

Making Updates Disk-I/O Friendly Using SSDs

Mohammad Sadoghi¹, Kenneth A. Ross^{1,2}, Mustafa Canim¹, Bishwaranjan Bhattacharjee¹

¹IBM T.J. Watson Research Center

²Columbia University

ABSTRACT

Multiversion databases store both current and historical data. Rows are typically annotated with timestamps representing the period when the row is/was valid. We develop novel techniques for reducing index maintenance in multiversion databases, so that indexes can be used effectively for analytical queries over current data without being a heavy burden on transaction throughput. To achieve this end, we re-design persistent index data structures in the storage hierarchy to employ an extra level of indirection. The indirection level is stored on solid state disks that can support very fast random I/Os, so that traversing the extra level of indirection incurs a relatively small overhead.

The extra level of indirection dramatically reduces the number of magnetic disk I/Os that are needed for index updates, and localizes maintenance to indexes on updated attributes. Further, we batch insertions within the indirection layer in order to reduce physical disk I/Os for indexing new records. By reducing the index maintenance overhead on transactions, we enable operational data stores to create more indexes to support queries. We have developed a prototype of our indirection proposal by extending the widely used Generalized Search Tree (GiST) open-source project, which is also employed in PostgreSQL. Our working implementation demonstrates that we can significantly reduce index maintenance and/or query processing cost, by a factor of 3. For insertions of new records, our novel batching technique can save up to 90% of the insertion time.

1. INTRODUCTION

In a multiversion database system, new records do not physically replace old ones. Instead, a new version of the record is created, which becomes visible to other transactions at commit time. Conceptually, there may be many rows for a record, each corresponding to the state of the database at some point in the past. Very old versions may be garbage-collected as the need for old data diminishes, in order to reclaim space for new data.

When indexing data, one typically indexes only the most recent version of the data, since that version is most commonly accessed. In such a setting, record insertions, deletions and updates trigger I/O to keep the indexes up to date. With a traditional index struc-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Articles from this volume were invited to present their results at The 39th International Conference on Very Large Data Bases, August 26th - 30th 2013, Riva del Garda, Trento, Italy.

Proceedings of the VLDB Endowment, Vol. 6, No. 11

Copyright 2013 VLDB Endowment 2150-8097/13/09... \$ 10.00.

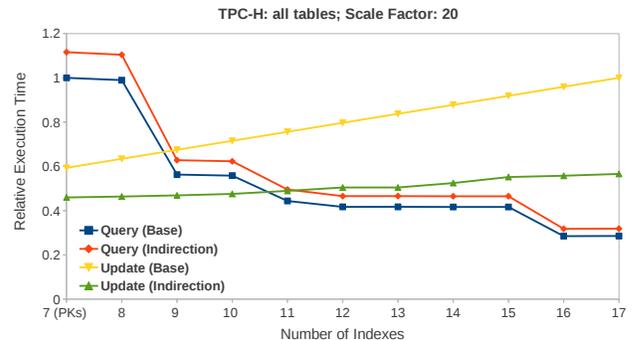


Figure 1: Effect of Adding Indexes in Operational Data Store

ture, the deletion of a record requires the traversal of each index and the removal of the row-identifier (RID) from the leaf node. The update of a record (changing one attribute value to another) creates a new version, triggering a traversal of all indexes to change the RIDs to the new version's RID. (In the case of the modified attribute, the position of the record in the index may also change.) For a newly inserted record, the new RID must be inserted into each index. Indexes may be large, and in aggregate much too large to fit in the RAM bufferpool. As a result, all of these index maintenance operations will incur the overhead of physical I/O on the storage device.

These overheads have historically been problematic for OLTP workloads that are update-intensive. As a result, OLTP workloads are often tuned to minimize the number of indexes available. This choice makes it more difficult to efficiently process queries and to locate records based on secondary attributes. These capabilities are often important for operational data stores [35]. For example, it is not uncommon to find tens of indexes to improve analytical and decision-making queries even in TPC benchmarks [20, 19] or enterprise resource planning (ERP) scenarios [14, 13].

Our goal is to reduce the overhead of index updates, so that indexes can be used effectively for analytical query processing without being a heavy burden on transaction throughput. The query vs. update dilemma is clearly captured in Figure 1, a preview of our experimental results. The execution time for analytical queries is reduced as more indexes are added. However, this reduction comes at the cost of increasing the update time in the Base approach. In contrast, by employing our technique (the Indirection approach) the incurred update cost is significantly smaller.

To address this dilemma, we utilize a solid state storage layer. Based on current technologies, solid state disks (SSDs) are orders of magnitude faster than magnetic disks (HDDs) for small random I/Os. However, per gigabyte, SSDs are more expensive than magnetic disks. It therefore pays to store the bulk of the data on magnetic disk, and reserve the SSD storage for portions of the data that

can benefit the most, typically items that are accessed frequently and randomly. In this paper we describe a prototype of our Indirection approach based on the widely used Generalized Search Tree (GiST) package [31, 29], which is also employed in industrial-strength open-source database management systems such as PostgreSQL [6], PostGIS [5], OpenFTS [4], and BioPostgres [1].

Unlike previous approaches [17, 37, 15, 18, 23], we do not propose to simply store “hot” data on SSDs. Instead, we change the data structures in the storage hierarchy to employ an extra level of indirection through solid state storage. Because the solid state memory is fast, the extra time incurred during index traversal is small, as we demonstrate experimentally. The extra level of indirection dramatically reduces the amount of magnetic disk I/O that is needed for index updates. Only SSD I/O is needed for deletions and updates, with the exception of indexes on changed attributes. We can also reduce the magnetic disk I/O overhead for insertions. While we describe our techniques in terms of SSDs, we are not limited to a disk-like form factor. In fact, alternative form factors (e.g., FusionIO auto-commit memory [27]) with smaller I/O granularities would provide even better performance because our proposed solid state updates are small.

Another potential advantage of our work – efficient index maintenance of multiversion databases – is to support multiversion concurrency control (MVCC). The MVCC model is reviving [6, 32, 38, 39] mostly due to the increased concurrency available in modern hardware such as large main memories and multicore processors. But this increased concurrency comes at the cost of increased locking contention among concurrent reads/updates queries, which could be alleviated using optimistic locking over a multiversion database.

1.1 Multiversion Databases

By keeping old data versions, a system can enable queries about the state of the database at points in the past. The ability to query the past has a number of important applications [50], for example, (1) a financial firm requires to retain any changes made to client information for up to five years in accordance with auditing regulations; (2) a retailer ensures to offer only one discount for each product at any given time; (3) a bank needs to retroactively correct an error for miscalculating the promised introductory interest rate. In addition to these business-specific scenarios, there is an inherent algorithmic benefit for retaining the old versions of the record and avoiding in-place, that is, to utilize efficient optimistic locking and latch-free data structures.

A simple implementation of a multiversion database would store the row-identifier (RID) of the old version within the row of the new version, defining a linked list of versions. Such an implementation allows for the easy identification of old versions of each row, but puts the burden of reconstructing consistent states at particular times on the application, which would need to keep timing information within each row.

To relieve applications of such burdens, a multiversion database system can maintain explicit timing information for each row. In a valid time temporal model [28] each row is associated with an interval [begin-time,end-time) for which it was/is current. Several implementation choices exist for such a model. One could store the begin-time with each new row, and infer the end-time as the begin-time of the next version. Compared with storing both the begin-time and end-time explicitly for each row, this choice saves space and also saves some write I/O to update the old version. On the other hand, queries over historical versions are more complex because they need to consult more rows to reconstruct validity intervals.

A *bitemporal* database maintains two kinds of temporal information, the system (i.e., transaction) time, as well as the application time (sometimes called “business time”).

In this work we do not commit to any one of these implementation options, each of which might be a valid choice for some workloads. For any of these choices, our proposed methods will reduce the I/O burden of index updates. Some of our techniques, such as the LIDBlock technique (Section 3.1), apply to both versioned and non-versioned databases.

1.2 Physical Organization

There are several options for the physical organization of a temporal database. A complete discussion of the alternatives is beyond the scope of this paper. We highlight two options that have been used in commercial systems, namely, the *history table* vs. the *single table* approach.

One organization option appends old versions of records to a *history table* and only keeps the most recent version in the main table, updating it in-place. Commercial systems have implemented this technique: In IBM DB2 it is called “System-period data versioning” [34], and it is used whenever a table employs transaction time as the temporal attribute. The Oracle Flashback Archive [47] also uses a history table. Such an organization clusters the history table by end-time, and does not impose a clustering order on the main table. Updates need to read and write the main table, and also write to the end of the history table. Because updates to the main table are in-place,¹ an index needs to be updated only when the corresponding attribute value changes. For insertions and deletions, all indexes need to be updated. In short, using the history table approach (1) the temporal ordering of the data is lost; (2) additional random I/Os are required to perform in-place updates of records; (3) the number of database objects (e.g., tables, indexes, and constraints) are potentially doubled, which increases the overall management and maintenance cost of database and slows down the query optimization runtime; and (4) less effective query plan is constructed for certain temporal range queries that are forced to union the history and the main tables instead of using range-partitioned table for maintaining the single table approach.

In this paper, we assume an organization in which there is a *single table* containing both current and historical data. Commercial systems that implement this technique include Oracle 11g where the concept is called “version-enabled tables” [46]. IBM’s DB2 also uses this approach for tables whose only temporal attribute is the application time. The single table approach is central to IBM DB2 with BLU Acceleration as well [3]. New rows are appended to the table, so that the entire table is clustered by begin-time. Updates need to read the table once and write a new version of the record to the end of the table.

We focus on applications that primarily use current data, but occasionally need to access older versions of the data. To support queries over current data, the most recent data may be extensively indexed. Older data may be less heavily indexed because it is queried less frequently, and is often more voluminous. Even within a single table, the system can offer an implementation option in which only the most recent version of a record appears in an index.

2. BASIC INDIRECTION STRUCTURE

Traditional index structures directly reference a record via a pointer known as a physical row-identifier (RID). The RID usually encodes

¹If one wanted to cluster the main table by a temporal attribute to improve temporal locality, then updates would not be in-place and additional indexes would need to be updated. Our proposed solution would reduce the burden of such index updates.

a combination of the database partition identifier, the page number within the partition, and the row number within the page. A RID index over current HDD-resident data is shown in Figure 2.

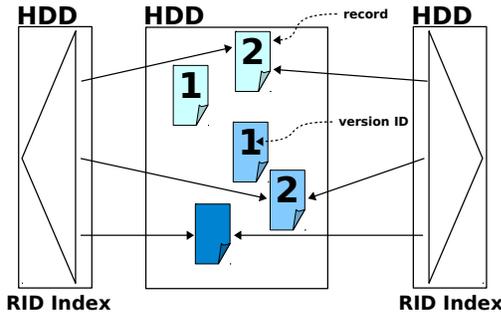


Figure 2: Traditional RID index structure

The choice of a physical identifier hinders the update performance of a multiversion database in which updates result in a new physical location for the updated record. Changes to the record induce I/O for every index, even indexes on “unaffected” attributes, i.e., attributes that have not changed. Random I/Os are required to modify HDD-resident leaf pages.

To avoid HDD I/O for indexes on unaffected attributes, we decouple the physical and logical representations of records spanning many versions. We distinguish between a physical row-identifier (RID) and a logical record identifier (LID). For any given record, there may be many RIDs for that record corresponding to the physical placement of all of the versions of that record. In contrast, the LID is a reference to the RID representing the most recent version of the record. For now, one can think of a table $LtoR(LID, RID)$ that has LID as the primary key. Indexes now contain LIDs rather than RIDs in their leaves.

Under our proposed Indirection technique, an index traversal must convert a LID to a RID using the $LtoR$ table. A missing LID, or a LID with a NULL RID in the $LtoR$ table are treated as deleted rows, and are ignored during search. By placing the $LtoR$ table on an SSD, we ensure that the I/O overhead for the extra indirection is relatively small.² Because the SSD is persistent, index structures can be recovered after a crash. Because we need only a few SSD bytes per record, it is possible to handle a large magnetic disk footprint with a much smaller solid state footprint. The new index design is demonstrated in Figure 3.

When an existing record is modified, a new version of that record is created. The $LtoR$ table is updated to associate the new row’s RID to the existing LID. That way, indexes on unchanged attributes remain valid. Only for the changed attribute value will index I/O be required for the indirection layer.

When a record is deleted, the (LID,RID) pair for this record in the $LtoR$ table is deleted. Index traversals ignore missing LIDs. Indexes can lazily update their leaves during traversal, when a read I/O is performed anyway. At that time, any missing LIDs encountered lead to the removal of those LIDs from the index leaf page. After a long period of activity, indexes should be validated off-line against the $LtoR$ table to remove deleted LIDs that have subsequently never been searched for.

When a new record is added, the new record is appended to the tail of the relation and its RID is fetched and associated to a new LID. The (LID, RID) pair for the new record is added to the $LtoR$

²Frequently accessed LIDs would naturally be cached in RAM by the database bufferpool manager, further reducing the overhead.

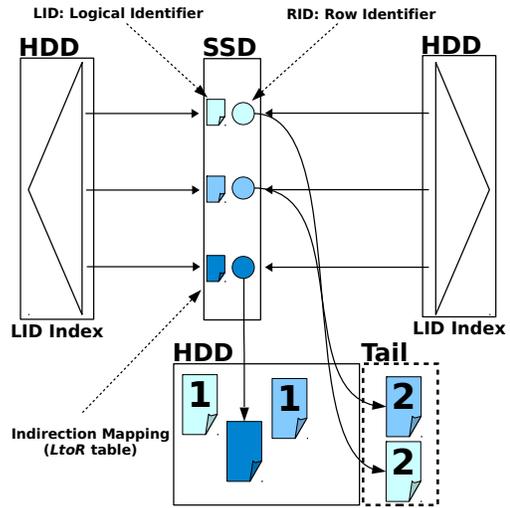


Figure 3: LID index using the indirection technique

table. All indexes are also updated with the new record LID accordingly. In Section 3, we discuss how to further improve record insertion and deletion.

3. ENHANCING INSERTIONS

We now develop techniques for improving the index performance of insertion. We define a batching structure called a LIDBlock, and employ yet another level of indirection.

3.1 LIDBlocks

To reduce the index overhead for insertions, we propose an SSD-resident auxiliary LIDBlock structure containing a fixed number of LIDs. The LIDs in a LIDBlock may be NULL, or may be valid LIDs from the $LtoR$ table. References to a LIDBlock are mapped to multiple LIDs through this extra level of indirection. Figure 4 shows the extended structure with LIDBlocks.

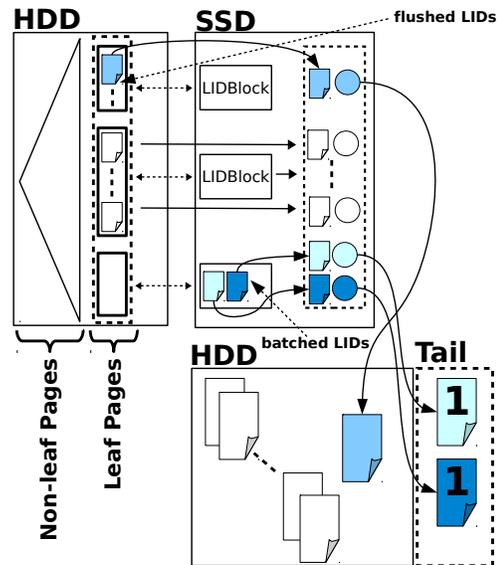


Figure 4: Indirection Technique with LIDBlock

The arrow from index leaf pages to LIDBlocks in Figure 4 could be implemented by keeping a block identifier (BID) within the leaf

page. The disadvantage of such a choice is that the leaf node of the index needs to be read from magnetic disk to locate the BID, requiring extra HDD I/O. Instead, we propose to store LIDBlocks within hash tables on the SSD. In the following sections, we describe more precisely how index pages refer to and traverse LIDBlocks.

3.2 The Dense Index Case

Consider first the case of a dense index, i.e., a secondary index where there are many records (more than a leaf-node’s worth) for each attribute value. For such indexes, we keep a list of LIDs for each value as before. In addition, we store a collection of LIDBlocks on the SSD in a hash table, hashed by the attribute value. Initially, each indexed value has a LIDBlock whose LIDs are all NULL.

When a new record is inserted, we need to modify the indexes to reflect the insertion. Suppose that the value of an indexed attribute is v , and that the index is dense. A LID is created for the record, and a suitable (LID,RID) pair is added to the *LtoR* table. The LIDBlock B for v is identified by accessing the hash table of LIDBlocks on the SSD. If there are unused (NULL) slots in B , one of the slots is overwritten with the LID of the new record. If there are no unused slots, then all LIDS in B and the LID of the new record are moved in bulk into the LID list in the index, amortizing the I/O cost.³

In this model, index traversal is slightly more complex: all the non-NULL LIDs in the LIDBlock for a value also need to be treated as matches. Deletions and attribute value updates may also need to traverse and modify a LIDBlock. There is some additional I/O, but only on solid state storage.

3.3 The Sparse Index Case

When there are few matches per index value, the organization above would need a very large number of LIDBlocks, most of which would be underutilized. Instead, for sparse indexes we maintain a single LIDBlock for an entire leaf node of the index. Rather than using a hash table hashed by attribute value, we use a hash table hashed by the address of the index leaf page. This address can be obtained using a partial traversal of the index, without accessing the leaf node itself. Since the internal nodes of a tree index occupy much less space than the leaf nodes, they are much more likely to be resident in the main memory bufferpool.

Searches have some additional overhead, because the shared LIDBlock would need to be consulted even for records that may not match the search condition. There would also be overhead during node splits and merges to maintain the LIDBlock structure.

The overhead of LIDBlocks on searches may be high for sparse indexes. For example, a unique index search would previously only have to look up one main file record. With a LIDBlock for a given key range, a search may need to read b of them, where b is the LIDBlock capacity. This example suggests an optimization: store both the LID and the key in the LIDBlock for sparse indexes. This optimization reduces the capacity of LIDBlocks, but significantly improves the magnetic disk I/O for narrow searches.

3.4 Revisiting Deletions and Updates

With LIDBlocks, it is now possible that an update or deletion occurs to a record whose LID is in a LIDBlock rather than the LID list. For deletions, one can simply set the LID to NULL in the LIDBlock. For updates that change the position of the record in the index, one needs to nullify the LID in the previous LIDBlock or LID-list, and insert it into the new LIDBlock.

³One can shift the burden of LIDBlock flushing outside of running transactions by triggering an asynchronous flush once a LIDBlock becomes full (or nearly full), rather than waiting until it overflows.

In this way, LIDBlocks also improve the I/O behavior of updates to indexes on changed attributes. In the original scheme, HDD-resident index leaf pages appropriate to the new attribute value needed to be updated with the LID of the updated record. In the enhanced scheme, writes are needed only on SSD-resident LIDBlocks most of the time. Magnetic disk writes are amortized over the number of records per LIDBlock.

4. EXTENDED EXAMPLE

In this section, we provide a detailed example to further illustrate the core of our Indirection proposal and motivate the analysis in Section 5.

Consider a table $R(A, B, C, D)$ with B-tree indexes on all four attributes. A is the primary key, and is, therefore, sparse. B is also sparse, while C and D are dense. R is stored on disk in a versioned fashion. For each row we show the RID of the previous version (if any) of the row; the previous-RID may or may not be explicitly stored. Suppose that at a certain point in time, the extension of R includes the rows given below. A flag indicating whether the row has been deleted is also included.

RID	A	B	C	D	Prev-RID	Deleted
345	100	3732	3	5	123	0
367	120	4728	3	6	NULL	0
369	130	2351	2	5	NULL	0
501	100	3732	2	5	345	0

Suppose the *LtoR* table for R is given by

LID	RID
10	367
11	369
13	501

Indexes use LIDs (e.g., 10, 11, 13) rather than RIDs to refer to rows, and only the most recent version is indexed. Given the above database, the immediate and deferred changes to the database in response to various update requests are described below. The I/O needed is summarized in square brackets (log I/O and I/O for index node splits/merges are ignored). Our I/O estimates assume that all internal index nodes are resident in the RAM bufferpool, and that all leaf nodes require I/O. These estimates also assume that all accesses to SSD-resident structures require I/O.⁴ The estimates also assume direct access to a page given the LID, as might be supported by a hash table. We assume that LIDBlocks contain b LIDs, and that for sparse indexes we are storing both the key and the LID in the LIDBlock. We now describe precisely each key operation, breaking down its steps into immediate and deferred actions. Immediate actions must be completed within the operation itself. Deferred actions are those that can happen later, such as when an operation causes a page to become dirty in the bufferpool but the actual I/O write comes later.

Update. We update the row with key 100 such that attribute D is changed from 5 to 6. The immediate actions are

- The LID of the row in question (i.e., 13) is identified using the index on A [1 HDD read].
- The entry (13, 501) in the *LtoR* mapping is read [1 SSD read].
- The row with RID 501 is read [1 HDD read].

⁴This might be too pessimistic, particularly for LIDBlocks that could be small enough to be cached in RAM.

- A new version of the record is created at the tail of R with a new value for D and a new RID (suppose 601).
- An index traversal for key 5 is initiated on attribute D . If LID 13 is present in the corresponding leaf, it is deleted; otherwise, LID 13 is located in the LIDBlock for key 5 and is removed from the LIDBlock. [1 HDD read, possibly 1 SSD read]
- A partial index traversal for key 6 is initiated on attribute D , and LID 13 is inserted into the corresponding LIDBlock [1 SSD read].

The required deferred actions are summarized as follows.

- The data page containing the row with RID 601 is dirty and will need to be flushed to the HDD. [1 HDD write amortized over all modifications to the page⁵].
- The entry (13, 501) in the $LtoR$ mapping is changed to (13, 601) [1 SSD write].
- The index page containing the key 5 will be dirty if LID 13 was in the leaf. The dirty page will need to be flushed to the HDD [1 HDD write, amortized over all modifications to the page]. If LID 13 was in the LIDBlock for key 5 then the dirty LIDBlock will need to be flushed to the SSD [1 SSD write].
- The LIDBlock for key 6 is dirty and will need to be flushed to the SSD [1 SSD write].

Insertion. Consider the insertion of a new record (140, 9278, 2, 6). The resulting immediate actions are:

- The absence of a previous version of the row is verified using the index on A , including the LIDBlock [1 HDD read, 1 SSD read].
- A new row is created at the tail of R , with a new RID (suppose 654). The previous-RID field is NULL.
- A new LID (suppose 15) is allocated.
- For the two dense indexes on C and D , the LIDBlocks for keys 2 and 6 (respectively) are identified, and LID 15 is inserted into each [2 SSD reads].
- For the two sparse indexes on A and B , the LIDBlocks for keys 140 and 9278 (respectively) are identified using partial index traversals, and LID 15 is inserted (paired with the key) into each [2 SSD reads].

The deferred actions are as follows.

- The data page containing the row with RID 654 is dirty and will need to be flushed to the HDD [1 HDD write, amortized over all modifications to the page].
- The entry (15, 654) is inserted into the $LtoR$ mapping [1 SSD read, 1 SSD write].
- The LIDBlocks for each of the four indexes are dirty, and need to be flushed [4 SSD writes].

⁵Amortization is expected to be high on the tail of a table.

- In the event that a LIDBlock fills (one time in b insertions), we need to convert all LIDs in the LIDBlock into regular index LIDs, and reset the LIDBlock to an empty state [4/b SSD writes, 3/b HDD reads (the leaf in the index on A has already been read), 4/b HDD writes].

Deletion. Now suppose deleting the row with key 100, which results in the following immediate actions:

- The LID of the row in question (i.e., 13) is identified using the index on A [1 HD read, possibly 1 SSD read].
- The pair (13, 501) in the $LtoR$ table is located [1 SSD read].
- The row with RID 501 is read and the deleted flag is set to 1 [1 HDD read].
- LID 13 is removed from the leaf node of the index on A .

The deferred actions for deleting row with key 100 are

- The data page containing the row with RID 501 is dirty and will need to be flushed to the disk. [1 HDD write, amortized over all modifications to the page].
- The pair (13, 501) in the $LtoR$ table is dropped [1 SSD write].
- The index leaf page for A containing the key 100 is dirty and will need to be flushed to the HDD [1 HDD write amortized over all modifications to the page].
- Whenever one of the other indexes is traversed and LID 13 is reached, LID 13 will be removed from the corresponding LID list [1 extra HDD write to modify the leaf, amortized over all modifications to the page].

Search. Suppose that the search returns m matches that all fit in one index leaf page of a sparse index. Since no write is involved, the search consists of only immediate actions:

- Traverse the index [1 HDD read].
- Read the LIDBlock for the leaf node [1 SSD read].
- Map LIDs to RIDs [m SSD reads].
- Read all matching records [m HDD reads (assuming an unclustered table)].

5. ANALYSIS

We analyze the performance of the Indirection method according to three criteria: (a) time for core operations; (b) SSD space requirements; (c) SSD wear-out thresholds.

5.1 Time Complexity

Table 4 summarizes the I/O complexity of the Base and Indirection methods. Most I/Os are random, meaning that SSD I/Os are much faster than HDD I/Os, by about two orders of magnitude on current devices. It is therefore worth investing a few extra SSD I/Os to save even one HDD I/O. Note that these estimates are pessimistic in that they assume that none of the SSD-resident data is cached in RAM. If commonly and/or recently accessed data was cached in RAM, many SSD read I/Os could be avoided, and many SSD writes could be combined into fewer I/Os.

The most striking observation is that with a small increase in SSD I/Os the costs of updates, insertions, and deletions are substantially reduced, while the cost of searches increases only slightly. With the Indirection technique, the immediate HDD cost is independent of the number of indexes.

Technique	Type	Immediate SSD	Deferred SSD	Immediate HDD	Deferred HDD
Base	Single-attr. update	0	0	$3 + k$	$\leq 2 + k$
	Insertion	0	0	$1 + k$	$\leq 1 + k$
	Deletion	0	0	$2 + k$	$\leq 1 + k$
	Search Uniq.	0	0	2	0
	Search Mult.	0	0	$\lceil m/I \rceil + m$	0
Indirection	Single-attr. update	2	2	3	≤ 2
	Insertion	$1 + k$	$2 + k + k/b$	1	$\leq 1 + (2k - 1)/b$
	Deletion	1	1	2	$\leq 1 + k$
	Search Uniq.	2	0	2	0
	Search Mult.	$\lceil m/I \rceil + m$	0	$\lceil m/I \rceil + m$	0

Table 1: Base vs. Indirection technique analysis. For single-attribute updates and deletions, we show the case where the LID was initially in the leaf rather than the LIDBlock. n is the number of attributes, k is the number of B-tree indexes, m is the number of results returned by a search, and I is the number of keys in an index leaf node.

5.2 SSD Space Complexity

We now estimate the space needed on the SSD for the *LtoR* table and the LIDBlocks. If N is the number of latest-version records in a table, then we need N (LID,RID) pairs in the *LtoR* table. In a typical scenario, table records may be 150 bytes wide, with 8 byte LIDs and 8 byte RIDs. Assuming a fudge-factor of 1.2 for representing a hash table, the *LtoR* table would constitute roughly 13% of the size of the current data in the main table.

We now consider the space consumption of LIDBlocks for sparse indexes; dense indexes would take less space. The number of LIDBlocks for a single sparse index is equal to the number of leaves in the index. With an 8 byte key, an 8 byte LID, an 8KB HDD page size and a 2/3 occupancy factor, a leaf can represent about 340 records. A LIDBlock contains b (LID,key) pairs, leading to a total space consumption of $16bN \times 1.2/340$ bytes per index. Even a small value of b , say 16, is sufficiently large to effectively amortize insertions. At this value of b , the LIDBlocks for a single index constitute less than 1% of the size of the current data in the main table.

Thus, even with 7 indexes per table, the SSD space required in a typical scenario is only about 20% of the HDD space required for the current data. Taking the historical data into account, the relative usage of SSD space would be even lower. Given that SSDs are more expensive per byte than HDDs, it is reassuring that a well-provisioned system would need much less SSD capacity than HDD capacity.

5.3 SSD Life-Expectancy

SSDs based on enterprise-grade SLC flash memory are rated for about 10^5 erase cycles before a page wears out [40]. SSDs based on phase-change memory are rated for even higher rates. SSDs internally implement wear-leveling algorithms that spread the load among all physical pages, so that no individual page wears out early.

In flash-based SSDs there is a write-amplification phenomenon in which internal movement of data to generate new erase units adds to the application write workload. This amplification factor has been estimated at about 2 for common scenarios [25].

To estimate the wear-out threshold, suppose that the Indirection method uses an SSD page size of 512 bytes. Then a device⁶ with a capacity of 8×10^7 KB can tolerate 8×10^{12} writes before wear-out, assuming a write-amplification of 2 and 10^5 erases per page.

Our most write-intensive operation is insertion, with about $1 + k$ SSD writes per insertion when there are k indexes.⁷ Assuming nine indexes, and ten insertions per transaction, we would need 100 SSD writes per transaction. Even running continuously at the high rate

⁶Device characteristics are based on an 80GB FusionIO device.

⁷Again, this estimate is pessimistic because it assumes no RAM caching of SSD pages.

of 800 transactions/second, the device would last more than three years.

6. EXPERIMENTAL EVALUATION

We present a comprehensive evaluation of our Indirection prototype based on the Generalized Search Tree (GiST) index package [31, 29]. Additionally, we provide complimentary kinds of evidence using DB2 [2] in two ways to further support the results demonstrated using our prototype. Since we are targeting operational data stores with both transactions and analytic queries, we base our evaluation on the TPC-H benchmark [54].

First, in Section 6.2 we try to answer the question “How would the performance of a state-of-the-art database change if it were to use our methods?” We employ the commercial database system DB2, but other database systems would have been equally good choices. Given a workload W , we construct a rewritten workload W' for the DB2 engine that simulates the I/O behavior of our technique for W . While the workload W' is not identical to W , we argue that the performance of DB2 on W' provides a conservative estimate of the performance of W on a (hypothetical) version of DB2 that implements the Indirection technique.

Second, in Section 6.3, we evaluate our Indirection technique by implementing it within the popular GiST index package [31, 29]. The GiST package have successfully been deployed into a number of well-known open-source projects including PostgreSQL [6], PostGIS [5], OpenFTS [4], BioPostgres [1], and YAGO2 [7]. All aspects of the method (insertions, deletions, modifications) have been implemented; we refer to the resulting prototype as LIBGiST^{mv}. We profile the I/O behavior of LIBGiST^{mv}, and create a detailed I/O and execution cost model for the Indirection technique.

Finally, in Section 6.4, we shift our focus to TPC-H style analytical query processing that is geared towards an operational data store. We provide evidence for the key tenet of this work, namely, *reducing the burden of index maintenance means that the system can afford more indexes, which in turn improves the performance of analytical query processing.*

6.1 Platform

Experiments were run on two machines. The first machine, dedicated to our DB2 experiments, was running Fedora 8 and was equipped with a Quad-core Intel Xeon CPU E5345 running at 2.33 GHz, having 4GB of RAM, three magnetic disks (7200 RPM SATA), and one 80GB Single Level Cell (SLC) FusionIO solid state drive.⁸ The configuration of our second machine running Ubuntu 10.4, used exclusively for our LIBGiST^{mv} experiments, was a 6-core Intel Xeon X5650 CPU running at 2.67GHz, having 32GB of RAM,

⁸Our SLC FusionIO can support up to 88,000 I/O operations per second at $50\mu\text{s}$ latency.

	Base	Indirection
Query	SELECT * FROM LINEITEM L WHERE ?	SELECT * FROM LINEITEM L, LtOR M WHERE ? AND L.LID = M.LID
Update	SELECT * FROM LINEITEM L WHERE ? INSERT INTO LINEITEM VALUES (?, ..., ?)	SELECT * FROM LINEITEM L, LtOR M WHERE ? AND L.LID = M.LID INSERT INTO LINEITEM VALUES (?, ..., ?) UPDATE LtOR SET RID = ? WHERE LID = ?

Table 2: Query re-writing to capture indirection mechanism

one magnetic disk (7200 RPM SATA), and one 80GB SLC FusionIO solid state drive.⁹

In our experiments, we used IBM DB2 version 9.7 [2]. We configured DB2 with adequate bufferpool size (warmed up prior to starting the experiments) to achieve an average 90% hit ratio on both data and index pages. For the DB2 experiments, we generate a TPC-H database [54] with scale factor 20. File system caching was disabled in all experiments.

For LIBGiST^{mv}, we extended LIBGiST v.1.0 to a multiversion generalized search tree C++ library that supports our Indirection techniques including LIDBlocks.

6.2 DB2 Query Re-writing Experiments

The goal of the query re-writing experiment is to study the I/O pattern for both the unmodified DB2 system (“Base”) and the Indirection approach. To evaluate Base for a query Q we simply run Q in the original schema S . To evaluate Indirection we rewrite Q into another query Q' . Q' is run in a schema S' containing an explicit LtOR table representing the LID-to-RID mapping. Ideally, the LtOR table is physically located on the SSD device; we empirically examine the impact of the location below. In S' , base tables are augmented with an additional LID column, where the value of LID is generated randomly. In S we build as many indexes as desired on the base tables. In S' we build a single index on the attribute selected in the query, typically the primary key of the base table.

6.2.1 Rewriting Queries

For queries, the rewriting simply adds the LtOR table to the FROM clause, with a LID equijoin condition in the WHERE clause. An example template of our query re-writing that simulates the indirection mechanism is shown in Table 2. The queries are written over TPC-H LINEITEM table and the indirection table, denoted by LtOR.

To see why the performance of Q' is a conservative estimate of the cost of the indirection technique, consider two cases for the query in Table 2. In the first case, some selection condition on an indexed attribute (or combination of conditions) is sufficiently selective that an access plan based on an index lookup is used. This case includes a point query specified as an equality condition on a key attribute. The Base plan for Q would include an index traversal and a RID-based lookup in the LINEITEM table. For Q' , we will also have an index traversal and a RID-based lookup of

LINEITEM, together with a LID-based lookup of the LtOR table. This is precisely the I/O pattern of the Indirection technique.

In the second case, the selection conditions are not very selective. In such a case, the system is likely to scan the base table in answering Q .¹⁰ To answer Q' in such a case requires a join of LINEITEM and LtOR in order to make sure that we only process the most recent versions of each record. This may actually an over-estimate of the cost needed by the Indirection technique, because the Indirection technique can also employ a scan without consulting the LtOR table.

6.2.2 Rewriting Updates

For updates, an extra UPDATE statement is added to keep the LtOR table current, as illustrated in Table 2. Since we are simulating a multiversion database, the Base method inserts a new row rather than updating an existing row. While the Base method pays the cost of inserting the new row into each index, we are slightly favoring the Base method by ignoring the cost of deleting the old row from the indexes. Depending on the implementation technique used for temporal attributes (Section 1.1) there may be additional I/O required to update the temporal attributes of the old row, but we ignore such I/O here.

For updates in which a single attribute value is modified, the Indirection method incurs just one index update and one update to the LtOR table. At first it may seem like there is more work done by the Indirection method simulation for the updates of Table 2, since the INSERT statements are the same and the Indirection method has an extra UPDATE statement. This impression is true only for the case in which there is one base table index in the base schema S . As soon as there are multiple indexes on base tables in S , the cost of the INSERT statement for the Base method exceeds that of the corresponding INSERT in the Indirection method because more indexes need to be updated.

Our profiling of update statements can easily be extended to delete statements, but we omit such profiling because the performance profile would be similar to that for updates. On the other hand, our rewriting does not model the LIDBlock technique, and thus cannot fully capture its performance advantages for insertions. The benefits of the LIDBlock technique will be evaluated in Section 6.3.

6.2.3 Results

All DB2 measurements are averages over 5 million randomly chosen queries/updates with the exception of our selectivity exper-

⁹While the FusionIO devices are high-end devices that are relatively expensive, we remark that recent SSDs such as the Intel 520 series can store about 500GB, cost about \$500, and can support 50,000 I/O operations per second at 85 μ s latency, more than sufficient for the workloads described here.

¹⁰We do not include old versions of records for these experiments; in a true multiversion database the base tables would contain some old record versions that would need to be filtered out during the scan, using the valid time attributes.

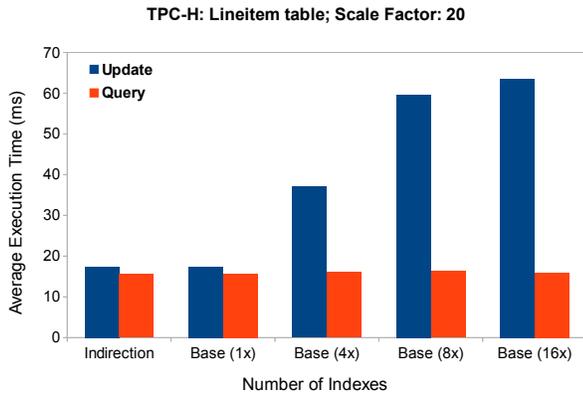


Figure 5: Varying the number of indexes vs. update/query time

iments and analytical queries, in which fewer operations were performed to complete the experiments in a reasonable time.

Effect of Indexes on Execution Time. Our first result confirms that adding the extra indirection layer has negligible overhead for query processing as shown in Figure 5. For this experiment, queries are point queries that are accessed by specifying the value of the primary key attribute. The query overhead is negligible because the indirection mapping from LID-to-RID requires only a single additional random SSD I/O, a delay that is orders of magnitude faster than the necessary random magnetic disk I/O for retrieving data pages holding entire records. Figure 5 also shows that the update execution time for the Base technique increases dramatically as the number of indexes is increased. With 16 indexes, Indirection outperforms Base by a factor of 3.6.

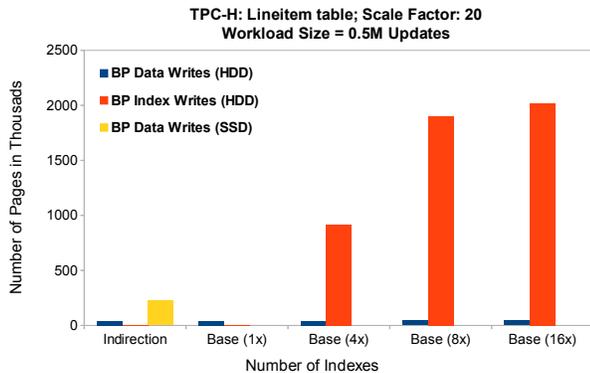


Figure 6: Varying the number of indexes vs. page writes

Digging deeper into the bufferpool behavior of DB2 also reveals that with a modest number of SSD page writes, the number of magnetic disk index writes are substantially reduced, as shown in Figure 6. This result demonstrates the effectiveness of Indirection in reducing the index maintenance cost for multiversion databases.

Effect of Query Selectivity and Index Clustering. Consider a range query over a single indexed attribute in the `LINEITEM` table. By varying the width of the range, we can vary the query selectivity.¹¹ Figure 7 shows the results for various selectivities, where the indexed attribute is the key of the `LINEITEM` table by which the table is clustered. On average the query overhead (for consulting the

¹¹The starting value of the range in our range queries are chosen randomly.

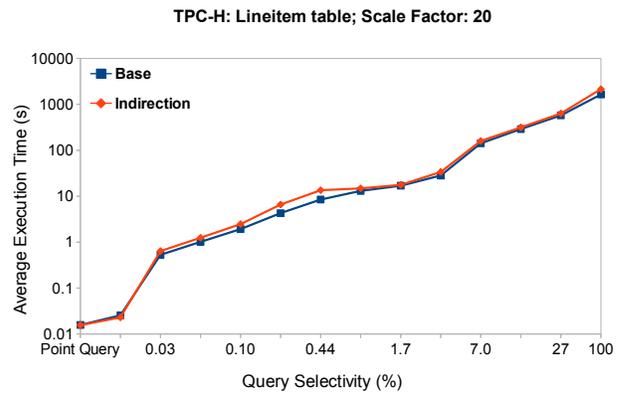


Figure 7: Varying the query selectivity vs. query execution time

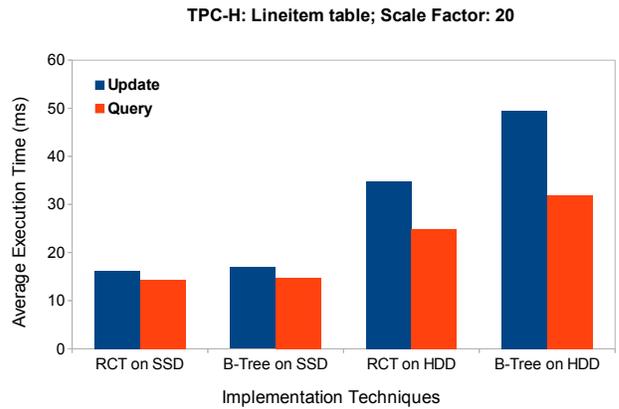


Figure 8: Varying the indirection implementation techniques

`LTOR` table) remains under 20% as query selectivity varied. There is a sudden jump at around 0.22% selectivity, but a closer examination of the DB2 query plans reveals that the sudden increase in execution time is attributable to the optimizer incorrectly switching from a nested-loops to a merge-sort join plan.¹²

The 20% overhead, while still small, is higher than the negligible overhead seen for point queries in Figure 5. We also observed that as the index clustering ratio decreases, the query processing gap between Indirection and Base decreases. For example, for the lowest clustering ratio index of the `LINEITEM` table (on the `SUPPKEY` attribute), the overhead drops to only 4% on average. These differences can be understood as a consequence of caching. With a high clustering ratio, one base table disk page I/O will be able to satisfy many access requests, meaning that each magnetic disk I/O is amortized over many records. On the other hand, every record will contribute an SSD I/O for the `LTOR` table since that table is not suitably clustered. For point queries and queries over an unclustered index, the ratio of SSD I/O operations to magnetic disk I/Os will be close to 1.

Indirection Mapping Implementation. Finally, we illustrate that the indirection random read/write access pattern is ideal for SSDs and not for magnetic disks. We tried two different implementations of the LID-to-RID mapping using either a DB2 range-clustered table (RCT) or a traditional B-Tree index hosted on either SSDs or HDDs. As shown in Figure 8, when an SSD is used, the

¹²This behavior can be avoided if optimizer hints are provided.

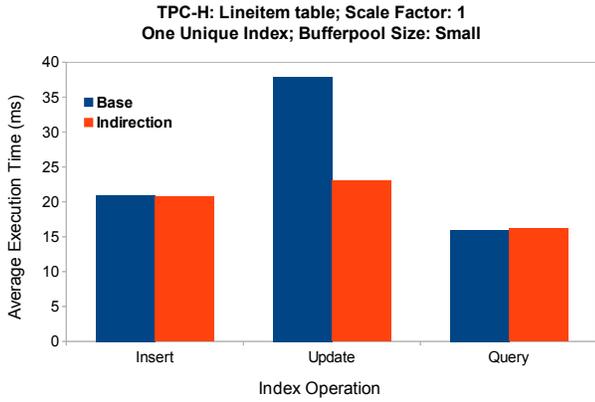


Figure 9: Insert/Update/Query execution time

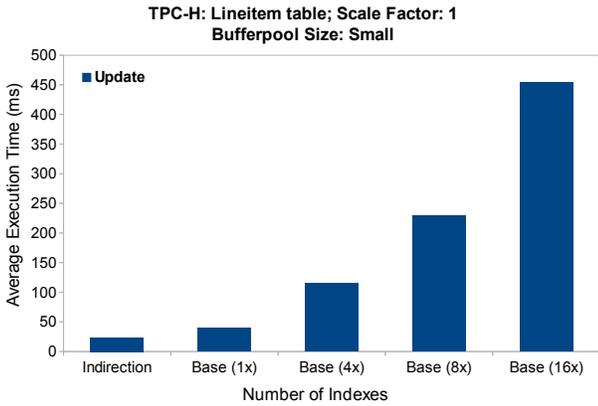


Figure 10: Varying the number of indexes vs. update time

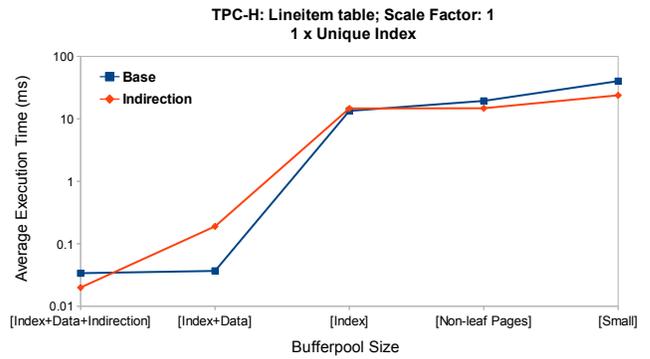
overall query and update cost is 1.9X and 2.7X lower, respectively, than on an HDD.

6.3 GiST Implementation

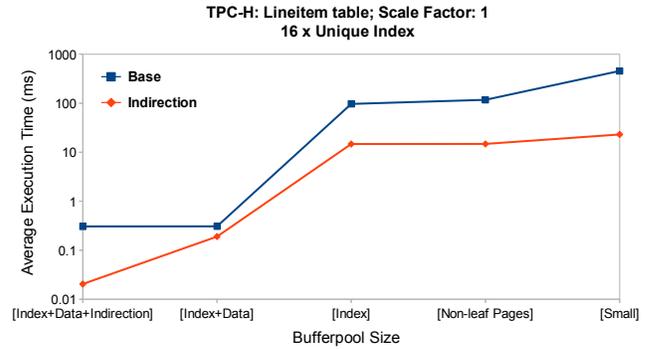
Our LIBGiST^{mv} codebase directly implements all aspects of our proposed approach, including the LIDBlock technique. We employ LIBGiST^{mv} as the basis for a systematic performance study in a controlled environment. All HDD and SSD reads and writes used *Direct I/O* to avoid extra copying of operating-system buffers. No bufferpool space was allocated for the indirection table, so requests to the indirection table always incur SSD I/Os.

In our prototype, we also extended the LIBGiST bufferpool in order to test a variety of memory configurations. We further enhanced the LIBGiST library to collect statistics on index operations, file operations, and a detailed bufferpool snapshots. All prototype experiments used the TPC-H schema with scale factor 1, and the workload consisted of random point queries and random insert and update queries (conceptually similar to the workload presented in Section 6.2). We focus on the I/O profile of index traversals and updates as encountered by queries, updates, and insertions. All results are averaged over 10^5 random queries/updates/insertions.

Comparison of Average Execution Time. We first isolate the query, update, and insert execution times for a single index with a *small bufferpool size*, enough pages to pin the root of the index and all pages in one path from the root to a leaf. Figure 9 shows that the insert and query times are virtually the same due to the negligible overhead introduced by the additional SSD read and write I/O operations. The “Update” column for the Indirection method in



(a) One Index



(b) Sixteen Indexes

Figure 11: Varying the bufferpool size vs. update time

Figure 9 reflects a traversal of the index to locate the prior version of the record, plus an update of the SSD-resident LID-to-RID mapping. The base method is more expensive because it has to perform additional writes to the magnetic disk.

Multiple Indexes. The update time shown in Figure 9 does not capture the true benefit of the Indirection method when there are multiple indexes. In such a case, the Indirection method needs to traverse only one index, and update one LID-to-RID mapping.¹³ In contrast, the base method needs to traverse and update HDD-resident pages for every index. Figure 10 shows that the performance improvement approaches 20X as the number of indexes increases from 1 to 16.

Varying the Bufferpool Size. We consider five categories of bufferpool sizes large enough to hold: (1) only few pages (small), (2) all index non-leaf pages, (3) the entire index, (4) the entire index and all data pages, or (5) everything including the index, data, and *LtoR* table. These sizes reflect possible use cases where the system has memory ranging from very small to very large. We explored an update workload under these five settings when having one (Figure 11(a)) or sixteen (Figure 11(b)) indexes on unaffected attributes. Note the logarithmic vertical scale.

In Figure 11(a), only when both the index and data pages are memory-resident, but the *LtoR* table is not, does the Indirection method perform poorly compared to the Base approach. This is not surprising, because even a fast SSD I/O is slower than access

¹³If an indexed attribute is updated, then extra I/O is needed for that index to relocate the index entry. Indexes that are neither traversed nor updated in the Indirection method are said to be “unaffected.”

TPC-H: Lineitem table; Scale Factor: 1
 Non-unique Index on lineitem.supkey; Average Number of Duplicates: 36
 Index Clustering Ratio = 3%; Bufferpool Size: Small

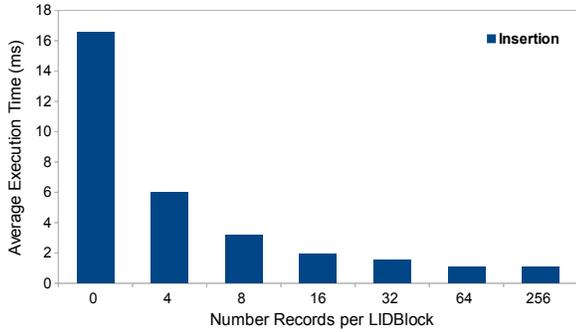


Figure 12: Varying the LIDBlock size vs. insertion time

to RAM. When an update operation results in some HDD I/O (either due to updating leaf pages or data pages), then Indirection is superior by up to 2X. When everything including the LID-to-RID mapping table is memory-resident, then Indirection continues to be superior because it does not touch most of the indexes.

Again, the true benefit of the Indirection technique surfaces as we increase the number of indexes. For example, when scaling the number of indexes to sixteen in Figure 11(b), Indirection typically wins by an order of magnitude. These experiments demonstrate that the Indirection technique provides significant benefits even if large portions of the database are memory-resident.

Varying the LIDBlock Size. So far our focus has been on improving index maintenance involving updates. We now demonstrate the power of our LIDBlock approach for insertions. In Figure 12, we vary the capacity of LIDBlock from no LIDs to 256 LIDs for the non-unique index defined on the *supkey* attribute of the *lineitem* table. When increasing the LIDBlock size to around 32 LIDs, we observed that the insertion cost is significantly reduced by up to 15.2X. This improvement is due to the amortization of a random leaf page update over many insertions, e.g., a LIDBlock size of 32 results in batching and flushing to the disk once every 32 insertions on average.

We can demonstrate a similar benefit with a non-unique index having many more duplicate entries, such as an index on the *quantity* attribute of the *lineitem* table, having 50 distinct values and 0.2M records per value. It is beneficial to allow larger LIDBlock sizes as shown in Figure 13, in which the insertion execution time is reduced by up to 4.9X. Unlike the previous case, the main reason for the speedup is not simple amortization of insertions; since there are so few key values the tree structure of the index is shallow and its traversal is already cheap due to caching. Instead, the speedup is due to having batches large enough to be resolved with one or two I/Os to update the LID lists for a key.

6.4 DB2 Operational Data Store Experiments

In this section, we study the effects of adding indexes in the context of an operational data store, in which small-scale analytical query processing is competing with transactional throughput. Our query workload is based on prior work [17] that modifies TPC-H queries so that they each touch less data. For our index workload, we rely on the DB2 Index Advisor recommendation given our query workloads defined over the entire TPC-H schema.

We first consider only the primary key indexes of the TPC-H relations. Subsequently, we add the remaining indexes recommended by DB2 Advisor one at a time, starting from the most to least beneficial index. After each round of index addition, we re-run

TPC-H: Lineitem table; Scale Factor: 1
 Non-unique Index on lineitem.quantity; Average Number of Duplicates: 0.2M
 Index Clustering Ratio = 25%; Bufferpool Size: Small

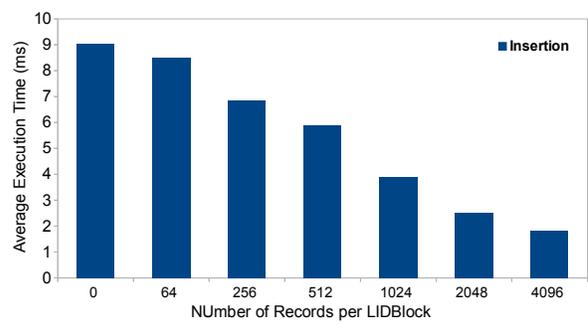


Figure 13: Varying the LIDBlock size vs. insertion time

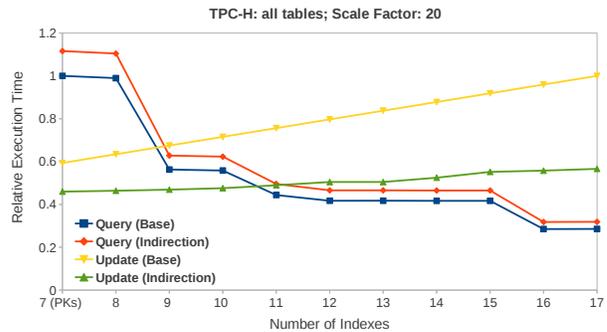


Figure 14: Effects of adding indexes on query/update time

our query workload. Likewise, after adding each index, we compute the update cost for a uniformly generated update workload, in which each non-primary-key attribute has an equal chance of being updated. The update cost is a normalized average execution time of updating indexes on DB2 and our on LIBGiST^{mv} with bufferpool size set to either small and non-leaf pages. The results are summarized in Figures 14 and 15.

Our first observation is that analytical query time is substantially reduced (by a factor of 3) as we add more indexes recommended by DB2 Advisor. More importantly, we observe that the additional indexes are more “affordable” for updates because our Indirection technique reduces the index maintenance overhead. In the base configuration, the index maintenance overhead increases linearly as more indexes are added, reducing transaction throughput. Our Indirection technique reduces the update cost by more than 40%.

Figure 15 shows a two dimensional plot of relative query time q versus relative update time u . On both axes, smaller is better. Each index configuration determines a point (u, q) on this graph, and one can choose a suitable configuration to achieve a desired query/update trade-off. In Figure 15, the Indirection technique dominates¹⁴ the Base method. This is a key result: The Indirection technique makes indexes more affordable, leading to lower query *and* update times.

To understand the importance of Figure 15 consider the following scenarios.

1. A DBA has an update-time budget of 0.6 units, and within that budget wants to optimize query processing time. According to Figure 15, the Indirection technique can achieve query performance of about 0.32 units under such conditions, while the Base method can achieve only 1.0 units, three times worse.

¹⁴Except when a very small weight is assigned to update time.

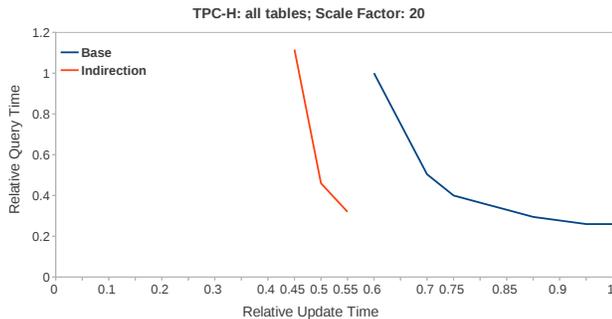


Figure 15: Effects of adding indexes on query/update time

2. A DBA has a query-time budget of 0.5 units, and within that budget wants to optimize update-time. The Indirection technique can achieve update performance of about 0.5 units under such conditions, while the Base method can achieve only 0.7 units, 40% worse.
3. A DBA wants to minimize the sum of update-time and query-time. The Indirection method can achieve a sum of about 0.87 at (0.55,0.32), whereas the best configuration in the Base method is 1.15 at (0.75,0.4), 32% worse.

7. RELATED WORK

There has been extensive work on storing high value data on SSDs. Some of these studies target the problem of data placement in relational databases to take better advantage of the SSD characteristics. In [17] database objects are either placed on HDDs or SSDs based on workload characteristics. As opposed to using SSDs and HDDs at the same level in the storage hierarchy, SSDs are also used as a second layer cache between main memory and HDDs [37, 15, 18, 23].

The use of an SSD cache with a write-back policy would not solve our problem as effectively as the Indirection technique. One would need space to hold entire indexes if one wants to avoid HDD index I/O for random updates and insertions. In contrast, our method avoids HDD I/O with a much smaller SSD footprint. The Indirection technique and SSD caching of HDD pages are complementary, and can be used in tandem. In fact, using Indirection improves the cache behaviour by avoiding reading/writing unnecessary pages (i.e., it avoids polluting the cache).

Adding a level of indirection is a commonly used programming technique relevant to a variety of systems problems [10]. The kind of indirection we propose in this paper is used in log-structured file systems [48, 33] and database indexing [39] but only at page granularity. We use indirection at record granularity. Furthermore, in [39], a latch-free B-Tree structure, called Bw-Tree, is proposed that virtualizes the concept of the page through page-level indirection (unlike our record-level indirection). But, more importantly, our Indirection technique is index-agnostic (not tuned for any particular index such as B-Tree) and addresses the general problem of index maintenance in multiversion databases, that is, to avoid updating indexes on unaffected attributes. Therefore, our approach could reduce the index maintenance of Bw-Tree as well. Another fundamental difference between our philosophy and Bw-Tree is that we consider de-coupling of the data from the index (in order to simplify the support of secondary indexes on the data), while Bw-Tree assumes a tight coupling of the index and data (i.e., storing the data as part of the index), essentially Bw-Tree is intended as an atomic record store similar to key-value stores [39].

The log-structured storage is also exploited in database management systems [55, 45, 41, 51], but no indirection layer is used; thus, the common merging and compaction operations in log-structured storage result in expensive rebuilding of all indexes. By employing our Indirection technique this index rebuilding can be avoided.

Page-level indirection tables are also used to improve the lifespan of SSDs. In [21], a system called CAFTL is proposed to eliminate duplicate writes on SSDs. By keeping a mapping table of blocks the redundant block writes on SSDs are eliminated.

In [56], Wu et al. proposes a software layer called BFTL to store B-tree indexes on flash devices efficiently. IUD operations cause significant byte-wise operations for B-tree reorganization. The proposed layer reduces the performance overhead of these updates on the flash device by aggregating the updates on a particular page.

In [24], Dou et al. propose specialized index structures and algorithms that support querying of historical data in flash-equipped sensor devices. Since the sensor devices have limited memory capacity (SRAM) and the underlying flash devices have certain limitations, there are challenges in maintaining and querying indexes.

The deferral of index I/Os is used in several recent papers on improving index performance [14, 44, 8]. In those papers, changes are accumulated at intermediate nodes, and propagated to children in batches. Searches need to examine buffers for keys that match. This line of work is complementary but similar to our LIDBlock method in that both techniques buffer insertions to amortize physical I/O.

SSDs are used to support online updates in data warehouses [11]. Incoming updates are first cached in SSDs and later merged with the older records on HDDs to answer queries. In [11], data records are accessed primarily through table scans rather than indexes.

Many specialized indexes for versioned and temporal data have been proposed. A comprehensive survey of temporal indexing methods is provided in [49]. Tree based indexes on temporal data include the multiversion B-tree [12], Interval B-tree [9], Interval B+tree [16], TP-Index [52], Append-only Tree [30] and Monotonic B+tree [26]. Efficiently indexing data with branched evolution is discussed by Jouni et al. [36], who build efficient structures to run queries on both current and historical data.

Specialized transaction time database systems such as Immortal DB [42, 43] provide high performance for temporal applications. Lomet et al. [43] describe how a temporal indexing technique, the TSB-tree, is integrated into SQL Server. The paper also describes an efficient page layout for multiversion databases.

8. CONCLUSIONS AND FUTURE WORK

The multiversion temporal database market is growing [22]. A temporal database simplifies application development and deployment by pushing the management of temporal logic into database engines. By adopting temporal technologies, the development cost can be reduced by a factor of 10 [22]. This success has led major database vendors (including Oracle [46], IBM [34], and TeraData [53]) to provide support for multiversion temporal data.

We tackle a key challenge of multiversion databases: providing good update performance and good query performance in a single system. Transaction throughput and analytical query processing often have conflicting requirements due to the high index maintenance cost for transactions. Our efficient index maintenance using Indirection makes indexes more “affordable,” substantially improving the available configuration choices. Our evaluation demonstrates a query cost reduction by a factor of 3 without an increase in update cost. The batching of insertions using our LIDBlock technique can save up to 90% of the insertion time.

There remain several open issues that we plan to address in future work. One primary concern in multiversion databases is space. If every version of a record is copied in full, the space needed to store the database will grow dramatically. There is obvious potential to compress records, taking advantage of commonality between old and new record versions. Even so, the trade-offs are not straightforward. One does not want to burden transactions or queries on current data with extensive compression/decompression tasks, so a background process that leaves the current record uncompressed but incrementally compresses the historical records may be a good idea. Keeping many versions of a record on a common page may help compression, but would be incompatible with other clustering orders.

9. REFERENCES

- [1] BioPostgres: Data management for computational biology. <http://www.biopostgres.org/>.
- [2] IBM DB2 Database for Linux, UNIX, and Windows. www.ibm.com/software/data/db2/linux-unix-windows/.
- [3] IBM DB2 with BLU Acceleration. www.ibm.com/software/data/db2/linux-unix-windows/db2-blu-acceleration/.
- [4] OpenFTS: Open source full text search engine. <http://openfts.sourceforge.net/>.
- [5] PostGIS: Geographic information systems. <http://postgis.refractions.net/>.
- [6] PostgreSQL: Open source object-relational database system. <http://www.postgresql.org/>.
- [7] YAGO2: High-quality knowledge base. <http://www.mpi-inf.mpg.de/yago-naga/yago/>.
- [8] D. Agrawal, D. Ganesan, R. K. Sitaraman, Y. Diao, and S. Singh. Lazy-adaptive tree: An optimized index structure for flash devices. *PVLDB*, 2(1):361–372, 2009.
- [9] C.-H. Ang and K.-P. Tan. The interval B-tree. *Inf. Process. Lett.*, 53(2):85–89, Jan. 1995.
- [10] R. Arpaci-Dusseau and A. Arpaci-Dusseau. *Operating Systems: Three Easy Pieces*. Arpaci-Dusseau Books, 0.5 edition, 2012.
- [11] M. Athanassoulis, S. Chen, A. Ailamaki, P. B. Gibbons, and R. Stoica. MaSM: efficient online updates in data warehouses. In *SIGMOD Conference*, pages 865–876, 2011.
- [12] B. Becker, S. Gschwind, T. Ohler, B. Seeger, and P. Widmayer. An asymptotically optimal multiversion B-Tree. *VLDB J.*, 5(4):264–275, 1996.
- [13] B. Bhattacharjee, L. Lim, T. Malkemus, G. Mihaila, K. Ross, S. Lau, C. McArthur, Z. Toth, and R. Sherkat. Efficient index compression in DB2 LUW. *Proc. VLDB Endow.*, 2(2):1462–1473, Aug. 2009.
- [14] B. Bhattacharjee, T. Malkemus, S. Lau, S. Mckeough, J.-A. Kirton, R. V. Boeschoten, and J. Kennedy. Efficient bulk deletes for multi dimensionally clustered tables in DB2. In *VLDB*, pages 1197–1206, 2007.
- [15] B. Bhattacharjee, K. A. Ross, C. A. Lang, G. A. Mihaila, and M. Banikazemi. Enhancing recovery using an SSD buffer pool extension. In *DaMoN*, pages 10–16, 2011.
- [16] T. Bozkaya and M. Özsoyoğlu. Indexing valid time intervals. *Lecture Notes in Computer Science*, 1460:541–550, 1998.
- [17] M. Canim, B. Bhattacharjee, G. A. Mihaila, C. A. Lang, and K. A. Ross. An object placement advisor for DB2 using solid state storage. *PVLDB*, 2(2):1318–1329, 2009.
- [18] M. Canim, G. A. Mihaila, B. Bhattacharjee, K. A. Ross, and C. A. Lang. SSD bufferpool extensions for database systems. *PVLDB*, 3(2):1435–1446, 2010.
- [19] S. Chaudhuri and V. Narasayya. Automating statistics management for query optimizers. *IEEE Trans. on Knowl. and Data Eng.*, 13(1):7–20, Jan. 2001.
- [20] S. Chaudhuri and V. R. Narasayya. An efficient cost-driven index selection tool for microsoft SQL server. In *Proceedings of the 23rd International Conference on Very Large Data Bases*, VLDB '97, pages 146–155, San Francisco, CA, USA, 1997. Morgan Kaufmann Publishers Inc.
- [21] F. Chen, T. Luo, and X. Zhang. CAFTL: A content-aware flash translation layer enhancing the lifespan of flash memory based solid state drives. In *FAST*, pages 77–90, 2011.
- [22] S. Chen. Time travel query or bi-temporal. In *DB2 for z/OS Technical Forum*, 2010.
- [23] J. Do, D. Zhang, J. M. Patel, D. J. DeWitt, J. F. Naughton, and A. Halverson. Turbocharging DBMS buffer pool using SSDs. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, SIGMOD '11, pages 1113–1124, New York, NY, USA, 2011. ACM.
- [24] A. J. Dou, S. Lin, and V. Kalogeraki. Real-time querying of historical data in flash-equipped sensor devices. In *IEEE Real-Time Systems Symposium*, pages 335–344, 2008.
- [25] G. Drossel. Methodologies for calculating SSD usable life. In *Storage Developer Conference*, 2009.
- [26] R. Elmasri, G. T. J. Wu, and V. Kouramajian. The time index and the monotonic B+–tree. In *Temporal Databases*, pages 433–456, 1993.
- [27] Fusion-io breaks one billion IOPS barrier. <http://www.fusionio.com/press-releases/fusion-io-breaks-one-billion-iops-barrier/>.
- [28] H. Garcia-Molina, J. D. Ullman, and J. Widom. *Database Systems: The Complete Book*. Prentice Hall Press, Upper Saddle River, NJ, USA, 2 edition, 2008.
- [29] The GiST indexing project. <http://gist.cs.berkeley.edu/>.
- [30] H. Gunadhi and A. Segev. Efficient indexing methods for temporal relations. *IEEE Trans. on Knowledge and Data Eng.*, 5(3):496, June 1993.
- [31] J. M. Hellerstein, J. F. Naughton, and A. Pfeffer. Generalized search trees for database systems. In *Proceedings of the 21th International Conference on Very Large Data Bases*, VLDB '95, pages 562–573, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc.
- [32] F. D. Hinshaw, C. S. Harris, and S. K. Sarin. Controlling visibility in multi-version database systems, 2007. US 7305386 Patent, Netezza Corporation.
- [33] D. Hitz, J. Lau, and M. Malcolm. File system design for an NFS file server appliance. In *Proceedings of the USENIX Winter 1994 Technical Conference on USENIX Winter 1994 Technical Conference*, WTEC'94, pages 19–19, Berkeley, CA, USA, 1994. USENIX Association.
- [34] DB2 10 for z/OS. ftp://public.dhe.ibm.com/software/systemz/whitepapers/DB210_for_zOS_Upgrade_ebook.pdf.
- [35] W. H. Inmon. *Building the Operational Data Store*. John Wiley & Sons, Inc., New York, NY, USA, 2nd edition, 1999.
- [36] K. Jouini and G. Jomier. Indexing multiversion databases. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, CIKM '07, pages 915–918, New York, NY, USA, 2007. ACM.
- [37] W.-H. Kang, S.-W. Lee, and B. Moon. Flash-based extended cache for higher throughput and faster recovery. *PVLDB*, 5(11):1615–1626, 2012.
- [38] P.-A. Larson, S. Blanas, C. Diaconu, C. Freedman, J. M. Patel, and M. Zwilling. High-performance concurrency control mechanisms for main-memory databases. *Proc. VLDB Endow.*, 5(4):298–309, Dec. 2011.
- [39] J. J. Levandoski, D. B. Lomet, and S. Sengupta. The Bw-Tree: A B-tree for new hardware platforms. In *Proceedings of the 2013 IEEE 29th International Conference on Data Engineering*, ICDE '13, Washington, DC, USA, 2013. IEEE Computer Society.
- [40] A. Leventhal. Flash storage memory. *Commun. ACM*, 51(7):47–51, July 2008.
- [41] Y. Li, B. He, Q. Luo, and K. Yi. Tree indexing on flash disks. In *Proceedings of the 2009 IEEE International Conference on Data Engineering*, ICDE '09, pages 1303–1306, Washington, DC, USA, 2009. IEEE Computer Society.
- [42] D. Lomet, R. Barga, M. F. Mokbel, G. Shegalov, R. Wang, and Y. Zhu. Immortal DB: transaction time support for SQL server. In *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, SIGMOD '05, pages 939–941, New York, NY, USA, 2005. ACM.
- [43] D. Lomet, M. Hong, R. Nehme, and R. Zhang. Transaction time indexing with version compression. *Proc. VLDB Endow.*, 1(1):870–881, Aug. 2008.
- [44] E. Omiecinski, W. Liu, and I. F. Akyildiz. Analysis of a deferred and incremental update strategy for secondary indexes. *Inf. Syst.*, 16(3):345–356, 1991.
- [45] P. E. O'Neil, E. Cheng, D. Gawlick, and E. J. O'Neil. The log-structured merge-tree (LSM-Tree). *Acta Inf.*, 33(4):351–385, 1996.
- [46] Oracle database 11g workspace manager overview. <http://www.oracle.com/technetwork/database/twp-appdev-workspace-manager-11g-128289.pdf>.
- [47] Oracle total recall/flashback data archive. <http://www.oracle.com/technetwork/issue-archive/2008/08-jul/flashback-data-archive-whitepaper-129145.pdf>.
- [48] M. Rosenblum and J. K. Ousterhout. The design and implementation of a log-structured file system. *ACM Trans. Comput. Syst.*, 10(1):26–52, Feb. 1992.
- [49] Salzberg and Tsotras. Comparison of access methods for time-evolving data. *CSURV: Computing Surveys*, 31, 1999.
- [50] C. M. Saracco, M. Nicola, and L. Gandhi. A matter of time: Temporal data management in DB2 for z/OS, 2010.
- [51] R. Sears and R. Ramakrishnan. bLSM: a general purpose log structured merge tree. In *SIGMOD Conference*, pages 217–228, 2012.
- [52] H. Shen, B. Chin, and O. H. Lu. The TP-Index: A dynamic and efficient indexing mechanism for temporal databases. In *In Proceedings of the Tenth International Conference on Data Engineering*, pages 274–281. IEEE, 1994.
- [53] R. T. Snodgrass. A case study of temporal data, 2010. Teradata Corporation.
- [54] TPC-H, decision support benchmark. <http://www.tpc.org/tpch/>.
- [55] H. T. Vo, S. Wang, D. Agrawal, G. Chen, and B. C. Ooi. LogBase: A scalable log-structured database system in the cloud. *PVLDB*, 5(10):1004–1015, 2012.
- [56] C.-H. Wu, T.-W. Kuo, and L.-P. Chang. An efficient B-tree layer implementation for flash-memory storage systems. *ACM Trans. Embedded Comput. Syst.*, 6(3), 2007.