

XIndex: A Scalable Learned Index for Multicore Data Storage

Chuzhe Tang^{†‡}, Youyun Wang^{†‡}, Zhiyuan Dong^{†‡}, Gansen Hu^{†‡}

Zhaoguo Wang^{†‡}, Minjie Wang[◇], Haibo Chen^{†‡}

[†] Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University

[‡] Shanghai Key Laboratory of Scalable Computing and Systems, Shanghai Jiao Tong University

[◇] Department of Computer Science, New York University

Abstract

We present XIndex, a concurrent ordered index designed for fast queries. Similar to a recent proposal of the learned index, XIndex uses learned models to optimize index efficiency. Comparing with the learned index, XIndex is able to effectively handle concurrent writes without affecting the query performance by leveraging fine-grained synchronization and a new compaction scheme, Two-Phase Compaction. Furthermore, XIndex adapts its structure according to runtime workload characteristics to support dynamic workload. We demonstrate the advantages of XIndex with both YCSB and TPC-C (KV), a TPC-C variant for key-value stores. XIndex achieves up to 3.2× and 4.4× performance improvement comparing with Masstree and Wormhole, respectively, on a 24-core machine, and it is open-sourced¹.

CCS Concepts • Information systems → Data structures; • Theory of computation → Concurrent algorithms.

1 Introduction

The pioneering study on the learned index [15] opens up a new perspective on how machine learning can re-sculpt the decades-old system component, indexing structure. The key idea of the learned index is to use learned models to approximate indexes. It trains the model with records' keys and their positions, then uses the model to predict the position with the given key.

To deliver high lookup performance, the learned index uses simple learned models such as linear models or single-layer neural networks. To accommodate to the limited capacity of simple models (inability to well fit complex functions),

the learned index adds extra requirements on the data layout. For instance, to learn an ordered index, it requires the data to be both ordered and contiguous, so the key-position mapping is easier to learn. Using simple learned models, the learned index² performs 1.5-3× better than B-tree [15].

However, the current study of the learned index is still preliminary and lacks practicability in a broad class of real-world scenarios because of two limitations. First, it does not support any modifications, including inserts, updates, or removes. Second, it assumes the workload has a relative static query distribution³ — it assumes all data are uniformly accessed. Nevertheless, making the learned index practical for dynamic workloads with writes is not an easy task, because its high performance is tied closely to both data distribution and query distribution, especially for ordered index. First, the learned index requires simple data distribution by enforcing ordered and contiguous data layout. Thus, it needs to reconstruct the layout and retrain the model to handle every write request. Although there are proposals [9, 18] which try to handle writes for the learned index efficiently, none of them can ensure the correctness in the face of concurrent operations. Second, the learned index is sensitive to the changes of the data and query distribution at runtime. Its current design employs several learned models, and each is in charge of a portion of data. However, the prediction error of every model varies. At the same time, queries in real-world workloads tend to be skew, where some “hot” keys are much more frequently queried than others [7, 10, 16, 25]. As a result, when the models that index those hot keys have large errors, queries can incur high overhead (Section 2.2).

In this paper, we present XIndex, a new fully-fledged concurrent index structure inspired by the learned index. While XIndex leverages learned models to speed up the lookup, it can handle concurrent writes efficiently with good scalability. Moreover, it is designed to adapt its structure deterministically at runtime and decouple its efficiency from runtime workload characteristics. Specifically, this paper makes the following contributions:

¹<https://ipads.se.sjtu.edu.cn:1312/opensource/xindex.git>

PPoPP '20, February 22–26, 2020, San Diego, CA, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '20)*, February 22–26, 2020, San Diego, CA, USA, <https://doi.org/10.1145/3332466.3374547>.

²In following paragraphs, we refer to the learned range index proposed in [15] as “the learned index”.

³The query distribution describes the access frequencies of keys among queries within a specific workload. By contrast, data distribution describes the keys and their lookup positions within a dataset.

A scalable and concurrent learned index, XIndex.

With understandings of the learned index’s pros and cons, XIndex range-partitions data into different groups, each attached with learned models for lookup and a delta index to handle inserts. To achieve high performance, XIndex exploits a combination of innovative methods (e.g., Two-Phase Compaction) and classic techniques (e.g., read-copy-update (RCU) [21], optimistic concurrency control [3, 4, 20]).

Data structure adjustment according to runtime workload characteristics. Unlike B-tree, whose structure is decided by the fanout, XIndex adapts its structure to runtime workload characteristics through structure update operations, such as group split and group merge. Users can configure the expected behaviors through parameters such as error bound threshold and delta index size threshold.

Implementation and evaluation with macro and micro benchmarks. We implement XIndex and compare it against state-of-the-art structures (The learned index [15], Wormhole [24], Masstree [20] and stx::Btree [1]). The benchmarks we used include different microbenchmarks, the YCSB benchmark, and the TPC-C (KV) benchmark, which is a TPC-C variant for key-value stores. The experimental results show that, with 24 CPU cores, XIndex achieves up to 3.2× and 4.4× performance improvement comparing with Masstree and Wormhole, respectively.

The rest of the paper is organized as follows: Section 2 describes the background and motivation; Section 3 gives the design of XIndex; Section 4, 5, 6 present the concurrent support, the runtime structure adjustment strategy, and optimizations accordingly; Section 7 shows the evaluation results; Section 8 provides a short discussion on alternative design choices and limitations; Section 9 summarizes related works; Section 10 concludes this paper.

2 Background & Motivation

2.1 The learned index

The insight of the learned index is that range indexes can be viewed as functions mapping keys to data positions. For fixed-length key-value records, assuming they are sorted in an array, this function is effectively a cumulative distribution function (CDF) of keys’ distribution. Given the CDF F , the position of a record is $\lfloor F(key) \times N \rfloor$, where N is the total number of records.

The core idea behind the learned index is to approximate the CDF with machine learning models, such as deep neural networks, and predict record positions using models. In order to provide the correctness guarantee despite prediction errors, the learned index stores the maximal and minimal prediction errors of the model. After training the model, the learned index calculates the errors by taking the difference between the actual and predicted positions of each key and stores the maximum and minimum. For a record in the sorted array, its actual position must fall in $[\text{pred}(key) +$

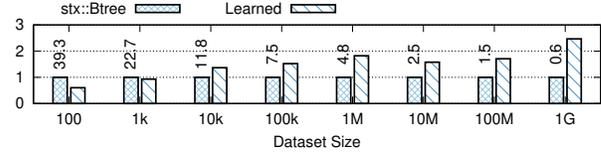


Figure 1. The learned index throughput **normalized to stx::Btree**. The numbers are stx::Btree’s absolute throughputs in MOPS. The learned index uses 10k models in 2-staged RMI where both stages use linear models. stx::Btree uses the default fanout, 16.

$\min_err, \text{pred}(key) + \max_err]$, where $\text{pred}(key)$ is the predicted position. Therefore, the learned index uses binary search within the range to locate the record. We use error bound, $\log_2(\max_err - \min_err + 1)$, to express the cost of lookup. The learned index will be more effective with a smaller error bound. In contrast with common machine learning practices where model generalization matters, the learned index expects the model to overfit to reduce errors over existing data.

However, using a single model to learn the entire CDF falls short in prediction accuracy due to the complexity of CDFs. To improve the prediction accuracy and reduce the error bound, the paper proposes a *staged model architecture*, termed Recursive Model Indexes (RMI). RMI contains multiple stages of models, which resemble the multi-level structure of B-tree. The model at each internal stage predicts which model to be activated at the next stage; the model in the leaf stage directly predicts the CDF values. With RMI architecture, each leaf stage model only approximates a small part of the complete CDF, a much simpler function to learn, which in turn reduces the model error bounds.

We evaluated the learned index with different dataset sizes under a *normal* distribution and compared it against stx::Btree [1], an open-sourced B-tree (Figure 1). The learned index can outperform stx::Btree with large datasets ($\geq 10k$) due to small binary search costs, while with small datasets, its performance is limited by the model computation cost. For example, when the dataset size is 100, the learned index spends much time on model computation (20 ns out of 42 ns). In contrast, stx::Btree only needs 25 ns to traverse two nodes for each query. When the dataset size increases, the learned index’s binary search cost increases much slower than stx::Btree’s query time and its model computation cost is constant. For example, when dataset size increases from 1M to 10M, the learned index’s binary search time only grows 37% (68 ns to 94 ns) and the error bound increases from 4.7 to 6.6. However, stx::Btree’s query time increases by 92% (207 ns to 399ns). Therefore, the learned index has better performance than stx::Btree with large datasets.

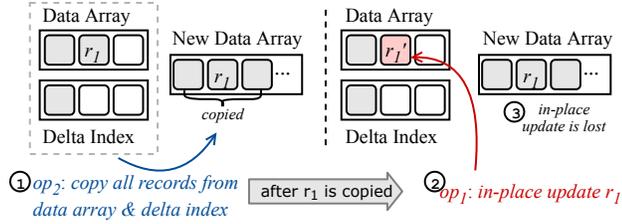


Figure 2. Consistency issue with concurrent operations.

2.2 The issues

Despite of the performance advantage of the learned index, there are two issues that limit its practicability.

First, the learned index does not provide an efficient method to handle writes, especially under concurrent scenarios. Based on the current design, an intuitive solution is to buffer all writes in a delta index, then periodically compact it with the learned index. The compaction includes merging the data into a new sorted data array and retraining the models. Though straightforward, this method suffers from severe slowdown for queries. One reason is that each request has to first go through the delta index before looking up the learned index. Considering building the learned index with 200M records, and using Masstree to be the delta index, with a workload of 10% writes, the query latency increases from 530ns to 1557ns due to the cost of searching Masstree. Another reason is that concurrent requests are blocked by the compaction, which is time-consuming. It takes up to 30 seconds to compact a delta index of 100k records with the learned index with 200M records.

A possible improvement for the above method is performing updates in-place with a non-blocking compaction scheme. When we perform updates to existing records in-place, then only newly inserted records are in the delta index. Thus, a query can only lookup the delta index when it fails to find a matching record in the learned index. Meanwhile, to avoid blocking query requests, we can compact the data asynchronously with background threads. However, the correctness issue arises if we simply use these two methods together – the effect of updates might be lost due to the data race with background compaction. Let us consider this example (Figure 2), where operation op_1 updates record r_1 in-place and operation op_2 concurrently merges the delta index with the learned index into a new data array. With the following interleaving, op_1 's update to r_1 will be lost due to the concurrent compaction: 1) op_2 starts the compaction and copies r_1 to the new array; 2) op_1 updates r_1 in the old array; 3) op_2 finishes the compaction, updates the data array, and retrains the model.

Second, the learned index's performance is tied closely to workload characteristics, including both data and query distributions. This is because the lookup efficiency depends on the error bounds of specific leaf stage models activated

Systems	Workloads			
	Skewed 1	Skewed 2	Skewed 3	Uniform
stx::Btree	1.84	1.86	1.83	1.16
learned index	1.57	3.71	1.41	2.38
Error bound	15.71	5.87	19.52	6.95

Table 1. Performance of stx::Btree and the learned index under different query distributions on the *osm* dataset. Throughputs are shown in MOPS. Error bound refers to the average error bound weighted by models' access frequencies.

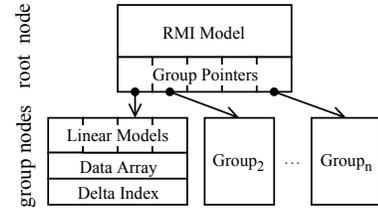


Figure 3. The architecture of XIndex.

for the queries. Meanwhile, the error bounds of different models vary. As a result, the learned index can have worse performance than B-tree with certain workloads. Table 1 shows the performance of the learned index and stx::Btree under both uniform and skewed query distributions on the *osm* dataset (details in Section 7). Under the uniform query distribution, all keys have the same chance to be accessed. Under the skewed query distribution, 95% queries access 5% hot records, and the hot records of each workload reside in different ranges. “Skewed 1” chooses hot keys from the 94th to 99th percentiles of the sorted data array. “Skewed 2” chooses from the 35th to 40th, while “Skewed 3” chooses from the 95th to 100th.

The learned index has better performance than stx::Btree under the workloads of “Skewed 2” and “Uniform”, but is outperformed under “Skewed 1” and “Skewed 3”. This is because under workload “Skewed 1” and “Skewed 3”, the learned index has much higher average error bounds on the frequently accessed records, which hinders the query performance. The underlying cause is that the learned index only minimizes each model's error individually, lacking the consideration for model accuracy differences. Similar results can be observed in other workloads as well (Section 7.3).

3 XIndex Data Structure

3.1 Overview

XIndex adopts a two-layer architecture design (Figure 3). The top layer contains a *root* node which indexes all *group* nodes in the bottom layer. The data is divided into groups by range partitioning. The root node uses a learned RMI model to index the groups. Each group node uses learned linear models to index its data. For writes, XIndex performs

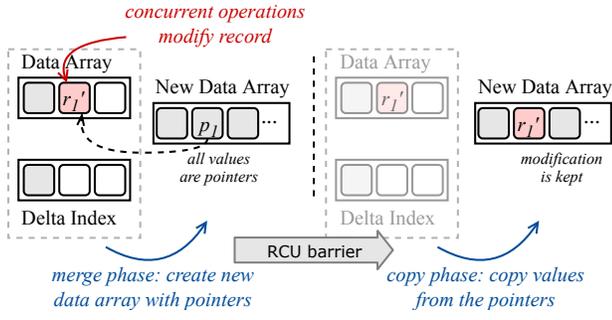


Figure 4. Two-phase compaction prevents concurrent operations being lost.

updates in-place on existing records, and associates each group with a delta index to buffer insertions.

XIndex introduces a new compaction scheme, Two-Phase Compaction (Section 3.4), to compact the delta index conditionally. The compaction is performed in the background and does not block any concurrent operations. The compaction has two phases: the *merge phase* and the *copy phase*. In the merge phase, XIndex merges the current data array and delta index into a new data array. Instead of directly copying the data, XIndex maintains data references in the new data array. Each reference points to records being compacted, residing in either the old data array or the delta index. After ensuring no accesses on the old data array through an RCU barrier, XIndex performs the copy phase. It replaces each reference in the new data array with the real value. Considering previous example (Figure 2) with Two-Phase Compaction in Figure 4, after the merge phase, the new data array contains references (e.g., p_1) to each record (e.g., r_1). If there is a concurrent writer which updates r_1 to r_1' , the writer can safely proceed as the record is already referenced in the new data array. After an RCU barrier, no thread will access the old data array anymore. XIndex replaces p_1 with r_1' in the copy phase.

XIndex is able to adjust its structure according to runtime workload characteristics (Section 5). At runtime, if some group incurs high prediction error, XIndex adds more linear models in that group with “*model split*” to improve the inference accuracy. If a group has too many models or its delta index is too large, XIndex performs *group split* — replacing the group with two new groups, each containing the half data of the old group. XIndex also performs *model merge* and *group merge*, if the merging does not affect the prediction accuracy. Furthermore, if there are too many groups, XIndex retrains the RMI model of the root node and may adjust its structure to improve the accuracy.

Algorithm 1: Structures

1 struct root_t :	13 struct group_t :
2 rmi_t rmi;	14 key_t pivot;
3 uint32_t group_n;	15 bool_t buf_frozen;
4 key_t pivots[];	16 uint16_t model_n;
5 group_t* groups[];	17 uint32_t array_size;
6	18 model_t
7 struct record_t :	19 models[MAX_MODEL_N];
8 key_t key;	20 record_t data_array[];
9 val_t val;	21 buffer_t* buf;
10 uint64_t // composite 8B	22 buffer_t* tmp_buf;
11 is_ptr : 1, removed : 1	23 group_t* next;
12 lock : 1, version : 61;	

3.2 Layout

XIndex maintains three basic structures — *record_t*, *root_t*, and *group_t*, for the record, the root node, and the group node, respectively (Algorithm 1).

The *record_t* is the basic representation of the data. It includes the key (*key*), the record data (*val*), and some metadata. The *is_ptr* flag indicates whether *val* is the actual value or a memory reference. The *removed* flag is set when a record is logically removed. The *lock* and *version* are concurrency control information, which ensures execution exclusiveness of concurrent operations.

The *root_t* contains the groups’ information and an RMI model. The group information includes each group’s address (*groups*), their smallest keys (*pivots*), and the total number of groups (*group_n*). The RMI model (*rmi*) is used to predict the group with a given key. It is trained with elements in *pivots* and their indexes, $\{(pivots[i], i) \mid i = 0, \dots, group_n - 1\}$. In the current design, XIndex uses a two-stage RMI architecture solely consisting of linear models. The number of models in its second stage is adjustable at runtime (Section 5).

The *group_t* has three basic components: the data, the models, and the delta index. For the data, all records indexed by the group is continuously stored in *data_array*. Each group uses at least one linear model to index the record in *data_array*, and the models are maintained in *models*. The *model_t* includes parameters of the linear model and the smallest key of the model’s belonging data range. The *buf* is the delta index, which buffers all insertions. During compaction, *buf_frozen* is set to be true and *buf* is frozen. The *tmp_buf* serves as a temporary delta index, which buffers all insertions temporarily during the compaction. The *next* pointer is used by group split operation (Section 3.5). For optimization purposes, the *group_t* maintains its smallest key in a separate variable, *pivot*.

3.3 Basic operations

XIndex provides basic index interfaces — *get*, *put*, *remove*, and *scan* (Algorithm 2). All operations first use the root to

Algorithm 2: Get and Put

```

1 get(root, key):
2   group ← get_group(root, key)
3   pos ← get_position(group, key)
4   val ← EMPTY
5   if pos ≠ EMPTY
6     val ← read_record(group.data_array[pos])
7   if val = EMPTY
8     val ← get_from_buffer(group.buf, key)
9   if val = EMPTY ∧ group.tmp_buf ≠ NULL
10    val ← get_from_buffer(group.tmp_buf, key)
11   return val
12
13 put(root, key, val):
14  retry:
15   group ← get_group(root, key)
16   pos ← get_position(group, key)
17   if pos ≠ EMPTY
18     if update_record(group.data_array[pos], val) = TRUE
19       return
20   if group.buf_frozen = FALSE
21     insert_buffer(group.buf, key, val)
22   else
23     if update_in_buffer(group.buf, key, val) = TRUE
24       return
25   if group.tmp_buf = NULL
26     goto retry
27   insert_buffer(group.tmp_buf, key, val)

```

find the corresponding group (Lines 2 and 15), then look up the position of the requested record in *data_array* with the given key (Lines 3 and 16). After then, their procedures diverge.

To find the corresponding group (*get_group*), XIndex first predicts a group number with the RMI model in the root node (*root.rmi*). Then, it corrects the group number with binary search the *root.pivots* within an error-bounded range. After finding a candidate group, it needs to check the group's *next* pointer further. If the pointer is not NULL, it follows the pointer to find the corresponding group by comparing *group.pivot* with the target key. Checking the *next* is necessary, this is because some newly created group may be linked to a group's *next*, and not indexed by the root yet (Section 3.5).

After finding the group *group*, XIndex tries to look up the record within its *data_array*. It first finds the correct linear model for the prediction. It scans the *group.models* and uses the first model whose smallest key is not larger than the target key. Then, it uses the model to predict a position in the *group.data_array*. Last, it corrects the position with binary search in a range bounded by the model's error.

Algorithm 3: Two-Phase Compaction

```

1 compact(group):
2   /* phase 1 */
3   group.buf_frozen ← TRUE
4   rcu_barrier()
5   group.tmp_buf ← allocate new delta index
6   new_group ← allocate new group
7   new_group.data_array ← merge(group.data_array,
8   group.buf)
9   new_group.buf ← group.tmp_buf
10  train new_group's models with new_group.data_array
11  init new_group's other fields
12  old_group ← group
13  atomic_update_reference(group, new_group)
14  rcu_barrier()
15  /* phase 2 */
16  for each record in new_group.data_array
17    replace_pointer(record)
18  rcu_barrier()
19  reclaim old_group's memory

```

After looking up *data_array*, the procedures diverge. For *get*, if XIndex finds a record matching the requested key (Line 5) in *data_array*, then it tries to read a consistent value with helper function *read_record* (Line 6). An EMPTY result indicates a logically removed record. In this case, the *get* proceeds to search *buf* (Line 7-8), then search the temporary delta index if *tmp_buf* is not NULL (Line 9-10). A *get* request returns as soon as a non-EMPTY result is fetched, otherwise it returns EMPTY.

For *put* and *remove*, similar to *get*, if a matching record is found inside *data_array* (Line 17), XIndex first tries to update/remove the record in-place (Line 18). If XIndex cannot perform update/remove in-place, then it proceeds to operate on *buf* (Line 21) and optionally *tmp_buf* (Line 27) only if the *frozen_buf* flag is true (Line 20). For *scan*, XIndex first locates the smallest record that is \geq requested key, and then consistently reads *n* consecutive records.

We elaborate on the details of *put*, *remove*, and helper functions in conjunction with concurrent background operations in Section 4 since most subtleties stem from consistency consideration.

3.4 Compaction

To ensure consistency in face of concurrent operations (Section 2.2), XIndex divides the compaction into two phases, *merge phase* and *copy phase* (Algorithm 3).

In the merge phase, XIndex merges a group's *data_array* and *buf* into a new sorted array where values are pointers to existing records. XIndex first sets the old group's *buf_frozen* flag to stop newly issued *puts* inserting to *buf* (Line 3). Then XIndex initializes *tmp_buf* to buffer insertions during compaction (Line 5). Afterwards, it creates a new group (Line 6)

Algorithm 4: Group Split

```

1 split(group):
2   /* step 1 */
3    $g'_a, g'_b \leftarrow$  allocate 2 new group
4    $\{g'_a, g'_b\}.data\_array, buf \leftarrow group.\{data\_array, buf\}$ 
5    $g'_a.pivot \leftarrow group.pivot$ 
6    $g'_b.pivot \leftarrow group.data\_array[group.array\_size / 2]$ 
7    $g'_a.next \leftarrow g'_b$ 
8   init other fields of  $g'_a$  and  $g'_b$ 
9    $old\_group \leftarrow group$ 
10  atomic_update_reference(group,  $g'_a$ )
11   $\{g'_a, g'_b\}.buf\_frozen \leftarrow TRUE$ 
12  rcu_barrier()
13   $\{g'_a, g'_b\}.tmp\_buf \leftarrow$  allocate new delta indexes
14  /* step 2.1, merge phase */
15   $g_a, g_b \leftarrow$  allocate 2 new groups
16   $tmp\_array \leftarrow$  merge(
17     $old\_group.data\_array, old\_group.buf$ )
18   $\{g_a, g_b\}.data\_array \leftarrow$  split( $tmp\_array, g'_b.pivot$ )
19   $\{g_a, g_b\}.buf \leftarrow \{g'_a, g'_b\}.tmp\_buf$ 
20  train  $\{g_a, g_b\}$ 's models with  $\{g_a, g_b\}.data\_array$ 
21   $\{g_a, g_b\}.pivot \leftarrow \{g'_a, g'_b\}.pivot$ 
22   $g_a.next \leftarrow g_b$ 
23  init  $g_a$ 's and  $g_b$ 's other fields
24  atomic_update_reference(group,  $g_a$ )
25  rcu_barrier()
26  /* step 2.2, copy phase */
27  for each record in  $\{g_a, g_b\}.data\_array$ 
28    replace_pointer(record)
29  rcu_barrier()
30  reclaim  $\{old\_group, g'_a, g'_b\}$ 's memory

```

and merges the old group's *data_array* and *buf* into the new group's *data_array* (Line 7). In *new_group.data_array*, the value of each record is the reference to the corresponding record in either *group.data_array* or *group.buf*, and the *is_ptr* flag of each record is set to be TRUE. During merging, XIndex skips the logically removed records. The old group's *tmp_buf* is reused as the new group's *buf* (Line 8). After training linear models (Line 9) and initializing the remaining metadata of the new group (Line 10), XIndex atomically replaces the old group with the new one by changing the group reference in root's *groups* (Line 12).

In the copy phase, XIndex replaces each reference in the new group's *data_array* with the latest record value (Line 16). The replacement is performed atomically with helper function *replace_pointer* (Algorithm 5). XIndex uses *rcu_barrier* (Line 17) to wait for each worker to process one request, so the old group will not be accessed after the barrier. Then it can safely reclaim the old group's memory resources (Line 18).

3.5 Structure update

XIndex adapts its structure to dynamic workloads at runtime (Section 5) with model split/merge, group split/merge, and root update operations.

Model split/merge. XIndex supports splitting and merging models within a group to improve lookup efficiency. For model split, XIndex first clones the group node. Both group nodes reference the same *data_array* and *buf*. Then, it increments the new node's *model_n*, evenly reassigns the group's data to each model, and retrains all models. At last, XIndex atomically updates the group reference in root's *groups* to the new group. For model merge, it essentially performs a reverse procedure of model split.

Group split. To avoid blocking other operations, XIndex uses two steps to split a group's data evenly into two groups (Algorithm 4).

In step 1, XIndex creates two logical groups. They share the data and delta index, but each has its own temporary delta index. As a result, both groups can buffer the insertion in their temporary delta indexes during the split. In detail, XIndex creates g'_a and g'_b (Line 3). They share the same *data_array* and *buf* with the old group (Line 4) but have different *pivot* keys (Line 5-6). XIndex links g'_b to g'_a 's *next* field (Line 7) and replaces the old group with g'_a in root's *groups* (Line 10). Last, XIndex sets the *buf_frozen* flag (Line 11) and allocates *tmp_buf* for g'_a and g'_b (Line 13).

In step 2, XIndex physically divides the data into two groups. Similar to the compaction, this step has two phases. In the merge phase, XIndex first merges the old group's *data_array* and *buf* into *tmp_array* (Line 17). Then, it splits the *tmp_array* with the key of $g'_b.pivot$, and initializes two new groups, g_a and g_b accordingly (Line 18). It also reuses the *tmp_buf* of g'_a (g'_b) as the *buf* of g_a (g_b) [Line 19]. In the copy phase, for each group, the references in *data_array* are replaced with real values (Line 28). Last, XIndex links g_b at $g_a.next$ (Line 22) and replaces g'_a with g_a in root's *groups* (Line 24).

Group merge. XIndex merges two consecutive groups' data into one new group to reduce the cost to lookup groups. Similar to the group split, data is merged in two phases. In the merge phase, both groups' *data_arrays* and *bufs* are merged together while inserts are buffered in a single shared *tmp_buf*. In the copy phase, the merged references are replaced with concrete values. Finally, among the two consecutive groups in root's *groups*, the former is replaced with the new group and the latter is marked as NULL, which will be skipped by *get_group*. For brevity, we omit the pseudocode for group merge.

Root update. XIndex flattens root's *groups* to reduce pointer access cost, retrains, and conditionally adjusts the RMI model to improve prediction accuracy. For root update, XIndex creates a new root node with a flattened *groups* and

Algorithm 5: Helper functions

```

1 read_record(rec):
2 while TRUE
3   ver ← rec.ver
4   removed, is_ptr, val ← rec.{removed, is_ptr, val}
5   if ¬rec.lock ∧ rec.version = ver
6     if removed
7       val ← empty
8     else if is_ptr
9       val ← read_record(DEREF(val))
10    return val
11
12 update_record(rec, val):
13   lock(rec.lock)
14   succ ← FALSE
15   if rec.is_ptr
16     succ ← update_record(DEREF(rec.val), val)
17   else if ¬rec.removed
18     rec.val ← val
19     succ ← TRUE
20   rec.version++
21   unlock(rec.lock)
22   return succ
23
24 replace_pointer(rec):
25   lock(rec.lock)
26   ref.val ← read_record(DEREF(rec.val))
27   if ref.val = empty
28     rec.removed ← TRUE
29   rec.is_ptr ← FALSE
30   rec.version++
31   unlock(rec.lock)

```

retrains the RMI model. After a new root is initialized, XIndex replaces the global root pointer atomically.

4 Concurrency

XIndex achieves high scalability on the multicore platform using Two-Phase Compaction, along with classic techniques such as fine-grained locking [2, 20, 24], optimistic concurrency control [3, 4, 20], and RCU [21]. We first discuss the coordination between writers that ensures execution exclusiveness (Section 4.1), then discuss how readers can always fetch a consistent result with concurrent writers (Section 4.2). Afterwards, we discuss the interleaving with concurrent background operations (Section 4.3) and provide a proof sketch of the correctness condition (Section 4.4). The formal proof can be found in Appendix A.

For brevity, we treat *remove* as a special *put*, which updates existing records' *removed* flag. We further omit group merge and root update in the discussion, as the reasoning resembles compaction's and group split's. In XIndex, compaction and

structure updates are performed by dedicated background threads sharing no conflicts, thus avoiding concurrency issues due to their interleavings.

4.1 Writer-writer coordination

XIndex ensures that conflicting writers, *put/removes* with the same key, will execute exclusively with the per-record lock in *data_array* and the concurrent delta index. All writers first try to update a matching record in *data_array* (Line 18, Algorithm 2), and the per-record lock is acquired to prevent interleaving with concurrent writers (Line 13, Algorithm 5). If updating *data_array* is not feasible, writers then operate on the delta index (Line 21, Algorithm 2), protected by a single read-write lock in the basic version. We improve its scalability with fine-grained concurrency control as an optimization (Section 6).

4.2 Writer-reader coordination

XIndex ensures readers can always fetch a consistent result in face of concurrent writers with locks and versions in *data_array* and the concurrent delta index. A *get* first tries to read a value from *data_array* (Line 6, Algorithm 2). It snapshots the version number before reading the value (Line 3, Algorithm 5). After the value is fetched (Line 4, Algorithm 5), *get* validates if the lock is being held (to detect concurrent writer) *and* if the current version number matches the snapshot (to detect inconsistent or stale result)[Line 5, Algorithm 5]. If the validation fails, the *get* repeats the procedure until a successful validation, so the result is consistent and the latest. If reading from *data_array* is not feasible, it then tries to read from the delta index (Line 8, Algorithm 2). The concurrent delta index with a single read-write lock ensures the fetched result is consistent.

4.3 Interleaving with background operations

With the presence of background operations, XIndex ensures that the effects of writers are preserved and can always be correctly observed by readers. Space constraints preclude a full discussion, but we mention two important conditions: 1) no successful *put* will be lost, and 2) no duplicate records (records with the same key) will be created⁴.

To ensure no lost *put*, the key is to perform data movement in two phases, the merge phase and the copy phase, to preserve concurrent modifications. During the merge phase, *all* records in the old group's *data_array* and *buf* can be correctly referenced in the new group's *data_array*. This is because both the *data_array* and *buf* of the old group are read- and update-only, as the *buf_frozen* flag forbids insertions (Line 3, Algorithm 3 and Line 11, Algorithm 4). In the copy phase, those references can be atomically replaced with

⁴Duplicate records do not directly violate correctness, as long as XIndex enforces a freshness ordering, $data_array \succcurlyeq buf \succcurlyeq tmp_buf$, where *data_array* has the latest version. However, doing so requires non-trivial efforts.

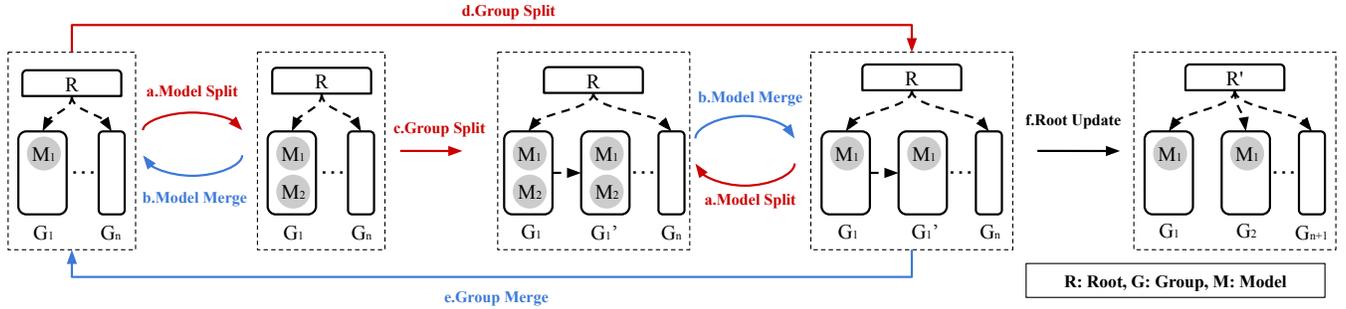


Figure 5. Dynamic adjustment procedure illustration.

latest values by *replace_pointer*, since it uses the per-record lock to coordinate with concurrent writers (Line 25, Algorithm 5). To ensure all concurrent writers use the same lock, XIndex places *rcu_barrier* before the copy phase (Line 13, Algorithm 3 and Line 25, Algorithm 4), which waits for each writer (and reader) to process one request. Therefore, later conflicting *puts* will not reference the old group’s *data_array*. In addition, concurrent inserts are preserved in the shared temporary delta index (Line 8, Algorithm 3 and Line 19, Algorithm 4).

To ensure no duplicate records, XIndex avoids insertions to different delta indexes, namely *buf* and *tmp_buf*. XIndex only initializes *tmp_buf* until all writers observe a frozen *buf* using the *rcu_barrier* (Line 3-5, Algorithm 3 and Line 11-13, Algorithm 4). Therefore, whenever *tmp_buf* is used to serve requests, the *buf* is sure to be read- and update-only.

4.4 Proof sketch

The correctness condition of XIndex can be described as “a *get(k)* must observe the latest committed *put(k, v)*”, namely, linearizability [13]. XIndex formally provides the correctness condition by ensuring the following inductive invariants. I_1) If there is a *put(k, v)* committed, then there is exactly one record with key k in XIndex; I_2) If there is a record with key k in XIndex, then its value equals the value of the last committed *put(k, v)*; and I_3) If there is a record with key k in XIndex when a *get* commits, then the *get* returns the value of the record.

For I_1 , in addition to *no duplicate records* guarantee we discussed in Section 4.3, XIndex ensures that a new record will be created by *put* if no record currently exists yet. This is obvious as such *put* will invoke *insert_buffer* (Lines 21 and 27, Algorithm 2), and the concurrent delta index will handle the record creation. For I_2 , the key is to ensure that no *put* will be lost, as we discussed in Section 4.3. For I_3 , the key is to let *get* and *put* have the same lookup order (*data_array* → *buf* → *tmp_buf*). Since only the last place (*buf* when *tmp_buf* is null, otherwise *tmp_buf*) is insertable, a *get* returning an empty result only indicates that the *put* that

Operations	Trigger Condition
a. Model Split	error bound $> e$ and # of models $< m$
b. Model Merge	error bound $\leq e \times f$ and # of models > 1
c. Group Split	error bound $> e$ and # of models $= m$
d. Group Split	buf $> s$
e. Group Merge	# of models $= 1$ and error bound $\leq e \times f$ and buf $\leq s \times f$
f. Root Update	when groups are created and/or removed

Table 2. Conditions for structure update operations.

creates the record has not yet finished. Therefore, a *get* can fetch the value correctly.

5 Adjusting XIndex at Runtime

To reduce the performance variation, XIndex adjusts its structure according to runtime workload characteristics. The basic idea is to keep both error bound and delta index size small with structure update operations (model split/merge, group split/merge, and root update). Several background threads periodically check error bound and delta index size of each group and perform corresponding operations accordingly (Table 2 and Figure 5).

First, XIndex leverages model split to lower the error bound and model merge to reduce the cost of traversing a group’s *models* array. Specifically, when a model’s error bound is greater than the error bound threshold (e , specified by the user) and the model number of the corresponding group is less than the model number threshold (m , specified by the user), XIndex will do model split (Figure 5-a). When a model’s error bound is less than or equal to $e \times f$ and the model number of the corresponding group is greater than one, XIndex will perform model merge (Figure 5-b). $f \in (0, 1)$ is a tolerance factor specified by the user.

When a model’s error bound is greater than e , but the model number of the corresponding group equals m , XIndex will perform group split (Figure 5-c). Besides, if a group’s delta index size is greater than the delta index size threshold (s , specified by the user), XIndex will also split the group (Figure 5-d). To reduce the cost of locating a group in the

root, XIndex performs group merge (Figure 5-e) when the following conditions hold – for two neighboring groups, 1) they both only have one model and the model’s error bounds are less than or equal to $e \times f$; and 2) their delta index sizes are both less than or equal to $s \times f$.

XIndex periodically updates the root to reduce the access cost. Specifically, XIndex first checks all groups and perform model split/merge and group split/merge accordingly. If there is any group created or removed, XIndex then performs root update (Figure 5-f). During root update, if the average error bound is greater than e , XIndex will increase the number of 2nd stage models of root’s 2-stage RMI⁵. If the average error bound is less than or equal to $e \times f$, XIndex reduces models.

6 Optimization

Scalable delta index. In the basic version, XIndex uses `stx::Btree` protected by a global read-write lock as its delta index. This limits the scalability when concurrent writers insert records to the same group. One possible solution is directly using Masstree as the delta index. However, Masstree provides unnecessary functionalities such as supporting variable length, multi-column values, and epoch-based memory reclamation. Thus, we implement a scalable delta index with a simplified design – each index node has a version to ensure that a *get* request can always fetch consistent content of the node and a lock to protect node update and split.

Sequential insertion. Sequential insertion is a common pattern in real-world workloads, such as periodically checkpointing. For such cases, the user can provide hints to XIndex so that XIndex can pre-allocate space to allow appending records directly to *data_array* and conditionally retrain models. Specifically, each group maintains an additional *capacity* field and a per-group lock. Only when XIndex detects the sequential insertion pattern, will it use the lock to coordinate concurrent sequential insertions. Otherwise, the lock is not used, so the scalability of XIndex is intact. Since many sequential insertion workloads have relatively static data distribution, XIndex only retrains models when the current model cannot generalize to newly appended data, namely, when the error bound exceeds the threshold.

7 Evaluation

We evaluate XIndex with complex workloads as well as micro-benchmarks of different characteristics and compare it against state-of-the-art systems.

Benchmarks. We develop a TPC-C (KV) benchmark by implementing TPC-C benchmark with only *get* and *put* operations, which is the same as [20]. We assign 8 distinct warehouses to each thread as their local warehouses for evaluation. Since XIndex does not support transactions, in order to avoid the impact of transaction abortions due to conflicts, we eliminate the conflicts by manipulating each thread to

⁵The number of models stops increasing when it reaches a given limit.

Name	Description
linear	Linear dataset with added noises
normal	Normal distribution ($\mu = 0, \sigma = 1$), scaled to $[0, 1 \times 10^{12}]$
lognormal	Lognormal distribution ($\mu = 0, \sigma = 2$), scaled to $[0, 1 \times 10^{12}]$
osm	Longitude values of OpenStreetMap locations scaled to $[0, 3.6 \times 10^9]$

Table 3. Datasets. For the *linear* dataset, we first generate keys $\{i \times A \mid i = 1, 2, \dots\}$, then add a uniform random bias ranging in $[-A/2, A/2]$ for each key, where $A = 1 \times 10^{14} / \text{dataset size}$. All keys are integers.

execute remote transactions on one of their own local warehouses. TPC-C (KV) benchmark can evaluate index systems under data and query distribution of real-world database workload while not requiring transaction support. YCSB includes six representative workloads (A-F) with different access patterns: update heavy (A), read mostly (B), read-only (C), read latest (D), short ranges (E) and, read-modify-write (F). For YCSB, besides its default data distribution, we also evaluate with a real-world dataset OpenStreetMap [6]. For microbenchmarks, we evaluate the performance under workloads with fixed read-write ratios (Section 7.2), under dynamic workloads (Section 7.3). We also analyze different factors that affect the performance in Section 7.4. All datasets used are listed in Table 3. The default dataset size is 200M unless otherwise noted, and each record has 8 bytes key and 8 bytes value by default.

Counterparts. `stx::Btree` [1] is an efficient, but thread-unsafe B-tree implementation. Masstree [20] is a concurrent index structure that hybrids B-tree and Trie. When the key size is 8 bytes, Masstree can be regarded as a scalable concurrent B-tree. Wormhole [24] is a concurrent hybrid index structure that replaces B-tree’s inner nodes with a hash-table encoded Trie. The learned index [15] is the original learned index. “learned+ Δ ” is the learned index attached with a Masstree as delta index, which buffers all writes.

Configuration & Testbed. We implement XIndex in C++, and configure 1 out 12 threads as dedicated background threads. Background thread(s) sleeps one second after it has checked all groups and root to perform compaction and structure update accordingly. Throughout all experiments, the error bound threshold (e) is 32, the delta index size threshold (s) is 256, the tolerance factor (f) is $\frac{1}{4}$, and the model number threshold (m) is 4. For the learned index, we test different configurations and pick the best one – 250k models in the 2nd stage.⁶ For “learned+ Δ ”, we use the same background threads as XIndex for compaction. For `stx::Btree`, Masstree, and Wormhole, we directly run their source code with the

⁶The candidates’ model number ranges from 50k to 500k (step is 50k).

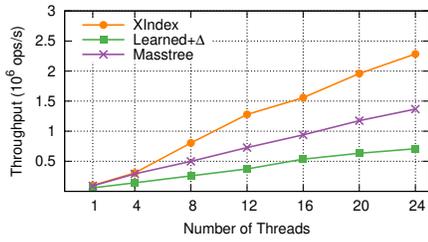


Figure 6. TPC-C (KV) throughput.

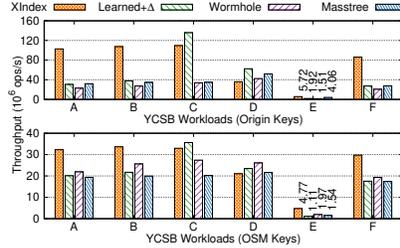


Figure 7. YCSB throughput.

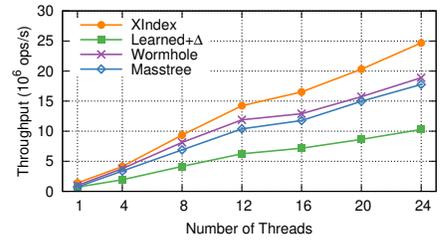


Figure 8. Read-write throughput.

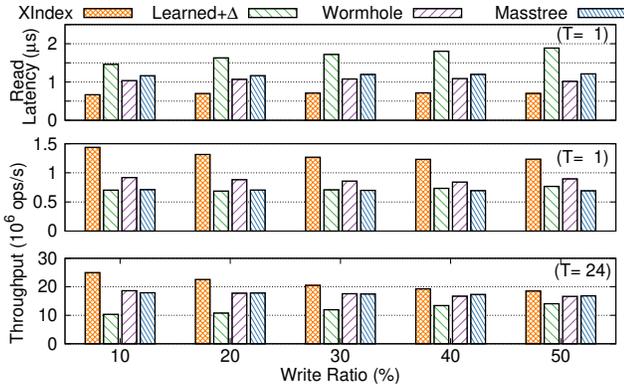


Figure 9. Read-write throughput and read latency with the normal dataset. T indicates the number of threads.

default setting. For each experiment, we first warmup all the systems and present steady-state results. The experiments run on a server with two Intel Xeon E5-2650 v4 CPUs, and each CPU has 12 cores. The hyperthreading is disabled during evaluation.

7.1 Performance Overview

TPC-C (KV). Figure 6 shows the performance comparison with different numbers of threads. Wormhole is excluded because the Wormhole implementation we use does not support multiple tables, while TPC-C (KV) requires them. XIndex outperforms Masstree by up to 67% with 24 threads. First, the data generated in TPC-C (KV) are multidimensional linear mappings. Therefore, the learned models can obtain a good approximation. Second, 63% of the write operations update existing records. Thus they can be efficiently executed in-place. Lastly, 34% of the write operations perform sequential insertion, which can be improved by the optimization (Section 6).

YCSB. We use both the default data distribution as well as the *osm* dataset, and 24 threads for the experiment. As shown in Figure 7, for workload A, B, E, and F, XIndex demonstrates superior performance advantage. This is because these workloads are read- and update-intensive. For workload C, which is read-only, XIndex is worse than “learned+Δ” by 19% because XIndex has model computation cost both in the root

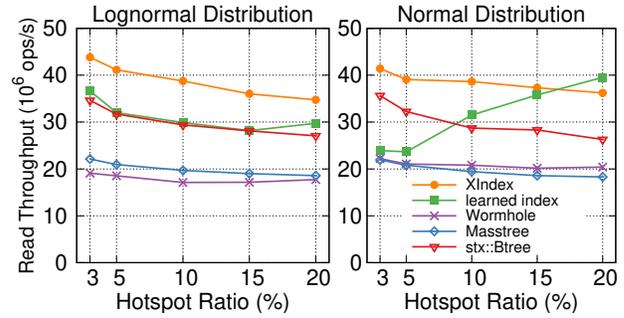


Figure 10. 24-thread read throughput in skewed query distribution.

and groups. For workload D, XIndex performance is worse than the other systems by up to 30%. The reason is that workload D tends to read recently inserted records that might not have been compacted, which brings overheads for read operations. With the *osm* dataset, the results are similar. However, because of the complex real-world data distribution, the advantage of XIndex is reduced.

7.2 Performance with writes

To further evaluate the performance of writes, we configure workloads with different read-write ratios. The ratio among different type of writes are constant (1:1:2 for insert, remove, and update) to keep the dataset size stable.

Scalability. Figure 8 shows the scalability with 10% writes using the *normal* dataset. Overall, XIndex achieves the best performance among all systems. With 24 threads, XIndex scales to 17.6× of its single-thread performance, which is 30% higher than Wormhole. “learned+Δ” has the worst performance because of its inefficient compaction, which severely degrades the read performance.

Varying write ratios. Figure 9 shows both throughput and latency with different write ratios with a single thread and 24 threads. XIndex has the best performance for all the listed write ratios, though the advantage tends to diminish with larger write ratios. For latency, XIndex achieves the lowest latency as most requests (80%) can be served without accessing delta indexes.

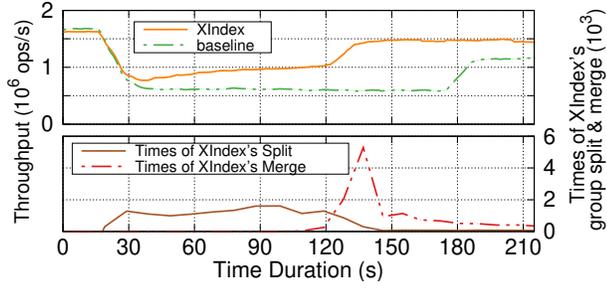


Figure 11. Read-write throughput and group split/merge frequency under dynamic workload.

7.3 Performance of dynamic workload

Query distribution. For dynamic workload, we first evaluate the performance when the query distribution is non-uniform. To control the skewness, we make the workload’s 90% queries access hotspot of different sizes (the hotspot ratio). All hotspots are ranges that start from the same key but end differently. The smaller the hotspot is, the more skewed the query distribution is. Figure 10 shows the throughput with different skewness levels under the *normal* and *lognormal* datasets. All systems except for the learned index see a performance improvement when the skewness level rises since the skewed query distribution brings a more friendly memory access locality. However, due to the learned index’s large error bound in the hotspot, the learned index can perform even worse than *stx::Btree* and *Wormhole*. For the *lognormal* dataset, when the hotspot ratio decreases under 5%, the increase of hot models’ error bounds slows down, thus we can observe a slight performance improvement of the learned index due to improved locality.

Data distribution. We then evaluate XIndex under the workload, whose data distribution and read-write ratio will be dynamically changed at runtime. As a baseline, we also run XIndex without background group split and merge. Figure 11 shows the throughput and the number of XIndex’s group split/merge under this workload using one worker thread and one background thread.

Both XIndex and baseline are initialized with 50M *normal* dataset, and the initial read-write ratio is 90:10. In the beginning, they have similar performance. At the 20th second, the write ratio becomes 100% (half inserts and half removes), and we remove all existing keys and insert new keys with 50M *linear* dataset. From this point, both throughput of XIndex and baseline begin to degrade due to the increase of write ratio and the dramatical changes of data distribution. While for XIndex, background threads begin to do group split to reduce the error of the group and delta index size, so we can see the throughput starts to increase at the 30th second.

At the 120th second, XIndex finishes dataset shifting, and at the 170th second, baseline ends shifting. Afterwards, the

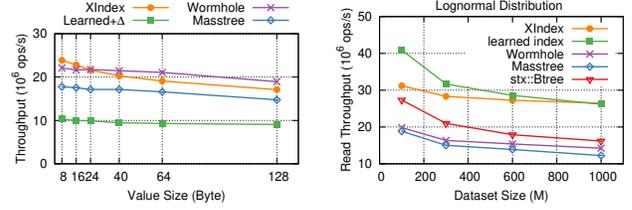


Figure 12. Read-write throughput of various value sizes.

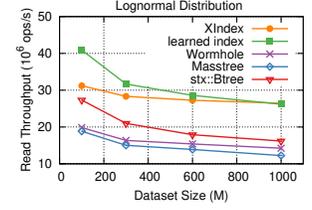


Figure 13. Read throughput of various dataset sizes.

read-write ratio is 90:10, and keys follow the *linear* distribution. XIndex’s background thread detects that both the delta index size and the error bound of groups are small after the shifting, so it invokes lots of group merge operations to reduce the number of groups. Overall, XIndex shows up to 140% performance improvement during and after the change of workload.

7.4 Other factors

Value size. We evaluate the performance of XIndex with different value sizes under the *normal* dataset with 24 threads. The read-write ratio is 90:10, and the value contains 8-128 random generated bytes. The result is shown in Figure 12. With the increase of value size, the performance of all systems is reduced due to the large memory consumption. Nevertheless, XIndex has the largest performance drop. This is because the overhead of data copying during compaction (128B’s overhead is 13.5× larger than 8B’s).

Dataset size. Figure 13 shows the performance of XIndex with different dataset sizes under the *lognormal* dataset using 24 threads. As dataset size increases, both the learned index and XIndex show a large performance advantage over other systems. However, the performance of the learned index degrades significantly because its error grows as the size increases. In contrast, XIndex adjusts its structure to maintain small model error bounds. Therefore, for large dataset sizes, XIndex can achieve similar performance with the learned index.

8 Discussion

Inline values v.s. separated values. XIndex contiguously stores keys and values in *data_array* (inline values). Another popular approach [19, 22] is to store values in separate storage and only the pointers along with keys (separated values). For small values, our approach has an advantage in reducing DRAM accesses, since both the key and value can reside in one cache line. For large values, separating value can reduce compaction cost, since only pointers will be copied. However, both approaches require Two-Phase Compaction as the record’s metadata, such as *removed* and *lock*, should be inline and might be changed during compaction. As separating the metadata may incur high compaction cost, since it needs

additional memory accesses for each record to check if the value has been logically removed.

Limitations. First, when the dataset is small, other index structures (e.g., `stx::Btree`) can outperform XIndex for the model computation cost and the cost to traverse the two-layer structure. Second, with long keys, the overhead of model training and inference will increase significantly, which can affect the efficiency of XIndex. For 64-byte keys, there can be up to 50% performance degradation compared with 8-byte keys. We leave the design of a more flexible structure for variable-scale workloads and reducing the cost of long keys as future work.

9 Related Works

There have been works that extend or build systems upon the learned index. Many works extend the learned index to support writes. ALEX [9] achieves this by reserving slots in a sorted array for inserting new data. It allocates a new array and retrains the model synchronously when there are not enough reserved slots. AIDEL [18] handles insertions by attaching a sorted list for each record in the sorted array. When the list is too long, it copies the data into a new sorted array and retrains the model synchronously. Both data structures are not designed for concurrent scenarios and operations are blocked during rearranging data and retraining models. PGM-index [11] extends the learned index to optimize the structure with respect to given space-time trade-offs. It recursively constructs the index structure and provides an optimal number of linear models. Comparing with PGM-index, XIndex adjusts its structure at runtime, does not assume an already known query distribution. SageDB [14] is a database that proposes to leverage the learned index for data indexing as well as for speeding up sorting and join. FITing-Tree [12] indexes data with a hybrid of B-tree and piece-wise linear function, making it a variant of the learned index. It supports insertions and provides strict error guarantees. Comparing with FITing-Tree, XIndex is a fully-fledged concurrent index structure and adapts its structure to both data and query distribution at runtime.

Classic concurrency techniques have long been used in concurrent data structures. Masstree [20] is a trie-like concatenation of B-trees and uses fine-grained locking and optimistic concurrency control to achieve high performance under multi-core scenarios. It carefully crafts its protocol to improve efficiency for reader-writer coordination. Wormhole [24] is a variant of B-tree that replaces B-tree's inner nodes with a hash-table encoded Trie. It uses per-node read-write locks to coordinate accesses to leaf nodes and uses a combination of locking and the RCU mechanism to perform internal node updates. Bonsai tree [5] is a concurrent balanced tree. It allows reads to proceed without locks in parallel with writes by using RCU mechanism, though a single write lock is still required to coordinate writes. HOT [2]

is a trie-based index structure which aiming to reduce the height of the trie. It uses per-node locks to coordinate writes and uses copy-on-write (COW) to allow reads to proceed with no synchronization overhead. The Bw-Tree [8, 17] is a completely lock-free B-tree and achieves its lock-freedom via COW and compare-and-swap (CAS) techniques. Building upon existing works, XIndex leverages fine-grained locking and optimistic concurrency control to coordinate to individual records and uses the RCU mechanism to eliminate interference with queries and writes due to background compaction and structures updates.

Dynamic data and query distributions are common in real-world workloads. While XIndex strikes to reduce performance variation between records, many works distinguish hot and cold data and further optimize the performance for hot data. Hybrid index structure [25] uses different storage schemes for hot keys and cold keys. Storage systems such as H-Store [7], COLT [23] are designed to detect the hotness and manage data accordingly in a self-tuning process.

10 Conclusion

In this paper, we introduced XIndex, a concurrent and flexible index structure based on the learned index. XIndex achieves high performance on the multicore platform via a combination of the innovative Two-Phase Compaction and a number of classical concurrency techniques. Furthermore, it can dynamically adjust its structure according to the runtime workloads to maintain competitive performance. Extensive evaluations demonstrate that XIndex can have a performance advantage by up to 3.2× and 4.4×, compared with Masstree and Wormhole, respectively. XIndex is publicly available at <https://ipads.se.sjtu.edu.cn:1312/opensource/xindex.git>.

Acknowledgments

We thank the authors of [15], Jinyang Li, Yueying Li, and the anonymous reviewers for their constructive feedback and suggestions. This work is supported in part by the National Natural Science Foundation of China (No. 61902242, 61925206), the HighTech Support Program from Shanghai Committee of Science and Technology (No. 19511121100), and a grant from Huawei Technologies. Zhaoguo Wang (zhaoguowang@sjtu.edu.cn) is the corresponding author.

References

- [1] Timo Bingmann. 2008. *STX B+ Tree C++ Template Classes*. <http://panthema.net/2007/stx-btree>
- [2] Robert Binna, Eva Zangerle, Martin Pichl, Günther Specht, and Viktor Leis. 2018. HOT: A Height Optimized Trie Index for Main-Memory Database Systems. In *Proceedings of the 2018 International Conference on Management of Data* (Houston, TX, USA) (*SIGMOD '18*). Association for Computing Machinery, New York, NY, USA, 521–534. <https://doi.org/10.1145/3183713.3196896>

- [3] Nathan G. Bronson, Jared Casper, Hassan Chafi, and Kunle Olukotun. 2010. A Practical Concurrent Binary Search Tree. In *Proceedings of the 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Bangalore, India) (PPoPP '10). Association for Computing Machinery, New York, NY, USA, 257–268. <https://doi.org/10.1145/1693453.1693488>
- [4] Sang K. Cha, Sangyong Hwang, Kihong Kim, and Keunjoo Kwon. 2001. Cache-Conscious Concurrency Control of Main-Memory Indexes on Shared-Memory Multiprocessor Systems. In *Proceedings of the 27th International Conference on Very Large Data Bases (VLDB '01)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 181–190.
- [5] Austin T. Clements, M. Frans Kaashoek, and Nickolai Zeldovich. 2012. Scalable Address Spaces Using RCU Balanced Trees. In *Proceedings of the Seventeenth International Conference on Architectural Support for Programming Languages and Operating Systems* (London, England, UK) (ASPLOS XVII). Association for Computing Machinery, New York, NY, USA, 199–210. <https://doi.org/10.1145/2150976.2150998>
- [6] OpenStreetMap contributors. 2019. *OpenStreetMap*. <https://aws.amazon.com/public-datasets/osm> Accessed: 2019-4-24.
- [7] Justin DeBrabant, Andrew Pavlo, Stephen Tu, Michael Stonebraker, and Stan Zdonik. 2013. Anti-Caching: A New Approach to Database Management System Architecture. *Proc. VLDB Endow.* 6, 14 (sep 2013), 1942–1953. <https://doi.org/10.14778/2556549.2556575>
- [8] Cristian Diaconu, Craig Freedman, Erik Ismert, Per-Ake Larson, Pravin Mittal, Ryan Stonecipher, Nitin Verma, and Mike Zwilling. 2013. Hekaton: SQL Server's Memory-Optimized OLTP Engine. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data* (New York, New York, USA) (SIGMOD '13). Association for Computing Machinery, New York, NY, USA, 1243–1254. <https://doi.org/10.1145/2463676.2463710>
- [9] Jialin Ding, Umar Farooq Minhas, Jia Yu, Chi Wang, Jaeyoung Do, Yinan Li, Hantian Zhang, Badrish Chandramouli, Johannes Gehrke, Donald Kossmann, David Lomet, and Tim Kraska. 2020. ALEX: An Updatable Adaptive Learned Index. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data* (Portland, OR, USA) (SIGMOD '20). Association for Computing Machinery, New York, NY, USA, 969–984. <https://doi.org/10.1145/3318464.3389711>
- [10] Ahmed Eldawy, Justin Levandoski, and Per-Åke Larson. 2014. Trekking through Siberia: Managing Cold Data in a Memory-Optimized Database. *Proc. VLDB Endow.* 7, 11 (jul 2014), 931–942. <https://doi.org/10.14778/2732967.2732968>
- [11] Paolo Ferragina and Giorgio Vinciguerra. 2020. The PGM-Index: A Fully-Dynamic Compressed Learned Index with Provable Worst-Case Bounds. *Proc. VLDB Endow.* 13, 8 (apr 2020), 1162–1175. <https://doi.org/10.14778/3389133.3389135>
- [12] Alex Galakatos, Michael Markovitch, Carsten Binnig, Rodrigo Fonseca, and Tim Kraska. 2019. FITing-Tree: A Data-Aware Index Structure. In *Proceedings of the 2019 International Conference on Management of Data* (Amsterdam, Netherlands) (SIGMOD '19). Association for Computing Machinery, New York, NY, USA, 1189–1206. <https://doi.org/10.1145/3299869.3319860>
- [13] Maurice P. Herlihy and Jeannette M. Wing. 1990. Linearizability: A Correctness Condition for Concurrent Objects. 12, 3 (jul 1990), 463–492. <https://doi.org/10.1145/78969.78972>
- [14] Tim Kraska, Mohammad Alizadeh, Alex Beutel, Ed H. Chi, Ani Kristo, Guillaume Leclerc, Samuel Madden, Hongzi Mao, and Vikram Nathan. 2019. SageDB: A Learned Database System. In *Proceedings of the 9th Biennial Conference on Innovative Data Systems Research* (Asilomar, CA, USA) (CIDR '19). <http://cidrdb.org/cidr2019/papers/p117-kraska-cidr19.pdf>
- [15] Tim Kraska, Alex Beutel, Ed H. Chi, Jeffrey Dean, and Neoklis Polyzotis. 2018. The Case for Learned Index Structures. In *Proceedings of the 2018 International Conference on Management of Data* (Houston, TX, USA) (SIGMOD '18). Association for Computing Machinery, New York, NY, USA, 489–504. <https://doi.org/10.1145/3183713.3196909>
- [16] Justin J. Levandoski, Per-Åke Larson, and Radu Stoica. 2013. Identifying hot and cold data in main-memory databases. In *Proceedings of the 29th IEEE International Conference on Data Engineering* (Brisbane, Australia) (ICDE '13), Christian S. Jensen, Christopher M. Jermaine, and Xiaofang Zhou (Eds.). IEEE Computer Society, 26–37. <https://doi.org/10.1109/ICDE.2013.6544811>
- [17] Justin J. Levandoski, David B. Lomet, and Sudipta Sengupta. 2013. The Bw-Tree: A B-tree for new hardware platforms. In *Proceedings of the 29th IEEE International Conference on Data Engineering* (Brisbane, Australia) (ICDE '13), Christian S. Jensen, Christopher M. Jermaine, and Xiaofang Zhou (Eds.). IEEE Computer Society, 302–313. <https://doi.org/10.1109/ICDE.2013.6544834>
- [18] Pengfei Li, Yu Hua, Pengfei Zuo, and Jingnan Jia. 2019. A Scalable Learned Index Scheme in Storage Systems. *CoRR* (2019). arXiv:1905.06256
- [19] Lanyue Lu, Thanumalayan Sankaranarayanan Pillai, Hariharan Gopalakrishnan, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. 2017. WiscKey: Separating Keys from Values in SSD-Conscious Storage. *ACM Trans. Storage* 13, 1, Article 5 (mar 2017), 28 pages. <https://doi.org/10.1145/3033273>
- [20] Yandong Mao, Eddie Kohler, and Robert Tappan Morris. 2012. Cache Craftiness for Fast Multicore Key-Value Storage. In *Proceedings of the 7th ACM European Conference on Computer Systems* (Bern, Switzerland) (EuroSys '12). Association for Computing Machinery, New York, NY, USA, 183–196. <https://doi.org/10.1145/2168836.2168855>
- [21] Paul E. McKenney and John D. Slingwine. 1998. Read-Copy Update: Using Execution History to Solve Concurrency Problems. In *Parallel and Distributed Computing and Systems*. Las Vegas, NV, 509–518.
- [22] Chris Nyberg, Tom Barclay, Zarka Cvetanovic, Jim Gray, and Dave Lomet. 1994. AlphaSort: A RISC Machine Sort. In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data* (Minneapolis, Minnesota, USA) (SIGMOD '94). Association for Computing Machinery, New York, NY, USA, 233–242. <https://doi.org/10.1145/191839.191884>
- [23] Karl Schnaitter, Serge Abiteboul, Tova Milo, and Neoklis Polyzotis. 2007. On-Line Index Selection for Shifting Workloads. In *Proceedings of the 23rd International Conference on Data Engineering Workshops* (Istanbul, Turkey) (ICDE '07). IEEE Computer Society, 459–468. <https://doi.org/10.1109/ICDEW.2007.4401029>
- [24] Xingbo Wu, Fan Ni, and Song Jiang. 2019. Wormhole: A Fast Ordered Index for In-Memory Data Management. In *Proceedings of the Fourteenth EuroSys Conference 2019* (Dresden, Germany) (EuroSys '19). Association for Computing Machinery, New York, NY, USA, Article 18, 16 pages. <https://doi.org/10.1145/3302424.3303955>
- [25] Huan Chen Zhang, David G. Andersen, Andrew Pavlo, Michael Kaminsky, Lin Ma, and Rui Shen. 2016. Reducing the Storage Overhead of Main-Memory OLTP Databases with Hybrid Indexes. In *Proceedings of the 2016 International Conference on Management of Data* (San Francisco, California, USA) (SIGMOD '16). Association for Computing Machinery, New York, NY, USA, 1567–1581. <https://doi.org/10.1145/2882903.2915222>

A Proof of Linearizability

We paraphrase the inductive invariants (I1, I2 and I3, Section 4.4) into Lemma 1, 5 and 6, and the correctness condition (Section 4.4) into Theorem 1. We show that Theorem 1 can be proved by Lemma 1, 5 and 6, and finally we prove the linearizability (Theorem 2) based on Theorem 1.

In the following proof, we do not discuss concurrent compaction and structure update as they are performed on a group exclusively in the implementation. We assume `get_group`

always returns the correct group such that its pivot key is less than or equal to given key and with its next group's pivot key is greater than given key. Since remove can be treated as a special case of put (set removed flag instead of change value, Section 4), we ignore it in the discussion. Since model split/merge and group merge can be seen as a simplified version of group split, we also omit them in the discussion.

Definition 1. (commit/linearization points) We say a put commits if, depending on the implementation states, put executes line 18, 21, 23 or 27 of Algorithm 2 (the linearization points). Specifically, when executing line 21, 23, 27, the linearization point is in between the invocation. For the vanilla (temporary) delta index implementation using a big lock, it is the point when the lock is aquired. put commits at.

- line 18 when it updates a record in sorted array;
- line 21 when it inserts/updates a record in delta index;
- line 23 when it updates a record in frozen delta index; and.
- line 27 when it inserts/updates a record in temporary delta index.

We say a get commits if, depending on the implementation states, get executes line 5, 8 or 10 of Algorithm 2 (the linearization points). Specifically, when executing line 10, 12, the linearization point is in between the invocation. For the vanilla (temporary) delta index implementation using a big lock, it is the point when the lock is aquired. get commits at.

- line 5 when it reads a record in sorted array;
- line 8 when it reads a record in delta index;
- line 8 when it fails to read a record in delta index and temporary delta index is not initialized;
- line 10 when it reads a record in temporary delta index; and.
- line 10 when it fails to read a record in temporary delta index and temporary delta index is initialized.

For two operations a and b, we use $a \ll b$ denote that a commits before b commits.

Definition 2. (conflicting operations) If operation a and b are invoked with the same key and one of them is put, then we say a and b are conflicting operations.

Definition 3. (phase transition) We say a compaction or group_split is performing phase transition if line 12-13, Algorithm 3 or line 10-12, 23-24, Algorithm 4 is being executed by a background thread.

It is obvious that, for given key, at most two groups are visible (returned by get_group) to threads during compaction or group_split performing phase transition. We term the group before phase transition as the *old group* and the created one as *new group*.

Definition 4. (record copy) A record copy in XIndex is a pair of key and value that.

- is stored in either **data_array**, **buf** or **tmp_buf**; and.
- the **data_array*/*buf*/*tmp_buf** can be accessed via the global root pointer and group pointers observed by an operation (i.e., returned by get_group).

If a record copy contains key k, then we say it is a record copy of k.

Lemma 1. (I1) For any possible concurrent execution of XIndex, for any key k, if there is put committed on k, then there is exactly one record copy of k in XIndex.

Proof of Lemma 1.

- 1 If there is put committed on k, then there is at least one record copy of k in XIndex.
 - 1.1 If there is no record copy of k in XIndex before put commits, then a new record copy of k will be created.
 - 1.1.1 If there is no record copy of k in XIndex before put commits, then put commits at line 21 or line 27, Algorithm 2.
 - 1.1.1.1 *data_array* is un-insertable after initialization.

By the implementation of get_position that returns empty when there is no record copy of k in **data_array**.
 - 1.1.1.2 *buf* cannot be in-place updated.

By the implementation of try_update_in_buffer that fails when there is no record copy of k in the buffer.

 - 1.1.1.3 Q.E.D.
 - 1.1.1.2 If put commits at line 21 or line 27, Algorithm 2, then a new record copy of k will be created.

By the implementation of insert_buffer that creates a new record copy of k when there is no record copy of k in the buffer.
 - 1.1.3 Q.E.D.
 - 1.2 Other operation does not remove record copies.

By OBVIOUS.
 - 1.3 Q.E.D.
- 2 If there is put committed on k, then there is at most one record copy of k in XIndex.
 - 2.1 It suffices to prove that for any key k, 1) two record copies of k can not exist within one group; and 2) two record copies of k can not exist in two different groups.

By OBVIOUS.
 - 2.2 Two record copies of k can not exist within one group.
 - 2.2.1 Two record copies of k can not exist in the different locations of a group.
 - 2.2.1.1 **CASE 1:** Two record copies of k can't exist in *data_array* and *buf* respectively.
 - 2.2.1.1.1 Record copy in *buf* is created by insert_buffer.

By line 21 and 27, Algorithm 2; line 8, Algorithm 3; line 4, 18, Algorithm 4.

- 2.2.1.1.2 **CASE 1:** Record copy created by inserting to *buf*.
 - 2.2.1.1.2.1 In case of 2.2.1.1.2, when put creates a record copy in *buf*, there are two record copy in *data_array* and *buf* respectively when put commits.
 - 2.2.1.1.2.1.1 Record copy in *data_array* is created before record copy in *buf*.
 - 2.2.1.1.2.1.1.1 *data_array* can't be inserted after initialization.
By the implementation of *data_array* that provides no insert interface.
 - 2.2.1.1.2.1.1.2 Only after initializing a group, put can insert new records to *buf*.
By the implementation that we first initialize *buf* of new group at line 8, Algorithm 3 and line 4, 18, Algorithm 4 before making group visible to threads at line 12-13, Algorithm 3 and line 10-12, 23-24, Algorithm 4.
 - 2.2.1.1.2.1.1.3 Q.E.D.
 - 2.2.1.1.2.1.2 Q.E.D.
 - 2.2.1.1.2.2 If put creates a record copy in *buf*, there is no record copy in *data_array* when put commits.
By Lemma 2.
 - 2.2.1.1.2.3 2.2.1.1.2.1 and 2.2.1.1.2.2 contradicts.
By OBVIOUS.
 - 2.2.1.1.2.4 Q.E.D.
- 2.2.1.1.3 **CASE 2:** Record copy created by inserting to *tmp_buf* of old note and later become the *buf* of new group.
 - 2.2.1.1.3.1 In case of 2.2.1.1.3, when put creates a record copy in *tmp_buf*, there are two record copy in *data_array/buf* and *tmp_buf* of the old group respectively when put commits.
 - 2.2.1.1.3.1.1 Record copy in *data_array/buf* of old group is created before record copy in *buf*.
 - 2.2.1.1.3.1.1.1 The record copy in new group's *data_array* is merged from old group's *data_array* or *buf*.
By line 7, Algorithm 3; line 17, Algorithm 4.
 - 2.2.1.1.3.1.1.2 The old group's *data_array/buf* is not insertable after old group's *tmp_buf* is initialized.
By line 3-4, Algorithm 3; line 11-12, Algorithm 4.
 - 2.2.1.1.3.1.1.3 put inserts record copy in *tmp_buf* after *tmp_buf* is initialized.
By OBVIOUS.
 - 2.2.1.1.3.1.1.4 Q.E.D.
 - 2.2.1.1.3.1.2 Q.E.D.
 - 2.2.1.1.3.2 If put creates a record copy in *tmp_buf*, there is no record copy in *data_array* or *buf* when put commits.
By Lemma 3.
 - 2.2.1.1.3.3 2.2.1.1.3.1 and 2.2.1.1.3.2 contradicts.
By OBVIOUS.
 - 2.2.1.1.3.4 Q.E.D.
 - 2.2.1.1.4 Q.E.D.
- 2.2.1.2 **CASE 2:** Two record copies of *k* can't exist in *data_array* and *tmp_buf* respectively.
 - 2.2.1.2.1 Record copy in *tmp_buf* is created by *insert_buffer*.
By line 21 and 27, Algorithm 2; line 8, Algorithm 3; line 4, 18, Algorithm 4.
 - 2.2.1.2.2 In case of 2.2.1.2, when put creates a record copy in *tmp_buf*, there are two record copy in *data_array* and *tmp_buf* respectively when put commits.
 - 2.2.1.2.2.1 Record copy in *data_array* is created before record copy in *tmp_buf*.
 - 2.2.1.2.2.1.1 *data_array* can't be inserted after initialization.
By the implementation of *data_array* that provides no insert interface.
 - 2.2.1.2.2.1.2 Only after initializing a group, put can insert new records to *tmp_buf*.
By the implementation that we always initialize *tmp_buf* to a group that is already initialized (line 8, Algorithm 3 and line 18, Algorithm 4).
 - 2.2.1.2.2.1.3 Q.E.D.
 - 2.2.1.2.2.2 Q.E.D.
 - 2.2.1.2.3 If put creates a record copy in *tmp_buf*, there is no record copy in *data_array* when put commits.
By Lemma 3.
 - 2.2.1.2.4 2.2.1.2.2 and 2.2.1.2.3 contradicts.
By OBVIOUS.
 - 2.2.1.2.5 Q.E.D.
- 2.2.1.3 **CASE 3:** Two record copies of *k* can't exist in *tmp_buf* and *buf* respectively.
 - 2.2.1.3.1 Record copy in *buf* and *tmp_buf* are created by *insert_buffer*.
By line 21 and 27, Algorithm 2; line 8, Algorithm 3; line 4, 18, Algorithm 4.
 - 2.2.1.3.2 In case of 2.2.1.3, when put creates a record copy in *tmp_buf*, there are two record copy in *buf* and *tmp_buf* respectively when put commits.
 - 2.2.1.3.2.1 If there is only one record copy, then the record copy in *buf* can't be created

- after the put that created record copy in *tmp_buf* commits
- 2.2.1.3.2.1.1 *buf* is not insertable after *tmp_buf* is initialized.
By line 3-5, Algorithm 3; line 10-11, 13, Algorithm 4.
 - 2.2.1.3.2.1.2 The put that created record copy in *tmp_buf* commits after *tmp_buf* is initialized.
By OBVIOUS.
 - 2.2.1.3.2.1.3 Q.E.D.
 - 2.2.1.3.2.2 Q.E.D.
 - 2.2.1.3.3 If put creates a record copy in *tmp_buf*, there is no record copy in *buf* when put commits.
By Lemma 3.
 - 2.2.1.3.4 2.2.1.3.2 and 2.2.1.3.3 contradicts.
By OBVIOUS.
 - 2.2.1.3.5 Q.E.D.
 - 2.2.1.4 Q.E.D.
 - 2.2.2 Two record copies of *k* can not exist in the same location of a group.
 - 2.2.2.1 **CASE 1:** two record copies of *k* can't both exist in *data_array*
 - 2.2.2.1.1 Two record copies of *k* with the same key can't exist within one *data_array* when initialization.
 - 2.2.2.1.1.1 *data_array* is constructed using the record copies from *data_array* and *buf* of another group.
By line 7, Algorithm 3; line 17, Algorithm 4.
 - 2.2.2.1.1.2 two record copies of *k* can't exist in *data_array* and *buf* of a group respectively.
By 2.2.1.
 - 2.2.2.1.1.3 Q.E.D.
 - 2.2.2.1.2 *data_array* can't be inserted after initialization.
By that *data_array* is implemented using fixed-length array.
 - 2.2.2.1.3 Q.E.D.
 - 2.2.2.2 **CASE 2:** two record copies of *k* can't both exist in *buf*
By the property of delta index implementation that repeating call to *insert_buffer* only updates previous record copy's value.
 - 2.2.2.3 **CASE 3:** two record copies of *k* can't both exist in *tmp_buf*
By the property of delta index implementation that repeating call to *insert_buffer* only updates previous record copy's value.
 - 2.2.2.4 Q.E.D.
 - 2.2.3 Q.E.D.
 - 2.3 Two record copies of *k* can not exist in two different groups.
 - 2.3.1 It suffices to prove that, during phase transition, two record copies of *k* can't exist in the old group and the new group respectively.
 - 2.3.1.1 Only during phase transition will there be two group visible (returned by *get_group*) to threads.
By OBVIOUS.
 - 2.3.1.2 Q.E.D.
 - 2.3.2 During phase transition, two record copies of *k* can't exist in the old group and the new group respectively.
 - 2.3.2.1 If two record copies of *k* exist in the old group and the new group respectively, they must be created during phase transition.
 - 2.3.2.1.1 Before phase transition, there is one group contains *k*.
By OBVIOUS.
 - 2.3.2.1.2 Two record copies of *k* can not exist within one group.
By 2.
 - 2.3.2.1.3 Q.E.D.
 - 2.3.2.2 If two record copies of *k* are created during phase transition, they are created by put.
 - 2.3.2.2.1 Only background threads and put creates new record copies.
By OBVIOUS.
 - 2.3.2.2.2 Background thread does not create new record copy during phase transition.
By line 12-13, Algorithm 3; 10-12, 23-24, Algorithm 4.
 - 2.3.2.2.3 Q.E.D.
 - 2.3.2.3 put can not created two record copies of *k* in the old group and the new group during phase transition.
 - 2.3.2.3.1 **CASE 1:** compaction's phase transition & *group_split*'s 2nd phase transition.
 - 2.3.2.3.1.1 If put creates two record copies of *k* in the old group and the new group during phase transition, two record copies will exist in both groups.
 - 2.3.2.3.1.1.1 All record copies are referenced in new group.
By line 7-8, Algorithm 3; 17-18, Algorithm 4.
 - 2.3.2.3.1.1.2 Q.E.D.
 - 2.3.2.3.1.2 Two record copies of *k* can not exist within one group.
By 2.2.
 - 2.3.2.3.1.3 Q.E.D.
 - 2.3.2.3.2 **CASE 2:** *group_split*'s 1st phase transition.

- 2.3.2.3.2.1 **CASE 1:** The duplicate record copy in new group is in *tmp_buf*.
 - 2.3.2.3.2.1.1 The duplicate record copy in old group is in *data_array* or *buf*.
By OBVIOUS.
 - 2.3.2.3.2.1.2 The old group's *data_array* and *buf* are shared with the new group.
By line 4, Algorithm 4.
 - 2.3.2.3.2.1.3 There are two record copies of *k* in the new group.
By 2.3.2.3.2.1.1 and 2.3.2.3.2.1.2.
 - 2.3.2.3.2.1.4 There cannot be are two record copies of *k* in the new group.
By 2.2.
 - 2.3.2.3.2.1.5 Q.E.D.
- 2.3.2.3.2.2 **CASE 2:** The duplicate record copy in new group is in *data_array* or *buf*.
 - 2.3.2.3.2.2.1 two record copies will exist in both old and new groups.
 - 2.3.2.3.2.2.1.1 The old group's *data_array* and *buf* are shared with the new group.
By line 4, Algorithm 4.
 - 2.3.2.3.2.2.1.2 Q.E.D.
 - 2.3.2.3.2.2.2 There cannot be are two record copies of *k* in the new group.
By 2.2.
 - 2.3.2.3.2.2.3 Q.E.D.
- 2.3.2.3.2.3 Q.E.D.
- 2.3.2.3.3 Q.E.D.
- 2.3.2.4 Q.E.D.
- 2.3.3 Q.E.D.
- 2.4 Q.E.D.
- 3 Q.E.D.

Lemma 2. For any possible concurrent execution of XIndex, if a put or get operating on the key *k* commits at *buf*, then there is no record copy of *k* in *data_array* at the time the request commits.

Proof of Lemma 2.

- 1 **CASE 1:** the request is a put request.
 - 1.1 If put inserts/updates a record copy in *buf*, there is no record copy in *data_array* when put executes line 18.
By the fact that otherwise put would then commit at line 18, Algorithm 2, instead of creating a record in *buf*.
 - 1.2 If there is no record copy in *data_array* when put executes line 18, there won't be a record copy in *data_array* after put executes line 18.
 - 1.2.1 *data_array* can't be inserted after initialization.
By the implementation of *data_array* that provides no insert interface.
 - 1.2.2 put accesses *buf* after *data_array* is initialized.
By OBVIOUS.

- 1.2.3 Q.E.D.
- 1.3 put commits after it executes line 18.
By OBVIOUS.
- 1.4 Q.E.D.
- 2 **CASE 2:** the request is a get request.
 - 2.1 If get commits at line 8, there is no record copy in *data_array* when get executes line 6.
By the fact that otherwise get would then commit at line 6, Algorithm 2.
 - 2.2 If there is no record copy in *data_array* when get executes line 6, there won't be a record copy in *data_array* after get executes line 6.
 - 2.2.1 *data_array* can't be inserted after initialization.
By the implementation of *data_array* that provides no insert interface.
 - 2.2.2 get accesses *buf* after *data_array* is initialized.
By OBVIOUS.
 - 2.2.3 Q.E.D.
 - 2.3 get commits after it execute line 6.
By OBVIOUS.
 - 2.4 Q.E.D.
- 3 Q.E.D.

Lemma 3. For any possible concurrent execution of XIndex, if a put or get operating on the key *k* commits at *tmp_buf* (reads value from *buf* or insert/update value in *tmp_buf*), then there is no record copy of *k* in *data_array* or *buf* at the time the request commits.

Proof of Lemma 3.

- 1 **CASE 1:** the request is a put request.
 - 1.1 If put commits at line 27, there is no record copy in *data_array* when put executes line 18 or 22.
By the fact that otherwise put would then commit at line 18, Algorithm 2, instead of creating a record in *buf*.
 - 1.2 If there is no record copy in *data_array* when put executes line 18 or 22, there won't be a record copy in *data_array* after put executes line 18 or 22.
 - 1.2.1 *data_array* can't be inserted after initialization.
By the implementation of *data_array* that provides no insert interface.
 - 1.2.2 put accesses *tmp_buf* after *data_array* is initialized.
By OBVIOUS.
 - 1.2.3 Q.E.D.
 - 1.3 If put commits at line 27, there is no record copy in *buf* when put executes line 21 or 23.
By the fact that otherwise put would then commit at line 21 or 23, Algorithm 2, instead of creating a record in *tmp_buf*.

- 1.4 If there is no record copy in *data_array* when put executes line 21 or 23, there won't be a record copy in *data_array* after put executes line 21 or 23.
 - 1.4.1 no put can insert new record copy to *buf*.
 - 1.4.1.1 *buf* is frozen before *tmp_buf* is initialized.
 - By line 3-5, Algorithm 3; line 10-12, Algorithm 4.
 - 1.4.1.2 *tmp_buf* is initialized.
 - By the assumption that put will commit at line 27, Algorithm 2.
 - 1.4.1.3 Q.E.D.
 - 1.4.2 Q.E.D.
 - 1.5 put commits after it executes line 18, 22, 21 and 23.
 - By OBVIOUS.
 - 1.6 Q.E.D.
 - 2 **CASE 2:** the request is a get request.
 - 2.1 If get commits at line 10, there is no record copy in *data_array* when get executes line 6.
 - By the fact that otherwise get would then commit at line 6, Algorithm 2.
 - 2.2 If there is no record copy in *data_array* when get executes line 6, there won't be a record copy in *data_array* after get executes line 6.
 - 2.2.1 *data_array* can't be inserted after initialization.
 - By the implementation of *data_array* that provides no insert interface.
 - 2.2.2 get accesses *tmp_buf* after *data_array* is initialized.
 - By OBVIOUS.
 - 2.2.3 Q.E.D.
 - 2.3 If get commits at line 10, there is no record copy in *buf* when get executes line 10.
 - By the fact that otherwise get would then commit at line 8, Algorithm 2.
 - 2.4 If there is no record copy in *data_array* when get executes line 8, there won't be a record copy in *data_array* after get executes line 8.
 - 2.4.1 no put can insert new record copy to *buf*.
 - 2.4.1.1 *buf* is frozen before *tmp_buf* is initialized.
 - By line 3-5, Algorithm 3; line 10-12, Algorithm 4.
 - 2.4.1.2 *tmp_buf* is initialized.
 - By the assumption that get will commit at line 10, Algorithm 2.
 - 2.4.1.3 Q.E.D.
 - 2.4.2 Q.E.D.
 - 2.5 get commits after it executes line 6 and 8.
 - By OBVIOUS.
 - 2.6 Q.E.D.
 - 3 Q.E.D.
- Lemma 4.** For any possible concurrent execution of XIndex, for any conflicting operations a and b, either $a \ll b$ or $b \ll a$.
- Proof of Lemma 4.**
- 1 **CASE 1:** a and b are puts.
 - 1.1 It suffices to prove that any conflicting put a,b must not commit simultaneously 1) within one group; and 2) in two different groups.
 - By OBVIOUS.
 - 1.2 any conflicting put a, b operating on key k can't commit simultaneously within one group.
 - 1.2.1 **CASE 1:** a, b must not commit at the same LP simultaneously.
 - 1.2.1.1 **CASE 1:** put commits at line 18, Algorithm 2.
 - By atomicity of commit ensured by lock, line 21, Algorithm 2.
 - 1.2.1.2 **CASE 2:** put commits at line 21, 23 or 27, Algorithm 2.
 - By the implementation of (temporary) delta index that ensures atomicity.
 - 1.2.1.3 Q.E.D.
 - 1.2.2 **CASE 2:** a, b must not commits at the different LP simultaneously.
 - 1.2.2.1 **CASE 1:** a commits at line 18, b commits at line 21 or 23 or 27.
 - 1.2.2.1.1 a record copy of k exists in *data_array* at the time a commits.
 - By the fact that otherwise a will continue execution and commit after line 18, Algorithm 2.
 - 1.2.2.1.2 no record copy of k exists in *data_array* at the time b commits.
 - By Lemma 2, 3.
 - 1.2.2.1.3 1.2.2.1.1 contradicts with 1.2.2.1.2.
 - By OBVIOUS.
 - 1.2.2.1.4 Q.E.D.
 - 1.2.2.2 **CASE 2:** a commits at line 21, b commits line 23 or line 27.
 - 1.2.2.2.1 *buf_frozen* is false at the time a commits.
 - By the fact that otherwise a would fail at branch condition at line 20 and cannot commit at line 21.
 - 1.2.2.2.2 all threads must have observed *buf_frozen* is true at the time b commits.
 - 1.2.2.2.2.1 *tmp_buf* have been initialized at the time b commits.
 - By b commits at line 23, line 25-26 of Algorithm 2.
 - 1.2.2.2.2.2 *tmp_buf* is only initialized after all threads have observed *buf_frozen* is true.

- 1.3.2.1.1.2.3.2.1.2 Q.E.D.
- 1.3.2.1.1.2.3.2.2 Q.E.D.
- 1.3.2.1.1.2.3.3 1.3.2.1.1.2.3.1 contradicts with 1.3.2.1.1.2.3.2.
By OBVIOUS.
- 1.3.2.1.1.2.3.4 Q.E.D.
- 1.3.2.1.1.2.4 Q.E.D.
- 1.3.2.1.1.3. **CASE 2:** b commits at line 21.
 - 1.3.2.1.1.3.1 a can only commit at line 18 or line 23 or line 27.
 - 1.3.2.1.1.3.1.1 All threads must have observed *buf_frozen* is true in the old group.
By *mfence* and *rcu_barrier* and line 3-4 of Algorithm 3.
 - 1.3.2.1.1.3.1.2 b can't commit at line 21.
By 2.2.1.2.1.1, line 27-21 of Algorithm 2.
 - 1.3.2.1.1.3.1.3 Q.E.D.
 - 1.3.2.1.1.3.2 **CASE 1:** a commits at line 18, 23.
 - 1.3.2.1.1.3.2.1 There exist a record copy of k in *data_array* or *buf* of the old group.
By a commits at *data_array* or *buf* of the old group.
 - 1.3.2.1.1.3.2.2 No record copy of k exists in *data_array* or *buf* of the old group.
 - 1.3.2.1.1.3.2.2.1 If there is a record copy of k exists in *data_array* or *buf* of the old group, then there is a reference in *data_array* of the new group.
By line 7 of Algorithm 3.
 - 1.3.2.1.1.3.2.2.2 If there is a reference in *data_array* of the new group, b will not commit at line 21.
By Lemma 2.
 - 1.3.2.1.1.3.2.2.3 Q.E.D.
 - 1.3.2.1.1.3.2.3 1.3.2.1.1.3.2.1 contradicts with 1.3.2.1.1.3.2.1.
By OBVIOUS.
 - 1.3.2.1.1.3.2.4 Q.E.D.
 - 1.3.2.1.1.3.3 **CASE 2:** a commits at line 27.
 - 1.3.2.1.1.3.3.1 *buf* of the new group and *tmp_buf* of the old group refers to the same delta index instance.
By line 8 of Algorithm 3.
 - 1.3.2.1.1.3.3.2 a,b can't commit simultaneously on the same delta index.
By the property of delta index implementation.
 - 1.3.2.1.1.3.3.3 Q.E.D.
 - 1.3.2.1.1.3.4 Q.E.D.
 - 1.3.2.1.1.4 Q.E.D.
 - 1.3.2.1.2 Q.E.D.
- 1.3.2.2 **CASE 2:** group split phase 1 \rightarrow phase 2.
 - 1.3.2.2.1 **CASE 1:** a, b can't commit in \mathbf{ga}^* , \mathbf{gb}^* simultaneously.
 - 1.3.2.2.1.1 \mathbf{ga}^* and \mathbf{gb}^* are created to have different key interval.
By line 5-6 of Algorithm 4.
 - 1.3.2.2.1.2 If a,b commit at \mathbf{ga}^* and \mathbf{gb}^* respectively, then they are not conflicting.
By 3.2.2.1.1.
 - 1.3.2.2.1.3 Q.E.D.
 - 1.3.2.2.2 **CASE 2:** a,b can't commit simultaneously in \mathbf{ga}^* (the new group) and the old group respectively.
 - 1.3.2.2.2.1 b can only commit at line 18 or line 21.
 - 1.3.2.2.2.1.1 *buf_frozen* is false in the old group.
By that group split will not set old group to frozen, Algorithm 4.
 - 1.3.2.2.2.1.2 Q.E.D.
 - 1.3.2.2.2.2 **CASE 1:** *buf_frozen* is false in \mathbf{ga}^* .
 - 1.3.2.2.2.2.1 *buf_frozen* is false, *tmp_buf* is null in both \mathbf{ga}^* and the old group.
By 1.3.2.2.2.2, *tmp_buf* is only initialized after *buf_frozen* is true.
 - 1.3.2.2.2.2.2 a can only commit at line 18 or line 21.
By 1.3.2.2.2.2.1.
 - 1.3.2.2.2.2.3 \mathbf{ga}^* and the old group share the same *data_array* and *buf*
By line 4 of Algorithm 4.
 - 1.3.2.2.2.2.4 a, b operate on the same *data_array* and *buf*.
By 1.3.2.2.2.2.3.
 - 1.3.2.2.2.2.5 a, b can't commit simultaneously.
By 1.3.2.2.2.2.4, 1.3.2.2.2.2.1.
 - 1.3.2.2.2.2.6 Q.E.D.
 - 1.3.2.2.2.3 **CASE 2:** *buf_frozen* is true in \mathbf{ga}^* .
 - 1.3.2.2.2.3.1 a can only commit at line 18, line 21 or line 23
By that *tmp_buf* is not initialized in \mathbf{ga}^* .
 - 1.3.2.2.2.3.2 **CASE 1:** a commits line 18.
 - 1.3.2.2.2.3.2.1 **CASE 1:** b commits at line 18
 - 1.3.2.2.2.3.2.1.1 a, b commit at the same *data_array*
By a, b both commit at line 18.
 - 1.3.2.2.2.3.2.1.2 a, b can't commit simultaneously.
By that records in *data_array* are protected by the per record lock.
 - 1.3.2.2.2.3.2.1.3 Q.E.D.
 - 1.3.2.2.2.3.2.2 **CASE 2:** b commits at line 21

- 1.3.2.2.3.2.2.1 There exist a record copy of k is in $data_array$ at the time a commits
By that a commits in $data_array$, and g_a' and the old group share the same $data_array$.
- 1.3.2.2.3.2.2.2 b update the record copy of k in $data_array$
By line 17-18 of Algorithm 2.
- 1.3.2.2.3.2.2.3 b will not commit at line 21
By 1.3.2.2.3.2.2.2.
- 1.3.2.2.3.2.2.4 Q.E.D.
- 1.3.2.2.3.2.3 Q.E.D.
- 1.3.2.2.3.3 **CASE 2:** a commits at line 21 or line 23
 - 1.3.2.2.3.3.1 No record copy of k in $data_array$ at the time a commits
By that a commits at buf
 - 1.3.2.2.3.3.2 b can only commit at line 21
By 1.3.2.2.3.3.1.
 - 1.3.2.2.3.3.3 a, b commits at the same buf
By that g_a' and the old group share the same $data_array$.
 - 1.3.2.2.3.3.4 a, b will not commit simultaneously.
By 1.3.2.2.3.3.3.
 - 1.3.2.2.3.3.5 Q.E.D.
- 1.3.2.2.3.4 Q.E.D.
- 1.3.2.2.4 Q.E.D.
- 1.3.2.2.3 **CASE 3:** a, b can't commit in g_b' , the old group simultaneously.
By symmetry of b , and 1.3.2.2.2.
- 1.3.2.2.4 Q.E.D.
- 1.3.2.3 **CASE 3:** group split phase 2 \rightarrow phase 3.
 - 1.3.2.3.1 $g_a' \rightarrow g_a$ and $g_b' \rightarrow g_b$ are two compaction process.
By design of structure update.
 - 1.3.2.3.2 a, b can't commit simultaneously.
By 1.3.2.1.
 - 1.3.2.3.3 Q.E.D.
- 1.3.2.4 Q.E.D.
- 1.3.3 Q.E.D.
- 1.4 Q.E.D.
- 2 **CASE 2:** a is put and b is get.
 - 2.1 It suffices to prove that any conflicting put a , get b can't commit simultaneously 1) within one group; and 2) in two different groups.
By OBVIOUS.
 - 2.2 Any conflicting put a , get b operating on the same key k can't commit simultaneously within one groups.
 - 2.2.1 **CASE 1:** b commits at line 6 of Algorithm 2.
 - 2.2.1.1 **CASE 1:** a commits at line 18 of Algorithm 2.
 - 2.2.1.1.1 a, b can't commit simultaneously in $data_array$
By that record copy in $data_array$ is protected using OCC.
 - 2.2.1.1.2 Q.E.D.
 - 2.2.1.2 **CASE 2:** a commits at line 21 or line 23 or line 27 of Algorithm 2.
 - 2.2.1.2.1 No record copy of k exists in $data_array$ at the time a commits.
By the fact that a commits at buf or tmp_buf and lemma 2, 3.
 - 2.2.1.2.2 There exists a record copy of k in $data_array$ at the time b commits.
By b commits at $data_array$
 - 2.2.1.2.3 2.2.1.2.1 contradicts with 2.2.1.2.2.
By OBVIOUS.
 - 2.2.1.2.4 Q.E.D.
 - 2.2.1.3 Q.E.D.
 - 2.2.2 **CASE 2:** b commits at line 8 of Algorithm 2.
 - 2.2.2.1 **CASE 1:** a commits at line 18 of Algorithm 2.
 - 2.2.2.1.1 No record copy of k exists in $data_array$ at the time b commits.
By the fact that get b commit at buf and Lemma 2.
 - 2.2.2.1.2 There exists a record copy of k in $data_array$ at the time a commits.
By put a commits at $data_array$
 - 2.2.2.1.3 2.2.2.1.1 contradicts with 2.2.2.1.2.
By OBVIOUS.
 - 2.2.2.1.4 Q.E.D.
 - 2.2.2.2 **CASE 2:** a commits at line 21,23 of Algorithm 2.
 - 2.2.2.2.1 a, b both commit at buf
By a commits at line 21 or line 23, b commits at line 10.
 - 2.2.2.2.2 a, b can't commit simultaneously in the same buf
By the property of delta index implementation.
 - 2.2.2.2.3 Q.E.D.
 - 2.2.2.3 **CASE 3:** a commits at line 27 of Algorithm 2.
 - 2.2.2.3.1 No record copy of k exists in $data_array$ or buf at the time a commits.
By the fact that a commits at line 27 and Lemma 3.
 - 2.2.2.3.2 There exists a record copy of k in buf at the time b commits.
By b commits at buf

- 2.2.2.3.3 2.2.2.3.1 contradicts with 2.2.2.3.2.
By OBVIOUS.
- 2.2.2.3.4 Q.E.D.
- 2.2.2.4 Q.E.D.
- 2.2.3 **CASE 3:** b commits at line 10 of Algorithm 2.
 - 2.2.3.1 **CASE 1:** a commit at line 18 of Algorithm 2.
 - 2.2.3.1.1 No record copy of k exists in *data_array* and *buf* at the time b commits.
By the fact that that b commits at *tmp_buf* and lemma 3.
 - 2.2.3.1.2 There exists a record copy of k in *data_array* at the time a commits.
By that a commits at *data_array*
 - 2.2.3.1.3 2.2.3.1.1 contradicts with 2.2.3.1.2.
By OBVIOUS.
 - 2.2.3.1.4 Q.E.D.
 - 2.2.3.2 **CASE 2:** a commits at line 21 of Algorithm 2.
 - 2.2.3.2.1 *tmp_buf* is initialized.
By that b commits at *tmp_buf*
 - 2.2.3.2.2 *buf_frozen* is true for all threads.
By *mfence* and *rcu_barrier*, line 3-5 of Algorithm 2.
 - 2.2.3.2.3 a can't commit at line 21.
By line 27 of Algorithm 2, 2.2.3.2.2.
 - 2.2.3.2.4 Q.E.D.
 - 2.2.3.3 **CASE 3:** a commits at line 23 of Algorithm 2.
 - 2.2.3.3.1 No record copy of k exists in *data_array* and *buf* at the time b commits.
By the fact that b commits at *tmp_buf* and lemma 3.
 - 2.2.3.3.2 There exists a record copy of k in *buf*
By a commits at line 23.
 - 2.2.3.3.3 2.2.3.3.1 contradicts with 2.2.3.3.2.
By OBVIOUS.
 - 2.2.3.3.4 Q.E.D.
 - 2.2.3.4 **CASE 4:** a commits at line 27 of Algorithm 2.
 - 2.2.3.4.1 a,b both commit at *tmp_buf*
By that a commits at line 10 and b commits at line 27.
 - 2.2.3.4.2 a,b can't commit simultaneously.
By the property of delta index implementation.
 - 2.2.3.4.3 Q.E.D.
 - 2.2.3.5 Q.E.D.
- 2.2.4 Q.E.D.
- 2.3 Any conflicting put a, get b can't commit simultaneously on two groups.
 - 2.3.1 It suffices to prove that, during phase transition, two conflicting put a, get b can't commit simultaneously in the old group and the new group.
 - 2.3.1.1 Only during phase transition will there be two group visible (returned by *get_group*) to threads.
By OBVIOUS.
 - 2.3.1.2 Q.E.D.
 - 2.3.2 During phase transition, two conflicting put a, get b can't commit simultaneously in the old group and the new group respectively.
 - 2.3.2.1 **CASE 1:** compaction phase 1→phase 2.
 - 2.3.2.1.1 If a and b commit in the old group and the new group respectively, then there is a contradiction.
 - 2.3.2.1.1.1 **CASE 1:** put a commits in the old group, get b commits in the new group.
 - 2.3.2.1.1.1.1 get b can only commit at line 6 or line 8 in the new group.
 - 2.3.2.1.1.1.1.1 b can't commit at line 10 in the new group.
By the fact that *tmp_buf* is not initialized in the new group, line 9 of Algorithm 2.
 - 2.3.2.1.1.1.1.2 Q.E.D.
 - 2.3.2.1.1.1.2 **CASE 1:** get b commits at line 6 of the new group.
 - 2.3.2.1.1.1.2.1 put a can only commit at line 18 or line 23 or line 27 in the old group.
 - 2.3.2.1.1.1.2.1.1 a can't commit at line 21 in the old group.
 - 2.3.2.1.1.1.2.1.1.1 *buf_frozen* is true for all threads in the old group.
By *mfence* and *rcu_barrier*, line 3-4 of Algorithm 3.
 - 2.3.2.1.1.1.2.1.1.2 a can't insert *buf* in the old group.
By 2.3.2.1.1.1.2.1.1.1, line 20-22 of Algorithm 2.
 - 2.3.2.1.1.1.2.1.1.3 Q.E.D.
 - 2.3.2.1.1.1.2.1.2 Q.E.D.
 - 2.3.2.1.1.1.2.2 **CASE 1:** a commits at line 18 in the old group.
 - 2.3.2.1.1.1.2.2.1 a,b both operate on the same record copy in *data_array* of the old group.
 - 2.3.2.1.1.1.2.2.1.1 There exists a record copy of k in *data_array* of the old group at the time a commits.
By the fact that a commits in *data_array* of the old group.
 - 2.3.2.1.1.1.2.2.1.2 There is a reference to the record copy at the time a commits in *data_array* of the new group.

- 2.3.2.1.1.1.2.2.1.2.1 *data_array* of the new groups contains references of all the record copies in *data_array* and *buf* of the old group.
By line 7 of Algorithm 3.
- 2.3.2.1.1.1.2.2.1.2.2 Q.E.D.
- 2.3.2.1.1.1.2.2.1.3 b accesses the same record copy in the old group's *data_array* through reference in the new group's *data_array*
 - 2.3.2.1.1.1.2.2.1.3.1 If the record copy that b accesses is in *buf* of the old group, then Lemma 1 is violated.
By OBVIOUS.
 - 2.3.2.1.1.1.2.2.1.3.2 Q.E.D.
- 2.3.2.1.1.1.2.2.1.4 Q.E.D.
- 2.3.2.1.1.1.2.2.2 a,b can't commit simultaneously.
By that record copies in *data_array* are protected with OCC.
- 2.3.2.1.1.1.2.2.3 Q.E.D.
- 2.3.2.1.1.1.2.3 **CASE 2:** a commits at line 23 or line 27 in the old group.
 - 2.3.2.1.1.1.2.3.1 There exists a record copy of k in *data_array* of the old group at the time a commits.
By the fact that a commits in *data_array* of the old group.
 - 2.3.2.1.1.1.2.3.2 There is no record copy of k in *data_array* or *buf* of the old group at the time b commits.
 - 2.3.2.1.1.1.2.3.2.1 There is no reference to a record copy of k in *data_array* of the new group at the time b commits.
By the fact that a commits at *buf* or *tmp_buf* of the new group, Lemma 2 and Lemma 3.
 - 2.3.2.1.1.1.2.3.2.2 *data_array* of the new groups contains references of all the record copies in *data_array* and *buf* of the old group.
By line 7 of Algorithm 3.
 - 2.3.2.1.1.1.2.3.2.3 Q.E.D.
 - 2.3.2.1.1.1.2.3.3 2.3.2.1.1.1.2.3.1 contradicts with 2.3.2.1.1.1.2.3.2.
By OBVIOUS.
 - 2.3.2.1.1.1.2.3.4 Q.E.D.
- 2.3.2.1.1.1.2.4 Q.E.D.
- 2.3.2.1.1.1.3 **CASE 2:** get b commits at line 8 of the new group.
 - 2.3.2.1.1.1.3.1 put a can only commit at line 18 or line 23 or line 27 in the old group.
 - 2.3.2.1.1.1.3.1.1 a can't commit at line 21 in the old group.
 - 2.3.2.1.1.1.3.1.1.1 *buf_frozen* is true for all threads in the old group.
By mfence and rcu_barrier, line 3-4 of Algorithm 3.
 - 2.3.2.1.1.1.3.1.1.2 a can't insert *buf* in the old group.
By 2.3.2.1.1.1.3.1.1.1, line 20-22 of Algorithm 2.
 - 2.3.2.1.1.1.3.1.1.3 Q.E.D.
 - 2.3.2.1.1.1.3.1.2 Q.E.D.
 - 2.3.2.1.1.1.3.2 **CASE 1:** a commits at line 18 in the old group.
 - 2.3.2.1.1.1.3.2.1 There exists a record copy of k in *data_array* of the old group at the time a commits.
By that fact that a commits at *data_array* of the old group.
 - 2.3.2.1.1.1.3.2.2 There is no record copy of k in *data_array* or *buf* of the old group at the time b commits.
 - 2.3.2.1.1.1.3.2.2.1 There is no reference to a record copy of k in *data_array* of the new group.
By that fact that b commits at *buf* of the new group, Lemma 2.
 - 2.3.2.1.1.1.3.2.2.2 *data_array* of the new groups contains references of all the record copies in *data_array* and *buf* of the old group.
By line 7 of Algorithm 3.
 - 2.3.2.1.1.1.3.2.2.3 Q.E.D.
 - 2.3.2.1.1.1.3.2.3 2.3.2.1.1.1.3.2.1 contradicts with 2.3.2.1.1.1.3.2.2.
By OBVIOUS.
 - 2.3.2.1.1.1.3.2.4 Q.E.D.
 - 2.3.2.1.1.1.3.3 **CASE 2:** a commits at line 23 in the old group.
 - 2.3.2.1.1.1.3.3.1 There exists a record copy of k in *buf* of the old group at the time a commits.
By that fact that a commits at *buf* of the old group and *buf* is in-place updated.
 - 2.3.2.1.1.1.3.3.2 There is no record copy of k in *data_array* or *buf* of the old group at the time b commits.
 - 2.3.2.1.1.1.3.3.2.1 There is no reference to a record copy of k in *data_array* of the new group.

- By that fact that *b* commits at *buf* of the new group, Lemma 2.
- 2.3.2.1.1.1.3.3.2.2 *data_array* of the new groups contains references of all the record copies in *data_array* and *buf* of the old group.
 - By line 7 of Algorithm 3.
 - 2.3.2.1.1.1.3.3.2.3 Q.E.D.
 - 2.3.2.1.1.1.3.3.3 2.3.2.1.1.1.3.3.1 contradicts with 2.3.2.1.1.1.3.3.2.
 - By OBVIOUS.
 - 2.3.2.1.1.1.3.3.4 Q.E.D.
 - 2.3.2.1.1.1.3.4 **CASE 3:** *a* commits at line 27 in the old group.
 - 2.3.2.1.1.1.3.4.1 *a* commits at *tmp_buf* of the old group, *b* commits at *buf* of the new group.
 - By the fact that *a* commits at line 27, *b* commits at line 10.
 - 2.3.2.1.1.1.3.4.2 *buf* of the new group and *tmp_buf* of the old group refers to the same delta index instance.
 - By line 8 of Algorithm 3.
 - 2.3.2.1.1.1.3.4.3 *a, b* can't commit simultaneously on the same delta index.
 - By the property of delta index implementation.
 - 2.3.2.1.1.1.3.4.4 Q.E.D.
 - 2.3.2.1.1.1.3.5 Q.E.D.
 - 2.3.2.1.1.1.4 Q.E.D.
 - 2.3.2.1.1.2 **CASE 2:** put *a* commits in the new group, get *b* commits in the old group.
 - 2.3.2.1.1.2.1 put *a* can only commit at line 18 or line 21 in the new group.
 - 2.3.2.1.1.2.1.1 put *a* can't commit at line 23 or line 27 in the new group.
 - By the fact that *buf_frozen* is false in the new group and line 22, 26 of Algorithm 2.
 - 2.3.2.1.1.2.1.2 Q.E.D.
 - 2.3.2.1.1.2.2 **CASE 1:** put *a* commits at line 18 in the new group.
 - 2.3.2.1.1.2.2.1 **CASE 1:** get *b* commits at line 6 in the old group.
 - 2.3.2.1.1.2.2.1.1 *a, b* both operate on the same record copy in *data_array* of the old group.
 - 2.3.2.1.1.2.2.1.1.1 There exists a record copy of *k* in *data_array* of the old group at the time *b* commits.
 - By the fact that *b* commits at *data_array*
 - 2.3.2.1.1.2.2.1.1.2 There exists a reference to the record copy in the new group's *data_array* at the time *b* commits.
 - 2.3.2.1.1.2.2.1.1.2.1 *data_array* of the new groups contains references of all the record copies in *data_array* and *buf* of the old group.
 - By line 7 of Algorithm 3.
 - 2.3.2.1.1.2.2.1.1.2.2 Q.E.D.
 - 2.3.2.1.1.2.2.1.1.3 *a* accesses the record copy in *oldgroup's *dataarray** through the reference in *data_array* of the new group.
 - 2.3.2.1.1.2.2.1.1.3.1 If the reference points to a record copy in *old_group's buf*, then Lemma 1 is violated.
 - By OBVIOUS.
 - 2.3.2.1.1.2.2.1.1.3.2 Q.E.D.
 - 2.3.2.1.1.2.2.1.1.4 Q.E.D.
 - 2.3.2.1.1.2.2.1.2 *a, b* can't commit simultaneously.
 - By the fact that record copies in *data_array* is protected with OCC, 2.3.2.1.1.2.2.1.1.
 - 2.3.2.1.1.2.2.1.3 Q.E.D.
 - 2.3.2.1.1.2.2.2 **CASE 2:** get *b* commits at line 8 in the old group.
 - 2.3.2.1.1.2.2.2.1 *a, b* both operate on the same record copy in *buf* of the old group.
 - 2.3.2.1.1.2.2.2.1.1 There exists a record copy of *k* in *buf* of the old group at the time *b* commits.
 - 2.3.2.1.1.2.2.2.1.1.1 If there is no record copy of *k* in *buf* of the old group, then *b* will not commit at line 8, which introduces a contradiction.
 - 2.3.2.1.1.2.2.2.1.1.1.1 *getfrombuffer* at line 8 will return empty.
 - By the assumption that no record copy of *k* in *buf* of the old group.
 - 2.3.2.1.1.2.2.2.1.1.1.2 *tmp_buf* is initialized in the old group.
 - By line 5 of Algorithm 3.
 - 2.3.2.1.1.2.2.2.1.1.1.3 *b* will commit at line 10, which introduces a contradiction.
 - By 2.3.2.1.1.2.2.2.1.1.1.1 and 2.3.2.1.1.2.2.2.1.1.1.2 line 11 of Algorithm 2.
 - 2.3.2.1.1.2.2.2.1.1.1.4 Q.E.D.
 - 2.3.2.1.1.2.2.2.1.1.2 Q.E.D.
 - 2.3.2.1.1.2.2.2.1.2 There exists a reference to the record copy in *data_array*

- of the new group at the time b commits.
- 2.3.2.1.1.2.2.1.2.1 *data_array* of the new groups contains references of all the record copies in *data_array* and *buf* of the old group.
By line 7 of Algorithm 3.
- 2.3.2.1.1.2.2.1.2.2 Q.E.D.
- 2.3.2.1.1.2.2.1.3 a accesses the record copy in *oldgroup*'s *buf* through the reference in **dataarray** of the new group.
- 2.3.2.1.1.2.2.1.3.1 If the reference points to a record copy in *oldgroup*'s **dataarray**, then Lemma 1 is violated.
By OBVIOUS.
- 2.3.2.1.1.2.2.1.3.2 Q.E.D.
- 2.3.2.1.1.2.2.1.4 Q.E.D
- 2.3.2.1.1.2.2.2.2 a, b can't commit simultaneously.
By the property of delta index implementation, 2.3.2.1.1.2.2.2.1.
- 2.3.2.1.1.2.2.2.3 Q.E.D.
- 2.3.2.1.1.2.2.3 **CASE 3:** get b commits at line 10 in the old group.
- 2.3.2.1.1.2.2.3.1 There is no record copy of k in *data_array* or *buf* of the old group at the time b commits.
By the fact that b commits at *tmp_buf*, Lemma 3.
- 2.3.2.1.1.2.2.3.2 There is a record copy of k in *data_array* or *buf* of the old group at the time a commits
 - 2.3.2.1.1.2.2.3.2.1 There exists a reference in *data_array* of the new group that refers to a record copy of k at the time a commits.
By the fact that a commits at *data_array* of the new group.
 - 2.3.2.1.1.2.2.3.2.2 *data_array* of the new groups contains references of all the record copies in *data_array* and *buf* of the old group.
By line 7 of Algorithm 3.
 - 2.3.2.1.1.2.2.3.2.3 Q.E.D.
- 2.3.2.1.1.2.2.3.3 2.3.2.1.1.2.2.3.1 contradicts with 2.3.2.1.1.2.2.3.2.
By OBVIOUS.
- 2.3.2.1.1.2.2.3.4 Q.E.D.
- 2.3.2.1.1.2.2.4 Q.E.D.
- 2.3.2.1.1.2.3 **CASE 2:** put a commits at line 21 in the new group.
 - 2.3.2.1.1.2.3.1 **CASE 1:** get b commits at line 6 in the old group.
 - 2.3.2.1.1.2.3.1.1 There exist a record copy of k in *data_array* of the old group at the time b commits.
By the fact that b commits at *data_array* of the old group.
 - 2.3.2.1.1.2.3.1.2 There is no record copy of k in *data_array* or *buf* of the old group at the time a commits.
 - 2.3.2.1.1.2.3.1.2.1 There is no record reference in *data_array* of the new group that refers to a record copy of k at the time a commits.
By the fact that a commits at *buf* of the new group.
 - 2.3.2.1.1.2.3.1.2.2 *data_array* of the new groups contains references of all the record copies in *data_array* and *buf* of the old group.
By line 7 of Algorithm 3.
 - 2.3.2.1.1.2.3.1.2.3 Q.E.D.
 - 2.3.2.1.1.2.3.1.3 2.3.2.1.1.2.3.1.1 contradicts with 2.3.2.1.1.2.3.1.2.
By OBVIOUS.
 - 2.3.2.1.1.2.3.1.4 Q.E.D.
 - 2.3.2.1.1.2.3.2 **CASE 2:** get b commits at line 8 in the old group.
 - 2.3.2.1.1.2.3.2.1 There exist a record copy of k in *buf* of the old group at the time b commits.
 - 2.3.2.1.1.2.3.2.1.1 If there is no record copy of k in *buf* of the old group, then b will not commit at line 8, which introduces a contradiction.
 - 2.3.2.1.1.2.3.2.1.1.1 *getfrombuffer* at line 8 will return empty.
By the assumption that no record copy of k in *buf* of the old group.
 - 2.3.2.1.1.2.3.2.1.1.2 *tmp_buf* is initialized in the old group.
By line 5 of Algorithm 3.
 - 2.3.2.1.1.2.3.2.1.1.3 b will commit at line 10, which introduces a contradiction.
By 2.3.2.1.1.2.3.2.1.1.1 and 2.3.2.1.1.2.3.2.1.1.2, line 9 of Algorithm 2.
 - 2.3.2.1.1.2.3.2.1.1.4 Q.E.D.
 - 2.3.2.1.1.2.3.2.1.2 Q.E.D.
 - 2.3.2.1.1.2.3.2.2 There is no record copy of k in *data_array* or *buf* of the old group at the time a commits
 - 2.3.2.1.1.2.3.2.2.1 There is no record reference in *data_array* of the new

- group that refers to a record copy of k at the time a commits.
 - By the fact that a commits at buf of the new group.
 - 2.3.2.1.1.2.3.2.2 $data_array$ of the new groups contains references of all the record copies in $data_array$ and buf of the old group.
 - By line 7 of Algorithm 3.
 - 2.3.2.1.1.2.3.2.3 Q.E.D.
 - 2.3.2.1.1.2.3.2.3 2.3.2.1.1.2.3.2.1 contradicts with 2.3.2.1.1.2.3.2.2.
 - By OBVIOUS.
 - 2.3.2.1.1.2.3.2.4 Q.E.D.
 - 2.3.2.1.1.2.3.3 **CASE 3:** get b commits at line 10 in the old group.
 - 2.3.2.1.1.2.3.3.1 a commits at buf of the new group, b commits at tmp_buf of the old group.
 - By the fact that a commits at line 21, b commits at line 10.
 - 2.3.2.1.1.2.3.3.2 buf of the new group and tmp_buf of the old group refers to the same delta index instance.
 - By line 8 of Algorithm 3.
 - 2.3.2.1.1.2.3.3.3 a, b can't commit simultaneously on the same delta index.
 - By the property of delta index implementation, 2.3.2.1.1.2.3.3.1, 2.3.2.1.1.2.3.3.2.
 - 2.3.2.1.1.2.3.3.4 Q.E.D.
 - 2.3.2.1.1.2.3.4 Q.E.D.
 - 2.3.2.1.1.2.4 Q.E.D.
 - 2.3.2.1.1.3 Q.E.D.
 - 2.3.2.1.2 Q.E.D.
 - 2.3.2.2 **CASE 2:** group split phase 1 \rightarrow phase 2.
 - 2.3.2.2.1 **CASE 1:** put a , get b can't commit in ga^* , gb^* simultaneously.
 - 2.3.2.2.1.1 g_a' and g_b' are created to have different key interval.
 - By line 5-6 of Algorithm 4.
 - 2.3.2.2.1.2 If a, b commit at g_a' and g_b' respectively, then they are not conflicting.
 - By 2.3.2.2.1.1.
 - 2.3.2.2.1.3 Q.E.D.
 - 2.3.2.2.2 **CASE 2:** put a , get b can't commit simultaneously in g_a' (the new group) and the old group respectively.
 - 2.3.2.2.2.1 **CASE 1:** put a commits in g_a' (the new group), get b commits in the old group.
 - 2.3.2.2.2.1.1 b can only commit at line 6 or line 8 in the old group.
 - 2.3.2.2.2.1.1.1 tmp_buf is not initialized in the old group.
- By the fact that groupsplit does not initialize the old group's $tmpbuf^*$.
- 2.3.2.2.2.1.1.2 b can't commit at line 10.
- By 2.3.2.2.2.1.1.1, line 9 of Algorithm 2.
- 2.3.2.2.2.1.1.3 Q.E.D.
- 2.3.2.2.2.1.2 **CASE 1:** b commits at line 6 in the old group.
- 2.3.2.2.2.1.2.1 a can only commit at line 18 or line 21 or line 23 in g_a'
- 2.3.2.2.2.1.2.1.1 tmp_buf is not initialized in g_a'
- By line 10-12 of Algorithm 4.
- 2.3.2.2.2.1.2.1.2 a can't commit at line 27 in g_a' .
- By 2.3.2.2.2.1.2.1.1, line 25-26 of Algorithm 2.
- 2.3.2.2.2.1.2.1.3 Q.E.D.
- 2.3.2.2.2.1.2.2 **CASE 1:** a commits at line 18 in g_a'
- 2.3.2.2.2.1.2.2.1 a commits at $data_array$ of g_a' , b commits a $data_array$ of the old group.
- By the fact that a commits at line 18, b commits at line 6.
- 2.3.2.2.2.1.2.2.2 a, b both operate on $data_array$ of the old group.
- By the fact that $data_array$ of g_a' refers to $data_array$ of the old group, 3.2.2.2.1.3.1.1.
- 2.3.2.2.2.1.2.2.3 a, b can't commit simultaneously in $data_array$
- By the fact that record copise in $data_array$ is protected by OCC, 2.3.2.2.2.1.2.2.2.
- 2.3.2.2.2.1.2.2.4 Q.E.D.
- 2.3.2.2.2.1.2.3 **CASE 2:** a commits at line 21 or line 23 in g_a'
- 2.3.2.2.2.1.2.3.1 There is no record copy of k in $data_array$ of the old group at the time a commits.
- 2.3.2.2.2.1.2.3.1.1 There is no record copy of k in $data_array$ of g_a' at the time a commits.
- By the fact that a commits at buf or tmp_buf of g_a' , Lemma 2 and Lemma 3.
- 2.3.2.2.2.1.2.3.1.2 $data_array$ of g_a' refers $data_array$ of the old group.
- By line 4 of Algorithm 4.
- 2.3.2.2.2.1.2.3.1.3 Q.E.D.
- 2.3.2.2.2.1.2.3.2 There exist a record copy of k in $data_array$ of the old group at the time b commits.
- By the fact that b commits at $data_array$ of the old group.

- 2.3.2.2.2.1.2.3.3 2.3.2.2.2.1.2.3.1 contradicts with 2.3.2.2.2.1.2.3.2.
 - By OBVIOUS.
- 2.3.2.2.2.1.2.3.4 Q.E.D.
- 2.3.2.2.2.1.2.4 Q.E.D.
- 2.3.2.2.2.1.3 **CASE 2:** b commits at line 8 in the old group.
 - 2.3.2.2.2.1.3.1 a can only commit at line 21 or line 23 in g_a'
 - 2.3.2.2.2.1.3.1.1 tmp_buf is not initialized in g_a'
 - By line 11-12 of Algorithm 4.
 - 2.3.2.2.2.1.3.1.2 a can't commit at line 27 in g_a' .
 - By 2.3.2.2.2.1.3.1.1, line 20 of Algorithm 2.
 - 2.3.2.2.2.1.3.1.3 There is no record copy of k in $data_array$ of the old group at the time b commits
 - By the fact that b commits at buf of the old group, Lemma 2.
 - 2.3.2.2.2.1.3.1.4 a can't commit at line 18 in g_a'
 - By 2.3.2.2.2.1.3.1.3.
 - 2.3.2.2.2.1.3.1.5 Q.E.D.
 - By 2.3.2.2.2.1.3.1.2, 2.3.2.2.2.1.3.1.4.
 - 2.3.2.2.2.1.3.2 a, b commits at the same buf
 - By that g_a' and the old group share the same buf , 2.3.2.2.2.1.3.1
 - 2.3.2.2.2.1.3.3 a, b can't commit simultaneously.
 - By 2.3.2.2.2.1.3.2, buf is concurrent safe index structure.
 - 2.3.2.2.2.1.3.4 Q.E.D.
- 2.3.2.2.2.1.4 Q.E.D.
- 2.3.2.2.2.2 **CASE 2:** put a commits in the old group, get b commits in $*g_a'$ (*the new group*)*
 - 2.3.2.2.2.2.1 a can only commit at line 18 or line 21 in the old group.
 - 2.3.2.2.2.2.1.1 a can't commit at line 23 in the old group.
 - By the fact that buf_frozen is false in the old group, line 20,22 of Algorithm 2.
 - 2.3.2.2.2.2.1.2 a can't commit at line 27 in the old group.
 - By the fact that tmp_buf is not initialized in the old group, line 25-26 of Algorithm 2.
 - 2.3.2.2.2.2.1.3 Q.E.D.
 - 2.3.2.2.2.2.2 **CASE 1:** a commits at line 18 in the old group.
 - 2.3.2.2.2.2.2.3.1 There exists a record copy of k in $data_array$ of the old group at the time a commits.
 - By the fact that a commits at $data_array$ of the old group.
 - 2.3.2.2.2.2.2.2 b will commit at line 6.
 - By 2.3.2.2.2.2.3.1, the old group and g_a' share the same $data_array$.
 - 2.3.2.2.2.2.2.3 a,b both operate on $data_array$ of the old group.
 - By the fact that $data_array$ of g_a' refers to $data_array$ of the old group, 2.3.2.2.2.2.2.2.
 - 2.2.3.2.2.2.2.4 a,b can't commit simultaneously in the same $data_array$.
 - By the fact that record copy in $data_array$ is protected by OCC, 2.3.2.2.2.2.2.2.
 - 2.2.3.2.2.2.2.4 Q.E.D.
 - 2.3.2.2.2.2.3 **CASE 2:** a commits at line 21 in the old group.
 - 2.3.2.2.2.2.3.1 There is no record copy of k in $data_array$ of the old group at the time a commits.
 - By the fact that a commits at buf of the old group.
 - 2.3.2.2.2.2.3.2 b will not commit at $data_array$.
 - By 2.3.2.2.2.3.1.
 - 2.3.2.2.2.2.3.3 a, b both commit at the same buf .
 - By 2.3.2.2.2.2.3.2, g_a' and the old group share the same buf .
 - 2.3.2.2.2.2.3.4 a, b can't commit simultaneously.
 - By 2.3.2.2.2.2.3.3, buf is a concurrent safe index structure.
 - 2.3.2.2.2.2.3.5 Q.E.D.
 - 2.3.2.2.2.2.4 Q.E.D.
 - 2.3.2.2.2.3 Q.E.D.
 - 2.3.2.2.3 **CASE 3:** a,b can't commit in g_b' , the old group simultaneously.
 - By symmetry, proved by 2.3.2.2.2.
 - 2.3.2.2.4 Q.E.D.
 - 2.3.2.3 **CASE 3:** group split phase 2 \rightarrow phase 3.
 - 2.3.2.3.1 $g_a' \rightarrow g_a$ and $g_b' \rightarrow g_b$ are two compaction process.
 - By design of structure update.
 - 2.3.2.3.2 a,b can't commit simultaneously.
 - By 2.3.2.1.
 - 2.3.2.3.3 Q.E.D.
 - 2.3.2.4 Q.E.D.
 - 2.3.3 Q.E.D.
 - 2.4 Q.E.D.
 - 3 Q.E.D.

Lemma 5. (I2) For any possible concurrent execution of XIndex, for any key k , if there is one record copy of k in XIndex, then its value equals the value of the last committed put on k .

Proof of Lemma 5.

- 1 It suffices to prove that 1) if a put(k,v) commits, then there is a one and only one record copy of k whose value equals this put's value; and 2) compaction and `group_split` do not make a record copy stale or lost.
 - 1.1 The value of a record copy is changed by put, compaction and `group_split`.
By OBVIOUS.
 - 1.2 Q.E.D.
- 2 If a put(k,v) commits, then there is a one and only one record copy of k whose value equals this put's value.
 - 2.1 If a put(k,v) commits, then there is a record copy of k whose value equals this put's value.
 - 2.1.1 If there is no record copy of k in XIndex before put commits, then a new record copy of k will be created.
By OBVIOUS.
 - 2.1.2 If there is a record copy of k in XIndex before put commits, then its value will be successfully updated to the put's value.
 - 2.1.2.1 put commits at either line 18, 21, 23 and 27, Algorithm 2.
By Definition 1.
 - 2.1.2.2 **CASE 1:** put commits at line 18.
 - 2.1.2.2.1 There is a record copy in `*data_array*`.
By OBVIOUS.
 - 2.1.2.2.2 The value of the record copy will be successfully updated to the put's value.
By the use of lock that ensures atomicity and exclusive access to the record copy and the implementation of `update_record` that updates the record copy's value.
 - 2.1.2.2.3 Q.E.D.
 - 2.1.2.2.3 **CASE 2:** put commits at line 21.
 - 2.1.2.2.3.1 There is a record copy in `*buf*`.
By OBVIOUS.
 - 2.1.2.2.3.2 The value of the record copy will be successfully updated to the put's value.
By the implementation of `insert_buffer` that ensures atomicity and updates the record copy's value.
 - 2.1.2.2.3.3 Q.E.D.
 - 2.1.2.2.4 **CASE 3:** put commits at line 23.
 - 2.1.2.2.4.1 There is a record copy in `*buf*`.
By OBVIOUS.
 - 2.1.2.2.4.2 The value of the record copy will be successfully updated to the put's value.
By the implementation of `try_update_in_buffer`

that ensures atomicity and updates the record copy's value.

- 2.1.2.4.3 Q.E.D.
- 2.1.2.5 **CASE 4:** put commits at line 27.
 - 2.1.2.5.1 There is a record copy in `*tmp_buf*`.
By OBVIOUS.
 - 2.1.2.5.2 The value of the record copy will be successfully updated to the put's value.
By the implementation of `insert_buffer` that ensures atomicity and updates the record copy's value.
 - 2.1.2.5.3 Q.E.D.
 - 2.1.2.6 Q.E.D.
 - 2.1.3 Q.E.D.
- 2.2 If a put(k,v) commits, there is a one and only one record copy of k in XIndex.
By Lemma 1.
- 2.3 Q.E.D.
- 3 compaction and `group_split` do not make a record copy stale or lost.
 - 3.1 During phase transition, a record copy will not be lost or stale.
 - 3.1.1 All record copies of the old group can be accessed in the new group.
 - 3.1.1.1 All record copies of the old group are referenced in the new group.
 - 3.1.1.1.1 **CASE 1:** compaction's phase transition.
 - 3.1.1.1.1.1 Old group's record copies in `data_array` and `buf` are referenced in new group's `data_array`.
By line 7, Algorithm 3.
 - 3.1.1.1.1.2 Old group's record copies in `tmp_buf` are referenced in new group's `buf`.
By line 8, Algorithm 3.
 - 3.1.1.1.1.3 Q.E.D.
 - 3.1.1.1.1.2 **CASE 2:** `group_split`'s phase transition.
 - 3.1.1.1.1.2.1 **CASE 1:** `group_split`'s first phase transition.
 - 3.1.1.1.1.2.1.1 Old group's record copies in `data_array` and `buf` are referenced in new group's `data_array` and `buf`.
By line 4, Algorithm 4.
 - 3.1.1.1.1.2.1.2 Q.E.D.
 - 3.1.1.1.1.2.2 **CASE 2:** `group_split`'s second phase transition.
 - 3.1.1.1.1.2.2.1 Old group's record copies in `data_array` and `buf` are referenced in new group's `data_array`.
By line 17, Algorithm 4.
 - 3.1.1.1.1.2.2.2 Old group's record copies in `tmp_buf` are referenced in new group's

buf.

By line 18, Algorithm 4.

- 3.1.1.1.2.2.3 Q.E.D.
- 3.1.1.1.2.3 Q.E.D.
- 3.1.1.1.3 Q.E.D.
- 3.1.1.2 Q.E.D.
- 3.1.2 Modification to record copies can be observed in both old and new group.
 - By 3.1.1.
 - 3.1.3 Q.E.D.
- 3.2 Outside of phase transition, a record copy will not be lost or stale.
 - 3.2.1 compaction and `group_split` do not remove or update (change value) existing record copies.
 - 3.2.1.1 It suffices to prove that compaction's copy phase and `group_split`'s third phase does not update existing record copies.
 - 3.2.1.1.1 Only compaction's copy phase and `groupsplit`'s third phase update references in `*dataarray*` to concrete value.
 - By the design of our algorithm.
 - 3.2.1.1.2 Q.E.D.
 - 3.2.1.2 compaction's copy phase and `group_split`'s third phase does not update (change value) existing record copies.
 - 3.2.1.2.1 Changing references in `data_array` to concrete values does not change record copies' value.
 - 3.2.1.2.1.1 References in `data_array` are atomically changed to concrete values during compaction's copy phase and `group_split`'s third phase.
 - By using lock to protect the change at line 14-31, Algorithm 5.
 - 3.2.1.2.1.2 The concrete value are copied to the place of previous reference.
 - By line 26, Algorithm 5.
 - 3.2.1.2.1.3 Q.E.D.
 - 3.2.1.2.2 Q.E.D.
 - 3.2.1.3 Q.E.D.
 - 3.2.2 Q.E.D.
 - 3.3 Q.E.D.
 - 4 Q.E.D.

Lemma 6. (I3) For any possible concurrent execution of XIndex, for any key k , for `get` with key k , if there is one record copy of k in XIndex when `get` commits, then `get` returns the value of the record copy.

Proof of Lemma 6.

- 1 If a record copy exists, it is in either `*data_array*`, `*buf*` or `*tmp_buf*` of the group returned by `get_group`.
 - 1.1 **CASE 1:** during phase transition of compaction and `group_split`.
 - 1.1.1 During phase transition of compaction and `group_split`, if a record copy exists in old group,

then it is referenced in the new group.

By line 7-8, Algorithm 3; line 4, 16-18, Algorithm 4.

- 1.1.2 Q.E.D.
- 1.2 **CASE 2:** outside of phase transition of compaction and `group_split`.
 - By OBVIOUS.
 - 1.3 Q.E.D.
- 2 **CASE 1:** record copy is in `*data_array*` when `get` commits.
 - 2.1 `get` can read the record copy from `*data_array*`.
 - 2.1.1 `get` will access `*data_array*`.
 - By line 3-6, Algorithm 2.
 - 2.1.2 `get` can read the record copy from `*data_array*`.
 - By the implementation of `read_record` that uses OCC to atomically read the record copy.
 - 2.1.3 Q.E.D.
 - 2.2 Q.E.D.
- 3 **CASE 2:** record copy is in `*buf*` when `get` commits.
 - 3.1 `get` reads the value of the record copy and commits at line 10, Algorithm 2.
 - 3.1.1 `get` will execute line 7-8, Algorithm 2.
 - 3.1.1.1 `get` cannot read record copy from `*data_array*`.
 - 3.1.1.1.1 Otherwise `get` commits at line 6 and `*data_array*` then contains the record copy, which contradicts 3.
 - By Lemma 1.
 - 3.1.1.1.2 Q.E.D.
 - 3.1.1.2 Q.E.D.
 - 3.1.2 When `get` executes line 7-8, Algorithm 2, it will read the value of the record copy.
 - 3.1.2.1 If `get` does not read the value of the record, then there is a contradiction.
 - 3.1.2.1.1 If `get` does not read the value of the record, it will check `tmp_buf`.
 - By line 9, Algorithm 2.
 - 3.1.2.1.2 **CASE 1:** `tmp_buf` is not initialized.
 - 3.1.2.1.2.1 `get` will commit and returns empty at line 8, Algorithm 2.
 - By Definition 1.
 - 3.1.2.1.2.2 `buf` does not contain record copy when `get` commit.
 - By 3.1.2.1.2.1, Definition 1 and the implementation of the concurrent delta buffer that ensures `read_buffer` returns empty when there is no record when commits.
 - 3.1.2.1.2.3 3.1.2.1.2.1 contradicts 3.
 - By OBVIOUS.
 - 3.1.2.1.2.4 Q.E.D.
 - 3.1.2.1.3 **CASE 2:** `tmp_buf` is initialized.
 - 3.1.2.1.3.1 `buf` does not contain record copy when `get` commit.

- 3.1.2.1.3.1.1 *get* failed to read record from *buf* during execution.
By 3.1.2.1.1.
- 3.1.2.1.3.1.2 *buf* is frozen and cannot be inserted during *get*.
 - 3.1.2.1.3.1.2.1 *tmp_buf* is initialized after all threads observe frozen *buf*.
By line 3-4, Algorithm 3; line 11-13, Algorithm 4.
 - 3.1.2.1.3.1.2.2 Q.E.D.
 - 3.1.2.1.3.1.3 Q.E.D.
- 3.1.2.1.3.2 3.1.2.1.3.1 contradicts 3.
By OBVIOUS.
- 3.1.2.1.3.3 Q.E.D.
- 3.1.2.1.4 Q.E.D.
- 3.1.2.2 Q.E.D.
- 3.1.3 Q.E.D.
- 3.2 Q.E.D.
- 4 **CASE 3:** record copy is in **tmp_buf** when *get* commits.
 - 4.1 No record copy of *k* in **data_array** and **buf**.
By Lemma 3.
 - 4.2 *get* reads the value of the record copy and commits at line 10, Algorithm 2.
 - 4.2.1 *get* will execute line 9-10, Algorithm 2.
 - 4.2.1.1 *get* will not get record from **data_array**
 - 4.2.1.1.1 Otherwise *get* commits at line 6 and **data_array** then contains the record copy, which contradicts 3.
By OBVIOUS.
 - 4.2.1.1.2 Q.E.D.
 - 4.2.1.2 *get* will not get record from **buf**
 - 4.2.1.2.1 Otherwise *get* commits at line 8 and **buf** then contains the record copy, which contradicts 3.
By OBVIOUS.
 - 4.2.1.2.2 Q.E.D.
 - 4.2.1.3 Q.E.D.
 - 4.2.2 When *get* executes line 9-10, Algorithm 2, it will read the value of the record copy.
 - 4.2.2.1 If *get* does not read the value of the record, then there is a contradiction.
 - 4.2.2.1.1 If *get* does not read the value of the record, *get* will commit and returns empty at line 10, Algorithm 2.
By Definition 1.
 - 4.2.2.1.2 If *get* will commit and returns empty at line 10, Algorithm 2, there is no record copy in **tmp_buf** when *get* commits.
By Definition 1 and the implementation of the concurrent delta buffer that ensures *read_buffer* returns empty when there is no record when commits.

- 4.2.2.1.3 There is a record copy in **tmp_buf** when *get* commits.
By 4.
- 4.2.2.1.4 Q.E.D.
- 4.2.2.2 Q.E.D.
- 4.2.3 Q.E.D.
- 4.3 Q.E.D.
- 4 Q.E.D.

Lemma 7. For any possible concurrent execution of XIndex *E*, for any key *k*, for any committed *get* operation $op_0 \in E$, let $S = \{op \mid op \in E, op \text{ is committed put before } op_0, op \text{ operates on } k\}$, if $|S| \neq 0$, then $\exists op \in S: \forall op' \in S/\{op\}, op' \ll op$.

Proof of Lemma 7.

- 1 If negation of Lemma 7 holds, then Lemma 4 is violated.
- 1.1 If negation of Lemma 7 holds, then for some possible concurrent execution of XIndex *E*, for some key *k*, for some committed *get* operation $op_0 \in E$, let $S = \{op \mid op \in E, op \text{ is committed put before } op_0, op \text{ operates on } k\}$, if $|S| \neq 0$, then $\forall op \in S: op' \in S/\{op\}, op' \ll op$.
By OBVIOUS.
- 1.2 If for some possible concurrent execution of XIndex *E*, for some key *k*, for some committed *get* operation $op_0 \in E$, let $S = \{op \mid op \in E, op \text{ is committed put before } op_0, op \text{ operates on } k\}$, if $|S| \neq 0$, then $\forall op \in S: \exists op' \in S/\{op\}, op' \ll op$; then there is a circle consist of puts $op_1 \in S, \dots \ll op_i \ll \dots \ll op_j \ll \dots \ll op_i \ll \dots$.
 - 1.2.1 If such circle does not exist, then the latest committed put exists.
By OBVIOUS.
 - 1.2.2 Q.E.D.
- 1.3 If such sequence exists, then $op_i \ll op_j$ and $op_j \ll op_i$.
By OBVIOUS.
- 1.4 If $op_i \ll op_j$ and $op_j \ll op_i$, then Lemma 4 is violated.
By OBVIOUS.
- 1.5 Q.E.D.
- 2 Q.E.D.

Definition 5. (latest committed put) For any possible concurrent execution of XIndex *E*, for any key *k*, for any committed *get* operation $op_0 \in E$, let $S = \{op \mid op \in E, op \text{ is committed put before } op_0, op \text{ operates on } k\}$. If $|S| \neq 0$, we say $op \in S$ is the latest committed put on *k* if $\forall op' \in S/\{op\}, op' \ll op$.

Theorem 1. (correctness condition) For any possible concurrent execution of XIndex *E*, for any key *k*, for *get* with key *k*, if there is put committed before *get* commits, then *get* returns the value of the last committed put on *k* before the *get*.

Proof of Theorem 1.

- 1 if there is put committed before get commits, then there is a record copy in XIndex.

By Lemma 1.

- 2 If there is a record copy in XIndex, then get returns the value of the latest committed put.

By Lemma 5, 6.

- 3 Q.E.D.

Definition 6. (sequential specification) XIndex x is an object supporting put and get operations. The state of XIndex is a set of records (key-value pairs) $x = \{<k_0, v_0>, \dots, <k_n, v_n>\}$ and is initially empty. The put and get operations induce the following state transitions of x with appropriate return values:

1. get(k), if there is a pair $<k_i, v_i>$ exists in x that $k_i=k$, returns v ; else returns empty; x always remains unchanged;
2. put(k, v), if there is a pair $<k_i, v_i>$ exists in x that $k_i=k$, changes x to $x \setminus \{<k_i, v_i>\} \cup \{<k_i, v>\}$; else changes x to $x \cup \{<k, v>\}$; put has no return value;

Theorem 2. (linearizability) For any possible concurrent execution of XIndex E , for any key k , the following holds:

- operations on k can be serialized by the linearization point; and.
- for operations a and b , if a finishes before b starts, then a is serialized before b .

Proof of Theorem 2.

- 1 Operations on k can be serialized by the linearization point.
 - 1.1 It suffices to prove that conflicting operations on k can be serialized by the linearization point. By OBVIOUS.
 - 1.2 Conflicting operations on k can be serialized by the linearization point.
 - 1.2.1 For conflicting operations a and b , a is either serialized before or after b .
 - 1.2.1.1 Conflicting operations a and b , either $a \ll b$ or $b \ll a$.
By Lemma 4.
 - 1.2.1.2 Q.E.D.
 - 1.2.2 The value get returns in the serialized execution is the same as its return value in the concurrent execution.
 - 1.2.2.1 **CASE 1:** get returns empty in the serialized execution.
 - 1.2.2.1.1 No put is serialized before get.
By the Definition 6.

- 1.2.2.1.2 If no put is serialized before get, then no put commits before get.

By the fact that operations are serialized according to linearization point.

- 1.2.2.1.3 If no put commits before get, get returns empty in concurrent execution.

By OBVIOUS.

- 1.2.2.1.4 Q.E.D.

- 1.2.2.2 **CASE 2:** get returns non-empty in the serialized execution.

- 1.2.2.2.1 get returns the value of the put serialized nearest before the get.

By the Definition 6.

- 1.2.2.2.2 If a put is serialized nearest before the get, then the put is the last committed put before the get.

By the fact that operations are serialized according to linearization point and Definition 5.

- 1.2.2.2.3 If there is put committed before get commits, get returns the value of the last committed put on k before the get in concurrent execution.

By Theorem 1.

- 1.2.2.2.4 Q.E.D.

- 1.2.2.3 Q.E.D.

- 1.2.3 Q.E.D.

- 1.3 Q.E.D.

- 2 For operations a and b , if a finishes before b starts, then a is serialized before b .

- 2.1 It suffices to prove that, for conflicting operations a and b , if a finishes before b starts, then a is serialized before b .

By OBVIOUS.

- 2.2 For conflicting operations a and b , if a finishes before b starts, then a is serialized before b .

- 2.2.1 If a finishes before b starts, then $a \ll b$.

- 2.2.1.1 a and b both commit after operation starts and before operation ends.

By Definition 1.

- 2.2.1.2 Q.E.D.

- 2.2.2 If $a \ll b$, then a is serialized before b .

By OBVIOUS.

- 2.2.3 Q.E.D.

- 2.3 Q.E.D.

- 3 Q.E.D.