

Efficient Data Ingestion and Query Processing for LSM-Based Storage Systems

Chen Luo
University of California, Irvine
clu08@uci.edu

Michael J. Carey
University of California, Irvine
mjcarey@ics.uci.edu

ABSTRACT

In recent years, the Log Structured Merge (LSM) tree has been widely adopted by NoSQL and NewSQL systems for its superior write performance. Despite its popularity, however, most existing work has focused on LSM-based key-value stores with only a single LSM-tree; auxiliary structures, which are critical for supporting ad-hoc queries, have received much less attention. In this paper, we focus on efficient data ingestion and query processing for general-purpose LSM-based storage systems. We first propose and evaluate a series of optimizations for efficient batched point lookups, significantly improving the range of applicability of LSM-based secondary indexes. We then present several new and efficient maintenance strategies for LSM-based storage systems. Finally, we have implemented and experimentally evaluated the proposed techniques in the context of the Apache AsterixDB system, and we present the results here.

PVLDB Reference Format:

Chen Luo, Michael J. Carey. Efficient Data Ingestion and Query Processing for LSM-Based Storage Systems. *PVLDB*, 12(5): 531-543, 2019.

DOI: <https://doi.org/10.14778/3303753.3303759>

1. INTRODUCTION

A wide range of applications, such as risk management, online recommendations, and location-based advertising, demand the capability of performing real-time analytics on high-speed, continuously generated data coming from sources such as social networks, mobile devices and IoT applications. As a result, modern Big Data systems need to efficiently support both fast data ingestion and real-time queries.

The Log-Structured Merge (LSM) tree [24] is a promising structure to support write-intensive workloads. It has been widely adopted by NoSQL and NewSQL systems [2, 3, 4, 7, 10, 14] for its superior write performance. Instead of updating data in-place, which can lead to expensive random I/Os, LSM writes, including inserts, deletes and updates, are first accumulated in memory and then subsequently flushed

to disk and later merged using sequential I/Os. A number of improvements have been proposed to optimize various aspects of the original LSM proposal [8, 16, 17, 28, 34].

Auxiliary structures, such as secondary indexes, are critical to enable the efficient processing of ad-hoc queries. Two types of LSM-based auxiliary structures have been used in practice to facilitate query processing, namely secondary indexes and filters. A secondary index is an LSM-tree that maps secondary key values to their corresponding primary keys. A filter, such as a Bloom filter [13] or a range filter [11] on secondary keys, is directly built into LSM-trees to enable data skipping for faster scans. While needed for queries, maintaining these structures during data ingestion comes with extra cost. Especially in the case of updates, both types of structures require accessing old records so that they can be properly maintained. Existing LSM-based systems, such as AsterixDB [1, 9], MyRocks [5], and Phoenix [6], employ an eager strategy to maintain auxiliary structures by prefacing each incoming write with a point lookup. This strategy is straightforward to implement and optimizes for query performance since auxiliary structures are always up-to-date. However, it leads to significant overhead during ingestion because of the point lookups.

An outage with respect to the general-purpose use of LSM indexing is that no particularly efficient point lookup algorithms have been proposed for the efficient fetching of the records identified by a secondary index search. While sorting fetch lists based on primary keys is a well-known optimization [19], performing the subsequent point lookups independently still incurs high overhead. This limits the range of applicability of LSM-based secondary indexes, requiring the query optimizer to maintain accurate statistics to make correct decisions for processing ad-hoc queries efficiently.

This paper focuses on efficient data ingestion and query processing techniques for general-purpose LSM-based storage systems. The first contribution of this paper is to propose and evaluate a series of optimizations for efficient index-to-index navigation for LSM-based indexes. We show how to leverage the internal structure of LSM to efficiently process a large number of point lookups. We further conduct a detailed empirical analysis to evaluate the effectiveness of these optimizations. The experimental results show that the proposed optimizations greatly improve the range of applicability of LSM-based secondary indexes. Even with relatively large selectivities, such as 10 - 20%, LSM-based secondary indexes can still provide performance better than or (worst case) comparable to that of a full scan.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 12, No. 5

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/3303753.3303759>

The second contribution of this paper is to study alternative maintenance strategies for LSM-based auxiliary structures. The key insight here is the applicability of a *primary key index*, which only stores primary keys, for use in index maintenance. For secondary indexes, we present a *validation* strategy that cleans up obsolete entries lazily using a primary key index. For filters, we introduce a *mutable-bitmap* strategy that allows deleted keys to be directly reflected for immutable data through mutable bitmaps by accessing primary keys instead of full records. Since primary keys are much smaller than full records, we show that the exploitation of a primary key index can greatly reduce required I/Os and significantly increase the overall ingestion throughput.

Finally, we have implemented all of the proposed techniques inside Apache AsterixDB [1, 9], an open-source LSM-based Big Data Management System. We have conducted extensive experiments to evaluate their impacts on ingestion performance and query performance.

The remainder of the paper is organized as follows: Section 2 provides background information and surveys the related work. Section 3 presents a general architecture for LSM-based storage systems and introduces various optimizations for efficient point lookups. Sections 4 and 5 describe in detail the proposed Validation strategy and the Mutable-bitmap strategy. Section 6 experimentally evaluates the proposed techniques. Finally, Section 7 concludes the paper.

2. BACKGROUND

2.1 Log-Structured Merge Trees

The LSM-tree [24] is a persistent index structure optimized for write-intensive workloads. In LSM, writes are first buffered into a memory component and then flushed to disk using sequential I/Os when memory is full. Each flush operation forms a new disk component. Once flushed, LSM disk components are immutable. Modifications (inserts, updates and deletes) are therefore handled by inserting new entries into memory. An insert or update simply adds a new entry with the same key, while a delete adds an “anti-matter” entry [10] indicating that a key has been deleted.

A query over LSM data has to reconcile the entries with identical keys from multiple components, as entries from the newer components override those from older components. As disk components accumulate, query performance tends to degrade since more components must be examined. To counter this, disk components are gradually merged according to a pre-defined *merge policy*. In general, two types of merge policies are used in practice [16, 17], both of which organize components into “levels”. The leveling merge policy maintains one component per level, and a component in a higher level will be exponentially larger than that of the next lower one. In contrast, the tiering merge policy maintains multiple components per level; these components are formed when a series of lower-level components are merged together using a merge operation.

2.2 Apache AsterixDB

Apache AsterixDB [9] is a parallel, semistructured Big Data Management System (BDMS) that aims to support ingesting, storing, indexing, querying and analyzing massive amounts of data efficiently. Here we briefly discuss storage management in AsterixDB [10], which relates to this paper.

The records of a dataset in AsterixDB are hash-partitioned based on their primary keys across multiple nodes. Each partition of a dataset uses a primary LSM-based B^+ -tree index to store the records. Secondary indexes, including LSM-based B^+ -trees, R-trees, and inverted indexes, are local to the primary index partition, as in most shared-nothing parallel databases. The memory components of all the indexes of a partition of a dataset share a memory budget, and thus they are always flushed together. A primary index entry is stored as the primary key plus the record, while a secondary index entry is composed of the secondary key and the associated primary key. Secondary index lookups are routed to all dataset partitions to fetch the corresponding primary keys. The returned primary keys are then sorted locally before retrieving the records in the local partitions in order to retrieve them in an efficient manner.

AsterixDB supports record-level transactions across multiple LSM indexes of a dataset to ensure that all secondary indexes are consistent with the primary index. It utilizes a no-steal/no-force buffer management policy with write-ahead-logging (WAL) to ensure durability and atomicity. For recovery and rollback, index-level logical logging and component shadowing are employed. Rollback for in-memory component changes is implemented by applying the inverse operations of log records in the reverse order. To perform crash recovery, committed transactions beyond the maximum LSN of disk components are replayed by examining log records. No undo is performed during recovery since the no-steal policy guarantees that disk components can only contain committed transactions.

2.3 Related Work

A number of improvements of the original LSM-tree index structure [24] have been proposed recently. bLSM [28] presented a spring-and-gear merge scheduler to reduce periodic write stalls. cLSM [18] is optimized for multi-core machines using non-blocking concurrency control mechanisms. Monkey [16] optimized the memory allocation of Bloom filters for LSM-trees. Dostoevsky [17] presented a lazy-leveling merge policy that can make better performance trade-offs. WisckKey [21] separated values from keys to reduce write amplification. All of these efforts have been focused on improving a single LSM index. Our work here is orthogonal to these optimizations of LSM-trees since we focus on LSM-based auxiliary structures.

DELI [31] presented a lazy secondary index maintenance strategy by repairing secondary indexes during the merge of primary index components. In addition to its eager strategy, AsterixDB [10] supports a deleted-key B^+ -tree strategy that attaches a B^+ -tree to each secondary index component that records the deleted keys in this component. Since DELI and the deleted-key B^+ -tree strategy are closely related to our work, we will further discuss them in detail in Section 4.1. Qadar et al. [26] conducted an experimental study of LSM-based secondary indexes. However, their study did not consider cleaning up secondary indexes in the case of updates.

Alsubaiee et al. [11] added the option of a range filter on LSM-based primary and secondary indexes for the efficient processing of time-correlated queries. Jia [20] exploited LSM range filters to accelerate primary key lookups for append-only and time-correlated workloads. However, the eager strategy for maintaining filters incurs high lookup cost during ingestion and reduces their pruning capabilities in the

presence of updates. Several commercial database systems have similarly supported range filter-like structures, such as zone maps in Oracle [35] and synopses in DB2 BLU [27], to enable data skipping during scans. However, these systems are not based on LSM-trees.

Several related write-optimized indexes [29, 32, 33] have been proposed to efficiently index append-only observational streams and time-series data. Though they provide high ingestion performance, updates and deletes are not supported. Our work targets general workloads that include updates and deletes as well as appends.

3. LSM STORAGE ARCHITECTURE

In this section, we present a general LSM-based storage architecture that will be the foundation for the rest of this paper. We also present the Eager strategy used in existing systems as well as a series of optimizations for index-to-index navigation given LSM-based indexes.

The storage architecture is depicted in Figure 1. Each dataset has a primary index, a primary key index, and a set of secondary indexes, which are all based on LSM-trees. All indexes in a dataset share a common memory budget for memory components, so they are always flushed together (as in AsterixDB). Each component in the figure is labeled with its *component ID*, which is represented as a pair of timestamps (minTS - maxTS). The component ID is simply maintained as the minimum and maximum timestamps of the index entries stored in the component, where the timestamp is generated using the local wall clock time when a record is ingested. Through component IDs, one can infer the recency ordering among components of different indexes, which can be useful for secondary index maintenance operations (as discussed later). For example, component IDs indicate that component 1-15 of Secondary Index 1 is older than component 16-18 of the primary index and that it overlaps component 1-10 of the primary index.

The primary index stores the records indexed by their primary keys. To reduce point lookups during data ingestion, as we will see later, we further build a primary key index that stores primary keys only. Both of these indexes internally use a B⁺-tree to organize the data within each component. Each primary or primary key disk component also has a Bloom filter [13] on the stored primary keys to speed up point lookups. A point lookup query can first check the Bloom filter of a disk component and search the actual B⁺-tree only when the Bloom filter reports that the key may exist. Secondary indexes use a composition of the secondary key and the primary key as their index key in order to efficiently handle duplicate secondary keys. A secondary index

query can first search the secondary index to return a list of matching primary keys and then perform point lookups to fetch records from the primary index.

In general, the primary index may have a set of filters for efficient pruning during scans¹. Without loss of generality, we assume that each primary index component may have a range filter that stores the minimum and maximum values of the component’s secondary filter key (denoted as [Ti, Tj] in Figure 1). During a scan, a component can be pruned if its filter is disjoint with the search condition of a query.

3.1 Data Ingestion with the Eager Strategy

During data ingestion, all storage structures must be properly maintained. Here we briefly review the Eager strategy commonly used by existing LSM-based systems such as AsterixDB [1, 9], MyRocks [5], and Phoenix [6].

To *insert* a record, its key uniqueness is first checked by performing a point lookup. As an optimization, the primary key index can be searched instead for efficiency. If the given primary key already exists, the record is ignored; otherwise, the record is recorded in the memory components of all of the dataset’s indexes. Any filters of the memory components must be maintained based on the new record as well.

To *delete* a record given its key, a point lookup is first performed to fetch the record. If the record does not exist, the key is simply ignored. Otherwise, anti-matter entries must be inserted into all of the dataset’s LSM indexes to delete the record. Any filters of the memory components must also be maintained based on the deleted record; otherwise, a query could erroneously prune a memory component and thus access deleted records.

To *upsert* a record, a point lookup is first performed to locate the old record with the same key. If the old record exists, anti-matter entries are generated to delete the old record from all of the dataset’s secondary indexes. The new record is then inserted into the memory components of all of the dataset’s LSM indexes. As an optimization, if the value of some secondary key did not change, the corresponding secondary index can be simply skipped for maintenance. Any filters of the memory components must be maintained based on both the old record (if it exists) and the new record.

As a running example, consider the initial LSM indexes depicted in Figure 2 for a UserLocation dataset with three attributes: UserID (the primary key), Location, and Time. The Location attribute stores the last known location of the user (in terms of states) and the Time attribute stores the time of the last known location (in terms of years). We have a secondary index on Location and a range filter on Time. Figure 3 shows the resulting LSM indexes after upserting a new record (101, NY, 2018) with an existing key 101. A point lookup is first performed to locate the old record (101, CA, 2015). In addition to adding the new record to all memory components, an anti-matter entry (-CA, 101) is added to the secondary index to eliminate the obsolete entry (which is dashed in the figure). Also, notice that the memory component’s range filter is maintained based on both 2015 and 2018 so that future queries will properly see that the old record with UserID 101 has been deleted.

We further use Figure 3’s example to illustrate how secondary indexes and filters are used in query processing. First

¹As suggested by [11], secondary indexes could have filters as well. However, in this paper we only focus on the use of filters on the primary index to support efficient scans.




Figure 1: Storage Architecture




Figure 2: Running Example

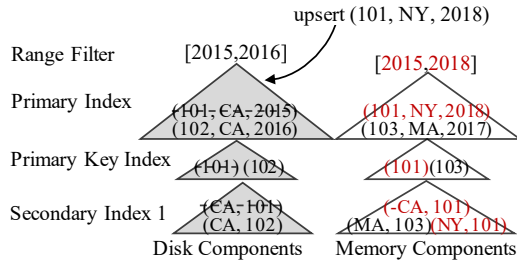


Figure 3: Upsert Example with Eager Strategy

consider a query Q1 that finds all user records with Location CA. Q1 first searches the secondary index to return a list of UserIDs. In this example only UserID 102 is returned since (CA, 101) is deleted by the anti-matter entry (-CA, 101). The primary index is then searched to fetch the record using UserID 102, which returns (102, CA, 2016). Note that without the anti-matter entry, the obsolete UserID 101 would be erroneously searched by Q1 as well. Consider another query Q2 that finds all records with Time < 2017. Q2 scans the primary index by first collecting a set of candidate components whose range filters overlap with the search condition Time < 2017. In this example, both components will be scanned by Q2 and only one record (102, CA, 2016) is returned. However, suppose that when upserting record (101, NY, 2018), the range filter of the memory component had only been maintained based on the new value 2018. In this case, the memory component would have been pruned and (101, CA, 2015) would be erroneously returned as well.

3.2 Efficient Index-to-index Navigation

Navigating from secondary indexes to the primary index is a fundamental operation for query processing. Traditionally, primary keys are sorted to ensure that the pages of the primary index will be accessed sequentially [19]. Here we discuss further optimizations to improve point lookup performance for LSM-trees. Note that some of the optimizations below are not new. Our contribution is to evaluate their effectiveness and integrate them to improve the range of applicability of LSM-based secondary indexes.

Batched Point Lookup. Even though primary keys are sorted, when searching multiple LSM components, it is still possible that index pages will be fetched via random I/Os since the sorted keys can be scattered across different components. To avoid this, we propose here a *batched point lookup* algorithm that works as follows: Sorted primary keys are first divided into batches. For each batch, all of the LSM components are accessed one by one, from newest to oldest. Specifically, for each key in the current batch that has not been found yet, it is searched against a primary

component by first checking the Bloom filter and then the B⁺-tree. A given batch terminates either when all components have been searched or all keys have been found. The batch search algorithm ensures that components’ pages are accessed sequentially, avoiding random I/Os when fetching their leaf pages. However, a downside is that the returned data records will no longer still be ordered on primary keys. We will experimentally evaluate this trade-off in Section 6.

Stateful B⁺-tree Lookup. To reduce the in-memory B⁺-tree search overhead, a stateful search cursor that remembers the search history from root to leaf can be used. Instead of always traversing from the root, the cursor starts from the last leaf page to reduce the tree traversal cost. One can further use exponential search [12] instead of binary search to reduce the search cost within each page. To search a key, this algorithm starts from the last search position and uses exponentially increasing steps to locate the range which this key resides in. The key is then located by performing a binary search within this range.

Blocked Bloom Filter. Finally, to reduce the overhead of checking the components’ Bloom filters, a cache-friendly approach called blocked Bloom filter [25] can be used. The basic idea is to divide the bit space into fixed-length blocks whose size is the same as the CPU cache line size. The first hash function maps a key to a block, while the rest of the hash functions perform the usual bit tests but within this block. This ensures that each Bloom filter test will only lead to one cache miss, at the cost of requiring an extra bit per key to achieve the same false positive rate [25].

4. VALIDATION STRATEGY

In this section, we propose the Validation strategy for maintaining secondary indexes efficiently. We first present an overview followed by detailed discussions of this strategy.

4.1 Overview

In a primary LSM-tree, an update can blindly place a new entry (with the identical key) into memory to mark the old entry as obsolete. However, this mechanism does not work with secondary indexes, as the index key could change after an update. Similarly, one cannot efficiently determine whether a given primary key is still valid based on a secondary index alone, as entries in a secondary index are ordered based on secondary keys. Extra work must be performed to maintain secondary indexes during data ingestion.

The Eager strategy maintains secondary indexes by producing anti-matter entries for each old record, which incurs a large point lookup overhead during data ingestion. An alternative strategy is to only insert new entries into secondary indexes during data ingestion while cleanup obsolete entries lazily so that the expensive point lookups can be avoided. This design further ensures that secondary indexes can only return false positives (obsolete primary keys) but not false negatives, which simplifies query processing. Although the idea of lazy maintenance is straightforward, two challenges must be addressed: (1) how to support queries efficiently, including both non-index-only and index-only queries; (2) how to repair secondary indexes efficiently to cleanup obsolete entries while avoiding making cleanup a new bottleneck.

There have been several proposals for lazy secondary index maintenance on LSM-based storage systems. DELI [31] maintains secondary indexes lazily, while merging the primary index components, without introducing any additional

structures. If multiple records with the same primary key are encountered during a merge, DELI produces anti-matter entries for obsolete records to clean up secondary indexes. However, this design does not completely address the above two challenges. First, non-index-only queries cannot be supported efficiently, as queries must fetch records to validate the search results. Second, DELI lacks the flexibility to repair secondary indexes efficiently. For update-heavy workload, it is desirable to repair secondary indexes more frequently to improve query performance. However, in DELI this requires constantly merging or scanning all components of the primary index, incurring a high I/O cost.

To support index-only queries, extra structures on top of secondary indexes should be maintained. AsterixDB supports a deleted-key B⁺-tree strategy that attaches a B⁺-tree to each secondary index component that records the deleted keys in this component. For index-only queries, validation can be performed by searching these deleted-key B⁺-trees without accessing full records. However, since these B⁺-trees are duplicated for each secondary index, and no efficient repairing algorithms have been described, cleaning up secondary indexes incurs a high overhead during merges.

To address the aforementioned challenges, we choose to use the primary key index to maintain all secondary indexes efficiently. The primary key index can be used for validating index-only queries by storing an extra timestamp for each index entry in all secondary indexes as well as in the primary key index. This timestamp is generated using the node-local clock time when a record is ingested and is stored as an integer with 8 bytes. We further propose an efficient index repair algorithm based on the primary key index, greatly reducing the I/O cost by avoiding accessing full records.

4.2 Data Ingestion

Under the Validation strategy, an *insert* is handled exactly as in the Eager strategy except that timestamps are added to the primary key index entries and secondary index entries. To *delete* a record given its key, an anti-matter entry is simply inserted into both the primary index and the primary key index. To *upsert* a record, the new record is simply inserted into all of the dataset’s LSM indexes. (Notice that when deleting or upserting a record, the memory component of the primary index has to be searched to find the location for the entry being added. As an optimization, then, if the old record happens to reside in the memory component, it can be used to produce local anti-matter entries to clean up the secondary indexes without additional cost.)

Consider the running example that began in Figure 2. Figure 4 shows the resulting LSM indexes after upserting the record (101, NY, 2018) under the Validation strategy. The primary key index and the secondary index each now contain an extra timestamp field (denoted as “ts”). To upsert the new record, we add the new record to all memory components without any point lookups. As a result, the obsolete secondary index entry (CA, 101, ts1) still appears to be valid even though it points to a deleted record. The Validation strategy can be naturally extended to support filters: Since no pre-operation point lookup is performed, the filters of the memory components are only maintained based on new records. As in the example, then, the range filter is only maintained based on 2018. As a result, a query that accesses an older component has to access all newer components in order not to miss any newer overriding updates.

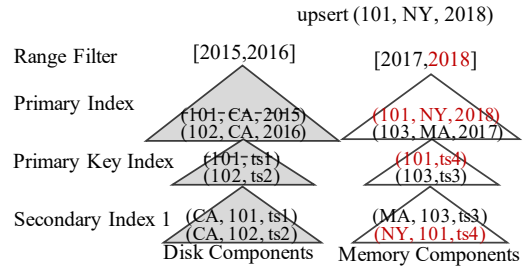


Figure 4: Upsert Example with Validation Strategy

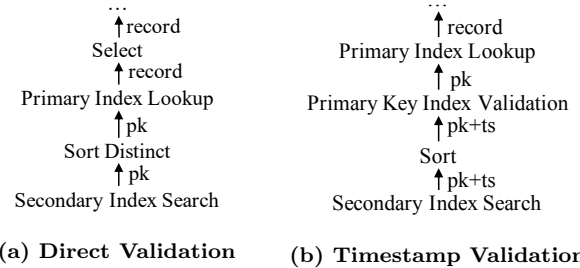


Figure 5: Query Validation Methods

4.3 Query Processing

In the Validation strategy, secondary indexes are not always guaranteed to be up-to-date and can thus return obsolete entries to queries. For correctness, queries have to perform an extra validation step to ensure that only valid keys are eventually accessed. Here we present two validation method variations suitable for different queries (Figure 5).

The Direct Validation method (Figure 5a) directly performs point lookups to fetch all of the candidate records and re-checks the search condition. A sort-distinct step is performed first to remove any duplicate primary keys. After checking the search condition, only the valid records are returned to the query. However, this method rules out the possibility of supporting index-only queries efficiently.

To address these drawbacks, the Timestamp Validation method (Figure 5b) uses the primary key index to perform validation. A secondary index search returns the primary keys plus their timestamps. Point lookups against the primary key index are then performed to validate the fetched primary keys. Specifically, a key is invalid if the same key exists in the primary key index but with a larger timestamp. The valid keys are then used to fetch records if necessary.

Consider the example in Figure 4, and a query that wants to find all records with Location CA. The secondary index search returns primary keys with their timestamps (101, ts1) and (102, ts2). Direct Validation performs point lookups to locate records (101, NY, 2018) and (102, CA, 2016). The first record will be filtered out because its Location is not CA anymore. Timestamp Validation would perform a point lookup against the primary key index to filter out UserID 101 since its timestamp ts1 is older than ts4.

4.4 Secondary Index Repair

Since secondary indexes are not cleaned up during ingestion under the Validation strategy, obsolete entries could accumulate and degrade query performance. To address




Figure 6: Repaired Timestamp Example

```

1: Create search cursor on merging components
2: position ← 0
3: while cursor.hasNext() do
4:   entry ← cursor.getNextEntry()
5:   add entry to new component
6:   add (pkey, ts, position) to sorter
7:   position ← position + 1
8: initialize bitmap with all 0s
9: sorter.sort()
10: for sorted entries (pkey, ts, position) do
11:   validate pkey against primary key index
12:   if pkey is invalid then
13:     mark position of bitmap to 1

```

Figure 7: Pseudo Code for Merge Repair

this, we propose performing index repair operations in the background to clean up obsolete index entries. Index repair can either be performed during merge time, which we will call *merge repair*, or scheduled independently from merges, which we will call *standalone repair*.

The basic idea of index repair is to validate each primary key in a component by searching the primary key index. For efficiency, the primary keys in a given component of a secondary index should only be validated against newly ingested keys. To keep track of the repair progress, we associate a *repaired timestamp* (repairedTS) with each disk component of a secondary index. During a repair operation, all primary key index components with maxTS no larger than the repairedTS can be pruned.² A repaired component receives a new repairedTS computed as the maximum timestamp of the unpruned primary key index components. To clarify this, consider the example in Figure 6. Component 1-15 of the secondary index has an initial repairedTS of 15. To repair this component, we only need to search components 11-18 and 19 of the primary key index, while component 1-10 can be pruned. After the repair operation, the new repairedTS of this component would be 19.

A naive implementation of merge repair would simply validate each primary key by performing a point lookup against the primary key index. For standalone repair, one could simply produce a new component with only valid entries without merging. However, this implementation would be highly inefficient because of the expensive random point lookups.

To handle this, we propose a more efficient merge repair algorithm (Figure 7). The basic idea is to first sort the primary keys to improve point lookup performance, and further use an immutable bitmap to avoid having to sort entries back into secondary key order. The immutable bitmap of a disk component indicates whether each of the component’s index entries is (still) valid or not, and thus it stores one bit per entry. Without loss of generality, we assume that a bit being 1 indicates that the corresponding index entry is invalid.

²This optimization applies to Timestamp Validation as well.




Figure 8: Merge Repair Example

Initially, we create a scan cursor over all merging components to obtain all valid index entries, i.e., entries where their immutable bitmap bits are 0. These entries are directly added to the new component (line 5). Meanwhile, the primary keys, with their timestamps and positions in the new component, are streamed to a sorter (line 6). These sorted primary keys are then validated by searching the primary key index. As an optimization, if the number of primary keys to be validated is larger than the number of recently ingested keys in the primary key index, we can simply merge scan the sorted primary keys and the primary key index. If a key is found to be invalid, that is, if the same key exists in the primary key index with a larger timestamp, we simply set the corresponding position of the new component’s bitmap to 1 (lines 12-13). Standalone repair can be implemented similarly, except that only a new bitmap is created.

To illustrate the immutable bitmap and the repair process, consider the example secondary index in Figure 4. Suppose we want to merge and repair all components of the secondary index. The process is shown in Figure 8. The index entries scanned from its old components are directly added to the new component, and in the meanwhile they are sorted into primary key order. The sorted primary keys are validated using the primary key index. During validation, the key 101 with timestamp ts1 is found to be invalid since the same key exists with a larger timestamp ts4. This index entry has ordinal position 1 (denoted “(1)” in the figure), so first bit in the component’s bitmap is set to 1. Note that this invalid entry will be physically removed during the next merge.

Bloom Filter Optimization. It is tempting to use Bloom filters to further optimize the index repair operation. The idea is that if the Bloom filters of the primary key index components do not contain a key, which implies that the key has not been updated, then the key can be excluded from sorting and further validation. However, if implemented directly, this would provide little help since a dataset’s various LSM-trees are merged independently. Consider the example LSM indexes back in Figure 1. Suppose we want to merge and repair the two disk components of Secondary Index 1. Since the disk components of the primary key index have already been merged into a single component beforehand, its Bloom filter would always report positives, which would provide no help and actually cause some extra overhead.

To maximize the effectiveness of the Bloom filter optimization, one must ensure that during each repair operation the unpruned primary key index components are always strictly newer than the keys in the repairing component(s). For this purpose, we can use a correlated merge policy [11] to synchronize the merge of all secondary indexes with the primary key index in order to ensure that their components

are always merged together. Furthermore, all secondary indexes must be repaired during every merge.

The Bloom filter optimization improves repair efficiency and is thus suitable for update-heavy workloads that require secondary indexes to be frequently repaired. However, for workloads that contain few updates, it might be better to just schedule repair operations during off-peak hours, and thus the Bloom filter optimization may not be suitable. To alleviate the tuning effort required from the end-user, we also plan to develop auto-tuning techniques in the future.

5. MUTABLE-BITMAP STRATEGY

In this section, we present the Mutable-bitmap strategy designed for maintaining a primary index with filters.

5.1 Overview

The key difficulty of applying filters to the LSM-tree is its out-of-place update nature, that is, updates are added to the new component. In the case of updates, the filter of the new component must be maintained using old records so that queries would not miss any new updates. The Eager strategy performs point lookups to maintain filters using old records, incurring a high point lookup cost during data ingestion. The Validation strategy skips point lookups but requires queries to access all newer components for validation, halving the pruning capabilities of filters.

The Mutable-bitmap strategy presented below aims at both maximizing the pruning capabilities of filters and reducing the point lookup cost during data ingestion. The first goal can be achieved if old records from disk components can be deleted directly. However, if we were to place new updates directly into the disk components where old records are stored, a lot of complexity would be introduced on concurrency control and recovery. Instead, our solution is to add a mutable bitmap to each disk component to indicate the validity of each entry using a very limited degree of mutability. We present efficient solutions to address concurrency control and recovery issues, exploiting the simple semantics of mutable bitmaps, that is, writers only change bits from 0 to 1 to mark records as deleted³. To minimize the point lookup cost, the maintenance of mutable bitmaps is performed by searching the primary key index instead of accessing full records. To achieve this, we synchronize the merges of the primary index and the primary key index using the correlated merge policy (as in Section 4.4). An alternative implementation is to add a key-only B⁺-tree to each primary index component, combining two indexes together.

5.2 Data Ingestion

For ease of discussion, let us first assume that there are no concurrent flush and merge operations, the handling of which are postponed to Section 5.3. An *insert* is handled exactly as in the Eager/Validation strategy, and no bitmaps are updated because no record is deleted. To *delete* a record given its key, we first search the primary key index to locate the position of the deleted key. If the key is found and is in a disk component, then its corresponding bit in that component’s bitmap is set (mutated) to 1. An anti-matter key is also added to the memory component, for two reasons. First, we view the mutable bitmap as an auxiliary structure built on top of LSM that should not change the semantics

³Aborts internally change bits from 1 to 0.




Figure 9: Upsert Example with Mutable-bitmap Strategy

of LSM itself. Second, if the Validation strategy is used for secondary indexes, inserting anti-matter entries ensures that validation can be performed with only recently ingested keys. To *upsert* a record, the primary key index is first searched to mark the old record as deleted if necessary, and the new record is added to the index memory components. Any filters of the memory components are only maintained based on the new record, not the old one.

Consider the running example from Figure 2. Figure 9 shows the resulting LSM-trees after upserting a new record (101, NY, 2018). The secondary index is not shown since it can be maintained using either the Eager or Validation strategy. In this case, the primary index and the primary key index are synchronized and their components share a mutable bitmap to indicate the validity of their records. To upsert the new record, the primary key index is searched to locate the position of the old record and the bitmap is mutated to mark the old record as deleted. The range filter of the memory component is maintained based only on the new record, avoiding unnecessary widening based on 2015.

We now discuss concurrency control and recovery issues for the mutable bitmaps. We assume that each writer acquires an exclusive (X) lock on a primary key throughout the (record-level) transaction. To prevent two writers from modifying the same bitmap byte, one can use latching or compare-and-swap instructions. For recovery, we use an additional update bit in the log record for each delete or upsert operation to indicate whether the key existed in a disk component. To abort a transaction, if the update bit is 1, we simply perform a primary key index lookup (without bitmaps) to unset the bit from 1 to 0. To unify the recovery of bitmaps with LSM-trees, we use a no-steal and no-force policy for bitmaps as well. A modified bitmap page is pinned until the transaction terminates to prevent dirty pages from being flushed. Regular checkpointing can be performed to flush dirty pages of bitmaps. Upon recovery, committed transactions are simply replayed to bring bitmaps up-to-date based on the last checkpointed LSN. Again, a log record is replayed on the bitmaps only when its update bit is 1.

5.3 Concurrency Control for Flush/Merge

Mutable bitmaps introduce concurrency control issues for flush and merge operations. This is because concurrent writers may need to modify the bitmaps of the components that are being formed by flush or merge operations. This problem bears some similarities with previous work on online index construction [23, 30], which builds new indexes concurrently with updates. Here we propose two concurrency control methods that differ in how new updates are applied. Due to space limitations, we only sketch the basic idea of the proposed methods, leaving their details to [22].

In the *Lock* method, new updates are directly applied to the new component. The component builder, which builds the new component during a flush or merge operation, keeps tracks of the progress of scanned keys. If a key to be deleted was already scanned by the component builder, then the delete is applied to the new component directly; otherwise, it is only applied to the old component, which will then be ignored by the component builder. To avoid accessing the same key, the component builder has to acquire a shared lock on each scanned key. The *Side-file* method avoids this locking overhead by buffering the new updates in a side-file and applying them after the new component has been fully built. To avoid interference from concurrent updates, the component builder has to create snapshots of mutable bitmaps at the beginning of a flush or merge operation.

6. EXPERIMENTAL EVALUATION

In this section, we evaluate the proposed techniques in the context of Apache AsterixDB [1]. We are primarily interested in evaluating the effectiveness of the various point lookup optimizations (Section 6.2) and the ingestion and query performance of the Eager, Validation, and Mutable-bitmap strategies discussed in this paper (Sections 6.3 - 6.4). We also evaluate the proposed secondary index repair method in detail (Section 6.5). See [22] for the additional evaluation of the two proposed concurrency control methods for the Mutable-bitmap strategy.

6.1 Experimental Setup

Since our work focuses on a partitioned database architecture, all experiments were performed on a single node with a single dataset partition. The overall performance of multiple partitions generally achieves near-linear speedup since both data ingestion and query processing are performed at each partition locally. The node has a 4-core AMD Opteron 2212 2.0GHZ CPU, 8GB of memory, and two 7200 rpm SATA hard disks. We used one disk for transactional logging and one for the LSM storage. We allocated 5GB of memory for the AsterixDB instance. Within that allocation, the disk buffer cache size is set at 2GB, and each dataset is given a 128MB budget for its memory components. Each LSM-tree has two memory components to minimize stalls during a flush. We used a tiering merge policy with a size ratio of 1.2 throughout the experiments, similar to the one used in other systems [2, 3]. This policy merges a sequence of components when the total size of the younger components is 1.2 times larger than that of the oldest component in the sequence. We set the maximum mergeable component size at 1GB to simulate the effect of disk components accumulating within each experiment period. Each LSM-tree is merged independently. The Bloom filter false positive rate setting is 1% and the data page size is set at 128KB to accommodate sequential I/Os. Finally, whenever a scan is performed, a 4MB read-ahead size is used to minimize random I/Os. To verify the applicability of the proposed techniques on SSDs, we repeated some key experiments on a node with 4-core Intel i7 2.70GHZ CPU, 16GB memory, and a 500GB SSD. All configurations on the SSD node were the same, except that the buffer cache size is set to 4GB and disk page size is reduced to 32KB.

For the experimental workload, although YCSB [15] is a popular benchmark for key-value store systems, it is not suitable for our evaluation since it does not have secondary

keys nor secondary index queries. As a result, we implemented a synthetic tweet generator for our evaluation, as in other LSM secondary indexing work [10, 26]. Each tweet has several attributes such as ID, message_text, user_id, location, and creation_time etc., and its size is about 500 bytes with variations due to the variable length (450 to 550 bytes) of the randomly generated tweet messages. Among these attributes, the following three are related to our evaluation. First, each tweet has an ID as its primary key, which is a randomly generated 64-bit integer. Second, the user_id attribute, which is a randomly generated integer in the range 0 to 100K, is used for formulating secondary index queries with various controlled selectivities. Finally, each tweet has a creation_time attribute, which is a monotonically increasing timestamp used to test the range filter.

6.2 Point Lookup Optimizations

We first studied the effectiveness of the point lookup optimizations discussed in Section 3.2. To this end, we performed a detailed experimental analysis as follows: On top of a naive lookup implementation (denoted as “naive”) which only sorts the primary keys, we enabled the batched point lookup with a batch size of 16MB (denoted as “batch”), stateful B^+ -tree lookup (denoted as “sLookup”), and blocked Bloom filter (denoted as “bBF”) optimizations one by one. We also studied another optimization proposed by Jia [20] (denoted as “pID”). Its basic idea is to propagate and use the component IDs of the secondary index components in which keys are found to prune primary index components during point lookups⁴.

To prepare the experiment dataset, we inserted 80 million tweet records with no updates. The resulting primary index has about 30GB of data and 20 disk components, while the secondary index on the user_id attribute only has 3GB of data and less than 10 disk components. We then evaluated the query performance for different controlled selectivities based on the user_id attribute. For each query selectivity, queries with different range predicates were executed until the cache was warmed and the average stable running time is reported. The running time of different selectivities is reported in Figures 10a and 10b. For low query selectivities (Figure 10a), which select a small number of records, batching slightly improves query performance by reducing random I/Os, while other optimizations are less effective since the query time is dominated by disk I/Os. For high query selectivities (Figure 10b), we also included the full scan time as a baseline. The upper line, at about 750s, represents the full scan time on the experimental dataset. Since disk read-ahead caused long I/O wait times, we further evaluated an optimized case on a dataset prepared with sequential primary keys (lower line). In that case, sequential I/Os and the OS’s asynchronous read-ahead were fully exploited to minimize the I/O wait time. As the query selectivity increases, the running time of the naive point lookup implementation grows quickly, since accessing multiple LSM components leads to random I/Os. The batched point lookup is the most effective at avoiding random I/Os. For large selectivities, the stateful B^+ -tree lookup and blocked Bloom filter optimizations further reduce the in-memory search cost,

⁴Jia’s original technique actually propagated the ranges of the range filter built on a time-correlated attribute; this is equivalent to propagating component IDs in our setting.

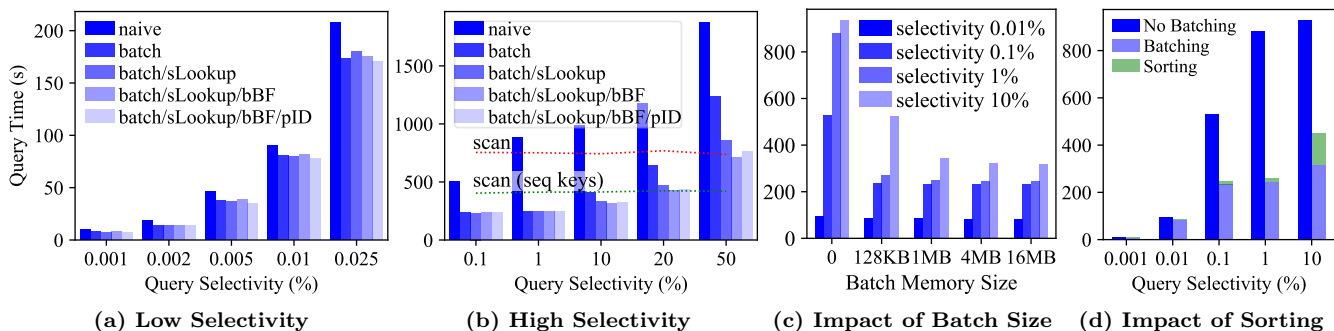


Figure 10: Effectiveness of Point Lookup Optimizations

since disk I/O cost is bounded as most pages must be accessed. Note that for large selectivities, a full scan starts to outperform a secondary index search since most pages of the primary index must be accessed while a secondary index search incurs the extra cost of accessing the secondary index and sorting primary keys. However, these optimizations together greatly improved the range of applicability of LSM-based secondary indexes, allowing query optimizers, especially rule-based ones, to have more confidence when choosing secondary indexes. Contrary to Jia [20], we found here that propagating component IDs provides little benefit. The reason is that Jia [20] considered an append-only temporal workload with hundreds of (filtered) LSM components in a dataset, so the propagation of component IDs led to the skipping of a large number of Bloom filter tests. However, in a more general setting, skipping a smaller number of Bloom filter tests per key would not make so big a difference.

We further evaluated batched point lookup in terms of the available batching memory and sorting overhead. In each experiment that follows, stateful B^+ -tree lookup and blocked Bloom filter were enabled by default. The running time of different query selectivities under different batch sizes is shown in Figure 10c. For selective queries, a small batch size such as 128KB already provides optimal performance since a batch of keys are often distributed over a large number of pages; for non-selective queries, a few megabytes suffice to provide optimal performance as well. Finally, we evaluated the additional sorting overhead since batching destroys the primary key ordering of the final results. The running time of different query selectivities is shown in Figure 10d. Even when sorting is performed, batching still improves the overall query performance. This is because the point lookup step needs to fetch records distributed across a large number of pages, while the resulting records can often fit into a small number of pages and can be sorted efficiently.

To summarize, batched point lookup is the most effective optimization for reducing random I/Os when accessing LSM components. Stateful B^+ -tree lookup and blocked Bloom filter are mainly effective for non-selective queries at further reducing the in-memory search cost.

6.3 Ingestion Performance

We next evaluated the ingestion performance of the different maintenance strategies, focusing on the following key questions: (1) What is the effectiveness of building a primary key index, which reduces the point lookup cost, for improving the overall ingestion throughput, since maintaining the primary key index itself incurs extra cost? (2) Does

repairing secondary indexes become a new bottleneck, especially for a dataset with multiple secondary indexes? (3) What is the effectiveness of building a single primary key index and the proposed repair algorithm (Section 4.4) for improving the ingestion throughput as compared to the deleted-key B^+ -tree strategy supported by AsterixDB?

To answer these questions, we used insert and upsert workloads in the following evaluation. Delete workloads were omitted since the cost of deletes would be similar to upserts; that is, an upsert is logically equivalent to a delete plus an insert. Unless otherwise noted, each experiment below was run for 6 hours with the goal being to ingest as much data as possible without concurrent queries. Each dataset has a primary index with a component-level range filter on the creation_time attribute, a primary key index, and a secondary index on the user_id attribute.

6.3.1 Insert Workload

For inserts, we evaluated two methods for enforcing key uniqueness, using either the primary index or the primary key index. In the former case, the primary key index was omitted to eliminate its overhead. This workload was controlled by a duplicate ratio, which is the ratio of duplicates among all records. Duplicates were randomly generated following a uniform distribution over all the past keys.

Figure 11 shows the insert throughput under different duplicate ratios on both hard disks and SSDs. Note that the experiment on hard disks was run for 12 hours so that even the primary key index cannot be totally cached. As the result shows, without the primary key index, ingestion performance degraded quickly once the dataset could not be totally cached; point lookups during ingestion incur a large overhead on both hard disks and SSDs. Building a primary key index greatly improves the ingestion throughput since the keys are much smaller and can be better cached to reduce disk I/Os. Even though the ingestion throughput drops when the primary key index cannot be totally cached, building a primary key index is still helpful by increasing the cache hit ratio. The duplicate-heavy workload also results in a higher ingestion throughput when using the primary key index since duplicate keys are simply excluded from insertion into the storage.

6.3.2 Upsert Workload

The upsert workload was controlled by an update ratio, which is the ratio of updates (records with past ingested keys) among total records. Updates are randomly generated by following either a uniform distribution, that is, all




Figure 11: Insert Ingestion Performance

past keys are updated equally, or a Zipf distribution with a theta value 0.99 as in the YCSB benchmark [15], that is, recently ingested keys are updated more frequently. Unless otherwise specified, the update ratio was chosen as 10% and the updates followed a uniform distribution.

For the Validation strategy, we evaluated two variations to measure the overhead of repairing secondary indexes. In the first variation, repair was totally disabled to maximize ingestion performance. In the second variation, merge repair was enabled but without the Bloom filter optimization (Section 4.4) to evaluate the worse case repairing overhead. For the Mutable-bitmap strategy, the secondary index was maintained using the Validation strategy without repair to minimize the ingestion overhead due to secondary indexes. The Side-file method was used for concurrency control to minimize the locking overhead.

Basic Upsert Ingestion Performance. In this experiment, we compared the proposed Validation and Mutable-bitmap strategies with the Eager strategy in terms of the upsert ingestion performance under different update ratios, ranging from no updates to 50% updates following either the uniform or Zipf distribution. The experimental results are shown in Figure 12. The Eager strategy, which ensures that secondary indexes are always up-to-date, has the worst ingestion performance because of the point lookups to maintain secondary indexes using the old records. On the other hand, the Validation strategy without repairing has the best ingestion performance since secondary indexes are not cleaned up at all. The result also shows that repairing secondary indexes incurs only a small amount of extra overhead. Of course, since secondary indexes are not always up-to-date, the Validation strategy sacrifices query performance, which we will evaluate later in Section 6.4. The Mutable-bitmap strategy also has much better ingestion performance than the Eager strategy since it only searches the primary key index instead of accessing full records. Updates generally have a small impact on the overall ingestion throughput, since updated records must be inserted into memory and subsequently flushed and merged (as for inserted records). The Eager and Mutable-bitmap strategies both benefit from skewed update workloads since most of the updates only touch recent keys, reducing disk I/O incurred by the point lookups. These results again confirm that it is helpful to build a primary key index to reduce the point lookup cost.

Ingestion Impact of Merge Operations. In this set of experiments, we evaluated the impact of merge operations on upsert ingestion performance since more merges can reduce the point lookup cost by reducing the number of components. To evaluate this, we varied the maximum

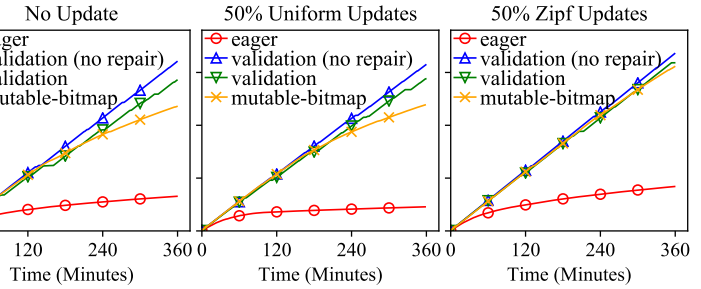


Figure 12: Upsert Ingestion Performance

mergeable component size to control the merge frequency. The resulting ingestion throughput is shown in Figure 13a. Except that more merges negatively impact the ingestion throughput of all strategies, point lookups still incur a large amount of overhead during data ingestion and the same relative performance trends still hold.

Ingestion Impact of Secondary Indexes. We also evaluated the scalability of the strategies by adding more secondary indexes. We excluded the Mutable-bitmap strategy, since it is unaffected by secondary indexes, but instead included the deleted-key B⁺-tree strategy for comparison. Figure 13b shows the ingestion throughput for different numbers of secondary indexes. Maintaining more secondary indexes negatively impacts the ingestion performance of all the strategies since more LSM-trees must be maintained. It also has a larger impact on the Validation strategies. The reason is that the bottleneck of the Eager strategy lies in the point lookups, while for the Validation strategy, its bottleneck is flushing and merging the LSM-trees. This impact is more obvious after adding the fourth secondary index, since our experiment machine has four cores and more secondary indexes result in higher contention among all merge threads. The experiment also shows the scalability of the proposed index repair operations since it introduces just a small overhead on data ingestion. Comparing to the current AsterixDB deleted-key B⁺-tree strategy, the negative impact of index repair operations has been greatly reduced by using a single primary key index and the efficient repair algorithm presented in Section 4.4.

6.4 Query Performance

We next evaluated the query performance of the different maintenance strategies, focusing on the following two aspects. First, we evaluated the overhead of the Validation

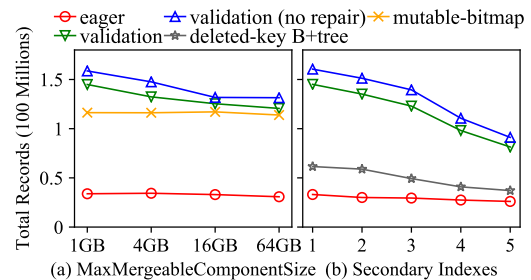


Figure 13: Impact of Merge Operations and Secondary Indexes on Upsert Ingestion Performance

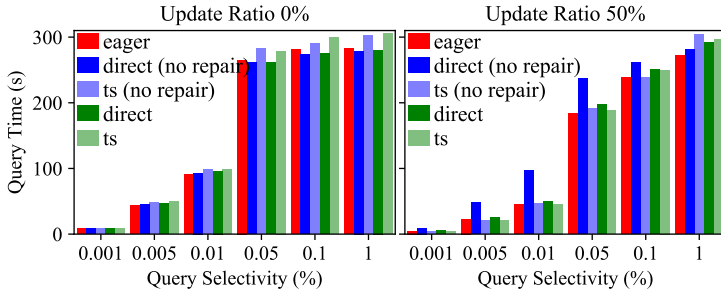


Figure 14: Non-Index-Only Query Performance

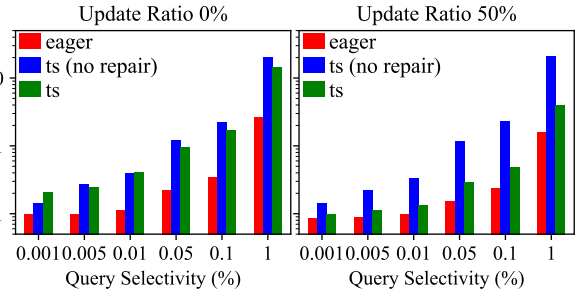


Figure 15: Index-Only Query Performance

strategy with both non-index-only and index-only queries compared to the read-optimized Eager strategy and the benefit of repairing secondary indexes. Second, we evaluated the pruning capabilities of filters resulting from different maintenance strategies. As in Section 6.2, each dataset that follows was prepared by upserting 80 million records with different actual update ratios (0% or 50%).

6.4.1 Secondary Index Query Performance

We first evaluated the overhead of the Validation strategy on secondary index range queries. We again considered two variations of the Validation strategy, depending on whether merge repair was enabled, to evaluate the benefit of repairing secondary indexes. Each query for a given selectivity was repeated with different range predicates until the cache was warmed and the average stable time is reported.

Non-Index-Only Query Performance. For non-index-only queries, we enabled batching (with 16MB memory), stateful B⁺-tree search, and blocked Bloom filter to optimize the subsequent point lookups. For the Validation strategy, we evaluated both Direct Validation (denoted as “direct”) and Timestamp Validation (denoted as “ts”). The running time of non-index-only queries is shown in Figure 14. In general, the Validation strategy has comparable performance to the Eager strategy for non-index-only queries. For the append-only workload (left side), the Direct Validation method has similar performance to the Eager strategy, since secondary indexes do not contain any obsolete entries. However, the Timestamp Validation method leads to some extra overhead because of the validation step. The results are different for the update-heavy workload (right side). When the secondary index is not repaired, the Direct Validation method leads to high overhead for selective queries, as a lot of I/O is wasted searching for obsolete keys. However, the extra overhead diminishes when the selectivity becomes larger, such as 0.1 - 1%, since even searches for valid keys must access almost all pages of the primary index. The Timestamp Validation method is helpful for reducing wasted I/O by filtering out the obsolete keys using the primary key index. With merge repair, most of the obsolete keys would be removed, and thus both validation methods would have comparable query performance to the Eager strategy.

Due to space limitations, we omitted one experiment with a smaller cache memory (512MB) to evaluate the Timestamp Validation method when the primary key index cannot be totally cached. We found that this has little impact on the Timestamp Validation method since the primary key index is much smaller than the primary index, resulting in only a small amount of extra I/O during validation.

Index-Only Query Performance. Figure 15 shows the query time (log scale) for index-only queries. We omit the Direct Validation method since it must fetch records for validation and thus has same performance as for non-index-only queries. In general, the Validation strategy performs worse (3x - 5x) than the Eager strategy because of its extra sorting and validation steps. Even without obsolete entries (left side), the validation step still leads to extra overhead. Still, Timestamp Validation provides much better performance than Direct Validation would by accessing the primary key index. Finally, merge repair is helpful for index-only queries by increasing repaired timestamps to prune more primary key index components during validation, as shown in the append-only case (left side), and by cleaning up obsolete entries for the update-heavy case (right side).

6.4.2 Range Filter Query Performance

We next evaluated the effectiveness of the maintenance strategies for filters. In particular, we used the range filter in our evaluation. Recall that the range filter was built on the creation_time attribute, which is a monotonically increasing timestamp. The range of the creation_time attribute of all tweet records in the experiment dataset spanned about 2 years. Each query in this experiment has a range predicate on the creation_time attribute, and it is processed by scanning the primary index with the component-level pruning provided by the range filters. We evaluated two types of queries, queries that access recent data (with creation_time > T) and that access old data (with creation_time < T). Each query was repeated 5 times with a clean cache for each run, and the average query time is reported.

The query times with range filters are summarized in Figure 16. For the queries that access recent data, all the strategies provide effective pruning capabilities. The Mutable-bitmap strategy further improves scan performance since LSM components are accessed one by one and the reconcil-

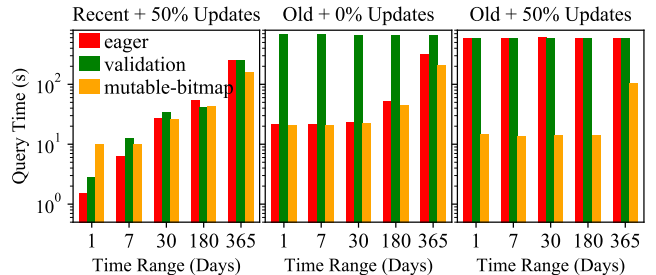


Figure 16: Query Performance of Range Filters

iation step is no longer needed. In contrast, when accessing old data, the Validation strategy provides no pruning capability since all newer components must be accessed to ensure correctness. The Eager strategy is only effective for the append-only case, but recall that its point lookups lead to high cost during data ingestion. The Mutable-bitmap strategy provides effective pruning capabilities via the use of Mutable bitmaps under all settings, and does so with only a small amount of overhead on data ingestion.

6.5 Index Repair Performance

We then evaluated the index repair performance of the proposed Validation strategy (referred as *secondary repair*) in detail, as well as the proposed Bloom filter optimization (denoted as “bf”). For comparison, we also evaluated the index repair method proposed by DELI [31] (referred as *primary repair*). Recall that DELI repairs secondary indexes by merging or scanning primary index components to identify obsolete records, while our approach uses a primary key index to avoid accessing full records.

In each of the following experiments, we upserted 100 million tweet records with merge repair enabled. Since the two methods trigger repair operations differently, the resulting secondary indexes may have different amounts of obsolete entries during data ingestion. To enable a fair comparison, instead of measuring the overall ingestion throughput, we measured the index repair performance directly as follows. For every 10 million records ingested, we stopped the ingestion and triggered a full repair operation to bring all secondary indexes up-to-date. This provides the trend of index repair performance as data accumulates. For our secondary repair method, standalone repair was performed to exclude the extra overhead due to merges.

Basic Repair Performance. We first evaluated the index repair performance under different update ratios (0% and 50%). For primary repair [31], we evaluated two variations depending whether a full merge is performed as a by-product of the repair operation. The index repair performance over time is depicted in Figure 17, where each data point represents the time to complete a repair operation. For primary repair, a full merge leads to extra overhead for append-only workloads, but improves subsequent repair performance for update-heavy workloads. The secondary repair method proposed here always outperforms the primary repair methods because only it only accesses the primary key index, significantly reducing disk I/O. The Bloom filter optimization further improves repair performance by reducing the volume of primary keys to be sorted and validated.

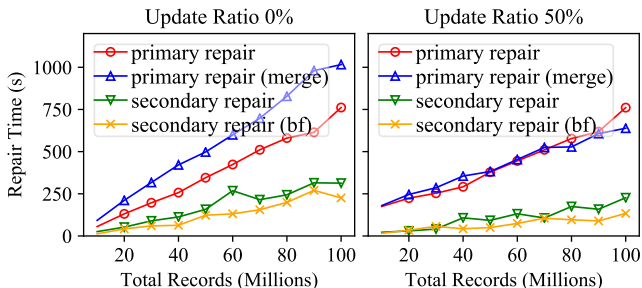


Figure 17: Repair Performance with Varying Update Ratio

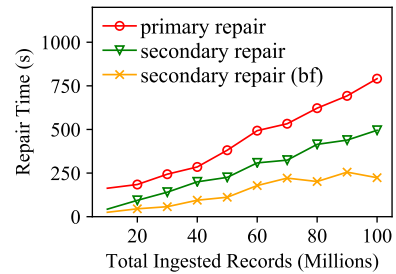


Figure 18: Repair Impact of Secondary Indexes

Impact of Secondary Indexes. We next used 5 secondary indexes to evaluate the scalability of the index repair methods. The update ratio was set to 10%. For the secondary repair method, each secondary index was repaired in parallel using multi-threading. The index repair time is depicted in Figure 18. Having more secondary indexes negatively impacts the performance of both methods. For the primary repair method, more anti-matter entries must be inserted into more secondary indexes. For the secondary repair method, more secondary index entries must be scanned and sorted for validation. The result does show that the proposed secondary index repair method is easily parallelizable, since it only performs a small amount of I/O and most operations are CPU-heavy. Furthermore, the Bloom filter optimization reduces the negative impact of having more secondary indexes, as it sorts fewer keys, significantly reducing the I/O overhead during the sort step.

7. CONCLUSION

In this paper, we have presented techniques for efficient data ingestion and query processing for LSM-based storage systems. We described and evaluated a series of optimizations for efficient point lookups, greatly improving the range of applicability for LSM-based secondary indexes. We further presented new and efficient strategies for maintaining LSM-based auxiliary structures, including secondary indexes and filters. The Validation strategy defers secondary index maintenance to the background using a primary key index, significantly improving ingestion performance by eliminating ingestion-time point lookups. This leads to a small overhead for non-index-only queries, but a relatively high overhead for index-only queries because of the extra validation step. The Mutable-bitmap strategy maximizes the pruning capabilities of filters and improves ingestion performance by only accessing the primary key index, instead of full records, during data ingestion.

We plan to extend our work in several directions. First, we hope to extend the Validation strategy to support more efficient processing of index-only queries, further improving its applicability. Second, we plan extend the proposed strategies to let queries drive the maintenance of auxiliary structures and to develop auto-tuning techniques so that the system could dynamically adopt the optimal maintenance strategies for a given workload.

Acknowledgments. This work is supported by NSF awards CNS-1305430, IIS-1447720, and IIS-1838248 along with industrial support from Amazon, Google, and Microsoft and support from the Donald Bren Foundation (via a Bren Chair).

8. REFERENCES

- [1] AsterixDB. <https://asterixdb.apache.org/>, 2018.
- [2] Cassandra. <http://cassandra.apache.org/>, 2018.
- [3] HBase. <https://hbase.apache.org/>, 2018.
- [4] LevelDB. <http://leveldb.org/>, 2018.
- [5] MyRocks. <https://http://myrocks.io/>, 2018.
- [6] Phoenix. <https://phoenix.apache.org/>, 2018.
- [7] RocksDB. <http://rocksdb.org/>, 2018.
- [8] M. Y. Ahmad and B. Kemme. Compaction management in distributed key-value datastores. *PVLDB*, 8(8):850–861, 2015.
- [9] S. Alsubaiee et al. AsterixDB: A scalable, open source BDMS. *PVLDB*, 7(14):1905–1916, 2014.
- [10] S. Alsubaiee et al. Storage management in AsterixDB. *PVLDB*, 7(10):841–852, 2014.
- [11] S. Alsubaiee et al. LSM-based storage and indexing: An old idea with timely benefits. In *International ACM Workshop on Managing and Mining Enriched Geo-Spatial Data*, pages 1–6, 2015.
- [12] J. L. Bentley and A. C.-C. Yao. An almost optimal algorithm for unbounded searching. *Information Processing Letters*, 5(3):82–87, 1976.
- [13] B. H. Bloom. Space/time trade-offs in hash coding with allowable errors. *CACM*, 13(7):422–426, 1970.
- [14] F. Chang et al. Bigtable: A distributed storage system for structured data. *ACM TOCS*, 26(2):4:1–4:26, 2008.
- [15] B. F. Cooper et al. Benchmarking cloud serving systems with YCSB. In *ACM SoCC*, pages 143–154, 2010.
- [16] N. Dayan et al. Monkey: Optimal navigable key-value store. In *ACM SIGMOD*, pages 79–94, 2017.
- [17] N. Dayan and S. Idreos. Dostoevsky: Better space-time trade-offs for LSM-tree based key-value stores via adaptive removal of superfluous merging. In *ACM SIGMOD*, pages 505–520, 2018.
- [18] G. Golan-Gueta et al. Scaling concurrent log-structured data stores. In *EuroSys*, pages 32:1–32:14, 2015.
- [19] G. Graefe. Modern B-tree techniques. *Found. Trends Databases*, 3(4):203–402, 2011.
- [20] J. Jia. *Supporting Interactive Analytics and Visualization on Large Data*. PhD thesis, Computer Science Department, University of California, Irvine, 2017.
- [21] L. Lu et al. WiseKey: Separating keys from values in SSD-conscious storage. In *USENIX Conference on File and Storage Technologies*, pages 133–148, 2016.
- [22] C. Luo and M. J. Carey. Efficient data ingestion and query processing for LSM-based storage systems. *CoRR*, abs/1808.08896, 2018.
- [23] C. Mohan and I. Narang. Algorithms for creating indexes for very large tables without quiescing updates. In *ACM SIGMOD*, pages 361–370, 1992.
- [24] P. O’Neil et al. The log-structured merge-tree (LSM-tree). *Acta Inf.*, 33(4):351–385, 1996.
- [25] F. Putze et al. Cache-, hash-, and space-efficient Bloom filters. *J. Exp. Algorithmics*, 14:4:4.4–4:4.18, 2010.
- [26] M. A. Qader et al. A comparative study of secondary indexing techniques in LSM-based NoSQL databases. In *ACM SIGMOD*, pages 551–566, 2018.
- [27] V. Raman et al. DB2 with BLU acceleration: So much more than just a column store. *PVLDB*, 6(11):1080–1091, 2013.
- [28] R. Sears and R. Ramakrishnan. bLSM: A general purpose log structured merge tree. In *ACM SIGMOD*, pages 217–228, 2012.
- [29] M. Seidemann and B. Seeger. ChronicleDB: A high-performance event store. In *EDBT*, pages 144 – 155, 2017.
- [30] V. Srinivasan and M. J. Carey. Performance of on-line index construction algorithms. In *EDBT*, pages 293–309, 1992.
- [31] Y. Tang et al. Deferred lightweight indexing for log-structured key-value stores. In *International Symposium on Cluster, Cloud and Grid Computing*, pages 11–20, 2015.
- [32] S. Wang et al. Lightweight indexing of observational data in log-structured storage. *PVLDB*, 7(7):529–540, 2014.
- [33] S. Wang et al. Fast and adaptive indexing of multi-dimensional observational data. *PVLDB*, 9(14):1683–1694, 2016.
- [34] X. Wu et al. LSM-trie: an LSM-tree-based ultra-large key-value store for small data. In *USENIX Annual Technical Conference*, pages 71–82, 2015.
- [35] M. Ziauddin et al. Dimensions based data clustering and zone maps. *PVLDB*, 10(12):1622–1633, 2017.