

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
18 April 2002 (18.04.2002)

PCT

(10) International Publication Number  
**WO 02/31696 A1**

(51) International Patent Classification<sup>7</sup>: **G06F 17/30**

Scott [US/US]; 36 Waitohu Road, York Bay, Eastbourne, Wellington (US).

(21) International Application Number: PCT/NZ01/00206

(22) International Filing Date: 1 October 2001 (01.10.2001)

(74) Agents: **HAWKINS, Michael, Howard** et al.; Baldwin Shelston Waters, P.O. Box 852, Wellington (NZ).

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
507386 9 October 2000 (09.10.2000) NZ

(71) Applicant (for all designated States except US): **MAXIMUM AVAILABILITY LIMITED** [NZ/NZ]; 46 Mulgan Way, Browns Bay, Auckland (NZ).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AT (utility model), AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, CZ (utility model), DE, DE (utility model), DK, DK (utility model), DM, DZ, EC, EE, EE (utility model), ES, FI, FI (utility model), GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (utility model), SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

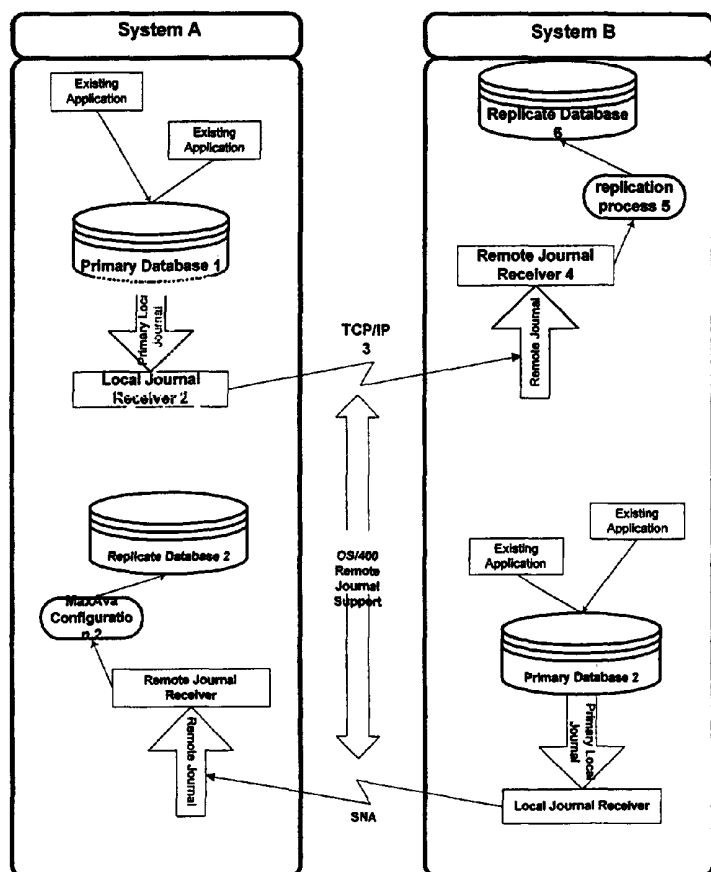
(72) Inventor; and

(75) Inventor/Applicant (for US only): **TARBELL, James,**

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR DATA PROCESSING



(57) Abstract: The invention relates to data processing methods and systems including: a method of database replication in which information strings are assigned to serialisation groups for processing. A method of memory management in which data is read from a storage space area whilst no data is written to it. A method of replicating a database in which a dynamic table is created to provided processing information for database members. A method of replicating a database wherein tasks are allocated to program components without program components interacting.



WO 02/31696 A1



patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

— with international search report

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

*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.*

## METHOD AND APPARATUS FOR DATA PROCESSING

### 5    Field of the Invention

The present invention relates to a method and apparatus for data processing. More particularly, but not exclusively, the invention relates to a method and apparatus for database replication.

10

### Background of the Invention

In a number of data processing applications fragments of data sent from a source system must be processed into a required data format on a target system.

15

In many instances it is desired to replicate a database on a target computer system from a database on a source system. This process may involve sending journal entries from the source database to allow updating of the target database. Databases may consist of one or more library, each of which contains one or more files, each file having one or more members. Each member consists of a table having one or more rows. A journal entry may contain an identifier of the library; file; file member and a row of changed data for the file member. This journal entry may be used by the target computer system to update its database.

20

25

It is important that database entries from a given table are updated in the correct sequence and that inter-related members are updated in the correct sequence. To ensure that journal entries are properly processed a receive process of the target computer system may compare an object name (library/file/member) with a database of objects stored on the target computer system. When a matching

30

object is located the processing information associated with that object may be used to process the journal entry.

The traditional approach has been to transfer journal entries, store  
5 them and replicate the database utilising a single engine. This approach is slow and complex.

It would be desirable for a database replication system to meet the following requirements:

10

1. Ensure that journal entries are serialised by database member (at a minimum), and by any user specified groupings.

15

2. Support an extremely large number of database apply processes so that database I/O (input/output) can be easily managed.

20

3. Process journal entries in a way which minimises the amount of system I/O (e.g. paging) between the time the entries are obtained from the journal and the time it is applied to the replica database.

25

4. The functions support any type of data packets, not just journal entries, to allow for future extensions to other types of replication (e.g. object, stream files etc).

5. The system hides the complexity of the memory management functions from other components.

30

It is an object of the present invention to provide a method and apparatus for information replication which meets these requirements or to at least provide the public with a useful choice.

Disclosure of the Invention

According to a first aspect of the invention there is provided a method  
5 of replicating information from a source system at a target system  
comprising the steps of:

- i) receiving information strings from the source system;  
and
- 10 ii) assigning the information strings to serialisation  
groups for processing such that inter-related information  
strings are processed in the same serialisation group.

15 The information strings may be journal entries from a source database  
which may be allocated to serialisation groups so that journal entries of  
the same type, or which are related to other journal entries, are  
processed in the same serialisation group.

20 According to a further aspect of the invention there is provided a  
method of managing memory space in a data transfer operation  
comprising the steps of:

- i/ defining a plurality of storage space areas;
- 25 ii/ writing data to a first storage space area; and
- iii/ reading data from the first storage space area whilst no data  
is written to the first storage space area.

30 The method may enable multiple simultaneous reads from other data  
storage space areas whilst information is written to only the first  
storage space area.

According to a further aspect of the invention there is provided a method of replicating a database from a source computer system at a target computer system comprising:

5

- i) receiving journal entries from the source computer system;
- ii) checking the journal entries to see if an entry exists in a  
10 dynamic index giving processing information relating to a database member to which the journal entry relates; and
- iii) if an entry exists in the dynamic table, processing the journal entry according to the associated processing  
15 information; or
- iv) if an entry does not exist in the dynamic index, looking up the related processing information for the database member in an assignment database, creating an entry  
20 and storing it in the dynamic index; and processing the journal entry according to the processing information.

The entry in the dynamic index may provide information as to whether the member needs to be processed with other members. A journal  
25 entry may be temporarily stored before being processed according to the processing information.

According to a further aspect of the invention there is provided a method of replicating a database from a source computer system to a  
30 target computer system comprising:

- i) receiving journal entries from the source computer system; and

- ii) allocating program components to process journal entries and update the target database, wherein a control program allocates tasks to program components and controls the program components substantially without program components interacting with one another.

Preferably the target computer system is a multi-processor computer system.

#### Brief Description of the drawings

The invention will now be described by way of example with reference to the accompanying drawings in which:

Figure 1: shows a schematic diagram of a source computer system which provides journal entries to a target computer system.

Figure 2: is a functional diagram illustrating the processes involved in database replication at a target computer system.

Figure 3: shows the mapping of storage space within the target computer system.

Figure 4: shows a flow diagram illustrating the process for allocating journal entries to serialisation groups.

#### Detailed Description of the Preferred Embodiment

The following description describes a database replication method where the source and target computer systems are IBM AS/400

computers operating under the OS/400 operating system. It will be appreciated that the method is applicable to other systems with appropriate modification.

5 Referring to figure 1, source system A contains a primary database 1. Primary database 1 may contain one or more library. Each library may contain one or more file. Each file may contain one or more members. Each member comprises a table having one or more rows. A unique library/file/member combination is referred to as an object.

10

When a row of any member of primary database 1 is modified a journal entry including the object name and the modified row is sent to local journal receiver 2. Local journal receiver 2 sends the journal entry via communications link 3 to a remote journal receiver 4 of a target

15 computer system B. A database replication process 5 receives the journal entries and modifies the contents of replica database 6 to maintain it in conformity with primary database 1.

Referring now to figure 2 the process and apparatus for replicating  
20 target database 6 of the target computer system will be described. To ensure proper replication of replica database 6, database members are updated in the replica database 6 in the same order as they are modified in the primary database 1. To achieve this a number of serialisation groups 8 are defined. Journal entries having the same  
25 object name are grouped into a common serialisation group so that they are updated in the correct order. Certain database members may have relationships with other database members (joins etc) and so may be assigned to a common serialisation group to ensure all inter-related members are updated in the correct sequence. A serialisation group  
30 may thus contain journal entries for a number of objects. The use of such serialisation groups enables database replication to be conducted in the appropriate sequence as well as facilitating efficient parallel processing.



Receive process 7 may either assign a received journal entry to a serialisation group, assign a journal entry to a default serialisation group or discard the journal entry. Serialisation group assignment is performed based upon an assignment database (MXSGMBAS) and a temporary OS/400 user index object. The journal entry assignment functions are provided via an ILE service program – which allows the underlying implementation to be modified without recompile/bind of the calling functions.

10

The assignment database MXSGMBAS contains all objects, their relationship with other objects (i.e. do they need to be grouped with other objects during processing) and their required manner of processing. Assignment of a journal entry to a serialisation group 8 could be conducted simply by comparing the object name of each received journal entry with the assignment database MXSGMBAS and assigning the journal entry to a serialisation group based upon the associated information. However, the assignment database MXSGMBAS contains many objects and considerable processing time is required to perform a database locate operation and extract the relevant processing information. According to the invention a member assignment (MBIX) index temporary object is used to store processing information for an object. This is an index of objects giving their associated serialisation group and related processing information (including a link to their associated control structures).

25

Referring now to figures 2 and 4 the serialisation group assignment will be described. When a journal entry is received in step 11 receive process 7 conducts a comparison in step 12 to see whether the object is present in the MBIX index. If so, operation proceeds to step 13 and a serialisation group number and database file index (DBFIDX) is returned and processing continues within the assigned serialisation group.

30

If the object name is not stored in the MBIX index then a full object name lookup is conducted in the MXSGMBAS database 9 in step 14.

If the lookup is successful then a serialisation group is returned, a

- 5 Database File Index (DBFIDX) is assigned which will point to the processing information stored in a dynamic array maintained by the associated serialisation group and an entry is added to the MBIX index in step 15. Each Database File Index (DBFIDX) is created simply by incrementing an index that is unique by serialisation group.

10

If a match is not achieved in step 14 then a generic name lookup is conducted in step 16. This involves a search by a library/file /\*all and then by library/\*all/\*all. If a generic match is achieved the full name is added to the MBIX table in step 17 and processing continues in steps

- 15 15 and 13 as before. If no match can be achieved the journal entry is discarded in step 18.

Accordingly, at startup, there will be no entries in the MBIX index 10.

As journal entries are processed, serialisation groups and the

- 20 processing information for objects will be added to MBIX index 10.

The serialisation group and processing information may be much more rapidly obtained from MBIX table 10 than from MXSGMBAS database 9

- 25 This method gives the following, significant, performance benefits:

1. The serialisation groups do not need to search for a member's related processing information. They simply maintain the processing information in a dynamic array with the Database
- 30 File Index as the means of access.

2. All operations relating to a particular member name may refer to the serialisation group and Database File Index value to uniquely identify the member (a "handle").

- 5 Referring now to figure 3 the method of memory management within the target computer system will be described. Storage object space is divided up into a number of storage units  $SU_1$ - $SU_N$ . Each storage unit has a storage unit header 20. The storage unit header 20 gives the number of serialisation groups which have journal entries in the storage  
10 unit. Each data segment consists of a storage entry header 21 and a storage entry 22. Storage entries are aligned on 16 byte boundaries with padding blocks 23 filling any space between an entry and a 16 byte boundary.
- 15 Journal entries are passed on from receive process 7 for storage in a storage object space 24. The journal entries from receive process 7 are stored in storage space object 24 in blocks 22. Each journal entry 22 has an associated storage entry header 21 (or handle) which contains information as to the displacement to the next journal entry in  
20 the storage unit for that serialisation group and an associated Database File Index (DBFIDX) containing the processing information for the member associated with the journal entry. The processing information is maintained in dynamic memory with the Database File Index as the means of access.
- 25 In normal operation journal entries are consecutively written to one storage unit until it is filled and then journal entries are written to the next available storage unit. Once writing to a storage unit has been completed journal entries may be read from the populated storage unit.
- 30 Partially filled storage units may be read out when system resources are not being otherwise utilised (i.e. no incoming journal entries need to be stored).

This approach means that memory locks are not required during reading and writing. During the writing process the receive process 7 has exclusive access to write to a storage unit. No locks are required during read operations and so journal entries may be simultaneously  
5 read to their associated serialisation group. The only locking required is to decrement the value held in storage unit header 20 when the last journal entry for a serialisation group is read out.

The available storage units queue (ASUQ) 25 controls the order in  
10 which free storage units are utilised. ASUQ 25 includes a last in first out (LIFO) buffer which stores addresses of free storage units. Journal entries of a serialisation group are read out of a storage unit until a null value is found in a storage entry header. As each storage entry 22 is read out the storage unit header 20 is decremented. When all journal  
15 entries are read out completely from a storage unit the storage unit header 20 will be decremented to zero and the storage unit number is returned to the ASUQ and is the first storage unit re-assigned when new journal entries must be written into storage space. In this way the most recently used storage units are maintained active to reduce  
20 the working set of storage units to a minimum.

When all journal entries in a storage unit have been read and the storage unit is released the entire address range of the storage unit may be purged without requiring writing of data to auxiliary storage.

25 Referring again to Figure 2 the manner of processing will be further described. Control process 19 oversees the replication process and controls processing in the receive process 7 and within the serialisation groups 8. In this manner processing can be conducted within each  
30 serialisation group without regard to processing within another serialisation group. By having the whole process controlled by an overarching control process 19 each serialisation group can conduct its

processing in isolation without regard to the complexity of the overall operation.

As each serialisation group receives journal entries for a member in  
5 sequence the updating of that member in the replica database 6 is sequential also. By processing linked members in a particular serialisation group processing is streamlined.

When a replica database 6 is to be made a primary database partially  
10 applied commits must be removed. Firstly, the control process 19 suspends receive process 7 and processing by serialisation groups 8. Control process 19 then identifies all "open" commit groups (e.g. commit IDs that have not yet received a commit or roll back journal entry). These are processed, serially, from the most recent (i.e. the  
15 commit group that has the most recent journal entry) to the oldest as follows:

- i) a receive process of receive process 7 receives the commit group's journal entries from journal receiver 26;
- 20 ii) all entries are assigned to a "default" serialisation group;
- iii) the entries are stored in storage unit 24 in the usual manner but are linked in reverse order (i.e. the head of the list is the last entry in the storage unit, with links moving backward until the first entry in the storage unit);
- 25 iv) if a storage unit is filled before that commit group's entries are complete, the storage unit is pushed onto LIFO queue TLQ 27 (instead of releasing it to the default serialisation group). Then a new storage unit is allocated (as normal) and entries continue to be stored;
- 30 v) when the commit group's available journal entries are completely received and stored in storage units, the storage

units are dispatched to the default serialisation groups in LIFO order. The result being that the serialisation group receives the journal entries in reverse order (from most recent to oldest);

- vi) the default serialisation group processes the entries as  
5 "reverse" entries (the entries include a flag to indicate that they are "reverse" entries). This results in all inserts being processed as deletes, updates being removed to their prior image and deletes being inserted etc. Only journal entries which had  
10 already been applied (e.g. during normal processing) to the database are processed;
- vii) the default serialisation group does not perform a commit on  
the "reverse" entries until it receives the "data commit group"  
journal entry. This ensures that if a failure is encountered  
during the "clean-up" the database is in a known state. This  
15 enables the "clean-up" to be restarted.

Once all of the "open" commit groups have been "removed" the control process 19 suspends the other processes and the replica database is ready to be used as the primary database.

20

This method allows rapid "clean-up" of partially applied commits which does not require processing capability of the system to be utilised unless a secondary database does in fact have to be made a primary database.

25

The method and apparatus of the invention provide a number of advantages as follows:

- 1. The allocation of storage unit blocks within a storage space  
30 object and control of read/writes avoids the need for locks and read/write concurrency issues.

2. The use of serialisation groups enables members to be updated in a serial manner and for inter-related members to be updated in correct chronology. Serialisation groups enable multiple streams of journal entries to be simultaneously processed whilst  
5 processing interrelated members together.
3. The use of the MBIX index greatly reduces lookup time for each journal entry. The use of storage entry headers 21 (handles)  
10 enables the next journal entry of a serialisation group to be located rapidly.
4. The use of a control process to oversee the operation of the receive process and processing within serialisation groups enables the sub-processes to process information efficiently  
15 without the need to interact with other processes.
5. Simple handling of commits where secondary database is to be made primary database.

20 Where in the foregoing description reference has been made to integers or components having known equivalents then such equivalents are herein incorporated as if individually set forth.

25 Although this invention has been described by way of example it is to be appreciated that improvements and/or modifications may be made thereto without departing from the scope or spirit of the present invention.

## CLAIMS:

1. A method of replicating information from a source system at a target system comprising the steps of:
  - 5 i) receiving information strings from the source system; and
  - ii) assigning the information strings to serialisation groups for processing such that interrelated information strings are processed in the same serialisation group.
- 10 2. A method as claimed in claim 1 wherein the information strings are journal entries from a source database.
- 15 3. A method as claimed in claim 2 wherein the journal entries are allocated to serialisation groups such that journal entries of the same type or related to other journal entries are processed in the same serialisation group.
- 20 4. A method as claimed in claim 2 or claim 3 wherein, when each database object is first encountered in a processing operation a temporary index object is generated containing the serialisation group and related processing information for that database object.
- 25 5. A method as claimed in claim 4 wherein, if a temporary index object already exists the serialisation group for a database object is assigned based upon the temporary index object data.
- 30 6. A method as claimed in claim 4 wherein if a temporary index object does not exist a full object name look-up of an assignment database is conducted, the serialisation group is obtained from the assignment database and a temporary index object is generated.



7. A method as claimed in any one of the proceeding claims wherein processing information for a database object is maintained in dynamic memory.
- 5
8. A method of managing memory space in a data transfer operation comprising the steps of:
- i) defining a plurality of storage space areas;
  - ii) writing data to a first storage space area; and
  - 10      iii) reading data from the first storage space area whilst no data is written to the first storage space area.
9. A method as claimed in claim 8 wherein the data is database journal entries.
- 15
10. A method as claimed in claim 9 wherein the journal entries are consecutively written to the first storage space area until the first storage space area is at least substantially filled and then data is written to the next storage space area.
- 20
11. A method as claimed in claim 9 or claim 10 wherein each journal entry has an associated header that records the displacement to the next journal entry having the same serialisation group.
- 25
12. A method as claimed in any one of claims 8 to 11 wherein data is written to a second storage space whilst data is read from the first storage space.
- 30
13. A method as claimed in any one of claims 8 to 12 wherein each storage space has a storage space header which records the

number of serialisation groups having journal entries within that storage space.

14. A method as claimed in claim 13 wherein when all data is read out of a storage space, the storage space header is  
5 decremented and the storage space is made available for data to be written to.

15. A method of replicating a database from a source computer system at a target computer system comprising:  
10 i) receiving journal entries from the source computer system;  
ii) checking the journal entries to see if any entry exists in a dynamic index giving processing information relating to a database member to  
15 which the journal entry relates; and  
iii) if an entry exists in the dynamic table, processing the journal entry according to the associated processing information; or  
iv) if an entry does not exist in the dynamic index,  
20 looking up the related processing information for the member in an assignment database; creating an entry and storing it in the dynamic index; and processing the journal entry according to the processing information.

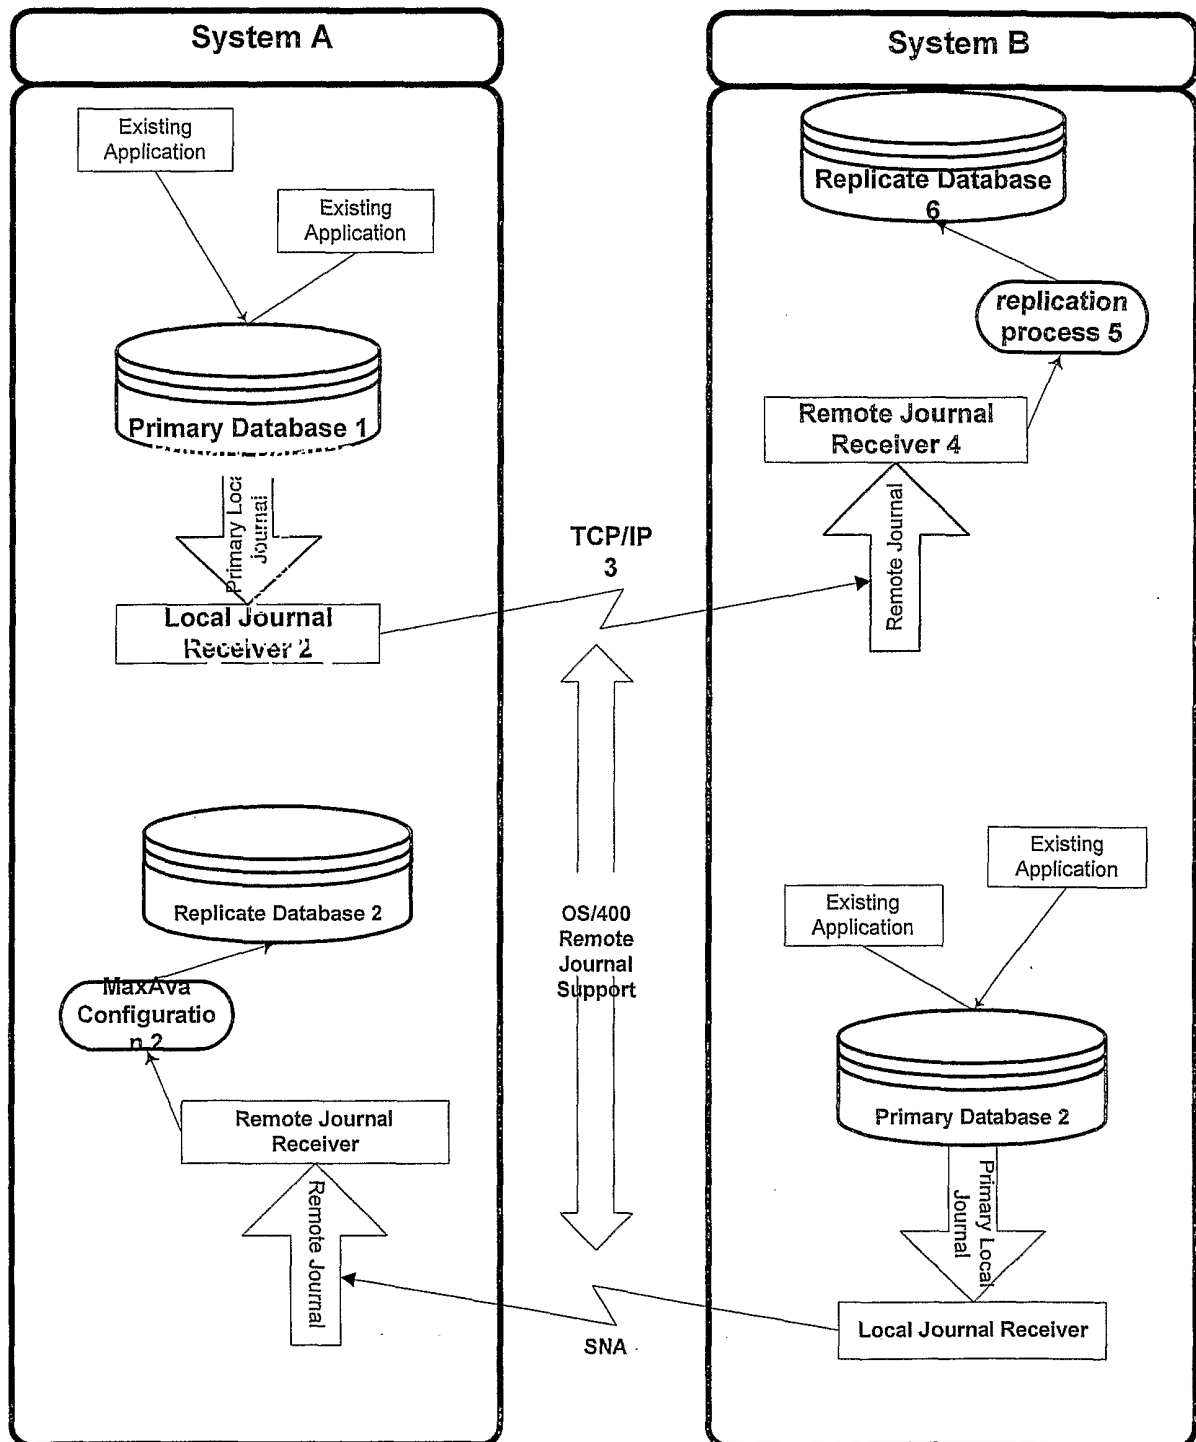
25 16. A method as claimed in claim 15 wherein the dynamic index includes information as to whether a database member needs to be processed with one or more other database member.

30 17. A method as claimed in claim 15 or claim 16 wherein journal entries are temporarily stored before the replication process is performed.

18. A method as claimed in any one of claims 15 to 17 wherein the memory management method of any one of claims 8 to 14 is employed.
- 5 19. A method as claimed in any one of claims 15 to 18 wherein the replication method of any one of claims 1 to 7 is employed.
20. A method of replicating a database from a source computer system to a target computer system comprising:
- 10 i) receiving journal entries from the source computer system; and
- ii) allocating program components to process journal entries and update the target database, wherein a control program allocates tasks to program
- 15 components and controls the program components substantially without program components interacting with one another.
21. A method as claimed in claim 20 wherein the target computer
- 20 system implements parallel processing.
22. A method as claimed in claim 20 or claim 21 wherein the target computer system is a multi-processor computer.
- 25 23. A computer system programmed to operate according to the method of any one of the preceding claims.
24. A computer program adapted to carry out the method of any one of claims 1 to 22.
- 30 25. A computer readable medium containing a computer program as claimed in claim 24.

26. A database replicated by the method of any one of claims 1 to 22.

1/4

**FIGURE 1**




Figure 2

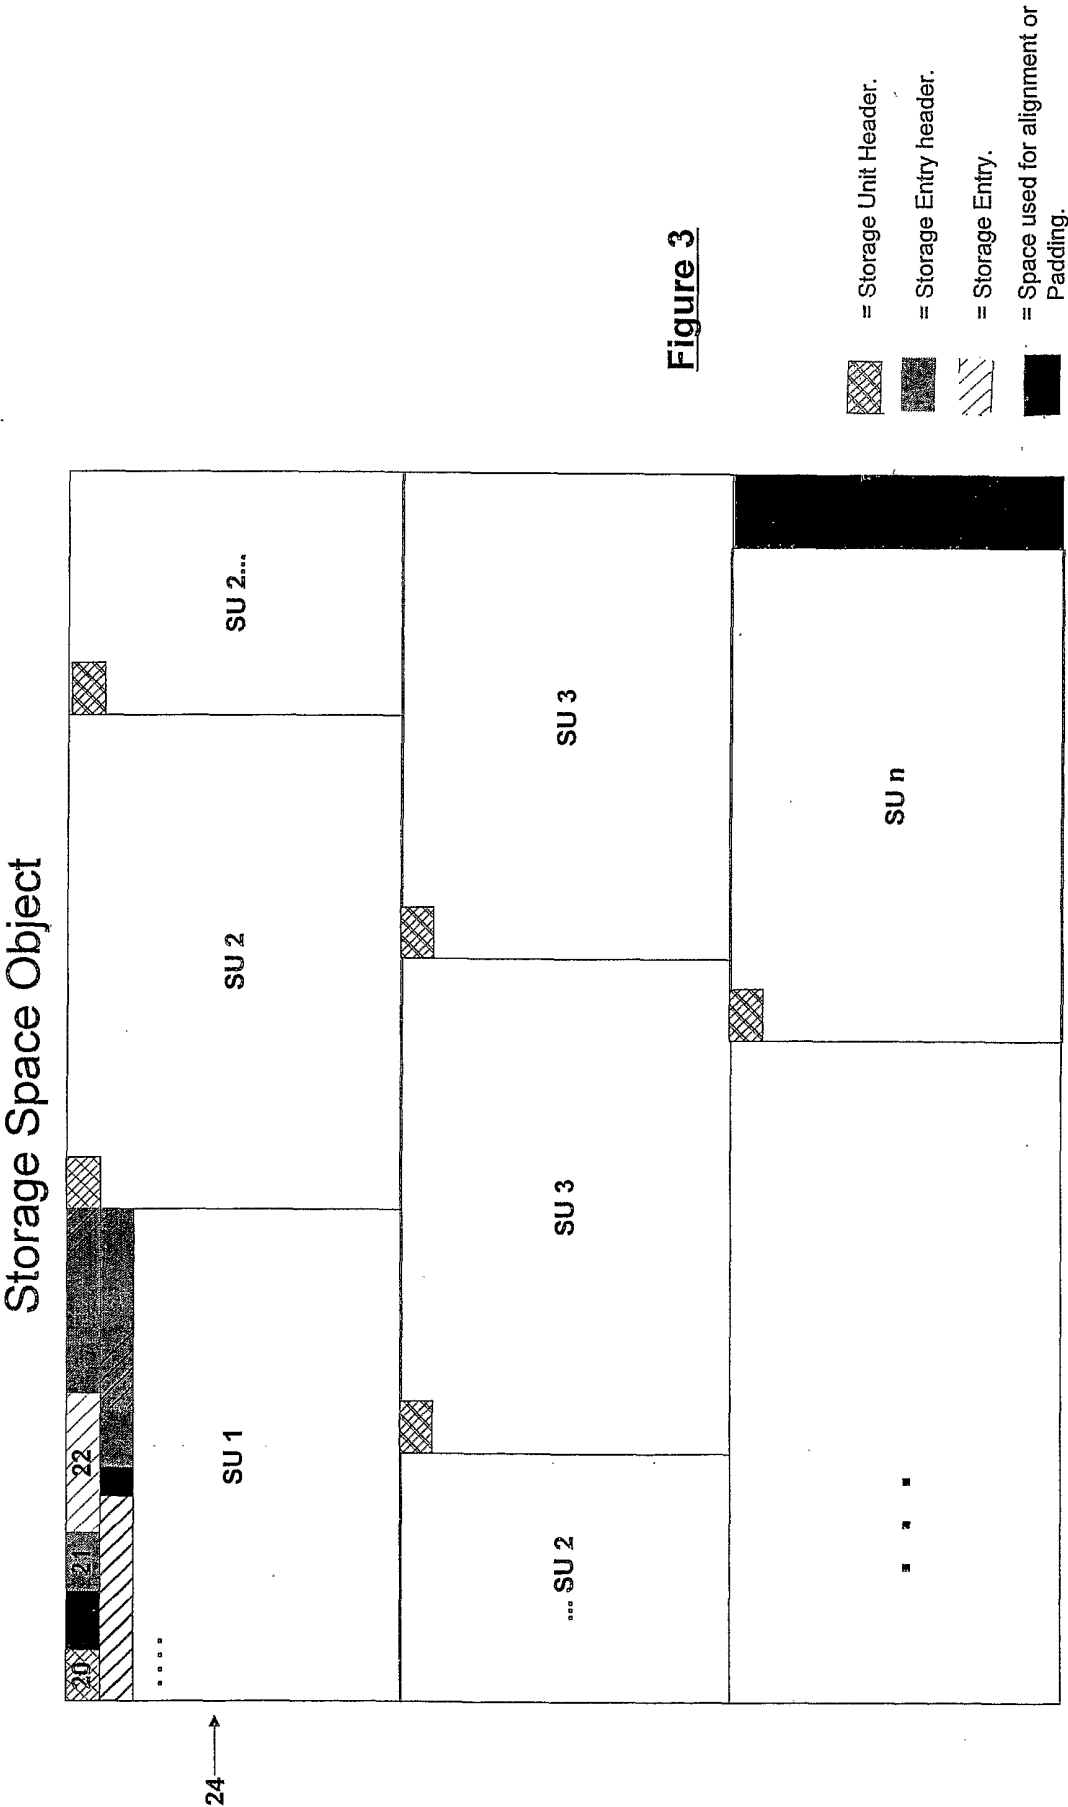
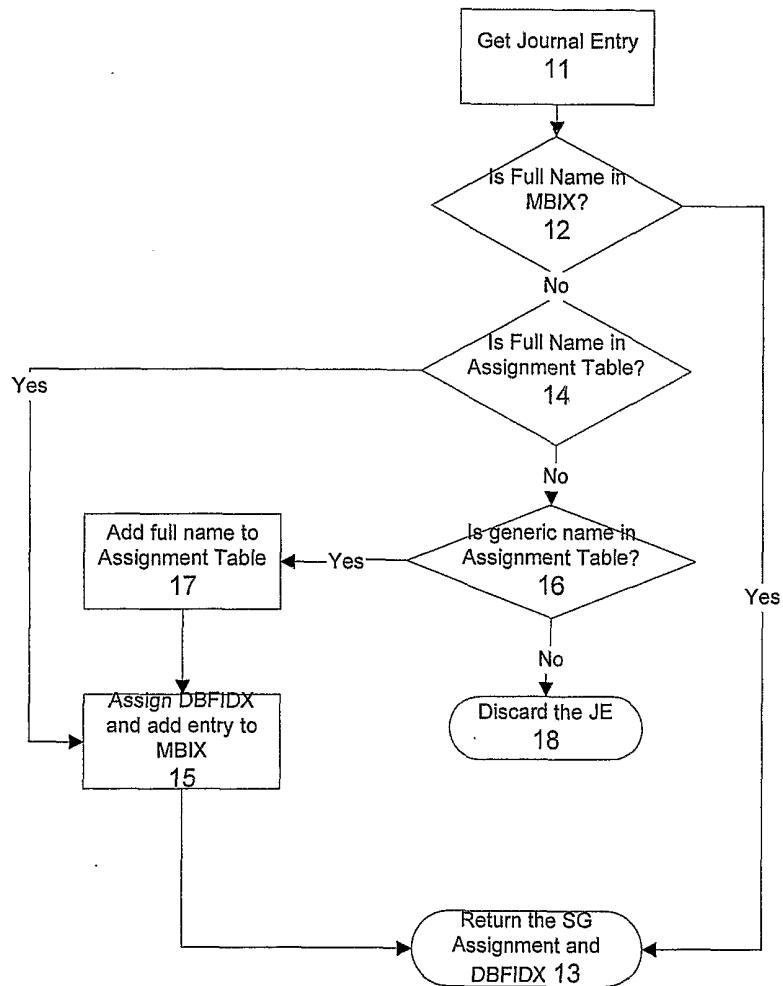


Figure 3

4/4

**Figure 4**



## INTERNATIONAL SEARCH REPORT

International application No.

PCT/NZ01/00206

<b>A. CLASSIFICATION OF SUBJECT MATTER</b>		
Int. Cl. <sup>7</sup> : G06F 17/30		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b>		
Minimum documentation searched (classification system followed by classification symbols)		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)		
KEYWORDS (WPAT, USPTO, INTERNET): REPLICAT+, DUPLICAT+, CLONE, MIRROR, SERIALI+, JOURNAL, STRING, GROUP, SOURCE, TARGET, INDEX, UPDAT+, MULTI+, +THREAD+, +TASK+,...		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	GB 2330220 A (INTERNATIONAL BUSINESS MACHINES CORPORATION) 14 April 1999 Page 3 line 10-20	8-9
P,A	US 6243715 B1 (BOGANTZ et al) 5 June 2001 See whole document	
A	US 6049809 A (RAMAN et al) 11 April 2000 See whole document	
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C <input checked="" type="checkbox"/> See patent family annex		
<p>* Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> <p>"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</p> <p>"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</p> <p>"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</p> <p>"&amp;" document member of the same patent family</p>		
Date of the actual completion of the international search 1 February 2002		Date of mailing of the international search report 13 FEB 2002
Name and mailing address of the ISA/AU AUSTRALIAN PATENT OFFICE PO BOX 200, WODEN ACT 2606, AUSTRALIA E-mail address: pct@ipaustalia.gov.au Facsimile No. (02) 6285 3929		Authorized officer  Stephen Lee Telephone No : (02) 6283 2205

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/NZ01/00206

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5950198 A (FALLS et al) 7 September 1999 See whole document	
A	GB 2327781 A (INTERNATIONAL BUSINESS MACHINES CORPORATION) 3 February 1999 See whole document	
A	US 5864851 A (BREITBART et al) 26 January 1999 See whole document	
A	An Efficient Scheme for Dynamic Data Replication (Swarup Acharya et al) September 1993, Tech Report CS-93-43 ( <a href="http://citeseer.nj.nec.com/cache/papers/cs/11460/ftp:zSzzSzftp.cs.brown.edu/zSpubzSztechreportszSz93zSzcs93-43.pdf/acharya93efficient.pdf">http://citeseer.nj.nec.com/cache/papers/cs/11460/ftp:zSzzSzftp.cs.brown.edu/zSpubzSztechreportszSz93zSzcs93-43.pdf/acharya93efficient.pdf</a> )	
A	Data Replication in Mariposa (Jeff Sidell et al) Sequoia 2000 Technical Report 95-60, University of California, Berkeley ( <a href="http://db.cs.berkeley.edu/papers/S2K-95-62.pdf">http://db.cs.berkeley.edu/papers/S2K-95-62.pdf</a> )	

**Box I Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)**

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos :  
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos :  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos :  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a)

**Box II Observations where unity of invention is lacking (Continuation of item 3 of first sheet)**

This International Searching Authority found multiple inventions in this international application, as follows:

See Extra sheet

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

**Remark on Protest** ☐ The additional search fees were accompanied by the applicant's protest.

☒ No protest accompanied the payment of additional search fees.

**Supplemental Box**

(To be used when the space in any of Boxes I to VIII is not sufficient)

**From Box II**

The international application does not comply with the requirements of unity of invention because it does not relate to one invention or to a group of inventions so linked as to form a single general inventive concept. In coming to this conclusion the International Searching Authority has found that there are four inventions:

The invention of independent claim 1 is for a method of database replication in which information strings are assigned to serialisation groups for processing.

It is considered that these features comprise a first "special technical feature".

The invention of independent claim 8 is for a method of memory management in which data is read from a storage space area whilst no data is written to it.

It is considered that these features comprise a second separate "special technical feature".

The invention of independent claim 15 is for a method of replicating a database in which a dynamic table is created to provided processing information for database members.

It is considered that these features comprise a third separate "special technical feature".

The invention of independent claim 20 is for a method of replicating a database wherein tasks are allocated to program components without program components interacting.

It is considered that these features comprise a fourth separate "special technical feature".

Since the above mentioned groups of claims do not share any special technical features identified, a "technical relationship" between the inventions, as defined in PCT rule 13.2 does not exist. Accordingly the international application does not relate to one invention or to a single inventive concept.

**INTERNATIONAL SEARCH REPORT**  
Information on patent family members

International application No.  
**PCT/NZ01/00206**

This Annex lists the known "A" publication level patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent Document Cited in Search Report		Patent Family Member	
GB	2330220	US	6226641
US	6243715	NONE	
US	6049809	NONE	
US	5950198	NONE	
GB	2327781	US	6253211
US	5864851	NONE	
END OF ANNEX			