

R-Tree Based Indexing of General Spatio-Temporal Data

Simonas Šaltenis and Christian S. Jensen

TR-45

A TIMECENTER Technical Report

Title R-Tree Based Indexing of General Spatio-Temporal Data

Copyright © 1999 Simonas Šaltenis and Christian S. Jensen. All rights reserved.

Author(s) Simonas Šaltenis and Christian S. Jensen

Publication History December 1999. A TIMECENTER Technical Report.

TIMECENTER Participants

Aalborg University, Denmark

Christian S. Jensen (codirector), Michael H. Böhlen, Curtis E. Dyreson, Heidi Gregersen, Dieter Pfoser, Simonas Šaltenis, Janne Skyt, Giedrius Slivinskas, Kristian Torp

University of Arizona, USA

Richard T. Snodgrass (codirector), Bongki Moon

Individual participants

Fabio Grandi, University of Bologna, Italy

Nick Kline, Microsoft, USA

Gerhard Knolmayer, University of Bern, Switzerland

Thomas Myrach, University of Bern, Switzerland

Kwang W. Nam, Chungbuk National University, Korea

Mario A. Nascimento, University of Alberta, Canada

Keun H. Ryu, Chungbuk National University, Korea

Michael D. Soo, amazon.com, USA

Andreas Steiner, TimeConsult, Switzerland

Vassilis Tsotras, University of California, Riverside, USA

Jef Wijzen, Vrije Universiteit Brussel, Belgium

Carlo Zaniolo, University of California, Los Angeles, USA

For additional information, see The TIMECENTER Homepage:

URL: <<http://www.cs.auc.dk/TimeCenter>>

Any software made available via TIMECENTER is provided “as is” and without any express or implied warranties, including, without limitation, the implied warranty of merchantability and fitness for a particular purpose.

The TIMECENTER icon on the cover combines two “arrows.” These “arrows” are letters in the so-called *Rune* alphabet used one millennium ago by the Vikings, as well as by their predecessors and successors. The Rune alphabet (second phase) has 16 letters, all of which have angular shapes and lack horizontal lines because the primary storage medium was wood. Runes may also be found on jewelry, tools, and weapons and were perceived by many as having magic, hidden powers.

The two Rune arrows in the icon denote “T” and “C,” respectively.

Abstract

Real-world objects are inherently spatially and temporally referenced, and many database applications rely on databases that record the past, present, and anticipated future locations of, e.g., people or land parcels. As a result, indices that efficiently support queries on the spatio-temporal extents of objects are needed. In contrast, past indexing research has progressed in largely separate spatial and temporal streams. In the former, focus has been on one-, two-, or three-dimensional space; and in the latter, focus has been on one or both of the temporal aspects, or dimensions, of data known as transaction time and valid time. Adding time dimensions to spatial indices, as if time was a spatial dimension, neither supports nor exploits the special properties of time. On the other hand, temporal indices are generally not amenable to extension with spatial dimensions.

This paper proposes an efficient and versatile technique for the indexing of spatio-temporal data with discretely changing spatial extents: the spatial aspect of an object may be a point or may have an extent; both the transaction time and valid time are supported; and a generalized notion of the current time, *now*, is accommodated for the temporal dimensions. The technique extends the previously proposed R^* -tree and borrows from the GR-tree, and it provides means of prioritizing space versus time, enabling it to adapt to spatially and temporally restrictive queries. Performance experiments were performed to evaluate different aspects of the proposed indexing technique, and are included in the paper.

1 Introduction

An increasing number of database applications manage spatiotemporal aspects of real-world, physical objects. Such objects have positions and extents in space, and these positions and extents may change as time passes. Example spatiotemporal objects range from people and vehicles to land parcels and hurricanes [6]. The past, present, and anticipated future positions of such objects are often of interest in population studies, urban planning, marketing, traffic management, vehicle navigation support, land management, and environmental studies. Some of these applications involve continuous movement, while others involve discrete movement. This paper proposes an efficient indexing technique for the latter kind of application.

Two temporal aspects of data are generally considered fundamental. The *valid time* of a fact is the time(s), past, present, or future, when the fact is true in the modeled reality, while the *transaction time* of a fact is the time(s) when the fact was or is current in the database [9, 16]. Data with both valid and transaction time associated is termed bitemporal. Full spatiotemporal support implies considering these two temporal aspects as well as two or three dimensions of space.

The previously proposed spatiotemporal indices [19, 13] assume only one time dimension, adopting one of two approaches. Either overlapping index structures are used that index spatial objects at different time instances and save space by sharing the unchanged parts of the indices [14, 23, 22]. Or time is added as another dimension to an existing spatial index [20].

The former approach may not generalize well to two time dimensions, and treating time as a spatial dimension has certain disadvantages. Specifically, time intervals associated with objects may be *now-relative*, meaning that their end points track the progressing current time. Consider the recording of addresses. The time a person resides at a given address may often extend from a known start time (the valid-time interval begin) to some unknown future time, which is captured by letting the valid-time interval end extend to the progressing current time. The same applies to the transaction time—the time a tuple is inserted into the database is known, but we do not know when the tuple will be deleted. This notion of *now* is peculiar to time and has no counterpart in space.

In order to support these aspects of time together with spatial dimensions, we base our new indexing technique on the efficient, spatial R^* -tree [1] and also attempt to reuse ideas presented in the GR-tree [2]. This latter index is itself based on the R^* -tree and is arguably the best index for general bitemporal data, which includes now-relative data. (The only competitive alternative known to us, when permitting now-relative bitemporal data, is the 4-R index [3].) The GR-tree represents the valid and transaction times of

a fact as a two-dimensional region and extends the \mathbb{R}^* -tree's structure and algorithms to support the kind of continuously growing regions that result from now-relative transaction and valid times. The \mathbb{R} -tree is used as the basis of the GR-tree and of the new indexing technique, because it partitions the data itself, as opposed to the embedding space, as do some other multidimensional access methods, e.g., Quadtrees. Unless transformations are employed, the growth of data regions makes it difficult to extend the latter type of access method to index now-relative data. The index presented here, called the \mathbb{R}^{ST} -tree ("spatio-temporal"), also handles more general now-relative data than does the GR-tree.

Section 2 presents the type of data that can be indexed with the new index. Section 3 solves the problem of choosing the minimum bounding regions to be used in the entries of the index, and algorithms for the tree operations are given in Section 4. Section 5 presents performance studies, and Section 6 concludes the paper.

2 Background

This section presents two example application areas illustrating the spatiotemporal data and queries supported, and it describes the temporal aspects of the data to be indexed.

2.1 Example Application Areas

Demographic data is used in applications such as advertisement, direct marketing, urban planning, and social studies. It is of great value to record the changing addresses, or locations, of the people in the database, and we will assume that the database records the history of the position (e.g., latitude and longitude) of each person's residence (and possibly place of work). This means that we are faced with two-dimensional point locations that may change discretely from time to time. With this data available, it is possible to answer a query such as "Who lived close to a chemical plant during the period when the plant leaked toxic materials into the environment?" or "Who moved out of a certain neighborhood during a specific period?"

Cadastral systems exemplify another kind of spatiotemporal application. Here the boundaries of land parcels are recorded together with the history of their change. And for legal reasons, all records must be maintained in an append-only fashion so that even mistaken records are retained. This is accomplished via transaction-time support. For indexing, we bound each land parcel with a minimum bounding rectangle and associate it with valid- and transaction-time intervals. An example query could be "Who owned some part of a specific piece of land sometime during the period from 1975 to 1980, as best known by 1990?" This is a spatiotemporal containment query that constrains all four dimensions: the two spatial dimensions, valid time, and transaction time.

The above examples illustrate the types of data and queries supported by the new index. According to the criteria for classification of spatiotemporal access methods proposed by Theodoridis et al. [19], the \mathbb{R}^{ST} -tree supports two-dimensional points and regions; it is bitemporal; both the cardinality and the positions of the spatial objects may change over time; the index is dynamic; and spatial, temporal, and spatiotemporal containment queries are supported. As mentioned, the index also supports now-relative data, which we proceed to characterize.

2.2 General Bitemporal Data

We adopt the standard four-timestamp format [17] for capturing valid and transaction time. With this format, each tuple is timestamped with four time attributes: VT^{h} and VT^{l} —the times when the tuple's information became and ceased to be true in the modeled reality; TT^{h} and TT^{l} —the times when the tuple became and ceased to be current in the database.

A tuple is now-relative if it is valid until the current time or is part of the current database state. This is captured using variables that denote the current time in the attributes VT^+ and TT^+ [4]. Variable UC (denoting ‘until changed’) is used in TT^+ , and variable NOW is used in VT^+ . Table 1 shows an example table with now-relative data. The time granularity is a month, and the current time (CT) is 9/97.

Table 1: The Demographic Relation

	Person	Position	TT^+	TT^+	VT^+	VT^+
(1)	John	Pos1	4/97	UC	3/97	5/97
(2)	Tom	Pos2	3/97	7/97	6/97	8/97
(3)	Jane	Pos3	5/97	UC	5/97	NOW
(4)	Julie	Pos4	3/97	7/97	3/97	NOW
(5)	Julie	Pos4	8/97	UC	3/97	7/97
(6)	Ann	Pos5	5/97	UC	3/97	NOW + 1
(7)	Scott	Pos6	4/97	UC	5/97	NOW - 2

Tuple 1 records that the information “John lived at Pos1” was true from 3/97 to 5/97 and that this was recorded during 4/97 and is still current. Tuple 3 records that “Jane lives at Pos3” from 5/97 until the current time, that we recorded this belief on 5/97, and that this remains part of the current state.

In the case of Tuple 3, the valid-time end being equal to NOW means that we currently do not believe that Jane will live at Pos3 next month (on 10/97). Sometimes such an assumption is too pessimistic. For example, there can exist a restriction that a person can only move with a month’s notice. Then we would believe that Jane would live at Pos3 also next month. To record this type of knowledge, Clifford et al. [4] propose to use $NOW + \Delta$ in the valid-time end attribute. The offset Δ may be any integer, positive or negative. For example, the latter case can be useful if a regulation states that a person has two months to report a change of address. Then we would not know with certainty that Jane currently lives at Pos3, but would be certain that she lived there two months ago. Tuples 6 and 7 exemplify the usage of positive and negative offsets.

Specific constraints apply to insertions, deletions, and modifications. When inserting a new tuple, the natural constraint $VT^+ \leq VT^+$ applies to valid time; and the constraints $TT^+ = CT$ and $TT^+ = UC$ apply to transaction time. Any *current* database tuple can be deleted or modified. Deleting a tuple, the TT^+ value UC is changed to the fixed value $CT - 1^1$, eliminating the tuple from the current state (e.g., Tuple 2), but retaining it in the database. A modification is modeled as a deletion followed by an insertion (e.g., an update led to Tuples 4 and 5).

The temporal aspects of a tuple can be represented graphically by a two-dimensional (“bitemporal”) region in the space spanned by transaction and valid time [9]. Figure 1 illustrates the *bitemporal regions* of the tuples from Table 1. A now-relative transaction-time interval yields a rectangle that “grows” in the transaction time direction as time passes (Tuple 1). Having both transaction- and valid-time intervals being now-relative yields a stair-shaped region growing in both the transaction time and valid time as time passes (Tuple 3).

The condition $VT^+ < TT^+ + \Delta$ yields a stair-shape with a high first step (Tuple 6). If, on the other hand, $VT^+ > TT^+ + \Delta$, the region has no extent for $VT^+ - (TT^+ + \Delta)$ time units from the time of its insertion; it appears only in the future. This is the case for Tuple 7, where, although $TT^+ = 4/97$, the actual region appeared only at 7/97; from 4/97 to 7/97, we supposed that Scott did not live, lives, or will live at Pos6. If at some time a tuple is deleted and thus stops being current, the bitemporal region ceases to grow (Tuples 2 and 4). If $VT^+ > TT^+ + \Delta$ and the tuple is deleted before its region appears, the region will never have an extent.

¹We use closed intervals and let $[TT^+, TT^+]$ denote the interval that includes TT^+ and TT^+ .

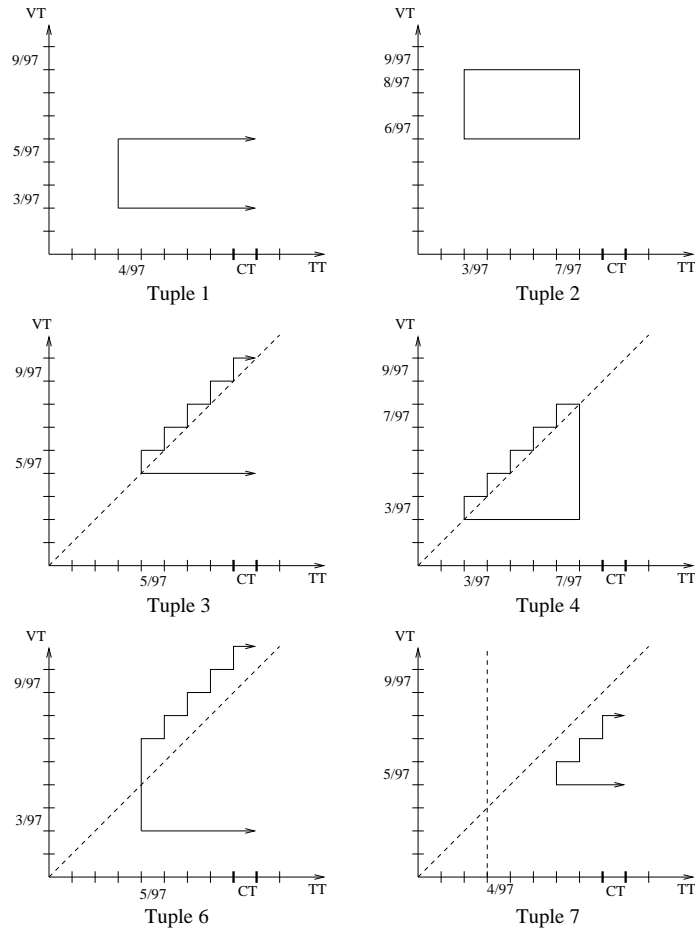


Figure 1: Bitemporal Regions of Tuples from Table 1

In general, we obtain six combinations of time attributes for which the bitemporal regions are qualitatively different. The combinations are listed in Table 2 where ‘tt1,’ ‘tt2,’ ‘vt1,’ and ‘vt2’ denote ground values that satisfy the constraints given above and the offset Δ is an integer.

The spatial representation of the bitemporal extents of tuples suggests the use of some spatial index as the basis for a bitemporal index. As mentioned in Section 1, this approach may also facilitate the incorporation of spatial dimensions into the resulting bitemporal index.

3 Index Structure

The new index has the same overall structure as the well-known R-tree [7] (and \mathbb{R}^* -tree [1]). As for the R-tree, each internal node is a record of index entries, each of which is a pair of a pointer to a node at the next level in the tree and a region that encloses all regions in the node pointed to. But unlike in the R-tree, the enclosing region is not simply a minimum bounding four-dimensional hyper-rectangle (in the sequel, we will simply use “rectangle” for both two and four dimensions). This section discusses the design of the bounding regions used in the \mathbb{R}^{ST} -tree.

The spatial and temporal dimensions may be considered separately, as a four-dimensional region is the product of its two two-dimensional spatial and bitemporal regions. For the spatial dimensions, the bounding

Table 2: Possible Combinations of Time Attributes

	TT^+	TT^-	VT^+	VT^-	Conditions	Examples
Case 1	tt1	UC	vt1	vt2		Tuples 1 and 5
Case 2	tt1	tt2	vt1	vt2		Tuple 2
Case 3	tt1	UC	vt1	$NOW + \Delta$	$(vt1 \leq tt1 + \Delta)$	Tuples 3 and 6
Case 4	tt1	tt2	vt1	$NOW + \Delta$	$(vt1 \leq tt1 + \Delta)$	Tuple 4
Case 5	tt1	UC	vt1	$NOW + \Delta$	$(vt1 > tt1 + \Delta)$	Tuple 7
Case 6	tt1	tt2	vt1	$NOW + \Delta$	$(vt1 > tt1 + \Delta)$	

region is simply a rectangle, as in the R-tree. We thus proceed to consider only the temporal regions, and we let the terms *minimum bounding rectangle* and *minimum bounding region* mean the projections of the corresponding four-dimensional region into the bitemporal hyper-plane.

In its leaf nodes, the R^{ST} -tree records the exact geometry of the bitemporal regions indexed (recall Section 2). The following format is used for a leaf-node index entry.

$$(TT^+, TT^-, VT^+, VT^-/\Delta, now-flag, \langle spatial\ part \rangle, \langle pointer \rangle)$$

Here, the first three components are the attributes introduced in the previous section, and they may obtain the same values as described there. The fourth and fifth components compactly encode the values of the VT^- attribute. A value of the form $NOW + \Delta$ is captured by setting the *now-flag* and storing Δ in VT^-/Δ ; other values are stored in this attribute, without the *now-flag* set. Variable UC is represented as a special, reserved value from the domain of timestamps.

Ideally, a non-leaf node entry should contain a minimum bounding region that encloses the regions pointed to as tightly as possible and that retains this property, should any of the enclosed regions be growing. This implies that the bounding regions must also grow. In addition, the bounding regions should be compact and simple to manipulate.

Selecting a type of bounding bitemporal regions that satisfies these requirements is not trivial. For example, consider the situation illustrated in Figure 2, where a growing stair-shape is placed together with a non-growing (“static” for short) rectangle in a minimum bounding rectangle growing in transaction time. One day, the stair-shape will outgrow its bounding rectangle, making the tree invalid.




Figure 2: A “Hidden,” Growing Stair-Shape

As a first step in designing bounding regions for the tree that also handle this problem, we consider four types of bounding regions that may be used.

Following the R-tree, the minimum bounding regions may be rectangles, which now may also be growing in transaction time only or in both directions. Such rectangles may then use UC for TT^\dagger and $NOW + \Delta$ for VT^\dagger . Then the TT^\dagger value of UC and the VT^\dagger value of $NOW + \Delta$ would represent a rectangle growing in both directions. In leaf node entries, this combination of timestamp values represents a growing stair-shape. This interpretation may also be chosen for non-leaf nodes. We consider these two first types of bounding regions in connection with the above-mentioned “hidden” stair-shapes.

When we employ bounding rectangles that grow in both directions, a bounding rectangle for a set of regions that includes one or more growing stair-shapes has initially its VT^\dagger set to $NOW + \Delta$, where Δ is large enough for the rectangle to enclose all rectangles above the stair-shapes. This is illustrated in Figure 3(a), where the regions from Figure 2 are being bounded. This approach has the disadvantage that the bounding rectangles are not minimum after some time, making the index unnecessarily cover areas, termed “dead space,” not covered by any data regions and also leading to increased overlap among bounding rectangles.

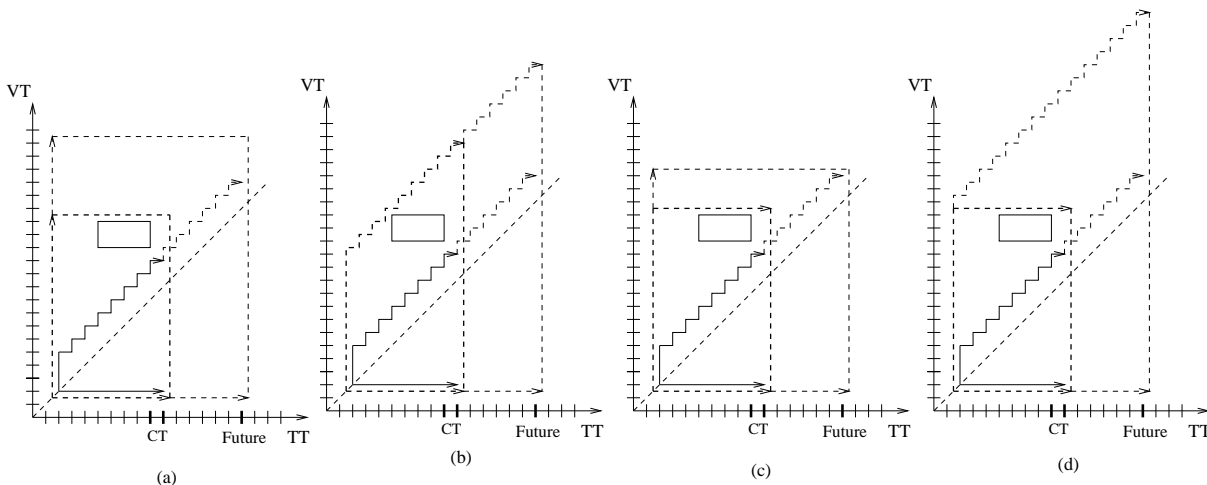


Figure 3: (a) A Bounding Rectangle, (b) a Bounding Stair-Shape, (c) a Bounding Rectangle Using the Flag *Hidden*, and (e) a Bounding Stair-Shape Using the Flag *Hidden*

The use of a growing stair-shape large enough to bound all regions is illustrated in Figure 3(b). At the current time, such a bounding stair-shape may have a larger area than the corresponding bounding rectangle (cf. Figures 3(a) and (b)). However, after some time, the area of the rectangle becomes larger than the area of the stair-shape. It can be proven that for any set of bitemporal regions that can be bounded with either a rectangle growing in both directions or a growing stair-shape, the rectangle will eventually outgrow the stair-shape.

The last two types of bounding regions derive from the GR-tree, where hidden stair-shapes are handled using a flag *hidden* in the non-leaf node entries. This enables regions with hidden stair-shapes and rectangles above the diagonal to initially be bounded with a rectangle growing only in transaction time. This idea may also be applied here, but due to the R^{ST} -tree’s more general bitemporal regions, the flag *hidden* should be accompanied by an offset Δ_{max} , which is the maximum offset of the hidden stair-shapes. Then, whenever the flag *hidden* of a rectangle is set and the current time plus Δ_{max} exceeds the VT^\dagger value for this rectangle, VT^\dagger may be adjusted to continue to bound the stair-shapes that have outgrown its rectangle.

In effect, we turn the rectangle, which grows only in transaction time, into a rectangle growing in both directions or into a growing stair-shape. In the first case, VT^\dagger is set to $NOW + \Delta_{max}$. This leads to a bounding rectangle that is minimal at all times, see Figure 3(c).

In the second case, where the original rectangle is replaced by a growing stair-shape, VT^+ is set to $NOW + \Delta$, where Δ is the difference between the old value of VT^+ and the value of TT^+ . The resulting growing stair-shape fully encloses the rectangle it replaces, see Figure 3(d). Creating a bounding stair-shape this large is necessary because we cannot know what regions the original rectangle is bounding without accessing the child node. Some of those regions may be located in the upper-left corner. The large bounding stair-shapes is the main disadvantage of this approach.

The two types of bounding regions that employ the *hidden* flag have the disadvantage of a larger and more complex entry structure. This might be justified if the flag is used frequently. But experiments with the GR-tree [2] show that the flag is not used very often (because the insertion algorithms obtain “good” groupings of regions).

In conclusion, we use the second kind of bounding region, illustrated in Figure 3(b), in the \mathbb{R}^T -tree. With this solution, the structure of a non-leaf node and a leaf node are exactly the same, and the kinds of bounding regions available is exactly the same as the kinds of bitemporal data regions possible. This homogeneity simplifies the algorithms associated with the tree. A bounding region (TT^+, TT^-, VT^+, VT^-) for a set of bitemporal regions r_i is computed as follows, where VT^+ is subsequently encoded in the VT^+/Δ and *now-flag* attributes given for leaf nodes earlier in this section.

$$\begin{aligned}
TT^+ &= \min_i(TT_{r_i}^+) \\
VT^+ &= \min_i(VT_{r_i}^+) \\
TT^- &= \begin{cases} UC & \text{if } \exists i (TT_{r_i}^- = UC) \\ \max_i(TT_{r_i}^-) & \text{otherwise} \end{cases} \\
VT^- &= \begin{cases} NOW + \Delta_b & \text{if } \exists i (TT_{r_i}^- = UC \wedge VT_{r_i}^- = NOW + \Delta) \vee \\ & (\forall i (TT_{r_i}^- \neq UC) \wedge \\ & \quad area(TT^+, TT^-, VT^+, NOW + \Delta_b) < area(TT^+, TT^-, VT^+, VT_b^-)) \\ VT_b^- & \text{otherwise} \end{cases}
\end{aligned}$$

Note that the argument to function *area* denotes a static stair-shape where the function is used first and a static rectangle were it is used the second time. Also

$$\begin{aligned}
\Delta_b &= \max(\max_{i: VT_{r_i}^- \neq NOW + \Delta} (VT_{r_i}^- - TT_{r_i}^-), \max_{i: VT_{r_i}^- = NOW + \Delta} (\Delta_{r_i})) \\
VT_b^- &= \max(\max_{i: VT_{r_i}^- \neq NOW + \Delta} (VT_{r_i}^-), \max_{i: VT_{r_i}^- = NOW + \Delta} (TT_{r_i}^- + \Delta_{r_i}))
\end{aligned}$$

The computations of TT^+ , TT^- , and VT^+ are trivial. The value of VT^- determines whether the region is a stair-shape or a rectangle. To avoid growing stair-shapes, we set VT^- to $NOW + \Delta_b$ only if at least one growing stair-shape exists among the regions, or if all the regions are static and a minimum bounding stair-shape (also static) has a smaller area than a minimum bounding rectangle.

It should be noted that growing stair-shapes with $VT^- > TT^- + \Delta$ (Figure 1, Tuple 7) are treated exactly the same as all other growing stair-shapes. When bounding such stair-shapes, the minimum bounding region is set so that it encloses all regions that will appear in the future, although it may not be minimal for some time.

4 Index Algorithms

This section describes the algorithms that serve to maintain the index structure just described. Section 4.1 offers a road map of the algorithms. Section 4.2 extends the heuristics used in the \mathcal{R} -tree insertion algorithm to support the new types of data regions. Sections 4.3–4.5 then describe the insertion algorithm, and Section 4.6 shows how the index can be tuned to support specific types of queries.

4.1 Road Map of Algorithms

The overall structure of the \mathcal{R}^{ST} -tree is the same as that of the \mathcal{R}^* -tree. As a result, the \mathcal{R}^* -tree’s search, deletion, and insertion algorithms may be used for the new index, provided that these algorithms employ a set of algorithms that manipulate the new kinds of regions presented in Section 3. This set includes an algorithm that determines whether a pair of regions overlap and algorithms that compute the area and margin of a region, the intersection of a pair of regions, and the minimum bounding region of a node. Indeed, the \mathcal{R}^{ST} -tree reuses the \mathcal{R}^* -tree’s search and deletion² algorithms, which have been adapted to use these low-level algorithms. This leaves only insertion unaccounted for.

The insertion algorithm is central because it is responsible for maintaining an efficient tree. The \mathcal{R} -tree’s insertion algorithm is likely to be inefficient for the new types of data and bounding regions in the \mathcal{R}^{ST} -tree. Before improving this algorithm, we describe its general working.

Given a new entry, the \mathcal{R}^* -tree’s insertion algorithm starts at the root node and traverses the tree downwards until it reaches a leaf node. At each visited node, it uses the ChooseSubtree algorithm to choose one subtree among the subtrees rooted at this node and then proceeds to the root of the chosen subtree. ChooseSubtree evaluates each subtree rooted at a node with respect to the new entry, and it chooses one subtree according to a set of heuristics.

If a new entry is to be inserted into a leaf node that is full, an overflow occurs. If during the insertion of the new entry, this is the first overflow at a given level in the tree, the algorithm RemoveTop is invoked; otherwise, the Split algorithm is invoked. RemoveTop³ removes p entries from a node and reinserts them. The entries to be removed are chosen according to heuristics. In the worst case, where all these entries are reinserted into the same node or they overflow some other node, the Split algorithm is used instead of RemoveTop. The split of a node can result in an overflow of a parent node. If this happens, the described procedure is repeated for the parent node.

The algorithm Split distributes the entries of an over-full node into two groups, forming two nodes. Because the number of possible distributions is exponential in the number of entries, only a subset of all possible distributions is considered by the Split algorithm. Again, the choice of the best distribution is made according to heuristics.

The \mathcal{R}^{ST} -tree also uses the overall procedure described above, but the heuristics that drive the decisions in the ChooseSubtree, Split, and RemoveTop algorithms are tailored to take into account the growth of bounding regions throughout the tree.

4.2 Heuristics and Time Parameterization

The \mathcal{R}^* -tree uses three basic heuristics in its insertion algorithm.

- H1. The volumes of bounding hyper-rectangles should be minimized.
- H2. The overlap (volume of intersection) among bounding rectangles should be minimized.
- H3. The margin of bounding rectangles should be minimized.

²It should be noted that the logical deletion presented in Section 2 is implemented in the index as a deletion of an old region followed by an insertion of a new one with a fixed TT^{\dagger} .

³This algorithm implements forced re-insertion, introduced in [1].

Recent analytical studies of the performance of R-trees validate these heuristics [15, 11, 18]. For a four-dimensional rectangle with side lengths given by (s_1, s_2, s_3, s_4) , the margin is given by the following sum, where $1 \leq i, j, k \leq 4$: $\sum_i s_i + \sum_{i < j} s_i s_j + \sum_{i < j < k} s_i s_j s_k$ [11].

In the R^{ST} -tree, the quantities of volume, overlap, and margin are functions of time, and the insertion algorithm should consider not only the current values of these, but also their future evolutions.

The evolution of the volume of a bounding region depends on the following four characteristics.

- i. The type of the region.
 1. Static regions (rectangles and stair-shapes).
 2. Growing rectangles.
 3. Growing stair-shapes.
- ii. The rate of growth of the region's volume.
- iii. The time period from now to the moment when the region starts to appear.
- iv. The volume of the region at the current time.

We say that an entry is of type t if the region in the entry is of type t , and a node is of type t if its bounding region is of type t .

To optimize an index for querying at all times after the current time, the characteristics of regions should be prioritized in the order given above. The types of the regions should be considered first. Highest priority should be given to static rectangles and static stair-shapes, independently of their volume because any growing rectangle or stair-shape will eventually outgrow a static rectangle or a static stair-shape. For the same reason, growing rectangles (which only grow in transaction time) should be given a higher priority than growing stair-shapes.

The second characteristic of a region is its rate of growth. Among two growing rectangles, the one that is narrower in the spatial dimensions or in valid time should be given highest priority because it grows by a smaller amount at each time unit. Similarly, a growing stair-shape with a larger projection onto the spatial dimensions will eventually build up more area than a stair-shape with a smaller projection.

Third, if two growing regions have the same type and growth rate, it may be possible to prioritize them according to the times when they appear. The later a region begins its existence, the more preferable it is.

Static regions are only compared using the fourth characteristic, exactly as is done in the R-tree.

Next, Heuristic H2, which minimizes overlap between regions, can also be extended in a similar way to address the growth of regions. This is done by considering the same four characteristics as above, but for the regions obtained by intersecting overlapping regions. Finally, the presented scheme can be applied similarly to the margin heuristic (H3).

It should also be observed that the presented prioritization scheme is based on the assumption that the index is queried for an indefinite time and that no further insertion or deletion operations are performed. In a realistic setting where the index is updated continuously a less strict prioritization of the four characteristics of regions may be desirable. For example, a very small growing rectangle is preferable over a very large static rectangle if the growing rectangle is updated and becomes static before it outgrows the static one.

The above prioritization scheme is implemented in the R^{ST} -tree in a relatively straightforward and flexible manner by introducing a *time parameter* p in the R^{ST} -tree's insertion algorithm, which then computes the areas of regions as of p time units in the future (for some p). If a sufficiently large time parameter is used when computing and comparing areas, we effectively obtain results similar to as with the above-described scheme. Relaxed prioritizations are achieved by using smaller time parameter values. What time-parameter values to use is investigated experimentally (Section 5.2) and depends mostly on the intensity of updates.

Experiments with the GR-tree show that its performance is substantially boosted when strict heuristics are used that look only at node types and group regions of the same type, trying single-mindedly to avoid introducing growing rectangles and, especially, growing stair-shapes. Experimental studies suggest that these heuristics not only do not work for spatiotemporal data, but even have a negative effect.

To understand why, consider the situation in a node split algorithm where there is a choice between having one “large” bounding, growing stair-shape and some static bounding region versus having two “smaller” bounding, growing stair-shapes. With only two temporal dimensions, the former choice is clearly preferable; the “large” stair-shape will most probably not be much larger than one of the two “smaller” stair-shapes, and these two overlap substantially. This is so because all growing stair-shapes follow the $VT^1 = \text{NOW}$ diagonal and have equal transaction-time end values. Introducing two spatial dimensions, the spatial extents of the two “smaller” stair-shapes may be non-overlapping and far apart in the spatial dimensions. Forcing them into one bounding stair-shape may produce a truly large growing stair-shape.

Summarizing, the time parameter is a simple and flexible way to extend the R^* -tree heuristics. Using the time parameter, the four characteristics of regions, be it bounding regions or the regions that occur as intersections between bounding regions, do not have to be inspected explicitly. The next sections explain in detail the workings of the ChooseSubtree, Split, and RemoveTop algorithms.

4.3 The ChooseSubtree Algorithm

The ChooseSubtree algorithm is used for deciding where to insert a new entry. The algorithm chooses one subtree among the subtrees rooted at a node and then repeats the procedure for that subtree until it reaches a leaf node. To optimize the overlap heuristic (H1), the R^* -tree’s ChooseSubtree algorithm considers the enlargement of the overlap between the bounding regions of the subtrees that would result from inserting the new entry in a subtree.

To determine this overlap enlargement, the overlaps between a subtree’s current minimum bounding region and the minimum bounding regions of all the other subtrees are determined. Then the subtree’s minimum bounding region is extended with the new entry, and the overlaps are determined anew. This is done for all subtrees, and ChooseSubtree proceeds with the subtree where including the new entry yields the *least overlap-area enlargement*. Ignoring the spatial dimensions, Figure 4, parts (a)–(c), give an example, where inclusion of a new entry in node 1 is chosen.

Ties are resolved by choosing the node whose minimum bounding region requires the *least area enlargement*, and further ties are resolved by choosing the node whose minimum bounding region has the *smallest area* with the new entry enclosed.⁴

The R^{ST} -tree’s ChooseSubtree algorithm performs a time-parameterized computation of overlap enlargement, area enlargement, and area. Figure 4, parts (d) and (e), show the result of applying ChooseSubtree with a later time parameter value. Now, node 2 is the best choice for the new entry, as it yields a smaller overlap area enlargement than inclusion in node 1, and its area enlargement is smaller than that of node 3.

4.4 The Split Algorithm

A node split algorithm can be characterized by the set of distributions of entries into two nodes that it considers, and by the criteria it employs for selecting one (the “best”) of these candidate distributions.

The subset of all possible distributions considered by the R^* -tree’s Split algorithm is selected as follows. Along each of the axes, entries of the over-full node are sorted according to their bottom and top values. Then, assuming two dimensions, for each of the four sortings, a total of $M - 2m + 2$ distributions are considered, where M is the maximum number of entries in the node, and m is the minimum allowed number of entries in the node. The i -th distribution is generated by assigning the first $m - 1 + i$ entries of the sorting to the first node and the rest to the other.

The best distribution is selected based on the three heuristics H1–H3, introduced in Section 4.2. For the pair of bounding rectangles resulting from a distribution of entries, we use *area-value* for the sum of their volumes, *overlap-area* for the volume of their intersection, and *margin-value* for the sum of their margins.

⁴For non-leaf nodes, overlap area enlargement is not considered—only area enlargement and area are considered.

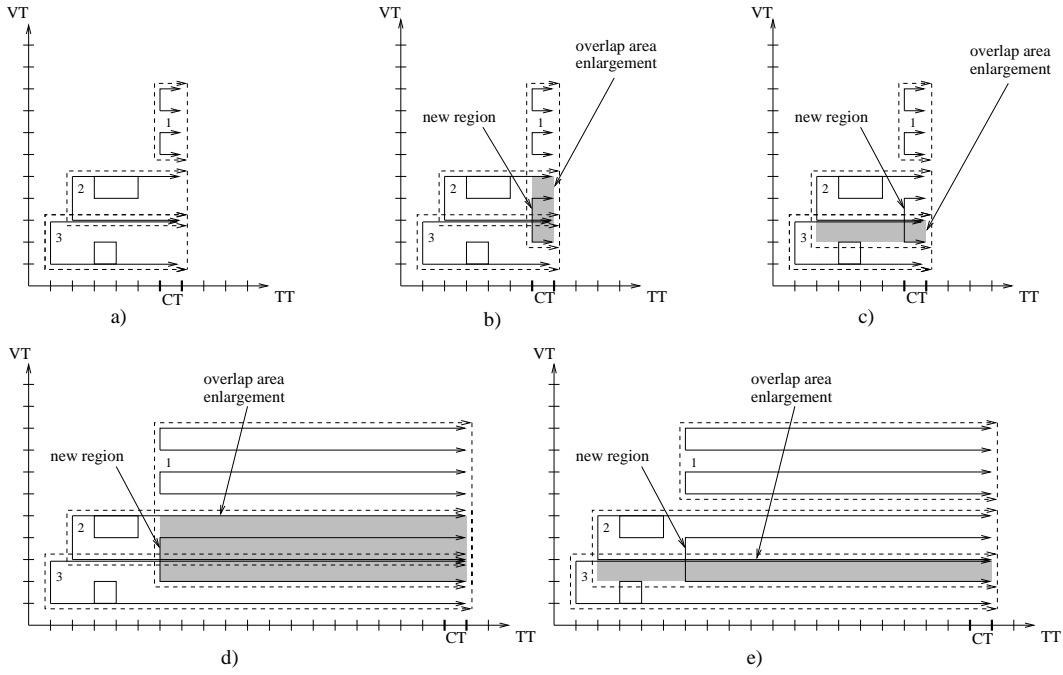


Figure 4: (a) Initial Bounding Regions, (b) Insertion of a New Entry into Node 1, (c) Insertion of a New Entry into Node 2, (d) Case (b) After a Period of Time, (e) Case (c) After a Period of Time

Using Heuristic H3 (minimum margin-value), one axis is selected. Then Heuristics H1 and H2 are used considering only the distributions along this axis.

The Original R^* -Tree Split Algorithm

S1 For each axis, sort the rectangles by their lower then by their upper value and determine all distributions as described above. Compute the sum of margin-values for all distributions for each axis.

S2 Let the axis with the minimum sum of margin-values be the split axis.

S3 Along the split axis, choose the distribution with the minimum overlap-value. Resolve ties by choosing the distribution with the minimum area-value.

The R^{ST} -tree makes this algorithm time-parameterized and introduces the generation of one more candidate distribution after Step S2. The new distribution is generated by trying to split entries of an over-full node so that the resulting bounding regions are of the best types possible. At the same time, the algorithm aims to not distribute entries of the same type into two different nodes. This procedure is described in detail next.

Each of the two nodes produced by a split may be bounded by a region of one of the three types defined in Section 4.2, leaving six possible pairs of types. These pairs are prioritized according to their “goodness,” as defined in Figure 5. Stated mathematically, a pair of bounding regions x_1 and x_2 is considered better than a pair of bounding regions y_1 and y_2 if: $(type(x_1) \neq type(y_1) \vee type(x_2) \neq type(y_2)) \wedge ((type(x_1) \leq type(y_1) \wedge type(x_2) \leq type(y_2)) \vee (type(x_1) < \max(type(y_1), type(y_2)) \wedge type(x_2) < \max(type(y_1), type(y_2))))$. The generation of the additional distribution is based on this ranking of pairs of node types.

Priority	Region 1	Region 2	Enabling conditions
1			$GR + GS = 0$
2			$GS = 0$ and $0 < GR \leq k$
3			$GS = 0$ and $GR > k$
4			$GS > 0$ and $GR + GS \leq k$
5			$0 < GS \leq k$ and $GR + GS > k$
6			$GS > k$

Type 1
 Type 2
 Type 3

GR: Number of growing rectangles
GS: Number of growing stair-shapes
 $k = M + 1 - m$

Figure 5: Pairs of Bounding-Region Types and Enabling Conditions

Generation of the Additional Candidate Distribution

S2.1 Among the six pairs of types of bounding regions, select the pair (t_1, t_2) such that, according to the conditions given in Figure 5, (a) it is possible to achieve this pair of bounding-region types when dividing the entries of the over-full node into two nodes, and (b) no other pair with a higher priority can be achieved. Let $t_1 \leq t_2$, and name the node bounded with a region of type t_1 and t_2 as N_1 and N_2 , respectively. Let S contain all entries of the over-full node.

S2.2 Move to N_2 all entries from S that cannot be put into N_1 because of the type of its bounding region.

S2.3 Let S_t denote all entries from S of type t . If $|S_1| \leq M - m + 1$, move S_1 into N_1 . Next, if $|S_2| + |N_1| \leq M - m + 1$, move S_2 into N_1 .

S2.4 If $|S| = 0$, goto **S3**; (N_1, N_2) is a new distribution.

S2.5 If $|N_1| = 0 \wedge |N_2| = 0$, goto **S3**; no new distribution was generated.

S2.6 If $|N_1| = 0$, pick a "seed" entry e from S for Guttman's quadratic *Distribute* algorithm [7] such that its inclusion into N_2 would enlarge that node's minimum bounding region the most. Put e into N_1 . Goto **S2.8**.

S2.7 If $|N_2| = 0$, pick a seed entry e from S and put it into N_2 .

S2.8 Apply Guttman's quadratic *Distribute* algorithm.

The above algorithm uses a time-parameterized version of Guttman's quadratic *Distribute* algorithm. Section 5 studies the impact on query performance of considering this additional distribution.

4.5 The Remove-Top Algorithm

The original R^* -tree RemoveTop algorithm sorts the entries of an over-full node by the distances of their centers from the center of the minimum bounding rectangle of the over-full node and then removes and reinserts a certain percentage of the entries with the largest distances.

A time-parameterized version of this algorithm could be used for the R^{ST} -tree, but performance experiments show that a RemoveTop algorithm based on the heuristic of volumes works better for spatiotemporal data. The R^{ST} -tree employs an algorithm of quadratic complexity in the number of entries in a node. This algorithm repeatedly removes the entry that, when removed, shrinks the time-parameterized volume of the minimum bounding region the most.

4.6 Prioritizing Space Versus Time

The insertion algorithm described in the previous sections aims to produce a tree that yields high performance of spatiotemporal intersection queries. The heuristics used so far have been based on the assumption that intersection queries are square on average, i.e., all the dimensions (temporal and spatial) are constrained by intervals of approximately the same length.

Due to the quite different semantics of the dimensions involved, this may not always be a good assumption. In some applications, most queries may be much more restrictive on the spatial dimensions than on the

temporal dimensions. For example, queries in a cadastral system may retrieve the current knowledge of the full history of ownership of some piece of land. In other applications, queries may be most restrictive on the temporal dimensions. Specifically, timeslice queries, which specify time points in the temporal dimensions, have very natural semantics and are often important. For example, a query in a demographic database may retrieve the population for a county, as it was two years ago, as currently best known.

Non-square queries may be due just to the different units of measurement used for the spatial and temporal dimensions. For example, if the granularity of time is one second and the granularity of space is one meter, and most queries are formulated in days and meters, the queries will be long in the temporal dimension. The solution to the problem of non-square queries also addresses the issues of different semantics and granularities of the spatial and temporal dimensions.

In the extreme, if a substantial number of queries offer no restrictions on either space or time, additional, separate temporal and spatial indices may be introduced for these queries. Whether the likely increase in query performance obtained from using these specialized indices justifies the extra cost of update and the extra storage space occupied will depend on the specific application requirements.

In order to obtain a versatile spatiotemporal index that may serve well the full spectrum queries—ranging from having only spatial constraints to having only temporal constraints—it is desirable to introduce a mechanism into the R^{ST} -tree that allows it to be tuned to support better either spatially or temporally restrictive queries.

In any R-tree-based index, one dimension can be prioritized over the others by intentionally favoring minimum bounding rectangles that are narrow in this dimension and long in the other dimensions. In Figure 6, a two-dimensional space is considered. The two sets of minimum bounding rectangles cover the same areas and do not overlap. Scenario (b) favors queries restrictive in the x dimension and not in the y dimension.

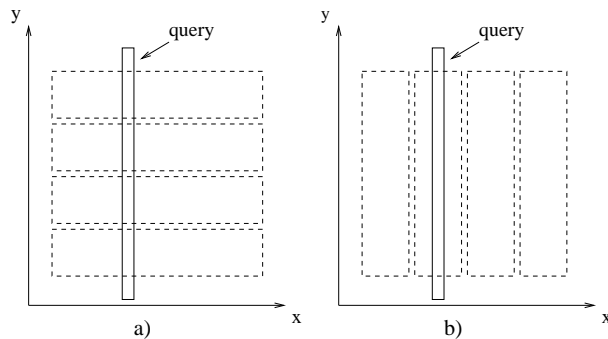


Figure 6: Different Geometries of Minimum Bounding Rectangles

Analytical studies also offer some insight into the effect of the shapes of the bounding regions on query performance [11, 15]. For a four-dimensional query with side lengths (q_1, q_2, q_3, q_4) , the estimated number of page accesses is proportional to $\sum_m (\prod_{i=1}^4 (s_{m,i} + q_i))$, where $s_{m,i}$ is the side length of bounding rectangle s_m in the i -th dimension and m ranges over all bounding rectangles of the tree. Expanding the product in this formula, it can be seen that the notion of margin, as defined in Section 4.2, plays a significant role. The formula suggests that non-rectangular queries may be supported better by weighting the components of the margin expression. However, in R^* -tree-based indices, the margin-value is used only as a secondary heuristic and only in the node split algorithm, and so a prioritization scheme based on margin will not be effective.

We propose a simple way to prioritize the dimensions in an R-tree-based index, which works with the existing tree algorithms. For each n dimensional rectangle, instead of considering the extents $(\Delta x_1, \Delta x_2,$

$\dots, \Delta x_n$), weighted extents $((\Delta x_1)^{\alpha_1}, (\Delta x_2)^{\alpha_2}, \dots, (\Delta x_n)^{\alpha_n})$ are used. If all α_i are equal to one, none of the dimensions are prioritized. The priority of dimension i is increased by setting α_i to a value greater than 1, and the priority of dimension i lowered by setting α_i to a value smaller than 1. Setting α_i to 0 makes the algorithms disregard the dimension.

Following this scheme the R^{ST} -tree uses a single parameter $\alpha \in [-1..1]$. The volume of a four-dimensional region r is then computed as follows.

$$volume(r) = \begin{cases} bitemporal_area(r)^{1+\alpha} \cdot spatial_area(r) & \text{if } \alpha \leq 0 \\ bitemporal_area(r) \cdot spatial_area(r)^{1-\alpha} & \text{otherwise} \end{cases}$$

where *bitemporal_area* is the area of the region’s time-parameterized bitemporal extent and *spatial area* is the area of its spatial extent.

The margin-value computation in the formula presented in Section 4.2 is updated similarly, so that the temporal and spatial extents are weighted with $1 - \alpha$ or $1 + \alpha$. It follows that α values less than 0 favor spatially selective queries, and α values greater than 0 favor temporally restrictive queries. As mentioned above, α values of -1 and 1 turn the R^{ST} -tree into a spatial and a temporal index, respectively. Section 5 investigates the effects of using different α values.

5 Performance Studies

Section 5.1 details the experimental setup and data generation. Then, Section 5.2 describes the effect of the introduction of the time parameter and the additional split distribution in the tree, and Section 5.3 studies the use of the α -parameter. Finally, Section 5.4 compares the R^{ST} -tree with its only competitor, the minimally adapted R^* -tree.

5.1 Experimental Setup and Data Generation

The studies are based on an implementation of the R^{ST} -tree based on the Generalized Search Tree Package, GiST [8]. The page size (and tree node size) is set to 4k bytes, and a page buffer of 200k bytes, i.e., 50 pages, is used [12] where the root of a tree is pinned and the least-recently-used page replacement policy is employed. Nodes changed during an index operation are marked as “dirty” in the buffer and are written to disk at the end of the operation or when they otherwise have to be removed from the buffer.

The performance studies are based on so-called workloads that intermix queries and update operations on the index, thus simulating index usage across a period of time. The remainder of this section describes the generation of workloads.

The workloads simulate the evolving spatiotemporal aspects of a set of objects. Each change to an object is characterized by three parameters: the time duration for which the previous spatial value of the object was valid (*VL*), the displacement of the center of the spatial value relative to its previous center, and the change of the area of the spatial value relative to its previous area. These parameters account for the spatial and valid-time aspects of objects and are adapted from the GSTD algorithm [21]. These spatial and valid-time values are augmented with transaction times when they are stored in the database, and GSTD is extended to make it possible to simulate the storage of these changing values of objects in the database. Insertions may be retroactive or predictive [10]. A retroactive insertion occurs when the valid-time begin of an inserted record is less than the time of insertion (transaction-time begin), and a predictive insertion occurs when the valid-time begin exceeds the time of insertion.

A *history* captures the current knowledge about the (spatial) evolution of a single object. An object’s history is composed of a number of triples of a spatial value, a valid-time interval, and a transaction-time

interval; the valid-time intervals meet, yielding at most one spatial value at each point in time, and for the transaction-time intervals, $TT^{\perp} = UC$.

The spatial value associated with the first valid-time interval in a history is taken from a set of spatial values generated using the A La Carte spatial data generator [5]. This set contains 2000 rectangles generated in an area spanning 20000 units in the x and y dimensions. The spatial density is set to 1, and the distribution of the bottom-left corners of the rectangles is uniform. Point data is obtained by taking the bottom-left corners of the rectangles.

The spatial value associated with any other valid-time interval in a history is generated from the spatial value associated with its predecessor valid-time interval, by extending this value by Δx and Δy , if it is a rectangle, and by adding $(\Delta x, \Delta y)$ to it if it is a point. The values of Δx and Δy are uniformly distributed over $[-SpChange, SpChange]$. This enables the modeling of shrinking/expanding rectangles and moving points.

The history of an object is *active* if it has a spatial extent with a valid time that includes the current time. For active histories, the interval having the largest valid-time begin is made now-relative. Histories that are not active contain no now-relative valid-time intervals. Valid-time interval lengths are uniformly distributed between 0 and VL . For now-relative intervals, the Δ -offset is normally distributed with mean 0 and deviation 100, and its valid-time begin is normally distributed with mean equal to the insertion time (transaction-time begin) and deviation $Dev = 500$.

The transaction-time intervals in the triples are obtained by simulating database modifications. To accomplish this, the workload generator maintains a list of the currently active histories. Each history has associated the time when the next insertion for the history will occur, along with the pair of the spatial value and valid-time interval to be inserted. The list is ordered on the insertion times.

An insertion is done by updating the valid-time end of the most recent entry in the history from being now-relative to being static, so that it meets the beginning of the new entry. This is accomplished by logically deleting the original now-relative data, inserting the data with the new static interval, and finally inserting the new, now-relative data.

After an insertion, a new, planned entry is generated together with its insertion time. This new entry's valid-time begin is normally distributed with deviation Dev and mean equal to its planned insertion time, which, in turn, is generated so that after inserting the new entry, the length of a previous interval will not exceed VL . This way, parameter VL captures how often objects change, and parameter Dev controls the correlation between transaction and valid time.

In addition to generating operations that insert new entries into active histories, index operations corresponding to one of three actions are generated at each time point.

1. Introducing a new object, thus starting a new history.
2. Ending an active history.
3. Updating a history.

In our experiments, workloads start at time $TStart = 2000$. For the first 2000 time units, at each time point, we introduce a new object. After this initialization, a history may be started, ended, and updated with probability 0.1, 0.1, and 0.8, respectively, at each time point.

A history is started by generating a new pair of a spatial value and a valid-time interval along with an insertion time. A history is ended by updating the valid-time end of the last interval from being now-relative to being static. A history is updated by either inserting a new valid-time interval somewhere between the beginning and the current end of the history, or deleting some existing valid-time interval from the history. The spatial part of an inserted interval is obtained by changing the spatial part of the interval in the history just before it using Δx and Δy . Special action is necessary to ensure that the valid-time intervals in a history continue to meet after an update. Following the insertion of a new interval, all intervals fully covered by it

are logically deleted, and partially covered intervals are updated. Following the deletion of an interval, the valid-time end of the previous interval, if it exists, is set to the valid-time end of the deleted interval.

Each workload used in an experiment contains 200,000 logical deletions and insertions, so that after running the workload, 200,000 data items are stored in the index. Depending on parameter VL , this results in workloads ranging approximately from 8000 to 34,000 time units in length. A workload also contains 10 queries for each 200 insertions or deletions. The queries are intersection queries represented as four-dimensional rectangles with square spatial and temporal projections. The valid-time begin of a query is uniformly distributed between 0 and the largest valid-time begin of an entry in the index. For half of the queries, the transaction-time end is set to the current time; for the others, it is uniformly distributed between $TStart$ and the current time. The four-dimensional volumes of queries are uniformly distributed between 1 and $4 \cdot 10^{12}$. The volumes are divided into bitemporal and spatial areas according to the ratio $T : S$, which ranges from 1 : 1000 for temporally restrictive queries to 1000 : 1 for spatially restrictive queries (see Table 3).

In addition, we experimented with bitemporal and spatial point queries as extreme cases of temporally and spatially restrictive queries. For these queries, the four-dimensional volume of a query varied from 1 to 10^8 . Another query type that is very likely to be seen in real-world applications is the transaction timeslice, which retrieves the history of some region in space as current in the database at some specific time point. In our experiments we constrain the region of space to a point.

The parameters that control the generation of workloads are given in Table 3. Standard values, used if a parameter was not varied in an experiment, are given in bold-face.

Table 3: Workload Parameters

Parameter	Description	Values Used
$SpChange$	Maximum value used for Δx and Δy .	0, 40, 200 , 1000, 5000
VL	Maximum valid-time interval length.	250, 500, 1000 , 2000, 4000
$T : S$	Ratio between the temporal and spatial parts of a query.	1 : 1000, 1 : 10, 1 : 1 , 10 : 1, 1000 : 1

5.2 Effects of the Time-Parameterization and the Additional Split Distribution

Compared to the R^* -tree, the R^{ST} -tree includes a time parameter and considers one additional distribution of the entries in an over-full node, when it is split. This section studies the effects of these additions.

It was pointed out in Section 4.2 that using a moderately large time-parameter value might yield a well performing tree in a realistic setting. In addition, the characteristics of the data may affect the utility of different time parameters.

A set of experiments was performed to determine the performance resulting from using different time parameters for different workloads. The results were that the performance is not very sensitive to differences in the time parameter value, but that a range of values yield quite similar performance; only the extreme values of 0 and “very large” tend to decrease the query performance (up to 100% for some workloads), especially for rectangle data (as opposed to point data). As a result, we fixed the time-parameter value at 5000 in all other experiments.

The experiments also show that what is a good time-parameter value is somewhat dependent on the average number of operations per time unit, which, in turn, depends on the VL parameter. This should be expected. The time-parameter value, used when splitting a node, should intuitively be dependent on for how long the resulting nodes will exist before being deallocated or reorganized by other splits or multiple insertions and deletions. The smaller the number of operations per time unit, the longer an average node survives, which tends to favor a large time-parameter value.

Switching the attention to the additional distribution considered in the split algorithm, a set of experiments aimed to determine how often this distribution is actually chosen by the split algorithm. If it is chosen quite often, it is interesting to see how its availability improves query performance. The experiments used workloads with different types of queries along with settings for the α parameter that favor the particular type of queries (see Section 5.3). Each workload was run with and without the additional distribution available.




Figure 7: (a) Percentage of Splits Choosing the Additional Distribution and (b) the Resulting Query Performance Improvement (in %) for Temporally Restrictive Queries

For α values favoring queries with $T : S = 1 : 1$ and spatially restrictive queries, i.e., $\alpha \leq 0$, the additional distribution is chosen very infrequently, and so its presence is insignificant. For the workloads with temporally restrictive queries, the additional distribution is chosen in up to 19% of all splits, in turn improving the query performance by 10% to 25% (see Figure 7).

5.3 Space-Time Prioritization

Section 4.6 presented an α -parameter, the use of which aimed at tailoring the index to different types of queries. A set of experiments was performed in order to investigate the effect of this parameter on the performance of different types of queries. Eight workloads with different types of queries were run using different α values.

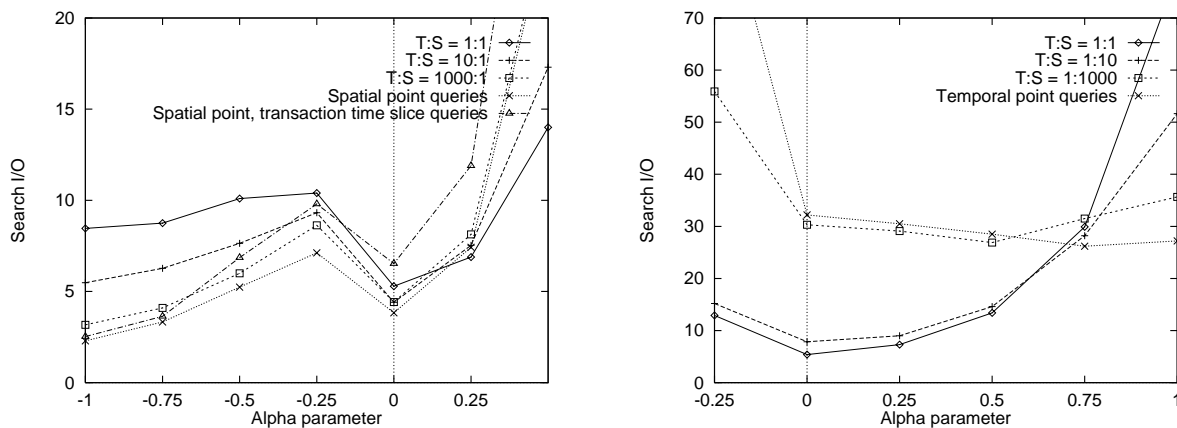


Figure 8: Query Performance for Varying α -Value, for Rectangle Data

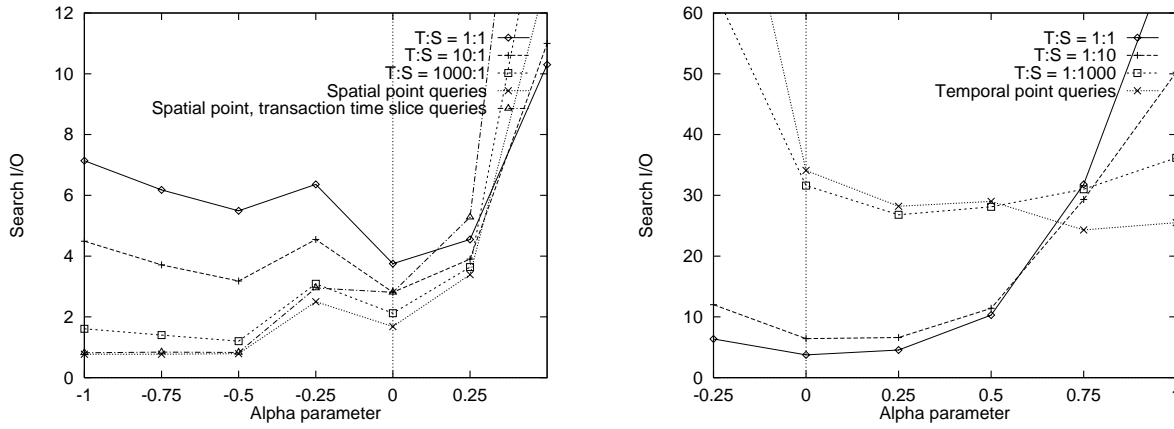


Figure 9: Query Performance for Varying α -Value, for Point Data

As shown in Figures 8 and 9, $\alpha = 0$ is best for workloads with neutral queries ($T : S = 1 : 1$). This is as expected.

Interestingly, for spatially restrictive queries, α values of -1.0 and 0 , for rectangle data, and -0.5 and 0 , for point data, outperform other values almost independently from $T : S$ ratio in the queries. For $\alpha < 0$, most of the bounding regions in a tree are growing stair-shapes. Only when $\alpha = 0$, the insertion algorithm starts to group growing rectangles, growing stair-shapes, and static regions into different nodes and this pays off even for spatially restrictive queries. Nevertheless, values larger than 0 penalize the spatially restrictive queries.

For temporally restrictive queries, depending on the $T : S$ ratio, the optimal α value ranges from 0 to 0.75 . The asymmetrical nature of these “spatial” and “temporal” results is as expected and reflects the very different geometries of bitemporal versus spatial regions. In contrast to average spatial regions, bitemporal regions tend to be quite long in transaction time. In addition, the distribution of regions in the bitemporal plane is different from their distribution in the spatial plane.

When the mix of queries posed against an \mathbb{R}^{ST} -tree does not exhibit a strong temporal restrictiveness or when it is not known which mix to expect, a neutral α value of 0 is preferable. A value of ca. 0.25 is preferable for temporally restrictive queries. For temporal point queries that are non-restrictive in the spatial dimensions, a value of ca. 0.75 shows the best performance.

5.4 Comparison with Straightforward \mathbb{R}^* -Tree-Based Indexing

Replacing all occurrences of UC and $\text{NOW} + \Delta$ in the data by the maximum time values for each dimension renders the data static and enables a standard four-dimensional \mathbb{R}^* -tree, followed by a filtering step, to index the spatiotemporal data considered here. The filtering step reverts to the correct temporal extent and re-applies the query predicate, thus eliminating false drops.

Figure 10 compares the \mathbb{R}^{ST} -tree to this straightforward solution, without considering the filtering step. In our experiments, which cover both rectangular data and point data, the \mathbb{R}^* -tree generally requires from approximately 50% to 100% more I/O operations than the \mathbb{R}^{ST} -tree for neutral queries and temporally restrictive queries (not shown in Figure 10); for spatially restrictive queries, the \mathbb{R}^* -tree generally requires from approximately 50% to 300% more I/O operations. These numbers are quite significant, especially given that the workloads used in these experiments are quite favorable towards the \mathbb{R}^* -tree.

The difference in performance among the two indices is highly dependent on the distribution of the queries. As more and more queries reach above the $\text{VT} = \text{TT}$ primary diagonal in the bitemporal plane, the more false drops will occur in the result obtained from the \mathbb{R}^* -tree; retrieving these incurs extra I/O

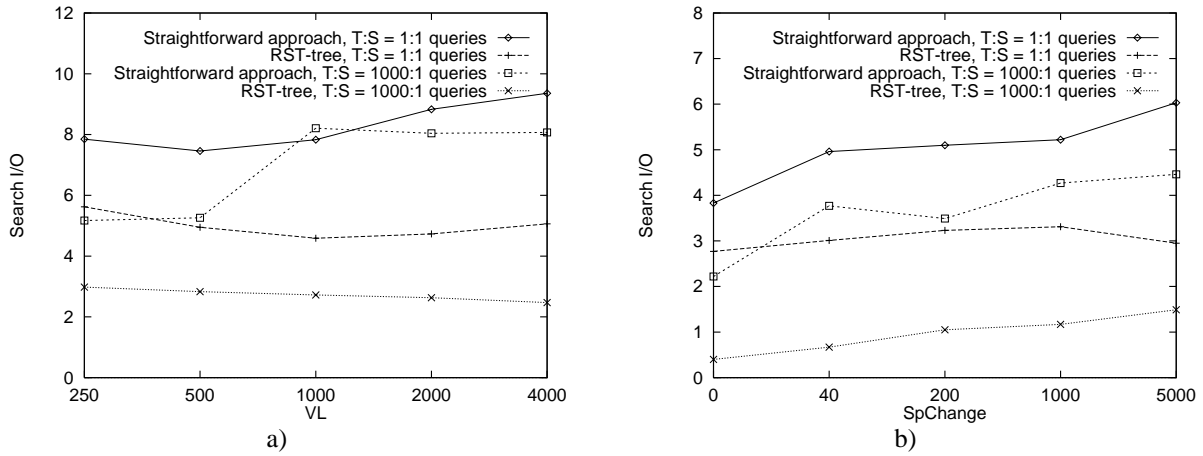


Figure 10: Query Performance for (a) Rectangle Data with Varying VL and (b) Point Data with Varying $SpChange$

not incurred by the R^{ST} -tree. The workloads used in the experiments reported here are very favorable towards the straightforward approach in this regard. For example, the workload used for the data points in Figure 10(a) for $VL = 4000$ produced only ca. 9% false drops, but still incurred 100% more I/O operations.

While not illustrated here, the performance of the R^* -tree is also highly dependent on the number of now-relative spatiotemporal extents; the more of these, the worse the performance. The R^{ST} -tree does not experience a similar effect.

6 Conclusions and Research Directions

This paper addresses the emerging need for efficient support for spatiotemporal data in databases by proposing a new, versatile spatiotemporal indexing technique, termed the R^{ST} -tree, for indexing discretely changing spatial extents—points or rectangles. The R^{ST} -tree differs from all previously proposed spatiotemporal indices in that it provides support for both transaction time and valid time, termed bitemporal support, and it accommodates general, now-relative transaction-time and valid-time intervals. The R^{ST} -tree is based on the R^* -tree and borrows from the GR-tree, which adapts the R^* -tree to accommodate a less general class of bitemporal data.

The paper details how to address the special features of the temporal dimensions by using novel bounding regions in the index that grow continuously with the advancement of time and that also take into account the spatial dimensions. Most prominently, these new bounding regions are accompanied by a new set of insertion algorithms. In addition, the tree offers the ability to weigh the spatial and temporal aspects of data, so that either queries that are very selective in the spatial or temporal dimensions are supported more efficiently. The GiST package is used in the implementation of the index. Performance comparisons with a straightforward, R^* -tree-derived spatiotemporal index demonstrate that the new bounding regions and algorithms significantly increase performance. It is also demonstrated that the mechanism for assigning weights makes it possible to better support spatially and temporally selective queries.

In future work, we plan to further investigate support for different types of spatiotemporal queries. Experiments show that an α -parameter of -1.0 seems to provide the best performance for spatially selective queries. A spatial index may be competitive in such cases. In general, given a spatiotemporal query, we would like to be able to determine which existing spatial, temporal, or spatio-temporal index is the better one and to know how the new index compares to the more specialized, existing indices. In addition, we hope to be able to obtain and experiment with real or partly real data, which may yield additional insights.

Acknowledgements

This research was supported in part by grants from the European Commission, the Danish Technical Research Council, the Danish Centre for IT Research, and the Nykredit corporation.

References

- [1] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger. The R*-tree: An Efficient and Robust Access Method for Points and Rectangles. *ACM SIGMOD'90*, pp. 322–331.
- [2] R. Bliujūtė, C. S. Jensen, S. Šaltenis, and G. Slivinskas. R-tree Based Indexing of Now-Relative Bitemporal Data. *VLDB'98*, pp. 345–356.
- [3] R. Bliujūtė, C. S. Jensen, S. Šaltenis, and G. Slivinskas. Light-Weight Indexing of General Bitemporal Data. *TIMECENTER TR-30*, 1998.
- [4] J. Clifford et al. On the Semantics of "Now" in Databases. *TODS*, 22(2):171–214, 1997.
- [5] O. Günther et al. Benchmarking Spatial Joins A La Carte. *SSDBM'98*, pp. 32–41.
- [6] R. H. Güting et al. A Foundation for Representing and Querying Moving Objects. *Informatik-Report 238*, FernUniversität Hagen, 1998.
- [7] A. Guttman. R-trees: A Dynamic Index Structure for Spatial Searching. *ACM SIGMOD'84*, pp. 47–57.
- [8] J. M. Hellerstein, J. F. Naughton, and A. Pfeffer. Generalized Search Trees for Database Systems. *VLDB'95*, pp. 562–573.
- [9] C. S. Jensen and R. Snodgrass. Semantics of Time-Varying Information. *Information Systems*, 21(4):311–352 (1996).
- [10] C. S. Jensen and R. Snodgrass. Temporal Specialization and Generalization. *IEEE TKDE*, 6(6):954–974 (1994).
- [11] I. Kamel and C. Faloutsos. On Packing R-trees. *CIKM'93*, pp. 490–499.
- [12] S. T. Leutenegger and M. A. Lopez. The Effect of Buffering on the Performance of R-Trees. *ICDE'98*, pp. 164–171.
- [13] M. A. Nascimento, J. R. O. Silva, and Y. Theodoridis. Access Structures for Moving Points. *TIME-CENTER TR-33*, 1998.
- [14] M. A. Nascimento and J. R. O. Silva. Towards Historical R-trees. In *ACM SAC'98*, pp. 235–240.
- [15] B.-U. Pagel, H.-W. Six, H. Toben, and P. Widmayer. Towards an Analysis of Range Query Performance in Spatial Data Structures. *PODS'93*, pp. 214–221.
- [16] R. T. Snodgrass and I. Ahn. A Taxonomy of Time in Databases. *ACM SIGMOD'85*, pp. 236–246.
- [17] R. T. Snodgrass. The Temporal Query Language TQuel. *ACM TODS*, 12(2):247–298, 1987.
- [18] Y. Theodoridis and T. Sellis. A Model for the Prediction of R-tree Performance. *PODS'96*, pp. 161–171.

- [19] Y. Theodoridis, T. Sellis, A. N. Papadopoulos, and Y. Manolopoulos. Specifications for Efficient Indexing in Spatiotemporal Databases. *SSDBM'98*, pp. 123–132.
- [20] Y. Theodoridis, M. Vazirgiannis, and T. Sellis. Spatio-temporal Indexing for Large Multimedia Applications. *Proc. 3rd IEEE Conf. on Multimedia Computing and Systems*, pp. 441–448, 1998.
- [21] Y. Theodoridis, J. R. O. Silva, and M. A. Nascimento. On the Generation of Spatiotemporal Datasets. *Proc. 6th Intl. Symposium on Spatial Databases*, pp. 147–164, 1999.
- [22] T. Tzouramanis, M. Vassilakopoulos, and Y. Manolopoulos. Overlapping Linear Quadtrees: A Spatio-Temporal Access Method. *Proc. 6th ACM-GIS Workshop*, pp. 1–7, 1998.
- [23] X. Xu, J. Han, and W. Lu. RT-tree: An Improved R-tree Index Structure for Spatiotemporal Databases. *Proc. 4th Intl. Symposium on Spatial Data Handling*, pp. 1040–1049, 1990.