

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

(19) World Intellectual Property Organization

International Bureau



(10) International Publication Number

WO 2016/044587 A1

(43) International Publication Date

24 March 2016 (24.03.2016)

(51) International Patent Classification:

G06K 9/00 (2006.01)

(21) International Application Number:

PCT/US2015/050677

(22) International Filing Date:

17 September 2015 (17.09.2015)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

14/488,867 17 September 2014 (17.09.2014) US

(71) Applicant: CIRCONUS, INC. [US/US]; 11830 West Market Street, Suite F, Fulton, Maryland 20759 (US).

(72) Inventor: SCHLOSSNAGLE, Theo Ezell; 11301 Liberty Street, Fulton, Maryland 20759 (US).

(74) Agent: KAMINSKI, Jeffri A.; Venable, LLP, P.O. Box 34385, Washington, District of Columbia 20043-9998 (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, 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, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: EFFICIENT TIME-SERIES HISTOGRAMS

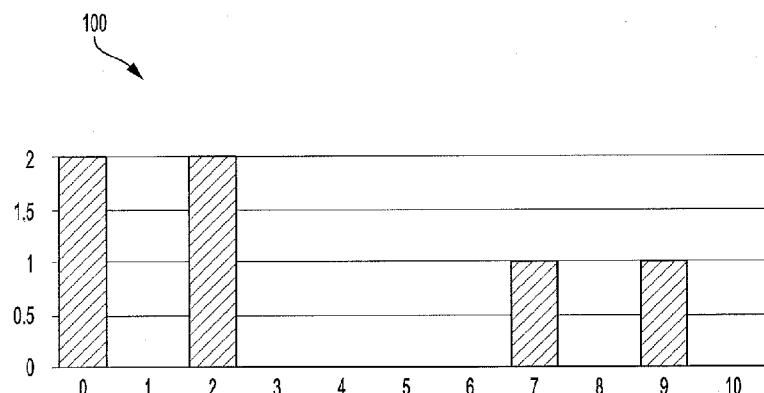


FIG. 1

(57) Abstract: A method for constructing histograms may include receiving a set of data values, wherein the data values are two-dimensional or three-dimensional and the data values contain at least one measurement over time; creating a plurality of bins in memory, wherein the plurality of bins are log-linear bins using 2 significant digits base 10; creating a zero bin representing the number 0; placing non-zero numbers of the set of data values into the plurality of bins; and generating one or more histograms using the bins.

WO 2016/044587 A1

EFFICIENT TIME-SERIES HISTOGRAMS

FIELD OF INVENTION

[0001] The present invention relates to efficient storing and displaying of time-series histogram data.

BACKGROUND OF THE INVENTION

[0002] A histogram is a graphical representation of data samples. In conventional settings, however, storing and retrieving time-series histograms may be inefficient. For example, when current systems retrieve time-series histograms, the histogram information stored in the system will often exceed the visual space needed to visualize the data. The data may exceed visual space over the time (x) axis, over the bin (y) axis, or both. Accordingly, efficient techniques are needed for storing and retrieving time-series histogram data.

SUMMARY

[0003] Embodiments of the invention may include a system, method, and/or computer-readable medium. In one embodiment, a method may exist for constructing histograms may comprise: receiving, by one or more processors, a set of data values, wherein the data values are two-dimensional or three-dimensional and the data values contain at least one measurement over time; creating, by the one or more processors, a plurality of bins in memory, wherein the plurality of bins are log-linear bins using 2 significant digits base 10; creating, by the one or more processors, a zero bin representing the number 0; placing, by the one or more processors, non-zero numbers of the set of data values into the plurality of bins; and generating, by the one or more processors, one or more histograms using the bins.

[0004] In another embodiment, a non-transitory computer-readable medium may exist comprising instructions, which when executed by one or more processors causes the one or more processors to perform operations for constructing histograms, the computer-readable medium comprising instructions to: receive a set of data values, wherein the data

values are two-dimensional or three-dimensional and the data values contain at least one measurement over time; create a plurality of bins in memory, wherein the plurality of bins are log-linear bins using 2 significant digits base 10; create a zero bin representing the number 0; place non-zero numbers of the set of data values into the plurality of bins; and generate one or more histograms using the bins.

5 [0005] In yet another embodiment, a system for constructing histograms may exist comprising: one or more processors for executing a plurality of instructions; a display device in communication with the one or more processors; and a storage device in communication with the one or more processors, the storage device holding the plurality of instructions, the plurality of instructions including instructions to: receive a set of data values, wherein the data values are two-dimensional or three-dimensional and the data values contain at least one measurement over time; create a plurality of bins in memory, wherein the plurality of bins are log-linear bins using 2 significant digits base 10; create a zero bin representing the number 0; place non-zero numbers of the set of data values into the plurality of bins; generate one or more histograms using the bins; and display the one or more histograms via the display device.

BRIEF DESCRIPTION OF THE DRAWINGS

10 [0006] The foregoing and other features and advantages of the invention will be apparent from the following, more particular description of various embodiments, as illustrated in the accompanying drawings wherein like reference numbers generally indicate identical, functionally similar, and/or structurally similar elements. The first digits in the reference number indicate the drawing in which an element first appears. Unless otherwise indicated, the accompanying drawing figures are not to scale.

15 [0007] FIG. 1 depicts a sample histogram of eleven units;

20 [0008] FIG. 2 depicts a sample time-series, with the measurement value on the y-axis and a time of measurement on the x-axis;

25 [0009] FIG. 3A-3D depict histograms using color density;

30 [00010] FIG. 4 depicts side-by-side vertical histograms representing measurements over time;

[00011] FIG. 5 depicts example fixed bin sizes;

[00012] FIG. 6 depicts an example of multiple bin aggregation;

[00013] FIG. 7 depicts an example flowchart describing processing performed in an illustrative embodiment of the present invention; and

[00014] FIG. 8 depicts an example computer processing system that may be used in

5 implementing an embodiment of the present invention.

DETAILED DESCRIPTION

[00015] Illustrative embodiments are discussed in detail below. While specific embodiments are discussed, it should be understood that this is done for illustration purposes only. In describing and illustrating the embodiments, specific terminology is employed for the sake of clarity. However, the embodiments are not intended to be limited to the specific terminology so selected. A person skilled in the relevant art will recognize that other components and configurations may be used without departing from the spirit and scope of the embodiments. It is to be understood that each specific element includes all technical equivalents that operate in a similar manner to accomplish a similar purpose. The examples and embodiments described herein are non-limiting examples.

[00016] Furthermore, the embodiments detailed below may be combined into a new embodiment and/or various features of the embodiments described below may be combined to form a new embodiment.

20 [00017] All publications cited herein are hereby incorporated by reference in their entirety. As used herein, the term “a” refers to one or more. The terms “including,” “for example,” “such as,” “e.g.,” “may be” and the like, are meant to include, but not be limited to, the listed examples.

25 [00018] FIG. 1 depicts a sample histogram of eleven units or bins. A histogram is a dense representation of data samples where samples themselves may be grouped together (losing accuracy). For example, given a set of numbers 0.5, 0.7, 2.0, 2.9, 7 and 9.2, one may divide the space from 0 to 10 into units of 1. The histogram of FIG. 1 depicts 11 “bins” of size 1: [0,1), [1,2), [2,3), [3,4), [4,5), [5,6), [6,7), [7,8), [8,9), and [9,10). The number and size of the bins may change. There are two numbers in each of [0,1) and [2,3), and one number in each of [7,8) and [9,10) as shown in FIG. 1.

[00019] FIG. 2 depicts a sample time-series, with a measurement on the y-axis and a time the measurement was taken on the x-axis. A time series is a set of measurements over time. In this example, it is measured how long a specific action took (e.g. loading a web page) every minute for 30 minutes. A common visual representation (line graph) would be to place the measurement on the y-axis and the time on the x-axis as shown in FIG. 2.

5

[00020] Histograms may also display three dimensions of data in a two dimensional view. For example, time may be the x-axis, a measurement value may be the y-axis, and a third dimensions may be represented using color density....

[00021] FIG. 3A depicts a histogram using color density to represent values. There are 10 eighteen bins along the axis. The frequency the values appear in the bins is along with y-axis. Large populations in a given bin result in a more dense (e.g., darker) color representation.

[00022] FIG. 3B depicts another histogram using color density to represent values. FIG 15 3B depicts an example of a histogram using color density to represent values without using bar height. Thus, the frequency is removed from the y-axis. Population can be understood based on color density without explicit bar height.

[00023] FIG. 3C depicts the histogram of FIG. 3B where height is reduced to compress space.

[00024] FIG. 3D depicts the reduced height histogram of FIG. 3C rotated counter- 20 clockwise 90 degrees as a “vertical histogram.”

[00025] FIG. 4 depicts side-by-side vertical histograms representing measurements over time. For example, once the vertical color-density transformation of FIGS. 3A-3D is applied to a histogram, many vertical histograms over time may be represented by placing them side-by-side and a large amount of data may be visually displayed in a

25 limited space. This may be one technique to visually display a large amount of data in a limited space.

[00026] In conventional settings, storing and retrieving (e.g., for visualization) time-series histograms may be inefficient. For example, when retrieving time-series histograms the histogram information stored in the system often exceeds the visual space available to 30 visualize the data. The data may exceed the visual space over the time (x) axis or it may exceed the visual space over the bin (y) axis, or both.

[00027] For example, if histograms are stored on a one minute basis, over a two year period more than one million histograms will be stored. In a limited visual space, such as a computer monitor or other display, where only some small number of time-series histograms can be displayed, aggregation of the one-minute histograms into a larger time interval for visualization may be required. Assuming, for example, the visual space can display 1,500 vertical histograms. Then, groups of 720 one-minute histograms may be aggregated into a single 12-hour-histogram. The number of 12-hour-histograms over 2 years is approximately 1,460 which may fit within the designated visual space.

5 Accordingly, some of the data is not required for visualization and transmitting and reducing histograms that cannot be discerned when visualized is inefficient.

[00028] In another example, a data set may have measurements that are very small and very big such that differentiating small values in a visualization is not possible (e.g., measurements 1,000,000, 900,000, 0.1, and 0.13). Given limited visual space, a linear representation of values, 0.1 and 0.13 may be represented too close visually to be discerned. Accordingly, the difference between the 0.1 and 0.13 values may be indistinguishable as opposed to the difference between 1,000,000 and 900,000.

15 Transmitting bins uniquely that cannot be discerned when visualized may be inefficient.

[00029] Time-Realm Efficiency

[00030] Previously, data being collected was analyzed to optimize for error reduction and clean visualization of a single histogram and not a series of histograms over time. Accordingly, previous approaches lacked a consistent approach to binning. Consistency in binning was not required as histograms were not aggregated into larger time units.

20 [00031] However, in one embodiment, time-realm efficiency may be achieved. Bin selection may be the same to allow, for example, aggregating two histograms.

25 Accordingly, bin measurements may maintain (1) a fixed size bin description, (2) a universal compatibility regardless of measurement magnitude while maintaining a satisfactory and useful level of precision, and (3) an ability to aggregate while controlling bin space.

[00032] A fixed size bin description may be accomplished by sizing the bin [begin, end] into a fixed bit-width (e.g., 16 bits) representation. The fixed-bit width may, for example, allow for more efficient key lookup. The bin size may be fixed (e.g., two significant

5 digits base 10) regardless of the data values. This creates a uniform bin scheme that may be used, for example, across different data streams and keeps the same key space. In one embodiment, a fixed keyspace may be provided. The fixed keyspace is above and beyond fixed bins and may be a limitation to the number of total bins. The fixed key space may represent buckets between 1.0×10^{-127} and 9.9×10^{128} , 0, and between -
1.0x10^-127 and -9.9x10^128 and thereby limits the keys represent into 16 bits.

10 [00033] The fixed keyspace allows for aggregating, for example, the 720 one-minute histograms into a single 12-hour histogram both (1) efficiently and (2) in a fashion that tracks error consistently. Histogram may be aggregated more efficiently because the keyspaces between the histograms are identical. For example, two histograms may be simply be added together (e.g., bin by bin). If the keyspaces were different, mathematical conversions would be required and the effects on error would be even more complicated. Additionally, the uniform binning across various data streams also provides for highly compressible histograms. For example, restricted use of keyspace leads to better
15 compression.

20 [00034] Combining histograms with different bin selections may be an issue because different bin selections that overlap but are not equal have different error introduced when samples are placed within them. Previously, when a new aggregate bin was created, a technique to re-bin the samples using some selected distribution assumption must be performed. The technique may not be possible if the bins are either really large
25 (e.g., not useful) or they are small but samples can be entirely misplaced (e.g., the possible error would exceed the bin itself).

25 [00035] In one embodiment, universal compatibility regardless of measurement magnitude while maintaining a satisfactory and useful level of precision may be accomplished through the following. First, the number zero may be treated as an exact bin (e.g., all measurements that equal zero). A bin zero may be defined as containing all measurements that share equality with the value zero. Non-zero numbers may be represented in bins that are bounded numbers with two significant digits in base ten. This has characteristics of representing numbers of wide magnitude space (i.e., 10^{-128} to
30 10^{127}) logarithmically, while maintaining a linear granularity within each bin set aiding human understanding and visualization (e.g., a large set of bins such as, [1.0, 1.1), [1.1,

1.2),...[9.9,10) collectively times 10^n). Bins may be log-linear bins using 2 significant digits base 10, and, therefore, the maximal error is 5% and the average error is approximately 0.65%. As the base or significant digits changes, the error would change as well.

5 [00036] Bins may be from one point on a discrete number line to the next, with a closed interval on the point closest to zero and an open interval on the number furthest from zero (e.g. [1.7,1.8) and [-3400,-3500) and [8.2x10¹⁴,8.3x10¹⁴)). When two bins touch (e.g., two adjacent bins), for example, [1.0,1.1) and [1.1,1.2) the bins cannot both represent 1.1 (e.g., no overlap). If both bins were closed, both would represent the point where they touch. If both bins were open, neither would represent the point where they touch. Accordingly, one bin must be open and one bin must be closed. A convention is needed to determine which bin the value of 1.1 would go into. So, closed-interval includes the point, open does not. Accordingly, in one embodiment, the point closest to zero is closed. In Mathsquare, for example, brackets [] indicate closed (including) and parenthesis () indicate open (excluding).

10

15

[00037] FIG. 5 depicts a bin selection using two significant digits in base 10 and illustrates how the bottom 10% of each exponent range is actually represented by the adjacent smaller exponent. Additionally, the last exponent bin range in the figure depicts 0 as a special value wherein smaller measurements cannot be accounted for. This is a log-linear bin selection from 1.0x10⁻¹²⁷ to 9.9x10¹²⁸.

20

[00038] Given discrete representation, selected histograms may be aggregated without arbitrarily expanding the bin space, which may provide an ability to aggregate (e.g., via addition and subtraction) while controlling bin space. This may have the distinct real-world advantage of a highly controlled bin space for a physical system (e.g., a system having prescribed constraints). FIG. 6 depicts an example of multiple bin aggregation. For example, histograms in (a) added with histograms in (b) equals the histograms in (c).

25

[00039] Techniques shown below calculate and store (e.g., time) aggregated histogram data to decrease required information access on retrieval.

[00040] In one embodiment, the display may communicate the number of histograms desired and the storage system may dynamically aggregate requested data into lengthier periods of time. For example, displaying a two-year histogram time series where all

30

histograms are stored by minute would require 1,051,200 histograms returned to the visualization system. Instead, for example, the display may request 1,000 histograms, which may result in the storage systems combining 720 one-minute histograms inclusively between [00:00 and 11:59] into one 12-hour histogram and the 720 one-minute histograms inclusively between [12:00 and 23:59] into one 12-hour histogram for each day of the two years resulting in 1,460 histograms to the visualization system instead of 1,051,200 with conventional techniques. This new representation presents a waste of 46% as opposed to a waste of 105000% with conventional techniques.

5 [00041] In one embodiment, the display may understand horizontal resolution and determine how many pixels are available for rendering vertical histograms. This “horizontal viewport size” may be the number of histograms requested (e.g., a number equal to or less than the number of horizontal pixels). The histogram delivery technique may attempt to aggregate the number of bins to a level where the number of histograms returned for the requested time range roughly approximates but is less than the “horizontal viewport size” (e.g., within a factor of 4 to be roughly approximate).

10 [00042] In one embodiment, to determine the number of bins requested, available visual space may be pulled or requested from the display itself through various techniques. In a web browser, for example, the number of horizontal and vertical pixels available in the viewport may be accessible by calling a Javascript routine.

15 [00043] In one embodiment, the number of histograms may be less than the number of horizontal pixels, as any number greater may be lost during visualization. The number of bins desired may be less than the number of vertical pixels, as any number greater may be lost during visualization.

20 [00044] In one embodiment, the number of histograms is roughly half the number of horizontal pixels. For example, with 1000 pixels of horizontal space, roughly 500 histograms may be created. And, with 600 pixels of vertical space, 300 bins may be used.

25 [00045] In one embodiment, fine granularity histograms may be automatically aggregated into successively wider granularity histograms to accelerate visualization of data over arbitrary periods of time (e.g., storing 1 minute, 5 minute, 1 hour, and 12 hour aggregations). For example, each time a one-minute histogram is stored, the containing

five-minute histogram is recalculated and stored, causing the containing one-hour histogram to be recalculated and stored, causing the containing 12-hour histogram to be recalculated and stored, etc. Accordingly, histogram data is pre-calculated as data is received. The largest feasible pre-calculated histogram accumulates is selected above. A user may provide the largest feasible pre-calculated histogram based on how often the user believes they would be used. For example, a user may provide the largest feasible pre-calculated histogram based on how long they believe they will retain data and how frequently long time periods will be visualized.

5 [00046] Bin-realm Efficiency

10 [00047] Another advantage of the techniques described herein is an ability to represent histograms in a smaller information footprint based on the visualization parameters.

15 [00048] Human beings often interpret logarithmically scaled data incorrectly and tend to operate more effectively with linearly scaled data. Given a wide magnitude space and a linear representation, there may be too much data to display (e.g., 45,391 bins. When visualizing data, the resolution at which data may be usefully displayed may be determined. For example, on displays available today, humans cannot distinguish in excess of 1,000 histogram bins without external apparatus. The technique described below provides a method of dynamic rebinning that may optimize the histogram for visual delivery in linearly scaled representation where bins otherwise not visible on-display may be combined into a single bin.

20 [00049] In one embodiment, specific user-supplied minimum and maximum visual limits may be provided. For example the number of pixels and size of a display may impact the visual limits. The greater of the absolute value of the minimum and maximum visual limits is taken as the focal maximum. In one example, there may be 90 bins like

25 [1.0,1.1), [1.1,1.2),... [9.9,10) each of which times 10^n . The logarithmic bin set (e.g., the 90 bins for a particular value of n) containing the focal maximum is preserved (e.g., unchanged, these bins will not be combined with other bins for efficiency). Smaller logarithmic bin sets (e.g., logarithmic bins of 10^x (where $x < n$ from the preserved set)) may be preserved until the total number of bins meets or exceeds the available vertical visual space (e.g., with 90 bins in each logarithmic set and if a set and the surrounding sets are preserved, 270 bins may be created) and the smallest logarithmic bin set (e.g.,

based on the number of vertical pixels available, the smallest bin with a sample in it, or a user-supplied minimum) is taken as the focal minimum. If $\text{minimum} \leq 0 \leq \text{maximum}$, a special 0 bin is preserved. EXAMPLE: Given a data set: 0.02, 0.04, 1.8, 11, 76, 873, we would represent this in a log-linear histogram as one sample in each of the following buckets: [0.02,0.03), [0.04,0.05), [1.8,1.9), [11,12), [76,77), and [870,880). The focal maximum of the data set is 880 (the top of the largest bin). A user-supplied maximum may override this, let us assume a user-supplied focal maximum of 100. Assuming a vertical space of 1000 pixels, we can display the bins [11,12) and [76,77) in 10 pixels of vertical space each (bins [10,11) through [99,100) consuming 900 pixels of vertical space total). This leaves 100 pixels for bins less than 10. Bins [1.0,1.1) through [9.9,10), consume one pixel of vertical space each totally 90 pixels. This leave 10 pixels to represent all smaller bins. Accordingly, each of the bins [0.02,0.03) through [0.99,1.0) would be assigned 0.1 pixels which is not possible visualize setting our focal minimum to 1.0. Accordingly, we will combine all bins between [0.2,1.0), into a single bin containing the sum of all contained samples and display it as a single bin.

[00050] FIG. 7 depicts an example flowchart describing processing performed in an illustrative embodiment of the present invention. In 710, a data set may be received or accessed. From 710, flow may move to 720.

[00051] In 720, minimum and maximum visual limits may be received. The focal maximum may be determined from the larger of the absolute value of minimum visual limit or maximum visual limit. From 720, flow may move to 730.

[00052] In 730, a determination may be made if data minimum (e.g., the smallest value in a data set) is greater than 0. If so, a new bin may be created using [data minimum, focal minimum) and the covered bins may be discarded as in step 780. A new bin may be created and all data in bins between [data minimum, focal minimum) may be placed in the newly created bin. For example, a new bin may be created using (10*focal maximum, data maximum) and the covered bins may be discarded as in step 780. From 730, flow may move to 740.

[00053] In 740, a determination may be made if data maximum is less than 0. If so, a new bin may be created using (-1 * focal minimum, data maximum) and the covered bins may be discarded as in step 780, and a new bin may be created using (data minimum, -10 *

focal maximum) and the covered bins may be discarded as in step 780. From 740, flow may move to 750.

5 [00054] In 750, a determination may be made if data minimum is less than or equal to 0 and data minimum is greater than or equal to $-1 * \text{focal minimum}$. If so, (1) the 0 bin is preserved, (2) a new bin may be created using (data minimum, 0) and the covered bins may be discarded as in step 780, (3) a new bin may be created using (0, focal minimum) and the covered bins may be discarded as in step 780, and (4) a new bin may be created using ($10 * \text{focal maximum}$, data maximum) and the covered bins may be discarded as in step 780. From 750, flow may move to 760.

10 [00055] In 760, a determination may be made if data maximum is less than or equal to focal minimum and data maximum is greater than or equal 0. If so, (1) the 0 bin is preserved, (2) a new bin may be created using (0, data maximum) and the covered bins may be discarded as in step 780, (3) a new bin may be created using ($-1 * \text{focal minimum}$, 0) and the covered bins may be discarded as in step 780, and (4) a new bin may be created using (data minimum, $-10 * \text{focal maximum}$) and the covered bins may be discarded as in step 780. From 760, flow may move to 770.

15 [00056] In 770, if none of the determinations of 730, 740, 750, or 760 apply, then (1) the 0 bin is preserved, (2) a new bin may be created using ($-1 * \text{focal minimum}$, 0) and the covered bins may be discarded as in step 780, (3) a new bin may be created using (data minimum, $-10 * \text{focal maximum}$) and the covered bins may be discarded as in step 780, (4) a new bin may be created using (0, focal minimum) and the covered bins may be discarded as in step 780, and (5) a new bin may be created using ($10 * \text{focal maximum}$, data maximum). From 770, flow may end. Note, each of the determinations of 730, 740, 750, and 760 may be exclusive.

20 [00057] In one embodiment, the stored histograms may be reevaluated over time for minimum and maximum values. These minimum and maximum values may be used as if they were supplied as described above in 720. For example, if user or viewer selected limits were not received, the minimum and maximum values may be calculated from the dataset itself.

25 [00058] FIG. 8 depicts an example computer system that may be used in implementing an illustrative embodiment of the present invention. Specifically, Figure 8 depicts an

illustrative embodiment of a computer system 800 that may be used in computing devices such as, e.g., but not limited to, standalone, client, server devices, or system controllers. Figure 8 depicts an illustrative embodiment of a computer system that may be used as a client device, a server device, a controller, etc. The present invention (or any part(s) or function(s) thereof) may be implemented using hardware, software, firmware, or a combination thereof and may be implemented in one or more computer systems or other processing systems. In fact, in one illustrative embodiment, the invention may be directed toward one or more computer systems capable of carrying out the functionality described herein. An example of a computer system 800 is shown in Figure 8, depicting an illustrative embodiment of a block diagram of an illustrative computer system useful for implementing the present invention. Specifically, Figure 8 illustrates an example computer 800, which in an illustrative embodiment may be, e.g., (but not limited to) a personal computer (PC) system running an operating system such as, e.g., (but not limited to) MICROSOFT® WINDOWS® NT/98/2000/XP/Vista/Windows 7/Windows 8, etc. available from MICROSOFT® Corporation of Redmond, WA, U.S.A. or an Apple computer executing MAC® OS or iOS from Apple® of Cupertino, CA, U.S.A or a smartphone running iOS, Android, or Windows mobile, for example. However, the invention is not limited to these platforms. Instead, the invention may be implemented on any appropriate computer system running any appropriate operating system. In one illustrative embodiment, the present invention may be implemented on a computer system operating as discussed herein. An illustrative computer system, computer 800 is shown in Figure 8. Other components of the invention, such as, e.g., (but not limited to) a computing device, a communications device, a telephone, a personal digital assistant (PDA), an iPhone, a 3G/4G wireless device, a wireless device, a personal computer (PC), a handheld PC, a laptop computer, a smart phone, a mobile device, a netbook, a handheld device, a portable device, an interactive television device (iTV), a digital video recorder (DVR), client workstations, thin clients, thick clients, fat clients, proxy servers, network communication servers, remote access devices, client computers, server computers, peer-to-peer devices, routers, web servers, data, media, audio, video, telephony or streaming technology servers, etc., may also be implemented using a computer such as that shown in Figure 8. In an illustrative embodiment, services may be provided on demand using,

e.g., an interactive television device (iTV), a video on demand system (VOD), via a digital video recorder (DVR), and/or other on demand viewing system. Computer system 800 may be used to implement the network and components as described above.

[00059] The computer system 800 may include one or more processors, such as, e.g., but not limited to, processor(s) 804. The processor(s) 804 may be connected to a communication infrastructure 806 (e.g., but not limited to, a communications bus, cross-over bar, interconnect, or network, etc.). Processor 804 may include any type of processor, microprocessor, or processing logic that may interpret and execute instructions (e.g., for example, a field programmable gate array (FPGA)). Processor 804 may comprise a single device (e.g., for example, a single core) and/or a group of devices (e.g., multi-core). The processor 804 may include logic configured to execute computer-executable instructions configured to implement one or more embodiments. The instructions may reside in main memory 808 or secondary memory 810. Processors 804 may also include multiple independent cores, such as a dual-core processor or a multi-core processor. Processors 804 may also include one or more graphics processing units (GPU) which may be in the form of a dedicated graphics card, an integrated graphics solution, and/or a hybrid graphics solution. Various illustrative software embodiments may be described in terms of this illustrative computer system. After reading this description, it will become apparent to a person skilled in the relevant art(s) how to implement the invention and/or parts of the invention using other computer systems and/or architectures.

[00060] Computer system 800 may include a display interface 802 (e.g., the HMI) that may forward, e.g., but not limited to, graphics, text, and other data, etc., from the communication infrastructure 806 (or from a frame buffer, etc., not shown) for display on the display unit 801. The display unit 801 may be, for example, a television, a computer monitor, a touch sensitive display device, or a mobile phone screen. The output may also be provided as sound through a speaker.

[00061] The computer system 800 may also include, e.g., but is not limited to, a main memory 808, random access memory (RAM), and a secondary memory 810, etc. Main memory 808, random access memory (RAM), and a secondary memory 810, etc., may be a computer-readable medium that may be configured to store instructions configured to

implement one or more embodiments and may comprise a random-access memory (RAM) that may include RAM devices, such as Dynamic RAM (DRAM) devices, flash memory devices, Static RAM (SRAM) devices, etc.

[00062] The secondary memory 810 may include, for example, (but is not limited to) a hard disk drive 812 and/or a removable storage drive 814, representing a floppy diskette drive, a magnetic tape drive, an optical disk drive, a compact disk drive CD-ROM, flash memory, etc. The removable storage drive 814 may, e.g., but is not limited to, read from and/or write to a removable storage unit 818 in a well-known manner. Removable storage unit 818, also called a program storage device or a computer program product, may represent, e.g., but is not limited to, a floppy disk, magnetic tape, optical disk, compact disk, etc. which may be read from and written to removable storage drive 814. As will be appreciated, the removable storage unit 818 may include a computer usable storage medium having stored therein computer software and/or data.

[00063] In alternative illustrative embodiments, secondary memory 810 may include other similar devices for allowing computer programs or other instructions to be loaded into computer system 800. Such devices may include, for example, a removable storage unit 822 and an interface 820. Examples of such may include a program cartridge and cartridge interface (such as, e.g., but not limited to, those found in video game devices), a removable memory chip (such as, e.g., but not limited to, an erasable programmable read only memory (EPROM), or programmable read only memory (PROM) and associated socket, and other removable storage units 822 and interfaces 820, which may allow software and data to be transferred from the removable storage unit 822 to computer system 800.

[00064] Computer 800 may also include an input device 803 which may include any mechanism or combination of mechanisms that may permit information to be input into computer system 800 from, e.g., a user or operator. Input device 803 may include logic configured to receive information for computer system 800 from, e.g. a user or operator. Examples of input device 803 may include, e.g., but not limited to, a mouse, pen-based pointing device, or other pointing device such as a digitizer, a touch sensitive display device, and/or a keyboard or other data entry device (none of which are labeled). Other

input devices 803 may include, e.g., but not limited to, a biometric input device, a video source, an audio source, a microphone, a web cam, a video camera, and/or other camera.

[00065] Computer 800 may also include output devices 815 which may include any mechanism or combination of mechanisms that may output information from computer system 800. Output device 815 may include logic configured to output information from computer system 800. Embodiments of output device 815 may include, e.g., but not limited to, display 801, and display interface 802, including displays, printers, speakers, cathode ray tubes (CRTs), plasma displays, light-emitting diode (LED) displays, liquid crystal displays (LCDs), printers, vacuum fluorescent displays (VFDs), surface-conduction electron-emitter displays (SEDs), field emission displays (FEDs), etc. Computer 800 may include input/output (I/O) devices such as, e.g., (but not limited to) input device 803, communications interface 824, connection 828 and communications path 826, etc. These devices may include, e.g., but are not limited to, a network interface card, onboard network interface components, and/or modems.

[00066] Communications interface 824 may allow software and data to be transferred between computer system 800 and external devices or other computer systems.

Computer system 800 may connect to other devices or computer systems via wired or wireless connections. Wireless connections may include, for example, WiFi, satellite, mobile connections using, for example, TCP/IP, 802.15.4, high rate WPAN, low rate WPAN, 6loWPAN, ISA100.11a, 802.11.1, WiFi, 3G, WiMAX, 4G and/or other communication protocols.

[00067] In this document, the terms “computer program medium” and “computer readable medium” may be used to generally refer to media such as, e.g., but not limited to, removable storage drive 814, a hard disk installed in hard disk drive 812, flash memories, removable discs, non-removable discs, etc. In addition, it should be noted that various electromagnetic radiation, such as wireless communication, electrical communication carried over an electrically conductive wire (e.g., but not limited to twisted pair, CAT5, etc.) or an optical medium (e.g., but not limited to, optical fiber) and the like may be encoded to carry computer-executable instructions and/or computer data that embodiments of the invention on e.g., a communication network. These computer program products may provide software to computer system 800. It should be noted that

a computer-readable medium that comprises computer-executable instructions for execution in a processor may be configured to store various embodiments of the present invention. References to "one embodiment," "an embodiment," "example embodiment," "various embodiments," etc., may indicate that the embodiment(s) of the invention so described may include a particular feature, structure, or characteristic, but not every embodiment necessarily includes the particular feature, structure, or characteristic.

5 [00068] Further, repeated use of the phrase "in one embodiment," or "in an illustrative embodiment," do not necessarily refer to the same embodiment, although they may. The various embodiments described herein may be combined and/or features of the

10 embodiments may be combined to form new embodiments.

[00069] Unless specifically stated otherwise, as apparent from the following discussions, it is appreciated that throughout the specification discussions utilizing terms such as "processing," "computing," "calculating," "determining," or the like, refer to the action and/or processes of a computer or computing system, or similar electronic computing device, that manipulate and/or transform data represented as physical, such as electronic, quantities within the computing system's registers and/or memories into other data similarly represented as physical quantities within the computing system's memories, registers or other such information storage, transmission or display devices.

[00070] In a similar manner, the term "processor" may refer to any device or portion of a device that processes electronic data from registers and/or memory to transform that electronic data into other electronic data that may be stored in registers and/or memory. A "computing platform" may comprise one or more processors.

[00071] Embodiments of the present invention may include apparatuses for performing the operations herein. An apparatus may be specially constructed for the desired purposes, or it may comprise a general purpose device selectively activated or reconfigured by a program stored in the device.

[00072] Embodiments may be embodied in many different ways as a software component. For example, it may be a stand-alone software package, or it may be a software package incorporated as a "tool" in a larger software product, such as, for example, a scientific modeling product. It may be downloadable from a network, for example, a website, as a stand-alone product or as an add-in package for installation in an existing software

25

30

application. It may also be available as a client-server software application, or as a web-enabled software application. It may also be part of a system for efficient manipulation and display of histograms. Computer system 800 may be used to create a general purpose computer. A general purpose computer may be specialized by storing programming logic that enables one or more processors to perform the techniques indicated herein and one or more of the steps of FIG. 6. Computer system 800 or multiple embodiments of computer system 800 may be used to perform the functions described above.

5

[00073] Embodiments of the present invention may include apparatuses for performing the operations herein. An apparatus may be specially constructed for the desired purposes, or it may comprise a general purpose device selectively activated or reconfigured by a program stored in the device.

10

[00074] Embodiments of the present invention may improve the functioning of computer system 800. For example, embodiments of the present invention processing speed for creating and manipulating histograms.

15

[00075] Embodiments may be embodied in many different ways as a software component. For example, it may be a stand-alone software package, or it may be a software package incorporated as a "tool" in a larger software product. It may be downloadable from a network, for example, a website, as a stand-alone product or as an add-in package for installation in an existing software application. It may also be available as a client-server software application, or as a web-enabled software application.

20

[00076] While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of the present invention should not be limited by any of the above-described illustrative embodiments, but should instead be defined only in accordance with the following claims and their equivalents. The embodiments of the present invention that have been described above may contain features that may be removed or combined between the described embodiments to derive additional embodiments.

25

What is claimed is:

1. A method for constructing histograms comprising:

5 receiving, by one or more processors, a set of data values, wherein the data values are two-dimensional or three-dimensional and the data values contain at least one measurement over time;

10 creating, by the one or more processors, a plurality of bins in memory, wherein the plurality of bins are log-linear bins using 2 significant digits base 10;

15 creating, by the one or more processors, a zero bin representing the number 0;

10 placing, by the one or more processors, non-zero numbers of the set of data values into the plurality of bins; and

generating, by the one or more processors, one or more histograms using the bins.

2. The method of claim 1, wherein:

15 the one or more histograms are generated based on 1 minute, 5 minutes, 1 hour, and 12 hour time increments.

3. The method of claim 2, wherein

20 measurement values are received each minute; and

each minute, a one-minute histogram is created, a five-minute histogram is recalculated, a one-hour histogram is recalculated, and a twelve-hour histogram is recalculated.

4. The method of claim 1, further comprising:

25 receiving minimum and maximum visual limits;

calculating a focal maximum by taking the greater number from the absolute value of the minimum and maximum visual limits;

calculating a focal minimum by taking the lowest number from of the absolute value of the minimum and maximum visual limits;

30 calculating a data minimum from the set of data values; and

calculating a data maximum from the set of data values.

5. The method of claim 4, wherein generating one or more histograms comprises:

(a) when the data minimum is greater than 0:

5 creating a first bin of the plurality of bins, wherein the first bin ranges from the data minimum to the focal minimum and discard covered bins, and

creating a second bin of the plurality of bins, wherein the second bin ranges from $10 * \text{focal maximum}$ to the data maximum and discard covered bins;

(b) when the data maximum is less than 0:

10 creating the first bin ranging from $-1 * \text{the focal minimum}$ to the data maximum and discard covered bins, and

creating the second bin ranging from the data minimum to $-10 * \text{the focal maximum}$ and discard covered bins;

(c) when the data minimum is less than or equal to 0 and the data minimum is greater than or equal to $-1 * \text{the focal minimum}$:

15 preserving the zero bin,

creating the first bin ranging from the data minimum to 0 and discard covered bins,

creating the second bin ranging from 0 to the focal minimum and discard covered bins, and

20 creating a third bin of the plurality of bins, wherein the third bin ranges from $10 * \text{the focal maximum}$ to the data maximum and discard covered bins;

(d) when the data maximum is greater than or equal to 0 and the data maximum is less than or equal to the focal minimum:

25 preserving the zero bin,

creating the first bin ranging from 0 to the data maximum and discard covered bins,

creating the second bin ranging from $-1 * \text{the focal minimum}$ to 0 and discard covered bins, and

30 creating the third bin ranging from the data minimum to $-10 * \text{the focal maximum}$ and discard covered bins; and

when (a)-(d) are not met:

preserving the zero bin,
creating the first bin ranging from -1 * the focal minimum to 0 and discard covered bins,
5 creating the second bin ranging from the data minimum to -10 * focal maximum and discard covered bins,
create the third bin ranging from 0 to the focal minimum and discard covered bins, and
..... creating a fourth bin of the plurality of bins, wherein the fourth bin ranges from 10 * the focal maximum to the data maximum and discard covered bins.

10

6. A non-transitory computer-readable medium comprising instructions, which when executed by one or more processors causes the one or more processors to perform operations for constructing histograms, the computer-readable medium comprising instructions to:

15 receive a set of data values, wherein the data values are two-dimensional or three-dimensional and the data values contain at least one measurement over time;
create a plurality of bins in memory, wherein the plurality of bins are log-linear bins using 2 significant digits base 10;
create a zero bin representing the number 0;
20 place non-zero numbers of the set of data values into the plurality of bins; and generate one or more histograms using the bins.

7. The non-transitory computer-readable medium of claim 6, wherein:
the one or more histograms are generated based on 1 minute, 5 minutes, 1 hour,
25 and 12 hour time increments.

8. The non-transitory computer-readable medium of claim 7, wherein
measurement values are received each minute; and
each minute, a one-minute histogram is created, a five-minute histogram is
30 recalculated, a one-hour histogram is recalculated, and a twelve-hour histogram is
recalculated.

9. The non-transitory computer-readable medium of claim 6, further comprising instructions to:

receive minimum and maximum visual limits;

5 calculate a focal maximum by taking the greater number from the absolute value of the minimum and maximum visual limits;

calculate a focal minimum by taking the lowest number from the absolute value of the minimum and maximum visual limits;

calculate a data minimum from the set of data values; and

10 calculate a data maximum from the set of data values.

10. The non-transitory computer-readable medium of claim 9, wherein generating one or more histograms further comprises instructions to:

(a) when the data minimum is greater than 0:

15 create a first bin of the plurality of bins, wherein the first bin ranges from the data minimum to the focal minimum and discard covered bins, and

create a second bin of the plurality of bins, wherein the second bin ranges from $10 * \text{focal maximum}$ to the data maximum and discard covered bins;

(b) when the data maximum is less than 0:

20 create the first bin ranging from $-1 * \text{the focal minimum}$ to the data maximum and discard covered bins, and

create the second bin ranging from the data minimum to $-10 * \text{the focal maximum}$ and discard covered bins;

(c) when the data minimum is less than or equal to 0 and the data minimum is greater than or equal to $-1 * \text{the focal minimum}$:

25 preserve the zero bin,

create the first bin ranging from the data minimum to 0 and discard covered bins,

30 create the second bin ranging from 0 to the focal minimum and discard covered bins, and

create a third bin of the plurality of bins, wherein the third bin ranges from 10 * the focal maximum to the data maximum and discard covered bins;

(d) when the data maximum is greater than or equal to 0 and the data maximum is less than or equal to the focal minimum:

5 preserve the zero bin,

create the first bin ranging from 0 to the data maximum and discard covered bins,

create the second bin ranging from -1 * the focal minimum to 0 and discard covered bins, and

10 create the third bin ranging from the data minimum to -10 * the focal maximum and discard covered bins; and

when (a)-(d) are not met:

preserve the zero bin,

create the first bin ranging from -1 * the focal minimum to 0 and discard covered bins,

15

create the second bin ranging from the data minimum to -10 * focal maximum and discard covered bins,

create the third bin ranging from 0 to the focal minimum and discard covered bins, and

20 create a fourth bin of the plurality of bins, wherein the fourth bin ranges from 10 * the focal maximum to the data maximum and discard covered bins.

11. A system for constructing histograms comprising:

one or more processors for executing a plurality of instructions;

25 a display device in communication with the one or more processors; and

a storage device in communication with the one or more processors, the storage device holding the plurality of instructions, the plurality of instructions including instructions to:

30 receive a set of data values, wherein the data values are two-dimensional or three-dimensional and the data values contain at least one measurement over time;

create a plurality of bins in memory, wherein the plurality of bins are log-linear bins using 2 significant digits base 10;

5 create a zero bin representing the number 0;

place non-zero numbers of the set of data values into the plurality of bins;

generate one or more histograms using the bins; and

display the one or more histograms via the display device.

12. The system of claim 11, wherein:

10 the one or more histograms are generated based on 1 minute, 5 minutes, 1 hour, and 12 hour time increments.

13. The system of claim 12, wherein

15 measurement values are received each minute; and
each minute, a one-minute histogram is created, a five-minute histogram is recalculated, a one-hour histogram is recalculated, and a twelve-hour histogram is recalculated.

14. The system of claim 11, wherein the plurality of instructions further include instructions to:

20 receive minimum and maximum visual limits;
calculate a focal maximum by taking the greater number from the absolute value of the minimum and maximum visual limits;
calculate a focal minimum by taking the lowest number from the absolute value of the minimum and maximum visual limits;
25 calculate a data minimum from the set of data values; and
calculate a data maximum from the set of data values.

15. The system of claim 14, wherein generating one or more histograms further comprises instructions to:

30 (a) when the data minimum is greater than 0:

create a first bin of the plurality of bins, wherein the first bin ranges from the data minimum to the focal minimum and discard covered bins, and

create a second bin of the plurality of bins, wherein the second bin ranges from $10 * \text{focal maximum}$ to the data maximum and discard covered bins;

5 (b) when the data maximum is less than 0:

create the first bin ranging from $-1 * \text{the focal minimum}$ to the data maximum and discard covered bins, and

create the second bin ranging from the data minimum to $-10 * \text{the focal maximum}$ and discard covered bins;

10 (c) when the data minimum is less than or equal to 0 and the data minimum is greater than or equal to $-1 * \text{the focal minimum}$:

preserve the zero bin,

create the first bin ranging from the data minimum to 0 and discard covered bins,

15 create the second bin ranging from 0 to the focal minimum and discard covered bins, and

create a third bin of the plurality of bins, wherein the third bin ranges from $10 * \text{the focal maximum}$ to the data maximum and discard covered bins;

20 (d) when the data maximum is greater than or equal to 0 and the data maximum is less than or equal to the focal minimum:

preserve the zero bin,

create the first bin ranging from 0 to the data maximum and discard covered bins,

25 create the second bin ranging from $-1 * \text{the focal minimum}$ to 0 and discard covered bins, and

create the third bin ranging from the data minimum to $-10 * \text{the focal maximum}$ and discard covered bins; and

when (a)-(d) are not met:

preserve the zero bin,

30 create the first bin ranging from $-1 * \text{the focal minimum}$ to 0 and discard covered bins,

create the second bin ranging from the data minimum to -10 * focal maximum and discard covered bins,

create the third bin ranging from 0 to the focal minimum and discard covered bins, and

5 create a fourth bin of the plurality of bins, wherein the fourth bin ranges from 10 * the focal maximum to the data maximum and discard covered bins.

1/8

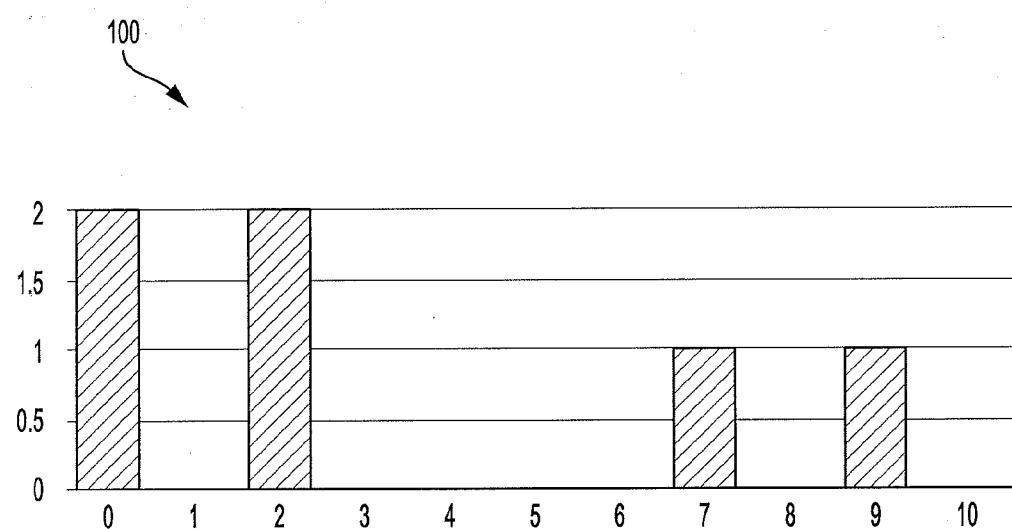


FIG. 1

2/8

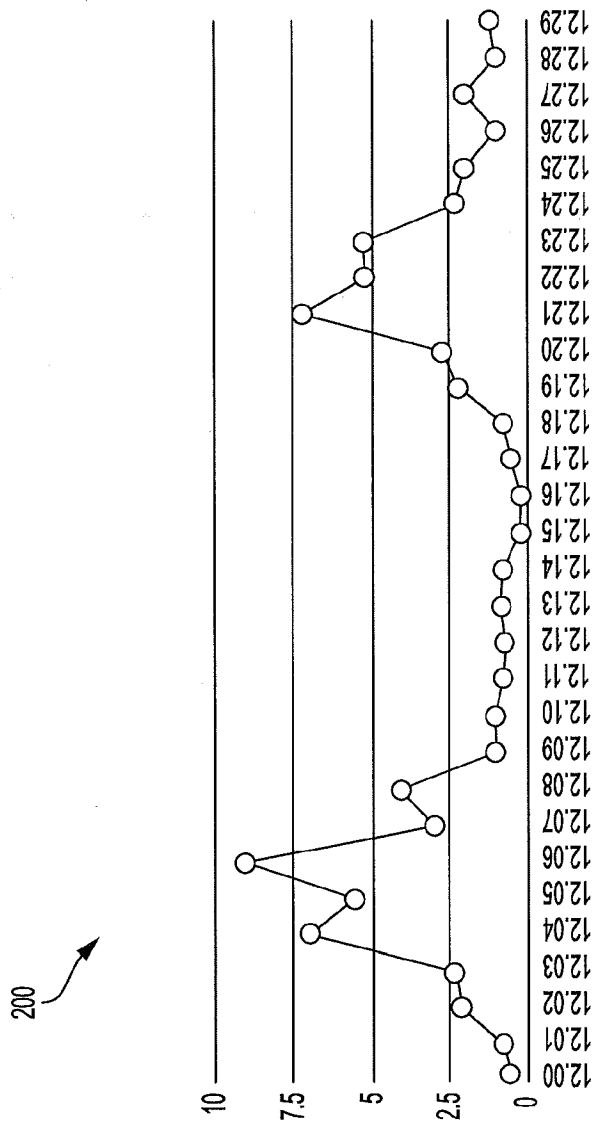
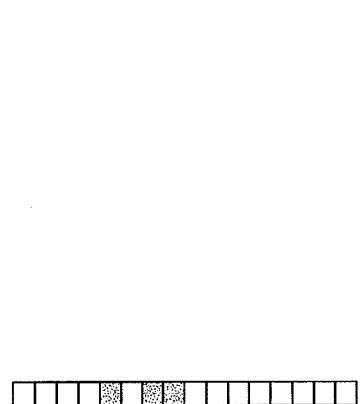
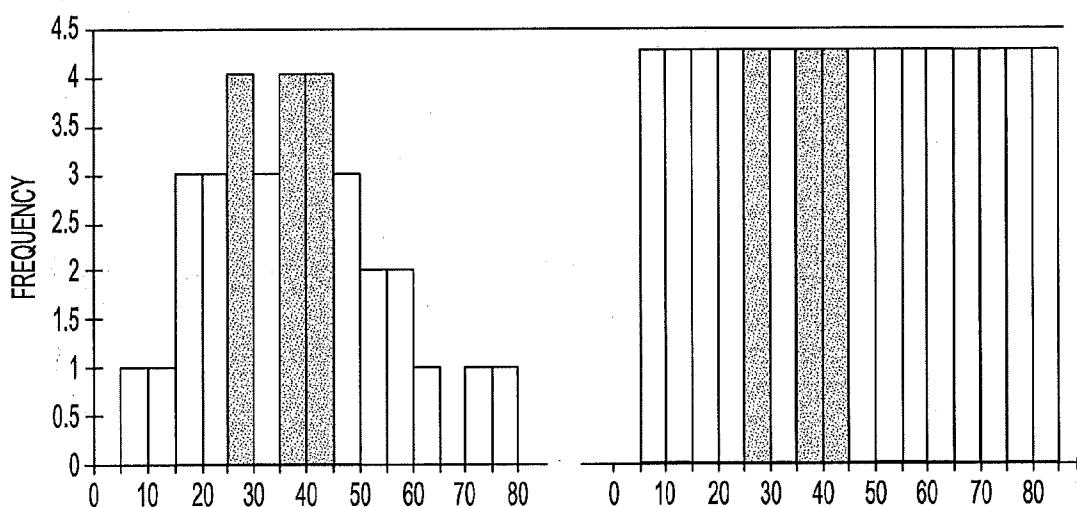


FIG. 2

3/8



4/8



FIG. 4

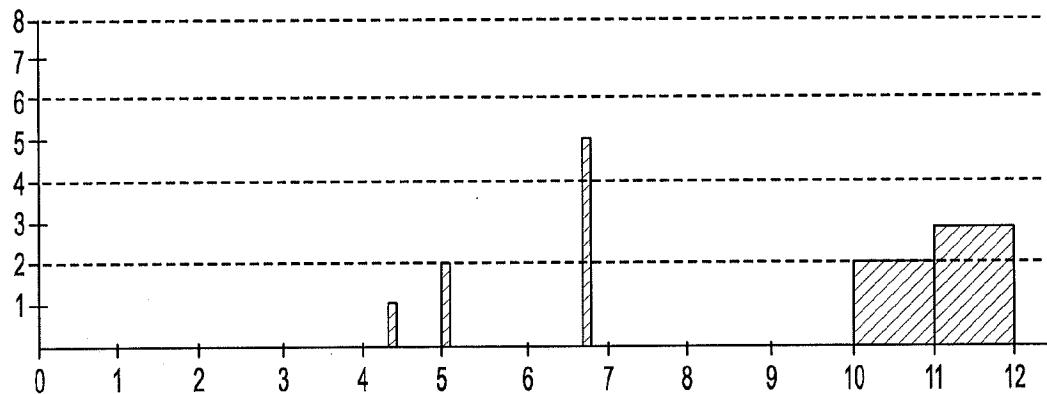
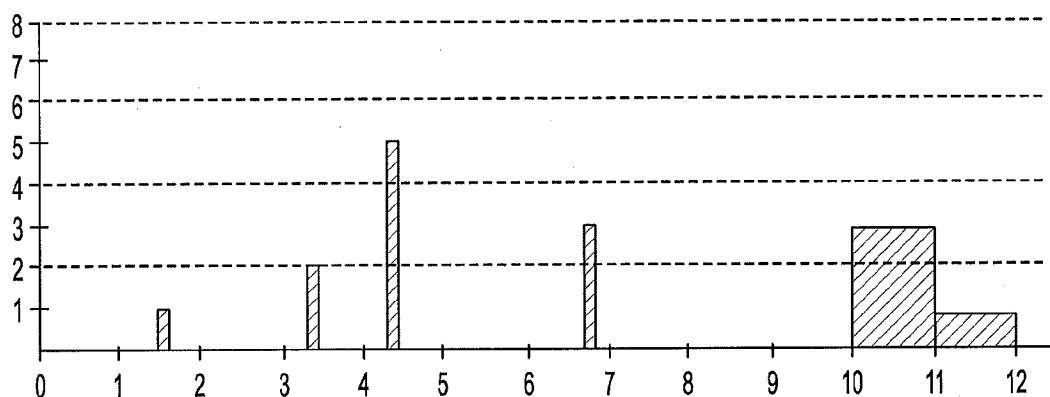
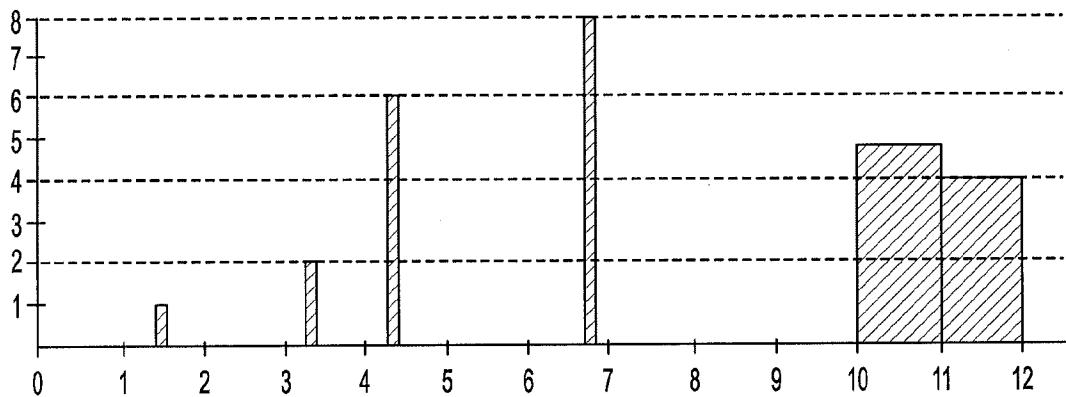
5/8



[0.1.0e127]	1.0e127	1.1e127	1.2e127	1.3e127	1.4e127	1.5e127	1.6e127	1.7e127	-	9.7e127	9.8e127	9.9e127
[0.1.0e126]	1.0e126	1.1e126	1.2e126	1.3e126	1.4e126	1.5e126	1.6e126	1.7e126	-	9.7e126	9.8e126	9.9e126
0-10	[10-11)	[11-12)	[12-13)	[13-14)	[14-15)	[15-16)	[16-17)	-	[97-98)	[98-99)	[99-100)	
0-1	[1.0-1.1)	[1.1-1.2)	[1.2-1.3)	[1.3-1.4)	[1.4-1.5)	[1.5-1.6)	[1.6-1.7)	-	[9.7-9.8)	[9.8-9.9)	[9.9-10)	
0-0.1	[0.10-0.11)	[0.11-0.12)	[0.12-0.13)	[0.13-0.14)	[0.14-0.15)	[0.15-0.16)	[0.16-0.17)	-	[0.97-0.98)	[0.98-0.99)	[0.99-0.1)	
0	1.0e-127	1.1e-127	1.2e-127	1.3e-127	1.4e-127	1.5e-127	1.6e-127	1.7e-127	-	9.7e-127	9.8e-127	9.9e-127

FIG. 5

6/8

a**b****c****FIG. 6**

7/8

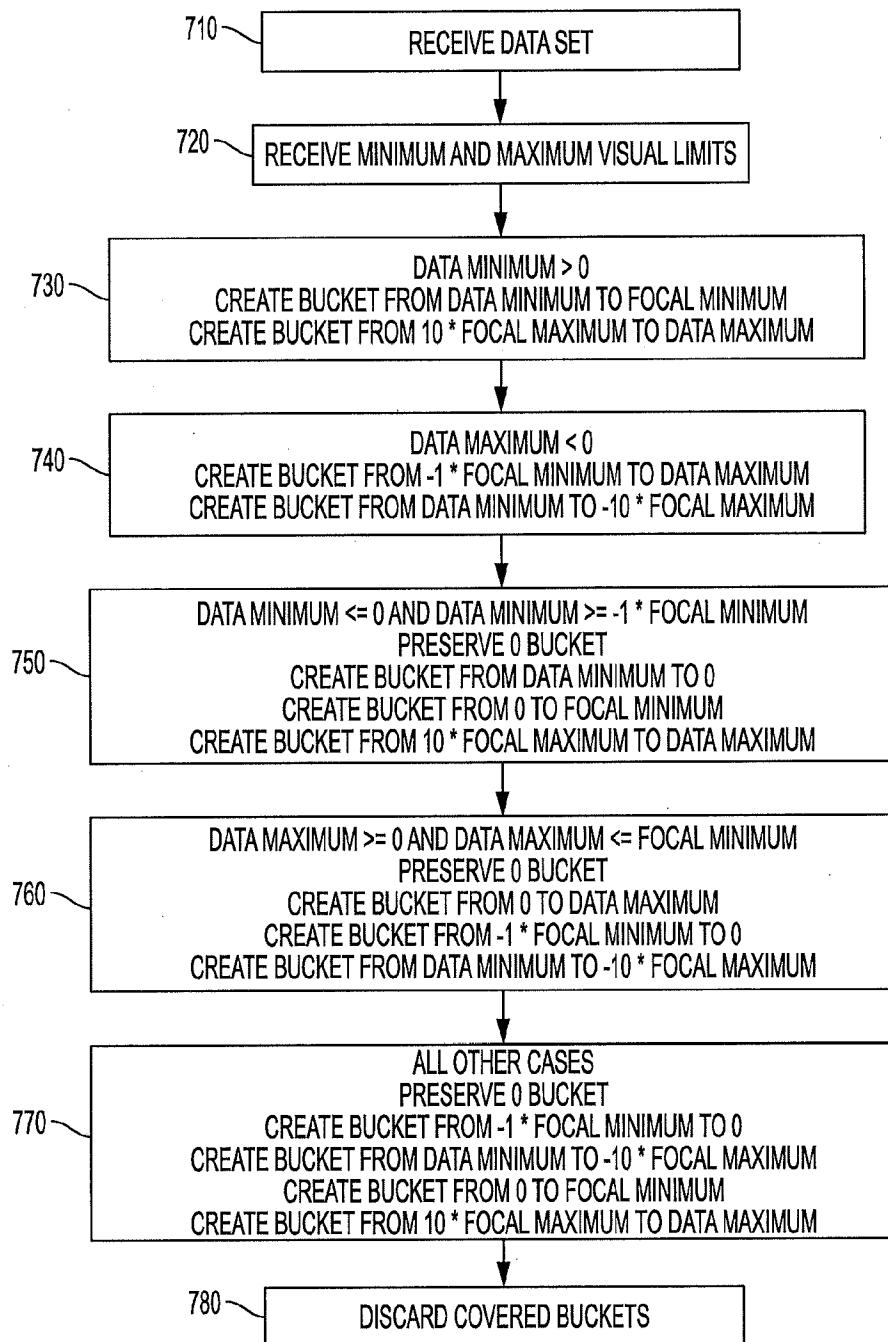


FIG. 7

8/8

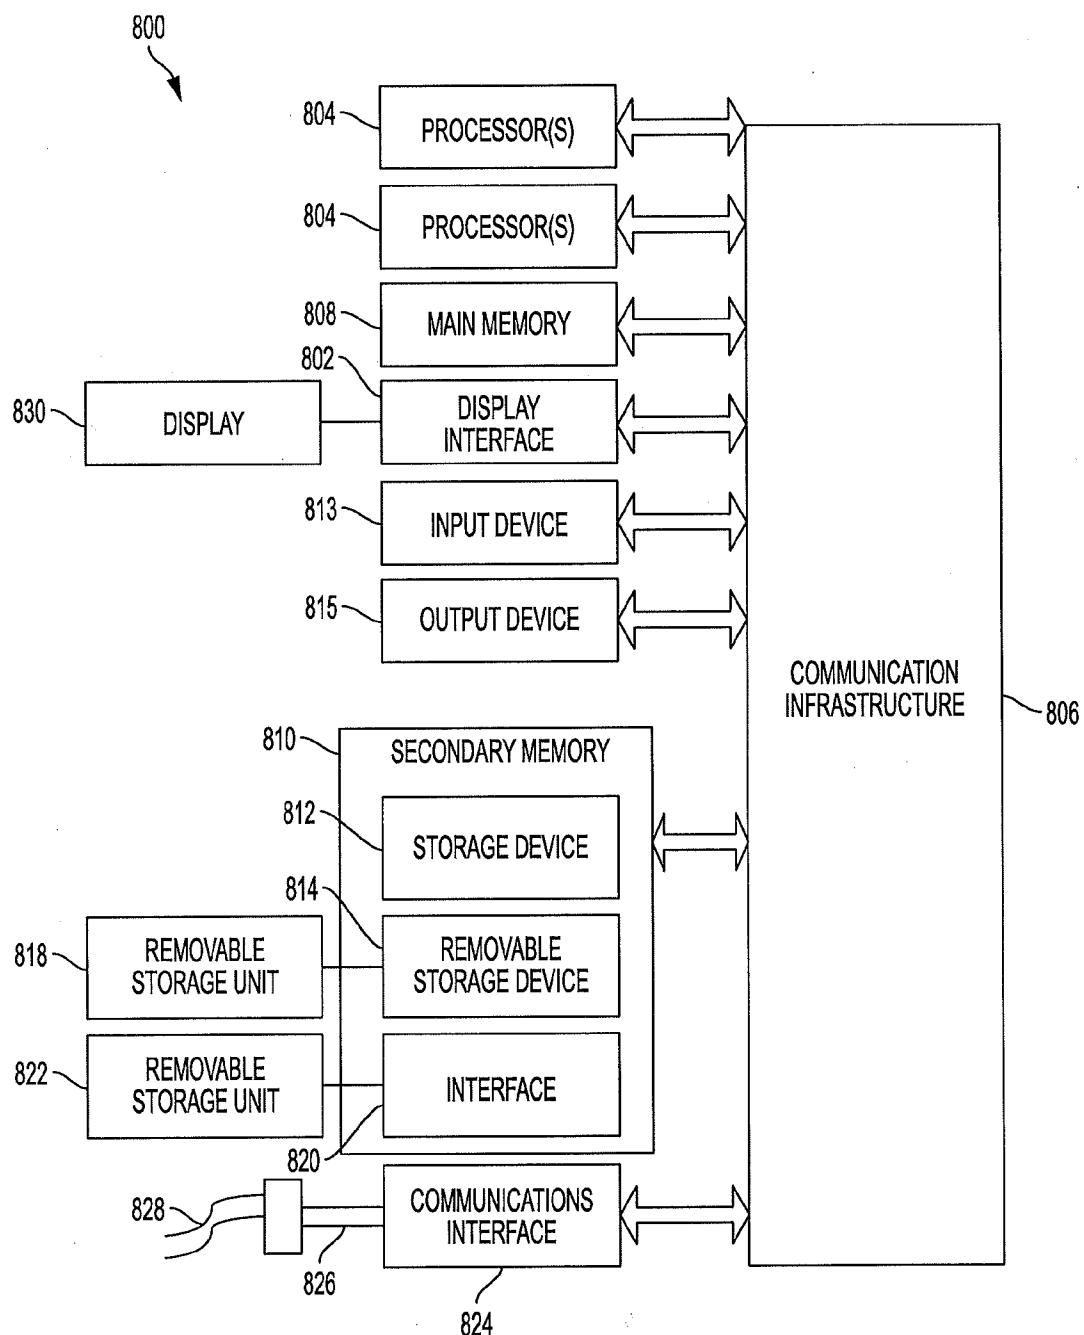


FIG. 8

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 15/50677

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06K 9/00 (2015.01)

CPC - G06T5/40

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)
CPC: G06T5/40; IPC(8): G06K 9/00 (2015.01); USPC: 382/168Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
CPC: G06T5/40, G06T5/009, H04N1/4074, G06K9/4647, G06T2207/10016; USPC: 382/168, 702/180, 358/522, 707/746 (keyword limited; terms below)Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
PatBase; Google Scholar
Search Terms: construct, generate, create, histogram, two-dimensional, three-dimensional, time, temporal, bin, zero, log-linear, base 10, significant digits

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X -- Y	US 2008/0027774 A1 (JAMESON) 31 January 2008 (31.01.2008), entire document, especially abstract and para [0078], [0092], [0174], [0249], [0323], [0376], [0379], [0478], [0515], [0517]-[0518], [0555], [0599], [0608], [0614]-[0626], [0702], [0757]-[0767], [0933], Fig. 11A, Fig. 13, Fig. 14, Fig. 15, Fig. 37, Fig. 48, Fig. 50, Fig. 65, Fig. 72.	1, 6, 11 ----- 2-3, 7-8, 12-13
A	US 2011/0063479 A1 (BITO et al.) 17 March 2011 (17.03.2011), entire document, especially abstract and para [0092], [0195]-[0200].	4-5, 9-10, 14-15
A	US 2004/0260517 A1 (DING et al.) 23 December 2004 (23.12.2004), entire document.	1-15

Further documents are listed in the continuation of Box C.

* Special categories of cited documents:	
"A" document defining the general state of the art which is not considered to be of particular relevance	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"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
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

02 November 2015 (02.11.2015)

Date of mailing of the international search report

18 DEC 2015

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450
Facsimile No. 571-273-8300

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774