



SP-GiST: An Extensible Database Index for Supporting Space Partitioning Trees

WALID G. AREF

IHAB F. ILYAS

Department of Computer Sciences, Purdue University, West Lafayette, IN 47907-1398, USA

aref@cs.purdue.edu

ilyas@cs.purdue.edu

Abstract. Emerging database applications require the use of new indexing structures beyond B-trees and R-trees. Examples are the k-D tree, the trie, the quadtree, and their variants. They are often proposed as supporting structures in data mining, GIS, and CAD/CAM applications. A common feature of all these indexes is that they recursively divide the space into partitions. A new extensible index structure, termed *SP-GiST* is presented that supports this class of data structures, mainly the class of space partitioning unbalanced trees. Simple method implementations are provided that demonstrate how *SP-GiST* can behave as a k-D tree, a trie, a quadtree, or any of their variants. Issues related to clustering tree nodes into pages as well as concurrency control for *SP-GiST* are addressed. A dynamic minimum-height clustering technique is applied to minimize disk accesses and to make using such trees in database systems possible and efficient. A prototype implementation of *SP-GiST* is presented as well as performance studies of the various *SP-GiST*'s tuning parameters.

Keywords: space-partitioning trees, spatial databases, extensible index, generalized search trees, clustering

1. Introduction

Emerging database applications require the use of new indexing structures beyond B+-trees. The new applications may need different index structures to suit the big variety of data being supported, e.g., video, image, and multidimensional data. Typical applications are cartography, CAD, GIS, telemedicine, and multimedia applications. For example, the quadtree (Finkel and Bentley, 1974; Klinger, 1971) is used in the Sloan Digital Sky Survey to build indexes for different views of the sky (a multi-terabyte database archive) (Szalay et al., 2000), the linear quadtree (Gargantini, 1982) is used in the recently released Oracle spatial product (Oracle, 1999), the trie data structure is used in (Aref et al., 1995) to index handwritten databases, and the pyramid multi-resolution data structure (Tanimoto and Pavlidis, 1975) is used in the Microsoft TerraServer (Barclay et al., 2000) which is an online atlas, currently being developed that combines around eight terabytes of image data. The reader is referred to (Berchtold et al., 1998; Oracle, 1999; Esperanca and Samet, 1996; Faloutsos and Gaede, 1996; Faloutsos et al., 1997; Gaede and Gunther, 1998; Güting, 1994; Orenstein and Manola, 1988; Samet, 1990a; Seeger and Kriegel, 1990; Sellis et al., 1997) for additional database applications that use a variety of spatial and non-traditional tree structures.

The support of the National Science Foundation under Grant IIS-0093116 is gratefully acknowledged.

Having a single framework to cover a wide range of these tree structures, although is very attractive from the point of view of database system implementation, is hindered by two main problems. The first problem is *the storage/structure characteristics* of spatial trees. Most of the *unbalanced* spatial tree structures are not optimized for I/O, which is a crucial issue for database systems. Quadrees, tries, and k-D trees can be so *skinny and long*. Unless the problem of appropriately clustering the tree nodes into pages is addressed properly, this would lead to many I/O accesses before getting the required query answer. Compare this to the B+-tree, that in most cases has a height of 2–3 levels, and to the R-tree (Guttman, 1984) and its variants, the R*-tree (Beckmann et al., 1990) and the R+-tree (Sellis et al., 1987) that play an important role as spatial database indexes, e.g., see Brinkhoff et al. (1996), DeWitt et al. (1994), and Papadias et al. (1998). The second problem is the *implementation effort* of building indexes. Hard wiring the implementation of a full fledged index structure with the appropriate concurrency and recovery mechanisms into the database engine is a non-trivial process. Repeating this process for each spatial tree that can be more appealing for a certain application requires major changes in the DBMS core code. After all, one may still need a new structure that will cause, rewriting/augmenting significant portions of the DBMS engine to add the new tree index. The *Generalized Search Tree (GiST)* (Hellerstein et al., 1995), was introduced in order to provide single implementation for B-tree-like indexes, e.g., the B+-tree (Knuth, 1973), the R-tree (Guttman, 1984), and the RD-tree (Hellerstein and Pfeffer, 1994). Although practically useful, the class of unbalanced spatial indexes, e.g., the quadtree, the trie, and the k-D tree, is not supported by GiST because of the structure characteristics mentioned.

One important common feature of the quadtree, the trie, and the k-D tree family of indexes is that at each level of the tree, the underlying space gets partitioned into disjoint partitions. For example, in the case of a two-dimensional quadtree, at each level of decomposition, the space covered by a node is decomposed into four disjoint blocks. Similarly, in the case of the trie (assuming that we store a dictionary of words), the space covered by a node in the trie gets decomposed into 26 disjoint regions (each region corresponds to one letter of the alphabet). The k-D tree exhibits similar behavior. We use the term *space-partitioning* trees to represent the class of hierarchical data structures that decomposes a certain space into disjoint partitions. The number of partitions and the way the space is decomposed differ from one tree to the other.

In this paper we study the common features among the members of the spatial space partitioning trees aiming at developing a framework that is capable of representing the different tree structures and overcoming the difficulties that prevent such useful trees from being used in database engines. The DBMS will then be able to provide a large number of index structures with simple method plug-ins. As demonstrated in the paper, for the framework of space partitioning trees, we furnish in the DBMS (only once) the common functionalities such as the insertion, deletion, and updating algorithms, concurrency control and recovery techniques and I/O access optimization. For example, in a multimedia or a data mining application, we may then freely choose the best way to index each feature depending on the application semantics. By writing the right extensions to the extensible single implementation, a quadtree, a trie, a k-D tree, or other spatial structures can be made available without messing with the DBMS internal code.

The rest of the paper is organized as follows. Section 2 presents the class of space-partitioning trees. In Section 3, the SP-GiST framework is presented. Section 3 also includes a description of SP-GiST external user interface, and illustrates the realization of various tree structures using it. This includes a realization of the k-D tree, the trie, the Patricia trie, and several variants of the quadtree. Section 4 gives the implementation of the internal methods of SP-GiST. Concurrency control and recovery for SP-GiST are discussed in Section 5. Node clustering in SP-GiST is presented in Section 6. The pseudo code of the clustering algorithm is given in Appendix A. Implementation and experimental results for the various tuning parameters of SP-GiST are given in Section 7. Section 8 contains some concluding remarks.

2. The class of space partitioning trees

The term *space-partitioning* tree refers to the class of hierarchical data structures that recursively decomposes a certain space into disjoint partitions. It is important to point out the difference between data-driven and space-driven decompositions of space. If the principle of decomposing the space is dependent on the input data, it is called *data-driven* decomposition, while if it is dependent solely on the space, it is called *space-driven* decomposition. Examples of the first category are the k-D tree (Bentley, 1975) and the point quadtree (Klinger, 1971). Examples of the second category are the trie index (de la Briandais, 1959; Fredkin, 1960), the fixed grid (Nievergelt et al., 1984), the universal B-tree (Bayer, 1997), the region quadtree (Finkel and Bentley, 1974), and other quadtree variants (e.g., the MX-CIF quadtree (Kedem, 1982), the bintree, the *PM quadtree* (Samet and Webber, 1985), the PR quadtree (Orenstein, 1982) and the PMR quadtree (Nelson and Samet, 1986)).

There are common underlying features among these spatial data structures. The term *quadtrie* was introduced in (Samet, 1990a) to reflect the structure similarity between the trie and the quadtree. Similarly, the k-D tree and the MX quadtree have many structural similarities, e.g., both structures recursively partition the space into a number of disjoint partitions. On the other hand, the two trees differ in the number of partitions to divide the space and also in the decomposition principle. The decomposition is data-driven in the case of the k-D tree, while it is space-driven in the case of the MX quadtree.

The structural and behavioral similarities among many spatial trees create the class of space-partitioning trees. In contrast, the differences among these trees enable their use in a variety of emerging applications. The nature of spatial data that the application is dealing with, as well as the types of queries that need to be supported, aid in deciding which space-partitioning tree to use.

Space-partitioning trees can be differentiated on the following basis:

• Structural differences

- SD_1 : The type of data they represent.
- SD_2 : The decomposition fan-out (the number of partitions).
- SD_3 : The resolution of the underlying space.
- SD_4 : Allowing single-arc nodes.
- SD_5 : The use of buckets.

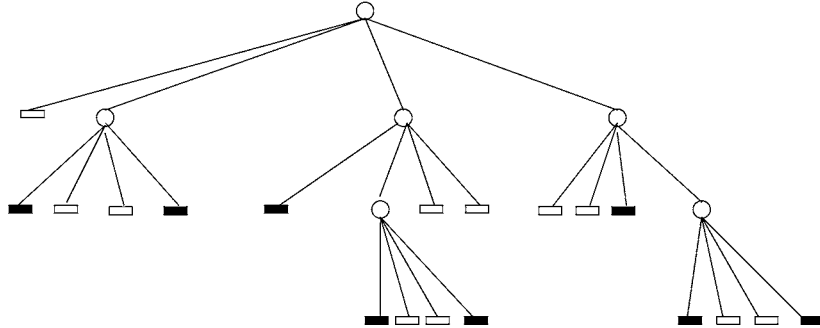


Figure 1. An example PR quadtree.

• Behavioral differences

- BD_1 : The decomposition principle (data or space driven partitioning).

The structural differences or design options can be viewed as *Shape Parameters* for the realized tree. For example, in the realization of the PR quadtree, or more precisely the PR-*quadtree*, the represented data is “point” (SD_1). The decomposition depends on the space not on the data inserted (compare to the k-D tree) (BD_1). Each time a partitioning of the space quadrant into four equal quadrants (SD_2 and SD_4) takes place to divide the quadrant that has two points so that each point is attached with one quadrant. The decomposition resolution is “variable” in the sense that the partitioning stops whenever one data point resides in the quadrant (SD_3). Figure 1 shows an example of the PR quadtree. At the leaf level, nodes can be “white” (i.e., contains no data) or “black” (i.e., contains one data point (SD_5)).

Using the same analogy, we can analyze the structure and behavior of the trie. The data represented in a trie is of type “word” (SD_1). The decomposition of the trie is space-dependent (BD_1), as we always decompose the space into 26 partitions (SD_2); one partition for each letter of the alphabet. In one variant of the trie, the resolution is “not variable” (SD_3) as we need to decompose the space until we consume all the letters of the inserted word (refer to figure 2(a) for illustration). This is in contrast to stopping the decomposition only when a space partition uniquely identifies the inserted word (see figure 2(b)). The same analysis can be applied to realize other quadtree and trie variants, the k-D tree, and the bin-tree.

In the following sections, we will introduce a general framework, termed *SP-GiST*, which we can use to implement a big collection of space-partitioning trees. *SP-GiST* has one core implementation as well as user plug-ins that reflect the required structural and behavioral characteristics. The existence of such a framework will facilitate the adaption of this class of space-partitioning trees into database engines.

3. SP-GiST framework interface

SP-GiST is a general index framework that covers a wide range of tree indexes representing the large class of *space-partitioning search trees* represented in Section 2.

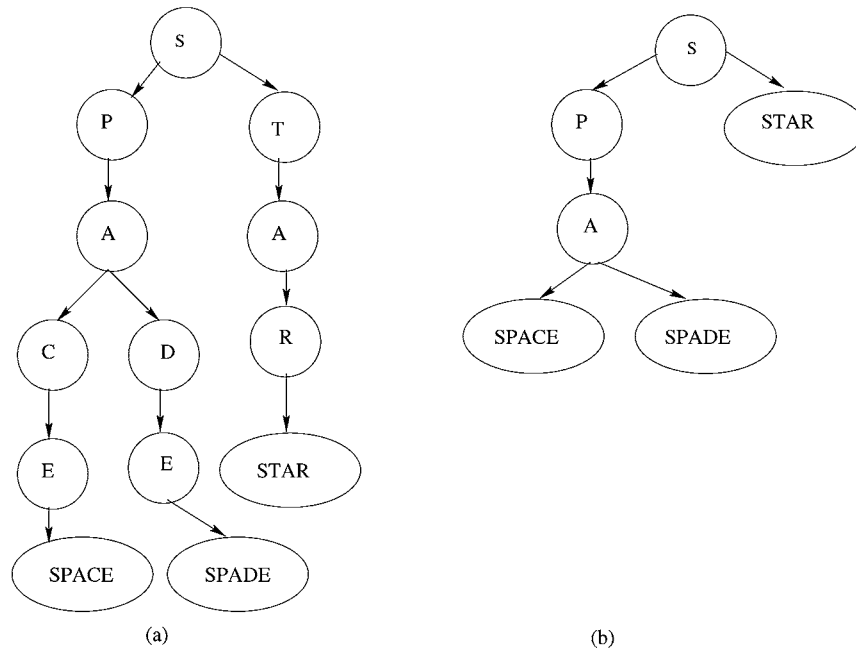


Figure 2. Two variants of the trie data structure : (a) Resolution is not variable (b) resolution is variable.

The *structural* characteristics of space-partitioning trees that distinguish them from other tree classes are: (1) Space-partitioning trees decompose the space recursively, each time, a fixed number of *disjoint* partitions is produced. (2) Space-partitioning trees are unbalanced trees (3) Space-partitioning trees suffer from limited fan-out, e.g., the quadtree has only a fan-out of four. So, space-partitioning trees can be *skinny* and *long*. (4) Two different types of nodes exist in a space-partitioning tree, namely, index nodes (internal nodes) and data nodes (leaf nodes). The framework reflects these facts by having two main parts; the *internal* tree methods that reflect the *similarities* among all members of the class of space-partitioning trees, and the *external* interface that enables us to identify the features specific to a particular tree reflecting the *differences* listed in Section 2.

By specifying user access methods as in GiST (Hellerstein et al., 1995), SP-GiST has interface parameters and methods that allow it to represent the class of space-partitioning tree indexes and reflect the structural and behavioral differences among them.

3.1. Interface parameters

The user can realize a particular space-partitioning tree using the following interface parameters:

- *Leaf Shrink*: Data is inserted at the first available leaf node. Decomposition will not depend on the maximum possible resolution. In this strategy, no index node will have one leaf node as we decompose only when there is no room for the newly inserted data item.
- *Tree Shrink*: The internal nodes are merged together to eliminate all single child internal nodes. This strategy is adapted from structures like the *Patricia* trie that aim at reducing the height of the tree as much as possible.

For example, in the case of *PathShrink* = “Never Shrink”, when storing the word “implementation” in the trie, the word will be stored in a leaf after a 14-nodes path, one level per input character. On the other hand, in the case of *PathShrink* = “Leaf Shrink”, the input word may be stored in a leaf after the three-node path “i”, “m”, “p”, and “lementation”, since based on the current words in the trie, splitting up to the letter “p” makes a unique leaf entry for the word “implementation”. Finally, in the case of *PathShrink* = “Tree Shrink”, the input word may be stored in a leaf after a three-node path “i”, “mp”, l, “ementation”. Since the only child of the index node “m” is the index node “p”, both nodes are merged together to reduce path length.

PathShrink is the way the framework uses to map the structural differences SD_3 and SD_4 . As shown, many variants of spatial tree can be realized according to these structural differences.

- *NodeShrink*: This parameter determines if empty partitions should be kept in the tree or not. When *NodeShrink* is set to *true*, the resulting space partitioning tree is best described as a *forest*, because only the partitions that have subtrees are represented. When *NodeShrink* is set to *false* all the partitions are kept. In case of shrinking index nodes, the nodes become variable-length nodes. Figure 4 shows the effect of the parameter *NodeShrink* on the trie where the words in the leaf nodes are the only data inserted. Note that *NodeShrink* has no effect on the path length of the tree (as opposed to *PathShrink*).
- *BucketSize*: This parameter gives the maximum number of data items a data node can hold. It also represents the *Split Threshold* for data nodes. For example, quadtrees have the notion of a bucket size that determines when to split a node (e.g., as in the PMR quadtree (Nelson and Samet, 1986)). The use of buckets (SD_5) is an attractive design option for many database applications where we are concerned about storing multiple data items per bucket for storage performance efficiency.

3.2. External methods (behavior)

The external methods are the second part of the SP-GiST interface that allows the user to specify the behavior of each tree. The main purpose is to map the behavioral difference BD_1 in Section 2. Note the similarity between the names of the first two methods and the ones introduced in the GiST framework (Hellerstein et al., 1995) although they are different in their functionalities.

Let $E : (p, ptr)$ be an entry in an SP-GiST node, where p is a node predicate or a leaf data key and ptr is a pointer. When p is a node predicate, ptr points to the child node corresponding to its predicate. When p is a leaf data key, ptr points to the data record associated with this key.

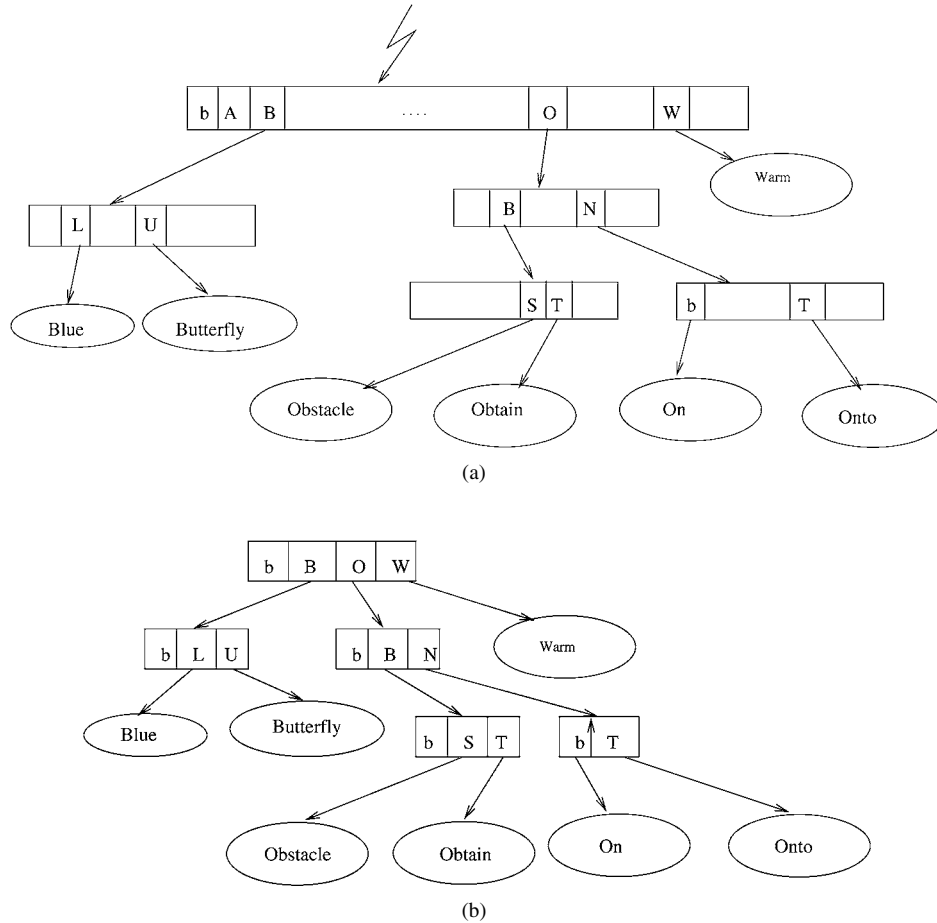


Figure 4. The effect of the parameter *NodeShrink*: (a) *NodeShrink* = *false*, and (b) *NodeShrink* = *true*.

- *Consistent(Entry E, Query Predicate q, level)*: A Boolean function that is false when $(E.p \wedge q)$ is guaranteed unsatisfiable, and is true otherwise. This method will be used by the tree search method as a navigation guide through the tree. Argument *level* is used in order to determine consistency depending on the current decomposition level. For example, in a quadtree, a query of a data point (x, y) is consistent only with the entry that points to the quadrant containing this point.
- *PickSplit(P, level, splitnodes, splitpredicates)*: Returns Boolean, where *P* is a set of *BucketSize*+1 entries that cannot fit in a node. *PickSplit* defines a way of splitting the entries into a number of partitions equal to *NumberOfSpacePartitions* and returns a Boolean value indicating whether further partitioning should take place or not. The parameter *level* is used in the splitting criterion because splitting will depend on the current decomposition level of the tree. For example, in a trie of English words, at level

i , splitting will be according to the i th character of each word in the over-full node. `PickSplit` will return the entries of the split nodes in the output parameter `splitnodes`, which is an array of buckets, where each bucket contains the elements that should be inserted in the corresponding child node. The predicates of the children are also returned in `splitpredicates`.

- `Cluster()`: This method defines how tree nodes are clustered into disk pages. The method is explained in more detail in Section 6.

The interface methods realize the *behavioral* design options listed in Section 2. Methods `Consistent` and `PickSplit` determine if the tree follows the space-driven or the data-driven partitioning. For example, in a k-D tree, which is a data-driven space partitioning tree, method `Consistent` compares the coordinates of the query point (the point to be inserted or searched for) against the coordinates of the point attached to the index node. The values of these coordinates are determined based on data that is inserted earlier into the k-D tree. On the other hand, method `Consistent` for a space-driven space partitioning tree, e.g., the trie, will only depend on the letters of the newly inserted word. The comparison is performed against the letter associated with the index node entry, which is space-dependent, and is independent of the previously inserted data.

We can also show that method `PickSplit` completes the specification of the behavioral design option by specifying the way to distribute node entries among the produced partitions. Examples of `PickSplit` for various tree structures are given in the following section.

3.3. Realization of space-partitioning trees

Using the SP-GiST interface, given in the previous sections, we demonstrate how to realize some commonly used space-partitioning indexes. More specifically, we present the realization of the k-D tree, variants of the quadtree, the trie, and the Patricia trie.

The k-D tree: k-D trees (Bentley, 1975) are a special kind of search trees, useful for answering range queries about a set of points in the k-dimensional space. The k-D tree uses a data-driven decomposition of the space (see Section 2). The tree is constructed by recursively *partitioning* the space into two sub-spaces with respect to one of the dimensions at each tree level.

The k-D tree insertion algorithm for the two-dimensional case (i.e., $k = 2$) with points in the xy plane is as follows: The algorithm selects any point and draws a line through it, parallel to the y -axis. This line partitions the plane vertically into two sub-planes. Another point is selected and is used to horizontally partition the sub-plane in which it lies. In general, a point that falls in a region created by a horizontal partition will divide this region vertically, and vice versa. This division process induces a binary tree structure, (e.g., see figure 5).

The realization of the k-D tree is given in Table 1. `PathShrink` is set to “Leaf Shrink” because we put each input point at the first available place depending on the previously inserted points. Each node will hold only one point, (`BucketSize = 1`). `NodeShrink` is set to `false`, so each index node will have a slot for the left subtree and a slot for the right subtree.

Table 1. Realization of the k-D Tree using SP-GiST.

Parameters	PathShrink = Leaf Shrink NodeShrink = False BucketSize = 1 NumberOfSpacePartitions = 2 Node Predicate = “left”, “right”, or blank. Key Type = Point
Consistent(E, q, level)	IF (level is odd AND q.x satisfies E.p.x) OR (level is even AND q.y satisfies E.p.y) RETURN TRUE ELSE RETURN FALSE
PickSplit(P, level)	Put the old point in a child node with predicate “blank” Put the new point in a child node with predicate “left” or “right”, RETURN FALSE

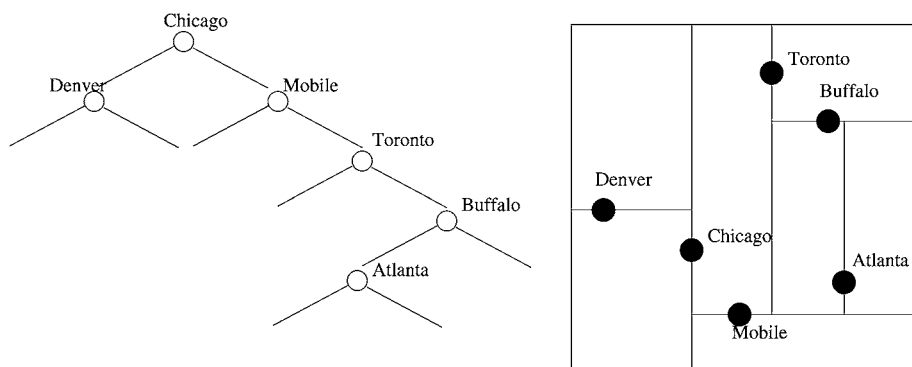


Figure 5. An example k-D tree.

We have only two space partitions for the “right” and “left” to a point ($NumberOfSpacePartitions = 2$).

The Quadtree: The term *quadtree* describes a class of hierarchical data structures whose common property is the recursive decomposition of space into quadrants. The quadtree can be realized by SP-GiST. In the next subsections, examples of various types of quadtrees are presented for point data, rectangles, and polygonal data. Note that for all the variants, the number of space partitions is equal to four ($NumberOfSpacePartitions = 4$), with a bucket size of B items ($BucketSize = B$). *NodeShrink* is set to *false*, so each index node will have a slot for each partition even if it is empty. Setting *NodeShrink* to *true* would realize a quadtree with all *white* nodes eliminated (see figure 6) (Samet, 1990a).

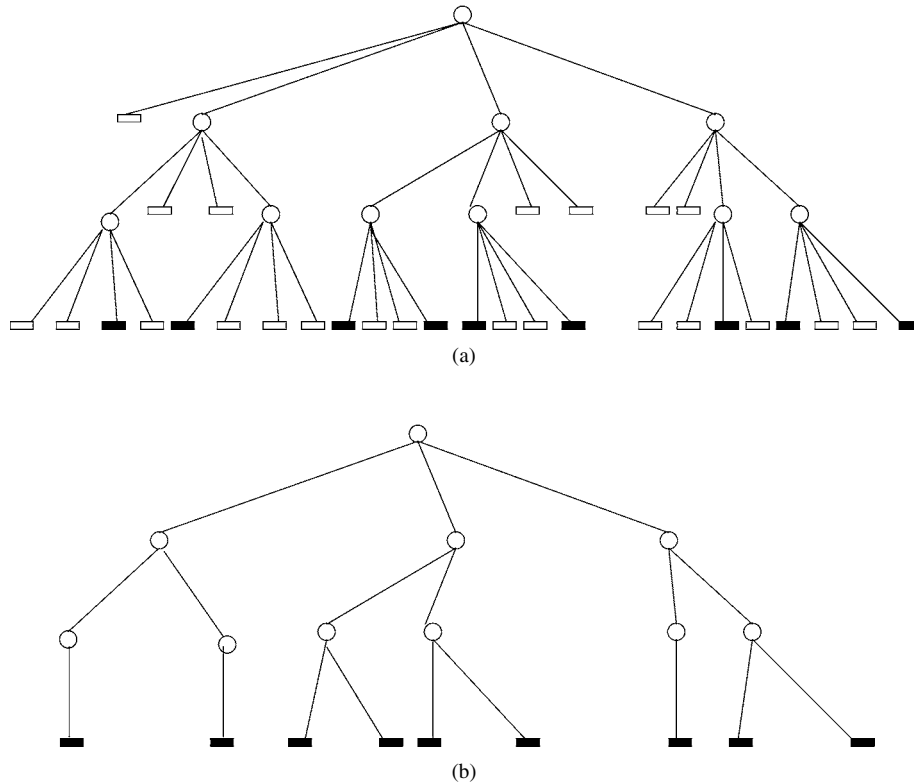


Figure 6. An example MX quadtree: (a) *NodeShrink = false*, and (b) *NodeShrink = true*.

The quadtree can be viewed as a trie structure in two dimensions—with only two possible characters in each dimension, in *trie* terminology, or even a one dimensional trie with only a four character alphabet set. Thus in the literature, *space-driven* quadtrees are often called *quadtries* (Samet, 1990b).

When we treat data points as nonzero elements in a square matrix, the resulting data structure is called the *MX quadtree* (MX for matrix). In the MX quadtree, leaf nodes are black or empty (white) corresponding to the presence or absence, respectively, of data points in the appropriate position in the matrix. Each point in an MX quadtree corresponds to a 1×1 square. Figure 6 gives an example of an MX quadtree. Notice that data nodes of the MX quadtree all appear at the same level. The number of space decompositions is predefined depending on the desired space resolution.

For the MX quadtree, realized in Table 2, *PathShrink* is set to “Never Shrink”. Therefore, the tree is expanded through successive splitting to the maximum space resolution. Thus, *PickSplit* will not be invoked as each point will fall in one node.

The MX quadtree is applicable as long as the domain of data points is discrete and finite. If this is not the case, the data points cannot be represented using an MX quadtree since the minimum separation between the data points will be unknown. This leads to

Table 2. Realization of the MX quadtree using SP-GiST.

Parameters	PathShrink = Never Shrink NodeShrink = False BucketSize = B NumberOfSpacePartitions = 4 Node Predicate = Quadrant represented by (x1, y1, x2, y2) where (x1, y1) are the values of the coordinates of the top left corner and (x2, y2) are the values of the coordinates of the bottom right corner Key Type = Point
Consistent(E, q, level)	IF (q coordinates inside E.quadrant) RETURN TRUE ELSE RETURN FALSE
PickSplit(P, level)	RETURN FALSE

Table 3. Realization of the PR quadtree using SP-GiST.

Parameters	PathShrink = Leaf Shrink NodeShrink = False BucketSize = B NumberOfSpacePartitions = 4 Node Predicate = Quadrant represented by (x1, y1, x2, y2) where (x1, y1) are the values of the coordinates of the top left corner and (x2, y2) are the values of the coordinates of the bottom right corner Key Type = Point
Consistent(E, q, level)	IF (q coordinates inside E.quadrant) RETURN TRUE ELSE RETURN FALSE
PickSplit(P, level)	Partition and allocate data points into quadrants according to the locations of the data points IF any partition is still over-full RETURN TRUE ELSE RETURN FALSE

the idea of associating data points with quadrants and hence realizing the *PR quadtree* (P for point and R for region) (Orenstein, 1982). Now, each data point maps to a quadrant and not to a 1×1 square as in the *MX quadtree*. Figure 1 gives an example of the *PR quadtree*.

The *PR quadtree* can be realized using SP-GiST by setting *PathShrink* to “Leaf Shrink” as we put each input point at the first available leaf node. The leaf node is not necessarily of size 1×1 . Realization of the *PR quadtree* using SP-GiST is given in Table 3.

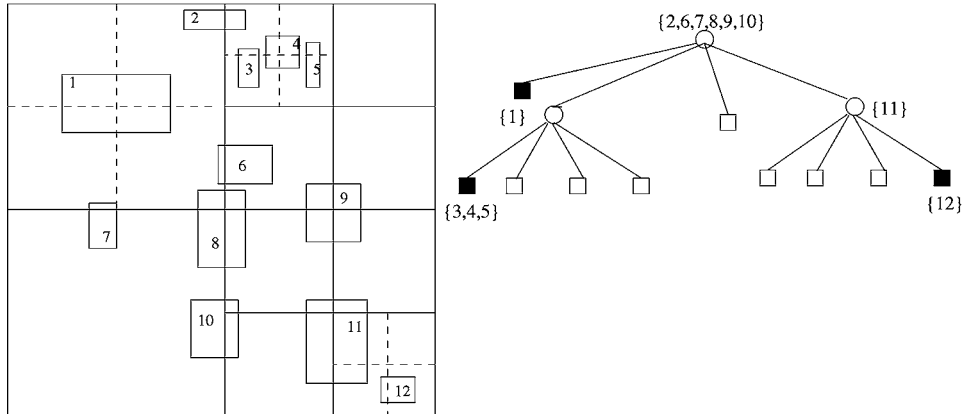


Figure 7. An example MX-CIF quadtree.

The *MX-CIF quadtree* is a quadtree variation for storing rectangles. It associates each rectangle, say R , with the quadtree node corresponding to the smallest block that contains R in its entry. Rectangles can be associated with both leaf and non-leaf nodes. The subdivision ceases whenever a node's block contains no rectangles. Figure 7 gives an example MX-CIF quadtree. Notice that more than one rectangle can be associated with a given node.

The MX-CIF quadtree can be realized by SP-GiST, as given in Table 4. PickSplit is not applicable here, because according to the MX-CIF insertion algorithm, there is not much choice as to where a rectangle gets inserted.

Table 4. Realization of the MX-CIF quadtree using SP-GiST.

Parameters	PathShrink = Leaf Shrink NodeShrink = False BucketSize = B NumberOfSpacePartitions = 4 Node Predicate = Quadrant represented by $(x1, y1, x2, y2)$ where $(x1, y1)$ are the values of the coordinates of the top left corner and $(x2, y2)$ are the values of the coordinates of the bottom right corner Key Type = Rectangle
Consistent (E, q, level)	IF (Node predicate is the minimum bounding quadrant of q AND the E.p is Blank) RETURN TRUE IF (E.p contains q) RETURN TRUE ELSE RETURN FALSE
PickSplit(P, level)	RETURN FALSE

Table 5. Realization of the PMR quadtree using SP-GiST.

Parameters	PathShrink = Leaf Shrink NodeShrink = False BucketSize = B NumberOfSpacePartitions = 4 Node Predicate = Quadrant represented by (x1, y1, x2, y2) where (x1, y1) are the values of the coordinates of the top left corner and (x2, y2) are the values of the coordinates of the bottom right corner Key Type = Line Segment represented by end points
Consistent (E, q, level)	IF (inserted line intersects E.quadrant) RETURN TRUE ELSE RETURN FALSE
PickSplit(P, level)	Partition the line segments according to their intersections with quadrants RETURN FALSE

Another quadtree variant is the *PMR quadtree* (Nelson and Samet, 1986) that is used to store polygonal maps. The key is of type line segment in the PMR quadtree, where line segments serve as the building block to construct polygons.

The PMR quadtree is an edge-based data structure. A line segment is stored in a PMR quadtree by inserting the line segment into the nodes corresponding to all the blocks that it intersects. If the bucket capacity is exceeded, the node's block is split *once, and only once*, into four equal quadrants. Thus, bucket capacity is really a splitting threshold. The PMR quadtree can be realized using SP-GiST, as given in Table 5.

The Trie: A *trie* (de la Briandais, 1959; Fredkin, 1960) is a tree in which the branching at any level is determined by only a portion of the key as in figure 4(a). The trie contains two types of nodes; index and data nodes. In the trie of figure 4(a), each index node contains 27 link fields. In the figure, index nodes are represented by rectangles, while data nodes are represented by ovals.

All characters in the key values are assumed to be one of the 26 letters of the alphabet. A blank is used to terminate a key value. At level 1, all key values are partitioned into 27 disjoint classes depending on their first character. Thus, LINK(T, i) points to a subtrie containing all key values beginning with the *i*th character (T is the root of the trie). On the *j*th level the branching is determined by the *j*th character. When a subtrie contains only one key value, it is replaced by a node of type data. This node contains the key value, together with other relevant information such as the address of the record with this key value, etc. The trie can be realized using SP-GiST, as given in Table 6. Notice that *PathShrink* is set to "Leaf Shrink" (refer to Section 3.1). *NodeShrink* is set to *false* in this realization of the trie. Another option is to set the *NodeShrink* to *true* to realize the forest trie (Knuth, 1973), as discussed in Section 3.1.

Table 6. Realization of the Trie using SP-GiST.

Parameters	PathShrink = Leaf Shrink NodeShrink = False BucketSize = B NumberOfSpacePartitions = 26 Node Predicate = letter or Blank Key Type = String
Consistent(E, q, level)	IF (q[level] == E.letter) OR (E.letter == BLANK AND level > length(q)) RETURN TRUE ELSE RETURN FALSE
PickSplit(P, level)	Partition the data strings in P according to the character values at position "level" IF any data string has length < level, insert data string in Partition "blank" IF any of the partitions is still over-full RETURN TRUE ELSE RETURN FALSE

The regular trie suffers from the problem of long skinny paths of single arc nodes. For example, for a trie with a bucket size of 2, inserting the three words "abate", "abacus", and "abort" will cause the node to split. Since we are at the first level, the split will depend on the first character in each word. Since all the words have "ab" as their first and second characters, splitting must continue until the third character, resulting in a skinny trie (see figure 8).

The *Patricia trie* (Morrison, 1968; Knuth, 1973) is a special trie structure that addresses this problem. It has the property that all nodes that have only one arc are merged with their

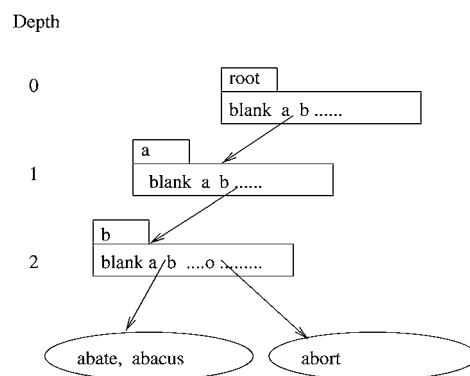


Figure 8. An example trie with *BucketSize* = 2: the final tree after inserting the three words.

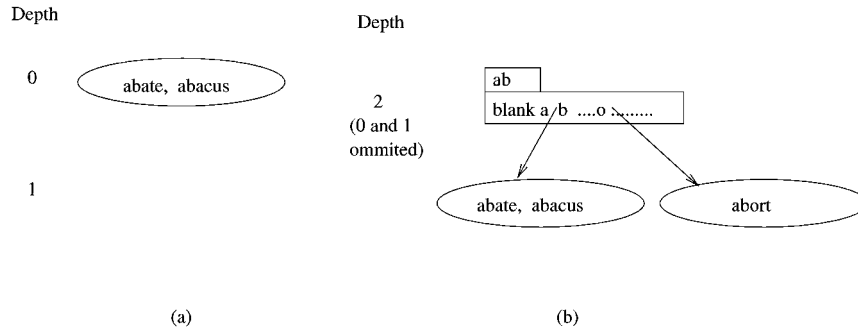


Figure 9. An example Patricia trie with *BucketSize* = 2: (a) after inserting the first two words, and (b) after inserting the third word.

parent nodes. To avoid false matches, each node in the Patricia trie must have either a count of the number of eliminated nodes or a pointer to the eliminated symbols. In the previous example, (refer to figure 9), the Patricia trie will split only once, thus eliminating the single arc nodes and storing the eliminated symbols (“ab”) in the parent node.

In the Patricia trie, *PathShrink* is set to “Tree Shrink”. When splitting a node, we search for the common prefix of all words. The common prefix is returned as the predicate of the parent node, while splitting is performed based on the next letter after that prefix. The realization of the Patricia trie using SP-GiST is given in Table 7.

4. SP-GiST internal methods

The methods for insertion, deletion, and search in SP-GiST are internal operations that are implemented inside the SP-GiST index engine. These methods are used in conjunction with the external methods to realize specific space-partitioning trees. The user of SP-GiST provides only the external methods, while the internal methods are hard coded into the SP-GiST index engine. The internal methods are general for the class of space-partitioning trees, and their behavior is tuned by making use of the user-defined external methods and parameters.

The internal methods are designed to accommodate for the space-partitioning, recursive decomposition, bucket sizes, insertion resolution, and node clustering (refer to the structural and behavioral characteristics of space-partitioning trees, given in Section 2).

Recall that unlike the GiST structure, SP-GiST has to support two distinct types of nodes; index and data nodes. **Index nodes (non-leaf nodes)** hold the various space partitions at each level. Each entry in an index node is a root of a subtree that holds all the entries that lie in this partition. The space partitions are disjoint. Besides having a slot for each space partition, the index node contains an extra *blank* slot to point to data nodes attached to the partition represented by this node. On the other hand, **data nodes (leaf nodes)** hold the key data and other pointer information to physical data records. We can think of data nodes as *Buckets* of data entries. Thus, a splitting strategy determined by *PickSplit* will be applied to split over-full data nodes.

Table 7. Realization of the Patricia Trie using SP-GiST.

Parameters	PathShrink = Tree Shrink NodeShrink = False BucketSize = B NumberOfSpacePartitions = 26 Node Predicate = letter or Blank Key Type = String
Consistent(E, q, level)	IF (q[level] == E.letter) OR (E.letter == BLANK AND level > length(q)) RETURN TRUE ELSE RETURN FALSE
PickSplit(P, level)	Find a Common prefix among words in P Update level = level+length of the common prefix Let P predicate = the common prefix Partition the data strings in P according to the character values at position "level" IF any data string has length < level, insert data string in Partition "blank" IF any of the partitions is still over full RETURN TRUE ELSE RETURN FALSE

The insert algorithm, given in Table 8, depends on the following interface parameters and external methods:

1. Parameter *PathShrink* specifies how deep we should proceed with the space decomposition.
2. Method *Consistent* specifies which branch to follow.
3. Method *PickSplit* to split over-full nodes. The return value of *PickSplit* tells us when we should stop the splitting process.

Method *Insert* begins by checking Parameter *PathShrink*. If *PathShrink* is set to "Never Shrink", method *Insert* performs a successive creation of index nodes to the maximum space resolution. If the parameter is set to "Leaf Shrink" or "Tree Shrink", the insertion algorithm searches for the first leaf node with a predicate that is *Consistent* with the key to be inserted. In the case of "Tree Shrink", some eliminated index nodes may be needed while locating the leaf. Hence, an internal split is performed to "expand" the eliminated index nodes. If the leaf node is over-full, then method *PickSplit* will be invoked continuously to distribute the entries among non over-full children or until it reaches the maximum resolution of the underlying space. Notice that method *Insert* invokes method *cluster* to dynamically re-cluster the nodes properly after insertion. Node clustering is further explained in Section 6.

Table 8. SP-GiST insertion algorithm.

```

1.  INSERT (TreeNode root, Key, level)
2.    CurrentNode = root /* Initially root is null */
3.    IF PathShrink is "Never Shrink" THEN
4.      LOOP WHILE level < SpaceResolution AND level < Key length
5.        IF node is NULL THEN E = Create a new node of type INDEX
6.        FOR each slot i in the index node LOOP
7.          IF (Consistent(E[i],key,level)) THEN index = i
8.          IF None is consistent /*due to NodeShrink*/
9.          THEN Create the missing index slot w.r.t level
10.         index = the position of the new slot
11.         CurrentNode = E[index].ptr /* the child pointed by entry E[index]*/
12.         level = level + 1
13.     IF CurrentNode is INDEX node /* pick a child to go */
14.       Compare the key with the CurrentNode predicate
15.       IF no match AND PathShrink is "Tree Shrink"
16.       THEN get the common prefix between the two
17.         Change CurrentNode predicate to the common prefix
18.         Create a new INDEX node with the rest of the old node predicate
19.         Let CurrentNode be the new index node
20.       FOR each slot i in the index node LOOP
21.         IF (Consistent(E[i], key, level)) THEN index = i
22.         IF None is consistent /*due to NodeShrink*/
23.         THEN Create the missing index slot w.r.t level
24.         index = the position of the new slot
25.         CurrentNode = CurrentNode[index].ptr
26.         INSERT (CurrentNode, key, level + 1) /* recursive */
27.     IF CurrentNode is full THEN /* DATA node and may need to be split*/
28.       LOOP WHILE PickSplit(node, level)
29.         n = Create new node of type INDEX
30.         Create Children for the split entries
31.         Parent(n) = Parent(CurrentNode)
32.         Adjust branches of 'n' to point to the new children
33.         level = level + 1
34.     ELSE insert the key in CurrentNode /* not a full node */
35.     Cluster() /* to recluster the tree nodes in pages */

```

Table 9. SP-GiST search algorithm.

1.	SEARCH (TreeNode root, Key, level)
2.	Found = false
3.	CurrentNode = root /* Initially root is null */
4.	LOOP WHILE level < SpaceResolution AND CurrentNode is an index node
5.	Compare the key with the CurrentNode predicate
6.	IF no match AND PathShrink is “Tree Shrink”
7.	THEN Found = FALSE
8.	break
9.	FOR each slot i in the index node LOOP
10.	IF (Consistent(E[i], key, level)) THEN index = i
11.	IF None is consistent /*due to NodeShrink*/
12.	THEN Found = FALSE
13.	break
14.	CurrentNode = E[index].ptr /* the child pointed by entry E[index]*/
15.	level = level + 1
16.	IF CurrentNode is NOT NULL /* leaf node */
17.	Search for the key among leaf node entries
18.	IF Key is in the leaf node THEN Found = TRUE
19.	RETURN Found

Method *Search* in SP-GiST is exactly similar to that of the GiST scheme (see Hellerstein et al., 1995), and is given in Table 9 for completeness. Method *Search* uses method *Consistent* as the main navigation guide. Starting from the root, the algorithm will check the search item against all branches using the method *Consistent* till reaching leaf nodes (data nodes in SP-GiST).

The algorithm for method *Delete* in SP-GiST uses logical deletion. Deleted items are marked deleted and are not physically removed from the tree. This will save the effort of reorganizing the tree after each deletion, specially for data-driven space-partitioning trees. A rebuild is used from time to time as a *clean* procedure.

5. Concurrency and recovery in SP-GiST

Concurrency and recovery in GiST have been addressed in Chakrabarti and Mehrotra (1999a) and Kornacker et al. (1998). In Kornacker et al. (1998), the authors provide general algorithms for concurrency control in tree-based access methods as well as a recovery protocol and a mechanism for ensuring repeatable read isolation (Gray, 1978). They suggest the use of *Node Sequence Number* (NSN) for concurrency control, first introduced in Kornacker and Banks (1995).

For SP-GiST, a split (only at the leaf level) transforms a data node into an index node. Data is then distributed among new leaf nodes *rooted at that split node*. This fact simplifies the concurrency control mechanism significantly. For example, consider the case when a

search for a key is interleaved with an insertion that causes the splitting of the target node. By the time the search reaches the target node, it can not falsely conclude the non-existence of the searched key, e.g., in contrast to a B-Tree scenario, because the new node is an index node. In that case, no right links need to be maintained between leaves as the search will need to continue *deeper in the tree* not on the siblings level. Thus, no special sequence number is needed for the concurrent operation to know that the node in question has been split. The operation will directly continue working with the child nodes.

Phantom protection in GiST has also been addressed in two different techniques. Predicate locking (Eswaran et al., 1976) is used in Kornacker et al. (1998) while the authors in Chakrabarti and Mehrotra (1999a), propose a dynamic granular locking approach (GL/GiST) to phantom protection. We adopt the granular locking technique since it is more preferable and less expensive than predicate locking. The fact that a “Containment Hierarchy” exists in space-partitioning trees, represented by SP-GiST, makes the algorithm introduced in Chakrabarti and Mehrotra (1998, 1999a) highly applicable and much simpler. Hence, in SP-GiST, because the node predicates form a containment hierarchy, we simply use the node predicates for granular locks.

The main difference in SP-GiST is that a page may contain multiple SP-GiST nodes. A clustering algorithm will hold the mapping between nodes and pages. In this context, we assume that the node size is smaller than or equal to the page size. Hence the problem transforms to locking at a finer granularity. Treating nodes clustered in pages as records, granular locks (Gray and Reuter, 1993) are used. The recovery technique used in Kornacker et al. (1998) is directly applicable to SP-GiST.

6. Node clustering in SP-GiST

Node clustering means choosing the group of nodes that will reside together in the same disk page. Considering physical storage of the tree nodes, a direct and simple implementation of a node is to assign a disk page for each node. However, for very sparse nodes, this simple assignment will not be efficient for database use. We provide to the user a default node clustering method that is shown to perform well in the dynamic case (Diwan et al., 1996). However, we allow the user to override the default clustering method and provide a different node clustering policy that is more suitable for the type and nature of the operations to be performed on the constructed index. This will enhance the query response time of SP-GiST. We propose the interface method *Cluster* for this purpose.

Introducing new nodes in the tree structure, e.g., due to insertions, will internally invoke the dynamic clustering algorithm defined in *Cluster* to reconstruct the tree disk page structure and reflect the change. However, for unexperienced users or for typical database applications, SP-GiST has a default node clustering algorithm that achieves minimum height and hence minimum I/O access. The dynamic clustering algorithm in Diwan et al. (1996) is a good clustering algorithm and we use it as our default in SP-GiST. The pseudo code and a brief outline of the clustering algorithm is given in Appendix A.

The user can choose other clustering algorithms that reflect the application semantics, specially for non-traditional data types as in multimedia or video databases. Some alternatives are: (1) *Fill-Factor Clustering*: Tries to keep each page half-full for space utilization

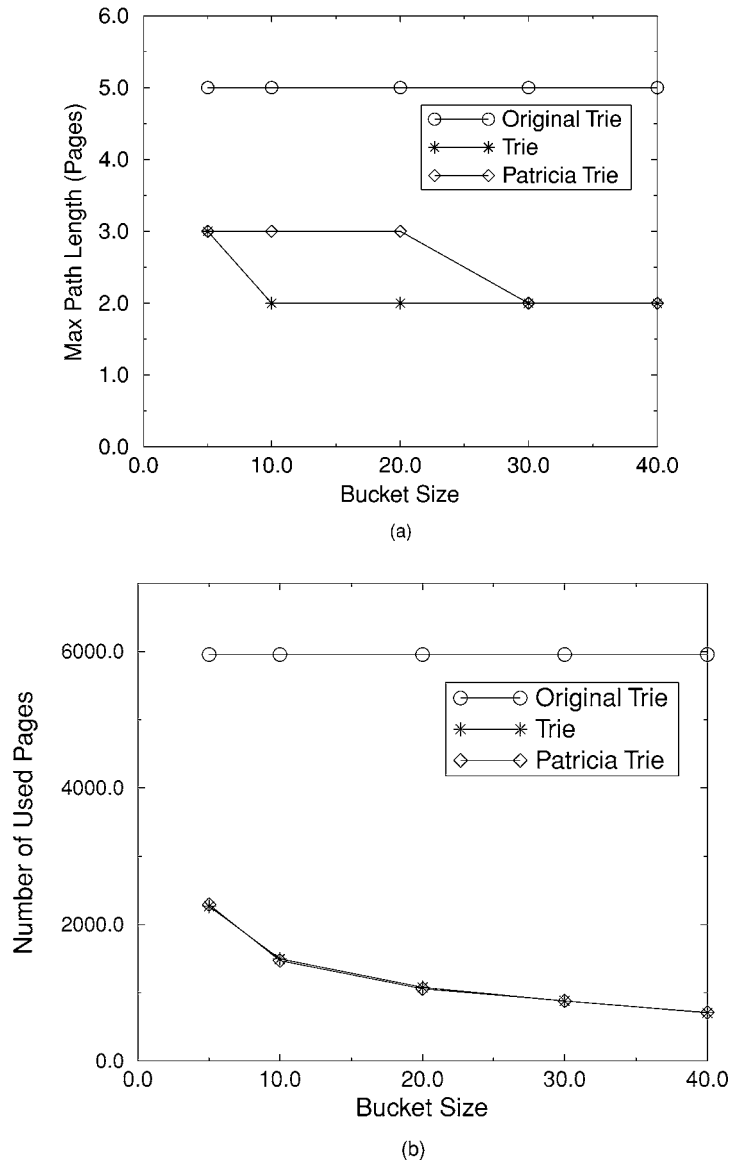
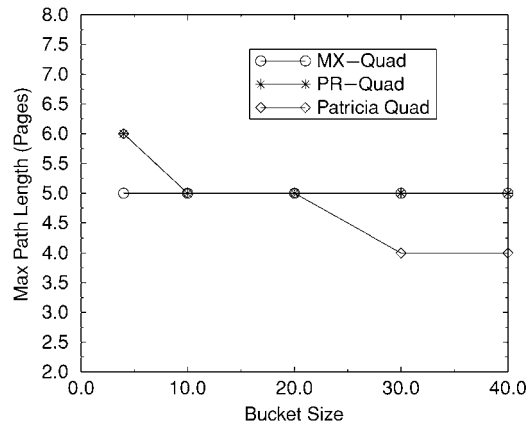
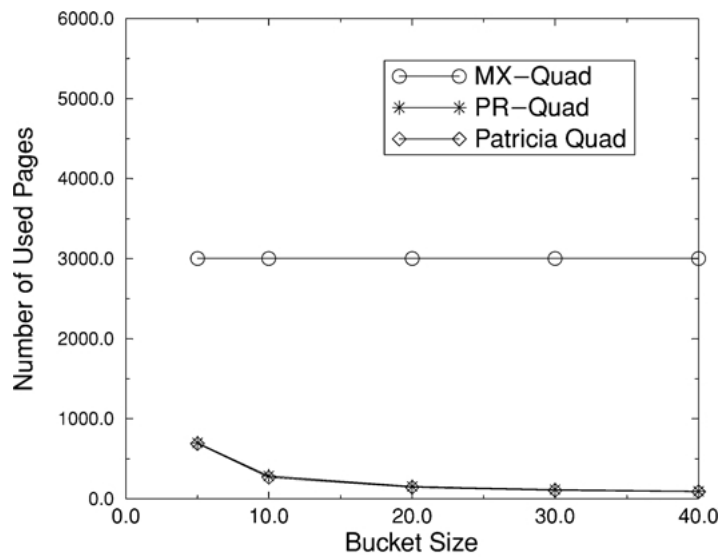


Figure 10. Effect of BucketSize on the maximum path length and number of used pages in the trie for different settings of PathShrink.

efficiency. (2) *Deep Clustering*: Chooses the longest linked subtree from the collection of page nodes to be stored together in the same page. This clustering method will enhance performance for depth-first traversal of trees. (3) *Breadth Clustering*: Chooses the maximum number of siblings of the same parent to be stored together in the same page.



(a)



(b)

Figure 11. Effect of *BucketSize* on the maximum path length and number of used pages in the quadtree for different settings of *PathShrink*.

7. Implementation and experimental results

We implemented SP-GiST using C++ on SunOS 5.6 (Sparc). As a proof of concept, using SP-GiST, we implemented the extensions for some data structures namely, the MX quadtree, the PR quadtree, the trie, and the Patricia trie. The implementation has proven the feasibility of representing space-partitioning trees using the interface proposed by SP-GiST and the settings in the tables in Section 3.3. We performed experiments on various settings of the tunable interface parameters; *BucketSize* and *PathShrink*. In our implementation we adopt

the *minimal height* clustering technique in Diwan et al. (1996). Results show that applying this clustering technique reduces the path length in terms of pages significantly.

As explained in Section 3.1, the interface parameter *PathShrink* can take one of three values; “Never Shrink”, “Leaf Shrink”, or “Tree Shrink”. For the trie, setting *PathShrink* to “Never Shrink” will have the effect of realizing the original trie, where splitting is performed to the maximum resolution of the space, leading to a long sparse tree. Setting *PathShrink* to “Leaf Shrink” will realize a common variant of the trie where data can be put in the first available node. On the other hand, if *PathShrink* is set to “Tree Shrink”, it will realize the Patricia implementation of the trie where no single-arc nodes are allowed.

Figure 10 gives the effect of this parameter on the trie data structure for various settings of *BucketSize* for a dataset of 10000 records with “string” keys. As expected, for the trie and the Patricia trie, the path length and the number of pages improve as the bucket size increases since less splitting takes place. On the other hand, the bucket size does not have an effect on the original trie. In this case, splitting will take place not because of the bucket size limit but to decompose the space to the maximum resolution. In the case of the original trie, each record will fall in a single node regardless of the setting of the bucket size.

For the quadtree, the same argument holds. Experimental results for point datasets of 10000 points are given in figure 11. In this case, setting *PathShrink* to “Never Shrink” will have the effect of realizing the MX quadtree while setting it to “Leaf Shrink” will realize the PR quadtree where data can be put in the first available node. Experiments with setting *PathShrink* to “Tree Shrink” show the realization of another variant of quadtree, where all *white nodes* are eliminated (Samet, 1990b), making it more attractive for databases and solving the problem of long degenerate quadtrees when the workload is highly skewed.

8. Conclusions

SP-GiST is a generalized space-partitioning tree implementation of a wide range of tree data structures that are not I/O-optimized for databases. This makes it possible to have single tree index implementation to cover various types of trees that suit different applications. Emerging database applications will require the availability of various index structures due to the heterogeneous collection of data types they deal with. SP-GiST is an interesting choice for multimedia databases, spatial databases, GIS, and other modern database systems. We have shown how to augment SP-GiST with parameters and methods that will enable the coverage of this class of space-partitioning trees. Clustering methods were also addressed to realize the use of these structures in practice in non-traditional database applications.

Recovery and concurrency for SP-GiST are addressed to enable the realization of SP-GiST in commercial database systems. Experiments proved the concept of SP-GiST and provided insight on the effect of the tunable interface parameters on the tree structure and performance.

Appendix A : Minimum height clustering algorithm

The clustering algorithm in Diwan et al. (1996) re-clusters the tree nodes into disk pages after updates to an already clustered state, and hence is dynamic, and guarantees *minimal*

height mapping after deleting or inserting in the tree. The pseudo code of the algorithm is given in Table A.1. The algorithm begins by removing all deleted nodes from the disk pages. All new nodes or affected roots of subtrees are kept in a set of affected nodes. Processing all the affected nodes starts bottom-up (no node is processed until all its children are processed).

Table A.1. Minimum height clustering algorithm.

1.	PROCEDURE ReCluster-Bottom Up(TreeNode root)
2.	$S = \{\}$;
3.	FOR each node n in delete-list (List of deleted nodes) DO
4.	Remove n from its current page
5.	IF n is the last node in the page
6.	THEN delete the page
7.	decluster(root)
8.	/* S is now the set of nodes that are affected */
9.	LOOP WHILE there are nodes in S that are not yet processed
10.	Choose an affected node P that is either a leaf or
11.	all of whose children are either not in S or have been processed
12.	process-node(P)
1.	PROCEDURE decluster(node n)
2.	add n to S
3.	IF (n is not a new inserted node) THEN
4.	remove n from its current page
5.	IF n is the last node in the page
6.	THEN delete the page
7.	FOR each child n_l of n DO
8.	IF (n_l is a new inserted node or if the subtree from n_l is modified)
9.	THEN decluster(n_l)
10.	ELSE
11.	IF (n_l is in the same page as n)
12.	THEN move n_l and all its descendants in the same page as n to a new page
1.	PROCEDURE process-node(TreeNode P)
2.	IF P is a leaf node
3.	THEN create a new page C containing node P
4.	ELSE Let $P_1 \dots P_n$, be the children of P
5.	Let $C_1 \dots C_n$ be the pages containing $P_1 \dots P_n$, respectively
6.	Let $P_{i1} \dots P_{im}$, be the children among the above whose page height is the greatest
7.	IF node P and the contents of the pages $C_{i1} \dots C_{im}$ can be merged in 1 page
8.	THEN merge the contents of $C_{i1} \dots C_{im}$, into a new page C
9.	delete $C_{i1} \dots C_{im}$
10.	ELSE create a new page C containing only P

The algorithm tries to put the longest path of nodes together in the same page. The authors in Diwan et al. (1996) have shown that the algorithm achieves minimal height mapping. They suggested some heuristics of merging sparse pages to achieve minimum fill factor of at least 50%.

References

- Aref, W.G., Barbará, D., and Vallabhaneni, P. (1995). The Handwritten Trie: Indexing Electronic Ink. In *Proceedings of SIGMOD'95*, San Jose, California.
- Barclay, T., Gray, J., and Slutz, D. (2000). Microsoft Terraserver: A Spatial Data Warehouse. In *Proceedings of SIGMOD'00*, Dallas, TX.
- Bayer, R. (1997). The Universal B-Tree for Multidimensional Indexing: General Concepts. In *Lecture Notes in Computer Science*, Vol. 1274. Berlin:Springer-Verlag.
- Beckmann, N., Kriegel, H.P., Schneider, R., and Seeger, B. (1990). The R*-Tree: An Efficient Robust Access Method for Points and Rectangles. *SIGMOD Record*, 19(2).
- Bentley, J.L. (1975). Multidimensional Binary Search Trees used for Associative Searching. *Communications of the ACM*, 19, 509–517.
- Berchtold, S., Boehm, C., and Kriegel, H.-P. (1998). Improving the Query Performance of High-Dimensional Index Structures by Bulk Load Operations. In *Lecture Notes in Computer Science*, Vol. 1377. Berlin:Springer-Verlag.
- Brinkhoff, T., Kriegel, H.-P., and Seeger, B. (1996). Parallel Processing of Spatial Joins using R-Trees. In *Proceedings of ICDE'96*, New Orleans, Louisiana.
- Chakrabarti, K. and Mehrotra, S. (1998). Dynamic Granular Locking Approach to Phantom Protection in R-Trees. In *Proceedings of ICDE'98*, Orlando, Florida, USA (pp. 446–454).
- Chakrabarti, K. and Mehrotra, S. (1999a). Efficient Concurrency Control in Multidimensional Access Methods. In *Proceedings of SIGMOD'99*, Philadelphia, Pennsylvania, USA (pp. 25–36).
- Chakrabarti, K. and Mehrotra, S. (1999b). The Hybrid Tree: An Index Structure for High Dimensional Feature Spaces. In *Proceedings of ICDE'99*, Sydney, Australia.
- Oracle Corporation. (1999). Oracle Spatial (Data Sheet). <http://www.oracle.com/database/documents/spatial-ds.pdf>.
- de la Briandais, R. (1959). File Searching using Variable Length Keys. In *Proceedings of the Western Joint Computer Conference* (pp. 295–298).
- DeWitt, D.J., Kabra, N., Luo, J., Patel, J.M., and Yu, J.-B. (1994). Client-Server Paradise. In *Proceedings of VLDB'94*, Santiago, Chile (pp. 558–569).
- Diwan, A.A., Rane, S., Seshadri, S., and Sudarshan, S. (1996). Clustering Techniques for Minimizing External Path Length. In *Proceedings of VLDB'96*, Mumbai, India (pp. 342–353).
- Esperanca, C. and Samet, H. (1996). Spatial Database Programming using Sand. In *Proceedings of the Seventh International Symposium on Spatial Data Handling*.
- Eswaran, K.P., Gray, J.N., Lorie, R.A., and Traiger, I.L. (1976). The Notions of Concurrency and Predicate Locks in a Data Base System. *Communications of the ACM*, 19(11).
- Faloutsos, C. and Gaede, V. (1996). Analysis of n -Dimensional Quadtrees using the Hausdorff Fractal Dimension. In *Proceedings of VLDB'96* (pp. 40–50).
- Faloutsos, C., Jagadish, H.V., and Manolopoulos, Y. (1997). Analysis of the n -Dimensional Quadtree Decomposition for Arbitrary Hyperrectangles. *TKDE*, 9(3), 373–383.
- Finkel, R.A. and Bentley, J.L. (1974). Quad Trees: A Data Structure for Retrieval on Composite Key. *Acta Informatica*, 4(1), 1–9.
- Fredkin, E. (1960). Trie Memory. *Communications of the ACM*, 3, 490–500.
- Gaede, V. and Gunther, O. (1998). Multidimensional Access Methods. *ACM Computer Surveys*, 30(2), 170–231.
- Gargantini, I. (1982). An Effective Way to Represent Quadtrees. *Communications ACM*, 25(12), 905–910.
- Gray, J.N. (1978). Notes on Data Base Operating Systems. In R. Bayer, R.M. Graham, and G. Seegmüller (Eds.), *LNCS, 'Operating Systems, an Advanced Course', (Heidelberg, FRG and NewYork NY, USA:Springer Verlag, Vol. 60.*

- Gray, J.N. and Reuter, A. (1993). *Transaction Processing: Concepts and Techniques*. Data Management Systems. San Mateo, CA: Morgan Kaufmann Publishers.
- Güting, R.H. (1994). An Introduction to Spatial Database Systems. *VLDB Journal*, 3(4), 357–399.
- Guttman, A. (1984). R-Trees: A Dynamic Index Structure for Spatial Searching. In *Proceedings of SIGMOD'84* (pp. 47–57).
- Hellerstein, J.M., Naughton, J.F., and Pfeffer, A. (1995). Generalized Search Trees for Database System. In *Proceedings of VLDB'95*.
- Hellerstein, J.M. and Pfeffer, A. (1994). The RD-Tree: An Index Structure for Sets. Technical Report, University of Wisconsin Computer Science.
- Kedem, G. (1982). The quad-CIF Tree: A Data Structure for Hierarchical on-line Algorithms. In *ACM IEEE Nineteenth Design Automation Conference Proceedings*, Los Alamitos, CA, USA (pp. 352–357).
- Klinger, A. (1971). Pattern and Search Statistics. In S. Rustagi (Ed.), *Optimizing Methods in Statistics* (pp. 303–337).
- Knuth, D.E. (1973). *The Art of Computer Programming*, Vol. 3, Sorting and Searching. Reading, MA: Addison-Wesley.
- Kornacker, M. and Banks, D. (1995). High-Concurrency Locking in R-Trees. In *Proceedings of VLDB'95*, Zurich, Switzerland.
- Kornacker, M., Mohan, C., and Hellerstein, J.M. (1998). Concurrency and Recovery in Generalized Search Trees. *Proceedings of SIGMOD'98* (pp. 62–72).
- Lin, K.-I., Jagadish, H.V., and Faloutsos, C. (1994). The TV-Tree: An Index Structure for High-Dimensional Data. *VLDB Journal*, 3(4), 517–542.
- Morrison, D.R. (1968). PATRICIA—Practical Algorithm to Retrieve Coded in Alphanumeric. *J. Assoc. Comput. Mach.*, 15(4), 514–534.
- Nelson, R.C. and Samet, H. (1986). A Consistent Hierarchical Representation for Vector Data. *Computer Graphics*, 20(4), 197–206. SIGGRAPH '86 Proceedings.
- Nievergelt, J., Hinterberger, H., and Sevcik, K. (1984). The Grid File: An Adaptable Symmetric Multi-Key File Structure. *ACM Transactions on Database Systems*, 9(1), 38–71.
- Orenstein, J.A. (1982). Multidimensional Tries used for Associative Searching. *Information Processing Letters*, 14(4), 150–157.
- Orenstein, J.A. and Manola, F. (1988). PROBE Spatial Data Modeling and Query Processing in an Image Database Application. *IEEE Transactions on Software Engineering*, 14(5), 611–629.
- Papadias, D., Mamoulis, N., and Delis, V. (1998). Algorithms for Querying by Spatial Structure. In *Proceedings of VLDB'98*, New York City, New York, USA (pp. 546–557).
- Samet, H. (1990a). *Applications of Spatial Data Structures*. Reading, MA: Addison-Wesley.
- Samet, H. (1990b). *The Design and Analysis of Spatial Data Structure*. Reading, MA: Addison-Wesley.
- Samet, H. and Webber, R.E. (1985). Storing a Collection of Polygons using Quadrees. *ACM Transactions on Graphics*, 4(3), 182–222.
- Seeger, B. and Kriegel, H.-P. (1990). The Buddy-Tree: An Efficient and Robust Access Method for Spatial Data Base Systems. In *Proceedings of VLDB'90*, Queensland, Australia.
- Sellis, T., Roussopoulos, N., and Faloutsos, C. (1987). The R+ -Tree: A Dynamic Index for Multi-Dimensional Objects. In *Proceedings of VLDB'87*, Brighton, UK.
- Sellis, T.K., Roussopoulos, N., and Faloutsos, C. (1997). Multidimensional Access Methods: Trees Have Grown Everywhere. In *Proceedings of VLDB'97* (pp. 13–14).
- Szalay, A., Kunszt, P., Thakar, A., Gray, J., Slutz, D., and Brunner, R. (2000). Designing and Mining Multi-terabyte Astronomy Archives: The Sloan Digital Sky Survey. In *Proceedings of SIGMOD'00*, Dallas TX (pp. 451–462).
- Tanimoto, S. and Pavlidis, T. (1975). A Hierarchical Data Structure for Picture Processing. *Computer Graphics and Image Processing*, 4(2), 104–119.