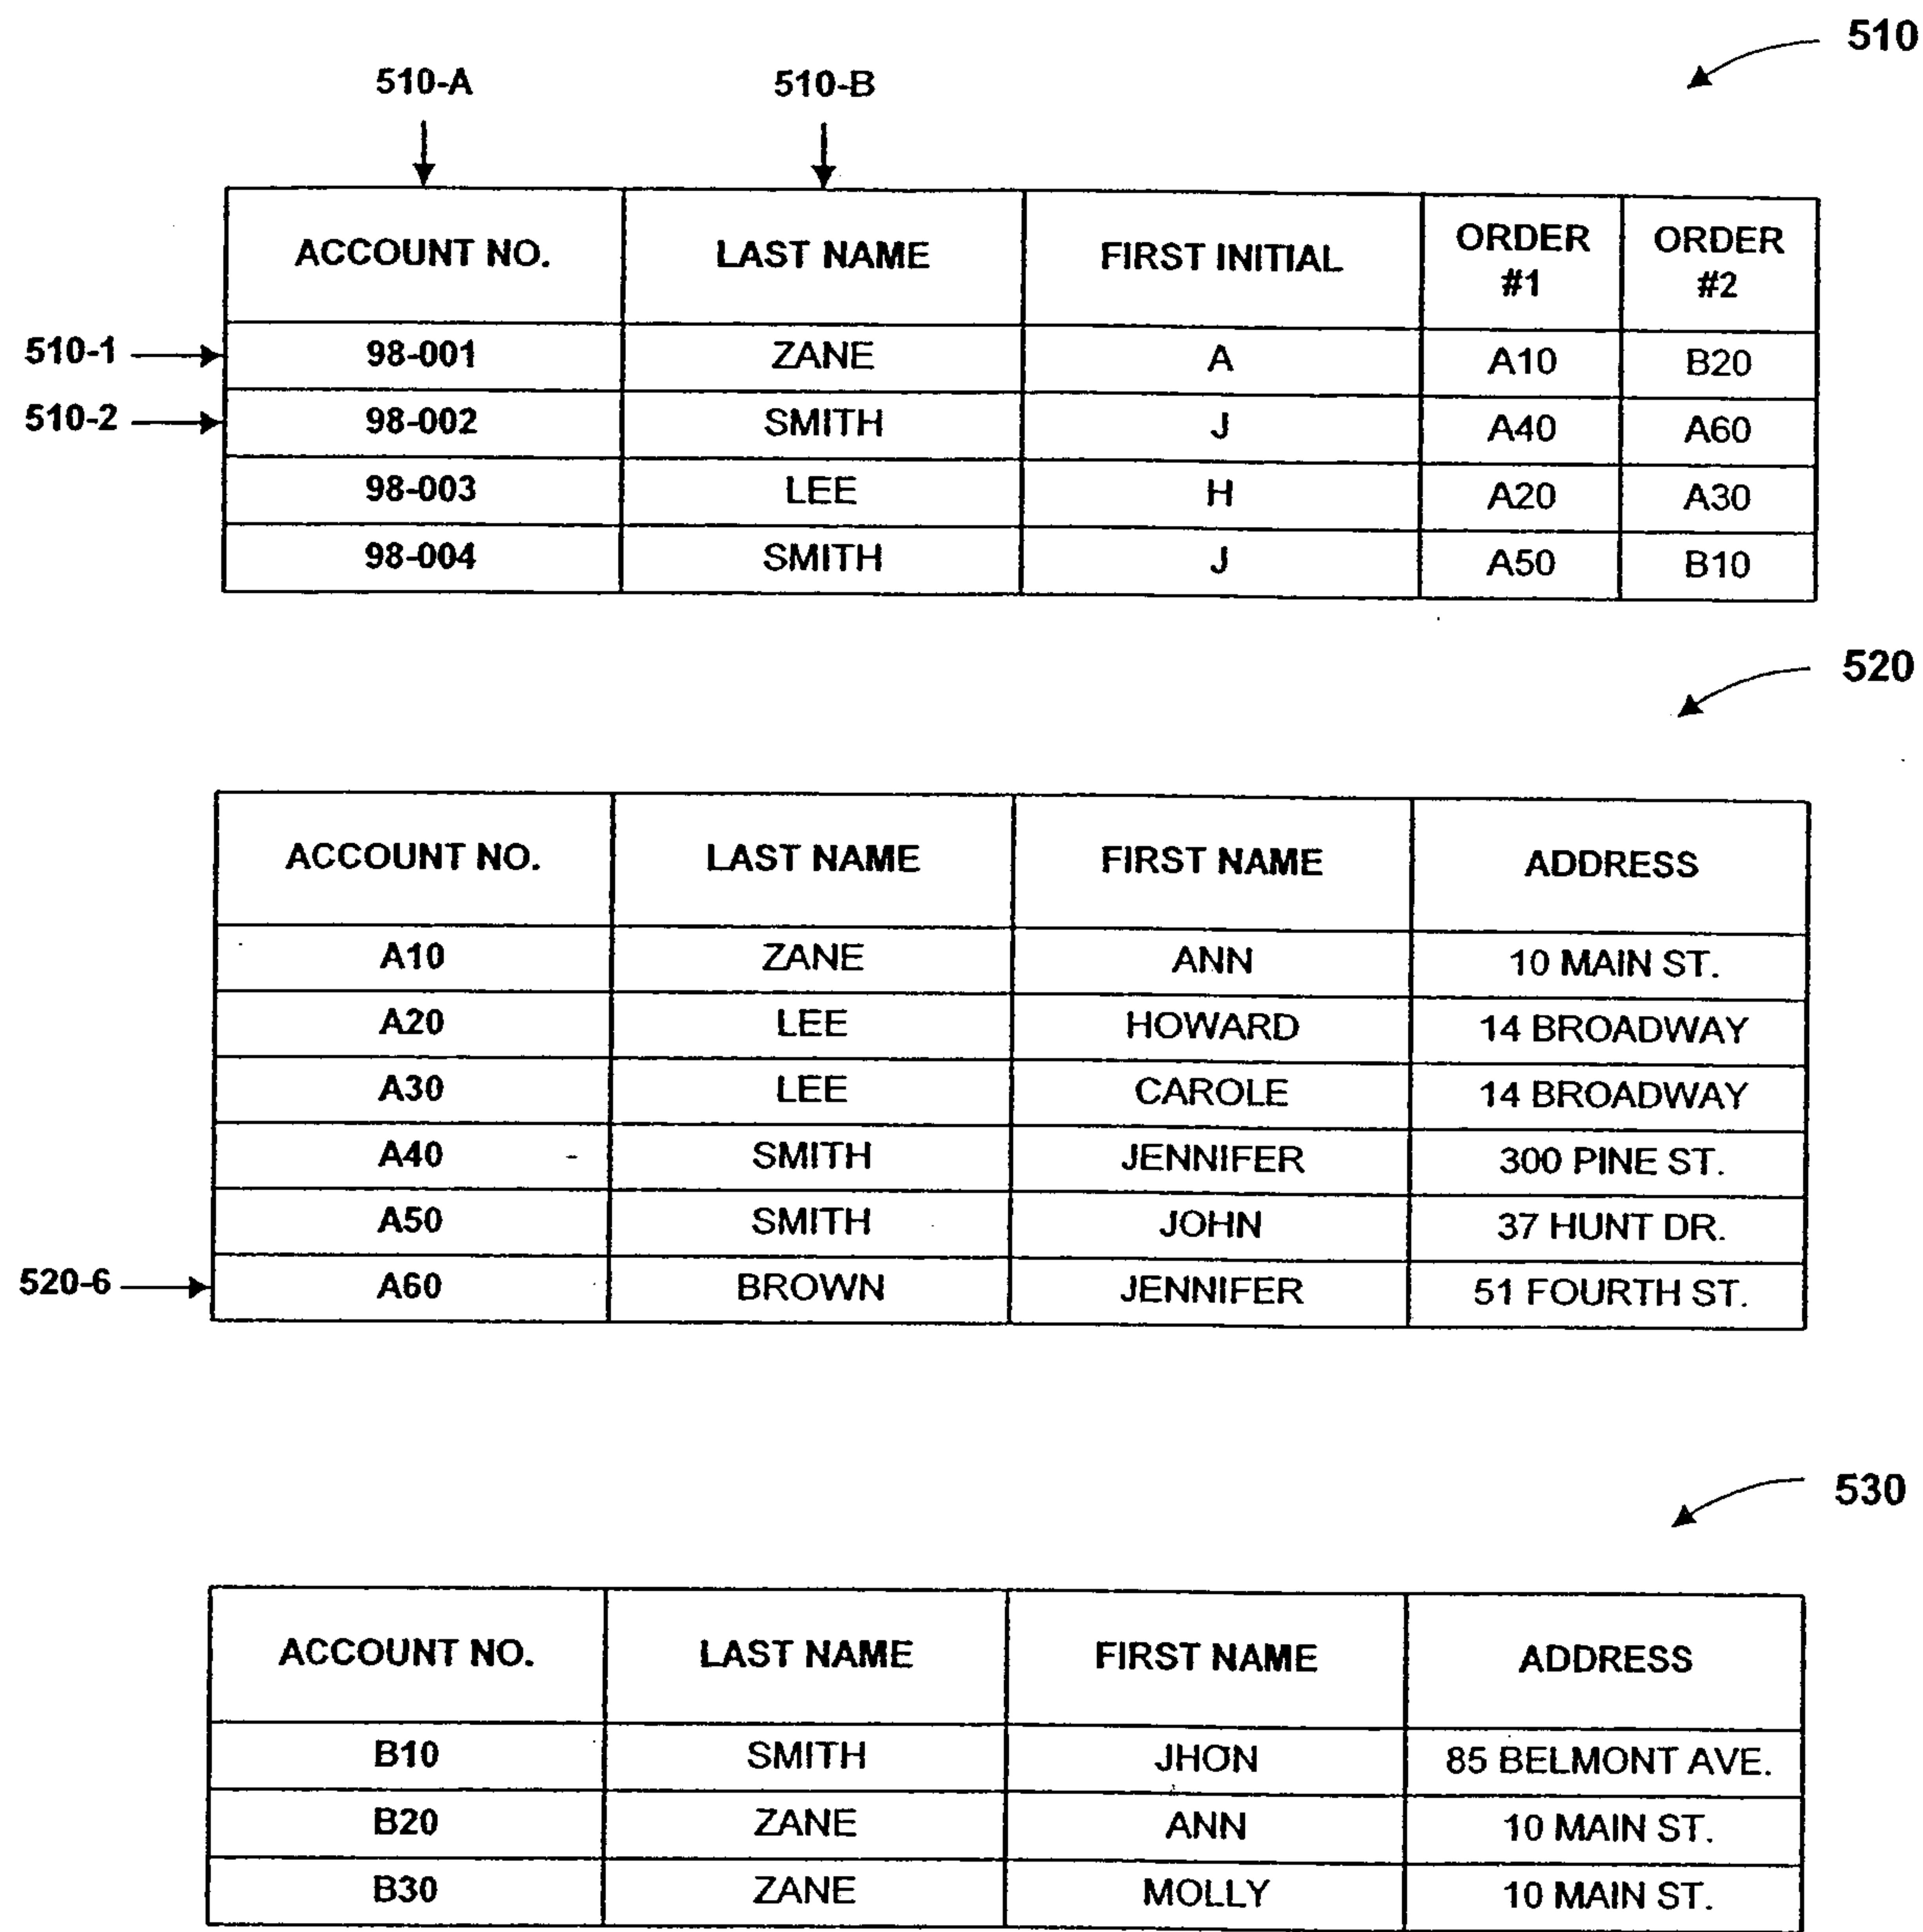


FIGURE 5

6/12

610 PIDN	620 LAST NAME	630 FIRST NAME	640 ACCT NO.	650 ACCT. #1	660 ACCT. #2	670 ADDRESS
1	ZANE	A	98-001	A10	B20	^
2	SMITH	J	98-002	A40	A60	^
3	LEE	H	98-003	A20	A30	^
4	SMITH	J	98-004	A50	B10	^
5	ZANE	ANN	^	A10	^	10 MAIN ST.
6	LEE	HOWARD	^	A20	^	14 BROADWAY
7	LEE	CAROLE	^	A30	^	14 BROADWAY
8	SMITH	JENNIFER	^	A40	^	300 PINE ST.
9	SMITH	JOHN	^	A50	^	37 HUNT DR.
10	BROWN	JENNIFER	^	A60	^	51 FOURTH ST.
11	SMITH	JHON	^	B10	^	85 BELMONT AVE.
12	ZANE	ANN	^	B20	^	10 MAIN ST.
13	ZANE	MOLLY	^	B30	^	10 MAIN ST.

FIGURE 6

7/12

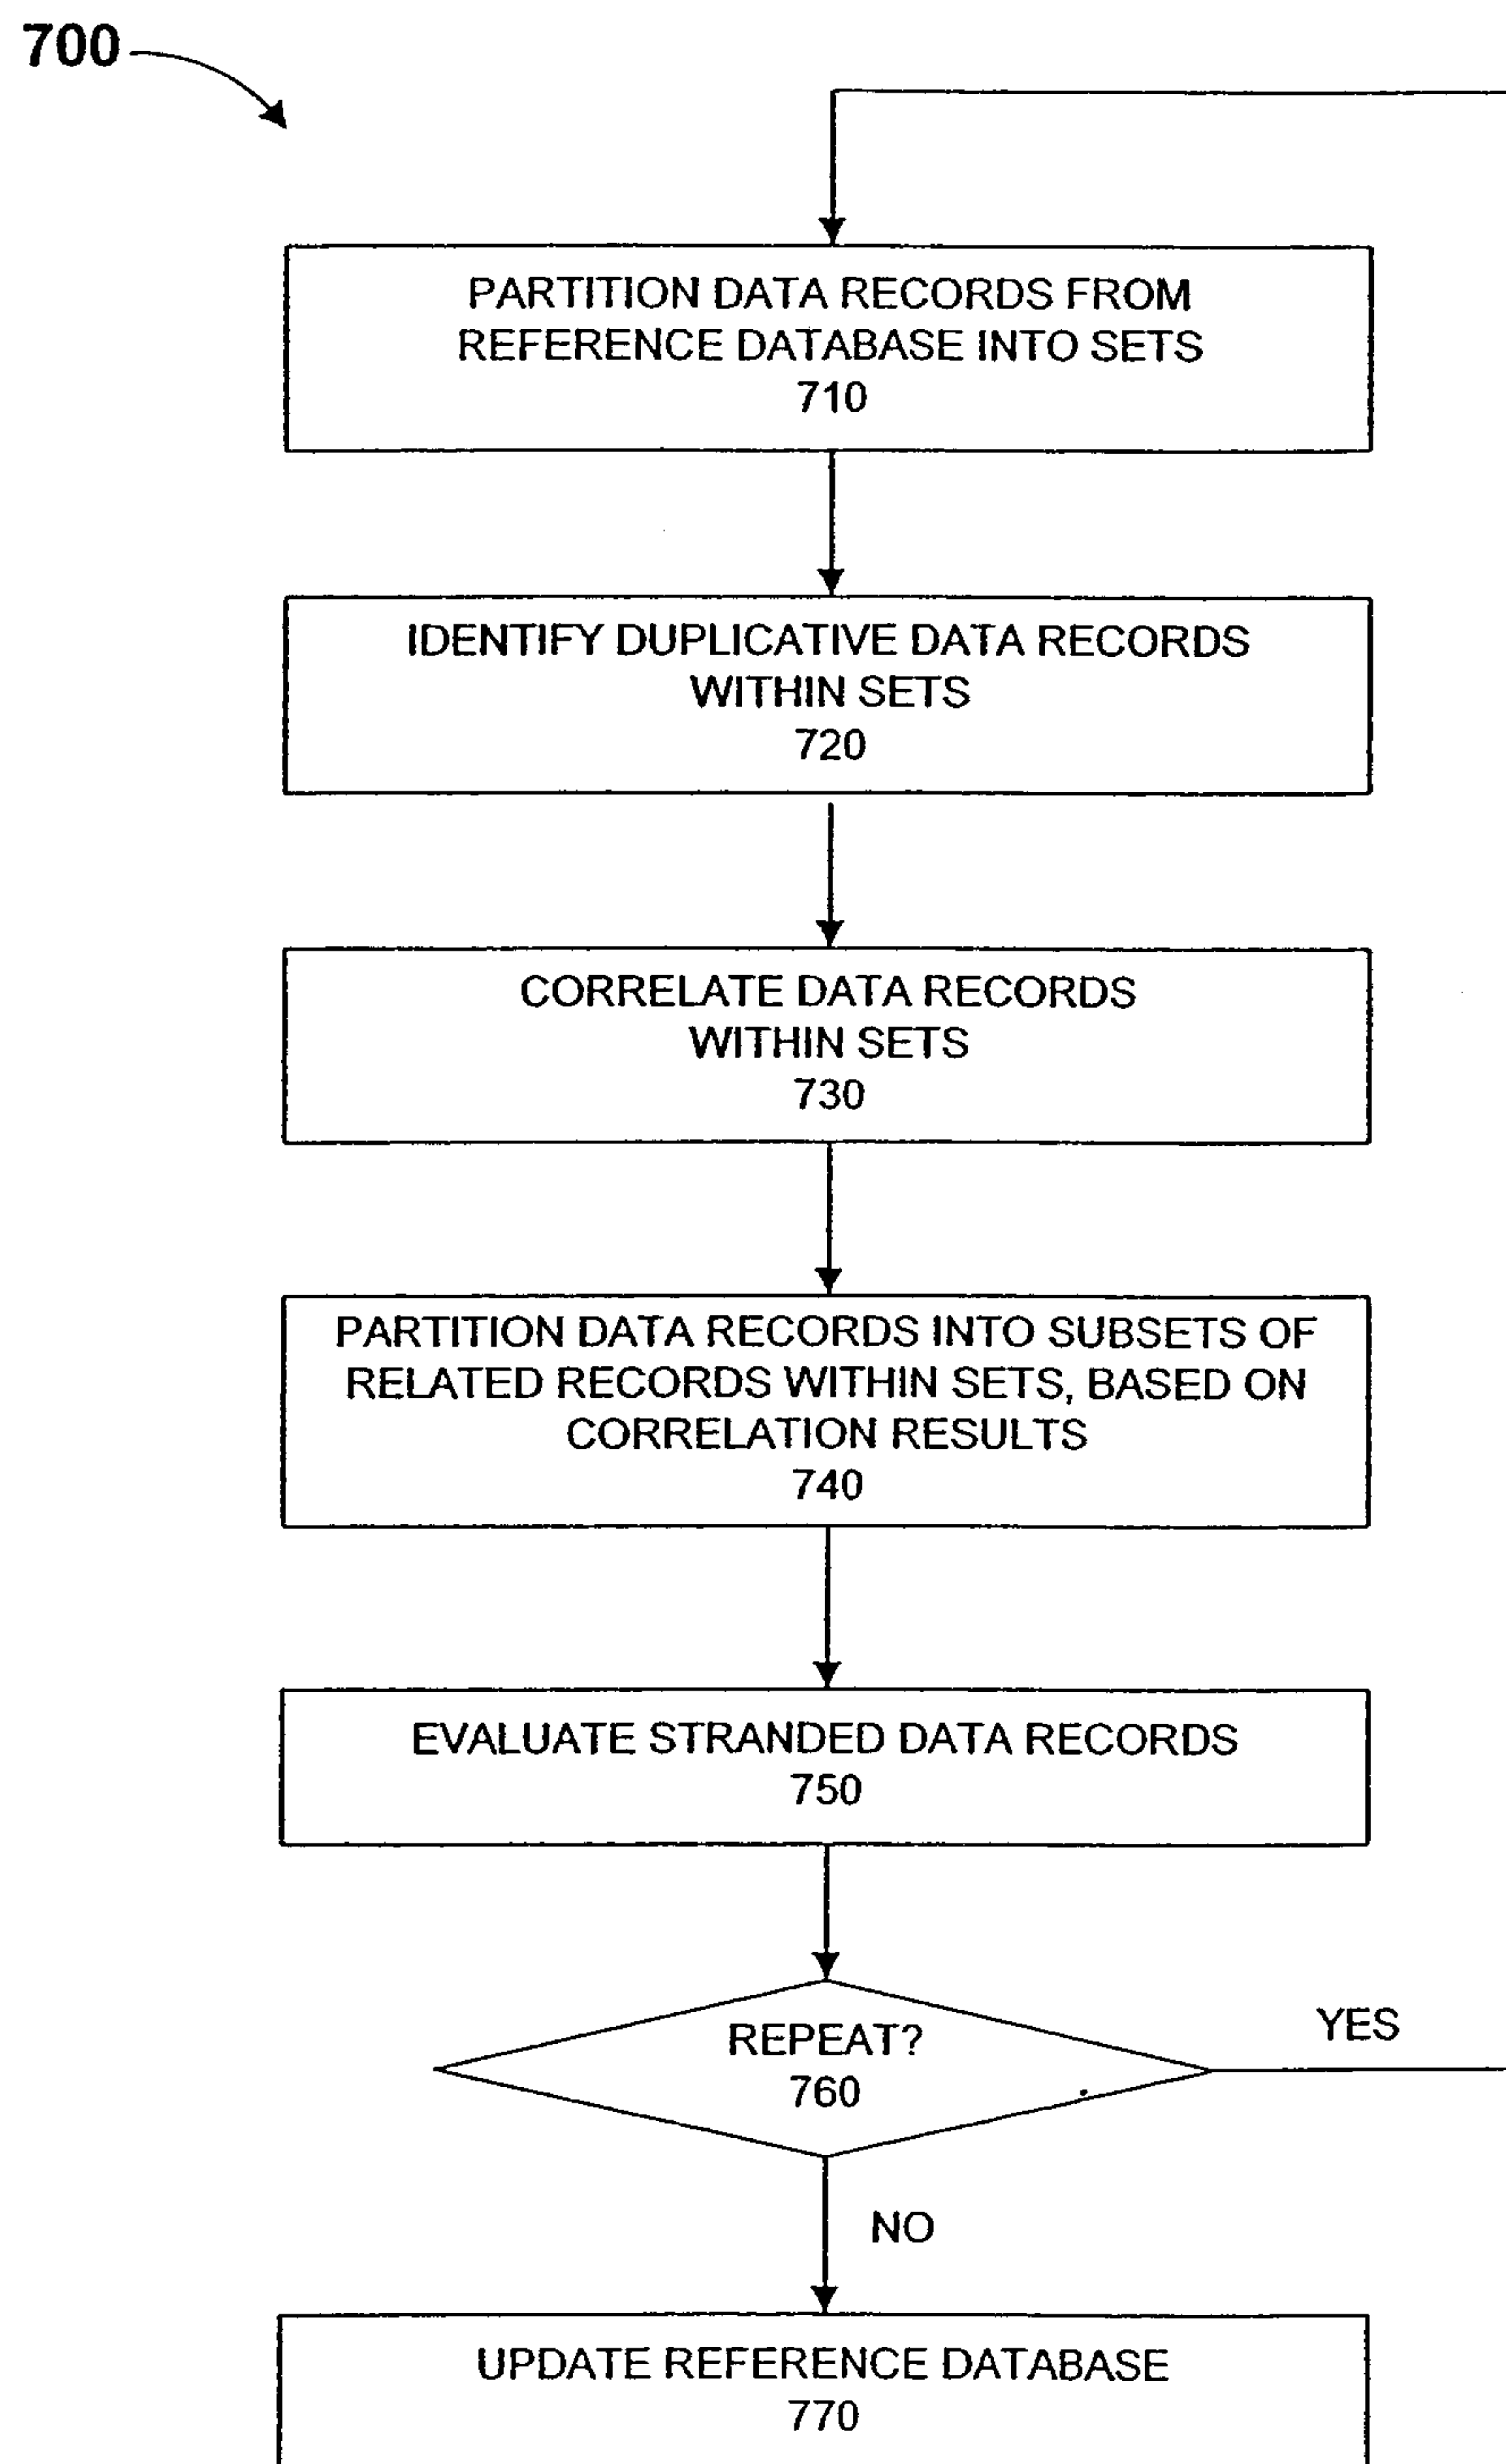


FIGURE 7

8/12

810

LAST NAME	PIDNs
BROWN	10
LEE	3, 6, 7
SMITH	2, 4, 8, 9, 11
ZANE	1, 5, 12, 13

820

LAST NAME	PIDNs
BROWN	10
LEE	3, 6
LEE	7
SMITH	2, 8
SMITH	4, 11
SMITH	9
ZANE	1, 5, 12
ZANE	13

FIGURE 8

9/12


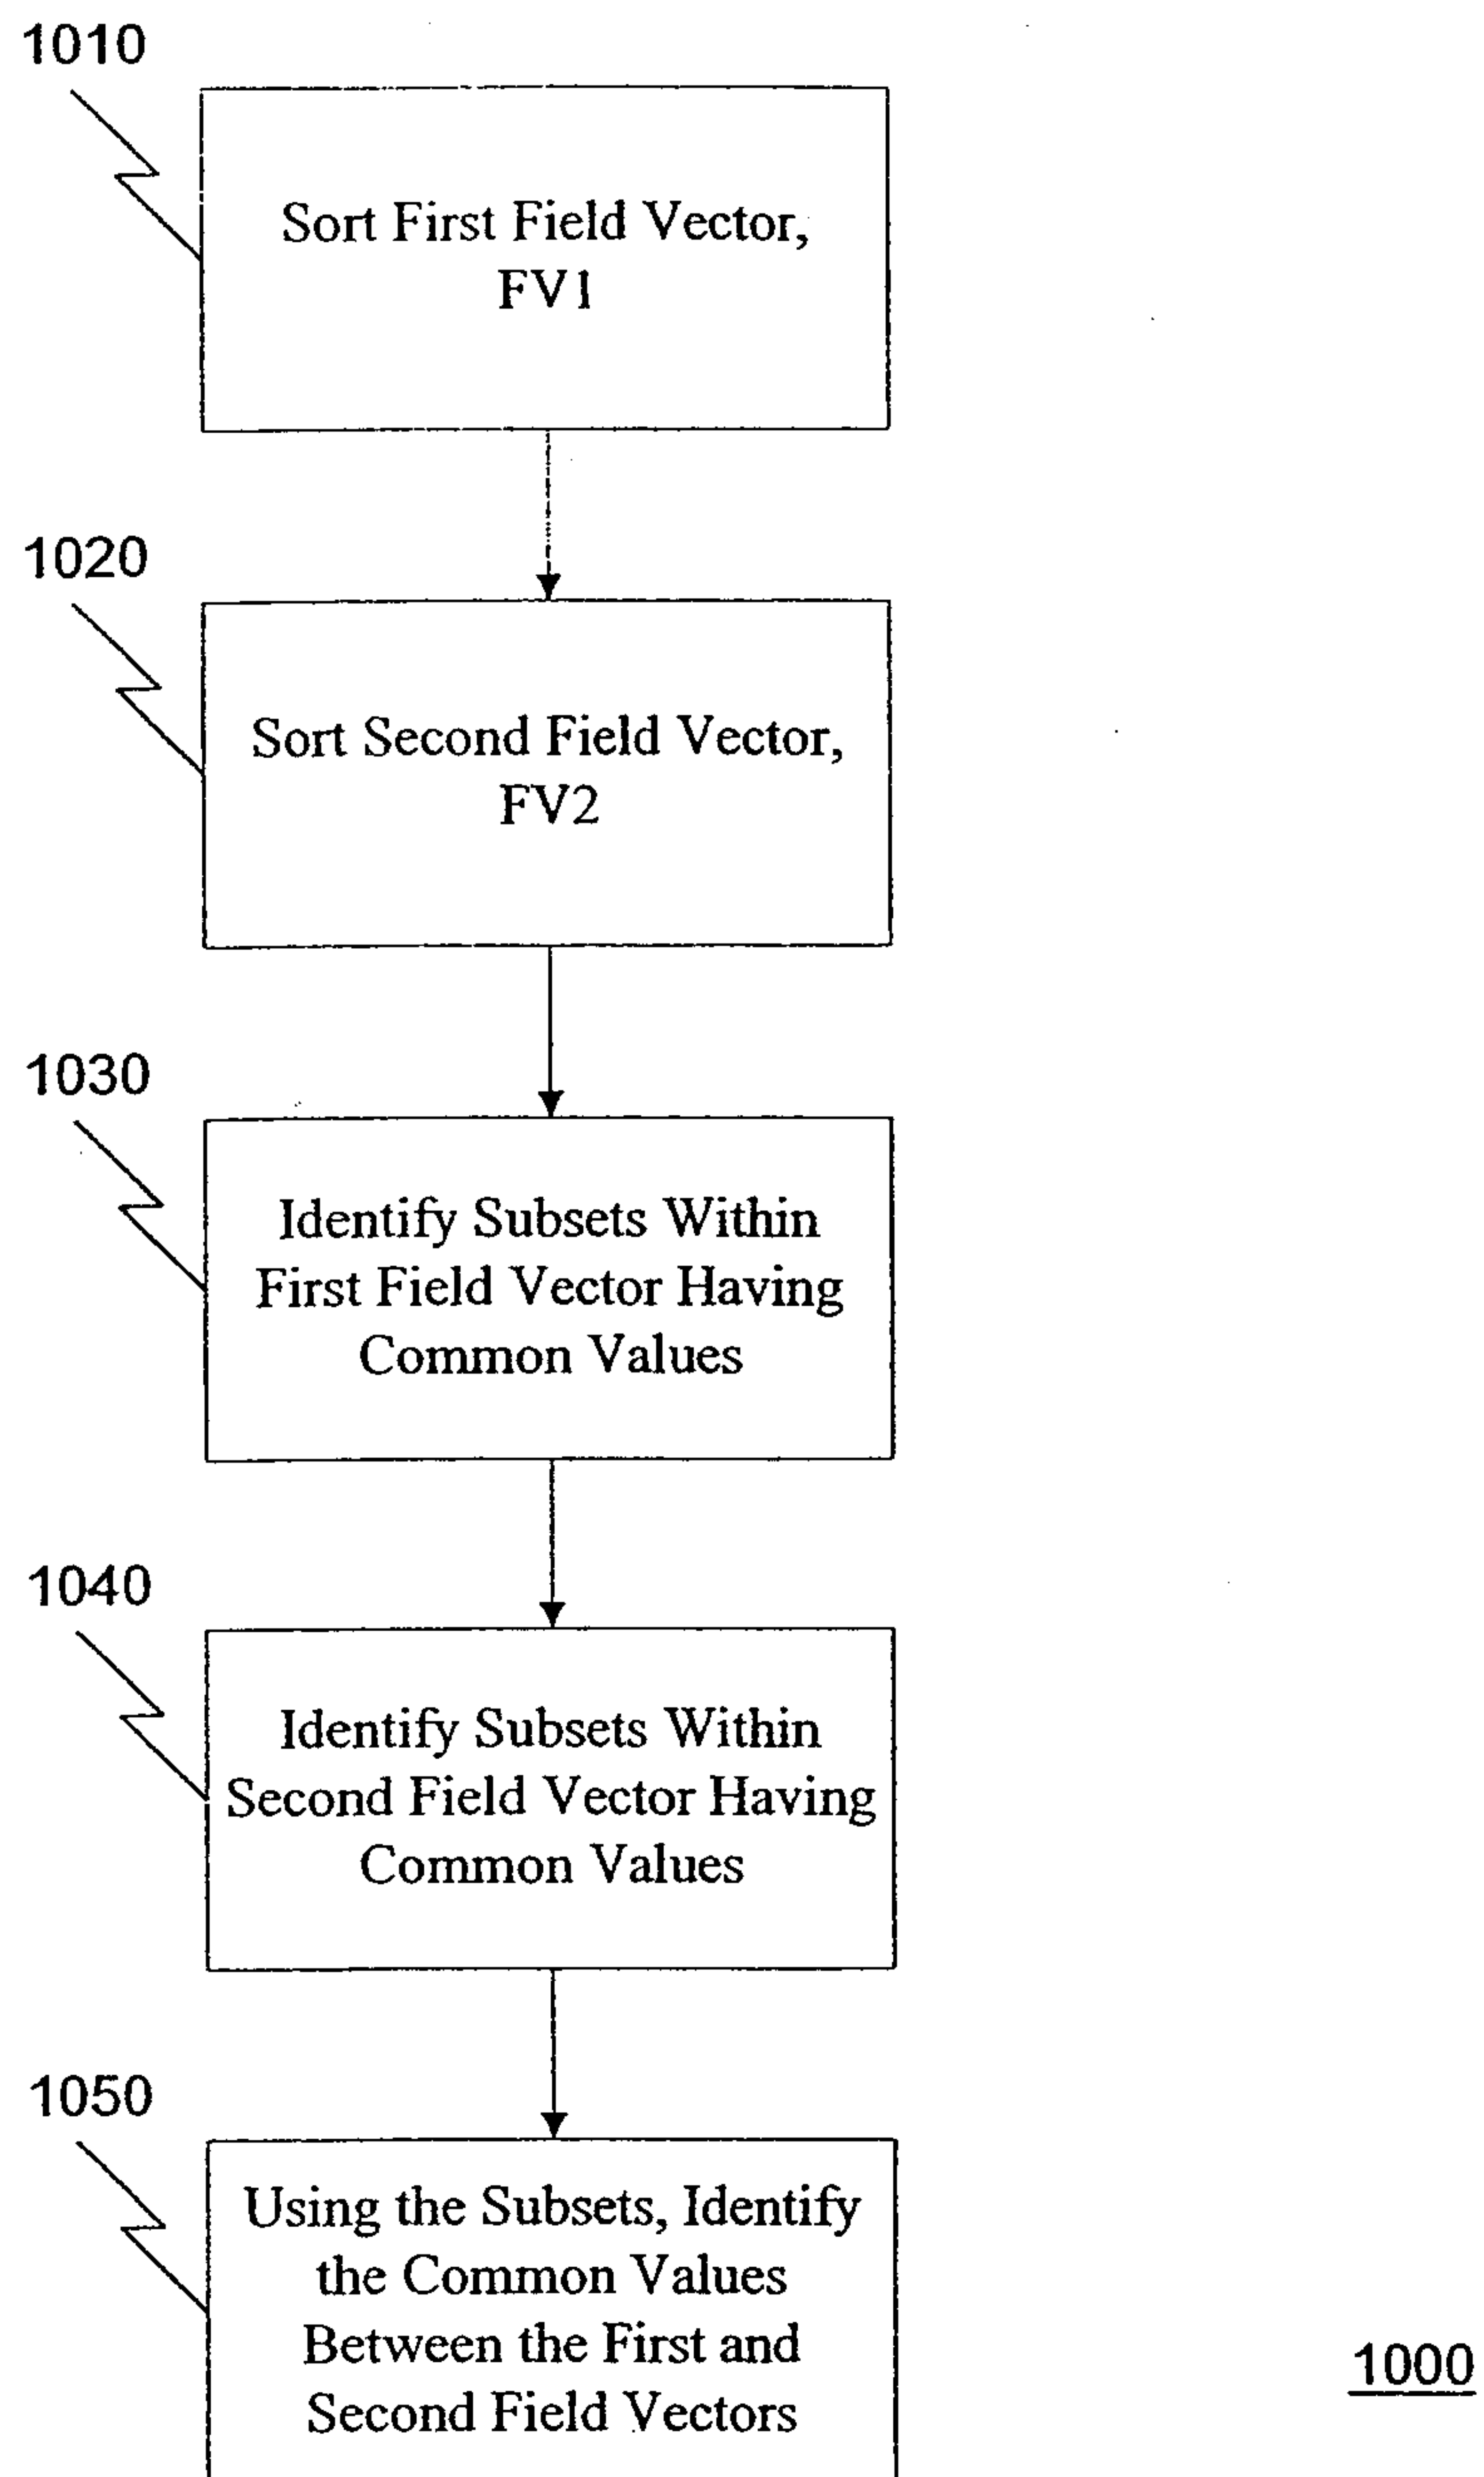


FIGURE 9

10/12**FIGURE 10**

11/12

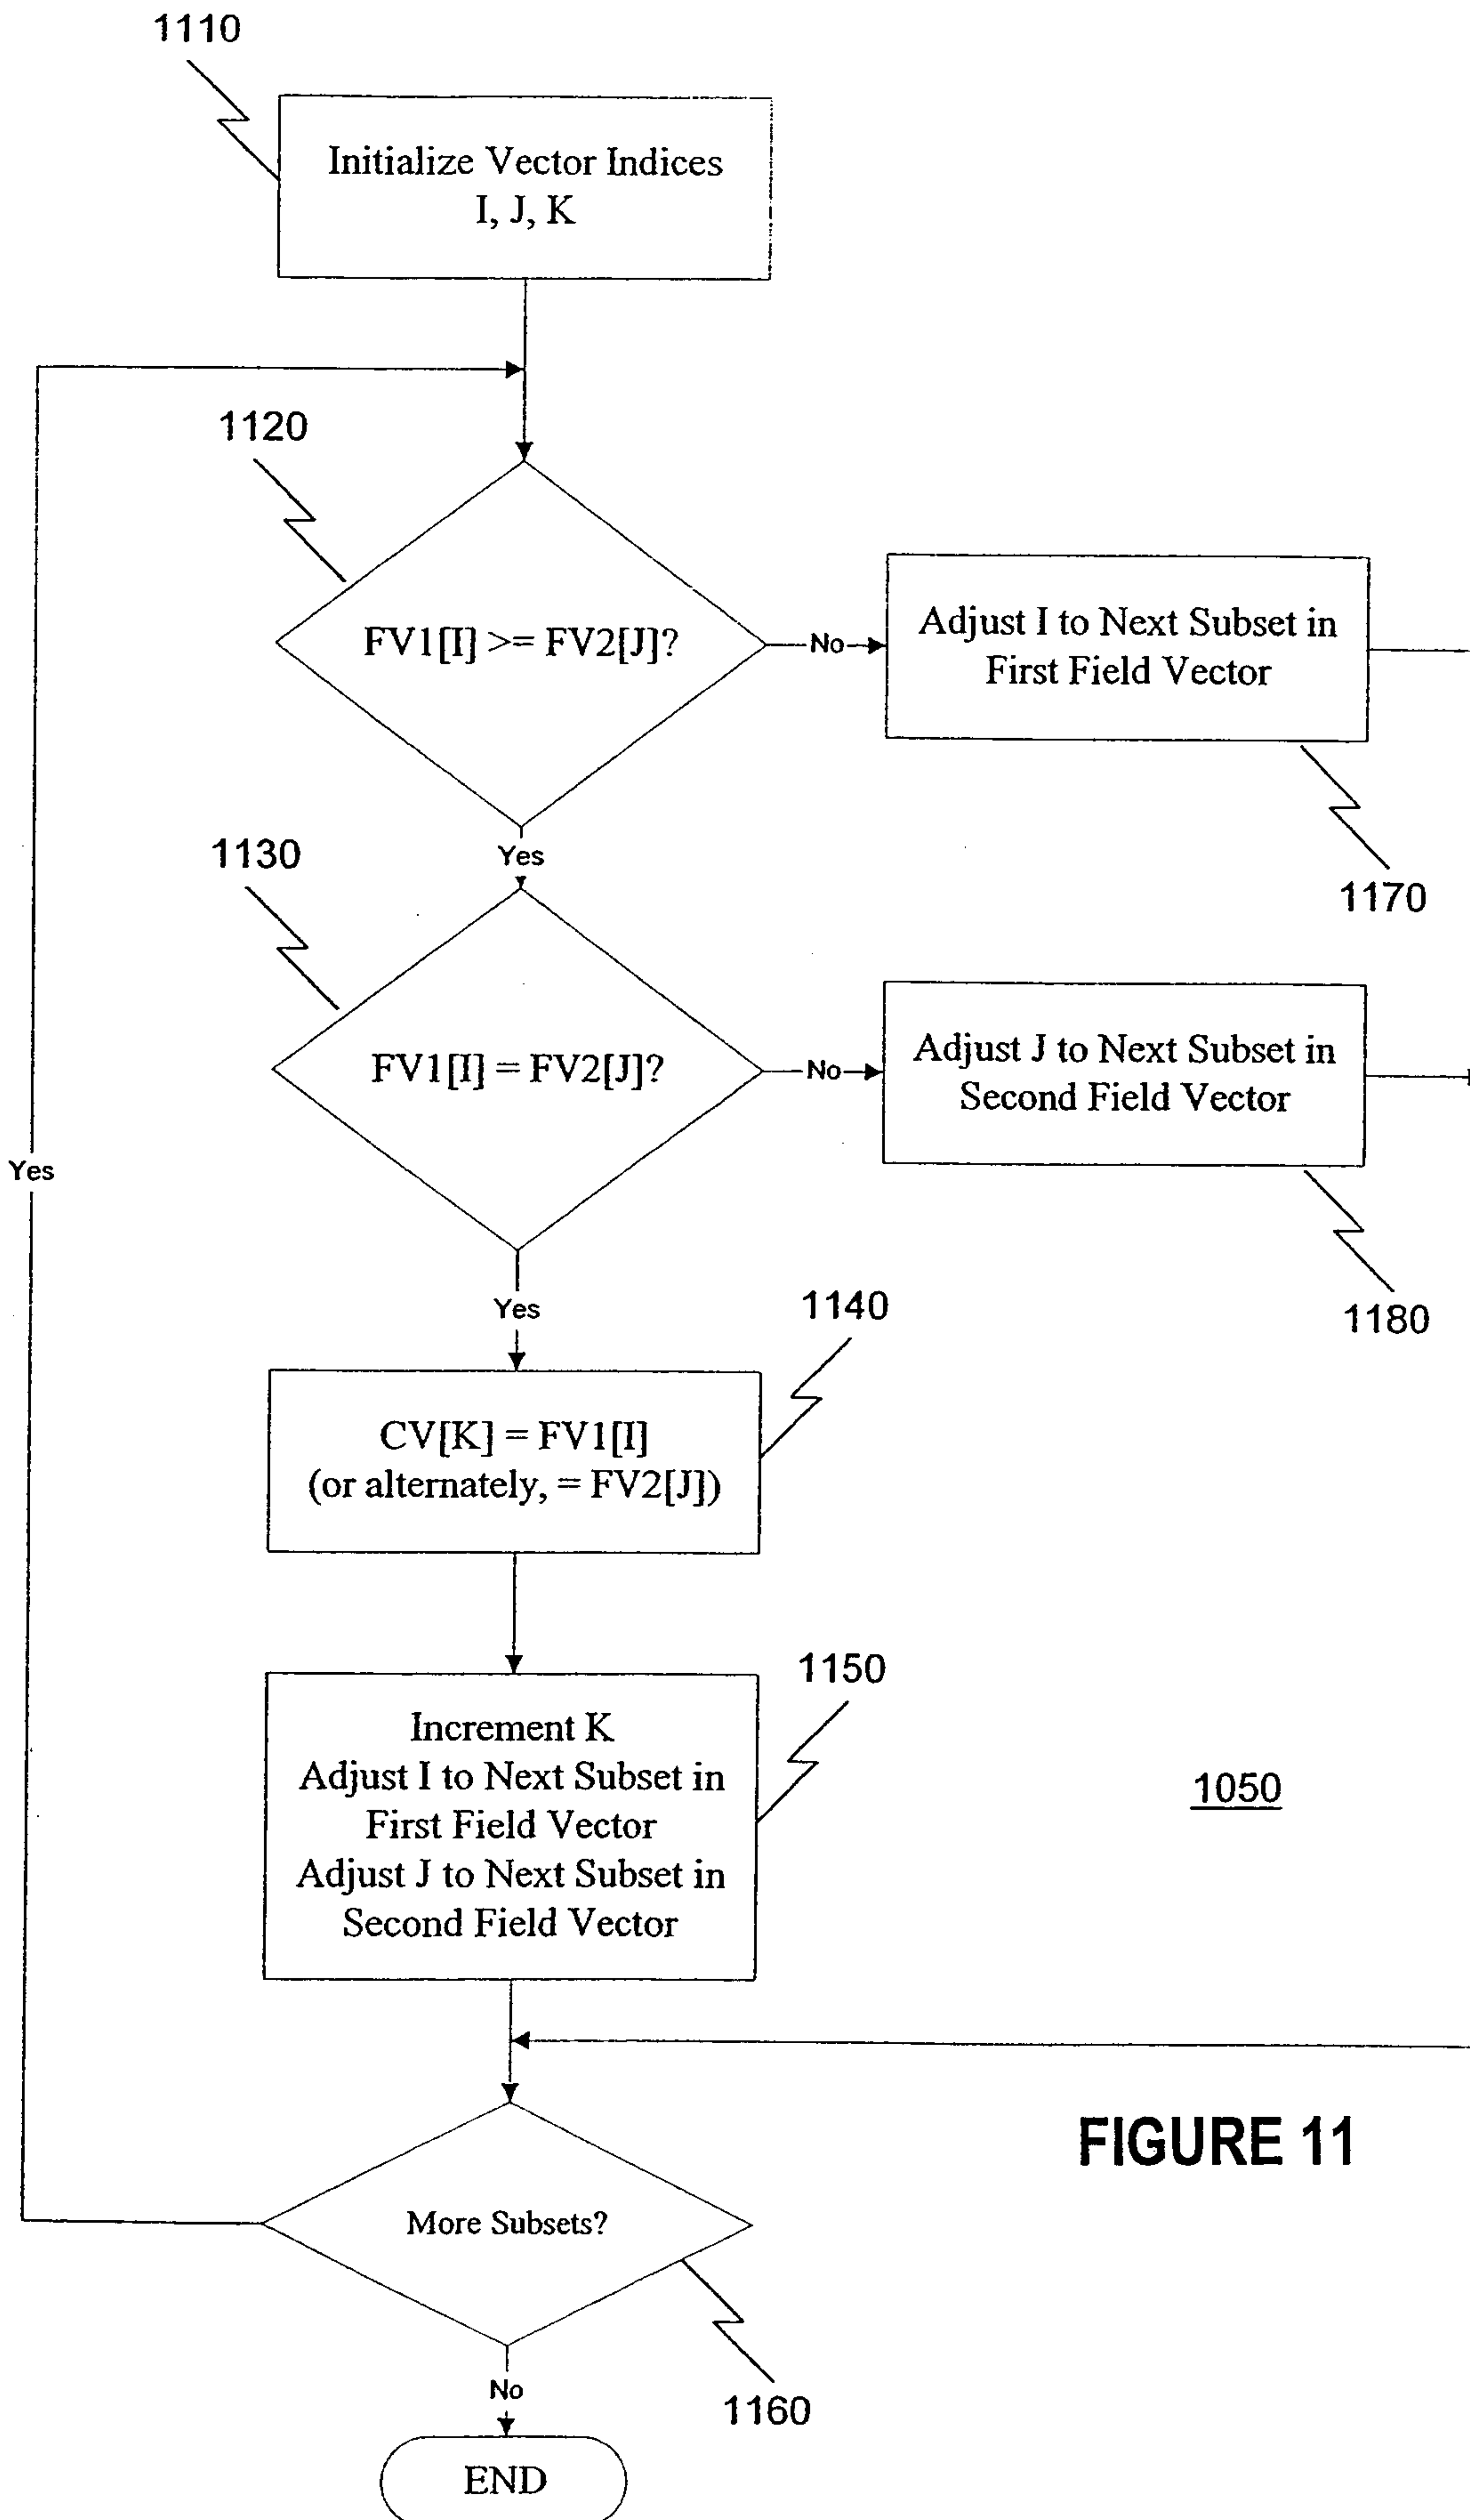


FIGURE 11

12/12

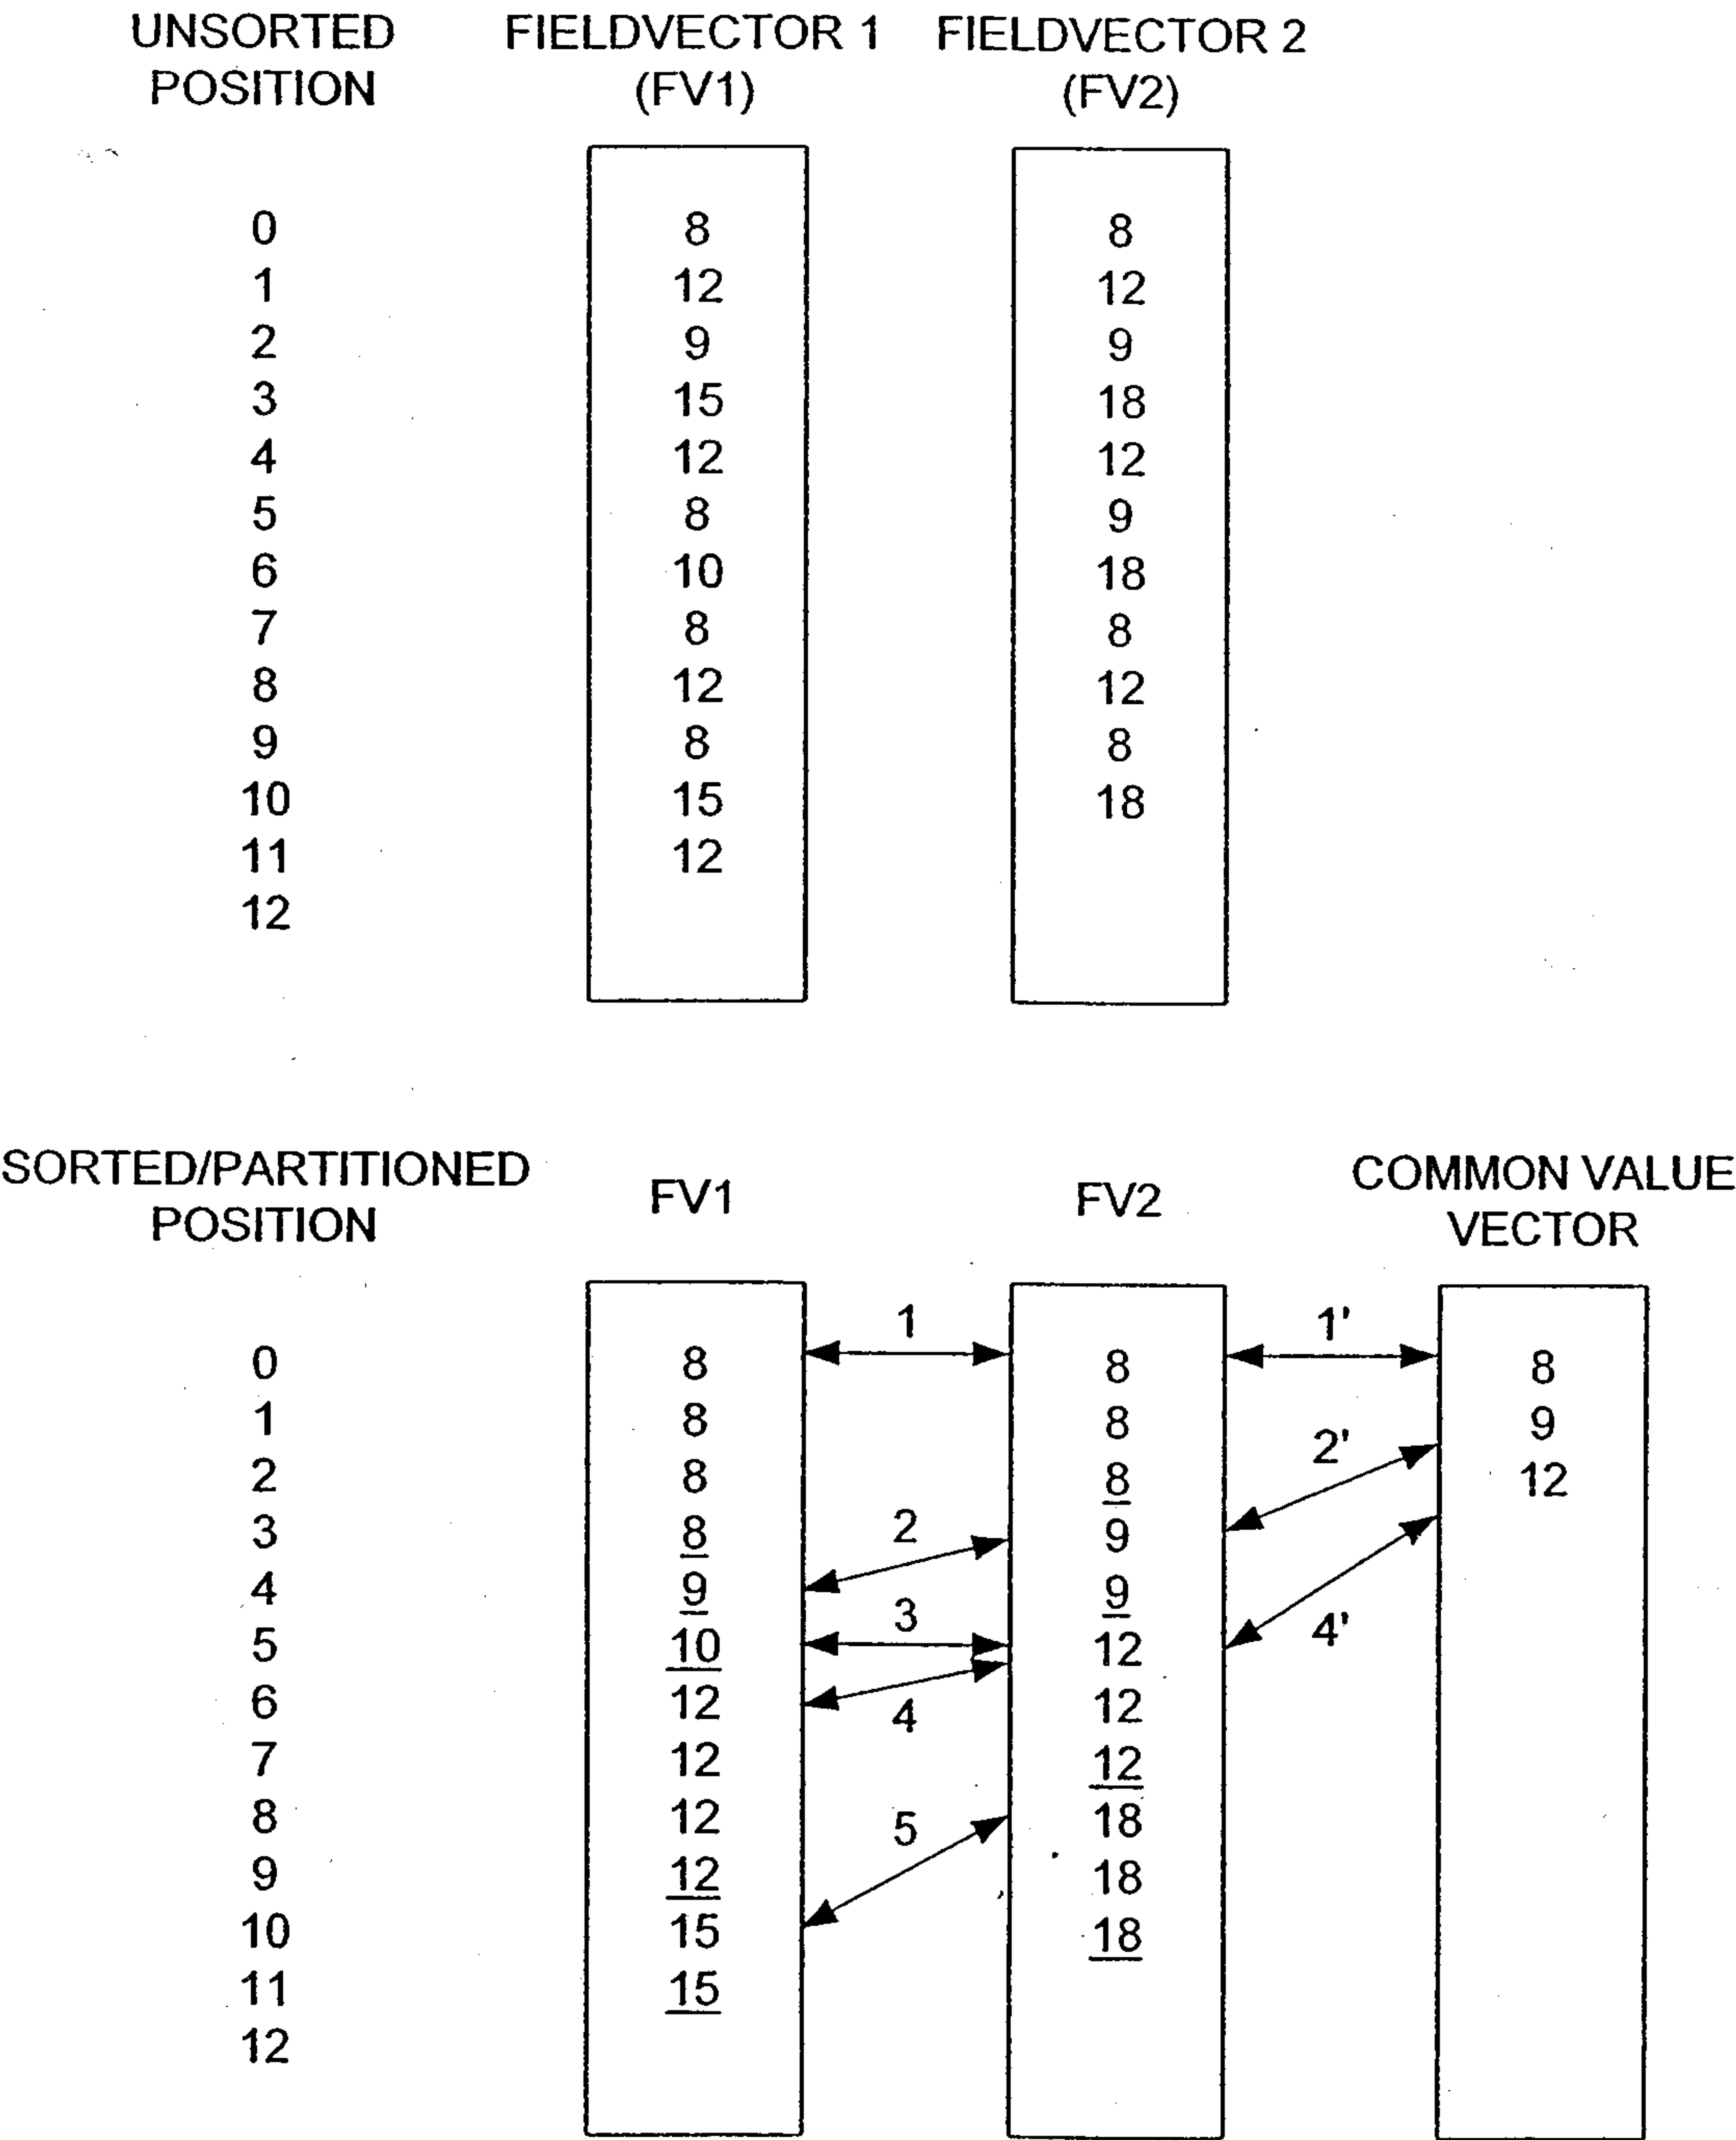


FIGURE 12

