



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F	A2	(11) International Publication Number: WO 98/40804
		(43) International Publication Date: 17 September 1998 (17.09.98)
<p>(21) International Application Number: PCT/US98/03569</p> <p>(22) International Filing Date: 24 February 1998 (24.02.98)</p> <p>(30) Priority Data: 60/039,173 26 February 1997 (26.02.97) US</p> <p>(63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Application US 60/039,173 (CIP) Filed on 26 February 1997 (26.02.97)</p> <p>(71) Applicant (for all designated States except US): SIEBEL SYSTEMS, INC. [US/US]; 1855 South Grant Street, San Mateo, CA 94402 (US).</p> <p>(72) Inventors; and</p> <p>(75) Inventors/Applicants (for US only): BRODERSEN, Robert, S. [US/US]; 17 Spinaker Drive, Redwood City, CA 94065 (US). CHATTERJEE, Prashant [US/US]; 21281 Glenmont Drive, Saratoga, CA 95070 (US). LIM, Peter, S. [US/US]; 917 Governors Bay Drive, Redwood City, CA 94065 (US).</p>		<p>(74) Agents: GOLDMAN, Richard, M.; Cooley Godward LLP, Five Palo Alto Square, 3000 El Camino Real, Palo Alto, CA 94306-2155 (US) et al.</p> <p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).</p> <p>Published <i>Without international search report and to be republished upon receipt of that report.</i></p>
(54) Title: DISTRIBUTED RELATIONAL DATABASE		
(57) Abstract		
<p>Method of and apparatus for collecting, storing, and retrieving data in a database management system. The database management system has an enterprise server and at least one workgroup user client. The method includes creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on the workgroup user client, and creating a transaction file corresponding to the transaction log in an outbox of the workgroup user client. The workgroup user client transaction log is read, skipping those transactions which originate at the enterprise server, data files are created corresponding to the entries. The data files corresponding to transactions originating at the workgroup user client are read to an inbox on the enterprise server, thus updating the transactions into an enterprise database on the enterprise server.</p>		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

DISTRIBUTED RELATIONAL DATABASE

INTRODUCTIONTechnical Field

This invention relates to a system and method for providing updates to a network of partially replicated relational database systems, and, more particularly, for providing an efficient means for computing the visibility to a client on the network of a transaction processed against the database.

Background

Relational databases are a commonly employed data structure for representing data in a business or other environment. A relational database represents data in the form of a collection of two-dimensional tables. Each table comprises a series of cells arranged in rows and columns. Typically, a row in a table represents a particular observation. A column represents either a data field or a pointer to a row in another table.

For example, a database describing an organizational structure may have one table to describe each position in the organization, and another table to describe each employee in the organization. The employee table may include information specific to the employee, such as name, employee number, age, salary, etc. The position table may include information specific to the position, such as the position title ("salesman", "vice president", etc.), a salary range, and the like. The tables may be related by, for example, providing in each row of the employee table a pointer to a particular row in the position table, coordinated so that, for each row in the employee table, there is a pointer to the particular row in the position table that describes that employee's position. A relational database management system (RDBMS) supports "joining" these tables in response to a query from a user, so that the user making a query about, for example, a particular employee, may be provided with a report of the selected employee, including not only the information in the employee table, but also the information in the related position table.

Relational databases may be much more complex than this example, with several tables and a multiplicity of relations among them.

With the widespread use of inexpensive portable computers, it is

advantageous to replicate a database onto a portable computer for reference at locations remote from the central computer. The replicated database may then be referenced by the user of the portable computer, without requiring reference to the main database, which may be maintained at a central location inconvenient to the user of the portable computer. However, there are a number of difficulties with the use of a replicated database.

One disadvantage is that a full copy of the central database may require more data storage than is desired or economical. For example, a salesman working in the field may need to refer to the database for information regarding sales opportunities in his sales area, but have no need to refer to any information regarding sales opportunities outside of his area. One possible approach to reduce the amount of required data storage is to simply replicate only that portion of the database that is needed by the user. However, this approach does not recognize that the criteria to determine which portions of the data are required is likely to vary over time. For example, the salesman may have a new city added to his territory. Under conventional approaches, the salesman would need to re-replicate his local copy of the database, this time selecting data including the added city. Such a practice is inconvenient, subject to error, and time-consuming.

A further disadvantage to a replicated database is the difficulties encountered in attempting to update data using the replicated copy. A change made to the replicated database is not made to the central database, leading to a discrepancy between the information that is stored in the replicated copy of the database and the information that is stored in the central database. Although it is possible to journal modifications made to the replicated copy and apply an identical modification to the central database, one problem that this approach faces is the possibility of colliding updates; that is, where a user of a replicated copy makes a change to data that is also changed by a user of the central copy or by the user of another replicated copy.

It is therefore desirable to provide a capability to maintain one or more partially-replicated copies of a central database, in such a way that the degree of replication may be easily changed without requiring a refresh of the entire replicated database, and that permits updates to be coordinated among users of the central database and users of the partially replicated databases.

SUMMARY OF THE INVENTION

The present invention is directed to a method of maintaining a

partially replicated database in such a way that updates made to a central database, or to another partially replicated database, are selectively propagated to the partially replicated database. Updates are propagated to a partially replicated database if the owner of the partially replicated database is deemed to have visibility to the data being updated. Visibility is determined by use of predetermined rules stored in a rules database. In one aspect of the invention, the stored rules are assessed against data content of various tables that make up a logical entity, known as a docking object, that is being updated.

In another aspect of the invention, the stored rules are assessed against data content of one or more docking objects that are not necessarily updated, but that are related to a docking object being updated. In one embodiment, the visibility attributes of the related docking objects are recursively determined.

In yet another aspect of the invention, changes in visibility are determined to enable the central computer to direct the nodes to insert the docking object into its partially replicated database. Such changes in visibility are determined so as to enable the central computer to direct a node to remove a docking object from its partially replicated database.

In a further aspect of the invention, the predetermined rules are in declarative form and specify visibility of databased upon structure of the data without reference to data content.

In still another aspect of the invention, the transactions made to the database are ordered and processed in such a way as to reduce the computational resources required to calculate the visibility of the transactions.

A further aspect of our invention is a method of collecting, storing, and retrieving data in a database management system having an enterprise server, with at least one docking client, and at least one workgroup user client. The docking client is interposed between the enterprise server and the workgroup user client. The method of the invention includes creating a transaction in a local database resident on the workgroup user client, entering the transaction into a transaction log resident on the workgroup user client, and creating a transaction file corresponding to the transaction in an outbox of the workgroup user client. The transaction file is copied to an inbox in the docking client identified to the workgroup user client and updated into an agency database resident on the docking client. The agency database includes a transaction log. The docking client transaction log is read. To be noted is that those transactions which originated at the enterprise server are not read, in order to avoid looping. Files are created corresponding to the entries in the docking client

transaction log, and data files corresponding to transactions originating at the workgroup user client are copied to an inbox on the enterprise server corresponding to the docking server. The transactions are updated into an enterprise database on the enterprise server. An additional aspect of this embodiment of our invention is an article of manufacture that is a computer usable medium having computer readable program embodied therein for causing the above method to be effected. Also encompassed within the scope of our invention is a machine readable program storage device tangibly embodying the above program product.

A still further aspect of our invention is a method of collecting, storing, and retrieving data in a database management system with an enterprise server and at least one workgroup user client. The method comprises: creating a transaction in a local database resident on the workgroup user client, entering the transaction into a transaction log resident on the workgroup user client, creating a transaction file corresponding to the transaction in an outbox of the workgroup user client; and reading the workgroup user client transaction log. As noted above, in reading the workgroup user client transaction log, those transactions which originate at the enterprise server are ignored in order to avoid looping. Data files are created corresponding to the entries in the workgroup user client transaction log, and the data files corresponding to transactions originating at the workgroup user client are copied to an inbox on the enterprise server. This results in updating the transactions into an enterprise database on the enterprise server. An important additional aspect of this embodiment of our invention is an article of manufacture that is a computer usable medium having computer readable program embodied therein for causing the above method to be effected. Also encompassed within the scope of our invention is a machine readable program storage device tangibly embodying the above program product.

A still further aspect of our invention is a method of merging data from one database into another database in a database management system. The database management system has an enterprise server and at least one workgroup user client. The method starts by creating a transaction in a local database resident on the workgroup user client, entering the transaction into a transaction log resident on the workgroup user client, creating a transaction file corresponding to the transaction log entry in an outbox of the workgroup user client, and reading the workgroup user client transaction log. Those transactions which originate at the enterprise server are skipped to avoid looping. Data files are created corresponding to the entries in the workgroup user client transaction log, and the data files corresponding to

transactions originating at the workgroup user client are read to an inbox on the enterprise server. This is followed by transferring a update log file from one database to the other database to merge the transactions.

5

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 depicts an overview of the operation of one embodiment of the present invention.

10

Figure 2 depicts a database schema that shows the relationship of the various components that make up a Docking Object.

Figure 3 depicts steps performed by an update manager to update a database.

15

Figure 4 depicts steps performed by a Docking Manager to transmit and/or receive one or more transaction logs.

Figure 5 depicts the steps performed by a merge processor to merge transaction log records into an existing database.

20

Figure 6 depicts the steps performed by a log manager to prepare a partial transaction log.

25

Figure 7 depicts the steps performed by a visibility calculator for calculating visibility for a docking object as invoked by a log manager.

Figure 8 depicts the steps performed to synchronize a partially replicated database in response to a change in data visibility.

30

DESCRIPTION OF SPECIFIC EMBODIMENTS

Overview

Figure 1 depicts an overview of the operation of one embodiment of the present invention. Figure 1 depicts a central computer system 1 and three remote computer systems (or "nodes") 21-a, 21-b, and 21-c. Each of nodes 21-a, 21-b and 21-c are depicted in various states of communication with central computer system 1, as will be more fully explained. Central computer system 1 includes a central database 3, a docking manager 5, a merge processor 7 and a log manager 9. Central computer system 1 additionally optionally includes update manager 11 responsive to user input 13.

40

Node 21-a is a remote computer system, such as a mobile client such as a laptop computer. Node 21-a includes a partially replicated

45

remote database 23-a, update manager 31-a responsive to user input 33-a, docking manager 25-a and merge manager 27-a. In operation, update manager is responsive to user input 33-a to make changes to remote database 23-a as directed by the operator of node 21-a. Updates made
5 are recorded, or journaled, in node update log 35-a.

At some point at the convenience of the operator of node 21-a, node docking manager 35-a is activated, and enters into communication with central docking manager 5. Update log 35-a is taken as input by
10 node docking manager 25-a, and provided to central docking manager 5. Central docking manager 5 creates a received node update log 19, which contains all the information that had been recorded in update log 35-a. Optionally, partial log 17-a is taken as input by central docking manager 5 and provided to node docking manager 25-a, as more fully
15 described herein.

At some point in time, at the convenience of the operator of central computer system 1, merge processor 7 is activated. Merge processor 7 takes as input received node update log 19, and applies the
20 updates described therein to central database 3. In the process of applying the updates from received node update log 19, merge processor journals the updates applied to central update log 15. Optionally, update manager 11, responsive to user input 12 makes additional changed to central database 3 as directed by the operator of central computer
25 system 1. The updates made by update manager 11 are additionally journaled in central update log 15.

At some point in time, at the convenience of the operator of central computer system 1, log manager 9 is activated. Log manager 9
30 takes as input central update log 15 and produces as output a set of partial logs 17-a, 17-b and 17-c according to visibility rules as will be further described herein. Each of partial logs 17-a, 17-b and 17-c corresponds to one of nodes 21-a, 21-b and 21-c. When a node docking manager such as node docking manager 25-a enters into communication
35 with central docking manager 5 and optionally requests transmission of its corresponding partial log, central docking manager 5 takes as input the appropriate partial log, such as partial log 17-a, and presents it to node docking manager 25-a. Node docking manager 25-a then replicates partial log 17-a as merge log 37-a.
40

At some point in the future, at the convenience of the operator of node 21-a, merge processor 27-a is activated. Merge processor 27-a takes as input merge log 37-a, and applies the updates described therein to partially replicated database 23-a.
45

In addition to node 21-a, Figure 1 also depicts two additional nodes 21-b and 21-c. Node 21-b is depicted in communication with central computer 1. However, unlike node 21-a, the operator of node 21-b has requested only to send his updates to central computer system 1, and has not requested to be presented with changes made elsewhere to be made to his partially replicated database 23-b. This may be, for example, if the operator has an urgent update that must be made as soon as possible, but does not have the time to receive updates from other nodes. Accordingly, Figure 1 shows only transmission of node update log 35-a from node docking manager 25-b to central docking manager 5, and no transmission from central docking manager 5 to node docking manager 25-b. Accordingly, the merge manager for node 21-b is not activated and is not shown.

Likewise, node 21-c is depicted as not in communication with central computer system 1. Accordingly, the docking manager for node 21-c is not activated and is not shown.

By the cycle described above, updates made by each of nodes 21-a, 21-b and 21-c are presented to central computer system 1, permitting central database 3 to be updated accordingly. In addition, each of the updates made by each of the nodes 21-a, 21-b and 21-c, as well as updates made on central computer system 1, are routed back to each of nodes 21-a, 21-b, and 21-c, thereby keeping each of partial databases 23-a, 23-b and 23-c in synchronization with each other and with central database 3.

Database Structure

The synchronization of central database 3 with node databases 23-a, 23-b and 23-c is performed using a construct called a Docking Object. A Docking Object consists of Member Tables (including one Primary Table), Visibility Rules, Visibility Events, and related Docking Objects.

A Member Table is a table of the relational database that makes up a docking object. When a docking object is propagated from central database 3 to one of node databases 23-a, 23-b or 23-c, the propagation takes the form of an insertion into each of the Member Tables associated with the particular docking object. Similarly, when a docking object is scheduled to be removed from a database, that removal consists of deleting records from the member tables associated with the docking object. For example, a docking object that represents a sales opportunity may include tables that represent the opportunity itself (e.g., named "S_OPTY"), the product whose sale is represented by the opportunity (e.g., named "S_OPTY_PROD"), the contact for the

opportunity (e.g., named "S_OPTY_CONTACT"), etc. Each of these tables is said to be a member table of the "Opportunity Docking Object."

5 A Primary Table is a Member Table that controls whether a particular instance of a Docking Object is visible to a particular node. The Primary Table has a Primary Row-ID value that is used to identify a row of the Primary Table being updated, deleted or inserted. For example, the "Opportunity Docking Object" may have as a primary table the table S_OPTY. The row-id of that table, i.e., S_OPTY.row_id,
10 is the Primary Row-ID for the Opportunity Docking Object.

A Visibility Rule is a criterion that determines whether a particular instance of a Docking Object is "visible" to a particular node 21. If a Docking Object is visible to a particular node, that
15 node will receive updates for data in the Docking Object. Visibility Rules are of two types, depending on the field RULE_TYPE. A Visibility Rule with a RULE_TYPE of "R" is referred to as an SQL Rule. An SQL Rule includes a set of Structured Query Language (SQL) statements that is evaluated to determine if any data meeting the criteria specified in
20 the SQL statements exists in the Docking Object. If so, the Docking Object is visible to the node. A Visibility Rule with a RULE_TYPE of "O" is referred to as a Docking Object Rule. A Docking Object Rule specifies another Docking Object to be queried for visibility. If the specified Docking Object is visible, then the Docking Object pointing
25 to it is also visible.

A Related Docking Object is a Docking Object that is propagated or deleted when the Docking Object under consideration is propagated or deleted. For example, an Opportunity Docking Object may have related
30 Docking Objects representing the sales contacts, the organizations, the products to be sold, and the activities needed to pursue the opportunity. When an Opportunity Docking Object is propagated from Central Database 3 to one of node databases 23, the related docking objects are also propagated.

35

Figure 2 depicts a database schema that shows the relationship of the various components that make up a Docking Object. The schema is a meta-database, in that it does not describe the data being accessed in the database. Rather, the schema is a separate database that defines
40 the structure of the database being accessed. That is, it is a database comprising tables that describe the relationships and data contexts of another database.

Each of the tables shown in Figure 2 is a table in a relational database, and as such is in row-column form. Many columns represent
45 fields that are common to all the illustrated tables. Such fields

include for example, a ROW_ID to identify a particular row in the table, as well as fields to tack the date and time that a row was created and last modified, and the identity of the user who created or modified the row. In addition, each table contains fields specific to that table, and which are described in detail below.

Table S_DOBJ 61 describes the Docking Objects in an application. Table S_DOBJ 61 includes the fields OBJ_NAME and PRIMARY_TABLE_ID. Field OBJ_NAME defines the name of the Docking Object being described. Field PRIMARY_TABLE_ID is used to identify the primary table associated with this Docking Object.

Table S_DOBJ_INST 63 describes whether a particular instance of a Docking Object, described by table S_DOBJ 61, is present on a particular node's database. Table S_DOBJ_INST 63 includes the fields NODE_ID, DOBJ_ID and PR_TBL_ROW_ID. Field NODE_ID points to a particular node table 65. Field DOBJ_ID points to the Docking Object to which the Docking Object instance applies. Field PR_TBL_ROW_ID is used to select a particular row in the Primary Table of the Docking Object. This value identifies the Docking Object instance.

Table S_REL_DOBJ 67 describes the related Docking Objects of a particular Docking Object, described by table S_DOBJ 61. Table S_REL_DOBJ 67 includes the fields DOBJ_ID, REL_DOBJ_ID, and SQL_STATEMENT. Field DOBJ_ID identifies the Docking Object that owns a particular related Docking Object. Field REL_DOBJ_ID identifies the related Docking Object that is owned by the Docking Object identified by DOBJ_ID. Field SQL_STATEMENT is an SQL statement that may be executed to obtain the Primary ID value of the related Docking Object.

Table S_DOBJ_TBL 69 describes the member tables of a particular Docking Object, described by table S_DOBJ 61. Table S_DOBJ_TBL 69 includes the fields DOBJ_ID, TBL_ID, and VIS_EVENT_FLG. Field DOBJ_ID identifies the Docking Object that contains the member table described by the row. Field TBL_ID identifies the particular table in the database that is the member table described by the row. Field VIS_EVENT_FLG is a flag that indicates whether a change to this Docking Object can result in a visibility event. A value of "Y" indicates that a change can result in a visibility event; a value of "N" indicates that it cannot.

Table S_DOBJ_VIS_RULE 71 contains the visibility rules associated with a particular Docking Object. S_DOBJ_VIS_RULE 71 contains the fields DOBJ_ID, RULE_SEQUENCE, RULE_TYPE, SQL_STATEMENT and CHECK_DOBJ_ID. Field DOBJ_ID identifies the Docking Object with which

a particular visibility rule is associated. Field RULE_SEQUENCE is a sequence number that indicates the sequence, relative to other visibility rules in table S_DOBJ_VIS_RULE 71, in which the particular visibility rule should be run. RULE_TYPE specifies whether the particular visibility rule is of type "R," indicating an SQL visibility rule or of type "O," indicating a Docking Object visibility rule.

If RULE_TYPE is equal to "R," field CHECK_DOBJ_ID is not meaningful, and field SQL_STATEMENT contains an SQL statement that is evaluated using the Primary ROW-ID of the primary table associated with this Docking Object and a particular Node 21. If the SQL statement returns any records, the Docking Object is deemed to be visible to the Node 21 for which visibility is being determined.

If RULE_TYPE is equal to "O," both field CHECK_DOBJ_ID and field SQL_STATEMENT are meaningful. Field CHECK_DOBJ_ID specifies a docking object whose visibility should be determined. If the specified docking object is deemed to be visible, then the docking object associated with the visibility rule is also visible. Field SQL_STATEMENT contains a SQL statement that, when executed, returns the Row-ID of the docking object identified by CHECK_DOBJ_ID that corresponds to the docking object instance associated with the visibility rule.

Table S_APP_TBL 73 is an Application Table that describes all the tables used in a particular application. It is pointed to by table S_DOBJ_TBL 69 for each member table in a docking object, and by table S_DOBJ for the primary table in a docking object. S_APP_TBL 73 points to table S_APP_COL 75, which is an Application Column Table that describes the columns of data in a particular application. S_APP_TBL 73 points to table S_APP_COL 75 directly through a primary key and indirectly through such means as a Foreign Key Column Table 81, User Key Column Table 83, and Column Group Table 85. The relationship of an Application Table, Application Column Table, Foreign Key Column Table, User Key Column Table and Column Group Table are well known in the art and are not further described.

Update Processing

Update processing is an important aspect of our invention. A further aspect of our invention is update processing, including a method of collecting, storing, and retrieving data in a database management system having an enterprise server, with at least one docking client, and at least one workgroup user client. The docking client is interposed between the enterprise server and the workgroup user client. The method of the invention includes creating a transaction in a local database resident on the workgroup user client, entering the transaction into a transaction log resident on the

workgroup user client, and creating a transaction file corresponding to the transaction in an outbox of the workgroup user client. The transaction file is copied to an inbox in the docking client identified to the workgroup user client and updated into an agency database resident on the docking client. The agency database includes a transaction log. The docking client transaction log is read. To be noted is that those transactions which originated at the enterprise server are not read, in order to avoid looping. Files are created corresponding to the entries in the docking client transaction log, and data files corresponding to transactions originating at the workgroup user client are copied to an inbox on the enterprise server corresponding to the docking server. The transactions are updated into an enterprise database on the enterprise server. An additional aspect of this embodiment of our invention is an article of manufacture that is a computer usable medium having computer readable program embodied therein for causing the above method to be effected. Also encompassed within the scope of our invention is a machine readable program storage device tangibly embodying the above program product.

A still further aspect of our invention is a method of collecting, storing, and retrieving data in a database management system with an enterprise server and at least one workgroup user client. The method comprises: creating a transaction in a local database resident on the workgroup user client, entering the transaction into a transaction log resident on the workgroup user client, creating a transaction file corresponding to the transaction in an outbox of the workgroup user client; and reading the workgroup user client transaction log. As noted above, in reading the workgroup user client transaction log, those transactions which originate at the enterprise server are ignored in order to avoid looping. Data files are created corresponding to the entries in the workgroup user client transaction log, and the data files corresponding to transactions originating at the workgroup user client are copied to an inbox on the enterprise server. This results in updating the transactions into an enterprise database on the enterprise server. An important additional aspect of this embodiment of the invention is an article of manufacture that is a computer usable medium having computer readable program embodied therein for causing the above method to be effected. Also encompassed within the scope of the invention is a machine readable program storage device tangibly embodying the above program product.

A still further aspect of our invention is a method of merging data from one database into another database in a database management system. The database management system has an enterprise server and at least one workgroup user client. The method starts by creating a

transaction in a local database resident on the workgroup user client, entering the transaction into a transaction log resident on the workgroup user client, creating a transaction file corresponding to the transaction log entry in an outbox of the workgroup user client, and
5 reading the workgroup user client transaction log. Those transactions which originate at the enterprise server are skipped to avoid looping. Data files are created corresponding to the entries in the workgroup user client transaction log, and the data files corresponding to transactions originating at the workgroup user client are read to an
10 inbox on the enterprise server. This is followed by transferring an update log file from one database to the other database to merge the transactions.

Figure 3 depicts steps performed by an update manager 31 such as
15 update manager 31-a, 31-b or 31-c in updating a database, such as a node database 23-a, 23-b or 23-c, responsive to user input. Execution of update manager 31 begins in step 101. In step 103, the update manager 31 accepts from the user input 33 in the form of a command requesting that the data in database 23 be altered. The request may be
20 in the form of a request to delete a row of a table, to add a row to a table, or to change the value of a cell at a particular column of a particular row in a table. In step 105, using a well-known means, the update manager 31 applies the requested update to database 23. In step 107, the update manager 31 creates a log record describing the update
25 and writes it to update log 35.

The contents of a log record describe the update made. Each log record indicates the node identifier of the node making the update, an identification of the table being updated, and an identification of the
30 type of update being made, i.e., an insertion of a new row, a deletion of an existing row, or an update to an existing row. For an insertion, the log record additionally includes an identifier of the row being inserted, including its primary key and the values of the other columns in the row. For a deletion, the log record identifies the primary key
35 of the row being deleted. For an update, the log record identifies the primary key of the row being updated, the column within the row being updated, the old value of the cell at the addressed row and column, and the new value of the cell.

40 After writing a log record in step 107, the update processor exits for this update. The foregoing description of the update processing preferably includes additional steps not material to the present invention, for example, to assure authorization of the user to make the update, to stage and commit the write to the database to allow
45 for rollback in the event of software or hardware failure, and the

like. These steps are well-known in the art and are not described further.

An update manager 11 executing in central computer system 1 operates in an analogous manner, except that it updates central database 3 and writes its log records to central update log 11.

Docking Processing

Figure 4 depicts steps performed by a Docking Manager 25 such as Docking Manager 25-a, 25-b or 25-c to transmit and/or receive one or more transaction logs. Docking Manager 25 is invoked by the user of a remote node such as node 21-a, 21-b or 21-c, whereby the user requests that the node dock with central computer 1 to upload an update log such as update log 35-a to central computer 1, to download a partial log such as partial log 17-a, or both. Execution of Docking Manager 25 begins in step 121. In step 123, Docking Manager 25 connects with central computer 1 under the control of Central Docking Manager 5. This connection can be any connection that enables data exchange. It is anticipated that the most common form of a connection is a telephone line used in conjunction with a modem, but other forms of data connection, such as a Local Area Network or a TCP/IP connection may also be used. Step 125 checks to see whether the user has requested that node update log 35-a be uploaded to the Central Computer 1. If so, execution proceeds to step 127. If not, step 127 is skipped and control is given to step 129. In step 127, Docking Manager 25 uploads its update log to central computer 1. The upload may be accomplished with any known file transfer means, such as XMODEM, ZMODEM, KERMIT, FTP, ASCII transfer, or any other method of transmitting data. In step 129, Docking Manager 25 checks to see whether the user has requested that a partial log such as partial log 17-a be downloaded from Central Computer 1. If so, execution proceeds to step 131. If not, step 131 is skipped and control is given to step 133. In step 131, Docking Manager 25 downloads its partial log from central computer 1. The download may be accomplished with any known file transfer means, such as XMODEM, ZMODEM, KERMIT, FTP, ASCII transfer, or any other method of transmitting data. In step 133, having completed the requested data transfer, Docking Manager 25 exits.

Merge Processing

Merge processing is performed by a processor such as node merge processor 27-a, 27-b, or 27-c, or central merge processor 7. The merge process serves to update its associated database with a transaction that has been entered by a user of a computer remote from the computer where merge processing is being performed. Merge processing is analogous to update processing and is similar in form to update processing as previously disclosed with reference to figure 3, with

three differences. First, the input to a merge processor is not an update entered directly by a user, but rather is a log file that is obtained from a computer remote from the computer where the merge is executing. A second difference is that, as shown by in Figure 1, merge processing does not produce a log when performed at a node. The function of a log on a node is to record a transaction for propagation to Central Computer system 1 and thence to other nodes as required. A transaction that is the subject of a merge in a node has been communicated to Central Computer System 1, and there is no need to re-communicate it.

A third difference is that merge processing must be capable of detecting and resolving multiple conflicting transactions. For example, assume that a field contains the value "Keith Palmer." Assume further that a user at node 27-a enters a transaction to update that field to "Carl Lake," and a user at node 27-b enters a transaction to update the same field to "Greg Emerson." Without collision detection, data among various nodes may become corrupt. When the transaction for user 27-a is merged, the field is updated from "Keith Palmer" to "Carl Lake." Without collision handling, when the transaction for node 27-b is merged, the field would be updated to "Greg Emerson," and the central database would then be out of synch with the database of node 27-a. Furthermore, when merge processing is performed on each of nodes 27-a and 27-b, each node will update its database with the other's transactions, leaving at least one node out of synch with the other node and with central database.

Therefore, merge processing must also have a means of detecting collisions and correcting them. In the above example, a simple way to detect and correct a collision is to compare the value in the database to the value that the merge log reflects as being the previous value in the node database. If the two values do not match, Merge processor 7 may reject the transaction and generate a corrective transaction to be sent to the node from which the conflicting transaction originated. In the above example, when the transaction for node 27-b was presented to merge processor 7, merge processor 7 would compare "Keith Palmer," the prior value of the field as recorded by node 27-b to "Carl Lake," the present value of the field as recorded in central database 3. Detecting the mismatch, merge processor 7 may then generate a transaction to change the value "Greg Emerson" to "Carl Lake," and write that transaction to update log 15. In a subsequent docking operation, that transaction would be routed back to node 27-b to bring its database 23-b in synch with the other databases.

The above is one example of a collision and a resulting corrective action. Other types of collisions include, for example, an update to a row that has previously been deleted, inserting a row that has previously been inserted, and the like. Merge processing must

detect and correct each of these collisions. This may be performed using any of a number of well-known methods, and is not discussed further.

Figure 5 depicts the steps performed by merge processor such as central merge processor 7. Although it depicts merge processor 7 writing to central database 3 and to transaction log 15, it is equally representative of a node merge processor such as node merge processor 27-a, 27-b or 27-c updating a node database 23-a, 23-b or 23-c. Merge processing begins at step 141. In step 143, merge processor 7 finds the first unprocessed transaction on received log 19. In step 147, merge processor 7 selects a transaction from received log 19. In step 149, merge processor 149 attempts to update database 3 according to the transaction selected in step 147. In step 151, merge processor 7 determines whether the database update of step 149 failed due to a collision. If so, merge processor proceeds to step 153, which generates a corrective transaction. Following the generation of the corrective transaction, the merge processor returns to step 149 and again attempts to update database 3. If no collision was detected in step 151, execution proceeds to step 157. In step 157, merge processing checks to see if it is executing on central computer 1. If so, step 155 is executed to journal the transaction to log 15. In any case, either if step 157 determines that the merge processing is being performed on a node or after step 155, execution proceeds to step 159. Step 159 checks to see if any transactions remain to be processed from log 19. If so, execution repeats from step 147, where the next transaction is selected. If not, merge processing exits in step 161.

Log Management

Figure 6 depicts the steps to be performed by log manager 9 to prepare a partial transaction log such as partial transaction log 17-a, 17-b, or 17-c. The procedure depicted in Figure 6 is executed for each node available to dock with central computer system 1. Log manager 9 begins execution in step 171. In step 173, Log Manager 9 finds the first unprocessed transaction for the node whose partial transaction log is being prepared. In step 175, log manager 9 selects a transaction for processing. In step 177, log manager 9 checks to see whether the selected transaction originated on the same node for which processing is being performed. If so, there is no need to route the transaction back to the node, and control proceeds to step 179. Step 179 checks to see whether there are any transactions remaining to be processed. If so, control is given again to step 175. If not, control passes to step 189, which records the last transaction that was processed for this node, and then exits at step 191. If the transaction originates in other than the same node as the node for which processing is being performed, control is given to step 181. Step 181 calls a visibility calculator to determine whether the

selected transaction is visible to the node being processed. The Visibility calculator routine is described in detail further herein. In step 183, merge processor 9 checks to see whether the visibility calculator determined that the transaction is visible. If it is not visible, control is passed to step 179, which performs as disclosed above. If the transaction is visible, control is passed to step 185. Step 185 writes a record for this transaction to the partial transaction log for the node being processed, for example, partial transaction log 17-a for node 21-a. In step 187, the log manager 9 records the last transaction that was processed for this node, and then passes control to step 179, which determines whether to select additional transactions or exit, as disclosed above.

Visibility Calculation

Figure 7 depicts a flowchart describing the process a visibility calculator for calculating visibility for a docking object as invoked by step 181 of log manager 9. The visibility calculator is called with the node-id of the node for which visibility is being calculated, the docking object for which the visibility is being calculated, and the row-id of the docking object whose visibility id being calculated. The visibility calculator uses this information, in conjunction with information obtained from meta-data stored in the schema depicted in Figure 2, to determine whether a particular transaction that updates a particular row of a particular docking object is visible to a particular node.

The Visibility calculator begins execution at step 201. In step 203, the visibility calculator makes a default finding that the transaction is not visible. Therefore, unless the visibility calculator determines that a transaction is visible, it will exit with a finding of no visibility. In step 205, the visibility calculator selects the first visibility rule associated with the docking object. This is done by finding the table S_DOBJ_VIS_RULE 71 associated with the current Docking Object as pointed to by table S_DOBJ 61. In step 205, the visibility calculator selects the row of table S_DOBJ_VIS_RULE 71 with the lowest value for field RULE_SEQUENCE.

In step 207, the Visibility Calculator checks the field RULE_TYPE for a value of "R." The value of "R" indicates that the rule is a SQL visibility rule. If so, the Visibility Calculator proceeds to step 209. In step 209 the Visibility Calculator obtains a SQL statement from field SQL_STATEMENT and executes it. An example of such an SQL statement might be:

```
SELECT 'X' FROM S_OPTY_EMP
WHERE OPTY_ID = :PrimaryRowId
AND EMP_ID = :NodeId;
```

This SQL statement causes a query to be made of application table S_OPTY_EMP. The query selects any records meeting two criteria. First, the records selected must have a field OPTY_ID, which is a row

id or key, equal to the Primary Row-ID of the Docking Object whose visibility is being determined. Second, the records selected must have a field EMP_ID, which may be for example, an identifier of a particular employee, equal to the NodeId of the node for whom visibility is being
5 determined. In ordinary language, this SQL statement will return records only if a row is found in a table that matches employees to opportunities, where the opportunity is equal to the one being updated, and the employee to whom the opportunity is assigned is the operator of the node.

10 This is a simplistic example, provided for maximum comprehension. More complex SQL statements are possible. For example, the rule:

```
SELECT 'X' FROM
  &Table_Owner.S_ACCT_POSTN ap
  &Table_Owner.S_EMP_POSTN ep
15 WHERE ap.POSITION_ID = ep.POSITION_ID
      AND ep.EMP_ID = :NodeId;
```

This rule queries the tables S_ACCT_POSTN (which relates a particular account with a particular position in the organization that is responsible for the account) and S_EMP_POSTN (which relates what
20 employee corresponds to a particular position). The condition "ap.POSITION_ID = ep.POSITION_ID" requires finding a row in the account-to-position table that has the same position as a row in the employee-to-position table. The condition "ep.EMP_ID = :NodeId" further requires that the selected row in the employee-to-position
25 table also have an Employee ID equal to the ID of the user of the Node for which visibility is being determined. In ordinary language, this condition allows visibility if the employee occupies the position that has responsibility for the account in the docking object being updated.

There is no particular limit to the complexity of the conditions
30 in the SQL statement used to evaluate visibility. Particular implementations of SQL may impose limitations, and resource considerations may make it desirable to use less complex statements, but these limitations are not inherent in the invention.

Step 211 evaluates whether the execution of SQL_STATEMENT in step
35 209 returned any records. If records were returned, this indicates that the Node for which visibility is being checked has visibility to the docking object being processed. Accordingly, if records are returned, the Visibility Calculator proceeds to step 213. In step 213, the transaction is marked visible. Because no further rules need to be
40 evaluated to determine visibility, the visibility calculator proceeds to step 228. Step 228 synchronizes the databases by determining whether the calculated visibility requires the insertion or deletion of a docking object into a particular node's partially replicated database. This may occur, for example, if a node is determined to have
45 visibility to a docking object due to a change to a related docking object. For example, an owner of a node may be assigned to a

particular activity that is related to a particular sales opportunity. As a result, the node should be provided with a copy of the object representing the sales opportunity.

Figure 8 depicts the steps performed to synchronize a partially replicated database in response to a change in data visibility. Execution begins in step 241. In step 243, the Visibility Calculator references the visibility just calculated for a docking object. If the Docking Object is visible, execution proceeds to step 245. Step 245 references the S_DOBJ_INST table, to verify that a row exists for the Docking Object for the current node. If a row exists, this indicates that the node in question already has a copy of the referenced Docking Object, and the routine proceeds to step 255, where it exits. If, however, no row exists for the Docking Object at the node being processed, this indicates that the node in question does not have a copy of the Docking Object on its partially replicated database. The routine then proceeds to step 247, where a transaction is generated to direct the node to insert the Docking Object into its partially replicated database.

If step 243 determines that the Docking Object is not visible, execution proceeds to step 249. Step 249 references the S_DOBJ_INST table, to verify that no row exists for the Docking Object for the current node. If step 243 determines that no row exists in the S_DOBJ_INST table for the current docking object for the current row, this indicates that the node in question does not have a copy of the referenced Docking Object, and the routine proceeds to step 255, where it exits. If, however, a row exists for the Docking Object at the node being processed, this indicates that the node in question does have a copy of the Docking Object on its partially replicated database. The routine then proceeds to step 251, where a transaction is generated to direct the node to delete the Docking Object from its partially replicated database.

Referring again to Figure 7, following the data synchronization routine of step 228, the Visibility Calculator proceeds to step 229, where it exits. Referring to Figure 6, as previously described, the resulting finding of visibility is available to be checked by the log manager in step 183 to determine to write the transaction.

Referring again to figure 7, if step 211 determines that no records were returned by the execution of the SQL statement in step 209, execution proceeds with step 215. Step 215 checks to see whether there are any remaining visibility rules to be assessed. If not, the visibility calculator proceeds to step 228 to synchronize the database, and then to step 229, where it exits. In this case, the default mark of no visibility that was set in step 203 remains set. This value will also be used by the log manager as shown in Figure 6, step 183, to determine not to write the transaction.

Referring again to Figure 7, if rules remain to be assessed, control proceeds to step 217, which selects the next rule to be processed. Control is then given again to step 207 to begin processing the new rule.

5 The preceding text provided a description of the processing or SQL visibility rule; that is, visibility rules of type "R." If step 207 determines that the visibility rule is not of type "R," the visibility rule is of type "O." Type "O" indicates a docking-object visibility rule. In such a case, the docking object being processed
10 will be considered to be visible if it is related to a particular related docking object that is visible. If field RULE_TYPE is not equal to "R," then, execution proceeds to step 221. Step 221 determines the related Docking Object whose visibility must be determined to determine whether the current docking object is visible.
15 The related Docking Object identifier is obtained from field CHECK_DOBJ_ID in table S_DOBJ_VIS_RULE 71. In step 223, the Visibility Calculator determines which row in the related Docking Object must be queried for visibility. In order to determine this, the Visibility Calculator obtains a predetermined SQL statement from the field
20 SQL_STATEMENT and executes it. The SQL statement is a query that select one or more rows of the Docking Object that, for example, correspond to the docking object for which the Visibility Calculator was invoked.

For example, assume that it is desired to indicate that a record
25 for a sales opportunity should be visible if the Node has visibility to any sales quote made for that sales opportunity. This may be accomplished using the following SQL statement:

```
SELECT "ID" FROM
      &Table_Owner.S_DOC_QUOTE
30      WHERE OPTY_ID=:Primary RowId
```

This SQL statement accesses a table S_DOC_QUOTE that contains all sales quotes. The WHERE clause specifies retrieval of all rows where the Opportunity ID of the row is equal to the Row-ID of the opportunity for which visibility is being calculated. The Visibility manager
35 retrieves the specified Row-Ids, thereby identifying the rows of the S_DOC_QUOTE table whose visibility must be checked.

Having determined the a related docking object and the row-ID of that related docking object upon whose visibility the visibility of the current docking object depends, the Visibility Calculator proceeds to
40 step 225. In step 225, the Visibility Calculator recursively invokes itself to determine visibility of the related docking object. The recursively invoked Visibility Calculator operates in the same manner as the Visibility Calculator as called from the Log Manager 9, including the capability to further recursively invoke itself. When
45 the recursive call concludes, it returns a visibility indicator for the related Docking Object, and control proceeds to step 227. In step 227,

the Visibility calculator determines whether the related Docking Object was determined to have been visible. If so, the Visibility Calculator proceeds to step 213 to mark the originally current Docking Object as visible, and then to step 228 to synchronize the database and then to step 229 to exit. If the related Docking Object was not determined to be visible, control proceeds to step 215 to determine whether additional visibility rules remain to be assessed.

The Visibility Calculator, in conjunction with the Log Manager is therefore able to determine what subset of update transaction data is required to be routed to any particular node. This operation serves to reduce the transmission of unneeded data from the Central Computer 1 to the various nodes such as nodes 21-a, 21-b and 21-c that utilize partially replicated databases, and to reduce the system resources such as disk space needed to store, and the CPU time needed to process, what would otherwise be required to maintain a fully replicated database on each remote node.

The operation of the log manager 9 in conjunction with the Visibility Calculator herein described will be apparent from reference to the description and to the drawings. However, as a further aid in the description of these facilities, a pseudocode representation of these facilities is hereto attached as Appendix A. Further features, embodiments and aspects of the invention are apparent from reference to Appendices B through H.

CONCLUSION

Various modifications to these embodiments will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other embodiments without the use of inventive faculty. Thus, the present invention is not intended to be limited to the embodiments shown herein, but is to be accorded the widest scope consistent with the principles and novel features disclosed herein.

All publications and patent applications mentioned in this specification are herein incorporated by reference to the same extent as if each individual publication or patent application was specifically and individually indicated to be incorporated by reference.

The invention now being fully described, it will be apparent to one of ordinary skill in the art that many changes and modifications can be made thereto without departing therefrom.

APPENDIX A

Writing User Transaction Log File for a Given Laptop Node

This program will be called by a server-side process that processes transaction log entries for all Laptop Nodes. For each Laptop Node, the calling process building the UserTrxnLogFileName and calling Program 1.

Input Parameters

- 10 • LaptopNodeId - node_id of the destination laptop
- UserTrxnLogFileName - full path of the file where txns will be written
- MaxBatchTxns - number of txns between commits and updates to the S_DOCK_STATUS table
- 15 • MaxTxns - number of txns to process in this session. Use this parameter to limit processing.

Main Algorithm

```

20  -- Check parameters
   IF (MaxTxns < 1 || MaxBatchTxns < 1) THEN
       Invalid Parameter
   END IF

25  -- Get last LOG_EXTRACT number for the Laptop from
   S_DOCK_STATUS
   last_txn_commit_number = UTLDStatGetLogNum(LaptopNodeId);

   -- Initialize Variables
30  NumTxns = 0;           -- Total number of txns processed
   NumBatchTxns = 0;      -- Total number of txns written in the
   current batch

   -- Read Docking Object and Table definitions into memory
35  structures
   StartDictApi ();

   -- Open the User Log Txn file
   Open User Log Txn file

40  -- Select and process new txns in S_DOCK_TRANSACTION_LOG
   -- where txn_commit_number > last_txn_commit_number
   FOR each new txn LOOP

45      -- Stop processing if reach MaxTxns
       IF NumTxns = MaxTxns THEN
           break;
       END IF;

50      -- Prevent circular txns. Do not send the txn back to the
       originating laptop
       IF txn.OriginNodeId = LaptopNodeId THEN
           Goto next transaction
       END IF;

55

```

```

-- Process all other types of transactions

-- This is the visibility calculator!
-- This routine also processes implicit visibility events
5  -- Later: Data Merge can call this function to check
whether a txn is
-- still visible when merging txns into a laptop or
server database.
CheckVisibility (LaptopNodeId, LogRecordType, TableName,
10 TransRowId);
IF txn is visible THEN
-- Write transactions to UserTxnLog file depending on
the
-- type of LogRecordType.
15 Write the txn to the user log file
++NumBatchTxns
END IF;

-- Finished processing the txn
20 -- Commit (if needed)
IF NumBatchTxns = MaxBatchTxns THEN
-- Assume that separate process comes around and
deletes
-- Txns in S_DOCK_TRANSACTION_LOG that have been
25 processed
-- for all nodes. So, no need to delete the txns from
the log.
Update last LOG_EXTRACT number for Laptop in
S_DOCK_STATUS
30 Commit;
NumBatchTxns = 0
END IF;

++NumTxns
35 End Loop; /* Each transaction in the Txn Log table */

-- Commit
Update last LOG_EXTRACT number for Laptop in S_DOCK_STATUS
Commit;
40
-- Close log file (if needed)
IF UserTxnLogFileP != NULL THEN
Close File;
END IF;
45
StopDictApi ();

```

Check Visibility Routines

```

50 -- Check if a record in the txn log is visible to a
LaptopNodeId
BOOL CheckVisibility (LaptopNodeId, LogRecordType,
TableName, TransRowId)
{
55 -- SQLStatements routed based on the destination list

```



```

    IF LogRecordType in ('SQLStatement') THEN
        IF Laptop Node in destination list THEN
            return TRUE;
        END IF;
5
    -- Shadow and Multi Record LogRecordTypes are routed to
    all nodes
    -- No visibility events with these LogRecordTypes.
    ELSIF LogRecordType in ('ShadowOperation',
10 'MultiRecordDelete',
        'MultiRecordUpdate') THEN
        return TRUE;

    -- Simple Deletes need more processing
15 ELSIF LogRecordType in ('Simple Delete') THEN
        IF (table.visibility in ('Enterprise', 'Limited')) THEN
            return TRUE;
        END IF;

    -- Simple Inserts and Simple Updates need more processing
    -- CheckTxnVisibility() also processes implicit
    visibility events
    ELSIF LogRecordType in ('Simple Insert', 'Simple Update')
    THEN
25     IF (table.visibility = 'Enterprise') THEN
        return TRUE;
        ELSIF table.visibility = 'Limited' THEN
            IF CheckTxnVisibility (LaptopNodeId, Table, RowId)
            THEN
30                 return TRUE;
                END IF;
            END IF;
        END IF;
    }
35
    -- Check if a record in the txn log is visible to a
    LaptopNodeId
    static BOOL CheckTxnVisibility (LaptopNodeId, Table, RowId)
    {
40     BOOL bVisible = FALSE;

        Find the Table in the Dictionary;
        IF Table not found THEN
            Error: Table not defined
45     END IF;

        FOR all docking objects that the table belongs to LOOP
            -- Generate SQL to get PrimaryId values of the Docking
            Object
50         GeneratePrimaryIdSQL (Table, RowId, DockingObject);
            FOR each PrimaryId value retrieved LOOP
                CheckObjectVisibility (LaptopNodeId, PrimaryTable,
                PrimaryRowId)
                IF object is visible THEN
55                 -- Because CheckObjectVisibility() also processes
                implicit

```

```

-- visibility events, we must loop through ALL
docking objects
-- even if we already know that the Txn is visible.
-- Exception: if the table has VIS_event_FLG = 'N'
5  -- then we can return immediately.
IF Table.visibilityEventFLG = 'N' THEN
    return TRUE;
ELSE
    bVisible = TRUE;
10  END IF;
END IF;
END LOOP;
END LOOP;

15  return bVisible;
}

-- Check if an instance of a docking object is visible to
20 the laptop user.
-- Also processes implicit visibility events!
BOOL CheckObjectVisibility (LaptopNodeId,
DockingObjectName, PrimaryRowId)
{
25  FOR each visibility rule for the Docking Object LOOP
    IF RuleType = RuleSQL THEN
        Run the select SQL statement using PrimaryRowId;
        IF any rows returned THEN
            -- row is visible
30            -- Process an implicit Download Object
            DownloadObjectInstance (LaptopNodeId,
PrimaryTableName,
PrimaryRowId);
            return TRUE;
35            END IF;
        ELSIF RuleType = CheckDockingObject THEN
            Run the ParametersQL using PrimaryRowId to get
newPrimaryRowId
            FOR each record retrieved by ParameterSQL LOOP
40                -- RECURSIVE!
                CheckObjectVisibility (LaptopNodeId,
CheckDockingObjectName,
newPrimaryRowId);
                IF rc = TRUE THEN
45                    -- Process an implicit Download Object
                    DownloadObjectInstance (LaptopNodeId,
PrimaryTableName,
PrimaryRowId);
                    return TRUE;
50                    END IF;
                END LOOP;
            END IF;
        END LOOP;
55  -- Object is not visible.

```

```
-- Process an implicit Remove Object
RemoveObjectInstance (LaptopNodeId, PrimaryTableName,
PrimaryRowId);
```

```
5   return FALSE;
}
```

Generate SQL Statement to Get PrimaryId

```
10  -- Generate the SELECT SQL statement to get the PrimaryId
    value of
    -- the docking object for the given MemberTable
    --
    -- SQL statement looks like:
15  --   SELECT tp.<row_id>
    --       FROM <table_owner>.<Table> t1,
    --              <table_owner>.<PKTable> t2,
    --              ... one or more intermediate tables between
    the table
20  --              and the PrimaryTable
    --              <table_owner>.<PKTable> tN
    --              <table_owner>.<PrimaryTable> tp
    --   WHERE t1.ROW_ID = :row_id /* row_id in transaction
    log */
25  --       /* join to PK table t2 */
    --       AND t1.<FKColumn> = t2.<PKColumn>
    --       AND <t1 FKCondition>
    --       /* any number of joins until reach the table
    that joins
30  --       to the PrimaryTable */
    --       /* join from t2 to tN */
    --       AND t2.<FKColumn> = tN.<PKColumn>
    --       AND <t2 FKCondition>
    --       /* join to the PrimaryTable */
35  --       AND tN.<FKColumn> = tp.<PKColumn>
    --       AND <tN FKCondition>
    --
    -- Note that there may be one or more paths from the Member
    Table
40  -- to the Primary Table. We need to generate a SQL select
    statement
    -- for each of the paths and UNION the statements together.
    --
    -- This function assumes that there are no loops in the
45  definition.
    --
    -- These SQL statement do not change for each Table in a
    Docking Object,
    -- so we can calculate them one and store them in memory.
50  --
    struct
    {
        CHAR* selectList;
        CHAR* fromClause;
55  CHAR* whereClause;
```

```

    UINT numTables; /* also the number of joint to reach the
    Primary Table */
    } GenStmt;

5  GeneratePrimaryIdSQL (Table, DockingObject)
    {
        /* there may be more than one SQL statement, so we have
        a dynamic
10     array of SQL statements. Each element in the array is
        a path
        from the Table to the Primary Table*/
        DynArrId GenStmtArr;
        GenStmt newGenStmt;

15     CHAR* sqlStmt;

        DynArrCreate (GenStmtArr);

20     -- Create the first element and initialize
        newGenStmt = malloc();
        newGenStmt.numTables = 1;
        newGenStmt.selectList = "SELECT row_id";
        newGenStmt.fromClause = "FROM <Table> t1";
25     newGenStmt.whereClause = "WHERE t1.ROW_ID = :row_id";
        DynArrAppend (GenStmtArr, &newGenStmt);

        /* Recursively follow FKs to the PrimaryTable */
        Build the select, from and where clause simultaneously
30     */
        AddPKTable (Table, DockingObject, GenStmtArr, 0);

        -- Union all the paths together
        numStmts = DynArrSize (GenStmtArr);
35     FOR all elements in the array LOOP
        tmpSqlStmt = GenStmtArr[j].selectList ||
GenStmtArr[j].fromClause || GenStmtArr[j].whereClause;
        sqlStmt = sqlStmt || 'UNION' || tmpSqlStmt;
        END LOOP;

40     DynArrDestroy (GenStmtArr);

        IF sqlStmt = NULL THEN
            Error: no path from Table to Primary Table.
45     END IF;
    }

    -- Recursively follow all FKs to the Primary Table
50 AddPKTable (Table, DockingObject, GenStmt, InputStmtNum)
    {
        UINT numFKS = 0;
        UINT StmtNum;
        GenStmt newGenStmt;

55     FOR all FKs for the table LOOP

```

```

        IF PKTable is a Member Table of the Docking Object THEN
            -- If there's more than one FK, then there is more
            than one path
            -- out of the current table.
5           -- Copy the SQL stmt to a new DynArrElmt to create a
            new path
                IF numFKs > 0 THEN
                    -- Create a new element and copy from
                    GenStmt[InputStmtNum]
10                    newGenStmt = malloc();
                        new Gen S t m t . n u m T a b l e s =
                    GenStmt[InputStmtNum].numTables;
                        new Gen S t m t . s e l e c t L i s t =
                    GenStmt[InputStmtNum].selectList;
15                    new Gen S t m t . f r o m C l a u s e =
                    GenStmt[InputStmtNum].fromClause;
                        new Gen S t m t . w h e r e C l a u s e =
                    GenStmt[InputStmtNum].whereClause;
                    DynArrAppend (GenStmtArr, &newGenStmt);
20                    StmtNum = DynArrSize (GenStmtArr);

                    -- Put a check here for infinite loops
                    IF StmtNum == 20 THEN
                        Error: Probably got an Infinite loop?
25                        END IF;
                    ELSE
                        StmtNum = InputStmtNum;
                        END IF;

30                -- Append the new PKTable to the fromClause and
                    whereClause
                    GenStmt[StmtNum].fromClause =
                        GenStmt[StmtNum].fromClause || ",\n <Table>
                    t<numTables + 1>";
35                    GenStmt[StmtNum].whereClause =
                        GenStmt[StmtNum].whereClause ||
                            "AND t<numTables>.<FKColumn> = t<numTables +
                    1>.<PKColumn>" ||
                            "AND <FKCondition for Table if any>";
40                    ++GenStmt.numTables;

                    -- PKTable is the Primary Table then Done.
                    IF PKTable = PrimaryTable THEN
                        RETURN;
45                    ELSE
                        AddPKTable (PKTable, DockingObject, GenStmt,
                    StmtNum);
                        END IF;

50                -- Only count FKs to other member tables in the same
                    Docking Object
                    ++numFKs;

                    END IF;
55                END LOOP;

```

```

    RETURN;
}

```

Process Visibility Events

```

5
-- Download an Object Instance to a Laptop
-- This function also downloads all Related Docking Object
instances.
BOOL DownloadObjectInstance (LaptopNodeId,  ObjectName,
10 PrimaryRowId)
{
    -- Check if the object instance is already downloaded to
the laptop
    Find the object instance in the S_DOBJ_INST table
15    IF exists on laptop THEN
        return TRUE;
    END IF;

    -- Register object instance in S_DOBJ_INST table
20
    -- Write Download Object records to the Txn Log
    FOR each member table of the docking object LOOP
        Generate SQL select statement to download records
        Write each retrieved record to the User Txn Log file
25    END LOOP;

    -- Download records for Parent Object instances
    FOR each RelatedDockingObject LOOP
        Run      ParameterSQL      to      get      newPrimaryId      of
30 RelatedDockingObjects
        FOR each newPrimaryId retrieved LOOP
            -- Check if the instance of the object is visible to
the laptop user
            CheckObjectVisibility (LaptopNodeId,  ObjectName,
35 PrimaryRowId)
            IF visible THEN
                DownloadObjectInstance (LaptopNodeId,
                                         RelatedDockingObject,
newPrimaryRowId);
40                END IF;
            END LOOP;
        END LOOP;

        return TRUE;
45 }

-- Remove an Object Instance to a Laptop
-- This function also removes all Related Docking Object
50 instances.
BOOL RemoveObjectInstance (LaptopNodeId,  ObjectName,
PrimaryRowId)
{
    -- Check if the object instance is already downloaded to
55 the laptop

```

```
Find the object instance in the S_DOBJ_INST table
IF does not exist on laptop THEN
    return TRUE;
END IF;

5  -- Delete the object instance from S_DOBJ_INST table

    -- Write Remove Object records to the Txn Log
    FOR each member table of the docking object LOOP
10  Generate SQL select statement to get records to delete
        Write each retrieved record to the User Txn Log file
    END LOOP;

    -- Remove for Parent Object instances
15  FOR each RelatedDockingObject LOOP
        Run ParameterSQL to get newPrimaryId of
        RelatedDockingObjects
        FOR each newPrimaryId retrieved LOOP
            -- Check if the instance of the object is visible to
20  the laptop user
                CheckObjectVisibility (LaptopNodeId, ObjectName,
                PrimaryRowId)
                IF not visible THEN
                    RemoveObjectInstance (LaptopNodeId,
25  RelatedDockingObject,
                    newPrimaryRowId);
                END IF;
            END LOOP;
        END LOOP;
30  return TRUE;
    }
```

We claim:

1. A method of collecting, storing, and retrieving data in a database management system having an enterprise server, at least one docking client, and at least one workgroup user client, said docking client interposed between said enterprise server and said workgroup user client, said method comprising:
 - (a) creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on said workgroup user client, and creating a transaction file corresponding thereto in an outbox of said workgroup user client;
 - (b) copying said transaction file to an inbox in the docking client identified to the workgroup user client and updating said transaction file into an agency database resident on said docking client, said agency database including a transaction log; and
 - (c) reading said docking client transaction log, skipping those transactions which originate at the enterprise server, creating data files corresponding to the entries therein, copying data files corresponding to transactions originating at the workgroup user client to an inbox on the enterprise server corresponding to the docking server, and updating the transactions into an enterprise database on the enterprise server.
2. The method of claim 1 comprising connecting said workgroup user client to said docking client and thereafter copying said transaction file to said docking client.
3. The method of claim 1 comprising updating transactions

into the enterprise database with a field corresponding to the docking client.

4. The method of claim 1 wherein said enterprise server includes an enterprise database, a docking manager, a update processor, and a log manager.

5. The method of claim 4 wherein said enterprise server further includes an update manager.

6. The method of claim 1 wherein said workgroup user client is a mobile client.

7. The method of claim 1 wherein said workgroup user client includes a partially replicated database, an update manager, a docking manager, and a update manager.

8. The method of claim 7 wherein said work group user client further includes an update log.

9. The method of claim 7 comprising updating said partially replicated database through said update manager.

10. The method of claim 9 wherein said work group user further includes an update log, said method further comprising updating said update log.

11. The method of claim 1 wherein said work group user includes a node docking manager.

12. The method of claim 11 comprising activating the node docking manager of the work group user and transmitting contents of the work group user update log to the docking manager of the enterprise server.

13. The method of claim 12 comprising applying the update log contents to the enterprise database in the enterprise server.
- 5 14. The method of claim 13 comprising sending a command from the update manager of the work group user to the update manager of the enterprise server to alter data therein, updating the enterprise database, creating a log record describing the update, and writing the log
10 record to the update log.
- 15 15. The method of claim 14 wherein the update log record contains the node identifier of the workgroup user making the update, identification of the enterprise database table updated, and the type of update made.
16. The method of claim 1 comprising preparing partial output logs of enterprise database contents according to visibility rules in the enterprise server.
- 20 17. The method of claim 1 comprising transferring data that meets query criteria from one database to another database.
- 25 18. The method of claim 1 comprising synchronizing the work group user database and the enterprise database.
- 30 19. The method of claim 18 comprising synchronizing the work group user database with the enterprise database using a docking object, said docking object comprising at least one member table, visibility rules, and visibility events.
- 35 20. The method of claim 19 wherein the member table comprises a table of a relational database, and synchronization comprises inserting an entry into or deleting an entry from a member table, said member

table being associated with the docking object.

21. The method of claim 19 wherein the member table contains visibility rules.

5

22. The method of claim 19 wherein the member table contains an identification of a row of the member table being updated.

10

23. The method of claim 19 wherein the docking object contains a statement to determine if any data meeting the criteria contained in the statement exists in the docking object.

15

24. The method of claim 19 wherein the docking object contains related docking objects and the method comprises propagating or deleting a related docking object when the docking object is propagated or deleted.

20

25. The method of claim 1 comprising updating data from one database to another database.

25

26. The method of claim 25 comprising transferring a update log file from one database to the other database.

30

27. The method of claim 25 comprising comparing a value in the database to a value that the update log reports as the previous value in the database to detect a collision.

35

28. The method of claim 25 comprising updating the database when the value in the database matches the value that the update log reports as the previous value in the database.

29. The method of claim 25 comprising generating an error signal when the value in the database does not match the value that the update log reports as the previous value in the database.

5

30. A method of collecting, storing, and retrieving data in a database management system having an enterprise server and at least one workgroup user client, said method comprising: creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on said workgroup user client, and creating a transaction file corresponding thereto in an outbox of said workgroup user client; and reading said workgroup user client transaction log, skipping those transactions which originate at the enterprise server, creating data files corresponding to the entries therein, copying data files corresponding to transactions originating at the workgroup user client to an inbox on the enterprise server, and updating the transactions into an enterprise database on the enterprise server.

10

15

20

31. The method of claim 30 comprising updating workgroup user transactions having a field corresponding thereto into the enterprise database.

25

32. The method of claim 30 wherein said enterprise server includes an enterprise database, a docking manager, a update processor, and a log manager.

30

33. The method of claim 32 wherein said enterprise server further includes an update manager.

35

34. The method of claim 30 wherein said workgroup user client is a mobile client.

35. The method of claim 30 wherein said workgroup user client includes a partially replicated database, an update manager, a docking manager, and an update manager.
- 5 36. The method of claim 35 wherein said work group user client further includes an update log.
- 10 37. The method of claim 36 comprising updating said partially replicated database through said update manager.
- 15 38. The method of claim 37 wherein said work group user further includes an update log, said method further comprising updating said update log.
- 20 39. The method of claim 30 comprising preparing partial output logs of enterprise database contents according to visibility rules in the enterprise server.
40. The method of claim 30 comprising transferring data that meets query criteria from one database to another database.
- 25 41. The method of claim 30 comprising synchronizing the work group user database and the enterprise database.
- 30 42. The method of claim 41 comprising synchronizing the work group user database with the enterprise database using a docking object, said docking object comprising at least one member table, visibility rules, and visibility events.
- 35 43. The method of claim 42 wherein the member table comprises a table of a relational database, and synchronization comprises inserting an entry into or deleting an entry from a member table, said member

table being associated with the docking object.

44. The method of claim 43 wherein the member table contains visibility rules.

5

45. The method of claim 42 wherein the member table contains an identification of a row of the member table being updated.

10 46. The method of claim 42 wherein the docking object contains a statement to determine if any data meeting the criteria contained in the statement exists in the docking object.

15 47. The method of claim 42 wherein the docking object contains related docking objects and the method comprises propagating or deleting a related docking object when the docking object is propagated or deleted.

20

48. The method of claim 30 comprising updating data from one database to another database.

25 49. The method of claim 48 comprising transferring a update log file from one database to the other database.

30 50. The method of claim 48 comprising comparing a value in the database to a value that the update log reports as the previous value in the database to detect a collision.

35 51. The method of claim 50 comprising updating the database when the value in the database matches the value that the update log reports as the previous value in the database.

52. The method of claim 50 comprising generating an error signal when the value in the database does not match the value that the update log reports as the previous value in the database.

5

53. A method of merging data from one database into another database in a database management system having an enterprise server and at least one workgroup user client, said method comprising: creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on said workgroup user client, and creating a transaction file corresponding thereto in an outbox of said workgroup user client; and reading said workgroup user client transaction log, skipping those transactions which originate at the enterprise server, creating data files corresponding to the entries therein, copying data files corresponding to transactions originating at the workgroup user client to an inbox on the enterprise server, and transferring a update log file from one database to the other database to merge the transactions.

10

15

20

25

54. The method of claim 53 comprising comparing a value in the database to a value that the merge log reports as the previous value in the database to detect a collision.

30

55. The method of claim 53 comprising updating the database when the value in the database matches the value that the merge log reports as the previous value in the database.

35

56. The method of claim 53 comprising generating an error signal when the value in the database does not match the value that the merge log reports as the previous

value in the database.

57. An article of manufacture comprising a computer usable medium having computer readable program code means embodied therein for causing collecting, storing, and retrieving data in a database management system having an enterprise server, at least one docking client, and at least one workgroup user client, said docking client interposed between said enterprise server and said workgroup user client, the computer readable program means in said article of manufacture comprising:

(a) computer readable program code means for causing a computer to effect creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on said workgroup user client, and creating a transaction file corresponding thereto in an outbox of said workgroup user client;

(b) computer readable program code means for causing a computer to effect copying said transaction file to an inbox in the docking client identified to the workgroup user client and updating said transaction file into an agency database resident on said docking client, said agency database including a transaction log; and

(c) computer readable program code means for causing a computer to effect reading said docking client transaction log, skipping those transactions which originate at the enterprise server, creating data files corresponding to the entries therein, copying data files corresponding to transactions originating at the workgroup user client to an inbox on the enterprise server corresponding to the docking server, and updating the transactions into an

enterprise database on the enterprise server.

58. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform method steps for collecting, storing, and retrieving data in a database management system having an enterprise server, at least one docking client, and at least one workgroup user client, said docking client interposed between said enterprise server and said workgroup user client, said method steps comprising:

(a) creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on said workgroup user client, and creating a transaction file corresponding thereto in an outbox of said workgroup user client;

(b) copying said transaction file to an inbox in the docking client identified to the workgroup user client and updating said transaction file into an agency database resident on said docking client, said agency database including a transaction log; and

(c) reading said docking client transaction log, skipping those transactions which originate at the enterprise server, creating data files corresponding to the entries therein, copying data files corresponding to transactions originating at the workgroup user client to an inbox on the enterprise server corresponding to the docking server, and updating the transactions into an enterprise database on the enterprise server.

59. An article of manufacture comprising a computer usable medium having computer readable program code means embodied therein for causing collection, storage, and

retrieval of data in a database management system having an enterprise server and at least one workgroup user client, the computer readable program means in said article of manufacture comprising:

- 5 (a) Computer readable program code means for causing a computer to effect creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on said workgroup user client, and creating a transaction file
10 corresponding thereto in an outbox of said workgroup user client; and
- (b) Computer readable program code means for causing the computer to effect reading said workgroup
15 user client transaction log, skipping those transactions which originate at the enterprise server, creating data files corresponding to the entries therein, copying data files corresponding to transactions originating at the workgroup user
20 client to an inbox on the enterprise server, and updating the transactions into an enterprise database on the enterprise server.

60. A program storage device readable by a machine, tangibly embodying a program of instructions
25 executable by a machine to perform method steps for collecting, storing, and retrieving data in a database management system having an enterprise server and at least one workgroup user client, said method steps comprising:
- 30 (a) creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on said workgroup user client, and creating a transaction file corresponding thereto in an
35 outbox of said workgroup user client; and
- (b) reading said workgroup user client transaction log, skipping those transactions which originate

at the enterprise server, creating data files corresponding to the entries therein, copying data files corresponding to transactions originating at the workgroup user client to an inbox on the enterprise server, and updating the transactions into an enterprise database on the enterprise server.

61. An article of manufacture comprising a computer usable medium having computer readable program code means embodied therein for causing merging data from one database into another database in a database management system having an enterprise server and at least one workgroup user client, the computer readable program means in said article of manufacture comprising:

(a) computer readable program code means for causing a computer to effect creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on said workgroup user client, and creating a transaction file corresponding thereto in an outbox of said workgroup user client; and

(b) computer readable program code means for causing the computer to effect reading said workgroup user client transaction log, skipping those transactions which originate at the enterprise server, creating data files corresponding to the entries therein, copying data files corresponding to transactions originating at the workgroup user client to an inbox on the enterprise server, and transferring a update log file from one database to the other database to merge the transactions.

62. A program storage device readable by a machine,

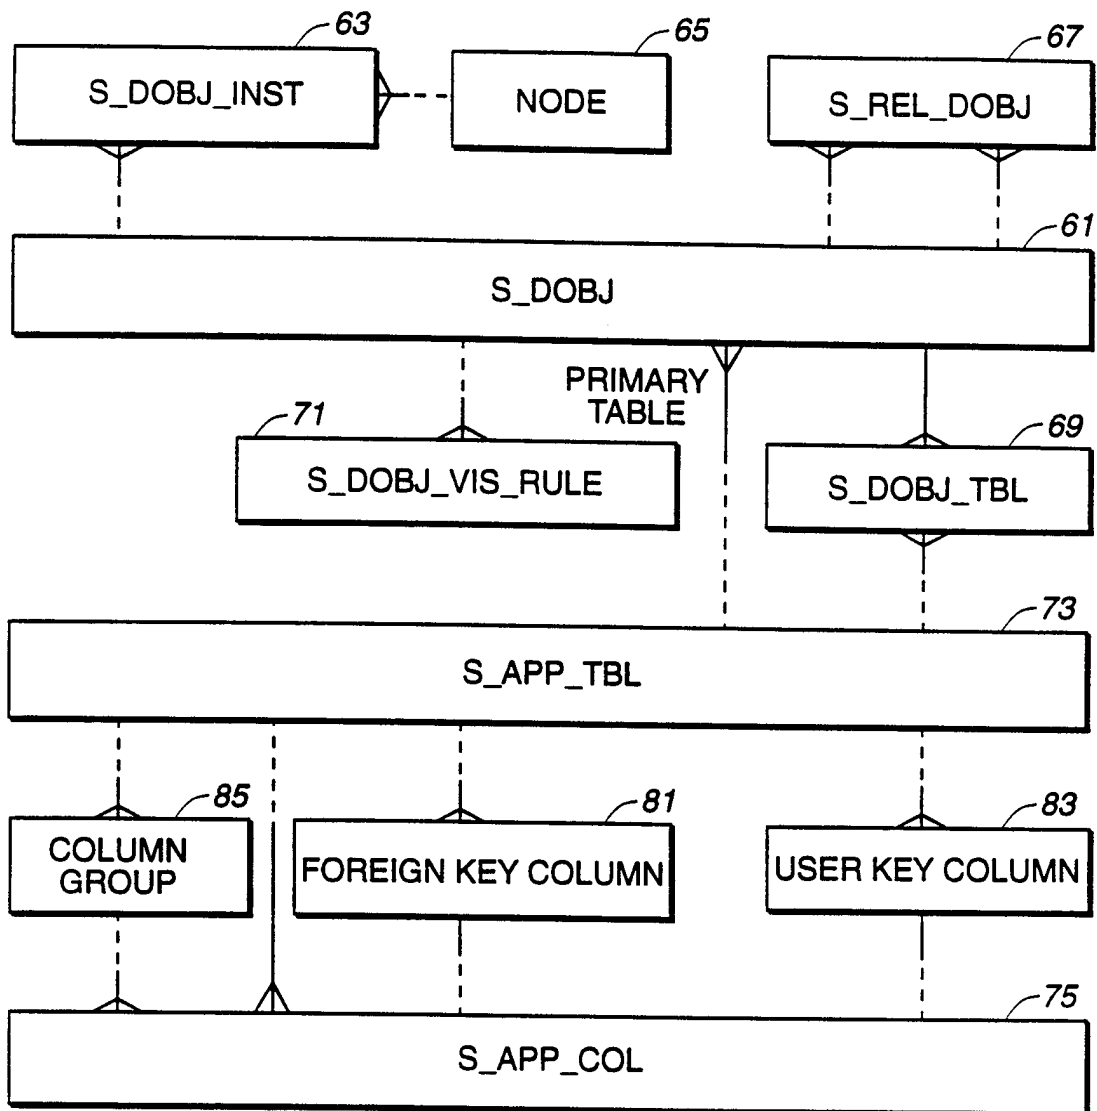
tangibly embodying a program of instructions executable by a machine to perform method steps for merging data from one database into another database in a database management system, said database management system having an enterprise server and at least one workgroup user client, said method steps comprising:

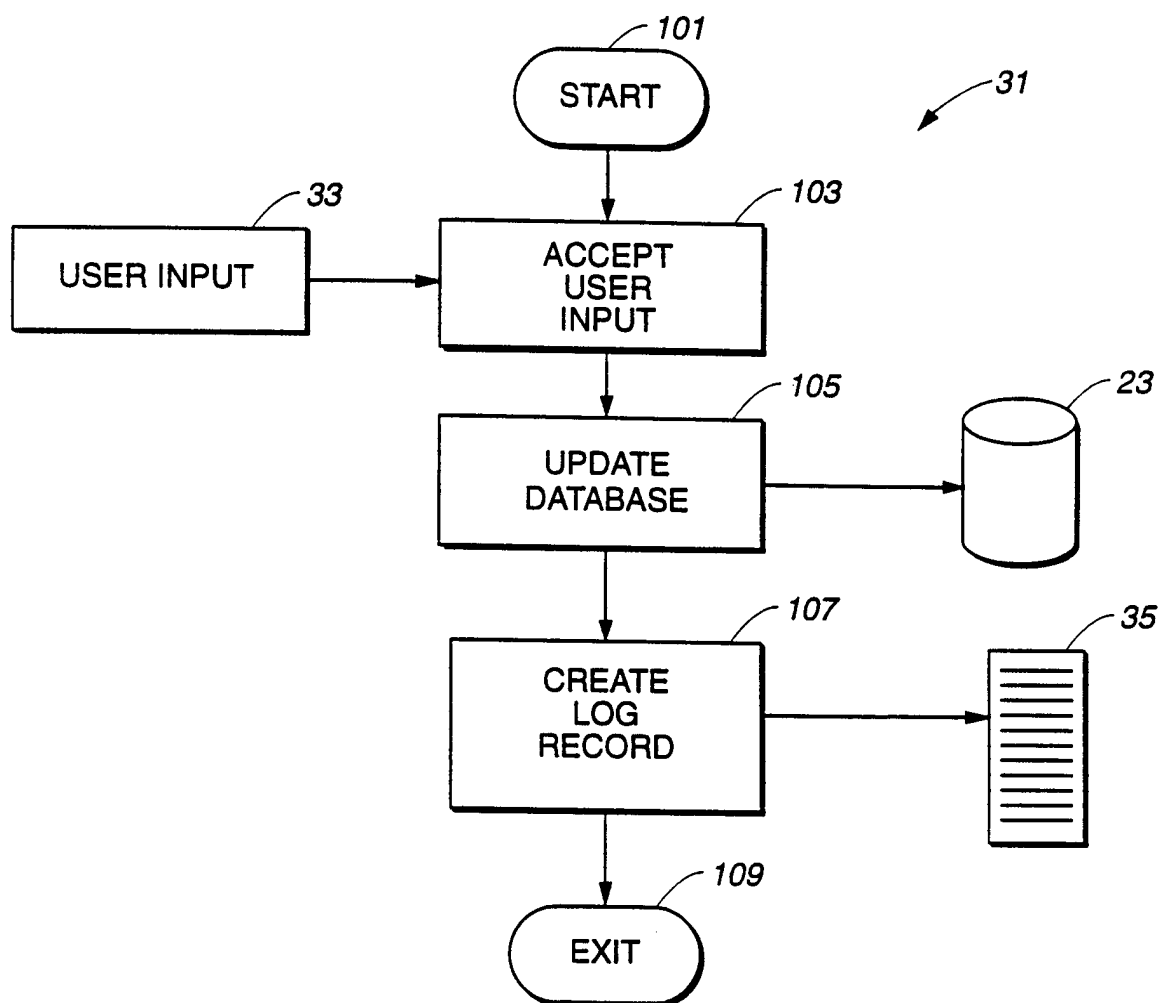
- (a) creating a transaction in a local database resident on said workgroup user client, entering the transaction into a transaction log resident on said workgroup user client, and creating a transaction file corresponding thereto in an outbox of said workgroup user client; and
- (b) reading said workgroup user client transaction log, skipping those transactions which originate at the enterprise server, creating data files corresponding to the entries therein, copying data files corresponding to transactions originating at the workgroup user client to an inbox on the enterprise server, and transferring an update log file from one database to the other database to merge the transactions.

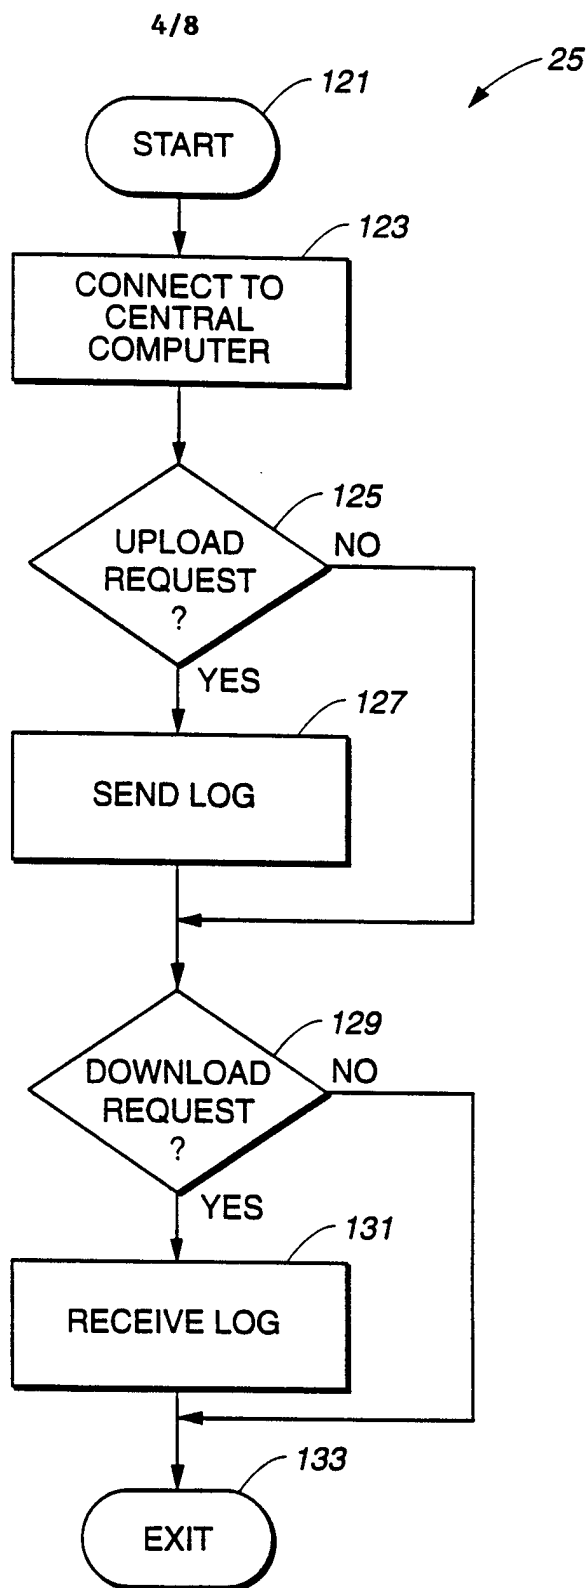
SUBSTITUTE SHEET (RULE 26)

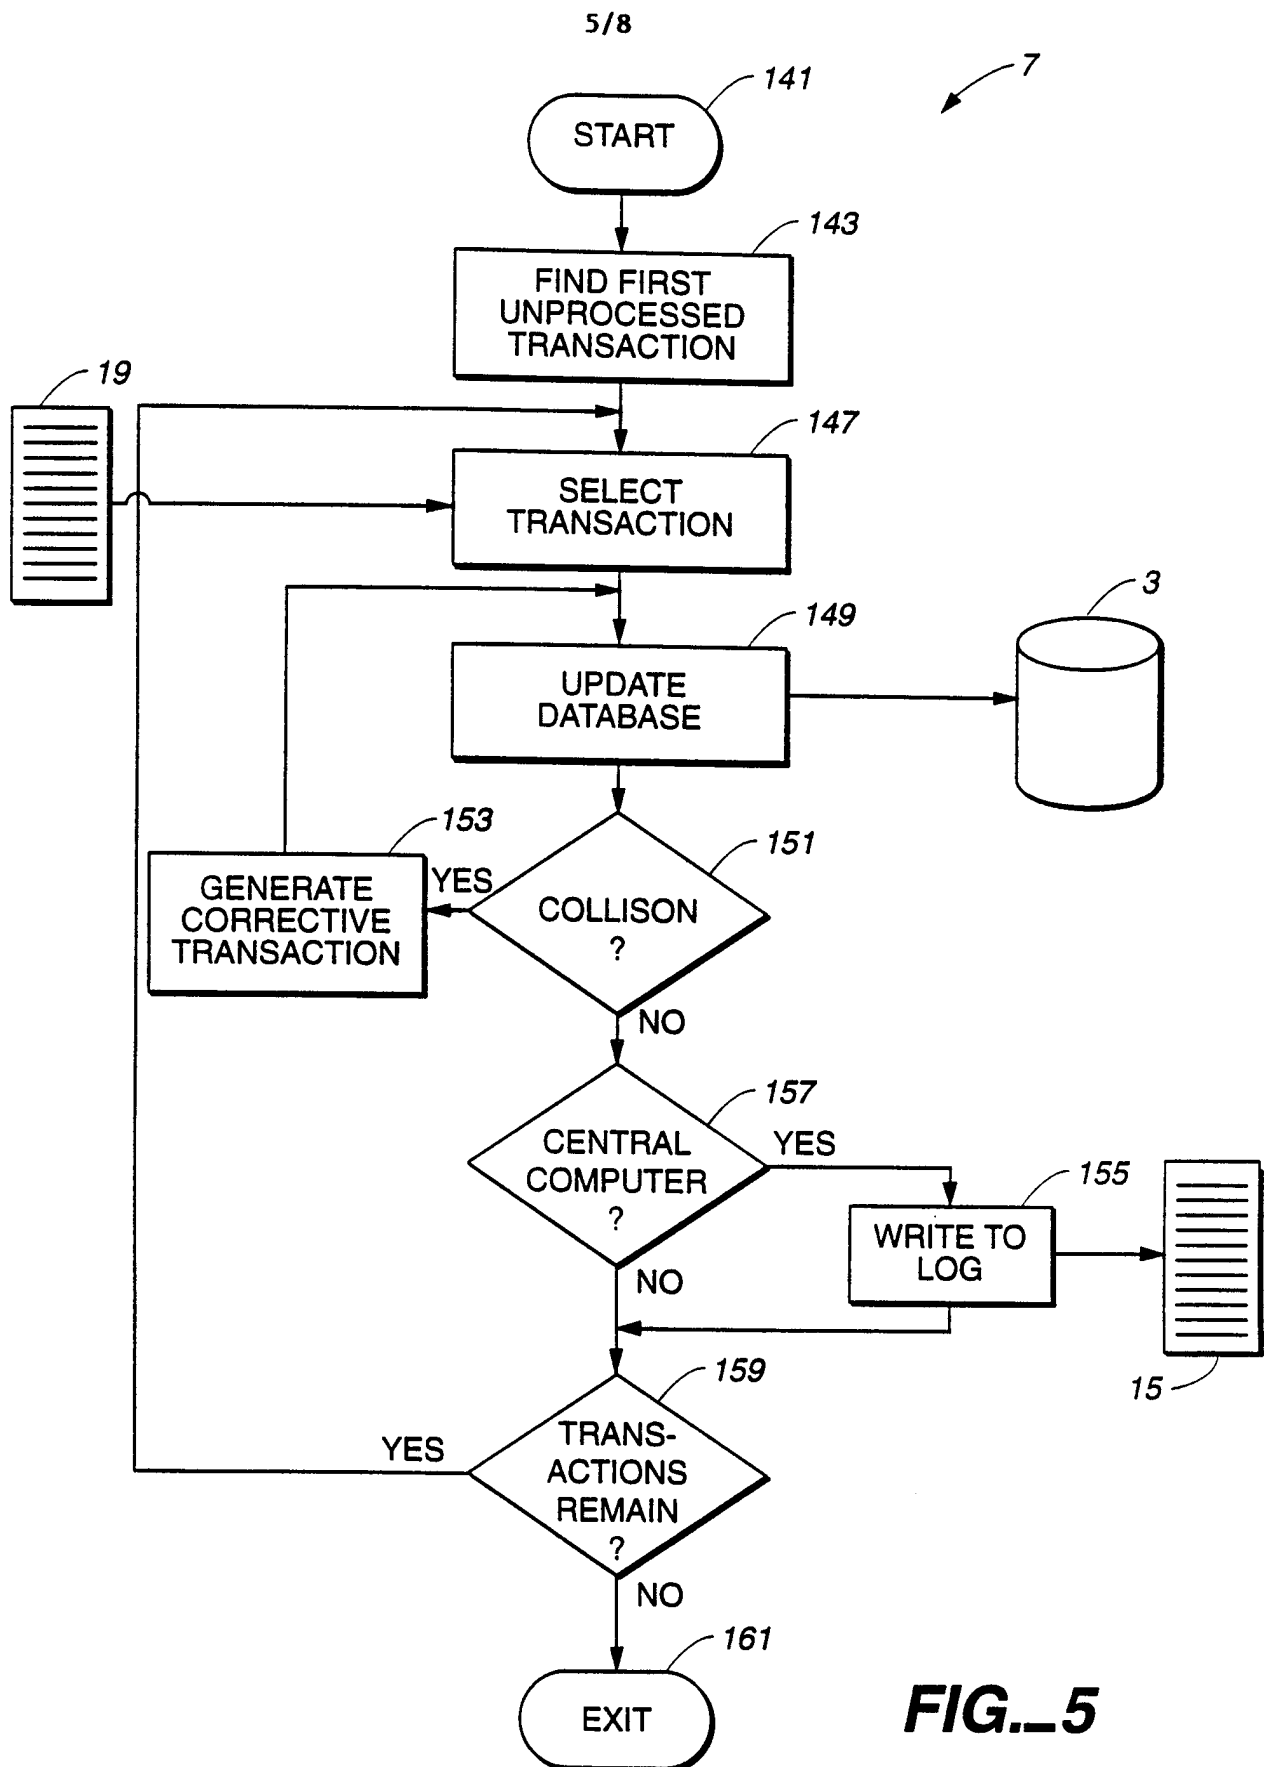


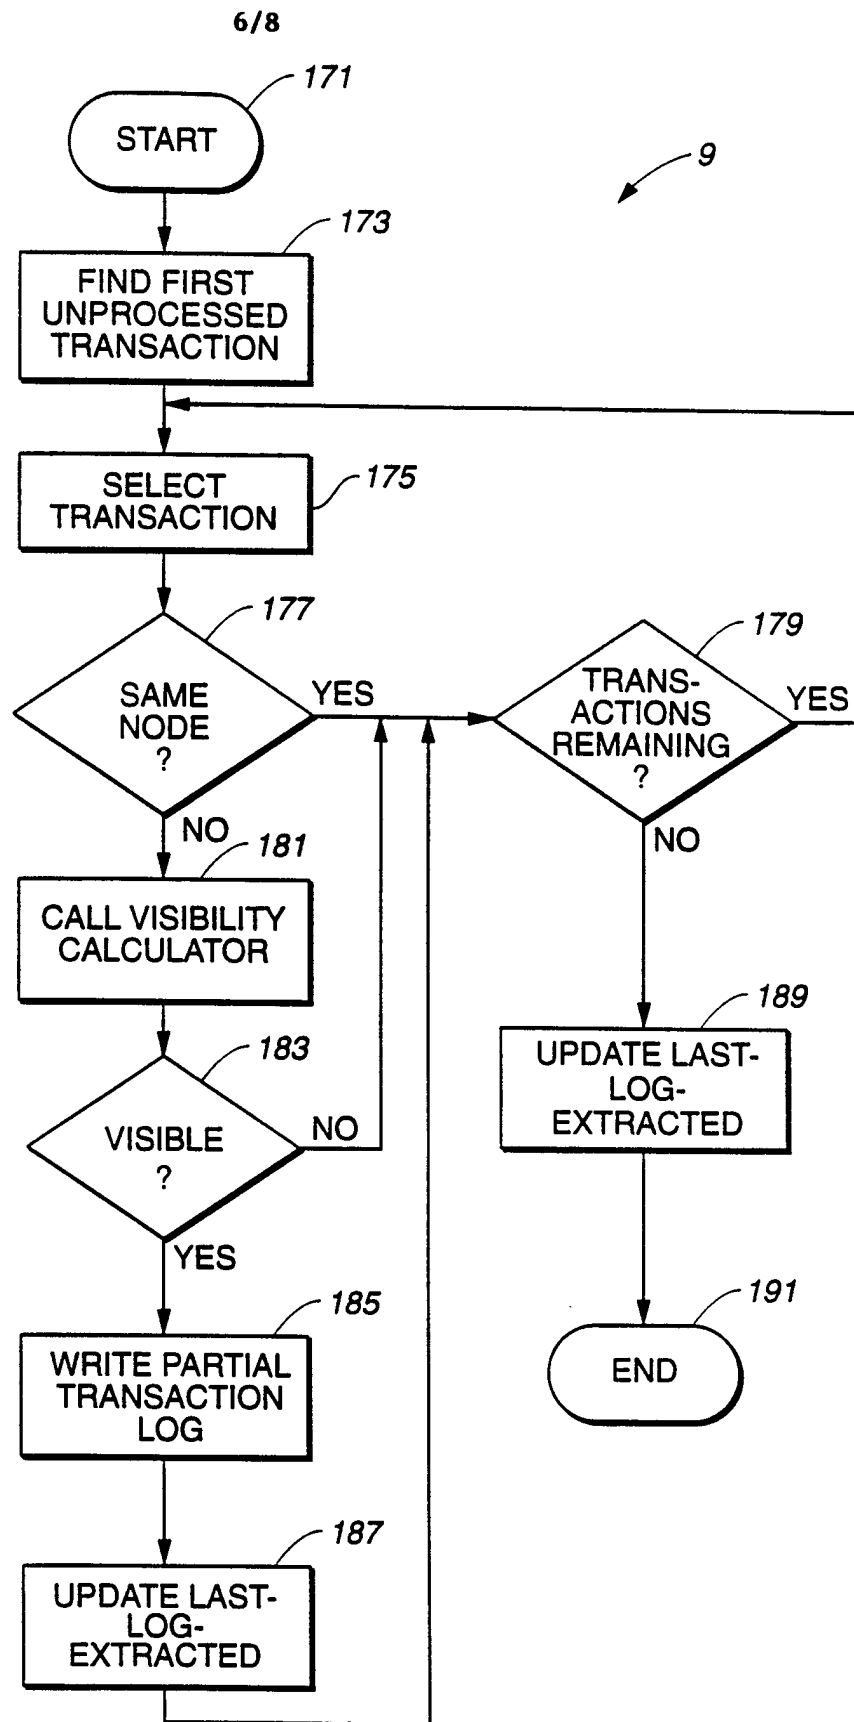
2/8

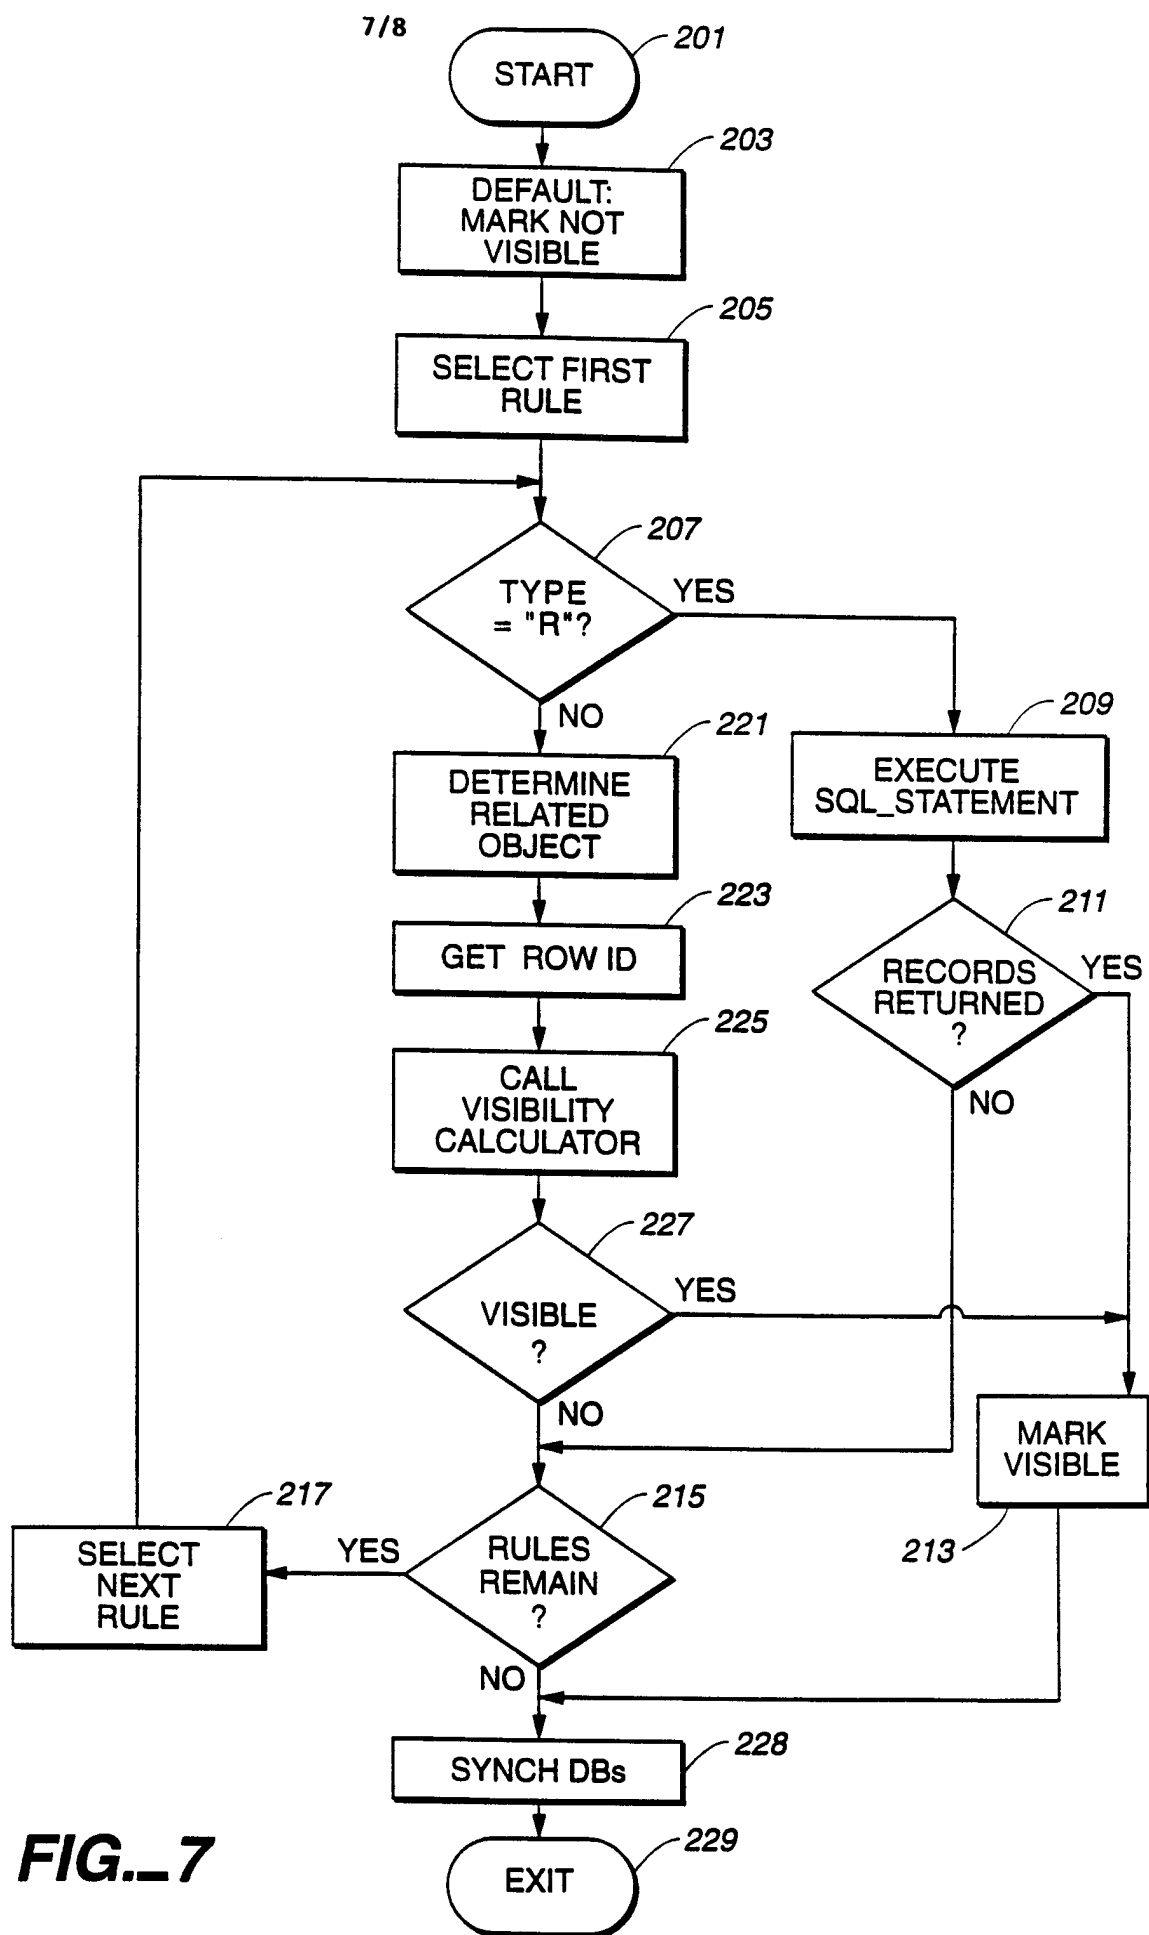
**FIG. 2**

**FIG. 3**

**FIG. 4**



**FIG. 6**

**FIG. 7**

8/8

