

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
15 March 2007 (15.03.2007)

PCT

(10) International Publication Number
WO 2007/028248 A1

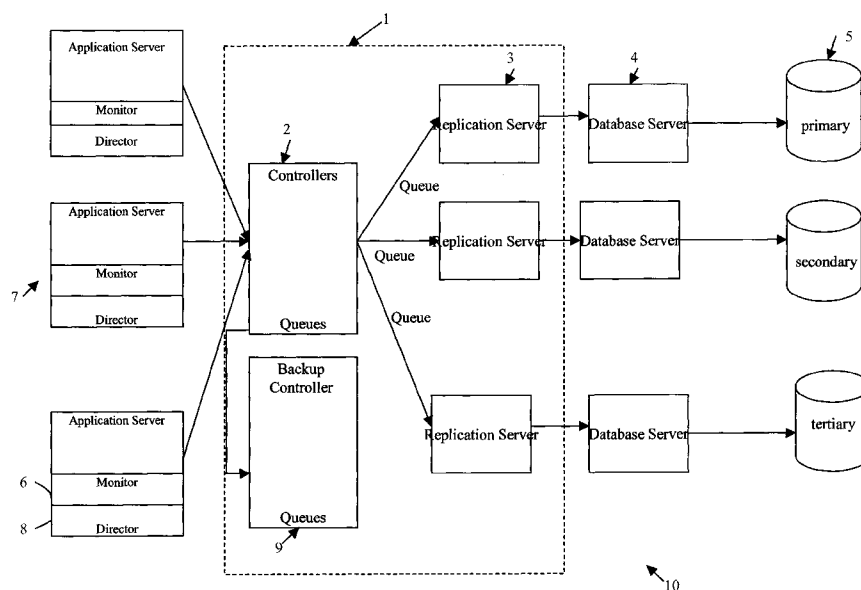
- (51) International Patent Classification:
G06F 17/00 (2006.01) *G06F 17/30* (2006.01)
- (21) International Application Number:
PCT/CA2006/001474
- (22) International Filing Date:
8 September 2006 (08.09.2006)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/221,752 9 September 2005 (09.09.2005) US
- (71) Applicant (for all designated States except US): **AVOKIA INC.** [CA/CA]; 36 Toronto Street, Suite 500, Toronto, Ontario M5C 2C5 (CA).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **WONG, Frankie** [CA/CA]; 1022 Mountcastle Crescent, Pickering, Ontario L1V 5H9 (CA). **YU, Xiong** [CA/CA]; 2 Hollybrook Crescent, North York, Ontario M2J 2H6 (CA). **WANG, Elaine** [CA/CA]; 35 Marchview Avenue, Aurora, Ontario L4G 7W5 (CA).
- (74) Agents: **TISDALL, Grant** et al.; Gowling Lafleur Henderson LLP, Suite 1600, 1 First Canadian Place, 100 King Street West, Toronto, Ontario M5X 1G5 (CA).

- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD AND APPARATUS FOR SEQUENCING TRANSACTIONS GLOBALLY IN A DISTRIBUTED DATABASE CLUSTER



(57) Abstract: A system and method for receiving and tracking a plurality of transactions and distributing the transactions to at least two replication queues over a network. The system and method comprise a global queue for storing a number of the received transactions in a first predetermined order. The system and method also comprise a sequencer coupled to the global queue for creating a copy of each of the transactions for each of said at least two replication queues and for distributing in a second predetermined order each said copy to each of said at least two replication queues respectively, said copy containing one or more of the received transactions.

WO 2007/028248 A1

**METHOD AND APPARATUS FOR SEQUENCING TRANSACTIONS GLOBALLY
IN A DISTRIBUTED DATABASE CLUSTER**

This application is a Continuation-In-Part of U.S. Patent Application 11/221,752, filed September 9, 2005, the contents of which are herein incorporated by reference.

FIELD OF THE INVENTION

[0001] This invention relates generally to the sequencing and processing of transactions within a cluster of replicated databases.

BACKGROUND OF THE INVENTION

[0002] A database has become the core component of most computer application software nowadays. Typically application software makes use of a single or multiple databases as repositories of data (content) required by the application to function properly. The application's operational efficiency and availability is greatly dependent on the performance and availability of these database(s), which can be measured by two metrics: (1) request response time; and (2) transaction throughput.

[0003] There are several techniques for improving application efficiency based on these two metrics: (1) Vertical scale up of computer hardware supporting the application – this is achieved by adding to or replacing existing hardware with faster central processing units (CPUs), random access memory (RAM), disk adapters / controllers, and network; and (2) Horizontal scale out (clustering) of computer hardware supporting the application – this approach refers to connecting additional computing hardware to the existing configuration by interconnecting them with a fast network. Although both approaches can address the need of reducing request response time and increase transaction throughput, the scale out approach can offer higher efficiency at lower costs, thus driving most new implementations into clustering architecture.

[0004] The clustering of applications can be achieved readily by running the application software on multiple, interconnected application servers that facilitate the execution of the application software and provide hardware redundancy for high availability, with the application software actively processing requests concurrently. However current database clustering technologies cannot provide the level of availability and redundancy in a similar active-active configuration. Consequently database servers are primarily configured as active-standby, meaning that one of the computer systems in the cluster does not process

application request until a failover occurs. Active-standby configuration wastes system resources, extends the windows of unavailability and increases the chance of data loss.

[0005] To cluster multiple database servers in an active-active configuration, one technical challenge is to resolve update conflict. An update conflict refers to two or more database servers updating the same record in the databases that they manage. Since data in these databases must be consistent among them in order to scale out for performance and achieve high availability, the conflict must be resolved. Currently there are two different schemes of conflict resolution: (1) time based resolution; and (2) location based resolution. However, neither conflict resolution schemes can be enforced without some heuristic decision to be made by human intervention. It is not possible to determine these heuristic decision rules unless there is a thorough understanding of the application software business rules and their implications. Consequently, most clustered database configurations adopt the active-standby model, and fail to achieve high performance and availability at the same time. There is a need for providing a database management system that uses an active-active configuration and substantially reduces the possibility of update conflicts that may occur when two or more databases attempt to update a record at the same time.

[0006] The systems and methods disclosed herein provide a system for globally managing transaction requests to one or more database servers and to obviate or mitigate at least some of the above presented disadvantages.

SUMMARY OF THE INVENTION

[0007] To cluster multiple database servers in an active-active configuration, one technical challenge is to resolve update conflict. An update conflict refers to two or more database servers updating the same record in the databases that they manage. Since data in these databases must be consistent among them in order to scale out for performance and achieve high availability, the conflict must be resolved. Currently there are two different schemes of conflict resolution: (1) time based resolution; and (2) location based resolution. However, neither conflict resolution schemes can be enforced without some heuristic decision to be made by human intervention. Consequently, most clustered database configurations adopt the active-standby model, and fail to achieve high performance and availability at the same time. Contrary to current database configurations there is provided a system and method for receiving and tracking a plurality of transactions and distributing the transactions to at least two replication queues over a network. The system and method

comprise a global queue for storing a number of the received transactions in a first predetermined order. The system and method also comprise a sequencer coupled to the global queue for creating a copy of each of the transactions for each of said at least two replication queues and for distributing in a second predetermined order each said copy to each of said at least two replication queues respectively, said copy containing one or more of the received transactions.

[0008] One aspect provided is a system for receiving and tracking a plurality of transactions and distributing the transactions to at least two replication queues, the system comprising: a global queue for storing a number of the received transactions in a first predetermined order; and a sequencer coupled to the global queue for creating a copy of each of the transactions for each of said at least two replication queues and for distributing in a second predetermined order each said copy to each of said at least two replication queues respectively, said copy containing one or more of the received transactions.

[0009] A further aspect provided is a system for receiving a plurality of transactions from at least one application server, distributing the transactions to at least two replication queues and applying the transactions to a plurality of databases comprising: a director coupled to each of said at least one application server for capturing a plurality of database calls therefrom as the plurality of transactions; and a controller for receiving each of the plurality of transactions, the controller configured for storing the transactions within a global queue in a predetermined order, for generating a copy of each said transaction for each of said at least two replication queues, and for transmitting in the predetermined order each said copy to each of said at least two replication queues respectively.

[0010] A still further aspect provided is a method for receiving and tracking a plurality of transactions and distributing the transactions to at least two replication queues, the method comprising: storing a number of the received transactions in a first predetermined order in a global queue; creating a copy of each of the transactions for each of said at least two replication queues; and distributing in a second predetermined order each said copy to each of said at least two replication queues respectively, said copy containing one or more of the received transactions.

[0011] A still further aspect provided is a system for receiving and tracking a plurality of transactions and distributing the transactions to at least two replication queues, the system

comprising: means for storing a number of the received transactions in a first predetermined order; and means for creating a copy of each of the transactions for each of said at least two replication queues and for distributing in a second predetermined order each said copy to each of said at least two replication queues respectively, said copy containing one or more of the received transactions.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] Exemplary embodiments of the invention will now be described in conjunction with the following drawings, by way of example only, in which:

[0013] Figure 1A is a block diagram of a system for sequencing transactions;

[0014] Figure 1B is a block diagram of a transaction replicator of the system of Figure 1A;

[0015] Figure 1C, 1D and 1E show an example operation of receiving and processing transactions for the system of Figure 1A;

[0016] Figure 1F shows a further embodiment of the transaction replicator of the system of Figure 1A;

[0017] Figure 2 is a block diagram of a director of the system of Figure 1A;

[0018] Figure 3 is a block diagram of a monitor of the system of Figure 1A;

[0019] Figure 4 is an example operation of the transaction replicator of Figure 1B;

[0020] Figure 5 is an example operation of a global transaction queue and a replication queue of Figure 1B;

[0021] Figure 6 is an example operation of the transaction replicator of Figure 1B for resolving gating and indoubt transactions; and

[0022] Figure 7 is an example operation of a replication server of Figure 1B.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0023] A method and apparatus for sequencing transactions in a database cluster is described for use with computer programs or software applications whose functions are designed primarily to replicate update transactions to one or more databases such that data in these databases are approximately synchronized for read and write access.

[0024] Referring to Figure 1A, shown is a system 10 comprising a plurality of application servers 7 for interacting with one or more database servers 4 and one or more databases 5 via a transaction replicator 1. It is understood that in two-tier applications, each of the application 7 instances represents a client computer. For three-tiered applications, each of the application 7 instances represents an application server that is coupled to one or more users (not shown). Accordingly, it is recognized that the transaction replicator 1 can receive transactions from applications 7, application servers 7, or a combination thereof.

[0025] Referring to Figures 1A and 1B, the transaction replicator 1 of the system 10, receives transaction requests from the application servers 7 and provides sequenced and replicated transactions using a controller 2 to one or more replication servers 3, which apply the transactions to the databases 5. By providing sequencing of transactions in two or more tiered application architectures, the transaction replicator 1 helps to prevent the transaction requests from interfering with each other and facilitates the integrity of the databases 5. For example, a transaction refers to a single logical operation from a user application 7 and typically includes requests to read, insert, update and delete records within a predetermined database 5.

[0026] Referring again to Figure 1A, the controller 2 can be the central command center of the transaction replicator 1 that can run for example on the application servers 7, the database servers 4 or dedicated hardware. The controller 2 may be coupled to a backup controller 9 that is set up to take over the command when the primary controller 2 fails. The backup controller 9 is approximately synchronized with the primary controller such that transaction integrity is preserved. It is recognized that the controller 2 and associated transaction replicator 1 can also be configured for use as a node in a peer-to-peer network, as further described below.

[0027] Referring again to Figure 1A, when a backup and a primary controller are utilized, a replica global transaction queue is utilized. The backup controller 9 takes over control of transaction replicator 1 upon the failure of the primary controller 2. Preferably, the primary and backup controllers are installed at different sites and a redundant WAN is recommended between the two sites.

[0028] As is shown in Figure 1B, the controller 2 receives input transactions 11 from a user application 7 and provides sequenced transactions 19 via the replication servers 3, the

sequenced transactions 19 are then ready for commitment to the database servers 4. The controller 2 comprises a resent transaction queue 18 (resent TX queue), an indoubt transaction queue 17 (indoubt TX queue), a global transaction sequencer 12 (global TX sequencer), a global TX queue 13 (global TX queue) and at least one global disk queue 14. The global queue 13 (and other queues if desired) can be configured as searchable a first-in-first out pipe (FIFO) or as a first-in-any-out (FIAO), as desired. For example, a FIFO queue 13 could be used when the contents of the replication queues 15 are intended for databases 5, and a FIAO queue 13 could be used when the contents of the replication queues 15 are intended for consumption by unstructured data processing environments (not shown). Further, it is recognized that the global disk queue 14 can be configured for an indexed and randomly accessible data set.

[0029] The transaction replicator 1 maintains the globally sequenced transactions in two different types of queues: the global TX queue 13 and one or more replication queues 15 equal to that of the database server 4 instances. These queues are created using computer memory with spill over area on disks such as the global disk queue 14 and one or more replication disk queues 16. The disk queues serve a number of purposes including: persist transactions to avoid transaction loss during failure of a component in the cluster; act as a very large transaction storage (from gigabytes to terabytes) that computer memory cannot reasonably provide (typically less than 64 gigabytes). Further, the indoubt TX queue 17 is only used when indoubt transactions are detected after a certain system failures. Transactions found in this queue have an unknown transaction state and require either human intervention or pre-programmed resolution methods to resolve.

[0030] For example, in the event of a temporary communication failure resulting in lost response from the global TX sequencer 12 to a transaction ID request, the application resends the request which is then placed in the resent TX queue 18. Under this circumstance, there can be two or more transactions with different Transaction ID in the global TX queue 13 and duplicated transactions are removed subsequently.

[0031] In normal operation, the controller 2 uses the global TX queue 13 to track the status of each of the input transactions and to send the committed transaction for replication in sequence. It is recognized that monitoring of the status of the transactions can be done by the director 8, the controller 2, or combination thereof.

[0032] Referring to Figures 1C, 1D, and 1E, shown is an example operation of the system 10 for receiving and processing a new transaction. In one embodiment, the new transaction is placed in the global queue 13 at commit time, e.g. when the transaction ID (represented by references K, L) is issued, thus denoting to the director 8 (or other database 5 status monitoring entity) that the transaction transmit request is recordable to signify the application 7 is allowed to commit its transmit request (associated with the new transaction) to the database 5. Commit time can be defined to include the step of: 1) the transmit request (associated with the application 7 and the new transaction) are recorded at the director 8; 2) thus providing for passing of the new transaction (e.g. a SQL statement) to the controller 2 by the director 8; 3) the controller 2 then issues the transaction ID (e.g. a commit token K, L) coupled to the new transaction; and 4) the new transaction along with the issued transaction ID (e.g. K, L) are added to the transaction sequence held in the global queue 13.

[0033] For example, upon receiving a new transaction, the sequencer 12 assigns a new transaction ID to the received transaction. The transaction ID is a globally unique sequence number for each transaction within a replication group. In Figure 1C, the sequence ID for the newly received transaction is "K". Once the controller 2 receives the transaction, the transaction and its ID are transferred to the global TX queue 13 if there is space available. Otherwise, if the global TX queue 13 is above a predetermined threshold and is full, for example, as shown in Figure 1C, the transaction K and its ID are stored in the global disk queue 14 (Figure 1D).

[0034] Before accepting any new transactions in the global TX queue, the sequencer distributes the committed transactions from the global TX queue 13 to a first replication server 20 and a second (or more) replication server 23 for execution against the databases. As will be discussed, the transfer of the transactions to the replication servers can be triggered when at least one of the following two criteria occurs: 1) a predetermined transfer time interval and 2) a predetermined threshold for the total number of transactions within the global TX queue 13 is met. However, each replication server 20, 23 has a respective replication queue 21, 24 and applies the sequenced transactions, obtained from the global queue 13, at its own rate to the respective database servers 4 over the network.

[0035] For example, when a slower database server 4 is unable to process the transactions at the rate the transactions are distributed by the controller 2, the transactions in the corresponding replication queue are spilled over to the replication disk queues. As shown

in Figures 1C and 1D, transaction F is transferred from the global TX queue 13 to the first and second replication servers 20, 23. The first replication server 20 has a first replication queue 21 and a first replication disk queue 22 and the second replication server 23 has a second replication queue 24 and a second replication disk queue 25. The replication queues are an ordered repository of update transactions stored in computer memory for executing transactions on a predetermined database. In this case, since the second replication queue 24 is above a predetermined threshold (full, for example) transaction F is transferred to the second replication disk queue 25. Referring to Figure 1D and Figure 1E, once space opens up in the second replication queue 24 as transaction J is applied to its database server, the unprocessed transaction F in the second replication disk queue 25 is moved to the second replication queue 24 for execution of the transaction request against the data within its respective database. In the case where both the replication disk queue and the replication queues are above a preselected threshold (for example, full), an alert is sent by the sequencer 12 and the database 5 is marked unusable until the queues become empty.

[0036] Referring to Figure 1F, shown is the replication server 20 further configured for transmission of the transaction contents 300 of the replication queue 21 (and replication disk queue 22 when used) to two or more database servers 4 that are coupled to respective databases 5. Accordingly, the replicated transactions 300 queued in the replication queue 21 may also be executed concurrently (i.e. in parallel) through multiple concurrent database connections 304 to the second or additional databases 5, for facilitating performance increases in throughput of the replicated transactions 300 against the secondary and/or tertiary databases 5. It is recognised that the replication server 20 coordinates the emptying of the replication queue 21 and disk queue 22 using sequential and/or parallel transmission of the replicated transactions 300 contained therein. The working principle is that when selected ones of the replicated transactions 300 are updating mutually exclusive records R_i , the selected replicated transactions 300 have no sequential dependency and can be executed concurrently using the multiple concurrent database connections 304. The system allows concurrent execution of transactions on the primary database, as described above. So naturally these transactions executed concurrently on the primary database can be assured exclusivity by the respective database engine/servers 4 through locking, and can be executed concurrently as the replicated transactions 300 on the secondary databases 5 accordingly.

[0037] Further, it is recognised that each of the replicated transactions 300 include one or more individual statements 302(e.g. SQL statement or database record access requests) for execution against the respective database 5. For example, each of the statements 302 in a respective replicated transaction 300 can be used to access different records Ri (e.g. R1 and R2) for the databases 5. The replication server 20 can be further configured for concurrent transmission of individual statements 302, from the same or different ones of the replicated transactions 300, for execution against the same or different databases 5 using the one or more concurrent database connections 304. For example, the SQL statements 302 in one of the replicated transactions 300 may be executed concurrently with the SQL statements 302 from another of the replicated transactions 300 in the replication queue 21. The replication server 20 has knowledge of the contents (one or more individual statements 302)of the replicated transactions 300 to assist in selection (e.g. accounting for execution order and/or which record Ri affected) of which transactions 300 to apply in parallel using the multiple concurrent database connections 304, i.e. have no sequential dependency . This knowledge can be represented in the transaction IDs associated with the replicated transactions 300 and/or the individual statements 302, for example.

[0038] In view of the above, it is also recognised that the replication server 20 can coordinate the transmission of the replicated transactions 300 and/or the individual statements 302 from multiple replication queues 21 to two or more databases 5, as desired.

[0039] The core functions of the controller 2 can be summarized as registering one or more directors 8 and associating them with their respective replication groups; controlling the replication servers' 3 activities; maintaining the global TX queue 13 that holds all the update transactions sent from the directors 8; synchronizing the global TX queue 13 with the backup controller 9(where applicable); managing all replication groups defined; distributing committed transactions to the replication servers 3; tracking the operational status of each database server 4 within a replication group; providing system status to a monitor 6; and recovering from various system failures.

[0040] The registry function of the controller 2 occurs when applications are enabled on a new application server 7 to access databases 5 in a replication group. Here, the director 8 on the new application server contacts the controller 2 and registers itself to the replication group. Advantageously, this provides dynamic provisioning of application servers to scale up system capacity on demand. The registration is performed on the first database call made by

an application. Subsequently the director 8 communicates with the controller 2 for transaction and server status tracking.

[0041] The replication server control function allows the controller 2 to start the replication servers 3 and monitors their state. For example, when an administrator requests to pause replication to a specific database 5, the controller then instructs the replication server to stop applying transactions until an administrator or an automated process requests it.

[0042] The replication group management function allows the controller 2 to manage one or more groups of databases 5 that require transaction synchronization and data consistency among them. The number of replication groups that can be managed and controlled by the controller 2 is dependent upon the processing power of the computer that the controller is operating on and the sum of the transaction rates of all the replication groups.

Director

[0043] Referring to Figure 2, shown is a block diagram of the director 8 of the system 10 of Figure 1A. The director can be installed on the application server 7 or the client computer. The director 8 is for initiating a sequence of operations to track the progress of a transaction. The director 8 comprises a first 27, a second 28, a third 29 and a fourth 30 functional module. According to an embodiment of the system 10, the director 8 wraps around a vendor supplied JDBC driver. As discussed earlier, the director 8 is typically installed on the application server 7 in a 3-tier architecture, and on the client computer in a 2-tier architecture. As a wrapper, the director 8 can act like an ordinary JDBC driver to the applications 7, for example. Further, the system 10 can also support any of the following associated with the transaction requests, such as but not limited to:

1. a database access driver/protocol based on SQL for a relational database 5 (ODBC, OLE/DB, ADO.NET, RDBMS native clients, etc...);
2. messages sent over message queues of the network ;
3. XML (and other structured definition languages) based transactions; and
4. other data access drivers as desired.

[0044] As an example, the first module 27 captures all JDBC calls 26, determines transaction type and boundary, and analyzes the SQLs in the transaction. Once determined to be an update transaction, the director 8 initiates a sequence of operations to track the progress

of the transaction until it ends with a commit or rollback. Both DDL and DML are captured for replication to other databases in the same replication group.

[0045] The second module 28 collects a plurality of different statistical elements on transactions and SQL statements for analyzing application execution and performance characteristics. The statistics can be exported as comma delimited text file for importing into a spreadsheet.

[0046] In addition to intercepting and analyzing transactions and SQL statements, the director's third module 29, manages database connections for the applications 7. In the event that one of the databases 5 should fail, the director 8 reroutes transactions to one or more of the remaining databases. Whenever feasible, the director 8 also attempts to re-execute the transactions to minimize in flight transaction loss. Accordingly, the director 8 has the ability to instruct the controller 2 as to which database 5 is the primary database for satisfying the request of the respective application 7.

[0047] Depending on a database's workload and the relative power settings of the database servers 4 in a replication group, the director 8 routes read transactions to the least busy database server 4 for processing. This also applies when a database server 4 failure has resulted in transaction redirection.

[0048] Similarly, if the replication of transactions to a database server 4 becomes too slow for any reason such that the transactions start to build up and spill over to the replication disk queue 16, the director 8 redirects all the read transactions to the least busy database server 4. Once the disk queue becomes empty, the director 8 subsequently allows read access to that database. Accordingly, the fill/usage status of the replication disk queues in the replication group can be obtained or otherwise received by the director 8 for use in management of through-put rate of transactions applied to the respective databases 5.

[0049] For example, when the director 8 or replication servers 3 fails to communicate with the database servers 4, they report the failure to the controller 2 which then may redistribute transactions or take other appropriate actions to allow continuous operation of the transaction replicator 1. When one of the database servers 4 cannot be accessed, the controller 2 instructs the replication server 3 to stop applying transactions to it and relays the database lock down status to a monitor 6. The transactions start to accumulate within the queues until the database server 3 is repaired and the administrator or an automated process

instructs to resume replication via the monitor 6. The monitor 6 may also provide other predetermined administrative commands (for example: create database alias, update parameters, changing workload balancing setting).

Monitor

[0050] Referring again to Figure 1A, the monitor 6 allows a user to view and monitor the status of the controllers 2, the replication servers 3, and the databases 5. Preferably, the monitor 6 is a web application that is installed on an application or application server 7 and on the same network as the controllers 2.

[0051] Referring to Figure 3, shown is a diagrammatic view of the system monitor 6 for use with the transaction replicator 1. The system monitor 6 receives input data 32 from both primary and backup controllers 2, 9 (where applicable), replication servers 3, the database servers 4 and relevant databases 5 within a replication group. This information is used to display an overall system status on a display screen 31.

[0052] For example, depending on whether the controller is functioning or a failure has occurred, the relevant status of the controller 2 is shown. Second, the status of each of the replication servers 3 within a desired replication group is shown. A detailed description of the transaction rate, the number of transactions within each replication queue 15, the number transactions within each replication disk queue 16 is further shown. The monitor 6 further receives data regarding the databases 5 and displays the status of each database 5 and the number of committed transactions.

[0053] The administrator can analyze the above information and choose to manually reroute the transactions. For example, when it is seen that there exists many transactions within the replication disk queue 16 of a particular replication server 3 or that the transaction rate of a replication server 3 is slow, the administrator may send output data in the form of a request 33 to distribute the transactions for a specified amount of time to a different database server within the replication group.

[0054] Referring to Figure 4, shown is a flow diagram overview of the method 100 for initializing and processing transactions according to the invention. The global TX sequencer 12 also referred to as the sequencer hereafter and as shown in Figure 1B, is the control logic of the transaction replicator 1.

[0055] When the controller 2 is started, it initializes itself by reading from configuration and property files the parameters to be used in the current session 101. The global TX Queue 13, indoubt TX queue 17 and resent TX queue 18 shown in Figure 1B, are created and emptied in preparation for use. Before accepting any new transactions, the sequencer 12 examines the global disk queue 14 to determine if any transactions are left behind from previous session. For example, if a transaction is found on the global disk queue 14, it implies at least one database in the cluster is out of synchronization with the others and the database must be applied with these transactions before it can be accessed by applications. Transactions on the global disk queue 14 are read into the global TX queue 13 in preparation for applying to the database(s) 5. The sequencer 12 then starts additional servers called replication servers 3 that create and manage the replication queues 15. After initialization is complete, the sequencer 12 is ready to accept transactions from the application servers 7.

[0056] The sequencer 12 examines the incoming transaction to determine whether it is a new transaction or one that has already been recorded in the global TX queue 102. For a new transaction, the sequencer 12 assigns a Transaction ID 103 and records the transaction together with this ID in the global TX queue 13. If the new transactions ID is generated as a result of lost ID 104, the transaction and the ID are stored in the resent TX queue 109 for use in identifying duplicated transactions. The sequencer 12 checks the usage of the global TX queue 105 to determine if the maximum number of transactions in memory has already been exceeded. The sequencer 12 stores the transaction ID in the global TX queue 13 if the memory is not full 106. Otherwise, the sequencer 12 stores the transaction ID in the global disk queue 107. The sequencer 12 then returns the ID to the application 108 and the sequencer 12 is ready to process another request from the application.

[0057] When a request from the application or application server 7, comes in with a transaction that has already obtained a transaction ID previously and recorded in the global TX queue 13, the sequencer 12 searches and retrieves the entry from either the global TX queue 13 or the disk queue 110. If this transaction has been committed to the database 111, the entry's transaction status is set to "committed" 112 by the sequencer 12, indicating that this transaction is ready for applying to the other databases 200. If the transaction has been rolled back 113, the entry's transaction status is marked "for deletion" 114 and as will be described, subsequent processing 200 deletes the entry from the global TX queue. If the

transaction failed with an indoubt status, the entry's transaction status is set to "indoubt" 115. An alert message is sent to indicate that database recovery may be required 116. Database access is suspended immediately 117 until the indoubt transaction is resolved manually 300 or automatically 400.

[0058] Referring to Figure 5, shown is a flow diagram of the method 200 for distributing transactions from the global TX queue 13 according to the invention. The global TX queue 13 is used to maintain the proper sequencing and states of all update transactions at commit time. To apply the committed transactions to the other databases, the replication queue 5 is created by the sequencer 12 for each destination database. The sequencer 12 moves committed transactions from the global TX queue to the replication queue based on the following two criteria: (1) a predetermined transaction queue threshold (Q threshold) and (2) a predetermined sleep time (transfer interval).

[0059] For a system with sustained workload, the Q Threshold is the sole determining criteria to move committed transactions to the replication queue 201. For a system with sporadic activities, both the Q Threshold and transfer interval are used to make the transfer decision 201, 213. Transactions are transferred in batches to reduce communication overhead. When one or both criteria are met, the sequencer 12 prepares a batch of transactions to be moved from the global TX queue 13 to the replication queue 202. If the batch contains transactions, the sequencer 12 removes all the rolled back transactions from it because they are not to be applied to the other databases 204. The remaining transactions in the batch are sent to the replication queue for processing 205. If the batch does not contain any transaction 203, the sequencer 12 searches the global TX queue for any unprocessed transactions (status is committing) 206. Since transactions are executed in a same order of occurrence, unprocessed transactions typically occur when a previous transaction has not completed, therefore delaying the processing of subsequent transactions. A transaction that is being committed and has not yet returned its completion status is called gating transaction. A transaction that is being committed and returns a status of unknown is called indoubt transaction. Both types of transactions will remain in the state of "committing" and block processing of subsequent committed transactions, resulting in the transaction batch being empty. The difference between a gating transaction and an indoubt transaction is that gating transaction is transient, meaning that it will eventually become committed, unless there is a system failure that causes it to remain in the "gating state" indefinitely. Therefore when the

sequencer 12 finds unprocessed transactions 207 it must differentiate the two types of “committing” transactions 208. For a gating transaction, the sequencer 12 sends out an alert 209 and enters the transaction recovery process 300. Otherwise, the sequencer 12 determines if the transaction is resent from the application 210, 211, and removes the resent transaction from the global TX queue 211. A resent transaction is a duplicated transaction in the global TX queue 13 and has not been moved to the replication queue 15. The sequencer 12 then enters into a sleep because there is no transaction to be processed at the time 214. The sleep process is executed in its own thread such that it does not stop 200 from being executed at any time. It is a second entry point into the global queue size check at 201. When the sleep time is up, the sequencer 12 creates the transaction batch 202 for transfer to the replication queue 203, 204, 205.

[0060] Referring to Figure 6, shown is a flow diagram illustrating the method 300 for providing manual recovery of transactions 116 as shown in Figure 100. There are two scenarios under which the sequencer 12 is unable to resolve gating transactions and indoubt transactions caused by certain types of failure and manual recovery may be needed. First, a gating transaction remains in the global TX queue 13 for an extended period of time, stopping all subsequent committed transactions from being applied to the other databases. Second, a transaction status is unknown after some system component failure. The sequencer 12 first identifies the transactions causing need resolution 301 and send out an alert 302. Then the transaction can be manually analyzed to determine whether the transaction has been committed or rolled back in the database 304 and whether any manual action needs to be taken. If the transaction is found to have been rolled back in the database, the transaction entry is deleted manually from the global TX queue 305. If the transaction has been committed to the database, it is manually marked “committed” 306. In both cases the replication process can resume without having to recover the database 500. If the transaction is flagged as indoubt in the database, it must be forced to commit or roll back at the database before performing 304, 305 and 306.

[0061] Referring again to Figure 6, the process 400 is entered when an indoubt transaction is detected 115 and automatic failover and recovery of a failed database is performed. Unlike gating transactions that may get resolved in the next moment, an indoubt transaction is permanent until the transaction is rolled back or committed by hand or by some heuristic rules supported by the database. If the resolution is done with heuristic rules, the

indoubt transaction will have been resolved as “committed” or “rolled back” and will not require database failover or recovery. Consequently the process 400 is only entered when an indoubt transaction cannot be heuristically resolved and an immediate database failover is desirable. Under the automatic recovery process, the database is marked as “needing recovery” 401, with an alert sent out 402 by the sequencer 12. To help prevent further transaction loss, the sequencer 12 stops the generation of new transaction ID 403 and moves the indoubt transactions to the indoubt TX queue 404. While the database is marked “needing recovery” the sequencer 12 replaces it with one of the available databases in the group 405 and enables the transaction ID generation 406 such that normal global TX queue processing can continue 200. The sequencer 12 then executes a user defined recovery procedure to recover the failed database 407. For example, if the database recovery fails, the recovery process is reentered 408, 407.

[0062] Referring to Figure 7, shown is a flow diagram illustrating the processing of committed transactions by the replication servers 3 and the management of transactions in the replication queue 15 according to the present invention. Replication queues 15 are managed by the replication servers 3 started by the sequencer 12. One of the replication servers 3 receives batches of transactions from the sequencer 12. The process 500 is entered if a new batch of committed transactions arrives or at any time when queued transactions are to be applied to the databases.

[0063] If the process is entered because of new transactions 501, the batch of transactions are stored in the replication queue in memory 508, 509, or in replication disk queue 511 if the memory queue is full. Replication disk queue capacity is determined by the amount of disk space available. If the disk is above a predetermined threshold or is full for example 510, an alert is sent 512 by the sequencer 12 and the database is marked unusable 513 because committed transactions cannot be queued up anymore.

[0064] If the process is entered in an attempt to apply transactions in the replication queue to the databases, the replication server first determines whether there is any unprocessed transaction in the replication queue in memory 502. If the memory queue is empty but unprocessed transactions are found in the replication disk queue 503, they are moved from the disk queue to the memory queue in batches for execution 504, 505. Upon successful execution of all the transactions in the batch they are removed from the replication queue by the replication server and another batch of transactions are processed 501. If there

are transactions in the replication disk queue 16, the processing continues until the disk queue is empty, at which time the replication server 3 waits for more transactions from the global TX queue 501. During execution of the transactions in the replication queue 15, error may occur and the execution must be retried until the maximum number of retries is exceeded 507, then an alert is sent 512 with the database marked unusable 513. However, even though a database is marked unusable, the system continues to serve the application requests. The marked database is inaccessible until the error condition is resolved. The replication server 3 stops when it is instructed by the sequencer during the apparatus shutdown process 118, 119 and 120 shown in Figure 4.

[0065] It will be evident to those skilled in the art that the system 10 and its corresponding components can take many forms, and that such forms are within the scope of the invention as claimed. For example, the transaction replicators 1 can be configured as a plurality of transaction replicators 1 in a replicator peer-to-peer (P2P) network, in which each database server 4 is assigned or otherwise coupled to at least one principal transaction replicator 1. The distributed nature of the replicator P2P network can increase robustness in case of failure by replicating data over multiple peers (i.e. transaction replicators 1), and by enabling peers to find/store the data of the transactions without relying on a centralized index server. In the latter case, there may be no single point of failure in the system 10 when using the replicator P2P network. For example, the application or application servers 7 can communicate with a selected one of the database servers 7, such that the replicator P2P network of transaction replicators 1 would communicate with one another for load balancing and/or failure mode purposes. One example would be one application server 7 sending the transaction request to one of the transaction replicators 1, which would then send the transaction request to another of the transaction replicators 1 of the replicator P2P network, which in turn would replicate and then communicate the replicated copies of the transactions to the respective database servers 4.

[0066] Further, it is recognized that the applications/ application servers 7 could be configured in an application P2P network such that two or more application computers could share their resources such as storage hard drives, CD-ROM drives, and printers. Resources would then accessible from every computer on the application P2P network. Because P2P computers have their own hard drives that are accessible by all computers, each computer can act as both a client and a server in the application P2P networks (e.g. both as an application 7

and as a database 4). P2P networks are typically used for connecting nodes via largely ad hoc connections. Such P2P networks are useful for many purposes, such as but not limited to sharing content files, containing audio, video, data or anything in digital format is very common, and realtime data, such as Telephony traffic, is also passed using P2P technology. The term "P2P network" can also mean grid computing. A pure P2P file transfer network does not have the notion of clients or servers, but only equal peer nodes that simultaneously function as both "clients" and "servers" to the other nodes on the network. This model of network arrangement differs from the client-server model where communication is usually to and from a central server or controller. It is recognized that there are three major types of P2P network, by way of example only, namely:

1) Pure P2P in which peers act as clients and server, there is no central server, and there is no central router;

2) Hybrid P2P which has a central server that keeps information on peers and responds to requests for that information, peers are responsible for hosting the information as the central server does not store files and for letting the central server know what files they want to share and for downloading its shareable resources to peers that request it, and route terminals are used as addresses which are referenced by a set of indices to obtain an absolute address; and

3) Mixed P2P which has both pure and hybrid characteristics.

Accordingly, it is recognized that in the application and replicator P2P networks the applications/ application servers 7 and the transaction replicators 1 can operate as both clients and servers, depending upon whether they are the originator or receiver of the transaction request respectively. Further, it is recognized that both the application and replicator P2P networks can be used in the system 10 alone or in combination, as desired.

[0067] In view of the above, the spirit and scope of the appended claims should not be limited to the examples or the description of the preferred versions contained herein.

WE CLAIM:

1. A system for receiving and tracking a plurality of transactions and distributing the transactions to at least two replication queues over a network, the system comprising:

a global queue for storing a number of the received transactions in a first predetermined order; and

a sequencer coupled to the global queue for creating a copy of each of the transactions for each of said at least two replication queues and for distributing in a second predetermined order each said copy to each of said at least two replication queues respectively, said copy containing one or more of the received transactions.

2. The system according to claim 1, wherein the predetermined orders are selected from the group comprising: the first predetermined order is the same as the second predetermined order; and the first predetermined order is different from the second predetermined order.

3. The system according to claim 2 in which the sequencer distributes each said copy at a predetermined time interval.

4. The system according to claim 2 in which the sequencer distributes each said copy when the number of the transactions within the global queue exceeds a predetermined value.

5. The system according to claim 2 in which the sequencer distributes each said copy upon the earlier of:

a predetermined time interval; and

the number of the transactions within the global queue exceeds a predetermined value.

6. The system according to claim 5 in which each of the transactions comprises an update transaction and a unique transaction id assigned by the sequencer.

7. The system according to claim 6 further comprising a global disk queue in communication with the global queue for receiving and storing the transactions when the global queue is above a global threshold.

8. The system according to claim 7 wherein each of said at least two replication queues have a corresponding replication disk queue for receiving and storing the transactions from the global queue when the corresponding replication queue is above a replication threshold.
9. The system according to claim 8 in which the global queue receives the transactions from the global disk queue and other than receives the transactions from said at least one application server when the global disk queue is other than empty.
10. The system according to claim 5 further comprising an indoubt transaction queue in communication with the sequencer for storing the transactions identified as having unknown status by a database server during system failures.
11. The system according to claim 6 wherein the update transaction comprises at least one of a read, insert, update or delete request for at least one database in communication with at least one of said at least two replication queues.
12. The system according to claim 6 further comprising a resent transaction queue for storing the transactions when the transactions repeated the request for the transaction id.
13. The system according to claim 2, wherein the global queue is configured for receipt of the received transactions from a network entity selected from the group comprising: an application; and an application server.
14. The system according to claim 2, wherein the global queue is a searchable first-in first-out pipe.
15. The system according to claim 14 further comprising the sequencer configured for assuring the order of transactions in the global queue remain consistent with their execution order at a database server coupled to at least one of the replication queues.
16. The system according to claim 14, wherein the global disk queue is configured for storing an indexed and randomly accessible data set.

17. The system according to claim 2, wherein the global queue and sequencer are hosted on a network entity selected from the group comprising: a central control server and a peer-to-peer node.

18. A system for receiving a plurality of transactions from at least one application server, distributing the transactions to at least two replication queues and applying the transactions to a plurality of databases comprising:

a director coupled to each of said at least one application server for capturing a plurality of database calls therefrom as the plurality of transactions; and

a controller for receiving each of the plurality of transactions, the controller configured for storing the transactions within a global queue in a predetermined order, for generating a copy of each said transaction for each of said at least two replication queues, and for transmitting in the predetermined order each said copy to each of said at least two replication queues respectively.

19. The system according to claim 18 further comprising at least two replication servers including said at least two replication queues wherein each of said at least two replication servers is coupled to each of the databases; wherein the director routes each of the transactions to one or more of the databases relative to the workload and transaction throughput.

20. The system according to claim 19 further comprising a backup controller for receiving the transactions from said at least one application server upon failure of the controller, the backup controller including a backup global queue wherein the backup global queue is substantially synchronized with the controller and the backup global queue is a copy of the global queue.

21. A method for receiving and tracking a plurality of transactions and distributing the transactions to at least two replication queues over a network, the method comprising:

storing a number of the received transactions in a first predetermined order in a global queue;

creating a copy of each of the transactions for each of said at least two replication queues; and

distributing in a second predetermined order each said copy to each of said at least two replication queues respectively, said copy containing one or more of the received transactions.

22. The method according to claim 21 wherein the step of distributing each said copy occurs at a predetermined time interval.

23. The method according to claim 21 wherein the step of distributing each said copy occurs when the number of the transactions within the global queue exceeds a predetermined number.

24. The method according to claim 21 wherein the step of distributing each said copy occurs upon the earlier of: a predetermined time interval; and the number of the transactions within the global queue exceeds a predetermined number.

25. The method according to claim 24, wherein each of the transactions comprises an update transaction and a unique transaction id assigned by the sequencer.

26. The method according to claim 24 further comprising the step of receiving and storing the transactions within a global disk queue when the global queue storage capacity reaches a global threshold.

27. The method according to claim 21 further comprising the steps of:

determining whether the global disk queue is other than empty; and

receiving the transaction from the global disk queue rather than receiving the transactions from said at least one application server when the global disk queue is other than empty.

28. The method according to claim 21 further comprising the step of storing the transactions within an indoubt transaction queue during system failures.

29. The method according to claim 25 wherein the update transaction comprises at least one of a read, insert, update or delete request for at least one database in communication with at least one of said at least two replication queues.

30. The method according to claim 24 further comprising the steps of:

determining when at least one of said at least two replication queues are above a replication threshold, each of said at least two replication queues having a corresponding replication disk queue;

storing a number of the transactions within said corresponding replication disk queue based upon the determination; and

sending an alert to notify when said at least two replication queues and said corresponding replication disk queue capacity reach a preselected threshold.

31. The method according to claim 30 further comprising the step of : redirecting the transactions to at least one of said at least two replication queues being below said preselected threshold, based on receiving the alert.

32. A system for receiving and tracking a plurality of transactions and distributing the transactions to at least two replication queues over a network, the system comprising:

means for storing a number of the received transactions in a first predetermined order; and

means for creating a copy of each of the transactions for each of said at least two replication queues and for distributing in a second predetermined order each said copy to each of said at least two replication queues respectively, said copy containing one or more of the received transactions.

33. The system of claim 1 further comprising a replication server configured for using at least two respective concurrent database connections between the first of said at least two replication queues and a secondary database for transmission of the copied transaction contents to the secondary database.

34. The system of claim 1 further comprising a replication server configured for using at least a first concurrent database connection between the first of said at least two replication queues and a secondary database and a second concurrent database connection between the first of said at least two replication queues and a tertiary database for transmission of the respective copied transaction contents to the secondary and tertiary databases.

35. The system of claim 33, wherein the replication server is further configured for selecting individual statements from at least one of the copied transactions in the first of said at least two replication queues and configured for using the at least two respective concurrent database connections between the first of said at least two replication queues and the secondary database for transmission of the individual statements to the secondary database.

36. The system of claim 34, wherein the replication server is further configured for selecting individual statements from at least one of the copied transactions in the first of said at least two replication queues and configured for using the first and second concurrent database connections for transmission of the respective individual statements to the secondary and tertiary databases.




Figure 1A





Figure 1B






Figure 1E




Figure 1F




Figure 2

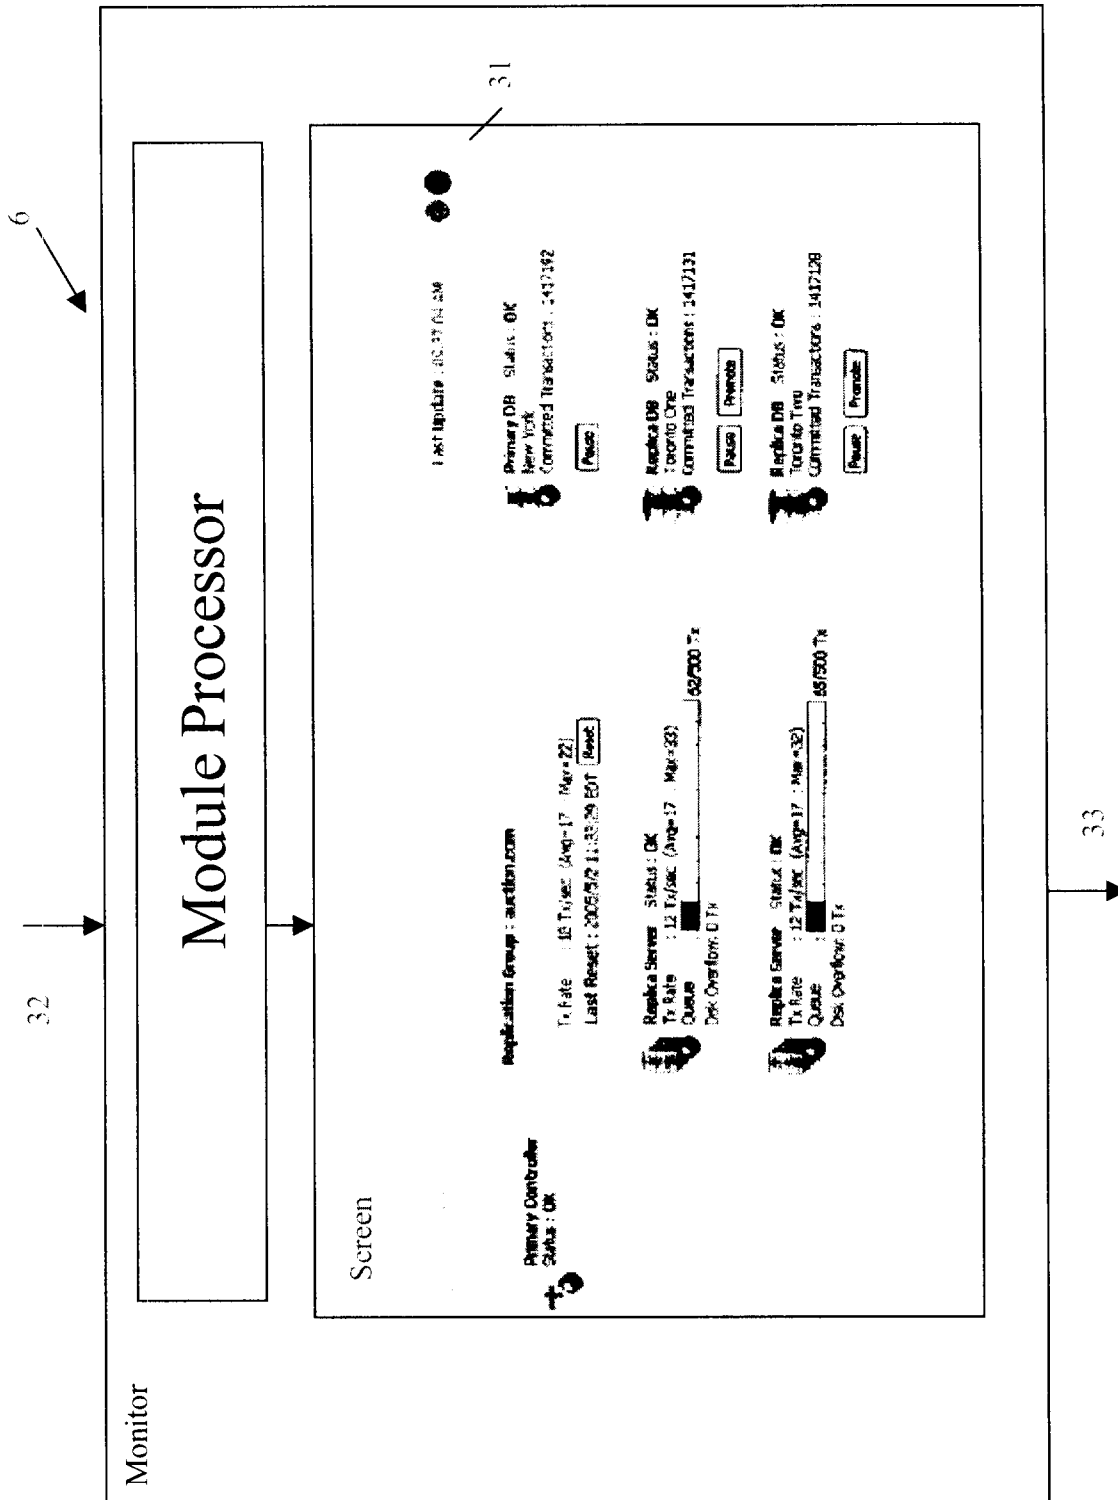


Figure 3

Figure 4

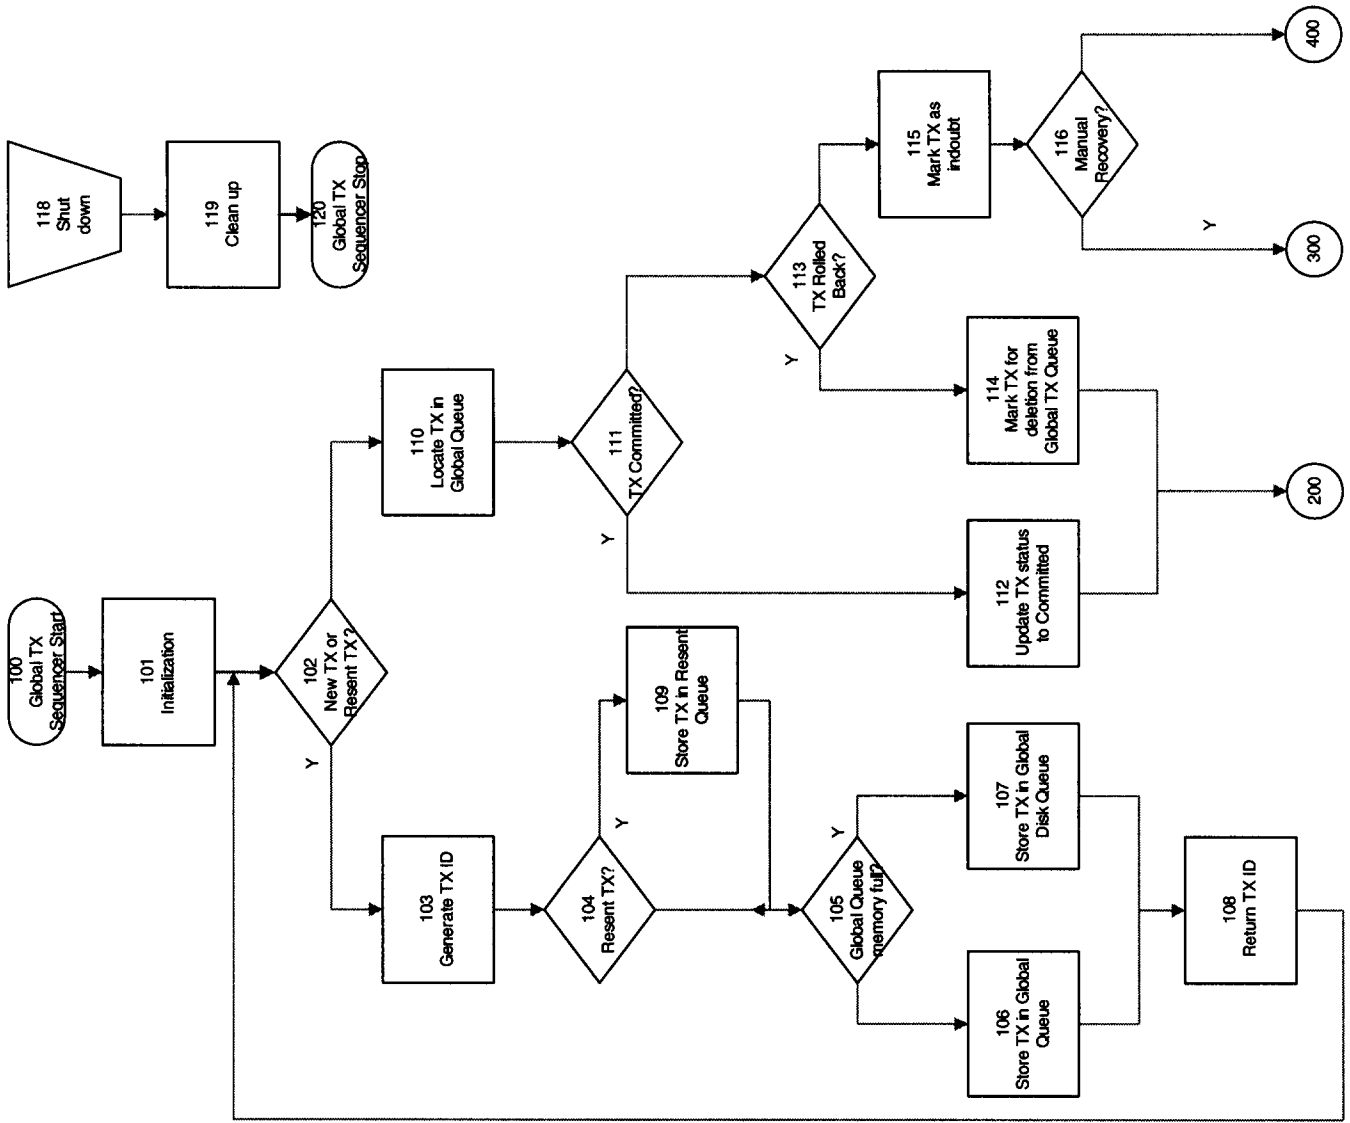
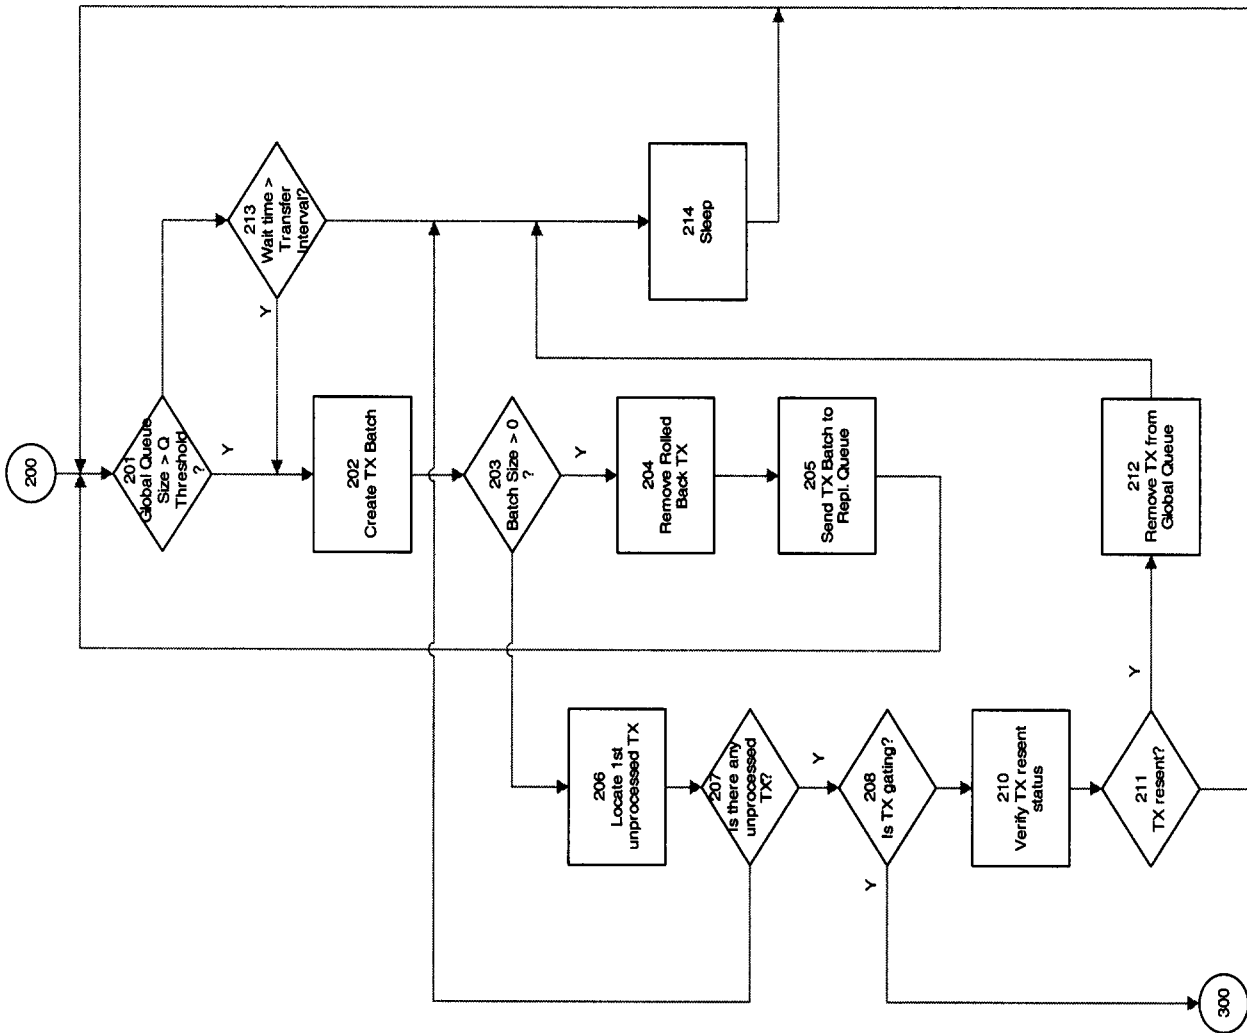


Figure 5



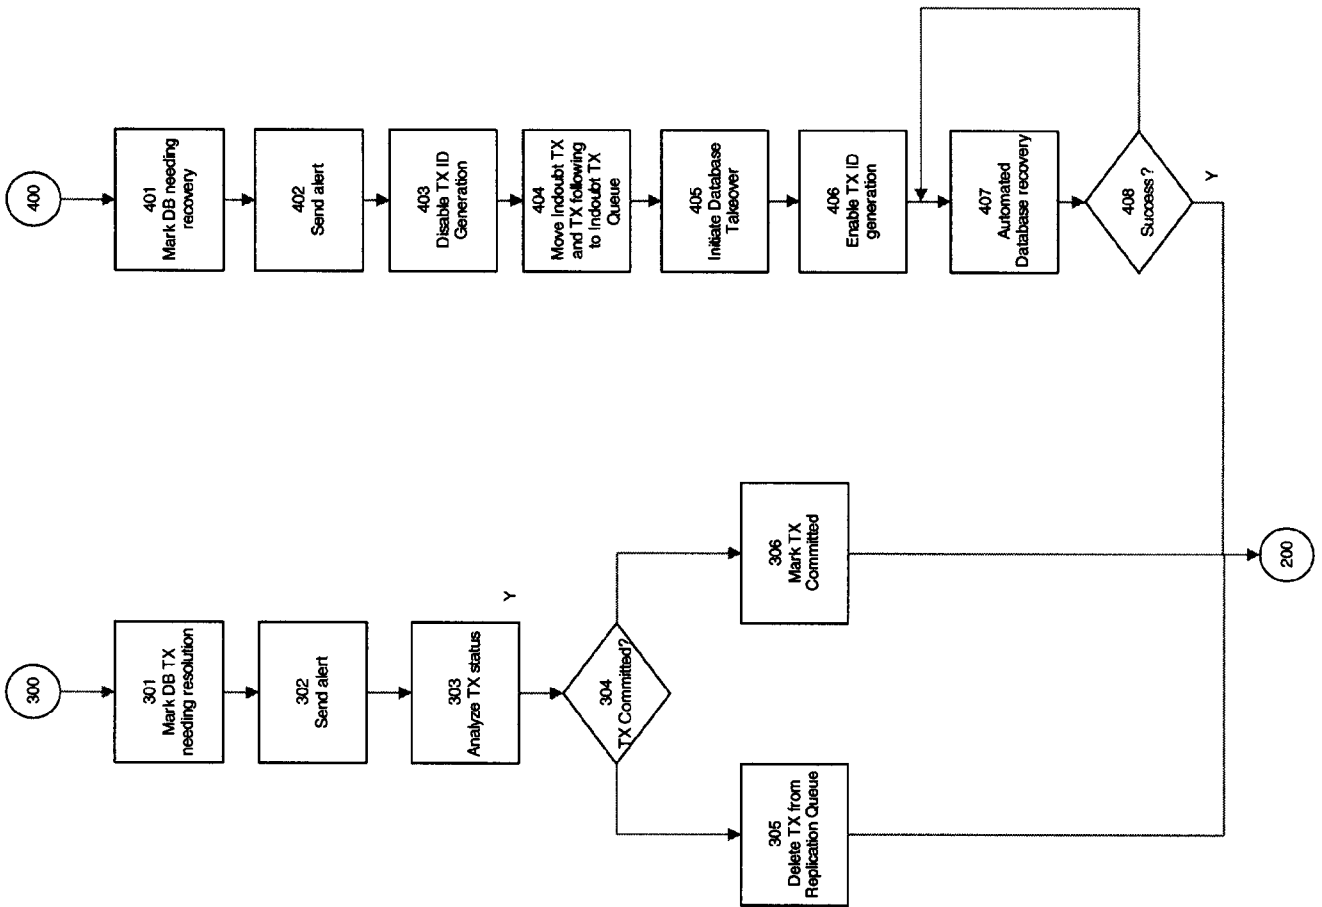
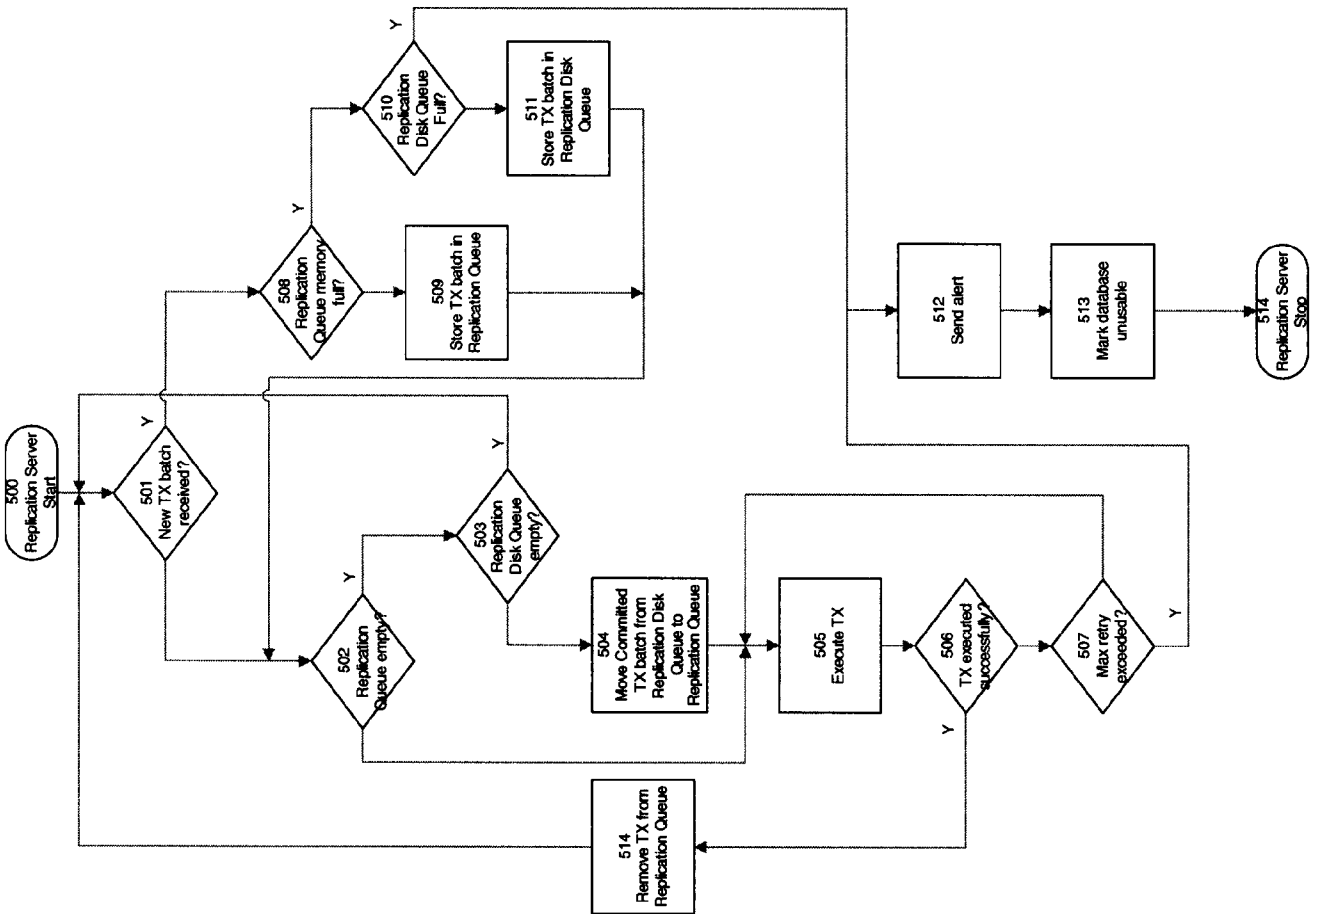


Figure 6

Figure 7



INTERNATIONAL SEARCH REPORT

International application No.
PCT/CA2006/001474

A. CLASSIFICATION OF SUBJECT MATTER IPC: G06F 17/00 (2006.01), G06F 17/30 (2006.01) According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) IPC: G06F 17/00 (2006.01), G06F 17/30 (2006.01) Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched Electronic database(s) consulted during the international search (name of database(s) and, where practicable, search terms used) USPTO, Delphion, Canadian Patent Database, Abstracts of Japan, Esp@cenet using keywords: receiving, tracking, distributing, transactions, replication, queue		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 20030212738 (Wookey et al.) 13 November 2003 (13-11-2003) (abstract, paragraph[0043]-[0331])	1-36
A	US 20020133491 (Sim et al.) 19 September 2002 (19-09-2002) (see whole document)	1-36
A	US 20010032282 (Marietta et al.) 18 October 2001 (18-10-2001) (abstract, paragraph[0011]-[0028])	1-36
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
* Special categories of cited documents :		
"A"	document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E"	earlier application or patent but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L"	document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O"	document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed	
Date of the actual completion of the international search 19 December 2006 (19-12-2006)		Date of mailing of the international search report 22 December 2006 (22-12-2006)
Name and mailing address of the ISA/CA Canadian Intellectual Property Office Place du Portage I, C114 - 1st Floor, Box PCT 50 Victoria Street Gatineau, Quebec K1A 0C9 Facsimile No.: 001(819)953-2476		Authorized officer Saadia Saifuddin 819-934-2671

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/CA2006/001474

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
US2003212738	13-11-2003	NONE	
US2002133491	19-09-2002	AT296514T T AU1465902 A AU2444802 A AU3643502 A CN1481635 A CN1481636 A DE60111072D D1 DE60111072T T2 EP1330907 A2 EP1364510 A2 EP1368948 A2 HK1054826 A1 US6857012 B2 US6970939 B2 US7047287 B2 US7058014 B2 US7076553 B2 US2002083118 A1 US2002083187 A1 US2003031176 A1 US2005198238 A1 WO0235359 A2 WO0235799 A2 WO0239323 A2	15-06-2005 06-05-2002 06-05-2002 21-05-2002 10-03-2004 10-03-2004 30-06-2005 26-01-2006 30-07-2003 26-11-2003 10-12-2003 06-01-2006 15-02-2005 29-11-2005 16-05-2006 06-06-2006 11-07-2006 27-06-2002 27-06-2002 13-02-2003 08-09-2005 02-05-2002 02-05-2002 16-05-2002
US2001032282		US6678773 B2 US6754752 B2 US6862283 B2 US7031258 B1 US7106742 B1	13-01-2004 22-06-2004 01-03-2005 18-04-2006 12-09-2006