

## Light-Weight Indexing of General Bitemporal Data

Rasa Bliujūtė      Christian S. Jensen    Simonas Šaltenis    Giedrius Slivinskas  
Nykredit Data A/S      Department of Computer Science, Aalborg University  
rasa@nykredit.dk      {csj, simas, giedrius}@cs.auc.dk

### Abstract

*Most data managed by existing, real-world database applications is time referenced. Often, two temporal aspects of data are of interest, namely valid time, when data is true in the mini-world, and transaction time, when data is current in the database, resulting in so-called bitemporal data. Like spatial data, bitemporal data thus has associated two-dimensional regions. Such data is in part naturally now-relative: some data is true until the current time, and some data is part of the current database state. So, unlike for spatial data, bitemporal data regions may grow continuously. Existing indices, e.g.,  $B^+$ - and  $R$ -trees, typically do not contend well with even small amounts of now-relative data.*

*In contrast, the 4-R index presented in the paper is capable of indexing general bitemporal data efficiently. The different kinds of growing data regions are transformed into stationary regions, which are then indexed by  $R^*$ -trees. Queries are also transformed to counter the data transformations, yielding a technique with perfect precision and recall. Performance studies indicate that the technique is competitive with the best existing index; and unlike this existing index, the new technique does not require extension of the DBMS kernel.*

### 1. Introduction

Two temporal aspects of data are fundamental—valid time and transaction time [10]. The valid time of a database tuple is the time when the tuple is true in the modeled reality, the mini-world. A tuple's transaction time is the time during which the tuple is current in the database. These temporal aspects of data are essential in a wide range of existing, real-world applications, including medical, financial, travel, and multimedia applications. For example, real-world (non-textbook!) banking databases do record at which times the (past and present) balances of an account apply. Transaction time is essential in applications where trace-ability or accountability are important. Data with both valid and transaction time associated is termed bitemporal.

Although several dozen data models and temporal query languages have been proposed and although the new SQL standard has an associated Temporal Part [21], the major DBMSs provide little support for temporal data management. As a result, each new database application is designed to solve anew and in an ad-hoc fashion temporal data management problems that could be solved by the DBMS. This paper proposes an efficient bitemporal indexing technique that can be implemented as a layer on top of an existing DBMS, by an independent third-party developer. In this sense, the index is light-weight.

There has recently been proposed a number of indices for temporal data (e.g., [3, 12, 13, 15, 16]; see [23] for a survey). The majority are for transaction-time data, and only few support valid-time data. Significantly less research has been done on creating indices for bitemporal data. Spatial indices are obvious candidates for indexing bitemporal data, due to the similarities between bitemporal and spatial data: the combined valid and transaction time of a tuple can be treated as a region in two-dimensional space. Several existing proposals [3, 12, 13] are based on the  $R^*$ -tree [1].

The bitemporal indices generally fall short in efficiently supporting now-relative data [6], data for which the end of the valid time or/and transaction time tracks the progressing current time. Now-relative data occurs naturally in most real-world databases. For example, consider the recording of a new employee in a company's database. The time when the employee starts working (valid-time interval begin) is known, but it is unknown when the employee will leave. This is captured by letting the valid-time end extend to the progressing current time. The same applies to transaction time. The transaction-time interval begin of a tuple is the time when it is inserted into the database. Since we do not know when the tuple will stop being current, the transaction-time end extends to the current time. Two of the existing indices, the 2-R index and the Bitemporal  $R$ -tree [13], efficiently support now-relative transaction time, but not now-relative valid time.

Only the GR-tree [3] supports both now-relative valid time and now-relative transaction time, and thus general bitemporal data, efficiently. But no existing DBMSs sup-

port this index, and adding it to a DBMS such as DB2, Oracle, or Sybase would require an extension of the DBMS’s kernel. The extensibility features of these systems do not offer adequate support for such extensions. Perhaps most prominently, obtaining adequate concurrency control remains problematic [5]. The indexing technique proposed in this paper achieves performance that is comparable to the GR-tree’s, and it may be implemented on-top of any DBMS that supports R-trees, as does, e.g., Informix.

The reliance of R-trees on (minimum) bounding rectangles does not combine well with growing regions. This shortcoming is overcome by applying transformations to the growing bitemporal data regions that render them stationary and thus amenable to R-tree indexing. Growing regions come in four kinds, and each kind of regions has its own transformation and is indexed with its own R\*-tree. The resulting index is termed the 4-R index. Queries on the index are transformed into four separate queries, one for each tree. The approach may be seen as a generalization of the approach underlying the 2-R index [12] or as an “extreme” case of the GR-tree.

To provide a proper understanding of the properties and behavior of the 4-R index, a substantial portion of the paper covers results of quite extensive performance studies of the index, thus providing detailed insight into the specifics of the index and its constituent trees.

The presentation is structured as follows. First, Section 2 briefly describes important concepts and introduces two-dimensional bitemporal regions. Section 3 surveys the existing work related to the indexing of bitemporal data. The 4-R index, including data and query transformations, is described in Section 4. Section 5 presents performance studies. The final section offers conclusions and directions for future work.

## 2. Bitemporal Data and Its Representation

As a foundation for understanding the challenges of indexing bitemporal data, this section first describes in more detail the nature of bitemporal data, then characterizes the different kinds of two-dimensional bitemporal data regions.

As mentioned in the previous section, valid time captures when a tuple is true in the modeled reality, and transaction time captures when a tuple is current in the database [10]. These two temporal aspects are orthogonal in that each could be recorded independently, and each has specific properties associated with it. The valid time of a tuple can be in the past or in the future and can be changed freely. In contrast, the transaction time of a tuple cannot extend beyond the current time and cannot be changed.

TQuel’s four-timestamp format [19] (4TS) is the most popular format for bitemporal data representation. With this format, each tuple has a number of non-temporal attributes

and four time attributes:  $VT^+$  and  $VT^-$ —the times when the tuple’s information became and ceased to be true in the modeled reality;  $TT^+$  and  $TT^-$ —the times when the tuple became and ceased to be current in the database.

A tuple is now-relative if its information is valid until the current time or if the tuple is part of the current database state. This is represented in the 4TS format by the use of variables, which denote the current time, for the time attributes  $VT^-$  and  $TT^-$  [6]. The variable UC (“until changed”) is used for  $TT^-$ , and the variable NOW is used for  $VT^-$ . Table 1 exemplifies bitemporal data. The time granularity is a month, and the current time is assumed to be 9/97.

	Employee	Department	$TT^+$	$TT^-$	$VT^+$	$VT^-$
1	John	Advertising	4/97	UC	3/97	5/97
2	Tom	Management	3/97	7/97	6/97	8/97
3	Jane	Sales	5/97	UC	5/97	NOW
4	Julie	Sales	3/97	7/97	3/97	NOW
5	Julie	Sales	8/97	UC	3/97	7/97
6	Michelle	Management	5/97	UC	3/97	NOW

**Table 1. The EmpDep Relation**


Tuple 1 records that the information “John works in Advertising” 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 works in Sales” from 5/97 until the the current time, that we recorded this fact on 5/97, and that this remains part of the current database state.

Specific constraints apply to insertions, deletions, and modifications of tuples. When inserting a new tuple, the constraints  $VT^+ \leq VT^-$  and  $VT^+ \leq \text{‘current time’}$  if  $VT^-$  is equal to NOW apply to valid time; and the constraints  $TT^+ = \text{‘current time’}$  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 ‘current time’ – 1<sup>1</sup>, making the tuple not current anymore (e.g., Tuple 2); tuples are not physically deleted. A modification is modeled as a deletion followed by an insertion (e.g., an update led to Tuple 4 and Tuple 5).

The temporal aspect of a tuple can be represented graphically by a two-dimensional (“bitemporal”) region in the space spanned by valid and transaction time [10]. Cases 1–4 in Figure 1 illustrate the *bitemporal regions* of Tuples 1–4, respectively.

A now-relative transaction-time interval yields a rectangle that “grows” in the transaction-time direction as time passes (Tuple 1, Case 1). Having both transaction- and valid-time intervals being now-relative yields a stair-shaped region growing in both transaction time and valid time as time passes (Tuple 3, Case 3). If, at some time, a tuple stops

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



**Figure 1. Bitemporal Regions**

being current, the bitemporal region stops growing (Tuples 2 and 4; Cases 2 and 4).

Information may be recorded in the database after it becomes true in the modeled reality. In such situations, if both the transaction- and valid-time intervals are now-relative (Tuple 6), the bitemporal region would be a growing stair-shape with a high first step. It is also possible to record information in the database before it becomes true in the modeled reality. In this case, the valid-time end must be a ground value (Tuple 2); otherwise, the valid-time end, which would extend to the current time, would initially be smaller than the valid-time start, violating the second insertion constraint of the valid time.

Stated generally, we obtain four combinations of time attributes for which the bitemporal regions are qualitatively different, as illustrated in Figures 1 and 2 where ‘tt1’, ‘tt2’, ‘vt1’, and ‘vt2’ denote ground values that satisfy the constraints given above. Stair-shapes having first steps of different heights are not treated as being qualitatively different.

	$TT^+$	$TT^-$	$VT^+$	$VT^-$
Case 1	tt1	UC	vt1	vt2
Case 2	tt1	tt2	vt1	vt2
Case 3	tt1	UC	vt1	NOW (tt1=>vt1)
Case 4	tt1	tt2	vt1	NOW (tt1=>vt1)

**Figure 2. Combinations of Time Attributes**

We have set the context for using spatial indices for indexing bitemporal data. Some indices for bitemporal data that are based on this approach already exist; we discuss them next.

### 3. Overview of the Existing Bitemporal Indices

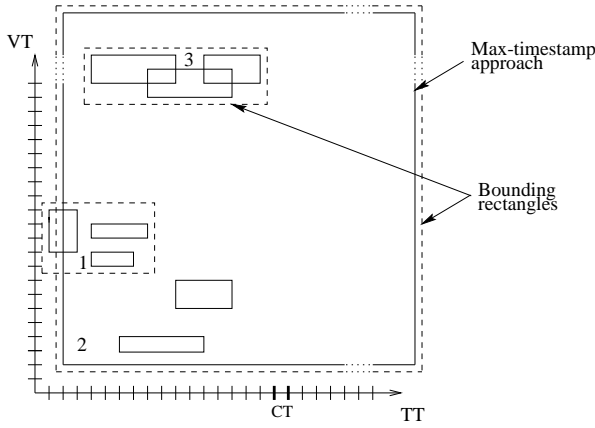
References [2, 23] provide comprehensive surveys of indices for different types of temporal data. This section focuses solely on the indexing of bitemporal data.

All existing indices for now-relative bitemporal data are based on the idea that bitemporal data can be viewed as a special case of spatial data (recall Figure 1) and that spatial indices can be utilized to index bitemporal data.


Many indices have been developed for spatial data with extent (i.e., non-point data) [17]. One of the most robust such indices is the R-tree [7] in its different variants—e.g., the  $R^+$ -tree [22], the  $R^*$ -tree [1], and the Hilbert R-tree [11]. In the R-tree, entries of a leaf-level node store minimum bounding rectangles of spatial regions together with pointers to the data tuples containing those regions. Entries of a non-leaf node store minimum bounding rectangles of child nodes together with pointers to those child nodes. The minimum bounding rectangle of a child node is the rectangle that bounds all entries of that child node. All variants of the R-tree try to minimize the overlap between the minimum bounding rectangles of the nodes at each level of the tree and to minimize the dead space in the bounding rectangle of each node (the dead space is the space in the minimum bounding rectangle not occupied by any of the enclosed rectangles). Minimizing overlap reduces the I/O-incurring branching of search into several subtrees. Minimizing dead space reduces the probability that queries unnecessarily access disk pages, eventually finding no qualifying data.

The  $R^*$ -tree is promising for indexing now-relative bitemporal data, but does not accommodate growing bitemporal regions. The straightforward approach to accommodating growing regions, the *maximum-timestamp approach*, is to represent these regions using the maximum possible transaction- and valid-time values. As a consequence, their minimum bounding rectangles are static, but the result is also excessive overlap and dead space, as illustrated in Figure 3(a).

Kumar et al. [12, 13] propose a new index to handling now-relative transaction time, but do not address now-relative valid time—data regions with an open valid time still must be represented using a maximum-timestamp value. In their index, the 2-R index, they use two R-trees. The front R-tree indexes all growing (i.e., current) rectangles, while the back R-tree indexes all static (i.e., logically deleted) rectangles. Observing that all growing rectangles are in the front tree and that they all extend to the progressing current time, Kumar et al. show that storing only the non-growing transaction-time begin value together with their fixed valid-time interval in the front tree is adequate to support now-relative transaction time. The 2-R index contends well with now-relative transaction time and performs



(a)



(b)

**Figure 3. Indexing Growing Bitemporal Regions (a) Using Maximum-Timestamp Values and (b) Using the GR-Tree**

better [3] than the single, maximum-timestamp R\*-tree, but it suffers from the penalty that two trees often have to be searched in a single query, resulting in more disk accesses and diminishing the advantages gained from the decreased overlap. The problem of representing now-relative valid time also remains open in the 2-R index. As discussed already, using the maximum valid-time value for NOW is not promising.

Instead of replacing UC and NOW with maximum time values, the GR-tree [3] extends the R\*-tree [1] to allow for the storage of these variables in tree nodes. The GR-tree thus accommodates bitemporal regions (shown in Figure 1) and uses minimum bounding regions that can be either static or growing and either rectangles or stair-shapes (see Figure 3(b)). It has been shown [3] that the GR-tree outperforms indices based on the previously described approaches by a factor of 3 or more. The GR-tree is an efficient index, but its implementation requires access to and extension of the DBMS kernel.

Based on the above, we may conclude that two different approaches exist for indexing now-relative bitemporal data. The first approach is to create a new index. The GR-tree, although based on the R\*-tree, exemplifies this approach. The second approach is to transform now-relative bitemporal data, in this way eliminating the variables UC and NOW and obtaining static data, and then apply an existing “off-the-shelf” spatial index. This paper explores the latter approach.

The main purpose of the data transformations employed here is to render growing regions stationary. Transformations are used widely in indexing, but always with different purposes.

## 4. The 4-R Index

In the 4-R index, we (1) divide bitemporal data regions into four classes, (2) perform transformations of the regions in each of the classes, thus eliminating any variables, and (3) use separate R-trees for indexing the transformed data regions of each class.

Section 4.1 describes the data transformation. In order to use the index when answering queries, the queries must also be transformed; Section 4.2 presents the query transformation. Section 4.3 concerns the implementation of the 4-R index.

### 4.1. Data Transformations

In the interest of generality and for our purpose, it is appropriate to model a bitemporal tuple as a pair of a bitemporal region and a tuple identifier. This corresponds to the information captured at the leaf level of a secondary index such as the R-tree.

The main goal of the transformation of now-relative bitemporal data is to eliminate the variables UC and NOW, so that the data can be indexed with R-trees. To this end, we distinguish between four types of bitemporal data, depending on whether  $TT^\dagger$  is or is not equal to UC and whether  $VT^\dagger$  is or is not equal to NOW. For each type, the transformed data is variable-free. The bitemporal region of a transformed, variable-free bitemporal tuple is always static. Before defining the data transformation, we define the domain of bitemporal data and the domain of variable-free bitemporal data.

**Definition 4.1** Let the domain of timestamp values be  $T$  and the domain of tuple identifiers be  $ID$ . We then define  $D^B$ , the domain of bitemporal tuples, and  $D^S$ , the domain of variable-free bitemporal tuples, as follows.

$$D^B \triangleq \{(TT_r^+, TT_r^-, VT_r^+, VT_r^-, id_r) \in T \times T \cup \{UC\} \times T \times T \cup \{NOW\} \times ID \mid$$

$$\begin{aligned} & (\text{TT}_r^- = \text{UC} \vee \text{TT}_r^+ \leq \text{TT}_r^-) \wedge \\ & (\text{VT}_r^- = \text{NOW} \vee \text{VT}_r^+ \leq \text{VT}_r^-) \end{aligned}$$

$$D^S \triangleq \{ \langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{VT}_r^-, id_r \rangle \in T \times T \times T \times T \times ID \mid \text{TT}_r^+ \leq \text{TT}_r^- \wedge \text{VT}_r^+ \leq \text{VT}_r^- \} \quad \square$$

The “ $r$ ” subscripts are used to clearly separate the data rectangles from the query rectangles that will be introduced in the next section.

The data transformation defined next transforms a bitemporal tuple into a variable-free bitemporal tuple augmented by a transformation type.

**Definition 4.2** Let  $R \subseteq D^B$  and define  $Type = \{1, 2, 3, 4\}$ . Then the 4-R data transformation  $\mathcal{T}_D : 2^{D^B} \rightarrow 2^{D^S \times Type}$  is defined as follows.

$$\mathcal{T}_D(R) \triangleq \{ \mathcal{T}_r(\langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{VT}_r^-, id_r \rangle) \mid \langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{VT}_r^-, id_r \rangle \in R \},$$

where

$$\mathcal{T}_r(\langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{VT}_r^-, id_r \rangle) \triangleq \begin{cases} \langle \text{TT}_r^+, \text{TT}_r^+, \text{VT}_r^+, \text{VT}_r^+, id_r, 1 \rangle & \text{if } \text{TT}_r^- = \text{UC} \wedge \text{VT}_r^- = \text{NOW} \\ \langle \text{TT}_r^+, \text{TT}_r^+, \text{VT}_r^+, \text{VT}_r^-, id_r, 2 \rangle & \text{if } \text{TT}_r^- = \text{UC} \wedge \text{VT}_r^- \neq \text{NOW} \\ \langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{VT}_r^+, id_r, 3 \rangle & \text{if } \text{TT}_r^- \neq \text{UC} \wedge \text{VT}_r^- = \text{NOW} \\ \langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{VT}_r^-, id_r, 4 \rangle & \text{if } \text{TT}_r^- \neq \text{UC} \wedge \text{VT}_r^- \neq \text{NOW} \end{cases} \quad \square$$

Four trees, numbered R1 through R4, are created and populated with result tuples according to their types. Figure 4, described next, illustrates the intuition behind the mapping.

Tree R1 indexes the regions that, before the transformation, had non-fixed valid- and transaction-time end values. Knowing that such regions extend to the current time in transaction time ( $\text{TT}^- = \text{UC}$ ) and to the line  $\text{VT} = \text{TT}$  ( $\text{VT}^- = \text{NOW}$ ), it suffices to represent the regions by two-dimensional points  $\langle \text{TT}_r^+, \text{TT}_r^+, \text{VT}_r^+, \text{VT}_r^+ \rangle$ , represented in turn by  $(\text{TT}^+, \text{VT}^+)$  in the index.

Next, tree R2 indexes regions that had fixed valid-time end values, but non-fixed transaction-time end values prior to the transformation. Because such regions are rectangles that grow in transaction time and extend to the current time in this dimension ( $\text{TT}^- = \text{UC}$ ), we can represent these regions by two-dimensional intervals  $\langle \text{TT}_r^+, \text{TT}_r^+, \text{VT}_r^+, \text{VT}_r^- \rangle$ , represented in the index as  $(\text{TT}^+, \text{VT}^+, \text{VT}^-)$ .

Tree R3 is devoted to regions that, before the transformation, had fixed transaction-time end values, but non-fixed valid-time end values. These regions are all stair-shapes that extend to the line  $\text{VT} = \text{TT}$  ( $\text{VT}^- = \text{NOW}$ ). These may be represented by two-dimensional intervals  $\langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{VT}_r^- \rangle$ , captured in the index as  $(\text{TT}^+, \text{TT}^-, \text{VT}^+)$ .

Finally, tree R4 accommodates originally static data regions for which there is no need for transformation.

The next step is to explore search in the four trees of the 4-R index that accommodate the transformed data. Queries must be transformed as well.

## 4.2. Query Transformations

We investigate the most common type of index query, namely the rectangular intersection query (denoted by “ $-//R/R$ ” in the terminology of Tsotras et al. [24]). This type of query includes point queries as well as the different kinds of range queries supported by the valid- and transaction-time timeslice operators frequently present in temporal query languages. Let  $\langle \text{TT}_q^+, \text{TT}_q^-, \text{VT}_q^+, \text{VT}_q^- \rangle \in T \times T \times T \times T$  denote the argument rectangle of an intersection query, where  $T$  is the domain of timestamps. We will make the reasonable assumptions that  $\text{TT}_q^+ \leq \text{TT}_q^-$ ,  $\text{VT}_q^+ \leq \text{VT}_q^-$ , and  $\text{TT}_q^- \leq \text{CT}$ , where  $\text{CT}$  is the value of the current time. The following definition gives the result of an intersection query on bitemporal data.

**Definition 4.3** Define  $\langle \text{TT}_{r_1}^+, \text{TT}_{r_1}^-, \text{VT}_{r_1}^+, \text{VT}_{r_1}^- \rangle \cap \langle \text{TT}_{r_2}^+, \text{TT}_{r_2}^-, \text{VT}_{r_2}^+, \text{VT}_{r_2}^- \rangle$  by  $(\text{TT}_{r_1}^+ \leq \text{TT}_{r_2}^+) \wedge (\text{TT}_{r_1}^- \geq \text{TT}_{r_2}^-) \wedge (\text{VT}_{r_1}^+ \leq \text{VT}_{r_2}^+) \wedge (\text{VT}_{r_1}^- \geq \text{VT}_{r_2}^-)$ . Also let  $q = \langle \text{TT}_q^+, \text{TT}_q^-, \text{VT}_q^+, \text{VT}_q^- \rangle$  and  $R \subseteq D^B$ . Then an intersection query  $\text{Intersect}^B$  on  $R$  with query rectangle  $q$  and current time value  $\text{CT}$  as parameters is defined as follows.

$$\begin{aligned} \text{Intersect}^B[q, \text{CT}](R) \triangleq & \{ id_r \mid \langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{VT}_r^- \rangle \in R \wedge \\ & ((\text{TT}_r^- = \text{UC} \wedge \text{VT}_r^- = \text{NOW} \wedge \\ & \quad \text{TT}_q^- \geq \text{VT}_q^+ \wedge q \cap \langle \text{TT}_r^+, \text{CT}, \text{VT}_r^+, \text{CT} \rangle) \vee \\ & (\text{TT}_r^- = \text{UC} \wedge \text{VT}_r^- \neq \text{NOW} \wedge \\ & \quad q \cap \langle \text{TT}_r^+, \text{CT}, \text{VT}_r^+, \text{VT}_r^- \rangle) \vee \\ & (\text{TT}_r^- \neq \text{UC} \wedge \text{VT}_r^- = \text{NOW} \wedge \\ & \quad \text{TT}_q^- \geq \text{VT}_q^+ \wedge q \cap \langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{TT}_r^- \rangle) \vee \\ & (\text{TT}_r^- \neq \text{UC} \wedge \text{VT}_r^- \neq \text{NOW} \wedge \\ & \quad q \cap \langle \text{TT}_r^+, \text{TT}_r^-, \text{VT}_r^+, \text{VT}_r^- \rangle) \} \end{aligned} \quad \square$$

The first line restricts result tuple identifiers to be in the argument tuples. As for Definition 4.2, each of the next four disjuncts is devoted to one type of bitemporal region. The first disjunct identifies the subset of qualifying growing stair-shapes, the second identifies qualifying growing rectangles, the third, static stair-shaped regions, and the fourth,

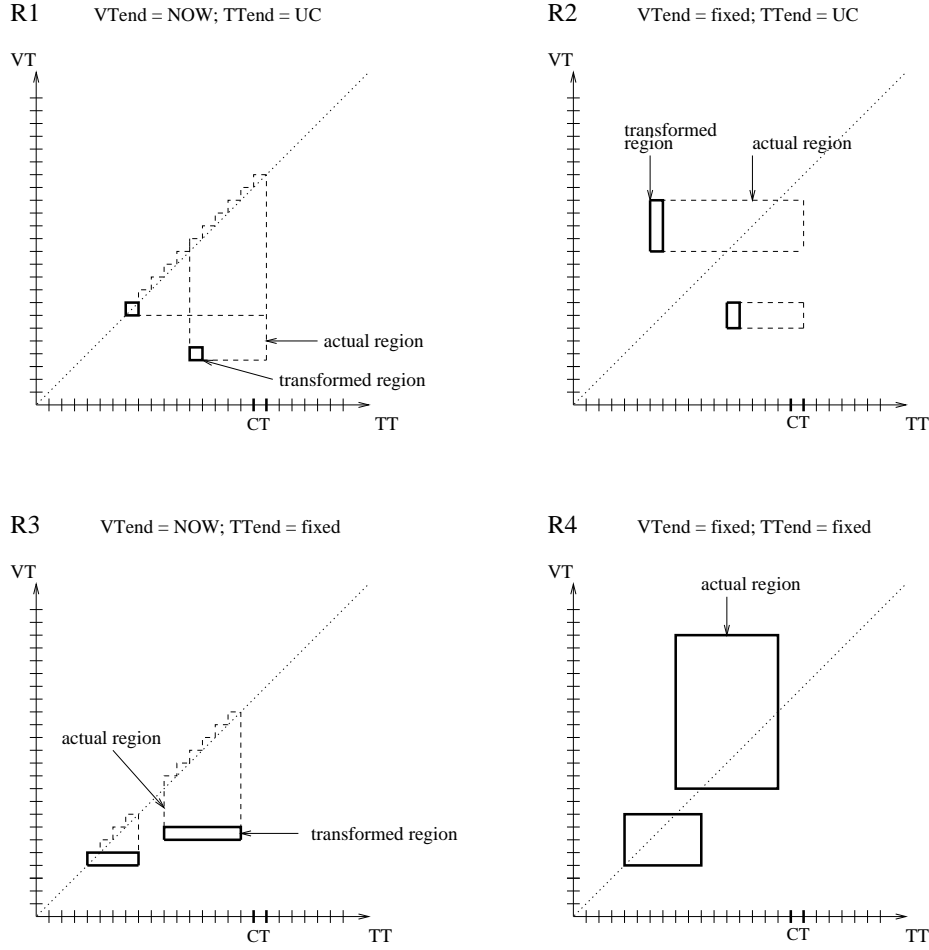


Figure 4. Data Storage in the Four Trees of the 4-R Index

static rectangles. In cases of stair-shaped regions, in addition to checking the intersection between the bounding rectangle of a region and a query rectangle, it is specified that the lower-right corner of a query must be below or on the  $VT = TT$  line (condition  $TT_q^- \geq VT_q^+$ ).

Similarly to Definition 4.3, we next define the rectangular intersection query on variable-free bitemporal data. The result of this query is independent of the current time.

**Definition 4.4** Let  $q = \langle TT_q^+, TT_q^-, VT_q^+, VT_q^- \rangle$  and  $S \subseteq D^S$ . Then an intersection query  $Intersect^S$  on  $S$  with a query rectangle  $q$  as a parameter is defined as follows.

$$Intersect^S[q](S) \triangleq \{id_r \mid \langle TT_r^+, TT_r^-, VT_r^+, VT_r^-, id_r \rangle \in S \wedge q \cap \langle TT_r^+, TT_r^-, VT_r^+, VT_r^- \rangle\} \quad \square$$

With Definitions 4.3 and 4.4 in place, we are now in a position to define the 4-R query transformation,  $\mathcal{T}_q$ , that goes with the data transformation given in the previous section.

The transformation maps an intersection query on the original data to two or four corresponding queries on the transformed data.

**Definition 4.5** Initially define:

$$\begin{aligned} R &\subseteq D^B \\ S &= \mathcal{T}_D(R) \\ S_i &= \{ \langle TT_r^+, TT_r^-, VT_r^+, VT_r^-, id_r \rangle \mid \\ &\quad \langle TT_r^+, TT_r^-, VT_r^+, VT_r^-, id_r, i \rangle \in S \}, i = 1, 2, 3, 4 \\ q &= \langle TT_q^+, TT_q^-, VT_q^+, VT_q^- \rangle \\ q_1 &= \langle 0, TT_q^-, 0, VT_q^- \rangle \\ q_2 &= \langle 0, TT_q^-, VT_q^+, VT_q^- \rangle \\ q_3 &= \langle \max(TT_q^+, VT_q^+), TT_q^-, 0, VT_q^- \rangle \\ q_4 &= \langle TT_q^+, TT_q^-, VT_q^+, VT_q^- \rangle \end{aligned}$$

Then the query transformation  $\mathcal{T}_q : [2^{D^B} \rightarrow 2^{ID}] \rightarrow [2^{D^S \times Type} \rightarrow 2^{ID}]$  is defined as follows.

$$\mathcal{T}_q(Intersect^B[q, CT])(S) \triangleq \begin{cases} \bigcup_{i=1,2,3,4} Intersect^S[q_i](S_i) & \text{if } TT_q^- \geq VT_q^+ \\ \bigcup_{i=2,4} Intersect^S[q_i](S_i) & \text{if } TT_q^- < VT_q^+ \end{cases} \quad \square$$

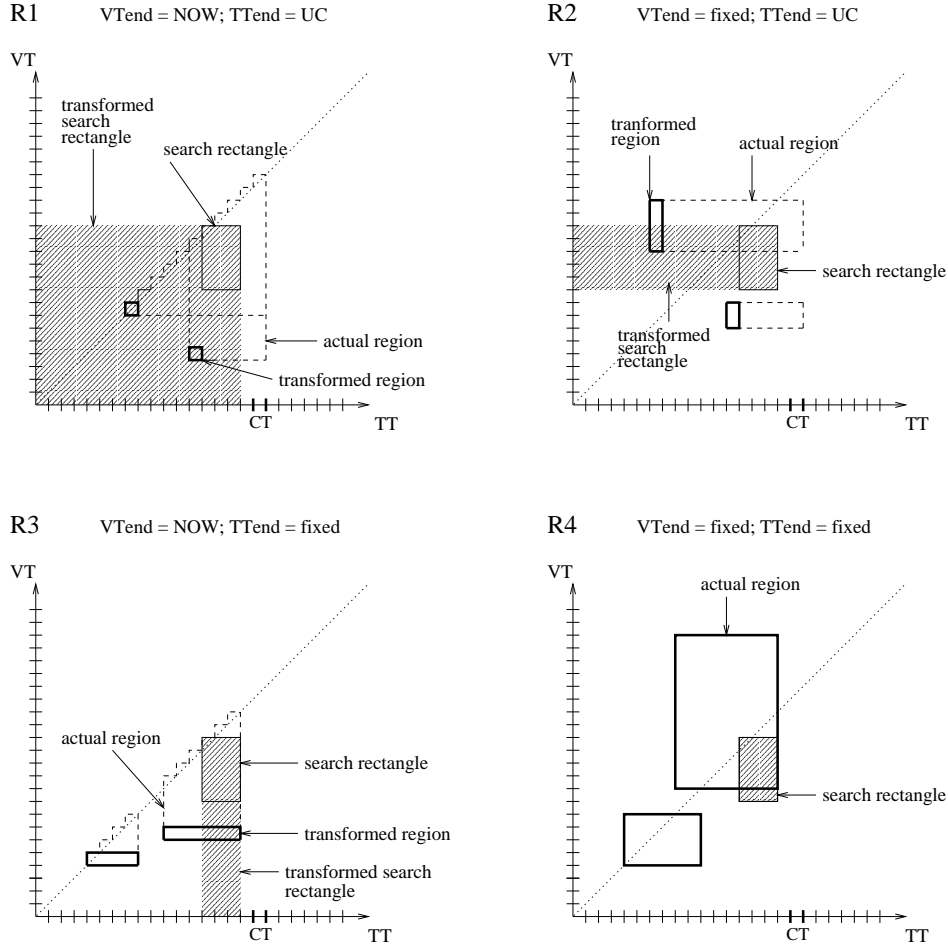


Figure 5. Search in the Four Trees of the 4-R Index

The search spreads to two or four trees and is performed differently in each tree. Figures 5 and 6 illustrate the original search rectangle and the corresponding transformed search rectangle in each tree. A discussion of the search in each of the four trees follows Theorem 4.1 that states that the combination of the 4-R query and 4-R data transformation yields perfect precision and recall, i.e., is correct. The proof is given in the extended version of the paper [4].

**Theorem 4.1** For each  $q = \langle TT_q^+, TT_q^-, VT_q^+, VT_q^- \rangle$  and each data set  $R \subseteq D^B$ ,

$$Intersect^B[q, CT](R) = \mathcal{T}_q(Intersect^B[q, CT])(\mathcal{T}_D(R)). \quad \square$$

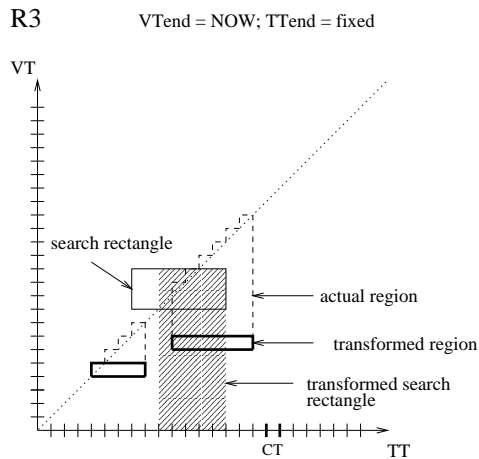
Searching tree R1, the argument search rectangle is enlarged to cover the space spanned from the origin of the transaction and valid time to the argument rectangle's top-right corner. Tree R1 contains no data points above the line  $VT = TT$  because the original regions encoded by the points in this tree extend only to the line  $VT = TT$ . Thus, the transformed search rectangle could also be reduced to not extend

above the line  $VT = TT$  without affecting correctness. But because the area avoided by this reduction contains no data, we simply use the unreduced rectangle.

When searching tree R2, we should look not only for data intervals overlapping with the original search rectangle, but also for intervals to the left of the search rectangle.

The search-rectangle transformation in tree R3 is similar, but now the argument search rectangle is extended downwards instead of to the left, and there is a subtle complication. When part of a search rectangle is below the line  $VT = TT$  and another part is above, only the part below this line should be extended downwards; extending the entire rectangle would yield “false drops” and would thus jeopardize precision. This is illustrated in Figure 6, where extending the entire search rectangle would yield one false drop. In Definition 4.5, function “max” handles this case.

Since tree R4 indexes untransformed bitemporal rectangles, transformation of the argument search rectangle is unnecessary.



**Figure 6. Search Specifics for Tree R3**

Several noteworthy special cases occur when querying the four trees. In trees R1 and R3, the indexed points and intervals reside only below the line  $VT = TT$ , and the regions encoded by these points or intervals also do not extend above this line. Thus, search in these trees is only performed when at least some part of a search rectangle goes below  $VT = TT$ .

Another special case is the current-time transaction-timeslice query ( $TT_q^+ = TT_q^- = CT$ ), which is expected to occur frequently in practice. Current data resides only in trees R1 and R2, so this timeslice query may be restricted to these two trees. If, in addition, a current time transaction-timeslice query is above the  $VT = TT$  line, the only tree to be searched is R2. As a final special case, if such a current-time transaction timeslice has no constraints on valid time, all tuples indexed by trees R1 and R2 should simply be returned, and no search is required.

### 4.3. Implementation

An important advantage of the 4-R index is that it may reuse preexisting and efficient R-tree-based indexes such as the R\*-tree without any modifications because only points, intervals, and rectangles are indexed—there are no growing bitemporal regions. This means that the 4-R index may be implemented in a layer on top of any DBMS that supports R-trees, with no need for extending the DBMS’s query optimizer and recovery and concurrency control subsystems. This renders the index a very practical one.

In addition, the data transformation of bitemporal shapes into intervals and, especially, points may positively impact performance because these simpler shapes take less space and thus lead to trees with higher fanouts. The query transformation, where argument rectangles are always enlarged, may have the reverse effect on the performance.

The layer implementing the 4-R technique on top of four R\*-trees is responsible for three main tasks.

**insertions** The layer determines into which of the four trees a new data region has to be inserted, and it passes the correct insertion statement to the appropriate tree.

**deletions** The layer determines the tree that may index a given data region, and it passes the correct deletion statement to that tree.

**queries** The layer translates queries into four corresponding queries, each of which is passed to the appropriate tree; and the layer returns the combined result obtained from searching the four trees.

Update is modeled as a deletion followed by an insertion. The algorithms implemented by the layer follow the theory developed earlier in this section and are outlined in Figure 7. Although somewhat dependent on the specific underlying DBMS, the efforts required to implement these algorithms are quite small.

## 5. Performance Experiments

In this section, we compare the performance of four indices: the 4-R index, the GR-tree, the R-tree, termed 1-R, and the 2-R index, with the latter two using the maximum-timestamp approach. We first consider data and query generation, then discuss the performance results.

### 5.1. Data and Query Generation

The four indices were implemented using the Generalized Search Tree Package, GiST [8]. The numbers of I/O operations are measured using simulation. The page size is set to 1024 bytes, and one tree node is stored in one page. Thus, one node read or write corresponds to one page access (one I/O operation). A buffer of 100 pages is allocated for each index<sup>2</sup>. Omitting a buffer may lead to quantitatively and qualitatively incorrect conclusions [14]. The root is always kept in the buffer; for the other nodes, the least-recently-used page replacement policy is employed. If a node is changed during an insertion or a deletion, its page is changed in the buffer and marked as “dirty.” Dirty pages are written to disk at the end of the operation or when they have to be removed from the buffer.

To fairly compare the search and update performance of the four indices, the same data has to be inserted into the indices, and the same queries have to be run on them. We use so-called *workloads* to simulate the construction and usage of an index for a certain period—the *index life-time*. In

<sup>2</sup>For the 4-R, 4 buffers of 25 pages are allocated, and for the 2-R, 2 buffers of 50 pages are allocated.



$$\begin{aligned}
\text{insert}(TT^+, TT^-, VT^+, VT^-) &= \begin{cases} \text{insertR1}(TT^+, VT^+) & \text{if } TT^- = UC \wedge VT^- = \text{NOW} \\ \text{insertR2}(TT^+, VT^+, VT^-) & \text{if } TT^- = UC \wedge VT^- \neq \text{NOW} \\ \text{insertR3}(TT^+, TT^-, VT^+) & \text{if } TT^- \neq UC \wedge VT^- = \text{NOW} \\ \text{insertR4}(TT^+, TT^-, VT^+, VT^-) & \text{if } TT^- \neq UC \wedge VT^- \neq \text{NOW} \end{cases} \\
\text{delete}(TT^+, TT^-, VT^+, VT^-) &= \begin{cases} \text{deleteR1}(TT^+, VT^+) & \text{if } TT^- = UC \wedge VT^- = \text{NOW} \\ \text{deleteR2}(TT^+, VT^+, VT^-) & \text{if } TT^- = UC \wedge VT^- \neq \text{NOW} \\ \text{deleteR3}(TT^+, TT^-, VT^+) & \text{if } TT^- \neq UC \wedge VT^- = \text{NOW} \\ \text{deleteR4}(TT^+, TT^-, VT^+, VT^-) & \text{if } TT^- \neq UC \wedge VT^- \neq \text{NOW} \end{cases} \\
\text{search}(TT^+, TT^-, VT^+, VT^-) &= \begin{cases} \text{searchR1}(0, TT^-, 0, VT^-) \cup \\ \text{searchR2}(0, TT^-, VT^+, VT^-) \cup \\ \text{searchR3}(\max(TT^+, VT^+), TT^-, 0, VT^-) \cup \\ \text{searchR4}(TT^+, TT^-, VT^+, VT^-) \text{ if } TT^+ \neq CT \wedge TT^- \geq VT^+ \\ \\ \text{searchR2}(0, TT^-, VT^+, VT^-) \cup \\ \text{searchR4}(TT^+, TT^-, VT^+, VT^-) \text{ if } TT^+ \neq CT \wedge TT^- < VT^+ \\ \\ \text{searchR1}(0, TT^-, 0, VT^-) \cup \\ \text{searchR2}(0, TT^-, VT^+, VT^-) \quad \text{if } TT^+ = CT \wedge [VT^+, VT^-] \neq [t_{min}, t_{max}] \wedge TT^- \geq VT^+ \\ \\ \text{searchR2}(0, TT^-, VT^+, VT^-) \quad \text{if } TT^+ = CT \wedge [VT^+, VT^-] \neq [t_{min}, t_{max}] \wedge TT^- < VT^+ \\ \\ \mathbf{R1} \cup \mathbf{R2} \quad \quad \quad \text{if } TT^+ = CT \wedge [VT^+, VT^-] = [t_{min}, t_{max}] \end{cases}
\end{aligned}$$

**Figure 7. Algorithms for Insertion, Deletion, and Search in the 4-R Index**

our experiments, a workload typically contains 60,000 update operations. An update operation is either an insertion or a (logical) deletion. First, 4000 consecutive insertions are performed, and at each later point in time an insertion occurs with probability  $Ins$  and a deletion occurs with probability  $1 - Ins$ .

Several parameters are used for generating the data to be inserted into the indices. The valid-time interval length is uniformly distributed between 0 and the maximum valid-time interval length,  $VL$ . Alternatively, the valid-time end can be NOW. The percentage of data to be inserted into an index and having valid-time end equal to NOW is denoted as  $PNow$ . We choose the valid-time begin to be strongly bounded to the data-insertion time. Specifically, it is normally distributed with a mean equal to the insertion time and with some deviation  $Dev$  that specifies how densely regions inserted into the trees are distributed around the  $y = x$  axis. If  $Dev$  is big and the valid-time end is fixed, the regions are scattered throughout the valid-time universe.  $Dev$  influences how far below the  $y = x$  axis the stair-shapes can start. In the experiments reported here, we use value 5000 for  $Dev$ .

A workload also contains queries intermixed with the update operations. We perform bitemporal range, timeslice, and point queries. Parameter  $Qcur$  denotes the percentage

of “current” queries, i.e., current timeslice and point queries or bitemporal range queries that extend to the current time. Parameters  $Qrange$ ,  $Qslice$ , and  $Qpoint$  denote the percentages of, respectively, bitemporal range, timeslice, and point queries that will be run. Parameter  $QmaxI$  denotes the maximum valid-time range for bitemporal range queries and timeslice queries, and the maximum transaction-time range for bitemporal range queries.

We use *overlap* as the query predicate, meaning that regions that overlap with the given query window qualify for the result.

The data and query generation parameters described above are termed *workload parameters* and are summarized in Table 2. In different experiments, the values of these parameters are varied. The values used are given together with the description of each concrete experiment. If the value of some parameter is omitted, the “standard” value is used.

We intermix queries with update operations in the workload with the aim of measuring search performance throughout the entire index life-time. In the experiments, for each used workload we compute the average I/O cost of update and search operations present in that workload. In each experiment, we also compute pagination, overlap, and dead space in the trees.

The pagination is the percentage of the allocated space

Parameter	Description	Values used	“Standard” value
<i>PNow</i>	percentage of data with valid-time end equal to NOW	0, 20, 40, 60, 80, 100	60
<i>Ins</i>	percentage of insertions	50, 60, 70, 80, 90, 100	70
<i>VL</i>	maximum valid-time interval length	50, 100, 500, 1000, 3000, 5000	500
<i>Qcur</i>	percentage of current queries	0, 25, 50, 75, 100	65
<i>Qrange</i>	percentage of range queries	25, 100	25
<i>Qslice</i>	percentage of timeslice queries	50, 100	50
<i>Qpoint</i>	percentage of point queries	25, 100	25
<i>Qmaxl</i>	maximum valid- and transaction-time intervals given in a query	1, 100, 300, 500, 1000, 3000	300

**Table 2. Workload Parameters**

that is utilized for data, i.e., it is the percentage of each node’s space that on the average is filled with entries. The dead space of a node is the difference between the area of the minimum bounding region of that node and the area of the union of the minimum bounding regions of all entries of that node. The dead space of a tree level is the sum of the dead space of each node at that level. The overlap in a node is the difference of the sum of all areas of that node’s entries and the area of the union of all the entries in the node. The overlap in a tree level is the sum of the overlap values of all nodes in that level. The measurement units for dead space and overlap are quadratic time-points, i.e., the same as those for the simple area of a bitemporal region.

## 5.2. Comparison of the Four Bitemporal Indices

In this section, we first discuss general issues in addition to specifics of the four indices that influence their performance. Then we present the results of concrete experiments with the bitemporal indices. For each experiment, the workload parameters are given, search and update performance results are illustrated, and discussions are offered that cover also index pagination, dead space, and overlap.

### 5.2.1. General Observations

The number of I/O operations performed during a search depends on index pagination, dead space, and overlap. If nodes are poorly filled, many nodes have to be accessed to retrieve a number of entries. Dead space leads to accessing nodes in vain, eventually finding no qualifying entries. Overlap between nodes leads to I/O-incurring branching of search into several subtrees.

In the 1-R and 2-R, dead space and overlap are excessive because they depend on the maximum-timestamp value, which must be very large in order to exceed any fixed time value used throughout the existence of an index.

As mentioned earlier, for some queries, only two trees in the 4-R need to be searched. In addition, several special cases of the 4-R occur for specific values of the workload parameters. If the value of the parameter *Ins*, the percentage of insertions, is 100%, regions will never be deleted, and

thus trees R3 and R4, devoted to non-current data, will be empty. If the value of the parameter *PNow*, the percentage of data with valid-time end equal to NOW, is 100%, trees R2 and R4 will be empty because they are devoted to data with fixed valid-time intervals. If *PNow* is 0% then trees R1 and R3 will be empty. Thus, it is appropriate to perform specific experiments to investigate properties of the individual trees of the 4-R index.

### 5.2.2. Experiments With Various Data

First, experiments were conducted to find out how varying percentages of data having valid-time end equal to NOW influence search and update performance in the four indices. Different workloads were constructed with varying values of parameter *PNow*. The search and update performance of the four indices with such data are given in Figures 8(a) and 9(a), respectively.

In general, the GR-tree has the best search performance. The 4-R index does not loose much and even outperforms the GR-tree when *PNow* is 100%, while the two maximum-timestamp-approach-based indices are clearly worse. The problem of the latter two is the big overlap and dead space caused by the huge rectangles representing growing bitemporal regions. One of the reasons why the GR-tree is better than the 4-R index is that its pagination (65%) is better than the pagination of the 4-R trees (50% in R3 and 60% in the other trees). Another reason is tree R2 in the 4-R index. This tree indexes data with non-fixed transaction-time intervals (i.e., current data) by physically storing only “vertical lines.” Because of the sequential nature of transaction time, the R\*-tree algorithms in many cases fail to group these vertical lines into nodes with quadratic minimum bounding rectangles. Often minimum bounding rectangles of the nodes in R2 are long in the valid-time direction. On the other hand, the transformed queries for R2 are not very long in the valid-time direction, but extend to the very beginning of transaction time. Such queries access a lot of nodes with minimum bounding rectangles that are long in valid time, but not many entries (some times none at all) from these nodes qualify for the answer.

However, when *PNow* is near 100%, the 4-R starts to




Figure 8. Search Performance, (a) Varying  $PNow$ , (b)  $Ins$ , and (c)  $VL$




Figure 9. Update Performance, (a) Varying  $PNow$ , (b)  $Ins$ , and (c)  $VL$

outperform the GR-tree. First, when  $PNow$  is large, data is concentrated in trees R1 and R3, and the bad performance of the poorly populated R2 does not seriously affect the overall 4-R index performance. Second, with 65% of current queries (we use this number as our standard), the majority of the qualifying data is “current” data that resides in R1 and R2. Since R2 is poorly populated, the major part of the qualifying data is retrieved from R1. Tree R1 of the 4-R index has very good selectivity, i.e., usually almost all entries of retrieved nodes qualify for the answer. Consequently, almost the minimal number of nodes have to be accessed during a search, and thus tree R1 significantly contributes to the good performance of the 4-R index.


When  $PNow$  is low, 1-R outperforms 2-R. One of the reasons is the worse pagination of the 2-R (55%, opposed

to 65% of the 1-R).

When  $PNow$  is large, the 2-R outperforms the 1-R. Then the majority of the rectangles in the 1-R extend to the maximum valid-time value; and to the maximum transaction-time value if data is current. The rectangles representing not current data (“short” in the transaction-time direction), however, are mixed with “current” ones in the tree nodes. Thus the selectivity in the 1-R when answering a current query (for which only current rectangles qualify) is not very high. The storing of old and current data in separate trees in the 2-R index shows the advantage.

Concerning updates, the 4-R and 2-R indices achieve better performance than one-tree indices.

Another set of experiments was carried out to investigate how the percentage of “current” data influences search and



**Figure 10. Search Performance, Varying  $Q_{cur}$  for (a) Range Queries, (b) Point Queries, and Varying  $Q_{maxI}$  for (c) Timeslice Queries**

update performance in the four indices. In this study, workloads generated with different values of parameter  $Ins$  were used. Results of the study are given in Figures 8(b) and 9(b).

When the  $Ins$  value increases, regions are more rarely deleted, current data remains current for a long period, and the amount of current data increases over time. Thus with 65% as the average percentage of current queries (meaning that mostly current data is of interest), when  $Ins$  increases, more data is retrieved from the indices, resulting in an increasing cost of a search. Overall, the GR-tree has the best performance, but the 4-R is quite close. The difference between the GR-tree and the 4-R index is more visible when  $Ins$  is close to 100%; the GR-tree then achieves better pagination (70%).

The 4-R and 2-R indices achieve the best update performance also in this study.

Experiments were also performed varying  $VL$  (see Figures 8(c) and 9(c)). All of the experiments show that 1-R and 2-R indices, although being simpler than 4-R, cannot compete with it in search performance for any workloads where  $PNow$  is not close to 0.

### 5.2.3. Experiments With Various Queries

For a database where the amount of current data increases over time (using  $Ins$  equal to 70%, this is the case), the later in the database lifetime a query is issued, the more results are retrieved, which in turn requires more I/O operations. We already noticed the impact of the current queries in the previous sections. In this section, we experiment with various kinds of queries. First, we experiment with different types of queries varying  $Q_{cur}$ , the percentage of current


queries. The results for range and point queries are shown in Figures 10(a) and 10(b). It can be seen that the performance of the 1-R, 2-R, and 4-R indices drops as expected with an increasing amount of current queries. The GR-tree, on the contrary, performs better as the number of current queries increases. Investigating this “phenomenon,” we found out the following about the four indices.

The GR-tree nodes containing current data have better pagination than nodes containing non-current data. This is due to the sequential nature of transaction time. “Current” nodes are filled up by continuously arriving new entries, while new entries are rarely inserted into the “old” nodes, which are left half-full after node splits (the GR-tree split algorithm usually separates older and newer entries). Thus, although the number of retrieved entries increases with an increasing amount of current queries, the qualifying (current) entries are packed into a smaller number of nodes.

In the 4-R index, while the pagination of the trees that contain current data (R1 and R2) is higher than the pagination of trees R3 and R4, the performance is adversely affected by the bad performance of R2, which in turn is caused by the geometry of the minimum bounding rectangles in this tree (see Section 5.2.2). 2-R performs better than 1-R when the percentage of current queries increases, because current and old data are mixed in the nodes of the 1-R, and the percentage of qualifying entries from each node is low for current queries.

The kind of queries used (range, timeslice, or point queries) practically does not change the results.

The influence of the size of timeslice and range queries was also tested (Figure 10(c) and 11(a)). As can be expected, bigger queries lead to more entries being retrieved,



**Figure 11. Search Performance, Varying  $Q_{maxI}$  and Range Queries with (a) Standard  $Ins$  and  $PNow$ , (b)  $Ins = 100\%$ ,  $PNow = 0\%$ , and (c)  $Ins = 100\%$ ,  $PNow = 100\%$**

thus requiring more efforts to perform. This is more pronounced for range queries, which can be expanded to any size in two directions, than for timeslice queries, which can be expanded only in the valid-time direction.

#### 5.2.4. Experiments With Specialized Data

We experimented with specialized data to understand the properties of the four trees in the 4-R index. We have already observed in the previous sections that tree R1 of the 4-R index has very good selectivity, and, in contrast, tree R2 has quite low selectivity.

When  $Ins$  is 100% (data is never deleted) and  $PNow$  is 0% (valid-time intervals are fixed), regions representing transformed bitemporal data are solely indexed by tree R2 in the 4-R. The low selectivity in this tree explains why the 4-R index is worse not only than the GR-tree, but also than the R-tree (1-R) (see Figure 11(b)). Note that results of the 2-R index and the 4-R index are the same because the front tree of the 2-R is the same as tree R2 of the 4-R index, since there is no data with now-relative valid-time intervals. The pagination in both the 2-R and the 4-R is 63%, while it is 74% in the GR-tree and 71% in the 1-R.

When  $Ins$  is 100% (data is never deleted) and  $PNow$  is 100% (valid-time end is always equal to NOW), regions representing transformed bitemporal data are indexed solely by tree R1 in the 4-R. In tree R1, a rectangle of the transformed query covers a huge area and therefore, during a search, most of the accessed entries qualify for the result. But tree R1 has low pagination (54%) and therefore the 4-R index loses to the GR-tree which has a pagination of 71% (see Figure 11(c)).

## 6. Conclusions

Almost none of the proposed indices for temporal data support both now-relative valid time and transaction time, which are accommodated by most of the temporal data models and are natural and meaningful for many kinds of applications. The straightforward R-tree based solution to indexing now-relative bitemporal data, the maximum-timestamp approach, is not effective. Another R-tree based index, the GR-tree, employs a special structure and algorithms to contend with now-relative valid time and transaction time. Although having good performance, the GR-tree is not currently available in any existing DBMSs.

This paper shows how efficient indexing of now-relative bitemporal data can be achieved with little effort by implementing a layer on top of a DBMS supporting R-trees. To enable the use of R-trees, the proposed technique applies transformations to bitemporal data regions. Four types of bitemporal regions are distinguished, and the transformed counterparts of the bitemporal regions of each of the four types are stored in separate R\*-trees. Each query is also transformed into (up to) four separate queries, one for each tree.

The proposed index, the 4-R index, may be seen as a generalization of the 2-R index [12], where transformation is used to support now-relative transaction time. In another sense, the 4-R index is a special case of the GR-tree. The insertion algorithm of the GR-tree separates bitemporal data regions of different types into different nodes, achieving a tree with groups of nodes storing bitemporal data regions of the same kind. Thus, the 4-R index is an extreme special case of the GR-tree, where such groups of nodes form four

different trees.

Detailed performance experiments show that the search performance of the 4-R index is comparable to that of the GR-tree, in most cases not differing more than 25% from the GR-tree. On the other hand, the 4-R index shows very good, steady update performance, surpassing the GR-tree's update performance.

The performance experiments also offer insight into possible improvements of the 4-R index. Minimum bounding rectangles in tree R2 that are unproportionally long in the valid-time direction adversely affect its performance. In contrast, because the transformed queries in R2 have relatively long transaction-time extents, the same kind of geometry for minimum bounding rectangles should result in less I/O during searches. The same observations hold for tree R3.

One approach to achieving the desired minimum bounding rectangles is to modify the split algorithms of the R\*-tree, but this would mean that "off-the-shelf" implementations of the R\*-tree cannot be used. To preserve the reusability of "off-the-shelf" technology, we are instead developing and experimenting with a number of the many possible variations of the data and query transformations presented in this paper, e.g., the transformation of the two-dimensional intervals from trees R2 and R3 into three-dimensional points and the accompanying query transformation.

The performance experiments used synthetically generated workloads to investigate the behavior of the indexing techniques under different, controlled conditions, defined by a number of controlled parameters. Experiments with real-world workloads may shed new light on the relative performance of the indices. In addition, the use of the indexing techniques in data warehousing and mining scenarios is a topic of future research.

## Acknowledgements

Curtis E. Dyreson provided insightful comments on an early draft of the paper. This research was supported in part by the Danish Technical Research Council through grant 9700780 and by grants from the Danish Centre for IT Research and the Nykredit Corporation.

## References

- [1] N. Beckmann et al. The R\*-tree: An Efficient and Robust Access Method for Points and Rectangles. *Proceedings of ACM SIGMOD*, pp. 322–331, 1990.
- [2] E. Bertino et al. Indexing Techniques for Advanced Database Systems. *Kluwer Academic Publishers*, 1997.
- [3] R. Bliujūtė, C. S. Jensen, S. Šaltenis, and G. Slivinskas. R-tree Based Indexing of Now-Relative Bitemporal Data. *Proceedings of VLDB*, pp. 345–356, 1998.
- [4] R. Bliujūtė, C. S. Jensen, S. Šaltenis, and G. Slivinskas. Light-Weight Indexing of General Bitemporal Data. *TIME-CENTER TR-30*, 1998. Extended version of this paper.
- [5] R. Bliujūtė, S. Šaltenis, G. Slivinskas, and C. S. Jensen. Developing a DataBlade for a New Index. *Proceedings of ICDE*, pp. 314–323, 1999.
- [6] J. Clifford et al. On the Semantics of "NOW" in Databases. *ACM TODS*, 22(2):171–214, 1997.
- [7] A. Guttman. R-trees: A Dynamic Index Structure for Spatial Searching. *Proceedings of ACM SIGMOD*, pp. 47–57, 1984.
- [8] J. M. Hellerstein, J. F. Naughton, and A. Pfeffer. Generalized Search Trees for Database Systems. *Proceedings of VLDB*, pp. 562–573, 1995.
- [9] C. S. Jensen et al. A Consensus Glossary of Temporal Database Concepts. In *Temporal Databases: Research and Practice*, O. Etzion, S. Jajodia, and S. Sripada (eds), Springer-Verlag, pp. 367–405, 1998.
- [10] C. S. Jensen and R. Snodgrass. Temporal Data Management. *IEEE TKDE*, 11(1):36–44, 1999.
- [11] I. Kamel and C. Faloutsos. Hilbert R-tree: An Improved R-tree Using Fractals. *Proceedings of VLDB*, pp. 500–509, 1994.
- [12] A. Kumar, V. J. Tsotras, and C. Faloutsos. Access Methods for Bitemporal Databases. In *Recent Advances in Temporal Databases*, J. Clifford, and A. Tuzhilin (eds), Springer-Verlag, pp. 235–254, 1995.
- [13] A. Kumar, V. J. Tsotras, and C. Faloutsos. Designing Access Methods for Bitemporal Databases. *IEEE TKDE*, 10(1):1–20, 1998.
- [14] S. T. Leutenegger and M. A. Lopez. The Effect of Buffering on the Performance of R-Trees. *Proceedings of ICDE*, pp. 164–171, 1998.
- [15] M. A. Nascimento, M. H. Dunham, and R. Elmasri. M-IVTT: An Index for Bitemporal Databases. *Proceedings of DEXA*, pp. 779–790, 1996.
- [16] M. A. Nascimento and M. H. Dunham. Indexing Valid Time Databases via B<sup>+</sup>-Trees. *IEEE TKDE*, 11(6):929–947, 1999.
- [17] H. Samet. *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, 1990.
- [18] R. T. Snodgrass. Temporal Databases. *IEEE Computer*, 19(9):35–42, 1986.
- [19] R. T. Snodgrass. The Temporal Query Language TQuel. *ACM TODS*, 12(2):247–298, 1987.
- [20] R. T. Snodgrass et al. The TSQL2 Temporal Query Language. *Kluwer Academic Publishers*, 1995.
- [21] R. T. Snodgrass et al. Adding Valid Time to SQL/Temporal. ANSI X3H2-96-501r2, ISO/IEC JTC 1/SC 21/WG 3 DBL-MAD-146r2, 1996.
- [22] T. K. Sellis, N. Roussopoulos, and C. Faloutsos. The R+-Tree: A Dynamic Index for Multi-Dimensional Objects. *Proceedings of VLDB*, pp. 507–518, 1987.
- [23] B. Salzberg and V. J. Tsotras. Comparison of Access Methods for Time-Evolving Data. *ACM Computing Surveys*, 31(2):158–221, 1999.
- [24] V. J. Tsotras, C. S. Jensen, and R. T. Snodgrass. An Extensible Notation for Spatiotemporal Index Queries. *ACM SIGMOD Record*, 27(1):47–53, 1998.