

A Fully On-disk Updatable Learned Index

Hai Lan¹, Zhifeng Bao¹, J. Shane Culpepper², Renata Borovica-Gajic³, Yu Dong⁴

¹RMIT University, ²The University of Queensland, ³The University of Melbourne, ⁴PingCAP

hai.lan@student.rmit.edu.au, zhifeng.bao@rmit.edu.au, s.culpepper@uq.edu.au

renata.borovica@unimelb.edu.au, yu.dong@pingcap.com

Abstract—While in-memory learned indexes have shown promising performance as compared to B+-tree, most widely used databases in real applications still rely on disk-based operations. From our experiments, we observe that directly applying the existing in-memory learned indexes into on-disk setting suffers from several drawbacks and cannot outperform a standard B+-tree in most cases. Therefore, we make the first attempt to show how the idea of learned index can benefit the on-disk index by proposing **AULID**, a fully on-disk updatable learned index that can achieve state-of-the-art performance across multiple workload types. The **AULID** approach combines the benefits from both traditional indexing techniques and the learned indexes to reduce the I/O cost – the main overhead under disk setting. Specifically, three aspects are taken into consideration in reducing I/O costs: (1) reduce the overhead in updating the index structure; (2) induce shorter paths from root to leaf node; (3) achieve better locality to minimize the number of block reads required to complete a scan. Five principles are proposed to guide the design of **AULID** which shows remarkable performance gains and meanwhile is easy to implement. Our evaluation shows that **AULID** has comparable storage costs to a B+-tree and is much smaller than other learned indexes, and **AULID** is up to 2.11x, 8.63x, 1.72x, 5.51x, and 8.02x more efficient than FITing-tree, PGM, B+-tree, ALEX, and LIPP.

I. INTRODUCTION

Nowadays, most widely used database systems still rely on on-disk indexing techniques for (at least) two reasons. First, the total index size may be larger than the main memory available – a consequence of growing data sizes in real applications [1]. Also, multiple indexes (not just one index) might be built to optimize workload-specific performance [2]; they are usually operationalized as a “secondary index”, where the leaf nodes should be included when calculating the total storage requirements. Second, main memory is also a precious resource for efficient query processing to store intermediate results, e.g., a hash table in a hash join [3]. If most of the available memory is used to hold the indexes, query performance could be significantly degraded.

When adapting the idea of learned indexes to a fully on-disk setting, most of their techniques are no longer useful since I/O costs are the dominant bottleneck. For example, when issuing a lookup query in a four-layer B+-tree, we find 93.6% of the total execution time is spent on I/O operations. Hence, reducing the number of block reads (and writes) is crucial.

An immediate question to ask then is how existing learned indexes perform on disk? To answer that, we adopt the

implementation of four state-of-the-art updatable learned indexes [4, 5, 6, 7] on disk in recent experimental study [8] and compared them against a standard B+-tree across six workload types commonly encountered by a database. Figures 1(a)-(b) present the normalized throughput on *COVID* and *FB*, which are representative of easy and hard datasets respectively [9]. We observe that although these learned indexes exhibit different strengths and weaknesses depending on the workload type and dataset distribution, *none of them outperforms or achieves competitive performance to the B+-tree across all workload types on any dataset*. This should come as no surprise to database designers, given that most research on learned indexes has focused on in-memory performance. The benefit of learned indexes in main memory and the shortage of current learned index on disk motivate us to develop a high-performing on-disk learned index.

A. Challenges in Building a Fully On-disk Learned Index

- **Challenge 1: a learned index cannot guarantee to reduce I/O costs when searching data on disk.** Figure 1(c) shows the average number of inner nodes, inner blocks, and total blocks per query for Lookup-Only and Scan-Only workloads on the *FB* dataset¹. For LIPP [7], the total number of fetched nodes is reported, and the number of nodes in the scan is highlighted in the bracket of the fifth row. When combining the results for a Lookup-Only workload and a Scan-Only workload in Figure 1(b), we observe that the performance rank is directly correlated with the number of fetched blocks. In contrast to in-memory indexes, reducing the search overhead for each step does not help on-disk indexes. Instead, reading or writing a block from/to disk is the main overhead. In a Lookup-Only workload, among all the learned indexes, only LIPP is more efficient than B+-tree. We also find that, in contrast to B+-tree, existing learned indexes have larger scan overheads, which means fetching the *next* item becomes more expensive. For example, to support a scan, LIPP traverses many nodes, which incurs a higher I/O cost and leads to poorer performance.

- **Challenge 2: large insertion overheads.** Current indexing techniques support an insertion using four steps: (1) find a slot to hold a new key-payload pair (**Search**); (2) do the insertion (**Insert**); (3) induce an index structure modifications operation (**SMO**) if necessary; (4) update various statistics (**Stats**), such

¹For a Scan-Only workload, we set the start key to the same key that was used in the Lookup-Only workload, and then we scan forward 99 keys. This ensures that ALEX, PGM, a FITing-tree, and a B+-tree fetch the same number of inner nodes and blocks needed for a lookup and a scan.

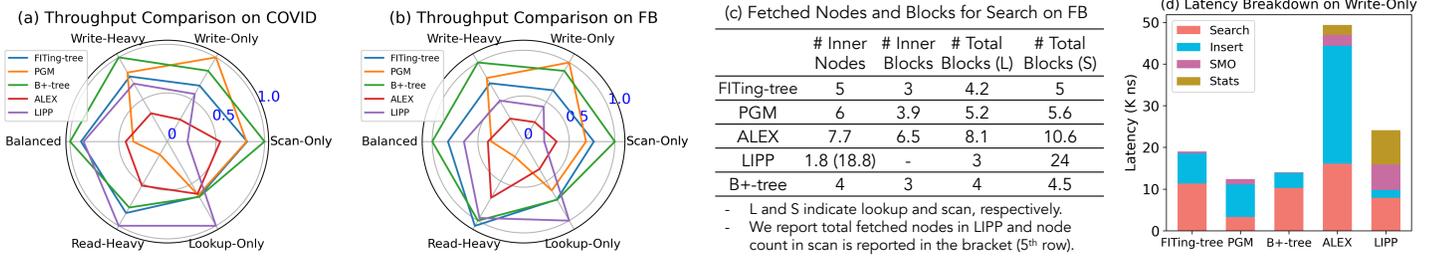


Fig. 1: Throughput Comparison and Analysis. Each index’s throughput in (a)-(b) is normalized by the largest under the same workload (higher is better), (c) is an analysis on the fetched blocks per query, and (d) is a latency breakdown per query.

as the total number of lookups and insertions, which determine when to induce SMOs. An SMO may create new nodes or re-construct an entire sub-tree during the insertion. This is necessary for the index to allocate empty slots, and for a learned index to benefit from future model-based operations. Figure 1(d) shows the average latency breakdown per query for a Write-Only workload on the *FB* dataset. We observe that existing learned indexes have several shortcomings. ALEX [4] and LIPP have a large overhead when updating statistics and performing an SMO. The FITing-tree [5], PGM [6], and ALEX incur a large overhead for the insertion. LIPP frequently induces an SMO to resolve conflicts between two keys, while ALEX re-writes large nodes (the leaf node) in every SMO. Both lead to large SMO overheads. The shift operations, which are used to obtain an empty slot to store a new key-payload pair in a FITing-tree, PGM, and ALEX, may span multiple blocks on disk – leading to more writes on disk.

B. AULID – Simple Is Better

In this paper, we show that the idea of learned index can benefit the on-disk index design by proposing AULID, an updatable learned index on disk. To address the above challenges, we propose five design principles for an on-disk learned index in §II-A. Then, we design AULID to meet them by leveraging the idea of B+-tree and learned indexes in three ways (an overview from the challenges to our design principles and solutions is in Figure 2):

❶ **Leaf Node Layout.** Instead of using a learned model to search items for all layers, we use model-based search only for inner nodes. This helps reduce the burden on leaf nodes in maintaining the benefits of model-based search. Specifically, we use a B+-tree styled layout for leaf nodes, which has a low overhead when updating the index. Since the majority of SMOs are on leaf nodes, a lightweight SMO mechanism for leaf nodes, as achieved with AULID, reduces the insertion overheads significantly. A B+-tree styled leaf node design also benefits scan operations in fetching the *next* item.

❷ **Inner Node Layout.** After building a B+-tree styled leaf node, the path from the root node to a leaf node should be shorter than that in a B+-tree. Otherwise, the learned index cannot outperform a B+-tree for on-disk operations. The results in Figure 1(a)-(b) inspired us to adopt the Fastest Minimum Conflict Degree (FMCD) algorithm in LIPP [7] to reduce the tree depth. Although it is not the best indexing

method for most workload types, it has the smallest number of fetched blocks for a lookup, making it suitable for on-disk indexes when attempting to reduce I/O costs. Moreover, a lookup is often the first step in other operations, e.g., scan and insert. Thus, a better performance of the lookup operation should boost their performance. However, for certain workload types and datasets, e.g., *OSM* (a hard dataset) [9], directly applying the aforementioned inner node layout still reveals several shortcomings – a larger storage size and a lower throughput. To overcome these shortcomings, we introduce two new inner node types and design processing algorithms upon the new layout. For example, with our design, AULID is 1.18x more efficient on the Lookup-Only workload while only taking 1.13x storage on *OSM* (as compared to a B+-tree).

❸ **Structural Modification Operations.** With a B+-tree styled leaf node, AULID already manages to achieve a lighter overhead in modifying the index structure due to a lower frequency in updating the inner nodes and a lighter overhead in updating the leaf node. However, the tree height in some region could grow and even become larger than that of a standard B+-tree. In turn, AULID will have a worse performance after lots of insertions. To avoid that from happening, we monitor the tree height of each branch and trigger a re-construct process to bound their tree height if needed.

In summary, we make the following technical contributions:

- To the best of our knowledge, AULID is the first approach to employ the ideas of learned index to a fully on-disk setting to replace a traditional B+-tree².
- We propose five principles to guide the design of AULID (§II-A), and carefully design the indexing structure (§II-B - II-C), the query processing algorithms (§III-A - III-C), and an SMO mechanism, to achieve efficient reads and writes (§III-D).
- We implement AULID in C++ [12], conduct comprehensive experiments across a wide range of datasets and workloads and compare it against a B+-tree and on-disk learned indexes [8]. Our evaluation shows that AULID has competitive storage costs to a B+-tree and is much smaller than most other learned indexes. Performance-wise, AULID achieves up to 2.11x, 8.63x, 1.72x, 5.51x, and 8.02x larger throughput than FITing-tree, PGM, B+-tree, ALEX, and LIPP, respectively. We also conduct an

²Two recent on-disk studies [1, 10] are built upon LSM tree [11] and suffer from poor read performance. Details are presented in §V.

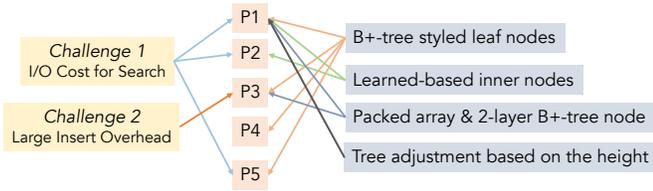


Fig. 2: An Overview of AULID: from Challenges to Design Principles and Solutions

in-depth evaluation on the benefits of the AULID design (§IV).

II. AN OVERVIEW OF AULID

We first introduce the principles and highlights of the AULID design – addressing the challenges discussed in §I-B. Then, we present the AULID layout. Figure 2 presents two challenges identified, the design principles used to resolve them, and how these principles are reflected in our proposal.

A. Design Principles

Based on the key properties on disk and learned indexes, we propose five principles to guide the design of AULID:

- **P1. Reducing the Tree Height of the Index.** Accessing each level in an index requires at least one disk access when an index is stored on disk. Reducing the tree height can reduce the number of disk access.
- **P2. Model-based Operations (Search and Insert).** An index with a reduced height usually has larger nodes in certain levels of the index. Model-based operations help AULID quickly find search keys in a specific part of the node, without the need to access the entire node on disk.
- **P3. Lightweight Structure Modification Operations.** Structure modification operations (SMOs) for the existing learned indexes incur a substantial amount of writes on disk. AULID should reduce the overhead of such SMO calls.
- **P4. Support Duplicate Index Keys.** Duplicate (i.e. non-unique) index keys are common in real systems. They can be supported using a linked list in a main memory setting [7], but not on disk, since it leads to additional disk reads.
- **P5. Better Scan Performance.** Existing learned indexes have their own limitations when supporting scans on disk (see Figure 1(a)-(b)). AULID must provide a lightweight method to fetch the *next* item efficiently.

B. Design Highlights

AULID uses a combination of existing and novel techniques to meet the above principles and achieve high performance on disk. AULID consists of inner nodes and leaf nodes, both of which are stored on disk. Leaf nodes, where most SMOs occur, are organized in a B+-tree manner. A low update cost at leaf nodes reduces the SMO overhead (P3). Moreover, AULID only uses the idea of learned index for inner nodes to index the maximum key of each leaf node, which leads to less frequent SMOs in updating the inner nodes (P3) and a low tree height in inner part (P1). Each leaf node is a

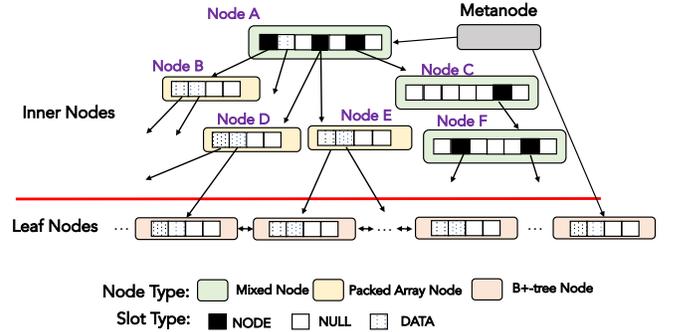


Fig. 3: AULID Index Structure

packed array – it stores pointers to its siblings and its size is equal to the *block size*. Using the packed array and links to siblings, AULID can support efficient scan operations (P5). We optimize our inner nodes based on properties of the disk drive. Fast lookup time with the learned model means that AULID can efficiently locate target leaf nodes (P1, P2). To achieve robust performance on different datasets (i.e., different distributions), we also introduce two new node types for the inner nodes, a packed array and a two-layer B+-tree, with the purpose of reducing the number of SMOs for non-leaf nodes (P3). By proposing a tailored mechanism to handle duplicate keys inserted in inner nodes, AULID manages to store duplicate values with reduced on-disk costs (P4). To maintain the performance gains achieved from the learned model, AULID adjusts the index structure based on the tree height and bounds the tree height during insertions (P1).

C. Node Structure

The index structure of AULID is presented in Figure 3. AULID is composed of two components: the inner nodes which store the route information to leaf nodes, and the leaf nodes which store the key-payload pairs.

1) *Metanode*: Metanode in AULID stores (1) the physical address of the root node, (2) the linear model of the root node, and (3) the physical address of the last leaf node, as well as the minimum and maximum keys of that node. We store the metanode in main memory, which requires only 80 bytes, a negligible main memory overhead.

2) *Inner Nodes*: AULID has two node types in the inner part, a mixed node type and a packed array node type. And there are three types of slot in the inner part: *NODE*, *NULL*, and *DATA*. The *NODE* slot stores the pointer to its child. The *NULL* slot is the empty slot and can be converted to *NODE* or *DATA*. The *DATA* slot stores the key-payload pair. Each mixed node has a model to predict the slots for a key search and can include three different slot types above. AULID stores the model in the parent node, combined with the physical address. If we store the model at the starting address of a mixed node, the large fanout for mixed nodes increases the chance that the predicted position and the model are located in different blocks. Thus, two blocks must be fetched from disk for each level in the tree. In contrast, when storing the model in the parent node, AULID only fetches one block per level.

The *NODE* slot in AULID can be further divided into three types, as shown in Figure 3: (1) a pointer to the packed array of fixed size (e.g. the first slot in Node A); (2) a small B+-tree node (e.g. the fourth slot in Node A). It contains at most four child nodes, each of which is a leaf node of a B+-tree. Such a B+-tree node and its children form a two-layer B+-tree; (3) a pointer to another mixed node (e.g. the sixth slot in Node A). For the first case, we introduce four different packed array types, each with a fixed size. The i^{th} packed array type can store 2^{i+2} items of *DATA* types, where $i \in [1, 4]$. A *DATA* slot in the inner nodes stores the physical address of a leaf node and the largest key it contains, i.e., the key-block pair. The second case is proposed to improve the performance for scenarios where the number of keys to be inserted into the same slot is greater than 64, but smaller than 1020 (which will be explained later in §II-C2). This indicates the *conflict degree* for the region. If we create a new mixed node to hold these keys, there can be key conflicts in the new node. This leads to a larger tree height (more than two layers). Conversely, a two-layer B+-tree can be used to hold the nodes and help AULID to bound the tree height for the region. Therefore, AULID is able to bound the number of fetched blocks. Also, the design of the packed array and the two-layer B+-tree can support more newly inserted key-block pairs (the routing knowledge to the leaf nodes) to be stored with low overheads. Using the packed array and a two-layer B+-tree, AULID achieves a better empty slot ratio, and this translates to smaller storage costs. In our implementation, we store metadata information for each node, e.g., its physical position, and node type, in its parent node. We use an 8-bit number to distinguish different nodes.

The B+-tree node contains only four child nodes for the following reasons: (1) The conflict degrees in most of the test datasets we have used (except for one)³ is less than 1000, which can be stored easily using a two-layer B+-tree. (2) A larger fanout requires more metadata (pivot keys and physical addresses) in a slot, and it increases the total storage cost significantly. A key-block pair occupies 16 bytes on disk in our implementation. Thus, a block with 4 KB can store 256 pairs. The first item records the item count for a two-layer B+-tree’s leaf node. Four children can store at most 1020 items.

3) *Leaf Nodes*: The leaf nodes have the same structure as a standard B+-tree. The *DATA* slot in the leaf node stores the key-payload pairs to be indexed. This layout design is based on the observation that most SMOs happen on leaf nodes as new key-payload pairs are added. Learned indexes need to read all of the items in a large leaf node and re-write them to disk to maintain the benefits of their unique structure, which incurs large I/O costs on disk. A lightweight SMO overhead for a leaf node design can help significantly reduce the number of SMO operations required (see the experiments in §IV-B2).

This simple design in the leaf nodes has many other benefits. First, the link between siblings, when using a packed storage layout, requires no additional utility structure to perform efficiently when scan operations must locate the start of a

query range. Second, the storage costs of the inner nodes can be significantly reduced by only storing the largest keys. In our experiments, AULID has a similar storage size and bulkload time to a B+-tree on disk, which is better than other learned indexes. Third, reducing the items inserted into the inner nodes also decreases the SMO frequency and the number of items that must be processed. Last, AULID can efficiently support duplicate index keys when using a B+-tree styled leaf node.

III. AULID OPERATIONS

A. Bulkload

AULID supports bulkload using two steps. In the first step, it creates leaf nodes to store the key-payload pairs using B+-tree styled leaf nodes. When building leaf nodes – with the exception on the last leaf node – AULID records the maximum key, and the physical address for each leaf node, i.e., the key-block pairs to be indexed in the inner nodes. For the final leaf node, AULID stores the minimum and maximum keys, as well as its physical address in a meta-node.

The second step builds the inner nodes for AULID over the key-block pairs. We first use the Fastest Minimum Conflict Degree (FMCD) algorithm in LIPP to generate a linear model for a node. Given the collection of keys to be indexed and the number of slots that can be used, FMCD aims to generate a linear model under which the maximum number of keys inserted into the same slot is minimized, i.e., the smallest “conflict degree”. Then, we insert the key using the resulting model. If only one key is inserted into a slot, this slot is labeled as *DATA* and used to store the key-block pair. Different from LIPP, AULID does not aggressively create a new node if more than one item is mapped to the same slot. Instead, we divide them into three cases depending on the size of the items that are mapped into the same slot: (1) If the size of the items mapped to one slot is smaller than 64, a packed array is created. (2) If the size is greater than 64 and less than 1020 (see explanation at the end of §II-C2), a two-layer B+-tree is created with at most four child nodes. (3) Otherwise, a new mixed node is created to hold the keys.

B. Lookup & Scan

1) *Lookup*: Given a search key, we first check whether it belongs to the last leaf node by comparing it with the minimum key and the maximum key that are stored in the meta-node. The overhead of this operation is negligible as the meta-node resides in main memory. If the key belongs to the last leaf node, the leaf node is read from disk, and then a binary search is initiated. Otherwise, the inner nodes are searched to find the leaf node address where the search key should reside.

When traversing from root to leaf node, five cases of model prediction can occur (assume a mixed node is the root node):

- *DATA* Slot: A leaf node is fetched based on the physical address contained in it. If the key in the *DATA* slot is less than the search key, then we fetch the successor.
- *NODE* Slot for a Packed Array: The packed array content is retrieved from disk, and a *DATA* slot is located to hold

³We have tested all the datasets proposed in a recent benchmark paper [9].

the search key. It is then processed in the same way as the *DATA* Slot case.

- *NODE* Slot for a B+-tree: Just as in a standard B+-tree, a child node is found which holds the search key (if it exists), and then it is fetched and processed in the same way as the *NODE* Slot for a Packed Array.
- *NODE* Slot for another mixed Node: The model from the node is used to predict which node to access next, and the search process is repeated.
- *NULL* Slot: Using the monotonic linear function from AULID and indexing the largest keys for each leaf node, to find the next *DATA* slot we must search forward. For example, given a search key, suppose the predicted position is the 5th slot in *Node A* of Figure 3, which is a *NULL* slot. AULID will scan forward to find the *next DATA* slot using *Node C* and *Node F*.

2) *Scan*: Given a query range $[u, v]$, we first call a lookup operation to locate the leaf node where u should reside, and the position of u in the node. Then, we scan forward until reaching the last key v . Using the links to sibling leaf nodes and the packed array, the *next* item can be quickly accessed without additional utility structure, such as bitmaps in ALEX to differentiate empty slot, or traversing many nodes in LIPP.

3) *Optimization for Reading*.: When storing only the largest key of each leaf node in the inner nodes, AULID could issue additional I/O requests for two reasons:

- *Issue 1*: When traversing the inner nodes, AULID may need to find the predecessor (the left sibling of a target node) and extra I/O is needed to fetch that target.
- *Issue 2*: If the predicted location is a *NULL* slot, a scan operation is triggered to locate the *next DATA* slot. E.g., if the predicted position is the 5th slot in *Node A* of Figure 3, AULID needs to access *Node C* and *Node F*. Due to large fanout of inner nodes, *Nodes A, C, and F* are stored in different blocks, which could incur additional I/O costs. Note that with a monotonic linear function in AULID, these two cases cannot happen in a single lookup simultaneously.

To address the first issue, if a *DATA* slot is found, instead of fetching the leaf node directly, we first check whether the key it contains is smaller than the search key. If so, we scan forward to find the next *DATA* slot. In contrast to a *NULL* slot, the scan operation is initiated only on the currently fetched block. The *DATA* slot found before is used if no new *DATA* slots are found in the same block. Otherwise, scanning forwards reads at least one additional block.

To address the second issue, AULID fulfills the preceding *NULL* slots for one *DATA* slot until reaching the previous *DATA* node during the bulkloading process. This operation has negligible overhead for bulkloading but only works with Read-Only workloads. For workloads involving a write operation, inserting a new key-block pair into the inner nodes will incur an update for empty slots and hence leads to increased latency during an insertion. The effectiveness of these two optimizations have been verified in our evaluation at § IV-D.

Algorithm 1: Insert(\mathcal{I}, k, p)

Input: \mathcal{I} : the AULID index, k : the key, p : the payload, T : maximum slot count in the leaf node

```

1 accessed_nodes = [];
2 leaf_node = GetLeafNode( $\mathcal{I}, k$ );
3 if leaf_node.size  $\geq T$  then
4   block_id, k_max = SplitNode(leaf_node);  $\triangleright$  max key and block
   id of left child;
5   e = FindEntry( $\mathcal{I}, k_{max}, accessed\_nodes$ );  $\triangleright$  first non mixed
   Node entry;
6   switch e.type do
7     case NULL do
8       insert (k_max, block_id) into e;
9     case DATA do
10      create a packed array PA;
11      insert (k_max, block_id) and (DATA.k, DATA.v) into
      PA;
12      update e as NODE to PA;
13     case B+-tree do
14      if e is full then
15        K = (k_max, block_id)  $\cup$  items in e;
16        P = BuildMixedNode(K);  $\triangleright$  Build a mixed node;
17        update e as NODE to P;
18      else
19        insert (k_max, block_id) into e ;
20     case Packed Array do
21      if e is full then
22        create a packed array or a B+-tree node P;
23        insert (k_max, block_id) and items in e into P;
24        update e as NODE to P;
25      else
26        insert (k_max, block_id) into e ;
27   StatsUpdate(accessed_nodes);  $\triangleright$  Update statistics for SMOs;
28   Adjust( $\mathcal{T}, accessed\_nodes$ );  $\triangleright$  See Section III-D;
29 else
30   InsertLeaf(k, p, leaf_node);

```

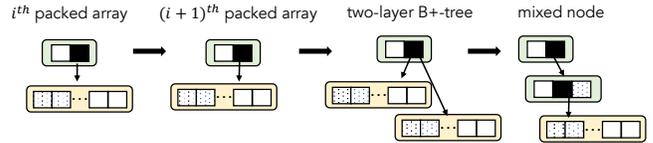


Fig. 4: AULID 's Insertion Process

C. Insert & Duplicate Index Keys Support

1) *Insert*: The full insertion process is presented in Algorithm 1, and the key process is depicted in Figure 4. Given a new key-payload pair (k, p) , a lookup operation is first called to locate the leaf node that should contain k and the position of k in the node (Line 2). If the item count in this leaf node is less than a predefined threshold T , (k, p) is added into this node (Line 30). Note that, if this is the last leaf node, and (k, p) was the first or last item in the node, the minimum key or maximum key is updated in the meta-node. If the item count exceeds the threshold, a splitting process is triggered (Line 4). Unlike a B+-tree, which keeps the smaller half of the items in the original node, AULID keeps the larger half of the items in the original node. Otherwise, the address of the last key is updated in the original inner node, which requires extra writes. After a new leaf node is created, the links to the sibling nodes are also updated.

The largest key and physical address of each leaf node are

indexed in the inner nodes (Lines 5-28). A lookup process is initiated to find the first non-mixed node slot to hold the new key. Just as in the search process, there are five different cases: 1) If we encounter a *NULL* slot, we insert the new key-block pair into it and the insertion is completed (Line 8). 2) If we encounter a *NODE* slot pointing to a mixed node, the model is fetched and the search process is repeated. 3) For a packed array, if full, a larger packed array type will be allocated. If the maximum supported packed array is already being used, it is converted into a two-layer B+-tree (Line 20-24). Otherwise, we insert the new key-payload pair into the empty slot and complete the process (Line 26). 4) If the B+-tree is not full, the process proceeds as in a standard B+-tree (Line 19). 5) Otherwise, it is converted into a mixed node (Lines 15-17).

After completing the insertion, the statistics of the mixed nodes are updated in the access path to guide later SMOs (Line 27). AULID records the number of items in a third layer or a deeper layer. Finally, we check if we need to initiate an SMO operation by calling the *Adjust* function (Line 28).

2) *Handling Duplicate Keys*: Duplicate keys are common in real databases. A linked list used in main memory is however not appropriate when on-disk, as it leads to additional I/O costs when fetching items from the list.

Using a B+-tree styled leaf node, AULID can efficiently store and search for duplicate keys in the leaf nodes. This process is the same as a standard B+-tree. If a duplicate key must be inserted into the inner nodes of AULID, one potential way is to directly insert the duplicate key into the inner nodes. AULID can handle key conflicts using the packed array/two-layer B+-tree proposed earlier. However, in this case, the maximum number of duplicate keys that can be supported is 163,840 for a block size of 4 KB⁴. Using the link between two sibling leaf nodes, another way is that we can only store the first leaf node’s address for the duplicate key. With the larger half of items stored in the original block after the leaf node is split, the address (block number) stored in the original slot in the inner nodes needs to be updated. The write overhead is the same as the last case while we can support an arbitrary number of duplicate keys.

D. Structural Modification of Inner Nodes

When a two-layer B+-tree node is converted into a new mixed node, certain regions may have a larger tree height. Therefore, the index structure must be carefully modified to bound the height of the branches in AULID’s inner nodes, e.g., at most three layers in our test datasets with up to 800M key-payload pairs. Otherwise, a B+-tree on disk will be the best. Packed arrays and two-layer B+-tree nodes do not have an impact on the tree height. Here, we focus on when and how to re-construct mixed nodes in AULID.

1) *When Should the Rebuilding Occur*: To bound the tree height of the inner nodes and avoid aggressive node updates, we introduce two new constraints to determine when a mixed node should be reconstructed.

⁴A B+-tree node with 4 children can store at most 640 leaf node addresses and each leaf node can store 256 items.

Criterion 1: the percentage of the items in a subtree rooted at node n in the third layer or a deeper layer ($l3_item$) is larger than α . This guarantees that no more than α leaf nodes can have a longer path than a B+-tree on disk.

Criterion 2: The number of current items rooted at node n is larger than β times the initial size.

In corner cases where a region has a high degree of conflict, a mixed node can have more than α items, even when it is initially being created. To avoid reconstructing this node frequently, we adjust it after observing a sufficient number of new items. A smaller value for α and β leads to a more frequent node reconstruction. By default, we set $\alpha = 0.05$ and $\beta = 1.2$ to balance the tree height, i.e., lookup performance and SMO overheads. We have verified the impact of different settings for α and β in §IV-D3.

2) *How to Reconstruct a Node*: If a mixed node meets both of the above criteria, all key-payload pairs stored in the inner nodes rooted at that node are collected, and then the bulkload process is called again to build a new mixed node.

E. Other Operations

To support a delete operation, AULID first locates the items to be deleted at the leaf node, and then deletes it in the same manner as a standard B+-tree. If no SMO is required (merging the sibling nodes), the delete operation is finalized. In this case, even if we delete the last key-payload pair in the leaf node, AULID still does not update the inner nodes. If a merge is required, a delete operation in the inner nodes is required. If the key-block pair to be deleted is in a mixed node, this slot is marked as an empty slot. If it is contained in the packed array or a two-layer B+-tree node, it will be removed. AULID will convert the packed array or a two-layer B+-tree node into a *DATA* slot if there is only one key-block pair remaining.

There are two types of update operations, updating the payload and updating the key. In the former, an in-place update is used⁵. For the latter, a delete operation and an insert operation are initiated.

IV. EXPERIMENTS

We have conducted extensive experiments to mainly answer three questions – **Q1**: How good is AULID as compared to other learned indexes and a B+-tree when disk-resident? **Q2**: How well does AULID scale to large datasets? **Q3**: Do the proposed index structure design and structural modification operation help improve the performance?

We start with the experimental setup in §IV-A. Then, we present our answers to **Q1-Q3**. To answer **Q1**, we compare AULID against five competitors across six different workload types and four different datasets. We demonstrate that AULID is superior in terms of throughput and storage cost in §IV-B. To answer **Q2**, in §IV-C we use another four datasets with 800M keys to study the performance of AULID on large datasets of varying hardness. Notably, most existing studies [4, 7, 13, 14] use at most 200M key-payload pairs to study index performance. Finally, to answer **Q3**, in §IV-D we compare AULID

⁵Currently, AULID supports key-payload pairs of a fixed length.

to its variants, with and without the proposed data structures and structure modification operations, in order to reveal the performance benefits of our proposed design choices.

A. Experimental Setup

1) *Baselines*: We implement a standard B+-tree and use the implementation [8] of the on-disk versions of the state-of-the-art updatable in-memory learned indexes – PGM [6], FITing-tree [5], ALEX [4], and LIPP [7].

2) *Datasets*: The most recent experimental study [9] on memory-resident learned indexes introduced 11 real datasets in their evaluation. Based on its profiling results, these datasets can be roughly divided into four categories (see Figure 2 in [9]): **C1**: Globally easy and locally easy, **C2**: globally normal and locally normal, **C3**: globally normal and locally hard, and **C4**: globally hard and locally normal. We select one dataset from each of these categories for our experiments: *COVID* (**C1**), *PLANET* (**C2**), *GENOME* (**C3**), and *OSM* (**C4**). Each dataset has 200M keys of type *uint64*. The performance of AULID and LIPP correlates to the *conflict degree* (the maximum number of keys being inserted into the same slot) in one dataset due to the usage of the FMCD algorithm. Consequently, datasets with a greater *conflict degree* are more challenging for AULID and LIPP. A summary of the conflict degrees of the tested datasets is presented in Table I.

TABLE I: Conflict Degree of Each Dataset

Dataset	COVID	PLANET	GENOME	OSM
<i>Conflict Degree</i>	27	22	585	4,106

To test the scalability of AULID on large datasets, we use *OSM* from [15], which contains 800M *uint64* keys. The generator proposed in [9] is used to generate another three datasets, each of which contain 800M *uint64* keys and has different levels of hardness (details presented in Section IV-C).

For all datasets, we generate a *uint64* payload for each key with key plus 1 as their value. The first four datasets require 2.98 GB of storage space on disk, and the last four datasets (used for the scalability testing) require 11.92 GB.

3) *Workloads*: We compare AULID against all baselines across six different workload types typically encountered in a database. **W1 - Lookup-Only** workload, where each index is built on 200M key-payload pairs and the workload consists of two different search key distributions: (1) 20,000 randomly sampled search keys (**W1-U**) and (2) 20,000 search keys sampled with a skewed Zipf distribution (**W1-S**). **W2 - Scan-Only** workload, where the start key of the search range is set to the same key in the lookup-only workload, and the search range is set to 100. The queries are issued on indexes prebuilt on the full dataset. **W3 - Write-Only** workload, where the initial index is built with 10M key-payload pairs that are randomly selected from a dataset, and then another 10M key-payload pairs are inserted. **W4 - Read-Heavy** workload includes 90% lookup queries and 10% write operations. **W5 - Balanced** workload consists of 50% lookup queries and 50% write operations. **W6 - Write-Heavy** workload includes 90% write operations and 10% lookup queries.

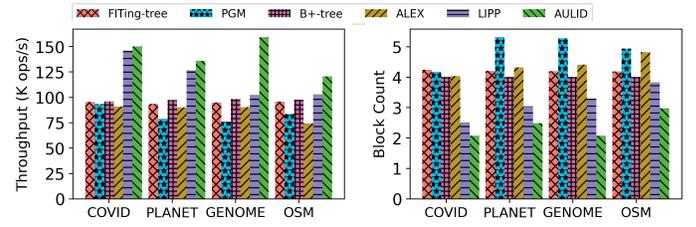


Fig. 5: Throughput and Block Count on Lookup-Only Workload (W1-U).

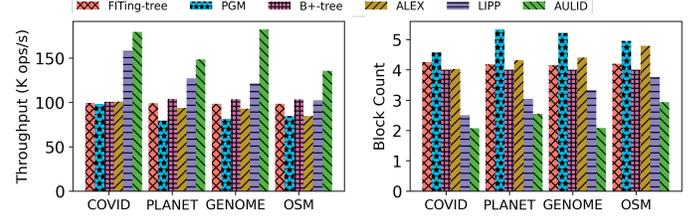


Fig. 6: Throughput and Block Count on Lookup-Only Workload (W1-S).

We refer to W4-W6 as mixed workloads, with the only difference between them being the ratio between reads and writes. For mixed workloads, the initial index is built over 10M key-payload pairs randomly sampled from a dataset, and then lookup queries and write operations are issued (10M queries in total). The search keys in all mixed workloads are randomly sampled from the existing keys of an index.

4) *Metrics & Environment*: The primary metric we measure is throughput. We also report the number of fetched blocks, the storage size of each index, and the tail latency. We conduct the experiments on a SATA HDD using a Red Hat Enterprise Server 7.9 on an Intel Xeon CPU E5-2690 v3 @ 2.60GHz with 256 GB memory and a 1TB HDD. The block size is 4 KB in all experiments.

B. Efficiency Comparisons on Disk

In this section, we compare AULID against four state-of-the-art learned indexes, and a B+-tree on disk. AULID outperforms all five indexes on every dataset and workload tested.

1) *Lookup-Only Workload*: Figure 5 and Figure 6 show the throughput and the average number of fetched blocks per query, for each index under different search key distributions. Overall, AULID is the most efficient indexing method. Specifically, on **W1-U**, it achieves up to 1.68x, 2.10x, 1.62x, 1.76x, and 1.55x higher throughput than the FITing-tree, PGM, B+-tree, ALEX, and LIPP, respectively. LIPP is the second most efficient index across all datasets. The performance of each index is correlated to the number of fetched blocks where more fetched blocks lead to a lower throughput. On **W1-S**, we observe the same behavior as on **W1-U**.

The performance of the FITing-tree, PGM, and B+-tree is similar across all datasets. ALEX however has the worst performance on *OSM*. The performance of AULID and LIPP vary across different datasets; specifically, the performance is related to the *conflict degree* of a dataset, where a higher number of conflicts usually leads to a greater tree height, and in turn more fetched blocks.

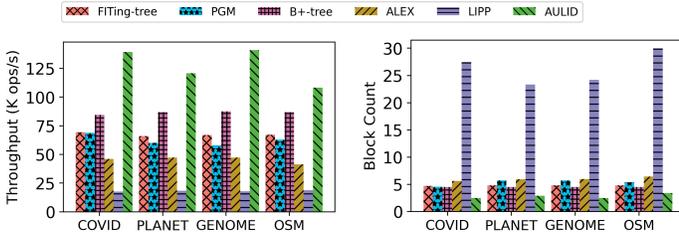


Fig. 7: Throughput on Scan-Only Workload (W2).

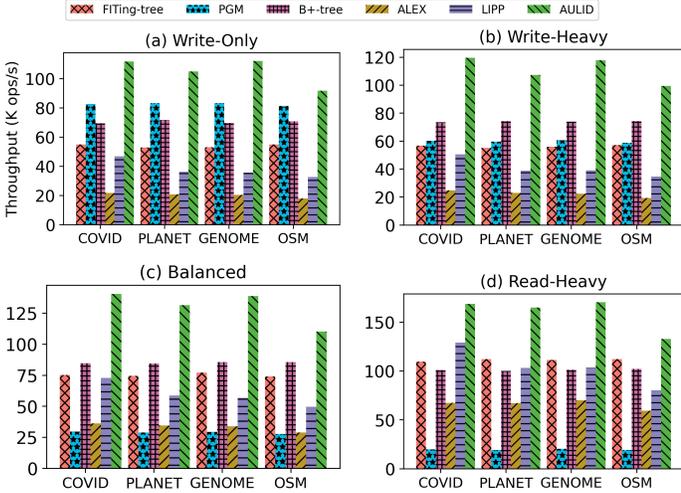


Fig. 8: Throughput of Mixed Workloads (W3-W6).

2) *Scan-Only Workload*: Figure 7 summarizes the throughput and the average number of fetched blocks for the Scan-Only workload. In terms of throughput, AULID outperforms FITting-tree, PGM, B+-tree, ALEX, and LIPP by up to 2.11x, 2.44x, 1.65x, 3.04x, 7.94x, respectively. Just as in the Lookup-Only workload, the performance of the scans is determined by the number of fetched blocks.

To support a scan query, all indexes first initiate the search process for a lookup query to locate the start key in the search range, and then scan forward until reaching the end key. Consequently, better performance in Lookup-Only workloads yields better performance in Scan-Only workloads. Using the packed array in leaf nodes and links between siblings, the B+-tree and AULID reap the benefits from efficient lookup queries, and are the two top-performing algorithms. In contrast, LIPP does not gain any benefits from lookup queries. LIPP only has one node type, where key-payload pairs, pointers to child nodes, and empty slots are all interleaved. Thus, when fetching the *next* item, LIPP may have to traverse multiple nodes. Since LIPP has a large fanout, there is a greater chance that these nodes are in different blocks. Also, we observe that the performance of ALEX decreases more quickly than the FITting-tree and PGM. With a gapped array in the leaf node, ALEX uses a bitmap to indicate whether a slot is occupied, and thus incurs additional I/O cost when fetching it.

3) *Write-Only and Mixed Workloads*: Figure 8 shows the throughput for the workloads that include write operations. AULID is still the best performer across all workloads and datasets. The superiority of AULID is attributed to three

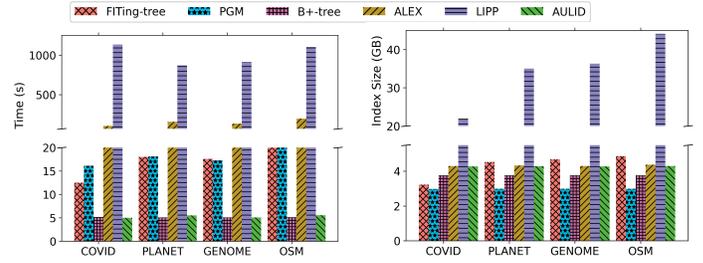


Fig. 9: Comparison of Bulkload Time and Storage Usage.

reasons: (1) a lower latency to locate where a new key-payload should be inserted, i.e., benefiting from the best lookup performance; (2) a lower SMO overhead on leaf nodes with the B+-tree styled leaf node design; and (3) a lower SMO overhead for the inner nodes, and fewer SMOs required. Based on our design of the packed array and the two-layer B+-tree nodes, most new key-block pairs can be stored without creating new mixed nodes. In the tested datasets, no dataset required AULID creation of new mixed nodes. Other learned indexes, however, require more SMOs, e.g., on the Write-Only workload, ALEX and LIPP require 45,897 and 4.5M SMOs on *GENOME*, respectively.

PGM outperforms other approaches on the Write-Only workload, but it performs worse with more reads. Better insertion support stemming from the LSM tree [11] allows PGM to be competitive for write operations. However, since multiple files are maintained as static PGM indexes, PGM may access more than one file for a lookup query, which increases the I/O cost. The performance gain from a faster lookup time can benefit the workloads containing more reads, e.g., the FITting-tree and LIPP on the Read-Heavy workload. However, as the number of writes increases, the SMO overhead and the cost of updating statistics (for ALEX and LIPP) [8] can outweigh the benefits gained from faster lookups.

4) *Bulkloads*: Figure 9 reports the bulkloading time and the on-disk index size after bulkloading. When calculating index sizes, we report the total size of the index file on disk. This ensures that the entire on-disk size is reported for a fair comparison in practice. In terms of bulkloading time, AULID is similar to the B+-tree, and both of them are significantly smaller than the other indexes. AULID also achieves similar storage cost to the B+-tree. The FITting-tree and LIPP have different storage sizes across different datasets. In the case of FITting-tree, harder datasets will create more leaf nodes (segments), and allocate additional buffers on disk for later key-payload pair insertions. For LIPP, a dataset with a larger degree of conflict will result in more nodes being created on disk, and in turn occupy more space. ALEX and LIPP have larger bulkloading times than the other methods due to model training, and more on-disk writes.

5) *Index Size*: Figure 10 presents the storage occupancy of all indexes after finishing workloads comprising writes (**W3-W6**). Overall, AULID achieves similar storage overheads to the B+-tree across all workload and dataset combinations. Among the rest of the competitors, PGM has the smallest storage index.

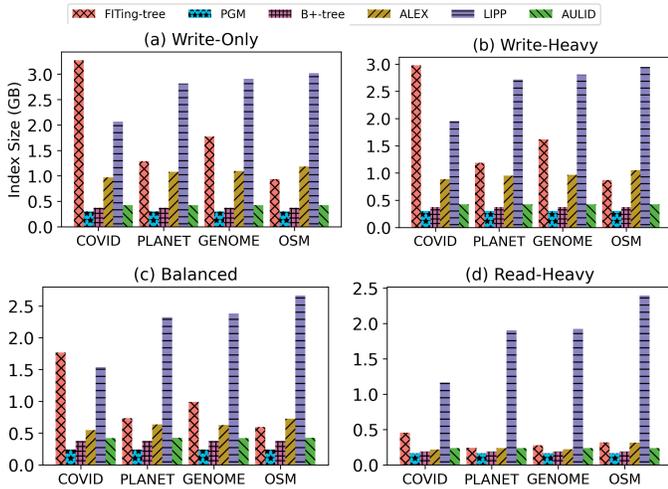


Fig. 10: Storage Occupancy of Mixed Workloads (W3-W6).

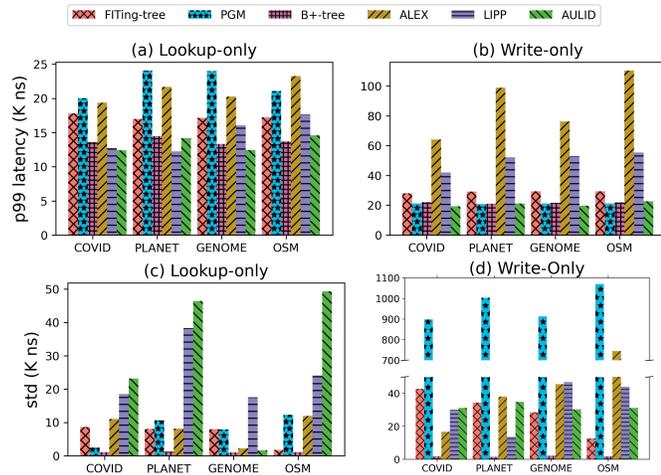


Fig. 11: Tail Latency on Lookup-Only (W1) and Write-Only (W3) Workloads.

This is attributed to the LSM tree used in PGM to support arbitrary insertions, i.e. after an index has been merged, we can delete it from disk. For LIPP, a dataset with a higher degree of conflict usually has a larger storage cost due to the creation of additional nodes. The FITing-tree has a large space occupancy, regardless of the dataset or workload.

6) *Tail Latency*: To study the robustness of each index, in Figure 11 we report the p99 latency and standard deviation on the Lookup-Only and Write-Only workloads. Overall, AULID has the smallest p99 latency in the Lookup-Only workload. AULID, PGM, and B+-tree have similar p99 latencies for the Write-Only workload – all of which are better than FITing-tree, ALEX, and LIPP. However, all learned indexes have a larger standard deviation than B+-tree across both workloads. Due to an unbalanced tree structure of LIPP and ALEX, accessing some regions may issue more I/O requests for the Lookup-Only case. When indexing only the largest key of each leaf node in the inner nodes, in a lookup, AULID may access more blocks to fetch the *next DATA* slot or read an extra block to locate the target leaf node as discussed in §III-B3. PGM

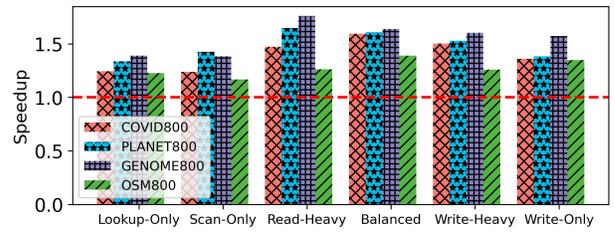


Fig. 12: Throughput Speedup on Large Datasets.

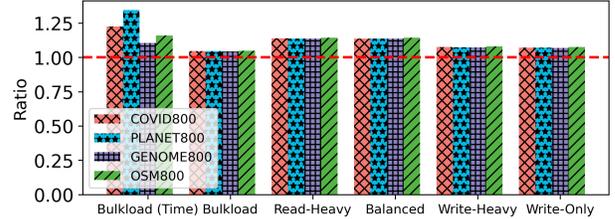


Fig. 13: Bulkload Time and Storage Usage Compared to B+-tree on Large Datasets.

will periodically merge items into a larger index. Heavy SMOs for certain queries result in a larger latency in the Write-Only workload, which in turn results in a larger variance.

C. Scalability Test

Setting. Since existing learned indexes perform worse than the B+-tree overall, in this section, we compare the scalability of AULID against the B+-tree only. To test the performance of AULID on datasets of different hardness, we include *OSM800* [15], and three other datasets of size 800M generated using the method from [9]. For each, we set the local hardness and global hardness to 4x of *COVID*, *PLANET*, and *GENOME* and name them as *COVID800*, *PLANET800*, and *GENOME800*, respectively.

For the Lookup-Only and Scan-Only workloads, we issue 800,000 queries over the index built on the entire dataset. Search keys are randomly sampled from the entire dataset. For the workloads that contain writes, we build an initial index containing 150M keys sampled from a dataset and issue a total of 50M operations, where the write ratio is the same as used for **W3-W6** in §IV-A3.

1) *Performance Speedup*: Figure 12 presents the throughput speedup of AULID compared to the B+-tree. AULID beats the B+-tree with up to 1.75x throughput gains on all tested workloads and datasets. Due to a carefully designed inner node structure and an SMO mechanism to bound the tree height, AULID is more efficient when locating the target leaf node, and also benefits scans and writes.

The superiority of AULID on large datasets also comes from the smaller SMO overhead for write operations. When indexing the largest key for each leaf node of the learned model, AULID reduces the number of SMOs needed to reap the benefits of model-based search. Moreover, a packed array and a two-layer B+-tree hold the new key-block pairs without increasing the tree height (See Figure 4), while incurring only small update overheads.

2) *Bulkload Time & Storage Usage*: The results of AULID and B+-tree are reported in Figure 13. To build an index for

800M key-payload pairs on disk, the B+-tree takes around 20s and AULID is competitive at 27s. Both are much more efficient than other learned indexes, even on small datasets. We also conduct an experimental study using different dataset scales and the results on COVID are shown in Table II. The bulkload time and storage usage of AULID are only slightly greater than the B+-tree across all dataset sizes.

TABLE II: Storage and Bulkload Time under Different Dataset Scales of COVID.

Metric	Index	200M	400M	600M	800M
Bulkload Time (s)	B+-tree	4.94	11.30	17.25	22.74
	AULID	6.26	13.47	19.41	27.82
Storage (GB)	B+-tree	3.76	7.52	11.27	15.03
	AULID	4.27	7.85	11.76	15.74

D. An In-depth Study of AULID Design

Next, we study how the design of inner nodes meets our proposed design principles and address the two challenges aforementioned in §I-A. Typically, basic operations include the lookup and insertion, which in turn define the studied workload types (**W1**, **W3-W6**) and are also the key step in **W2**. Specifically, we first study the impact of the AULID design on these two operations, and then investigate the impact of the adjustment strategy proposed in §III-D.

1) Impact of Different Designs on Lookup-Only Workloads:

To study the effectiveness of the proposed design, we compare AULID against LIPP-B+ – an approach which directly adopts LIPP as the inner nodes, and organizes the leaf nodes in the same vein as B+-tree. We report the throughput of the Lookup-only Workload (**W1**) in Table III. Across all datasets, AULID outperforms LIPP-B+ and fetches fewer blocks.

TABLE III: Throughput Comparison of AULID and LIPP-B+ on Lookup-Only Workload (W1).

Metric	Index	COVID	PLANET	GENOME	OSM
Thrpt	LIPP-B+	158,489	133,404	153,851	104,659
	AULID	164,897	141,543	163,825	123,749
Blocks	LIPP-B+	2.15	2.72	2.25	3.30
	AULID	2.07	2.50	2.07	2.96

The performance of AULID is attributed to our *read optimization strategies* (§III-B3), *packed array design*, and *two-layered B+-tree nodes* (§II-C2). The first helps reduce the number of blocks being fetched, and the latter two help reduce the tree height.

TABLE IV: Extra Fetched Blocks under Different Optimizations

Dataset	w/o Opt.	Fulfill	ScanFwd	Fulfill & ScanFwd
COVID	26,107	18,337	9,266	277
PLANET	59,711	52,619	30,090	21,027
GENOME	47,229	40,727	9,456	710
OSM	30,148	22,368	23,232	14,924

Benefits of Read Optimizations. Recall §III-B3, in AULID, there are two cases that may incur additional I/O costs: *Case*

1 - located in the predecessor of the target leaf node and require an additional block being fetched; *Case 2* - located in a *NULL* slot but need to scan forward until the next *DATA* slot is found. For the first case, AULID scans forward (**ScanFwd**) to determine whether the current block has another *DATA* slot. For the second case, AULID fulfills (**Fulfill**) the empty slot with the next *DATA* slot during a bulkloading process.

With the **Fulfill** optimization, AULID avoids extra blocks in *Case 2*. All of the fetched extra blocks in the **Fulfill** optimization are stemming from *Case 1* (the third column in Table IV). From Table IV, we observe that most additional fetched blocks are from *Case 2*. With the **ScanFwd** optimization, AULID significantly reduces the number of extra fetched blocks for *COVID*, *PLANET*, and *GENOME*. With the **Fulfill** optimization, AULID avoids fetching extra blocks in *Case 1*.

When enabling both operations, AULID can reduce the number of extra fetched blocks by at least 50%, particularly for *COVID* and *GENOME*. By default, we only enable the **ScanFwd** optimization, which can reduce 0.08, 0.15, 0.18, and 0.03 blocks per query for *COVID*, *PLANET*, *GENOME*, and *OSM*, respectively. Thus, improvements in AULID on *COVID* and *GENOME* in Table III are produced by **ScanFwd** optimization only.

TABLE V: Impact of Packed Array and Two-Layer B+-tree on the Average *DATA* Slot Height and Storage.

Metric	Index	COVID	PLANET	GENOME	OSM	OSM800
Avg. Height	LIPP-B+	1.00	1.60	1.01	2.29	2.28
	AULID	1.00	1.36	1.00	1.83	1.93
Storage (GB)	LIPP-B+	4.27	4.47	4.28	5.11	19.51
	AULID	4.27	4.28	4.27	4.29	15.77

Benefits of Data Structure Design. The design of the packed array and the two-layer B+-tree node in AULID helps further reduce the number of fetched blocks for the *PLANET* and *OSM* datasets due to the lower tree height than LIPP-B+.

Table V reports the average node heights after a bulkload on *COVID* and *GENOME*. AULID and LIPP-B+ both have the minimal average node height, where most *DATA* slots are located in the first level. However, for datasets with a larger *conflict degree*, AULID has a smaller average node height. This is because LIPP-B+ creates more nodes to eliminate the number of conflicts in each dataset with a larger *conflict degree*, which leads to a greater height. A lower tree height can be achieved with the packed data structures discussed above, and thus AULID requires less storage space than LIPP-B+ on the hard datasets as shown in Table V.

TABLE VI: Throughput Comparison of AULID and LIPP-B+ on Write-Only Workload (W3)

Index	COVID	PLANET	GENOME	OSM
LIPP-B+	111,017	100,430	109,631	76,865.4
AULID	111,669	104,816	107,661	91,707.1

2) *Packed Array, Two-Layer B+-tree Nodes, and B+-tree Styled Leaf Nodes for Write-Only Workloads.*: Table VI reports the throughput for the Write-Only workload (**W3**). We observe that AULID and LIPP-B+ have similar performance

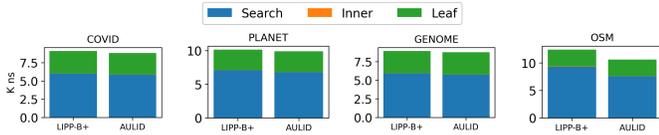


Fig. 14: Latency Breakdown of Write-Only Workload (W3).

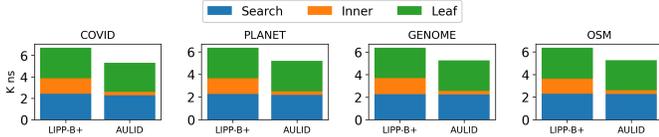


Fig. 15: Latency Breakdown of Append-Only Workload.

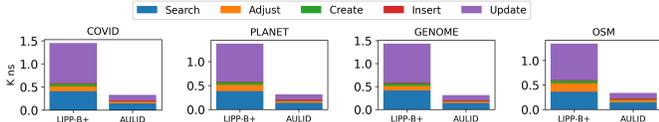


Fig. 16: Inner Nodes' Latency Breakdown of Append-Only Workload.

for *COVID*, *PLANET*, and *GENOME*, but AULID outperforms LIPP-B+ by 1.19x on *OSM*.

We break down the insertion process into three steps: (1) the search step (**Search**) to find the leaf node that will hold the new key-payload pair, (2) the insertion step on a leaf node (**Leaf**), and (3) the update step in the inner nodes (**Inner**). From Figure 14, we can see that the main overhead is brought by the first two steps. The overhead from indexing the new key-block pairs in the inner nodes (**Inner**) is negligible. AULID and LIPP-B+ update the inner nodes when a leaf node is split. When compared against LIPP, both have fewer SMO operations. For example, on *GENOME*, AULID and LIPP-B+ only require 49,038 SMOs on the leaf nodes, where the leaf nodes must be split; in contrast, LIPP requires 4.6M SMOs and most of them are caused by creating LIPP nodes to eliminate conflicts. On *OSM*, AULID has a more efficient **Search** step than LIPP-B+. The design of the packed array and the two-layer B+-tree nodes in AULID consistently results in shorter paths to leaf nodes, and hence smaller search cost.

TABLE VII: Throughput Comparison of AULID and LIPP-B+ on Append-Only Workload

Index	COVID	PLANET	GENOME	OSM
LIPP-B+	144,432	153,060	152,633	153,214
AULID	182,732	188,015	187,886	184,051

Hot Region Insertions. Another potential problem with LIPP is insertions in a hot region which occur in the inner nodes. This produces a high number of conflicts and triggers additional SMOs. We use an Append-Only workload to analyze this case. The throughput and latency breakdown for this case are shown in Table VII and Figure 15, respectively. To support new key-payload pairs, AULID and LIPP-B+ index the last leaf node in the meta-node. Thus, they have similar performance on the **Search** and **Leaf** in Figure VII.

AULID has a lower latency when updating the inner nodes (**Inner**) across all datasets. To study the overhead of the inner

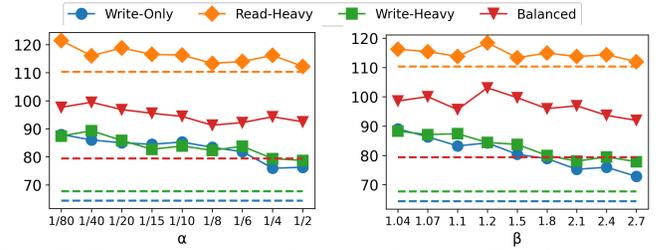


Fig. 17: Throughput under Different Settings. The dashed line indicates the throughput of the corresponding workload with the same color without any adjustments.

nodes, we further break down the process of the inner nodes into five steps: (1) search to find a slot to hold the new key-block pair (**Search**), (2) create a new node or convert the node type (**Create**), (3) insert the key-block pair into a *DATA* slot (**Insert**), (4) adjust the tree structure (**Adjust**), and (5) update the statistics for later adjustment operations (**Update**).

In Figure 16, we observe that, except for the **Insert**, AULID has a lower latency than LIPP-B+. The packed array and two-layer B+-tree node design of AULID yield a shorter path to target slots holding the new key-block pairs, which contributes to lower search and update times. Moreover, conflicts in the Append-Only workload require more operations when creating LIPP nodes and adjusting tree structures in LIPP-B+. For example, AULID only requires 1,163 and 37 SMOs in creating new nodes and adjusting the tree structure, while LIPP-B+ requires 38,837 and 1,625 SMOs, respectively.

3) *Adjustment Study:* Last, we study the effectiveness of index adjustments and provide an analysis on the parameters in AULID – α and β presented in §III-D. We use the *OSM* dataset, and build an initial index using 50M keys. Then, another 50M queries are issued with different write ratios, just as in **W3-W6**. To study the impact of α and β , we set $\beta = 1.07$ and $\alpha = 0.0025$ as the default. Using small default values makes us study the impact of other parameters in isolation. Figure 17 illustrates the throughput when using different settings. We also report the throughput without adjusting the index (the dashed lines). In Figure 17, we observe: (1) The index adjustment in AULID significantly improves the throughput, especially for workloads with writes. (2) Larger values of α and β usually result in worse performance, and the workloads with more writes are more sensitive to α and β . If index is not adjusted, certain regions in the inner nodes can result in longer paths to the leaf nodes. This incurs more reads to fetch leaf nodes and more writes to update statistics on disk. The default values of α (0.05) and β (1.2) in AULID result in good performance across all tested workloads.

E. Impact of Block Size

In Figure 18, we show the lookup-only workload performance under different block sizes. With the same search key distribution across all testing datasets, B+-tree has the same average fetched block count in all datasets. The number of blocks fetched by AULID does not change significantly when the block size is varied. The position prediction in AULID is highly accurate. When the block size is 16 KB, the height of

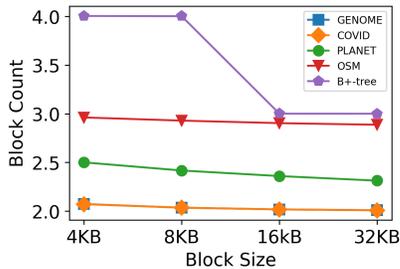


Fig. 18: Impact of Block Size.

B+-tree decreases for a larger fanout. Thus, the fetched block count is reduced significantly. Regardless, AULID fetches fewer blocks compared to the B+-tree across all settings.

V. RELATED WORK

Learned Indexes Outside Main Memory. The authors of [10, 1] studied how to use learned indexes on disk in a log-structured merge tree (LSM) [11]. A learned model is constructed for each Sorted Strings Table (SSTable), which is *immutable* after being created. The LSM framework supports efficient writes at the cost of reads. Two recent studies [14, 13] focus on the *larger than main memory setting*, where part of the index is pinned in main memory, and different caching strategies are proposed. TreeLine [14] uses the partitioning method proposed in PGM [6] to generate leaf nodes, and uses a B+-tree to index the nodes. The B+-tree in TreeLine is pinned in main memory, and a record-level caching strategy is used to cache frequently accessed items. FILM [13] builds a PGM index in main memory, and uses one bit for each item to indicate the item location – in main memory or on disk. Also, FILM introduces a global chain and a local chain to organize the segments in the last level and the items in the segment, respectively. Lan et al. [8] present a comprehensive experimental study of on-disk learned indexes. All of our observations align with their conclusions as well.

FILM and TreeLine have considered caching techniques. Based on the results shown in Figure 8 [14], record-level caching is the biggest overall contributor to the performance gains. As a standalone component in TreeLine, record-level caching could also be directly adopted by AULID. FILM however is designed for *append-only insertions*, and we allow *arbitrary insertions*. Supporting *append-only insertions* allows FILM to use PGM [6] as the index structure, which has the smallest index size when compared to other index structures, allowing it to use any remaining space for caching.

Learned Indexes in Main Memory. Kraska et al. [16] first proposed the idea of a learned index. A hierarchy of models, called RMI, was built to replace a B+ Tree on sorted 1-dimensional data. Given that RMI only supports lookup queries, subsequent studies [4, 6, 7, 5, 17, 18] develop tailored indexes to support index modification. The FITing-tree [5] replaces the last layer of a B+-tree with model-based search and supports insertions by introducing buffers for each segment. PGM [6] uses a similar idea to the FITing-tree but leverages model-based search for every layer based on an optimal partitioning algorithm [19]. PGM handles arbitrary insertions

in an LSM tree [11] style manner. Although FITing-tree and PGM leverage model-based search, additional binary search operations are needed. ALEX [4] inserts a key-payload pair using model prediction, and hence manages to have accurate predictions in the inner nodes without a binary search. For each data node (leaf node), ALEX uses a gapped array, and an exponential search is used to find the target position. Model predictions in LIPP [7] are accurate in every layer. LIPP has been shown to have better performance in practice than other learned indexes in most settings [9]. However, LIPP requires much more memory and is not efficient for a range query. Wu et al. [20] use Normalizing Flows [21] to transform a dataset so that it can be easily modeled with linear model, and extend the idea of LIPP. CARMI [22], a cache-aware learned index, designs novel partitioning algorithms and a prefetching mechanism. Several studies conduct experimental studies [15, 9, 23], theoretical analysis [24], and performance analysis [25]. There are also studies on read-only case [26, 27, 28] and on boosting B+-tree performance [29, 30]. Several other learned indexes have also been proposed for secondary indexes [31, 32], for multi-dimensional data [33, 34, 35], spatial data [36, 37, 38, 39], and string data [40, 41].

Learned Index With Concurrency Support. Among the in-memory learned indexes, XIndex [42], FINEdex [43], ALEX+[9], and LIPP+[9] can support concurrent operations. All of them use optimistic locking – associating a versioning lock to each node. Since the node size is larger than a single block in the leaf node, TreeLine [14] proposes several locking strategies based on node-level lock and block-level lock to support concurrent operations on LSM data structures. Aligned with the latest study on the larger than main memory case [13], we focus on the single-threaded setting in this work. In a fully on-disk learned index with varying sizes of nodes and blocks, the node-level lock and the block-level lock should be introduced simultaneously. Without a novel mechanism, one cannot achieve consistent scalability in a multi-core setting when using a block-level lock. Although TreeLine [14] has introduced these two locks, its inner nodes are pinned in main memory, which is much easier than the above case. We believe our work is the first step to solve this problem.

VI. CONCLUSION

We propose AULID, a simple yet efficient on-disk learned index. We propose five principles to build a learned index on disk, focusing on reducing the I/O cost, the main bottleneck for disk-resident indexes. We carefully craft the index structure, propose the query processing algorithms, and introduce an index adjustment mechanism to meet these principles. AULID outperforms all the baselines across a wide range of settings. Our evaluation shows that AULID has competitive storage cost to B+-tree (the smallest of the alternatives), and achieves up to 2.11x, 8.63x, 1.72x, 5.51x, and 8.02x higher throughput than FITing-tree, PGM, B+-tree, ALEX, and LIPP respectively.

Acknowledgements. This research is supported in part by ARC DP240101211, DP220101434, DE230100366, and L’Oréal for Women in Science Fellowship.

REFERENCES

- [1] H. Abu-Libdeh, D. Altinbükten, A. Beutel, E. H. Chi, L. Doshi, T. Kraska, X. Li, A. Ly, and C. Olston, "Learned indexes for a google-scale disk-based database," *CoRR*, vol. abs/2012.12501, 2020.
- [2] X. Zhou, L. Liu, W. Li, L. Jin, S. Li, T. Wang, and J. Feng, "Autoindex: An incremental index management system for dynamic workloads," in *ICDE*. IEEE, 2022, pp. 2196–2208.
- [3] L. M. Haas, M. J. Carey, M. Livny, and A. Shukla, "Seeking the truth about ad hoc join costs," *VLDB J.*, vol. 6, no. 3, pp. 241–256, 1997.
- [4] J. Ding, U. F. Minhas, J. Yu, C. Wang, J. Do, Y. Li, H. Zhang, B. Chandramouli, J. Gehrke, D. Kossmann, D. B. Lomet, and T. Kraska, "ALEX: an updatable adaptive learned index," in *SIGMO*, 2020, pp. 969–984.
- [5] A. Galakatos, M. Markovitch, C. Binnig, R. Fonseca, and T. Kraska, "Fiting-tree: A data-aware index structure," in *SIGMOD*, 2019, pp. 1189–1206.
- [6] P. Ferragina and G. Vinciguerra, "The pgm-index: a fully-dynamic compressed learned index with provable worst-case bounds," *Proc. VLDB Endow.*, vol. 13, no. 8, pp. 1162–1175, 2020.
- [7] J. Wu, Y. Zhang, S. Chen, Y. Chen, J. Wang, and C. Xing, "Updatable learned index with precise positions," *Proc. VLDB Endow.*, vol. 14, no. 8, pp. 1276–1288, 2021.
- [8] H. Lan, Z. Bao, J. S. Culpepper, and R. Borovica-Gajic, "Updatable learned indexes meet disk-resident DBMS - from evaluations to design choices," *Proc. ACM Manag. Data*, vol. 1, no. 2, pp. 139:1–139:22, 2023.
- [9] C. Wongkham, B. Lu, C. Liu, Z. Zhong, E. Lo, and T. Wang, "Are updatable learned indexes ready?" *Proc. VLDB Endow.*, vol. 15, no. 11, pp. 3004–3017, 2022.
- [10] Y. Dai, Y. Xu, A. Ganesan, R. Alagappan, B. Kroth, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau, "From wiskey to bourbon: A learned index for log-structured merge trees," in *OSDI*, 2020, pp. 155–171.
- [11] P. E. O’Neil, E. Cheng, D. Gawlick, and E. J. O’Neil, "The log-structured merge-tree (lsm-tree)," *Acta Informatica*, vol. 33, no. 4, pp. 351–385, 1996.
- [12] "Aulid’s code," <https://github.com/rmitbggroup/AULID>.
- [13] C. Ma, X. Yu, Y. Li, X. Meng, and A. Maolinyazi, "Film: a fully learned index for larger-than-memory databases," *Proc. VLDB Endow.*, vol. 16, no. 3, pp. 561–573, 2022.
- [14] G. X. Yu, M. Markakis, A. Kipf, P.-r. Larson, U. F. Minhas, and T. Kraska, "Treeline: An update-in-place key-value store for modern storage," *Proc. VLDB Endow.*, vol. 16, no. 1, p. 99–112, 2022.
- [15] R. Marcus, A. Kipf, A. van Renen, M. Stoian, S. Misra, A. Kemper, T. Neumann, and T. Kraska, "Benchmarking learned indexes," *Proc. VLDB Endow.*, vol. 14, no. 1, pp. 1–13, 2020.
- [16] T. Kraska, A. Beutel, E. H. Chi, J. Dean, and N. Polyzotis, "The case for learned index structures," in *SIGMOD*, 2018, pp. 489–504.
- [17] J. Ge, H. Zhang, B. Shi, Y. Luo, Y. Guo, Y. Chai, Y. Chen, and A. Pan, "SALI: A scalable adaptive learned index framework based on probability models," *CoRR*, vol. abs/2308.15012, 2023.
- [18] P. Li, H. Lu, R. Zhu, B. Ding, L. Yang, and G. Pan, "DILI: A distribution-driven learned index," *Proc. VLDB Endow.*, vol. 16, no. 9, pp. 2212–2224, 2023.
- [19] J. O’Rourke, "An on-line algorithm for fitting straight lines between data ranges," *Commun. ACM*, vol. 24, no. 9, pp. 574–578, 1981.
- [20] S. Wu, Y. Cui, J. Yu, X. Sun, T. Kuo, and C. J. Xue, "NFL: robust learned index via distribution transformation," *Proc. VLDB Endow.*, vol. 15, no. 10, pp. 2188–2200, 2022.
- [21] D. J. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *ICML*, vol. 37. JMLR.org, 2015, pp. 1530–1538.
- [22] J. Zhang and Y. Gao, "CARM: A cache-aware learned index with a cost-based construction algorithm," *Proc. VLDB Endow.*, vol. 15, no. 11, pp. 2679–2691, 2022.
- [23] Z. Sun, X. Zhou, and G. Li, "Learned index: A comprehensive experimental evaluation," *Proc. VLDB Endow.*, vol. 16, no. 8, pp. 1992–2004, 2023.
- [24] P. Ferragina, F. Lillo, and G. Vinciguerra, "Why are learned indexes so effective?" in *ICML*, ser. Proceedings of Machine Learning Research, vol. 119. PMLR, 2020, pp. 3123–3132.
- [25] M. Maltry and J. Dittrich, "A critical analysis of recursive model indexes," *Proc. VLDB Endow.*, vol. 15, no. 5, pp. 1079–1091, 2022.
- [26] A. Kipf, R. Marcus, A. van Renen, M. Stoian, A. Kemper, T. Kraska, and T. Neumann, "Radixspline: a single-pass learned index," in *aiDM@SIGMOD*. ACM, 2020, pp. 5:1–5:5.
- [27] A. Hadian and T. Heinis, "Shift-table: A low-latency learned index for range queries using model correction," in *EDBT*. OpenProceedings.org, 2021, pp. 253–264.
- [28] M. Stoian, A. Kipf, R. Marcus, and T. Kraska, "PLEX: towards practical learned indexing," *CoRR*, vol. abs/2108.05117, 2021.
- [29] A. Llaveschi, U. Sirin, A. Ailamaki, and R. West, "Accelerating b+ tree search by using simple machine learning techniques," in *Proceedings of the 1st International Workshop on Applied AI for Database Systems and Applications*, 2019.
- [30] A. Hadian and T. Heinis, "Interpolation-friendly b-trees: Bridging the gap between algorithmic and learned indexes," in *EDBT*. OpenProceedings.org, 2019, pp. 710–713.
- [31] A. Kipf, D. Horn, P. Pfeil, R. Marcus, and T. Kraska, "LSI: a learned secondary index structure," in *aiDM@SIGMOD*. ACM, 2022, pp. 4:1–4:5.
- [32] Y. Wu, J. Yu, Y. Tian, R. Sidle, and R. Barber, "Designing succinct secondary indexing mechanism by exploiting column correlations," in *SIGMOD*. ACM, 2019, pp. 1223–1240.

- [33] V. Nathan, J. Ding, M. Alizadeh, and T. Kraska, "Learning multi-dimensional indexes," in *SIGMOD*. ACM, 2020, pp. 985–1000.
- [34] J. Ding, V. Nathan, M. Alizadeh, and T. Kraska, "Tsunami: A learned multi-dimensional index for correlated data and skewed workloads," *Proc. VLDB Endow.*, vol. 14, no. 2, pp. 74–86, 2020.
- [35] A. Behrend, A. Dignös, J. Gamper, P. Schmiegelt, H. Voigt, M. Rottmann, and K. Kahl, "Period index: A learned 2d hash index for range and duration queries," in *SSTD*. ACM, 2019, pp. 100–109.
- [36] J. Qi, G. Liu, C. S. Jensen, and L. Kulik, "Effectively learning spatial indices," *Proc. VLDB Endow.*, vol. 13, no. 11, pp. 2341–2354, 2020.
- [37] P. Li, H. Lu, Q. Zheng, L. Yang, and G. Pan, "LISA: A learned index structure for spatial data," in *SIGMOD*. ACM, 2020, pp. 2119–2133.
- [38] S. G. Pai, M. Mathioudakis, and Y. Wang, "Towards an instance-optimal z-index," in *4th International Workshop on Applied AI for Database Systems and Applications (AIDB@ VLDB2022)*, 2022.
- [39] V. Pandey, A. van Renen, A. Kipf, J. Ding, I. Sabek, and A. Kemper, "The case for learned spatial indexes," in *AIDB@VLDB*, 2020.
- [40] B. Spector, A. Kipf, K. Vaidya, C. Wang, U. F. Minhas, and T. Kraska, "Bounding the last mile: Efficient learned string indexing," *CoRR*, vol. abs/2111.14905, 2021.
- [41] Y. Wang, C. Tang, Z. Wang, and H. Chen, "Sindex: a scalable learned index for string keys," in *APSys '20: 11th ACM SIGOPS Asia-Pacific Workshop on Systems, Tsukuba, Japan, August 24-25, 2020*. ACM, 2020, pp. 17–24.
- [42] C. Tang, Y. Wang, Z. Dong, G. Hu, Z. Wang, M. Wang, and H. Chen, "Xindex: a scalable learned index for multicore data storage," in *PPoPP*. ACM, 2020, pp. 308–320.
- [43] P. Li, Y. Hua, J. Jia, and P. Zuo, "Finedex: A fine-grained learned index scheme for scalable and concurrent memory systems," *Proc. VLDB Endow.*, vol. 15, no. 2, pp. 321–334, 2021.