

TEMPLATE FOR ENTRY
in Encyclopedia of Database Systems:

GRID FILE

Apostolos N. Papadopoulos	Yannis Manolopoulos	Yannis Theodoridis	Vassilis Tsotras
<i>Department of Informatics Aristotle University of Thessaloniki Thessaloniki, Greece papadopo@csd.auth.gr</i>	<i>Department of Informatics Aristotle University of Thessaloniki Thessaloniki, Greece manolopo@csd.auth.gr</i>	<i>Department of Informatics University of Piraeus Piraeus, Greece ythead@unipi.gr</i>	<i>Department of Computer Science and Engineering University of California at Riverside Riverside, CA, USA tsotras@cs.ucr.edu</i>

SYNONYMS

DEFINITION [250 or fewer words defining the entry title]

The Grid File is a multidimensional indexing scheme capable to efficiently index database records in a symmetrical manner, i.e. by avoiding the distinction between primary and secondary keys. The structure is dynamic and adapts gracefully to its contents under insertions and deletions. A single record retrieval costs two disk accesses at most (upper bound), whereas range queries and partial match queries are also executed efficiently. The Grid File can be thought of as a generalization of dynamic hashing (e.g., extendible hashing) in multiple dimensions.

HISTORICAL BACKGROUND [500 or fewer words describing when/why the concept or technique developed]

Until the 80's there have been proposed many file structures for the processing of single attribute queries, i.e. queries on the primary key or any secondary key for which a corresponding index has been built. Multi-attribute queries are the ones where the user seeks objects that satisfy constraints (such as equality or range) on several attributes. Such queries can be executed by accessing all the corresponding indices (if they exist) and combine the partial results, or resort to sequential scanning.

To speed up the processing of multiple attribute queries a better solution is to create an index that leads the search directly to the objects of interest. Such an index can be designed if we envision a data record with k attributes as a point in a k -dimensional space. A multi-attribute range query would then be a hyper-rectangle in this k -dimensional space and the answer to it would be all points inside this rectangle. Access methods that can handle multi-dimensional points are called *Point Access Methods* (PAMs). Bentley in 1975 proposed such a basic PAM, which is called k -dimensional tree or k -d tree [2]. The Grid File is yet another structure designed to handle similar cases, proposed by Nievergelt, Hinterberger and Sevcik in 1984 [9]. Since then, several variations have been proposed in the literature in an effort to optimize its space and time performance behavior.

SCIENTIFIC FUNDAMENTALS [illustration and elaboration of the entry title definition, and outline the key points]

The Grid File can be viewed as an access method comprising of two separate parts: (a) the *directory*, and (b) the *linear scales*. To conceive this, assume that we want to index an *Employee* file using two attributes, say *salary* and *dept* (extension to more dimensions is straightforward). The Grid File imposes a grid on the two-dimensional attribute space. Each cell in this grid corresponds to one data page. The data points that “fall” inside a given cell are stored in the cell's corresponding page. Each cell must thus store a

pointer to its corresponding page. This information is stored in the Grid File's *directory*. However, cells that are empty do not use a page. Rather, two or more cells can share a page (i.e., point to the same page). The grid adapts to the data density by introducing more divisions in areas where there are more points.

The information of how each dimension is divided (and thus how data values are assigned to cells) is kept through *linear scales*. There is one linear scale per dimension (indexed attribute). Each linear scale is a one-dimensional array that divides the values on a particular dimension in such a way that records (points) are uniformly distributed across cells. An example of a Grid File on the "Dept" and "Salary" attributes appears in Figure 1. The dotted lines indicate cells that share a data page.

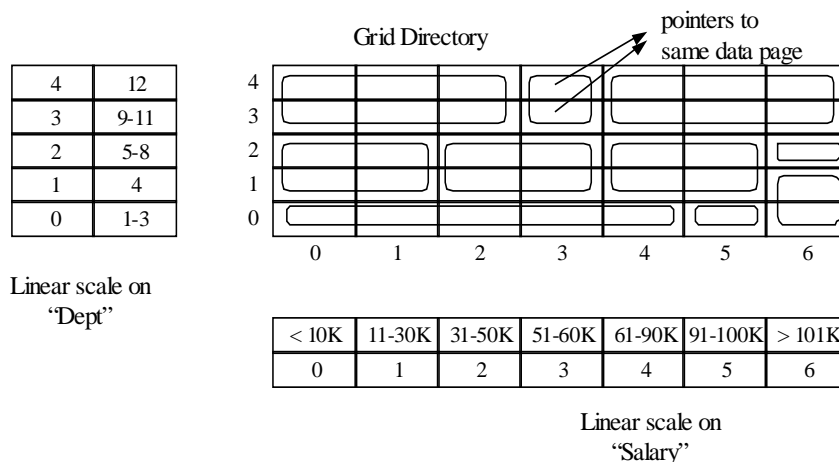


Figure 1. A Grid File.

Searching for a record with given attribute values involves two operations: (a) the Grid File's directory is searched to locate the cell that the record is hosted (b) the cell's pointer is followed to access the corresponding data page (say A) and (c) the record is searched only in data page A. If the record is found in A then the search terminates successfully, otherwise the search for the specific record is unsuccessful (i.e., the record does not exist). The Grid File can also address multi-dimensional range queries by selecting from each dimension's linear scale the appropriate cells. For example, such a query may ask for all employee records with the *salary* attribute ranging between 10K and 40K and the *dept* attribute ranging between 2 and 6. Again, the first step examines the directory and determines the cells that are intersected by the query range in both attributes, then the corresponding pointers to data pages are collected and finally the data pages are examined for relevant records. Apparently, the accessed cells may also contain some records outside the query range. These records are evidently eliminated from further consideration and they are not returned as part of the query answer.

Inserting a new record in this method is straightforward. First, the two linear scales are searched so as to map the record's *salary* and *dept* attribute values in each dimension. This mapping provides a cell in the directory. This cell is then accessed and using its pointer, the appropriate page, say A, for the new record is found. If this page has enough space to accommodate the new record the insertion process is complete. Otherwise, a new page B is allocated. If page A was pointed by more than one cells, the pointers of these cells are rearranged so as some will point to page A and some to page B (and the records of page A are redistributed accordingly between A and B). If page A was pointed by a single cell and overflows, a reorganization of the Grid File is needed. This reorganization will expand the directory and the scales by introducing a new column (or row) of cells.

In the sequel, we illustrate the insertion process by an example given in Figure 2. White dots correspond to existing records, whereas black dots are used to indicate new records being inserted to the Grid File. We assume that each data page can host at most three records. Practically, this number is larger in real

applications and depends on the size of the data page and the number of attributes. We assume that initially the Grid File is empty (does not contain any records). The first three records can be easily accommodated in the single data page *A* pointed by the single cell of the directory (corresponding to the whole data space), as it is illustrated in Figure 2(a). The next inserted record is *d*. However, the new record can not be hosted by data page *A* because its capacity is exceeded. Therefore, another data page *B* is allocated and records are distributed to the two data pages as it is shown in Figure 2(b). The next two insertions for records *e* and *f* do not cause any reorganization since the new records can be easily accommodated in the corresponding data pages pointed by the cells. This case is illustrated in Figure 2(c). Finally, the insertion of record *g* causes an overflow in data page *A*. The corresponding cell is split again using the other attribute and one more data page is allocated and records are distributed accordingly. The final shape of the Grid File is given in Figure 2(d).




Figure 2. Insertions in the Grid File.

Deletions are also supported, but they are handled differently. Initially, the deleted record is located using the directory and the corresponding data page is determined. If the record is found it is deleted from the data page. Instead of overflowing data pages deletions may cause the underutilization effect which means that several data pages may contain too few records. Therefore, appropriate merging operations are required to maintain the storage utilization of the Grid File at an acceptable level. For a detailed description of the methods used for merging as well as for splitting the reader is directed to reference [8].

The Grid File has a set of nice properties: (1) it is based on simple mechanisms for insertion, deletion and search, (2) it guarantees only two disk accesses for exact match queries (one for the directory and one for the data page), and, (3) it treats all indexed attributes symmetrically which leads to simple directory management policies. However, it has a set of serious disadvantages such as: (1) it introduces a space overhead for the directory, which can be large for high-dimensional spaces, (2) it has an extra update overhead, since a reorganization affects many cells and not only the cell with the overflowing page, and, (3) it suffers from performance degradation if the attributes are correlated, since the uniform scheme for performing splits is not adequate to guarantee performance efficiency.

Towards improving the behavior of the Grid File several research efforts have been performed. We briefly highlight some of them in the following lines. One of the first variations, the BANG File, has been proposed by Freeston [4]. The BANG File is based on a self-balanced tree-based directory which better reflects changes of the data distribution. To achieve better storage utilization the Twin Grid File has been proposed by Hulflesz, Six and Widmayer in [6]. The new scheme is as efficient as the original Grid File during range query processing but shows significant improvements regarding storage utilization. Blanken et al proposed the Generalized Grid File [3] which offers fast access for single attribute queries. The Multilevel Grid File [13] is another research effort to improve the performance of the original structure for exact-match, partial-match and range queries. This new scheme uses multiple grid levels and succeeds in better directory management and more efficient query processing than the original structure.

In addition to the variations proposed in the literature, there are efforts to use the Grid File in a parallel environment, towards more efficient data management. In [14] the authors study the problem of partitioning a Grid File to multiple disk devices towards more efficient search. When a data page split is

performed, the new data page is carefully allocated to a disk. Since disks can be accessed in parallel, several data pages can be read simultaneously during range query processing, offering significant performance improvements in comparison to a single-disk system. More complex queries on Grid Files, like spatial joins, have been also parallelized [7] towards reduced query response times. A different approach has been followed by [8]. The authors have proposed a method to load a Grid File in parallel. The data file is initially partitioned to the available processors using dynamic programming and sampling, and then each processor builds its own part of the Grid File.

KEY APPLICATIONS [\[current and potential users\]](#)

Spatial Databases

In Spatial Databases it is commonly required to join spatial data sets or perform nearest neighbor searches. Several algorithms have been proposed for such operations by adopting the Grid File as the underlying access method [1].

Data Mining

The Grid File can be, also, used for clustering data sets to identify correlation characteristics of the underlying value space. This stems from its ability to group patterns into blocks and cluster them with respect to the blocks by a topological neighbor search algorithm [12].

Data Warehouses

The Grid File can be used for efficient data cube storage in warehouses [11].

FUTURE DIRECTIONS*

The Grid File has eventually come up as a popular theoretical access method. However, although it has been widely honored in theory, in practice it has not been used by the database industry. Also, the fact that an entry for the Grid File exist in Wikipedia [5] shows that the subject is rather stabilized.

EXPERIMENTAL RESULTS *

A detailed performance evaluation of the Grid File can be found in [9], where the authors offer a detailed experimental section studying the properties of the structure regarding capacity of data pages, directory size and evaluation of splitting and merging policies. Moreover, interesting experimental results can be found in [3, 6] which compare the original Grid File with the corresponding variation proposed in each work.

CROSS REFERENCES [\[Other topics in the Encyclopedia which may be of interest to the reader of this entry. It is encouraged to redirect the readers to an overview entry in the subject area\]](#)

EXTENDIBLE HASHING

K-D TREES

MULTIDIMENSIONAL INDEXING

RANGE QUERY

SPATIAL JOIN

RECOMMENDED READING* [Between 5 and 15 citations to important literature, e.g., in journals, conference proceedings, and websites]

- [1] Becker L., Hinrichs K., Finke U. (1993): A New Algorithm for Computing Joins with Grid Files. ICDE 1993: 190-197.
- [2] Bentley, J.L. (1975): Multidimensional Binary Search Trees Used for Associative Searching. *Communications of the ACM*, 18(9):509-517.
- [3] Blanken H.M., Ijbema A., Meek P., van den Akker B. (1990): The Generalized Grid File: Description and Performance Aspects. ICDE Conference 1990: 380-388.
- [4] Freeston M. (1987): The BANG File: A New Kind of Grid File. SIGMOD Conference 1987: 260-269.
- [5] Grid File. Lemma at Wikipedia. http://en.wikipedia.org/wiki/Grid_file.
- [6] Hutflesz A., Six H.-W., Widmayer P. (1988): Twin Grid Files: Space Optimizing Access Schemes. SIGMOD Conference 1988: 183-190.
- [7] Kim J.-D, Hong B.-H. (1999): Parallel Spatial Join Algorithms using Grid Files. DANTE Conference 1999: 226-234.
- [8] Li J., Rotem D., Srivastava J. (1993): Algorithms for Loading Parallel Grid Files. SIGMOD Conference 1993: 347-356.
- [9] Nievergelt, J., Hinterberger, H., Sevcik K.K. (1984): The Grid File: an Adaptable, Symmetric Multikey File Structure, *ACM Transactions on Database Systems*, 9(1):38-71.
- [10] Lim, Y., Kim, M. (2004): A Bitmap Index for Multidimensional Data Cubes. DEXA Conference 2004: 349-358.
- [11] Luo, C., Hou, W.C., Wang, C.F., Wang, H., Yu X. (2006): Grid File for Efficient Data Cube Storage. *Computers and their Applications*, pp.424-429.
- [12] Schikuta, E., Erhart, M. (1997): The BANG-Clustering System: Grid-Based Data Analysis, *IDA*, p.513, 1997.
- [13] Whang K.-Y., Krishnamurthy R. (1991): The Multilevel Grid File - A Dynamic Hierarchical Multidimensional File Structure. DASFAA Conference 1991: 449-459.
- [14] Zhou Y., Shekhar S., Coyle M. (1994): Disk Allocation Methods for Parallelizing Grid Files. ICDE Conference 1994: 243-252.