



US009697250B1

(12) **United States Patent**  
**Ward**

(10) **Patent No.:** **US 9,697,250 B1**  
(45) **Date of Patent:** **\*Jul. 4, 2017**

(54) **SYSTEMS AND METHODS FOR HIGH-SPEED SEARCHING AND FILTERING OF LARGE DATASETS**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **MOONSHADOW MOBILE, INC.**, Eugene, OR (US)

5,276,868 A 1/1994 Poole  
5,655,129 A \* 8/1997 Ito ..... G06F 17/30625  
704/10

(72) Inventor: **Roy W. Ward**, Eugene, OR (US)

5,710,915 A 1/1998 McElhiney  
5,721,899 A 2/1998 Namba  
5,737,732 A 4/1998 Gibson et al.  
5,829,004 A 10/1998 Au  
5,978,797 A 11/1999 Yianilos

(73) Assignee: **Moonshadow Mobile, Inc.**, Eugene, OR (US)

(Continued)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 124 days.  
This patent is subject to a terminal disclaimer.

OTHER PUBLICATIONS

U.S. Appl. No. 14/607,038, filed Jan. 27, 2015, Ward et al.  
(Continued)

Primary Examiner — Tyler Torgrimson

(74) Attorney, Agent, or Firm — David S. Alavi

(57) **ABSTRACT**

A binary data file embodies an inline tree data structure storing fields of a hierarchical dataset. The inline tree comprises first-level binary string segments, each comprising substantially contiguous second-level binary string segments, corresponding to subranges of first and second subsets of data fields. Size is reduced by substituting: binary string indices for alphanumeric strings; a data clump index for a set of correlated/anticorrelated strings; field masks for unoccupied data fields. A dedicated conversion program generates the inline tree from conventional database formats, which is read entirely into RAM to be searched/filtered by a dedicated search/filter program. Small size (<2 bytes/field/record) and contiguous arrangement enables searching/filtering of >10<sup>6</sup> records (>100 data fields) in <500 nano-seconds/record/core. Recursive subdivision of selection field ranges can guide searches that include those selection fields. One example includes geographic searching/filtering of records that include latitude and longitude fields.

(21) Appl. No.: **14/678,982**

(22) Filed: **Apr. 4, 2015**

**Related U.S. Application Data**

(63) Continuation of application No. 13/326,326, filed on Dec. 15, 2011, now Pat. No. 9,002,859.

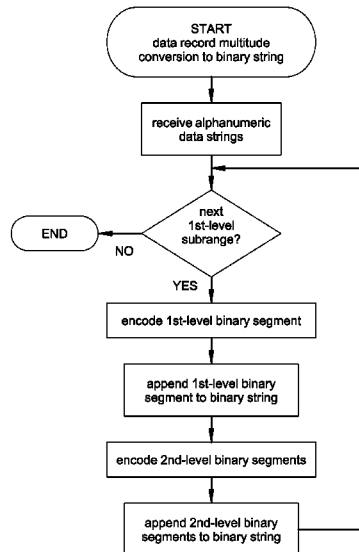
(60) Provisional application No. 61/424,063, filed on Dec. 17, 2010.

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

(52) **U.S. Cl.**  
CPC .. **G06F 17/30424** (2013.01); **G06F 17/30327** (2013.01)

(58) **Field of Classification Search**  
None  
See application file for complete search history.

**31 Claims, 13 Drawing Sheets**



(56)

References Cited

U.S. PATENT DOCUMENTS

6,012,057 A 1/2000 Mayer et al.  
 6,175,835 B1 1/2001 Shadmon  
 6,341,280 B1 1/2002 Glass  
 6,374,251 B1 4/2002 Fayyad et al.  
 6,499,033 B1 12/2002 Vagnozzi  
 6,560,610 B1 5/2003 Eatherton et al.  
 6,581,058 B1 6/2003 Fayyad et al.  
 6,675,171 B2 1/2004 Tikkanen et al.  
 6,721,723 B1 4/2004 Gibson et al.  
 6,728,732 B1 4/2004 Eatherton et al.  
 7,827,218 B1 11/2010 Mittal  
 8,977,656 B2 3/2015 Ward  
 8,990,204 B1 3/2015 Ward  
 9,002,859 B1 4/2015 Ward  
 2001/0042186 A1 11/2001 Iivonen et al.  
 2002/0035660 A1 3/2002 Tikkanen et al.  
 2002/0087510 A1 7/2002 Weinberg et al.  
 2003/0018620 A1 1/2003 Vishnubhotla  
 2003/0061189 A1 3/2003 Baskins et al.  
 2003/0061227 A1\* 3/2003 Baskins ..... G06F 17/30961  
 2003/0105745 A1 6/2003 Davidson et al.  
 2004/0059725 A1 3/2004 Sharangpani et al.  
 2004/0107184 A1 6/2004 Longshaw  
 2004/0111439 A1 6/2004 Richardson et al.  
 2004/0117396 A1 6/2004 Avadhanam et al.  
 2004/0133590 A1 7/2004 Henderson et al.  
 2004/0193619 A1 9/2004 Venkatachary et al.  
 2004/0205517 A1 10/2004 Lampert et al.  
 2004/0243576 A1 12/2004 Shrivastava et al.  
 2004/0267710 A1 12/2004 Cotarmanac'h et al.  
 2005/0027744 A1 2/2005 Avadhanam et al.  
 2005/0055351 A1 3/2005 Barton et al.  
 2005/0091223 A1 4/2005 Shaw et al.  
 2005/0171959 A1 8/2005 Deforche et al.  
 2005/0240604 A1 10/2005 Corl et al.  
 2005/0251331 A1 11/2005 Kreft  
 2006/0271540 A1 11/2006 Williams  
 2006/0282457 A1 12/2006 Williams  
 2006/0288024 A1 12/2006 Braica  
 2007/0192548 A1 8/2007 Williams  
 2007/0255748 A1 11/2007 Ferragina et al.  
 2008/0019317 A1 1/2008 Vellanki et al.  
 2008/0086464 A1 4/2008 Enga  
 2009/0077113 A1 3/2009 Fidaali et al.

2009/0138790 A1 5/2009 Larcheveque et al.  
 2010/0011125 A1 1/2010 Yang et al.  
 2010/0023515 A1 1/2010 Marx  
 2010/0057792 A1 3/2010 Ylonen  
 2010/0085223 A1 4/2010 Hendrickson  
 2010/0127902 A1 5/2010 Schneider  
 2010/0146004 A1 6/2010 Sim-Tang  
 2010/0185609 A1 7/2010 Wright et al.  
 2010/0281082 A1 11/2010 Ylonen  
 2010/0332561 A1 12/2010 Prahlad et al.  
 2011/0016153 A1 1/2011 Atta et al.  
 2012/0005239 A1 1/2012 Nevin  
 2012/0016908 A1 1/2012 Leung et al.

OTHER PUBLICATIONS

U.S. Appl. No. 13/733,890, filed Jan. 4, 2013, Ward.  
 U.S. Appl. No. 13/743,613, filed Jan. 17, 2013, Ward.  
 U.S. Appl. No. 14/489,430, filed Sep. 17, 2014, Ward.  
 Office action dated Oct. 31, 2014 in co-owned U.S. Appl. No. 13/347,646.  
 Notice of Allowance dated Oct. 12, 2015 in co-owned U.S. Appl. No. 13/347,646.  
 Office action dated Jun. 12, 2014 in co-owned U.S. Appl. No. 13/733,890.  
 Notice of Allowance dated Jul. 16, 2014 in co-owned U.S. Appl. No. 13/740,278.  
 Office action dated Sep. 12, 2014 in co-owned U.S. Appl. No. 13/743,613.  
 International Search Report and Written Opinion dated Aug. 27, 2012 in co-owned App No. PCT/US2012/020841.  
 Gabow, Harold N. Data structures for weighted matching and nearest common ancestors with linking. U. of Colorado, Boulder, Dept of Computer Science, 1990.  
 Acharya et al; Adaptive Algorithms for Cache-efficient Trie Search; ALENEX'99, LNCS 1619, pp. 296-311, Goodrich & McGeoch (Eds); Springer-Verlag (1999).  
 Bagwell, Fast and Space Efficient Trie Searches, 2000.  
 Office action dated Mar. 25, 2013 in parent U.S. Appl. No. 13/326,326.  
 Office action dated Oct. 23, 2013 in parent U.S. Appl. No. 13/326,326.  
 Notice of Allowance dated Feb. 27, 2015 in parent U.S. Appl. No. 13/326,326.

\* cited by examiner

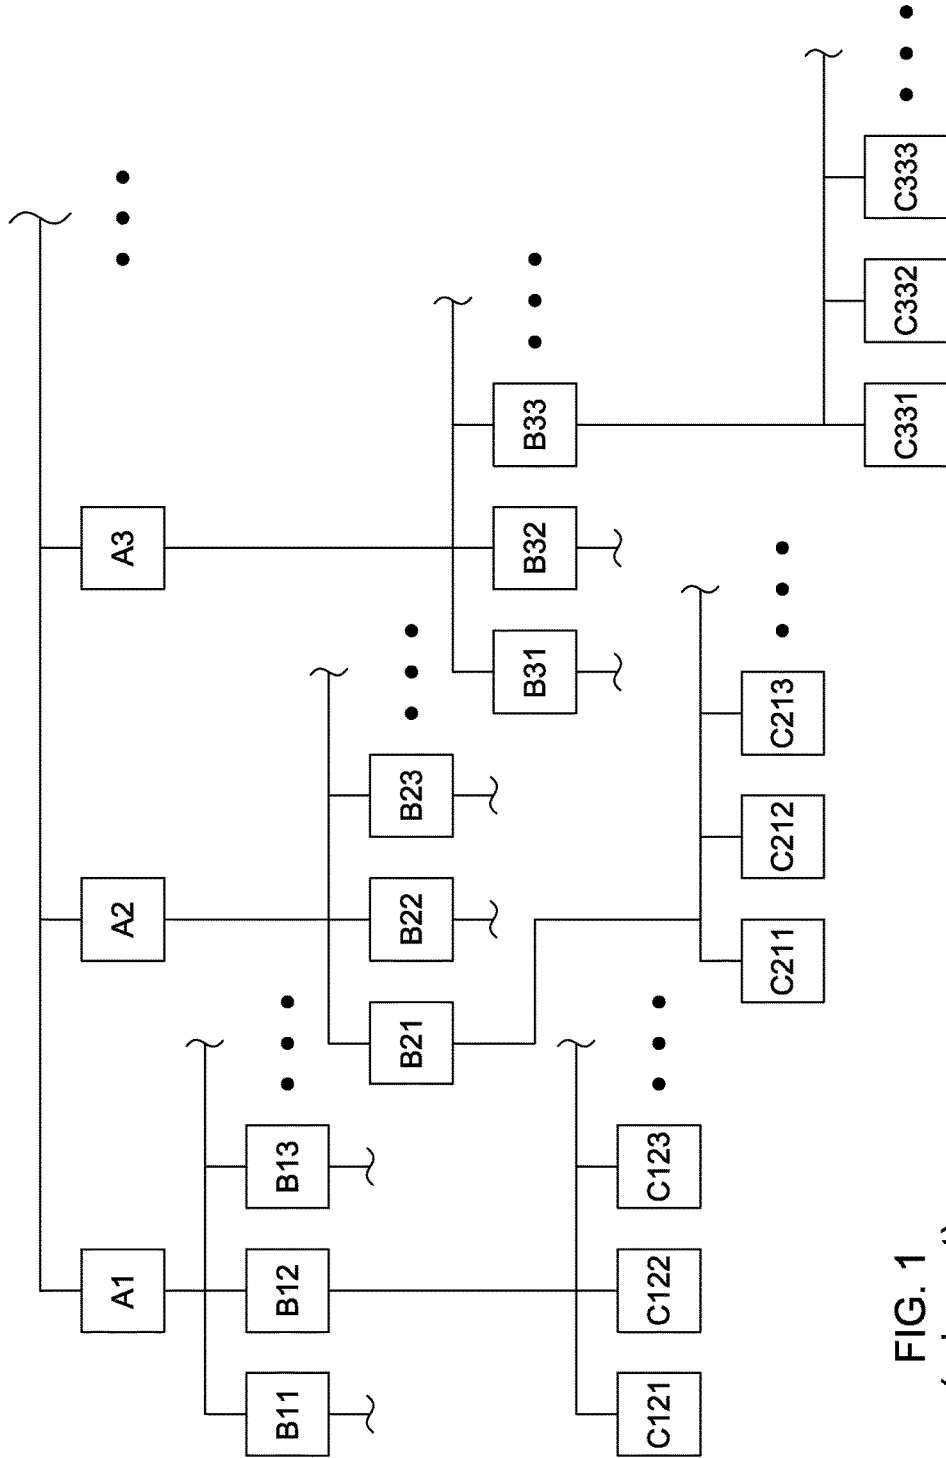


FIG. 1  
(prior art)




FIG. 2  
(prior art)


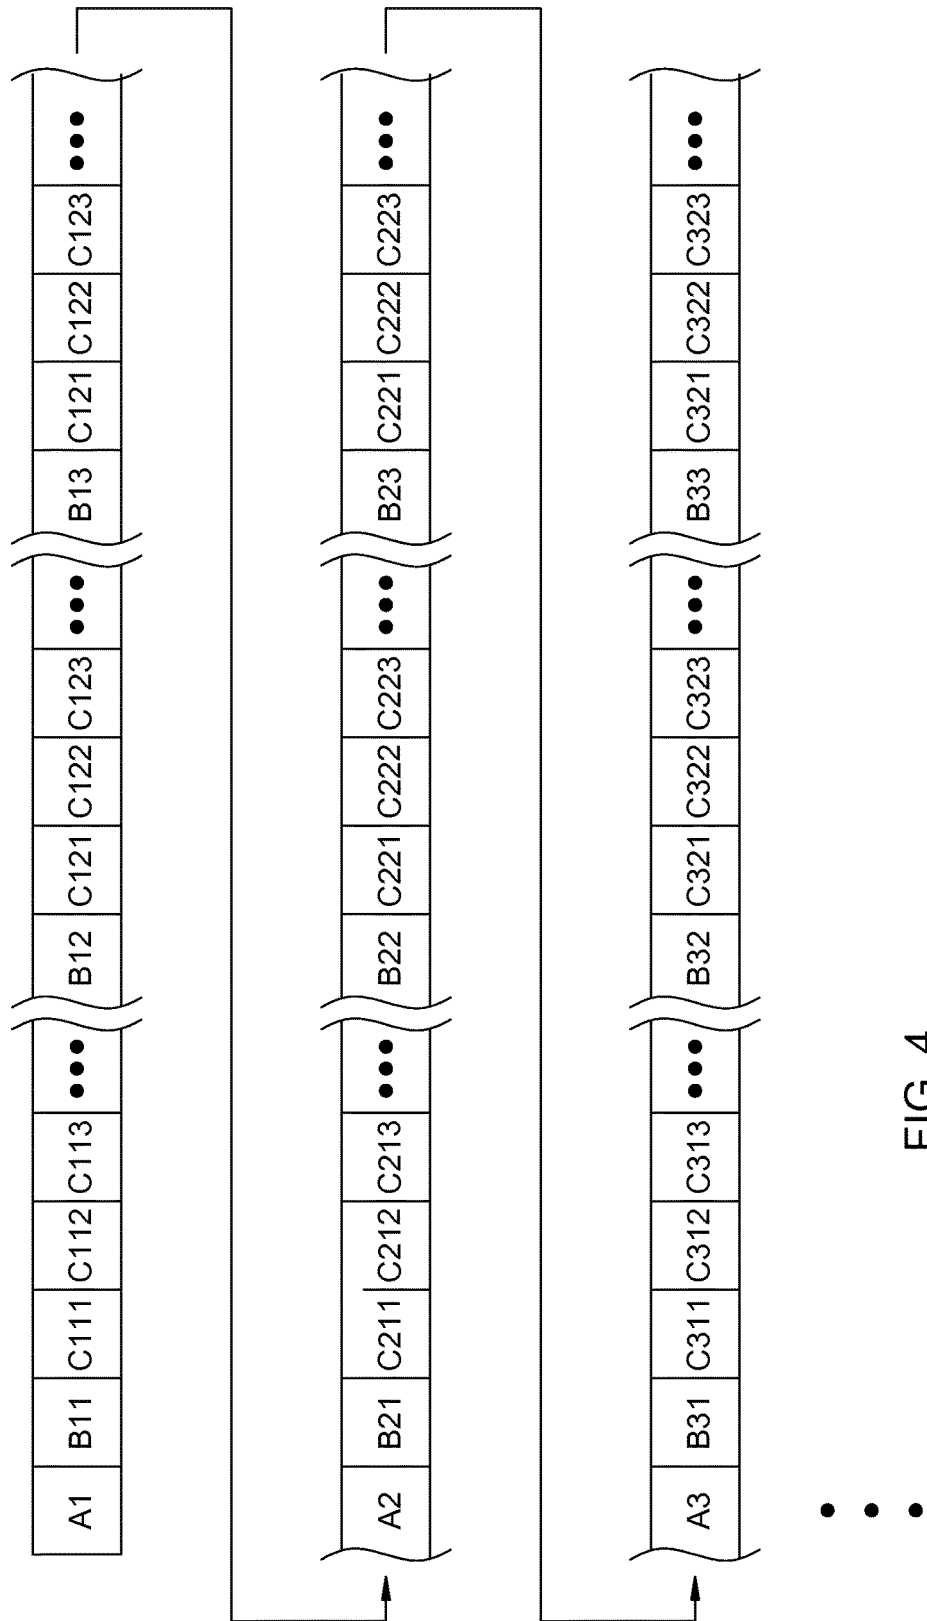


FIG. 3  
(prior art)



<u>aux table</u>	
...001	Street
...010	Boulevard
...011	Avenue
...100	Lane
...101	Road
...	...

FIG. 5B

<u>string table</u>	
...	...
...000	Smith
...001	Jones
...010	Doe
...	...
...010	John
...011	Mary
...	...
...101	Oak
...110	Main
...111	Elm
...	...
...111	Portland
...000	Eugene
...001	Corvallis
...	...

FIG. 5A

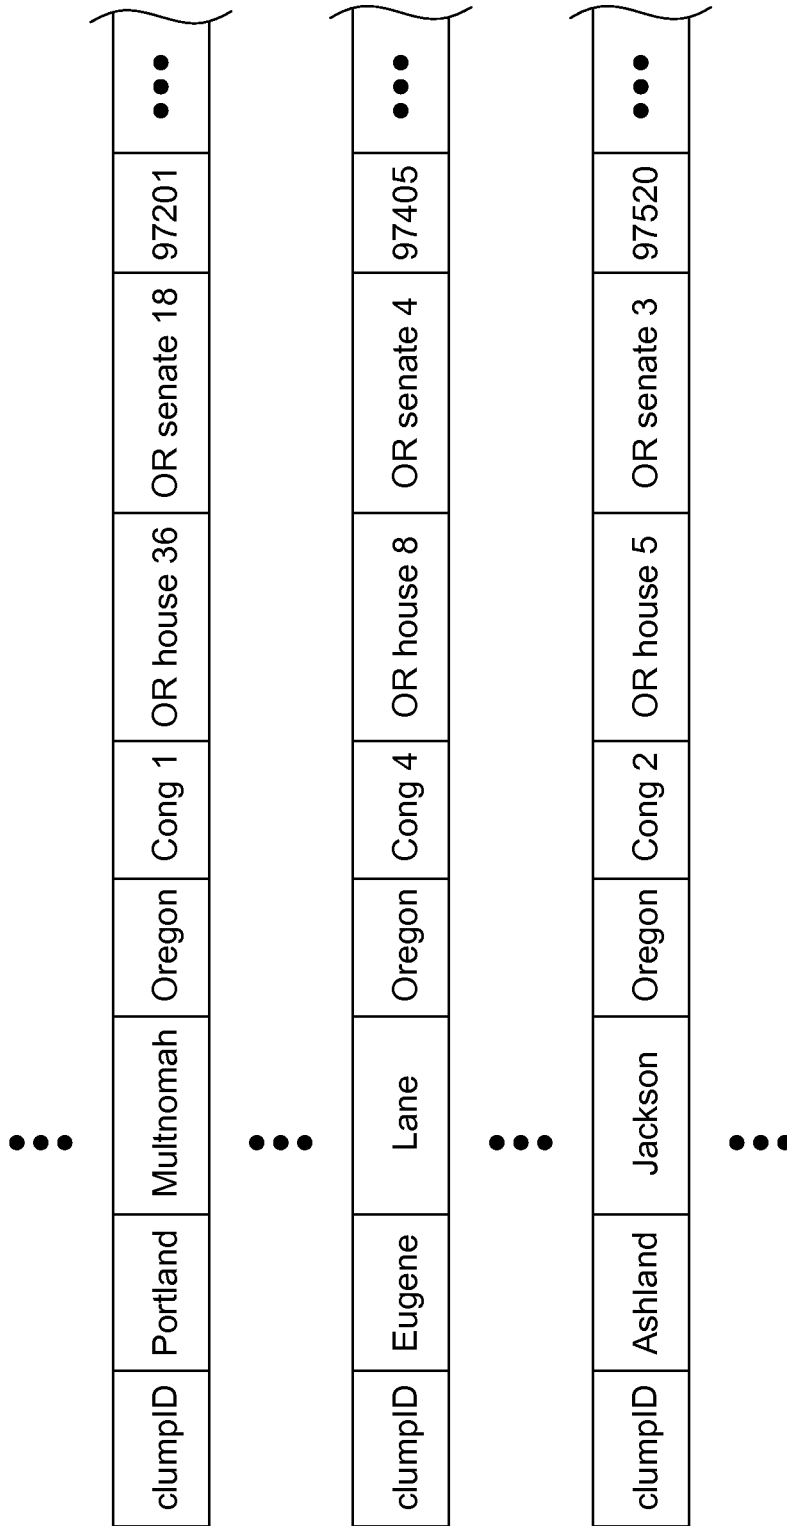


FIG. 5C



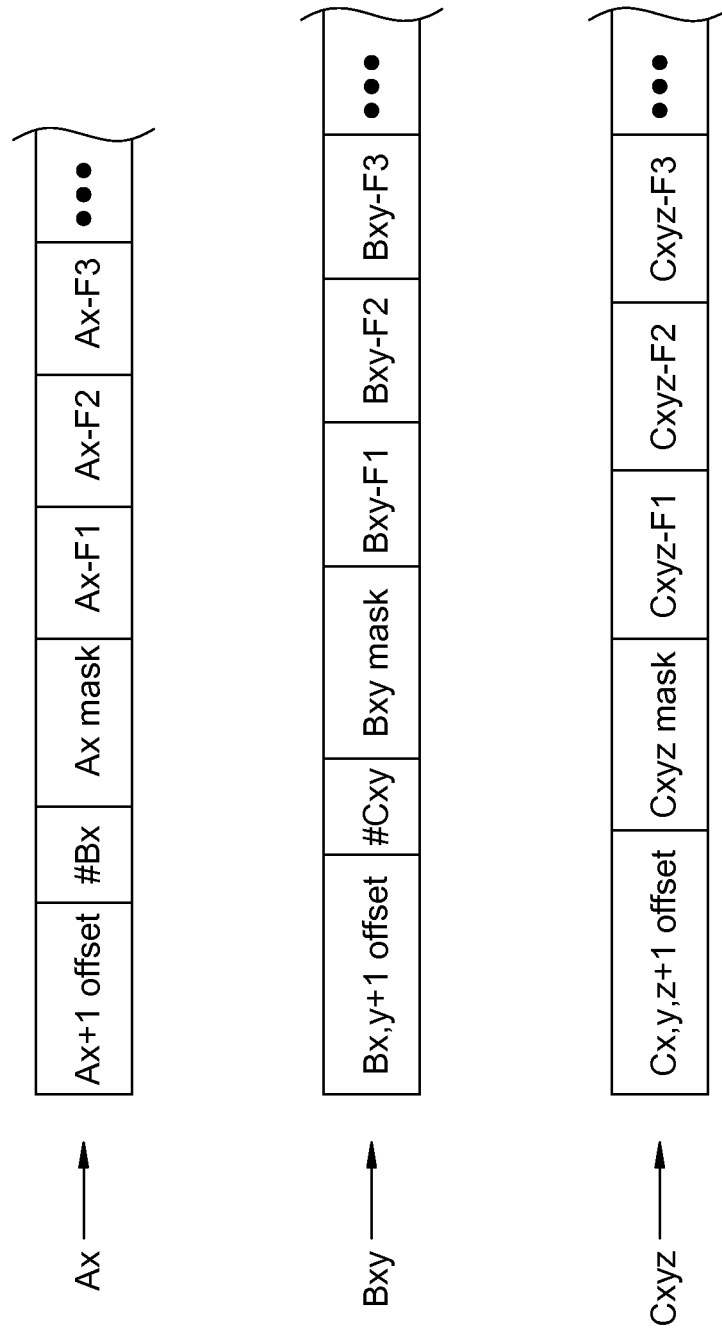


FIG. 7



FIG. 8

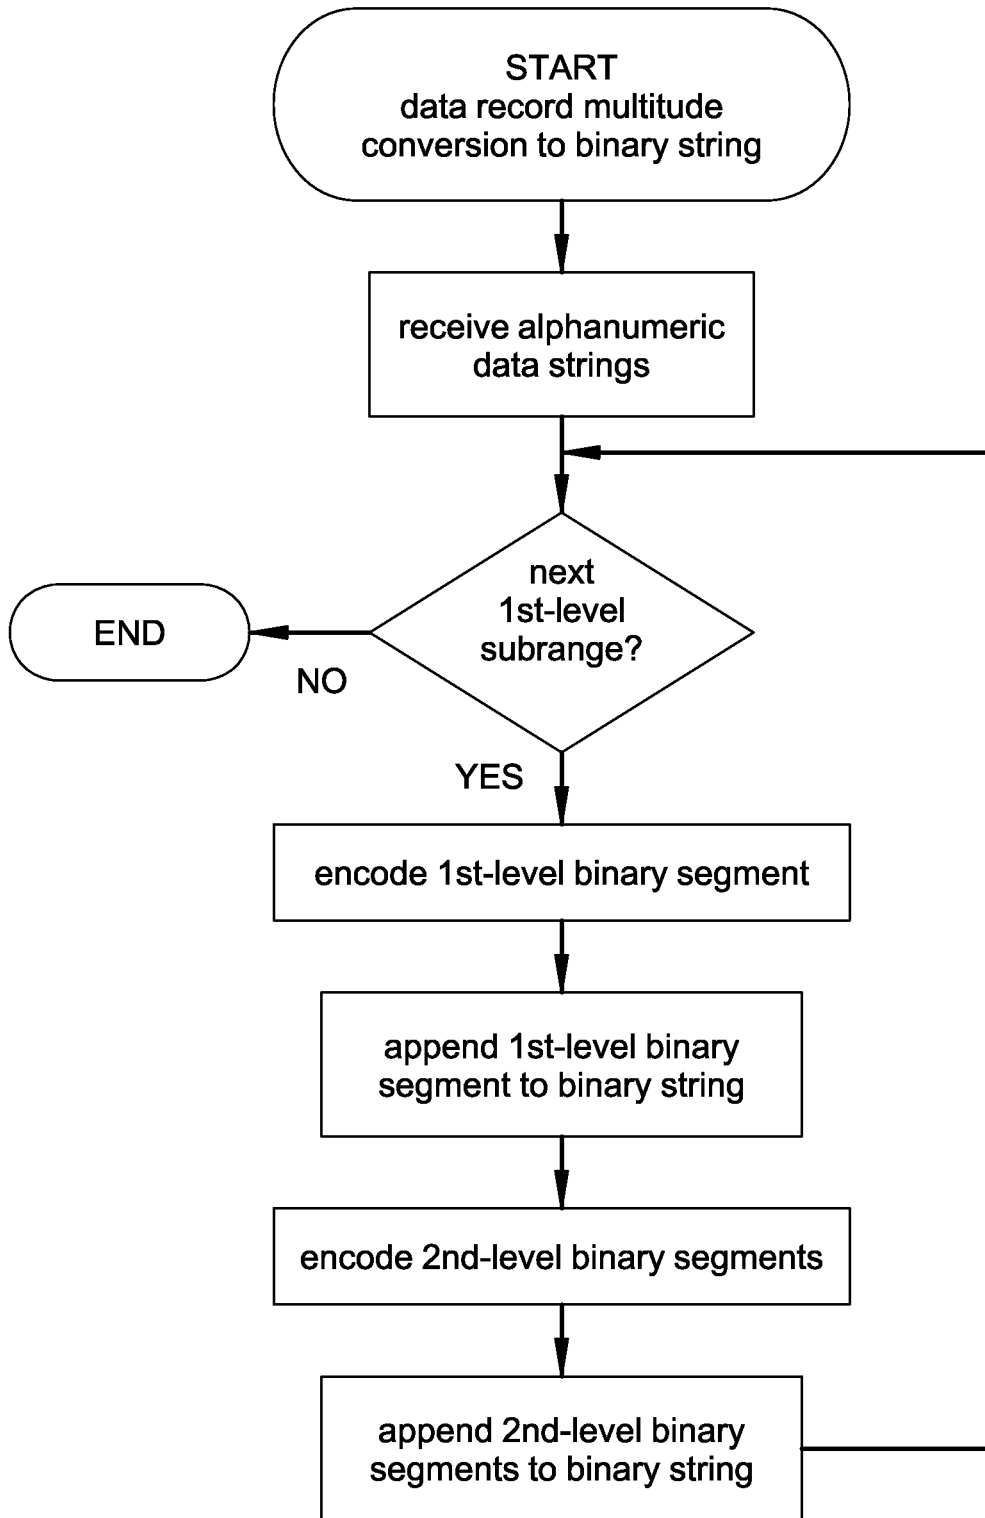


FIG. 9

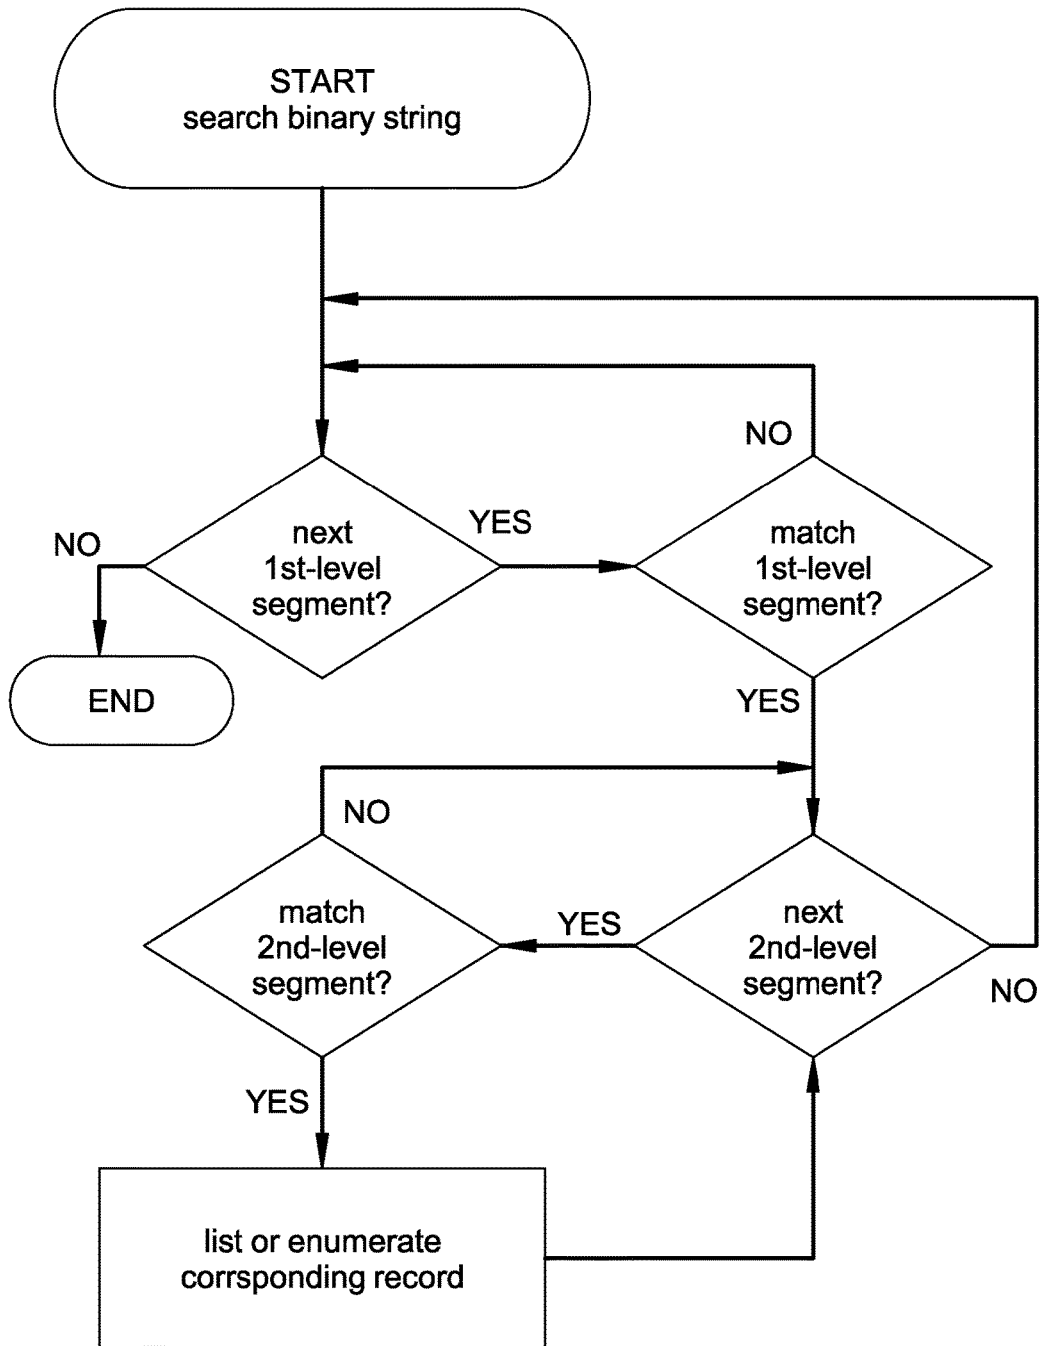


FIG. 10

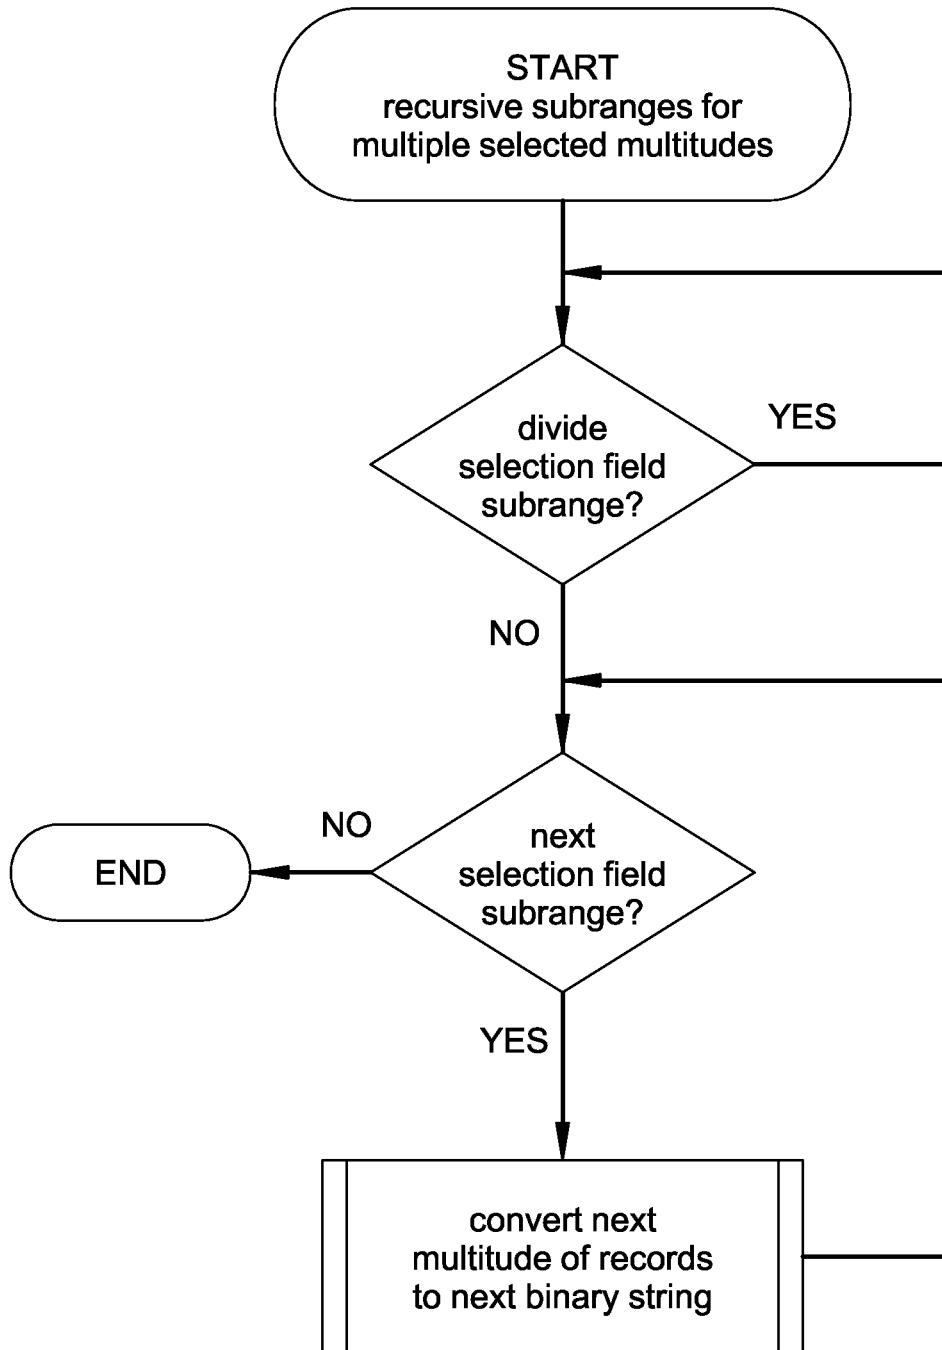


FIG. 11

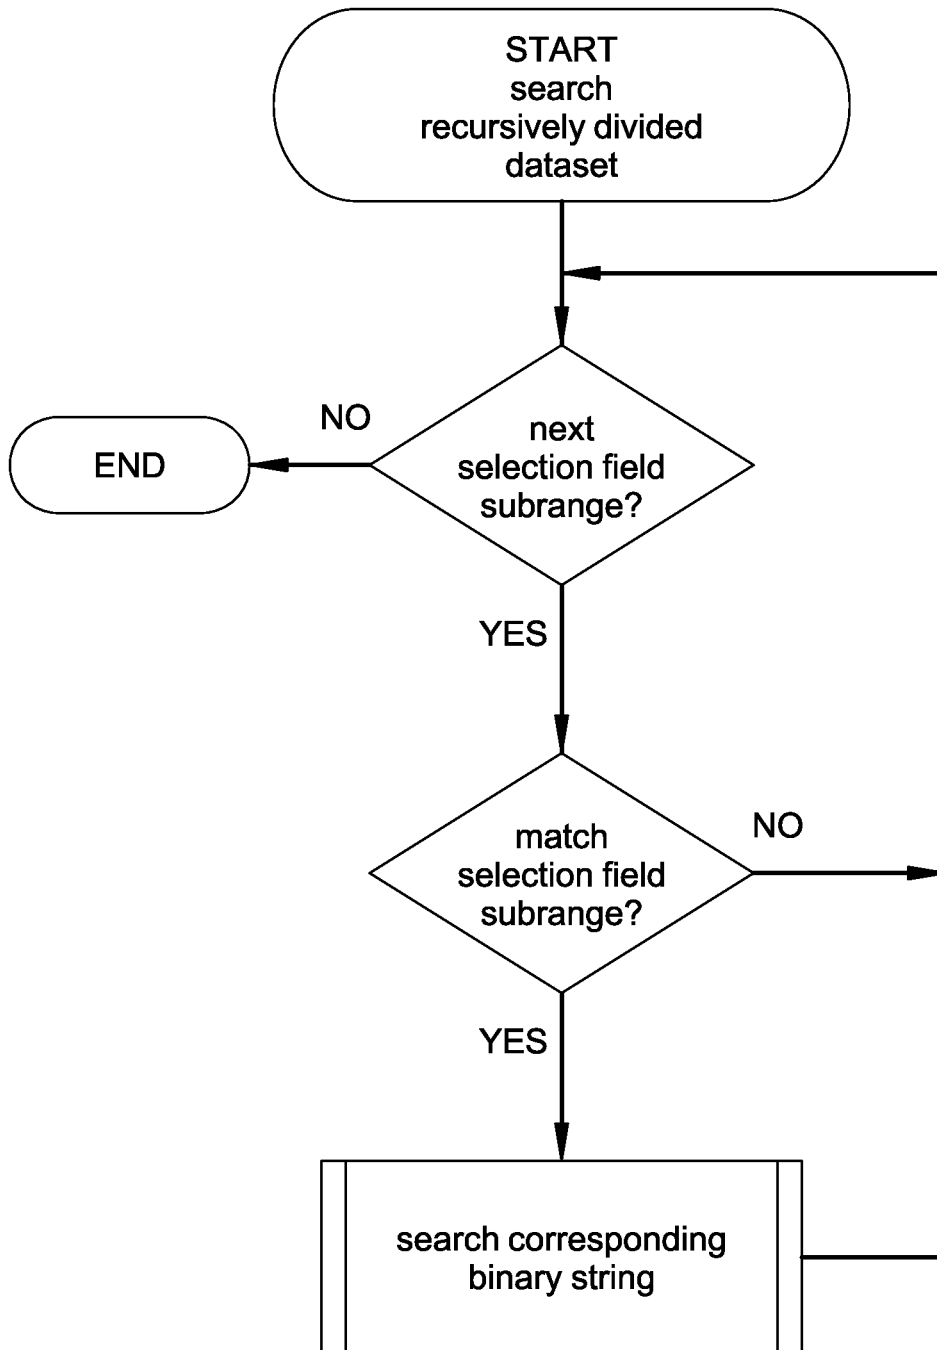


FIG. 12

1

## SYSTEMS AND METHODS FOR HIGH-SPEED SEARCHING AND FILTERING OF LARGE DATASETS

### BENEFIT CLAIMS TO RELATED APPLICATIONS

This application, is a continuation of non-provisional application Ser. No. 13/326,326 filed Dec. 15, 2011 in the name of Roy W. Ward (now U.S. Pat. No. 9,002,859), which claims benefit of provisional App. No. 61/424,063 filed Dec. 17, 2010 in the name of Roy W. Ward, said non-provisional and provisional applications being hereby incorporated by reference as if fully set forth herein.

### BACKGROUND

The field of the present invention relates to electronic data search and retrieval. In particular, systems and methods are disclosed herein for high-speed searching and filtering of large datasets.

Many situations exist in which very large amounts of data are generated or collected (e.g.,  $10^4$ ,  $10^5$ ,  $10^6$ , or more data records, each comprising multiple data fields). For data in a dataset to be of any practical use, indicia representing the dataset are stored according to a data structure arranged so that particular pieces of information can be located and retrieved from the dataset. In the pre-digital past, such data structures often comprised printed alphanumeric indicia on suitable media (often including an accompanying printed index), and data search and retrieval were manual functions performed by humans. The introduction of electronic data storage and search capabilities around the middle of the last century revolutionized the ability to store large datasets, and to search for and retrieve specific information from those stored datasets.

Today, alphanumeric indicia representative of a dataset are typically stored according to digital, electronic data structures such as an electronic spreadsheet or an electronic relational database. A spreadsheet (also referred to as a flat file database) can be thought of as a single table with rows and columns, with each row corresponding to a specific data record, and with each column corresponding to a specific data field of that data record. In a simple example (one that will be used repeatedly within the instant specification), each data record can correspond to a registered voter in a dataset of all registered voters in a particular state, e.g., Oregon. The data fields in each data record can include, e.g., last name, first name, middle name or initial, age, gender, marital status, race, ethnicity, religion, other demographic information, street address (likely divided into multiple data fields for street number, street name, and so on), city, state, zip code, party affiliation, voting history, county, U.S. house district, state senate/house district, school district, other administrative districts, and so on.

A relational database typically comprises multiple tables, each comprising multiple records with multiple fields, and relations defined among various fields in differing tables. In the registered voter example given above, a "voter" table might include voter records with name and demographic information in corresponding fields, and an "address" table might include address records that includes street address and district information in corresponding fields. A field in the voter table can include a pointer to the corresponding address in the address table, defining a one-to-many relationship between each address and one or more correspond-

2

ing voters. Other tables and relationships can be defined (including many-to-many relationships and so-called pivot tables to define them).

Electronic spreadsheets and electronic relational databases have become standard methods for storage of digital datasets. They offer nearly unlimited flexibility in arranging the data, for updating the data, for adding new data, and for sorting, searching, filtering, or retrieving data. However, it has been observed that for a very large dataset (e.g.,  $>10^6$  or more records, or even as few as  $>10^4$  or  $>10^5$  records), spreadsheets and databases tend to become unwieldy to store, access, and search. In particular, search and retrieval of information from such a large electronic dataset can become so slow as to render it essentially useless for certain data retrieval applications.

It would be desirable to provide systems and methods that enable high-speed search and retrieval of information from large electronic datasets that substantially exceed search and retrieval speeds from conventional electronic data structures (e.g., conventional spreadsheets and databases), so as to enable data search and retrieval applications that are too slow for practicable use with those conventional data structures.

### SUMMARY

One or more binary data files embodies an inline tree data structure for storing data fields of hierarchically organized data records in a dataset. The inline tree data structure comprises first-level binary string segments, each comprising one or more constituent, substantially contiguous second-level binary string segments. The first- and second-level binary string segments are arranged according to data sub-ranges of corresponding first and second subsets of data fields in the data records. The size of the inline tree data structure is reduced by substituting (i) binary string indices for alphanumeric strings in some data fields, (ii) a data clump index for a set of sufficiently correlated or anticorrelated alphanumeric strings in multiple other data fields, or (iii) field masks for unoccupied data fields. The resulting file size can be less than about 2 bytes per field per record (e.g., a dataset of one million records having 100 fields each can be stored in less than about 200 MB).

A dedicated, specifically adapted conversion program generates the inline tree data structure from data records in a more conventional database format. The inline tree data structure can be stored on any computer-readable medium, and is read entirely into RAM to be searched (with or without filtering on one or more filter data fields). A dedicated, specifically adapted search and filter program is employed, which can list or enumerate the retrieved data records. The small size and contiguous arrangement of the inline tree data structure enables searching and filtering of a million or more data records (each including over 100 data fields) in less than about 500 nanoseconds per record per processor core. A recursive subdivision of the range of designated selection data fields and arrangement of the inline tree data structure according to that recursive subdivision can significantly speed up searches of data records that include those selection fields. One example includes geographic searching and filtering of data records that include latitude and longitude used as selection data fields.

Objects and advantages pertaining to electronic data search and retrieval may become apparent upon referring to the exemplary embodiments illustrated in the drawings and disclosed in the following written description or appended claims.

This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates schematically a hierarchical arrangement of a generic dataset.

FIG. 2 illustrates schematically the arrangement of indicia corresponding to the dataset of FIG. 1 in an exemplary conventional flat file database.

FIG. 3 illustrates schematically the arrangement of indicia corresponding to the dataset of FIG. 1 in an exemplary conventional relational database.

FIG. 4 illustrates schematically the arrangement of indicia corresponding to the dataset of FIG. 1 in an exemplary inline tree binary data structure according to the present disclosure.

FIGS. 5A, 5B, and 5C illustrate schematically examples of tables establish correspondence between binary data strings in the data structure of FIG. 4 and alphanumeric data strings in the dataset of FIG. 1.

FIGS. 6A and 6B illustrate schematically examples of binary data fields masks incorporated into the data structure of FIG. 4.

FIG. 7 illustrates schematically detailed exemplary arrangements of binary data strings in the data structure of FIG. 4.

FIG. 8 illustrates schematically a set of selection rectangles superimposed on a map.

FIG. 9 is a flow diagram of an example method for forming a data structure according to the present disclosure.

FIG. 10 is a flow diagram of an example method for searching a data structure according to the present disclosure.

FIG. 11 is a flow diagram of an example method for forming a data structure according to the present disclosure.

FIG. 12 is a flow diagram of an example method for searching a data structure according to the present disclosure.

The embodiments shown in the Figures are exemplary, and should not be construed as limiting the scope of the present disclosure or appended claims.

#### DETAILED DESCRIPTION OF EMBODIMENTS

In many examples of an electronic dataset, the data comprise a multitude of alphanumeric data records, and each one of those data records in turn comprises a corresponding alphanumeric data string in each of multiple data fields. In many instances, the dataset is hierarchical and can be organized according to a multilevel tree structure. Each node of such a tree structure typically represents a one-to-many relationship between (i) a single value (or perhaps a single subrange of values) in one or more data fields at one level of the tree and (ii) one or more values or subranges in one or more other data fields at the next level of the tree.

A dataset of all registered voters in the state of Oregon will be used repeatedly as an example in the present disclosure. The systems and methods disclosed or claimed herein are not, however, limited to that dataset or to datasets of that general type, but can be applied to any dataset in which the data can be arranged according to data structures exemplified herein. The registered voter dataset includes records for about  $1.9 \times 10^6$  individual voters at about  $1.0 \times 10^6$  distinct

addresses. There are several dozen possible data fields for each voter and about 100 possible data fields for each address. A conventional spreadsheet or flat file database containing the Oregon registered voter dataset is about 2 GB (gigabytes) in size when stored on a computer hard disk.

FIG. 1 illustrates schematically an exemplary generic tree structure for organizing data into a three-level hierarchy (levels designated by A, B, and C in FIG. 1). One example of a data hierarchy for the registered voter example might comprise streets (A1, A2, A3, etc), addresses (B11, B12, B13, etc on street A1; B21, B22, B23, etc on street A2; and so on for other addresses Bxy on other streets Ax), and voters (voters C111, C112, C113, etc at address B11; voters C121, C122, C123, etc at address B12; and so on for other voters Cxyz at other addresses Bxy). A terminal node of the tree structure (i.e., at the end of a branch; Cxyz in the example of FIG. 1, or a single voter in the voter dataset) can be referred to as a "leaf node" or simply a "leaf," and corresponds to an individual data record within the dataset. Each data record comprises data strings in corresponding data fields that designate the leaf node and its associated attributes, and can also include data strings in corresponding data fields that designate the higher level nodes to which the leaf node is connected (and attributes associated with those higher level nodes). A hierarchical data tree can include as many levels as needed or desired (which can vary by branch of the tree), and can include as many nodes as needed or desired at any given level. In a further example, the entire hierarchical data arrangement of FIG. 1 can itself constitute a terminal node or intermediate nodes of a larger tree structure (discussed further below). In addition to the registered voter example, other specific examples of data that can be advantageously organized according to hierarchical tree can include sales data, e.g., organized by customers (A), orders (B), and payments (C), or geopolitical data, e.g., organized by continents (A), countries (B), states or provinces (C), and cities (D). Those and any other suitable examples shall fall within the scope of the present disclosure or appended claims.

For convenience of description in the present specification and claims, stored electronic indicia and the underlying data they represent may be referred to interchangeably. It should be noted that the data themselves are an abstraction, and that the representative indicia are the objects that are electronically stored, handled, arranged in a data structure, searched, retrieved, or otherwise manipulated in the methods and systems disclosed or claimed herein. Use of the term "data" in the present disclosure shall be understood to indicate the representative indicia if appropriate in a given context.

One conventional electronic data structure that can be employed to store the data represented in FIG. 1 is an electronic spreadsheet in which electronic indicia representing the data are organized into rows and columns (i.e., a flat file database, with "rows" and "columns" defined in the usual way). Several rows of such a spreadsheet are illustrated schematically in FIG. 2. Each row of the spreadsheet corresponds to one data record of the dataset, hence to one of the "leaf nodes" of the tree of FIG. 1 (e.g., Cxyz). The columns of the spreadsheet correspond to data fields Cxyz-F1, Cxyz-F2, etc for data record Cxyz, corresponding data fields Bxy-F1, Bxy-F2, etc for node Bxy (the corresponding node at the next higher level in the hierarchy), and data fields Ax-F1, Ax-F2, etc for node Ax (the corresponding node two levels higher in the hierarchy). Additional fields would be required for additional levels. Note that there is space reserved in the spreadsheet for every possible data field for every data record, regardless of whether a given data record

has data in that field. Note also that data for the higher-level nodes are repeated in each data record that corresponds to a leaf node connected to that higher-level node.

Another conventional electronic data structure that can be employed to store the data represented in FIG. 1 is an electronic relational database in which electronic indicia representing the data are organized into tables, as illustrated schematically in FIG. 3. Each table record in the "C" table represents a corresponding "leaf node" Cxyz and includes an identifier field Cxyz-ID, corresponding data fields Cxyz-F1, Cxyz-F2, etc., and a field for an identifier Bxy-ID of the corresponding node Bxy in the next higher level. Each table record in the "B" table represents a corresponding node Bxy and includes a field for the identifier Bxy-ID, corresponding data fields Bxy-F1, Bxy-F2, etc., and a field for an identifier Ax-ID of the corresponding node Ax in the next higher level. Each table record in the "A" table represents a corresponding node Ax and includes a field for the identifier Ax-ID and corresponding data fields Ax-F1, Ax-F2, etc. Each table diagram of FIG. 3 is understood to represent multiple different table records of the illustrated contents, as is understood by those skilled in database administration. The dotted lines connecting certain fields of different tables represent one-to-many relationships established within the relational database structure (e.g., one Ax to one or more Bxy's; one Bxy to one or more Cxyz's). Note that, as with the spreadsheet data structure of FIG. 2, space is reserved for every possible field for every data record. However, unlike the spreadsheet example of FIG. 1, data fields common to multiple data records need not be stored repeatedly for every leaf node. For example, the relationship between the Bxy-ID fields in the "B" and "C" tables enables storage of each of the Bxy-Fi fields only once, in the "B" table. The example of FIG. 3 is a relatively simple example of a relational database structure that includes only one-to-many relationships; more complicated examples might include more tables and many-to-many relationships that require so-called "pivot tables."

As noted above, conventional electronic data structures, e.g., spreadsheets and databases, offer great flexibility in terms of adding, removing, or modifying data records, establishing relationships between data fields in different records, and enabling a wide variety of sorts, searches, filters, or queries of the dataset. However, to provide such flexibility, the data structures become quite large and increasingly inefficient as the number of records in the dataset increases, partly due to the data required to define the data structure (i.e., "overhead") and partly due to space reserved for data fields that are empty. To boost speed, relational databases often include search indices, but those further increase the overall size of the data structure. The significant fraction of the impact of the large size of the data structure on the speed at which that structure can be sorted or searched arises from the manner in which large data structures are handled by the computer or server.

In typical use, only a portion of a large dataset can be loaded into the random-access memory (RAM) of a computer or server. A significant fraction of the time required to execute a sort or search of a large dataset is taken up by locating a needed segment of the dataset stored on a disk and pulling that segment into RAM and then into the processor's memory registers for processing, as opposed to the actual processing time once the data is in the processor registers. That sequence must be successively repeated until the entire dataset has been processed. Even worse, in many instances a given segment of the dataset is pulled into RAM more than once during each search operation. One reason for this lies

in the way that data is typically handled by a computer processor. In typical conventional computer processors, data is retrieved into RAM or into a memory cache on the processor in fixed-size segments (e.g., 512 bytes into RAM, or 64 bytes into the cache). To retrieve a particular data field during a search operation, for example, the processor retrieves such a segment of the data that includes the desired field, but that typically also contains other data fields that are not of interest at that time. However, in the course of the entire search operation, it is likely that those other fields will be needed. If so, then the same segment of the data must be retrieved again, perhaps multiple times, to eventually retrieve all of the data fields in that segment.

To significantly speed up certain search, sort, or filter operations on a large dataset, an alternative data structure has been developed and is the subject of the present disclosure. That data structure is illustrated schematically in FIG. 4. The objectives of the data structure of FIG. 4 are to enable dramatic reduction in the overall size of the stored data structure (among other reasons, to allow it to be stored in RAM in its entirety, even if it includes millions or tens of millions of records), and to reduce the number of times a given segment of the data is retrieved from RAM into the processor cache or registers (preferably reduced to a single such retrieval per data segment). For a dataset having a million records of 100 fields each, size reductions by factors of about 5 to 10 or more can be achieved and have been observed, relative to the same dataset in a conventional data structure. For simple search, sort, or filter operations on that dataset, speed enhancements by factors of about 5 to 100 or more can be achieved and have been observed, relative to similar operations performed on the same dataset in a conventional data structure.

The data structure of FIG. 4 can be referred to as an "inline tree" data structure in which the branches and leaves of the tree of FIG. 1 are separated and arranged sequentially. There is no row/column arrangement as in a spreadsheet, nor is there any table arrangement as in a relational database. The data structure of FIG. 4 can be regarded as a single, continuous string of binary indicia representing a single line of characters or digits; a preferred format is a single string of binary digits, as will be explained further below. Within the binary indicia, binary fields represent the alphanumeric data fields in the underlying dataset in a way that reduces their size, and no space is reserved for data fields that contain no data. Data fields are arranged so as to increase the likelihood (i) that when one data segment is pulled into the processor cache for processing, the next segments to be processed have been pulled in along with it, and (ii) that all fields in that segment will be processed after it is first pulled into the processor cache, so that it does not need to be pulled into the processor cache again.

In the hierarchical data of FIG. 1, the data fields Ax-F1, Ax-F2, etc can be referred to as first-level fields. Each node Ax can be defined by specifying, for each data field Ax-Fi, a subrange of data strings (equivalently, data values) that appear in that field in one or more data records. Note that a given subrange can comprise a single string, or a null string (i.e., no string stored in the field). Each node Ax therefore corresponds to a first-level subset of data records in the dataset, wherein the first-level subset includes only those data records for which the data string of each first-level data field Ax-Fi falls within the corresponding subrange. Similarly, each of the data fields Bxy-F1, Bxy-F2, etc can be referred to as second-level fields. Each node Bxy can be defined by specifying, for each field Bxy-Fi, a subrange of data strings (equivalently, data values) that appear in that

field in one or more data records (again, a given subrange can comprise a single string or a null string). Each node Bxy therefore corresponds to a second-level subset of data records within the corresponding first-level subset, wherein the second-level subset includes only those data records for which the data string of each second-level data field Bxy-Fi falls within the corresponding subrange. The foregoing description can be generalized to third-level data field(s) and data record subset(s), fourth-level data field(s) and data record subset(s), and so on.

The general arrangement of the inline tree data structure is illustrated schematically in FIG. 4. Each block in the diagram corresponds to a substantially contiguous binary string, each of which represents one or more data fields that in turn correspond to the branch nodes or leaf nodes of the underlying data (FIG. 1). For example, the binary strings labeled Ax (i.e., A1, A2, A3, etc) include strings representing the values in the data fields Ax-F1, Ax-F2, Ax-F3, etc for the corresponding first-level subsets of the data records. Similarly, the binary strings labeled Bxy include strings representing the values in the data fields Bxy-F1, Bxy-F2, etc, for the corresponding second-level subsets of the data records, and the binary strings labeled Cxyz include strings representing the values in the data fields Cxyz-F1, Cxyz-F2, etc for each corresponding data record.

The binary strings Ax, Bxy, and Cxyz can be arranged in the inline tree so that each first-level subset of data records is represented by binary indicia that comprise a substantially contiguous first-level binary string segment, e.g., binary strings A1, B1y, and C1yz together form a substantially contiguous first-level binary string segment that represents a corresponding first-level subset of data records, binary strings A2, B2y, and C2yz together form another substantially contiguous first-level binary string segment that represents a different corresponding first-level subset of the data records, and so on. Each binary string Ax acts as a header for its corresponding substantially contiguous first-level binary string segment.

Within each first-level binary string segment (whether contiguous or not), the binary strings Bxy and Cxyz are arranged in the inline tree so that each second-level subset of data records is represented by binary indicia that comprise a substantially contiguous second-level binary string segment, e.g., binary strings B11 and C11z together form a substantially contiguous second-level binary string segment that represents a corresponding second-level subset of data records, binary strings B23 and C23z together form another substantially contiguous second-level binary string segment that represents a different corresponding second-level subset of the data records, and so on. Each binary string Bxy acts as a header for its corresponding substantially contiguous second-level binary string segment. The effect of the contiguous arrangement of the second-level binary string segments (and the first-level binary string segments, in some instances) is discussed further below.

Several techniques are employed to drastically reduce the computer memory required to store the inline tree data structure of FIG. 4. As discussed further below, that size reduction leads to significantly faster search and filter operations on the dataset, as well as being desirable in its own right.

A first technique for dataset size reduction includes substitution of a numerical index for each alphanumeric string stored in a data field (sometimes referred to in computer science as string interning). The data in the fields Ax-Fi, Bxy-Fj, and Cxyz-Fk are conventionally represented by alphanumeric data strings, i.e., letters and numbers, and the

data structures are arranged to store in each field any possible alphanumeric string up to a maximum permitted character length. If the maximum character length is, for example, 32 characters, then there are  $36^{32} \approx 6 \times 10^{49}$  possible alphanumeric strings that can be stored in each field (e.g., using any letter or number but not symbols or punctuation marks). Each alphanumeric string stored in the conventional way (i.e., as numbers and letters requiring 1 byte per character plus overhead) would require at least 33 bytes of storage. In any practical circumstance, however, only a tiny fraction of those possible alphanumeric strings actually occur in the dataset. Recognizing that fact allows the size of the inline tree data structure to be substantially reduced relative to conventional spreadsheet or database structures.

To achieve that size reduction, the dataset is analyzed and every unique alphanumeric string that actually occurs in the dataset is identified, enumerated, and stored (only once) in a master string table of any suitable type or format. An example is illustrated schematically in FIG. 5A, in which the enumeration is via a four-byte index (only the last three bits of each index are shown), enabling enumeration of up to  $2^{32} \approx 4.3 \times 10^9$  different alphanumeric strings. In the registered voter example, strings might include every first, middle, or last name, every street name, every city, county, or state name, every party affiliation, every district name, and many dozens of other voter attributes. In an actual dataset of over  $1.9 \times 10^6$  registered voters (each with several dozen possible attributes) and about  $10^6$  addresses (each with about 100 possible attributes) in the state of Oregon, the master string table includes only about 300,000 unique entries (actually slightly less). In the inline tree structure, instead of storing binary indicia that represent alphanumeric strings in the conventional way (e.g., requiring one byte per character plus overhead, or at least 33 bytes per string), the corresponding four-byte numerical index is stored instead, reducing the space required for storing those fields by about a factor of 8. Another advantage of using a string table is that an arbitrary upper limit on the size of the strings need not be imposed. Arbitrarily long strings can be included in the string table without affecting the size of the inline tree data structure.

The string index technique can be further exploited for certain data fields in which only a very limited set of unique alphanumeric strings occur. For example, there are only limited choices for the type of street that appears in a street name, e.g., Street, Boulevard, Avenue, Lane, Road, etc. That field can be replaced by a one-byte index in the inline tree data structure (allowing indexing of up to 256 such strings; only the last three bits are shown) and a corresponding auxiliary string table (illustrated schematically in FIG. 5B). Another example is party affiliation, which can also be replaced by a one byte index in the inline tree data structure (currently there are fewer than 256 recognized political parties) and a corresponding auxiliary string table. Other examples include gender, marital status, street direction, and so on. Any suitable index size or combination of index sizes can be employed (e.g., one-byte, two-byte, three-byte, etc; need not be restricted to a number of whole bytes, i.e., fractional bytes could be used).

Using a master string table and storing a binary index in the inline tree data structure, it is still possible to store any possible alphanumeric string (up to a specified maximum length). Storing the alphanumeric string only once (in the master string table) and storing the corresponding binary indices in the inline tree data structure results in substantial reduction of the size of resulting data file. It should be noted that string indexing can be implemented to reduce the size of data structures other than the inline tree data structure of

FIG. 4. In particular, string indexing can be employed (alone or in combination with other size-reducing techniques, including those disclosed herein) to reduce the size of an inline data structure that is not necessarily arranged according to a hierarchical tree organization scheme, or to reduce the size of a conventional flat file or relational database, or other data structure.

A second technique for dataset size reduction exploits the overlap properties of various attributes in the dataset. In the registered voter example, there are numerous address-related attributes (about 100) that are geographically constrained. These include attributes such as congressional district, state house and senate districts, school, water, or other administrative districts, zip code, county, city, ward, precinct, and so on. Assuming 100 attributes and an average of 10 alternatives per attribute (a conservative estimate), then there are about  $10^{100}$  possible combinations of those attributes. However, many of those combinations include mutually exclusive combinations, e.g., an address in a state senate district in the northeast corner of the state cannot also lie within a school district in the southwest corner of the state, or an address in a county in the southeast corner of the state cannot also lie within a city in the northwest corner of the state. In a specific example, analysis of the registered voter dataset for Oregon reveals that only about 7000 unique combinations of about 100 address-related attributes actually occur among the roughly  $10^6$  unique addresses in the dataset, which affords another opportunity for massively reducing the size of the inline tree data structure of FIG. 4. Each of those combinations shall be referred to herein as an “attribute clump,” “record clump,” “data clump,” or simply as a “clump.” Note that a given clump might include a “null” entry for one or more of the clumped attributes. Attribute clumping enables the substitution into the inline tree data structure of a single clump index per address (e.g., two-byte, four-byte, or other suitable size) to replace alphanumeric strings (e.g., 33 bytes each) or four-byte indices (if the numerical string index technique described above has been implemented) in the nearly 100 data fields per address. A “clump table” can be employed to store the correspondence between the clump index (which can be referred to as a composite data string, because one data string takes the place of a combination of multiple data field values; the clump index can be any suitable size) and the specific alphanumeric strings associated with that clump (exemplary partial entries in such a clump table are illustrated schematically in FIG. 5C). The resulting overall size reduction of the data structure can be enormous (e.g., a reduction of over 3 GB out of about 3.5 GB for an dataset including 100 32-character alphanumeric fields for  $10^6$  addresses, or a reduction of about 400 MB out of about 600 MB for a dataset including 100 4-byte-indexed fields, as described above, for  $10^6$  addresses).

It should be noted that, in the registered voter example dataset, the “street” designation (i.e., the Ax nodes) do not typically correspond to entire streets. To facilitate compression of the data using clumping, each street is divided into segments so that each segment falls within only a single geographic clump. When a “street” is referred to as a level in the hierarchical dataset, it is actually these street segments that are referred to. The clump index becomes one of the fields Ax-Fi of each first-level binary string segment.

The attribute clumping described above is not restricted to geographically constrained, address-related attributes. Any attributes of a given data record can be advantageously clumped in a similar manner, if there is a sufficiently high degree of correlation or anti-correlation between specific

field values in the corresponding fields. For example, in a dataset pertaining to recorded music, certain artists are unlikely to perform in certain genres (e.g., unlikely to have “Philharmonic” and “heavy metal” in the same data record).

In another example, in a dataset pertaining to sales, purchasers of certain products might be quite likely to purchase certain other products (e.g., purchasers of camping gear are likely to also purchase hiking boots).

A single clump encompassing all data records (i.e., no attributes clumped) results in no reduction in size of the data structure; one data record per clump (i.e., all attributes clumped) also results in no size reduction. Between those extremes, one or more optimum subsets of attributes can be found for minimizing the size of the stored data structure using clumping, and various suitable subsets of attributes can be employed for significantly reducing the size of the data structure. Use of such optimum or suitable subsets to reduce the size of the inline tree data structure by clumping shall fall within the scope of the present disclosure or appended claims. The choice of which attributes to clump together depends on the nature of the particular dataset, and the degree of correlation (or anti-correlation) between field values in the corresponding data fields. A certain amount of trial and error may be required for finding a suitable subset of attributes to clump to achieve a needed or desired reduction in the size of the data structure. It is typically but not necessarily the case that clumping is most advantageous when it includes only attributes specific to only the first-level subsets of the data records in a hierarchical dataset (e.g., only address-specific fields in the registered voter example, or only the fields Ax-Fi in the generic example of FIG. 4). It should be noted that attribute clumping can be implemented to reduce the size of data structures other than the inline tree data structure of FIG. 4. In particular, attribute clumping can be employed (alone or in combination with other size-reducing techniques, including those disclosed herein) to reduce the size of an inline data structure that is not necessarily arranged according to a hierarchical tree organization scheme, or to reduce the size of a conventional flat file or relational database, or other data structure.

A third technique for dataset size reduction includes the use of so-called field masks to eliminate the need for space in the data structure for fields that contain no data. For fields that have not been clumped, the corresponding attributes must be stored in the inline tree data structure (as a one-, two-, or four-byte index, for example, as described above). However, not every data record has a specific value stored in every possible field, i.e., some data fields are “nulled.” In conventional data structures such as those illustrated in FIGS. 2 and 3, those nulled data fields take up space as if they were filled. In the inline tree structure of FIG. 4, each binary string Ax, Bxy, and Cxyz includes a field mask near its beginning that specifies which fields are occupied. Each of the binary strings Ax, Bxy, and Cxyz in the inline tree data structure can include a field mask for its corresponding fields Ax-Fi, Bxy-Fi, and Cxyz-Fi, respectively. Examples are illustrated schematically in FIGS. 6A and 6B, in which a one-byte field mask is used to indicate the presence or absence of data in each of eight data fields Ax-F1 . . . Ax-F8.

In FIG. 6A, the one-byte field mask comprises the binary string 10011000, and is followed by values Ax-F1, Ax-F4, and Ax-F5 (in the form of one-, two-, or four-byte indices as described above, for example; a similar field mask could be employed for alphanumeric strings instead). Having 1’s in the 1<sup>st</sup>, 4<sup>th</sup>, and 5<sup>th</sup> bits of the masks and 0’s in the others indicates that the succeeding data values are for the 1<sup>st</sup>, 4<sup>th</sup>, and 5<sup>th</sup> data fields, respectively. No space is required to save

null fields for the 2<sup>nd</sup>, 3<sup>rd</sup>, 6<sup>th</sup>, 7<sup>th</sup>, or 8<sup>th</sup> fields. Instead, the five “0” bits in the mask are stored, which occupy negligible space compared to the 20 bytes potentially required to store the corresponding null fields. Similarly, in FIG. 6B the one-byte field mask comprises the binary string 01010110 followed by values (indices) for Ax-F2, Ax-F4, Ax-F6, and Ax-F7. The only space required for Ax-F1, Ax-F3, Ax-F5, and Ax-F8 are the four 0 bits in the field mask, indicating no data is stored for those fields. The size of the field mask is preferably made large enough to accommodate all data fields in a given binary string that might not contain data. It should be noted that field masking can be implemented to reduce the size of data structures other than the inline tree data structure of FIG. 4. In particular, field masking can be employed (alone or in combination with other size-reducing techniques, including those disclosed herein) to reduce the size of an inline data structure that is not necessarily arranged according to a hierarchical tree organization scheme, or to reduce the size of a conventional flat file or relational database, or other data structure.

FIG. 7 illustrates schematically details of exemplary binary strings Ax, Bxy, and Cxyz. In each example, the binary string begins with an indicator of the location in the inline tree data structure of the next string, e.g., an indicator of the location and Ax+1 occurs at the beginning of the Ax binary string, an indicator of the location of Bx,y+1 occurs at the beginning of the binary string Bxy, and an indicator of the location of Cx,y,z+1 occur at the beginning of the string Cxyz. Those indicators typically take the form of a relative offset (from the current location) or an absolute offset (from the beginning of the binary data file that contains the inline tree data structure). The offset indicators allow the binary strings to assume some characteristics of a linked list, in that each binary string has within it an indicator directing a search program to the next analogous binary string in the data structure to be processed. Note that certain strings may instead include (i) an indicator of the location in the inline tree data structure of a string of a different type (e.g., and indicator in the string Bxy of the position of the string Ax+1), (ii) an indicator that it is the last string in the list, or (iii) another type of indicator. Location indicators in the binary header portions speed up the searching by allowing entire string segments to be skipped, e.g., if the clump index does not match the filter criteria, then the other data fields within that clump need not be searched or filtered.

Next in each exemplary binary string is an indicator of the number of nodes in the next level, e.g., each Ax binary string can include binary digits #Bx indicating the number of B-level nodes correspond to the node Ax (i.e., how many second-level binary string segment are contained within that first-level binary string segment), and each Bxy binary string similarly can include binary digits #Cxy indicating the number of C-level nodes in the next level. Next in each binary string is a field mask, described above, followed by strings representing data in those fields that are indicated by the field mask as containing data. A field for a clump index is included in the appropriate binary strings if data clumping has been employed. Recall that the data strings are not the literal alphanumeric strings, but instead are one-, two-, or four-byte indices (or other suitable binary indices) that correspond to alphanumeric strings according to FIGS. 5A-5C.

The inline tree data structure of FIG. 4 differs profoundly from the conventional data structures of FIGS. 2 and 3 in several important ways. The use of string indexing, clumping, and field masks allow for significant reduction of the size of the stored data structure. Implementing all three

techniques can cut the size by a factor of 10 or more. For example, the Oregon registered voter data set (about  $1.6 \times 10^6$  voters with up to about 25 attributes each at about  $10^6$  addresses with up to about 100 attributes each) can be stored in an inline tree data structure in about 160 MB. A flat file database storing the same data is about 1.5 GB, and a relational database storing that data is about 3 GB (varying depending on the number of different search indices set). Another reason for the size reduction is the substantial lack of so-called “overhead” in the binary file in which the inline tree data structure is stored. In a conventional flat file or relational database, at least one overhead byte is required for each alphanumeric string that is stored. In addition, additional storage is required to store the underlying table structure and relations of a relational database, even before fields of those tables are populated. In contrast, the binary file is just a string of bytes that are not recognizable as a dataset until “decoded” by a search program specifically tailored to the inline tree data structure. Note that a similarly tailored “conversion” program is required to generate the inline tree data structure.

One reason the size reduction is significant is that it enables the entire dataset to be loaded into RAM on a computer or server having reduced memory requirements. The entire 160 MB inline tree data structure can be readily loaded into a computer or server with a relatively ordinary 4 to 8 GB of RAM without significantly burdening the system, whereas the conventional flat file or relational database version of the dataset would severely tax such a system (if it could be loaded at all—a 3 GB database loaded into a 4 GB machine would leave scant resources for the operating system and other vital computer functions). On the other hand, the comparatively small size of the inline tree data structure can enable much larger datasets (e.g.,  $10^8$  voters) to be loaded entirely into RAM in high-end machines having 32 or 64 GB of RAM, wherein the equivalent conventional flat file or relational database simply could not be loaded entirely into RAM on any currently practicable computer or server. Even as hardware capabilities increase, the inline tree data structure will always enable use of a less powerful, less expensive machine to search a dataset of a given size, or searching of larger datasets using a machine of a given memory size.

The size reduction of the data structure is desirable in its own right, as it enables datasets of a given size to be handled by smaller, less powerful computing devices, enables computing devices of given size and power to handle larger datasets, enables faster loading or rebooting of the dataset, or reduces time or cost associated with transmitting, reading, writing, or storing the dataset. Those benefits of size reduction can be realized to varying degrees by applying one or more of the techniques disclosed herein to any suitable data structure, including the inline tree data structure disclosed herein, an inline data structure that is not necessarily arranged according to a hierarchical tree organization scheme, a conventional flat file or relational database, or other data structure. Using the techniques disclosed herein in combination, the reduced data structure size typically can be less than about 5 bytes per field per record or less than about 3 bytes per field per record, often less than about 2 bytes per field per record (e.g., a dataset of one million records having 100 fields each can be stored in less than about 200 MB), or sometimes less than about 1 byte per field per record. Contrast those sizes with 1 to 2 kilobytes per field per record typically required for conventional data structures.

The profoundly reduced size of the inline tree data structure does not come without a cost, however. Flat file

and relational databases excel in their flexibility, enabling ready addition, deletion, or modification of data records in the dataset, often in real time while the database is “live.” A wide variety of search, sort, filter, and retrieval functions can be readily implemented, adapted, or modified, for example using standardized Structured Query Language (SQL). However, as already discussed above, such conventional data structures quickly become impractically slow when they contain large numbers of individual data records. “Large” can mean  $10^6$  records or more in some instances, or may mean as few as  $10^5$  data records or even only  $10^4$  data records in other instances.

The inline tree data structure, on the other hand, cannot be readily modified; if the underlying dataset changes, the inline tree data structure typically must be generated anew by the dedicated conversion program (a relatively slow process). A separate “update” or “override” file or table can be appended to or used with the inline tree data structure, but significantly degrades search and filter speed as it accumulates data records and is therefore not an optimal solution. The inline tree data structure is specifically arranged to perform only one basic task—extremely rapid, filtered search of the data records in a large dataset, for listing or enumeration. Particular data records cannot be randomly accessed or addressed within the inline tree data structure, nor can SQL be used to formulate queries. However, the inline tree data structure can be traversed by a customized search program extremely rapidly, during which a running list or count is kept of those data records matching one or more specified filter criteria. The intermixing of differing data field types within a single inline structure (e.g., the Ax-Fi, Bxy-Fi, and Cxyz-Fi fields all in the same inline structure) is quite unusual and counterintuitive to most database engineers, but that intermixing in part enables the high-speed filtering of the large dataset. That search program is specifically tailored and adapted to the specific arrangement of the inline tree data structure, as is described further below, and the speed of the search is facilitated by the specific arrangement of the binary strings that represent the data records. The available filtering is dictated in part by the clumping, indexing, and masking described above. Differing searches or differing filtering capabilities can require (i) a different inline tree data structure to be generated (using a different, dedicated conversion program) from the same underlying data records and (ii) a different, dedicated search program to be employed. Once generated, the inline tree data structure cannot be readily modified or added to. If the underlying data records are modified or updated, an entirely new inline tree data structure is preferably generated to incorporate those changes.

Another novel feature of the inline tree data structure is that, as a simple sequence or binary indicia (i.e., bytes), the binary file containing the data structure stored on a hard disk quite closely resembles the copy of that file that is read into RAM. That close correspondence has the desirable effect that little if any processing of the file is required when it is first loaded into RAM in preparation for searching. Consequently, the binary file loads into RAM very quickly (e.g., less than 2 second to load the dataset for 1.9 million registered voters). Contrast that with the commonplace experience of waiting several (or many) seconds for, e.g., an ordinary word processor file to load when it is opened; that file’s form when stored on disk differs substantially from its form in RAM, and significant processing (as well as time) is required to achieve the conversion between the two. That processing is substantially eliminated for the binary file of the inline tree data structure. Once the entire binary file is

loaded into RAM, it continues to reside there as long as the user desires to perform searches of the dataset. Fast loading into RAM can be important, however, in a public server-based system in which reliability is important. Rapid loading into RAM can enable fast reboot of the system in the event of an error or crash. Redundant servers can be employed to enhance reliability, or to enable serial updating of the binary data file without making the dataset unavailable during reprocessing of the updated dataset from its original data structure.

An important objective of the size reduction enabled by the inline tree data structure is to increase search speed. By making the data structure fit entirely into RAM, time-consuming seek and retrieval from a hard disk is eliminated. But that is only a first step in accelerating the search speed. As data is used in a computer, the processor’s L1 and L2 memory caches are scanned for the next needed data. If the needed data are present there, they are retrieved from there into the registers and processed (processing means searching in the context of the present disclosure). If not, then the necessary data are retrieved from RAM into the registers and caches.

In typical interactions between a computer processor and the computer’s RAM and disk storage, data is typically retrieved in uniformly sized portions, which get smaller as the data moves from disk to RAM to the registers. Retrieval speeds also increase substantially as the data moves from disk to RAM to the registers. For example, data typically can be read from the disk into RAM in 512 byte portions, with an access time on the order of 10 milliseconds per read. That data is divided into multiple 64 byte portions that are stored in the processor’s L2 cache, and into multiple 16 to 64 byte portions that are stored in the processor’s L1 cache. Those portions align with conventional partitioning of the RAM into uniformly sized segments (e.g., 64-byte alignment of the data in RAM). Data in RAM can be read into L2 cache, L1 cache, and a register with about a 20-50 nanosecond read time; data in the L2 cache can be read into L1 cache and a register with about a 5-10 nanosecond read time; data from the L1 cache can be read into a register with about a 1-2 nanosecond read time.

By virtue of the substantially contiguous, sequential arrangement of the second-level binary string segments (and the first-level binary string segments in some instances), each 64-byte segment read from RAM typically needs to be accessed from RAM only once during any given search, because after it is read into the processor L1 or L2 cache, it remains there until it is processed in its entirety. Once that has occurred, that particular 64-byte portion need not be read from RAM again during that search. Due to the sequential arrangement of the data in the inline tree data structure, slow reads (i.e., from RAM) need only occur relatively infrequently. Each such read brings the next contiguous 64-byte portion of the data into the processor, which is in turn processed substantially in its entirety. The majority of individual reads into the processor registers are from the processor caches, and those reads are significantly faster than reading from RAM. Each 64-byte portion read into cache memory is completely processed before the next 64-byte portion is read. Because of the small size of the binary file, each such read into cache memory enables processing of at least 16 data fields (for fields represented by four-byte indices) or over 100 data fields (when a clump index is read in the voter example).

Contrast this to typical processing of a conventional data structure. The use of alphanumeric data strings limits to about two the number of data fields processed per read from

cache memory. Because there is no intentional sequential arrangement of the bytes read from RAM, it is quite likely that for any given read of 512 bytes only a fraction are relevant to the data fields being processed at that moment. For example, reading multiple attributes for a given voter record typically requires reads from multiple different tables in a relational database, which virtually guarantees that the needed data strings will have to read separately from the hard disk; each of those reads likely includes data from those tables relevant to other voters that are not needed immediately (if at all). The remaining bytes are not used immediately and are eventually written over. However, at some later time during the search process, those unused bytes will be needed and read from RAM again, along with surrounding bytes that, again, might not be needed (and may be needed later, or may already have been processed after an earlier read). Not only is the conventional data structure larger (and therefore inherently slower to read and process), but the reading process also includes significant fractions of useless or redundant reads. Such inefficiencies can be negligible when processing a few hundred or a thousand records, but their cumulative effect becomes readily apparent when attempting to process  $10^4$ ,  $10^5$ ,  $10^6$ , or even larger numbers of records.

As an example of the speed achievable, the example voter database (about  $1.9 \times 10^6$  voter records with about 25 data fields per voter located among about  $10^6$  addresses with about 100 data fields per address) can be searched and filtered at a rate of about 250-400 nanosecond per voter record on a conventional server using a single-core processor running at a standard clock speed (e.g., about 2-3 GHz; usually less than about 4 GHz). That speed is sufficient for the search to appear to a user to occur nearly in real time. A particular filter or set of filters can be selected (e.g., female Democrats aged 40-59 in the 4<sup>th</sup> Congressional district of Oregon) and the total number of voters meeting those criteria (about 35,000 out of about 1.9 million) appears in a fraction of a second. That search and filter speed is about 100 times faster than those achievable with the same data in a conventional relational database (e.g., meaning that the voter number that appeared in a fraction of a second using the inline tree data structure would take a minute or more to update using the conventional relational database). Even with extreme optimization efforts by an experienced database administrator that would be problematic to employ in a typical deployment environment (e.g., consolidation of the relational database into a single flat table, reallocation of computing resources to give the search program unconditional priority over all other computer processes), searching and filtering the conventional data structure thus optimized is still about ten times slower than searching and filtering the inline tree data structure. Search and filter speeds generally achievable using the inline tree data structure with 100 fields per record (using a processor running at a conventional clock speed) can typically be less than about 500 nanoseconds per record per processor core (e.g., a dataset of one million records having 100 fields each can be searched and filtered in less than about 0.5 seconds with a single-core processor running at a standard clock speed), often less than about 400 nanoseconds per record per processor core, or sometimes less than about 300 nanoseconds per record per processor core. Contrast those speeds with 2000 to 5000 nanoseconds per record per core for the extremely optimized conventional database described above, and even slower speeds for conventional data structures that have not had expert optimization or do not have unconditional priority over other computer processes.

A customized binary file generation process is needed to convert a dataset from a conventional data structure (e.g., flat file or relational database) into an inline tree data structure. In contrast to the high-speed search program, the conversion program is typically quite slow, taking on the order of 10 minutes to process  $10^5$  data records. However, that conversion process is performed only infrequently (e.g., to update the data) compared to the frequency of the search and filter process (many times per minute if data is being visualized and manipulated on a map, as described above). A suitable conversion process typically is embodied as a computer program operating on one or more computers, computer systems, or servers, which include one or more processors and include or are otherwise operatively coupled to one or more computer-readable media of any suitable type. Any suitable hardware or hardware-plus-software implementation can be employed for performing the conversion process, which includes: (i) receiving from a first computer-readable storage medium the electronic dataset comprising a multitude of alphanumeric data records arranged according to a conventional data structure; and (ii) using one or more computer processors programmed therefor and operatively coupled to the first storage medium, generating and storing binary indicia of the data set as one or more binary data files on a second computer-readable storage medium operatively coupled to the one or more computer processors, wherein the binary indicia are arranged according to an inline tree data structure as described above.

The generated and stored data structure can also include a string table, any needed or desired auxiliary tables, or a clump table as described above, and the generation process can include, inter alia, analyzing the original dataset and extracting a list of all occurring strings, assigning indices to the strings, writing indicia of the strings and the corresponding indices in the string table, analyzing the data fields to determine combinations of data fields suitable for clumping, identifying the clumps that occur in the dataset, assigning clump indices, and/or writing indicia of the clumps and the corresponding indices into a clump table. It should be noted that the string table, clump table, or auxiliary table are used primarily during the dataset conversion process, for translating requested search filters prior to a search, or for providing a list of retrieved data records (e.g., actual names and addresses of voters meeting the filter criteria). Those ancillary tables are not needed or accessed during the actual search process.

For example, in the voter example, a user request to count (i) male (ii) Republicans (iii) age 45-59 (iv) in Lane County, Ore. might be translated into a search performed by a dedicated search program that counts instances of (i) Cxyz-F4=1, (ii) Cxyz-F8=2, (iii) Cxyz-F5=4, and (iv) clump index=2134 through 2857. Generating a list of those voters might include translating (i) Cxyz-F1=0011 . . . 001 (4-byte), (ii) Cxyz-F2=1101 . . . 110 (4-byte), (iii) Bxy-F1=1110 . . . 000 (4-byte), (iv) Bxy-F2=10101101 (1-byte), (v) Bxy-F3=0001 . . . 011 (4-byte), (vi) Bxy-F4=00011011 (1-byte), and (vii) clump index=2390 into (ii) John (i) Doe, (iii) 1250 (iv) East (v) 17th (vi) Avenue, (vii) Eugene Ore. 97403. Those field numbers, alphanumeric strings, and binary strings are merely one possible example. Myriad examples of searches employing various combinations of filter criteria can be employed within the scope of the present disclosure or appended claims. Any suitable assignment or allocation of field numbers or strings can be employed within the scope of the present disclosure or appended claims.

A suitable search or filtering process typically is embodied as a computer program operating on one or more computers, computer systems, or servers, which include one or more processors and include or are otherwise operatively coupled to one or more computer-readable media of any suitable type. The computers, systems, or servers that perform the search or filtering functions need not be, and often are not, the same as those that performed the data conversion process. In both cases (convert and search/filter), the computer, server, or system can be a stand-alone machine or can comprise one or more machines connected by a local- or wide-area network (LAN or WAN) or the Internet. Any suitable hardware or hardware-plus-software implementation can be employed for searching or filtering, which includes: (a) receiving an electronic query for data records, or an enumeration thereof, having data strings in one or more of the first data fields that fall within a corresponding specified search subranges for those data fields; (b) in response to the query of part (a), with a computer processor programmed therefor, automatically electronically interrogating the first-level binary string segments to identify one or more first-level binary string segments that indicate one or more data records that have data strings within the specified search subranges queried in part (a); (c) in response to the query of part (a), with a computer processor programmed therefor, automatically electronically interrogating the one or more first-level binary string segments identified in part (b) to identify one or more data records indicated therein that have data strings within the one or more specified search subranges queried in part (a); and (d) generating a list or an enumeration of the one or more data records identified in part (c).

Data fields in the original dataset are selected for determining a suitable hierarchical arrangement for the data structure. In some instances, suitable choice will be readily apparent, e.g., if the original dataset is arranged in a series of data tables arranged as a series of one-to-many relationships (as in FIG. 3). In other instances, several choices for a suitable hierarchy might be possible, and one might be selected on the basis of the nature of searches to be performed (e.g., choosing streets as the highest level nodes in the voter data example lends itself to geographic searching or filtering). In an exemplary sales dataset, organizing the dataset with customers as the highest-level nodes might facilitate searching and filtering based on customer-related data fields, while organizing the dataset with products as the highest-level nodes might facilitate searching or filtering based on product-related data fields. Once the hierarchy is selected and defined, data fields not assigned to clumps are assigned to corresponding levels in the hierarchy, and field masks are defined for each level of the hierarchy.

The "selecting," "defining," and similar steps are performed by suitably adapting the dataset conversion program to arrange the inline tree data structure in the desired way. That can be done by direct manual alteration of the conversion program, by indirect alteration of the conversion program using a suitably adapted graphical or text user interface, or by automated alteration of the conversion program based on an automated analysis of the original dataset.

With a suitably adapted conversion program, the original dataset typically is read from a computer-readable medium and processed to produce the corresponding inline tree data structure and its accompanying tables (string, auxiliary, clump). The conversion program works its way through the original, conventional data structure, e.g., to read the alphanumeric strings from the original data fields and store the corresponding binary indices in sequence in the inline tree,

to determine which alphanumeric data fields are populated and store the corresponding field masks in sequence in the inline tree, or to determine to which clump a data record belongs and store the corresponding clump index in sequence in the inline tree. The one or more binary data files embodying the inline tree data structure can be stored on any suitable computer-readable medium as it is generated or after it is generated. In many instances the binary data file is stored on a hard drive or other persistent storage medium, where it is ready to be loaded into RAM or other medium directly accessible to the computer processor that performs the search. In preparation for searching, the binary data file is loaded into RAM in its entirety, as described above, where it can be accessed and retrieved into the processor's caches or registers as described above. The binary data file can be loaded into RAM "on demand" (i.e., in response to a search request) or preferably can reside in RAM in anticipation of one or more subsequent search requests.

One application of the inline tree data structure described herein is high-speed visualization of the geographically linked data overlaid on a map. As the map is panned across a viewing window, or as a user-defined polygon is manipulated on the map, the numbers of data records with geo-coordinates within the window or polygon (both total and filtered according to any one or more desired data fields) are enumerated by searching the inline tree data structure for records having suitable geo-coordinates. In the registered voter example, the number of voters (total or filtered) is updated in near real time as the viewing window or polygon is manipulated (a fractional-second lag is observed when a few dozen simultaneous filters are employed). Each update of those numbers represents a complete reprocessing of the entire dataset (ca. 1.9 million different voter records) and enumeration of those data records that fall within the window or polygon and match the selected filter criteria; that speed is quite remarkable. Such speeds could never be reproduced using a conventional data structure in a typical computing environment. At best, a user would have to wait at least a few seconds up to nearly a minute for each update. Such geographic data visualization is just one example of a completely new use of the dataset that is enabled by the substantially increased search and filter speed, and represents a new and useful result provided by systems and methods disclosed or claimed herein.

An additional technique can be applied that can accelerate certain types of searching and filtering of the dataset, and includes recursive attribute subdivision. The term "recursive subdivision" as used herein shall denote the repeated subdivision of intervals within a range of a particular data field; in the present context, "recursive" does not refer to the computer programming technique commonly referred to as recursion. Recursive attribute subdivision is suitable for attributes that include a range of values that can be readily subdivided into subranges and that can be readily correlated with other attributes of a data record. Recursive attribute subdivision is not necessarily applicable to every dataset. A two-dimensional example of recursive attribute subdivision is recursive geographic subdivision that can be applied to, e.g., the exemplary registered voter dataset (or to any other dataset that includes geo-location data). In the voter example, every address is (or can be) associated with a unique set of geo-coordinates (typically latitude and longitude). For example, every attribute clump (which in this example included address-related attributes) can be assigned subranges of latitudes and longitudes so that every address within the clump falls within the corresponding subranges.

One or more recursively divided attributes can serve as corresponding designated selection fields for the data records of the dataset, facilitating searching and filtering of the on the basis of those selection fields. In the voter dataset, geo-coordinates (or subranges thereof) of each street, address, or clump can serve as designated selection fields to facilitate searching and filtering based on geographic location.

An example of recursive geographic subdivision is illustrated schematically in FIG. 9 for the registered voter dataset. A map of Oregon is shown recursively divided into quartered rectangles. In fact, the boundaries of each "rectangle" are lines of constant latitude or longitude on the approximately spherical surface of the earth, so each "rectangle" is actually the intersection of a spherical lune (bounded by lines of constant longitude) and a spherical zone (bounded by lines of constant latitude); the resulting area will nevertheless be referred to as a rectangle herein. Each rectangle can be specified by its latitude and longitude (beginning and end points, or beginning points and ranges); each voter record includes an address that is (or can be) associated with a latitude and longitude. During generation of the binary data file, the number of voter addresses within each rectangle is determined and compared to a selected record limit (e.g., an absolute numerical limit of 1000 voters; in another example, a relative numerical limit of  $1/1000$  of the total number of voters; any suitable absolute or relative numerical limit can be employed, e.g., 200 voters or 500 voters). If the number of voter records with geo-coordinates within a given rectangle is greater than 1000 voters, then that rectangle is quartered and the process is repeated for the four resulting smaller rectangles. When a given rectangle is found to contain fewer than 1000 voters, there is no further subdivision of that rectangle; it is a so-called "terminal" rectangle. The result is a branched, multilevel "tree" of nested rectangles that cover the geographic area of interest, with smaller rectangles covering areas of higher population density and larger rectangles covering areas of lower population density. Each "leaf" of the recursively subdivided geographic "tree" corresponds to one of the undivided terminal rectangles on the map, each of which encompasses fewer than 1000 voters. Each of those terminal rectangles corresponds to subranges of latitude and longitude (i.e., selection field subranges). In FIG. 9, the rectangles shown illustrate this principle only qualitatively. To actually encompass less than 1000 voters each, the rectangles in even moderately densely populated areas would be much smaller than those shown, and would in fact appear as a mass of dots at the scale shown.

More generally, instead of a numerical record limit to terminate the recursive subdivision, some other criterion or combination of criteria can be employed. For example, in the registered voter example, the subdivision can be terminated when a rectangle encompasses less than a specified maximum number of addresses or streets, or when a minimum geographic area is reached. Any suitable criteria can be employed.

During further generation of the binary data file, the subset of addresses on a given street that fall within a given terminal rectangle are considered as candidates to define a corresponding "street" (more accurately, a street segment; a "street" is an example of a first-level subset of data records within the voter dataset). If data clumping has been employed, and if addresses on the candidate street fall within different clumps, the candidate street is further divided into segments having addresses falling within only one data clump. A street in this example dataset is therefore a subset

of addresses that (i) fall within the same one of the geographic terminal rectangles, and (ii) fall within the same data clump.

A binary selection header string can be included in the binary data file. Such a header can comprise a linked list of the corresponding selection field subranges, in a manner that reflects the recursively subdivided "tree" structure described above (e.g., a selection field subrange of a given rectangle can be linked in the list to selection field subrange of one of its sub-rectangles). Each terminal record in the linked list (i.e., each "leaf") corresponds to one of the terminal rectangles, and can indicate a location within the inline tree data structure of a first-level header of a corresponding first-level binary string. In the more concrete example of the voter dataset, the binary selection header comprises a linked list of latitude and longitude subranges for the recursively subdivided rectangles. Each terminal record in the linked list (designating one of the undivided, terminal rectangles) indicates the location of one of the street-level binary headers in the inline tree data structure. The subset comprising streets that fall within a given terminal rectangle can be arranged in the binary data file as a linked list of first-level binary street segments that fall within that rectangle. The terminal record of the linked list of those streets can indicate the next record in the linked list of latitude/longitude subranges. That pattern can be repeated until all rectangles and streets segments are linked.

The structure described in the foregoing paragraph can enable extremely efficient searching and filtering based on geographic location. The search program can be adapted to first search the linked list of latitude/longitude subranges and compare those to a user-selected viewing window or polygon on a map. Any rectangles that do not overlap the window or polygon can be skipped over without searching or filtering any of the corresponding street, address, or voter fields. The recursively subdivided tree structure can be thought of as guiding the search and filter processes to those portions of the data structure where pertinent data records are to be found.

The inline tree data structure and recursive subdivision based on geographical coordinates can enable near-real-time visualization or manipulation of extremely large datasets overlaid on a map (e.g.,  $>10^6$  data records). That new and useful result can be applied to a virtually endless variety of geographically-linked data. Just a few of myriad examples include data pertaining to voters, populations, demographics, economics, taxation, government administration, law enforcement, education, political polling, campaigns, or elections, media distribution or consumption (print, radio, video, Internet), telecommunications, real estate, insurance, transportation and shipping (land, sea, or air), fleet management (autos, trucks, buses, trains, transit vehicles, boats or ships, aircraft, and so on), product or material marketing, sales, or distribution (wholesale or retail), manufacturing, supply chains, raw materials (water, forests, mineral deposits, fossil fuel deposits), agriculture, medical or epidemiologic data, wildlife monitoring or management, astronomical data (using galactic latitude and longitude), power generation or transmission, manmade or natural disasters, disaster response or logistics, and so on.

Other types of datasets can be arranged according to recursively divided subranges of data strings in one or more designated selection fields. Such arrangements can enable rapid searching and filtering of data records having attributes falling within designated selection field subranges. Any recursive subdivision of data field subranges of any desired dimensionality using any suitable set of one or more chosen

selection fields shall fall within the scope of the present disclosure or appended claims. For example, in a dataset pertaining to people and organized by last name, recursive subdivision based on letters in each name can be employed, with varying numbers of letters defining each recursively divided subrange as needed. It should be noted that recursive subdivision of designated selection field subranges can be implemented to facilitate searching of data structures other than the inline tree data structure of FIG. 4. In particular, such recursive subdivision can be employed (alone or in combination with other techniques, including size-reducing techniques disclosed herein) to guide searching and filtering of an inline data structure that is not necessarily arranged according to a hierarchical tree organization scheme, or to guide searching and filtering of a conventional flat file or relational database, or other data structure.

The systems and methods disclosed herein can be implemented as or with general or special purpose computers or servers or other programmable hardware devices programmed through software, or as hardware or equipment "programmed" through hard wiring, or a combination of the two. A "computer" or "server" can comprise a single machine or can comprise multiple interacting machines (located at a single location or at multiple remote locations). Computer programs or other software code, if used, can be implemented in temporary or permanent storage or in replaceable media, such as by including programming in microcode, machine code, network-based or web-based or distributed software modules that operate together, RAM, ROM, CD-ROM, CD-R, CD-R/W, DVD-ROM, DVD±R, DVD±R/W, hard drives, thumb drives, flash memory, optical media, magnetic media, semiconductor media, or any future storage alternatives. The one or more binary data files embodying the inline tree data structure can also be stored on any suitable computer-readable medium, including those listed above, but as disclosed herein the inline tree data structure is preferably loaded entirely into a computer-readable medium that is directly accessible to a computer processor executing a search of the data structure, e.g., a computer random access memory (RAM).

It is intended that equivalents of the disclosed exemplary embodiments and methods shall fall within the scope of the present disclosure or appended claims. It is intended that the disclosed exemplary embodiments and methods, and equivalents thereof, may be modified while remaining within the scope of the present disclosure or appended claims.

In the foregoing Detailed Description or in the Drawings, various features may be grouped together in several exemplary embodiments for the purpose of streamlining the disclosure. This method of disclosure is not to be interpreted as reflecting an intention that any claimed embodiment requires more features than are expressly recited in the corresponding claim. Rather, as the appended claims reflect, inventive subject matter may lie in less than all features of a single disclosed exemplary embodiment. Thus, the appended claims are hereby incorporated into the Detailed Description, with each claim standing on its own as a separate disclosed embodiment. However, the present disclosure shall also be construed as implicitly disclosing any embodiment having any suitable set of one or more disclosed or claimed features (i.e., sets of features that are not incompatible or mutually exclusive) that appear in the present disclosure or the appended claims, including those sets that may not be explicitly disclosed herein. It should be

further noted that the scope of the appended claims do not necessarily encompass the whole of the subject matter disclosed herein.

For purposes of the present disclosure and appended claims, the conjunction "or" is to be construed inclusively (e.g., "a dog or a cat" would be interpreted as "a dog, or a cat, or both"; e.g., "a dog, a cat, or a mouse" would be interpreted as "a dog, or a cat, or a mouse, or any two, or all three"), unless: (i) it is explicitly stated otherwise, e.g., by use of "either . . . or," "only one of," or similar language; or (ii) two or more of the listed alternatives are mutually exclusive within the particular context, in which case "or" would encompass only those combinations involving non-mutually-exclusive alternatives. For purposes of the present disclosure or appended claims, the words "comprising," "including," "having," and variants thereof, wherever they appear, shall be construed as open ended terminology, with the same meaning as if the phrase "at least" were appended after each instance thereof.

In the appended claims, if the provisions of 35 USC §112 ¶6 are desired to be invoked in an apparatus claim, then the word "means" will appear in that apparatus claim. If those provisions are desired to be invoked in a method claim, the words "a step for" will appear in that method claim. Conversely, if the words "means" or "a step for" do not appear in a claim, then the provisions of 35 USC §112 ¶6 are not intended to be invoked for that claim.

The Abstract is provided as required as an aid to those searching for specific subject matter within the patent literature. However, the Abstract is not intended to imply that any elements, features, or limitations recited therein are necessarily encompassed by any particular claim. The scope of subject matter encompassed by each claim shall be determined by the recitation of only that claim.

What is claimed is:

1. A computer-implemented method comprising:

receiving from one or more computer-readable storage media electronic indicia of a multitude of alphanumeric data records of a hierarchical dataset, each data record including alphanumeric data strings for multiple corresponding defined first-level and second-level data fields; and

using one or more computer processors programmed therefor and operatively coupled to the one or more storage media, generating and storing binary indicia of the multitude of alphanumeric data records as one or more binary data files stored on one or more computer-readable storage media operatively coupled to the one or more computer processors,

wherein:

the one or more data files form a single, continuous binary string including multiple first-level binary string segments and, following each first-level binary string segment, one or more corresponding second-level binary string segments;

for the first-level data fields, each range of data strings for the first-level data fields is divided into multiple corresponding first-level subranges, and the multitude of data records comprises multiple first-level subsets of the data records, wherein each first-level subset includes only those data records for which each data string of the first-level data fields falls within a corresponding one of the first-level subranges;

each first-level binary string segment encodes the data strings of the first-level data fields of a corresponding one of the first-level subsets of the data records;

23

each second-level binary string segment encodes data strings of the second-level data fields of that first-level subset of the data records which corresponds to the immediately preceding first-level binary string; and arrangement of the one or more binary files enables access to the first-level and second-level binary strings in only the order in which each appears in the single, continuous binary string, without enabling random access to the binary indicia of the data strings of the multitude of data records.

2. The method of claim 1 wherein:

each data record further includes alphanumeric data strings for multiple corresponding defined third-level data fields;

the single, continuous binary string further includes, following each second-level binary string segment, one or more corresponding third-level binary string segments;

for the second-level data fields, each range of data strings for the second-level data fields is divided into multiple corresponding second-level subranges, and each one of the multiple first-level subsets of the data records comprises multiple second-level subsets of the data records, wherein each second-level subset includes only those data records for which each data string of the second-level data fields falls within a corresponding one of the second-level subranges;

each second-level binary string segment encodes the data strings of the second-level data fields of a corresponding one of the second-level subsets of the data records;

each third-level binary string segment encodes data strings of the third-level data fields of that second-level subset of the data records which corresponds to the immediately preceding second-level binary string; and arrangement of the one or more binary files enables access to the first-level, second-level, and third-level binary strings in only the order in which each appears in the single, continuous binary string, without enabling random access to the binary indicia of the data records of the multitude of data records.

3. The method of claim 1 wherein each of the first-level binary string segments includes a corresponding first-level header portion and each of the second-level binary string segments includes a corresponding second-level header portion, wherein:

one or more of the first-level header portions indicate which of several of the first-level data fields have data strings encoded by the corresponding first-level binary string segment; and

one or more of the second-level header portions indicate which of several of the second-level data fields have data strings encoded by the corresponding second-level binary string segment.

4. The method of claim 1 wherein:

each combination of specific data strings, for a selected subset of the defined first-level data fields, that occurs in the dataset is assigned a corresponding composite data string; and

the selected subset of the defined data fields is encoded in a corresponding one of the first-level binary string segments by binary indicia of only the composite data string.

5. The method of claim 4 wherein the selected subset of the defined data fields includes the first set of data fields.

6. An article comprising a tangible, non-transitory computer-readable medium encoded to store one or more binary data files generated and stored by a method comprising:

24

receiving from one or more computer-readable storage media electronic indicia of a multitude of alphanumeric data records of a hierarchical dataset, each data record including alphanumeric data strings for multiple corresponding defined first-level and second-level data fields; and

using one or more computer processors programmed therefor and operatively coupled to the one or more storage media, generating and storing binary indicia of the multitude of alphanumeric data records as one or more binary data files stored on the tangible, non-transitory computer-readable medium,

wherein:

the one or more data files form a single, continuous binary string including multiple first-level binary string segments and, following each first-level binary string segment, one or more corresponding second-level binary string segments;

for the first-level data fields, each range of data strings for the first-level data fields is divided into multiple corresponding first-level subranges, and the multitude of data records comprises multiple first-level subsets of the data records, wherein each first-level subset includes only those data records for which each data string of the first-level data fields falls within a corresponding one of the first-level subranges;

each first-level binary string segment encodes the data strings of the first-level data fields of a corresponding one of the first-level subsets of the data records;

each second-level binary string segment encodes data strings of the second-level data fields of that first-level subset of the data records which corresponds to the immediately preceding first-level binary string; and

arrangement of the one or more binary files enables access to the first-level and second-level binary strings in only the order in which each appears in the single, continuous binary string, without enabling random access to the binary indicia of the data strings of the multitude of data records.

7. The article of claim 6 wherein the one or more binary data files have a total size that is less than about 2 bytes per field per data record of the multitude.

8. The article of claim 6 wherein the tangible, non-transitory computer-readable medium is directly accessible to a computer processor.

9. The article of claim 6 wherein the tangible, non-transitory computer-readable medium comprises random access memory.

10. The article of claim 9 wherein the one or more binary data files have a total size that is less than about 2 bytes per field per data record of the multitude.

11. A computer-implemented method for using the article of claim 6 to search the one or more binary data files stored on the tangible, non-transitory computer-readable medium, the method comprising:

(a) receiving an electronic query for data records, or an enumeration thereof, having certain data strings in one or more of the first-level or second-level data fields that fall within corresponding specified search subranges for those data fields;

(b) in response to the query of part (a), with a computer processor programmed therefor, automatically electronically interrogating each of the first-level binary string segments, in only the order in which each appears in the single, continuous binary string, to identify one or more first-level binary string segments

25

that encode data strings falling within a corresponding one of the specified search subranges queried in part (a);

- (c) in response to the query of part (a), with a computer processor programmed therefor, after each identification of part (b) and before interrogation of the next first-level binary string segment, automatically electronically interrogating, in only the order in which each appears in the single, continuous binary string, those corresponding second-level binary string segments that follow, in the single, continuous binary string, the corresponding first-level binary string segment, to identify one or more second-level binary string segments that encode data strings falling within a corresponding one of the specified search subranges queried in part (a); and
- (d) generating a list or an enumeration of the one or more data records corresponding to the second-level binary string segments identified in part (c).

**12.** The method of claim **11** wherein: each combination of specific data strings, for a selected subset of the defined first-level data fields, that occurs in the dataset is assigned a corresponding composite data string; and

the selected subset of the defined data fields is encoded in a corresponding one of the first-level binary string segments by binary indicia of only the composite data string.

**13.** The method of claim **11** wherein part (b) includes, for each first-level binary segment, skipping to the next first-level binary string segment, without interrogating the intervening second-level binary string segments, if the corresponding first-level binary string segment does not encode a data string that falls at least overlaps a corresponding one of the specified search subranges.

**14.** The method claim **11** wherein the tangible, non-transitory computer-readable medium is directly accessible to the computer processor of part (b) or (c).

**15.** The method of claim **14** further comprising loading sequentially, according to location within the single, continuous binary string, the first-level binary string segments into a processor cache computer memory.

**16.** The method of claim **15** further comprising loading sequentially, according to location within the single, continuous binary string, into a processor cache computer memory, those second-level binary string segments interrogated in part (c).

**17.** The method of claim **11** wherein the tangible, non-transitory computer-readable medium comprises random-access memory.

**18.** The method of claim **11** wherein part (b) includes, for each first-level binary segment, skipping to the next first-level binary string segment according a length or position indicated by a first-level header portion of the corresponding first-level binary string segment, without interrogating the intervening second-level binary string segments, if the corresponding first-level binary string segment does not encode a data string that at least overlaps a corresponding one of the specified search subranges.

**19.** The method of claim **11** wherein the one or more binary data files indicate at least 100,000 data records, and the interrogations of parts (b) and (c) are performed in less than 500 nanoseconds per data record per processor core.

**20.** The method of claim **11** wherein the one or more binary data files indicate at least 1,000,000 data records, and the interrogations of parts (b) and (c) are performed in less than 500 nanoseconds per data record per processor core.

26

**21.** The method of claim **11** further comprising:

- (e) receiving an electronic indicator of one or more designated filter data fields and corresponding filter subranges for each filter data field;
- (f) with a computer processor programmed therefor, automatically identifying one or more of the data records identified in part (c) that have a data string in each of the filter data fields that falls within the corresponding indicated filter subrange; and
- (g) generating a list or an enumeration of one or more data records identified in part (f).

**22.** A computer-implemented method comprising: receiving from one or more computer-readable storage media electronic indicia of multitudes of alphanumeric data records of a hierarchical dataset, each data record including alphanumeric data strings for multiple corresponding defined first-level and second-level data fields;

for each of one or more designated selection data fields of each data record, using one or more computer processors programmed therefor and operatively coupled to the one or more storage media, automatically recursively subdividing into multiple selection field subranges a corresponding range of data strings of the selection data fields, wherein each different combination of selection field subranges indicates a different corresponding selected multitude of data records of the hierarchical dataset, wherein each corresponding selected multitude includes only those data records for which each data string of the selection data fields falls within a corresponding selection field subrange;

for each one of the selected multitudes of data records, using one or more computer processors programmed therefor and operatively coupled to the one or more storage media, generating and storing binary indicia of the selected multitude of alphanumeric data records as one or more binary data files stored on one or more computer-readable storage media operatively coupled to the one or more computer processors; and using one or more computer processors programmed therefor and operatively coupled to the one or more storage media, generating and storing binary indicia of a binary header string in one or more of the binary data files,

wherein:

for each one of the selected multitudes of data records, the one or more data files form a single, continuous binary string including multiple first-level binary string segments and, following each first-level binary string segment, one or more corresponding second-level binary string segments;

for the first-level data fields, each range of data strings for the first-level data fields is divided into multiple corresponding first-level subranges, and the multitude of data records comprises multiple first-level subsets of the data records, wherein each first-level subset includes only those data records for which each data string of the first-level data fields falls within a corresponding one of the first-level subranges;

each first-level binary string segment encodes the data strings of the first-level data fields of a corresponding one of the first-level subsets of the data records;

each second-level binary string segment encodes data strings of the second-level data fields of that first-level subset of the data records which corresponds to the immediately preceding first-level binary string; and

27

arrangement of the one or more binary files enables access to the first-level and second-level binary strings in only the order in which each appears in the single, continuous binary string, without enabling random access to the binary indicia of the data strings of the multitude of data records;

the binary selection header string comprises a linked list of indicia of multiple combinations of selection field subranges, and the linked list is linked in a manner that reflects the recursive subdivision of the selection field subranges; and

each terminal record in the linked list indicates a corresponding one of the selected multitudes of data records.

23. The method of claim 22 wherein (i) the dataset comprises a set of data records linked to geographical coordinates, (ii) the designated selection fields are latitude and longitude, (iii) the first-level data fields are related to street segments, and (iv) the second-level data fields are related to street addresses.

24. The method of claim 22 wherein the recursive subdivision is terminated based on a selected absolute or relative numerical limit on a number of data records in each selected multitude.

25. An article comprising a tangible, non-transitory computer-readable medium encoded to store one or more binary data files generated and stored by a method comprising:

receiving from one or more computer-readable storage media electronic indicia of multitudes of alphanumeric data records of a hierarchical dataset, each data record including alphanumeric data strings for multiple corresponding defined first-level and second-level data fields;

for each of one or more designated selection data fields of each data record, using one or more computer processors programmed therefor and operatively coupled to the one or more storage media, automatically recursively subdividing into multiple selection field subranges a corresponding range of data strings of the selection data fields, wherein each different combination of selection field subranges indicates a different corresponding selected multitude of data records of the hierarchical dataset, wherein each corresponding selected multitude includes only those data records for which each data string of the selection data fields falls within a corresponding selection field subrange;

for each one of the selected multitudes of data records, using one or more computer processors programmed therefor and operatively coupled to the one or more storage media, generating and storing binary indicia of the selected multitude of alphanumeric data records as one or more binary data files stored on one or more computer-readable storage media operatively coupled to the one or more computer processors; and

using one or more computer processors programmed therefor and operatively coupled to the one or more storage media, generating and storing binary indicia of a binary header string in one or more of the binary data files,

wherein:

for each one of the selected multitudes of data records, the one or more data files form a single, continuous binary string including multiple first-level binary string segments and, following each first-level binary string segment, one or more corresponding second-level binary string segments;

for the first-level data fields, each range of data strings for the first-level data fields is divided into multiple cor-

28

responding first-level subranges, and the multitude of data records comprises multiple first-level subsets of the data records, wherein each first-level subset includes only those data records for which each data string of the first-level data fields falls within a corresponding one of the first-level subranges;

each first-level binary string segment encodes the data strings of the first-level data fields of a corresponding one of the first-level subsets of the data records;

each second-level binary string segment encodes data strings of the second-level data fields of that first-level subset of the data records which corresponds to the immediately preceding first-level binary string; and

arrangement of the one or more binary files enables access to the first-level and second-level binary strings in only the order in which each appears in the single, continuous binary string, without enabling random access to the binary indicia of the data strings of the multitude of data records;

the binary selection header string comprises a linked list of indicia of multiple combinations of selection field subranges, and the linked list is linked in a manner that reflects the recursive subdivision of the selection field subranges; and

each terminal record in the linked list indicates a corresponding one of the selected multitudes of data records.

26. A computer-implemented method for using the article of claim 25 to search the one or more binary data files stored on the tangible, non-transitory computer-readable medium, the method comprising:

(a) receiving an electronic query for data records, or an enumeration thereof, having data strings in one or more designated selection data fields, and certain data strings in one or more of the first-level or second-level data fields, that fall within corresponding specified search subranges for each of those data fields;

(b) in response to the query of part (a), with a computer processor programmed therefor, automatically electronically interrogating the binary selection header string to identify one or more of the selected multitudes of data records that have designated selection data field data strings within the one or more specified search subranges queried in part (a);

(c) in response to the query of part (a), with a computer processor programmed therefor, for each selected multitude identified in part (b), automatically electronically interrogating each one of the first-level binary string segments of the corresponding single, continuous binary string, in only the order in which each appears in that single, continuous binary string, to identify one or more first-level binary string segments that encode data strings falling within the one or more specified search subranges queried in part (a);

(d) in response to the query of part (a), with a computer processor programmed therefor, after each identification of part (c) and before interrogation of the next first-level binary string segment, automatically electronically interrogating, in only the order in which each appears in the single, continuous binary string, those corresponding second-level binary string segments that follow, in the single, continuous binary string, the corresponding first-level binary string segment, to identify one or more second-level binary string segments that encode data strings falling within the one or more specified search subranges queried in part (a); and

(e) generating a list or an enumeration of one or more data records identified in part (d).

29

27. The method of claim 26 wherein (i) the dataset comprises a set of data records linked to geographical coordinates, (ii) the designated selection fields are latitude and longitude, (iii) the first-level data fields are related to street segments, and (iv) the second-level data fields are related to street addresses.

28. The method of claim 26 wherein the tangible, non-transitory computer-readable medium is directly accessible to the computer processor of part (b) and the computer processor of part (c).

29. The method of claim 28 wherein the tangible, non-transitory computer-readable medium comprises random access memory.

30. A computer system comprising one or more computer processors and one or more computer-readable storage media operatively coupled to one or more of the computer processors, wherein the computer system is structured, connected, and programmed to perform a method comprising:

receiving from one or more of the computer-readable storage media electronic indicia of a multitude of alphanumeric data records of a hierarchical dataset, each data record including alphanumeric data strings for multiple corresponding defined first-level and second-level data fields; and

using one or more of the computer processors, generating and storing binary indicia of the multitude of alphanumeric data records as one or more binary data files stored on one or more of the computer-readable storage media,

wherein:

the one or more data files form a single, continuous binary string including multiple first-level binary string segments and, following each first-level binary string segment, one or more corresponding second-level binary string segments;

for the first-level data fields, each range of data strings for the first-level data fields is divided into multiple corresponding first-level subranges, and the multitude of data records comprises multiple first-level subsets of the data records, wherein each first-level subset includes only those data records for which each data string of the first-level data fields falls within a corresponding one of the first-level subranges;

each first-level binary string segment encodes the data strings of the first-level data fields of a corresponding one of the first-level subsets of the data records;

each second-level binary string segment encodes data strings of the second-level data fields of that first-level subset of the data records which corresponds to the immediately preceding first-level binary string; and arrangement of the one or more binary files enables access to the first-level and second-level binary strings in only

30

the order in which each appears in the single, continuous binary string, without enabling random access to the binary indicia of the data strings of the multitude of data records.

31. An article comprising a tangible, non-transitory medium encoding computer-readable instructions that, when applied to a computer system comprising at least one processor, instruct the computer system to perform a method comprising:

receiving from one or more computer-readable storage media electronic indicia of a multitude of alphanumeric data records of a hierarchical dataset, each data record including alphanumeric data strings for multiple corresponding defined first-level and second-level data fields; and

using one or more computer processors programmed therefor and operatively coupled to the one or more storage media, generating and storing binary indicia of the multitude of alphanumeric data records as one or more binary data files stored on one or more computer-readable storage media operatively coupled to the one or more computer processors,

wherein:

the one or more data files form a single, continuous binary string including multiple first-level binary string segments and, following each first-level binary string segment, one or more corresponding second-level binary string segments;

for the first-level data fields, each range of data strings for the first-level data fields is divided into multiple corresponding first-level subranges, and the multitude of data records comprises multiple first-level subsets of the data records, wherein each first-level subset includes only those data records for which each data string of the first-level data fields falls within a corresponding one of the first-level subranges;

each first-level binary string segment encodes the data strings of the first-level data fields of a corresponding one of the first-level subsets of the data records;

each second-level binary string segment encodes data strings of the second-level data fields of that first-level subset of the data records which corresponds to the immediately preceding first-level binary string; and

arrangement of the one or more binary files enables access to the first-level and second-level binary strings in only the order in which each appears in the single, continuous binary string, without enabling random access to the binary indicia of the data strings of the multitude of data records.

\* \* \* \* \*