

MASTER'S THESIS

# Main-Memory Operation Buffering for Efficient R-Tree Update

DE4 project group d635a  
Department of Computer Science  
Aalborg University

Supervisor  
Simonas Šaltenis  
simas@cs.aau.dk

Laurynas Biveinis  
lauras@cs.aau.dk

June 8, 2007

## Abstract

Emerging communication and sensor technologies enable new applications of database technology that require database systems to efficiently support very high rates of spatial-index updates. Previous works in this area exhibit one or more of the following drawbacks: they require the availability of large amounts of main memory, they do not exploit all the main memory that is indeed available, or they fail to support some of the standard index operations.

Assuming a setting where the index updates need not be written to disk immediately, we propose an R-tree-based indexing technique that does not exhibit any of these drawbacks. This technique exploits the buffering of update operations in main memory as well as the grouping of operations to reduce disk I/O. In particular, operations are performed in bulk so that multiple operations are able to share I/O. The emptying of buffers is controlled by means of an analytical cost model.

The thesis reports on empirical studies that demonstrate that the cost model is effective in controlling the buffer emptying. The studies also show that, in terms of update I/O performance, the proposed technique improves on state of the art in settings with frequent updates: it is significantly more efficient than an LRU-cached R-tree and a previous proposal for an update-efficient R-tree.

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
<b>2</b>	<b>Problem Setting</b>	<b>5</b>
2.1	Data and Update Operations . . . . .	5
2.2	The R-tree . . . . .	6
<b>3</b>	<b>Data Structure and Algorithms</b>	<b>7</b>
3.1	The $R^R$ -tree Data Structure . . . . .	7
3.2	Insertion and Deletion Operations . . . . .	8
3.3	EmptyBuffer Algorithm . . . . .	10
3.4	GroupUpdate Algorithm . . . . .	10
3.5	IntegrateChild Algorithm . . . . .	12
3.6	GetPushDownOps Algorithm . . . . .	13
3.7	Example . . . . .	14
3.8	Search Algorithm . . . . .	16
<b>4</b>	<b>Analytical Cost Modeling</b>	<b>17</b>
4.1	Frequency of EmptyBuffer Invocations . . . . .	17
4.2	EmptyBuffer Cost Model . . . . .	19
<b>5</b>	<b>Experimental Evaluation</b>	<b>21</b>
5.1	Experimental Setup . . . . .	21
5.2	Workloads . . . . .	22
5.3	Cost Model Validation . . . . .	23
5.3.1	Frequency of Annihilation . . . . .	23
5.4	Comparing the Indexes . . . . .	24
5.4.1	Performance in a Persistent Setting . . . . .	27
5.5	Exploring the $R^R$ -tree Algorithms . . . . .	27
5.5.1	Memory Division Experiments . . . . .	27
5.5.2	Effect of Query Piggybacking . . . . .	29
5.5.3	Group Size Statistics . . . . .	29
5.5.4	Tree Quality . . . . .	32
<b>6</b>	<b>Related Work</b>	<b>33</b>
<b>7</b>	<b>Conclusion and Research Directions</b>	<b>35</b>
7.1	Research Directions . . . . .	36

## List of Figures

1	R-tree with maximum fanout 3, minimum fanout 2 . . . . .	7
2	Possible states of the index for the $\langle id, p \rangle$ tuple . . . . .	8
3	The UPDATE algorithm . . . . .	9
4	The EMPTYBUFFER algorithm . . . . .	10
5	The GROUPUPDATE algorithm . . . . .	11
6	The INTEGRATECHILD algorithm . . . . .	12
7	The INSERTSUBTREE algorithm . . . . .	13
8	The GETPUSHDOWNOPS algorithm . . . . .	13
9	The GROUPOPERATIONS algorithm . . . . .	14
10	Performing a set of operations on the $R^R$ -tree . . . . .	15
11	The SEARCH algorithm . . . . .	16
12	The PIGGYBACK algorithm . . . . .	17
13	Cost model validation, 5% buffer . . . . .	23
14	Cost model validation, 20% buffer . . . . .	24
15	Validation of annihilation assumptions . . . . .	24
16	Scalability with respect to available main memory . . . . .	25
17	I/O performance versus the update/query ratio . . . . .	26
18	Scalability with respect to workload duration . . . . .	26
19	Scalability with respect to disk page size . . . . .	27
20	Scalability with respect to main memory in a persistent setting .	28
21	I/O performance with varying main-memory divisions in a per- sistent setting . . . . .	28
22	I/O performance for varying main-memory divisions . . . . .	29
23	Frequency of piggybacking . . . . .	30
24	Group size distributions, whole buffer emptying . . . . .	30
25	Group size distributions, best threshold emptying . . . . .	31
26	Group size distributions, largest groups emptying . . . . .	31
27	Root level group size distributions . . . . .	32
28	Disk R-tree quality . . . . .	33

# 1 Introduction

As sensor, communication, and computing technologies continue to evolve, it is becoming increasingly feasible and attractive to monitor, record, and query data that capture the evolving states of continuously changing real-world phenomena.

As an example that is representative of this new class of update-intensive applications, consider advanced location-based services, e.g., in the context of intelligent transport systems, where the positions of a large population of GPS-equipped moving objects are tracked on a server. To maintain accurately the current positions of such objects, the positions must be updated frequently. Experiments with GPS data from vehicles traveling in semi-urban environments show that to know the positions of the vehicles being tracked with an accuracy of 200 meters requires an update from each vehicle on average every 15 seconds [21].

The efficiency of the processing of updates determines the maximum update throughput, which in turn determines the maximum number of objects that can be tracked with a given accuracy.

Spatial indexes were originally developed with applications in mind that were characterized by relatively static data, and focus was on the efficient support for queries. In particular, the predominant spatial index, the R-tree [3, 11], supports queries efficiently, but updates are far from being efficient enough to support applications such as the one that involves the tracking of a large population of moving objects [15].

To improve the efficiency of spatial-index updates, a number of approaches have recently been proposed. The so-called bottom-up approaches aim to reduce the need for the expensive top-down tree traversals involved in R-tree deletions. They do so by offering direct leaf-level access [14]. Another approach is to try to avoid performing deletions on disk, by accumulating them in main-memory [24]. Yet another type of approach is to buffer and group operations in main memory with the purpose of enabling operations to share I/O's [1, 15].

With few exceptions [25], previous approaches assume a *persistent* index, such that, upon finishing the processing of an update, the update is recorded in the index on disk. This means that an update costs at least one I/O. In the types of applications we consider, even this cost is too high to support a high throughput of updates. The high update rates offer an opportunity to soften the index-persistence requirement: even if some recent updates are lost due to a system crash, after only a short waiting period, fresh updates will arrive for the objects affected. We note that existing approaches can achieve sub-one I/O updates by simply using a write-back disk page buffer, where dirty pages are not committed to disk after each operation. However, this use of the available main-memory is inferior when the objective is to speed up updates—the thesis offers experimental results that demonstrate this.

Assuming the softer persistency requirements, this thesis explores the efficient use of all available main memory to reduce the I/O costs of R-tree update operations as much as possible. The main contribution is a data structure, named the  $R^R$ -tree, that consists of a main-memory R-tree and a conventional disk-based R-tree. The former enables efficient search on a main memory buffer of update operations, i.e., both insertions and deletions.

When the operation buffer gets full, all operations in the buffer are clustered into groups. The operations that belong to groups above a certain size are then performed in bulk on the disk R-tree in order to minimize the I/O. The thesis

includes an analytical study that aims to identify the minimum size for the groups with operations that should be performed in bulk.

An empirical study is reported. A setting is assumed in which the positions of moving objects are known only with some predefined accuracy. Thus, a moving object issues an update when its current position deviates by the predefined accuracy from the position it reported most recently. This setting improves on those where accurate positions are assumed. The empirical study demonstrates that the analytical modeling is effective in choosing the appropriate group size, and it offers insights into the performance properties of the  $R^R$ -tree. In particular, the study shows that the  $R^R$ -tree is able to significantly outperform an  $R$ -tree that uses an LRU buffer, which is the current state of the art. The study also shows that the  $R^R$ -tree significantly outperforms a recent proposal for an update-efficient  $R$ -tree.

In contrast to previous research on I/O efficient update of  $R$ -trees, the  $R^R$ -tree does not *require* some minimum amount of main memory to work. Thus, if no main memory is available, the  $R^R$ -tree works (largely) as a regular  $R$ -tree. If all data fits in main memory, the  $R^R$ -tree works as a main-memory  $R$ -tree. This behavior makes the proposal quite flexible.

Additionally, the proposed data structure is compatible with other main-memory buffering techniques, e.g. LRU buffering. Consequently it is possible to tune the data structure for a particular (not necessarily hyper-dynamic) workload by dividing the available main memory between several buffers.

Finally, because the  $R$ -tree heuristics are treated as black-boxes, the proposed approach is not limited to  $R$ -trees, but is easily extensible to the wider class of so-called grow-post trees, which includes variants of the  $R$ -tree [16].

The thesis starts with a formulation of the problem setting in Section 2. The proposed data structure, algorithms, and their properties are described in Section 3, and a cost model is presented in Section 4. The results of an experimental evaluation are reported in Section 5. Section 6 then covers related work. Finally, Section 7 concludes and identifies directions for future research.

## 2 Problem Setting

### 2.1 Data and Update Operations

Consider a location-based service or fleet management application where a server tracks the current positions of a population of moving objects. It is a fundamental observation that, due to the discrete sampling of continuous movements and the inherent inaccuracy of positioning technologies, the positions of the objects cannot be known accurately at the server. Instead, the position of an object  $o$  may be assumed to be known with an accuracy  $thr_o$  [21, 22].

The inaccuracy was only recently taken into account in moving-object indexing research [13, 12]. In particular, two approaches may be used that take into account the inaccuracy of positions while still achieving perfect recall of queries, meaning that queries are guaranteed to return all qualifying objects, possibly in addition to false positives. With *data enlargement*, the position of an object  $o$  is represented by a circle with radius  $thr_o$ , which can be replaced by a bounding square for indexing purposes. With the less flexible *query enlargement*, a global maximum accuracy threshold  $thr$  is used, and each query region

is enlarged according to this threshold.

We assume data enlargement; thus, the database stores pairs consisting of an object id,  $objId$ , and the object’s current position,  $p$ , represented by a square, which as a special case can degenerate to a point. Such an  $\langle objId, p \rangle$  pair is termed an *identification tuple*.

We classify database update operations as object insertions, object deletions, and object updates. When an object reports its position for the first time, and insertion occurs: a new identification tuple  $\langle objId, p_{new} \rangle$  is inserted into the index. When an object stops being tracked by the system, a deletion occurs: the identification tuple  $\langle objId, p_{old} \rangle$  is deleted from the index. Updates are the most common operations in the workloads we consider. These consist of a pair of an index deletion and an index insertion that deletes the old identification tuple  $\langle objId, p_{old} \rangle$  and inserts a new identification tuple  $\langle objId, p_{new} \rangle$ . In the following, we consider only index-level operations, so an (index) update operation denotes either an insertion or a deletion.

To simplify the presentation, we assume that the sequences of operations are valid. Thus, the operations concerning an object  $objId$  come in the following order:  $ins\langle objId, p_1 \rangle, del\langle objId, p_1 \rangle, ins\langle objId, p_2 \rangle, \dots, del\langle objId, p_{i-1} \rangle, ins\langle objId, p_i \rangle$ .

As mentioned in the introduction, we do not require the index to be fully persistent. It is assumed that after a crash, the index is rebuilt by waiting until the objects report their positions, which will happen in a relatively brief time interval in a scenario with frequent updates. Alternatively, the proposed index can be used as a secondary index on otherwise persistent data. The index can then be rebuilt from scratch in the event of a crash.

Although the proposed data structure supports the same queries as an R-tree, for brevity, we focus on range queries that, given a query rectangle, return all objects such that their true positions intersect the rectangle.

## 2.2 The R-tree

The R-tree is a height-balanced tree [11]. Each node consists of a set of entries. The identification tuples are stored as entries in the leaf nodes. The  $objId$  values of these tuples may then be interpreted as pointers to additional data for the objects. Non-leaf node entries take the same form as the leaf-node entries, but an identifier now points to a tree node at the next level of the tree, and the position is a minimum bounding rectangle (MBR) that contains the (approximations of the) positions of all objects stored in the subtree pointed to.

The size of a node is that of a disk page. This size determines the maximum fan-out of the index, which is the maximum number of entries that fit in a node. The actual fan-out of any non-root node is between the maximum and some value that is no larger than the ceiling of the maximum divided by 2. The actual fan-out of the root is at least 2. The R-tree is equipped with update algorithms that ensure these properties. The algorithms eliminate underfull nodes and reinsert their entries, and they split overfull nodes.

Figure 1 shows an example of an R-tree that indexes the positions of 9 objects.

The prototypical query supported by the R-tree is the range query: given a rectangle, this query retrieves all objects with positions that may overlap with this rectangle.

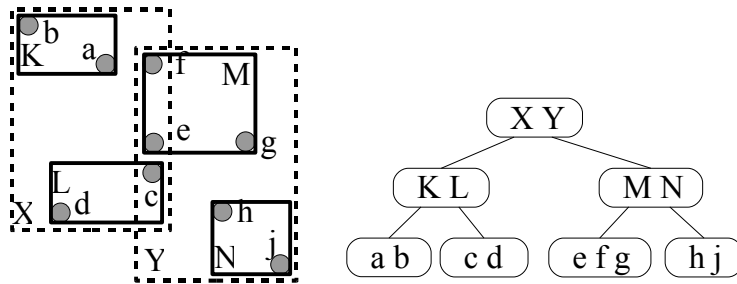


Figure 1: R-tree with maximum fanout 3, minimum fanout 2

The R\*-tree [3] is a variant of the R-tree that has the same data structure as the R-tree but is equipped with update algorithms that render it more efficient than the R-tree.

### 3 Data Structure and Algorithms

Section 3.1 introduces the data structure and its properties. Sections 3.2–3.6 describe the update algorithms and their subroutines, simultaneously proving selected properties. Section 3.7 illustrates the workings of these algorithms by means of an example. Finally, Section 3.8 covers range search.

#### 3.1 The R<sup>R</sup>-tree Data Structure

The R<sup>R</sup>-tree data structure consists of two R-trees: a disk-based tree and an *operation-buffer* tree in main memory. The disk-based tree is a standard R-tree with the identification tuples introduced in Section 2 as leaf-node entries. In contrast, the operation-buffer tree stores operations yet to be performed on the disk-based tree. Its leaf-node entries thus consist of identification tuples augmented with *deletion flags* that indicate whether an entry represents an insertion or a deletion.

The operation buffer has room for  $C_{max}$  pending insertion and deletion operations, and the current number of operations in the buffer is denoted by  $C$ ,  $C \leq C_{max}$ . Since flushing of the operation buffer to the disk-based tree is done only at certain points in time, the buffer might contain operations that invalidate some of the data in the disk tree. The property below characterizes the relation between the data on disk and the operations in the buffer.

**Property 1** *For any given identification tuple  $\omega = \langle objId, p \rangle$  from the sequence of operations, exactly one of the following four cases is true at any given time. If the last operation concerning  $\omega$  was a deletion,*

1.  $\omega$  is not in the disk tree, and no entry with  $\omega$  is in the buffer, or
2.  $\omega$  is in the disk tree and a deletion entry with  $\omega$  is in the buffer.

*If the last operation concerning  $\omega$  was an insertion,*

3.  $\omega$  is not in the disk tree, and an insertion entry with  $\omega$  is in the buffer, or

4. there is an entry with  $\omega$  in the disk tree, and there is no entry with  $\omega$  in the buffer.

Note that if no operations concerning  $\omega$  have yet been performed, case 1 of the property holds trivially.

Figure 2 illustrates the four cases and the possible transitions between them. Note that the property is not formulated with respect to  $objId$  or  $p$  alone. There may naturally exist several objects in the index that share the same location (equal  $p$  values). Further, several obsolete locations for the same object (equal

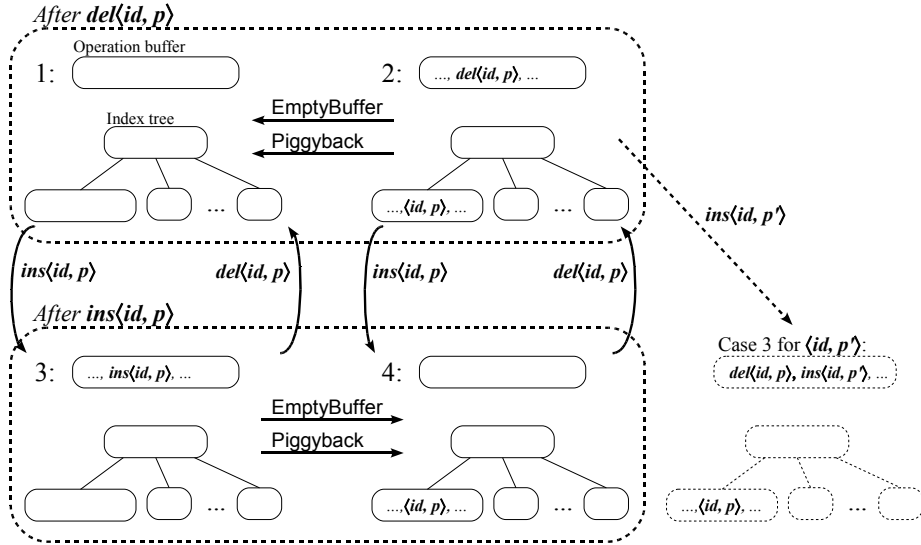


Figure 2: Possible states of the index for the  $\langle id, p \rangle$  tuple

$objId$  values) might be recorded in the index together with the object's current location. This case is illustrated in the bottom-right corner of the figure, where  $p'$  is the current position of the object  $id$  and  $p$  is an old position.

### 3.2 Insertion and Deletion Operations

Insertions and deletions are implemented by the same UPDATE algorithm, using the  $delFlag$  parameter to distinguish the two types of operations (see Figure 3). This algorithm first checks whether the operation buffer ( $buf$ ) has room for the current update operation and empties part of the buffer if necessary (EMPTYBUFFER). Then it checks if the opposite operation already exists in the buffer by querying it. An opposite operation is defined as the operation with the same  $p$  and  $objId$ , but with the opposite deletion flag value. If such an operation is found, the previous operation is removed from the buffer and the incoming operation is ignored. We call this an *annihilation*. Otherwise, if the opposite operation is not found, the incoming operation is inserted into the operation buffer.

The lemma that follows implies that the annihilation algorithm is correct with respect to Property 1.



```

UPDATE(objId, p, delFlag)
1  if buf is full
2    then EMPTYBUFFER
3  oppositeEntry  $\leftarrow$   $\langle$ objId, p,  $\neg$ delFlag $\rangle$ 
4  if buf.exists(oppositeEntry)
5    then buf.remove(oppositeEntry)
6    else buf.insert( $\langle$ objId, p, delFlag $\rangle$ )

```

Figure 3: The UPDATE algorithm

**Lemma 1** *Property 1 holds after invocation of the UPDATE algorithm if  $C < C_{max}$  holds as a precondition (the buffer is not full and EMPTYBUFFER is not called).*

**Proof** Proof by induction. In the case of a newly created, empty tree, Property 1 holds trivially. Assume that Property 1 holds before an UPDATE. Next assume that we are about to perform the first update (i.e., an insertion) for  $\langle objId, p \rangle$  or that the most recent update for this pair was a deletion. As illustrated in Figure 2, the index is then either in state 1 or in state 2, the update to be performed next is an insertion, and this update will transform the index to state 3 or state 4. If the most recent update was an insertion, the index is either in state 3 or in state 4, the update to be performed next is a deletion, and this update will transform the index to state 1 or state 2.

Processing an insertion when the index is in state 1 is trivial. The incoming operation is added to the buffer, transforming the index into state 3. Similarly, a deletion in state 4, is simply added to the buffer, transforming the index into state 2.

Processing a deletion when the index is in state 3 is the most common case of annihilation. The incoming deletion causes the removal of the corresponding entry in the buffer, transforming the state of the structure into state 1.

Processing an insertion when the index is in state 2 yields a relatively rare case of annihilation. For example, this may occur when a stationary object repeatedly sends updates. In this case, an insertion arrives when there is already a deletion in the buffer for the same object at the same position. The annihilation removes both operations, transforming the tree into state 4. **Lemma is proved.**

While the annihilation in state 3 is intuitive, to better understand the intuition of annihilation in state 2, consider what would happen if the involved operations were allowed to complete on the disk tree one by one. The deletion from the buffer first removes the data from the tree; then, later, the insertion creates an entry that is identical to the one removed. The combined effect is that the tree transforms into state 4, the effect of which is identical to that of the annihilation.

### 3.3 EmptyBuffer Algorithm

The EMPTYBUFFER algorithm, presented in Figure 4<sup>1</sup>, is responsible for committing some of the operations in the buffer to disk. Note that although all operations are removed from the buffer in line 2 of the pseudo code, some of them are usually put back into the buffer in the GROUPUPDATE algorithm, described in Section 3.4. In addition, with the appropriate bookkeeping, the implementation of the algorithm needs not materialize the *opList* as is done in the pseudo code.

```
EMPTYBUFFER()  
1 opList ← buf.GetAllOperations  
2 buf.remove(opList)  
3 siblings ← GROUPUPDATE(rootNode, opList)  
4 while  $|siblings| > 1$   
5 do rootNode ← newNodeWithEntries(siblings)  
6   siblings ← GROUPSPLIT(rootNode)
```

Figure 4: The EMPTYBUFFER algorithm

The GROUPUPDATE algorithm performs the argument operations on the root node of the disk tree and returns a set of root-level sibling nodes, i.e., subtrees of the same height as the disk tree, including the original root. If this set contains other nodes than the original root, the tree is grown by creating a new root. It is possible that there were so many sibling entries that the new root overflows. In that case, the new root is split, which grows the tree further.

The GROUPSPLIT algorithm is used at line 6 (and additionally as a subroutine of GROUPUPDATE) to check whether the node is overflowing. If not, it returns the original node; otherwise, the node is repeatedly split by the regular R\*-tree split algorithm [3] until none of the resulting nodes are overflowing, in which case all the nodes created are returned. Although more advanced implementations of GROUPSPLIT are possible [4], we do not consider them. This is because the node is usually split no more than once, as the number of buffer entries processed is typically a small fraction of the size of the tree disk.

### 3.4 GroupUpdate Algorithm

The GROUPUPDATE, presented in Figure 5, is the core of the proposed indexing technique. It executes the specified operation list (a sublist of the operations formerly present in the buffer) on the specified subtree.

If the subtree is a leaf node then the operation list is executed on the node trivially by means of EXECUTEENTRIES, which just adds insertion entries to the node and removes existing entries matched by deletion entries. If the subtree is rooted at a non-leaf node, GETPUSHDOWNOPS returns a mapping from its child nodes to the individual operation sublists that should be executed on these nodes. As will be discussed later, some of the operations from *opList* can also be put back into the buffer and are thus not processed during this buffer emptying.

---

<sup>1</sup>For brevity, we do not distinguish between nodes and entries pointing to these nodes in the pseudo code descriptions of the algorithms. Consequently, all MBR update and node reading and writing operations are implicit.

```

GROUPUPDATE(node, opList)
1  if node is leaf
2    then EXECUTEENTRIES(node, opList)
3    else pushG  $\leftarrow$  GETPUSHDOWNOPS(opList, node)
4          regroup  $\leftarrow$  false
5          for each  $\langle$ child, opSublist $\rangle$  in pushG
6            do if regroup
7              then return GROUPUPDATE(node, pushG)
8              pushG  $\leftarrow$  pushG  $\setminus$  opSublist
9              node.remove(child)
10             newCh  $\leftarrow$  GROUPUPDATE(child, opSublist)
11             if newCh = {newChild}
12               then regroup
13                    $\leftarrow$  INTEGRATECHILD(node, newChild)
14             else node.add(newChildren)
15 if |node| = 1 and node is non-leaf
16   then node  $\leftarrow$  node[1]
17 return GROUPSPLIT(node)

```

Figure 5: The GROUPUPDATE algorithm

Each mapping in *pushG* is then executed in turn by a recursive GROUPUPDATE call (line 10). This results in a list of new child nodes in *newCh*. The rest of the algorithm is concerned with integrating these nodes into the tree. In the simplest case, there are two or more such nodes. Their entries are trivially added to the current node. If only one child node is returned, further and more complicated treatment may be necessary in case the child is underfull or has shrunk in height. This is handled by INTEGRATECHILD and is discussed in the next section. This algorithm may change other sibling child entries, invalidating the remaining, non-processed mappings in *pushG*. This is signaled by the *regroup* flag, which forces a restart of GROUPUPDATE. It is guaranteed that at least one mapping is completed before such a restart can occur; thus, the algorithm is guaranteed to terminate. Note that the restart may occur with arbitrarily few remaining operations. As confirmed later experimentally, this may cause very small groups appear on the upper tree levels.

After performing all the operations, the current node might contain only one entry, in which case it is replaced by its child in the tree, thus shrinking the tree in height. The node may also end up overfull. It is then passed to GROUPSPLIT, which returns a list of non-overflowing nodes. This list is returned as the result of the algorithm.

The next lemma states that the UPDATE algorithm complies with Property 1 when EMPTYBUFFER is called. The proof assumes that those operations that reach the leaves of the disk tree are correctly routed there by the GETPUSHDOWNOPS algorithm.

**Lemma 2** *Property 1 holds after invocation of the UPDATE algorithm if  $C = C_{max}$  holds as a precondition (the buffer is full).*

**Proof** According to the Lemma 1, Property 1 holds as a precondition. For every identification pair  $\omega$  that is contained in the buffer, the tree might be in either state 2 or state 3. EMPTYBUFFER will either leave the  $\omega$  in the buffer, or the operation will be performed correctly on the disk tree and removed from the buffer. In the latter case, the index will move from state 2 to state 1, or from state 3 to state 4 (see Figure 2). After buffer emptying, the incoming operation will be processed as in Lemma 1. **Lemma is proved.**

A consequence of Lemmas 1 and 2 is the following theorem. The proof is trivial and thus omitted.

**Theorem 1** *After invocation of the UPDATE algorithm, Property 1 holds.*

### 3.5 IntegrateChild Algorithm

The INTEGRATECHILD algorithm is responsible for adding an arbitrarily-looking R-tree to the current node. Several cases must be handled. One case is where the root of the new tree is the only child of the current node. Then the current node is removed and the child node takes its place. More commonly, the root of the new tree has siblings in the current node and is either underfull or normal. The contents of an underfull node are merged into the sibling nodes by the MERGESUBTREES call, whereas a normal node is inserted into tree by INSERTSUBTREE, as discussed shortly.

```

INTEGRATECHILD(node, newCh)
1  if |newCh|  $\neq$  0
2    then if |node| = 0
3        then node  $\leftarrow$  newCh
4        else if newCh is underfull
5            then
6                return MERGESUBTREES(node, newCh)
7            else
8                return INSERTSUBTREE(node, newCh)
9  return mustNotRegroup

```

Figure 6: The INTEGRATECHILD algorithm

Subroutines INSERTSUBTREE and MERGESUBTREES serve similar purpose—putting a subtree into the main tree—and have similar implementations. The principal difference is that INSERTSUBTREE handles a subtree whose root is a normal (not underfull) node, whereas MERGESUBTREES handles the case of an underfull node.

The INSERTSUBTREE (Figure 7) checks whether the specified node is at the right level for insertion of the subtree, i.e., the distance to the leaf level from the node in the main tree must be equal to the subtree height.

If so, a new entry for the subtree is created and inserted into the node. This case is the most common, and it does not incur any additional I/O. However, if the specified node is not at the right level to accommodate the subtree, then, using the minimum bounding rectangle of the subtree, the standard R\*-tree CHOOSESUBTREE algorithm is called to identify an appropriate node one level

```

INSERTSUBTREE(node, subtree)
1  if subtree is one level below node
2    then node.add(subtree)
3  else childNode ← CHOOSESUBTREE(node, subtree)
4      INSERTSUBTREE(childNode, subtree)
5      childNodes ← GROUPSPLIT(childNode)
6      node.add(childNodes)
7      if |childNodes| > 1
8        then return mustRegroup
9  return mustNotRegroup

```

Figure 7: The INSERTSUBTREE algorithm

below, and INSERTSUBTREE is called recursively. After returning from the recursive call, GROUPSPLIT is called on the potentially overflowing child node, and the resulting entries are inserted into the parent node.

Since GROUPSPLIT may replace one child node with two new child nodes, it invalidates any parent GROUPUPDATE mappings for this node. This situation is signaled to the caller by the return value.

The MERGESUBTREES algorithm is very similar to INSERTSUBTREE. The only difference is that when the node at the right level is found, it is merged with the underfull node passed as an argument to the algorithm.

Finally, to complete the definition of the update algorithms, the GETPUSH-DOWNOPS algorithm is discussed next. It is used in algorithm GROUPUPDATE to route groups of operations down the disk tree.

### 3.6 GetPushDownOps Algorithm

Given an operation list and a node of the disk tree, the GETPUSHDOWNOPS algorithm (Figure 8) determines which of the operations should be pushed down and to which children of the node. First, the GROUPOPERATIONS subroutine

```

GETPUSHDOWNOPS(node, opList)
1  mapping ← GROUPOPERATIONS(node, opList)
2  if node is the root
3    then for each  $\langle chNode, sublist \rangle$  in mapping
4      do if |sublist| < k
5        then
6          mapping ← mapping \  $\{ \langle chNode, sublist \rangle \}$ 
7          buf.insert(sublist)
8  return mapping

```

Figure 8: The GETPUSHDOWNOPS algorithm

is called to establish the mapping between the operations and the child nodes according to the usual R\*-tree rules (see Figure 9). The mapping returned is a list of  $\langle childNode, opSublist \rangle$  tuples. Each insertion is mapped to a single child node according to the R\*-tree's CHOOSESUBTREE algorithm [3]. Each deletion is

```

GROUPOPERATIONS(node, opList)
1  opMapping  $\leftarrow$   $\emptyset$ 
2  for each op in opList
3  do if op is a deletion
4      then matchingChildren  $\leftarrow$  Covering(node, op)
5          for each chNode in matchingChildren
6              do opMapping  $\leftarrow$  opMapping  $\cup$  {op  $\rightarrow$  chNode}
7      else childNode  $\leftarrow$  CHOOSESUBTREE(node, op)
8          opMapping  $\leftarrow$  opMapping  $\cup$  {op  $\rightarrow$  childNode}
9  return opMapping

```

Figure 9: The GROUPOPERATIONS algorithm

mapped to all child nodes whose bounding rectangle includes it (as returned by the *Covering* method in line 4 of the algorithm).

As shown in Figure 8, if the node is not the root, the algorithm just returns the mapping obtained from the call to GROUPOPERATIONS. If the node is the root, the algorithm removes from the mapping all operation sublists that have fewer than  $k$  operations. Operations in these sublists are put back into the buffer. The idea using an operation threshold for the buffer emptying is to avoid spending expensive I/O on the small groups, so that as many operations as possible share the same I/O.

Choosing a threshold  $k$  is not trivial. If a large value for  $k$  is chosen, this will cause only very few groups to be emptied, leading in turn to the need for very frequent buffer emptying. Choosing a small  $k$  value also has obvious drawbacks. The analytical cost modeling in Section 4 provides a way to choose an optimal value of  $k$ .

Note that because some sublists of operations are put back into the buffer, one copy of a deletion may be put back into the buffer while another is sent down the tree. For this reason, if the deletion succeeds to delete an entry in the index tree, the EXECUTEENTRIES algorithm (called from GROUPUPDATE) removes the corresponding copy of the deletion operation from the buffer if there is one.

Two special cases are not shown in the pseudo code. If all sublists of operations have fewer than  $k$  operations, one largest sublist is sent down the tree. Finally, a special and rare case occurs when only groups, consisting solely of copies of deletions, are sent down the tree. If none of these deletions are successful, no operations are removed from the buffer. If this occurs, we empty the whole buffer.

Having described all the algorithms needed for maintaining the  $R^R$ -tree, we proceed to illustrate the workings of the algorithms.

### 3.7 Example

Figure 10 shows a small example where a number of update operations are performed, causing a buffer emptying. The leftmost part of the figure shows the positions of eight objects stored in an  $R$ -tree. The four empty circles ( $a_1, a_2, c_2, j$ ) represent insertions that are in the buffer or future updates that will be dis-

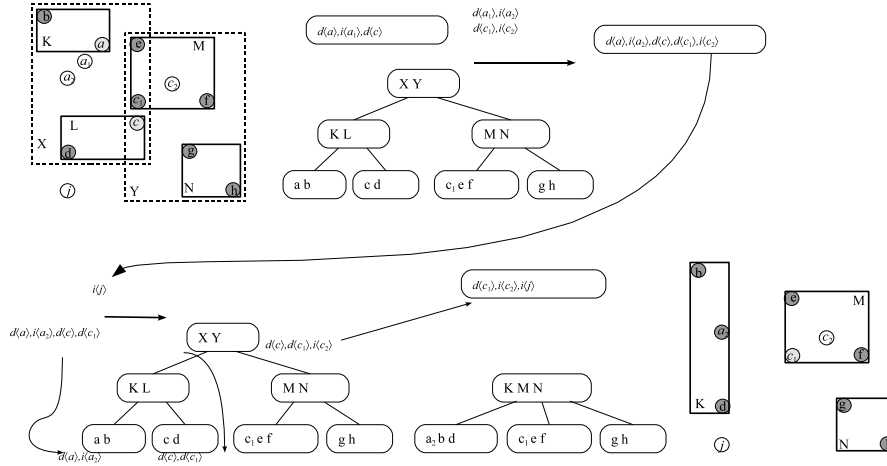


Figure 10: Performing a set of operations on the  $R^R$ -tree

cussed in the following. The light gray circles  $(a, c)$  represent the positions that are in the disk tree, but that have corresponding deletion entries in the buffer. In the example, object  $a$  moves from  $a$  to  $a_1$  to  $a_2$ , and object  $c$  moves from  $c$  to  $c_1$  to  $c_2$ . We assume a disk R-tree with a maximum fan-out of 3 and a minimum fan-out of 2. The buffer size is  $C = 5$  and the operation threshold is  $k = 4$ .

Note that the initial state of the tree holds two positions,  $c$  and  $c_1$ , for object  $c$ . This happened after the part of the buffer containing the insertion of  $c_1$  was emptied, while the deletion of  $c$  was left in the buffer. Next,  $a$  and  $c$  update their positions, resulting in the following sequence of updates:  $del\langle a_1 \rangle$ ,  $ins\langle a_2 \rangle$ ,  $del\langle c_1 \rangle$ ,  $ins\langle c_2 \rangle$ . The deletion of  $a_1$  and the corresponding insertion already in the buffer annihilate one another. The remaining three operations are inserted into the buffer.

The insertion of a new object  $j$  triggers a buffer emptying. The figure shows how the GROUPOPERATIONS algorithm divides the operations into two groups: one for node  $X$  and one for node  $Y$ . Note that, due to the overlap between the MBRs of  $X$  and  $Y$ ,  $del\langle c \rangle$  and  $del\langle c_1 \rangle$  are copied into both groups. Because the group for node  $Y$  has fewer than  $k$  elements, its entries are put back into the buffer. The operations in the other group proceed down the tree to nodes  $K$  and  $L$ , as shown in the figure. Note that in node  $L$ ,  $del\langle c_1 \rangle$  is unsuccessful, but  $del\langle c \rangle$  is successful, and its copy is removed from the buffer.

The performed operations result in a number of structural changes to the disk tree. Node  $L$  is rendered underfull and is merged with node  $K$ . As a result, node  $X$  contains only one child; thus,  $X$  is replaced by its child, reducing the height of the subtree to one. The shorter subtree is then inserted into node  $Y$  using the INSERTCHILD algorithm. Finally, the single-entry root is removed, finishing the buffer emptying. At last, the operation that caused the buffer emptying,  $ins\langle j \rangle$ , is inserted into the buffer. The resulting disk tree and buffer are shown at the rightmost end of the figure.

Note that the demonstrated shrinking (and growing) of subtrees are very

rare for realistic node sizes and workloads that mostly contain deletion-insertion pairs.

### 3.8 Search Algorithm

The SEARCH algorithm in Figure 11 is an adjusted standard R-tree search algorithm. It operates by first querying the disk tree. The resulting answer set is modified to take the buffer data into account. Thus, objects that have corresponding deletions in the buffer are removed from the set, and any objects from the buffer that overlap with the query rectangle are included into the set. As

```

SEARCH(rectangle)
1  leafNodes ← GETOVERLAPPINGLEAFNODES(rectangle)
2  answer ← ∅
3  for each leaf in leafNodes
4  do answer ← answer ∪ SEARCHNODE(leaf, rectangle)
5     PIGGYBACK(node)
6     for each entry in answer
7     do if there is a matching deletion in the operation buffer
8         then answer ← answer \ {entry}
9     return answer ∪ buf.SearchInsertions(rectangle)

```

Figure 11: The SEARCH algorithm

the following theorem shows, the resulting set is the correct answer.

**Theorem 2** *The SEARCH algorithm, performed after a series of calls to the UPDATE algorithm, returns correct results.*

**Proof** Assume that the current position of object *objId* is *pc*. In a series of update operations concerning object *objId*, the last operation for any identification tuple  $\langle \textit{objId}, p \rangle$  ( $p \neq pc$ ) is a deletion. According to Theorem 1, the index is in state 1 or in state 2 with respect to these identification tuples (see Figure 2). Trivially, the query will not return any such identification tuples corresponding to state 1. Identification tuples corresponding to state 2, even if chosen in line 4 of the algorithm, are subsequently eliminated by matching entries from the buffer in line 7.

Now let the last operation for  $\langle \textit{objId}, pc \rangle$  be an insertion. According to Theorem 1, the index is in state 3 or in state 4 with respect to this identification tuple (the tuple is either in the buffer or in the disk tree). Thus, if *pc* satisfies the query, it will be returned either from the disk tree (line 4) or the buffer (line 9). **Lemma is proved.**

In addition to querying, the SEARCH algorithm takes the opportunity to perform some of the operations in the buffer on the leaf nodes of the disk tree that are read during querying. This is done by calling PIGGYBACK presented in Figure 12.

Potentially all the insertion entries that are contained in the node MBR and all the deletion entries that can remove corresponding node entries might be performed and thus removed from the buffer. However, the number of operations



actually performed is limited by the requirement to keep the node from neither overflowing nor underflowing, to avoid the cost and complexity of having to adjust the tree.

```

PIGGYBACK(node)
1  matchingBufEntries ← SEARCHBUFFER(MBR(node))
2  i ← getNumOfInsertions(matchingBufEntries)
3  d ← getNumOfDeletions(matchingBufEntries)
4  newSize ← |node| + i - d
5  if newSize >  $S_{max}$ 
6      then i ← i - (newSize -  $S_{max}$ )
7      else if newSize <  $S_{min}$ 
8          then d ← d - ( $S_{min}$  - newSize)
9  EXECUTEFROMBUFFER(matchingBufEntries, node, i, d)

```

Figure 12: The PIGGYBACK algorithm

The operations are executed by means of an algorithm called EXECUTEFROMBUFFER, which is similar to EXECUTEENTRIES used in GROUPUPDATE, except that it takes two additional parameters that specify thresholds for the number of insertions and deletions that can be executed. It also removes the completed operations from the buffer.

In Figure 12,  $S_{min}$  and  $S_{max}$  is the minimum and the maximum number of allowed entries in an R-tree node, respectively.

Because such *piggybacking* corresponds to the emptying of part of the buffer, the reasoning in the proof of Lemma 2 also applies to show that PIGGYBACK does not violate Property 1. Thus, query-result correctness is not compromised. Also, note that piggybacking is designed to be a performance optimization that can be disabled without any effects in the correctness of SEARCH or any other algorithm.

## 4 Analytical Cost Modeling

As mentioned in Section 3.6, choosing the right value for the operation threshold ( $k$ ) is not trivial. The cost of buffer emptying has to be balanced against the frequency of the buffer emptying. This section presents a cost model that has as its goal to enable estimation of an optimal value for  $k$ .

This goal may be achieved without having to model the performance accurately in absolute terms. This in turn enables us to make a number of simplifying assumptions. We explain these assumptions at the beginning of Sections 4.1 and 4.2.

We first model how often the buffer is emptied and then estimate the I/O cost of the buffer emptying. The notation used in this section is summarized in Table 1.

### 4.1 Frequency of EmptyBuffer Invocations

To enable the estimation of the frequency of EMPTYBUFFER invocations, we make the following assumptions. We assume that the number of moving objects

$C$	Current number of operations in the buffer
$C_{max}$	Capacity of the buffer
$O$	Number of objects
$k$	Operation threshold
$S$	Average fan-out of disk tree nodes
$R$	Fan-out of the root
$X$	Sequence of update operations
$W$	Expected number of buffer entries
$K$	I/O costs associated with buffer emptying

Table 1: Cost Model Notation

being tracked is  $O$  and that  $C_{max} < O$ . We also assume that the index is in *steady state*, such that the set of objects being tracked neither grows nor shrinks. A workload then consists of pairs of deletions and insertions. As we assume a general setting where updates are relatively much more frequent than queries, we assume that the effects of piggybacking on the buffer can be ignored.

With these assumptions, every incoming update operation is an insertion or a deletion with equal probability.

Now, we will determine how many operations that can be entered into the buffer before the buffer gets full and an invocation of `EMPTYBUFFER` occurs. To do that, we need to consider the effects of annihilation because these reduce the size of the buffer.

First, observe that the second case of annihilation, where an incoming insertion annihilates with a deletion in the buffer (the transition from state 2 to state 4 in Figure 2), is rare. This case occurs only when an object sends duplicate updates in rapid succession.

Next, the first case of annihilation, where an incoming deletion annihilates with an insertion in the buffer, is more frequent. Its probability depends on the current number of insertions residing in the buffer. Thus, if all the  $C$  operations in the buffer are insertions, which is highly unlikely, the probability is  $C/O$ . At the other extreme, the smallest number of insertions in a buffer storing  $C$  operations is  $C/2$ . This case occurs when no annihilations have occurred since the last emptying of the buffer. In this case, the probability of an incoming deletion annihilating with an insertion is  $C/(2O)$ . Thus, the probability of an annihilation occurring is between  $C/(2O)$  and  $C/O$ . We will assume that annihilation is not very frequent and thus use the probability  $C/(2O)$  in the following.

Because the probability of an incoming operation being a deletion is  $1/2$ , the probability of removing an operation from the buffer is  $P_{dec}(C) = C/(4O)$ . The complementary probability of adding one operation to the buffer is then  $P_{inc}(C) = 1 - C/(4O)$ .

The values of these two probabilities vary in the ranges from  $\min(P_{inc}(C)) = P_{inc}(C_{max}) = 1 - C_{max}/(4O)$  to  $\max(P_{inc}(C)) = P_{inc}(0) = 1$  and correspondingly from  $\min(P_{dec}(C)) = P_{dec}(0) = 0$  to  $\max(P_{dec}(C)) = P_{dec}(C_{max}) = C_{max}/(4O)$ . To simplify the continued analysis, we assume that these probabilities are constant and use their average values (i.e., their values when the buffer

is half full):

$$\begin{aligned} P_{inc} &= P_{inc} \left( \frac{C_{max}}{2} \right) = 1 - \frac{C_{max}}{8O}, \\ P_{dec} &= P_{dec} \left( \frac{C_{max}}{2} \right) = \frac{C_{max}}{8O}. \end{aligned}$$

We now know with what probability an incoming operation will increment or decrement the buffer size by one.

Assume that an empty buffer receives an incoming workload of  $X$  operations. In such a workload, the expected number of buffer-size increasing operations will be equal to  $P_{inc}|X| = |X| - (|X|C_{max})/(8O) = (8O|X| - C_{max}|X|)/(8O)$ , and the expected number of buffer-size decreasing operations will be equal to  $C_{max}|X|/(8O)$ . Thus, if the buffer does not become full in the process, the expected final number of entries in the buffer after processing the workload will be equal to

$$\begin{aligned} W(X, O, C_{max}) &= \frac{8O|X| - C_{max}|X|}{8O} - \frac{C_{max}|X|}{8O} \\ &= \frac{|X|(4O - C_{max})}{4O}. \end{aligned} \quad (1)$$

In the worst case, the buffer is filled every  $C_{max}$  operations. In the average case, we assume an initially empty buffer and derive the expected buffer-filling workload size  $|X|$  from Equation 1 as follows:

$$\begin{aligned} C_{max} &= W(X, O, C_{max}) = \frac{|X|(4O - C_{max})}{4O} \\ \Rightarrow |X|(O, C_{max}) &= \frac{4C_{max}O}{4O - C_{max}}. \end{aligned} \quad (2)$$

Note that we have indicated that  $|X|$  is a function of  $O$  and  $C_{max}$ .

Observe that

$$\lim_{O \rightarrow \infty} |X|(O, C_{max}) = C_{max} \text{ and } \lim_{O \rightarrow C_{max}} |X|(O, C_{max}) = \infty.$$

This confirms the intuition that annihilations are less likely to happen when the number of objects is relatively large compared to the size of the buffer, so that an incoming deletion is less likely to find a matching insertion in the buffer. Further, if all the objects almost fit in the buffer, annihilations happen all the time, and buffer emptying is rare. The next section models the I/O cost of buffer emptying.

## 4.2 EmptyBuffer Cost Model

Updates incur I/O only when `EMPTYBUFFER` is invoked. To estimate the cost of `EMPTYBUFFER`, we make several simplifying assumptions.

We assume that incoming operations are uniformly distributed in the data space; thus, an operation from the buffer is equally likely to be routed by the R-tree algorithms to any of the subtrees of the root. Next, we assume that a deletion is only sent to one subtree at the root level. This is a reasonable

for R-trees storing point or small-region data (which is the case in our setting) because this results in MBRs with little overlap.

Finally, we assume that each object has only one position recorded in the disk tree. Thus, cases such as the one shown for object  $c$  in Figure 10 are very rare.

To estimate the cost of buffer emptying, we assume that those sublists of operations that are sent to the subtrees of the root read and write all the nodes in these subtrees. Note that, due to the operation threshold parameter in GET-PUSHDOWNOPS, only relatively large sublists are sent down the tree. If the number of uniformly distributed operations in a sublist is equal to the number of leaf nodes in the subtree, all the nodes of the subtree will be touched.

The cost of buffer emptying can now be expressed as the cost of reading and writing the whole R-tree minus the cost of reading and writing those subtrees that do not receive sublists of operations. In the following, we compute these two costs.

If the average fanout of a node is  $S$ , the height of the tree is  $h = \lceil \log_S O \rceil$ . The total I/O cost of reading and writing the tree is then

$$K_{tree}(S, O) = 2 + 2S + 2S^2 + \dots + \left\lceil \frac{2O}{S} \right\rceil \approx \frac{2(O-1)}{S-1}. \quad (3)$$

To approximate the sizes of the subtrees not touched during buffer emptying, we model the distribution of the sizes of the operation sublists at the root node of the disk tree.

In general, distributing  $n$  objects among  $S$  groups with uniform probability results in groups of objects, whose sizes follow the discrete binomial distribution  $B(n, 1/S)$  [18]. This distribution can be approximated by the normal distribution  $N(n/S, (1-1/S)n/S)$ , provided that  $n/S > 10$  and  $n(1-1/S) > 10$ .

In our setting, the size of each group is approximated by the normal distribution  $N(C_{max}/R, (1-1/R)C_{max}/R)$ , where  $R$  is the current root node fanout, which in most cases is different from  $S$ . The approximation is then valid if  $C_{max}/R > 10$ , which is true for all buffers that are not unreasonably small.

Every root node operation sublist will be of size  $k$  or less with probability  $F(k)$ , where  $F$  is the cumulative distribution function of the normal distribution. Thus, the expected number of sublists that are too small and are not sent down the tree is equal to  $R \times F(k)$ . The height of the child subtrees of the root, which are not touched by the buffer emptying, is  $\lceil \log_S(O/R) \rceil$ . The cost of reading and writing one such subtree is thus equal to

$$K_1(R, S, O) = 2 \sum_{i=0}^{\lceil \log_S \frac{O}{R} \rceil - 1} S^i \approx 2 \left( \frac{\frac{O}{R} - 1}{S - 1} \right).$$

The total amount of I/O operations saved by not pushing down sublists with fewer than  $k$  operations is equal to

$$K_{saved}(k, R, S, O) = K_1(R, S, O)R \times F(k) = \frac{2(O-R)F(k)}{S-1}.$$

Thus, the total cost of EMPTYBUFFER becomes

$$\begin{aligned} K_{total}(k, R, S, O) &= K_{tree}(S, O) - K_{saved}(k, R, S, O) \\ &= \frac{2(O-1) - 2(O-R)F(k)}{S-1}. \end{aligned}$$

Because entries of the small operation sublists are returned back to the buffer, buffer emptying is performed more often. The expected total number of entries in the too small sublists is equal to

$$C_{ret}(R, k) = R \int_0^k x f(x) dx,$$

where  $f$  is the probability mass function of the normal distribution.

The expected number of updates between EMPTYBUFFER invocations is decreased by the number of entries returned to the buffer. Using Equation 2,

$$\begin{aligned} |X|_T(k, O, R, C_{max}) &= |X|(O, C_{max}) - C_{ret}(R, k) \\ &= \frac{4C_{max}O}{4O - C_{max}} - R \int_0^k x f(x) dx. \end{aligned}$$

Finally, putting the pieces together, we obtain the following expected update cost:

$$\begin{aligned} K_U(k, R, S, O, C_{max}) &= \frac{K_{total}(k, R, S, O)}{|X|_T(k, O, R, C_{max})} \\ &= \frac{2(C_{max} - 4O)(1 + O(F(k) - 1) - F(k)R)}{(S - 1)(4C_{max}O + R(C_{max} - 4O) \int_0^k x f(x) dx)}. \end{aligned}$$

Because the derivative of  $K_U(k, R, S, O, C_{max})$  with respect to  $k$  does not have a closed form, we use numerical methods to find a value of  $k$  that minimizes  $K_U$  for given settings of the other parameters. In Section 5, we compare the results of the numerical computations with the results of empirical experiments. Note that numerical computation of an optimal  $k$  needs to be performed only when the values of parameters  $R, S, O$ , and  $C_{max}$  change significantly.

## 5 Experimental Evaluation

In this section we describe the results of performance experiments with the R<sup>R</sup>-tree and two competing proposals. The settings for the experiments and the used workloads are described first. This is followed by the validation of the cost model, then we present the results of performance comparison with the competing proposals and finally explore selected properties of the R<sup>R</sup>-tree.

### 5.1 Experimental Setup

Three indexing techniques are covered in the experiments: the R<sup>R</sup>-tree, the RUM-tree [24], and the R\*-tree. The XXL library [6] implementation of the R\*-tree was used. The R<sup>R</sup>-tree was also implemented on top of the XXL library [6]. The RUM-tree implementation was kindly provided by its inventors [24]. We modified the implementation to add an LRU buffer.

Unless stated otherwise, a 4 kilobyte page size (and thus tree node size) is used both for disk trees and the buffer tree of the R<sup>R</sup>-tree. Also when not stated otherwise, the size of the operation buffer in the R<sup>R</sup>-tree is chosen so that, on average, it is able to store the number of operations that is equal to the 5% of the number of tracked objects. For most of the experiments with the R<sup>R</sup>-tree,

no disk-page buffer is used. When a page buffer is used, the LRU write-back page replacement policy is applied so that dirty pages are written to disk only when they are evicted from the buffer.

## 5.2 Workloads

To generate workloads for most of the experiments, the generator from the COST benchmark [12] is used. This generator simulates a number of moving objects that update their positions using the shared-prediction based update approach. More specifically, point-based tracking is used, meaning that an object updates its position when its current position reaches a distance to its most recently reported position that is equal to an agree-upon accuracy threshold *thr*. The objects' positions are inserted into the index as bounding squares of the circles with radius *thr*.

We use two kinds of workloads generated by the COST generator. In the uniform workloads, an object starts from a random position and moves in a random direction and at a random speed between 0 and 180 km/h in-between updates. In the non-uniform workloads, simulated objects move in a road network that is a complete graph connecting 20 randomly placed intersections. Three classes of objects are generated with maximum speeds of 45, 90, and 180 km/h.

Each of the generated workloads is a mix of insertions, deletions, and queries. First, all objects are inserted. Then, they send updates (deletion-insertion pairs) according to the update approach described above. The I/O performance is measured also only after the index reaches steady state, i.e., all objects are inserted.

Range queries are distributed uniformly in the workload and in the data space. Each query range is a rectangle covering 0.02% of the data space. Table 2 shows the default values of the workload parameters used for all of the experiments (unless the value of the parameter is varied in the experiments). Note that in this table (and later in the graphs) an update denotes an individual

Parameter	Default Value
Number of moving objects	100,000
Size of the dataspace	100km × 100km
Accuracy threshold	200m
Number of updates	400,000
Queries per update	1 : 20,000

Table 2: Workload Parameters

deletion or insertion, not counting the initial insertions. Thus, when the ratio of the number of queries to the number of updates is equal to 1/20,000, this means that one query is performed for every 10,000 deletion-insertion pairs.

We have chosen to use 400,000 operations in the workloads because the use of workloads with larger numbers of operations does not significantly change the performance numbers. A performance graph later in this section demonstrates this.

The experiments that validate selected properties of cost model use an uniform workload which is generated in a different way. In this workload, referred to as a *pure uniform* workload, notions of direction, speed and tracking accuracy

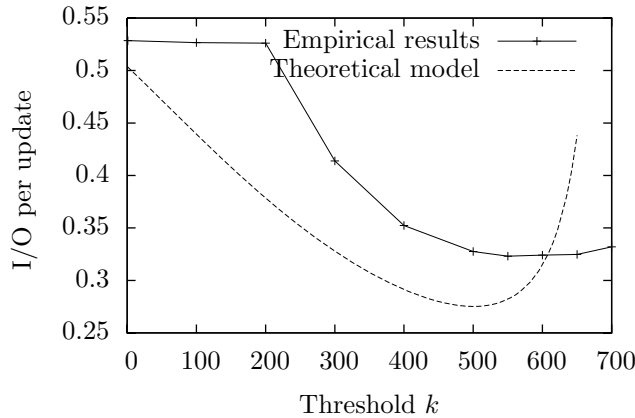


Figure 13: Cost model validation, 5% buffer

do not exist. The next object to issue an update for is chosen with uniform probability from the set of all objects and its new coordinates are chosen from the space with uniform probability, with no impact from the old coordinates. Thus, in such workload, the data objects “jump around” the space in a completely random order. A dataset generated in this way comes closest to the theoretical assumptions of the cost model, thus enabling us to validate conclusions of the model. Also, when this dataset is used, the buffer emptying strategy is always set to empty everything.

This workload for our experiments contains 400,000 operations for 100,000 moving objects.

### 5.3 Cost Model Validation

The first set of experiments were run with the goal of validating the cost model presented in Section 4 in a most general way. For this purpose, uniform workloads were run with different values of the operation threshold  $k$  and two settings for the operation buffer size: 5% and 20% of the number of moving objects. Note that when  $k = 1$ , the buffer is emptied completely. Figures 13 and 14 plot the update performance of the R<sup>R</sup>-tree together with curves obtained numerically from the cost model. Observe that, in each of the graphs, the minima of both curves almost coincide. This validates the effectiveness of the cost model. The curves of the cost model go sharply up with large values of  $k$  because, for these values of  $k$ , the assumptions of the cost model do not hold.

#### 5.3.1 Frequency of Annihilation

The goal of this experiment is to validate the cost model assumption that probability of annihilation is between  $C/O$  and  $C/2O$  on average for every update operation. The pure uniform dataset is used for this experiment. The results are presented in Figure 15. They show that while cost model successfully captures linearity and rate of change of the annihilation probability function, it is sub-

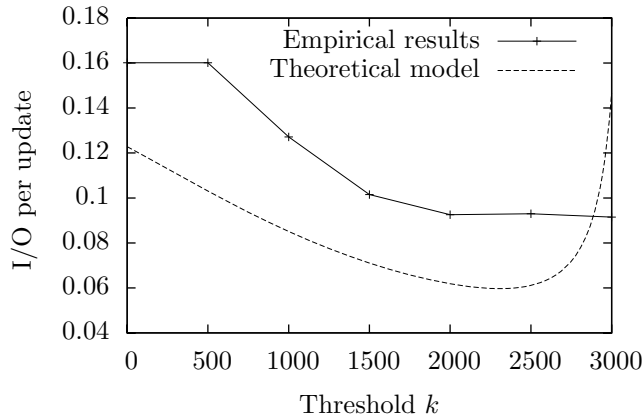


Figure 14: Cost model validation, 20% buffer

stantially off-mark with respect to its absolute value. Consequently, providing a better estimation of annihilation probabilities is one of the future work items.




Figure 15: Validation of annihilation assumptions

For the rest of the experiments, we use non-uniform workloads.

## 5.4 Comparing the Indexes

To ensure that the comparison of the three indexes is fair and thus obtains meaningful results, in each of the experiments, we allocate exactly the same amount of main memory to each index. All the available main memory is used for the page buffer for the  $R^*$ -tree.

The  $R^R$ -tree uses all of the available memory for the operation buffer in most experiments. The exception is that in the experiments with high ratios of



queries, the  $R^R$ -tree divides memory between the operation buffer and the page buffer. The division used in each case was chosen through a number of performance experiments. The RUM-tree uses a part of the available main memory for its update memo, and the rest is used for the page buffer. Experiments with three different configurations of the RUM-tree were run to determine which one performs the best: 20% inspection ratio with clean-on-touch, 0% inspection ratio with clean-on-touch, and 0% inspection ratio without clean-on-touch. The configuration with 20% inspection ratio and clean-on-touch was chosen. This result is in correspondence with the findings of the inventors of the RUM-tree.

Figure 16 shows how the indexes compare for different sizes of available main memory. Note that the amount of main memory is given in terms of the average number of objects that can fit in the memory organized into an R-tree. The amount of main memory is varied from 0.1% to 100% of the number of moving objects. The  $R^R$ -tree clearly outperforms the other two approaches.

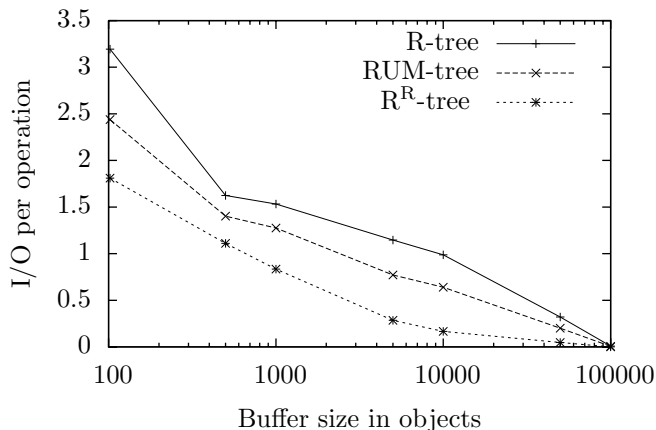


Figure 16: Scalability with respect to available main memory

For example, when the amount of main-memory is equal to 10% of the number of moving objects, the update performance of the  $R^R$ -tree is almost four times better than that of the RUM-tree and more than seven times better than that of the  $R^*$ -tree.

Figure 17 shows how the different approaches perform when the ratio of updates and queries is changed. Again the  $R^R$ -tree is better under all settings except when there are more queries than updates. In this setting, where the query performance is the dominant factor, the RUM-tree performs better than the  $R^R$ -tree, because the  $R^R$ -tree uses a smaller page buffer than the RUM-tree.

The next two sets of experiments explore the performance of the indexes under different workload sizes and different disk page sizes. The graphs in Figure 18 show that the performance of the  $R^R$ -tree does not degrade with an increasing length of the workloads. The graphs in Figure 19 demonstrate that, as expected, the I/O performance is improved when the disk page size is increased.




Figure 17: I/O performance versus the update/query ratio




Figure 18: Scalability with respect to workload duration

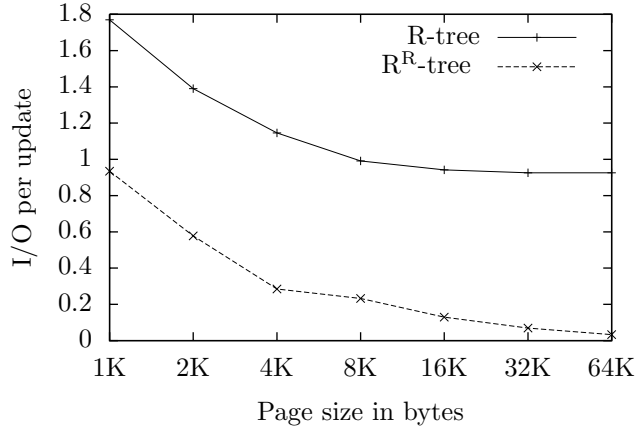


Figure 19: Scalability with respect to disk page size

#### 5.4.1 Performance in a Persistent Setting

An interesting experiment is to try the  $R^R$ -tree outside its intended setting and to measure its performance assuming full persistency. In this case it is assumed that every incoming operation is written to the log immediately and LRU buffers for the  $R^R$ -tree,  $R$ -tree and RUM-tree are write-through instead of write-back, i.e. the pages that become dirty are immediately written to the disk.

Two experiments were performed in such setting. Figure 20 presents I/O performance relation to the available main memory, where all three indexing solutions demonstrate comparable performance. Figure 21 presents I/O performance relation to the memory division between  $R^R$ -tree buffer and LRU cache. It shows that available memory should be allocated to the  $R^R$ -tree buffer in this setting as well, but the resulting performance is not a significant improvement over the  $R$ -tree with LRU cache, as shown by the rightmost data point in the figure. To conclude, the  $R^R$ -tree is not as well suited for the full-persistency setting as it is for the partial-persistency one.

## 5.5 Exploring the $R^R$ -tree Algorithms

### 5.5.1 Memory Division Experiments

In the next set of experiments, we explore selected properties of the  $R^R$ -tree. The goal of the first experiment is to determine how effective the operation buffer is at increasing the update performance in comparison to the use of an LRU write-back page buffer. In this experiment, part of the available memory is allocated to the operation buffer, and the rest is allocated to the LRU page buffer. The total amount of main memory is kept constant, but the fraction allocated to the operation buffer is changed. The experiment is run with two workloads, one with few queries and one with an equal amount of queries and object updates. The I/O performance is averaged over all three types of operations: insertions, deletions, and queries.

Figure 22 shows very clearly that when most of the operations are updates,




Figure 20: Scalability with respect to main memory in a persistent setting

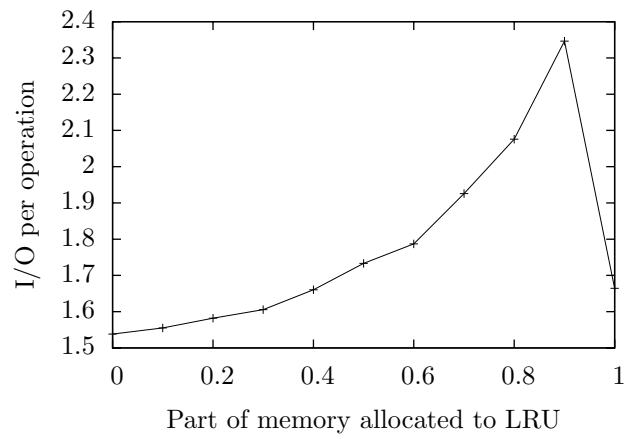


Figure 21: I/O performance with varying main-memory divisions in a persistent setting

all of the available main memory should be allocated to the operation buffer.

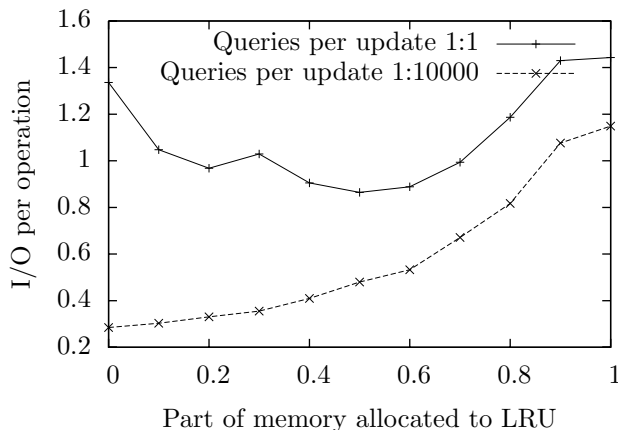


Figure 22: I/O performance for varying main-memory divisions

On the other hand, reducing the size of the page buffer increases the cost of queries. Thus, when the workload contains a lot of queries, part of the memory should be allocated to the operation buffer and part to the page buffer.

### 5.5.2 Effect of Query Piggybacking

The effect of query piggybacking (as described in Section 3.8) was studied for operation buffer sizes of 5% and 20% of the number of moving objects (Figure 23). The experiments show that piggybacking starts to have a noticeable effect when there is a substantial fraction of queries in the workload. For example, for the 5% buffer size and 200 updates per query in the workload, only about 4.9% of updates are performed by the PIGGYBACK algorithm during the processing of queries. When the number of updates per query drops to 20, about half of the updates are performed by the PIGGYBACK algorithm. The corresponding percentages of updates for the 20% buffer size are 20.2% and 98%.

### 5.5.3 Group Size Statistics

To better understand how the buffer emptying strategy affects I/O performance and to provide insights into future research directions, we discuss the actual group sizes and their frequencies encountered in EMPTYBUFFER invocations.

The group size data is provided for three distinct levels, assuming a tree with the height of three. Level 2 data shows the group size distribution at the buffer prior to executing any of the operations. If no grouping is performed before executing the operations, like in the case of full buffer emptying, then all groups at this level are considered to be the size of the buffer. Level 1 contains the actual groups that were executed on the disk tree root. Level 1 and 2 data are closely related, with important difference that level 1 includes deletion splits and GROUPUPDATE restart information. Finally, level 0 data contains the actual groups that were received by the leaf level nodes.




Figure 23: Frequency of piggybacking

Empirical group size distributions for the simplest strategy of emptying whole buffer every time are presented in Figure 24. These distributions can be characterized by a large number of very small (less than 10 operations) groups at the leaf level and relatively large groups at the upper levels.




Figure 24: Group size distributions, whole buffer emptying

The group size distributions for the threshold-based buffer emptying strategy is presented in Figure 26. Compared to the whole buffer emptying strategy, the group size distribution at the non-leaf level has shifted right significantly and the number of very small groups at the leaf level has been slightly reduced. An apparent anomaly is the appearance of small number of very small groups at the level 1. Such groups are caused by the GROUPUPDATE algorithm restart.

The distributions for the largest-group buffer emptying strategy (Figure 26) are provided solely for completeness as they do not differ significantly from the threshold-based case. This is in line with performance findings.




Figure 25: Group size distributions, best threshold emptying

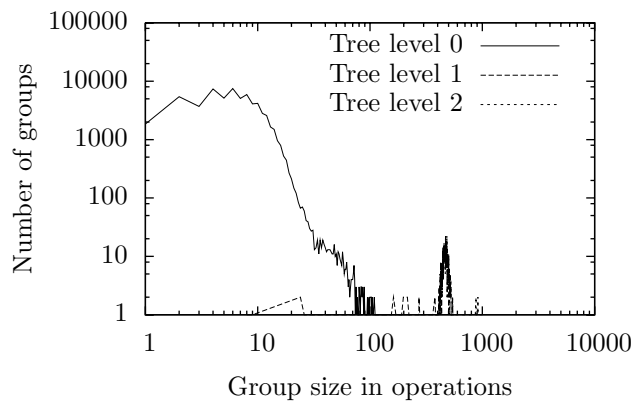


Figure 26: Group size distributions, largest groups emptying

All three strategies influence only the root level groups directly. Thus Figure 27 presents the root level distributions separately from the rest. The presented data confirms that the non-trivial buffer emptying strategies successfully increase group sizes at the root level. Moreover, threshold-based and largest-groups based emptying have roughly the same distributions and consequently almost identical performance.




Figure 27: Root level group size distributions

A perhaps more important observation is that all the emptying strategies discussed exhibit a large number of very small groups at the leaf level of the tree. The advanced buffer emptying strategies achieve better performance mainly by increasing groups at root level, which has only a limited impact on the overall group size and average I/O per operation. Compared to this, potential performance improvement coming from avoiding small groups at the leaf level looks very promising as a future work item.

#### 5.5.4 Tree Quality

The  $R^R$ -tree insertion algorithm does not use some of the ways of the standard  $R^*$ -tree insertion algorithm to reorganize the tree. Namely, the re-insertion of underfull nodes is not done. Thus the goal of this experiment is to measure how much of the *tree quality* is lost due to this simplification. The tree quality is considered to be high when the number of I/O operations spent for queries is low.

The workload for this experiment was prepared in a specific way. Additional 10,000 range queries were appended to the pure uniform workload and the performance data is gathered only for these queries. Moreover, the  $R^R$ -tree buffer is forcibly completely emptied before executing these queries so that all data resides on disk. For the same purpose no LRU buffer is used for the R-tree. The results of this experiment in Figure 28 indicate that about 8% of search performance is lost due to different insertion algorithm. Another conclusion is that the different operation group sizes witnessed in settings with different buffer sizes do not directly affect tree quality.



Special attention should be paid to the rightmost data point that indicates performance with buffer size of 100,000 objects. In this setting whole data for the R<sup>R</sup>-tree fits in the main memory and the forced buffer emptying creates the whole disk R-tree at once. Thus it can be viewed as an evaluation of the EMPTYBUFFER algorithm as a means to do bulk-loading for the R-tree. And in this case the quality of the resulting tree is as good as that of the R-tree created by series of standard R-tree insertion operations.

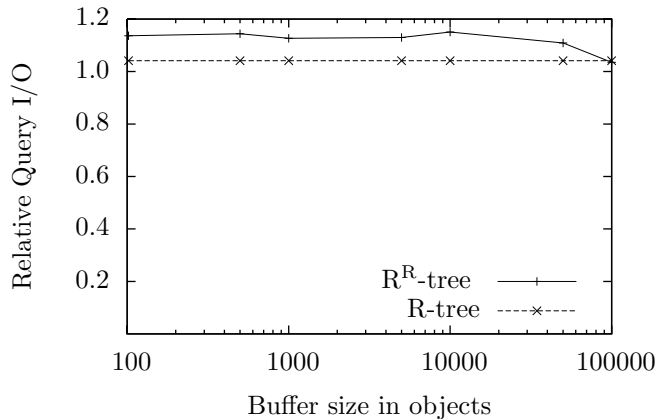


Figure 28: Disk R-tree quality

## 6 Related Work

We proceed to consider relevant previous work on the efficient update of general disk-based data structures and, particularly, on efficient updates in R-trees.

Techniques for the bulk-loading of data structures are relevant for the problem of efficient updates. Notably, Choubey et al. [5] have proposed algorithms that aim to perform many insertions on an existing index structure in one go. However, these algorithms are not directly applicable in our setting because they focus on the search performance of the resulting structure, whereas the performance of the bulk-loading itself is of secondary importance. In addition, they do not consider deletions.

An approach to dealing with frequent updates of continuous variables such as positions is to model the positions as linear functions of time instead of the standard constant functions of time. When representing the position of an object by a linear function as, e.g., proposed by Wolfson et al. [23], the modeled position stays close to the actual for longer periods of time. It has been shown that, for a range of reasonable accuracy thresholds, the number of updates needed to maintain the positions of vehicles is reduced by almost a factor of 3 [21]. The use of such linear functions on the top of R-tree was first proposed by Šaltenis et al. [19] and subsequently explored by others. These proposals are orthogonal to our proposal, and it is possible to combine them to obtain the combined benefits.

In a survey, Graefe [9] presents a general overview of techniques for speeding up B-tree updates. One of the techniques covered—the buffering of insertions in a separate data structure—is used in our proposal. The survey also describes a differential file design that enables the buffering of deletions and updates in addition to insertions. Differential indexes and on-line index construction were researched in the context of B-trees and relaxed index persistence [17, 20]. These approaches are similar to the buffer of the  $R^R$ -tree that buffers all incoming updates (i.e., deletions and insertions). The idea of a non-logged B-tree is also mentioned, which directly corresponds to our notion of relaxed persistence or use as a secondary index. Finally, in recent work, Graefe [10] briefly discusses how so-called merged B-trees can be used to conveniently store a main index and to-do lists that contain unprocessed keys.

The problem of frequent updates in R-tree was recently tackled by Lee et al. [14], who present the most competitive proposal for bottom-up updates of R-trees. The bottom-up approach avoids potentially expensive top-down traversals and exploits the update locality of continuous variables. However, an auxiliary data structure, required in order to access the bottom level, uses a significant amount of space and is disk based. Because of this, updates still take at least 3 I/Os. Additionally, the proposal includes a main-memory data structure that compactly stores a summary of the tree. This structure helps save I/O in the case of non-local updates, but uses a fixed amount of main memory. In contrast, our proposed technique is top-down and mitigates the drawbacks of top-down traversals by performing operations in batches. Thus, there is no need for direct leaf-level access, and our proposal uses all available main memory.

Arge has proposed a general tree buffering technique that associates a buffer with every non-leaf node [1]. Operations are then not performed immediately, but are placed in the buffers. Once a buffer gets full, its contents are moved to the buffers at the next level in its subtree. Main memory is used for the buffer emptying. A drawback of this approach is that it cannot answer queries immediately. Arge et al. have also applied the technique to R-trees [2], and Van den Bercken and Seeger [7] have explored similar techniques. Our proposal shares the idea of buffering and performing updates lazily and in batches. A key difference is that the  $R^R$ -tree does not use disk-based buffers for the intermediate nodes—it uses only a single main-memory buffer that can be compared to the root buffer in Arge’s approach. Another difference is that our approach supports efficient, instantaneous queries.

Lin and Su combine the buffer-tree and bottom-up techniques [15]. They attach disk-based insertion buffers to intermediate nodes and perform insertions top-down in batches. To avoid multiple partial traversals for a single deletion, deletions are gathered in a main-memory deletion buffer that is applied lazily and directly to the leaf-level nodes. A main-memory based leaf-level access table facilitates this. The major drawback of this approach is that the size of the main-memory based leaf-level access table is of the same order of magnitude as the size of the database. For example, for R-trees, the size of this access table is one third of the size of the index.

A recent approach by Xiong and Aref [24] significantly lowers the average update cost by performing deletions in main memory. However, insertions are performed using the ordinary R-tree algorithm, so there are no further performance gains for these. A main-memory data structure, called update memo, is responsible for keeping information about deletions as well as about the latest

and obsolete data entries. This approach is similar to ours, in that deletions are performed in main memory.

The LUGrid approach [25] is the only related approach that indexes spatial data and makes the same persistence assumptions we do. The approach uses a disk-based grid file, a main-memory memo structure for keeping track of deletions, and a main-memory grid for storing unprocessed insertions. Similar to our proposal, the LUGrid flushes the insertions from a main-memory grid cell when it becomes full. In contrast to our proposal, the LUGrid cannot store objects with extents, the circular object-positions of moving objects cannot be indexed.

The AGILE method [8] addresses the issue of balancing conflicting performance requirements for update and query operations in a more general setting of information filtering. This approach proposes augmenting existing data structures in a way that selectively increases update performance at the cost of losing query accuracy. This proposal mentions implementing AGILE on the of the R-tree as a part of their future work. The outline of such implementation suggests buffers on internal nodes that store frequently updated objects. Detailed comparison with our approach is impossible at this time as the AGILE has not been actually implemented on top of the R-tree yet.

While the related approaches described above strive to solve problems similar to ours, with the exception of the B-tree approaches and the LUGrid, they cannot be directly compared to our proposal because of the different persistence assumptions. Those approaches assume a persistent index; thus, an update requires always at least one I/O. In contrast to this, the  $R^R$ -tree does not have to write every incoming operation to disk. As a result, better performance than one I/O per update can be achieved.

## 7 Conclusion and Research Directions

Motivated by emerging database applications that involve the monitoring of large collections of continuous variables and thus are characterized by high rates of updates, this thesis presents a novel data structure, called the  $R^R$ -tree, that supports updates more efficiently than existing proposals while also supporting queries. In contrast to related work, this data structure is capable of efficiently using any amount of main memory, and it supports the same settings and operations as an ordinary R-tree. Another core difference to the related work is a relaxed persistence assumption: we argue that relaxed persistence of an index is appropriate in a setting with hyper-dynamic data.

The  $R^R$ -tree builds on two main ideas. First, *operation buffering* in main memory enables execution of the majority of update operations quickly, and it allows some updates to annihilate one another, thus avoiding any I/O at all. Second, by *grouping* operations on buffer emptying, the index enables all operations that travel to the same node of the tree to share I/O.

A strong point of this approach is its orthogonality to the tree-like data structure being buffered. This makes it possible to adapt the approach to other types of trees. This yields a quite general contribution. Furthermore, this approach is orthogonal to other means of exploiting the available main memory. This enables adaptation of the  $R^R$ -tree to different workloads, by combining the operation buffer with some page-cache buffer, e.g., an LRU cache.

The empirical performance study includes favorable comparisons with the conventional LRU-cached R-tree and with the state-of-the-art RUM-tree. The study also offers general insight into the benefits of LRU caching in the thesis' setting.

## 7.1 Research Directions

Several interesting directions for future research exist. We expect it to be very interesting to apply the proposed techniques to other indexes in the class of grow-post trees, e.g., the TPR-tree. Next, we note that this thesis' focus has been on I/O efficiency. However, the  $R^R$ -tree does use an in-memory R-tree to efficiently support querying. Thus, the proposal should be competitive with respect to CPU performance. Still, studies of CPU performance and the investigation of techniques that reduce the CPU cost are in order. Finally, it may perhaps be possible to improve the quality of the disk tree by tuning the group split strategy, e.g., by adopting a clustering technique.

**Acknowledgments** I would like to thank my supervisor Simonas Šaltenis for his guidance throughout my work on this paper, Christian S. Jensen who was closely involved in supervision throughout whole project and Kristian Torp for his insightful comments. Also I am thankful to the staff and fellow students in the Database and Programming Technologies Research Group for creating friendly study environment.

## References

- [1] Lars Arge. The buffer tree: A new technique for optimal I/O-algorithms (extended abstract). In Selim G. Akl, Frank K. H. A. Dehne, Jörg-Rüdiger Sack, and Nicola Santoro, editors, *WADS*, volume 955 of *Lecture Notes in Computer Science*, pages 334–345. Springer, 1995.
- [2] Lars Arge, Klaus Hinrichs, Jan Vahrenhold, and Jeffrey Scott Vitter. Efficient bulk operations on dynamic R-trees. *Algorithmica*, 33(1):104–128, 2002.
- [3] Norbert Beckmann, Hans-Peter Kriegel, Ralf Schneider, and Bernhard Seeger. The  $R^*$ -tree: an efficient and robust access method for points and rectangles. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 19(2):322–331, June 1990.
- [4] Sotiris Brakatsoulas, Dieter Pfoser, and Yannis Theodoridis. Revisiting R-tree construction principles. *Lecture Notes in Computer Science*, 2435:149–162, 2002.
- [5] Rupesh Choubey, Li Chen, and Elke A. Rundensteiner. GBI: A generalized R-tree bulk-insertion strategy. In Ralf Hartmut Güting, Dimitris Papadias, and Frederick H. Lochovsky, editors, *Advances in Spatial Databases, 6th International Symposium, SSD'99, Hong Kong, China, July 20-23, 1999, Proceedings*, volume 1651 of *Lecture Notes in Computer Science*, pages 91–108. Springer, 1999.

- [6] Jochen Van den Bercken, Björn Blohsfeld, Jens-Peter Dittrich, Jürgen Krämer, Tobias Schäfer, Martin Schneider, and Bernhard Seeger. XXL — a library approach to supporting efficient implementations of advanced database queries. In Peter M. G. Apers, Paolo Atzeni, Stefano Ceri, Stefano Paraboschi, Kotagiri Ramamohanarao, and Richard T. Snodgrass, editors, *VLDB*, pages 39–48. Morgan Kaufmann, 2001.
- [7] Jochen Van den Bercken and Bernhard Seeger. An evaluation of generic bulk loading techniques. In *VLDB '01: Proceedings of the 27th International Conference on Very Large Data Bases*, pages 461–470, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.
- [8] Jens-Peter Dittrich, Peter M. Fischer, and Donald Kossmann. Agile: adaptive indexing for context-aware information filters. In *SIGMOD '05: Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, pages 215–226, New York, NY, USA, 2005. ACM Press.
- [9] Goetz Graefe. B-tree indexes for high update rates. *SIGMOD Rec.*, 35(1):39–44, 2006.
- [10] Goetz Graefe. Master-detail clustering using merged B-tree indexes. Manuscript, 19 pages, 2007.
- [11] Antonin Guttman. R-trees: a dynamic index structure for spatial searching. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 14(2):47–57, 1984.
- [12] Christian S. Jensen, Dalia Tiešytė, and Nerius Tradišauskas. The COST benchmark-comparison and evaluation of spatio-temporal indexes. In Mong-Li Lee, Kian-Lee Tan, and Vilas Wuwongse, editors, *DASFAA*, volume 3882 of *Lecture Notes in Computer Science*, pages 125–140. Springer, 2006.
- [13] Christian S. Jensen, Dalia Tiešytė, and Nerius Tradišauskas. Robust B+-tree-based indexing of moving objects. *MDM*, 00:12, 2006.
- [14] Mong-Li Lee, Wynne Hsu, Christian S. Jensen, Bin Cui, and Keng Lik Teo. Supporting frequent updates in R-trees: A bottom-up approach. In *VLDB*, pages 608–619, 2003.
- [15] Bin Lin and Jianwen Su. Handling frequent updates of moving objects. In Otthein Herzog, Hans-Jörg Schek, Norbert Fuhr, Abdur Chowdhury, and Wilfried Teiken, editors, *CIKM*, pages 493–500. ACM, 2005.
- [16] David B. Lomet. Grow and post index trees: Roles, techniques and future potential. In *SSD '91: Proceedings of the Second International Symposium on Advances in Spatial Databases*, pages 183–206, London, UK, 1991. Springer-Verlag.
- [17] Kerttu Pollari-Malmi, Eljas Soisalon-Soininen, and Tatu Ylönen. Concurrency control in B-trees with batch updates. *IEEE Transactions on Knowledge and Data Engineering*, 8(6):975–984, 1996.

- [18] Martin Raab and Angelika Steger. "balls into bins" - a simple and tight analysis. In *RANDOM '98: Proceedings of the Second International Workshop on Randomization and Approximation Techniques in Computer Science*, pages 159–170, London, UK, 1998. Springer-Verlag.
- [19] Simonas Saltenis, Christian S. Jensen, Scott T. Leutenegger, and Mario A. Lopez. Indexing the positions of continuously moving objects. In Weidong Chen, Jeffrey F. Naughton, and Philip A. Bernstein, editors, *SIGMOD Conference*, pages 331–342. ACM, 2000.
- [20] V. Srinivasan and Michael J. Carey. Performance of on-line index construction algorithms. In *EDBT '92: Proceedings of the 3rd International Conference on Extending Database Technology*, pages 293–309, London, UK, 1992. Springer-Verlag.
- [21] Alminas Čivilis and Stardas Pakalnis. Techniques for efficient road-network-based tracking of moving objects. *IEEE Transactions on Knowledge and Data Engineering*, 17(5):698–712, 2005. Senior Member-Christian S. Jensen.
- [22] Ouri Wolfson, A. Prasad Sistla, Sam Chamberlain, and Yelena Yesha. Updating and querying databases that track mobile units. *Distrib. Parallel Databases*, 7(3):257–387, 1999.
- [23] Ouri Wolfson, Bo Xu, Sam Chamberlain, and Liqin Jiang. Moving objects databases: Issues and solutions. In Maurizio Rafanelli and Matthias Jarke, editors, *SSDBM*, pages 111–122. IEEE Computer Society, 1998.
- [24] Xiaopeng Xiong and Walid G. Aref. R-trees with update memos. In *ICDE '06: Proceedings of the 22nd International Conference on Data Engineering (ICDE'06)*, page 22, Washington, DC, USA, 2006. IEEE Computer Society.
- [25] Xiaopeng Xiong, Mohamed F. Mokbel, and Walid G. Aref. Lugrid: Update-tolerant grid-based indexing for moving objects. In *MDM*, page 13. IEEE Computer Society, 2006.