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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
6 August 2015 (06.08.2015)

WIPO | PCT

(10) International Publication Number
WO 2015/116221 A1

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

(21) International Application Number:
PCT/US2014/014329

(22) International Filing Date:
31 January 2014 (31.01.2014)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).

(72) Inventors: **CHEN, QiFan**; 14231 Tandem Boulevard, Austin, Texas 78728 (US). **KOSURU, Ramakumar**; 14231 Tandem Boulevard, Austin, Texas 78728 (US). **LAKSHMINARAYAN, Choudur**; 14231 Tandem Boulevard, Austin, Texas 78728 (US).

(74) Agents: **SHOOKMAN, Jeb A** et al.; Hewlett-Packard Company, Intellectual Property Administration, 3404 E. Harmony Road, Mail Stop 35, Fort Collins, Colorado 80528-9599 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *as to the identity of the inventor (Rule 4.17(i))*

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

Published:

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

(54) Title: MANAGING DATABASE WITH COUNTING BLOOM FILTERS

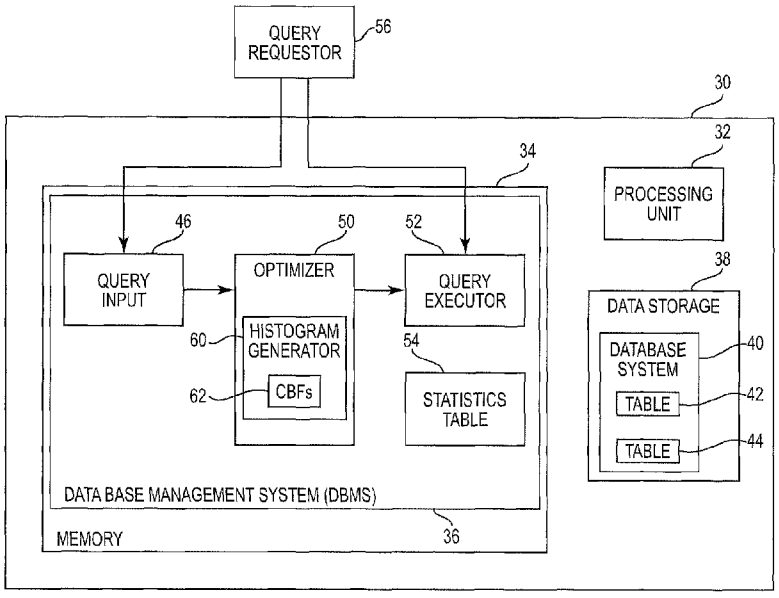


Fig. 1

(57) Abstract: A method of managing a database including creating an initial counting bloom filter (CBF) instance having an array of counters and hash functions that map an inserted value to the array of counters, and designating the initial CBF instance as a current CBF instance, and sequentially inserting each value of a sample data set of a table column into the hash functions of the current CBF instance and incrementing counters of the array of counters to which the value is mapped. The method further includes, prior to inserting each value into the hash functions of the current CBF instance, when a number of counters of the array of counters having non-zero values is at least at a threshold level, designating the current CBF instance as an old CBF instance, creating a new CBF instance having an array of counters and hash functions that map an inserted value to the array counters, and designating the new CBF instance as the current CBF instance.

WO 2015/116221 A1

MANAGING DATABASE WITH COUNTING BLOOM FILTERS

Background

[0001] Relational database systems store tables of data arranged in rows and columns, and which are typically linked by relationships that simplify data storage and make queries of the stored data more efficient. Structured query language (SQL) is a standardized language for creating and operating on relational databases. When receiving a query, such as from a system external to the database system (e.g. a web browser), a database system typically employs a query optimizer to construct a query plan that provides the most efficient way to execute the query (i.e. an optimal query plan).

[0002] To produce an optimal query plan, the query optimizer uses statistics associated with the involved data (e.g. a column of data), such as the number of unique values and cardinality estimates, for example. Cardinality estimates, in-turn, rely on histogram statistics of the involved table columns. In Enterprise Data Warehouse environments, content of the tables changes frequently. As a result, it is also necessary to frequently update database statistics so that the statistics accurately reflect the content changes and enable the query optimizer to continue providing optimal query plans.

Brief Description of the Drawings

[0003] Figure 1 is a schematic block diagram illustrating an example of a system for managing a database.

[0004] Figure 2 is a schematic diagram illustrating an example of a counting bloom filter instance.

[0005] Figure 3 is a schematic diagram illustrating an example operation of a counting bloom filter instance.

[0006] Figure 4 is a schematic diagram illustrating an example operation of a counting bloom filter instance.

[0007] Figure 5 is a schematic diagram illustrating an example operation of a counting bloom filter instance.

[0008] Figure 6 is a graph illustrating an example of an equal-width histogram.

[0009] Figure 7 is a graph illustrating an example of an equal-height histogram.

[0010] Figure 8 is a flow diagram illustrating an example of a method for managing a database.

[0011] Figure 9 is a flow diagram illustrating an example of a method for managing a database.

[0012] Figure 10 illustrates tables showing relative database management performance according to the present disclosure.

[0013] Figure 11 a schematic block diagram illustrating an example of a computer system.

Detailed Description

[0014] In the following detailed description, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration specific examples in which the disclosure may be practiced. It is to be understood that other examples may be utilized and structural or logical changes may be made without departing from the scope of the present disclosure. The following detailed description, therefore, is not to be taken in a limiting sense, and the scope of the present disclosure is defined by the appended claims. It is to be understood that features of the various examples described herein may be combined, in part or whole, with each other, unless specifically noted otherwise.

[0015] Figure 1 is a schematic block diagram illustrating an example of a system 30 for managing a database in which Counting Bloom Filters (CBFs) for providing histogram statistics in fast and memory-efficient fashion, according to embodiments of the present disclosure, can be implemented. System 30 includes a processing unit 32, and a memory 34. Processing unit 32 can include at least one processing core. Memory 34 can be used to store data and machine executable instructions, and can be implemented as a computer readable medium, such, for example, random access memory, non-volatile memory, etc. Memory 34 can include a database management system (DBMS) 36 that accesses a database 40 stored in data storage 38. Data storage 38 can be implemented as a computer readable medium such as, for example, a hard disk system, a solid state drive system, random RAM (volatile or non-volatile), etc. Database 40, in one example, can be implemented as a relational database including tables, such as tables 42, 44, having data arranged in rows and columns which can be queried using SQL.

[0016] DBMS 36, according to the illustrated example, includes a query input 46 (e.g. a queue), a query optimizer 50, a query executor 52, and a statistics table 54. In one example, query input 52 can receive a query from a query requestor 56. Query requestor 56 can be, for example, a system external to DBMS 36 or to both DBMS 36 and system 2, such as an application executing on another computer. For example, query requestor 56 could be implemented as a web browser. In one example, the received query can be implemented as SQL relational logic (e.g. includes relational operators and/or predicates and literals).

[0017] According to one example, query input 46 provides the received query to query optimizer 50 to generate a query plan. To generate an optimal query plant, query optimizer 50 employs heuristics and logics, and relies on statistics for a given attribute, such as a number of unique values and, particularly, cardinality estimates. Cardinality estimates, in-turn, rely on histogram statistics of involved data, such as a column of data from one of tables of database 40. Such database statistics are maintained and accessed by optimizer 50, such as in statistics table 54. In Enterprise Data Warehouse environments, contents of the tables of database 40, such as tables 42 and 44, are frequently updated. To

reflect the content changes, the statistics for the tables, including histogram statistics, must also be frequently updated.

[0018] A histogram for a column of data, such as a column of data from one of the tables 42, 44, typically splits the values in the column in distinct ranges, with each distinct range having approximately the same number of rows (an equal-height histogram). For each distinct range, the number of unique values, the so-called unique entry count (UEC), the row count (the total number of rows), a list of skew values, and other parameters are maintained.

[0019] Histogram statistics for a set of columns are determined by randomly sampling rows from the table, sorting each column in these rows, and building the histogram intervals. Presently, updating the statistics for a table is performed using two approaches, depending on the size of the sample set. If the sample is small enough, the entire sample is read into the memory and the data is sorted for each column. This is sometimes referred to as the "direct sample" method. If the sample is too large to be read into the memory, or too many columns are demanding histograms, the so-called "sample table" approach is employed, where the sampled rows are saved in a sample table. The sample table is then batch-processed one column, or one subset of columns, at a time. However, such an approach takes minutes, or even hours, to complete.

[0020] The large amount of time required to update the histogram statistics for the larger sample sets leads to a scalability problem as the amount of data in the EDW tables grows. A significant amount of time during the Extract-Transform-Load (ETL) phase of the EDW is spent building histograms. In some instances, the updating of the histogram statistics may not be completed in a given ETL window, resulting in the EDW operating with stale statistics. If the sampling rate is reduced so that the faster "direct sample" method can be employed, the statistics will be less accurate. Either way, the quality of the query plan provided by optimizer 50 suffers.

[0021] According to one example, optimizer 50 includes a histogram generator 60 which, in accordance with the present disclosure, employs a series of Counting Bloom Filters (CBFs) 62 to identify the unique values and the

frequencies thereof in a table column, and applies an optimized sort to select boundary values for and to construct an equal-height histogram. Histogram generator 60, in accordance with the present disclosure, constructs an equal-height histogram in times which are comparable to the "direct" sample approach while using substantially less memory space.

[0022] As described above, histogram generator 60 employs a series of CBFs, or CBF instances 62, to identify unique values and their occurrence frequencies in a table column. As will be described in greater detail below, according to one example, each CBF instance includes a sequential array of bins, each bin including a multi-bit (e.g. a 3-bit counter, a 4-bit counter, etc.), a list data structure, and an overflow table, and is configured by a set of parameters. In one example, the set of parameters includes a number of hash functions (k), an estimate of a number of unique values in the sample data set from the column of the table (n), an average frequency of occurrences of a value before it is moved to an overflow table (a), and a false positive probability (p). It is noted that each CBF instance may be uniquely defined and may include a different set of parameters, different number of hash functions and different hash functions, and bin arrays having differing numbers of bins, for example.

[0023] The parameter n is the number of distinct values in the sample data set. However, in one example, parameter n may be estimated by using an enumerated estimate of the distinct values from the pilot sample. In some instances, the parameter n may be estimated over multiple samples. According to some embodiment, the estimation of n may be carried out by using an estimation tool such as Jackknife, Schlosser or a hybrid combination thereof.

[0024] The parameter, n (i.e., an estimate of the number of unique values in a sampled data set, such as in the table column), is used by histogram generator 60 to determine the number of bins, m, in the sequential array of bins. If the number of bins is set too small, the CBF will under-identify the number of unique values in a table column (i.e. false positives will be high). If the number of bins is set too large, even though the probability of false positives produced by the CBF will be reduced, the sequential array of bins will be sparsely populated and memory space will have been wasted.

[0025] When operating to provide histogram statistics for a column of data, such as a column of data from table 42, histogram generator 60 obtains a sample data set from a number of rows of the table column. The number of rows in a table, such table 42, can vary from the hundreds to the billions. In one example, the number of rows in the sample data set is a percentage of the total number of rows in the table column. In one example, the percentage is 5% for a small table and 1% for a large table. In one example, the number of rows in the sample data set has a cap, such as at several million rows.

[0026] Histogram generator 60 creates a first CBF instance based on the parameters described above, including using the parameter "n" (i.e. an estimate of the number of unique values in the sample data set) to determine to the number of bins, m, in the array of bins. In one example, the number of bins, m, in the array is determined by the equation $m = (n \ln p)/(\ln 2)^2$; where p is the false positive probably of the CBF instance.

[0027] In one example, histogram generator 60 uses an estimate of the number of unique values, n, in the sample data set based on a number of unique values in the table column which is known from previous statistical analyses and is stored in statistics table 54, for instance. In one example, when the number of unique values in the table column is known from previous analyses, the number is scaled down or up in proportion to the size of the sample data set to determine the value of n. In one example, when a previously determined value of the number of unique values in the table column is not known, histogram generator 50 performs a pilot sampling of the sample data set and estimates the number of unique values in the pilot sample using an estimating tool, such as by using a Jackknife--Schlosser estimators and hybrid formulations thereof. The number of unique values from the pilot sample is then scaled up to estimate the number of unique values, n, for the sample data set.

[0028] In one embodiment, the pilot sample is a fixed number of rows of the sample data set, such as a range from 1,000 to 2,000 rows of the sample data set (e.g. the first 1,000 rows of the sample data set).

[0029] Figure 2 illustrates an example of a CBF instance 70 that can be used by histogram generator 60. CBF instance 70 includes an array 72 of bins, such as

bin 74, with the bins being sequentially labeled as bins 1 through m, with each bin having a multi-bit counter, such as multi-bit counter 76 of bin 74. Together, the counters of the array of bins 72 form an array of counters, each counter of the array of counters corresponding to a bin. In the illustrated example, each bin has a 4-bit counter indicated as being set at zero. According to one example, the number of bits histogram generator 60 employs for the counter for each bin is selected such that the counter can store at least an average frequency of the unique values as calculated from the estimation performed on the pilot sample. In addition to the bin array 72, and a parameter set, as described above, CBF instance 70 further includes a list data structure 78 and an overflow table 80, each of which be described in greater detail below.

[0030] To identify all unique values and their corresponding occurrence frequencies in the sample data set of the table column, histogram generator 60 begins by successively inserting the values of the sample data set into the first CBF instance. Each inserted value becomes the argument to each of the *k*-hash functions of the first CBF instance, such as CBF instance 70, with each of the *k*-hash functions mapping the inserted value to one of the *m* array positions or bins of the bin array 72 and incrementing the corresponding multi-bit counter accordingly.

[0031] It is noted that, initially, all of the multi-bit counters, such as multi-bit counter 76 of bin 74, is set to zero. If the counter of at least one bin to which the inserted value is mapped by the *k*-hash functions is incremented from "zero", the inserted value is determined to be a unique value, *v*, and is added to a list data structure of unique values associated with the CBF instance. If the counter of all the bins to which the inserted value is mapped by the *k*-hash functions already have non-zero values when incremented, the value is determined to already have been identified as being a unique value and is not added to the associated list data structure of the CBF instance.

[0032] Figures 3-5 illustrate an example of this process where CBF instance 70 includes three hash functions (k=3), referred to as $hash_1$, $hash_2$, and $hash_3$, and where all of the counters of the bin array are initially set to zero (as illustrated by Figure 2). Referring to Figure 3, initially, a first value, *x*, becomes the argument

to each of the three hash functions. According to the illustrated example, $hash_1(x)$ maps to bin number 2, $hash_2(x)$ maps to bin number 7, and $hash_3(x)$ maps to bin number m-1, with the corresponding counter of each bin being incremented by one. Since the counter of at least one of the bins to which the first value "*x*" is mapped is incremented from a value of "zero", the inserted value "*x*" is determined to be a unique value. As such, the unique value "*x*" is added to a list data structure, as illustrated at 78, with the counters of the bin array indicating the unique value "*x*" as having a frequency of "1".

[0033] At Figure 4, the next value "*y*" of the sample data set is inserted into each of the three hash functions of CBF instance 70. According to the illustrated example, $hash_1(y)$ maps to bin number 7, $hash_2(y)$ maps to bin number 4, and $hash_3(y)$ maps to bin 8, with the corresponding counter of each bin being incremented by one. Again, since the counter of at least one of the bins to which the value "*y*" is mapped is incremented from a value of "zero", the inserted value "*y*" is deemed to be a unique value. As such, the unique value "*y*" is added to the list data structure, as illustrated at 78, with the counters of the bin array indicating the unique value "*y*" as having a frequency of "1".

[0034] At Figure 5, the next value of the sample date set, which is again "*x*", is inserted into each of the three hash functions of CBF instance 70. As before, with reference to Figure 3, $hash_1(x)$ maps to bin number 2, $hash_2(x)$ maps to bin number 7, and $hash_3(x)$ maps to bin number m-1, with the corresponding counter of each bin being incremented by one. However, in this case, since the counter of each of the bins to which the value "*x*" are incremented from non-zero values, the value "*x*" is determined to have already been identified as being a unique value. As such, the value "*x*" is not added to the list data structure 78 as a new unique value, but the counters of the bin array now indicate that unique value "*x*" as having a frequency of "2" which, as will be described in greater detail below, is the minimal value of all counters associated with the hashes of the value "*x*".

[0035] According to one example, this process is repeated for until all values of the sample data set have been inserted into the CBF instance. However, according to one example, in order to avoid over-populating the bin array of a

single CBF instance, which increases the probability of the occurrence of false positives (resulting in the number of unique values being underestimated by the CBF), histogram generator 60 uses a series of CBF instances to determine the number of unique values and their corresponding occurrence frequencies in the sample data set of the table column.

[0036] For example, according to one embodiment, before a value from the sample data set is inserted into a current CBF instance being employed, such as CBF instance 70, histogram generator 60 determines the number of bins 74 of the bin array 72 having multi-bit counters with non-zero values. If the number of bins having counters with non-zero values is below a selected threshold, say 50%, for example, the current value from the sample data set is inserted into the current CBF instance. If the number of bins having counters with non-zero values is at or above the selected threshold, histogram generator 60 creates a new CBF instance and employs the new CBF instance as the "current" CBF instance, wherein the new current CBF instance has counters for all bins which are set at zero. The bin array 72 of the "old" CBF instance is maintained in its final state and holds the frequency information of the unique values identified from the sample data set as listed in the corresponding list data structure.

[0037] By creating and employing a new CBF instance as the current CBF instance when the number of bits set in the bin array of the previous CBF instance have reached a selected threshold, the occurrence of "false positives" (i.e. the incorrect determination that a value has already been identified) is reduced and a more accurate identification of the unique values and their frequencies from the sample data set is achieved.

[0038] According to one example, the new CBF instance is created using the same parameters as the previous, or "old", CBF instance. In one example, the parameter "n" (i.e. an estimate of the number of unique value in the sample data set of the table column) is re-estimated prior for the creation of the new CBF instance.

[0039] According to one embodiment, whenever more than one CBF instance exists, prior to inserting a next value from the sample data set into the current CBF instance, the value is first inserted into at least the most recent "old" CBF

instance. For example, if the current CBF instance is the third CBF instance being employed, before inserting the value into the third CBF instance, the value if first inserted into the second CBF instance (i.e. the most recent "old" CBF instance). If the value has already been identified by the most recent "old" CBF instance, the counters in the bins of the most recent "old" CBF instance to which the value is mapped are incremented (i.e. the frequency is incremented), and the value is not applied to the current CBF instance (as it has already been identified and been added to the list data structure of unique values of the most recent "old" CBF instance). If the value has not already been identified by the most recent "old" CBF instance, the value is inserted into the current CBF instance and processed accordingly.

[0040] According to one example, before inserting a value from the sample data set into the current CBF instance, the value is first inserted into and compared to the bin arrays of all old CBF instances. For example, if the current CBF instance is the third CBF instance being employed by histogram generator 60, before inserting the value into the third CBF instance, the value if first inserted into both the first and the second CBF instances.

[0041] In one example, the value is inserted into and compared to the bin array of only the most recent "old" CBF instance. For example, if the current CBF instance is the third CBF instance being employed by histogram generator 60, the value is inserted into and compared to the bin array of only the second CBF instance.

[0042] According to one example, values from the sample data set are inserted into only the current CBF instance and are not inserted into any old CBF instances.

[0043] According to one example, values from the sample data set are inserted into the most recent "old" CBF instance only when there is only one "old" CBF instance. For example, if the current CBF instance is the second CBF instance being employed by histogram generator 60, before inserting the value into the current CBF instance (i.e. the second CBF instance being employed), the value is first inserted into the most recent "old" CBF instance (i.e. the first CBF instance). However, once more than two CBF instances have been employed,

values from the sample data set are inserted only into the current CBF instance and are not inserted into any old CBF instances. For example, if the current CBF instance is the third CBF instance being employed by histogram generator 60, the value from the sample data set is inserted only into the current CBF instance (i.e. the third CBF instance) and is not inserted into any "old" previous CBF instances, including the most recent "old" CBF instance (i.e. the second CBF instance).

[0044] According to one example, overflow table 80 is employed to track the frequencies of values when each of the bins of the array of bins 72 to which the value has been mapped ("hashed") have counters which have reached their maximum value. For example, with reference to Figure 3, if the 4-bit counter of each of the bins "2", "7", and "*m-1*" were each already set to the maximum value of "1111" when the value "*x*" was mapped to them by $hash_1(x)$, $hash_2(x)$, and $hash_3(x)$, histogram generator 60 adds the value "*x*" to overflow table 80 and thereafter tracks the occurrence frequency "*f*" of value "*x*" via a corresponding counter (e.g. a 5-bit counter) in overflow table 80. By tracking the occurrence frequencies "*f*" in overflow table 80 for what is typically a relatively small number of frequently occurring values memory space is saved by not providing all counters of bin array with a bit capacity necessary to count the frequencies of such frequently occurring values, capacity that would likely go unutilized. For example, rather than providing all counters in the bin array with 5-bits, counters of the bin array are provided with only 4-bit counters, and only those values that have frequencies exceeding the capacity of a 4-bit counter are moved to overflow table 80 and tracked with counters having additional bit length.

[0045] After all values of the sample data set have been inserted into the series of CBF instances (a series of at least one), all unique values occurring in the sample data set will have been identified and saved to the corresponding list data structures 78 of each of the CBF instances, and the occurrence frequencies of all unique values will have been accumulated by and stored in the counters of the bin arrays and overflow tables of the series of CBF instances. Based on the identified unique values and their occurrence

frequencies in the sample data set, histogram generator 60 then determines equal-height histogram statistics for the table column.

[0046] In order to do so, histogram generator 60 employs the identified unique values from the list data structures of the CBF instances and extracts the frequency data of the identified unique values from the counters and overflow tables of the corresponding CBF instances. For example, according to one embodiment, for each CBF instance of the series of CBF instances (wherein the series includes at least one CBF instance) histogram generator 60 applies each unique value from the CBF instance's corresponding list data structure to the hash functions to determine the bins of the bin array to which the unique value is mapped. The histogram generator 60 then takes the lowest valued counter of the bins to which the unique value is mapped (including values in any overflow tables) as the frequency of the unique value. For example, referring to Figure 5, when determining the frequency of unique value "x", histogram generator 60 looks to bin number 2 (having a counter value of 0010), bin number 7 (having a counter value of 0011), and bin number m-1 (having a counter value of 0010). Histogram generator 60 takes the lowest counter value of these bins (in this case, either bin number 2 or bin number m-1) as the frequency of unique value "*x*", in this example, a frequency of 2.

[0047] Histogram generator 60 then employs the identified unique values and their occurrence frequencies to determine histogram statistics for the table column. According to one example, prior to generating the histogram statistics, histogram generator 60 encodes the identified unique values, *v*, so that each encoded value, encode(*v*), is a fixed width value, and then pairs the encoded value, encode(*v*) with the corresponding frequency value, *f*. In one example, the histogram generator 60 encodes each identified value, *v*, so that the encoded value, encode(*v*) is an 8-byte value. In one embodiment, for SQL integer data types, the encoded value is the value itself, while for SQL CHAR, DATE, and timestamp data types, the encoding is a floating point value (e.g. an 8-byte floating point value) which preserves the sort order of the value in the entire domain of the table column. For example, if value "*x*" is greater than value "*y*", encode(*x*) is greater than encode(*y*). Encoding the identified unique

values enables histogram generator 60 to work with the data in a compact form, thereby reducing an amount of required memory.

[0048] Histogram generator 60 then works with the encoded values, encode(v), paired with the corresponding frequencies, f, to determine equal-height histogram statistics for the table column. In one example, to generate equal-height histogram statistics, histogram generator 60 first forms an equal-width histogram with all pairs (encode(*v*), *f*), and then converts the equal-width histogram to an equal-height histogram.

[0049] Figure 6 is a graph 90 illustrating an example of an equal-height histogram, *W*. Histogram *W* has a number of intervals, *b*, illustrated as intervals I_{W1} to I_{Wb}. In one example, the number of intervals, b, is set to the number of intervals of the equal-height histogram constructed from histogram *W* (as will be described by Figure 7 below). In one example, the total number of intervals, *b*, is typically a value from 50 to 100. In graph 90, the x-axis is the magnitude of the unique values, in this example, the encoded values, encode(*v*), and the y-axis is the frequency of the unique values, in this example, the total frequency of each interval.

[0050] . In one example, a width of each interval, W_I, is a constant found from the equation $W_I = (max-min)/b$; where "*max*" is the encoded value of the largest identified unique value, "*min*" is the encoded value of the smallest identified unique value, and "*b*" is the number of intervals in *W*. The smallest identified unique value, *min*, is the lower end or lower boundary of the first interval I_{w1}, and the largest identified unique value , *max*, is the upper end or upper boundary of the last interval I_{Wb}.

[0051] In one example, to populate histogram W, the encoded values, encode(*v*), and corresponding frequencies, *f*, are inserted into one of the intervals I_W based on equation Interval# = $(encode(v)-min)/W_i$; where, as described above, W_I is the width of each interval. After all of the unique values, encode(*v*) are inserted into histogram W, for any two intervals of histogram *W*, say interval I_x and interval I_y, where $x < y$, any encoded value in interval I_x is less than any encoded value in interval I_y. The combined or total frequency is determined for all encoded values in each interval and represents the height of

each interval, as illustrated along the y-axis. Histogram generator 60 then constructs an equal-height histogram from the equal-width histogram *W*.

[0052] Figure 7 is a graph 92 illustrating an example of an equal-height histogram, *H*. Histogram *H* has a number of intervals, *b*, illustrated as intervals I_{H1} to I_{Hb}. As described above, the number of intervals, b, is a constant and is typically a value in a range from 50 to 100. In graph 92, the x-axis is the magnitude of the unique values, in this example, the encoded values, encode(*v*), and the y-axis is the frequency of the unique values, in this example, the combined or total frequency of each interval. Each interval I_H has a height, *h*, which is set to a desired number of rows for each interval (a desired total frequency of the encode(*v*) values for each interval).

[0053] From equal-width histogram, *W*, the total frequency of each interval I_W is known. According to one example, beginning with the first interval, I_{W1}, of equal-height histogram, *W*, histogram generator 60 sequentially transfers the encode(*v*) values of each interval I_W into equal-height histogram *H* to sequentially fill the intervals I_{H1} to I_{Hb} to the desired height, *h*, also optionally sorting the encode(v) values within each interval I_W sequentially from smallest to largest (e.g. no sorting is necessary when all values in I_W can be moved to I_H).

[0054] As an example of how histogram generator 60 populates histogram H from histogram W, let a current interval of histogram *H* being filled from histogram *W* be denoted as I_H , and a number of remaining occurrences of encode(*v*) values to fill I_H to the desired height, *h*, be denoted as n_H. For an interval I_W in equal-width histogram, *W*, if the total frequency *F* of interval I_W is less than n_H, histogram generator 60 moves all values from I_W to I_H, subtracts *F* from n_H, and proceeds to the next interval in histogram *W*.

[0055] If the total frequency *F* of interval *IW* is greater than n_H, the encoded values within interval I_W are split, with a portion of the encoded values within interval I_W used to complete the filling of interval I_H and the remaining portion being used to fill subsequent interval(s) of histogram *H*. To complete filling interval I_H, a histogram generator 60 performs a sort of all encode(*v*) values in interval I_W and moves all occurrences of the *k* smallest encode(*v*) values to interval I_H, where the sum of all occurrences of the *k* smallest encode(*v*) values

(i.e. the total frequency of the *k* smallest encode(*v*) values) is at least equal to n_H, and where the sum of all occurrences of the *k-1* smallest encode(*v*) values is less than n_H. The upper end of the range of values for interval I_H (an upper boundary value for interval I_H) is the largest encode(*v*) value of the k smallest encode(*v*) values, with the value of the *k+1* encoded(*v*) value in interval I_W being the lower boundary value for next interval, I_{H+1}, of equal-height histogram *H*.

[0056] Histogram generator 60 carries out the above described process until any remaining values of the encode(*v*) values of the last interval, I_{Wb}, of equal-width histogram *W* have been inserted into the last interval, I_{Hb}, of equal-height histogram *H*. According to one example, after populating equal-height histogram *H*, histogram generator 60 estimates, for each interval I_H, the number of unique values, the row count (the total frequency), and other parameters and scales them up for the entire table column.

[0057] Figure 8 is a flow diagram illustrating an example of a method 100 for managing a database according to the present disclosure. Method 100 begins at 102 where an initial CBF instance is created, the initial CBF instance having an array of counters, and hash functions which map an inserted value to the array of counters. At 104, the initial CBF instance is designated as a current CBF instance.

[0058] At 106, method 100 sequentially inserts each value of a sample data set of a table column (a column of data from a table in the database) into the hash functions of the current CBF instance, and increments the counters of the array of counters to which the value is mapped by the hash functions. However, prior to inserting each value of the sample data set, if a number of counters in the array of counters having non-zero values is at least a threshold level (e.g. 50% of the counter in the array), method 100 includes designating the current CBF instance as an old CBF instance, creating a new CBF instance having an array of counters and hash functions that map an inserted value to the array of counters, and designating the new CBF instance as the current CBF instance.

[0059] Figure 9 is a flow diagram illustrating an example of a method 200 for managing a database in accordance with the present disclosure. Method 200 begins at 202 with obtaining a sample data set from a table column of a

database. At 204, an initial CBF instance is created, such as CBF instance 70 illustrated by Figure 2, which includes an array of counters, a set of hash functions which map an inserted value to the array of counter, a list data structure, and an overflow table. The initial CBF instance is created based on a plurality of parameters, such as a number of unique values in the sample data set, a number of hash functions, a false positive probability, and an average frequency of occurrences of a value before it is moved to the overflow table.

[0060] According to one example, a number of counters in the array of counters is based on an estimated number of unique values in the table column of data. In one example, the estimated number of unique values in the table column is known and stored from previous estimates for the table column. In one example, the estimated number of unique values in the table column in not known and is determined by taking a pilot sample of the sample data set and estimating the number of unique values in the table column by applying an estimation tool, such as a Jackknife family-Schlosser or a hybrid combination estimator, for example.

[0061] At 206, the initial CBF instance is designated as a current CBF instance. At 208, method 200 begins a process of sequentially inserting each value of the sample into the current CBF instance, beginning by setting a first value of the sample data set as a selected value.

[0062] At 210, method 200 queries whether a number of counters in the array of counters of the current CBF instance having non-zero values is above a threshold level (such as 50%, for example)? If the answer to the query at 210 is "NO", method 200 proceeds to 212. If the answer to the query at 210 is "YES", method 200 proceeds to 214, where the current CBF instance is designated as an "old" CBF instance. At 216, a new CBF instance is created based on the plurality of parameters as described above, but can have a number of counters in the array of counters and different hash functions, for example, than the initial CBF instance (i.e. each CBF instance can be unique). At 218, the new CBF instance is designated as the current CBF instance, and method 200 proceeds to 212.

[0063] At 212, method 200 inserts the selected value from the sample data set in into the current CBF instance, and the selected value is mapped to the corresponding array of counters by the hash functions. At 220, method 200 queries whether each of the counters of the array of counters to which the selected value is mapped are at a maximum value. For example, a 4-bit binary counter having a value of "1111" (i.e. a value of "15"), would be at a maximum value. If the answer to the query at 220 is "NO", method 200 proceeds to 222.

[0064] If the answer to the query at 220 is "YES", method 200 proceeds to 224, where the selected value is moved to an overflow table, such as overflow table 80 (see Figure 2) and thereafter has its occurrence tracked via a corresponding counter in the overflow table. Method 200 then proceeds to 222.

[0065] At 222, the counters to which the selected value are incremented, including any counters that have been moved to the overflow table. Method 200 then proceeds to 226, where it is queried as to whether any of the counters to which the selected values have been mapped were incremented from a value of "zero". If the answer to the query at 226 is "NO", method 200 proceeds to 228. If the answer to the query at 226 is "YES", method 200 proceeds to 230 where the selected value is deemed to be a unique value, and is added to the list data structure of unique values corresponding to the current CBF instance (see Figures 2-5, for example). Method 200 then proceeds to 228.

[0066] In one example, which is not illustrated in Figure 9, prior to proceeding to 212 from both the "NO" query at 210 and the designation of a new CBF instance as a current CBF instance at 218, method 200 first queries whether there any old CBF instances. If the answer is "NO", method 200 proceeds to 212. If the answer is "YES", method 200 then inserts the selected value into the hash functions of each of the old CBF instances to determine whether the selected value has already been identified as being a unique value by any of the old CBF instances. If the selected value has not previously been identified by any of the old CBF instances, method 200 proceeds to 212. If the selected value has already been identified by an old CBF instance, the counters to which the selected value is mapped in the old CBF instance which previously identified the selected value are incremented and method 200 proceeds to 228.

[0067] At 228, method 200 queries whether the selected value of the sample data set is the last value of the sample data set. If the answer to the query at 228 is "NO", method 200 proceeds to 232, where the next value of the sample data set is set at the selected value, and then returns to the query at 210.

[0068] If the answer to the query at 228 is "YES", method 200 has completed processing the entire sample data set through the series of CBF filter instances (a series of at least one CBF instance), and proceeds to 234. At 234, method 200 determines all unique values and their occurrence frequencies based on the unique values listed in the list data structures and from the counts stored in the array of counters of the current and old CBF instances.

[0069] At 236, based on the complete list of unique values and occurrence frequencies thereof, method 200 constructs and equal-height histogram for the table column. According to one example, as illustrated and described above by Figures 6 and 7, method 200 first pre-sorts the identified values into an equal-width histogram during a first pass through the list of unique values and frequencies, and then converts the equal-width histogram to an equal-height histogram. In one example, method 200 estimates the number of unique values of entries (the unique entry count (UEC)), the row count, and other parameters for each interval in the equal-height histogram (intervals I_H in Figure 7), and scales up to estimate the parameters for the entire table column.

[0070] By using a series of CBF instances to identify unique values and occurrences in a sample data set of a table column, and by pre-sorting the identified unique values and frequencies into an equal-width histogram, and then converting the equal-width histogram to an equal-height histogram, according to examples of method of the present disclosure, the time and memory space required to provide equal-height histogram statistics for a database are improved relative to present techniques.

[0071] Figure 10 illustrates tables "A" and "B" compares the performance of methods according to the present disclosure relative to present techniques, such as the "Direct Sample" method and the "Sample Table" method, which were described earlier with respect to Figure 1. Table "A" illustrates performance in terms of "elapsed time", and Table "B" illustrates performance in

terms of memory usage as measured in bytes. In both tables, the performance is with regard to sets of 7 and 10 columns from a database table (such as table 42 of Figure 1) having 13 columns and containing over 40 million rows of data. From Tables "A" and "B", it can be seen that a method using a series of CBF instances to identify unique values and occurrences and constructing equal-height histograms from unique values and occurrence histories pre-sorted into an equal-width histogram, according to the present application, is at least as fast as the "Direct Sample" and "Sample Table" techniques, and uses less memory than the "Direct Sample" technique.

[0072] Figure 11 is schematic block diagram illustrating an example of system 300 in which embodiments of the present disclosure can be implement, such as those described and illustrated above by Figures 1-10, including system 30, DBMS 36, and optimizer 50 having histogram generator 60 with CBF instances 62. System 300 can include various systems and subsystems, and can be a personal computer, a laptop computer, a workstation, a computer system, and application specific integrated circuit (ASIC), a server, a server blade center, and a server farm, among others, for example.

[0073] System 300 includes a system bus 302, a processing unit 304, a system memory 306, memory devices 308 and 310, a communications interface 312 (e.g. a network interface), a communication link 314, a display 316 (e.g. a video screen), and an input device 318 (e.g. a keyboard, mouse). System bus 302 can be in communication with processing unit 304 and system memory 306. Additional memory device 308 and 310, such as a hard disk drive, server, stand-alone database, or other non-volatile memory, can also be in communication with system bus 302. System bus 302 interconnects processing unit 304, memory devices 306, 308, 310, communication interface 312, display 316, input device 318, and additional ports, such as a universal serial bus (USB) port, for example.

[0074] Processing unit 304 can be a computing device and can include an ASIC. Processing unit 304 can execute instructions for performing operations, such as database management operation in accordance with examples of histogram

generator 60 as described by the present application. Processing unit 304 can include a processing core.

[0075] Additional memory device 306, 308, and 310 can store data, programs, instructions, and any other information necessary to operate a computer. Memories 306, 308, and 310 can be implemented as a computer-readable media (integrated and/or removable), such a memory card, disk drive, compact disk, or server accessible over a network. Memory devices 306, 308, and 310 can serve as databases or data storage, such as data storage 38 of Figure 1.

[0076] Histogram generator 60 can be implemented as hardware, firmware, software, or any combination thereof. For example, all of part of histogram generator 60 can be implemented as a computer storage medium storing computer executable instructions to perform a method. Computer storage media includes RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices. Any such computer storage media may be part of system 300 and employed as memory devices 306, 308, and 310.

[0077] Although specific examples have been illustrated and described herein, a variety of alternate and/or equivalent implementations may be substituted for the specific examples shown and described without departing from the scope of the present disclosure. This application is intended to cover any adaptations or variations of the specific examples discussed herein. Therefore, it is intended that this disclosure be limited only by the claims and the equivalents thereof.

CLAIMS

1. A method of managing a database comprising:

creating an initial counting bloom filter (CBF) instance having an array of counters and hash functions that map an inserted value to the array of counters;

designating the initial CBF instance as a current CBF instance;

sequentially inserting each value of a sample data set of a table column into the hash functions of the current CBF instance and incrementing counters of the array of counters to which the value is mapped, wherein prior to inserting each value, when a number of counters of the array of counters having non-zero values is at least at a threshold level:

designating the current CBF instance as an old CBF instance;

creating a new CBF instance having an array of counters and hash functions that map an inserted value to the array counters; and

designating the new CBF instance as the current CBF instance.

2. The method of claim 1, wherein prior to inserting each value, when the number of counters of the array of counters having non-zero values is below the threshold level, continue employing the current CBF instance without designating the current CBF instance as an old instance, without creating a new CBF instance, and without designating the new CBF instance as the current CBF instance.

3. The method of claim 1, further including:

identifying unique values and corresponding frequencies in the column table based on counter values of the counter arrays of the current and old CBF instances.

4. The method of claim 3, further including:

deeming a value of the sample data set to be a unique value when at least one of the counters to which the value is mapped by the hash functions is incremented from zero; and

placing the unique value in a list data structure corresponding to the current CBF instance

5. The method of claim 4, wherein after all values of the sample data set have been inserted into a CBF instance, the method further including:

for each list data structure, sequentially inserting each unique value into the hash functions of the corresponding CBF instance and taking the lowest value of the counters to which the unique value is mapped as the frequency of the unique value;

identifying all unique values and corresponding frequencies in the sample data set; and

extrapolating from the unique values and corresponding frequencies in the sample data set to determine frequencies of the unique values in the table column using an estimation tool.

6. The method of claim 3, further including constructing an equal-height histogram for the table column from the identified unique values and corresponding frequencies.

7. The method of claim 6, including:

constructing an equal-width histogram by sorting the identified unique values into a number of intervals, each interval having a value range of a same magnitude, such that when the intervals are taken in ascending order, any unique value in a given interval is greater than any unique value in an immediately preceding interval; and

converting the equal-width histogram to the equal-height histogram, the equal-height histogram having the same number of intervals as the equal-width histogram, by successively redistributing the unique values from the intervals of the equal-width histogram into intervals of the equal-height histogram such that

all occurrences of a given unique value are included in a same interval and such that a total number of unique values in each interval of the equal-height histogram, except for the final interval, is at least equal to a desired number.

8. The method of claim 7, prior to sorting the unique values to the equal-width histogram, encoding the unique values so that each unique value has a same length.

9. The method of claim 1, wherein when each counter of an array of counters to which a value is mapped is at a maximum value, the method including:

inserting the value in an overflow table of the CBF instance; and

incrementing a corresponding counter in the overflow table to track the occurrence frequency of the value thereafter.

10. The method of claim 1, wherein the number of counters in each array of counters of each CBF instance is based on an estimated number of unique values in the sample data set, and wherein the method includes:

taking a pilot sample of the sample data set;

applying an estimating tool to the pilot sample to estimate the number of unique values in the table column from the pilot sample.

11. A non-transitory computer readable storage medium having a set of machine readable instructions that, when executed, cause a database management system to:

receive a list of unique values and corresponding occurrence frequencies present in a sample data set of a table column of the database;

construct an equal-width histogram by sorting the identified unique values into a number of intervals, each interval having a value range of a same magnitude, such that when the intervals are taken in ascending order, any unique value in a given interval is greater than any unique value in an immediately preceding interval; and

convert the equal-width histogram to the equal-height histogram, the equal-height histogram having the same number of intervals as the equal-width histogram, by successively sorting the unique values from the intervals of the equal-width histogram into intervals of the equal-height histogram such that all occurrences of a given unique value are included in a same interval and such that a total number of unique values in each interval of the equal-height histogram, except for the final interval, is at least equal to a desired number.

12. The non-transitory computer readable storage medium of claim 11, the set of machine readable instructions further causing the database management system to encode the unique values in the list into fixed width values prior to constructing the equal-width histogram, wherein a lower end of the value range for the first interval of the equal-width histogram is that of the minimum unique value, and the upper end of the value range of the last interval of the equal-width histogram is that of the maximum unique value.

13. A database management system comprising:

a memory device to store computer executable instructions; and

a processing unit to access the memory device and execute the computer executable instructions, the computer executable instructions providing instructions to:

create an initial counting bloom filter (CBF) instance having an array of counters and hash functions that map an inserted value to the array of counters;

designate the initial CBF instance as a current CBF instance;

sequentially insert each value of a sample data set of a table column into the hash functions of the current CBF instance and incrementing counters of the array of counters to which the value is mapped, wherein prior to inserting each value, when a number of counters of the array of counters having non-zero values is at least at a threshold level:

designate the current CBF instance as an old CBF instance;

create a new CBF instance having an array of counters and hash functions that map an inserted value to the array counters; and

designate the new CBF instance as the current CBF instance.

14. The database management system of claim 13,the computer executable instructions providing instructions to: wherein prior to inserting each value into the hash functions of the current CBF instance, when there is more than one old CBF instance to insert the value into the current CBF, and when there is only one old CBF instance:

to insert the value into the hash functions of the old CBF instance to determine if the value has already been identified as being unique by the old CBF instance, and if so, incrementing counters of the array of counters of the old CBF instance to which the value is mapped by the hash functions and not inserting the value into the current CBF instance, and when not already identified as being unique by the old CBF instance to insert the value into the current CBF instance.

15. The database management system of claim 13, wherein the computer executable instruction further provide instructions to:

construct an equal-width histogram by sorting the identified unique values into a number of intervals, each interval having a value range of a same magnitude, such that when the intervals are taken in ascending order, any unique value in a given interval is greater than any unique value in any preceding interval; and

convert the equal-width histogram to the equal-height histogram, the equal-height histogram having the same number of intervals as the equal-width histogram, by successively sorting the unique values from the intervals of the equal-width histogram into intervals of the equal-height histogram such that all occurrences of a given unique value are included in a same interval and such that a total number of unique values in each interval of the equal-height histogram, except for the final interval, is at least equal to a desired number.

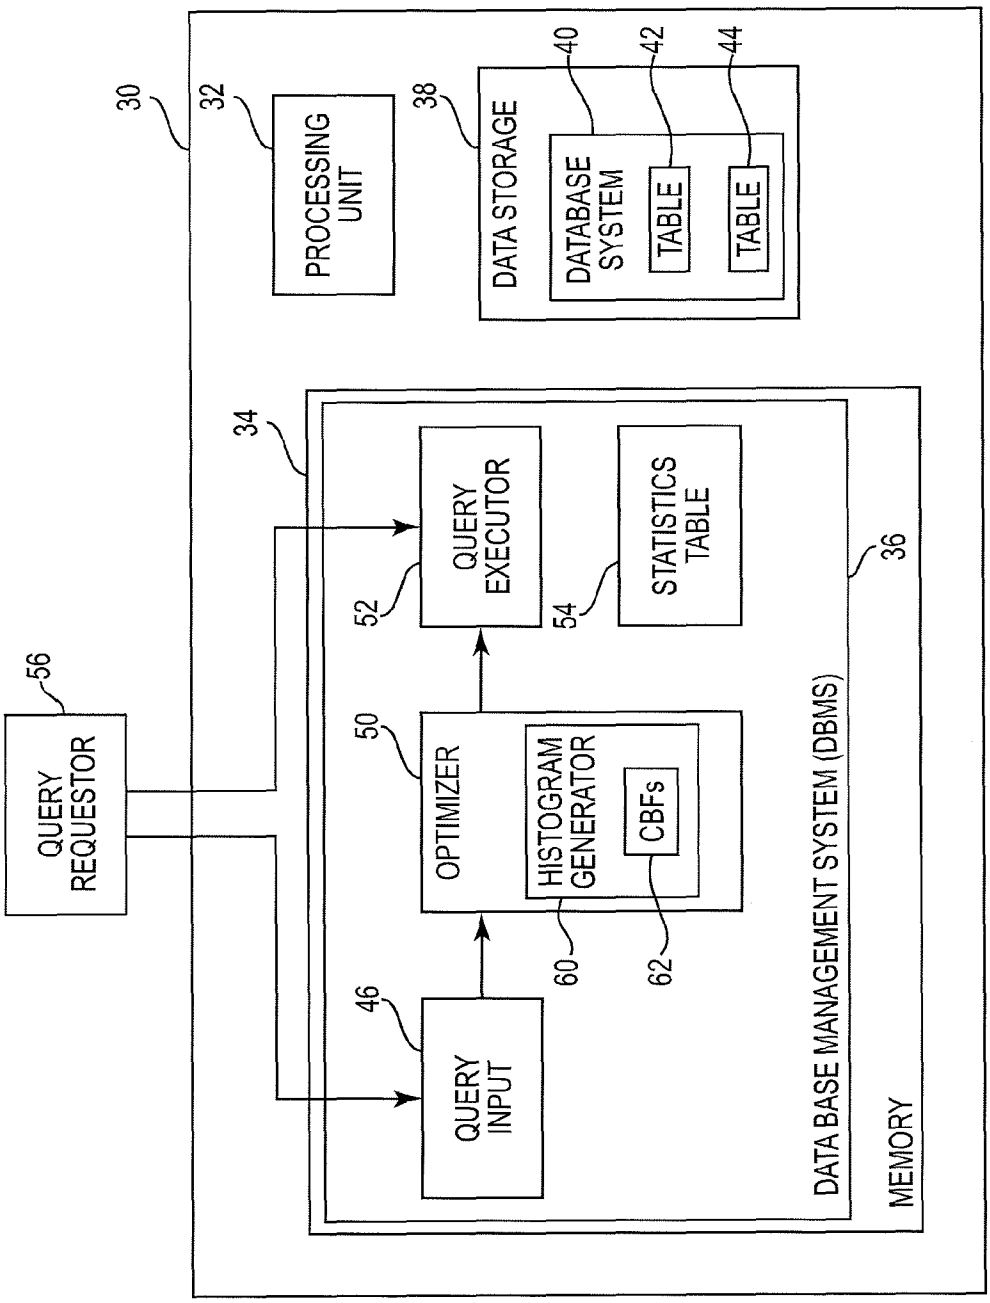

30
32
PROCESSING UNIT
38
DATA STORAGE
40
DATABASE SYSTEM
42
TABLE
44
TABLE
34
QUERY EXECUTOR
52
STATISTICS TABLE
54
36
56
QUERY REQUESTOR
50
OPTIMIZER
HISTOGRAM GENERATOR
CBFs
60
62
46
QUERY INPUT
DATA BASE MANAGEMENT SYSTEM (DBMS)
MEMORY

Fig. 1

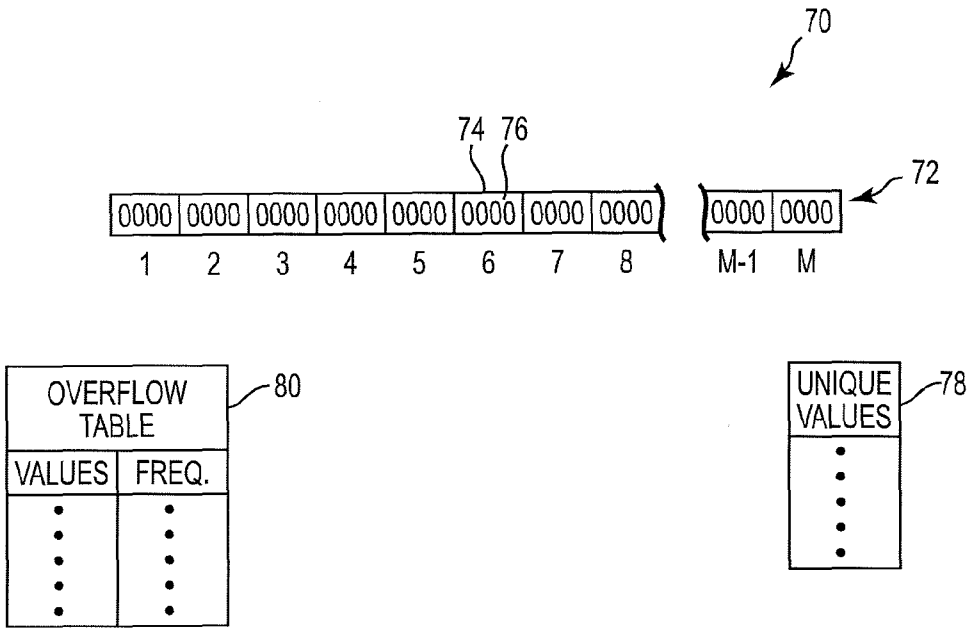

70
74
76
72
0000
0000
0000
0000
0000
0000
0000
0000
0000
0000
1
2
3
4
5
6
7
8
M-1
M
OVERFLOW TABLE
80
VALUES
FREQ.
UNIQUE VALUES
78

Fig. 2

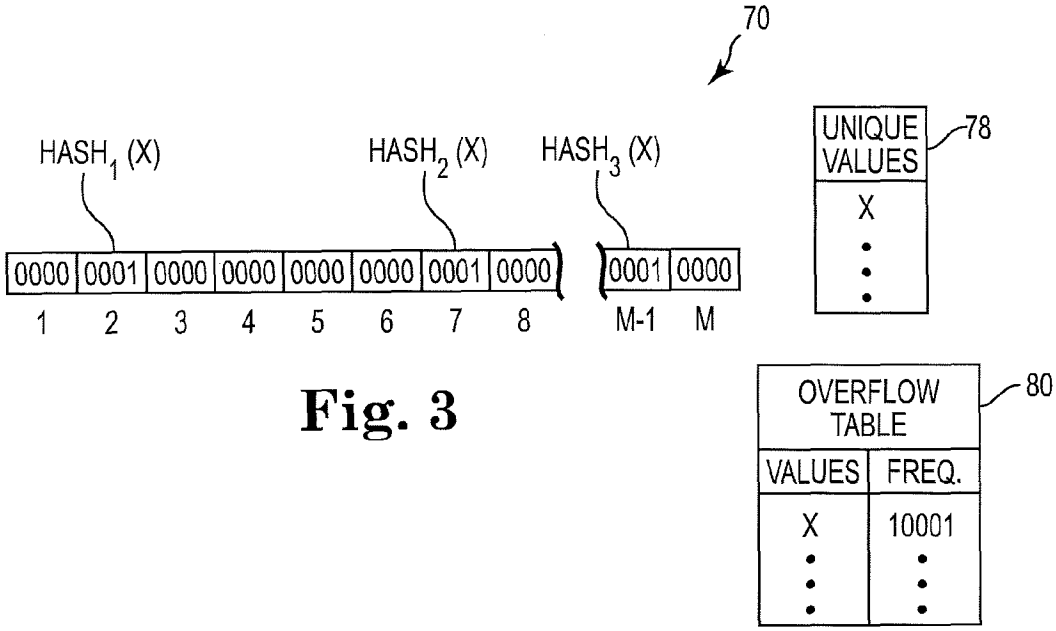


Fig. 3

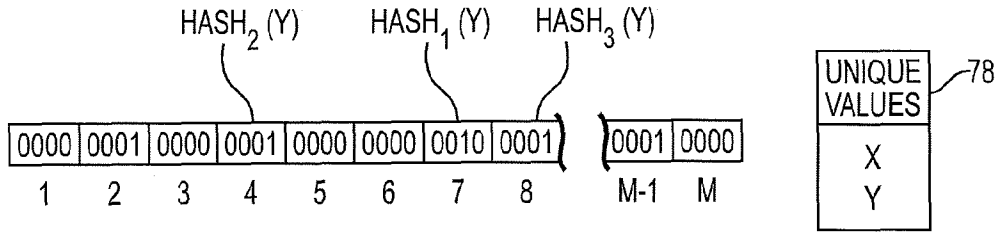


Fig. 4

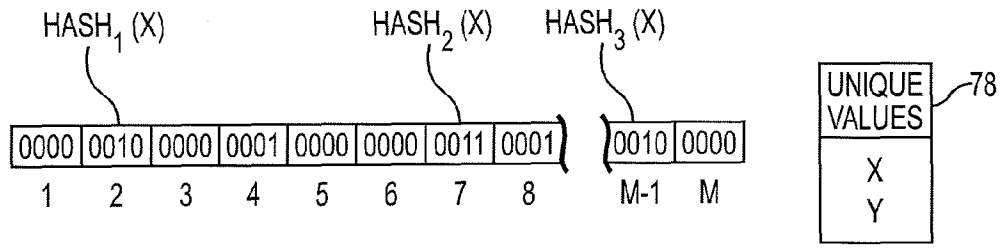


Fig. 5

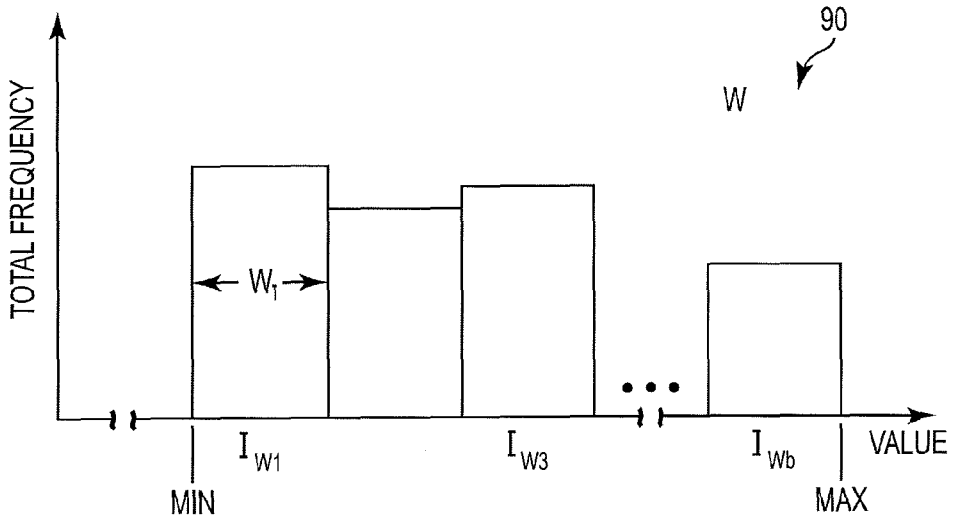

90
W
TOTAL FREQUENCY
W1
I W1
I W3
I Wb
VALUE
MIN
MAX

Fig. 6

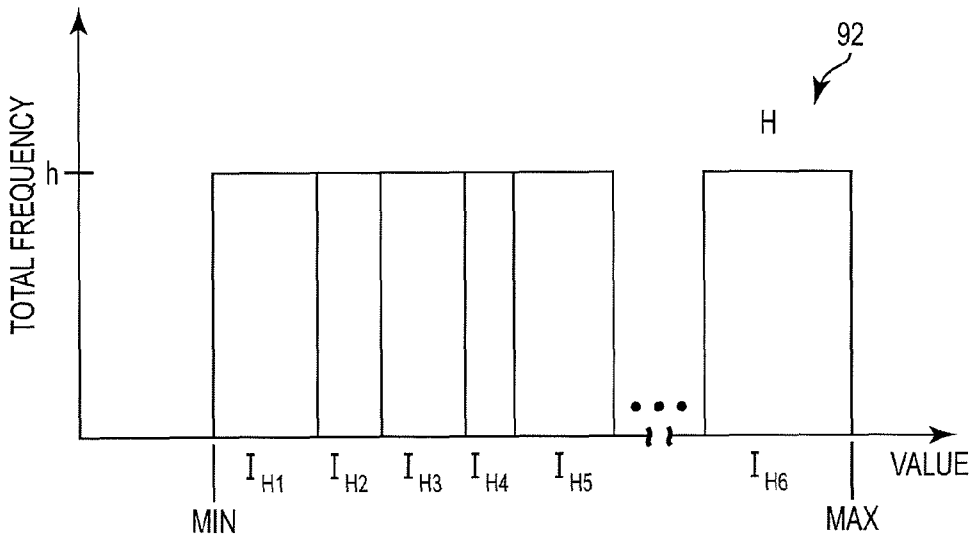

92
H
TOTAL FREQUENCY
h
I H1
I H2
I H3
I H4
I H5
I H6
VALUE
MIN
MAX

Fig. 7

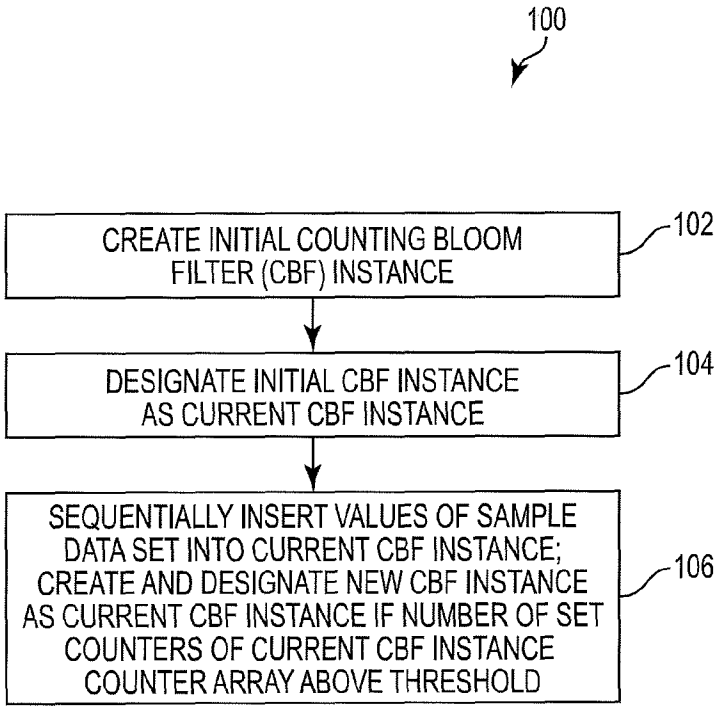

100
CREATE INITIAL COUNTING BLOOM FILTER (CBF) INSTANCE
102
DESIGNATE INITIAL CBF INSTANCE AS CURRENT CBF INSTANCE
104
SEQUENTIALLY INSERT VALUES OF SAMPLE DATA SET INTO CURRENT CBF INSTANCE; CREATE AND DESIGNATE NEW CBF INSTANCE AS CURRENT CBF INSTANCE IF NUMBER OF SET COUNTERS OF CURRENT CBF INSTANCE COUNTER ARRAY ABOVE THRESHOLD
106

Fig. 8

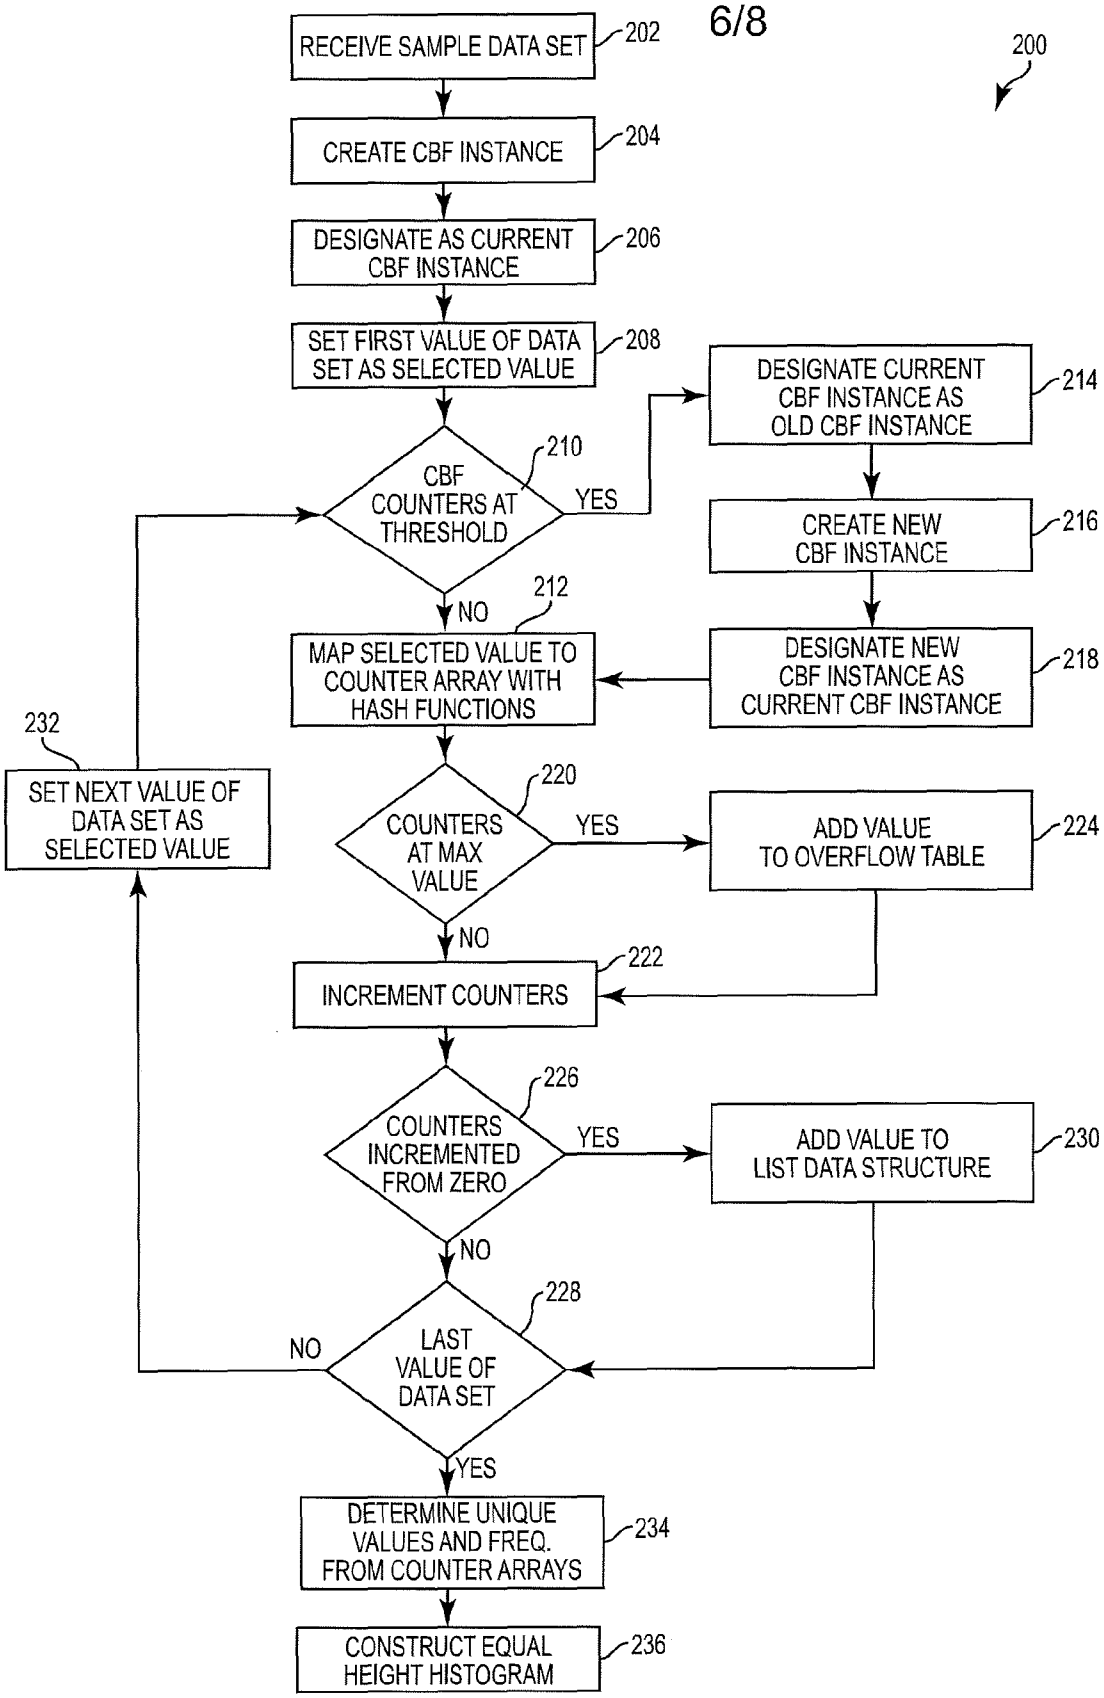

200
RECEIVE SAMPLE DATA SET
202
CREATE CBF INSTANCE
204
DESIGNATE AS CURRENT CBF INSTANCE
206
SET FIRST VALUE OF DATA SET AS SELECTED VALUE
208
CBF COUNTERS AT THRESHOLD
210
YES
NO
DESIGNATE CURRENT CBF INSTANCE AS OLD CBF INSTANCE
214
CREATE NEW CBF INSTANCE
216
DESIGNATE NEW CBF INSTANCE AS CURRENT CBF INSTANCE
218
212
MAP SELECTED VALUE TO COUNTER ARRAY WITH HASH FUNCTIONS
232
SET NEXT VALUE OF DATA SET AS SELECTED VALUE
COUNTERS AT MAX VALUE
220
YES
ADD VALUE TO OVERFLOW TABLE
224
NO
INCREMENT COUNTERS
222
COUNTERS INCREMENTED FROM ZERO
226
YES
ADD VALUE TO LIST DATA STRUCTURE
230
NO
LAST VALUE OF DATA SET
228
NO
YES
DETERMINE UNIQUE VALUES AND FREQ. FROM COUNTER ARRAYS
234
CONSTRUCT EQUAL HEIGHT HISTOGRAM
236

Fig. 9

TABLE A

DATA SET	DIRECT SAMPLE METHOD	SAMPLE TABLE METHOD	EXAMPLE PRESENT METHOD
SET 1 (7 COLUMNS)	3.16 SEC.	17.81 SEC.	2.8 SEC.
SET 2 (10 COLUMNS)	5.29 SEC.	27.03 SEC.	6.0 SEC.

TABLE B

DATA SET	DIRECT SAMPLE METHOD	EXAMPLE OF PRESENT METHOD	PERCENT REDUCTION
SET 1 (7 COLUMNS)	32,734,484 BYTES	3,827,396 BYTES	88.3%
SET 2 (10 COLUMNS)	53,175,252	39,649,239	35.4%

Fig. 10

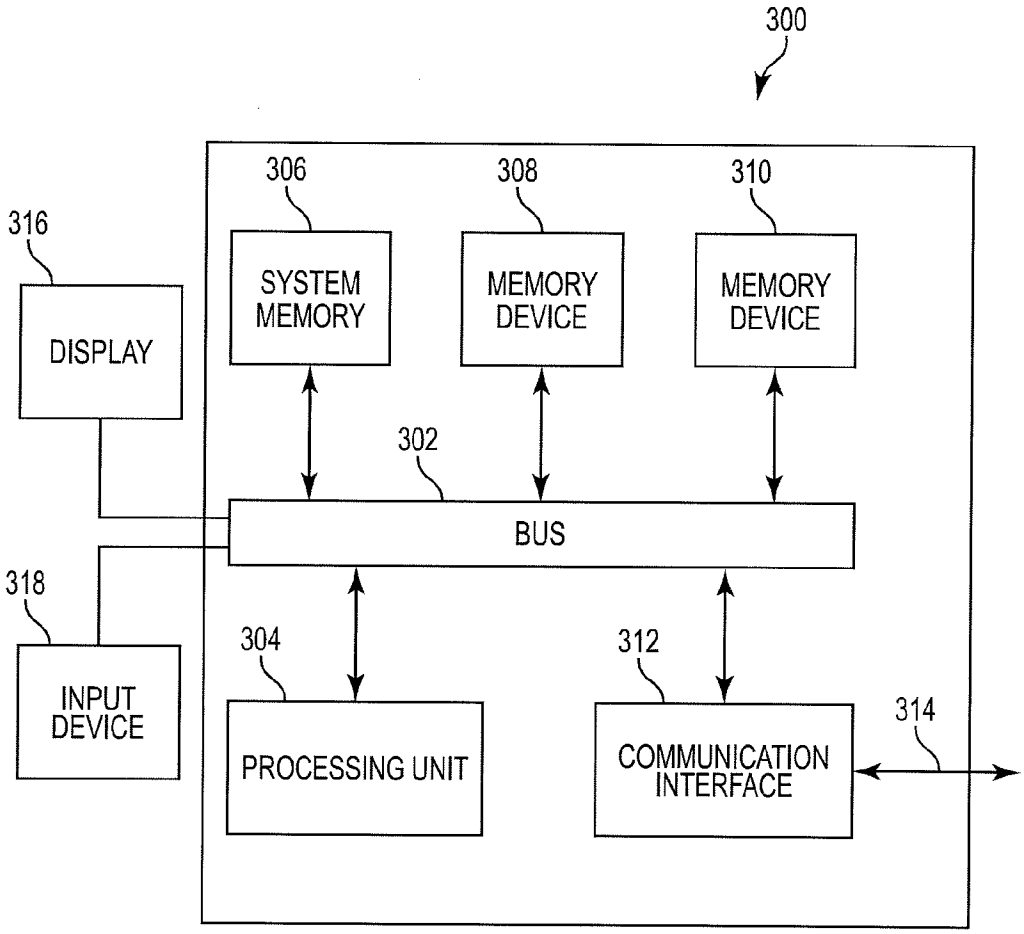

300
306
308
310
316
SYSTEM MEMORY
MEMORY DEVICE
MEMORY DEVICE
DISPLAY
302
BUS
318
304
312
314
INPUT DEVICE
PROCESSING UNIT
COMMUNICATION INTERFACE

Fig. 11

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2014/014329

A. CLASSIFICATION OF SUBJECT MATTER

G06F 17/30(2006.01)i

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F 17/30; G06F 7/00; H04L 12/26

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: counting bloom filter, threshold, create, construct, histogram, width and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	SABORIT, J. AGUILAR et al., "Dynamic Count Filters," ACM SIGMOD Record, Vol. 35, No. 1, pp. 26-32, March 2006 See pages 28-29 and figure 3.	1-10,13-15
A	LINARI, ALESSANDRO et al., "Efficient Peer-to-Peer Semantic Overlay Networks based on Statistical Language Models," In: Proceedings of the international workshop on Information retrieval in peer-to-peer networks, P2PIR ' 06, pp. 9-16, 11 November 2006 See page 12.	11-12
A	US 2013-0226972 A1 (KOSURU, RAMAKUMAR et al.) 29 August 2013 See paragraphs [0042]-[0043]; claims 1-15; and figure 7.	1-10,13-15
A	US 2008-0313132 A1 (HAO, FANG et al.) 18 December 2008 See paragraphs [0036]-[0042] and figure 3.	1-10,13-15
A	US 2011-0069632 A1 (CHEN, AIYOU et al.) 24 March 2011 See paragraphs [0048]-[0055] and figure 3.	1-10,13-15

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents,such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
15 October 2014 (15.10.2014)	**15 October 2014 (15.10.2014)**

Name and mailing address of the ISA/KR
International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea
Facsimile No. +82-42-472-7140

Authorized officer
NHO, Ji Myong
Telephone No. +82-42-481-8528

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2014/014329

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

Group I. Claims 1-10 and 13-15 directed to methods and systems of managing a database using new and old CBF instances.
Group II. Claims 11-12 directed to non-transitory computer readable storage medium having a set of machine readable instructions to make a total number of unique values in each interval of an equal-height histogram be equal to a desired number.

1. ☐ As all required addtional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. ☒ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of any additional fees.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.
PCT/US2014/014329

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013-0226972 A1	29/08/2013	None	
US 2008-0313132 A1	18/12/2008	US 7930547 B2	19/04/2011
US 2011-0069632 A1	24/03/2011	US 8134934 B2	13/03/2012