




(22) 1995/06/01

(43) 1996/12/02

(45) 2001/05/01

(72) BARU, Chaitanya, CA

(72) FECTEAU, Gilles, CA

(72) KIRTON, Jo-Anne, CA

(72) KOLLAR, Lubor J., CA


(72) KOO, Fred, CA

(73) IBM CANADA LTD. - IBM CANADA LIMITÉE, CA

(51) Int.Cl.⁶ G06F 17/30

(54) **METHODE ET DISPOSITIF DE DEGROUPE PARTIEL DANS
UN SYSTEME A BASES DE DONNEES PARALLELES**

(54) **METHOD AND APPARATUS FOR IMPLEMENTING PARTIAL
DECLUSTERING IN A PARALLEL DATABASE SYSTEM**



(57) The invention provides a method and apparatus for distributing data of a table across a parallel data base system having a plurality of interlinked database nodes. The table is partitioned in accordance with a predetermined partitioning arrangement and data of the table is distributed across a subset of the nodes in accordance with the partitioning arrangement. The partitioning arrangement can comprise horizontal or vertical partitioning. The partitioning arrangement can comprise a partitioning map defining the relationship of partitions of the table to said subset of nodes.

CA9-95-001

METHOD AND APPARATUS FOR IMPLEMENTING PARTIAL
DECLUSTERING IN A PARALLEL DATABASE SYSTEM

ABSTRACT

5

10

15

The invention provides a method and apparatus for distributing data of a table across a parallel data base system having a plurality of interlinked database nodes. The table is partitioned in accordance with a predetermined partitioning arrangement and data of the table is distributed across a subset of the nodes in accordance with the partitioning arrangement. The partitioning arrangement can comprise horizontal or vertical partitioning. The partitioning arrangement can comprise a partitioning map defining the relationship of partitions of the table to said subset of nodes.

METHOD AND APPARATUS FOR IMPLEMENTING PARTIAL
DECLUSTERING IN A PARALLEL DATABASE SYSTEM

Field of the Invention

5 This invention relates generally to the storing of data in parallel database systems and more particularly to a method of distributing data in a table across a subset of the nodes of the parallel database system. The invention is particularly useful in relational database systems.

10 Background of the Invention

Employing static partitioning for data in a shared nothing parallel data processing system, one method of exploiting parallel processing is to partition database tables across the nodes
15 (typically containing one or more processors and associated storage) of the system. This is referred to as "declustering" of the table. If a database table is partitioned across only a subset of the nodes of the system then that table is said to be "partially declustered". While the concepts of declustering and partial
20 declustering are known, no method previously existed for constructing or supporting partially declustered tables in a statically partitioned parallel database system.

Without partial declustering, the information in each table of the parallel database system would need to be spread across the
25 entire parallel database system which could result in significant inefficiency from excess communication overhead as a result of small tables being distributed across a large number of nodes in the parallel database system. This is referred to as "full declustering" and in addition to the above inefficiency does not
30 permit multiple databases to be divided up into groups of nodes that are subsets of the parallel system. It would be desirable in some cases to divide up the parallel database system into a number of groups of nodes, each group being assigned to a particular database with the tables of that particular database being
35 declustered across the particular group of nodes.

Without the ability to perform partial declustering in a parallel database system the user may not have the ability to

control the performance and manageability of the parallel database system.

Summary of the Invention

5 The invention herein overcomes the problems of the prior art in providing a method of partial declustering for a statically partitioned parallel database system.

10 The invention provides for and in a statically partitioned parallel database system which has a plurality of interlinked database node in which each of the nodes is associated with storage means and processor resources, a method of distributing data of a table across a subset of the nodes using partitioning of the table. Horizontal or vertical partitioning can be used.

15 More particularly, optimization of the parallel database system under the invention can be achieved by selecting the number of nodes for a subset based on the size of the table to be declustered, the capacity of storage resources at each node, the maximum number of partitions of the table, and the nature of operations to be performed on said table, such as performing a sequential scan, an index based scan or relational join operations.

20 Advantageously the method of the invention includes the use of a partitioning map which is generated to define the mapping of partitions of the table to the nodes in a subset.

25 The method of the invention may include the distribution of data to the nodes in accordance with a partitioning key value, partitioning method, and information stored in the partitioning map. The partitioning key comprises a set of user defined fields for the table. The partitioning key value is the value of a selected set of fields for a particular row of the table.

30 Advantageously the partitioning method may include hash partitioning, range partitioning, or round-robin partitioning, which is applied to the key value to obtain an index value to the partitioning map which defines the location in the partitioning map which provides the node number where the row is to be stored.

35 A useful embodiment of the invention includes one in which each of the nodes is associated with its own storage resources and its own processor or its own set of processors, ie. a shared

nothing system.

Further aspects of the invention provide a parallel data base system comprising a plurality of interlinked database nodes, each of the nodes having associated storage and processor resources.

5 The database system includes data distribution apparatus for distributing data of a table across the nodes, in which the apparatus includes a partitioner for partitioning the table in accordance with a predetermined partitioning arrangement and distributor for distributing data of said table across a subset of
10 said nodes in accordance with said partitioning arrangement.

The partitioning arrangement advantageously may comprise a partitioning map defining the relationship of partitions of the table to the subset of nodes.

15 Preferably, the partitioning map defines the mapping of horizontal partitions of the table to the subset of nodes.

Yet another aspect of the invention provides an article of manufacture (a computer program product) comprising a computer useable medium having computer readable program code routine embodied therein for causing a computer system to distribute data
20 of a table across a subset of nodes of a parallel database system. A computer readable program code routine in the article of manufacture comprises a computer readable program code partitioning routine for causing a computer system to partition the table in accordance with a predetermined partitioning arrangement, and
25 computer readable program code for causing a computer system to distribute data of the table across a subset of the nodes in accordance with the partitioning arrangement.

30 The partitioning arrangement preferably comprises a computer readable program code partitioning arrangement routine for causing a computer system to horizontally or vertically partition said table.

The partitioning arrangement is preferably a program routine for causing a computer to establish a partitioning map defining the relationship of partitions of the table to the subset of nodes.

35

Description of the Preferred Embodiment

The problem of determining the best storage strategy for the

tables in a given parallel database system is known to be a difficult problem (see S. Padmanabhan, "Data Placement in Shared-Nothing Parallel Database Systems", PhD Thesis, EECS Department, University of Michigan, Ann Arbor, 1992) and several approaches have been taken to solve this problem (see E. Ozkarahan and M. Ouksel, "Dynamic and Order Preserving Data Partitioning for Database Machines", Proceedings of 1985 VLDB International Conference, 1985; G. Copeland et al, "Data Placement in BUBBA", Proceedings of 1988 ACM SIGMOD Conference, Chicago, pp. 99-108, June 1988; S. Ghandeharizadeh and D.J. DeWitt, "MAGIC: A Multiattribute Declustering Mechanism for Multiprocessor Database Machines", IEEE Transactions on Parallel and Distributed Systems, Volume 5, May 1994).

Three important aspects of data placement include declustering, assignment and partitioning. Declustering refers to the technique of distributing the rows of a single table across multiple nodes. If the rows are stored across all the nodes of the parallel database system, then the table is said to be "fully declustered". If the rows are distributed across a subset of nodes, then the table is said to be "partially declustered". The number of nodes across which a table is declustered is referred to as the "degree of declustering" of the table. The term "table partition" refers to the set of rows of a given table that are all stored at one node of the shared-nothing system (therefore, the number of table partitions equals the degree of clustering).

After choosing the degree of declustering, it is important to solve the assignment problem, which is the problem of determining the particular set of nodes on which the table partitions are to be stored. The following issues arise during assignment. Given any two database tables, their assignment may be nonoverlapped, i.e., the two tables do not share any common nodes. Conversely, their assignment may be overlapped, in which case the two tables share at least one node. If both tables share exactly the same set of nodes, then the tables are said to be fully overlapped. Finally, the problem of partitioning refers to the problem of choosing a technique to assign each row of a table to a table partition. Common techniques are, round-robin, hash, and range partitioning.

In the last two, a set of columns (attributes) of the table are defined as the partitioning keys and their values in each row are used for hash or range partitioning.

5 While the invention herein is useful in a shared nothing parallel data processing system it is also useful in systems that share some or all resources. In a shared-nothing parallel database system implementing a relational database system, a single database can be stored across several computers (which includes a processor and storage) or nodes that do not share memory or disk storage. A
10 technique called "horizontal partitioning" is used to spread the rows of each table in the database across multiple nodes. The advantage of horizontal partitioning is that one can exploit parallel input/output capability to improve the speed at which data is read from storage units associated with the nodes. The
15 technique used to determine in which node a given row of a table is stored is called the "partitioning strategy". A number of suitable partitioning strategies exist, eg. key range, round robin, and hash partitioning. For example, in a hash partitioning strategy, a hash function is applied to values in a predetermined set of columns,
20 namely the partitioning key columns, as illustrated in Figure 5, and the resultant value is used as the node number at which the corresponding row is stored.

With horizontal partitioning, there is typically an optimal number of nodes beyond which it is not beneficial to partition
25 data, due to the overheads involved in storing, retrieving and collecting data that is spread across many nodes. Typically, smaller tables should be partitioned across fewer nodes, larger tables should be partitioned against more nodes. Thus, a database user should be allowed to exercise control on the number of nodes
30 on which a table is stored. Also, it is possible that an organization that owns a large parallel computer may administer it in such a way that different databases reside on different sets of nodes. Thus, a shared-nothing parallel database system advantageously should preferably provide mechanisms by which users
35 can control the number of nodes, and the specific nodes, on which the database tables are stored. This is achieved by the present invention and is illustrated in conjunction with the following

particular embodiment.

While embodiments of the invention are discussed in terms of horizontal partitioning, it will be realized by those skilled in the art referring to this specification, that vertical partitioning can be utilized to spread the columns of a table, or tables across multiple nodes and that the operations and description pertaining to rows can be applied to columns when using vertical partitioning.

Brief Description of the Drawings

The features of the invention will become more apparent by reference to the following description taken in conjunction with the accompanying drawings, in which:

Figure 1 is a data relationship diagram illustrating the data stored in catalogues (sets of tables describing data in the database) of the database to implement partial declustering of tables in the database system;

Figure 2 is an illustration of the tables contained in the catalogue of the parallel database system, representing information that the system keeps in order to implement partial declustering for user tables;

Figure 3 is an illustration of a series of steps performed when a user creates a partially declustered table in accordance with the invention;

Figure 4 is an illustration of the process by which a user inserts data into a partially declustered table already resident on the parallel database system.

Figure 5 is an illustration of partitioning keys and a partitioning map.

Figure 6 is an illustration of a parallel database system.

The present invention provides a method and apparatus for implementing partial declustering so that users can optimize the performance of database systems. In the particular embodiment discussed below, partial declustering is implemented using horizontal partitioning in the context of hash partitioning of database tables.

A parallel database system and the steps for creating a

computer program to implement the invention are described herein.

The specific implementation herein makes use of nodegroups in order to support partial declustering of hash partition database tables. Nodegroups are subsets each of which is uniquely identified, eg. by a user provided name of the nodes of a parallel database system. Nodegroups are defined within each database, in this example, by a user, although the processing system can provide default nodegroup definitions.

At the time of their creation, tables are created within existing nodegroups. As a result, the data in the table is partitioned only across the set of nodes defined in the corresponding nodegroup. In order to facilitate the invention, data definition instructions have been created to allow users to create, drop and alter nodegroup definitions. Relational databases typically use SQL for operation so this embodiment will be discussed with the premise that SQL is used. In the specific embodiment support for nodegroups, new SQL Data Definition Language (DDL) statements have been created. These are discussed below. Referring to Figure 2, new catalogue tables have also been introduced to store information related to nodegroups in normalized form. These will be discussed in greater detail subsequently in this description.

Data Structures

Figure 1 indicates the basic data structures used to implement partial declustering. The figure is basically an entity relationship diagram showing the relationship between various entities (i.e. the items within the boxes). All the entities are specific to each database, except the entities called "nodes". Databases implemented in the parallel database system have access to the nodes of the parallel database system. The entities that are specific to a database are tables, indexes and two new entities implemented in the present invention, namely "nodegroups" and "partitioning maps".

Nodes

In referring to Figure 6, in parallel database systems, nodes

1, represent a collection of computational resources including usually a processor 6 for processing, main memory 7, disk storage 8, and communication resources 9. The physical processor 6, which has its own main memory 7, and disks 8, and which can communicate with other processors, represents a node. It is also possible to implement multiple nodes in a single physical processor as long as each node manages its own memory disks and communications. In this case, such nodes will typically multiplex the use of a single physical processor or CPU. In the specific implementation herein, the shared-nothing parallel database system uses a known set of nodes across which data from all databases can be stored. Each node is uniquely identified by a node identifier in the embodiment herein. The set of nodes is common to all databases in the system, that is to say, all databases in the system conceptually have the ability to operate on each of the nodes, however, whether they will or not depends on the specific implementation chosen by the database user applying the method of the invention herein.

Nodegroups

Referring to Figure 1, the database object called NODEGROUP 2, is a named subset of the set of nodes in a shared-nothing parallel database system described herein. Each nodegroup in a given database is identified by a unique name. As indicated in Figure 1, the implementation of the invention herein supports a many-to-many (M-N) relationship between nodes and nodegroups. A nodegroup 2, may contain one or more nodes and each node can be a member of zero or more nodegroups. A nodegroup must contain at least one node. Figure 5 illustrates another nodegroup formed from nodes N1, N2, N3, N4.

Partitioning Maps

Referring again to Figure 1, a partitioning map 3 is a data structure associated with a nodegroup 2, which indicates the node on which a given row of a given table is stored. Each partitioning map has a unique partitioning map identification (PMID). As indicated in Figure 1, each nodegroup 2 is associated with one partitioning map 3 and each partitioning map 3 belongs only to one

nodegroup.

A partitioning map can be generated by allocating node numbers to the partitions using a round robin allocation scheme to assign node numbers in the partitioning map. For example, in a 20 node system if there are three nodes in a nodegroup eg. nodes 7, 11, 15 and assuming the partition map contains 4K entries then the partitioning map would be generated as 7, 11, 15, 7, 11, 15.....which would repeat to fill the entire 4K space. This of course, assumes a uniform distribution of data, allocating an equal number of partitions for each node. Rows of the database are mapped to the nodes in the nodegroup using the partitioning map.

The SQL DDL statement to form the nodegroup named MYNODEGROUP would be CREATE NODEGROUP MYNODEGROUP ON NODES (7, 11, 15).

15 Tables

Still referring to Figure 1, a database consists of a set of tables 4. A table 4 in the database is uniquely identified by the creator name and table name, in a typical implementation. Each table is created within a nodegroup. A nodegroup can contain zero or more (N) tables.

Indexes

A table 4 may have zero or more indexes 5 associated with it. Each index is uniquely identified by the name of the creator and the name of index in this implementation. Other identification methods are useful as well. Each index 5 is associated with a single table 4. Typically the index 5 consists of a fast access structure to access data within the table. This is well known in the art.

Physical Data Structures

The information about the above data structures, nodegroups, partitioning maps, tables and indexes, are stored in the form of catalogue tables, some of which are illustrated in Figure 2 as follows:

SYSNODEGROUP 10: contains information about nodegroups in the

whole processing system eg. the Name of the nodegroup, Definer (the name of the person who defined the nodegroup), Partitioning Map ID for nodegroup.

5 **SYSNODEGROUPDEF 11:** contains information about the relationship between nodes and nodegroups, i.e. which subset of nodes are in a given nodegroup.

10 **SYSPARTITIONMAP 12:** contains information about partitioning maps.

15 **SYSTABLE 13:** contains information about tables, eg. the name of the table, the name of the creator of the table, the name of the nodegroup the table was created in, the partitioning maps ID of the nodegroup.

SYSINDEXES 14: contains information about indexes.

Procedures for Updating Data Structures

20 Referring to Figure 3, the following steps are followed in creating a partially declustered table:

1. Create a nodegroup on the desired set of nodes.

25 (a) Insert a row of data for the nodegroup into SYSNODEGROUPS.

 (b) Insert a row of data for the corresponding partitioning map into SYSPARTITIONMAPS.

30 (c) Insert one row of data for each node in the nodegroup into SYSNODEGROUPDEF.

2. Referring to Figure 3(b), the following method is used to create a table in the specified node group:

 (a) Insert a row in SYSTABLES for this table.

35 (b) In the row corresponding to this table in SYSTABLES, include information, eg. nodegroup name, that links this

table to the specified nodegroup.

(c) In the row corresponding, store the partitioning key information for the table in SYSTABLES. The partitioning key information is used to decluster the table across the nodes in the nodegroup, as specified in the next section.

Inserting Rows into the Table

Referring to Figure 4 and 5, for each row to be inserted, use the partitioning key value (either directly or after hashing) as an index to the partitioning map data structure of the node group corresponding to this table. The partitioning map indicates the node in which this row is stored.

Referring to Figure 5 the distribution of data of a table across a subset of nodes (nodes 1, 2, 3, 4) in a parallel system is illustrated.

In this illustration A is a column of table 4 and is used as a partitioning key.

$H()$, (15) is a hash function used to obtain a partition number when applied to the partitioning key value.

Partitioning map 3 is an array in which each entry contains the node number of the node in which rows of data that hash to this partition (node) are stored.

In this illustration column B is not used in hashing, the partitioning of data from column B follows that of column A.

From a review of Figure 5 the distribution of data A_i and B_i ($i=1$ to 6) from table 4 nodes N1, N2, N3 and N4 respectively, is accomplished using partitioning map 3 as may be readily appreciated.

A further detailed example of the preferred embodiment of the invention using DDL statements is provided below, in which the aspects of partial declustering, overlapped assignment of database tables into multiple nodegroups, and hash partitioning are illustrated in a parallel database in which several nodegroups are implemented. The following nodegroup DDL statement creates three node groups:

CA9-95-001

12

```
CREATE NODEGROUP GROUP_1 ON NODES (1 TO 32, 40, 45, 48)
CREATE NODEGROUP GROUP_2 ON NODES (1, 3, 33)
CREATE NODEGROUP GROUP_3 ON NODES (1 TO 32, 40, 45, 48)
```

5 In the above example, GROUP_1 and GROUP_3 are two different nodegroups, even though they contain the same set of nodes, (nodes 1 to 32, 40, 45, and 48). Nodegroup GROUP_2 is partially overlapped with GROUP_1 and GROUP_3 (on nodes 1 and 3).

10 Extensions to CREATE TABLE DDL

 When creating a table, it is advantageous to specify the nodegroup on which the table will be declustered. In addition, it is also advantageous to specify the columns to be used for the partitioning key.

15 The following example illustrates the use of DDL extensions to the CREATE TABLE statement:

```
CREATE TABLE PARTS (Partkey integer, Partno Integer) IN GROUP_1
PARTITIONING KEY (partkey) USING HASHING
```

20 CREATE TABLE PARTSUPP (Partkey Integer, Suppkey Integer, PS_Descp char[50]) IN GROUP_1 PARTITIONING KEY (Partkey) USING HASHING

25 CREATE TABLE CUSTOMERS (Custkey Integer, C_Nation char[20]) IN GROUP_1 PARTITIONING KEY (Custkey) USING HASHING

```
CREATE TABLE SUPPLIERS (Suppkey Integer, S_Nation char [20]) IN
GROUP_1 PARTITIONING KEY (Suppkey) USING HASHING
```

30 CREATE TABLE ORDERS (Orderkey Integer, Custkey Integer, Orderdate date) IN GROUP_1 PARTITIONING KEY (Orderkey) USING HASHING

35 The partitioning key of tables PARTS and PARTSUPP is Partkey. All tables are partitioned across the set of nodes identified by the nodegroup, GROUP_1.

Choosing the Number of Nodes in a Nodegroup

When a set of tables is to be created in a nodegroup, the size of the nodegroup should typically be determined by the size of the tables in this set. The nodegroup should contain a minimum number of nodes such that there is sufficient amount of disk storage available to store the data belonging to all the tables. At the same time, the nodegroup should not contain so many nodes that some nodes end up without any data being stores on them. This would be wasteful of system resources.

Another consideration is the nature of database operations being performed on the tables. In theory, if a database operation is "perfectly scalable", then the time taken to perform this operation on a given table (or set of tables) decreases linearly as the number of nodes in the nodegroup increase (this is because each node has less data to process, as the number of nodes increase). However, in practice, due to data transfer, communications, and other synchronization overheads, most operations are not perfectly scalable. As a result, there may be an "optimal" number of nodes beyond which increasing the number of nodes may result in no change in performance or even a decrease in performance. The point at which this can happen must be determined on a case by case basis, based on the hardware capabilities of each node, the nature of the database operation, and the amount of data involved.

The database query optimizer present in most database systems can be used to assist in this determination as will be appreciated by those skilled in the art.

The absolute maximum for the number of nodes in a nodegroup is, of course, determined by the total number of nodes in the system.

Nodegroup example:

Suppose we have a database consisting of 3 tables:

Table A	500,000 rows	50 GBytes
Table B	50,000 rows	10 GBytes
Table C	100 rows	10 Kbytes

Suppose we have a system with 64 nodes and a little over 10 Gbytes of database disk storage per node.

A minimum of $(50+10+0.00001)$ GBytes / $(10 \text{ Gbytes/node}) = 6$ nodes is needed to store all the data.

5 Suppose that the typical database operations being performed on these tables are searching of tables A and B and joining of tables A, B, and C. Since the search and join operations in this case are scalable operations, we can increase the number of nodes beyond 6 and obtain a performance improvement for the operations
10 performed on Tables A and B. However, Table C is very small in size and 6 nodes is already too many for that table (assuming a disk page size of 4Kbytes, the entire Table C can fit on 2.5 pages for which a single node is sufficient). Also, suppose that the optimal number of nodes for performing searching and joins of
15 tables A and B is less than 64, say 40. Then these tables should be created in a nodegroup containing 40 out of the 64 nodes.

Thus, for the above database example, we would create two nodegroups, one with 40 nodes and other with 1 node (could be one of the 40 or not). Tables A and B would be created in the 40-node
20 nodegroup and Table C would be created in the 1-node nodegroup.

As will be well recognized by those skilled in the art to which this invention pertains, the invention may be practised in computer systems and in computer programs for the operation of computer systems.

CA9-95-001

15

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

1. For a statically partitioned parallel database system
5 comprising a plurality of interlinked database nodes, each of said nodes having associated storage and processor resources, a method of distributing data of a table across said nodes; comprising:

partitioning said table across a subset of nodes in accordance with a predetermined partitioning arrangement; and,

10 distributing data of said table across said subset of nodes in accordance with said partitioning arrangement.

2. For a statically partitioned parallel database system
15 comprising a plurality of interlinked database nodes, each of the nodes having associated storage and processor resources, a method of distributing data of a table across the nodes, comprising:

selecting a subset of nodes to be associated with each other, wherein the nodes to be associated with each other are selected based on a nature of operations to be performed on the table;

20 associating the subset of nodes with each other;

associating the table with the subset of nodes;

partitioning the table across the subset of nodes in accordance with a predetermined partitioning arrangement; and

25 distributing data of the table across the subset of nodes in accordance with the partitioning arrangement.

3. The method of claim 1 wherein said partitioning arrangement comprises horizontal or vertical partitioning.

CA9-95-001

16

4. The method of claim 1 wherein said partitioning arrangement comprises horizontal partitioning.

5. The method of claim 1 wherein said partitioning arrangement
5 comprises a partitioning map defining a relationship of partitions of said table to said subset of nodes.

6. The method of claim 1 or claim 2 wherein the partitioning is
10 selected from the group of at least one of a horizontal partitioning, vertical partitioning, or a partitioning map defining the relationship of partitions of the table to the subset of nodes.

7. The method of claim 4 wherein said partitioning arrangement
15 comprises a partitioning map defining a mapping of horizontal partitions of said table to said subset of nodes.

8. The method of claim 5 in which data is distributed to said
20 subset of nodes in accordance with said partitioning arrangement as applied to a partitioning key value, and information contained in said partitioning map;

wherein:

a partitioning key is a set of user defined fields for said
table; and

a partitioning key value is a value of a selected set of said
25 fields for a particular row of said table, in the case of horizontal partitioning; or, for a particular column of said table, in the case of vertical partitioning.

9. The method of claim 6 in which data is distributed to the

CA9-95-001

17

subset of nodes as applied to a partitioning key value, and information contained in a partitioning map, wherein:

a partitioning key is a set of user defined fields for the table; and

5 a partitioning key value is a value of a selected set of the fields for a particular row of the table, in the case of horizontal partitioning; or, for a particular column of the table, in the case of vertical partitioning.

10 10. The method of claim 1 in which data is distributed to said subset of nodes in accordance with a partitioning key value:

wherein a partitioning key is a set of user defined fields for said table; and

15 a partitioning key value is a value of a selected set of said fields for a particular row of said table, in the case of horizontal partitioning; or, for a particular column of said table in the case of vertical partitioning.

20 11. The method of claim 8 or claim 9 wherein said partitioning arrangement is applied to said key value to obtain an index value to said partitioning map which defines a location in said partitioning map which provides a node number where said particular row is to be stored, in the case of horizontal partitioning; or, where said particular column is to be stored, in the case of
25 vertical partitioning.

12. The method of claim 11 wherein said partitioning arrangement comprises a partitioning method selected from hash partitioning, range partitioning or round robin partitioning.

CA9-95-001

18

13. The method of claim 3 or claim 6 wherein said nodes selected for said subset are selected based on one or more of the following factors:

a size of said table;

a capacity of storage at each node;

a maximum number of horizontal or vertical partitions, respectively;

and a nature of operations to be performed on said table.

14. The method of claim 1 or claim 2 wherein each of the nodes in the subset of nodes is associated with its own storage resources and at least one processor.

15. The method of claim 1 as applied to a plurality of tables comprising:

partitioning said tables across subsets of nodes in accordance with predetermined partitioning arrangements; and

distributing data of said tables across said subsets of nodes in accordance with said partitioning arrangements.

16. For a statically partitioned parallel database system comprising a plurality of interlinked database nodes, each of said nodes having associated storage and processor resources, a method of distributing data of a plurality of tables across said nodes, comprising:

selecting subsets of nodes to be associated with each other, wherein the nodes to be associated with each other are selected based on a nature of operations to be performed on the tables;

associating the subsets of nodes with each other;

CA9-95-001

19

associating each of the tables with the subsets of nodes;
partitioning the tables across the subsets of nodes in
accordance with predetermined partitioning arrangements; and
distributing data of the tables across subsets of nodes in
5 accordance with the partitioning arrangements.

17. The method of claim 16 or claim 15 wherein each table is
partitioned in accordance with an unique partitioning arrangement.

10 18. The method of claim 15 or claim 16 wherein said subsets of
nodes are selectively non-overlapping or overlapping.

19. The method of claim 1 wherein said nodes are implemented in at
least one physical processor wherein each node is associated with
15 its own storage resource.

20. A statically partitioned parallel database system comprising
a plurality of interlinked database nodes, each of said nodes
having associated storage and processor resources, data
20 distribution means for distributing data of a table across said
nodes; comprising:

partitioning means for partitioning said table across a subset
of nodes in accordance with a predetermined partitioning
arrangement; and,

25 distributing means for distributing data of said table across
said subset of nodes in accordance with said partitioning
arrangement.

21. A statically partitioned parallel database system comprising

CA9-95-001

20

a plurality of interlinked database nodes, each of the nodes having associated storage and processor resources, data distribution means for distributing data of a table across the nodes, comprising:

means for selecting a subset of nodes to be associated with each other, wherein the nodes to be associated with each other are selected based on a nature of operations to be performed on the table;

means for associating the subset of nodes with each other;

means for associating the table with the subset of nodes;

partitioning means for partitioning the table across the subset of nodes in accordance with a predetermined partitioning arrangement; and,

distributing means for distributing data of the table across the subset of nodes in accordance with the partitioning arrangement.

22. The system of claim 20 or claim 21 wherein said partitioning arrangement comprises means for horizontally or vertically partitioning said table.

23. The system of claim 20 or claim 21 wherein said partitioning arrangement comprises a partitioning map defining a relationship of partitions of said table to said subset of nodes.

24. The system of claim 22 wherein said partitioning arrangement comprises a partitioning map defining a relationship of horizontal partitions of said table to said subset of nodes.

25. The system of claim 23 comprising means for distributing data

CA9-95-001

21

to said subset of nodes using said means for partitioning applied to a partitioning key value, and information contained in said partitioning map;

wherein:

5 a partitioning key is a set of user defined fields for said table; and,

a partitioning key value is a value of a selected set of said fields for a particular row of said table, in the case of horizontal partitioning or, for a particular column of said table,
10 in the case of vertical partitioning.

26. The system of claim 25 wherein said means for partitioning is operated on said key value to obtain an index value to said partitioning map which defines a location in said partitioning map
15 which provides a node number where said particular row is to be stored, in the case of horizontal partitioning; or, where said particular column is to be stored, in the case of vertical partitioning.

20 27. The system of claim 26 wherein said means for partitioning comprises partitioning means selected from hash partitioning means, range partitioning means or round robin partitioning means.

28. The system of claim 22 wherein said nodes selected for said
25 subset are selected based on one or more of the following factors:

a size of said table;

a capacity of storage at each node;

a maximum number of horizontal or vertical partitions respectively; and

CA9-95-001

22

a nature of operations to be performed on said table.

29. The system of claim 20 and claim 21 in which each of the nodes in the subset of nodes is associated with its own storage resources and at least one processor.

30. The system of claim 20 applied to a plurality of tables comprising:

means for partitioning said tables across subsets of nodes in accordance with predetermined partitioning arrangements; and

means for distributing data of said tables across said subsets of nodes in accordance with said partitioning arrangements.

31. The system of claim 21 applied to a plurality of tables comprising:

means for selecting subsets of nodes to be associated with each other, wherein the nodes to be associated with each other are selected based on a nature of operations to be performed on the tables;

means for associating the subsets of nodes with each other; means for associating each of the tables with the subsets of nodes;

means for partitioning the tables across the subsets of nodes in accordance with predetermined partitioning arrangements; and

means for distributing data of the tables across the subsets of nodes in accordance with the partitioning arrangements.

32. The system of claim 30 or claim 31 in which each partitioning arrangement is unique.

CA9-95-001

23

33. The system of claim 30 or claim 31 in which said subsets of nodes are selectively non-overlapping or overlapping.

34. The system of claim 20 wherein said nodes are implemented in a single physical processor wherein each node is associated with its own storage resources.

35. A computer program product comprising computer instructions recorded on a recording medium, said computer instructions capable of enabling a parallel data processing system to implement any of the database systems of claims 20, 21, 32 or 33 to distribute data of a table across the subset of nodes of the processing system to facilitate processing of the table data by the data processing system.

36. An article of manufacture comprising:

a computer usable medium having computer readable program code means embodied therein for causing a computer system to distribute data of a table across a subset of nodes of a parallel database system;

said computer readable program code means in said article of manufacture comprising:

computer readable program code partitioning means for causing a computer to partition said table across a subset of node in accordance with a predetermined partitioning arrangement; and

computer readable program code distributing means for causing a computer to distribute data of said table across said subset of nodes in accordance with said partitioning arrangement.

CA9-95-001

24

37. An article of manufacture comprising:

a computer usable medium having computer readable program code embodied therein for causing a computer system to distribute data of a table across a subset of nodes of a parallel database system;

5 the computer readable program code in the article of manufacture comprising:

computer readable program code for causing a computer to select a subset of nodes to be associated with each other, wherein the nodes to be associated with each other are selected based on a nature of operations to be performed on the table, associate the subset of nodes with each other, associate the table with the subset of nodes, partition the table in accordance with a predetermined partitioning arrangement, and

15 distribute data of the table across the subset of nodes in accordance with the partitioning arrangement.

38. The article of manufacture of claim 36 or claim 37 wherein said partitioning arrangement comprises computer readable program code for causing a computer to select from the group of at least one of a horizontal partitioning or a vertical partitioning of said table.

39. The article of manufacture of claim 36 or claim 37 where said partitioning arrangement comprises computer readable program code means for causing a computer to establish a partitioning map defining a relationship of partitions of said table to said subset of nodes.

40. The article of manufacture of claim 38 wherein said

CA9-95-001

25

partitioning arrangement comprises computer readable program code means for causing a computer to establish a partitioning map defining a mapping of horizontal partitions of said table to said subset nodes.

5

41. The article of manufacture of claim 39 or claim 40 comprising computer readable program code means for causing a computer to distribute data to said subset of nodes using said code partitioning means applied to a partitioning key value, and information contained in said partitioning map;

10

wherein:

a partitioning key is a set of user defined fields for said table; and

a partitioning key value is a value of a selected set of said fields for a particular row of said table, in the case of horizontal partitioning or, for a particular column of said table, in the case of vertical partitioning.

15

42. The article of manufacture of claim 39 comprising computer readable program code means for causing a computer to distribute data to said subset of nodes in accordance with a partitioning key value, and information contained in said partitioning map;

20

wherein:

a partitioning key is a set of user defined fields for said table; and

25

a partitioning key value is a value of a selected set of said fields for a particular row of said table, in the case of horizontal partitioning or, in the case of vertical partitioning, for a particular column of said table.

CA9-95-001

26

43. The article of manufacture of claim 41 or claim 42 comprising computer readable program code for causing a computer to utilize the partitioning code with the key value to obtain an index value to the partitioning map which defines a location in the partitioning map which provides a node number where the particular row is to be stored, in the case of horizontal partitioning; or, where the particular column is to be stored in the case of vertical partitioning.

44. The article of manufacture of claim 43 wherein the partitioning code comprises partitioning selected from hash partitioning, range partitioning or round robin partitioning.

45. The article of manufacture of claim 38 or claim 39 wherein said nodes selected for said subset are selected based on one or more of the following factors:

a size of said table;

a capacity of storage at each node;

a maximum number of horizontal or vertical partitions respectively; and

a nature of operations to be performed on said table.

46. The article of manufacture of claim 36 or claim 37 in which each of the nodes in the subset of nodes is associated with its own storage resources and at least one processor.

47. The article of manufacture of claim 36 or claim 37, for use with a plurality of tables, comprising:

computer readable program code for causing a computer to

CA9-95-001

27

partition the tables in accordance with predetermined partitioning arrangements; and

computer readable program code for causing a computer to distribute data of the tables across subsets of the nodes in accordance with the partitioning arrangements.

48. The article of manufacture of claim 47 in which each partitioning arrangement is unique.

49. The article of manufacture of claim 47 in which the subsets of nodes are relatively non-overlapping or overlapping.

50. The apparatus of claim 36 wherein said nodes are implemented in a single physical processor wherein each node is associated with its own storage resources.

51. A program storage device readable by a computer system, tangibly embodying a program of instructions executable by said system to perform the method steps for distributing data of a table across a subset of nodes of a statically partitioned parallel database system comprising a plurality of interlinked database nodes, said method steps comprising:

partitioning said table across a subset of nodes in accordance with a predetermined partitioning arrangement; and

distributing data of said table across said subset of nodes in accordance with said partitioning arrangement.

52. A program storage device readable by a computer system, tangibly embodying a program of instructions executable by the

CA9-95-001

28

system to perform the method steps for distributing data of a table across a subset of nodes of a statically partitioned parallel database system comprising a plurality of interlinked database nodes, the method steps comprising:

5 selecting a subset of nodes to be associated with each other, wherein the nodes to be associated with each other are selected based on a nature of operations to be performed on the table;

 associating the subset of nodes with each other;

 associating the table with the subset of nodes;

10 partitioning the table in accordance with a predetermined partitioning arrangement; and

 distributing data of the table across the subset of nodes in accordance with the partitioning arrangement.

15 53. The device of claim 51 wherein said partitioning arrangement comprises horizontal or vertical partitioning.

54. The device of claim 51 wherein said partitioning arrangement comprises horizontal partitioning.

20 55. The device of claim 51 wherein said partitioning arrangement comprises a partitioning map defining a relationship of partitions of said table to said subset of nodes.

25 56. The device of claim 54 wherein said partitioning arrangement comprises a partitioning map defining a mapping of horizontal partitions of said table to said subset of nodes.

57. The device of claim 51 further comprising distributing data to

CA9-95-001

29

said subset of nodes in accordance with said partitioning arrangement of claim 55 applied to a partitioning key value, and information contained in said partitioning map;

wherein:

5 a partitioning key is a set of user defined fields for said table; and

a partitioning key value is a value of a selected set of said fields for a particular row of said table, in the case of horizontal partitioning; or, for a particular column of said table,
10 in the case of vertical partitioning.

58. The device of claim 51 further comprising distributing data to said subset of nodes in accordance with a partitioning key value, wherein a partitioning key is a set of user defined fields for said
15 table; and

a partitioning key value is the value of a selected set of said fields for a particular row of said table, in the case of horizontal partitioning; or, for a particular column of said table, in the case of vertical partitioning.

20 59. The device of claim 57 further comprising applying said partitioning arrangement to said key value to obtain an index value to said partitioning map which defines a location in said partitioning map which provides a node number where said particular
25 row is to be stored, in the case of horizontal partitioning; or, where said particular column is to be stored, in the case of vertical partitioning.

60. The device of claim 59 wherein said partitioning method

CA9-95-001

30

comprises a partitioning method selected from hash partitioning, range partitioning method selected from hash partitioning, range partitioning or round robin partitioning.

5 61. The device of claim 53 wherein said nodes selected for said subset are selected based on one or more of the following factors:

a size of said table;

a capacity of storage at each node;

10 a maximum number of horizontal or vertical partitions, respectively; and

a nature of operations to be performed on said table.

15 62. The device of claim 53 in which each of said nodes is associated with its own storage resources and one or more processors.

63. The device of claim 53 as applied to a plurality of tables comprising:

20 partitioning said tables in accordance with predetermined partitioning arrangements; and

distributing data of said tables across subsets of said nodes in accordance with said partitioning arrangements.

25 64. The device of claim 63 further comprising partitioning each said table in accordance with an unique partitioning arrangement.

65. The device of claim 63 in which data is distributed to subsets of nodes that are non-overlapping.

CA9-95-001

31

66. The device of claim 63 in which data is distributed to subsets of nodes that are overlapping.

67. A memory for storing data for access by a program being
5 executed on a data processing system, comprising:

a data structure stored in the memory, the data structure including information resident in a data base used by the program and including:

10 a plurality of data objects stored in the memory, each of the data objects containing different information from the data base;

a data object for selecting a subset of nodes to be associated with each other, wherein the nodes to be associated with each other are selected based on a nature of operations to be performed on the table;

15 a data object for associating the subset of nodes with each other;

a data object for associating the table with the subset of nodes;

20 a data object for partitioning the table in accordance with a predetermined partitioning arrangement; and

a data object for distributing data of the table across the subset of the nodes in accordance with the partitioning arrangement.

2150745




Figure 1




Figure 2

2150745




Figure 3(a)




Figure 3(b)

2150745

INSERT ROW INTO TABLE

Obtain value of partitioning key for the row being inserted

Apply partitioning function, $H(PK)$, (hash-based, range-based, or others) to partitioning key

Use output of partitioning function as index into partitioning map.
Store row at node indicated at the corresponding position, i.e. at node, $PM[H(PK)]$

Figure 4




Figure 5




FIGURE 6

