



- (51) International Patent Classification:
G06F 17/00 (2006.01) *G06F 17/30* (2006.01)
- (21) International Application Number:
PCT/US2013/070083
- (22) International Filing Date:
14 November 2013 (14.11.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 11445 Compaq Center Drive West, Houston, Texas 77070 (US).
- (72) Inventors: **KUNO, Harumi**; 1501 Page Mill Rd., Palo Alto, California 94304-1100 (US). **GRAEFE, Goetz**; 521 Riverside Dr, Madison, Wisconsin 53704 (US).
- (74) Agents: **SEGARRA, Roosevelt V.** et al.; Hewlett-Packard Company, Intellectual Property Administration, 3404 E. Harmony Road, Mail Stop 35, Fort Collins, Colorado 80528 (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,

[Continued on next page]

(54) Title: ESTIMATING DATA

(57) Abstract: Disclosed herein are a system, non transitory computer-readable medium, and method for estimating data-base performance. A request for an estimate of data is read. The estimate is calculated based at least partially on a node located in a data structure.




Fig. 2





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.

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))

- (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,

ESTIMATING DATA

BACKGROUND

[0001] Database management software may utilize various hierarchical data structures to index and store data. One approach used to predict and improve database performance is to estimate the amount of data it would return in response to a given predicate. Such estimates may be used to optimize a query plan.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] FIG. 1 is an example system in accordance with aspects of the disclosure.

[0003] FIG. 2 is a flow diagram in accordance with aspects of the disclosure.

[0004] FIG. 3 is a working example in accordance with aspects of the disclosure.

[0005] FIG. 4 is a further working example in accordance with aspects of the disclosure.

DETAILED DESCRIPTION

[0006] As noted above, the performance of a database may be improved with an optimized query plan. The optimization of a query plan may be based upon estimates of the amount of data the query would process and an execution engine's ability to carry out the query. A query may comprise a predicate with criteria such that a response to the predicate contains data satisfying the criteria. Therefore, the performance of a query may be predicted by estimating the amount of data that would satisfy a predicate associated with the query. The data estimates may be obtained without actually executing the predicate since doing so may utilize a significant amount of resources. Thus, one way to estimate the amount of data returned in response to a predicate is to sample data from the database and generate histograms based on the samples. Unfortunately, the underlying data in a database may be very large and may be distributed unevenly. Therefore, estimates produced using a sampling or histogram approach may be off by multiple orders of magnitude.

[0007] One cause of inaccurate estimates may be that the samples are unrepresentative of the criteria. Another cause may be stale statistics or unnecessarily frequent updates to the data. When the underlying data changes, a

histogram may need to be refreshed to accurately reflect the change. Current policies for triggering refresh of statistics may be heuristic and these policies may fail to refresh statistics when needed or may refresh too frequently. A further cause of inaccurate estimates may be the limited precision of histograms generally. As noted above, the actual distribution of data may be much more uneven than what can be represented by a histogram.

[0008] In view of the foregoing, various examples disclosed herein provide a system, non transitory computer-readable medium and method to estimate the performance of a database. In one aspect, the estimate may be calculated based at least partially on child nodes or content of a node located in a hierarchical data structure. Such data structures may be used to index a database. Rather than generating histograms based on samples that may or may not be representative of the criteria, the database index, which may be much more reliable than samples, may be used to generate the estimate. The nodes of the index may be utilized as equal area histograms, since the data associated with the nodes may be balanced in order to accelerate the searches. Moreover, hierarchical database indices may allow estimates to be calculated at various levels of precision. The aspects, features and advantages of the application will be appreciated when considered with reference to the following description of examples and accompanying figures. The following description does not limit the application; rather, the scope of the application is defined by the appended claims and equivalents.

[0009] FIG. 1 presents a schematic diagram of an illustrative computer apparatus 100 depicting various components in accordance with aspects of the present disclosure. The computer apparatus 100 may include all the components normally used in connection with a computer. For example, it may have a keyboard and mouse and/or various other types of input devices such as pen-inputs, joysticks, buttons, touch screens, etc., as well as a display, which could include, for instance, a CRT, LCD, plasma screen monitor, TV, projector, etc. Computer apparatus 100 may also comprise a network interface (not shown) to communicate with other devices over a network using conventional protocols (e.g., Ethernet, Wi-Fi, Bluetooth, etc.). The computer apparatus 100 may also contain a processor 110, which may be any number of well known processors, such as processors from Intel® Corporation. In

another example, processor 110 may be an application specific integrated circuit ("ASIC"). Non-transitory computer readable medium ("CRM") 112 may store instructions that may be retrieved and executed by processor 110. As will be discussed in more detail below, the instructions may include an estimator 114. Non-transitory CRM 112 may be used by or in connection with any instruction execution system that can fetch or obtain the logic from non-transitory CRM 112 and execute the instructions contained therein.

[0010] Non-transitory CRM 112 may also comprise any one of many physical media such as, for example, electronic, magnetic, optical, electromagnetic, or semiconductor media. More specific examples of suitable non-transitory computer-readable media include, but are not limited to, a portable magnetic computer diskette such as floppy diskettes or hard drives, a read-only memory ("ROM"), an erasable programmable read-only memory, a portable compact disc or other storage devices that may be coupled to computer apparatus 100 directly or indirectly. Alternatively, non-transitory CRM 112 may be a random access memory ("RAM") device or may be divided into multiple memory segments organized as dual in-line memory modules ("DIMMs"). The non-transitory CRM 112 may also include any combination of one or more of the foregoing and/or other devices as well. While only one processor and one non-transitory CRM are shown in FIG. 1, computer apparatus 100 may actually comprise additional processors and computer readable mediums that may or may not be stored within the same physical housing or location. For example, non transitory CRM 112 may be a hard drive or other storage media located in a server farm of a data center. Accordingly, references to a processor, computer, or computer readable medium will be understood to include references to a collection of processors, computers, or mediums that may or may not operate in parallel.

[0011] Estimator 114 may comprise any set of instructions to be executed directly (such as machine code) or indirectly (such as scripts) by processor 110. In this regard, the terms "instructions," "scripts," and "applications" may be used interchangeably herein. The computer executable instructions may be stored in any computer language or format, such as in object code or modules of source code. Furthermore, it is understood that the instructions may be implemented in the form of

hardware, software, or a combination of hardware and software and that the examples herein are merely illustrative.

[0012] As will be discussed in more detail below, a database index may be arranged as a hierarchical, tree-like data structure that relies on splitting and merging to maintain its balance and optimize searches. In one example, estimator 114 may instruct processor 110 to read a request to estimate an amount of data that would fall within a scope of criteria in a predicate without actually executing the predicate. In another aspect, estimator 114 may instruct processor 110 to locate a node in the data structure whose child nodes or content fall within the scope of the criteria. In yet a further example, the estimate may be calculated based at least partially on the located node.

[0013] Working examples of the system, method, and non-transitory computer-readable medium are shown in FIGS. 2-4. In particular, FIG. 2 illustrates a flow diagram of an example method 200 for estimating database performance. FIGS. 3-4 each show a working example in accordance with the techniques disclosed herein. The actions shown in FIGS. 3-4 will be discussed below with regard to the flow diagram of FIG. 2.

[0014] In block 202 of FIG. 2, a request to estimate an amount of data in response to criteria in a predicate may be read. A node in a data structure may be located, as shown in block 204. Referring now to FIG. 3, an example data structure 300 is shown. While the example data structure of FIG. 3 is one type of B-tree data structure, it is understood that any type of balanced tree data structure may be used to implement the techniques disclosed herein including, but not limited to, B-plus trees, B-link trees or any other ad-hoc variation of a B-tree. Furthermore, other non B-tree data structures may also be used. The example B-tree in FIG. 3 is an index of an illustrative inventory database in which each product is associated with its own unique product number. For ease of illustration, the product numbers range from 0 to 500. The product numbers are used as keys in the index of the inventory database. Root node 304 may contain keys ranging from 0 to 500; thus, root node 304 contains the broadest range of keys. Interior node 306 has keys that range from 0 to 200 and interior node 308 has keys ranging from 201 to 500. Leaf node 310 contains the actual data of products associated with keys 30 to 140; leaf node 312

contains the actual data of products associated with keys 140 to 200; finally, leaf node 314 contains the actual data of products associated with keys 201 to 240. The data records maintained by the leaf nodes 310, 312, and 314 may be a linked list data structure of records.

[0015] A B-tree data structure may be configured to maintain its “balance” such that every path from the root node (e.g., root node 304) to a leaf node (e.g. leaf nodes 310, 312, and 314) is of the same length. Such balance may optimize a search for data beginning at the root node of the tree and ending at a leaf node. Particularly, a B-tree may be configured to maintain its balance when new data is added or deleted. That is, new data may be added or deleted in such a way that the balance is maintained. An interior node of a B tree may have a predetermined number of key entries that may be used as “separator keys” by a search engine. The keys allow a search engine to determine the correct path to the data sought after in the leaf nodes. The goal of database indices may be to optimize access time to data. As will be demonstrated below, a balanced index may also be used to provide accurate predicate response estimates.

[0016] In the example of FIG. 3, estimator 302 attempts to meet a request for an estimate of products whose associated keys are within the range of 201 to 500. In one example, estimator 302 may determine a level in the data structure hierarchy that corresponds to the scope of the criteria of the request. The node used to determine the estimate may be located at the level in the hierarchy of the data structure that corresponds to the scope. In the example of FIG. 3, interior node 308 and interior node 306 seem to be at a level corresponding to the scope of the criteria, with interior node 308 and interior node 306 holding keys ranging from 0 to 500; however, interior node 308 more precisely corresponds to the scope, since it holds keys ranging from 201 to 500. Estimator 302 may arrive at interior node 308 by reading the separator keys in root node 304 and following the link to interior node 308.

[0017] Referring back to FIG. 2, the requested estimated may be calculated based at least partially on the child nodes of the located node or on the content of the located node, as shown in block 206. As such, if the level in the hierarchy is not a leaf level, the located node may be an interior node whose child nodes fall within

the scope of the criteria. This is the case in the example of FIG. 3. In the simplest case, the located node may be the root node. However, if the level is the leaf level, the located node may be a leaf node whose content falls within the scope of the criteria. As noted above, leaf nodes are the lowest level nodes of the hierarchy that contain the actual data records. Referring back to FIG. 3, estimator 302 may count the number of keys within the 201 to 500 range that are associated with the child nodes of interior node 308. The number of keys maintained by interior node 308 may be returned as the estimate of products associated with keys within the 201 to 500 range. In one example, the estimate may be returned as a cardinality estimate. In another example, cardinality may be defined as a number of records or rows of data that would be returned in response to a predicate. In a further example, the estimate may be returned as a selectivity estimate. In yet a further example, a selectivity estimate may be defined as a percentage of data that would be returned in response to a predicate.

[0018] Referring back to FIG. 3, interior node 308 has one child node, which is a leaf node containing records associated with a narrower range of keys, 201 to 240. That is, in this example, there are no records of products associated with a key higher than 240. Thus, the estimate returned above would actually be an estimate of records associated with keys within the 201 to 240 range and not within the 201 to 500 range. Simply retrieving the number of keys located in interior node 308 may mislead a user into thinking that there is a broader range of product numbers in the database. Such misleading estimates may occur with conventional techniques. However, in yet another aspect of the present disclosure, a new level in the hierarchy of the data structure that corresponds to the scope with enhanced precision may be determined. Referring now to FIG. 4, a user may send a request to re-estimate the amount of products associated with product numbers within the range of 201 to 500 with enhanced precision. In this instance, estimator 302 may move to the next level in the hierarchy from interior node 308 to leaf node 314, as shown in FIG.4. In leaf node 314, estimator 302 may use the content of the leaf node to determine that the estimate does not include records associated with a product number higher than 240. The ability to enhance the estimate in this manner is one of the advantages of the techniques disclosed herein.

[0019] While the examples herein make reference to a single index data structure, it is understood that multiple indices may be used. The criteria of the estimate request may involve data located in various data sets each with its own index configuration. In this instance, a second node in a second data structure whose child nodes or content fall within the scope of the criteria may also be located. The estimate may be further calculated based at least partially on the child nodes or content of this second node. It is further understood that multiple indices may involve using a single index data structure multiple times (e.g., a self-join). Moreover, the examples herein are not limited to inner joins, but may also be applicable to semi-joins, outer joins, etc. Predicates across multiple indices may include conjunctive joins, disjunctive joins, unions, etc.

[0020] Advantageously, the above-described apparatus and method may provide more accurate estimates of database cardinality or selectivity by using the index of a database as its source of information. In this regard, balanced and well maintained index data structures may be used to provide estimates of cardinality or selectivity that are more precise than conventional techniques. In turn, users or consumers basing decisions on such estimates may be rest assured that the estimates are precise and reliable.

[0021] Although the disclosure herein has been described with reference to particular examples, it is to be understood that these examples are merely illustrative of the principles of the disclosure. It is therefore to be understood that numerous modifications may be made to the examples and that other arrangements may be devised without departing from the spirit and scope of the disclosure as defined by the appended claims. Furthermore, while particular processes are shown in a specific order in the appended drawings, such processes are not limited to any particular order unless such order is expressly set forth herein. Rather, processes may be performed in a different order or concurrently.

CLAIMS

1. A system comprising:
 - a database comprising an index, the index comprising a hierarchical data structure of interlinked nodes;
 - an estimator which upon execution instructs at least one processor to:
 - read a request to estimate an amount of data that would fall within a scope of a criteria in a predicate without executing the predicate;
 - locate a node in the data structure whose child nodes or content fall within the scope of the criteria; and
 - calculate the estimate based at least partially on the located node.
2. The system of claim 1, wherein the estimator upon execution instructs at least one processor to determine a level in a hierarchy of the data structure that corresponds to the scope of the criteria.
3. The system of claim 2, wherein the node of the data structure is located at the level in the hierarchy of the data structure that corresponds to the scope of the criteria.
4. The system of claim 2, wherein the estimator upon execution instructs at least one processor to determine a new level in the hierarchy of the data structure that corresponds to the scope with enhanced precision.
5. The system of claim 1, wherein the estimator upon execution locates a second node in a second data structure whose child nodes or content fall within the scope of the criteria and to further calculate the estimate based at least partially on the child nodes or the content of the second node.

6. A non-transitory computer readable medium having instructions therein which, if executed, cause at least one processor to:
 - read a request to estimate an amount of data in a B-tree data structure that would fall within a scope of criteria in a predicate without executing the predicate;
 - determine a level in the B-tree that corresponds to the scope;
 - locate a node in the B-tree data structure at the determined level whose child nodes or content fall within the scope of the criteria; and
 - estimate the amount of data based at least partially on the located node.
7. The non-transitory computer readable medium of claim 6, wherein the instructions therein upon execution instruct at least one processor to determine a new level in a hierarchy of the B-tree that corresponds to the scope with enhanced precision.
8. The non-transitory computer readable medium of claim 6, wherein the instructions therein upon execution instruct at least one processor to locate a second node in a second data structure whose child nodes or content fall within the scope of the criteria and to further calculate the estimate based at least partially on the second node.
9. The non-transitory computer readable medium of claim 6, wherein the estimate comprises a cardinality estimate.
10. The non-transitory computer readable medium of claim 6, wherein the estimate comprises a predicate selectivity estimate.
11. A method comprising
 - reading, using at least one processor, a request to estimate an amount of data in a data structure that falls within a scope of criteria in a predicate without executing the predicate;


determining, using at least one processor, a level in the data structure that corresponds to the scope;

if the level is not a leaf level, locating, using at least one processor, an interior node at the determined level whose child nodes fall within the scope of the criteria;


if the level is the leaf level, locating, using at least one processor, a leaf node at the leaf level whose content falls within the scope of the criteria; and

estimating, using at least one processor, the amount of data based at least partially on the interior node or the leaf node.

12. The method of claim 11, wherein the estimate comprises a cardinality estimate.
13. The method of claim 11, wherein the estimate comprises a predicate selectivity estimate.
14. The method of claim 11, locating, using at least one processor, a second interior node in a second data structure whose child nodes fall within the scope of the criteria and to further calculate the estimate based at least partially on the second interior node.
15. The method of claim 11, locating, using at least one processor, a second leaf node in a second data structure whose content falls within the scope of the criteria and to further calculate the estimate based at least partially on the second leaf node.

100***Fig. 1***

2/4

200***Fig. 2***




Fig. 3




Fig. 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/070083**A. CLASSIFICATION OF SUBJECT MATTER****G06F 17/00(2006.01)i, 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 17/00; G06F 15/347; G06F 7/08; G06F 7/00; G06F 17/30

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: database, index, hierarchical, B-tree, estimator, data amount, scope, criteria, predicate, 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.
Y	US 4774657 A (ANDERSON, MARK J. et al.) 27 September 1988 See column 1, lines 52-56; column 2, lines 26-32; column 3, lines 7-42; column 4, lines 11-20; column 6, lines 8-25; claims 1, 9, 12 and 16; and figures 1-2 and 8.	1-3, 5-6, 8-15
A		4, 7
Y	US 8108399 B2 (DUMITRU, MARIUS et al.) 31 January 2012 See column 5, lines 6-67; claim 1; and figures 1-2.	1-3, 5-6, 8-15
A		4, 7
A	US 7805410 B2 (SASAKI, SHIGERO) 28 September 2010 See column 4, line 22 - column 6, line 25; and figures 3-5.	1-15
A	US 7987200 B2 (FAUNCE, MICHAEL S. et al.) 26 July 2011 See column 12, line 35 - column 15, line 43; and figure 5.	1-15
A	US 2009-0299989 A1 (ZHANG, NING et al.) 03 December 2009 See paragraphs [0031]-[0048]; and figures 1A-1B.	1-15

☐ 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

21 August 2014 (21.08.2014)

Date of mailing of the international search report

21 August 2014 (21.08.2014)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

JANG, Ho Keun

Telephone No. +82-42-481-8187



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/070083

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 4774657 A	27/09/1988	EP 0252234 A3 EP 0252234 B1 JP 06-073136 B2 JP 1950065 C JP 62-288935 A	22/01/1992 12/01/1994 14/09/1994 10/07/1995 15/12/1987
US 8108399 B2	31/01/2012	US 2008-0288524 A1 WO 2008-150638 A2 WO 2008-150638 A3	20/11/2008 11/12/2008 19/03/2009
US 7805410 B2	28/09/2010	JP 2008-225575 A JP 5088668 B2 US 2008-0222090 A1	25/09/2008 05/12/2012 11/09/2008
US 7987200 B2	26/07/2011	US 2006-0106777 A1 US 2008-0052286 A1 US 2008-0065594 A1 US 7343369 B2 US 7483888 B2	18/05/2006 28/02/2008 13/03/2008 11/03/2008 27/01/2009
US 2009-0299989 A1	03/12/2009	None	