

(19) AUSTRALIAN PATENT OFFICE

(54) Title
An index structure to access hierarchical data in a relational database system

(51)⁶ International Patent Classification(s)
G06F 12/00 (2006.01) **OBMJP** **G06F**
G06F 17/30 (2006.01) 17/30
G06F 12/00 20060101ALI2007072
20060101AFI2005122 **1BMEP**
PCT/US02/30875

(21) Application No: 2002334721 (22) Application Date: 2002 .09 .26

(87) WIPO No: W003/030032

(30) Priority Data

(31) Number	(32) Date	(33) Country
60/326,052	2001 .09 .28	US
10/171,728	2002 .05 .28	US
60/378,800	2002 .05 .07	US

(43) Publication Date : 2003 .04 .14
(43) Publication Journal Date : 2003 .06 .26

(71) Applicant(s)
Oracle International Corporation

(72) Inventor(s)
Murthy, Ravi; Sedlar, Eric; Jalali, Neema; Agarwal, Nipun

(74) Agent/Attorney
PIZZEYS, GPO Box 1374, Brisbane, QLD, 4001

(56) Related Art
US 6236988 B1
WO 2000/049533 A

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

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
10 April 2003 (10.04.2003)

PCT

(10) International Publication Number
WO 03/030032 A2

- (51) International Patent Classification: **G06F 17/30**
- (21) International Application Number: PCT/US02/30875
- (22) International Filing Date:
26 September 2002 (26.09.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/326,052 28 September 2001 (28.09.2001) US
60/378,800 7 May 2002 (07.05.2002) US
10/171,728 28 May 2002 (28.05.2002) US
- (71) Applicant: **ORACLE INTERNATIONAL CORPORATION** [US/US]; 500 Oracle Parkway, M/S 5op7, Redwood Shores, CA 94065 (US).
- (72) Inventors: **JALALI, Neema**; 3416 Beresford Avenue, Belmont, CA 94002 (US). **SEDLAR, Eric**; 4270 Cesar Chavez Street, San Francisco, CA 94131 (US). **AGARWAL, Nipun**; 4768 Cheney Street, Santa Clara, CA 95054 (US). **MURTHY, Ravi**; 33227 Jamie Circle, Fremont, CA 94555 (US).
- (74) Agents: **BINGHAM, Marcel**; 1600 Willow Street, San Jose, CA 95125 et al. (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, EG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— without international search report and to be republished upon receipt of that report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*



WO 03/030032 A2

(54) Title: AN EFFICIENT INDEX STRUCTURE TO ACCESS HIERARCHICAL DATA IN A RELATIONAL DATABASE SYSTEM

	IRowID	NODE ID KEY	CHILD IDs	AccInfo
402 →	r1	1	{"b",r2,R2},{c",r3,R3}	a1
408 →	r2	2	{"d",R4},{e",R5}	a2
412 →	r3	3	{"f",R6},{a",R7}	a3
414 →				

(57) Abstract: Described is a hierarchical index that captures the hierarchical relationship of a hierarchy emulated by a relational database system. The hierarchical index is implemented, using a database table which contains rows that serve as entries of the hierarchical index. Another table has rows that are associated with nodes in the hierarchy. Each entry in the hierarchical index maps to a row that corresponds to a node in the hierarchy. A node in the hierarchy may be a parent node with one or more child nodes. In this case, the corresponding entry in the hierarchical index contains identifiers which identify other entries in the index, where the other entries correspond to rows associated with child nodes of the parent node.

**AN EFFICIENT INDEX STRUCTURE TO ACCESS HIERARCHICAL
DATA IN A RELATIONAL DATABASE SYSTEM**

RELATED APPLICATIONS

5 The present application is related to U.S. Patent Application No. 10/171,728,
filed by Eric Sedlar on May 28, 2002 and entitled "Maintenance of Hierarchical Index
in a Relational System".

The present application is also related to U.S. Patent Application No.
09/251,757, U.S. Patent No. 6,427, 123 filed by Eric Sedlar on February 18, 1999 and
entitled "Hierarchical Indexing for Accessing Hierarchically Organized Information
10 in a Relational System".

The present application claims priority from U.S. Provisional Patent
Application No. 60/326,052, filed on September 28, 2001, entitled "File Based Access
Provided With a Database System", by Eric Sedlar and Viswanathan Krishnamurthy;

15 The present application claims priority from U.S. Provisional Patent
Application No. 60/378,800, filed on May 7, 2002, entitled SQL Access to Data that
Provides a File System Abstraction, by Nipun Agarwal, Ravi Murthy, Eric Sedlar,
Sivasankaran Chandrasekar, Fei Ge, Syam Pannala, Neema Jalali and Muralidhar
Krishnaprasad.

FIELD OF THE INVENTION

20 The present invention relates to relational database systems and, more
specifically, to techniques for indexing hierarchical data within relational database
systems.

BACKGROUND OF THE INVENTION

25 Humans tend to organize information in categories. The categories in which
information is organized are themselves typically organized relative to each other in
some form of hierarchy. For example, an individual animal belongs to a species, the
species belongs to a genus, the genus belongs to a family, the family belongs to an
order, and the order belongs to a class.

30 With the advent of computer systems, techniques for storing electronic
information have been developed that largely reflected this human desire for
hierarchical organization. Conventional computer file systems, for example, are
typically implemented using hierarchy-based organization principles. Specifically, a
typical file system has directories arranged in a hierarchy, and documents stored in the

directories. Ideally, the hierarchical relationships between the directories reflect some intuitive relationship between the meanings that have been assigned to the directories. Similarly, it is ideal for each document to be stored in a directory based on some intuitive relationship between the contents of the document and the meaning assigned to the directory in which the document is stored.

Fig. 1 shows an example of a typical file system. The illustrated file system includes numerous directories arranged in a hierarchy. Two documents 118 and 122 are stored in the directories. Specifically, documents 118 and 122, both of which are entitled "Example.doc", are respectively stored in directories 116 and 124, which are respectively entitled "Word" and "App4".

In the directory hierarchy, directory 116 is a child of directory 114 entitled "Windows", and directory 114 is a child of directory 110. Similarly, directory 124 is a child of directory 126 entitled "VMS", and directory 126 is a child of directory 110. Directory 110 is referred to as the "root" directory because it is the directory from which all other directories descend. In many systems, the symbol "/" is used to refer to the root directory.

When electronic information is organized in a hierarchy, each item of information may be located by following a "path" through the hierarchy to the entity that contains the item. Within a hierarchical file system, the path to an item begins at the root directory and proceeds down the hierarchy of directories to eventually arrive at the directory that contains the item of interest. For example, the path to file 118 consists of directories 110, 114 and 116, in that order.

Hierarchical storage systems often allow different items to have the same name. For example, in the file system shown in Figure 1, both of the documents 118 and 122 are entitled "Example.doc". Consequently, to unambiguously identify a given document, more than just the name of the document is required.

A convenient way to identify and locate a specific item of information stored in a hierarchical storage system is through the use of a "pathname". A pathname is a concise way of uniquely identifying an item based on the path through the hierarchy to the item. A pathname is composed of a sequence of names, referred to as path elements. In the context of a file system, each name in the sequence of names is a "filename". The term "filename" refers to both the names of directories and the names of documents, since both directories and documents are considered to be "files".

Within a file system, the sequence of filenames in a given pathname begins with the name of the root directory, includes the names of all directories along the path from the root directory to the item of interest, and terminates in the name of the item of interest. Typically, the list of directories to traverse is concatenated together, with some kind of separator punctuation (e.g., '/', '\', or ':') to make a pathname. Thus, the pathname for document 118 is /Windows/Word/Example.doc, while the pathname for document 122 is /VMS/App4/Example.doc.

The relationship between directories (files) and their contained content varies significantly between different types of hierarchically organized systems. One model, employed by various implementations, such as Windows and DOS file systems, requires each file to have exactly one parent, forming a tree. In a more complicated model, the hierarchy takes the form of a directed graph, where files can have multiple parents, as in the UNIX file system in which hard links are used.

In contrast to hierarchical approaches to organizing electronic information, a relational database stores information in tables comprised of rows and columns. Each row is identified by a unique RowID. Each column represents an attribute or field of a record, and each row represents a particular record. Data is retrieved from the database by submitting queries to a database server that manages the database. The queries must conform to the database language supported by the database server.

Structured Query Language (SQL) is an example of a database language supported by many existing database management systems.

Each type of storage system has advantages and limitations. A hierarchically organized storage system is simple, intuitive, and easy to implement, and is a standard model used by most application programs. Unfortunately, the simplicity of the hierarchical organization does not provide the support required for complex data retrieval operations. For example, the contents of every directory may have to be inspected to retrieve all documents created on a particular day that have a particular filename. Since all directories must be searched, the hierarchical organization does nothing to facilitate the retrieval process.

A relational database system is well suited for storing large amounts of information and for accessing data in a very flexible manner. Relative to hierarchically organized systems, data that matches even complex search criteria may be easily and efficiently retrieved from a relational database system. However, the process of formulating and submitting queries to a database server is less intuitive

than merely traversing a hierarchy of directories, and is beyond the technical comfort level of many computer users.

In the past, hierarchically organized systems and relationally organized systems have been implemented in different ways that were not compatible.

5 However, some relationally organized systems incorporate features that allow the systems to emulate a hierarchically organized system. This type of emulation is especially desirable when the storage capability and flexibility of a relational system is needed, but the intuitiveness and ubiquity of the hierarchical system is desired.

10 One such feature is based on the connect-by clause defined by SQL. The connect-by clause allows a user to issue queries that request data based on a hierarchical organization. The data is returned by a relational database system in a way that reflects the hierarchical organization. The connect-by is used to specify the condition that defines the hierarchical relationship upon which the hierarchical organization is based.

15 However, using the connect-by clause to formulate queries has disadvantages. First, computing such queries can entail computing multiple join operations, a process that can be very expensive to the database server processing the queries. Use of the connect-by clause is also more burdensome to users. Incorporating a connect-by clause into queries further complicates the already complex task of formulating
20 queries.

Consequently, it is desirable to provide a mechanism that allows relational database systems to emulate hierarchically organized systems in ways that are more efficient than conventional mechanisms for this type of emulation. It is further desirable that this type of emulation be provided in a way that mitigates the
25 complexity of formulating queries that request and return hierarchically organized data.

The reference to any prior art in this specification is not, and should not be taken as an acknowledgement or any form of suggestion that the referenced prior art forms part of the common general knowledge in Australia.

30 SUMMARY OF THE INVENTION

In one broad form of the invention there is provided a computer-readable medium having an index, wherein:

said index indexes a table that includes one or more fields;

said index is arranged and ordered based on key values associated with at least one of said one or more fields;

said index contains a plurality of index entries, wherein each entry of said index entries:

5 maps to at least one corresponding row in said table, and
contains first access control data, this is not reflected in said key values, that defines user access privileges to access said at least one corresponding row;

wherein said index is updated by a database server in response to receiving a database statement, the database server using said index to execute said database
10 statement;

wherein said database statement specifies operations that involve a table; and
wherein said database statement does not reference said index.

In a another form of the invention, there is provided a method for resolving a path name that has a sequence of path elements, said method comprising the steps of:

15 accessing a first entry in an index that corresponds to a particular path element from said sequence of path elements;

determining, based on access control data in the first entry, whether a user may access a first item associated with the particular path element in a particular way;

20 wherein said index is updated by a database server in response to receiving a database statement, the database server using said index to execute said database statement; and

wherein said database statement does not reference said index.

In a related form of the invention, there is provided a computer-readable medium carrying one or more sequences of instructions used for resolving a path name
25 that has a sequence of path elements, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of the above method.

In a further form of the invention there is provided a method of accessing information managed by a database server, said method comprising the steps of:

30 receiving at a database server a statement that conforms to a database language supported by the database server, wherein said statement specifies operations that involve a table;

wherein said table contains rows that are associated with a hierarchy described by a hierarchical index;

wherein a certain entry in said hierarchical index is associated with a certain node in the hierarchy;

wherein said certain entry identifies another entry in said hierarchical index associated with a child node of said certain node; and

5 wherein said database statement does not reference said hierarchical index; and in response to receiving the database statement, the database server using said hierarchical index to execute said statement.

10 In related form of the invention there is provided a computer-readable medium carrying one or more sequences of instructions used for accessing information managed by a database server, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of the above method.

15 In the specification the term "comprising" shall be understood to have a broad meaning similar to the term "including" and will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps. This definition also applies to variations on the term "comprising" such as "comprise" and "comprises".

BRIEF DESCRIPTION OF THE DRAWINGS

20 The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

FIG. 1 is a block diagram illustrating a hierarchical file system.

FIG. 2 is a block diagram illustrating an information hierarchy.

25

FIG. 3 is a block diagram illustrating tables that may be used to capture the information hierarchy illustrated in FIG. 2 within a relational system according to an embodiment of the invention.

FIG. 4 is a block diagram of a database table that serves a hierarchical index according to an embodiment of the present invention.

FIG. 5 is a block diagram of a system on which embodiments of the invention may be implemented.

DETAILED DESCRIPTION OF THE INVENTION

A method and apparatus for accessing hierarchical information stored in a relational database system is described. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

OVERVIEW

Described herein is a new implementation of a hierarchical index that captures the hierarchical relationship of a hierarchy emulated by a relational database system. The hierarchical index is implemented, using a database table which contains rows that serve as entries of the hierarchical index. Another table has rows that are associated with nodes in the hierarchy. Each entry in the hierarchical index maps to a row that corresponds to a node in the hierarchy. A node in the hierarchy may be a parent node with one or more child nodes. In this case, the corresponding entry in the hierarchical index contains identifiers which identify other entries in the index, where the other entries correspond to rows associated with child nodes of the parent node.

In addition, the index contains information about how users may access the rows associated with the hierarchy. The information may be used to determine the access privileges of a user during the course of performing operations that involve the index, allowing both operations and the task of determining the access privileges to be performed more efficiently overall.

Finally, a database server may use the hierarchical index to execute statements like a native index supported by the database server. The type of statements that can be supported in this way include Data Definition Language ("DDL") statements, and

Data Manipulation Language ("DML") statements. Both types of statements are written in a database language, such as SQL.

SYSTEM OVERVIEW

FIG. 2 is a block diagram that illustrates a hierarchy 200 that is used in examples that are provided herein to facilitate an understanding of embodiments of the invention. Hierarchy 200 includes eight nodes. The highest node in the hierarchy is referred to as the "root" node. The nodes at the end of each branch in the hierarchy are "leaf" nodes. The nodes between the root node and the leaf nodes are "intermediate" nodes. In the illustrated hierarchy, nodes 1, 2, and 3 are intermediate nodes, and nodes 4, 5, 6, and 7 are leaf nodes.

In an information hierarchy, the nodes correspond to information. Typically, the piece of information associated with each node will have some form of name, and some type of content. For example, in a hierarchy that corresponds to a hierarchical file system, the nodes will typically correspond to files (where a "folder" or "directory" is one type of file). Each such file will have a name, and some form of content.

FIG. 3 is a block diagram of two tables 302 and 350 that may be used to represent hierarchy 200 in a relational database system. Table 302 includes one row for each node in hierarchy 200. A RowID pseudo-column RRowID has RowIDs that identify a row in table 302. Column NODE contains logical identifiers (herein "node ids") that uniquely identify a node in hierarchy 200. Column NODE may be a primary key that contains primary key values. Column DATA contains values representing data associated with a node. The row for a given node in table 302 includes the row's RowID, a node id identifying the node, and the data associated with the node. For example, row 304, which is identified by RowID R1, corresponds to node 1, data 306 associated with node 1 and its content. Rows in table 302 are herein referred to by their respective RowID.

Table 350 includes rows that each define a parent-child relationship in hierarchy 200. Columns PARENT and CHILD contain node identifiers. Column CHILD NAME contains the "child names" of a child for a particular parent-child relationship in hierarchy 200. For the particular parent-child relationship defined by a row in table 350, column PARENT contains a node id that identifies the parent node, column CHILD contains a node id that identifies the child node, and CHILD NAME contains the child name for that child under the particular parent-child relationship. Similarly, rows 354 and 356 respectively indicate that node 1 is the parent of nodes 2

and 3. CHILD NAME in row 354 specifies that the name of node 2 under the parent-child relationship represented by row 354 is "b."

Although not explicitly depicted in hierarchy 200, a node may have multiple parents in an hierarchy, and for each of these parent-child relationships, a different child name. For example, node 4 may be a child of node 1, and for this parent-child relationship, have a child name Z. Thus, the path represented by this parent-child relationship is "/a/Z". For the row of table 350 representing this parent-child relationship, PARENT contains 1, CHILD contains 4, and CHILD NAME contains Z.

Child name is one example of a link property, that is, a property specific to a parent-child relationship rather than the parent or the child. In another embodiment, table 350 may contain other columns for other link properties. For example, the link property may designate whether or not a parent-child relationship can be seen by any one else other than those with the highest level of access authority for a system (e.g. system administrators). Or, a link property may designate that a parent-child relationship is fixed, that is, it cannot be changed by end users. The rows in table 350 for such fixed parent-child relationships may be cached in volatile memory indefinitely because they very unlikely to be changed.

HIERARCHICAL INDEX

FIG. 4 shows a hierarchical index 402 that describes the hierarchical relationship of hierarchy 200 as represented by tables 302 and 350, according to an embodiment of the present invention. Index 402 is a table with multiple columns and multiple rows. Each row is an index entry. For each intermediate node in hierarchy 200, its corresponding entry in index 402 identifies the index entries of the child nodes of the intermediate node, and rows in table 302 that correspond to the child nodes. Index 402 contains no entries for leaf nodes.

Column IRowID in index 402 is a RowID psuedo-column that has RowIDs that identify an entry in index 402. Column NODE ID KEY contains the primary key values for index 402, the key values being the node ids in column NODE of table 302. Column Child IDs contains a collection of composite-ids, each composite-id containing a child name of a child node, a RowID identifying the entry, if any, in index 402 corresponding to the child node, and a RowID identifying the row in table 302 corresponding to the child node. Child IDs may be implemented as, for example, a column of the data type "character large binary object". The data type allows, for a given entry, many composite-ids to be stored in a column. Column AccInfo contains access information for accessing nodes and their corresponding rows in table 302.

For a given node and its corresponding entry in index 402, the entry contains composite-ids that identify the other entries in index 402 that correspond to the child nodes of the given node, but only if those child nodes are intermediate nodes. For example, in entry 408, column NODE ID KEY contains the node id value 1. Thus, entry 408 corresponds to node 1. Child IDs in entry 408 contains composite-ids {"b",r2,R2} and {"c",r3,R3}, which identify entries 412 and 414, each of which correspond to intermediate nodes 2 and 3, respectively. Node 2 is a child of node 1. Composite-id {"b",r2,R2} specifies that row R2 in table 302 corresponds to child node 2, and that its child name is "b". In entry 412, NODE ID KEY contains the value 2, and Child IDs contains {"d",R4},{ "e",R5}. Composite-ids {"d",R4} and {"e",R5} identify no entry in index 402, indicating that the corresponding child nodes are leaf nodes. Composite-id {"d",R4} identifies as a child node the node corresponding to row R4 in table 402, which is node 4. Node 4 is a leaf node.

Tables 302 and 350 and index 402 capture information of the hierarchy 200 in a relational format. While embodiments of the invention shall be described hereafter through the use of examples that refer to hierarchy 200, tables 302 and 350, and index 402, such embodiments are merely exemplary. The way a relational database system stores information about hierarchies may vary from implementation to implementation, and the present techniques are not limited to any particular implementation.

EXEMPLARY TRAVERSAL OF INDEX

Index 402 may be traversed to access nodes in hierarchy 200 in response to requests based on the nodes' position within hierarchy 200. For example, a query is issued requesting data associated with the child nodes of the node identified by the path "/a/b", node 2. Such a query may be formulated using operators described in U.S. Patent Application Serial No. _____, filed on equal day herewith, entitled OPERATORS FOR ACCESSING HIERARCHICAL DATA IN A RELATIONAL SYSTEM, by Nipun Agarwal, Ravi Murthy, Eric Sedlar, Sivasankaran Chandrasekar and Fei Ge (Attorney Docket No. 50277-1975). To get the child nodes, the index entry corresponding to node named "a", entry 408, is accessed. The entry in index 402 corresponding to next node in the path is determined by examining composite-ids in column Child IDs of entry 408. Composite-id {"b",r2,R2}, which contains a child name matching the next node in the path "/a/b", identifies RowID r2, the RowID of entry 412. Next, entry 412 is accessed. The composite-ids in Child IDs of entry 412 are {"d",R4} and {"e",R5}, which identify

rows in table 302 associated with node 2 as child nodes, rows R4 and R5. Rows R4 and R5 correspond to nodes 4 and 5, identified by the paths "/a/b/d" and "/a/b/f", respectively. These rows are accessed using their RowIDs to retrieve the requested data in the rows.

BENEFITS OF STORING MANY ROW IDS IN ONE CELL

An advantage of index 402 is that data that identifies index entries and rows in table 302 of the child nodes of a parent node can reside within one data block. A data block is an atomic unit of storage for storing database data managed by a database, and is the unit of data at the lowest level of granularity that may be loaded into volatile memory from static storage. Accessing a data block can be an expensive operation because the data block may have to be loaded from static storage into volatile memory. Significant efficiency gains can be realized for a process that accesses data blocks by reducing the number of data blocks needed to perform the process. Thus, the ability to store all the data needed to identify the index entries and rows in table 302 of the child nodes of a parent node is an advantageous feature, because then only one data block need be accessed to get this data.

This feature benefits in particular a process referred to as "path resolution" or as "resolving a path". Path resolution refers to the set of operations performed to determine what one or more entities are identified by a path. It is a common and important function performed by any system that accesses hierarchically organized data based on a path. Therefore, improving its efficiency is particularly important.

For example, to resolve the path "/a/b", the index entry corresponding to node a, entry 408, is accessed. The composite-ids in Child IDs of entry 408 are evaluated to determine that they include r2, which identifies the entry for node b. Thus, the path identifies a valid node, and row R4 in table 302 is a data structure to which path "/a/b" is resolved.

As demonstrated above, resolving path "/a/b" requires accessing one index entry and one data block for each level in the path, except the last. Thus, the number of data blocks accessed to resolve the path using index 402 is linearly proportional to the number of levels in the path.

This linear relationship exists when the RowIDs needed to identify all child node index entries can be stored in one data block, a condition that is likely true for hierarchically organized data typically stored on database servers. For example, a typical database server may have an average block size of 8 k ("kilobytes") and an average RowID size of 16 bytes. A data block can therefore store enough RowIDs to

identify the entries of approximately 500 child nodes. This threshold number of child nodes is unlikely to be exceeded for a given parent in most hierarchies represented by a database server.

Another advantageous feature of index 402 is that it is structured and managed as a table by a database server. This allows index 402 to be concurrently accessed using powerful native features that enable efficient concurrent access to database tables. An example of such a feature is row level concurrency. In row level concurrency, to access a row in a table a process need only lock the row. Another form of concurrency that can be less efficient is table level concurrency. In table level concurrency, to access a row in a table or any portion thereof, a process must lock the whole table. In general, multiple concurrent processes can be executed more efficiently if they can access the same data structure concurrently. Under row level concurrency, multiple processes can access a table concurrently if they access different rows in the table. A process that locks a row does not block a process that requires access to another row. However, with the table level concurrency, to access a row a process must lock the whole table, blocking other processes requiring access to any row in the table, even those processes that need to access a row not needed or accessed by the process locking the table.

PRE-COMMIT CACHE

Transactions executing on a database system change parent-child relationships and the rows in tables 302 and 350 that represent the relationships. In general, when a transaction changes a row, the row is locked before the transaction is committed. For example, if the parent-child relationship between node 1 and node 2 is changed by a transaction, then row 354 and entry 408 are locked. The lock of row 354 blocks other processes attempting to change the parent-child relationship between node 1 and node 2. However, the lock to entry 408 blocks processes attempting to not only change this parent-child relationship but also others, i.e. the parent child relationship between node 1 and node 3. For purposes of changing a parent-child relationship, the locking of rows on table 402 occurs at the parent level, while the lock of the rows on table 350 occur at a lower level of granularity at the parent-child relationship level.

To reduce the collateral blocking affect of changing rows in table 402 that are made to reflect parent-child relationship changes, and to improve concurrency, the locking of rows changed by a transaction is deferred until the transaction is about to commit. The changes to rows made by a transaction are tracked in a "pre-commit cache". Only when the transaction is about to commit are the changed rows locked,

reducing the overall time the rows are locked for a transaction and the collateral blocking affects that would otherwise occur.

USING ACCESS CONTROL DATA WHILE TRAVERSING HIERARCHICAL INDEX

Column AccInfo contains access control data used to determine user access privileges, that is, what way, if any, may a user, group of users, or class of users, access a node in hierarchy 200. For a particular entry in index 402, AccInfo contains data used to determine user privileges for a node and the row corresponding to the node. The data may take the form of data that explicitly specifies user access privileges, data that refers to a source of access control data, or a combination thereof.

According to an embodiment of the present invention, access control data is managed and maintained by a database server for accessing table 350 and other tables managed by the database server. Such access control data is referred to herein as table access control data. Table access control data for a table may be stored at least in part in, for example, one or more columns of the table.

The data stored in AccInfo reflects (i.e. is consistent with) table access control data for table 350. Thus, for a particular entry, the access control data stored in AccInfo indicates user access privileges that are consistent with user privileges specified by table access control data for the row. For example, if data in AccInfo for entry 408 indicates that a user can read but not write data in node 1, table access control data for row R1 indicates that the same user can read but not write data to row R1.

According to an embodiment of the present invention, access control data that governs access to hierarchical and relationally organized data can be implemented as described in U.S. Patent Application Serial No. _____, filed on equal day herewith, entitled MECHANISM FOR UNIFORM ACCESS CONTROL IN A DATABASE SYSTEM, by Ravi Murthy, Eric Sedlar, Nipun Agarwal, Sam Idicula, and Nicolas Montoya (Attorney Docket No. 50277-1980).

The user access privileges defined by access control data in AccInfo include, without limitation, the right to read content of a node, to write content of the node, to define user privileges for the node, to delete the node, and to traverse the node. The right to traverse a node refers to the ability to perform any type of access to the descendants of the node. For example, with respect to node 1, a user has a right to traverse node 1 but not read or write its content. The user may access the children of

node 1, which are node 2 and 3, but may not read the content of node 1, i.e. read row R1.

Storing access control data in AccInfo enables user access privileges to be determined more efficiently in the course of performing operations that involve traversing index 402, such as path resolution. Because the access control data is stored in the entries of index 402, which are already being accessed during traversal, there is no need to get access control data from another source, such as access control data stored in table 302, except for the final node in a path. For example, to resolve the path "/a/b", entry 408 is accessed but not entry 412, which contains access control information in AccInfo for node 2, the final node in the path. Instead, this information is obtained from table 302.

Path resolution may be performed as an atomic operation that includes not only identifying a node specified by a path but determining whether a user has a particular user privilege needed to perform a particular type of operation involving the node. Access control data in AccInfo allows this type of path resolution to be performed more efficiently. For example, while traversing hierarchical index 402 to resolve path "/a/b/c" for a user, entry 408 is accessed. Data in AccInfo is examined to determine that the user may traverse node 1. Next, entry 412 is accessed. Data in AccInfo for entry 412 is examined to determine that the user may not traverse node 2. Thus, the user cannot see any of the children of node 2, including the child node identified by the path, which is node 3. Path resolution is complete and does not need continued traversal of index 402. Thus, path resolution was performed while not only avoiding accessing table access control data for table 350, but also by avoiding completely traversing entries corresponding to each path element of the path.

Storing access control data in an hierarchical index is one example of a type of index that may contain access control information. Other types of indexes, such as B-tree indexes, may contain access control information about the items indexed. The access control information stored in the other types of indexes may be used to improve processes that involve both index traversal and access control. Therefore, the present invention is not limited to storing access control information in hierarchical indexes.

INTEGRATION OF HIERARCHICAL INDEX

An advantage of native indexes supported by a database server is that the database server may use them to execute database statements that do not specify whether and how to use the index. This ability relieves the user from onerous task of

formulating queries that specify operations needed to use the index. For example, a database server receives a request to execute a query that does not specify whether or how to use an index. The database server then generates an execution plan that defines operators, definitions of a set of operations that carry out a task, and an order in which to execute the operators. The operators include an operator for accessing a particular index. When generating an execution plan, the database server evaluates various factors that affect efficiency. Once an execution plan is generated, the database executes the plan's operators, including the operator for accessing the index.

When the database server is able to automatically use an index to execute a database statement without need for the database statement to specify whether or how to use the index, the index or its use is referred to as being "below the database language layer", or "below the SQL layer".

To support an index in a manner that is below the database command layer, software for a database server may be programmed to support the index in this way. Another way to enable this type of support is through the use of a mechanism referred to as extensible indexing. Extensible indexing is described in U.S. Patent No. 5,893,104, entitled *Extensible Indexing*, issued on April 6, 1999 to Jagannathan Srinivasan, Ravi Murthy, Chin Hong, Samuel DeFazio, and Anil Nori, the contents of which are herein incorporated by reference. Extensible indexing allows a database server, which does not have built-in support to support an index type, to extend its index capabilities to support a new index type, by registering the index type and index routines (e.g. object methods) invoked by the database server to use and support the indexes that are instances of the index type. In general, index routines include routines that create, delete, and modify the definition of indexes (DDL index routines), routines that modify the data in existing indexes (DML index routines), routines that retrieve data from existing indexes (query processing index routines), and routines that are called to generate and evaluate execution plans (query optimization index routines).

Extensible indexing allows the database server to perform, automatically and below the database language layer, the operations needed to use and support an index of a particular index type. For example, a database server receives a DDL statement to drop or truncate table 350. The DDL statement references table 302 but not index 402. When the database server executes the DDL statement, the index is automatically dropped or truncated by invoking and executing a DDL index routine. When a database server receives a query, it evaluates and generates an execution plan, and in

the course of doing so, invokes query optimization index routines that participate in evaluating whether and how to use index 402. Once it generates the execution plan, the database server executes the execution plan, invoking query processing index routines as needed to carry out the execution plan.

To create an index, a user issues a create index DDL statement. According to an embodiment of the present invention, a create index statement for a hierarchical index specifies as arguments a resource table and a link table. A resource table, such as table 302, contains content of nodes (logically, physically, or combination thereof) in a hierarchy, such as hierarchy 200. A link table, such as table 350, links rows that represent parent nodes to rows that represent child nodes of the parent nodes. A database server defines a table object type for resource tables (resource table type), and a table object type for link tables (link table type). A create index DDL command for a hierarchical index specifies a table that is an instance of the resource table type and a table that is an instance of the link table type. The link table type defines column attributes (e.g. PARENT and CHILD table 302) that hold node ids for mapping a parent node to a child node. The resource table type defines a column attribute for the node ids (e.g. NODE in table 302). A DDL index routine for creating a hierarchical index takes arguments of the type resource table type and link table type.

HARDWARE OVERVIEW

Figure 5 is a block diagram that illustrates a computer system 500 upon which an embodiment of the invention may be implemented. Computer system 500 includes a bus 502 or other communication mechanism for communicating information, and a processor 504 coupled with bus 502 for processing information. Computer system 500 also includes a main memory 506, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 502 for storing information and instructions to be executed by processor 504. Main memory 506 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 504. Computer system 500 further includes a read only memory (ROM) 508 or other static storage device coupled to bus 502 for storing static information and instructions for processor 504. A storage device 510, such as a magnetic disk or optical disk, is provided and coupled to bus 502 for storing information and instructions.

Computer system 500 may be coupled via bus 502 to a display 512, such as a cathode ray tube (CRT), for displaying information to a computer user. An input

device 514, including alphanumeric and other keys, is coupled to bus 502 for communicating information and command selections to processor 504. Another type of user input device is cursor control 516, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 504 and for controlling cursor movement on display 512. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

The invention is related to the use of computer system 500 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 500 in response to processor 504 executing one or more sequences of one or more instructions contained in main memory 506. Such instructions may be read into main memory 506 from another computer-readable medium, such as storage device 510. Execution of the sequences of instructions contained in main memory 506 causes processor 504 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

The term "computer-readable medium" as used herein refers to any medium that participates in providing instructions to processor 504 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 510. Volatile media includes dynamic memory, such as main memory 506. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 502. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 504 for execution. For

example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 500 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 502. Bus 502 carries the data to main memory 506, from which processor 504 retrieves and executes the instructions. The instructions received by main memory 506 may optionally be stored on storage device 510 either before or after execution by processor 504.

Computer system 500 also includes a communication interface 518 coupled to bus 502. Communication interface 518 provides a two-way data communication coupling to a network link 520 that is connected to a local network 522. For example, communication interface 518 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 518 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 518 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

Network link 520 typically provides data communication through one or more networks to other data devices. For example, network link 520 may provide a connection through local network 522 to a host computer 524 or to data equipment operated by an Internet Service Provider (ISP) 526. ISP 526 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 528. Local network 522 and Internet 528 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 520 and through communication interface 518, which carry the digital data to and from computer system 500, are exemplary forms of carrier waves transporting the information.

Computer system 500 can send messages and receive data, including program code, through the network(s), network link 520 and communication interface 518. In the Internet example, a server 530 might transmit a requested code for an application

program through Internet 528, ISP 526, local network 522 and communication interface 518.

The received code may be executed by processor 504 as it is received, and/or stored in storage device 510, or other non-volatile storage for later execution. In this manner, computer system 500 may obtain application code in the form of a carrier wave.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

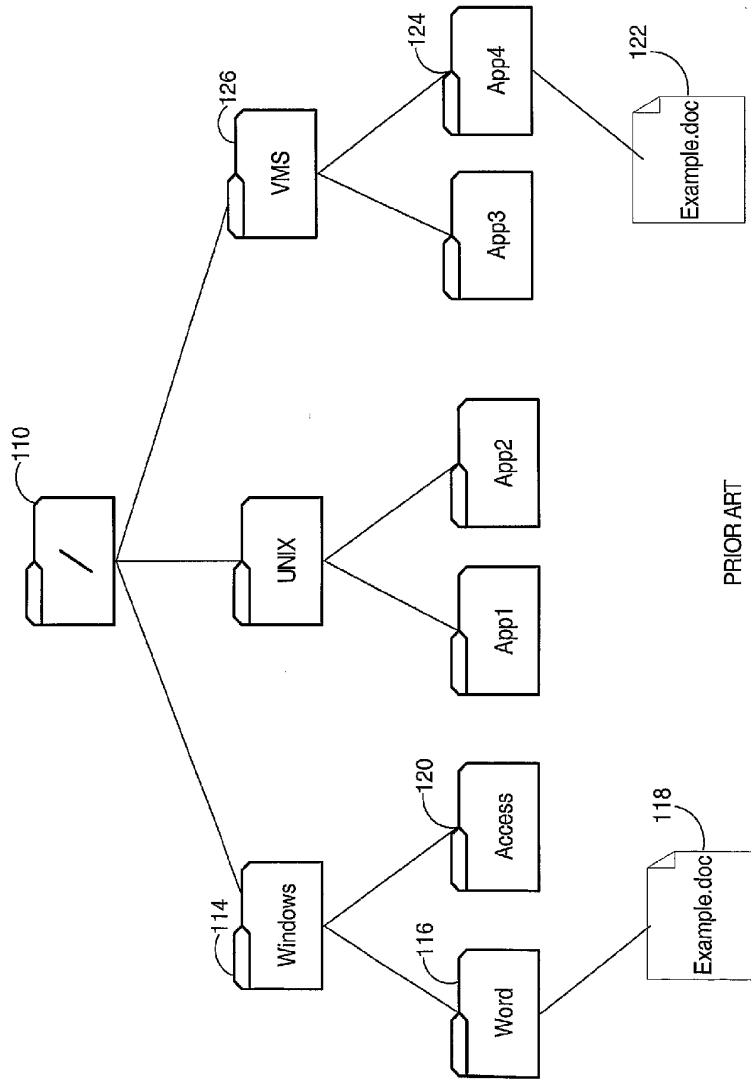
1. A computer-readable medium having an index, wherein:
said index indexes a table that includes one or more fields;
5 said index is arranged and ordered based on key values associated
with at least one of said one or more fields; and
said index contains a plurality of index entries, wherein each
entry of said index entries:
maps to at least one corresponding row in said table, and
10 contains first access control data, this is not reflected in
said key values, that defines user access privileges
to access said at least one corresponding row;
wherein said index is updated by a database server in response to receiving a
database statement, the database server using said index to execute
15 said database statement;
wherein said database statement specifies operations that involve a table; and
wherein said database statement does not reference said index.
2. The computer-readable medium of claim 1, wherein:
said table contains rows that are associated with nodes in a
20 hierarchy;
second access control data defines user access privileges for accessing the
rows; and
said first access control data reflects said user access privileges defined by said
second access control data.
- 25 3. The computer-readable medium of claim 1, wherein:
said plurality of index entries are associated with nodes in a
hierarchy; and
wherein said first access control data indicates that one or more
users have user access privileges for traversing one or
30 more nodes of said nodes.
4. A method used for resolving a path name that has a sequence of
path elements, comprising the steps of:

- 5
accessing a first entry in an index that corresponds to a particular
path element from said sequence of path elements;
determining, based on access control data in the first entry,
whether a user may access a first item associated with the
particular path element in a particular way;
wherein said index is updated by a database server in response to receiving a
database statement, the database server using said index to execute
said database statement; and
wherein said database statement does not reference said index.
- 10 5. The method of claim 4, wherein:
at least one path element follows said particular path element in
said sequence; and
the steps further include attempting to access an entry in said
index that corresponds to said at least one path element
only if said user may access said first item in said
particular way.
- 15
6. The method of claim 4, wherein said first item is a row in a table
that contains rows that correspond to nodes in a hierarchy.
- 20 7. The method of claim 4, wherein:
items correspond to nodes in a hierarchy and said first item
corresponds to a particular node in the hierarchy; and
wherein the step of determining whether a user may access a first
item includes determining whether the user is permitted to
access any descendant node of the particular node.
- 25 8. A method of accessing information managed by a database server, comprising
the steps of:
receiving at a database server a statement that conforms to a database language
supported by the database server, wherein said statement specifies
operations that involve a table;
30 wherein said table contains rows that are associated with a hierarchy described
by a hierarchical index;

- wherein a certain entry in said hierarchical index is associated with a certain node in the hierarchy;
- wherein said certain entry identifies another entry in said hierarchical index associated with a child node of said certain node; and
- 5 wherein said statement does not reference said hierarchical index; and in response to receiving the database statement, the database server using said hierarchical index to execute said statement.
9. The method of claim 8, wherein:
- 10 said database statement specifies a DROP operation for said table; and the step of using said hierarchical index includes performing said DROP operation on said hierarchical index.
10. The method of claim 8, wherein:
- 15 said database statement specifies a TRUNCATE operation for said table; and the step of using said hierarchical index includes performing said TRUNCATE operation on said hierarchical index.
11. The method of claim 9, wherein the steps further include generating an execution plan to execute said statement, said execution plan including an operator defining an operation that uses said hierarchical index.
12. A computer-readable medium carrying one or more sequences of instructions used for resolving a path name that has a sequence of path elements, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of:
- 20 accessing a first entry in an index that corresponds to a particular path element from said sequence of path elements;
- 25 determining, based on access control data in the first entry, whether a user may access a first item associated with the particular path element in a particular way;
- wherein said index is updated by a database server in response to receiving a database statement, the database server using said index to execute said database statement; and
- 30 wherein said database statement does not reference said index.

13. The computer-readable medium of claim 12, wherein:
at least one path element follows said particular path element in
said sequence; and
the steps further include attempting to access an entry in said
5 index that corresponds to said at least one path element
only if said user may access said first item in said
particular way.
14. The computer-readable medium of claim 12, wherein said first
item is a row in a table that contains rows the correspond to
10 nodes in a hierarchy.
15. The computer-readable medium of claim 12, wherein:
items correspond to nodes in a hierarchy and said first item
corresponds to a particular node in the hierarchy; and
wherein the step of determining whether a user may access a first
15 item includes determining whether the user is permitted to
access any descendant node of the particular node.
16. A computer-readable medium carrying one or more sequences of instructions
used for accessing information managed by a database server, wherein
20 execution of the one or more sequences of instructions by one or more
processors causes the one or more processors to perform the steps of:
receiving at a database server a statement that conforms to a database language
supported by the database server, wherein said statement specifies
operations that involve a table;
25 wherein said table contains rows that are associated with a hierarchy described
by a hierarchical index;
wherein a certain entry in said hierarchical index in associated with a certain
node in the hierarchy;
wherein said certain entry identifies another entry in said hierarchical index
associated with a child node of said certain node; and
30 wherein said statement does not reference said hierarchical index; and
in response to receiving the database statement, the database server using said
hierarchical index to execute said statement.

17. The computer-readable medium of claim 16, wherein:
said database statement specifies a DROP operation for said table; and
the step of using said hierarchical index includes performing said DROP
5 operation on said hierarchical index.
18. The computer-readable medium of claim 16, wherein:
said database statement specifies a TRUNCATE operation for said table; and
the step of using said hierarchical index includes performing said
TRUNCATE operation on said hierarchical index.
- 10 19. The computer-readable medium of claim 17, wherein the steps further include
generating an execution plan to execute said statement, said execution plan
including an operator defining an operation that uses said hierarchical index.
- 15 20. A method of accessing information managed by a database server, substantially
as hereinbefore described with reference to any one or more of FIGs 2 to 5 of
the accompanying drawings.
- 20 21. A method used for resolving a path name that has a sequence of path elements,
substantially as hereinbefore described with reference to any one or more of
FIGs 2 to 5 of the accompanying drawings.
- 25 22. A computer-readable medium carrying one or more sequences of instructions
used for the method of either claim 20 or claim 21, wherein execution of the
one or more sequences of instructions by one or more processors causes the
one or more processors to perform the steps of a respective method.



PRIOR ART
FIG. 1

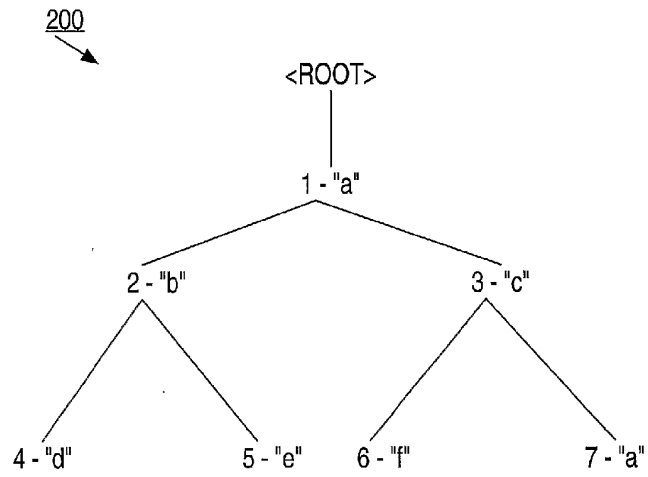


FIG. 2

FIG. 3

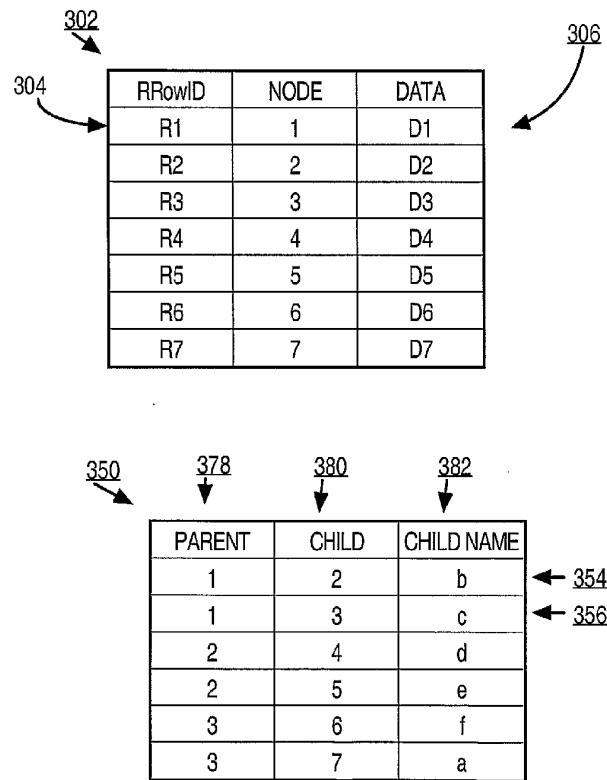


FIG. 4

402

	IRowID	NODE ID KEY	CHILD IDs	AcclInfo
408 →	r1	1	{"b",r2,R2},{"c",r3,R3}	a1
412 →	r2	2	{"d",R4},{"e",R5}	a2
414 →	r3	3	{"f",R6},{"a",R7}	a3

FIG. 5

