

# An Efficient Strategy for Tiling Multidimensional OLAP Data Cubes

Paulo Marques, Paula Furtado\*, Peter Baumann

{marques, furtado, baumann}@forwiss.tu-muenchen.de

Forwiss – Bavarian Research Center for Knowledge Based Systems

Orleansstrasse, 34, D-80801 München, Germany

## Abstract

Computing aggregates over selected categories of multidimensional discrete data (MDD) cubes is the core operation of many on-line analytical processing (OLAP) systems. In order to support efficient computations of these aggregates in a multidimensional OLAP (MOLAP) system, a careful design of the database storage architecture must be undertaken. In particular, tiling (i.e., subdivision of an MDD cube into blocks) plays a crucial role in the overall performance of the system. Nevertheless, to our knowledge, the current MOLAP systems only provide regular tiling. In this paper we present a more efficient tiling strategy for partitioning MDD cubes in the context of MOLAP systems. We argue that, by providing explicit semantic information about the categories along each dimension of the MOLAP data cubes, a more accurate and efficient tiling strategy – Directional Tiling – can be implemented. Finally, we report the performance results obtained by using the described approach.

## 1. Introduction

Multidimensional on-line analytical processing (MOLAP) systems store data using sparse arrays of a basic cell type. A careful design of the storage management system for these sparse arrays is very important in order to provide efficient access to the data. Storage management involves aspects such as index management, clustering, tiling and compression.

Tiling is the name given to the process of dividing a multidimensional discrete data array (MDD array) into sub-arrays, so it can be stored and accessed in a database management system (DBMS). Tiling is a specially important operation since, if not carefully done, severe performance penalties may incur later on, when accessing the data. Multidimensional discrete database management systems which support MOLAP usually take the approach suggested in [Sarawagi94] of

---

\* Ph.D. supported by a PRAXIS XXI scholarship.

dividing the MDD arrays into regular blocks. As far as we know from the literature, no attempts have been made to optimize tiling for MOLAP systems except for the traditional use of regular tiling. In our approach, we specifically optimize the tiling strategy for this application area, leading to significant increases in performance.

The approach and experiments described in this paper were implemented into the RasDaMan system ([Baumann97], [Furtado98]). This system is a multidimensional discrete database management system, which supports arbitrary base cell types and number of dimensions. Also an advanced query language for multidimensional operations has been developed in the context of the system. In this paper we focus on “directional tiling”. This tiling strategy was developed and implemented in RasDaMan for increased performance in MOLAP applications.

The rest of the paper is organized as follows: Section 2 explains the importance of using non-regular tiling in MOLAP systems and Section 3 discusses the implemented algorithm. We present performance results in Section 4, and finally the conclusions and future work in Section 5.

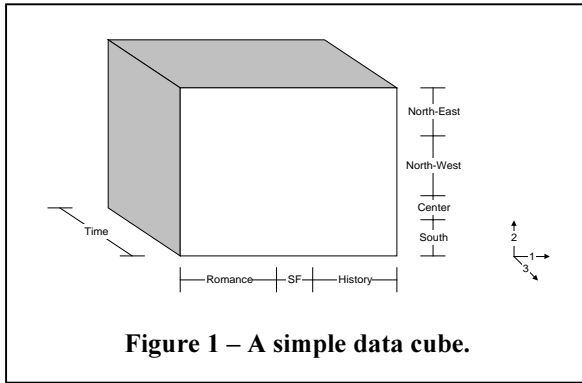
## **2. Motivation**

Multidimensional databases and in particular MOLAP systems are gaining increasing importance in today’s scientific and commercial communities. In order for those systems to be successful, it is critical that, not only they model and process the stored information in a better way than traditional systems do, but also that they do it efficiently.

Common approaches for storage management in multidimensional databases, namely, tiling, clustering, indexing and compression schemes, have been shown to deliver the efficiency needed for that success ([Sarawagi94], [Chen95], [Zhao98]). Even so, regarding tiling, most systems available today only implement regular aligned tiling.


It is our believe that, by providing the database management system with relevant semantic information about the data stored, more appropriate tiling schemes can be implemented, leading to an overall improvement of the systems’ performance.

In MOLAP systems computing Cube-By operations in specific categories ([Gray96], [Zhao97]), significant amounts of data must be retrieved. Although most algorithms focus on keeping the least possible amount of data in memory and retrieving as few blocks from disk as possible, a large quantity of irrelevant information may still have to be fetched when selecting sub-cubes from the original data due to inadequate tiling of the MDD arrays.



To make things clearer, let's consider the simple data cube in Figure 1. This cube represents the book sales of a publisher, over time, in a certain country. Dimension 1 contains the book titles, divided into three categories<sup>1</sup>: Romance, Science Fiction and History. This dimension does not, obviously, contain the same number of books in each category. In dimension 2, the numerous stores to which the publisher sells are represented, grouped into four regions: Northeast, Northwest, Center and South. Finally, dimension 3 represents time.

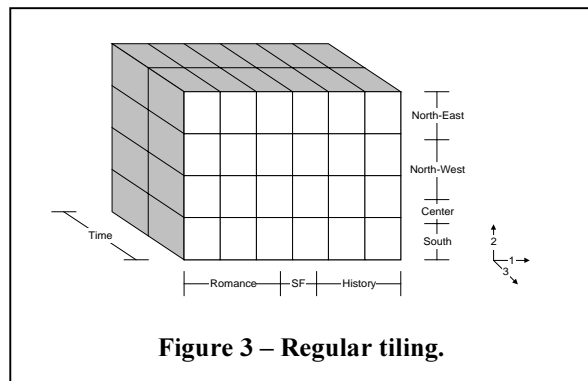
After having performed tiling, MOLAP systems store the generated blocks<sup>2</sup> in Binary Large Objects (BLOBs). Currently, systems only provide efficient access to full or linear parts of BLOBs. These BLOBs are stored in a linear fashion, and accessing multidimensional parts of them can be very inefficient ([Sarawagi94], [Furtado97]). This is illustrated in Figures 2a, 2b and 2c.



In Figure 2c, the pages marked in a darker shade must all be read in order to access the data requested by the user. Even so, most of those pages will contain almost no relevant data for the access taking place.

Let's suppose that regular tiling has been applied to the cube of Figure 1. A possible result of regular tiling on this cube is shown in Figure 3.

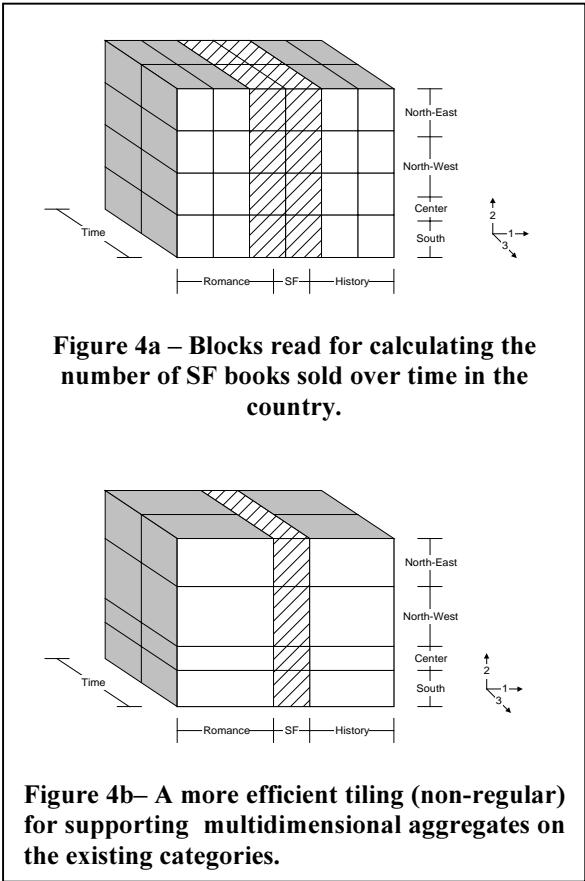
Using this form of tiling, if multidimensional aggregates are performed in order to find out, for instance, the total number of books sold in a certain category, over time, or the



<sup>1</sup> By category we mean a set of dimensional elements corresponding to the same parent element in a higher level of the dimension hierarchy. For example, "SF Books", in Figure 1, is a category of dimension 1 (books).

<sup>2</sup> The terms "block" and "tile" are used synonymously in the following.

total number of books sold in a region, by category over a specific time period, then a significant amount of irrelevant data to the operations taking place must be read. As an example, let's consider that the user wants to find out the total number of SF books sold in the whole country. To perform this multidimensional aggregate, all the blocks marked in a different pattern in Figure 4a must be read. What happens is that information about books in the Romance and History categories must be read, since they are stored in the same blocks as SF books are, although they are of no use to the computations requested by the user.



We argue that, by using non-regular tiling as the one shown in Figure 4b, a more efficient support of multidimensional aggregates on selected categories is possible.

In this figure, the borders of the generated blocks, after tiling, are aligned with the existing categories on each dimension. Thus, the percentage of data read from the blocks that is actually used in the computations is much higher.

Nevertheless, several issues have to be taken into account when performing non-regular tiling. For instance, the generated blocks should not be too large since it may degrade performance of access types that are not based on category aggregations. Thus, sub-tiling of the generated blocks may be required, in order to assure a good overall system performance. Also, the user may have “directional accessing patterns” and those

should be taken into account when performing tiling and especially when sub-tiling. These details will be further discussed in the next section.

### 3. The “Directional Tiling” Strategy

In order to perform non-regular tiling in a MOLAP data cube, the system has to have semantic information about the existing data in it. This semantic information consists of, for each dimension, linear ranges corresponding to the existing categories. Providing this information to the system is analogous to a database designer or administrator informing a relational database management system (RDBMS) on which columns an index should be created. Although this is not strictly necessary for the operation of the system, if done, the user benefits from increased performance, since the DBMS is able to take better decisions on how to store and access the data, minimizing the amount of information that has to be read from disk.

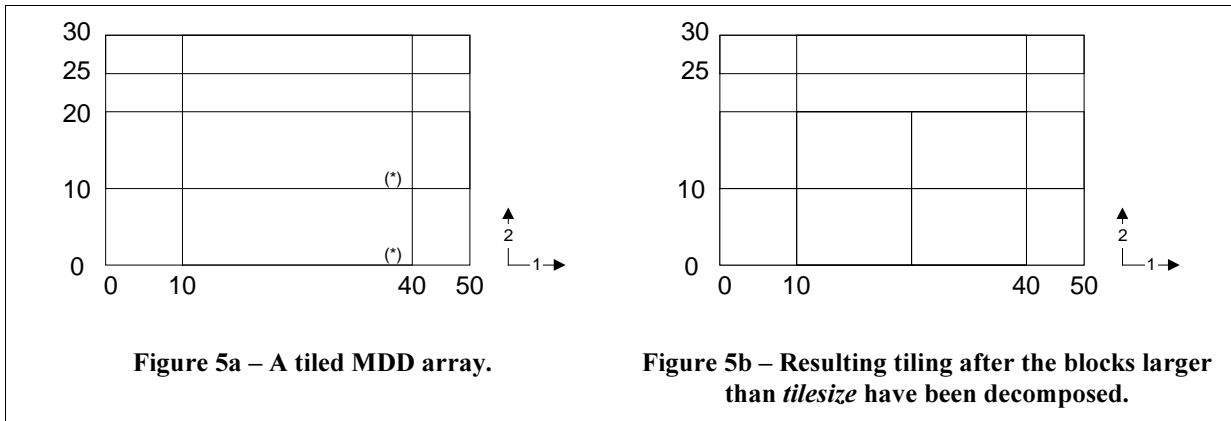
Thus, if a data cube of  $n$  dimensions is about to be loaded into the database, the user provides the following specification:

$$part_i = [l_{i,0}, l_{i,1}, \dots, l_{i,k_i}], i = 1..n.$$

For each dimension  $i$ , ranging from 1 to  $n$ , the user specifies a partition –  $part_i$ , which contains the limits of each category for that dimension –  $l_{i,0}$  to  $l_{i,k_i}$ , where  $k_i$  is the number of existing categories in the dimension<sup>3</sup>. Each partition is an ordered set having:  $l_{i,0} < l_{i,1} < \dots < l_{i,k_i}$ , where  $l_{i,0}$  corresponds to the lower limit and  $l_{i,k_i}$  corresponds to the higher limit of the domain in the  $i^{\text{th}}$  dimension.

The algorithm uses this partition information to decompose the data cube into blocks such that, in each dimension, the edges of the generated blocks coincide with the edges of the existing categories. After this decomposition, if any of the resulting blocks is larger than a specified size,  $tile\ size$ , then regular tiling is performed on those blocks.  $Tile\ size$  represents the maximum allowed size of a block in order that efficient access to it can still be performed. Typically,  $tile\ size$  is a multiple of the underlying storage system data unit (page size).

In Figure 5a a 2D MDD array is shown, with three categories along dimension 1 and four along dimension 2. The result of performing tiling on this array is twelve smaller blocks, which are stored in separate BLOBs. Let’s now assume that the two blocks marked with an (\*) on Figure 5a are larger than  $tile\ size$ . If so, regular tiling is applied to these blocks. Figure 5b shows the final tiling for this data cube.



Sub-tiling the blocks that are larger than a certain size is critical since there may exist data cubes that do not have any categories in certain dimensions, or have such large categories that performance would deteriorate, when executing operations not based on category aggregations. By monitoring the size of the resulting blocks and sub-tiling them if they are larger than a certain specified threshold, the algorithm guarantees that no large asymmetries in the resulting sizes arise and that, therefore, all regions of the data cube can be accessed efficiently.

<sup>3</sup> The ranges are shown in terms of integers because, as the system works with discrete multidimensional data, existing categories and elements in each dimension are mapped onto the  $\mathbf{Z}$  set.

When informing the system about the categories location in the data cube, the user is also free not to specify partitions for certain dimensions. When this happens, the user decides if such dimensions are to be considered “preferred access directions”, where tiling should not be performed, or just “ordinary directions” where tiling can be performed normally. If a partition specification is given by  $part_j = []$  then the corresponding dimension should be considered a “preferred access direction” and tiling should not be done along it. The reasoning behind this is that some applications may access the data in certain directional patterns. Figure 6a shows a 2D MDD array with three categories along dimension 2, and none along dimension 1. Let’s assume that the block marked with (\*) in this figure is larger than *tilesize*. Now, if the user knows that he typically accesses the information as complete rows along dimension 1, or even not complete rows but linearly, then this dimension should be specified as a “preferred access direction”. By doing this, sub-tiling will not be performed on that dimension, unless as a last resort<sup>4</sup> and will be done along whichever “non-preferred access directions” exist. In the example this is the case of dimension 2.




Figure 6b shows the result of sub-tiling of the (\*) block, where dimension 1 is specified as a “preferred access direction”, and Figure 6c shows a possible result of sub-tiling the (\*) block, if this dimension is not specified as a “preferred access direction”. In both figures the generated blocks from (\*) are shown in a different pattern. If tiling is done like shown in Figure 6c, and the user accesses the object linearly along dimension 1, then performance will be severely degraded. This is so because it is necessary to access both generated blocks for reading a full row. This is not the case with the sub-tiling shown in Figure 6b, where only accessing one block is required for the same kind of query.

If a certain dimension  $j$  does not have any categories associated with it, and the user does not want to specify it as a “preferred access direction”, then the partition specification will consist of the

<sup>4</sup> For instance, if the “volume” of the block being considered for tiling, without taking the current “preferred access direction” dimension into account, is already larger than *tilesize*.

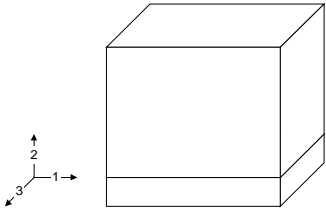
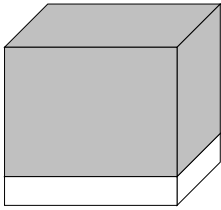
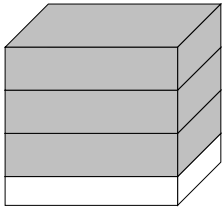
lower and upper bounds of the domain for that dimension:  $part_j = [d_{low,j}, d_{high,j}]$ . Here  $d_{low,j}$  and  $d_{high,j}$  are the lower and upper bounds of the domain for the  $j^{th}$  dimension.

When performing sub-tiling of the blocks larger than  $tilesize$ , the system tries to create  $n$ -dimensional blocks which are full domain cuts on the dimensions which are specified to be “preferred access directions” and cuts of  $edgesize$  in all other dimensions.  $Edgesize$  is given by:

$$edgesize = \left\lceil \sqrt[n-k]{\frac{tileze / cellsize}{\prod_{i=1}^k (d_{hogh,i} - d_{low,i} + 1)}} \right\rceil$$

In this expression,  $n$  is the number of dimensions of the MDD array,  $tilesize$  is the maximum size of a block, in bytes,  $cellsize$  is the size of a base cell of the array (also in bytes),  $k$  is the number of dimensions which are not “preferred access directions”<sup>5</sup>, and finally,  $d_{high,i}$  and  $d_{low,i}$  represent the higher and lower limits of the data cube domain in the  $i^{th}$  dimension.

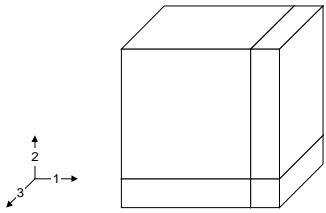
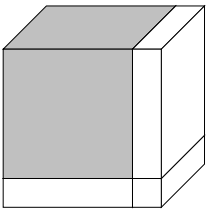
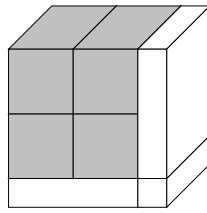
One important point that should be made is that the “non-preferred access directions” represent degrees of freedom in which the system can perform sub-tiling. If possible, no tiling should ever take place on the specified “preferred access directions”. Figures 7 and 8 illustrate the use of the “non-preferred access directions” to perform sub-tiling.

		
<ul style="list-style-type: none"> <li>– The user only specifies two categories in one of the dimensions (dimension 2).</li> </ul>	<ul style="list-style-type: none"> <li>– The darker block is larger than <math>tilesize</math>. Sub-tiling must take place.</li> <li>– No categories are specified along dimensions 1 and 3, which are considered “preferred”.</li> </ul>	<ul style="list-style-type: none"> <li>– As dimensions 1 and 3 are considered “preferred access directions” tiling is not done along them. The system only performs tiling along dimension 2.</li> </ul>

**Figure 7 – Tiling in a 3D data cube, with one degree of freedom.**

In Figure 7, two categories are specified in a three-dimensional domain. In this example, the system has only one degree of freedom for performing sub-tiling.

<sup>5</sup> This corresponds to the number of dimensions which have categories specified.

		
<ul style="list-style-type: none"> <li>- The user specifies four categories in the data cube. Two in dimension 1 and two in dimension 2.</li> </ul>	<ul style="list-style-type: none"> <li>- The darker block is larger than <i>tilsize</i>. Sub-tiling must be performed.</li> <li>- No categories are specified along dimension 3, which is considered “preferred”.</li> </ul>	<ul style="list-style-type: none"> <li>- As dimension 3 is considered a “preferred access direction”, tiling is not done along it. The system performs tiling along dimensions 1 and 2.</li> </ul>

**Figure 8 – Tiling in a 3D data cube, with two degrees of freedom.**

In Figure 8, for a similar domain, four categories and one “preferred access direction” are specified. In this case, the system has two degrees of freedom when performing sub-tiling on the existing data cube.

To conclude this section, we present the algorithm that performs directional tiling in the system<sup>6</sup>:

```

1. Directional_Tiling(DataCubeDomain, CategoriesSpecification, tilsize)
2. =====
3.
4. Result = {}
5. BlockIntervals = Create_Blocks(DataCubeDomain, CategoriesSpecification)
6.
7. For each b ∈ BlockIntervals
8.   If (b.size > tilsize) Then
9.     If (CategoriesSpecification.Unspecified_Dimensions > 0) Then
10.      edgesize = Calculate_Edgesize()
11.
12.      For i=1 to DataCubeDomain.Number_Dimensions Do
13.        If (CategoriesSpecification.part[i].Is_Not_Specified) Then
14.
15.
16.
17.
18.
19.
20.
21.
22.
23.
24.
25.
26.
27.
28.
29.
30.
31.
32.
33.
34.
35.
36.

```