

## The DC-Tree: A Fully Dynamic Index Structure for Data Warehouses

Martin Ester, Jörn Kohlhammer, Hans-Peter Kriegel  
University of Munich, Oettingenstr. 67, D-80538 München  
{ester | kriegel}@dbs.informatik.uni-muenchen.de jkohlham@crcg.edu

### Abstract

*In a data warehouse, updates are typically collected and performed periodically in a batch mode, e.g., over night. This standard approach of bulk incremental updates to data warehouses has some drawbacks. First, the average runtime for a single update is small but the total runtime for the whole batch of updates may become rather large. Second, the contents of the data warehouse is not always up to date. In this paper, we introduce the DC-tree, a fully dynamic index structure for data warehouses modeled as a data cube. This new index structure is designed for applications where the above drawbacks of the bulk update approach are critical. The DC-tree is a hierarchical index structure - similar to the X-tree - exploiting the concept hierarchies typically defined for the dimensions of a data cube. We conducted an extensive experimental performance evaluation using the TPC-D benchmark data. Our results demonstrate that the DC-tree yields a significant speed-up compared to the X-tree and the sequential search when processing general range queries on a data cube.*

### 1. Introduction

A data warehouse [3] is a collection of data from multiple sources, integrated into a common repository and extended by summary information (such as aggregate views) that is used primarily in organizational decision making. Often, a data cube is used to model a data warehouse and a relational database is used for its implementation. A data cube [2] consists of several independent attributes, grouped into *dimensions*, and some dependent attributes which are called *measures*. A data cube can be viewed as a  $d$ -dimensional array whose cells contain the measures for the respective subcube.

Typical queries on a data cube involve a lot of data records and tend to be very expensive. Therefore, it is a common approach ([7], [8], [9]) to materialize the results of many of the relevant queries in order to speed-up query processing. This approach, however, fails in a dynamic envi-

ronment where the queries are not known in advance and where the number of possible queries becomes very large. In such an environment general range queries should be supported. A *range query* [6] specifies a contiguous range for each of the dimensions of the data cube and applies a given aggregation operator to the set of selected cells. Several multi-dimensional index structures for data warehouses ([6], [12]) have been proposed which store some derived information to efficiently support general range queries.

Typically, a data warehouse is not updated immediately when insertions and deletions on the operational databases occur. Updates are collected and applied to the data warehouse periodically in a batch mode, e.g., each night [9]. Then, all derived information such as index structures has to be updated as well. This approach of bulk incremental updates, however, has two drawbacks:

(1) While the average runtime for one update is small, the total runtime for the whole batch of updates is rather large [12]. Bulk incremental updates require a considerable time window where the data warehouse is not available for OLAP. Global companies with branches all over the world, however, will more and more want to have their data warehouse available 24 hours a day.

(2) The contents of the data warehouse is not always up to date. In many applications this may not be necessary, but it may become critical in very dynamic applications such as stock markets or the WWW.

In this paper, we introduce the DC-tree, a fully dynamic index structure for data warehouses which avoids these two drawbacks. This new index structure is designed for applications where the above drawbacks of the well-known approaches are critical. This paper is organized as follows. Section 2 discusses related work. Section 3 presents the concepts of the DC-tree and section 4 discusses the major algorithms for constructing and querying DC-trees. We conducted an experimental evaluation of the DC-tree which is reported in section 5. Section 6 summarizes the contributions of this paper and outlines some directions for future research.

## 2. Related work

In this section, we briefly review related work from the area of data warehousing as well as from the area of spatial index structures.

The *data cube* [2] has been introduced as a multi-dimensional model for data warehouses. It is defined by several independent attributes, the *dimensions*, and some dependent attributes which are called *measures*. Each cell of a data cube contains the measures for the respective subcube. Often, a subset of the data cube is materialized to speed-up query processing. [7] presents an efficient algorithm which selects a nearly optimal subset of all cells for materialization. The proposed approach is static, i.e. it is useful only for the initial load of the cube but does not support incremental changes on dynamic updates of the data warehouse.

Several one-dimensional index structures have been proposed for efficiently processing queries in a data warehouse and, in particular, bitmap indices have become popular in this context. In a *bitmap index*, leaf pages of an index structure do not contain lists of record ids but bit vectors with one bit for each data record. For instance, [11] discusses several types of bitmap index structures suitable for different query types. [10] introduces *bitmap join indices* which precompute binary joins in a data warehouse. Bitmap indices, however, are static because on the insertion of a data record all index entries have to be updated. Furthermore, one-dimensional index structures build secondary indices which do not impact the clustering of the database. Therefore, they show poor performance for multi-dimensional range queries of the data cube.

Several multi-dimensional index structures for data cubes have been developed. [6] introduces a generalized quad-tree where the entries of the nodes consist of the description of a subcube and the materialized maximum measure value for this subcube. This index structure efficiently supports range-max queries but the branch-and-bound-optimization cannot be applied to other query types. For range-sum queries in dense data cubes, multi-dimensional prefix sums of the data cube are precomputed. For sparse data cubes, a set of non-intersecting subcubes is found and the prefix sum is only computed for these dense regions. An R\*-tree is used to manage the minimum bounding hyper-rectangles and the prefix sums of the dense subcubes. Unfortunately, no experimental performance results are reported.

[12] introduces the *Extended Datacube Model* (EDM) which allows to represent a data cube and its underlying relational data in a uniform way. An EDM is mapped to a cubetree and realized by a collection of packed R-trees. Furthermore, an algorithm is presented to perform bulk incremental updates of the cubetree. An experimental evaluation demonstrates that the cubetree supports multi-dimensional range queries very efficiently. The bulk update

experiments show that the required I/O time is very high due to the huge size of the data cubes.

Many applications require the management of *spatial data* (data with a location in a multi-dimensional space) such as points, lines and polygons. To speed up query processing in spatial databases, many *spatial index structures* have been developed to restrict the search to the relevant part of the space (*cf* [4] for a survey). The *R-tree* [5] generalizes the 1-dimensional B-tree to  $d$ -dimensional data spaces, i.e. an R-tree manages  $d$ -dimensional hyperrectangles instead of 1-dimensional numeric keys. An R-tree may organize extended objects such as polygons using *minimum bounding rectangles* (MBR) as approximations as well as point objects as a special case of rectangles. To answer a range query, starting from the root, the set of MBRs intersecting the query range is determined and then their referenced child nodes are searched until the data pages are reached.

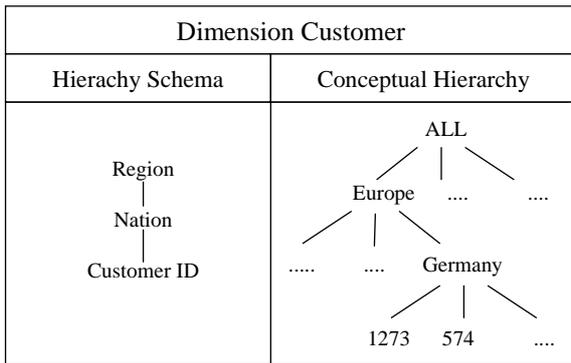
Since the overlap of the MBRs in the directory nodes grows with increasing dimension  $d$ , the R-tree and most other spatial index structures are efficient only for moderate values of  $d$ . Recently, several index structures such as the X-tree [1] have been designed which are also efficient for large values of  $d$ . If the standard topological split (considering properties of the MBRs such as their extension and their partitioning of dead space) results in high overlap, the X-tree tries to find an overlap-minimal split based on the split history. If the number of MBRs in one of the resulting partitions is below a given threshold the split would be too unbalanced and, therefore, the split algorithm terminates without providing a split. In this case, the current node is extended to become a so-called *supernode* with a multiple of the standard block size.

## 3. The DC-tree

In this section, we introduce the structure of the DC-tree, which is similar to that of the X-tree. To make use of the concept hierarchies for the dimensions, the MBRs are replaced by MDSs (minimum describing sequences).

### 3.1. Concept hierarchies

A *data cube* [2] consists of several functional attributes, grouped into *dimensions*, and some dependent attributes which are called *measures*. A data cube can be viewed as a  $d$ -dimensional array whose cells contain the measures for the respective subcube. If more than one functional attribute per dimension exists, these multiple attributes are organized by hierarchy schemata. A concept hierarchy is an instance of a hierarchy schema. Figure 1 shows an example for a dimension *Customer* and its functional attributes *Region*, *Nation* and *Customer ID*. *ALL* is the root of every concept hierarchy and denotes the union of all values in the concept hierarchy.



**Figure 1. Hierarchy Schema and Concept Hierarchy for dimension Customer**

We can define a partial ordering on a dimension using its concept hierarchy. This partial ordering is important for the split algorithms of the DC-tree and the definition of MDSs. The following definitions introduce these notions and the concept of a data cube itself more formally.

**Definition 1.** (Concept Hierarchy, Partial Ordering  $\leq$ , Hierarchy Level)

Let  $D_i, 1 \leq i \leq d$ , be sets of attribute values with  $ALL \in D_i$ . A *concept hierarchy* for  $D_i$  is a tree with the following properties:

- the nodes of the tree represent the elements of  $D_i$
- the root represents the special value  $ALL$
- the edges of the tree represent the “is-a” relationship between the two connected nodes.

Let  $a, b \in D_i, 1 \leq i \leq d$ . The *partial ordering*  $\leq$  for  $D_i$  is defined as follows:  $a \leq b$  if and only if  $a$  is equal to  $b$  or  $a$  is a (direct or indirect) son of  $b$  in the concept hierarchy of dimension  $i$ .

The *hierarchy level* of an attribute value is defined as the distance of the node with that attribute level from the leaves, i.e. the leaves have a hierarchy level of 0.

**Definition 2.** (Data Cube, Data Record)

We define a *data cube*  $D$  over the domains  $D_i, 1 \leq i \leq d$ , with  $m$  measures as follows: .

$$D \subseteq D_1 \times \dots \times D_d \times \mathfrak{R}^m$$

An element  $(a_1, \dots, a_d, x_1, \dots, x_m)$  of the datacube with  $a_i \in D_i, x_j \in \mathfrak{R}$  is called a *data record*.

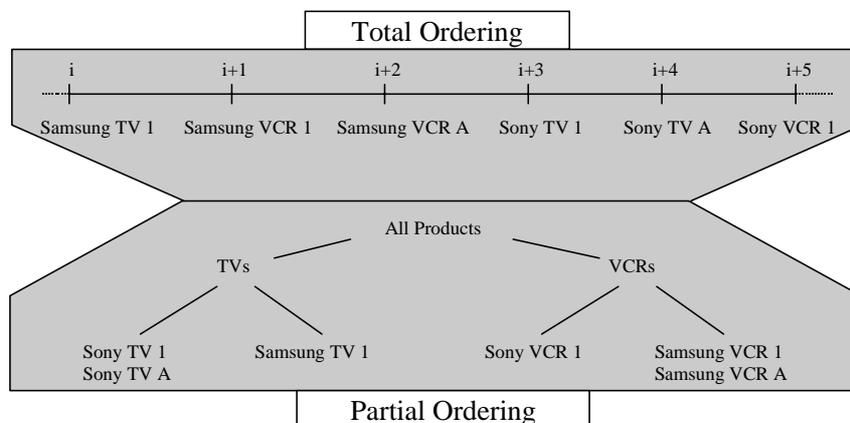
For example, in the concept hierarchy of figure 1  $Germany \leq Europe$  and  $a \leq ALL$  holds for each value  $a$ .

In order to apply a spatial access method such as the X-tree for indexing a data cube, one could define an arbitrary total ordering for each dimension. The advantages of a partial ordering compared to a total ordering are illustrated in figure 2 for the dimension *Product*. In a total ordering, an ID has to be assigned to every single product. A problem occurs when new products have to be inserted. For example, a new Samsung TV would receive an unfavourable ID, as it would naturally fit in between  $i$  and  $i+1$ . Furthermore, the attributes are not treated equally by a total ordering. The ordering in figure 2, for instance, would rather prefer range queries by makes than by product types.

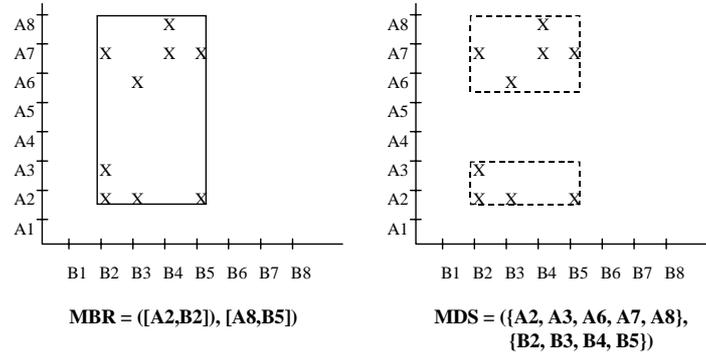
When using a partial ordering, the insertion of new products is natural because each leaf of the concept hierarchy is organized as a set. Although this structure prefers range queries by product types, a range query by makes can still be answered more efficiently by the leaf sets of the hierarchy than by a totally ordered structure. This holds because an algorithm knows which leaves contain the relevant data and does not have to scan the complete list of products.

The DC-tree stores one concept hierarchy per dimension and assigns an ID to every attribute value of a data record that is inserted. This assignment is performed to avoid the storage of long strings, but not to define a total ordering. Note that the DC-tree manages its concept hierarchies dynamically.

An ID is represented by a 32-bit integer. The highest four bits define the height of an ID in the concept hierarchy of its dimension to distinguish IDs from different levels.



**Figure 2. Total Ordering and Partial Ordering for Products**



**Figure 3. Sample MBR and MDS**

This representation of an ID has a fixed length and requires only 4 bytes. The DC-tree represents the concept hierarchies by means of dictionaries that store the ID of the father for each ID in one concept hierarchy. For the dimension *Customer* of figure 1, for example, every data record includes one value each for *Region*, *Nation* and *Customer ID*. The DC-tree assigns an ID to each value and updates its concept hierarchy for the dimension *Customer*.

### 3.2. Minimum describing sequences

*Minimum bounding rectangles (MBRs)*, used as the approximation method of the X-tree, are not appropriate for the DC-tree as demonstrated by the example in figure 3. For two partially ordered dimensions A and B, let an arbitrary total ordering be imposed on the dimensions given by the indices of the values  $A_i$  and  $B_j$ . On the lefthand side we have a MBR which assumes totally ordered dimensions. ([A2, A8], [B2, B5]) is a sufficient definition of this MBR because [A2, A8], e.g., implicitly includes each attribute value between A2 and A8. This does not hold if the dimensions are only partially ordered. Thus, we use so-called MDSs (*minimum describing sequences*) instead of MBRs.

On the righthand side of figure 3 the same data records are approximated by an MDS. Only attribute values occurring for at least one data record are included in the definition of this MDS. For instance, A1 and A4 are not contained in the first set of the MDS. Obviously, the MDS covers less dead space than the MBR. On the other hand, an MDS has to store more information and it has variable size. Note that all attribute values of a given dimension must belong to the same level of the concept hierarchy of that dimension. More formally, an MDS is defined as follows.

**Definition 3.** (Minimum Describing Sequence, MDS, Relevant Level)

Let  $D$  be a data cube with  $d$  dimensions  $D_i$ . Let  $S \subseteq D$  be a subcube of  $D$ , i.e. a set of data records  $(a_1, \dots, a_d, x_1, \dots, x_m)$  with  $a_i \in D_i, x_j \in \mathfrak{R}$ .

A *minimum describing sequence (MDS)* for  $S$  is a

sequence of entries  $(M_1, \dots, M_d)$  where

$M_i = (d_i, l_i)$  describes dimension  $i$  by a set of attribute values  $d_i \subseteq D_i$  which all belong to the *relevant level*  $l_i$  of the concept hierarchy of dimension  $i$ .

$d_i$  has to satisfy the following two properties:

1. (*coverage*) For all  $(a_1, \dots, a_d, x_1, \dots, x_m) \in S$  and for all  $i, 1 \leq i \leq d$ , there is some  $m_i \in d_i$  with  $a_i \leq m_i$ .
2. (*minimality*) If  $(N_1, \dots, N_d)$  with  $N_i = (e_i, l_i)$  satisfies the property of coverage, then for all  $i, 1 \leq i \leq d$ , and for all  $x \in d_i$ , there is some  $y \in e_i$  with  $x \leq y$ .

As an example, we examine the following data records (prior to an assignment of IDs according to chapter 3.1) for a data cube with dimensions *Customer*, *Supplier* and *Time* and one measure:

(Germany, North America, 1996, 310.27 \$)

(France, North America, 1997, 1245.80 \$)

For these data records, the MDS for the relevant levels 2, 2, and 2 is  $(\{\text{Germany, France}\}, \{\text{North America}\}, \{1996, 1997\})$  (see figure 9). For each dimension, only the attribute values of the relevant level of the concept hierarchy are depicted. If we chose the next higher concept level as the relevant level of the first dimension, we would obtain an MDS of  $(\{\text{Europe}\}, \{\text{North America}\}, \{1996, 1997\})$ .

The first MDS of a new DC-tree is the MDS  $(ALL, \dots, ALL)$ , i.e. the relevant level is initialized to the top level for each dimension. When performing a (hierarchy) split of a given node, one split dimension is selected and the relevant level of this dimension may be decreased by one for the MDSs of the two resulting subgroups. For example, the MDS  $(\{\text{Europe}\}, \{\text{North America}\}, \{1996, 1997\})$  may be split into two MDS  $(\{\text{Germany, France, Netherlands}\}, \{\text{North America}\}, \{1996, 1997\})$  and  $(\{\text{Switzerland, Greece, Italy}\}, \{\text{North America}\}, \{1996, 1997\})$ .

The last element of the above sample data record is the measure value according to the measure attribute of the data cube. The measure value is not part of the MDS, but is related to it and will be stored together with the MDS in each node of the DC-tree. The measure value for an MDS of a data node or a directory node is the aggregation (e.g.

the sum or the average) of the measure values of all data records covered by this MDS. Finally, the following definition introduces several notions for MDSs which are essential for the algorithms of the DC-tree.

**Definition 4.** (Size, Contains, Volume, Overlap, Extension)

Let  $D$  be a data cube with  $d$  dimensions  $D_i$ . Let  $M = (M_1, \dots, M_d)$ ,  $M_i \subseteq D_i$ , and  $N = (N_1, \dots, N_d)$ ,  $N_i \subseteq D_i$ , be MDSs. Let  $|S|$  denote the cardinality of a set  $S$ . The *size* of  $M$ , denoted as  $size(M)$ , is defined as:

$$size(M) = \sum_{i=1}^d |M_i|$$

$N$  *contains*  $M$  if for each dimension  $i$ ,  $1 \leq i \leq d$ , and for all  $m_i \in M_i$ , there is some  $n_i \in N_i$  with  $m_i \leq n_i$ .

The *volume* of  $M$ , denoted as  $volume(M)$ , is defined as

$$volume(M) = \prod_{i=1}^d |M_i|$$

The *overlap* of  $M$  and  $N$ , denoted as  $overlap(M, N)$ , is

$$defined as \quad overlap(M, N) = \prod_{i=1}^d |M_i \cap N_i|$$

The *extension* of  $M$  and  $N$ , denoted as  $extension(M, N)$ ,

$$is defined as \quad extension(M, N) = \prod_{i=1}^d |M_i \cup N_i|$$

Note that the definitions of *overlap* and *extension* assume that in each dimension  $i$  the elements of both  $M_i$  and  $N_i$  have to belong to the same level of the concept hierarchy. This is necessary because, e.g., the union of *American customers* and *North America* makes no sense.

## 4. Algorithms for the DC-tree

In this section, we present the major algorithms for constructing and querying DC-trees. While the insert algorithm is quite similar to that in the X-tree, we developed a completely new split algorithm.

### 4.1. Insert

Figure 4 shows the insert procedure for directory nodes. The data record to be inserted already contains the assigned IDs for its attribute values and its measure value.

After updating the measure value of the directory node, the choose-subtree algorithm selects a son *follow*, in which the data record will be further inserted. If *follow* had to be split as a result of this insertion, the directory node contains a new son and can now be overfilled itself. In this case, the split algorithm for directory nodes will be called. A successful split will then create a new brother of this directory node. If the split was not successful, a supernode will be created or, if the directory node has already been a supernode, the supernode will be enlarged.

```
int DWTDDirNode :: Insert (data *d)
{
    Update measure value;
    follow = choose_subtree (d);
    follow -> insert (d);
    If follow has been split :
    {
        Insert new son;
        If the capacity is reached :
        {
            Split (DWTDDirNode *NewBrother);
            If split was successful:
                return SPLIT;
            Else :
                return SUPERNODE; }
    }
    return NONE; } }
```

**Figure 4. Insert Algorithm for Directory Nodes**

### 4.2. Split

Again, the algorithm for directory nodes will illustrate the split procedure of the DC-tree (see figure 5). The directory split runs through all dimensions, until it finds an appropriate split or until it has examined every dimension.

The algorithm starts with selecting a split dimension by considering the hierarchy level of the elements of the MDS in the different dimensions. To induce a balanced structure of the DC-tree, the algorithm always selects the dimension with the highest hierarchy level of the elements of the MDS. For instance, if an MDS contains only the ALL value in one dimension and in all other dimensions attribute values from lower levels of their concept hierarchies, then the dimension with the ALL value will be selected as split dimension.

```
DWTDDirNode :: Split (DWTDDirNode *NewBrother)
{
    While found = FALSE and at least one dimension has not been processed :
    {
        Choose dimension by level of concept hierarchy;
        Adapt MDSs of entries to MDS of directory node;
        Hierarchy_split (Modified_MDSs[], Split_Dimension);
        If nodes are balanced and overlap is not too high :
            found = TRUE;}
    If non of the splits was successful :
    {
        Create supernode; }
```

**Figure 5. Split Algorithm for Directory Nodes**

Now the MDSs of the directory entries will be adapted to the MDS of the directory node. The reason for this procedure lies in the constraint for the operations with MDSs as described in chapter 3.2. All MDSs corresponding to the entries of a node have to be comparable to each other, i.e. in every dimension they have to contain elements of the same level in the corresponding concept hierarchy. Because the MDS of a directory node itself contains all MDSs of the entries of this directory node (see section 3.2), this directory MDS is the best choice for the adaptation of the MDSs.

The modified MDSs and the chosen split dimension are the parameters of the hierarchy split, an algorithm that is described in the next section. It results in two groups of MDSs. If these two groups are too unbalanced or their overlap is too high, another dimension will be selected. If no appropriate split is found for any dimension, a supernode is created. Note that the DC-tree allows to split a supernode just like a normal directory node. Such a split is performed if the directory node capacity multiplied by the number of blocks of the supernode is exceeded.

### 4.3. Hierarchy split

The hierarchy split is based on the quadratic split of [5], that was initially developed for the R-tree and its variants. The idea of this algorithm is to split a group of MBRs into two subgroups by first choosing two MBRs, so-called seeds, that should not be contained in the same subgroup. Then each remaining MBR is inserted into one of the subgroups, according to certain criteria such as the volume of the resulting group. Because each run of the while-loop considers each remaining MBR for the next insertion, the runtime of the algorithm is quadratic.

The hierarchy split shown in figure 6 has the same basic structure as the original quadratic split in the R-tree, but it exploits the partial ordering induced by the concept hierarchy. First, two seeds are being chosen from all MDSs. These two MDSs form the initial two groups of the algorithm. In every run of the while loop, the algorithm has to

make two decisions: Which MDS will be inserted next and into which group will this MDS be inserted?

The first decision is made by using the split dimension. The purpose of splitting along a split dimension is to obtain two groups with disjunct attribute values in the split dimension. The algorithm depicted in figure 6 considers the extension of the groups just in the split dimension (differing from [5]). It selects a group such that the new MDS and the MDS of the group share as many attribute values as possible in the split dimension. In the best case, the two resulting groups will contain disjoint attribute values in the split dimension.

The second decision is based on three criteria: overlap, extension and volume (in the order of their importance). Thus, the chosen MDS will be inserted into the group with the least resulting overlap of the two groups. Should the overlap be equal, the extension will be considered. If even the volume increase or the volume do not allow a decision, any of the groups will be chosen.

### 4.4. Range query

To demonstrate how the DC-tree makes use of the measure values stored in the entries of the tree, we will discuss the range query algorithm for directory nodes (see figure 7). The query range for this algorithm is specified by an MDS, the so-called *range\_MDS*.

The range query algorithm runs through every entry of the directory node. The for-loop makes the two MDSs comparable to each other. It is similar to the one in the split algorithm (see figure 5) but in this case here we do not know which of the two MDSs contains the higher level attribute values.

If the overlap between the *range\_MDS* and the MDS of the entry is empty, the entry is not relevant for the query and the result remains as it is. If the MDS of the entry is fully contained in the range, then the measure value stored in the son node referenced by the current directory entry is added to the result. Otherwise, if the MDS of the entry and the range overlap each other, we cannot use the measure

```

Hierarchy_Split (MBAs[], split dimension)
{
    Compute the covering MDS for each pair of MDSs;
    Initiate the two seeds by choosing the pair with the largest MDS;
    While remaining MDSs exist :
    {
        Compute the enlargement of the two groups in the split dimension
        for each remaining MDS;
        Choose the MDS by the greatest difference between the
        enlargements of the two groups in the split dimension;
        Insert this MDS into the group with the minimum resulting overlap
        between the groups;
        Resolve ties by choosing the minimum sum of extensions;
        Resolve further ties by choosing the minimum sum of volumes; } }

```

Figure 6. Hierarchy Split

```

double DWTDDirNode :: Range_Query (range_MDS)
{
    result = 0.0;
    For each directory entry :
    {
        For each dimension :
        {
            If the range_MDS and the MDS of the entry are not on the
            same level in the current dimension, adapt the MDS with
            the lower level to the one with the higher level; }
        If there is overlap between range_MDS and the MDS of the entry :
        {
            If the MDS of the entry is contained in the range_MDS :
                Result += measure value of entry;
            Else :
                Result += entry -> Range_Query (range_MDS); }
    }
    return result;}

```

**Figure 7. Range Query Algorithm**

value and have to recursively call the range query for the son node. Note, that in this case the aggregation *SUM* is being used within the algorithm. Any other aggregation, e.g. *AVERAGE*, would have to be treated accordingly.

## 5. Performance evaluation

The performance of the DC-tree is evaluated on a realistic data cube. We compare the DC-tree with the X-tree and the sequential search. The sequential search was evaluated on the same machine as the DC-tree and the X-tree. This allows further comparisons with other index structures which also have been compared with the sequential search.

### 5.1. Test environment

All performance tests use the database of the TPC Benchmark D [13]. The intended data cube is created by SQL select operations on the TPC-D database. The output of these operations is stored in a flatfile which functions as the insert file for the DC-tree and for the two other index structures being compared with the DC-tree.

As not all attributes in the TPC-D database were important for this performance evaluation, the corresponding database schema was simplified (see figure 8). Thus, the resulting data cube consists of four dimensions: *Supplier*, *Customer*, *Part* and *Time*. Figure 8 involves that each data record contains 14 attributes organized in hierarchy schemata as illustrated in figure 9. The 14th attribute is the measure attribute *Extended Price*.

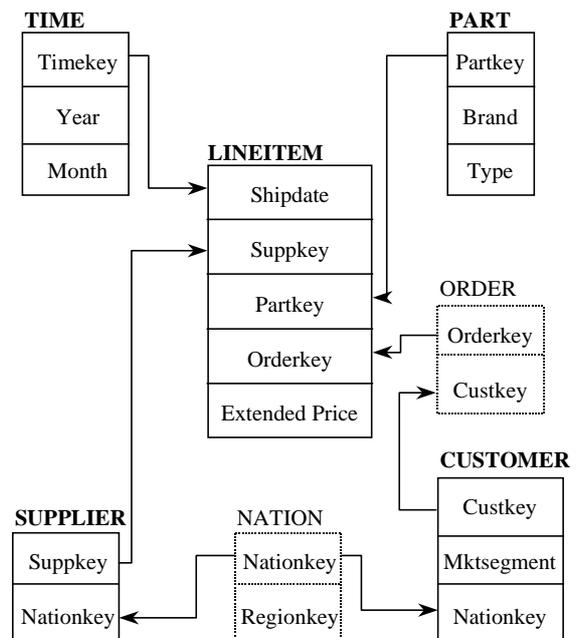
### 5.2. Generation of the range queries

To evaluate the results, the generation of the range queries and the processing of these queries in the X-tree and in the sequential search are important. Our algorithm for the generation of the range queries works as follows.

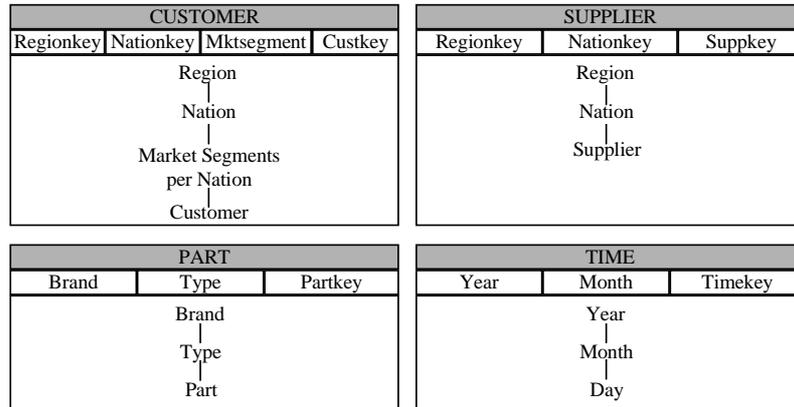
The algorithm randomly chooses a level in the concept hierarchy of each dimension that the set of the range\_MDS

in this dimension will be assigned to. For dimension *Customer*, e.g., the algorithm will randomly select Region, Nation, Market Segment or Customer (see figure 9). Depending on its choice, the range\_MDS will contain IDs of regions, nations, market segments or customers. The size of each set of the range\_MDS is limited by the selectivity. For instance, a selectivity of 25 % involves a range that contains up to 25 % of all attribute values of the chosen level in each dimension. These attributes values will again be randomly chosen.

The generation of a range query for a DC-tree is different from that for an X-tree, because the X-tree was developed for using MBRs and therefore cannot use MDSs and concept hierarchies. To use the existing range query algorithms of the X-tree for our test environment, we assigned a



**Figure 8. TPC-D Database Scheme**



**Figure 9. Dimensions of the Test Datacube**

dimension to each level of the concept hierarchies. Figure 10 shows the fourteen dimensions of the X-tree and the corresponding hierarchy levels of the DC-tree.

By using the total ordering of the IDs assigned to the attribute values by the insert procedure (see section 3.1), the range\_MDS can be converted to a range\_MBR for the X-tree. Thus, we can compare an index structure using totally ordered dimensions (the X-tree) with an index structure using partially ordered dimensions (the DC-tree).

The range query algorithm for the sequential search simply runs through every existing data record and determines whether this data record is contained in the range\_MDS or not. In the positive case, the measure value of the data record is added to the result.

### 5.3. Results

In this section, we present the results of the comparison between the DC-tree, the X-tree and the sequential search. The size of the underlying test data cubes ranges from 50,000 data records to 350,000 data records. The insertion time and the time per range query will be analyzed for different selectivities.

Figure 11 (a) shows the insertion time for the DC-tree and the X-tree for up to 350,000 data records. As the X-tree does not store concept hierarchies and avoids many computations the DC-tree has to do, the insertion time is significantly lower for the X-tree. However, figure 11 (b) depicts that the insertion of a single data record into the DC-tree takes only about 0.025 seconds on a HP C160 workstation

(64 PA-RISC) with 768 Megabyte RAM and HP UX 10.20. Thus, the dynamic insertion of data records has no significant impact on the runtime of a data warehouse and it is reasonable to keep the DC-tree up-to-date at all times.

Several tests with range queries of different selectivities were performed to compare the efficiency of the DC-tree to the X-tree and the sequential search. Figure 12 shows three comparative tests between the DC-tree and the X-tree for the selectivities 1 %, 5 % and 25 % as well as a comparison between the DC-tree and the sequential search. For all selectivities the time per query is determined as the average of 100 random queries.

To assure a fair comparison, the main memory available for the X-tree was restricted to the memory size that the DC-tree uses. In all performed tests the range queries are executed much faster on the DC-tree than on the comparative index structure. In fact, we obtain a speed-up of about 4.5 for range queries on the DC-tree compared with those on the X-tree.

When looking at the absolute numbers in figure 12 (a) - (c), range queries with selectivity 5 % are processed faster than the others. The reason for this fact lies in a trade-off between the level on which the DC-tree can completely answer a range query and the performance costs when executing the range queries with larger range-MDSs. The larger the query MDS, the higher is the probability that the MDS of an entry is fully contained in the query MDS. Thus, a larger query MDS yields a better performance. On the other hand, a larger query MDS involves more expensive computations of the overlap, because a large MDS consists of

CUSTOMER				SUPPLIER		
Regionkey	Nationkey	Mktsegment	Custkey	Regionkey	Nationkey	Suppkey
0	1	2	3	4	5	6

PART			TIME			
Brand	Type	Partkey	Year	Month	Timekey	Measure
7	8	9	10	11	12	13

**Figure 10. Dimensions of the X-tree for the Test Datacube**

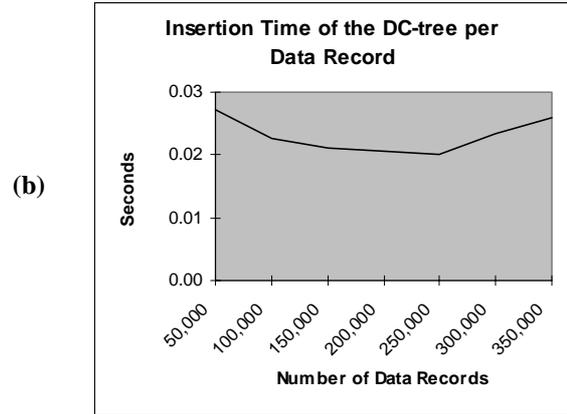
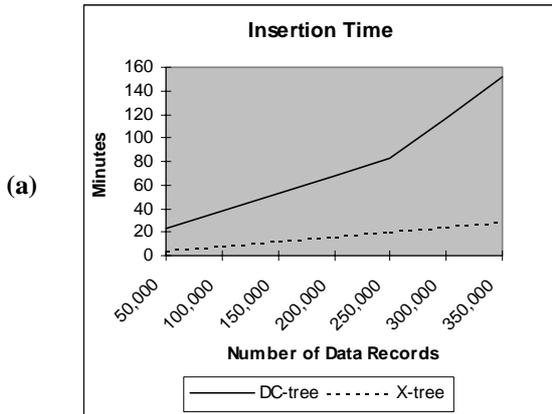


Figure 11. Insertion Time

large sets for the single dimensions. While the query MDSs of range queries with selectivity 1 % are in general too small to answer the query on a high level of the DC-tree and the query MDSs of range queries with a selectivity of 25 % involve very expensive computations, the range queries with selectivity 5 % seem to be the best compromise among the performed tests.

The comparison between the DC-tree and the sequential search, depicted in figure 12 (d), shows that the sequen-

tial search is no reasonable alternative. Even when using a selectivity of 25 % (the worst case for the DC-tree), we obtain a speed-up of 12.5 for range queries on the DC-tree compared to the sequential search.

Figure 13 depicts the node sizes, i.e. the average number of entries, for the two highest levels of the DC-tree below the root node. The node size increases linearly with increasing number of data records for the second highest level but the node size stabilizes at 15 entries on the highest

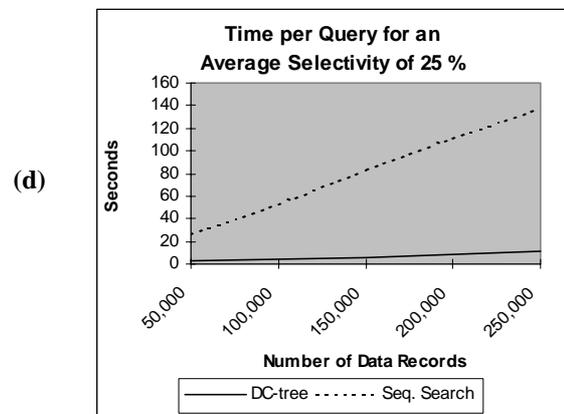
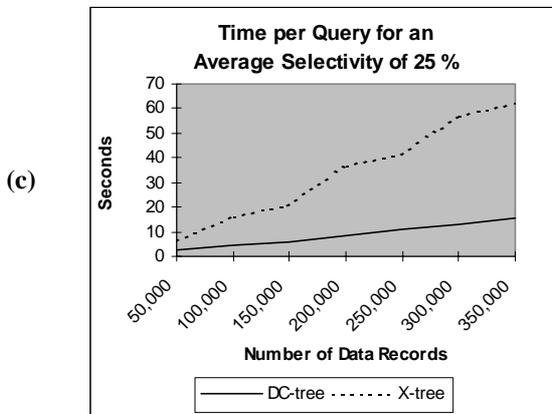
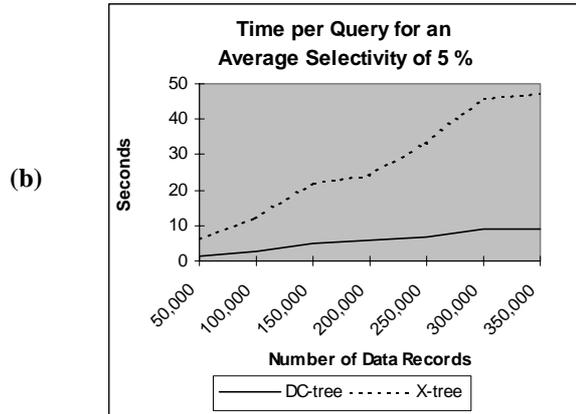
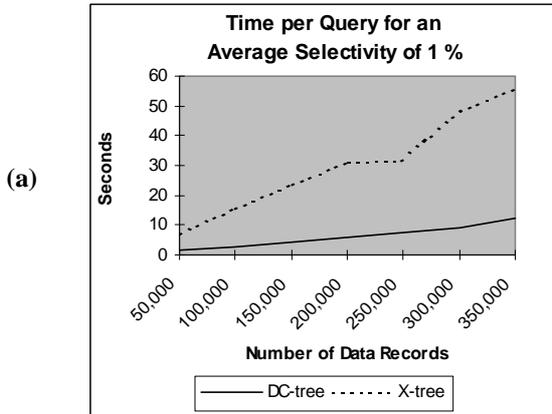


Figure 12. Query Time

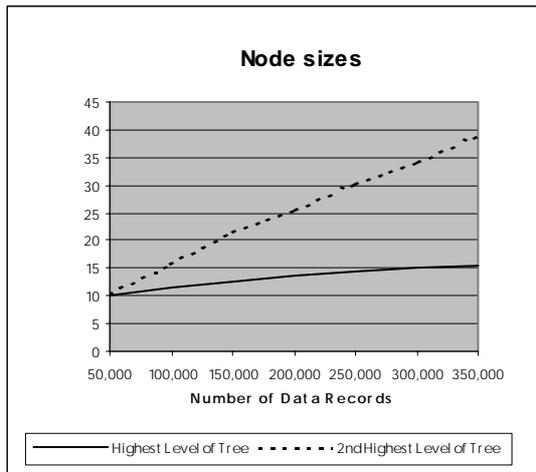


Figure 13. Node Sizes per Level

level. This effect is due to the way of splitting directory nodes. Every split results in two new directory nodes with MDSs that contain either less attribute values or attribute values of a lower level of the concept hierarchy in each dimension. Furthermore, the possible attribute values of a directory node's MDS are constrained by the MDS of its father node. Depending on the overall number of possible attribute values per dimension at some level of the DC-tree, it becomes hard to split a directory node simply because its MDS does not contain enough attribute values in any dimension, i.e. the MDS is already too special to be split further more. Due to the relatively small number of possible attribute values, the MDSs of the directory nodes already start to become rather special on the second highest level (see figure 13). Consequently, the split algorithm creates more and more supernodes or further enlarges existing supernodes, respectively. This results in an average number of entries of about 2.5 times the capacity of a regular directory node on the second highest level under the root node. This effect and possible improvements of the split algorithm will be further investigated in future work.

## 6. Conclusions

In this paper, we introduced the DC-tree, a fully dynamic index structure for data warehouses. This new index structure is designed for applications where the above drawbacks of the well-known approaches are critical. The DC-tree is a hierarchical index structure - similar to the X-tree - exploiting the concept hierarchies typically defined for the dimensions of a data cube. The DC-tree uses minimum describing sequences and the partial ordering of the attribute values induced by the concept hierarchies instead of minimum bounding rectangles and an artificial total ordering. Furthermore, for each minimum describing sequence in the directory the values of the measure attributes are materialized. Our experimental performance evaluation

on the TPC-D benchmark data demonstrated a significant speed-up compared to the X-tree and the sequential search when processing general range queries on a data cube.

Future work includes the following issues. The split algorithm of the DC-tree is rather expensive and, in particular, more expensive than the split algorithm of the X-tree. Therefore, alternative split algorithms should be investigated which have less than quadratic cost but nevertheless yield reasonably good splits. A data cube is typically implemented by using a relational DBMS. Thus, the DC-tree should be integrated into a commercial DBMS to evaluate its performance in this context.

## Acknowledgements

We thank Jörg Sander for fruitful discussions in early phases of this project.

## References

- [1] Berchtold S., Keim D. A., Kriegel H.-P.: "The X-Tree: An Index Structure for High-Dimensional Data", Proc. 22th Int. Conf. on Very Large Data Bases, Bombay, India, 1996, pp. 28-39.
- [2] Gray J., Bosworth A., Layman A., Pirahesh H.: "Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tabs and Subtotals", Proc. 12th Int. Conf. on Data Engineering, 1996, pp. 152-159.
- [3] Chaudhuri, S., Dayal, U.: "An Overview of Data Warehousing and OLAP Technology", ACM SIGMOD Record 26(1), March 1997.
- [4] Gaede V., Günther O.: "Multidimensional Access Methods", ACM Computing Surveys, Vol. 30, No. 2, 1998, pp. 170-231.
- [5] Guttman A.: 'R-trees: A Dynamic Index Structure for Spatial Searching', Proc. ACM SIGMOD Int. Conf. on Management of Data, Boston, MA, 1984, pp. 47-57.
- [6] Ho Ch., Agrawal R., Megiddo N., Srikant R.: "Range Queries in OLAP Data Cubes", Proc. ACM SIGMOD Int. Conf. on Management of Data, 1997, pp. 73-88.
- [7] Harinarayan V., Rajaraman A., Ullman J.D.: "Implementing Data Cubes Efficiently", Proc. ACM SIGMOD Int. Conf. on Management of Data, 1996, pp. 205-216.
- [8] Huyn N.: "Multiple-View Self-Maintenance in Data Warehousing Environments", Proc. 23rd Int. Conf. on Very Large Data Bases, Athens, Greece, 1997, pp. 26-35.
- [9] Mumick I. S., Quass D., Mumick B. S.: "Maintenance of Data Cubes and Summary Tables in a Warehouse", Proc. ACM SIGMOD Int. Conf. on Management of Data, 1997, pp. 100-111.
- [10] O'Neil P., Graefe G.: "Multi-Table Joins through Bitmapped Join Indices", SIGMOD Record 24(3), 1995, pp. 8-11.
- [11] O'Neil P., Quass D.: "Improved Query Performance with Variant Indices", Proc. ACM SIGMOD Int. Conf. on Management of Data, 1997, pp. 38-49.
- [12] Roussopoulos N., Kotidis Y., Roussopoulos M.: "Cubetree: Organization of and Bulk Incremental Updates on the Data Cube", Proc. ACM SIGMOD Int. Conf. on Management of Data, 1997, pp. 89-99.
- [13] Transaction Processing Council (TPC): "TPC Benchmark D", <http://www.tpc.org>, 1998.