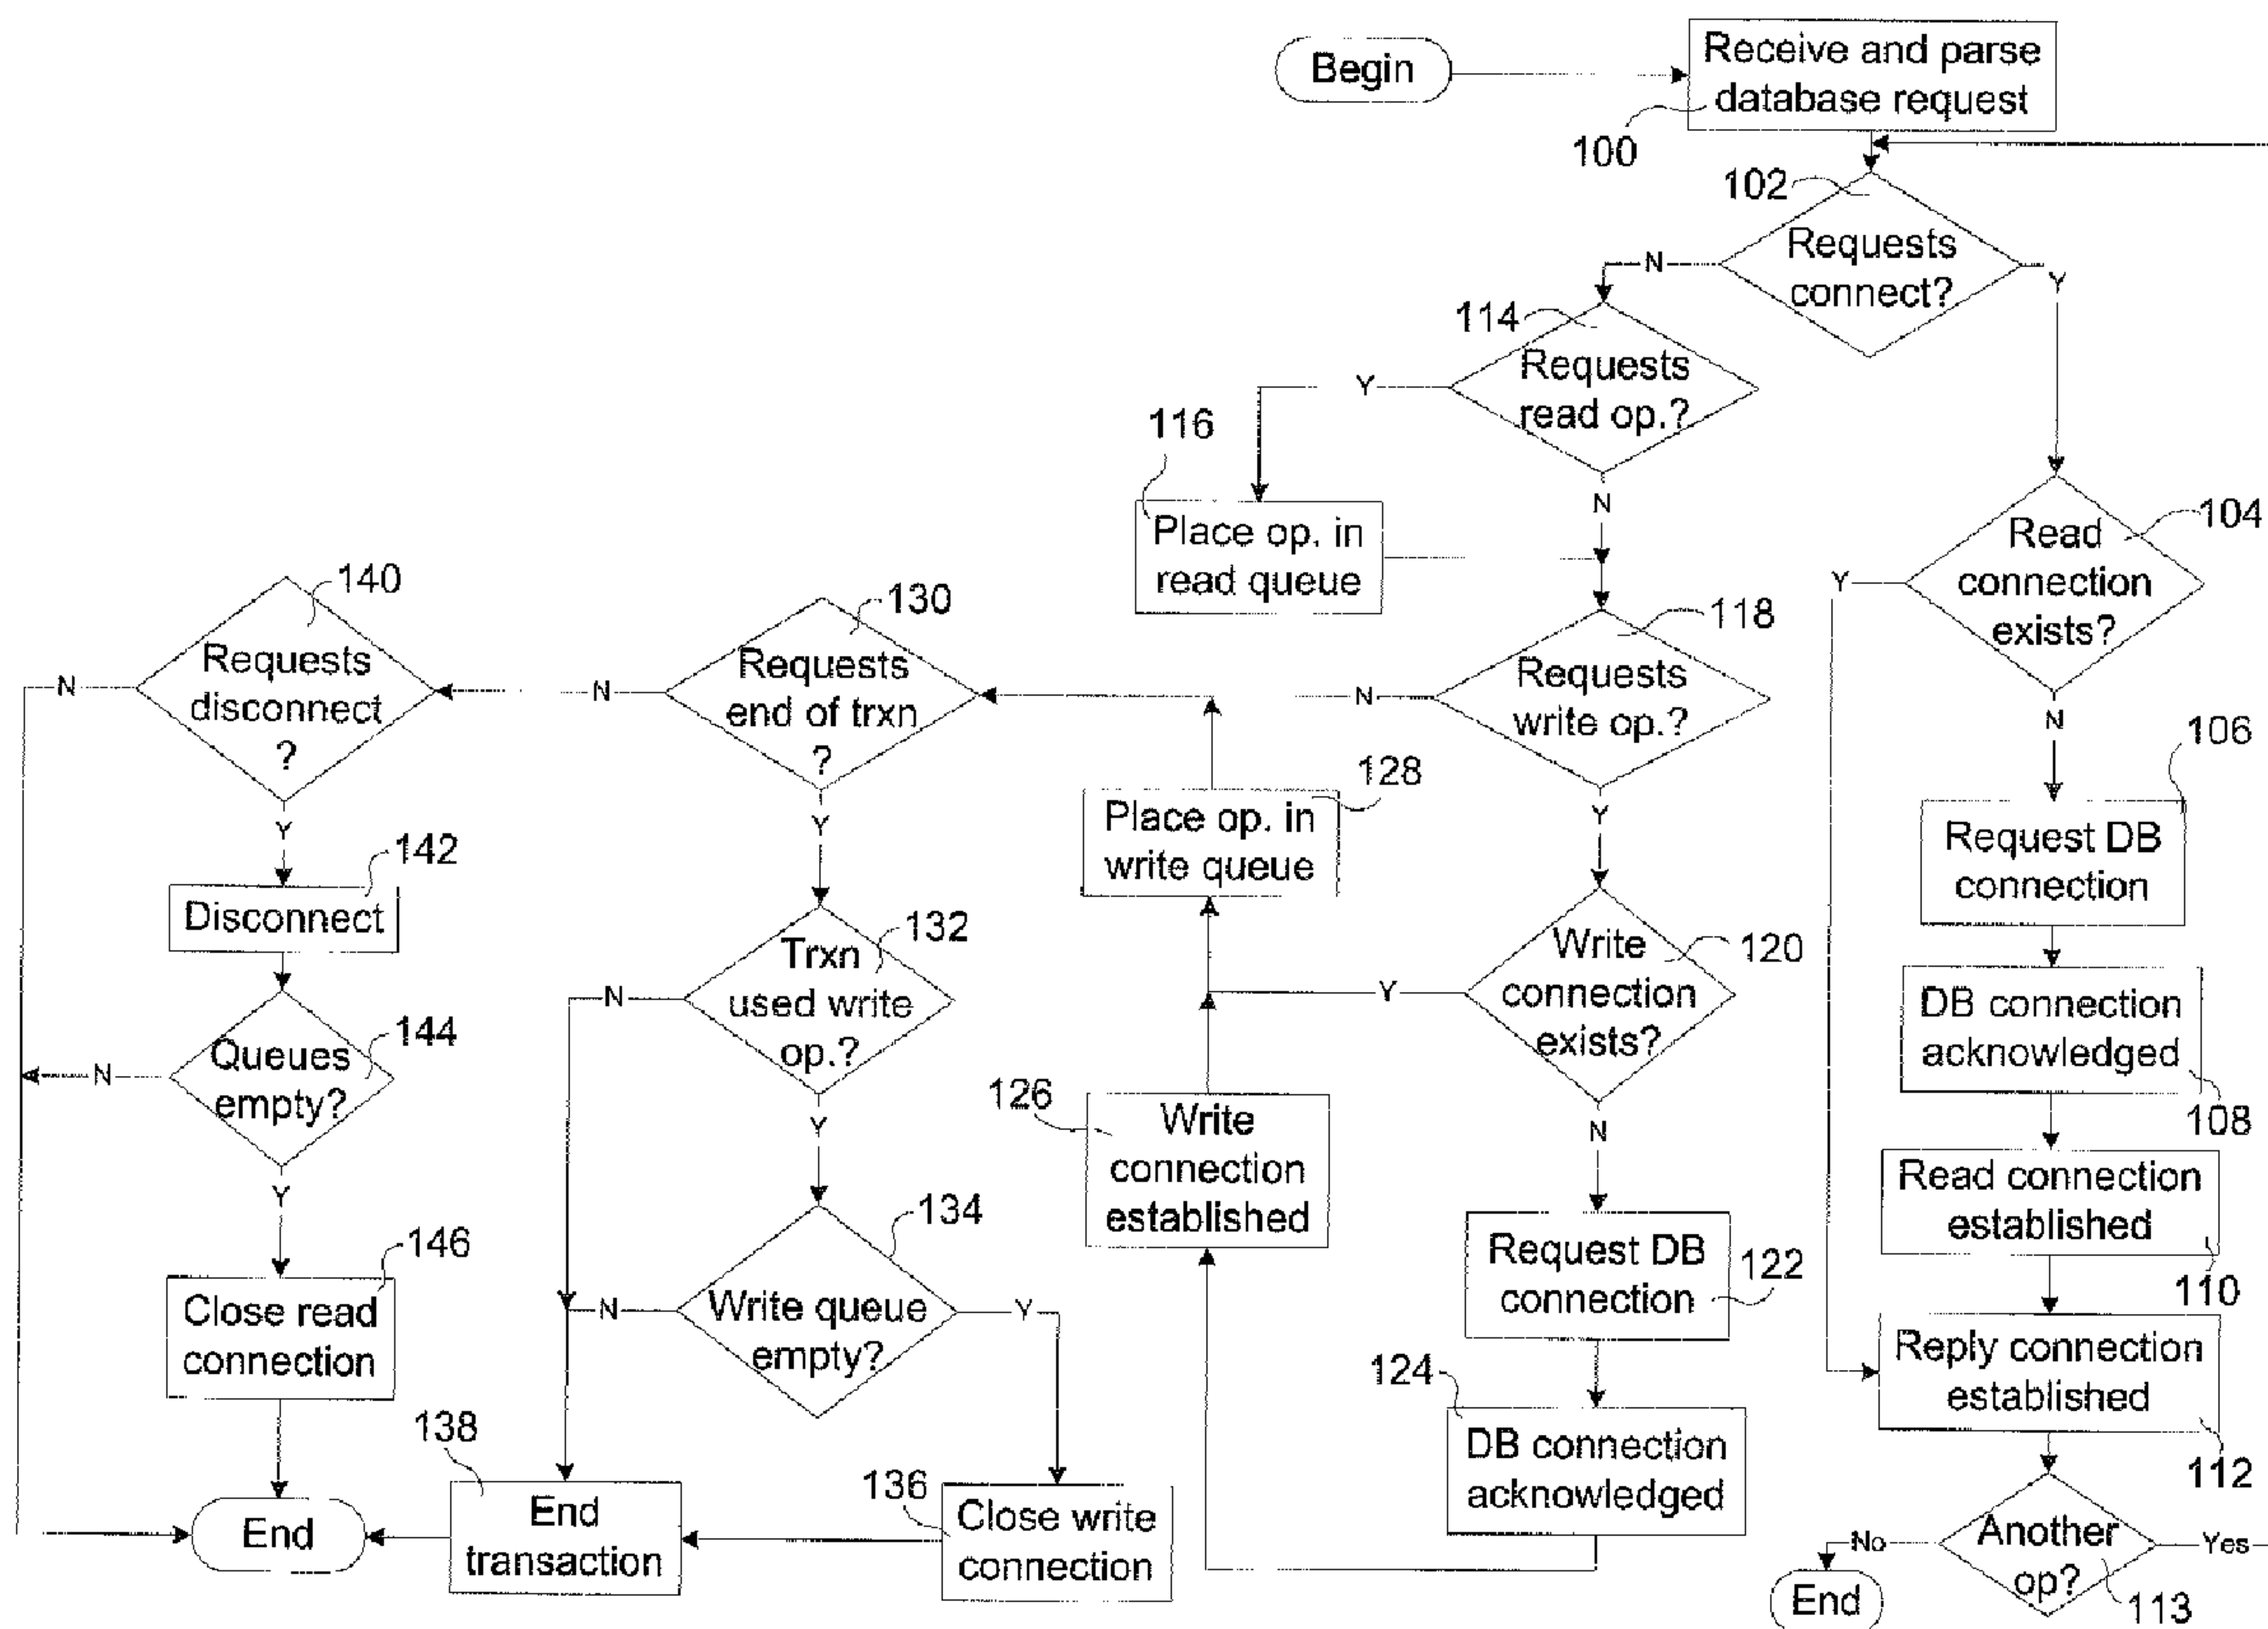




(22) Date de dépôt/Filing Date: 2003/04/08
(41) Mise à la disp. pub./Open to Public Insp.: 2004/10/08
(45) Date de délivrance/Issue Date: 2008/06/03

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01),
G06F 13/14 (2006.01), *G06F 7/00* (2006.01)
(72) Inventeurs/Inventors:
LIMOGES, JOSEPH SERGE, CA;
EVANS, DOMINIQUE J., CA;
GODINEZ, MARIO S., CA
(73) Propriétaire/Owner:
IBM CANADA LIMITED - IBM CANADA LIMITEE, CA
(74) Agent: CHAN, BILL W.K.

(54) Titre : METHODE ET SYSTEME D'ACCES A DES RESSOURCES
(54) Title: METHOD AND SYSTEM FOR RESOURCE ACCESS



(57) Abrégé/Abstract:

A method and database management system for connecting to a database involves reusing connections for operations performed in sequence, rather than creating independent connections for each requested client connection. An application server (or layer)

(57) **Abrégé(suite)/Abstract(continued):**

receives the requests, parses them to identify modes of operation. Separate connections are created to support each mode of operation so that the database receives the operations and requests in a format that is sorted by type, at a rate that can be adjusted to match capacity of the database.

METHOD AND SYSTEM FOR RESOURCE ACCESS**ABSTRACT**

A method and database management system for connecting to a database
5 involves reusing connections for operations performed in sequence, rather than creating
independent connections for each requested client connection. An application server (or
layer) receives the requests, parses them to identify modes of operation. Separate
connections are created to support each mode of operation so that the database receives
the operations and requests in a format that is sorted by type, at a rate that can be
10 adjusted to match capacity of the database.

METHOD AND SYSTEM FOR RESOURCE ACCESS

TECHNICAL FIELD

[0001] The invention generally relates to transaction distribution systems, and in particular relates to a database transaction distribution system.

5 BACKGROUND

[0002] Databases are used to store, and to provide access to, a large proportion of digital information. As current databases have capacity to store large volumes of data, and to provide many thousands of users with access to the digital information, they have been widely deployed by large and mid-sized businesses. The integrity of the digital information
10 in a database is paramount.

[0003] Systems for providing access to the digital information contained in a database are known in the art. Many of these database access systems include an application server associated with the database through which a set of clients may access the database. The application server may perform numerous operations including controlling access to the
15 database to prevent clients from performing unauthorized actions, converting client data into a format and protocol used for accessing the database, and sequencing the database requests.

[0004] Application servers relay client requests to the database via database connections without any regard for the operations they may contain. The application
20 servers receive the database requests from the client side of the application server, and provide for the processing of the database requests. Processing the database requests involves one or more database connections. The setup, and tear-down of connections involves connection control processing and signaling for each of the database connections. As there are many of these database connections concurrently maintained at the database,
25 significant processing and transmission resources are occupied at the database, the application server, and the communications equipment extending therebetween.

[0005] While the connection setup and release signaling and processing is a load that degrades the performance of the application server and the communications equipment, these penalties can be mitigated relatively inexpensively by adding or upgrading application servers, and increasing network capacity as required. However, the penalty
5 incurred at the database that receives, processes and generates one side of the connection setup and release signaling, is not likewise overcome by duplication. It is not possible to duplicate the digital information of the database for access purposes, without sacrificing consistency and integrity of the digital information. Duplication of the digital information leads to intractable problems with keeping the copies of the digital information consistent.
10 Thus an inherent bottle-neck exists where the modifications to the digital information takes place.

[0006] Accordingly, there is a need for a DBMS that reduces a connection control workload of a resource base, particularly if it provides for a more efficient use of resource base connections to reduce the number of connections to the resource base. Further a
15 method of delivering resource requests in a sorted format to a resource base, is needed.

SUMMARY

[0007] Therefore an object of the invention is to provide a DBMS that reduces a connection control workload associated with a database and concentrates resource requests onto connections between an application server and a database.

20 **[0008]** The invention provides for the reuse of connections for a sequence of database management operations that are serially exchanged with the database. This reduces the connection control processing and signaling load that may be experienced by the DBMS, as well as at the application server.

[0009] Further, in accordance with embodiments of the invention, properties of the
25 queries can be used to improve efficiency of the DBMS. In particular, there are two modes of requesting data from a database: a read access mode, and a write access mode. It happens in most DBMSs that a majority of the database requests are in read access mode, and that the read access mode requests are typically returned by the DBMS much faster, as

they involve a small fraction of the processing that write requests require. The write access mode requests have to obtain exclusive access to the objects they modify, and are subject to transaction control.

5 **[0010]** By separating read access mode operations from write access mode operations, responses to read operations are more expedient, without having to wait for write operations to complete. A further advantage of the separation by request type (such as the operation mode) is that the database that receives the operations over a connection, receives the data sequentially, sorted by type. Accordingly the operations are pooled and assigned to connections in a series.

10 **[0011]** The pooling of the operations provides an opportunity to more efficiently control database connections. Controlling database connections involves creating and taking down database connections, in accordance with predefined criteria. In particular, as there are typically more read than write access mode operations, it is advantageous to ensure that one read type connection is available for as long as there are database requests
15 being received by clients. This ensures that a connection exists when an operation is received.

[0012] According to an aspect of the invention, there is provided a method for directing a server to send database queries to a database management system configured to manage a database, the method including receiving a database request, parsing the database request
20 into operation requests for performing operations on the database, grouping together operation requests of a particular type, and forwarding operation requests of a particular type over an associated connection to the database.

[0013] According to another aspect of the present invention, there is provided a computer program product having a computer-readable medium tangibly embodying
25 computer executable code for directing a server to send database queries to a database management system configured to manage a database, the computer program product including code for receiving a database request, code for parsing the database request into operation requests for performing operations on the database, code for grouping together

operation requests of a particular type, and code for forwarding operation requests of a particular type over an associated connection to the database.

[0014] According to yet another aspect of the invention, there is provided a a server for sending database queries to a database management system configured to manage a database, the server including means for receiving a database request, means for parsing the database request into operation requests for performing operations on the database, means for grouping together operation requests of a particular type, and means for forwarding operation requests of a particular type over an associated connection to the database.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] Further features and advantages of the present invention will become apparent from the following detailed description, taken in combination with the appended drawings, in which:

[0016] FIG. 1 is a schematic diagram of a typical database access system;

[0017] FIG. 2 is a schematic block diagram illustrating principal components of an application server in accordance with the present invention;

[0018] FIG. 3 is a flow chart illustrating principal steps in processing a request from a client at an application server in accordance with the invention;

[0019] FIG. 4 is a flow chart illustrating principal steps involved in a method processing resource requests in accordance with the invention; and

[0020] FIG. 5 is a flow chart illustrating principal steps involved in a method of controlling database connections in accordance with the invention.

[0021] It should be noted that throughout the appended drawings, like features are identified by like reference numerals.

DETAILED DESCRIPTION

[0022] The following detailed description of the embodiments of the present invention is not intended to limit the implementation of the invention to any particular computer programming language. The present invention may be implemented in any computer programming language provided that the OS (Operating System) provides the facilities that may support the requirements of the present invention. A preferred embodiment is implemented in the C or C++ computer programming language (or other computer programming languages in conjunction with C/C++). Any limitations presented would be a result of a particular type of operating system, computer programming language, or data processing system and would not be a limitation of the present invention.

[0023] The invention provides a method for connection control, a method for processing a request, and a system for processing a resource request that improve connection usage between an application server and a digital resource storage and retrieval system, such as a database. The invention reduces the number of connections between the application server and the digital resource storage and retrieval system, thereby reducing the volume of connection control processing and signaling. The invention also provides for the serialized and sorted delivery of requests to the digital resource storage and retrieval system, in order to reduce the amount of processing demanded of the digital resource storage and retrieval system.

[0024] FIG. 1 schematically illustrates a typical database access system 10 for accessing a digital resource storage and retrieval system. The database access system 10 includes an application server 12 that provides a plurality of clients 14 with access to a database 16. Typically, the clients are workstations or stand-alone computers running software application programs that use available programming interfaces and other operating system- and software-provided functionality to connect to the application server 12, which exchanges data with the database 16. A data network 18 interconnects the clients 14 with the application server 12. The connection between the application server 12 and database 16 (i.e. the back-end connection) is not part of the network.

[0025] FIG. 2 schematically illustrates principal functional blocks of a resource processor of the application server 12, in accordance with an exemplary embodiment of the invention. The resource processor functionality of application server 12 includes a request processor 20 that is adapted to process requests from the clients 14 that are written to a receive buffer 22, from one of a plurality of logical input ports 24. The application server 12 also includes a transmit buffer 26 and a plurality of logical output ports 28 through which replies to respective requests are returned to the respective clients 14. The request processor 20 has a parser for inspecting requests and identifying operations to be applied at the database 16, and logic for determining whether these operations need to be run under transaction control. The request processor 20 further performs queue management. The request processor 20 may further perform access control, and other implementation specific operations as required, in a manner well known in the art.

[0026] To understand the request processing operation of the application server 12 in greater detail, an exemplary structure for database requests, and how they are applied at a database is described. Database requests (for example queries) may include a plurality of operations usually called statements, or sentences of a predefined query language, such as a structured query language (SQL). Compound database requests require a plurality of database connections. Each of the database requests specifies one or more operations. Each operation applies to a single database object (a page, a table, a row, a field, a character,...). A database object is a modifiable unit of a database. A page is the finest unit of the digital information that is independently retrievable from storage at the database.

[0027] The processing which an operation performs on an object depends on the keywords and terms of the operation. These operations can usually be classified into two kinds: those that run under transaction control: and, those that do not. In principle it is possible for a database controller to concurrently handle a plurality of different modes for applying the operations, and these modes could be handled with a straightforward application of the present embodiment. The separation of requests into operations respects atomicity of the operations, so that if one operation cannot be committed or aborted unless

a second operation is committed or aborted (due to potential database inconsistency), the two operations are together performed under transaction control.

[0028] Operations that run under transaction control are applied using processing controls that provide for documentation of the processing steps performed during the operation (e.g. annotations to database pages, logging, etc.). The documentation is used during database recovery when restarting after a failure, as will be well understood by those of ordinary skill in the art. Any operation that requires a change to be made to its database object usually runs under transaction control.

[0029] Conversely, operations that do not modify their respective database objects do not require transaction control for most applications. While this general rule is followed by most database systems, it should be noted that in some cases, read operations may need to be performed under transaction control, while some fields of a database (which do not need to be maintained in a consistent state) may not need to be modified under transaction control. Notwithstanding the possible exceptions, operations that require transaction control are for purposes of illustration and explanation herein termed write operations and those that do not are termed read operations.

[0030] Referring again to FIG. 2, the request processor 20 inspects database requests to retrieve the operations, and separates the read operations from write operations. The request processor 20 performs queue management for operation queues. More specifically, two queues (read queue 32 and write queue 34) are provided for temporarily storing the operations, until a connection is available. The read 32 and write 34 queues each store operations of the designated type that are waiting for transmission over database connections associated with the request type of the queue. Each queue is therefore associated with a respective disjoint set of database connections.

[0031] In operation, the request processor 20 places the operations retrieved from the client database requests into the appropriate queue, for retrieval and processing when access to the database becomes available. When a read database connection (a logical construct using one of the input/output (I/O) ports 36, in this case 36a) becomes available,

the request processor 20 identifies the next operation in the read queue 32, and sends the identified operation to the database 16 along the connection. Likewise, when either of write database connection 1 (using I/O port 36b) or write database connection 2 (using I/O port 36c) becomes available, the next operation in the write queue 34 is forwarded to the database 16. The I/O ports 36 are identical, however, the request processor 20 will have designated some I/O ports (e.g. 36a) as being reserved for read operations and others (e.g. 36b and c) as being reserved for write operations. The I/O ports 36 may use a different protocol and type (usually called the back-end connection) from those of the input 24 and output 28 ports that are connected to the clients 14 via network equipment. If the protocols and types are different, the request processor provides translation or reformatting as required. Further procedures for managing a queue of this kind regarding priority, scheduling, etc. may also be performed by the request processor 20.

[0032] The request processor 20 further includes a connection handler for requesting setup and tear-down of database connections. In accordance with the invention, connections may be created and taken down in a more efficient manner than on a client-connection basis. Generally speaking, the connection handler monitors the queues to determine if predefined criteria are met, and sets up or tears down database connections accordingly. The request processor 20 therefore may, depending upon certain criteria, throttle or accelerate the rate at which the operations are sent to the database 12 by adding or removing connections. The condition may relate to the number of operations in a corresponding queue, or some other traffic metric. For example, if the connection handler determines that an empirically-derived threshold (or threshold function) is exceeded in one of the queues, it creates or deletes connections, accordingly. The rate at which the operations are being responded to, and/or the wait time of the objects in queue may also, or alternatively, be monitored to make this determination.

[0033] Other rules for governing the creation and deletion of connections of respective types may also be used. For example, database capacity messages may be used to determine if the application server 12 is the cause of a bottleneck in the database access

system 10, or whether the database is being flooded. The response time of the database may be used to judge the capacity of the database.

[0034] FIG. 3 schematically illustrates, in flow chart form, principal steps that may be performed in processing an operation of a client request at the application server 12, in accordance with the invention. The process begins, and in step 100, the request is received, and parsed to identify the operations contained in the request. All similar operations are considered together. If it is determined in step 102 that one or more of the operations are connection requests, it is determined (in step 104), whether a read connection exists. Because it is assumed that the majority of operations requested are for read operations, in anticipation of such an operation, the procedure verifies that a read connection exists and creates a connection if one does not exist when a connection request is received. This assures that a read connection will have been set up for subsequent read operations, thus imposing a slight performance advantage for read operations, given that read operations are more frequent and less processing intensive than write operations.

[0035] The steps in creating the connection involve requesting a database connection (step 106), receiving acknowledgement of the connection request (step 108) and thereafter determining that the read connection is established (step 110). In step 112, once the existence of a read connection is ensured, a reply is sent to the client 14, to indicate that the connection is established. If it is determined in step 113 that no other operations are contained in the request, the process ends when the request is processed. Otherwise the procedure returns to step 102.

[0036] If, in step 102 it is determined that the received database request does not include any connection request or that these have already been processed, it may be determined (in step 114) that the request includes at least one read operation. If so, each read operation is placed in a read queue (step 116). If a plurality of operations are extracted from a single request, particularly if a plurality of read connections transport respective operations, care must be taken to ensure that the responses are properly associated to the corresponding operations in a manner known to those skilled in the art. Usually such mechanisms make use of operating systems features and lower level protocol processing.

Once the read operation(s) are placed in the read queue, the procedure advances to step 118, where it is determined if the database request contains a write operation.

[0037] If one or more write operations are identified (in step 118), it must be first determined whether or not a write connection exists (in step 120), as connection requests are treated as requiring a read connection to be created. If it is determined that no write connection exists, one is created in steps 122-126 in a manner similar to that described with respect to the steps for creating a read connection described above with reference to steps 106-110. Once a write connection is determined to exist, the write operation(s) of the request are placed in the write queue (step 128).

10 **[0038]** If no write operations are contained in the request, or once all of the write operations have been placed in the write connection queue, it is determined (in step 130) if the request contains a signal for the end of a group of operations called a transaction. Transactions may be distributed across a plurality of client requests, and an end transaction request signals the end of the group. In such a case, in accordance with the bias in favour of read connections, any write connections set up to handle the terminated transaction may be eligible to be taken down. Accordingly, in step 132, it is determined if the transaction (to be terminated) used a write operation, and (in step 134) whether the (corresponding) write queue is empty. If either of these is determined to be false, the transaction is ended (step 138), and the processing of the request is complete. If both (steps 132,134) are found true, then the corresponding write connection is taken down (step 136), and the transaction is ended (step 138). Consequently, the processing of the request is ended. On the other hand, the read connection (if established) is allowed to continue to exist.

25 **[0039]** If none of the previously identified types require further processing (or all operations associated therewith have been placed in respective queues), it is determined (in step 140) whether or not the request contains a disconnect operation. If it does not, the request may be deemed processed, and the process ends until another request is received. Otherwise, in step 142, the connection is taken down. If both the read and write queues are empty (as determined in step 144) the read connection is taken down (in step 146). Presumably, any write queues will have been terminated upon completion of a transaction.

Otherwise there are other client connections open, and the read connection is not taken down.

[0040] It will be appreciated by those skilled in the art that steps of determining if a queue is empty (134,144) may include applying a test of the read and write connections, 5 respective protocol buffers, or control information to determine if a request is being sent or waiting to be sent. Such a test may be applied only once the database has responded to any pending operations retrieved from the database request.

[0041] While the invention has thus far been described with reference to a database system, the invention may be applied in a broader set of environments. Substantially any 10 application server that receives requests for a resource, and supports connections to a source (i.e. a digital resource storage and retrieval system) for each of the requests can be improved by concentrating the requests onto fewer connections, and by so doing, serially providing the requests to the source. If there are different types of requests that are handled differently by the source, associated with different modes for requesting the resources, the 15 application server may sort the requests according to type, and issue the requests over connections associated with respective ones of the request types. Examples of the request modes include the previously described transaction control and its alternative, in relation to a database type of source. In other embodiments, different protocols or sets of applications associated with receiving and/or processing the requests, may have priority, time 20 sensitivity, client history, or any other application-related feature that impacts on how the request is handled by the source may be associated with a respective mode. Accordingly, the invention may have application to object request brokering systems, differentiated web services deployment, and to systems that manage requests for a service, processor capacity, data, or other managed resources that are received at an application layer or server, and 25 require differentiated handling of the requests at the source. If requests are arriving at the source too quickly or too slowly, the number of connections can also be increased/decreased accordingly.

[0042] FIG. 4 is a flow chart schematically illustrating principal steps involved in a method of request processing at an application server, in accordance with the invention.

The method involves receiving a request from a client (step 50), and parsing the request to identify any separate operations contained therein (step 52). In accordance with the illustrated embodiment the requests are parsed to potentially extract a plurality of operations, depending on the specific request received. Each of the operations are further
5 assumed to be one of a first mode and a second mode of requesting resources from the source. Once the operations are identified, they are inspected, in step 54, to determine which mode of requesting is associated with each operation. The number of modes may be other than two, however only two are illustrated. Operations associated with the first mode are issued (step 56) over mode 1 connections. As was the case in the database embodiment,
10 the operations may be queued, subjected to access control procedures and/or translated, prior to being sent to the source. The operation may further be preprocessed and/or sorted in accordance with known techniques for handling requests. The operations associated with the second mode of requesting are sent over mode 2 connections (step 58).

[0043] FIG. 5 schematically illustrates principal steps in a connection control process
15 that may be used in accordance with the invention to control a flow of requests through the application server that employs a technique for queuing the requests. The control process begins, and in step 80, it is determined whether the queues are empty, and whether there are any client connections through the application server to the source. If the request processing is halted, or there are no connections and empty queues, the control process has
20 no function, and so the procedure ends. If it is determined that the request processing is ongoing, then, with some predefined frequency, or in response to a predetermined event, a test is applied to determine if a condition is met for adding or deleting connections of one of the predefined types (step 82). As previously explained, the condition can be defined in terms, for example, of the rate at which requests, or operations, are received, the rate of
25 completion of requests at the application server, the volume in an associated queue, or the time taken by the source to respond to an operation, and therefore can be determined without requiring explicit feedback from the source. Alternatively, recourse to the source may be useful in determining whether to increase or decrease the rate of requests. If the condition is determined not to be met, the control process returns to step 80, otherwise, in

step 84, the addition or deletion of a selected connection of the identified type is effected. Once the addition or deletion is complete, the control process returns to step 80.

[0044] This control process allows for modifying the number of connections to the source in a flexible manner, with reduced connection signaling and processing overhead.
5 In the illustrated embodiment, the control process is separated from the request processing operations. Alternatively, FIG. 5 shows one way that the control process can be integrated with the request processing functions at an application server.

[0045] The invention therefore has been described with reference to a system and a method for controlling connections between a source and an application server function.
10 The method enables fewer connections between the source and application server than a number of client connections to the application server. The reduction in connections results in less connection control processing and signaling. The efficiency of supplying the source with request operations in a format that is sorted by type also provides for more efficient handling of the operations. Further still, the operations are supplied at a rate that is
15 convenient for the source, improving efficiency at the source.

[0046] The present invention can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in any combination thereof. Apparatus of the invention can be implemented in the computer program product tangibly embodied in a machine-readable storage device for execution by a programmable system; and
20 methods/actions can be performed by a programmable processor executing a programmer's instructions to perform functions of the invention by operating on input data and generating output. The invention can be implemented advantageously in one or more computer programs that are executable on a programmable system including at least one input device, and at least one output device. Each computer program can be implemented in a high-level
25 procedural or object-oriented programming language, or in assembly or machine language if desired; and in any case, the language can be a compiled or an interpreted language.

[0047] Suitable processors include, by way or example, both general and specific micro-processors. Generally, a processor will receive instructions and data from a read-

only memory and/or a random access memory. Generally, a computer will include one or more mass storage devices for storing data files; such devices include magnetic disks, such as internal hard disks and removable disks; magneto-optical disks; and optical disks. Storage devices suitable for tangibly embodying computer program instructions and data
5 include all forms of non-volatile memory, including by way of example semi-conductor memory devices such as EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM and DVD disks. Any of the foregoing can be supplemented by, or incorporated in, ASICs application-specific integrated circuits.

10 **[0048]** Programmable processing systems contained in the resource processor shown in FIG. 2 and the application servers shown in FIG. 1 are suitable for implementing or performing the apparatus or methods of the invention. The system may comprise a processor, a random access memory, a hard drive controller and an input/output controller coupled by a processor bus.

15 **[0049]** It will be apparent to those skilled in this art that various modifications and variations may be made to the embodiments disclosed herein, consistent with the present invention, without departing from the spirit and scope of the present invention.

[0050] For example, the data network 18 may be a public Internet, and some embodiments it may be a private intranet, or it may be a combination of the two. Further
20 the network may be of any known configuration or type that provides for the communication of data. The network may further include gateway servers, and the application server may serve a plurality of databases and other back end systems. Furthermore, in other embodiments, the application server 12, client 14 and database 16 may all be instantiated on one or more processors or processor systems requiring no
25 network.

[0051] An operating system associated with the request processor 20 handles exchange of request and response messages, assignment of ports to handle network connections, and other lower level communications functions. Depending on the

embodiment of the invention, the operating system may take on more or less of the responsibility for connection configuration and control. Thus, while the request processor 20 may be embodied as a software application that includes the connection handler, which requests the addition and removal of the database connections using lower
5 level service primitives, the connection handler may be embodied by lower level processing, or a combination of both.

[0052] It should be noted that while two queues have been illustrated in FIG. 2, one queue, a plurality of queues each associated with a priority value, a priority value and a connection type, or a specific connection, can be used in other embodiments. If each of a
10 plurality of read and write database connections is associated with a respective read or write queue, the assembling of database responses to formulate response messages to be sent to clients may be somewhat simplified. However, having only one read queue and one write queue regardless of the number of read and write database connections, simplifies the addition and deletion of connections, and more evenly distributes delay among the
15 operations. Other relationships between queues and connections may be suggested by hardware limitations or by traffic and service requirements or otherwise suggested for use in particular embodiments.

[0053] In FIG. 3, the receipt of a connection request in step 102 led to the assurance of the existence of a read queue. The assurance may be relied upon by a connection handler to
20 predict imminent usage and to create connections to handle the predicted demand. In other embodiments assurance of both read and write connections may be performed, especially if a time to setup a connection is substantial. Alternatively, the connection types assured may depend on a type, or a history of the client, for example. Furthermore, in some embodiments no connection is created until a read or write request is received. The
25 procedures for controlling the connections is preferably optimized to permit a highest usage with a minimum of underutilized capacity.

[0054] Other embodiments consistent with the present invention will become apparent from consideration of the specification and the practice of the invention disclosed therein.

[0055] Accordingly, the specification and the embodiments are to be considered exemplary only, with a true scope and spirit of the invention being disclosed by the appended claims.

The embodiments of the invention in which an exclusive property or privilege is claimed are defined as follows:

1. A computer-implemented method for directing a server to send a database query to a database management system that manages a database, the method
5 comprising:

receiving a database request;

parsing the database request into a plurality of operation requests for performing operations on the database;

10 separately grouping the operation requests into at least a first group and a second group, wherein the first group contains a sequence of serial operations of a first particular type, and wherein the second group contains a sequence of serial operations of a second particular type that is distinct from the first particular type;

wherein the first group includes read operations;

15 wherein the second group includes write operations, so that responses to the read operations do not await the completion of the write operations

establishing a dedicated first connection for the operations of the first group;

establishing a dedicated second connection for the operations of the second group;

20 forwarding the operation requests of the first group over the first connection to the database;

maintaining and reusing the first connection until all the operations of the first group have been forwarded to the database;

25 forwarding the operation requests of the second group over the second connection to the database; and

maintaining and reusing the second connection until all the operations of the second group have been forwarded to the database, sequentially, sorted by type, so that the read operations are pooled in series and assigned to the first connection, and the write operations are pooled in series and assigned to the second connection.

2. The method of claim 1, further comprising taking down the first connection upon exhaustion of all the operations of the first group.

3. The method of claim 1, further comprising faking down the second connection upon exhaustion of all the operations of the second group.

5 4. The method of claim 2, wherein taking down the first connection comprises taking down the first connection if there are insufficient operation requests of the first group in a queue associated with the operation requests.

5. The method of claim 3, wherein taking down the second connection comprises taking down the second connection if there are insufficient operation requests
10 of the second group in a queue associated with the operation requests.

6. A computer program product having instruction codes stored on a computer-usable medium, for directing a server to send a database query to a database management system that manages a database, the computer program product comprising:

15 a set of instruction codes for receiving a database request;

a set of instruction codes for parsing the database request into a plurality of operation requests for performing operations on the database;

20 a set of instruction codes for separately grouping the operation requests into at least a first group and a second group, wherein the first group contains a sequence of serial operations of a first particular type, and wherein the second group contains a sequence of serial operations of a second particular type that is distinct from the first particular type;

wherein the first group includes read operations;

25 wherein the second group includes write operations, so that responses to the read operations do not await the completion of the write operations;

a set of instruction codes for establishing a dedicated first connection for the operations of the first group;

a set of instruction codes for establishing a dedicated second connection for the operations of the second group;

a set of instruction codes for forwarding the operation requests of the first group over the first connection to the database;

a set of instruction codes for maintaining and reusing the first connection until all the operations of the first group have been forwarded to the database;

5 a set of instruction codes for forwarding the operation requests of the second group over the second connection to the database; and

a set of instruction codes for maintaining and reusing the second connection until all the operations of the second group have been forwarded to the database, sequentially, sorted by type, so that the read operations are pooled in series and assigned
10 to the first connection, and the write operations are pooled in series and assigned to the second connection.

7. The computer program product of claim 6, further comprising a set of instruction codes for taking down the first connection upon exhaustion of all the operations of the first group.

15 8. The computer program product of claim 6, further comprising a set of instruction codes for taking down the second connection upon exhaustion of all the operations of the second group.

9. The computer program product of claim 7, wherein the set of instruction codes takes down the first connection if there are insufficient operation requests of the
20 first group in a queue associated with the operation requests.

10. The computer program product of claim 8, wherein the set of instruction codes takes down the second connection if there are insufficient operation requests of the second group in a queue associated with the operation requests.

11. A computer-implemented database management system for directing a
25 server to send a database query to a database management system that manages a database, the database management system comprising:

means for receiving a database request;

means for parsing the database request into a plurality of operation requests for performing operations on the database;

means for separately grouping the operation requests into at least a first group and a second group, wherein the first group contains a sequence of serial operations of a first particular type, and wherein the second group contains a sequence of serial operations of a second particular type that is distinct from the first particular type;

5 wherein the first group includes read operations;

 wherein the second group includes write operations so that responses to the read operations do not await the completion of the write operations;

 establishing a dedicated first connection for the operations of the first group;

 establishing a dedicated second connection for the operations of the second
10 group;

 forwarding the operation requests of the first group over the first connection to the database;

 maintaining and reusing the first connection until all the operations of the first group have been forwarded to the database;

15 forwarding the operation requests of the second group over the second connection to the database; and

 maintaining and reusing the second connection until all the operations of the second group have been forwarded to the database, sequentially, sorted by type, so that the read operations are pooled in series and assigned to the first connection, and the
20 write operations are pooled in series and assigned to the second connection

12. The database management system of claim 11, further means for comprising taking down the first connection upon exhaustion of all the operations of the first group.

13. The database management system of claim 12, wherein the means for setting up and taking down the associated connections favor operation requests of a first type
25 over operation requests of a second type, further comprising means for taking down the second connection upon exhaustion of all the operations of the second group.

14. The database management system of claim 12, wherein the means for taking down the first connection takes down the first connection if there are insufficient operation requests of the first group in a queue associated with the operation requests.

15. The database management system of claim 13, wherein the means for taking down the second connection takes down the second connection if there are insufficient operation requests of the first group in a queue associated with the operation requests.




FIG. 1
Prior Art

FIG. 2




FIG. 3

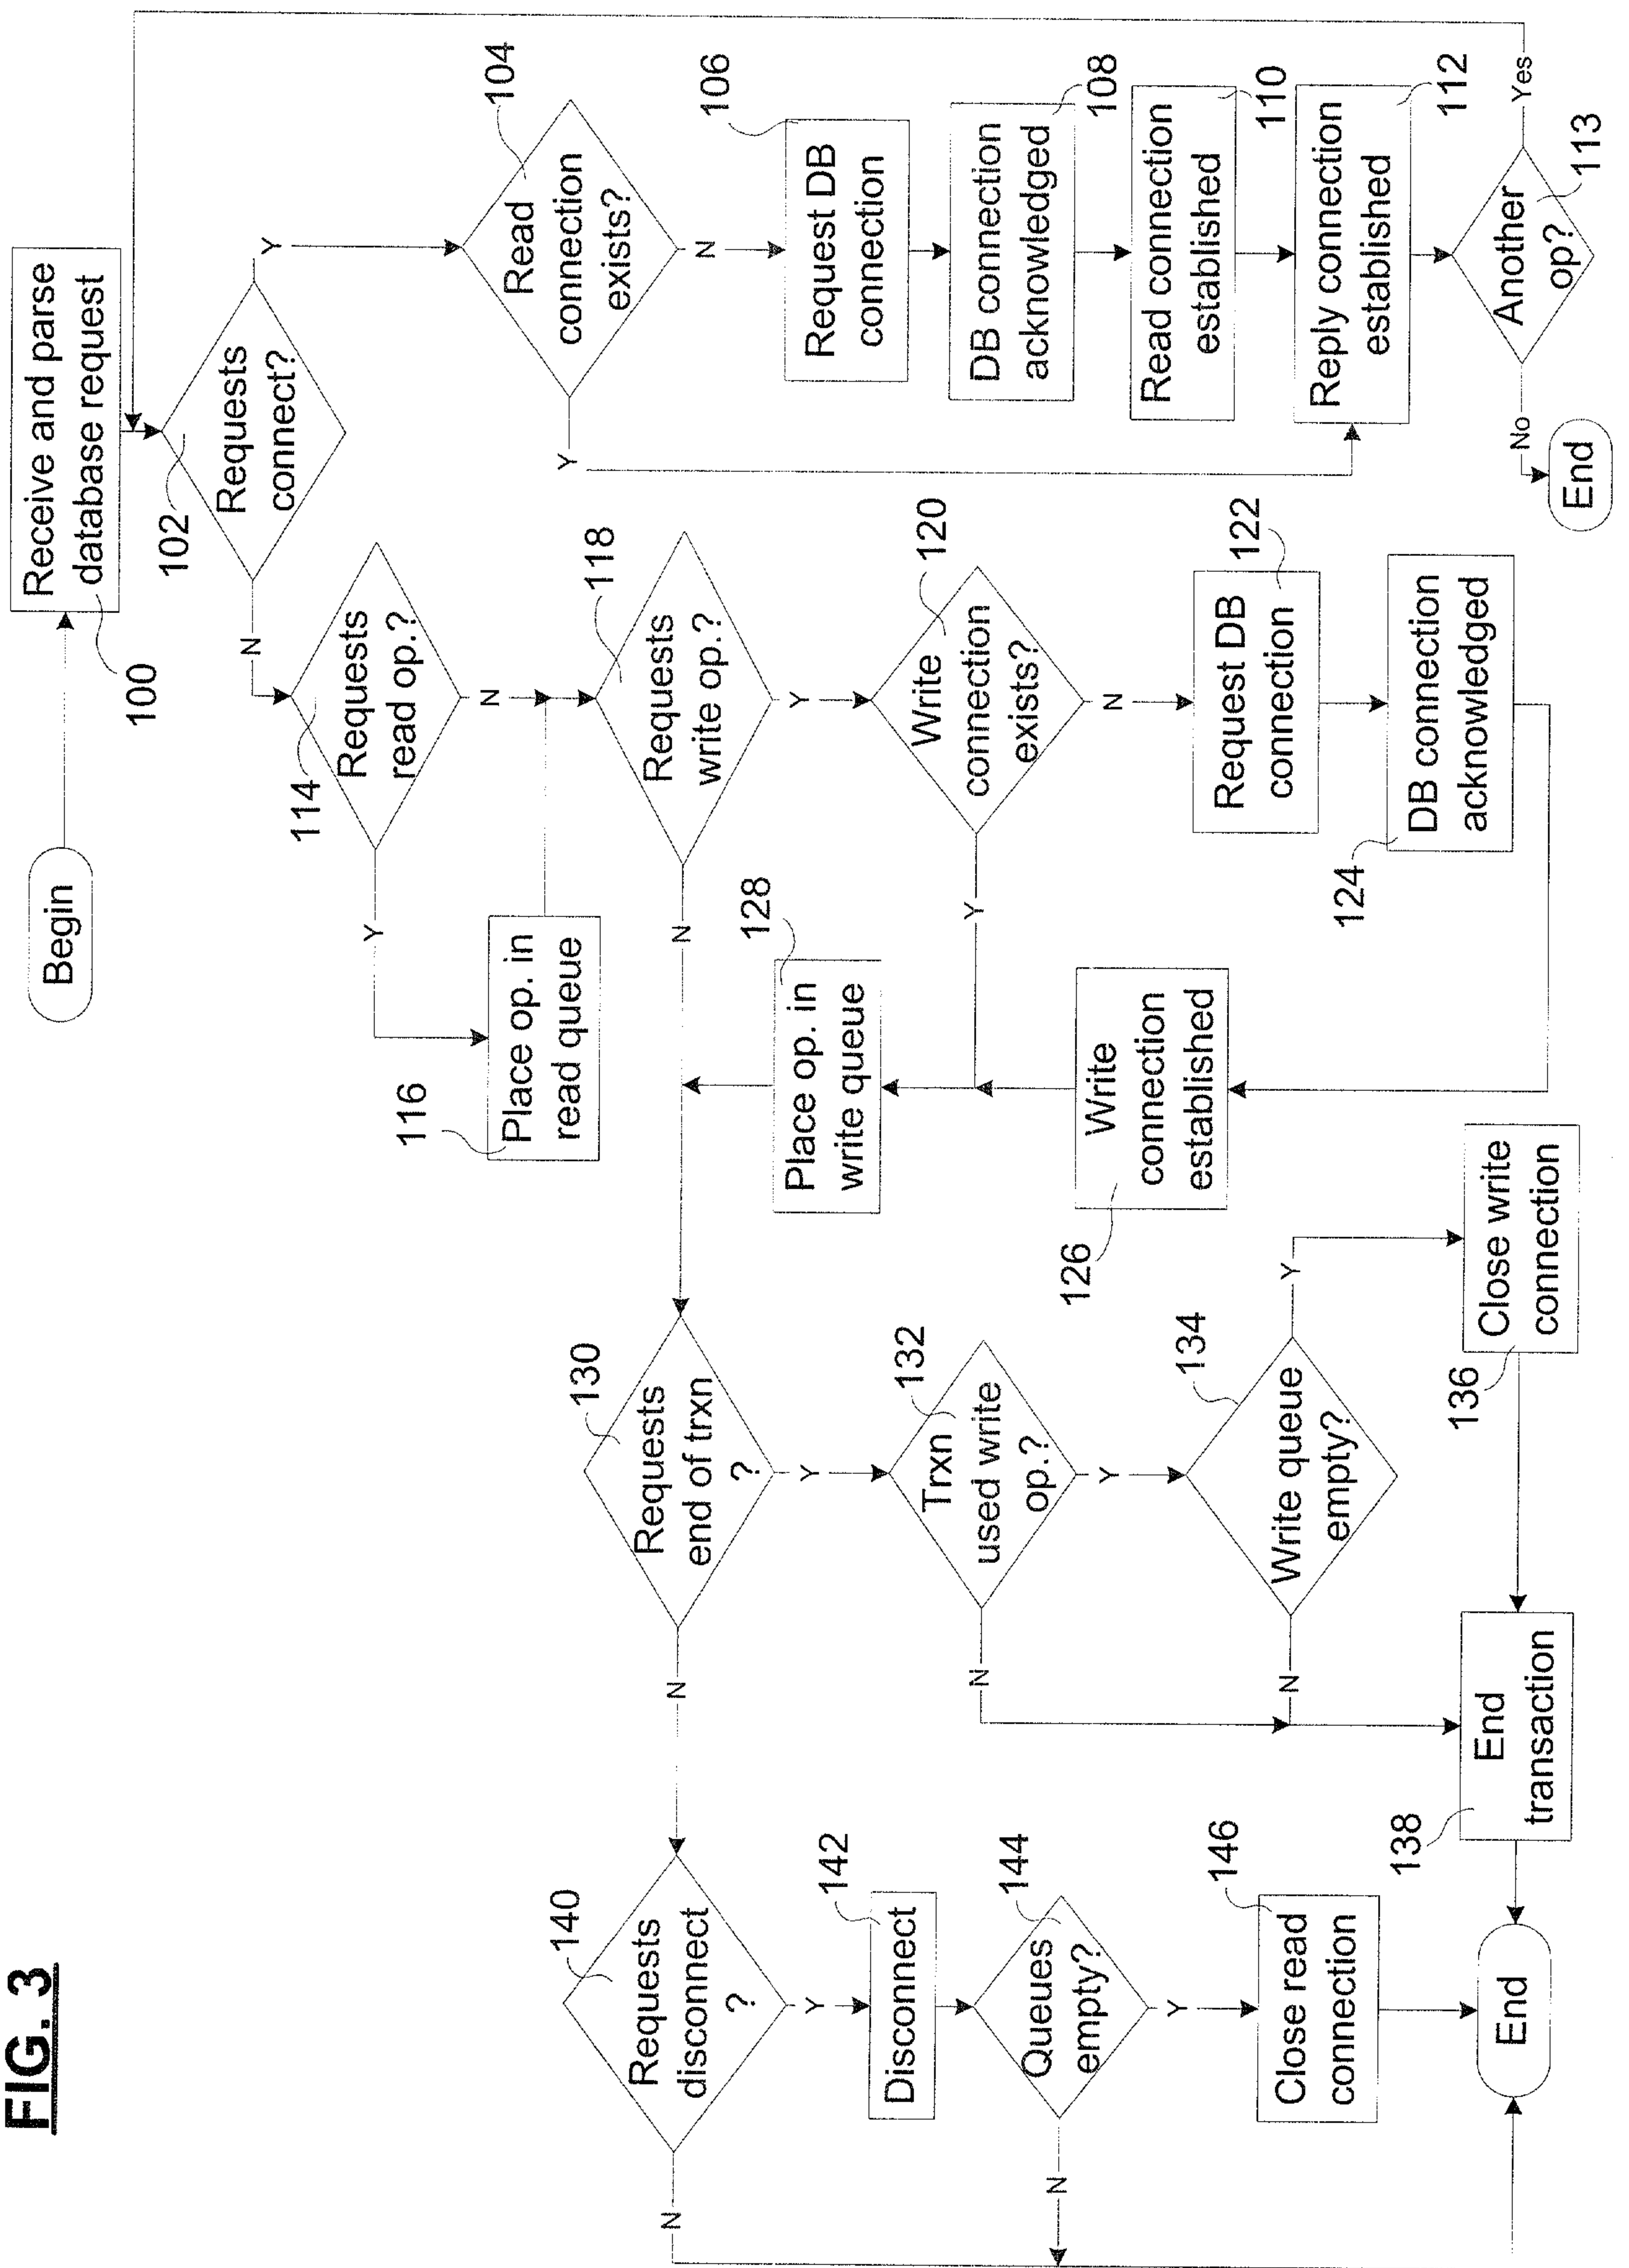


FIG. 4

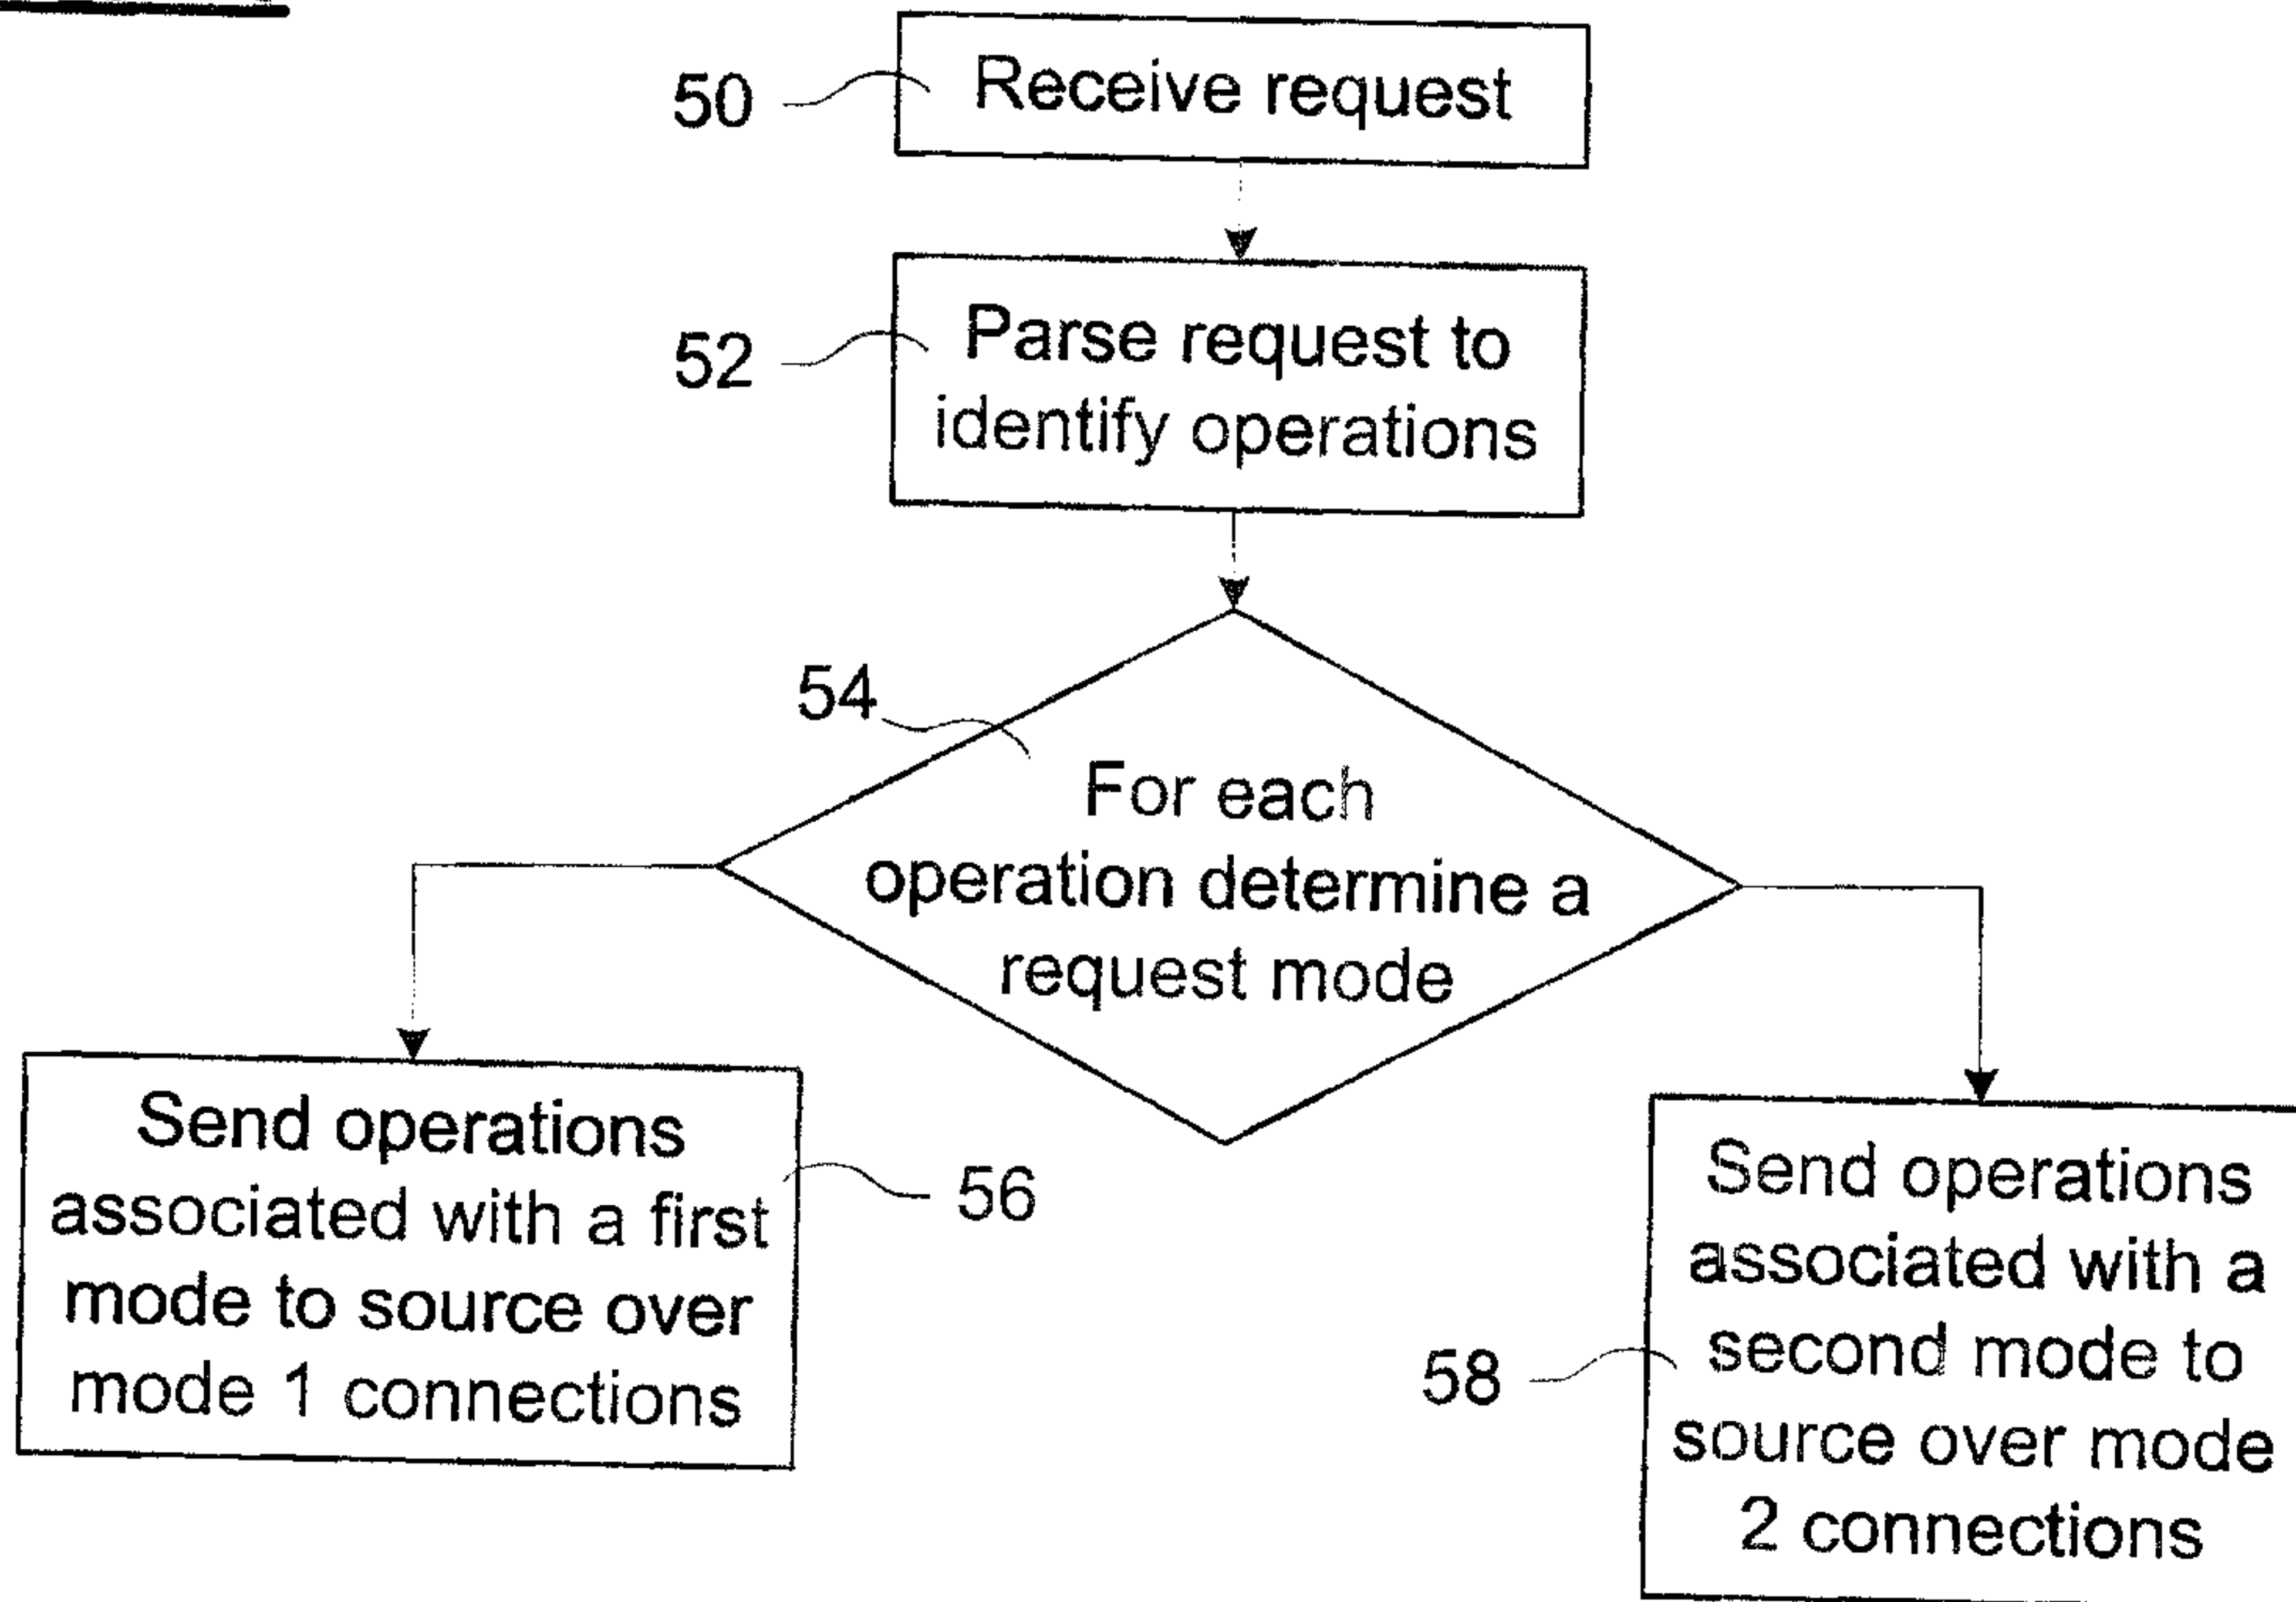


FIG. 5

