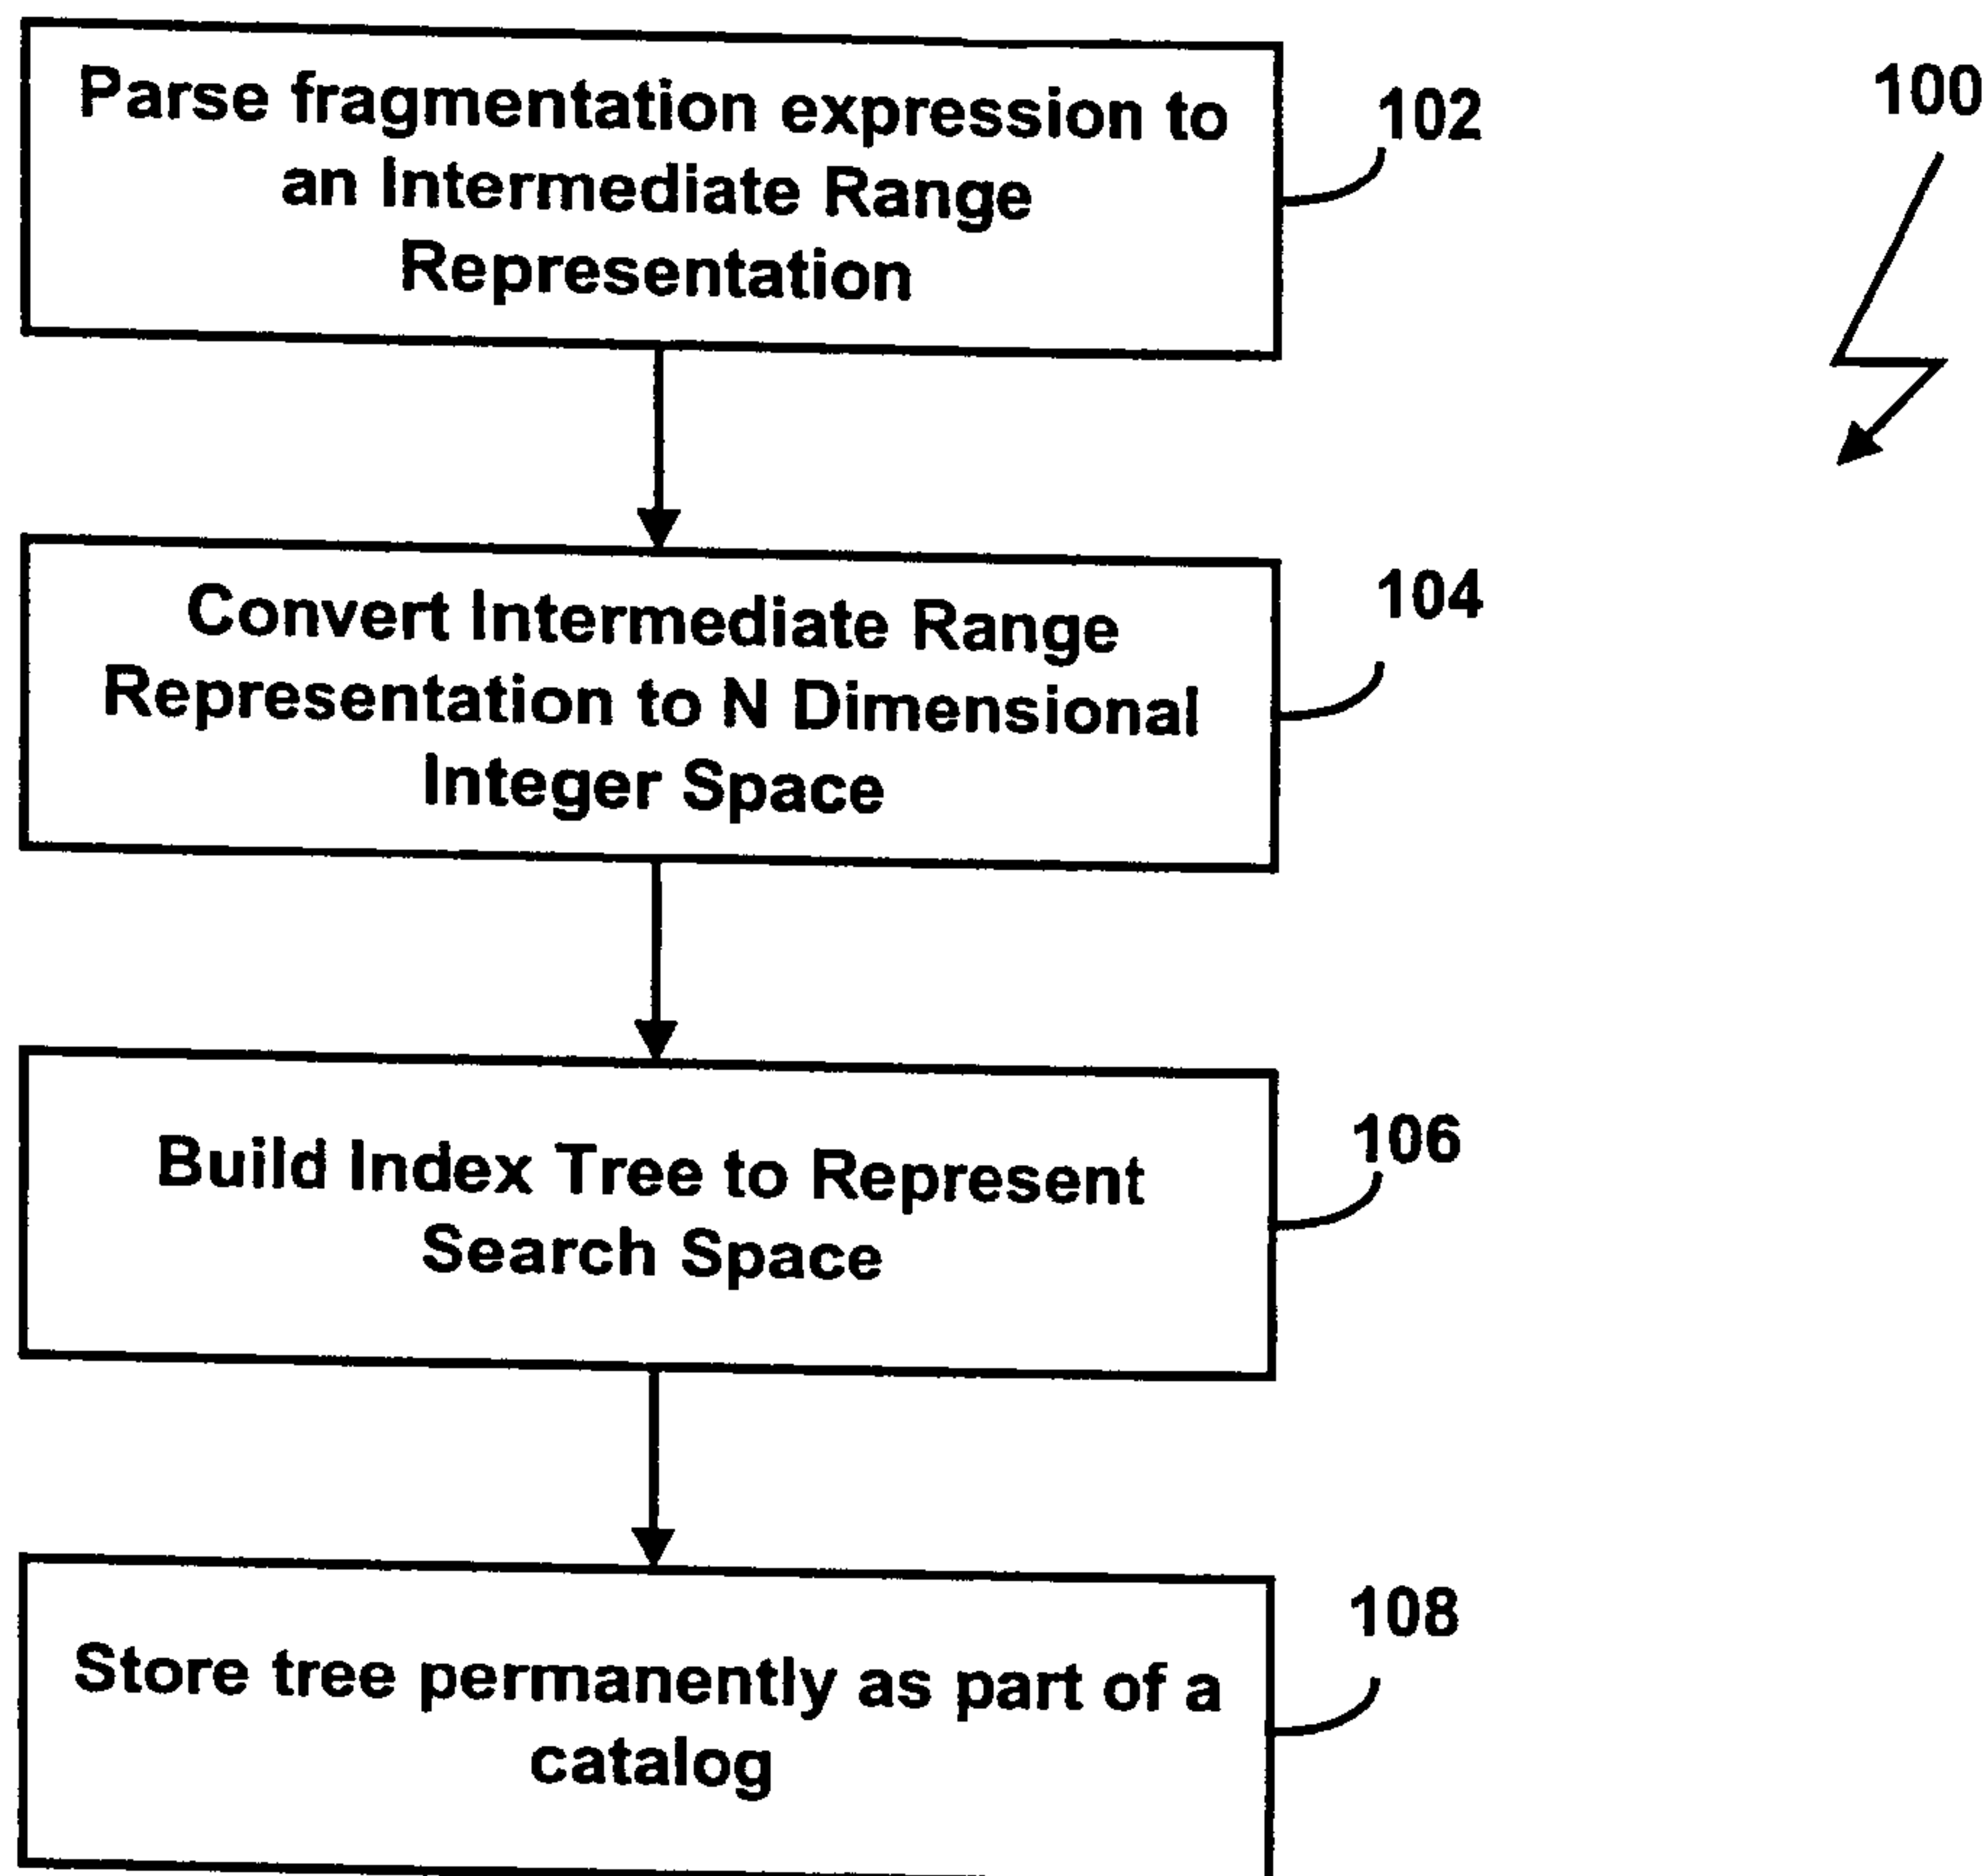




(86) Date de dépôt PCT/PCT Filing Date: 2000/11/06  
 (87) Date publication PCT/PCT Publication Date: 2001/05/10  
 (45) Date de délivrance/Issue Date: 2010/05/25  
 (85) Entrée phase nationale/National Entry: 2002/04/19  
 (86) N° demande PCT/PCT Application No.: US 2000/030638  
 (87) N° publication PCT/PCT Publication No.: 2001/033436  
 (30) Priorité/Priority: 1999/11/04 (US09/434,849)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01)  
 (72) Inventeur/Inventor:  
 ZOU, CHENDONG, US  
 (73) Propriétaire/Owner:  
 INTERNATIONAL BUSINESS MACHINES  
 CORPORATION, US  
 (74) Agent: CHAN, BILL W.K.

(54) Titre : PROCÉDE ET SYSTÈME D'ATTRIBUTION DYNAMIQUE DE ZONE EN FONCTION DE LOCALITE POUR UN RESEAU DE COMMUNICATION SANS FIL  
 (54) Title: SYSTEM FOR MANAGING RDBM FRAGMENTATIONS



(57) **Abrégé/Abstract:**

A relational database system (100) manages data fragments in a database by converting a query or fragmentation expression to an intermediate range representation (104); building an index tree data structure to represent a search space (106) associated with the data fragments; and using the index tree data structure (106) to locate a desired data fragment (108).

## (12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



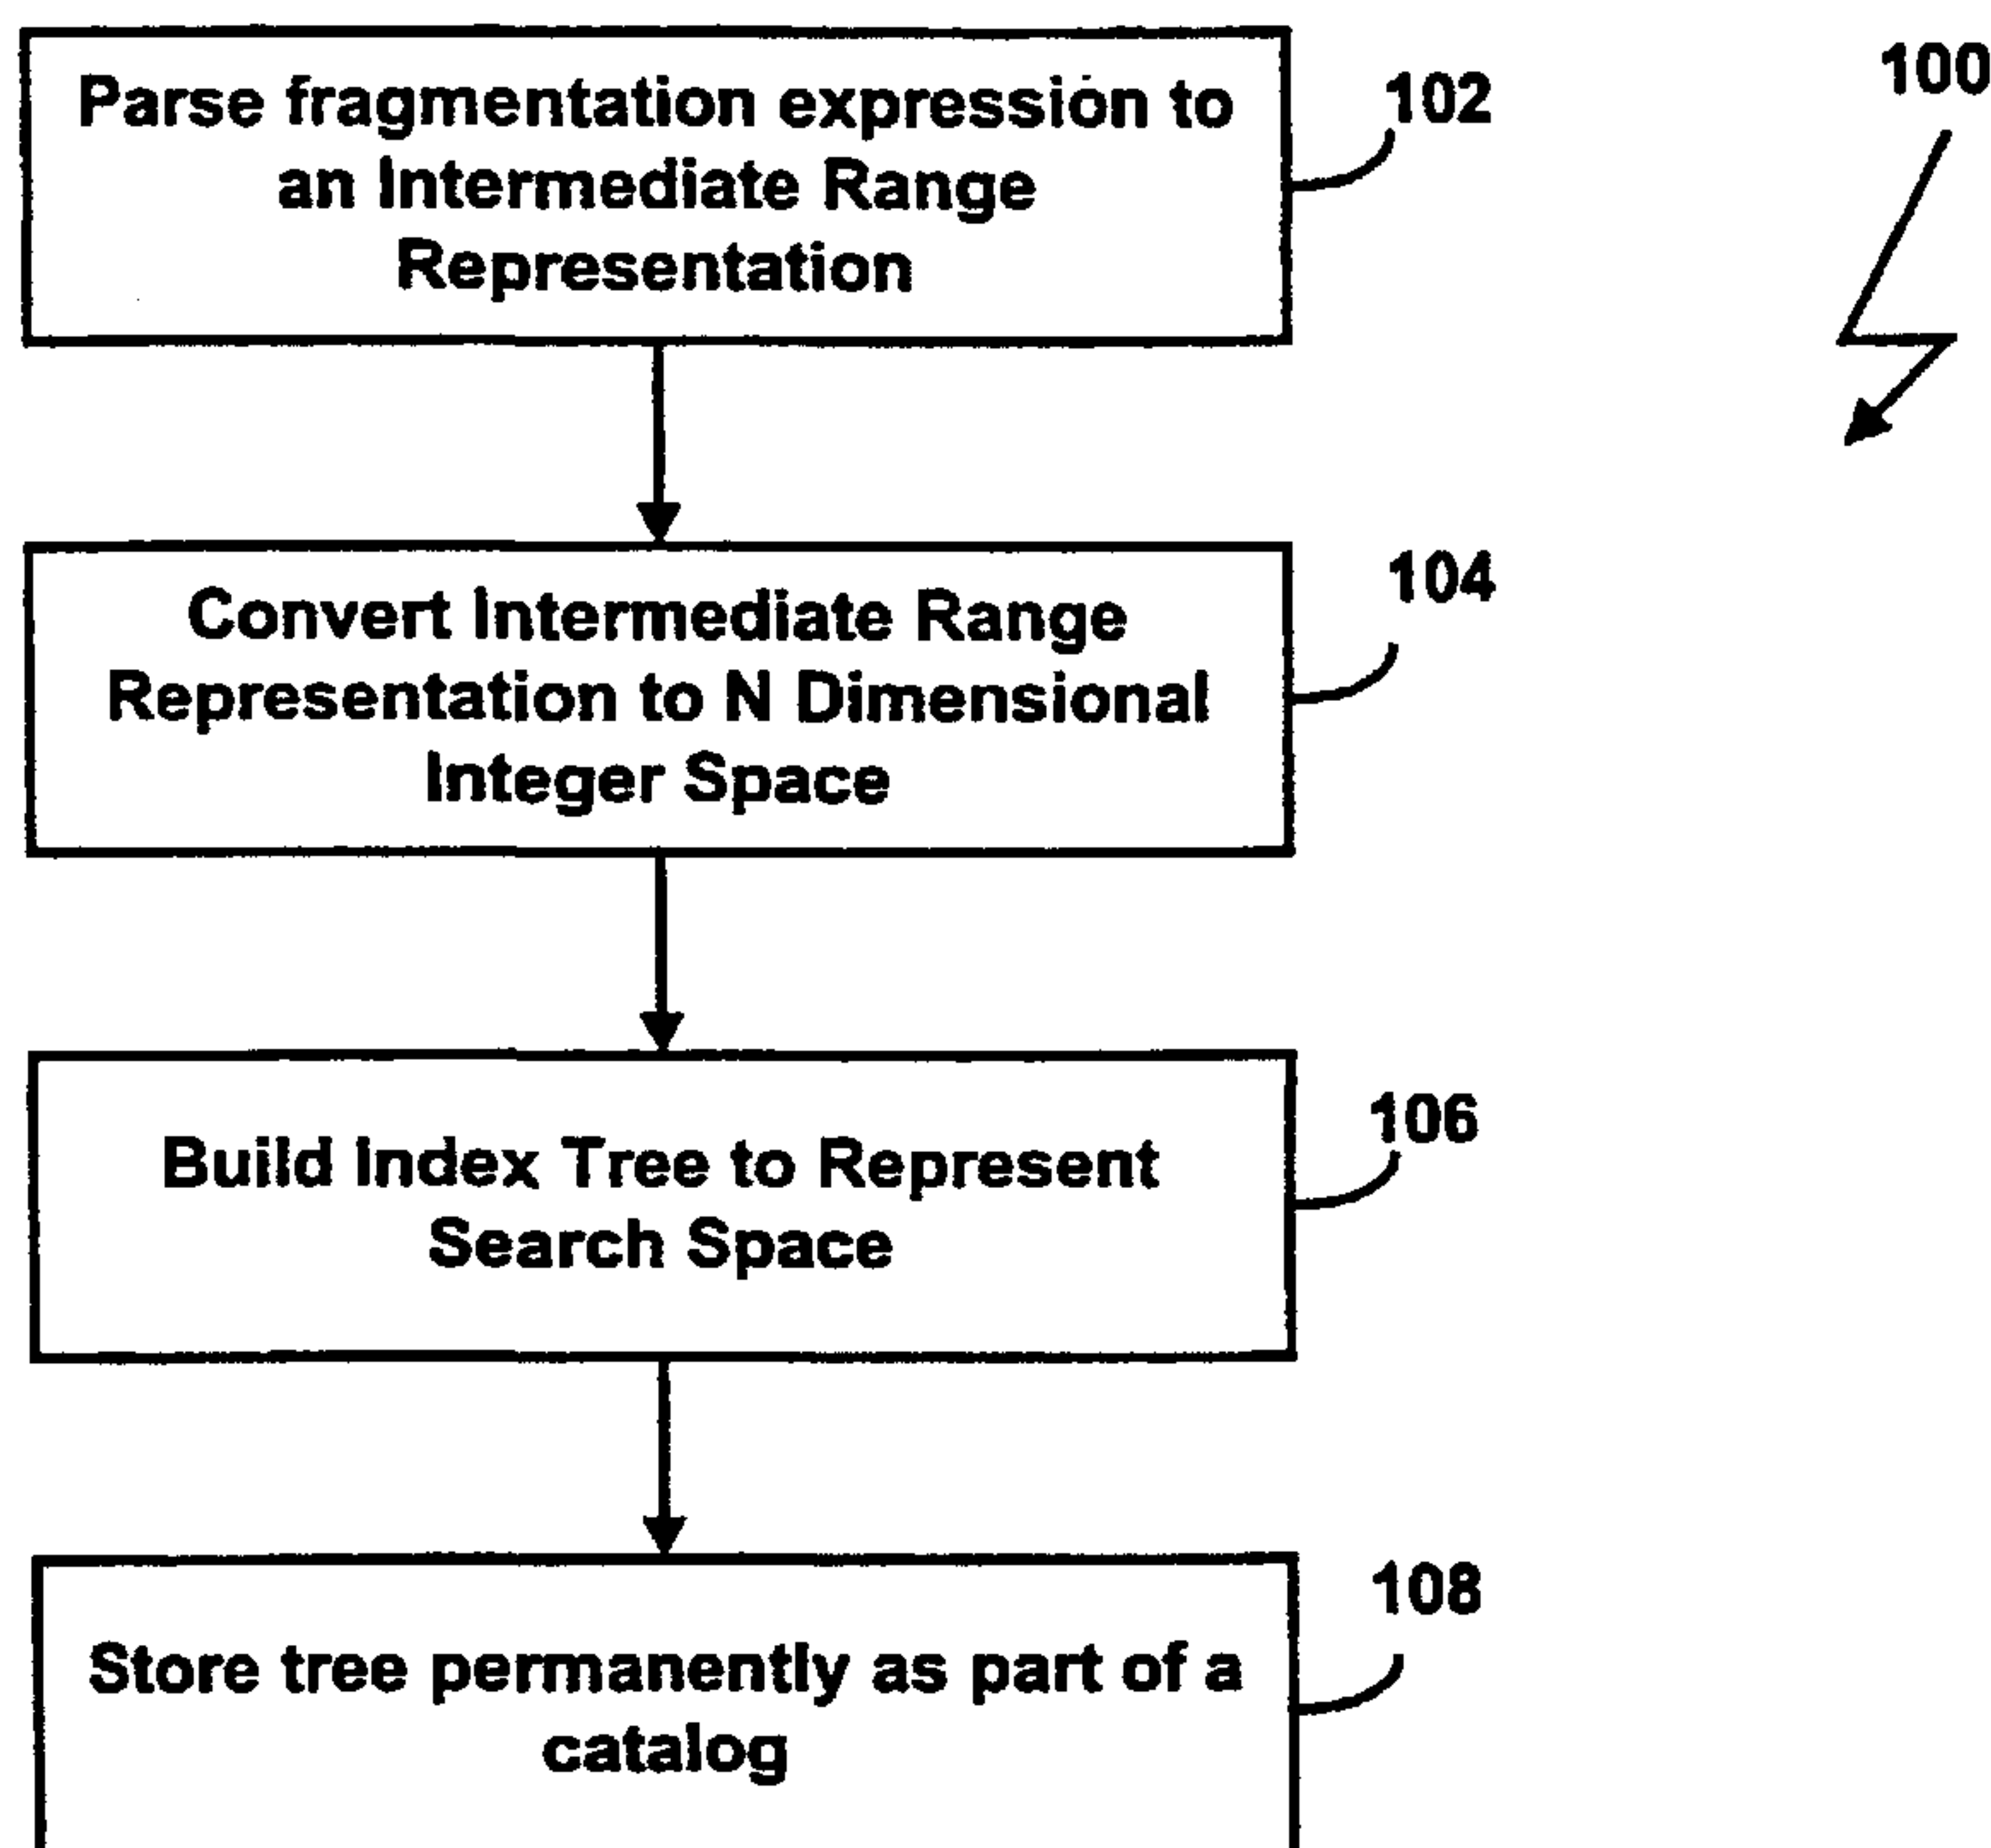
(43) International Publication Date  
10 May 2001 (10.05.2001)

PCT

(10) International Publication Number  
**WO 01/33436 A1**

- (51) International Patent Classification<sup>7</sup>: G06F 17/30 (72) Inventor: ZOU, Chendong; 1026 Foxchase Drive, Apartment 324, San Jose, CA 95123 (US).
- (21) International Application Number: PCT/US00/30638 (74) Agents: EGAN, William, J. et al.; Fish & Richardson P.C., 2200 Sand Hill Road #100, Menlo Park, CA 94025 (US).
- (22) International Filing Date:  
6 November 2000 (06.11.2000) (81) Designated States (*national*): AU, BR, CA, JP, MX.
- (25) Filing Language: English (84) Designated States (*regional*): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).
- (26) Publication Language: English
- (30) Priority Data:  
09/434,849 4 November 1999 (04.11.1999) US  
Published:  
— With international search report.
- (71) Applicant: INFORMIX SOFTWARE, INC. [US/US];  
4100 Bohannon Drive, Menlo Park, CA 94025 (US).  
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: SYSTEM FOR MANAGING RDBM FRAGMENTATIONS



(57) Abstract: A relational database system (100) manages data fragments in a database by converting a query or fragmentation expression to an intermediate range representation (104); building an index tree data structure to represent a search space (106) associated with the data fragments; and using the index tree data structure (106) to locate a desired data fragment (108).



WO 01/33436 A1

## SYSTEM FOR MANAGING RDBM FRAGMENTATIONS

### 5 Background

The present invention relates to indexing and managing data using fragmentation in relational databases.

10 A database is a collection of information. A relational database is a database that is perceived by its users as a collection of tables. Each table arranges items and attributes of the items in rows and columns, respectively. Each table row corresponds to an item (also referred to as a record or tuple), and each table column corresponds to an attribute of the item (referred to as a field or, more correctly, as an attribute type or field type).

15 Fragmentation is a technique used to increase database performance. A system supports data fragmentation if a given relation can be divided up into pieces or fragments. Data can be stored at the location where it is used most frequently. Moreover, two types of fragmentation can be used: horizontal fragmentation and vertical fragmentation, corresponding to relational operations of restriction and projection. The rules assigning a row to a fragment are defined by a database user or administrator and are part of a "fragmentation scheme". It is possible for a fragment of a given table to be empty if none of the rows of the table satisfy the fragmentation scheme's assignment rules for that fragment. "Fragment elimination" is a process by which a database system identifies fragments from a table that cannot participate in the result of a query and removes those fragments from consideration in processing the query, thus improving the performance of the database system.

25 Also, fragments may be stored independently on separate disks or on separate nodes in a computer cluster or network architecture. Logically, all fragments may be scanned simultaneously, thereby increasing the overall rate at which the complete table can be read, subject to the limitations of the physical storage of the data.

30 As more and more businesses are run from mission-critical systems that store information on database systems, increasingly higher demands are being placed on these database systems to provide enterprise-wide decision support and to provide timely on-line access to critical business information through database queries.



intermediate range representation to an integer range representation. The system also contains means for building an index tree data structure to represent a search space associated with the data fragments and means for using the index tree data structure to locate a desired data fragment.

5           Among the advantages of the invention are one or more of the following. The invention provides high performance for managing data fragments in a database. The invention can manage fragmentation schema of arbitrary number of columns. This property allows the invention to manage large databases where the number of columns used for fragmentation can become quite large.

10           During the construction of the tree, an integer array is used to represent the range structure associated with the fragments. When the tree is used to locate fragments, the index of the array is used to do the search. This is both simple and efficient, because integer comparisons are computationally "cheaper" than SQL type comparisons. In addition, the modeling and mapping from SQL data-type to integer  
15           contribute to the simplicity and efficiency of the invention in performing operations with data fragmentation. Overlapping fragments are pruned from the tree, thus improving search performance. Moreover, the height of the tree generated by the invention is minimized. Thus, searching the tree will be fast and efficient. This will make insertion into the database and query optimization (eliminating unnecessary  
20           fragments) operations more efficient. Also, the system has a low data storage requirement. Due to these advantages, the invention can efficiently manage data fragments in a database.

          Other features and advantages of the invention will become apparent from the following description and from the claims.

25

#### Brief Description of the Drawings

          Fig. 1 is a diagram illustrating a database engine in accordance with the present invention.

          Fig. 2A is a flowchart illustrating a process for creating a tree data  
30           structure.

Fig. 2B is a flowchart illustrating a process for applying the tree data structure in response to a query.

Fig. 3 is a flowchart illustrating a process for partitioning dataset in the data fragmentation process of Fig. 2.

5 Fig. 4. is a flowchart illustrating a process for searching a tree created in Fig. 2.

Fig. 5 is a flowchart illustrating a process for building the tree in accordance with Fig. 2.

10 Fig. 6 is a flowchart illustrating a process for optimizing data fragmentation in accordance with the present invention.

Figs. 7A-7E are diagrams illustrating exemplary operations of the system of Fig. 1.

15 Fig. 8 is a block diagram of a computer system platform suitable for implementing an embodiment of a database system in accordance with the present invention.

Fig. 9 is a block diagram of a computer and computer elements suitable for use in the database engine of FIG. 1.

#### Detailed Description

20 Fig. 1 represents an embodiment of a database engine 30 in accordance with the present invention. The database engine 30 has a user application 32 that communicates with a client application programming interface (API) 34. The client API 34 in turn communicates with an SQL parser 38 that parses an SQL query with one or more SQL expressions into its constituencies, which may be stored in a binary  
25 tree data structure representing the components of the query. The output of the SQL parser 38 is provided to a code generator 40 that outputs intermediate code from the SQL expression. The intermediate code represents an execution plan for satisfying the SQL query. The code generator 40 also eliminates redundant data and performs error-checking, such as confirming the validity of table and column names. This code is then  
30 optimized by an optimizer 42, which optimizes the code using a cost-based analysis for formulating a query execution plan. The resulting optimized code associated with the

execution plan is provided to a fragment manager 44 that handles references to fragmented tables in accordance to a predetermined fragmentation scheme. The fragment manager 44 deploys an index tree, whose creation and operation are discussed in more details below. The output of the fragment manager 44 is provided to an RSAM 46 for executing the execution plan. The RSAM 46 in turn controls one or more data storage device managers such as disk managers, for example. The client API 34, SQL parser 38, code generator 40, optimizer 42, fragment manager 44 and RSAM 46 reference a library 48 and call various low-level routines such as access methods that will be executed by a processor. The result of the execution plan is provided to the user application 32.

Figs. 2A and 2B are flowcharts of processes for creating and managing data fragments in accordance with the invention. In Fig. 2, a query fragmentation expression, as generated by the user application 32, is parsed into an intermediate range representation with a range representation (step 102). Next, the intermediate range representation is converted into an N-dimensional integer space, where N is the number of columns in the fragmentation expression or scheme (step 104). An index tree is then generated to represent the search space (step 106). The tree is stored permanently as part of a database catalog (step 108). Once the tree has been constructed, the tree can be used to locate the data fragment, as discussed in more detail below.

Referring now to Fig. 2B, the application of the tree to locate data is shown. First, the query fragmentation expression is parsed into an intermediate range representation with a range representation (step 110). Next, the intermediate range representation is converted into an N-dimensional integer space, where N is the number of columns in the fragmentation expression or scheme (step 112). The tree can then be used to locate data fragments (step 114).

The embodiment of Fig. 2A-2B models data fragmentation as a spatial partitioning problem. That is, each column is considered as a dimension in N-dimensional space. Then, the SKD tree is used as an indexing mechanism. Data fragmentation can be performed for a variety of expressions: >, >=, <, <=, =, and logical operators: IN, NOT, AND, OR, ISNULL. A series of combinations with

complicated expression can be handled. The embodiment maps different SQL data types to integer values. In this process, all data points (SQL data values) used in the fragmentation expressions are collected and stored as a multi-dimensional array. Each column's data points are stored in an array. The embodiment of Fig. 2 uses the array index 0 as the index point for a NULL value. This models the ISNULL expression as a minimal point. The system also use the upper bound of the array as the index point for “positive” infinity.

In one embodiment, two kinds of node are used in the tree: index node and data node. Data nodes are leaf nodes of the tree, which contains only data items. A data item in the tree is represented by an array of ranges (low, up) that describe the space. Each pair of (low, up) describes a range in one dimension of the space.

Index nodes are internal (non-leaf) nodes of the tree. Each index node contains the following fields:

Index dimension	The dimension of which the index key is in.
Index key	The index key value (k) of the index node.
Space covered	Each index is responsible for a subspace, this field describes the space covered by all the data items under this index node.
Index bitmap	This field is optional and represents fragments that are covered by this node. This makes marking a fragment as active much easier.
Left	This field points to the SKD-tree that represents the subspace where all data items' value in the D dimension are smaller than (or equal to) the index key K.
Right	This field points to the tree that represents the subspace where all data items' value the D dimension are larger than (or equal to) the index key K.
Equal flag	This flag indicates which subtree represents the subspace whose values in the D dimension are equal

to the index key K. It is either Left or Right. By default, the Equal flag is set to Right.

The process starts with a data set S. Each data item of S describes a subspace in the m-dimensional space and is of the form:

5 (lb<sub>1</sub>, low<sub>1</sub>, up<sub>1</sub>, ub<sub>1</sub>), ..., (lb<sub>m</sub>, low<sub>m</sub>, up<sub>m</sub>, ub<sub>m</sub>)

where each (lb<sub>i</sub>, low<sub>i</sub>, up<sub>i</sub>, ub<sub>i</sub>) describes a range in the I-th dimension, lb<sub>i</sub> and ub<sub>i</sub> indicates whether the range includes the low point and/or the up point (1 < I < m). In addition, a multi-dimensional array V is used where V[i] records the number of points and their values of the ith dimension in the data set S.

10 The data set is then evenly partitioned with a chosen index key value so that a balanced tree can be built. This is achieved by first selecting the dimension that has the most variance -- the most points -- as the indexing dimension. The median point of that dimension is then used as the index key value to partition the data set.

15 After the index key value has been selected, the data set is partitioned based on that key. Where the index dimension is D, the index key value is K and its associated Equal flag is E, the following operations are performed. For each data items' dimensional value d of the form (low, up, flag) in the data set, a partitioning routine is executed. Figure 3 shows the partition algorithm. The process of Fig. 3 partitions the data set into two parts, the left part L, and the right part R. The process  
20 120 calls a routine append to add a data item into a data set, and a routine split to split a data item. In splitting a data item, copies of the data item are made, and the data item is partitioned according to the partitioning dimension and key by modifying the corresponding columns' range in the new data items. The new data items need to have the right flags in the partitioning dimension.

25 Also, the data items in the leaf nodes of the tree should not overlap with each other. This is achieved by adding an overlapping checking when a data item is added into a partition. If the subspace represented by the data item is already completely covered by some data item in the partition already, then that data item is not added to the partition. This effectively eliminates the overlapping subspaces in the  
30 leaf nodes of the tree.

Turning now to Fig. 3, a process 120 for partitioning a dataset is shown. This process is needed in managing the fragmentation scheme since each record can only reside in one fragment. Hence, the dataset has to be partitioned. In the process 120, each data item D of the form (low,up,flag) in the dataset, the process 120 is  
5 executed. The process 120 partitions the dataset into two parts, a left part L and a right part R.

In Fig. 3, the up attribute of each data item D is compared against an index key value K (step 122). If the up attribute is less than K or if the up attribute is equal to K and the range does not including the up boundary (up\_boundary = 0), the data  
10 item is placed into a left partition (step 122). Alternatively, the process 120 compares a low attribute for each data item D against K (step 124). If the low attribute is greater than K or if the low attribute is equal to K and the range does not including the low boundary (low\_boundary = 0), the process 120 places the data item in a right  
15 partition (step 126). Alternatively, if the low attribute is equal to the up attribute or equal to K (step 128), the process 120 determines whether an equal flag E has been sent to indicate the left partition (step 130). If so, the data item is placed in the left partition (step 132). From step 130, if E is not equal to left partition, the process 120 sets E to the right partition and places the data into the partition using an append call (step 134). From step 122, 126, 132 or 134, the process 120 exits.

20 From step 128, if the low attribute is not equal to the up attribute, or if the low attribute is not equal to K, the process 120 splits the data item (step 136) and adds the data item to both the left and right partitions, respectively using an append function (step 138) and exits.

Pseudo-code for the flowchart of Fig. 3 is as follows:

25 If  $up < K$  or ( $up = k$  and non-inclusive), place the data item to the left partition with an `append(L, d)`.

Else if  $low > K$  or ( $low = k$  and non-inclusive), then it is put in the right partition with an `append(R, d)`.

30 Else if  $low = up = K$ : if Equal flag  $E = \text{Left}$ , then partition the data item to the left by call `append(L, d)`; otherwise, set E to Right and `append(R, d)`.

Otherwise,  $low \leq K \leq up$  but  $low$  is not equal to  $up$ . In this case, split the data item so that part of it goes to the left  $L$  and part of it goes to right  $R$ , as follows:

- (a) split the data item to two parts,  $ld$  and  $rd$  by calling the split routine,  $split(d, ld, rd, D, K, E)$ .
- (b) add them into the left and the right partition respectively,  $append(L, ld)$  and  $append(R, rd)$ .

Each data node contains only one data item in an implementation of the tree for managing fragmentation. As such, each record can only reside in one fragment. In this embodiment, the data set needs to be further partitioned. When the point count is less than 2, only certain index key values can be used. For example, ranges  $(x1, x2)$  and  $(x1, x2]$  can not be partitioned using the process 120 (The "]" represents the condition where the upper boundary is included). In this case, additional rules for further partitioning need to be applied. In this case, an appropriate right column needs to be selected. The column on which the ranges of different data items have different flags should be chosen as the indexing column. This is from the observation that two ranges  $(1 < x < 5, 1 < y < 5)$  and  $(1 < x < 5, 1 < y \leq 5)$  can only be partitioned into two if  $y$  is selected as the indexing column. The Equal flag needs to be set so that correct partitioning can be achieved. The heuristics is as follows:

```

/* maximum is two points in each column
* have to be careful about picking up the split key
* otherwise recursion won't terminate.
* suppose two points are  $x1 < x2$ , the follow scheme will
* be followed:
*       $x2$  (up)
*      /\
*      < / \ (>=)
*       $x1$ (lw)  [ $x2, x2$ ]
*      /\

```

\* <= / \ >  
\* [x1,x1] (x1, x2)  
\*/

5 Referring now to Fig. 4, a process 150 for searching the tree created in Fig. 2 is detailed. First, the process 150 starts at the root of the tree (step 152). The process 150 then compares a search key value with an index key (step 154). Next, the process 150 compares the key value against the index key (step 156). If the key value is greater than the index key, the process 150 follows a pointer and proceeds down the right tree (step 160). Alternatively, if the key value is less than or equal to the index  
10 key, the process 150 follows the pointer and goes down the left tree (step 158).

When searching the tree, the process starts from the root of the tree, compares the searched key value with the index key, and follows the pointer according the result of the comparison. At the data node level, the search point is tested to see if  
15 it is covered by the subspace in that node.

When searching a particular range, the process first compares the subspace searched with the space covered by an index node (starting from the root). If the searched subspace covers the space described by the index node, all fragments covered by the index node are activated using a bitmap field in the index node. Range  
20 searches may have to do partitioning, as described in Fig. 3, and thus may follow multiple paths in the tree. If the search space is not completely covered by the nodes in the tree after the process has searched down to the leaf level, if a remainder fragment exists, the remainder fragment is activated.

Referring now to Fig. 5, a process 180 for building the tree used in Fig. 2 is  
25 detailed. The basic assumption in the process 180 is that the data set is relatively static with few insertions, deletions and updates on the data set. The most often operation on the data set is look-up (search). The process selects a partition value (the index key value) so that data sets are evenly partitioned.

First, the process 180 determines whether the dataset contains no (zero)  
30 item (step 182). If so, the process 180 simply exits. Alternatively, if the dataset contains more than zero item, the process 180 checks whether the dataset contains

only one item (step 184). If so, the item is returned at the data node of the tree (step 186) and the process 180 exits. Alternatively, if more than one item exist in the dataset, the process 180 selects a predetermined dimension to evenly partition the dataset, as previously described (step 188). Next, it builds the index node (step 190).  
5 The process 180 then builds a left tree (step 192) and also builds a right tree (step 194). The index node is then filled (step 196) by constructing a bitmap in one embodiment. The completed tree is then returned (step 198) before the process of Fig. 4 exits.

Turning now to Fig. 6, an optimization process 200 is shown. First, the  
10 process 200 maps a user query space into an integer space (step 202). Next, the process 200 determines fragments associated with the overlapping integer space (step 204). Finally, the process 200 looks up the fragments that overlap the integer space (step 206). The mapping of the user query space into integer space is advantageous since only integer operations need to be performed. In contrast, the conventional user  
15 query space may include a variety of data types that can be complex.

Applying the operations of Figs. 2-6 to a table that is fragmented by range expressions on  $m$  columns with  $N$  fragments, on average the height of the tree is of  $O(\log N)$ . During a record insertion, the search time to find the correct fragment should be of  $O(m + \log N)$  because it takes  $O(\log N)$  to get down to a leaf node and  
20  $m$  comparisons to decide if the record belongs to that fragment. Thus, performance is enhanced using the tree. Without the tree, the system is expected to perform an average of  $N/2$  comparisons to find a fragment. It is to be noted that in a special case where only one dimension is partitioned, the tree becomes a binary search tree.

For range query elimination, performance is also improved due to the  
25 realization that, if the search space covers the space described by an index node, all fragments described by the index node using the bitmap in the index node are activated. This short-cut saves considerable search cost.

The invention will not be optimal under a worst case scenario where all ranges overlaps each other on all the dimensions, but no one is covered by the other.  
30 In this scenario, the construction of the tree will generate more splits, thus adversely

affecting the performance. As this is an unlikely situation, the tree should provide good performance overall.

Certain data structures used in one implementation of the invention will be discussed next. First, the structure associated with `dimension_t`, representing a range  
5 in a column, is shown below:

```
/* range in a column represented by relatively index of the points in that column*/
```

```
typedef struct dimension
```

```
{
```

```
10 char dm_lb;          /* 1 to include lower bound*/
```

```
char dm_ub;          /* 1 to include upper bound */
```

```
int dm_lw;          /* low index of the range */
```

```
int dm_up;          /* up index of the range */
```

```
} dimension_t;
```

15

Next, the data structure for `splist`, a linked list of subspaces is shown. Each subspace is an array of `dimension_t`. Each subspace represents a fragment.

/\* subspace represented by the range expression indexed by indices in columns \*/

typedef struct splist

```
{  
    long sp_fragid;  
5    dimension_t *sp_range;  
    struct splist *sp_next;  
} splist_t;
```

10 Next, skdtree\_t is the tree data structure. It has a flag to indicate whether  
the node is an index node or a data node. skd\_internal\_t is the index node data  
structure. skd\_data\_t is the data structure for data node.

typedef struct skdtree

```
{  
15    char skd_flag;          /* index node or data node */  
    void *skd_content;     /* either index or data node depends flag*/  
} skdtree_t;
```

```
#define SKD_INDEX          0x00
```

```
20 #define SKD_DATA          0x01
```

```
#define DATANODEP(node)   ((node->skd_flag) & SKD_DATA)
```

```
#define datanode(tree)    ((skdt_data_t *) tree->skd_content)
```

```
#define indexnode(tree)  ((skdt_internal_t *) tree->skd_content)
```

25

```
typedef struct skdt_internal
```

```
{
```

```
long    skd_colid;          /* the column id of the idx value  
                                * ranges from 0 to number of cols - 1 */
```

```
5    int    skd_value;      /* the index value of the point */
```

```
maptype4_t *skd_map;      /* space covered by this node */
```

```
dimension_t *skd_space;
```

```
short    skd_eq;          /* where equal value goes,  
                                * default(right:1) */
```

```
10   struct skdtree *skd_left;
```

```
struct skdtree *skd_right;
```

```
    } skdt_internal_t;
```

```
typedef struct skdt_data
```

```
15   {
```

```
        long skd_fragid;
```

```
        dimension_t *skd_space;
```

```
    } skdt_data_t;
```

```
20
```

/\* colarr\_t is used for Range Fragment Elimination and stores all the constant points in  
 /\* one column.\*/

typedef struct colarr

```

    {
5      int   col_id; /* column id of this column */
      int   col_dtttype; /* datatype of the column */
      int   col_npts; /* number of points (size) in colarr_t*/
      value_t *col_pts; /* the array of points on this column */
    } colarr_t;

```

10

frangrange\_t is the fragment elimination information data structure. It stores certain metadata information and a pointer to the tree. It is also the accessing point for fragment management information.

15

/\* This structure will be stored on the disk and will contain information  
 \* required to do Range Fragment Elimination.  
 \*/

typedef struct frangrange

```

{
20      int   fr_ncol; /* number of columns referenced in
                    * the fragment expressions.
                    */
      int   fr_remainder; /* set TRUE if REMAINDER */
      colarr_t *fr_col; /* one point array for each column */
25      skdtree_t *fr_skdt; /* skdtree of ranges (subspace) */
    } frangrange_t;

```

25

In order to save the storage and the time of comparison, all point values (value\_t structure) are converted to an integer value. The integer value is an index in the  
 30 colarr\_t->col\_pts array.

30

Moreover, `keyarr_t` is the data structure used during construction of the tree. It records the `base_key` of an key array and the number of keys in the array. This is used during partitioning, both the `base_key` and the `cnt` variables are modified according to the partitioning.

5

```
typedef struct keyarr
{
    int base_key;      /* the base key value */
    int cnt;          /* number of keys in the column
10                    * starting from base_key, base_key + 1, .. +cnt
                    */
} keyarr_t;
```

15

Turning now to Figs. 7A-7D, exemplary operations on a simplified expression with a logical operator is discussed next. In this example, the simplified expression is:

<simple\_exp> <logical operation> <simple\_exp>

where

20 <simple\_exp> => <column> <operation> <constant>

In this example, a table is created using the following command:

```
CREATE TABLE T( a char(10), b date)
```

25 In this case, the column a is of character type (char) and can store a maximum of ten characters, and column b stores date type of information. The database is fragmented by the following expressions for first, second and third database fragments db1, db2 and db3, respectively:

'cc' < a < 'ff' and '7-1-66' < '10-1-86' in db1

30 'dd' < a < 'LL' and '7-1-76' < b < '10-1-96' in db2

and the remainder is in db3

Referring now to Fig. 7A, the above expressions are converted into an intermediate representation. The representation is shown as a list of a two-dimensional array describing the range structure of the expressions. Note that, at this point, the two dimensional data point array is unsorted.

5 In Fig. 7B, the operation to map the intermediate representation to an integer space representation is performed. At this stage, the fragment data point array is sorted. Further, location 0 of the array is reserved for a NULL value and the data points are stored thereafter. Also, the upper bound value up stores the index value to the data point array. Since 'ff' is the third element in the sorted data point array, up stores a value of 3. The mapping operation is performed for the remaining elements in the dimension\_t array.

Figs. 7C and 7D illustrate the step of building the tree in this example. First, the system picks a column that has the maximum number of data points. In this case, a and b have the same number of data points. Hence, the system picks a as a default selection.

15 In Fig. 7C, the system splits the dimension\_t array with the index being set to column a, the key being 2 and Equal set to Right. The split creates one left partition and two right partitions. Since the left partition (lower bound = 1 and upper bound = 2) has only one data node, the system is done with the left partition.

20 With respect to the right partitions, the system picks column b since it now has most data points. After splitting with the index being set to column b the key being 2, and Equal set to Right., the tree has one left partition and two right partitions. Again, the left partition is left alone since it has one node. The right partitions are split again with the index being set to column a, the key being 3, and Equal set to Right.

25 Fig. 7D shows the result of the split of the right columns in step 3. In this case, the split results in another left partition and a right partition. Note that the partition at the bottom of Fig. 7D contains an overlapped partition. The bottom partition is then discarded. In this manner, overlapping fragments are pruned from the tree, thus improving search performance. Moreover, the height of the tree generated by the invention is minimized.

30 During operation, a user may execute a query such as:

Select from T when 'dd' < a 'ff' & '7-1-76 < b < '7-1-86'

This query is resolved into an intermediate form of Fig. 7E. In this example, the mapping process maps the intermediate representation to the integer value of 3 to represent the value of '7-1-86'. The SKD tree can then be rapidly searched using integer comparison operations on the index of the array rather than SQL comparison operations. This is both simple and efficient as compared to SQL type comparisons.

The tree thus provides high performance for managing data fragments in a database. Although the fragmentation schema represents only two columns in this example, the tree can manage large databases where the number of columns used for fragmentation can become quite large. The modeling and mapping from SQL data-type to integer contribute to the simplicity and efficiency of the system in performing operations with data fragmentation.

Fig. 8 illustrates a computer system 400 that is a suitable platform for supporting a relational database system and storing relational database tables, which will be referred to simply as tables. The computer system 400 includes one or more computers 402 (individually, computers 402a and 402b). Multiple computers may be connected by a link 404, which may be high-speed backbone that creates the cluster of computers, or a local or wide-area network connection linking the computers. The computers have one or more persistent data stores 406a-406e. Typically, each data store is a storage subsystem including one or more disk drives that operate independently of the disk drives of every other data store, which are controlled through disk controllers installed in the associated computer and operated under the ultimate control of the database system.

In the database system that will be described and used for illustrative purposes, a database definition initially resides in one database storage space in which the database is placed by operation of a "create database" command to the database system. A database initially includes a set of relational tables called the system catalogs (not shown). The system catalogs describe all aspects of the database, including the definitions of all tables and the fragmentation of all tables. As new tables

are created, with “create table” commands, for example, new data is added to the system catalogs to describe the new tables.

The system catalogs include a system fragments table for persistently storing information about the fragmentation of the database. Each fragment may be represented by an individual row in the system fragments table. When the system needs to refer to fragments, it can run queries against the system fragments table to obtain the necessary fragmentation information for any given table. One attribute of the system fragments table is the fragmentation method: a table that is fragmented using a referential fragmentation scheme, described later in this specification, will have an attribute value such as “reference” that identifies the fragment as one that was created with a referential fragmentation scheme. The referential key information that is used by a referential fragmentation scheme is also stored in a table in the system catalogs.

Each data store may store one or more fragments 408a-408i of one or more tables managed by the database system. It is generally advantageous not to split fragments across data storage subsystems that can be operated in parallel.

Shown in FIG. 9 is a block diagram of a computer 1002 suitable for use in the computer system platform described earlier with reference to FIG. 1. The invention may be implemented in digital electronic circuitry or in computer hardware, firmware, software, or in combinations of them. Apparatus of the invention may be implemented in a computer program product tangibly embodied in a machine-readable storage device for execution by a computer processor; and method steps of the invention may be performed by a computer processor executing a program to perform functions of the invention by operating on input data and generating output. Suitable processors 1020 include, by way of example, both general and special purpose microprocessors. Generally, a processor will receive instructions and data from a read-only memory 1022 and/or a random access memory 1021. Storage devices suitable for tangibly embodying computer program instructions include all forms of non-volatile memory, including by way of example semiconductor memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic tapes; magnetic disks such as internal hard disks and removable disks 1040; magneto-optical disks; and

CD-ROM disks. Any of the foregoing may be supplemented by, or incorporated in, specially-designed ASICs (application-specific integrated circuits).

Other embodiments are within the scope of one or more of the following claims.

5

**WHAT IS CLAIMED IS:**

1. In a relational database system, a computer-implemented method for managing data fragments in a database, where data fragments are the result of data fragmentation which divides a given relation of a relational database into pieces or fragments, so that the fragments can be stored where they are used most frequently, the method comprising:
  - 5 converting a query expression to a multi-dimensional intermediate range representation, which is a list of a two-dimensional array describing the range structure of the query expression;
  - 10 mapping the intermediate range representation to an N-dimensional integer space, where N is the number of columns in the query expression, to thereby create an integer range representation;
  - building an index tree data structure to represent a search space associated with the data fragments, using the integer range representation; and
  - 15 using the index tree data structure to locate a desired data fragment.
2. The method of claim 1, where the integer range representation is independent of Structured Query Language (SQL) data types.
- 20 3. The method of claim 1 or claim 2, further comprising partitioning a data set into even segments to balance the tree data structure.
4. The method of claim 1 or claim 2, further comprising:
  - 25 collecting data points used in one or more fragmentation expressions; and storing the data points in a multi-dimensional array.
5. The method of claim 1 or claim 2, further comprising using an upper bound of the array as an index point of positive infinity.
- 30 6. The method of claim 1 or claim 2, wherein the mapping step further comprises using a binary search to convert the intermediate range representation into the integer range representation.
7. The method of claim 1 or claim 2, wherein the index tree data structure represents a multi-dimensional search space.
  - 35
8. The method of claim 1 or claim 2, wherein the using step further comprises searching the index tree data structure with an  $O(\log(N))$  search complexity.
- 40 9. A relational database system, that carries out a method for managing data fragments in a database, where data fragments are the result of data fragmentation which divides a given relation of a relational database into pieces or fragments, so that the fragments can be stored where they are used most frequently, the relational database system comprising:
  - 45 means for converting a query expression to a multi-dimensional intermediate range representation, which is a list of a two-dimensional array describing the range structure of the query expression;

means for mapping the intermediate range representation to an N-dimensional integer space, where N is the number of columns in the query expression, to thereby create an integer range representation;

means for building an index tree data structure to represent a search space associated with the data fragments, using the integer range representation; and

means for using the index tree data structure to locate a desired data fragment.

10. The relational database system of claim 9 where the integer range representation is independent of Structured Query Language (SQL) data types.

11. The relational database system of claim 9 or claim 10 further comprising means for partitioning a data set into even segments to balance the tree data structure.

12. The relational database system of claim 9 or claim 10 further comprising:  
means for collecting data points used in one or more fragmentation expressions; and  
means for storing the data points in a multi-dimensional array.

13. The relational database system of claim 9 or claim 10 further comprising means for using an upper bound of the array as an index point of positive infinity.

14. The relational database system of claim 9 or claim 10 further comprising means for using a binary search to convert the intermediate range representation into the integer range representation.

15. The relational database system of claim 9 or claim 10 wherein the index tree data structure represents a multi-dimensional search space.

16. The relational database system of claim 9 or claim 10 wherein the using means further comprises an  $O(\text{Log}(N))$  search complexity means for searching the index tree data structure.

17. A computer program product stored on a computer-readable storage medium for, when run on a computer system, carrying out a method for managing data fragments in a database, where data fragments are the result of data fragmentation which divides a given relation of a relational database into pieces or fragments, so that the fragments can be stored where they are used most frequently, the method comprising:

converting a query expression to a multi-dimensional intermediate range representation, which is a list of a two-dimensional array describing the range structure of the query expression;

mapping the intermediate range representation to an N-dimensional integer space, where N is the number of columns in the query expression, to thereby create an integer range representation;

building an index tree data structure to represent a search space associated with the data fragments, using the integer range representation; and

using the index tree data structure to locate a desired data fragment.

18. The computer program product of claim 17, wherein the method further includes the integer range representation is independent of Structured Query Language (SQL) data types.
- 5 19. The computer program product of claim 17 or claim 18, wherein the method further comprises partitioning a data set into even segments to balance the tree data structure.
20. The computer program product of claim 17 or claim 18, wherein the method further comprises:  
10 collecting data points used in one or more fragmentation expressions; and  
storing the data points in a multi-dimensional array.
21. The computer program product of claim 17 or claim 18, wherein the method further comprises using an upper bound of the array as an index point of positive infinity.  
15
22. The computer program product of claim 17 or claim 18, wherein the method includes the mapping step further comprising a binary search to convert the intermediate range representation into the integer range representation.
- 20 23. The computer program product of claim 17 or claim 18, wherein the method includes the index tree data structure representing a multi-dimensional search space.
24. The computer program product of claim 17 or claim 18, wherein the method includes the using step further comprising searching the index tree data structure with an  
25  $O(\log(N))$  search complexity.

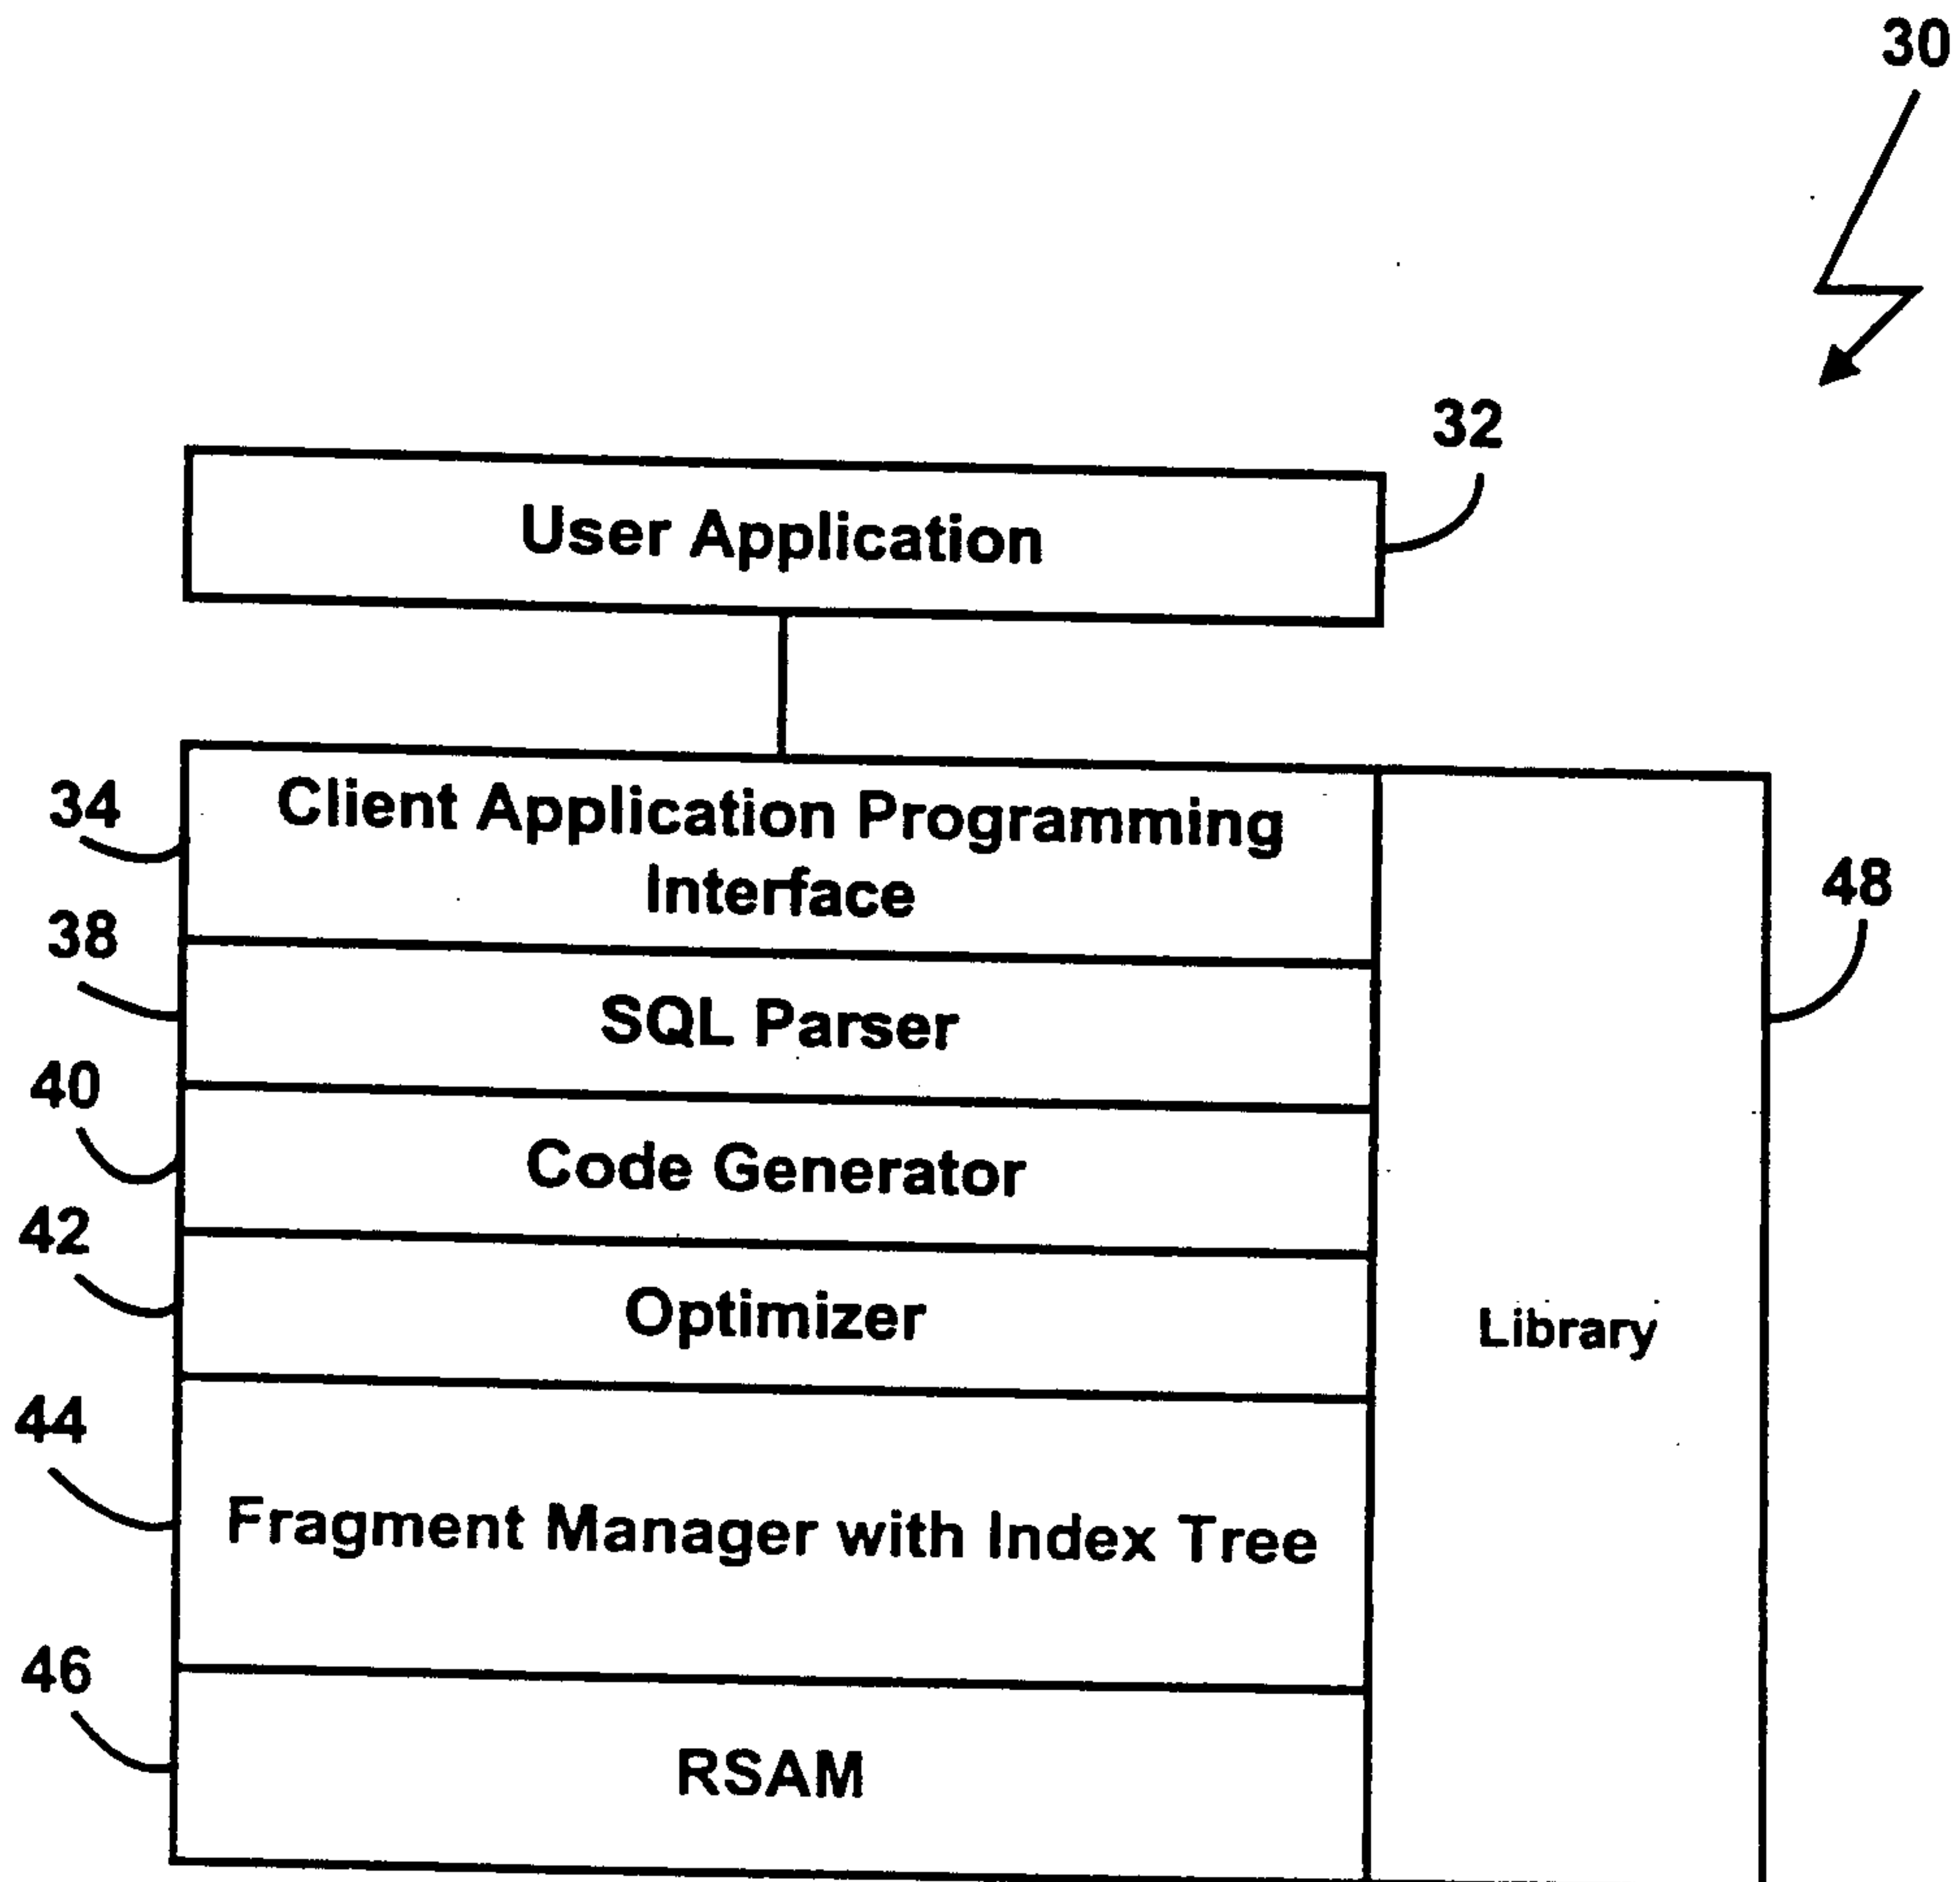


FIG. 1

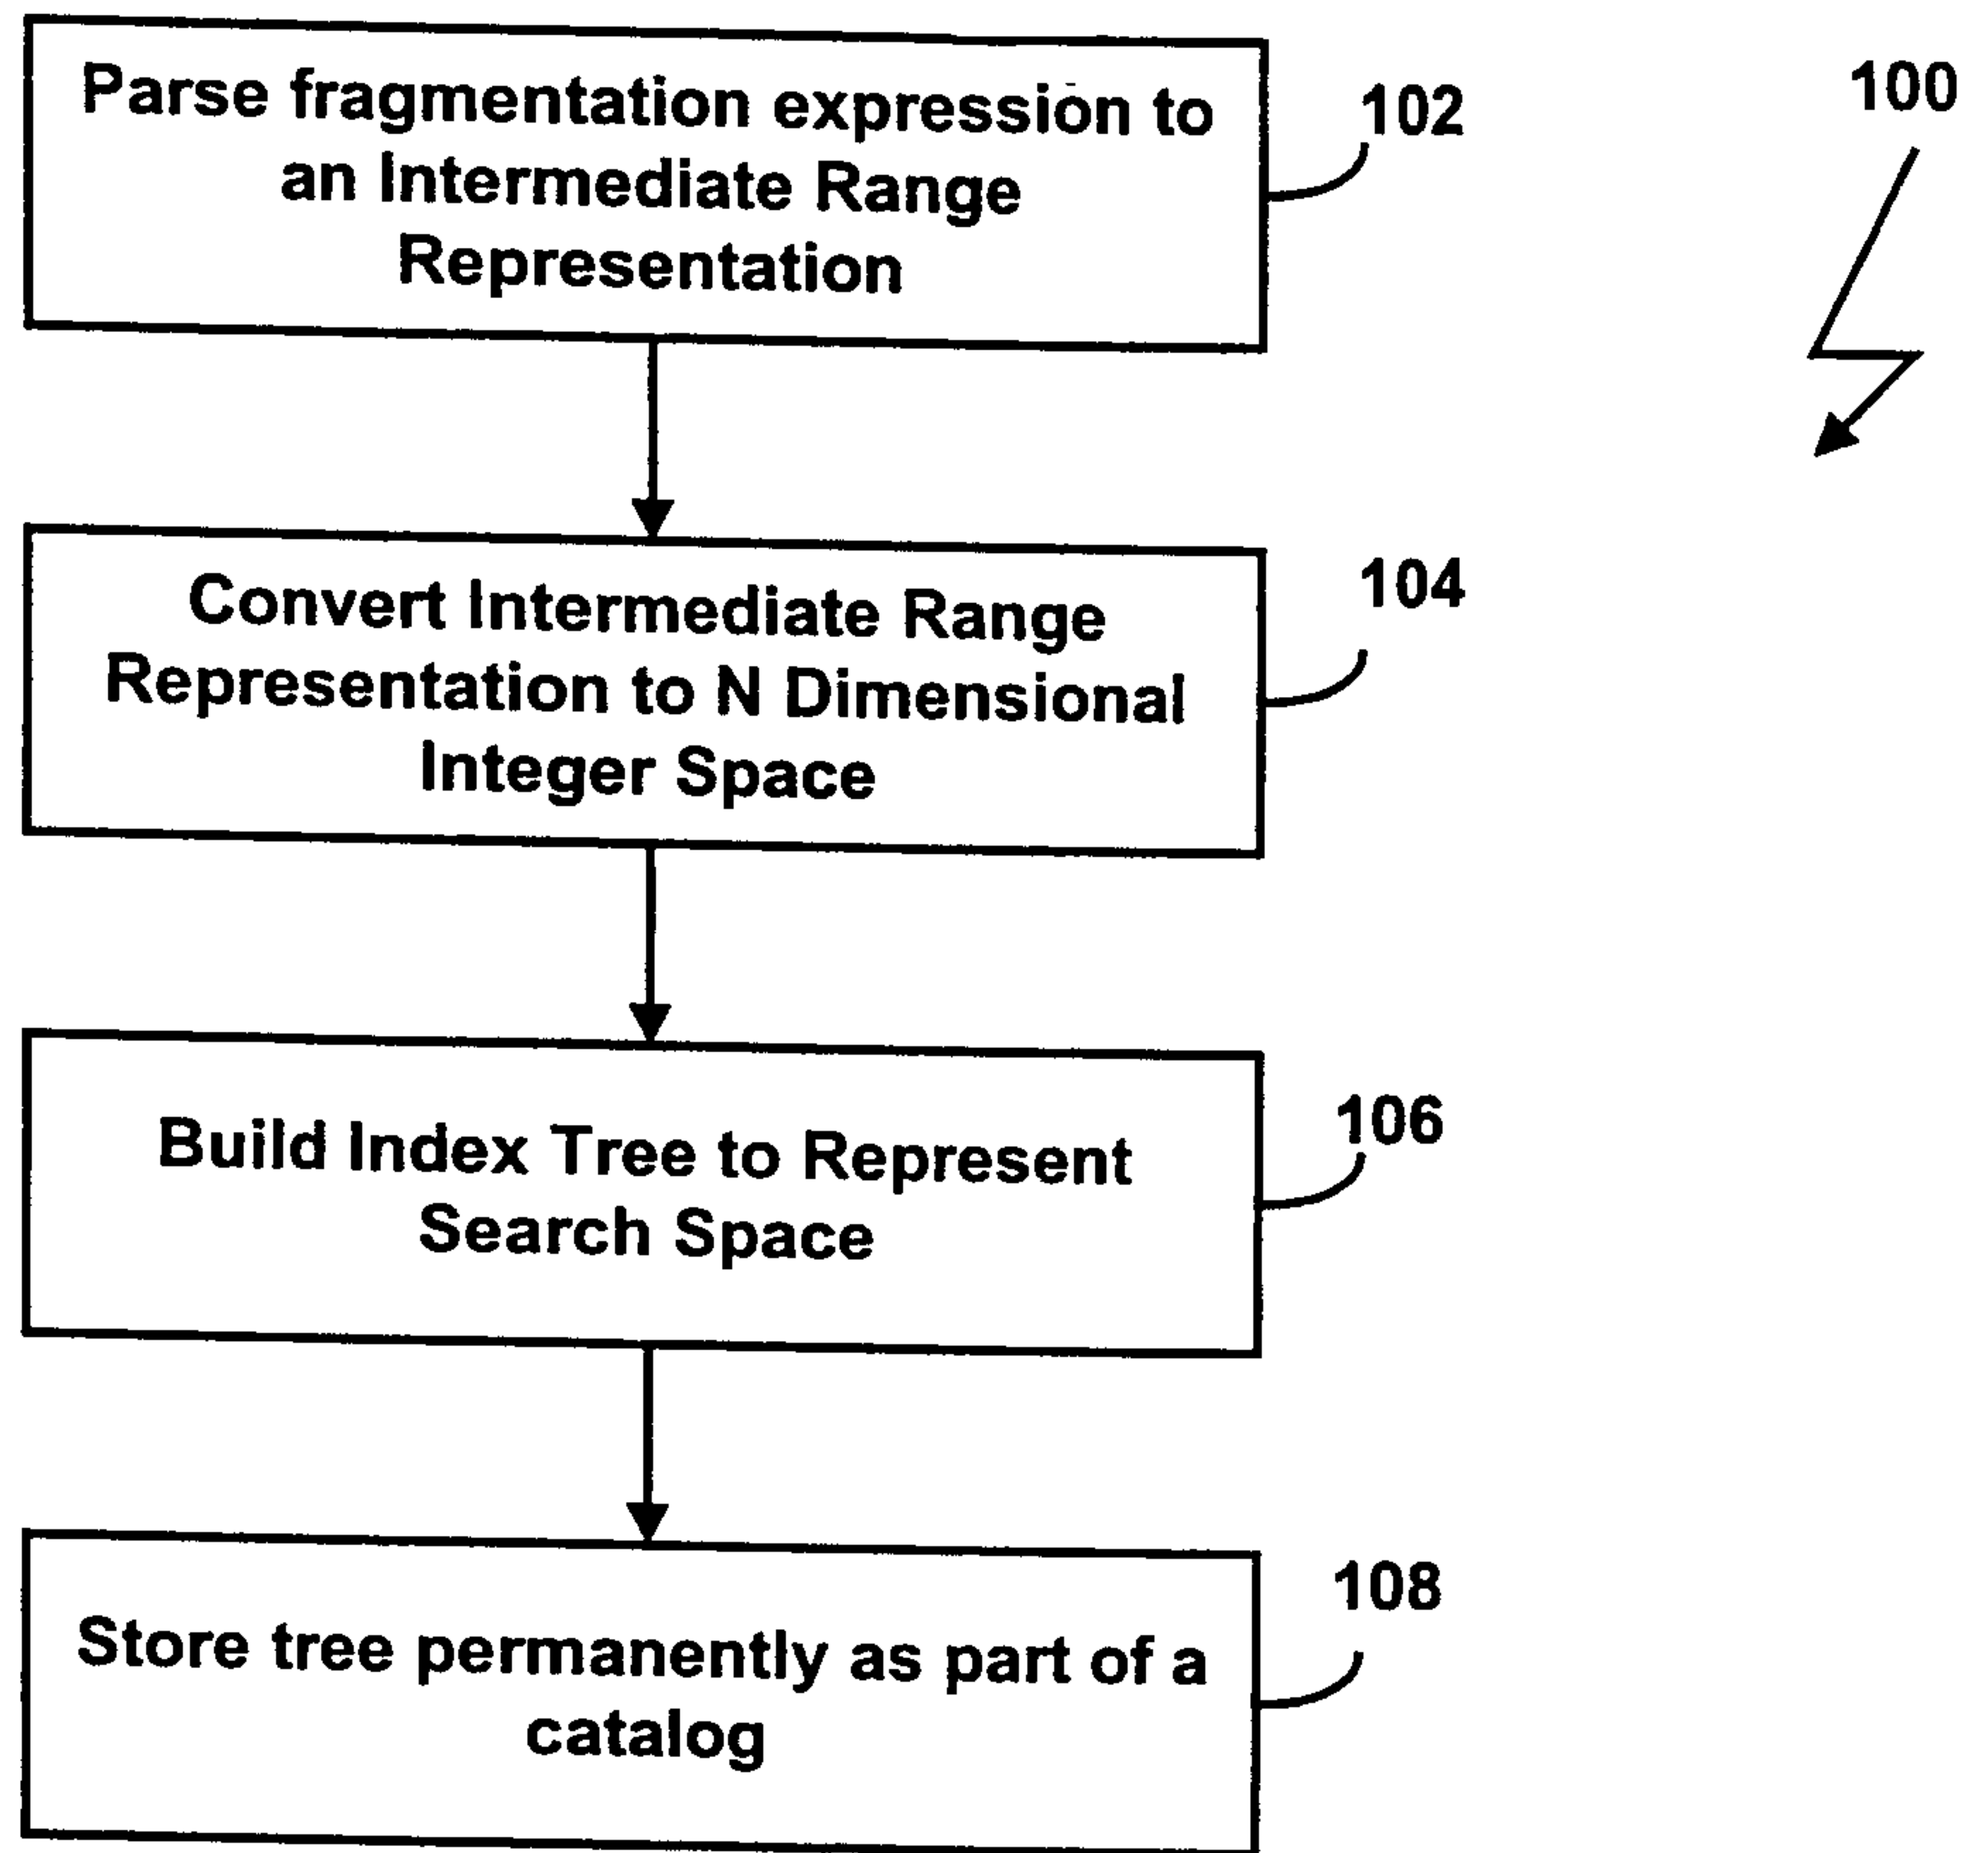


FIG. 2A

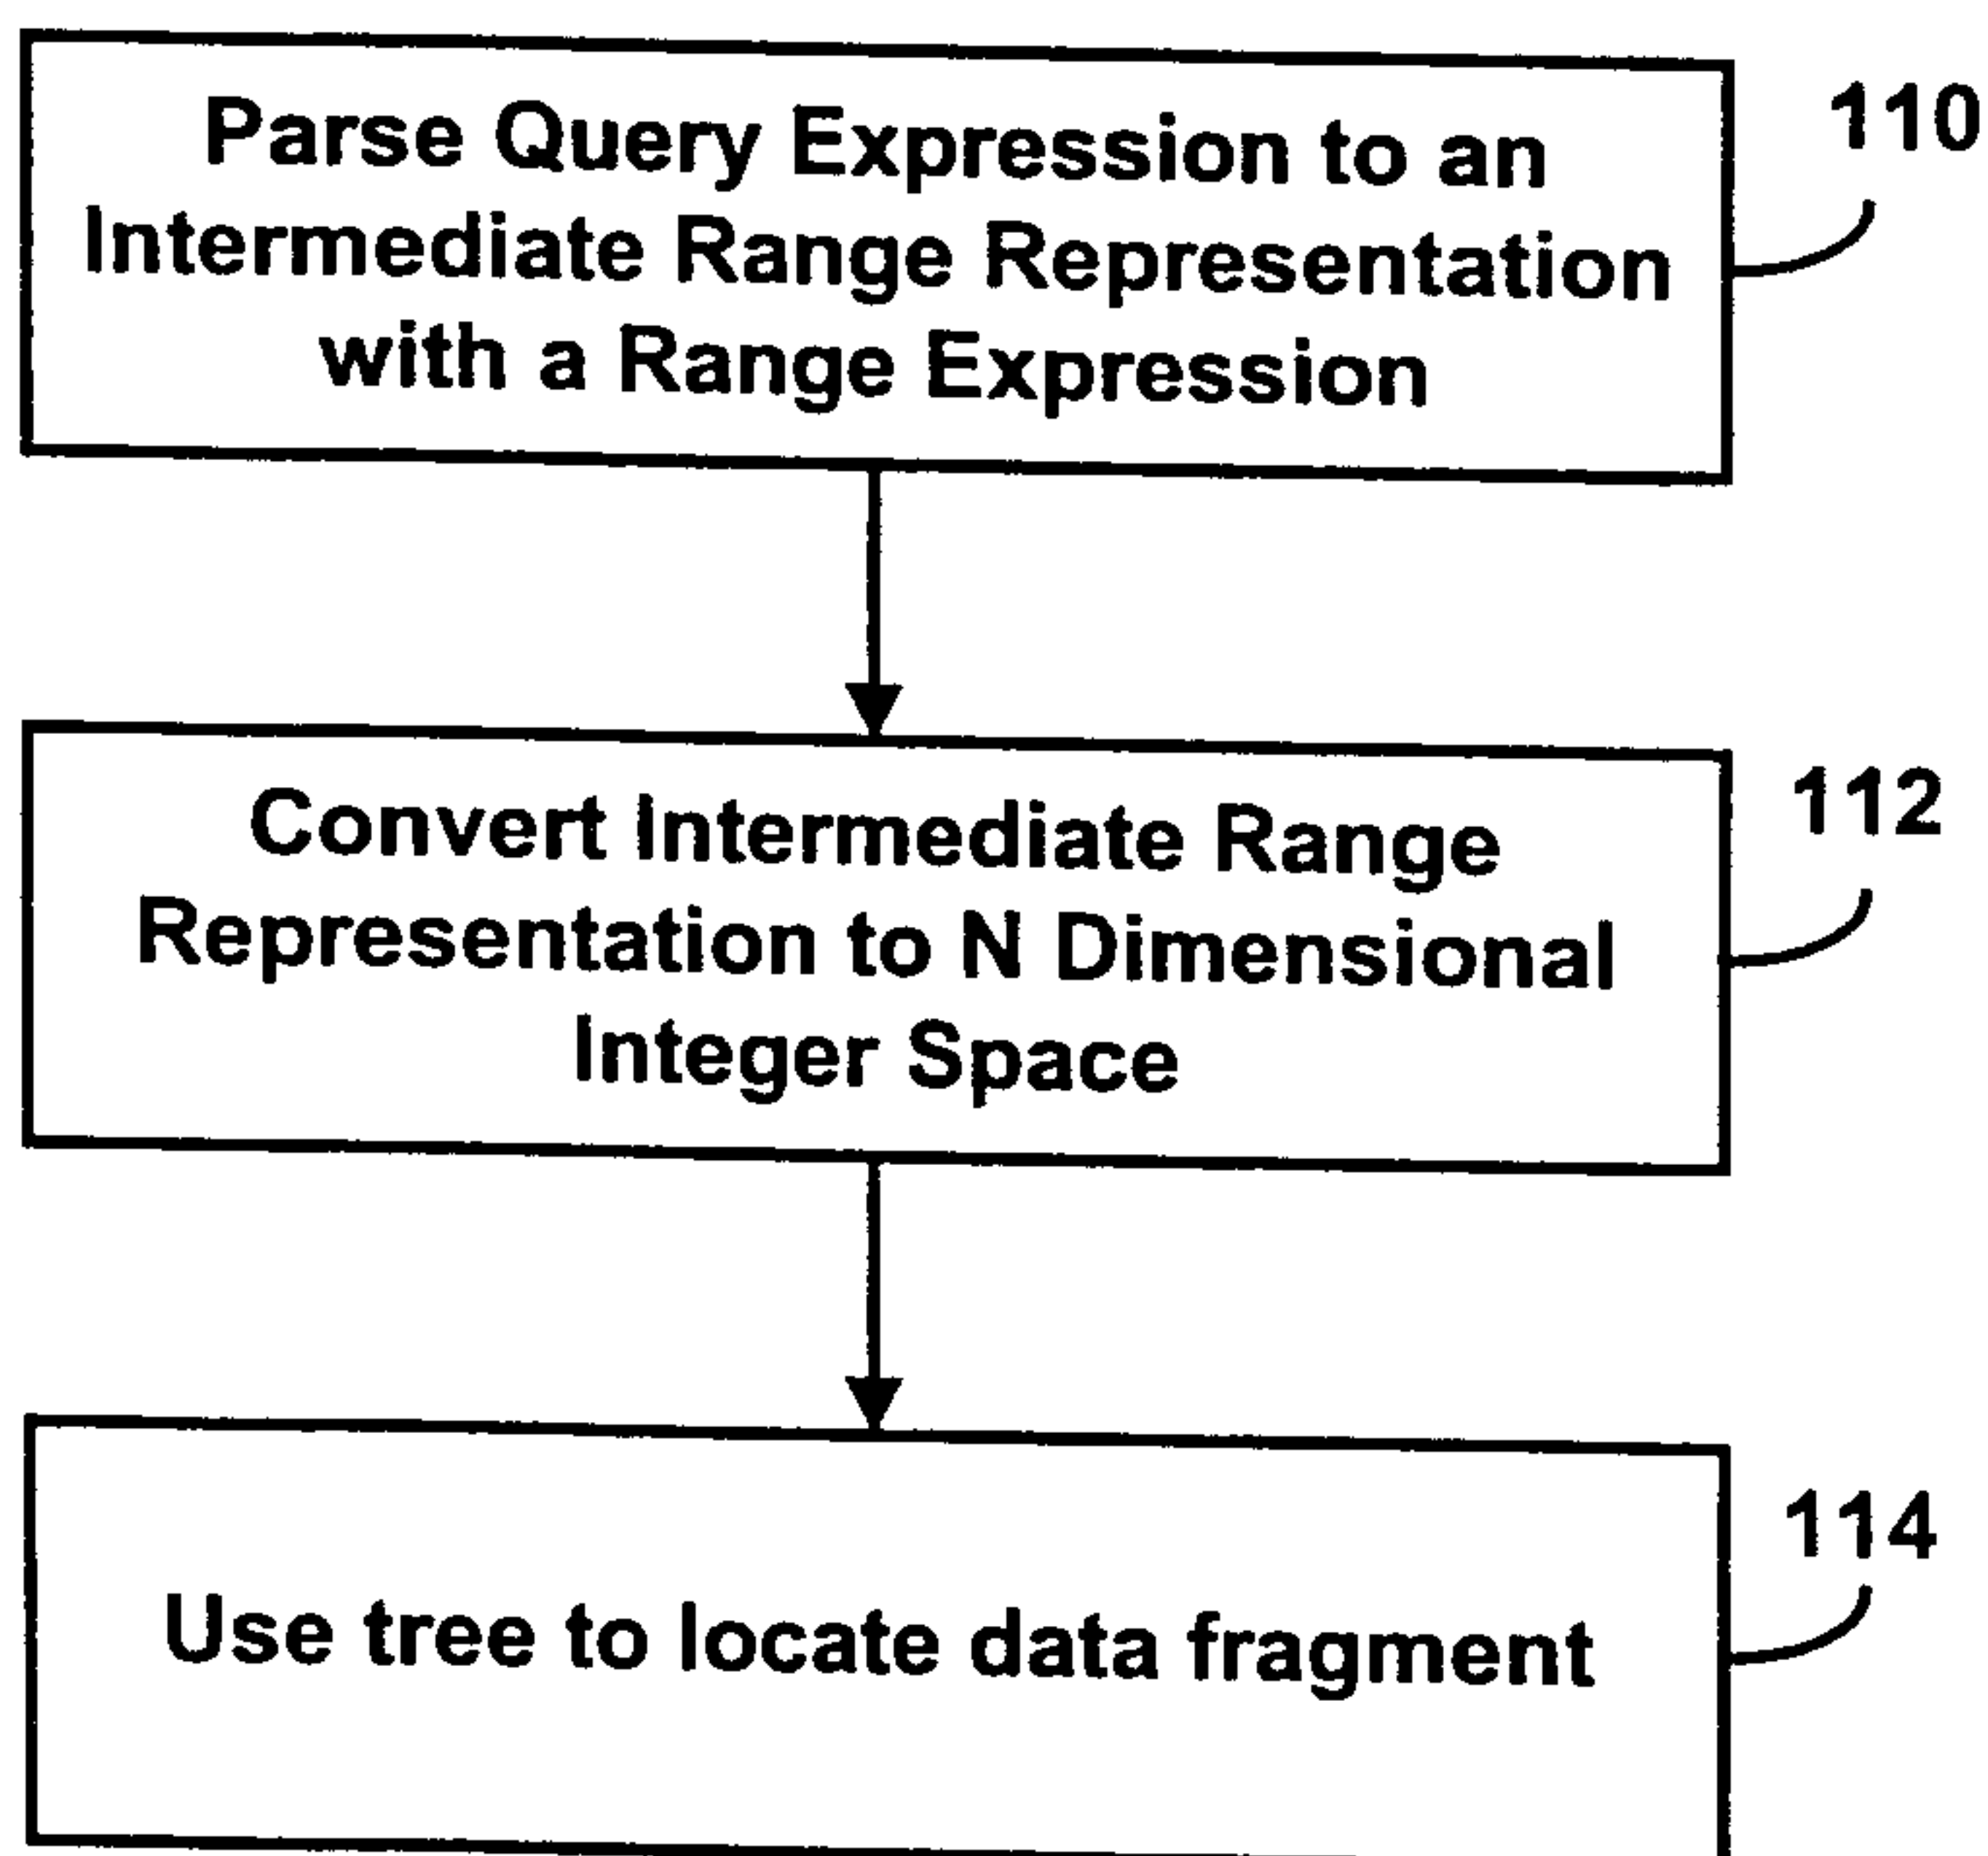


FIG. 2B

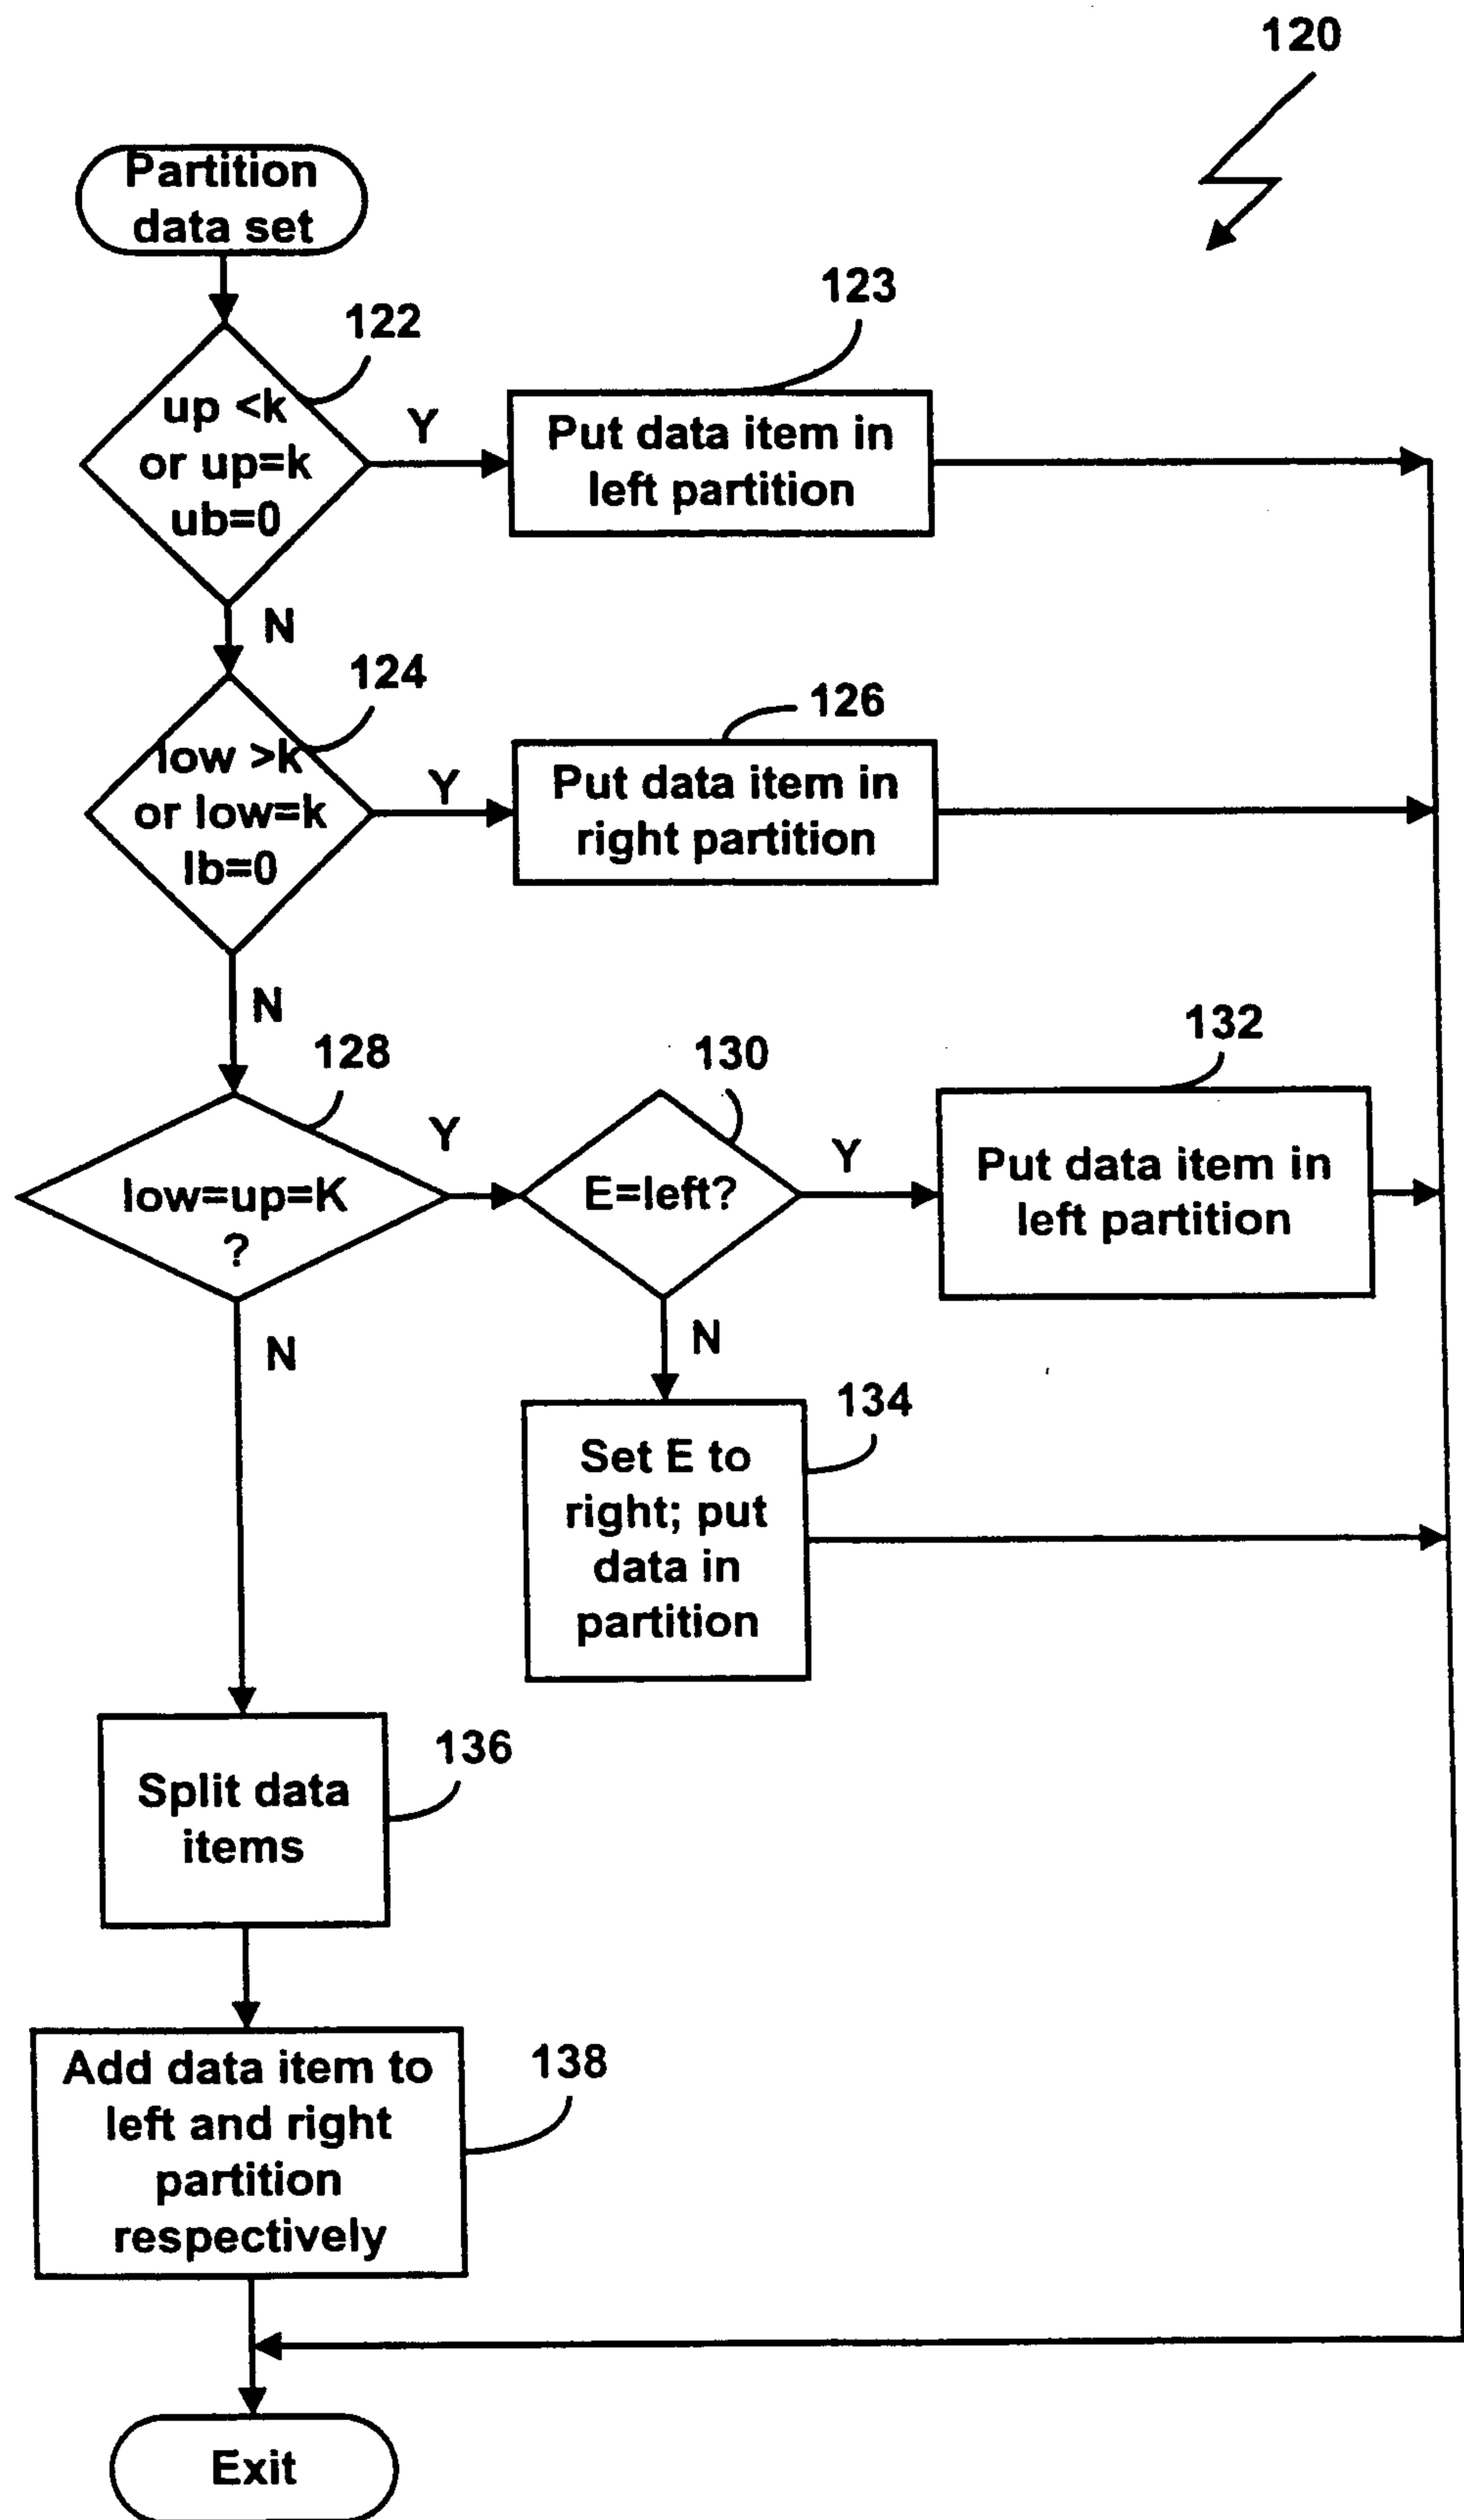


FIG. 3

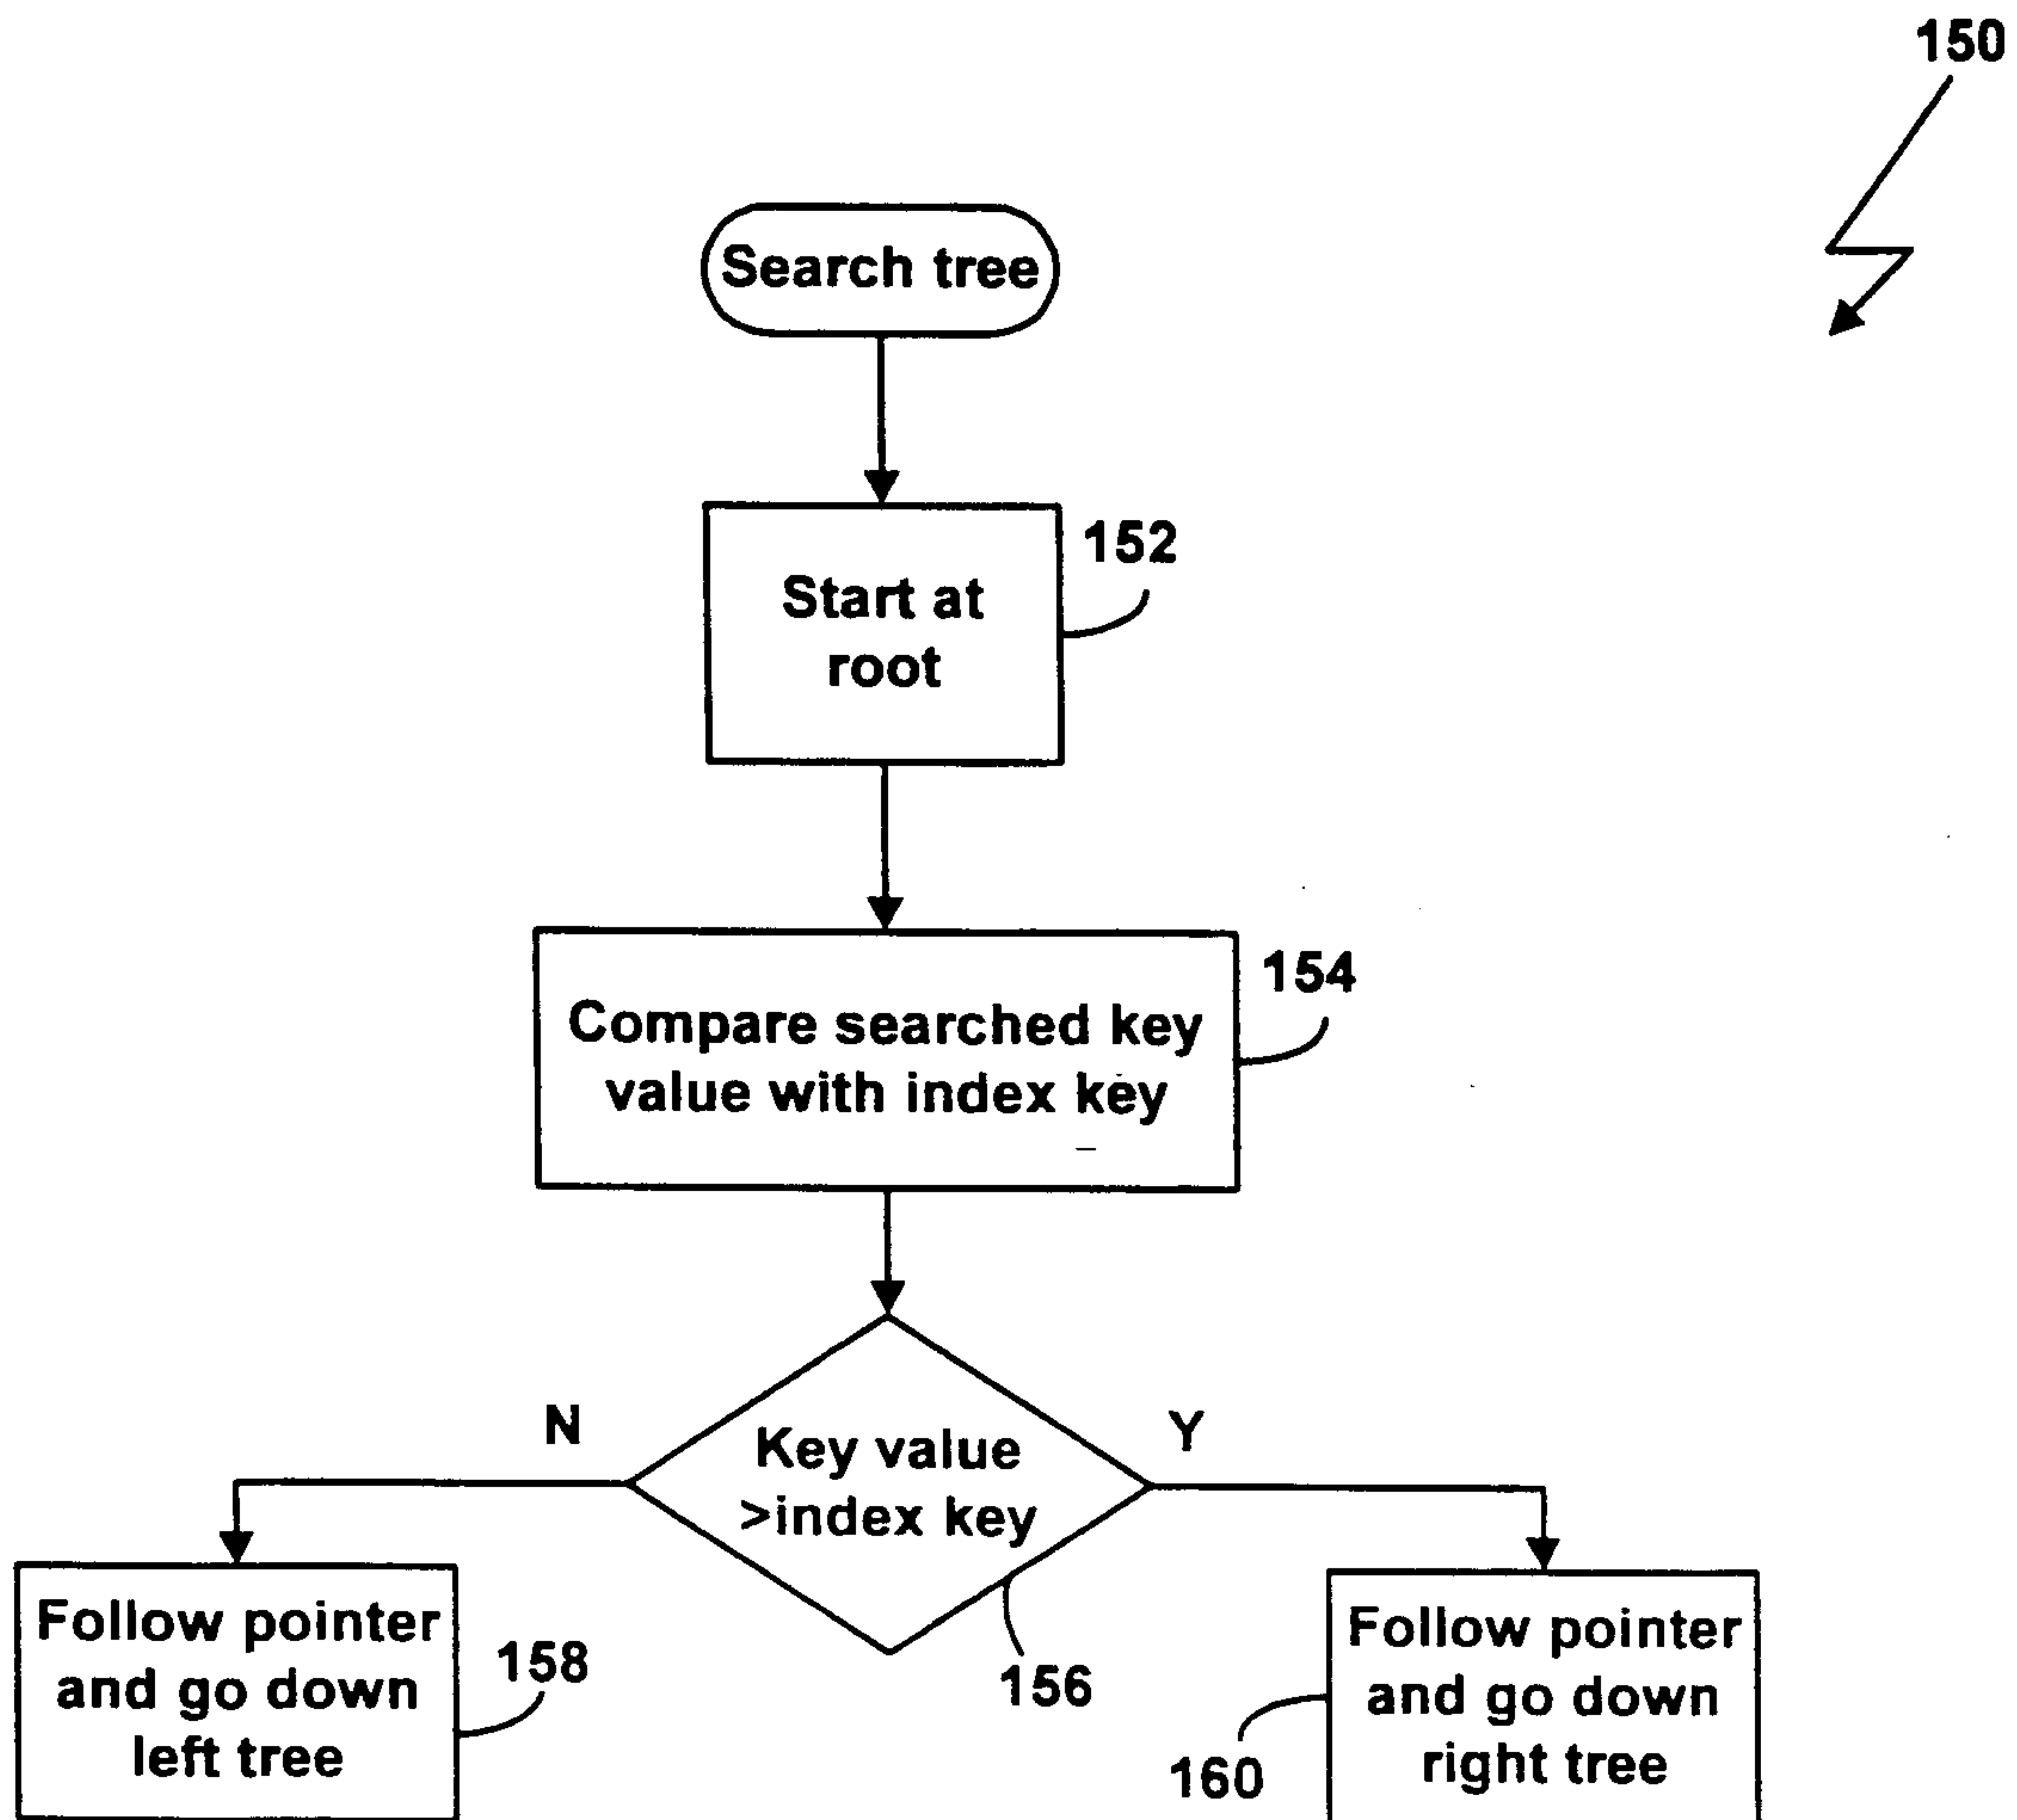


FIG. 4

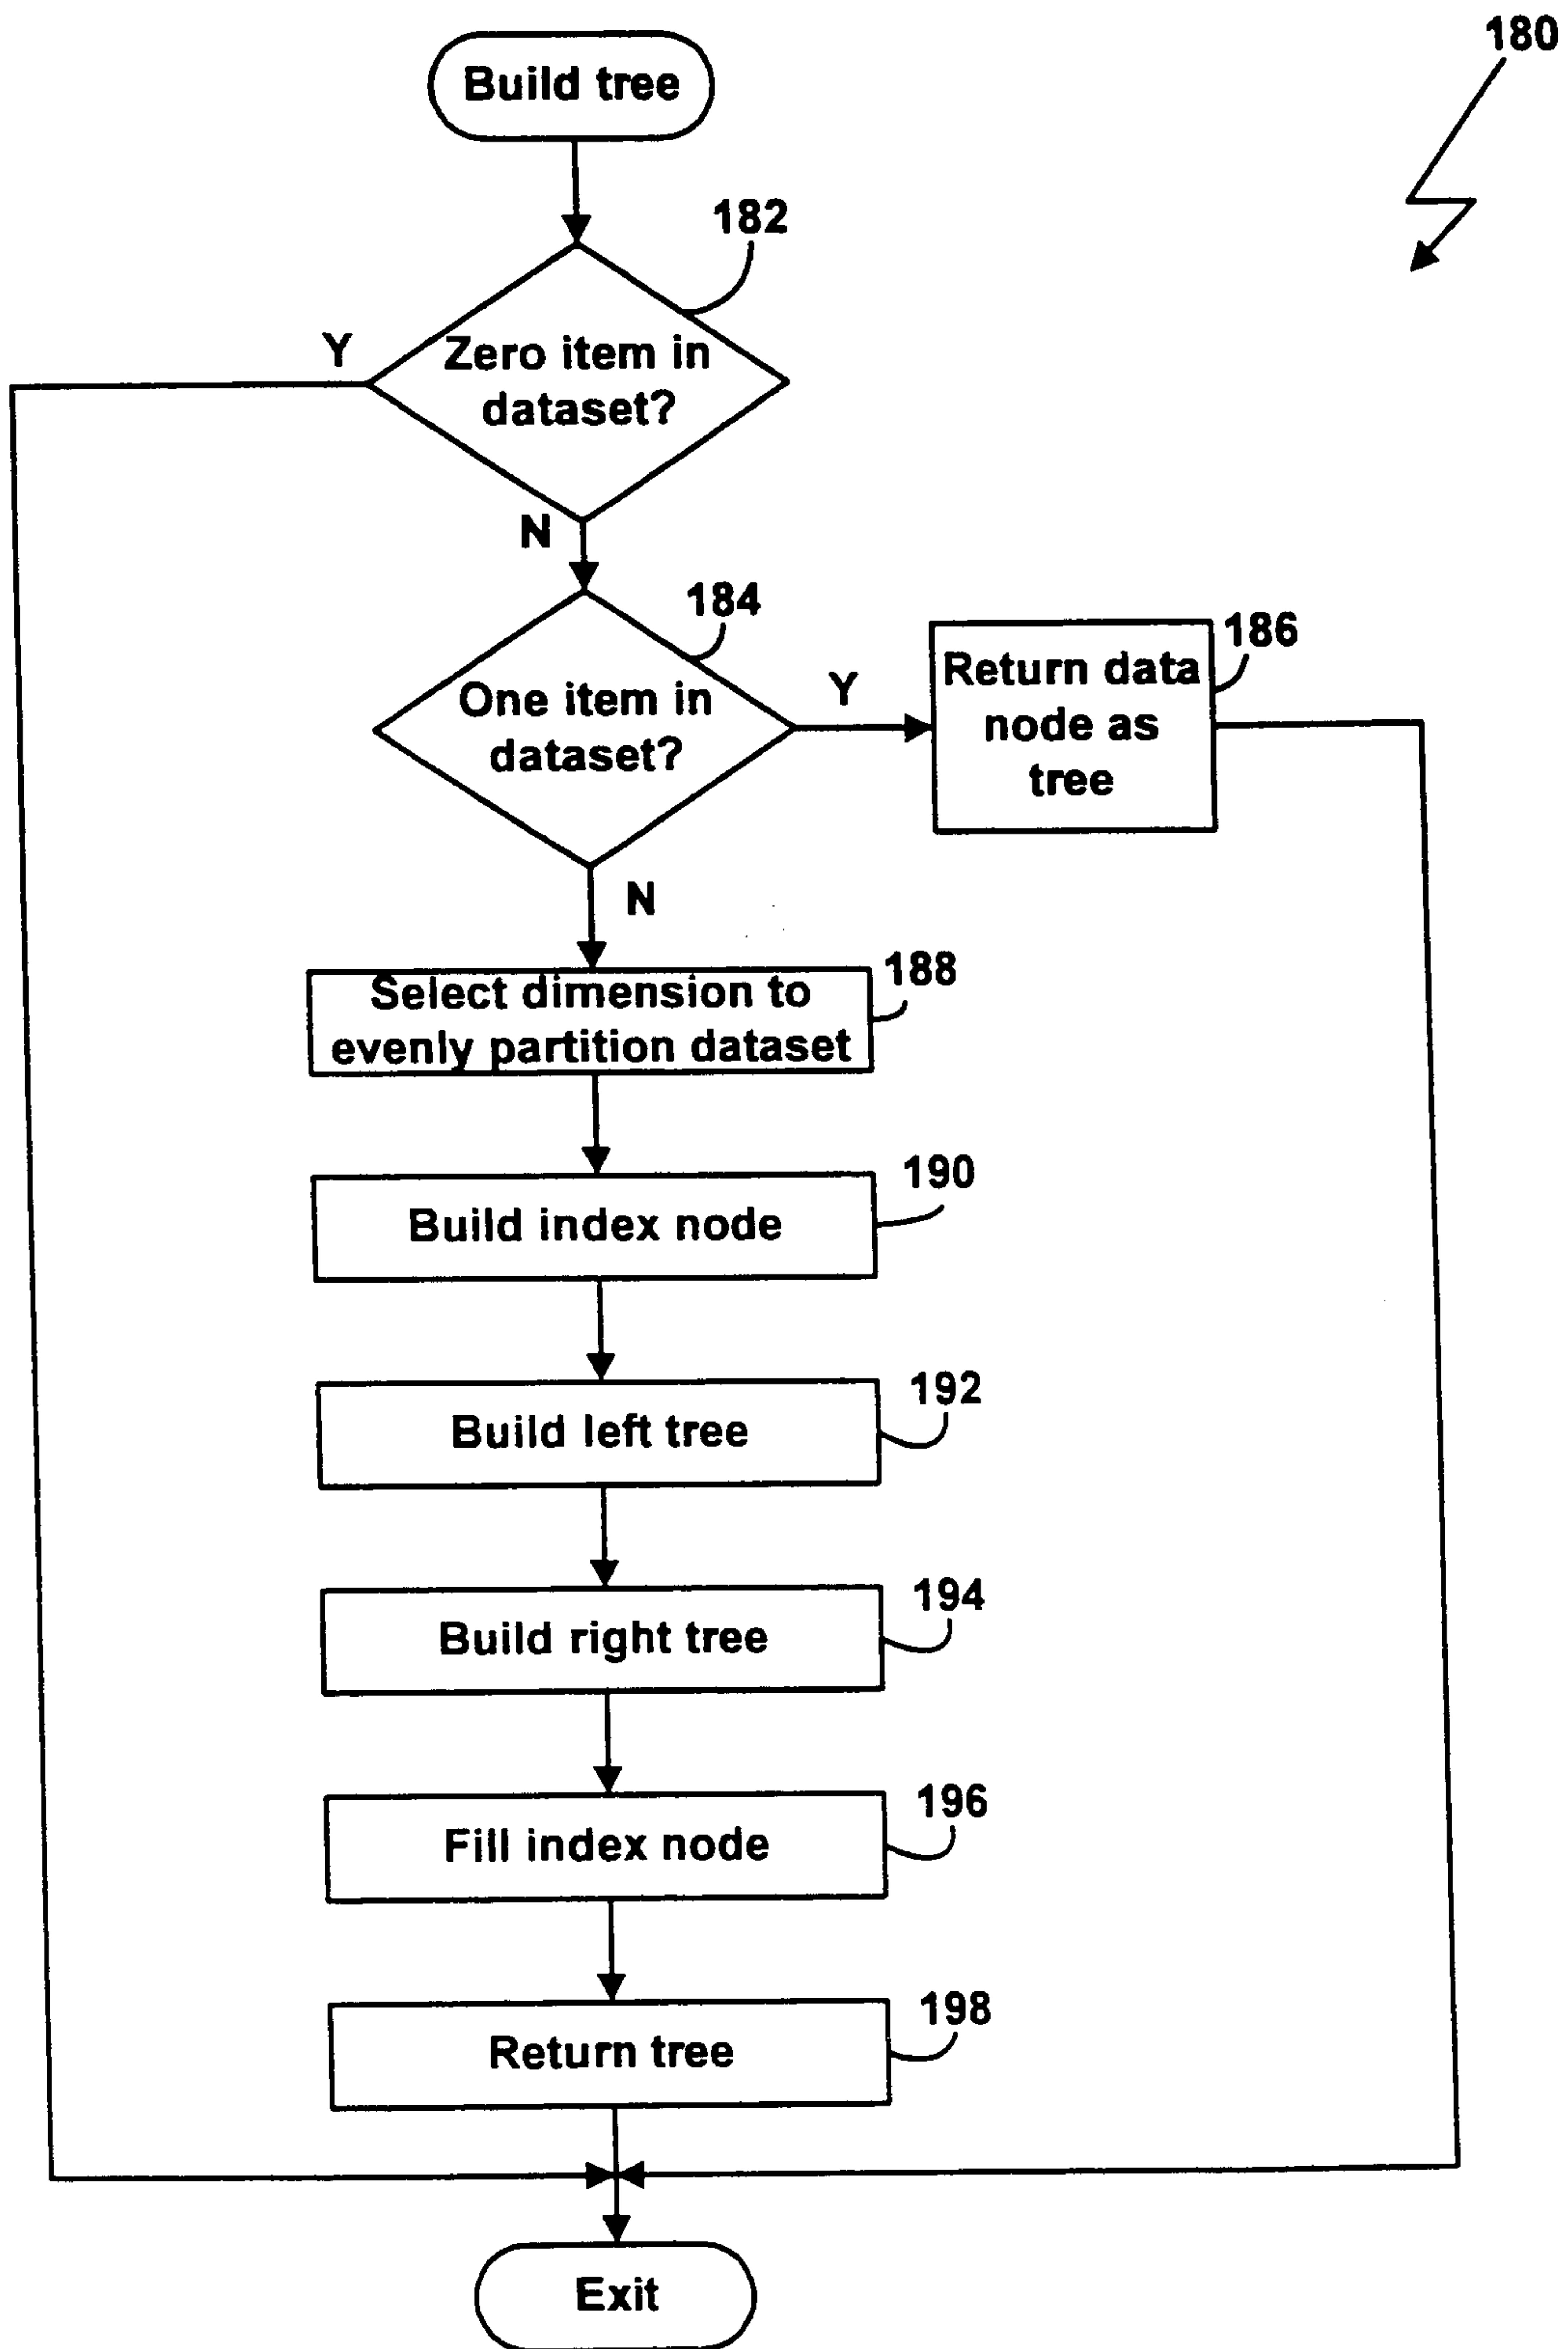


FIG. 5

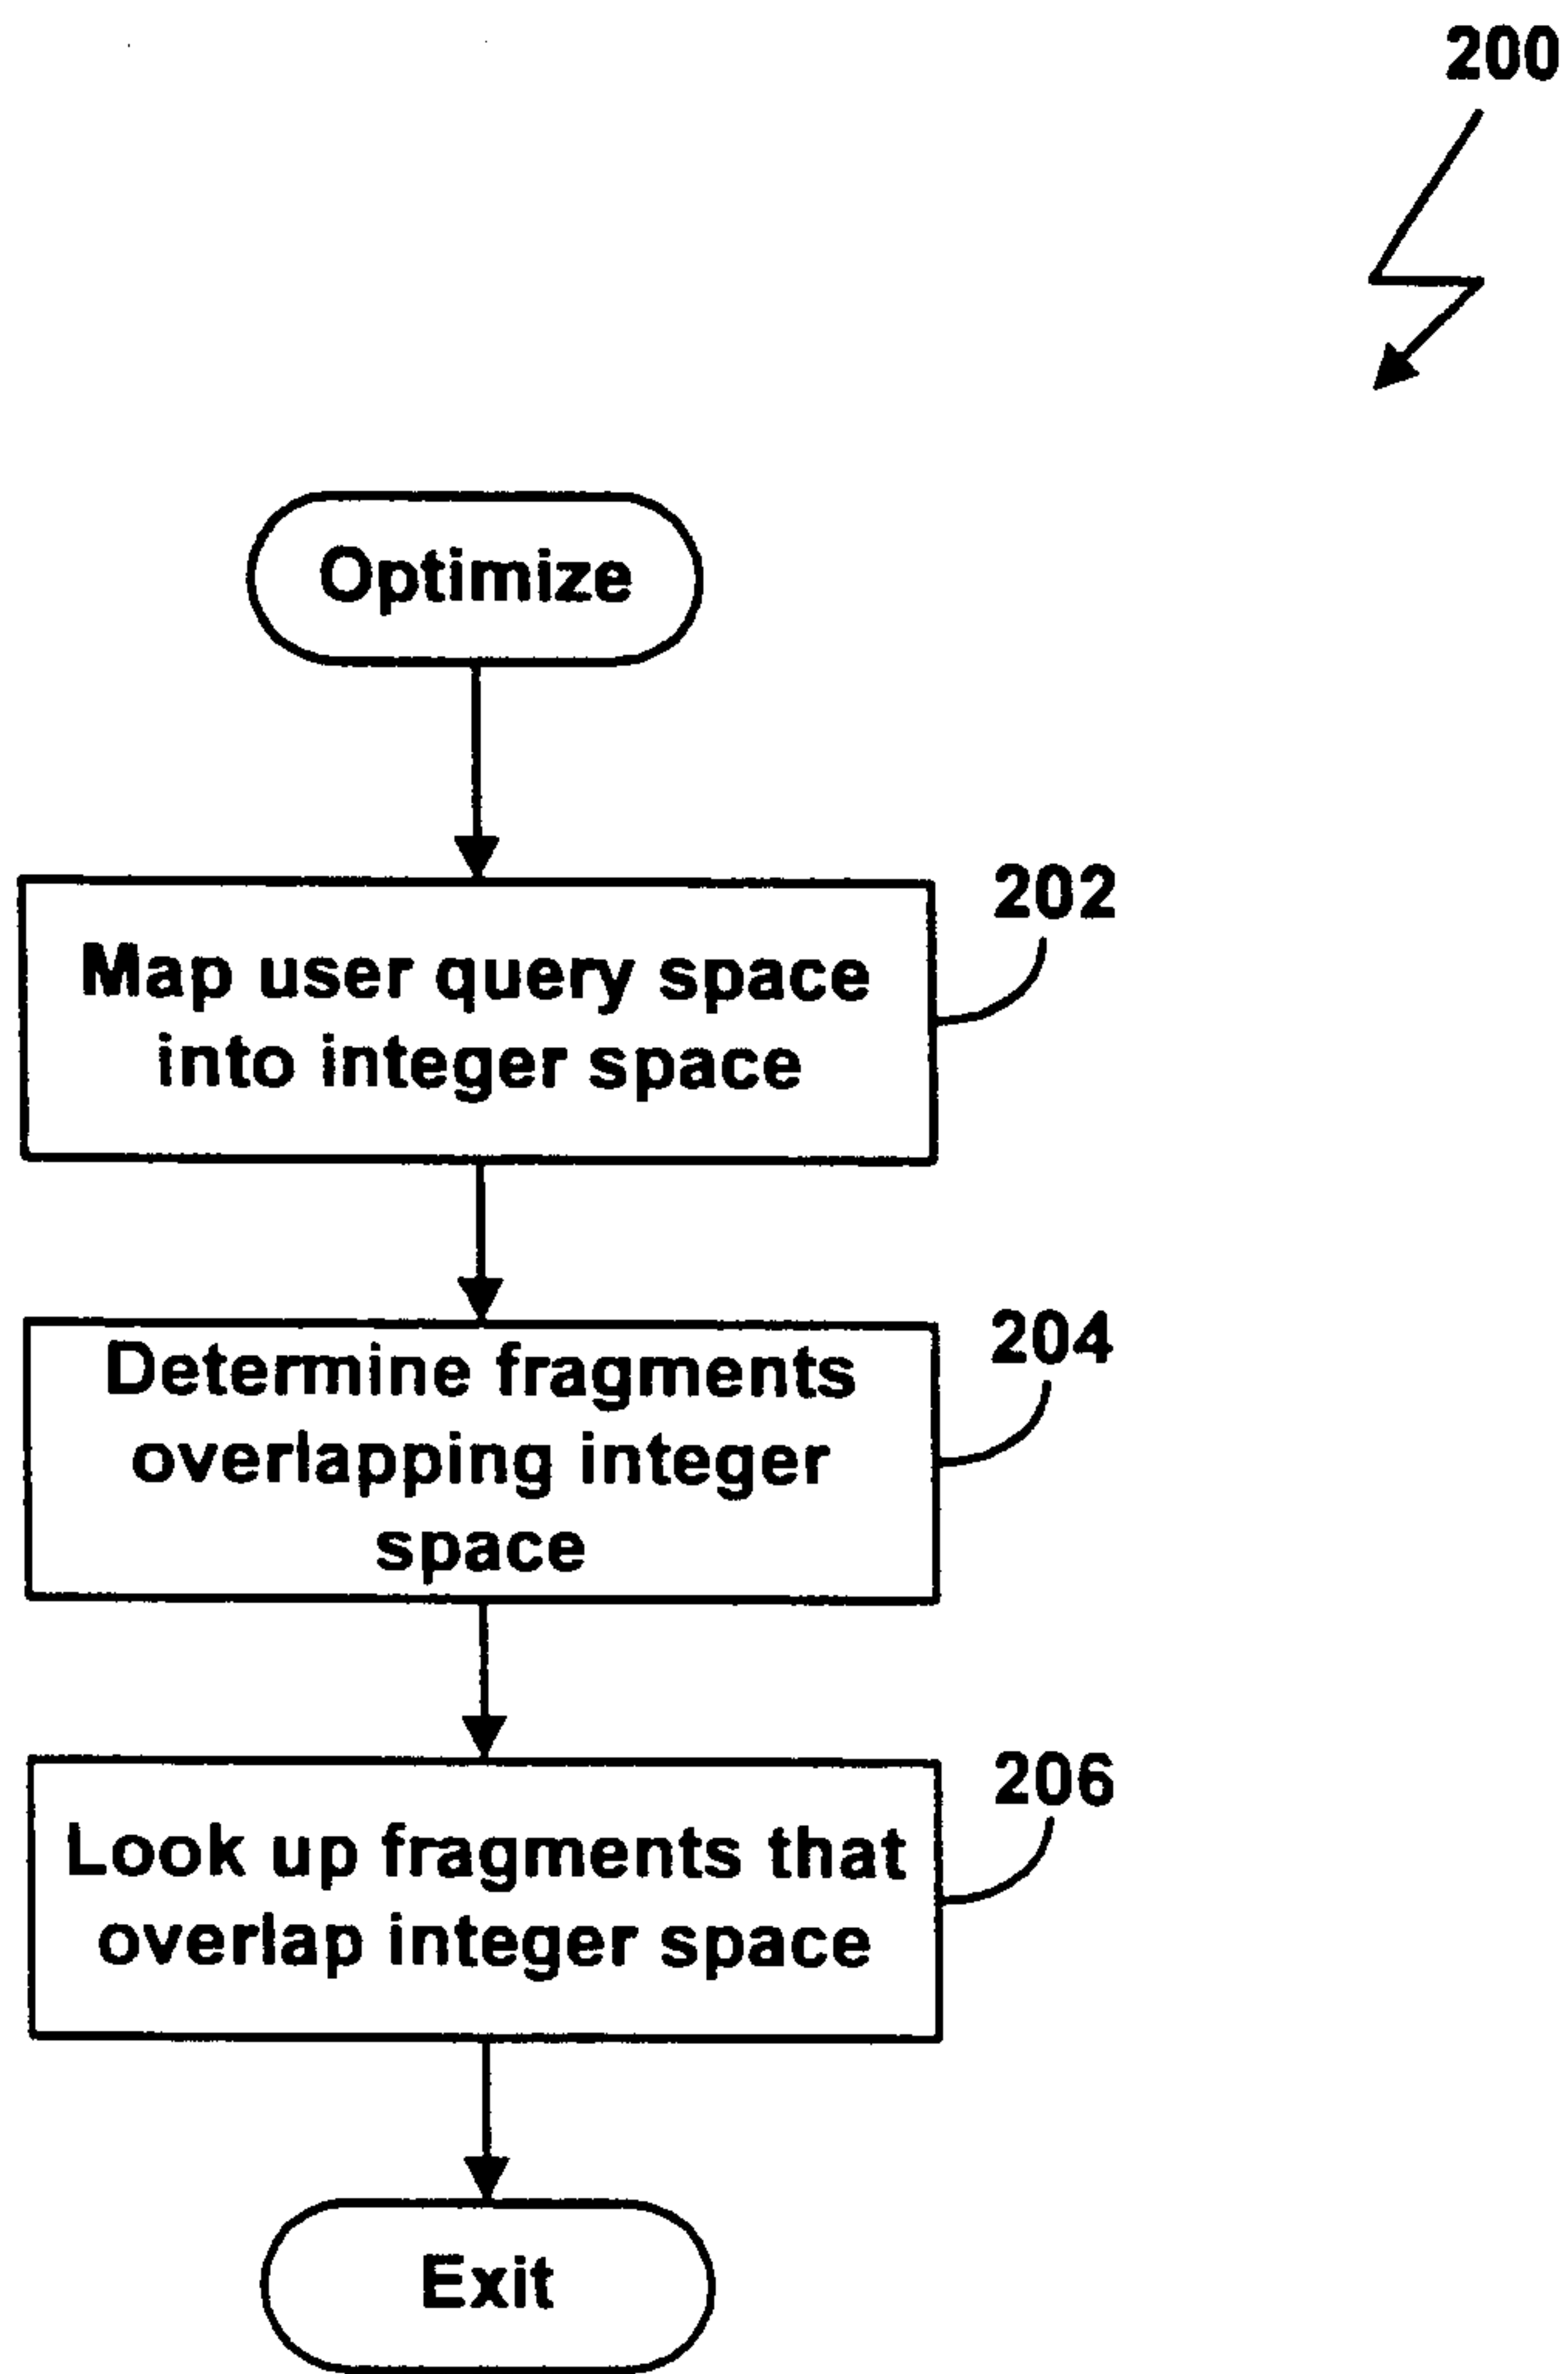


FIG. 6

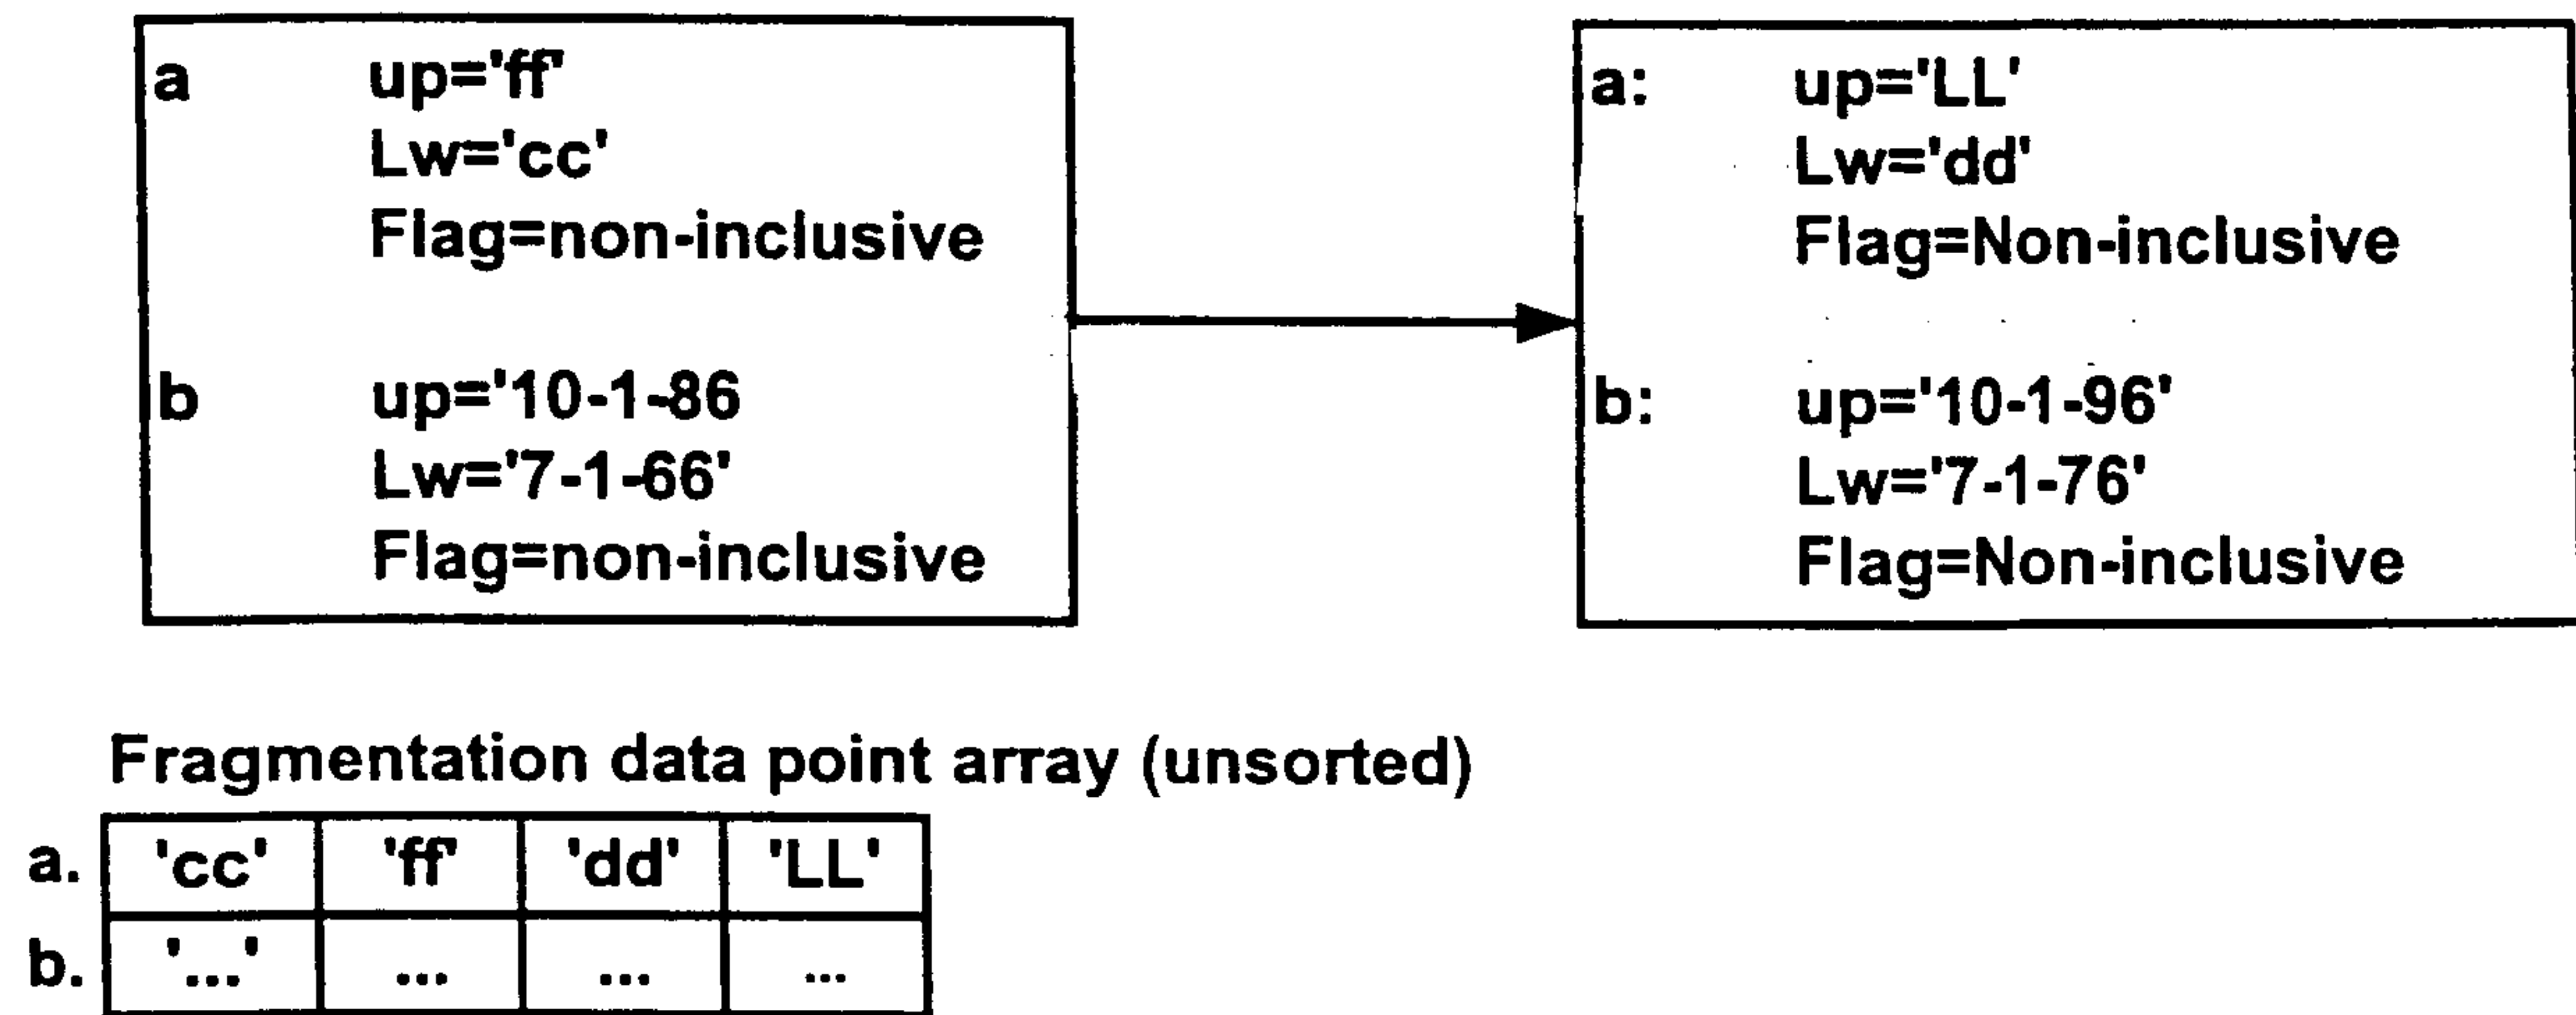


FIG. 7A

Fragment data point array (sorted)

a	Reserved for NULL	'cc'	'dd'	'ff'	'LL'
b	'...'	'...'	...	...	...

Dimension\_t array

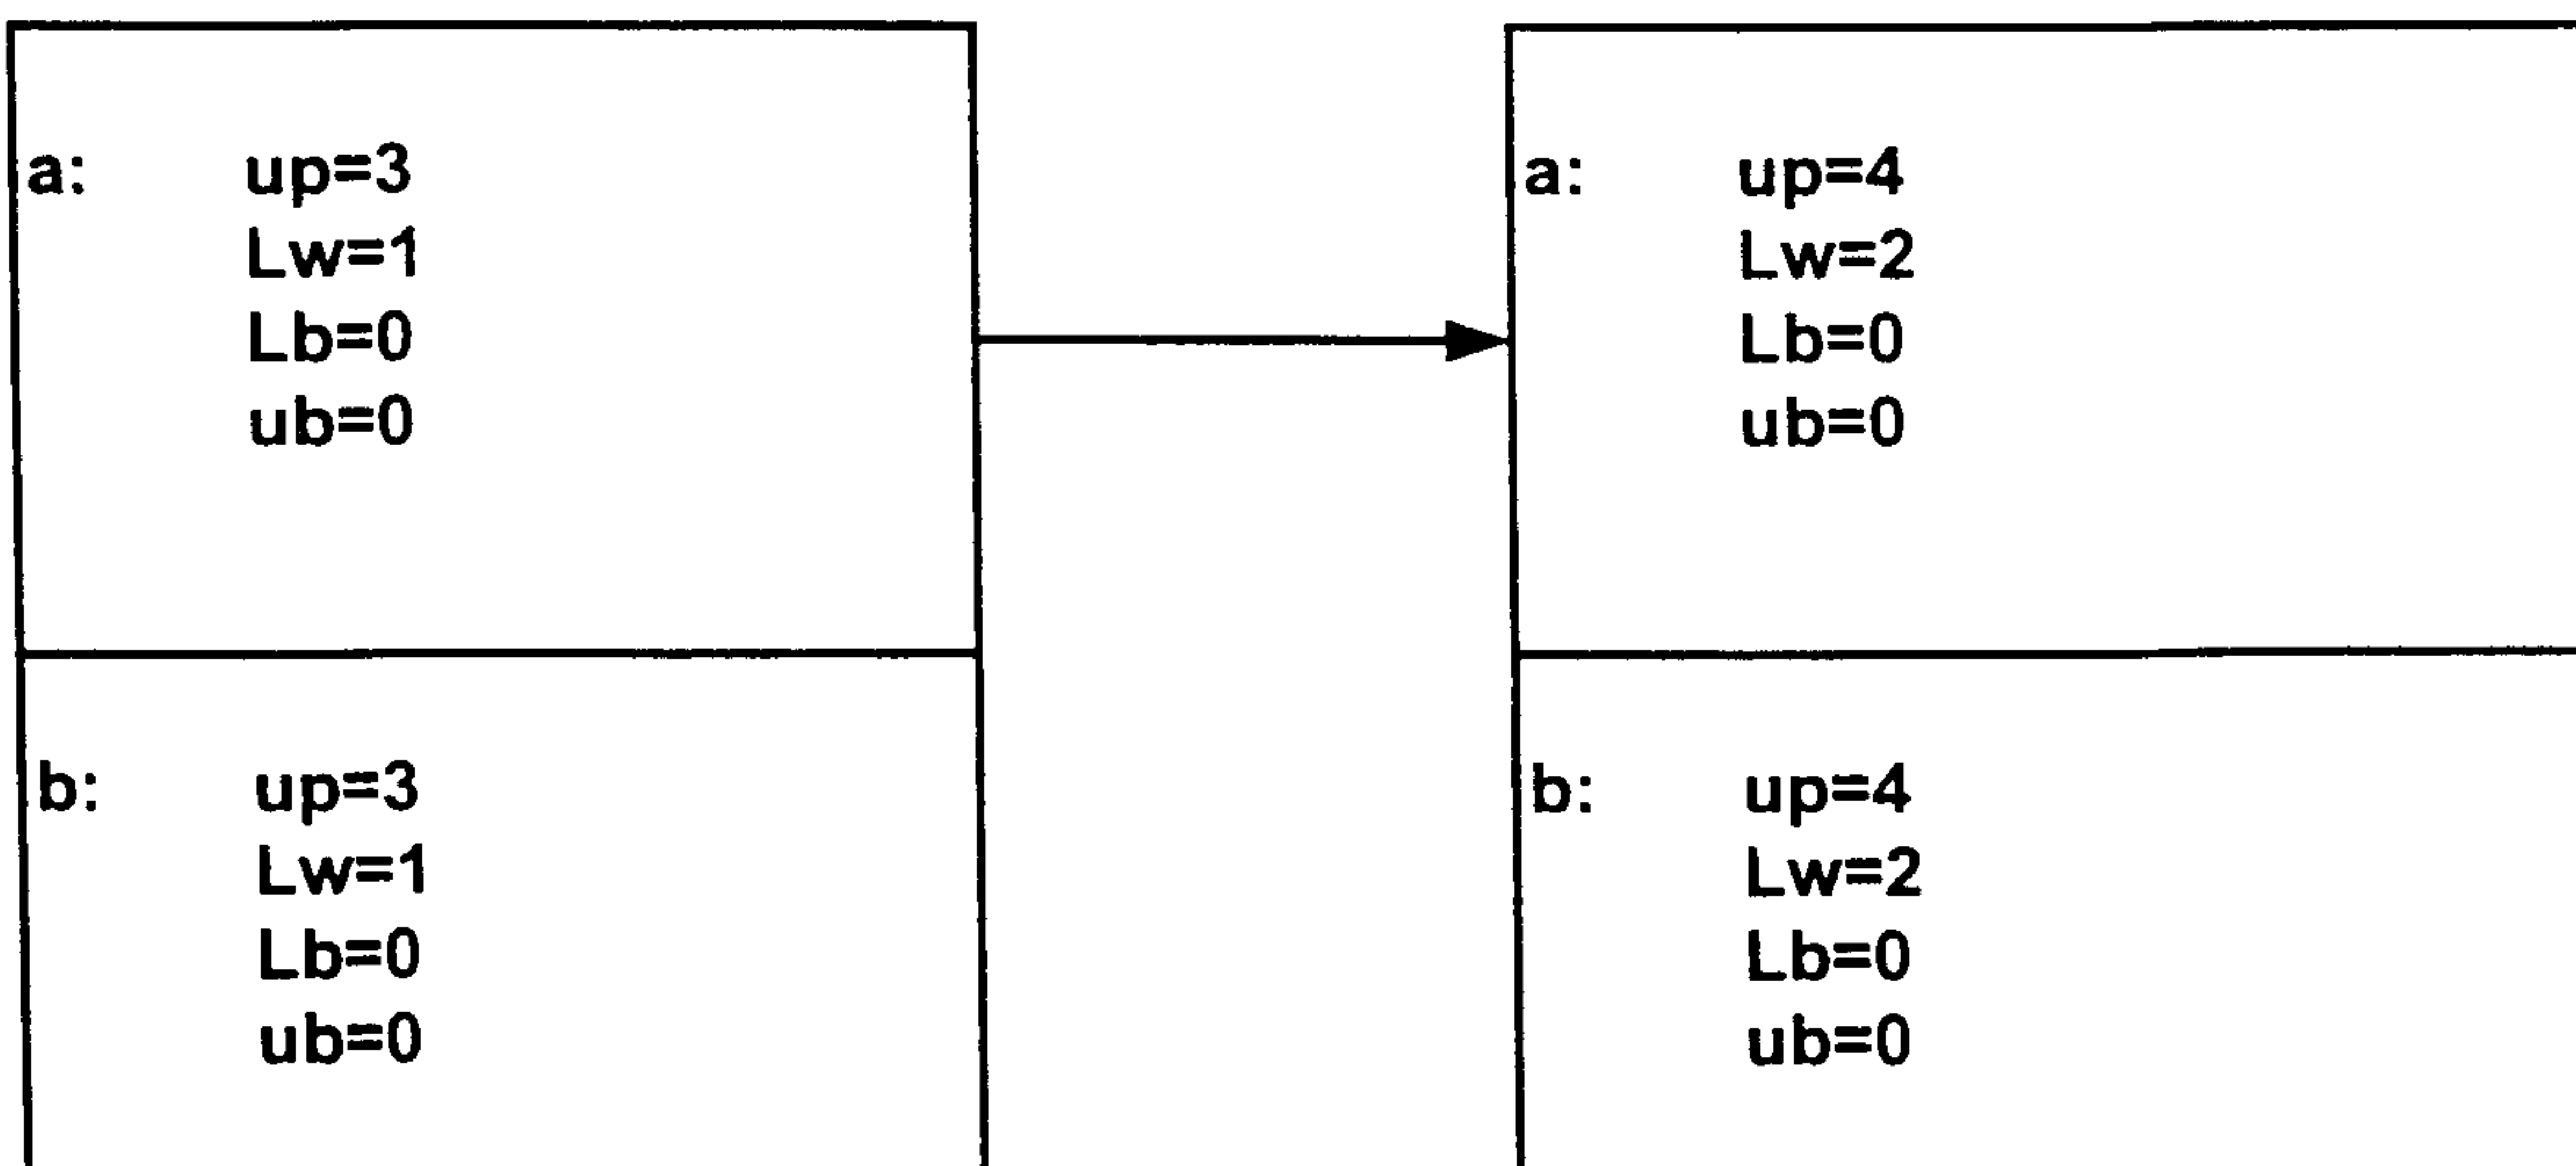


FIG. 7B

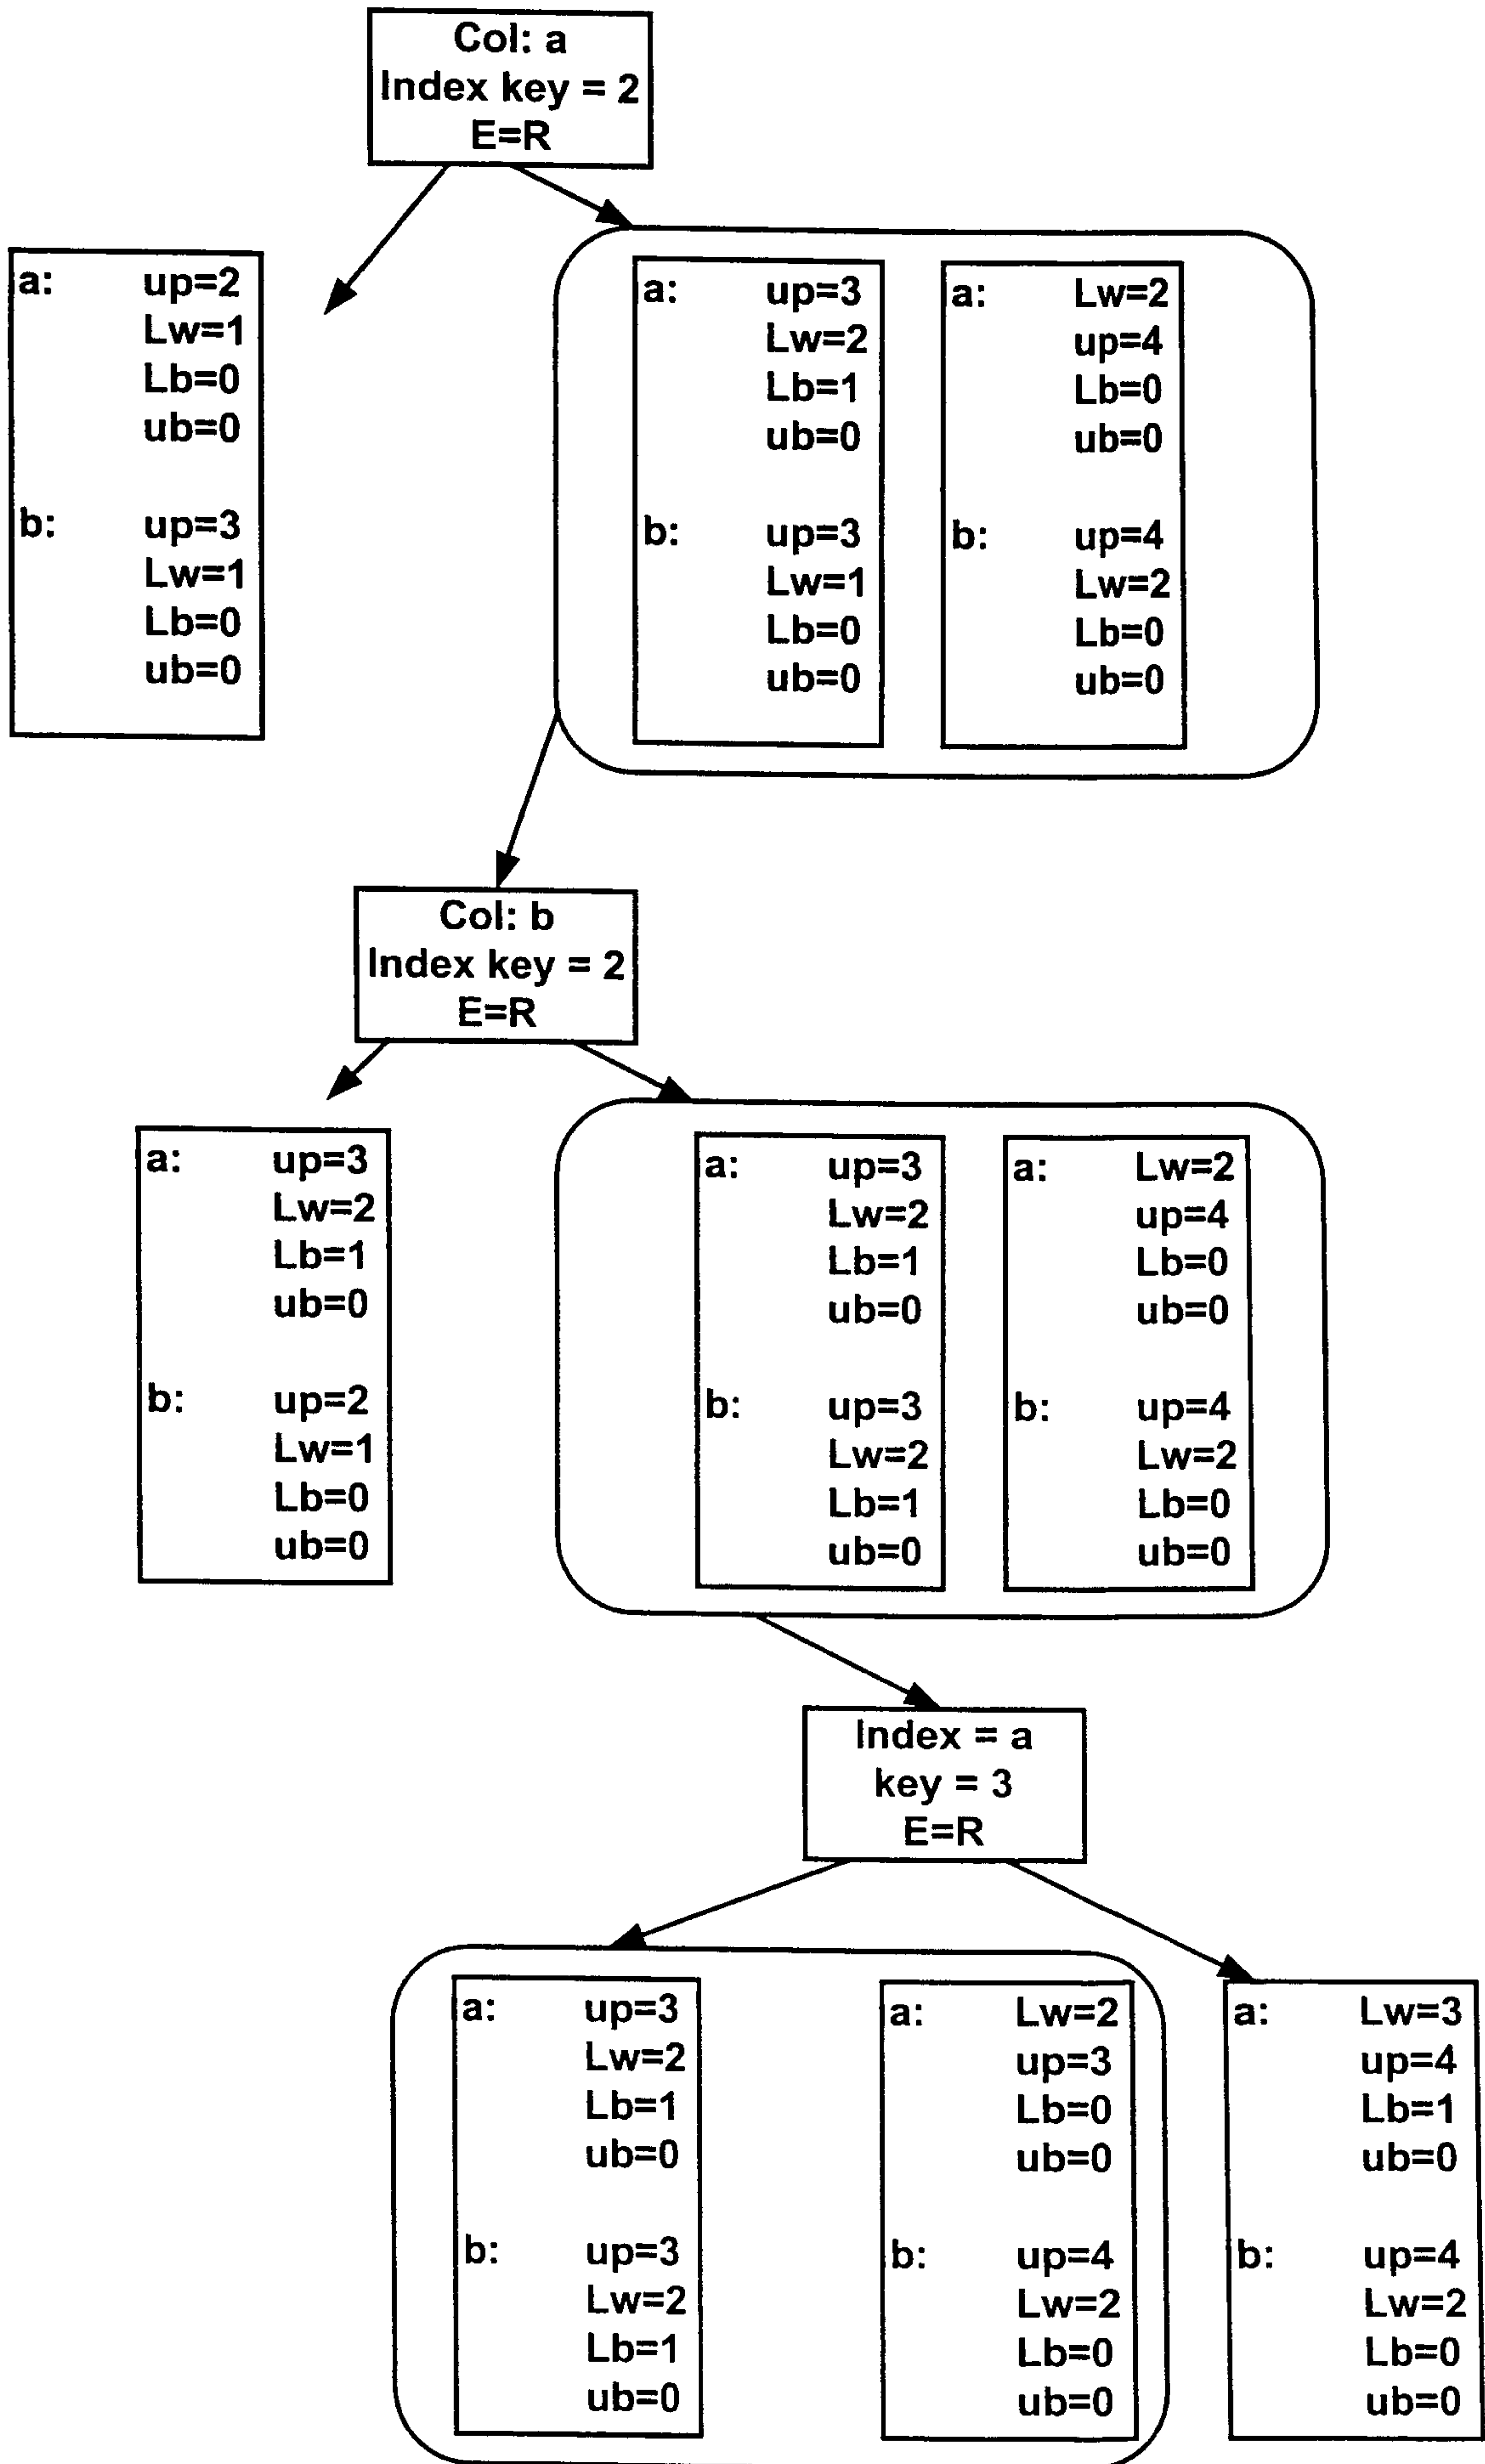


FIG. 7C

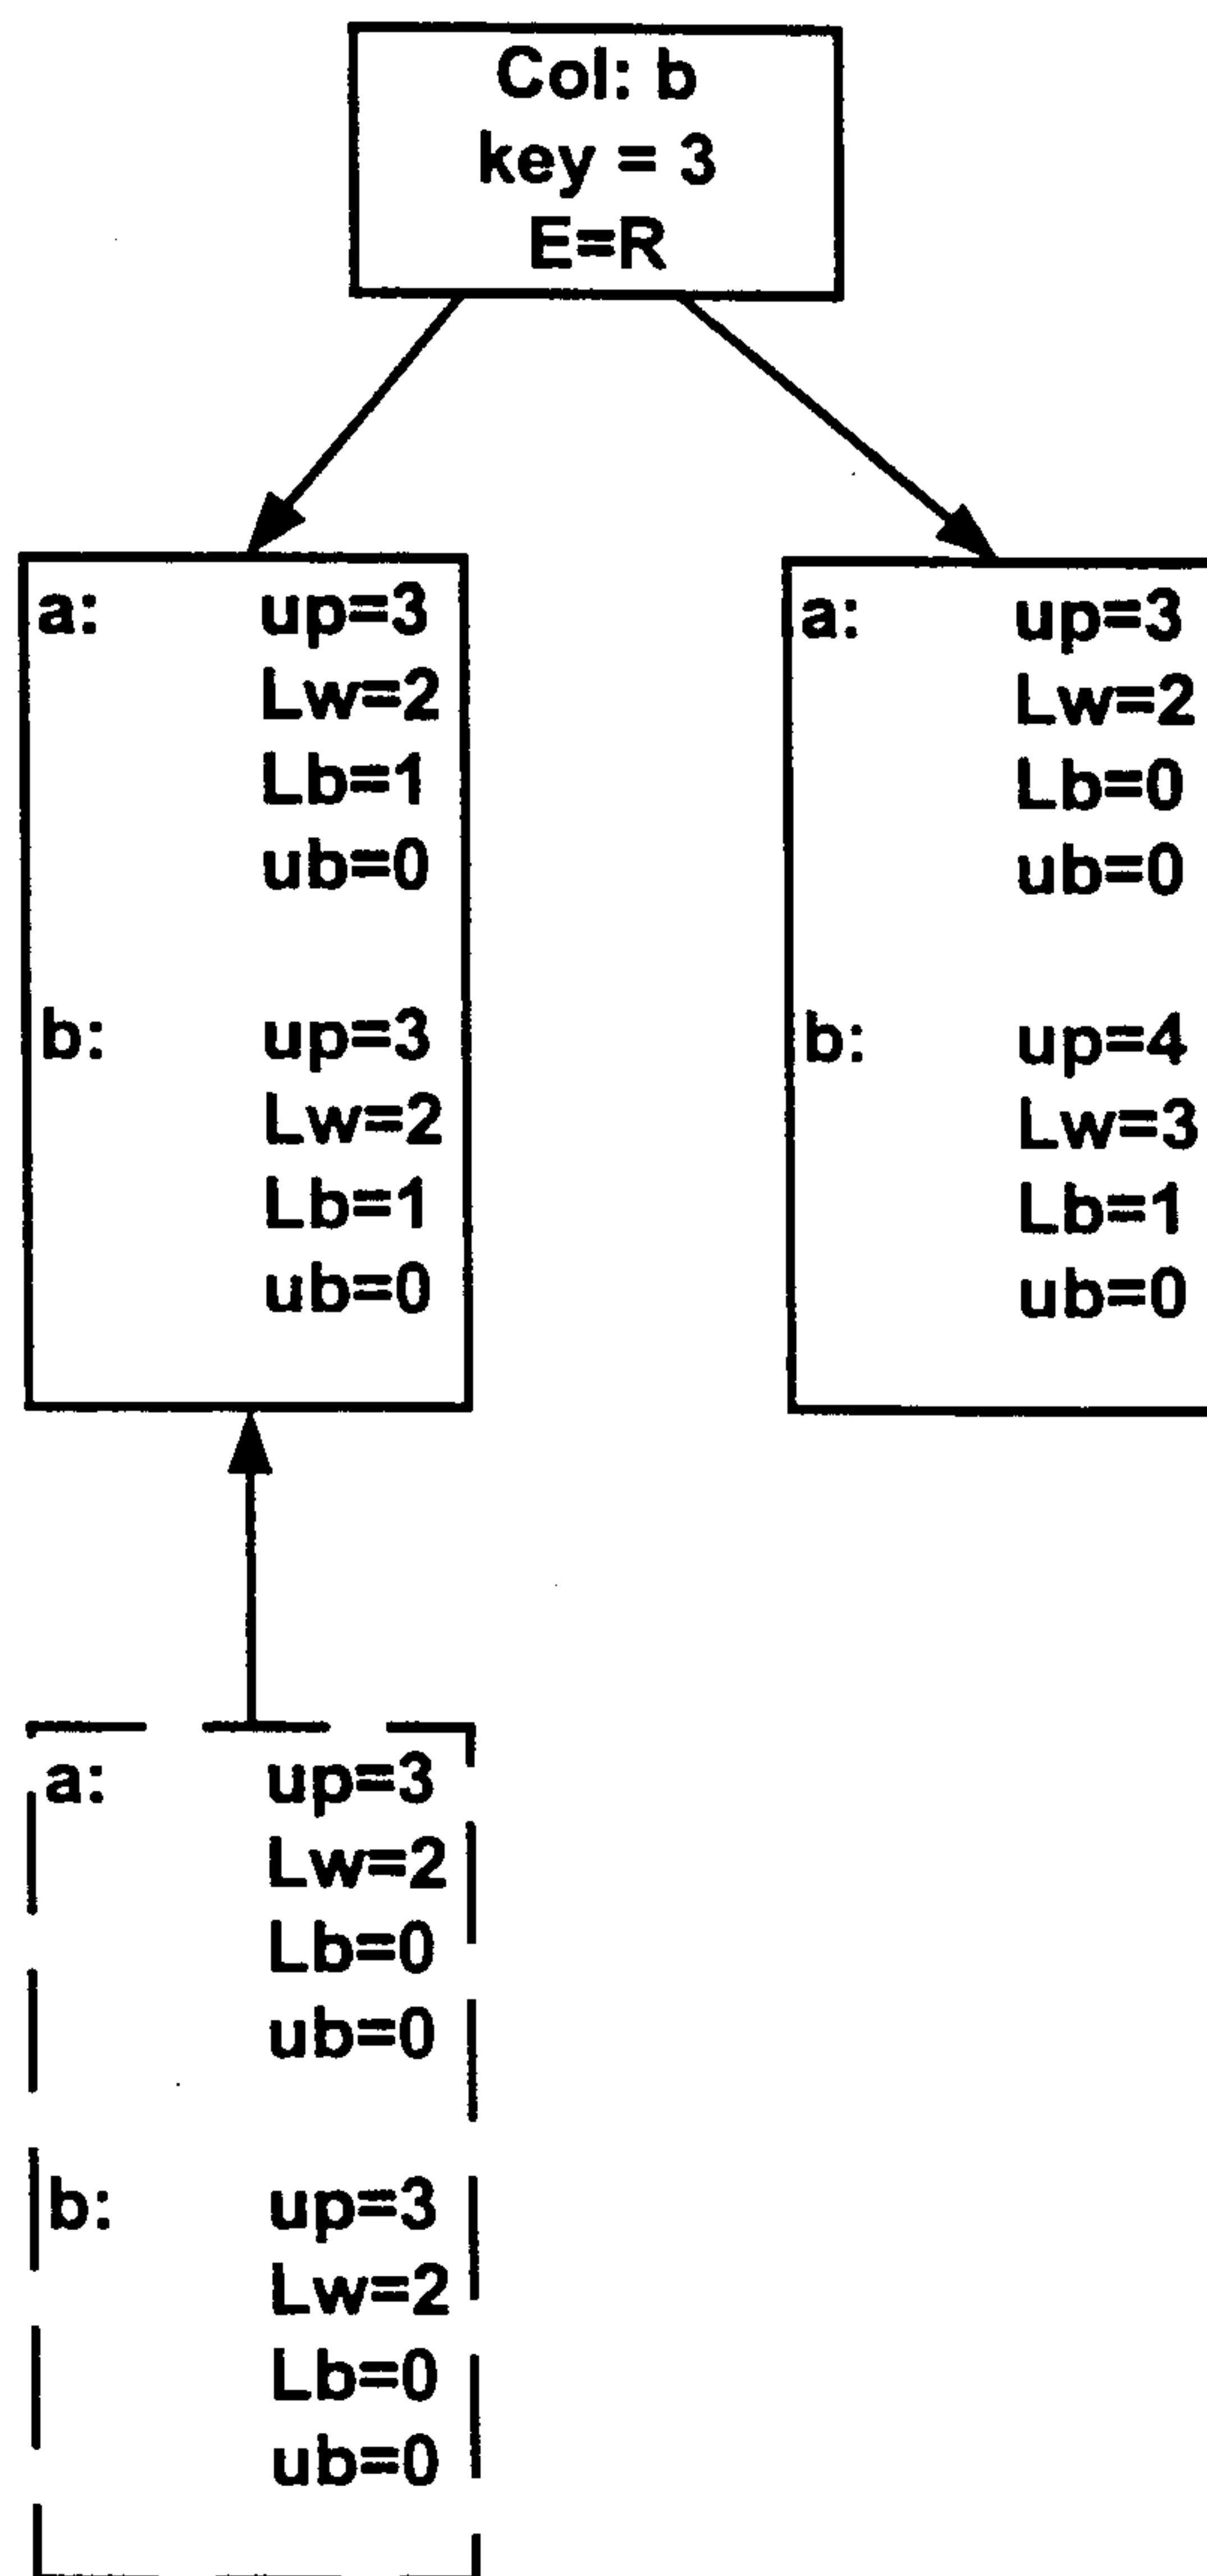
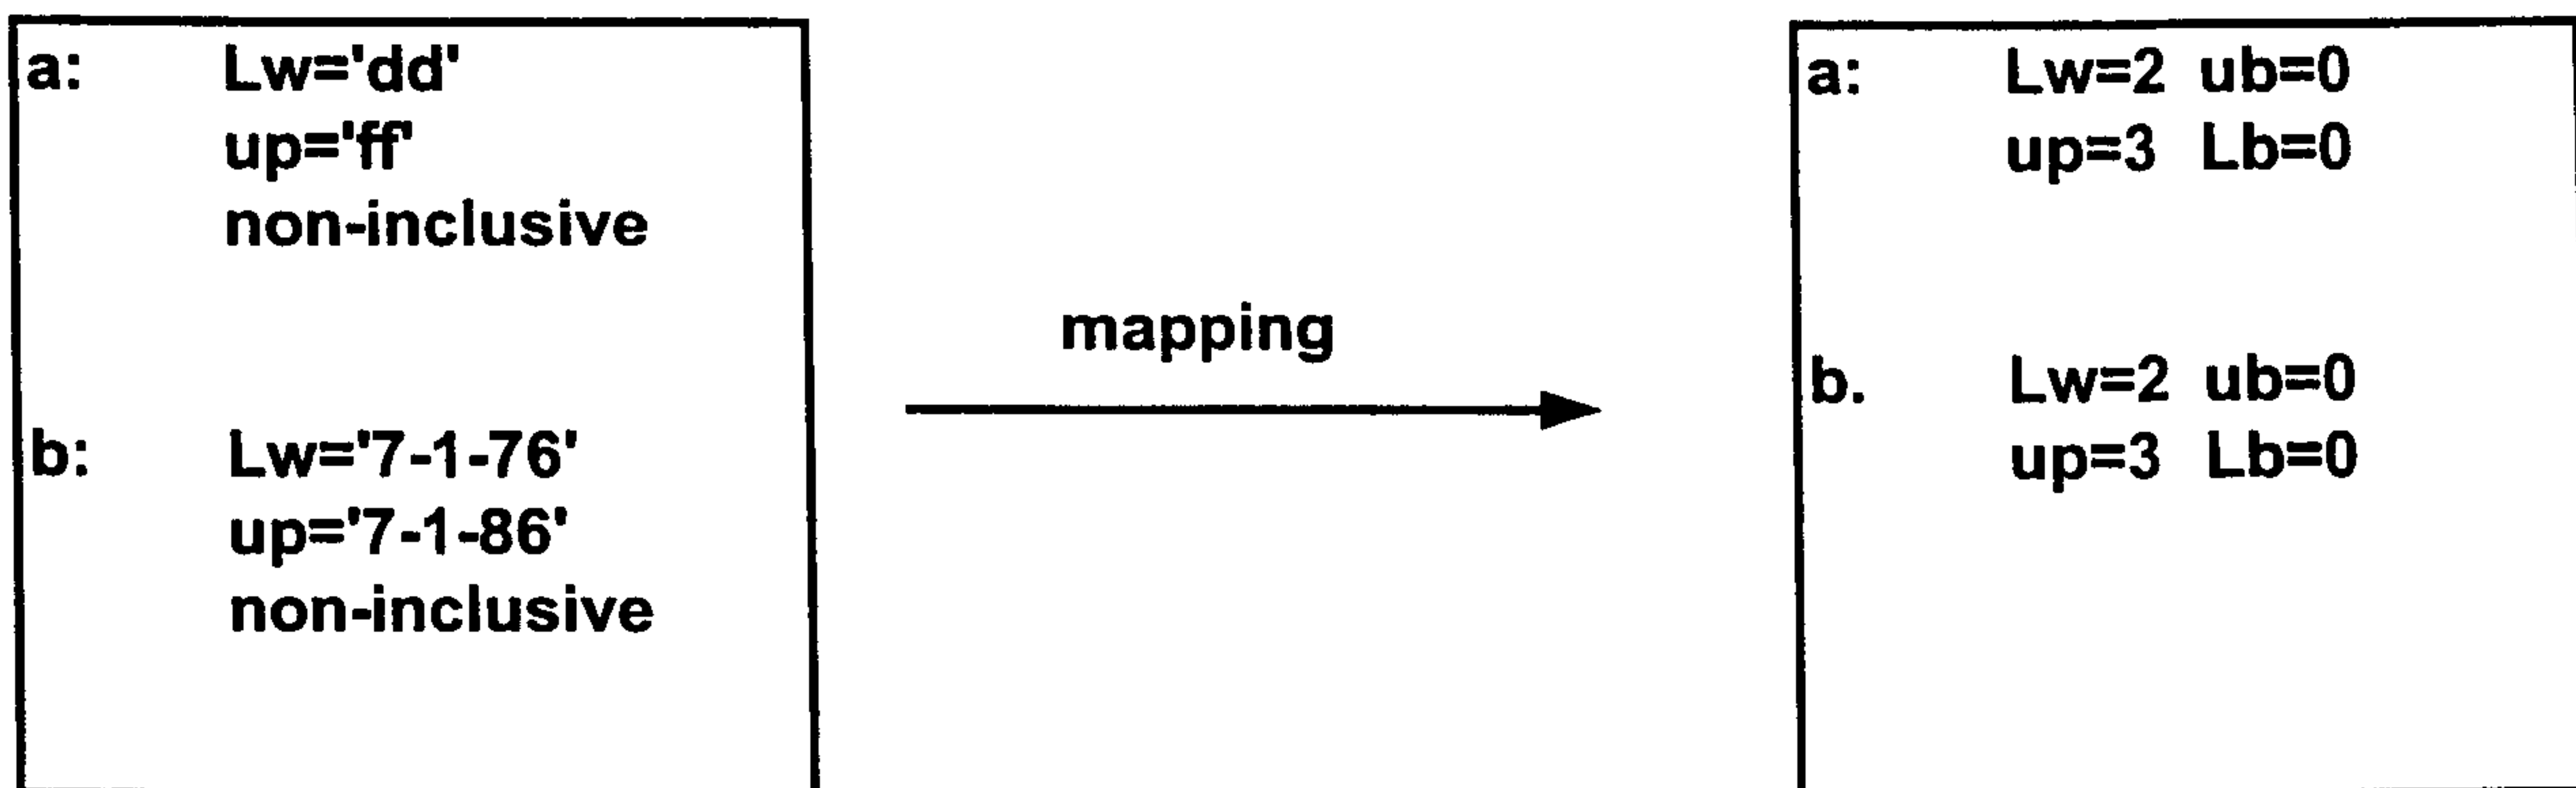


FIG. 7D



**FIG. 7E**

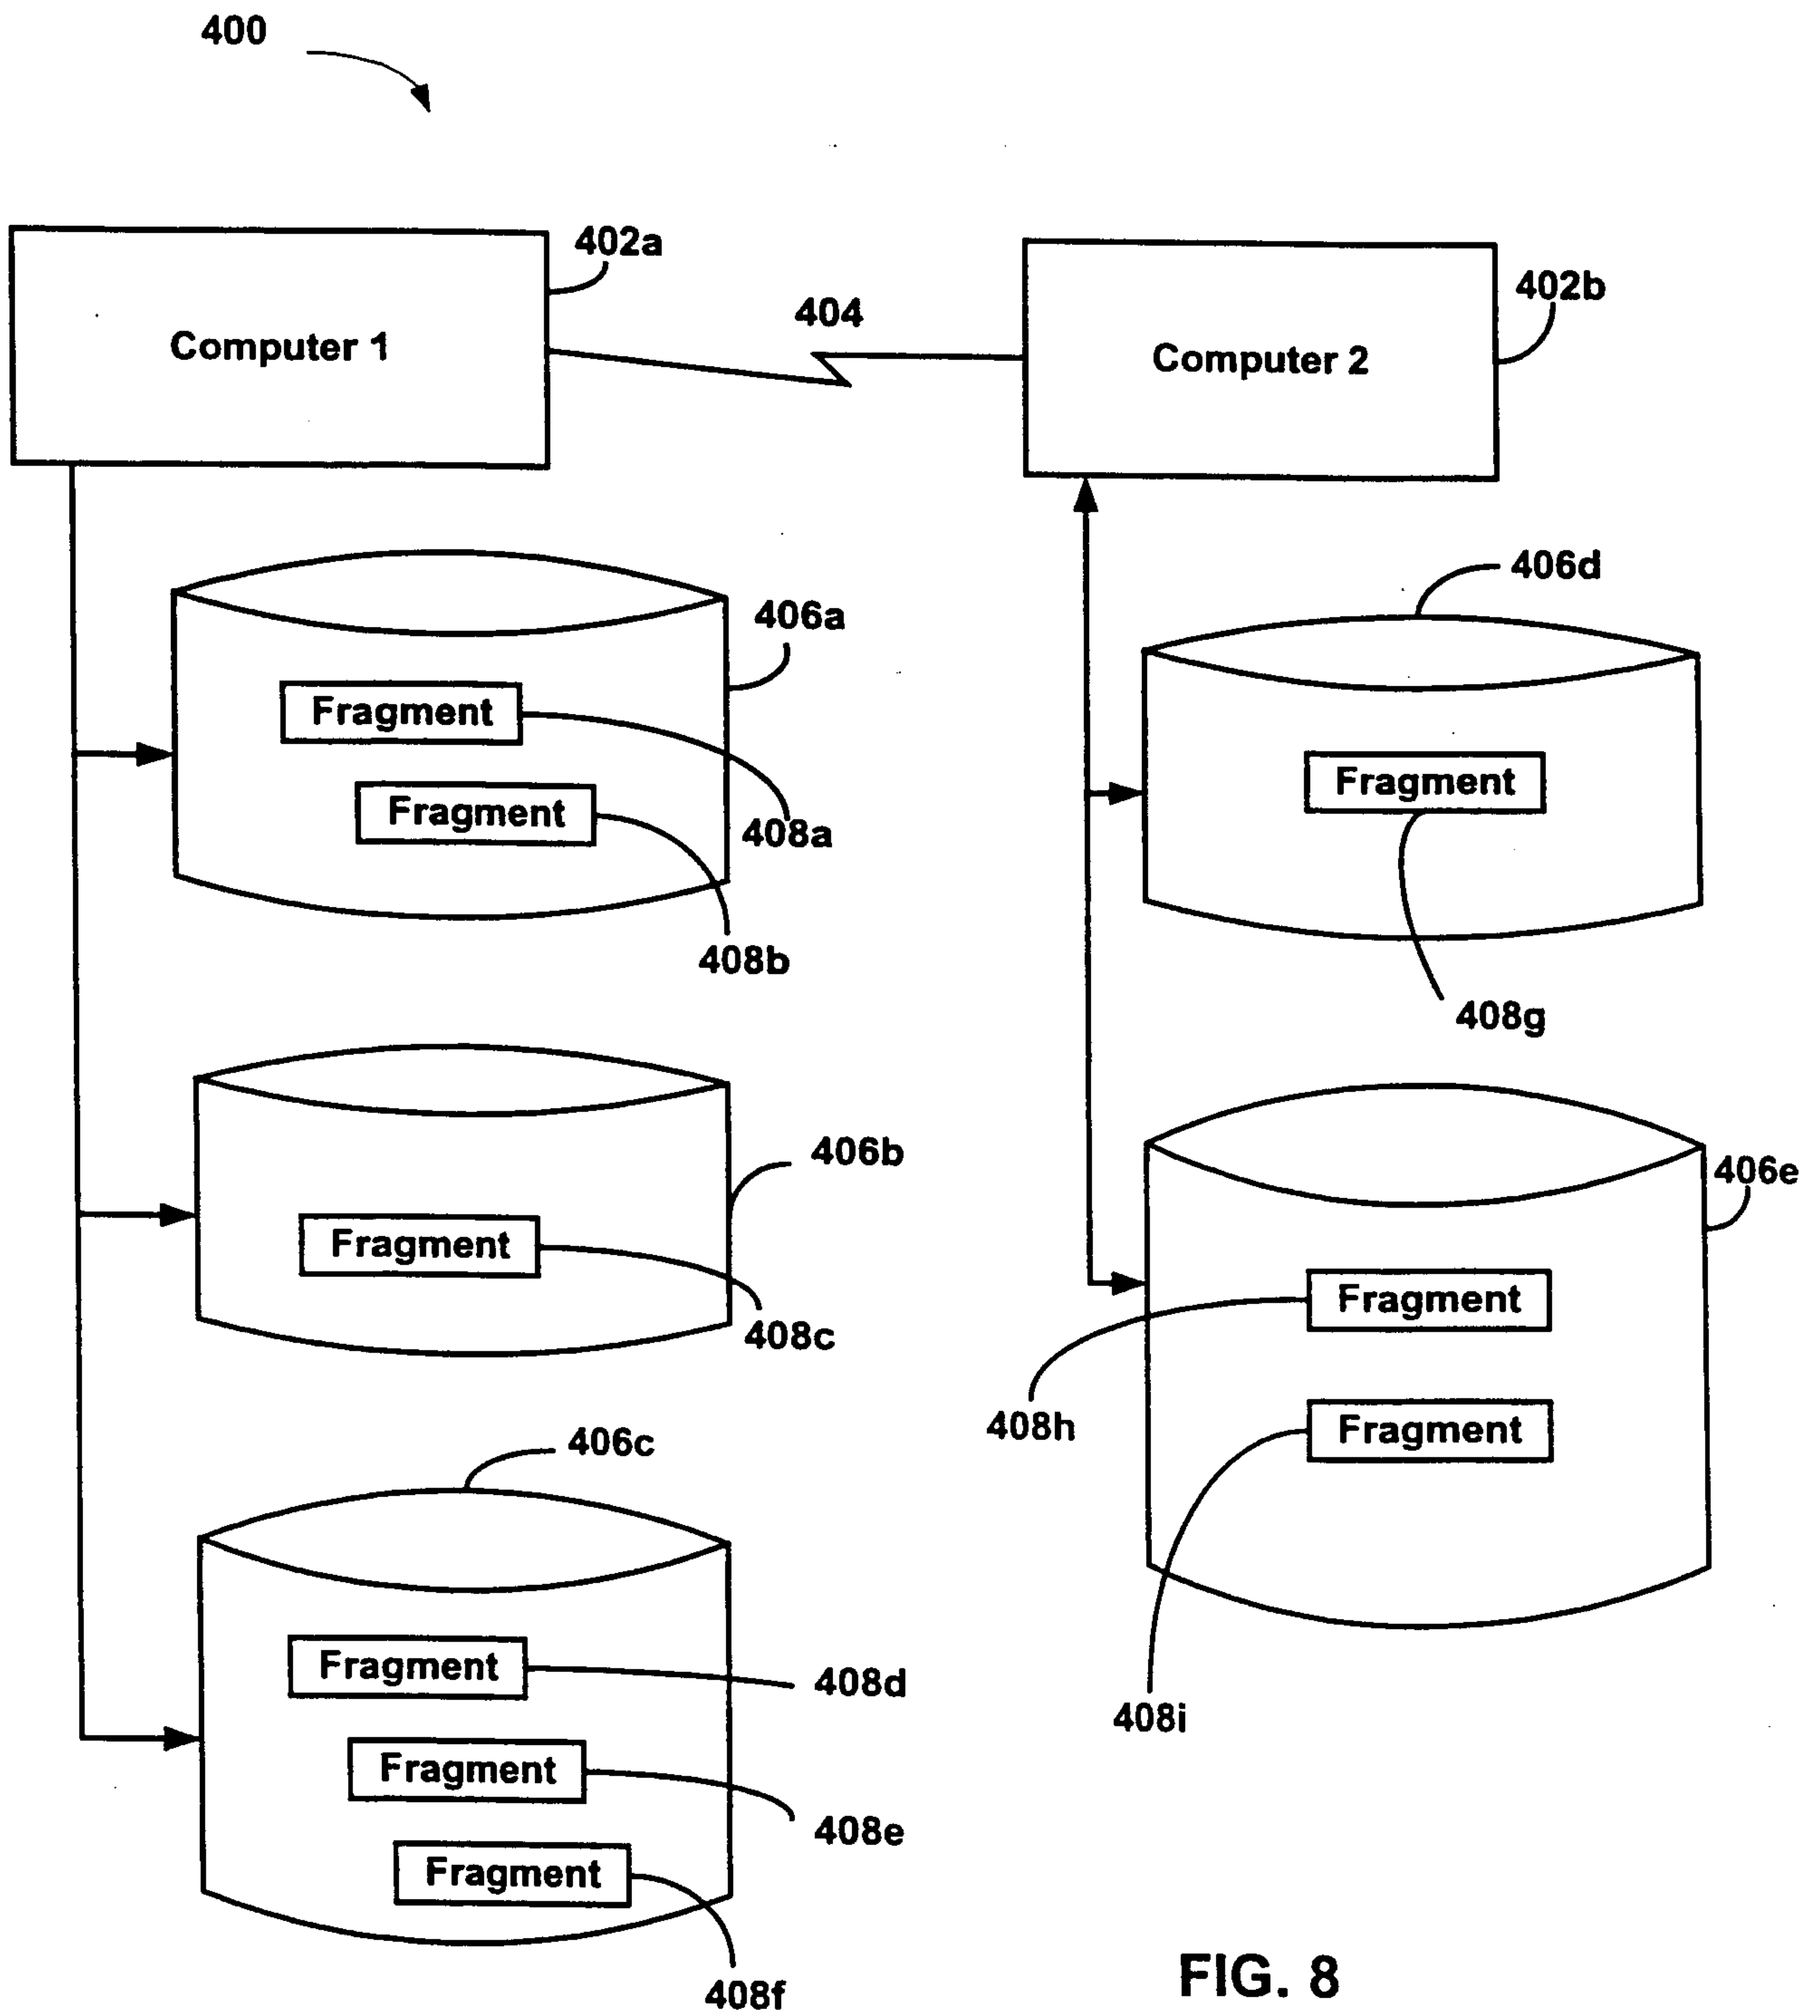

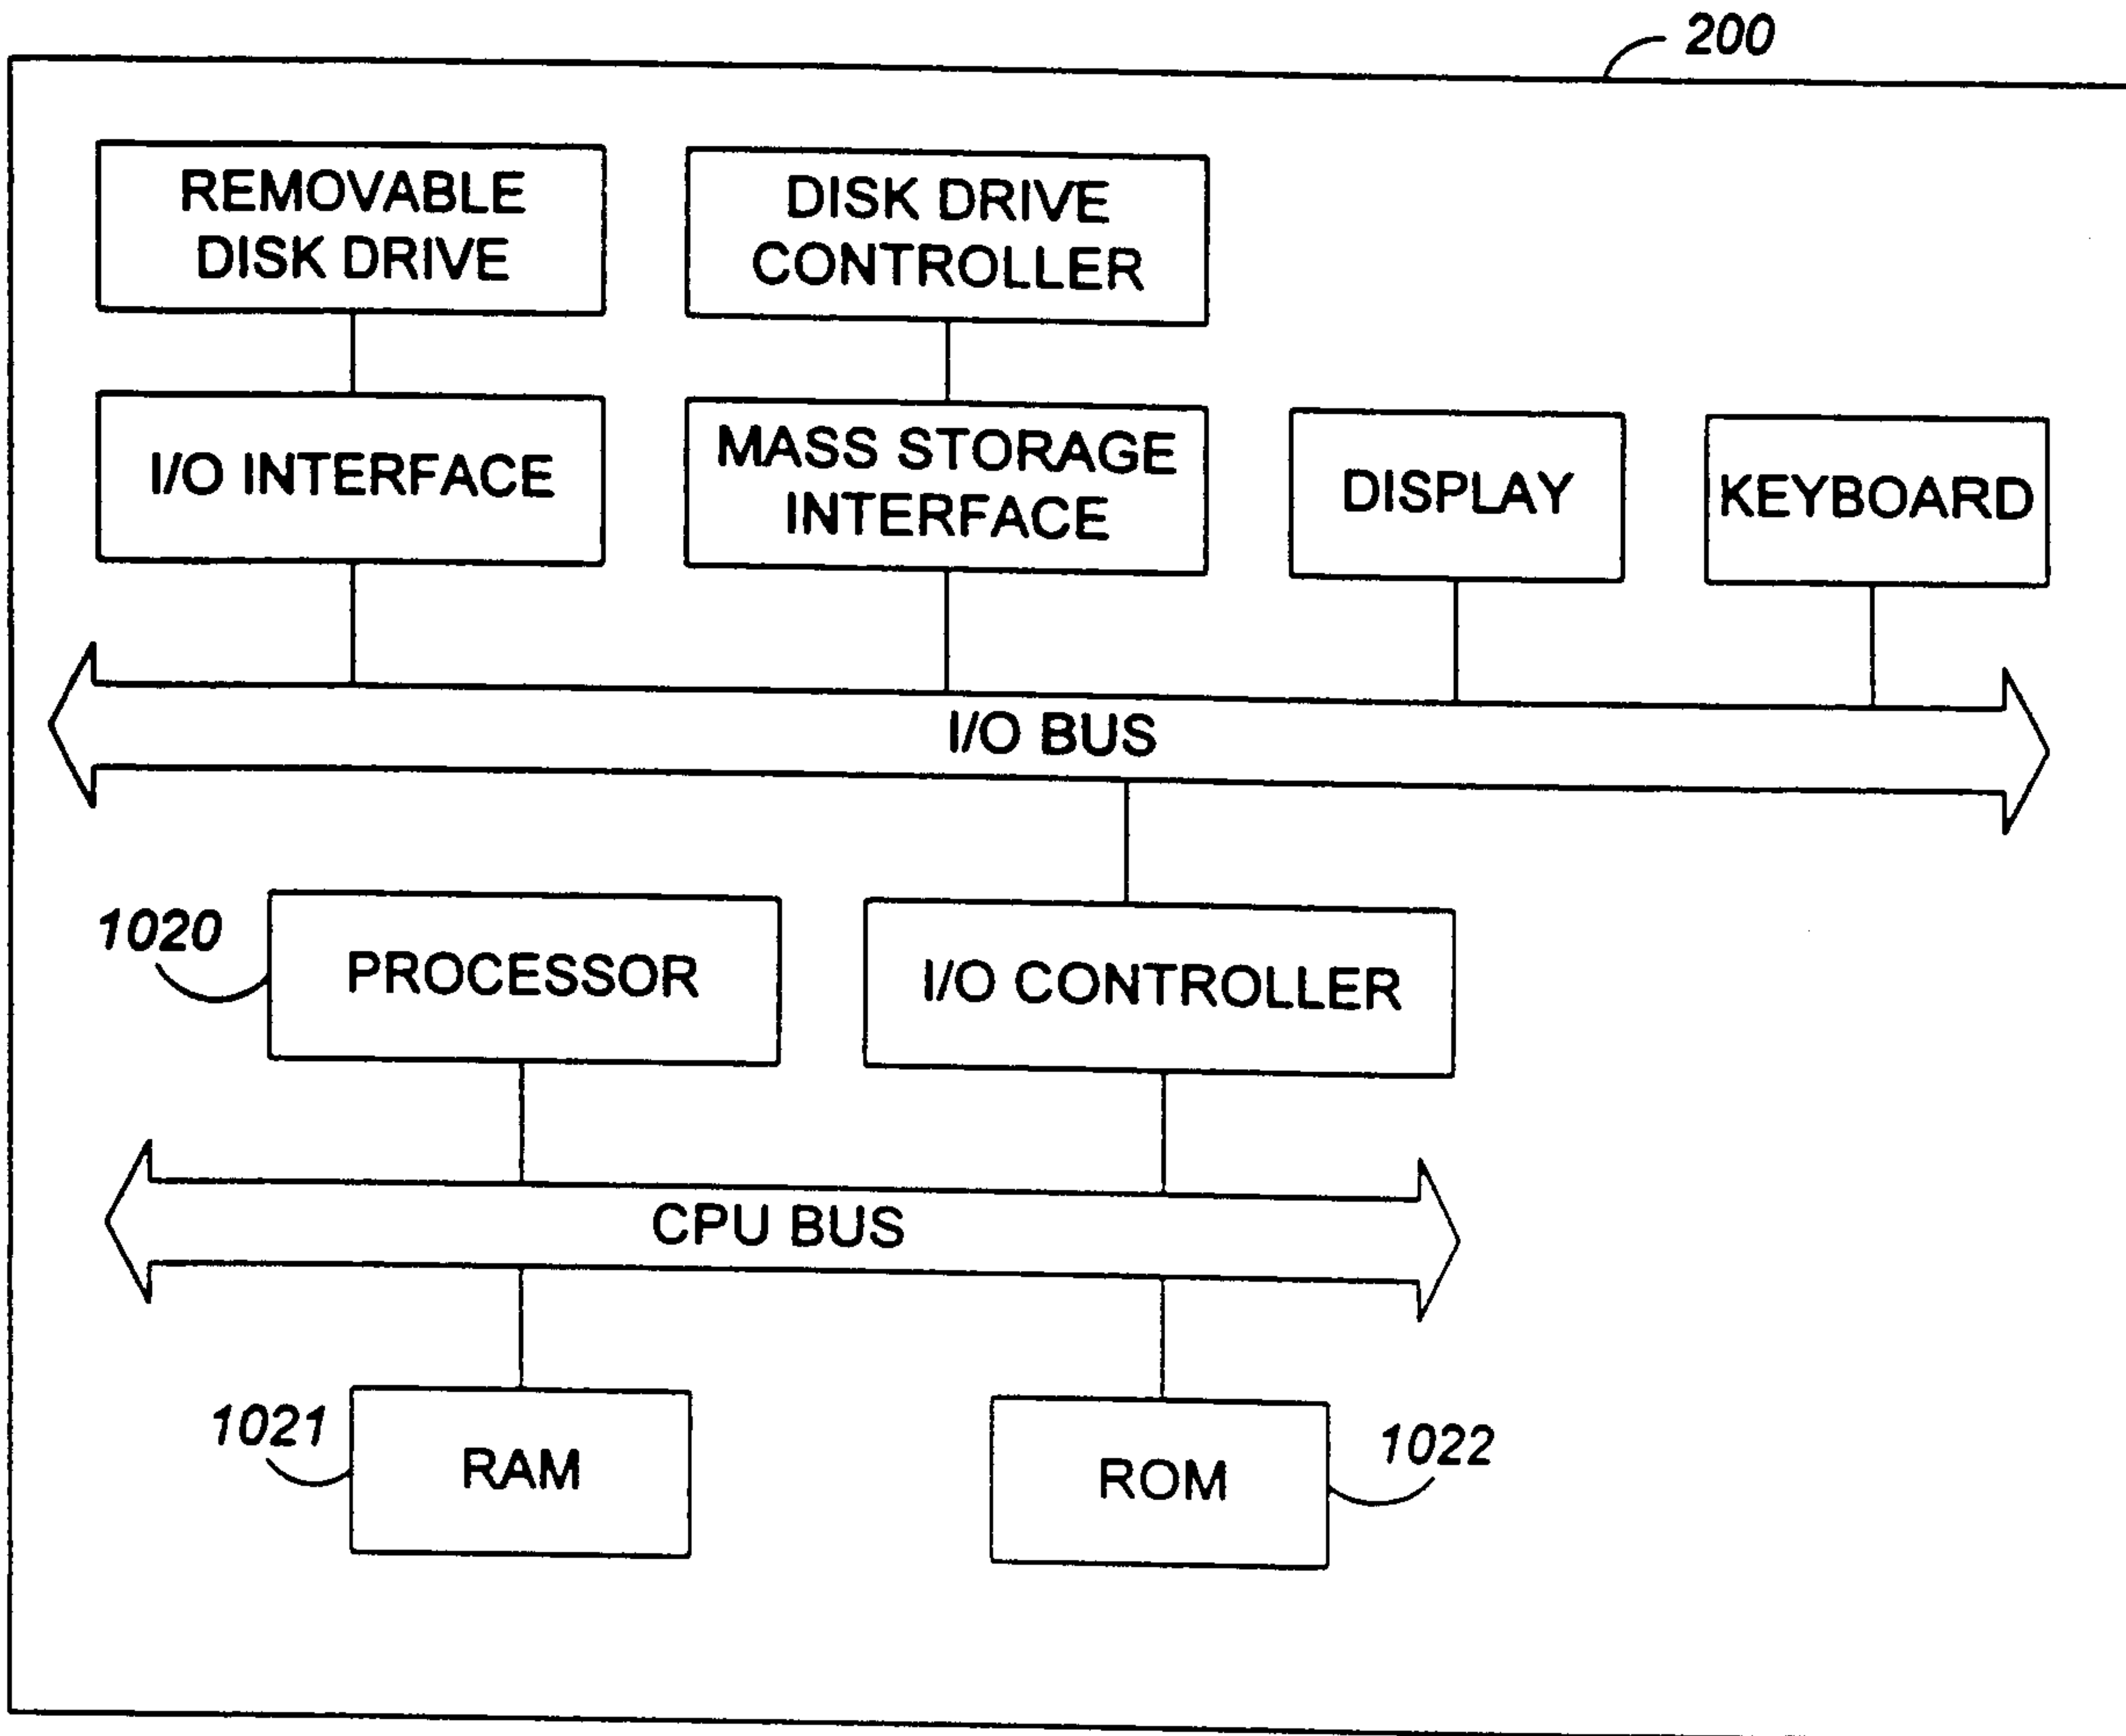


FIG. 8

FIG. 9



100

**Parse fragmentation expression to  
an Intermediate Range  
Representation**

102

**Convert Intermediate Range  
Representation to N Dimensional  
Integer Space**

104

**Build Index Tree to Represent  
Search Space**

106

**Store tree permanently as part of a  
catalog**

108

