

DyTIS: A Dynamic Dataset Targeted Index Structure Simultaneously Efficient for Search, Insert, and Scan

Jin Yang¹, Heejin Yoon¹, Gyeongchan Yun¹, Sam H. Noh^{2*}, Young-ri Choi¹

¹UNIST, ²Virginia Tech

¹{yangjin, heejin5178, rugyoon, ychoi}@unist.ac.kr, ²samhnoh@vt.edu

Abstract

Many datasets in real life are complex and dynamic, that is, their key densities are varied over the whole key space and their key distributions change over time. It is challenging for an index structure to efficiently support all key operations for data management, in particular, search, insert, and scan, for such dynamic datasets. In this paper, we present DyTIS (Dynamic dataset Targeted Index Structure), an index that targets dynamic datasets. DyTIS, though based on the structure of Extendible hashing, leverages the CDF of the key distribution of a dataset, and learns and adjusts its structure as the dataset grows. The key novelty behind DyTIS is to group keys by the natural key order and maintain keys in sorted order in each bucket to support scan operations within a hash index. We also define what we refer to as a dynamic dataset and propose a means to quantify its dynamic characteristics. Our experimental results show that DyTIS provides higher performance than the state-of-the-art learned index for the dynamic datasets considered.

CCS Concepts: • Information systems → Data structures; Data management systems.

Keywords: dynamic datasets, index structure, key distribution

ACM Reference Format:

Jin Yang¹, Heejin Yoon¹, Gyeongchan Yun¹, Sam H. Noh^{2*}, Young-ri Choi¹. 2023. DyTIS: A Dynamic Dataset Targeted Index Structure Simultaneously Efficient for Search, Insert, and Scan. In *Eighteenth European Conference on Computer Systems (EuroSys '23)*, May 8–12, 2023, Rome, Italy. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3552326.3587434>

*This research was conducted while at UNIST.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. *EuroSys '23, May 8–12, 2023, Rome, Italy*

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9487-1/23/05...\$15.00
<https://doi.org/10.1145/3552326.3587434>

1 Introduction

For in-memory data management systems, such as in-memory databases and key-value stores [4, 12, 34, 35, 56], the efficiency of the index structure is critical, strongly affecting the final latency of the systems [21, 44, 57, 62]. Moreover, even for disk-based data management systems where, previously, storage I/O latency had been the dominant performance bottleneck [5], the software layer (i.e., the index structure) becomes important with the recent advent of fast storage systems such as the Samsung Z-SSD [54], KIOXIA XL-FLASH [17], and CXL SSDs [33]. Traditional indexes such as B+-trees and hash indexes are excellent structures that have been embraced by data management systems. They are an integral part of systems today supporting three key operations, specifically, search, insert, and scan. Their limitations, however, are that each excels for one operation but may not support another efficiently. For example, while hash indexes perform superbly for individual key search, they do not support range queries well, while in contrast, the B+-tree is the other way around. Recently, learned index structures have been proposed that excel for all three operations [21, 26, 57]. However, there is a catch here as well. Learned indexes need to be trained, and training, in recent work, has been done in the form of bulk loading [21, 26, 57]. When the characteristics of a dataset are static and simple, they are excellent. However, once this assumption breaks, how much to train and when to train again become difficult hurdles to overcome, potentially resulting in degraded performance.

Consider that in a data management system, a *dataset* is defined by a key distribution (i.e., keys inserted to the system) and the insertion order of the keys. Many datasets in real life are actually complex and dynamic such that the key distribution is not uniform and also changes as keys are inserted. In this paper, we present DyTIS (Dynamic dataset Targeted Index Structure), an index structure that competitively supports all search, insert, and scan operations, does not require a training phase (i.e., bulk loading), and is especially effective for dynamic datasets. DyTIS, though based on Extendible hashing [25], leverages the CDF (Cumulative Distribution Function) of the key distribution of a dataset, and learns and adjusts its structure as the dataset grows. The key novelty behind DyTIS is to group keys by the natural key order and maintain keys in sorted order in each bucket to support scan operations within a hash index. A natural

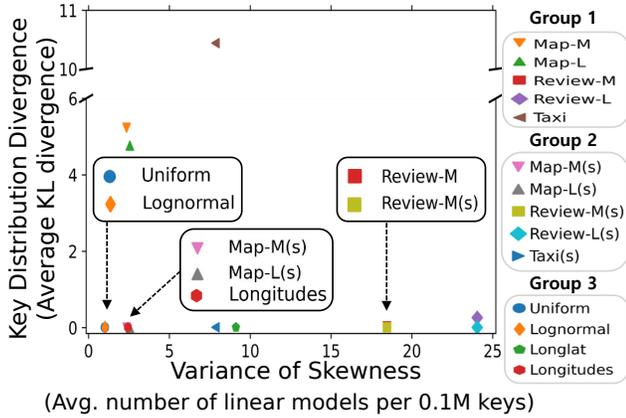


Figure 1. Dynamic characteristics of various datasets

question, then, is how to manage the hash index when the dataset becomes non-uniform, that is, unbalanced or skewed, which is typical of datasets. To efficiently handle such key distributions, DyTIS employs *remapping functions* that redistribute the non-uniform keys into a uniform distribution while preserving the natural order of the keys.

Aside from learned indexes, to the best of our knowledge, DyTIS is the first index structure that is simultaneously efficient for all search, insert, and scan operations. Unlike learned indexes, DyTIS focuses on realistic real-world datasets with various dynamic characteristics that have *complex key insertion order* such that while keys are inserted over time, the distribution of inserted keys continuously changes.

The main contributions of this work are as follows: First, we give our definition of a dynamic dataset and propose a means to quantify its dynamic characteristics. Second, we design an efficient index, DyTIS, which adjusts the CDF, as keys are inserted, to reflect the real key distribution. Finally, we implement DyTIS¹ and evaluate its performance by comparing it with three different dynamic indexing techniques over five different real-world datasets. Our experimental results demonstrate that DyTIS provides higher performance than the state-of-the-art learned index for the dynamic datasets considered.

2 Motivation and Design Direction

2.1 Dynamic Datasets

Dynamic Dataset: We first define what we mean by a dynamic dataset and show how this is quantified. A dynamic dataset is defined based on two observations; one, that the whole key space of a dataset can consist of small key ranges with different key densities [11, 45, 65] and two, that the key densities of key ranges, and thus, the key distribution of the dataset, change over time [42, 57]. The former is termed skewness, which refers to how dense keys are over a key

¹<https://github.com/unist-ssl/DyTIS>

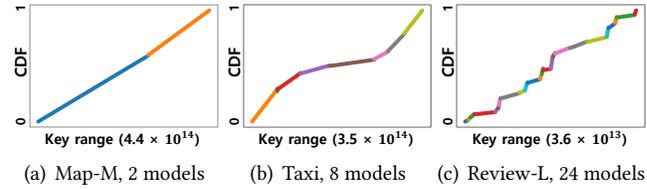


Figure 2. Variance of skewness for three datasets

range; if keys are dense (sparse), then skewness is high (low). The latter is termed distribution divergence, which captures the rate in which the key distribution of the dataset tends to change. We find that *variance of skewness* and *key distribution divergence* are the key factors that strongly affect the performance of indexes that leverage the key distribution of a dataset. This is because the number of linear models used to approximate the CDF and the retraining/adjusting of the linear models have strong impact on the performance of these indexes, which we discuss in detail below. Figure 1 shows how various datasets fare in terms of these dynamic characteristics. (The datasets are discussed in detail in Section 4.2.) Note that the Uniform dataset shows no skewness and no divergence in key distribution throughout the whole key space. We now discuss how these values are quantified and how they are to be interpreted.

Variance of Skewness: Variance of skewness is a metric we use to capture how skewness changes throughout the dataset. This represents the complexity of approximating the CDF of the key distribution of the key space. Considering that piecewise linear regression is used to approximate the CDF [64], if the skewness of each small key range is highly varied, a large number of linear models are needed. Figure 2 shows how the number of linear models varies for the Map-M, Taxi, and Review-L datasets that show low, medium, and high variance of skewness, respectively, where lines with different colors represent different linear models. Hereafter, for simplicity, we use skewness to refer to variance of skewness.

To quantify skewness for a dataset, we take the average number of linear models used to approximate the CDF for a fixed number of keys per key range. Note that we use a fixed number of keys, then average them to normalize their values as real-world datasets will all differ in size. The approximated CDF is measured by the maximum error-bounded Piecewise Linear Representation (PLR) technique [64]².

Key Distribution Divergence: For many real-world datasets, keys are inserted to a data management system in dynamic patterns. For example, in map datasets where keys are usually based on the longitudes and latitudes of locations [51], data with similar longitude and latitude values may be inserted consecutively as a bulk, while in sensing or weather

²<https://github.com/RyanMarcus/plr>. In this work, to construct the PLR of a dataset, the error bound is set such that the Uniform dataset only needs one linear model.

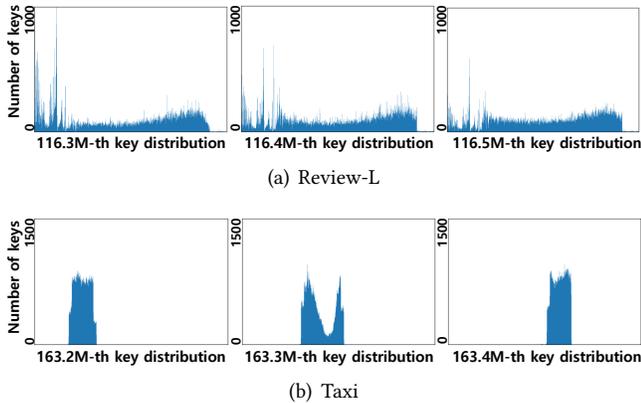


Figure 3. Key distribution divergence for two datasets

related datasets [9], data may be inserted with some periodic patterns such as diurnality and seasonality, and, as a final example, in datasets where a timestamp is used as a part of keys [28], keys may be inserted sequentially. Therefore, the distribution of data (i.e., inserted keys) typically changes over time.

Key Distribution Divergence (KDD) captures how vastly the distributions change. It is measured by adopting the Kullback–Leibler divergence (KL divergence) [41], which measures the difference between two probability distributions, where higher KL divergence value means that the two probability distributions are more different. To attain KDD, we divide the whole dataset into sub-datasets with a fixed number of keys and assume that each key in the sub-dataset is a random variable following the same discrete probability distribution. Then, the average of the KL divergence of every two consecutive sub-datasets is taken as the KDD for that dataset. To compute the KL divergence of two sub-datasets, we approximate the probability distribution of each sub-dataset using a histogram such that the key range of its histogram is determined by the minimum and maximum key values from the two sub-datasets. Note that for these metrics, the number of keys considered per key range and the size of a sub-dataset need to be chosen. While we choose 0.1 million for both, we also find the trend of the computed metrics are largely insensitive to this choice.

Figure 3 shows how the key distributions diverge over three consecutive sub-datasets for Review-L and Taxi that represent datasets with low and high KDD values, respectively. We observe that the three key distributions of Review-L are virtually all the same, while for those of Taxi, the three show quite different distributions even to the naked eye.

Revisiting Figure 1: We now revisit Figure 1 where there are three groups of datasets. Group 1 contains the dynamic datasets that we use throughout our experiments, while those in Group 2, marked with (s), are the shuffled version of those in Group 1. The shuffled datasets contain keys that are randomly shuffled so that the keys in the key space are inserted uniformly over time. We make the following observations

of Groups 1 and 2. First, real-world datasets (Group 1) show varying degrees of dynamism such as high skewness and low KDD (Review-M/L), to low skewness and moderate KDD (Map-M/L), to moderate skewness and high KDD (Taxi). Second, shuffled datasets (Group 2) have the effect of lowering the KDD value, that is, stabilizing the distribution of the keys. For datasets in Group 3, we observe little dynamic characteristics having both low skewness and KDD. We emphasize here that the target dataset of our study, Group 1, is different from Groups 2 and 3 used in previous studies [18, 21, 40].

2.2 DyTIS Design Focus and Philosophy

An index structure that will perform well for scan as well as insert and search operations has long been sought after. Recent developments seem to have finally found the solution in learned index structures [21, 26, 40, 57]. A learned index leverages the key distribution to build an index structure for keys stored in a sorted array. It uses machine learning algorithms such as neural networks and linear regression to approximate the CDF for the key distribution, and this CDF is used to predict the position of a given key in the sorted array. Of learned indexes, ALEX [21], a representative approach, maintains a data structure that looks like a B+-tree, where each internal node stores a linear model (i.e., an approximated CDF) and a pointer array to children nodes, and each leaf node (i.e., data node) stores a linear model and two arrays for keys and values. Unlike the original learned index [40] with the static recursive-model index (RMI), ALEX is dynamically adapted such that depending on the dataset, each data node can have a different depth. The index maintaining operations such as split, which is conceptually similar to the B+-tree split, or expansion, where the data node size increases with a scaled or retrained linear model, are selected by a cost model that is learned with bulk loading and, thereafter, re-learned during runtime. Through the adaptive RMI, ALEX can find a queried key by passing through multiple linear models in internal and data nodes.

Unfortunately, a couple of key constraints remain in learned indexes, especially for real-world datasets [11, 42, 45, 57, 65]. The first is that the model must first be pre-trained (i.e., bulk loaded). Currently, bulk loading is done manually in an ad-hoc manner as it is difficult to know how many keys should be used for learning so that the model sufficiently reflects the real distribution. This limitation naturally requires re-training of the model if the key distribution changes, so as not to degrade performance. Second, existing learned indexes leverage a hierarchy of models for either the whole key space [40, 57] (i.e., RMI) or certain key ranges with high key density [21] (i.e., adaptive RMI). Therefore, model re-training in a node may result in a cascade of retraining of other models such as the model of a parent or grandparent node, which will be an expensive operation. Moreover, the quality of the model in the upper hierarchy (such as the model of the root) will have stronger effect on the overall

performance. Third, considering that a large number of linear models must be built to approximate the CDF for a dataset with high skewness, existing learned indexes that allow only one linear model for a node [21, 40] need to create a large number of nodes for such a dataset, resulting in a complex index structure.

DyTIS targets dynamic datasets with high skewness and/or high KDD. When designing DyTIS, we took into account the valuable lessons from learned indexes. However, the following unique considerations also needed to be distilled into the design to efficiently support the dynamic nature of datasets.

1) Free of bulk loading: DyTIS learns and adjusts the CDF incrementally as keys are inserted without bulk loading overhead and also without performance degradation caused by inaccurate models built during bulk loading. **2) Local model re-training:** The cost of model retraining should be low to handle high KDD. DyTIS enables models for particular key ranges to be retrained locally and independently, making it more suitable for dynamic datasets. **3) Multiple models per node:** To accommodate key ranges with high skewness more efficiently, DyTIS allows a node to have multiple models.

3 DyTIS: The Index Structure

3.1 Extendible Hashing

Extendible hashing is a dynamic indexing structure that can *grow* and *shrink* as the workload dynamically changes [25]. It is unlike traditional hashing structures that require *rehashing* [20, 32, 48, 52, 53]. Figure 4 shows Extendible hashing, in which n is the number of bits used for a key K , and a *pseudo-key* $K' = h(K)$, where h is the hash function, is used for indexing [25]. In Extendible hashing, the hash table is organized with a *directory* and *buckets*. A directory (i.e., *dir*) is an array where prefixes, i.e., the most significant bits (MSBs), or suffixes, i.e., the least significant bits (LSBs), of the pseudo-key are used as an index. In Figure 4(a), two MSBs of a pseudo-key are used for the directory. The directory has a *global depth* GD that indicates how many MSBs are used as an index. Thus, GD determines the size of the directory.

Each entry in the directory points to a bucket that can store a fixed number of key-value pairs. Note that a value in each pair can be a pointer to an actual value associated with the key. When inserting a key-value pair, GD MSBs of its hash key are first used to find the directory entry, and then it is stored in the bucket pointed to by the directory entry (if there is free space). For example, Figure 4(a) shows the state after inserting the key “011101 ...” to the hashing index. Each bucket has a *local depth* LD , where $LD \leq GD$, indicating that this bucket contains all keys starting with LD MSBs of its associated directory index. Thus, two or more entries of the directory can point to the same bucket.

When a bucket finally overfills, the bucket must be split into two buckets. There are two situations when a bucket

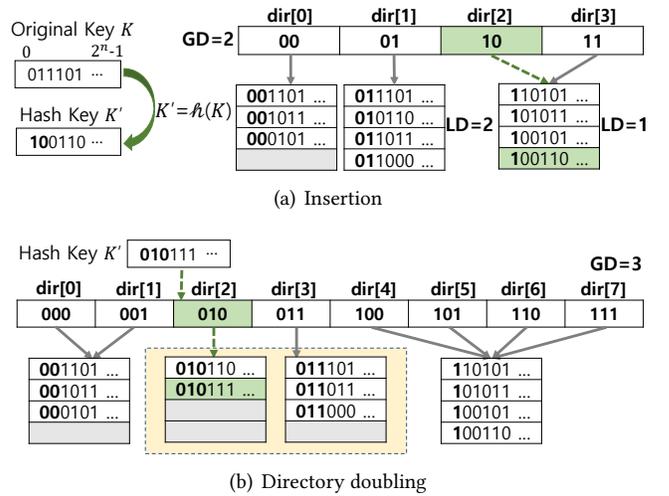


Figure 4. Extendible hashing

splits. The first is when $LD < GD$, which means that multiple directory entries were pointing to the bucket as in the case of the rightmost bucket in Figure 4(a). For this case, GD simply remains the same and LD of both buckets is increased by one, while the contents of the original bucket are adjusted according to the $LD + 1$ MSBs of the pseudo-keys. The second situation is when $GD = LD$. In this case, one bucket is pointed to by only one directory entry. Thus, the directory has to be expanded to accommodate the split, which is referred to as *directory doubling*. For this, the number of bits used as an index, that is, GD , must first be increased by one. Then, the dividing of the contents and the LD values are adjusted in the same manner as the first case. For example, as one more key is added to the bucket pointed to by *dir*[1] in Figure 4(a), the directory is doubled such that the values of LD for the buckets pointed to by *dir*[2] and *dir*[3] become three, while those for other buckets remain unchanged as shown in Figure 4(b).

CCEH [49] is a variant of Extendible hashing that uses a three-level structure of a directory, intermediate segments composed of a number of buckets, and buckets. CCEH utilizes MSBs of a pseudo-key as a segment index, while it utilizes LSBs of the key as a bucket index in the segment. Having intermediate segments is useful in reducing the overhead of directory doubling. In DyTIS, we adopt this three-level structure of CCEH, but we do not use LSBs of the keys since keys are stored in sorted order to support scan operations.

3.2 System Overview

DyTIS has an Extendible hash table structure, but uses *remapped* keys as pseudo-keys, rather than hash keys, to preserve the natural order of keys. It also leverages CDFs like learned indexes, but, as we will see, not in exactly the same manner.

Figure 5 presents the architecture of DyTIS as well as a walk-through example that we will use. DyTIS consists of multiple Extendible Hashing (EH) tables and has a key range

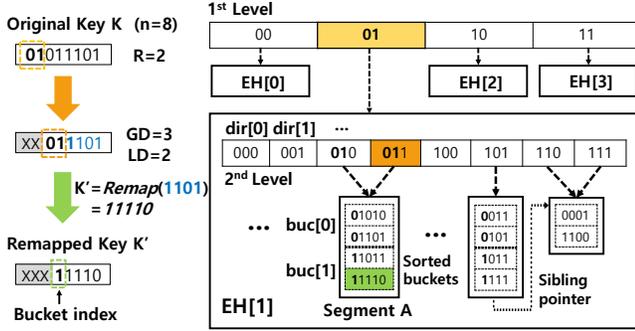


Figure 5. DyTIS architecture

of $[0, 2^n)$, where n is the number of bits used for a key. It has a two-level architecture so that the first level has an array of EHs. The first level *statically* divides the entire key range into 2^R sub-ranges based on R MSBs of the key so that each EH handles, in the second level, keys only in a given sub-range. This two-level structure allows each EH to deal with a smaller number of keys within the sub-range individually. Therefore, the range of keys for each EH becomes $[0, 2^{n-R})$ set by the $(n - R)$ LSBs from the keys. In Figure 5, $n = 8$ and $R = 2$, thus, each EH makes use of the six LSBs of the key.

An EH table in the second level has a three-level structure of a directory, segments, and buckets as in CCEH [49]. Each bucket is composed of an array of keys and an array of values such that a key and its value are stored in sorted order by the values of keys separately in the two different arrays [21]. Also, segments within the same EH maintain a sibling pointer to the next segment to accelerate scan performance. Note that buckets have a fixed size. In this three-level structure, each segment contains all keys starting with the same LD MSBs, where LD is the local depth of the segment. In the example, $LD = 2$ and thus, all keys with MSBs 01 are in Segment A. Consequently, the key range of a segment becomes $[0, 2^{n-R-LD})$ set by the $(n - R - LD)$ LSBs of the keys. In Figure 5, the four LSBs (for example, $1101_{(2)}$ for key $01011101_{(2)}$) are used resulting in a range of $[0, 2^4)$. Note that in Figure 5, instead of original keys, remapped keys, which are discussed next, are shown in Segment A.

We have, so far, only described how the bits (actually their positions) that comprise the key are used within DyTIS. Now we describe the key idea behind DyTIS, which is to use the raw keys themselves as pseudo-keys, rather than the hash keys, enabling efficient scan operations. However, the key predicament here is that non-uniform key distributions may result in a huge and imbalanced directory where a large number of keys exist in a certain range, while in the rest of the ranges, only a few keys exist so that many entries in the directory are pointing to the same bucket. Such situations will result in high directory doubling cost. To remedy this issue, DyTIS uses *remapping* functions that re-distribute keys with any distribution uniformly. That is, in each segment

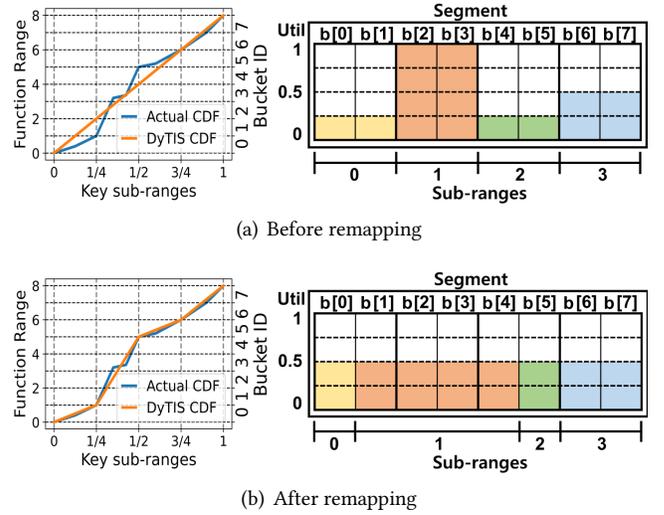


Figure 6. Adjusting the remapping functions

within the EH, if DyTIS deems the keys within the segment to be skewed, *remapping* spreads the keys evenly over the segment. This is depicted as $K' = \text{Remap}(1101)$ in Figure 5.

DyTIS leverages the CDF for a key distribution as the remapping function of each segment. We exploit the fact that since the CDF is a monotonic increasing function, each key K has a unique K' value, and thus, all the keys with any distribution mapped using its CDF are mapped to a uniform distribution. In DyTIS, the CDF for the remapping function is approximated as a set of linear functions (similarly to other learned index techniques [21, 28, 38, 57]). That is, given the key range of a segment, the key range is *statically* divided into 2^P sub-ranges based on P MSBs of the given keys, and each sub-range is associated with a linear function.

Initially, we assume that keys are uniformly distributed within each sub-range, initializing the linear function as $K' = K$, as shown in the left sub-figure of Figure 6(a) where there are four sub-ranges for a segment. As keys are inserted into EH, we incrementally and dynamically adjust the linear function (i.e., slope and intercept) for each sub-range to reflect the real distribution as in the left sub-figure of Figure 6(b). For a sub-range where more keys are inserted than predicted (for example, $[1/4, 1/2)$), we increase the slope of its linear function, while for a sub-range where less keys are inserted than predicted (for example, $[0, 1/4)$), we decrease the slope of the function. Moreover, intercept values of linear functions need to be modified such that the functions are connected to handle the entire range of the remapped keys. (Note that the right sub-figures of Figure 6 will be discussed in Section 3.3.)

Along with remapping functions, to handle possibly highly skewed key distributions without directory doubling or split, DyTIS also exploits segments with varying sizes such that DyTIS can dynamically adjust the number of buckets that a segment can hold. We now discuss how a key is rescaled, i.e.,

how the bucket index of a key is computed, within a segment with multiple buckets. Consider a segment where $K' = F(K)$, with K and K' being the raw key and the remapped key, respectively, and F is the scaled approximate CDF for the distribution to the function range (i.e., the range of y -axis of F). The function domain (i.e., the range of x -axis of F) is the same as the key range of the corresponding segment, which is $[0, 2^{n-R-LD})$, while, initially, when the segment has only one bucket, the function range is $[0, 2^{n-R-LD})$. As the number of buckets in the segment, B , increases, this range will proportionally extend to $[0, B \times 2^{n-R-LD})$. Accordingly, function range $[i \times 2^{n-R-LD}, (i+1) \times 2^{n-R-LD})$ corresponds to a bucket index i , where $0 \leq i < B$. Therefore, with a given remapped key K' , the bucket index, where K' is stored, can be computed by dividing K' with 2^{n-R-LD} , the size of the function domain. Note that in general, the number of buckets in a segment need not be a multiple of 2 as described in Section 3.3 and thus, only the quotient is used since the bucket index is an integer. Also, note that in Figure 6(a), the domain of the function is given as $[0, 1)$. A similar simplification is made for Figure 7 with the function domain and range. While this is to simplify the explanations, the actual function domain and range are as discussed above.

3.3 Individual Operations

Search: We explain how the search operation works using the walk-through example in Figure 5. In this figure, we have an 8 bit key (i.e., $n = 8$) and use two bits in the first level (i.e., $R = 2$). Assume we want to find key $K = 01011101_{(2)}$. We find $\text{EH}[1]$ using two MSBs ($01_{(2)}$) of K . Within $\text{EH}[1]$, the $n - R$ LSBs ($011101_{(2)}$) of K are used to compute the index in the directory, whose size is always a multiple of 2, using GD MSBs. As GD is 3, we use three MSBs, $011_{(2)}$, to find the directory index, which points to segment A. Since LD of segment A is 2, its key range is $[0, 2^4)$ and the range of the remapping function is $[0, 2^5)$ as it has two buckets, where $b[0]$ and $b[1]$ correspond to $[0, 10000_{(2)})$ and $[10000_{(2)}, 100000_{(2)})$, respectively. Then, we remap the four LSBs of K ($1101_{(2)}$) to a new key $K' = 11110_{(2)}$, and find the bucket index to be 1 by dividing K' by $2^4(10000_{(2)})$.

Once the bucket is located, an exponential search [21] is performed on the key array. If K is found, DyTIS reads the value of K stored at the same location in the value array and returns it. Otherwise, it returns “not exist”.

Scan: For a scan operation, a starting key K and a scan key range c are given. DyTIS first finds the index of the directory entry d that could contain K . It finds the starting position by searching for K or the smallest key larger than K in the segment of $\text{dir}[d]$. If not found in that segment, the starting position is the first key in the first bucket of the segment of $\text{dir}[d+1]$. It then linearly reads c keys, possibly continuing on to subsequent buckets, segments, or EHs, or until it reaches the end of the index. For scanning multiple segments within an EH, sibling pointers of segments are used.

Algorithm 1 Insert(*key* K , *segment* s , *bucket* b)

```

1: If  $b$  in  $s$  is not full, insert  $K$  to  $b$  and return;
2:  $LD =$  a local depth of  $s$ ;
3:  $U_s =$  utilization of  $s$ ;
4: if  $LD < GD$  then
5:   if  $U_s > U_t$  then
6:     Do Split;
7:   else
8:     Do Remapping;
9:     If it fails, do split;
10:  end if
11: else if  $LD == GD$  then
12:   if  $U_s > U_t$  then
13:     Do Expansion;
14:   else
15:     Do Remapping;
16:   end if
17:   if the above fails then
18:     Do Directory Doubling;
19:   end if
20: end if

```

Insertion: DyTIS inherits the basic schemes of split and directory doubling of Extendible hashing. DyTIS, though, employs additional schemes, remapping, and expansion, to handle non-uniform key distributions. DyTIS is required to use the MSBs of remapped keys so that keys can be clustered by the natural key order in this way. Note that a remapped key is used to find the bucket index but the raw key is stored in the bucket. For the operations that alter the indexing structure, we utilize a threshold U_t that determines whether the utilization of a segment is high or low.

Algorithm 1 shows the insertion algorithm of DyTIS. To insert a key K , DyTIS first finds bucket b of segment s where the key will be stored, as in the search operation. DyTIS checks if key K already exists. If so, it performs an in-place update for K such that the updated value will be stored at the same location in the value array of b . Otherwise, if b is not full, DyTIS inserts K and the value of K to b (Line 1), which could shift keys larger than K and their values to maintain the sorted order.

If b is full, then there is no space to insert K . Depending on the local depth of segment s , and the current utilization of s , DyTIS performs one of the following operations: split, remapping, expansion, and doubling. Note that at the beginning, DyTIS only performs basic schemes of Extendible hashing until EH reaches a certain local depth, having a number of inserted keys, before leveraging the key distribution to build an index structure.

Low utilization of segment s indicates that the key distribution of s is currently non-uniform (because the bucket of interest is full, yet the overall utilization of the segment is low). Therefore, in such cases regardless of $LD < GD$ or

$LD == GD$, DyTIS attempts to adjust the remapping functions of s to alleviate the skewed distribution of keys in s (Lines 8 and 15). While adjusting the remapping functions, it might increase the size of segment s .

On the other hand, high utilization of s means that the current remapping functions for s are distributing keys uniformly to a good degree, but s lacks space to store more keys. Therefore, when $LD < GD$, it performs a split (Line 6). On the other hand, when $LD == GD$, DyTIS performs an expansion, which doubles the size of s while scaling the remapping functions (Line 13). This has the same effect of directory doubling for keys in segment s . As there is a cap on the segment size that is dependent on its local depth, which is described in the last part of this section, the expansion and remapping of a segment can fail if DyTIS cannot increase the size of s due to this limit. To handle the failure, when $LD == GD$, it performs directory doubling (Line 18). When $LD < GD$, if remapping fails, it performs a split (Line 9). Note that remapping and expansion do not increase data movement (i.e., memory copy) as they are performed to avoid segment splits that incur data movement (which is discussed in Section 4).

Next, we discuss each of the operations in Algorithm 1.

Directory Doubling (Line 18): DyTIS uses the same doubling mechanism as Extendible hashing described previously.

Split (Lines 6 and 9): DyTIS uses a similar split mechanism as Extendible hashing, but with the following difference. When the original segment is split into two new segments, it is possible that one of the new segments may contain most of the keys, if the keys are skewed. Therefore, for each of the new segments, DyTIS first computes the segment size that will accommodate the keys of a sub-range from the original segment and then doubles its size, while keeping the slope(s) of the remapping function(s) of the sub-range for the new segment. For example, assume that the original segment has four buckets where one bucket is used to store keys from the left half of its key range and the other three buckets are used to store keys from the right half of its key range. Then the segments are doubled such that one segment will have two buckets, while the other will have six buckets.

Expansion (Line 13): DyTIS simply doubles the size while scaling the remapping functions (i.e., doubling the slope). Therefore, a certain key range that originally uses one bucket will use two buckets after the expansion.

Remapping (Lines 8 and 15): Recall that for a segment s , we divide its key range into multiple sub-ranges. Initially, a segment has one sub-range, but may be divided into multiple sub-ranges as the following discussion. Consider that we are trying to insert key K into bucket b in segment s , but have found that b is full, but the utilization of s is low. It is at this condition that we are performing remapping. Thus, DyTIS first checks if the sub-ranges of s are fine-grained enough such that each linear function for each sub-range, denoted r , truly reflects the real CDF of r . For example, Figure 7(a) illustrates a case where the DyTIS CDF in $[1/4, 1/2)$ does

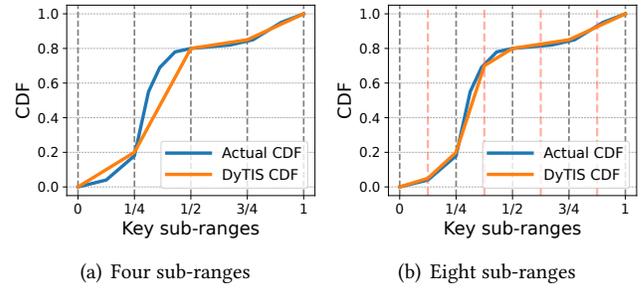


Figure 7. Dividing a sub-range into smaller ones

not approximate the real CDF accurately. In this sub-range, many keys exist in the left half of the sub-range while only a few keys exist in the right half. In such cases, the DyTIS CDF is too coarse and cannot represent the skewness of the sub-range. This results in the utilization of r being low and yet, the bucket in which K is to be inserted being full. To remedy this, remapping is employed. In remapping, we partition the key range of s into smaller sub-ranges until the target sub-range to which K will belong has utilization larger than U_t . Figure 7(b) shows the case where the four sub-ranges in Figure 7(a) are divided into eight sub-ranges. We observe that DyTIS CDF for $[1/4, 3/8)$ and $[3/8, 1/2)$ now becomes closer to the real CDF. For this, the remapping functions also need to be adjusted.

Adjusting the remapping functions, generally speaking, is done in the following manner. If there exists a sub-range whose utilization is low, then DyTIS simply steals some number of buckets from that sub-range and gives those buckets to the target sub-range t with high utilization.

Let us illustrate how remapping functions are adjusted using the example in Figure 6. In this example, $U_t = 0.5$. In the figure, a segment with 8 buckets ($b[0]$ to $b[7]$) is divided into four sub-ranges, having each sub-range correspond to a slope of 8 as in Figure 6(a). To be remapped as Figure 6(b), DyTIS steals two buckets, one each from sub-ranges 0 and 2 (i.e., $[0, 1/4)$ and $[1/2, 3/4)$) whose utilization is 0.25, and gives them to sub-range 2 (i.e., $[1/4, 1/2)$), altering the remapping functions as follows. The slopes of sub-ranges 0, 1, 2, and 3 are first computed as 4, 16, 4, and 8, and the intercepts of the sub-ranges are computed such that the remapping functions are started at $(0,0)$, $(1/4,1)$, $(1/2, 5)$, and $(3/4, 6)$, respectively. After remapping, sub-range 1 can use four buckets, lowering the utilization of the sub-range to 0.5 making it the same as the other sub-ranges.

Specifically, in DyTIS, we attempt to double the number of buckets for the target sub-range t . To do so, we compute how many buckets we can steal from other sub-ranges in segment s based on the utilization of a sub-range. For each sub-range (that is not t) whose utilization is less than U_t , we compute the minimum number of needed buckets. Based on this, we compute the number of buckets that this sub-range

can give to t . If we can steal all the needed buckets from other sub-ranges, we adjust the remapping functions accordingly.

The above process can fail if the utilization of other sub-ranges is also high and thus, DyTIS is prevented from stealing the needed buckets from other sub-ranges. In this case, DyTIS increases the size of segment s such that the number of buckets for sub-range t is doubled. Therefore, the slope of the remapping function for sub-range t is doubled and thus, the overall function range for segment s is increased proportionally to the number of buckets in s .

Once the new remapping functions are computed along with the size of the segment, DyTIS creates a new segment and copies each key from the old segment to the new segment using the new remapping functions. Once this is done, the segment location that points to the old segment, which is stored in the directory entry, is made to point to the new segment. Then, the old segment is deleted.

Deletion: To delete a key K , DyTIS searches for the bucket where K is stored and deletes it from the bucket. This will shift keys that are larger than K and also their values, if any. Similar to ALEX [21], if deletion causes high under-utilization of a bucket, the bucket can be merged with others, reducing the size of the segment. This process is similar to remapping but in the opposite direction.

Selecting a segment size: DyTIS basically imposes a limit on the segment size that is dependent on the local depth of the segment. This limit doubles when the local depth increases by one so that with a larger local depth, a larger number of keys can be stored in a segment. The effect of the segment size on performance will differ depending on the type of key distribution. With highly skewed keys, DyTIS will invoke remapping frequently and thus, a large segment size will increase remapping overhead. With uniform keys, then DyTIS leverages expansion and thus, a large segment size is beneficial as more keys can be handled without increasing the local or global depth. DyTIS optimizes performance considering both cases as follows. Assume DyTIS starts to perform remapping and expansion at local depth L . Until DyTIS reaches a certain local depth L' where $L' > L$, it keeps track of the occurrences of the expansion operations. At local depth L' , if the portion of segments where expansion is performed is large, this means that the keys are uniformly distributed. Thus, DyTIS increases the limit on the segment size at L' . Note that the value of L' needs to be selected so that DyTIS can collect sufficient information about the occurrences of split, remapping, and expansion operations, but, at the same time, early enough such that not too many doubling operations occur. For this, we empirically decide that $L' = L + 2$ and use this in our implementation.

3.4 Concurrency

With the advance of multi-core technologies, a server node commonly used for data management systems can leverage a large number of cores, which requires efficient concurrent

accesses to the index structure [14, 19, 30, 31, 44, 57, 58]. Providing efficient concurrency support is critical in sustaining performance scalability as the number of cores continues to increase [19, 24, 44]. Concurrency is supported in DyTIS so that it can be used for a multi-threaded system such as Memcached [1, 27]. Note that storage systems, developed for distributed clusters and/or multi-core servers may leverage multiple single-threaded engines for data access as in H-Store [6, 35] and Redis Cluster [3, 12]. Such systems may also use the single-threaded version of DyTIS that does not use locks instead of the multi-threaded one discussed in this section. Recall that we adopt the structure of CCEH [49], which supports concurrent accesses based on the two-level locking scheme adapted from Ellis [23]. Inherently, DyTIS adopts two levels of locking for each EH. At the high level, the directory array *dir* needs to be synchronized, that is, at the EH level, while at the lower level, we synchronize at the segment level. For segments, synchronization is done through a metadata container that we call a segment object, which contains information such as remapping functions of the segment and a pointer to a data array where key-value pairs are actually stored.

Using reader/writer locks, synchronization is administered only at the segment level when performing insert operations such as the normal insert (which inserts a key to a bucket without altering the indexing structure), remapping, and expansion, as these operations only change values of the information inside a segment object. Similarly, for search and scan operations, segment level locking is used. For scan, multiple segments, within or across EHS, are locked one by one over the scan range, and once done, all the locks are released.

For insert operations that change the structure such as the split, which creates a new segment object and then updates the segment object pointer to the new one, both segment and EH level locks are used to prevent other threads from performing directory doubling. Similarly, directory doubling or updating a sibling pointer between two segments require segment and EH level lock synchronization.

Note that CCEH leverages concurrency at finer grains of buckets within segments. We also explored this, but found that performance of DyTIS generally degrades. Our analysis shows that this is due to the overhead of additional memory for the fine-grained locks and the handling of segments with variable sizes, which is unlike CCEH that uses a fixed segment size.

4 Experimental Results

4.1 Methodology

For our experimental study, we use a machine with two Intel Core i9-9900K (8 core, 3.6Ghz) with 16MB L3 cache and 64GB DDR DRAM. Ubuntu 18.04 LTS with Linux kernel version 5.4 was installed on the machine. Hyperthreading is turned off during our experiments. To evaluate DyTIS, we compare

its performance with the STX B+-tree³ [2, 10] (hereafter, simply referred to as B+-tree), ALEX [7, 21], and XIndex [8, 57], which uses a two-level architecture for learned models, where the first level uses a learned RMI while the second level uses linear models and supports concurrent operations. For B+-tree, the fanout is set to 128 that shows the best performance in our setup. ALEX and XIndex require bulk loading of datasets. Thus, 70% of each dataset are bulk loaded for XIndex, while 10% or 70% are used for ALEX. Note that in the original ALEX paper [21], the authors use roughly 10~50% of the datasets to train the index. A discussion of the effect of bulk loading is given in Section 4.3. Also, unlike DyTIS and XIndex, the original ALEX and B+-tree do not support in-place updates. For our experiments, we modified ALEX and B+-tree to support in-place updates as suggested within the ALEX code⁴. For DyTIS, in the default setting, the array size in the first level is set to 2^9 (i.e., $R = 9$), U_t is set to 0.6, and the bucket size is 2KB. The local depth L_{start} where DyTIS starts remapping and expansion, is set to 6. For datasets with large expansion operations (which is determined by DyTIS dynamically), at local depth 8, a limit on the segment size $Limit_{seg}$ is increased to 128 times (from 2 times by default). We will discuss the effect of these parameter settings in Section 4.3. We run the experiments with a single thread, except for those in Section 4.5, and report the average of 3 runs.

4.2 Datasets

Table 1 shows the five real-world datasets with different key distribution characteristics used in our experiments where MM, ML, RM, RL, and TX denote the Map-M, Map-L, Review-M, Review-L, and Taxi datasets, respectively, and L, M, and H denote low, medium, and high skewness or KDD, respectively. For each record in the datasets, the size of the keys is configured to 8 bytes as done in prior works [21, 26, 40, 57] and also the size of the values is configured in the same way. In the table, for each dataset, we provide the number of keys (in millions: M), the key range size (i.e., the difference between the minimum and maximum key values), and the dataset size.

For the datasets, we adopt real-world datasets that have been widely used in existing indexing studies [21, 22, 28, 57, 62]. To generate unique integer keys, transformations are applied with selected fields in the real-world datasets. MM and ML are generated by combining the longitudes and latitudes of two different continents, respectively, from the OpenStreetMap [29, 51] (similarly to ALEX). RM and RL are generated from the Amazon review data [46]. From the original data, we select three fields, item ID, user ID, and review time, and generate unique keys by concatenating them. TX

³We chose this open-sourced B+-tree implementation as it was the choice of multiple earlier studies on efficient index structures [21, 45, 57].

⁴<https://github.com/microsoft/ALEX/blob/master/src/core/alex.h>, lines 1123 and 1124

Table 1. Datasets used in our experiments

Name	Description	Number of keys	Key range size ($\times 10^{18}$)	Dataset size	Skewness, KDD
MM	South America	356M	5.92	7.0GB	L,M
ML	Africa	903M	15.4	18GB	L,M
RM	Deduplicated data	82M	1.31	1.6GB	H,L
RL	Ratings only	228M	5.18	4.5GB	H,L
TX	Taxi trip in New York	325M	1.08	6.7GB	M,H

consists of pickup time and drop-off time fields of yellow taxi trips in New York City from 2017 to 2020 in the TLC Trip Record data [15]. Note that unlike ALEX [21], where the keys are randomly shuffled, our datasets are generated from the original datasets without shuffling and thus, the key distribution can dynamically shift over time as discussed in Section 2.1.

4.3 Results with Real-World Workloads

We evaluate the performance of the four different index structures, DyTIS, ALEX, XIndex, and B+-tree, using seven workloads that roughly correspond to workloads Load, A, B, C, D, E and F of YCSB [16]. For ALEX, two versions with 10% and 70% bulk loading, ALEX-10 and ALEX-70, are used. Using such workloads for real-world datasets is similar to that done for the ALEX study [21], where for each workload except Load, a batch of the workload is repeated for at least 60 seconds. The Load workload is composed of 100% inserts, the A workload, 50% reads and 50% updates, the B workload, 95% reads and 5% updates, the C workload, 100% reads, and the E workload, 95% scans where the range value is set to 100, and 5% inserts. Note that unlike in the ALEX study [21], we also include workloads F, which consists of 50% reads and 50% read-modify-write, and D'. D' is the same as the original workload D in that it includes 5% inserts, but is different in that the 95% reads are selected from the existing keys, and not from the latest keys⁵. For all workloads except Load, the keys are selected from the datasets using Zipfian distribution (with the default Zipfian constant in YCSB, 0.99). Note that we also ran all the experiments with uniform distribution as well, finding the results to be similar with Zipfian distribution.

For workloads A, B, C, and F, the Load workload is first executed to insert all the records in the dataset to the index. In this Load phase, for ALEX, 10%/70% are bulk loaded and 90%/30% are inserted, while for XIndex, 70% are bulk loaded and 30% are inserted. Then, we run each workload such that the number of operations performed is more than 50% of the dataset size for at least 60 seconds. For workloads D' and E that include inserting new keys, 80% of the keys are loaded. Then, we start measuring the throughput until all the keys in the dataset are inserted.

Figure 8 shows the throughput of seven workloads of YCSB (i.e., Load, A, B, C, D', E, and F), respectively, for DyTIS, ALEX-10, ALEX-70, XIndex, and B+-tree. From the results of Load, a

⁵Measuring the performance of repeated batches makes the exact modeling of the D workload complex.

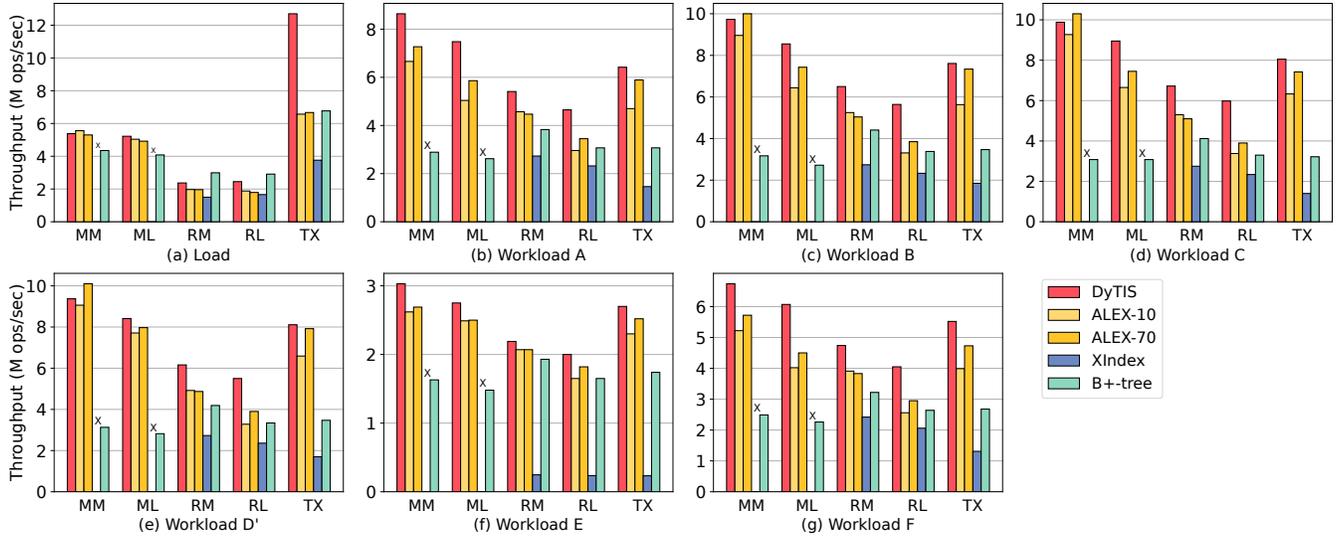


Figure 8. Throughput of seven real-world workloads over various datasets

purely insert workload, shown in Figure 8(a), we observe the following: First, DyTIS tends to provide better throughput than the other indexes for datasets with high KDD such as TX and for datasets with medium KDD and low skewness such as ML, as the local adjustment of its model is effective when the key distribution changes. Second, for datasets with high skewness such as RM and RL, DyTIS performance is lower (24.86% lower on average) than that of B+-tree because of remapping overhead. However, it is still higher than those of the learned indexes, ALEX-10, ALEX-70, and XIndex. Third, DyTIS shows better insertion performance than ALEX for more dynamic datasets of RM, RL and TX. Note that for ALEX and XIndex, the results do not include bulk loaded keys. Now, from the results of a purely search workload C shown in Figure 8(d), we observe that the search throughput of DyTIS is the highest, with DyTIS being higher than the other indexes except for MM, where ALEX-70 is 6.8% better than DyTIS.

The results of the above insert and search operations exemplify the key differences between DyTIS and ALEX and their performance implications. Recall that DyTIS uses a hash-based structure with segments, each with multiple models, while ALEX uses an adaptive RMI based structure with data nodes, each with a single model. Due to these structural differences, for insert operations, with DyTIS, the overhead for maintaining the index structure changes due to splits, expansions, and remappings is low when KDD is relatively high, such as for ML and TX, as the local model is effectively adjusted. This overhead, though, increases with higher skewness, such as for RM and RL. With ALEX-10, we find that similar structure maintenance overhead is, compared to DyTIS, 50% higher on average as it needs to perform many more expensive operations to retrain or create models for datasets

with high skewness or high KDD. For MM, ML, RM, RL, and TX with ALEX-10, the percent of those expensive operations over the total maintaining operations is 21%, 30%, 41%, 47%, and 24%, respectively. For TX, even though the size of TX is around 3 times smaller than ML, the absolute number of the maintaining operations and consequently, that of the expensive operations are around 10 and 8 times higher than those of ML. For search operations, we first note that to query a key, DyTIS always uses a linear model once, but ALEX uses at least two linear models, one at the root node and the other at a data node, with, possibly, more models in internal nodes of the tree. For the datasets, the average number of models used in ALEX-10 is up to $3.33\times$ (for RL) of that in DyTIS. Due to this, we find that the traversal overhead to traverse to a target bucket or data node for a given key with ALEX-10 is 22% higher on average than that with DyTIS. Finally, we find that the performance difference between DyTIS and ALEX for searching or inserting a key within a bucket or data node is not a dominant factor.

In Figure 8, we also observe that for MM, compared to DyTIS, the performance of ALEX-10 is higher for Load, while that of ALEX-70 is higher for workloads B, C, and D'. MM has low skewness and is smaller in size relative to ML. Also, its KDD is medium. Therefore, the built structure with ALEX-10 has minimum depth similar to the Uniform dataset, showing higher insert performance. Note that even though the index structure with ALEX-10 has minimum depth, its search performance is lower than DyTIS as the root node has a large number of children nodes (i.e. data nodes), which increases the traversal overhead. With ALEX-70, the effect of KDD for MM becomes much lower as it bulk loads a larger amount of the dataset compared to ALEX-10. Therefore, ALEX-70 ends up having the smallest number of nodes for the index among

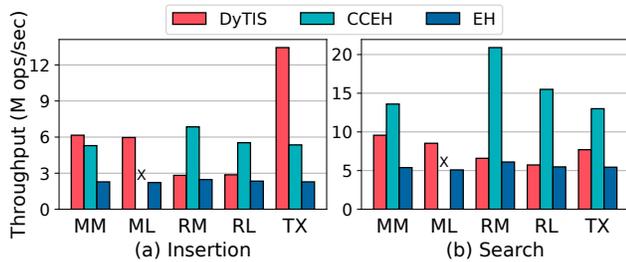


Figure 9. Comparison of DyTIS with CCEH and EH

all the datasets (which is 7% smaller than that with ALEX-10). Consequently, its traversal overhead becomes small, resulting in good performance as exemplified for workloads B, C, and D’.

From Figure 8, we also observe that aside from Load and MM with ALEX-70 as mentioned above, DyTIS performance is always the highest among all the indexes for all datasets as DyTIS provides good search and update performance. For workload E, we find that performance of B+-tree is lower than those of ALEX and DyTIS mainly due to its small data node size. The average data node size is $4.0\sim 295.1\times$ and $4.7\sim 319.2\times$ smaller than those of ALEX(-10 and -70) and DyTIS, respectively, forcing accesses to a much larger number of data nodes.

Overall, XIndex performs considerably worse than DyTIS or ALEX because it employs a delta index and a temporary delta index to handle key insertion along with a background compaction thread. Such additional structures and background thread overhead for merging data in the index and the delta index take their toll compared to DyTIS and ALEX, which are free of such structures and overhead. In the experiments, XIndex cannot load MM and ML (with 70% bulk loading) due to out-of-memory faults and thus, the subsequent workloads cannot run.

Next, we discuss the performance of DyTIS when using the datasets in Groups 2 and 3 of Figure 1, which are relatively simple ones. For the datasets in Group 2, which are the shuffled versions of those in Group 1, DyTIS shows the highest throughput among all the indexing techniques for all the YCSB workloads, except for Load with RM and RL, similarly to the original datasets, and for MM. The datasets in Group 3, which are used in the ALEX study [21], are much less dynamic than those in Group 1 such that their key distributions are easy to predict, which is ideal for traditional learned indexes. For the Uniform dataset of one billion keys, ALEX-10 shows 18.6% better throughput than DyTIS on average for the YCSB workloads. DyTIS, however, still provides up to $3.5\times$ (for workload C) higher throughput compared to B+-tree for all the workloads. For Longlat, which has the highest skewness in Group 3, DyTIS performs better (7.3~16.6%) than ALEX-10 for workloads A, E and F, while performing lower (2.5~6.4%) for workloads Load, B, C and D’.

Figure 9 shows the performance of DyTIS, CCEH, and Extendible Hashing (EH) described in Section 3.1. Figure 9 shows that DyTIS provides better insert and search performance than EH for all the datasets. It also shows that for insertion, CCEH and DyTIS give and take depending on the dataset, while for search performance, DyTIS is lower than CCEH by an average of 50%. Note that CCEH fails for ML due to out-of-memory. Generally speaking, DyTIS supports scans by replacing a hash function with a remapping function resulting in deteriorated search performance compared to CCEH, but which is still higher than B+-tree, ALEX-10, ALEX-70 (except for MM), and XIndex.

Effect of Bulk Loading on Learned Indexes We discuss how bulk loading affects the performance of ALEX and XIndex, which influences our first design consideration. As shown in Figure 8, ALEX-70 does not always provide better performance than ALEX-10. In particular, for RM, the performance of ALEX-10 is better than or similar to that of ALEX-70 for the YCSB workloads. Also, as shown in Figure 8(a), the load performance of ALEX-70 tends to be lower than that of ALEX-10. This is because as nodes are updated with splits and expansions due to insertion, the structure becomes more complex and more keys need to be moved. In the end, we find that the node size and depth of ALEX-70 are 337% larger and 26% deeper, on average, than ALEX-10 after bulk loading. To observe the influence of bulk loading further, we also evaluate ALEX with 30%, 50%, and 90% bulk loading for each dataset. Figure 10 shows the throughput of each workload normalized to that with ALEX-10. Note that for workloads D’ and E, 80% of the keys are loaded as in the earlier experiments except for ALEX-90 (which bulk-loads 90% of the dataset), and the number of operations for each of these workloads remains the same over various bulk loading percentages except for Load, and workloads D’ and E with ALEX-90. The key finding from the experiments is that no regularity can be found between load size and performance. For example, in the case of RM, as we increase bulk loading from 10% to 70%, performance generally decreases or remains similar, while for TX, ALEX-50 performs better than or similarly to ALEX-90, in most cases. Also, for MM and ML, ALEX-70/90 tend to provide better performance than ALEX with lower bulk loading (i.e., ALEX-10/30/50), while ALEX-10/30/50 show similar performance with each other. Overall, for our datasets, the performance difference over the different bulk loading percentages is as high as 21%, 31%, 36%, 40%, 42%, 14%, and 32% for the Load, A, B, C, D’, E, and F workloads, respectively. Our deeper analysis reveals that bulk loading is critical with ALEX as once the structure is built with a particular depth during bulk loading, ALEX vigorously deters increasing this depth.

Similar experiments were conducted for XIndex for TX, showing similar impact of bulk loading percentages on its performance to ALEX. We also find that for RM and RL, insertion failed with less than 70% bulk loading, conjecturing

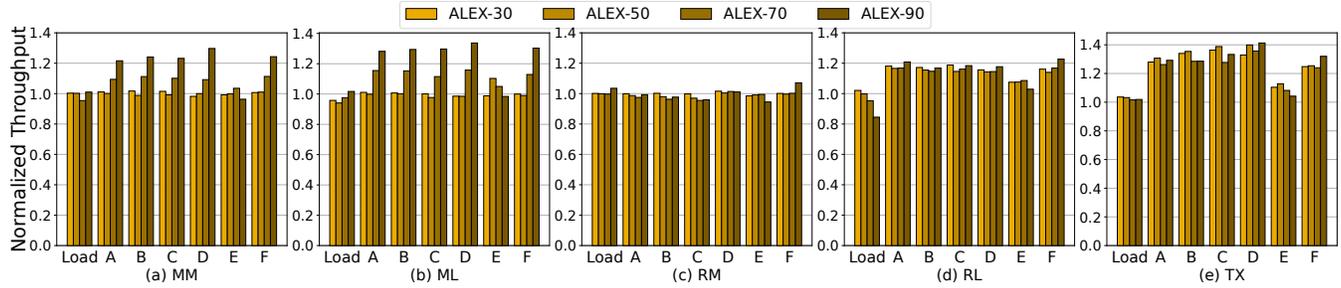


Figure 10. Throughput of ALEX over various bulk loading percentages normalized to ALEX-10

that memory access issues during insertion result in the failures. Thus, 70% bulk loading was chosen for the earlier experiments.

Insertion Breakdown and Tail Latency Analysis: We analyze the insertion execution time breakdown for DyTIS, that is, how much time it spends for each operation. We observe that for RM and RL with high skewness, remapping is leveraged the most. For TX with high KDD, it spends a relatively large portion of time for both remapping and expansion. We also analyze that the overhead of remapping is composed of memory copy overhead (58% on average over the five datasets) and remapping function adjustments (42%), and thus, it is proportional to the size of the segment. It is worth noting that considering a segment with size x , the overhead of remapping, which ends up having a size slightly larger than x , is similar to or only a bit larger than that of split, which ends up having a size of $2x$, for the segment. Furthermore, the expansion of x , which results in a segment of $2x$, has the same memory copy overhead as split, having similar overhead to remapping.

Table 2 presents the average, the 99th and the 99.99th percentile tail latency numbers of workloads Load and A, where best values are bold faced, for DyTIS, ALEX-10, ALEX-70, XIndex, and B+-tree. In the case of the Load workload, DyTIS is better than ALEX(-10 and -70) for dynamic datasets, RM, RL, and TX. However, B+-tree performs the best for most cases. For the 99.99th percentile tail latency, DyTIS performs worse than B+-tree due to the overhead for remapping large segments. However, we also see that the 99.99th tail latency of ALEX(-10 and -70) whose (re)training operations can be more expensive is $3.3\times$ and $3.1\times$ larger on average, respectively, than that of DyTIS. For workload A, DyTIS does better than ALEX-10 and ALEX-70 except for the average latency for TX with ALEX-70. Note that though not shown, for the other workloads, DyTIS provides better or similar performance compared to ALEX in most cases.

Memory Usage Analysis: We measure the maximum memory usage of each index structure using “dstat” where the amount of memory used for the benchmark execution and for bulk loading in the case of ALEX and XIndex is included, though, it is important to note that the memory needed for bulk loading is released once the bulk loading process is

done. We find that on average, the maximum memory usage of ALEX-10, ALEX-30, ALEX-50, ALEX-70, ALEX-90, and B+-tree is 25.9%, 26.3%, 27.0%, 22.6%, 5.3%, and 27.4% less than that of DyTIS for the YCSB workloads. DyTIS uses more memory because a segment consists of multiple buckets as in CCEH [49] and thus, each key must be stored in a particular bucket unlike ALEX and B+-tree. Interestingly, the memory usage of ALEX-90 increases (using only 5.3% less memory than DyTIS) as the sum of the memory needed for bulk loading and for the index structure when the bulk loading is done, which is the point of the maximum memory usage, becomes much larger. For XIndex, we observe that its memory usage is much higher than the others (e.g., $4.2\times$ higher than DyTIS).

Parameter Effect: We investigate how the control parameters affect DyTIS performance. We use throughput as our performance metric, averaged over all the datasets for the YCSB workloads, normalized to that of the default parameter value. With smaller B_{size} , insertion and search become cheaper as the overhead to shift keys for insert and to find a key within a bucket is reduced. However, more buckets need to be accessed for scan, degrading its performance. Over different B_{size} values, 1KB, 2KB (default), and 4KB, insertion, search, and scan throughput changes are $-16.3\sim 0.2\%$, $-10.1\sim 12.8\%$, and $-13.2\sim 2.8\%$, respectively. Larger L_{start} reduces remapping overhead, improving insertion performance, but generates more segments, degrading search and scan performance. On the other hand, smaller L_{start} degrades insertion performance due to structure changes based on a small number of inserted keys. Over different L_{start} values, 4, 6 (default), 8, and 10, insertion, search, and scan performance is affected by $-11.0\sim 7.3\%$, $-2.8\sim 0.3\%$, and $-5.5\sim 2.8\%$, respectively. Larger first level bit R increases the number of EHs having the effect of spreading keys and reducing rebalancing overhead, resulting in improved insertion performance. The performance change for insertion over different R values, 7, 9 (default), 11, and 13 is $-6.7\sim 6.2\%$. Lower utilization threshold U_t reduces remappings, improving insertion performance. However, this also decreases memory utilization. Over different U_t values, 0.5, 0.55, 0.6 (default), 0.65, and 0.7, insertion performance is affected by $-12.6\sim 6.8\%$. In addition, larger $Limit_{seg}$ causes degradation of insertion performance for datasets with high

Table 2. Average, 99th and 99.99th percentile tail latencies. na: not available (Best values are bold faced)

Average / 99th / 99.99th percentile latency (in ns)										
	Load					YCSB A				
	DyTIS	ALEX-10	ALEX-70	XIndex	B+-tree	DyTIS	ALEX-10	ALEX-70	XIndex	B+-tree
MM	223/624/43054	201/582/122640	210/573/131120	na/na/na	254/740/ 12050	185/541/797	207/627/954	197/610/920	na/na/na	361/800/2063
ML	233/643/51346	221/617/131380	224/583/119780	na/na/na	275/802/ 2253	211/586/828	257/712/1093	236/685/1044	na/na/na	401/842/2190
RM	477/1066/46086	542/1240/90494	538/1204/61704	668/1288/18195	370/810/2793	269/719/1008	274/743/1097	271/748/1434	376/892/2535	277/ 679/970
RL	460/1052/37576	560/1235/86570	578/1178/91276	618/1225/16448	374/852/2456	313/772/1473	396/971/2209	348/912/2178	445/1047/5806	339/773/1994
TX	97/140/26250	169/338/177400	165/324/172400	281/536/9360	184/307/ 1698	236/ 653/898	273/743/1287	230/674/950	697/1557/15996	340/770/2000

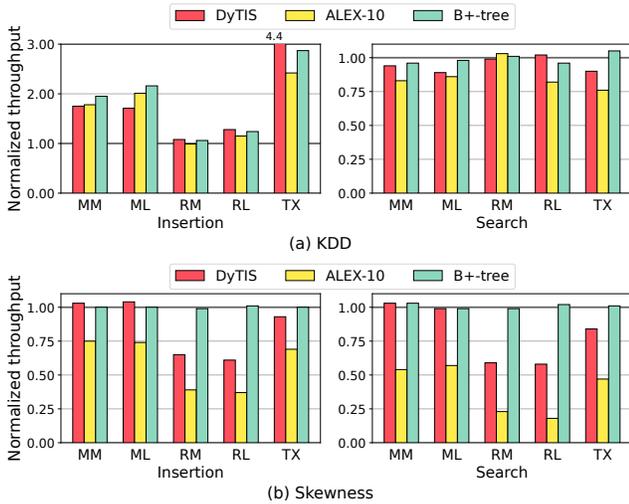


Figure 11. Influence of dynamic characteristics

skewness (like RM, RL and TX) due to frequent remappings, but for datasets with low skewness (like MM and ML), it has no effect on insertion performance while search and scan performance increases.

4.4 Effects of Dynamic Characteristics

In this section, we discuss how our second and third design considerations in Section 2.2 are manifested in supporting dynamic datasets by analyzing the effects of KDD and skewness of datasets on the performance of the indexes. We consider insertion and search, and use the results of Load and workload C of the previous section. To separately analyze the effect of KDD, we compare the performance of the original datasets where the key distribution changes over time with that of the shuffled versions that remove such changes. Figure 11(a) shows the insert and search performance of DyTIS, ALEX-10, and B+-tree for the original datasets, normalized to that for the shuffled versions. From these results, we observe the following: (1) In the case of insert, higher KDD has a more positive effect on performance for all the index structures. For RM and RL whose KDD is very small, the performance difference is relatively small compared to other datasets. For TX with high KDD, the performance improvement of DyTIS is as high as 339.76%. Similarly, MM and ML, with relatively high KDD, also show substantial performance differences. This is due to the spatial locality that exists in

the original dataset, and the indexes exploit this when inserting. (2) Let us now observe search, which is an operation that we call upon after the structure is constructed. First, we see that B+-tree is insensitive to KDD, which is logical as a balanced structure is constructed in the end, regardless of KDD. In contrast, we observe that DyTIS and ALEX-10 are affected somewhat, with ALEX-10 being more strongly so. Also, the performance degradation generally tends to be proportional to the KDD values of the datasets. The reason for this is that the structures were constructed in a more imbalanced manner due to the spatial locality (i.e., KDD) during insertion. However, ALEX-10 is more strongly affected as its structure becomes more complex with a larger number of linear models when the models in the upper hierarchy are inaccurate for a dataset with high KDD such as TX, unlike DyTIS which is based on remapping. One peculiar point is that of RL, which has relatively low KDD and yet, strongly influences ALEX-10. We find that though KDD of RL is not high, as RL also has very high skewness, this results in divergence affecting the built structure more strongly, increasing the number of linear models by 83,209 (20% increase).

We now analyze the effect of skewness, and to do so, we compare the performance for the shuffled versions of the original datasets with that for the Uniform datasets of the same size, which have no skewness. Figure 11(b) shows the performance for shuffled datasets normalized to that for the corresponding Uniform datasets. We observe the following: (1) For B+-tree, there is no performance effect caused by skewness of a dataset as the normalized performance for all the datasets is 1. (2) DyTIS is robust to low skewness, as we can see from the performance results for MM and ML. However, for datasets with moderate and high skewness, the performance of DyTIS degrades as DyTIS needs more models to accommodate such skewness. (3) Finally, we find ALEX-10 is sensitive to any degree of skewness. This is because, unlike DyTIS, ALEX-10 creates a node for every model resulting in a large number of nodes compared to the Uniform dataset. For RM and RL with high skewness, its performance degradation becomes severe as the number of nodes increases by 1341x on average, while, for DyTIS, there is only a 17x increase.

4.5 Concurrency

All results presented thus far were with single thread executions as ALEX and B+-tree are single thread solutions. In this section, we evaluate the performance of DyTIS with

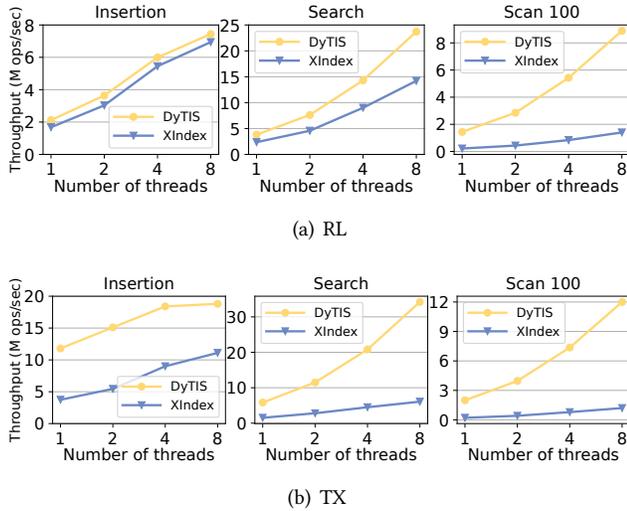


Figure 12. Throughput with different number of threads

XIndex over various numbers of threads. Figure 12 (a) and (b) show the throughput of DyTIS and XIndex on the RL and TX datasets, respectively, where requests for operations are assigned to threads in round-robin fashion. As shown, DyTIS provides performance better than that of XIndex for all operations. For TX, the insertion scalability of DyTIS increases only slightly when eight threads are used. The reason for this is that TX is highly dynamic such that a range of keys accessed within a certain time period concurrently by multiple threads is much smaller compared to other datasets. Thus, the benefit of higher concurrency becomes smaller. Note that though not shown, the performance trend of DyTIS and XIndex for RM is similar to that for RL.

5 Related Work

There are efforts to make an efficient updatable learned index [21, 26, 28, 57, 61]. LIPP [61]⁶ attempts to reduce the exponential search cost in the leaf node as well as to eliminate unbounded last-mile searches in ALEX. The PGM index has the structure of a linear piecewise regression model to predict the CDF of each segment where the data is partitioned [26]. The FITing-Tree also leverages piecewise linear functions [28], while APEX proposes a learned index that is optimized for persistent memory [43]. The RS index [38] that leverages a linear spine has been proposed as another way of approximating the CDF of the dataset. Flood [50] and Tsunami [22] were proposed as multi-dimensional learned indexes to address the limitation of existing learned indexes that consider one-dimensional data. However, Flood and

⁶In our setup, LIPP cannot build an index for 4 of the 5 datasets due to out-of-memory or type conversion errors. While we were able to insert for RM, we observed a huge number of key losses upon search.

Tsunami aim to optimize multi-dimensional read-only workloads and do not support insertion operations and concurrency. While most learned index designs consider only integer keys, there have been efforts to support variable-length string keys [55, 59]. Kipf et al. used real-world datasets for benchmarks for learned indexes but only focused on search performance over sorted datasets [37]. Simultaneously to our work, skewness of a dataset is similarly analyzed for learned indexes, where skewness is called “hardness” [60].

Bourbon [18] and SageDB [39] leverage learning of data patterns for database system components. Bourbon [18] replaces block indexes in Log-structured Merge Tree (LSM) with learned indexes, while SageDB [39] employs index structures and sorting algorithms based on learned models.

Hybrid indexing structures that use both the hash and B+-tree indexes have been investigated for large-scale memory-based data storage systems [36] and for DRAM-NVM memory systems [63]. HiKV proposes to use a hash index persisted to persistent memory, for single-key operations such as insert and search, while it also uses a B+-tree index, which resides in DRAM, to support scan operations [63]. Wormhole is an in-memory indexing scheme that makes use of the hash, B+-tree, and trie, focusing on string key optimizations [62].⁷

For concurrency, Masstree leverages fine-grained locking for updates and optimistic concurrency control [13, 47] for lookups [44]. Clevel hashing [14] supports high scalability by eliminating expensive lock-based concurrency control.

6 Conclusion

In this paper, we presented an index structure targeting dynamic datasets called DyTIS, which is simultaneously efficient for all search, insert, and scan operations. We also defined what dynamic datasets are and presented a means to quantify their characteristics. Our experimental results with real-world datasets demonstrate that DyTIS provides performance that is superior to ALEX, the state-of-the-art learned index.

Acknowledgement

We would like to thank our shepherd Angela Demke Brown and the anonymous reviewers for their invaluable comments. This research was partly supported by the MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program (IITP-2023-2021-0-01817) supervised by the IITP (Institute for Information & Communications Technology Planning & Evaluation) and National Research Foundation of Korea (NRF) (NRF-2022R1H1A2091297), and also by Samsung Electronics Co., Ltd. Young-ri Choi is the corresponding author.

⁷As it targets string keys, its performance is not high for our datasets that are composed of integer keys, as similarly discussed by Marcus et al. [45].

References

- [1] 2004. Memcached. <http://memcached.org/>.
- [2] 2007. STX B+-tree. <https://github.com/bingmann/stx-btree>.
- [3] 2009. Redis. <https://redis.io/docs/management/scaling/>.
- [4] 2011. Kyoto Cabinet: a straightforward implementation of DBM. <http://fallabs.com/kyotocabinet/>.
- [5] 2014. A fast and Lightweight Key-Value Store Library by Google. <http://code.google.com/p/leveldb>.
- [6] 2016. H-Store. <https://hstore.cs.brown.edu/>.
- [7] 2020. ALEX. <https://github.com/microsoft/ALEX/tree/57efb5>.
- [8] 2020. XIndex. <https://ipads.se.sjtu.edu.cn:1312/opensource/xindex>.
- [9] DF Andrews, AM Herzberg, DF Andrews, and AM Herzberg. 1985. Monthly mean sunspot numbers. *Data: A Collection of Problems from Many Fields for the Student and Research Worker* (1985), 67–74.
- [10] Timo Bingmann. 2007. STX B+ Tree C++ Template Classes. <https://panthema.net/2007/stx-btree/>
- [11] Zhichao Cao, Siying Dong, Sagar Vemuri, and David H.C. Du. 2020. Characterizing, Modeling, and Benchmarking RocksDB Key-Value Workloads at Facebook. In *Proceedings of 18th USENIX Conference on File and Storage Technologies (FAST)*. 209–223.
- [12] Josiah Carlson. 2013. *Redis in action*. Simon and Schuster.
- [13] Sang Kyun 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 VLDB Endowment*. 181–190.
- [14] Zhangyu Chen, Yu Huang, Bo Ding, and Pengfei Zuo. 2020. Lock-free Concurrent Level Hashing for Persistent Memory. In *Proceedings of the 2020 USENIX Annual Technical Conference (USENIX ATC)*. 799–812.
- [15] NYC Taxi & Limousine Commission. 2020. TLC Trip Record Data. <https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>.
- [16] Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. 2010. Benchmarking Cloud Serving Systems with YCSB. In *Proceedings of the ACM Symposium on Cloud Computing (SoCC)*. 143–154.
- [17] Toshiba Memory Corporation. 2019. Toshiba Memory Corporation Introduces XL-FLASH™ Storage Class Memory Solution. <https://www.kioxia.com/en-jp/about/news/2019/20190806-1.html>.
- [18] Yifan Dai, Yien Xu, Aishwarya Ganesan, Ramnathan Alagappan, Brian Kroth, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau. 2020. From WiscKey to Bourbon: A Learned Index for Log-Structured Merge Trees. In *Proceedings of 14th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 155–171.
- [19] Tudor David, Rachid Guerraoui, and Vasileios Trigonakis. 2015. Asynchronous Concurrency: The Secret to Scaling Concurrent Search Data Structures. *ACM SIGARCH Computer Architecture News* (2015), 631–644.
- [20] Martin Dietzfelbinger and Christoph Weidling. 2007. Balanced allocation and dictionaries with tightly packed constant size bins. *Theoretical Computer Science* 380 (June 2007), 47–68.
- [21] 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 (SIGMOD)*. 969–984.
- [22] Jialin Ding, Vikram Nathan, Mohammad Alizadeh, and Tim Kraska. 2020. Tsunami: A learned multi-dimensional index for correlated data and skewed workloads. In *Proceedings of the VLDB Endowment*. 74–86.
- [23] Carla Schlatter Ellis. 1983. Extendible Hashing for Concurrent Operations and Distributed Data. In *Proceedings of the 2nd ACM SIGACT-SIGMOD Symposium on Principles of Database Systems (PODS)*. 106–116.
- [24] Franz Faerber, Alfons Kemper, Per-Åke Larson, Justin Levandoski, Thomas Neumann, Andrew Pavlo, et al. 2017. Main Memory Database Systems. *Foundations and Trends® in Databases* (2017), 1–130.
- [25] Ronald Fagin, Jurg Nievergelt, Nicholas Pippenger, and H. Raymond Strong. 1979. Extendible Hashing—a Fast Access Method for Dynamic Files. *ACM Transactions on Database Systems (TODS)* (1979), 315–344.
- [26] Paolo Ferragina and Giorgio Vinciguerra. 2020. The PGM-index: a fully-dynamic compressed learned index with provable worst-case bounds. In *Proceedings of the VLDB Endowment*. 1162–1175.
- [27] Brad Fitzpatrick. 2004. Distributed Caching with Memcached. *Linux journal* (2004), 5.
- [28] Alex Galakatos, Michael Markovitch, Carsten Binnig, Rodrigo Fonseca, and Tim Kraska. 2019. FITing-Tree: A Data-aware Index Structure. In *Proceedings of the 2019 ACM SIGMOD International Conference on Management of Data (SIGMOD)*. 1189–1206.
- [29] Geofabrik. 2019. OpenStreetMap Data Extracts. <https://download.geofabrik.de/>.
- [30] Maurice Herlihy, Victor Luchangco, and Mark Moir. 2003. Obstruction-Free Synchronization: Double-Ended Queues as an Example. In *Proceedings of the 23rd International Conference on Distributed Computing Systems (ICDCS)*. 522–529.
- [31] Deukyeon Hwang, Wook-Hee Kim, Youjip Won, and Beomseok Nam. 2018. Endurable Transient Inconsistency in Byte-Addressable Persistent B+-Tree. In *Proceedings of the 16th Usenix Conference on File and Storage Technologies (FAST)*. 187–200.
- [32] LR Johnson. 1961. An Indirect Chaining Method for Addressing on Secondary Keys. *Commun. ACM* (1961), 218–222.
- [33] Myoungsoo Jung. 2022. Hello Bytes, Bye Blocks: PCIe Storage Meets Compute Express Link for Memory Expansion (CXL-SSD). In *Proceedings of the 14th ACM Workshop on Hot Topics in Storage and File Systems (HotStorage)*. 45–51.
- [34] Olzhas Kaiyrakhmet, Songyi Lee, Beomseok Nam, Sam H Noh, and Young-ri Choi. 2019. SLM-DB: Single-Level Key-Value Store with Persistent Memory. In *Proceedings of the 17th USENIX Conference on File and Storage Technologies (FAST)*. 191–205.
- [35] Robert Kallman, Hideaki Kimura, Jonathan Natkins, Andrew Pavlo, Alexander Rasin, Stanley Zdonik, Evan PC Jones, Samuel Madden, Michael Stonebraker, Yang Zhang, et al. 2008. H-Store: A High-Performance, Distributed Main Memory Transaction Processing System. In *Proceedings of the VLDB Endowment*. 1496–1499.
- [36] Ankita Kejriwal, Arjun Gopalan, Ashish Gupta, Zhihao Jia, Stephen Yang, and John Ousterhout. 2016. SLIK: Scalable Low-Latency Indexes for a Key-Value Store. In *Proceedings of the 2016 USENIX Annual Technical Conference (USENIX ATC)*. 57–70.
- [37] Andreas Kipf, Ryan Marcus, Alexander van Renen, Mihail Stoian, Alfons Kemper, Tim Kraska, and Thomas Neumann. 2019. SOSD: A Benchmark for Learned Indexes. In *Proceedings of NeurIPS Workshop on Machine Learning for Systems*.
- [38] Andreas Kipf, Ryan Marcus, Alexander van Renen, Mihail Stoian, Alfons Kemper, Tim Kraska, and Thomas Neumann. 2020. RadixSpline: A Single-Pass Learned Index. In *Proceedings of the Third International Workshop on Exploiting Artificial Intelligence Techniques for Data Management (aiDM)*. 1–5.
- [39] 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 (CIDR)*.
- [40] Tim Kraska, Alex Beutel, Ed H Chi, Jeffrey Dean, and Neoklis Polyzotis. 2018. The Case for Learned Index Structures. In *Proceedings of the 2018 ACM SIGMOD International Conference on Management of Data (SIGMOD)*. 489–504.
- [41] Solomon Kullback. 1997. *Information theory and statistics*. Courier Corporation.
- [42] Justin J. Levandoski, Per-Åke Larson, and Radu Stoica. 2013. Identifying Hot and Cold Data in Main-Memory Databases. In *Proceedings of 2013 IEEE 29th International Conference on Data Engineering (ICDE)*. 26–37.

- [43] Baotong Lu, Jialin Ding, Eric Lo, Umar Farooq Minhas, and Tianzheng Wang. 2021. APEX: A High-Performance Learned Index on Persistent Memory. In *Proceedings of the VLDB Endowment*. 597–610.
- [44] 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 (EuroSys)*. 183–196.
- [45] Ryan Marcus, Andreas Kipf, Alexander van Renen, Mihail Stoian, Sanchit Misra, Alfons Kemper, Thomas Neumann, and Tim Kraska. 2020. Benchmarking Learned Indexes. In *Proceedings of the VLDB Endowment*. 1–13.
- [46] Julian McAuley and Alex Yang. 2016. Addressing Complex and Subjective Product-Related Queries with Customer Reviews. In *Proceedings of the 25th International Conference on World Wide Web*. 625–635.
- [47] Paul E McKenney, Jonathan Appavoo, Andi Kleen, Orran Krieger, Rusty Russell, Dipankar Sarma, and Maneesh Soni. 2001. Read-copy update. In *AUUG Conference Proceedings*. 175.
- [48] Robert Morris. 1968. Scatter Storage Techniques. *Commun. ACM* (1968), 38–44.
- [49] Moohyeon Nam, Hokeun Cha, Young-ri Choi, Sam H. Noh, and Beomseok Nam. 2019. Write-Optimized Dynamic Hashing for Persistent Memory. In *Proceedings of the 17th USENIX Conference on File and Storage Technologies (FAST)*. 31–44.
- [50] Vikram Nathan, Jialin Ding, Mohammad Alizadeh, and Tim Kraska. 2020. Learning Multi-Dimensional Indexes. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD)*. 985–1000.
- [51] OpenStreetMap on AWS. 2021. Amazon AWS. <https://registry.opendata.aws/osm/>.
- [52] Rasmus Pagh and Flemming Friche Rodler. 2004. Cuckoo Hashing. *Journal of Algorithms* (2004), 122–144.
- [53] W. W. Peterson. 1957. Addressing for Random-Access Storage. *IBM Journal of Research and Development* (1957), 130–146.
- [54] SAMSUNG. 2017. Ultra-low latency with Samsung Z-NAND SSD. https://www.samsung.com/us/labs/pdfs/collateral/Samsung_Z-NAND_Technology_Brief_v5.pdf.
- [55] Benjamin Spector, Andreas Kipf, Kapil Vaidya, Chi Wang, Umar Farooq Minhas, and Tim Kraska. 2021. Bounding the Last Mile: Efficient Learned String Indexing. In *Proceedings of 3rd International Workshop on Applied AI for Database Systems and Applications (AIDB)*. 1–5.
- [56] Michael Stonebraker and Ariel Weisberg. 2013. The VoltDB Main Memory DBMS. *IEEE Database Engineering Bulletin* (2013), 21–27.
- [57] Chuzhe Tang, Youyun Wang, Zhiyuan Dong, Gansen Hu, Zhaoguo Wang, Minjie Wang, and Haibo Chen. 2020. XIndex: A Scalable Learned Index for Multicore Data Storage. In *Proceedings of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*. 308–320.
- [58] Josh Triplett, Paul E McKenney, and Jonathan Walpole. 2011. Resizable, Scalable, Concurrent Hash Tables via Relativistic Programming. In *Proceedings of the 2011 USENIX Annual Technical Conference (USENIX ATC)*. 11.
- [59] Youyun Wang, Chuzhe Tang, Zhaoguo Wang, and Haibo Chen. 2020. SIndex: A Scalable Learned Index for String Keys. In *Proceedings of the 11th ACM SIGOPS Asia-Pacific Workshop on Systems (APSys)*. 17–24.
- [60] Chaichon Wongkham, Baotong Lu, Chris Liu, Zhicong Zhong, Eric Lo, and Tianzheng Wang. 2022. Are Updatable Learned Indexes Ready? (2022), 3004–3017.
- [61] Jiacheng Wu, Yong Zhang, Shimin Chen, Yu Chen, Jin Wang, and Chunxiao Xing. 2021. Updatable Learned Index with Precise Positions. In *Proceedings of the VLDB Endowment*. 1276–1288.
- [62] Xingbo Wu, Fan Ni, and Song Jiang. 2019. Wormhole: A Fast Ordered Index for In-Memory Data Management. In *Proceedings of the 14th ACM European Conference on Computer Systems (EuroSys)*. 1–16.

- [63] Fei Xia, Dejun Jiang, Jin Xiong, and Ninghui Sun. 2017. HiKV: A Hybrid Index Key-Value Store for DRAM-NVM Memory Systems. In *Proceedings of the 2017 USENIX Annual Technical Conference (USENIX ATC)*. 349–362.
- [64] Qing Xie, Chaoyi Pang, Xiaofang Zhou, Xiangliang Zhang, and Ke Deng. 2014. Maximum error-bounded Piecewise Linear Representation for online stream approximation. In *Proceedings of the VLDB Endowment*. 915–937.
- [65] Adar Zeitak and Adam Morrison. 2021. Cuckoo Trie: Exploiting Memory-Level Parallelism for Efficient DRAM Indexing. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles (SOSP)*. 147–162.

A Artifact Appendix

A.1 Abstract

DyTIS (Dynamic dataset Targeted Index Structure) is an index that targets dynamic datasets. DyTIS, though based on the structure of Extendible hashing, leverages the CDF of the key distribution of a dataset, and learns and adjusts its structure as the dataset grows. The key novelty behind DyTIS is to group keys by the natural key order and maintain keys in sorted order in each bucket to support scan operations within a hash index.

The artifact is available as a public repository on Github at <https://github.com/unist-ssl/DyTIS>. The DOI of the artifact is <https://doi.org/10.5281/zenodo.7721895>.

A.2 Description & Requirements

A.2.1 Hardware dependencies. For our experimental study, we use a machine with two Intel Core i9-9900K (8 core, 3.6Ghz) with 16MB L3 cache and 64GB DDR DRAM.

A.2.2 Software dependencies.

- Ubuntu 18.04 LTS with Linux kernel version 5.4
- g++ 8.4
- libboost-dev 1.65.1

A.2.3 Benchmarks. We provide a sample dataset, *review-small.csv* (initial 1 million data of review-M), which is stored in Google Drive.

- *review-small.csv*: <https://bit.ly/3FjYJOZ>

A.3 Set-up

The instructions for set-up are referred to README.md in the repository of our source code.

- To install libboost: `sudo apt-get install libboost-all-dev`

A.4 Evaluation Workflow

A.4.1 Major claims.

- (C1): DyTIS efficiently supports all search, insert, and scan operations for dynamic datasets without requiring an up-front training phase (i.e., bulk loading).
- (C2): DyTIS shows the performance using seven workloads that roughly correspond to workloads Load, A,

B, C, D, E, and F of YCSB. This is proven by the experiment (E2) described in Section 4.3 whose results are illustrated in Figure 8.

A.4.2 Experiments.

Experiment (E1): Availability of source code [a few human-minutes]: It shows insert, search and scan performance of DyTIS.

[Preparation] Before running the experiment, the dataset to be used in the experiment should be downloaded.

[Execution] The experiment has been automated by the script (*scripts/run_benchmark.sh*). The script will automatically run the experiments of insert, search, and then scan workloads over a given dataset.

- To run benchmarks (insert, search and scan) on the review-small dataset in the path *data/* :
 - *./scripts/run_benchmark.sh data/review-small.csv*

[Results] All throughput and latency results are saved as a log file in *benchmark/result/*.

Experiment (E2): Real-world workloads [10 human-minutes]: It shows results with real-world workloads that roughly correspond to workloads Load, A, B, C, D, E, and F of YCSB.

[Preparation] Before running the experiment, the dataset to be used in the experiment should be downloaded and located in the proper path (e.g., in *data/review-small.csv*). If the dataset is not the aforesaid one (e.g., different dataset or different path), modify *dataset_files* properly in the script *scripts/run_ycsb_style_exp.sh*.

[Execution] The experiment has been automated by the script (*scripts/run_ycsb_style_exp.sh*). The script will automatically run the experiments of seven real-world workloads over a given dataset.

- To run real-world workloads:
 - *./scripts/run_ycsb_style_exp.sh*

[Results] All throughput and latency results for each workload are saved as a log file in *benchmark/result/*.