Order-Preserving Key Compression for In-Memory Search Trees

Huanchen Zhang
Carnegie Mellon University
huanche1@cs.cmu.edu

Xiaoxuan Liu
Carnegie Mellon University
xiaoxual@andrew.cmu.edu

David G. Andersen
Carnegie Mellon University
dga@cs.cmu.edu

Michael Kaminsky
BrdgAI
kaminsky@cs.cmu.edu

Kimberly Keeton
Hewlett Packard Labs
kimberly.keeton@hpe.com

Andrew Pavlo
Carnegie Mellon University
pavlo@cs.cmu.edu

ABSTRACT

We present the *High-speed Order-Preserving Encoder* (HOPE) for in-memory search trees. HOPE is a fast dictionary-based compressor that encodes arbitrary keys while preserving their order. HOPE's approach is to identify common key patterns at a fine granularity and exploit the entropy to achieve high compression rates with a small dictionary. We first develop a theoretical model to reason about order-preserving dictionary designs. We then select six representative compression schemes using this model and implement them in HOPE. These schemes make different trade-offs between compression rate and encoding speed. We evaluate HOPE on five data structures used in databases: SuRF, ART, HOT, B+tree, and Prefix B+tree. Our experiments show that using HOPE allows the search trees to achieve lower query latency (up to 40% lower) and better memory efficiency (up to 30% smaller) simultaneously for most string key workloads.

CCS CONCEPTS

• **Information systems → Data compression**.

ACM Reference Format:
Huanchen Zhang, Xiaoxuan Liu, David G. Andersen, Michael Kaminsky, Kimberly Keeton, and Andrew Pavlo. 2020. Order-Preserving Key Compression for In-Memory Search Trees. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD'20), June 14–19, 2020, Portland, OR, USA.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3318464.3380583

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.
SIGMOD'20, June 14–19, 2020, Portland, OR, USA
© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6735-6/20/06…$15.00
https://doi.org/10.1145/3318464.3380583

1 INTRODUCTION

SSDs have increased the performance demand on main-memory data structures; the growing cost gap between DRAM and SSD storage together with increasing database sizes means that main-memory structures must be both compact *and* fast. The $/GB ratio of DRAM versus SSDs increased from 10× in 2013 to 40× in 2018 [8]. Together with growing database sizes, database management systems (DBMSs) now operate with a lower memory to storage size ratio than before. DBMS developers in turn are changing how they implement their systems' architectures. For example, a major Internet company's engineering team assumes a 1:100 memory to storage ratio to guide their future system designs [22].

One challenge with assuming that memory is severely limited is that modern online transaction processing (OLTP) applications demand that most if not all transactions complete in milliseconds or less [46]. To meet this latency requirement, applications use many *search trees* (i.e., indexes, filters) in memory to minimize the number of I/Os on storage devices. But these search trees consume a large portion of the total memory available to the DBMS [16, 30, 50].

Compression is an obvious way to reduce the memory cost of a DBMS's search trees. Compression improves the cache performance of the search tree and allows the DBMS to retain more data in memory to further reduce I/Os. A system must balance these performance gains with the additional computational overhead of the compression algorithms.

Existing search tree compression techniques fall into two categories. The first is *block-oriented compression* of tree pages using algorithms such as Snappy [9] and LZ4 [7]. This approach is beneficial to disk-based trees because it minimizes data movement between disk and memory. For in-memory search trees, however, block compression algorithms impose too much computational overhead because the DBMS is unable to operate directly on the search tree data without having to decompress it first [50]. The second approach is to design a *memory-efficient data structure* that avoids storing unnecessary key information and internal

meta-data (e.g., pointers) [17, 33, 40, 50, 51]. Although these new designs are smaller than previous implementations, they still are a major source of DMBSs' memory footprints.

An orthogonal approach is to compress the individual input keys before inserting them into the search tree. Key compression is important for reducing index memory consumption because real-world databases contain many variable-length string attributes [42] whose size dominates the data structure's internal overheads. A common application of string compression is in columnar DBMSs [11], which often use dictionary compression to replace string values in a column with fixed-length integers. Traditional dictionary compression, however, does not work for in-memory search trees (e.g., OLTP indexes) for two reasons. First, the DBMS must continually grow its dictionary as new keys arrive. Second, key compression in a search tree must be order-preserving to support range queries properly.

We, therefore, present **High-speed Order-Preserving Encoder** (HOPE), a dictionary-based key compressor for in-memory search trees (e.g., B+trees, tries). HOPE includes six entropy encoding schemes that trade between compression rate and encoding performance. When the DBMS creates a tree-based index/filter, HOPE samples the initial bulk-loaded keys and counts the frequencies of the byte patterns specified by a scheme. It uses these statistics to generate dictionary symbols that comply with our theoretical model to preserve key ordering. HOPE then encodes the symbols using either fixed-length codes or optimal order-preserving prefix codes. A key insight in HOPE is its emphasis on encoding speed (rather than decoding) because our target search tree queries need not reconstruct the original keys.

To evaluate HOPE, we applied it to five in-memory search trees: SuRF [51], ART [33], HOT [17], B+tree [10], and Prefix B+tree [15]. Our experimental results show that HOPE reduces their query latency by up to 40% and saves their memory consumption by up to 30% at the same time for most string key workloads.

We make four primary contributions in this paper. First, we develop a theoretical model to characterize the properties of dictionary encoding. Second, we introduce HOPE to compress keys for in-memory search trees efficiently. Third, we implement six compression schemes in HOPE to study the compression rate vs. encoding speed trade-off. Finally, we apply HOPE on five trees and show that HOPE improves their performance and memory-efficiency simultaneously.

2 BACKGROUND AND RELATED WORK

Modern DBMSs rely on in-memory search trees (e.g., indexes and filters) to achieve high throughput and low latency. We divide these search trees into three categories. The first is the B-tree/B+tree family, including Cache Sensitive B+trees (CSB+trees) [45] and Bw-Trees [35, 47]. They store keys horizontally side-by-side in the leaf nodes and have good range query performance. The second category includes tries and radix trees [14, 17, 19, 23, 25, 29, 33, 41]. They store keys vertically to allow prefix compression. Recent memory-efficient tries such as ART [33] and HOT [17] are faster than B+trees on modern hardware. The last category is hybrid data structures such as Masstree [38] that combine the B+tree and trie designs in a single data structure.

Existing compression techniques for search trees leverage general-purpose block compression algorithms such as LZ77 [4], Snappy [9], and LZ4 [7]. For example, InnoDB uses the zlib library [3] to compress its B+tree pages/nodes before they are written to disk. Block compression algorithms, however, are too slow for in-memory search trees: query latencies for in-memory B+trees and tries range from 100s of nanoseconds to a few microseconds, while the fastest block compression algorithms can decompress only a few 4 KB memory pages in that time [7].

Recent work has addressed this size problem through new data structures [17, 40, 50, 51]. For example, the succinct trie in SuRF consumes only 10 bits (close to the theoretical optimum) to represent a node [51]. Compressing input keys using HOPE, however, is an orthogonal approach that one can apply to any of the above search tree categories to achieve additional space savings and performance gains.

One could apply existing field/table-wise compression schemes to search tree keys. Whole-key dictionary compression is the most popular scheme used in DBMSs today. It replaces the values in a column with smaller fixed-length codes using a dictionary. Indexes and filters, therefore, could take advantage of those existing dictionaries for key compression. There are several problems with this approach. First, the dictionary compression must be order-preserving to allow range queries on search trees. Order-preserving dictionaries, however, are difficult to maintain with changing value domains [37], which is often the case for string keys in OLTP applications. Second, the latency of encoding a key is similar to that of querying the actual indexes/filters because most order-preserving dictionaries use the same kind of search trees themselves [18]. Finally, dictionary compression only works well for columns with low/moderate cardinalities. If most of the values are unique, then the larger dictionary negates the size reduction in the actual fields.

Existing order-preserving frequency-based compression schemes, including the one used in DB2 BLU [43] and padded encoding [36], exploit the column value distribution skew by assigning smaller codes to more frequent values. Variable-length codes, however, are inefficient to locate, decode, and process in parallel. DB2 BLU, thus, only uses up to a few different code sizes per column and stores the codes of the same size together to speed up queries. Padded encoding, on

the other hand, pads the variable-length codes with zeros at the end so that all codes are of the same length (i.e., the maximum length of the variable-length codes) to facilitate scan queries. DB2 BLU and padded encoding are designed for column stores where most queries are reads, and updates are often in batches. Both designs still use the whole-key dictionary compression discussed above and therefore, cannot encode new values without extending the dictionary, which can cause expensive re-encodes of the column. HOPE, however, can encode arbitrary input values using the same dictionary while preserving their ordering. Such property is desirable for write-intensive OLTP indexes.

The focus of this paper is on compressing string keys. Numeric keys are already small and can be further compressed using techniques, such as null suppression and delta encoding [11]. Prefix compression and suffix truncation are common techniques used in B+trees. HOPE can provide additional benefits on top of these compression methods.

Prior studies considered entropy encoding schemes such as Huffman [27] and arithmetic coding [48] too slow for columnar data compression because their variable-length codes are slow to decode [11, 16, 18, 20, 37, 44]. This concern does not apply to search trees because non-covering index and filter queries do not reconstruct the original keys[1]. In addition, entropy encoding schemes produce high compression rates even with small dictionaries because they exploit common patterns at a fine granularity.

Antoshenkov et al. [12, 13] proposed an order-preserving string compressor with a string parsing algorithm (ALM) to guarantee the order of the encoded results. We introduce our string axis model in the next section, which is inspired by the ALM method but is more general. The ALM compressor belongs to a specific category in our compression model.

3 COMPRESSION MODEL

Different dictionary encoding schemes, ranging from Huffman encoding [27] to the ALM-based compressor [12], provide different capabilities and guarantees. For example, some can encode arbitrary input strings while others preserve order. In this section, we introduce a unified model, called the string axis model, to characterize the properties of a dictionary encoding scheme. This model is inspired by the ALM string parsing algorithm [13], which solves the order-preserving problem for dictionary-based string compression. Using the string axis model, we can construct a wide range of dictionary-based compression schemes that can serve our target application (i.e., key compression for in-memory search trees). We divide qualified schemes into four categories, each making different trade-offs. We then briefly describe six representative compression schemes supported by HOPE.

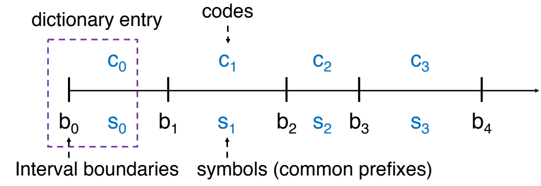

Figure 1: String Axis Model – All strings are divided into connected intervals in lexicographical order. Source strings in the same interval share a common prefix (s_i) that maps to code (c_i).

3.1 The String Axis Model

As shown in Figure 1, a *string axis* lays out all possible source strings on a single axis in lexicographical order. We can represent a dictionary encoding scheme using this model and highlight three important properties: (1) completeness, (2) unique decodability, and (3) order-preserving.

Let Σ denote the source string alphabet. Σ^* is the set of all possible finite-length strings over Σ. Similarly, let X denote the code alphabet and X^* be the code space. Typically, Σ is the set of all characters, and $X = \{0, 1\}$. A dictionary $D : S \to C, S \in \Sigma^*, C \in X^*$ maps a subset of the source strings S to the set of codes C.

On the string axis, a dictionary entry $s_i \to c_i$ is mapped to an interval I_i, where s_i is a prefix of all strings within I_i. The choice of I_i is not unique. For example, as shown in Figure 2, both [abcd, abcf) and [abcgh, abcpq) are valid mappings for dictionary entry abc→0110. In fact, any sub-interval of [abc, abd) is a valid mapping in this example. If a source string *src* falls into the interval I_i, then a dictionary lookup on *src* returns the corresponding entry $s_i \to c_i$.

We can model the dictionary encoding method as a recursive process. Given a source string *src*, one can lookup *src* in the dictionary and obtain an entry $(s \to c) \in D, s \in S, c \in C$, such that s is a prefix of *src*, i.e., $src = s \cdot src_{suffix}$, where "$\cdot$" is the concatenation operation. We then replace s with c in *src* and repeat the process[2] using src_{suffix}.

To guarantee that encoding always makes progress, we must ensure that every dictionary lookup is successful. This means that for any *src*, there must exist a dictionary entry $s \to c$ such that $len(s) > 0$ and s is a prefix of *src*. In other words, we must consume some prefix from the source string at every lookup. We call this property **dictionary completeness**. Existing dictionary compression schemes for DBMSs are usually not complete because they only assign codes to the string values already seen by the DBMS. These schemes

[1]The search tree can recover the original keys if needed: entropy encoding is lossless. Exploring lossy compression on search trees is future work

[2]One can use a different dictionary at every step. For performance reasons, we consider a single dictionary throughout the process in this paper.

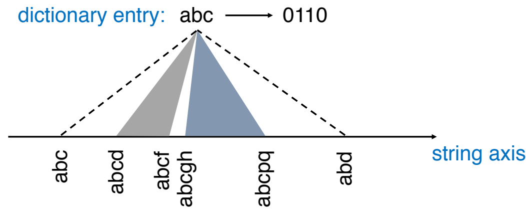

Figure 2: Dictionary Entry Example – All sub-intervals of [abc, abd) are valid mappings for dictionary entry abc ⟶ 0110.

cannot encode arbitrary strings unless they grow the dictionary, but growing to accommodate new entries may require the DBMS to re-encode the entire corpus [18]. In the string axis model, a dictionary is complete if and only if the union of all the intervals (i.e., $\bigcup I_i$) covers the entire string axis.

A dictionary encoding $Enc : \Sigma^* \rightarrow X^*$ is **uniquely decodable** if Enc is an *injection* (i.e., there is a one-to-one mapping from every element of Σ^* to an element in X^*). To guarantee unique decodability, we must ensure that (1) there is only one way to encode a source string and (2) every encoded result is unique. Under our string axis model, these requirements are equivalent to (1) all intervals I_i's are disjoint and (2) the set of codes C used in the dictionary are uniquely decodable (we only consider prefix codes in this paper).

With these requirements, we can use the string axis model to construct a dictionary that is both complete and uniquely decodable. As shown in Figure 1, for a given dictionary size of n entries, we first divide the string axis into n consecutive intervals $I_0, I_1, \ldots, I_{n-1}$, where the max-length common prefix s_i of all strings in I_i is not empty (i.e., $len(s_i) > 0$) for each interval. We use $b_0, b_1, \ldots, b_{n-1}, b_n$ to denote interval boundaries. That is, $I_i = [b_i, b_{i+1})$ for $i = 0, 1, \ldots, n-1$. We then assign a set of uniquely decodable codes $c_0, c_1, \ldots, c_{n-1}$ to the intervals. Our dictionary is thus $s_i \rightarrow c_i, i = 0, 1, \ldots, n-1$. A dictionary lookup maps the source string to a single interval I_i, where $b_i < src < b_{i+1}$.

We can achieve the **order-preserving** property on top of unique decodability by assigning monotonically increasing codes $c_0 < c_1 < \ldots < c_{n-1}$ to the intervals. This is easy to prove. Suppose there are two source strings (src_1, src_2), where $src_1 < src_2$. If src_1 and src_2 belong to the same interval I_i in the dictionary, they must share common prefix s_i. Replacing s_i with c_i in each string does not affect their relative ordering. If src_1 and src_2 map to different intervals I_i and I_j, then $Enc(src_1) = c_i \cdot Enc(src_{1_{suffix}})$, $Enc(src_2) = c_j \cdot Enc(src_{2_{suffix}})$. Since $src_1 < src_2$, I_i must precede I_j on the string axis. That means $c_i < c_j$. Because c_i's are prefix codes, $c_i \cdot Enc(src_{1_{suffix}}) < c_j \cdot Enc(src_{2_{suffix}})$.

For encoding search tree keys, we prefer schemes that are complete and order-preserving; unique decodability is implied by the latter property. Completeness allows the scheme to encode arbitrary keys, while order-preserving guarantees that the search tree supports meaningful range queries on the encoded keys.

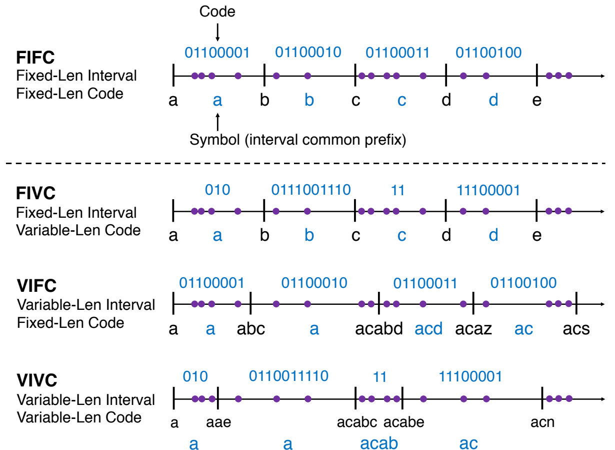

Figure 3: Compression Models – Four categories of complete and order-preserving dictionary encoding schemes.

3.2 Exploiting Entropy

For a dictionary encoding scheme to reduce the size of the corpus, its emitted codes must be shorter than the source strings. Given a complete, order-preserving dictionary $D : s_i \rightarrow c_i, i = 0, 1, \ldots, n-1$, let p_i denote the probability that a dictionary entry is accessed at each step during the encoding of an arbitrary source string. Because the dictionary is complete and uniquely decodable (implied by order-preserving), $\sum_{i=0}^{n-1} p_i = 1$. The encoding scheme achieves the best compression when the compression rate $CPR = \sum_{i=0}^{n-1} len(s_i)p_i \,/\, \sum_{i=0}^{n-1} len(c_i)p_i$ is maximized.

According to the string axis model, we can characterize a dictionary encoding scheme in two aspects: (1) how to divide intervals and (2) what code to assign to each interval. Interval division determines the symbol lengths ($len(s_i)$) and the access probability distribution (p_i) in a dictionary. Code assignment exploits the entropy in p_i's by using shorter codes (c_i) for more frequently-accessed intervals.

We consider two interval-division strategies: fixed-length intervals and variable-length intervals. For code assignment, we consider two types of prefix codes: fixed-length codes and optimal variable-length codes. We, therefore, divide all complete and order-preserving dictionary encoding schemes into four categories, as shown in Figure 3.

Fixed-length Interval, Fixed-length Code (FIFC): This is the baseline scheme because ASCII encodes characters in this way. We do not consider this category for compression.

Fixed-length Interval, Variable-length Code (FIVC): This category is the classic Hu-Tucker encoding [26]. If order-preserving is not required, both Huffman encoding [27] and

arithmetic encoding [48] also belong to this category[3]. Although intervals have a fixed length, access probabilities are not evenly distributed among the intervals. Using optimal (prefix) codes, thus, maximizes the compression rate.

Variable-length Interval, Fixed-length Code (VIFC): This category is represented by the ALM string compression algorithm proposed by Antoshenkov [12]. Because the code lengths are fixed (i.e., $len(c_i) = L$), the compression rate $CPR = \frac{1}{L}\sum_{i=0}^{n-1} len(s_i)p_i$. ALM applied the "equalizing" heuristic of letting $len(s_0)p_0 = len(s_1)p_1 = \cdots = len(s_n)p_n$ to maximize CPR. We note that the example in Figure 3 has two intervals with the same dictionary symbol. This is allowed because only one of the intervals will contain a specific source string, Also, by using variable-length intervals, we no longer have the "concatenation property" for the encoded results (e.g., $Code(ab) \neq Code(a) \cdot Code(b)$). This property, however, is not a requirement for our target application.

Variable-length Interval, Variable-length Code (VIVC): To the best of our knowledge, this category is unexplored by previous work. Although Antoshenkov suggests that ALM could benefit from a supplementary variable-length code [12], it is neither implemented nor evaluated. VIVC has the most flexibility in building dictionaries (one can view FIFC, FIVC, and VIFC as special cases of VIVC), and it can potentially lead to an optimal compression rate. We describe the VIVC schemes in HOPE in Section 3.3.

Although VIVC schemes can have higher compression rates than the other schemes, both fixed-length intervals and fixed-length codes have performance advantages over their variable-length counterparts. Fixed-length intervals create smaller and faster dictionary structures, while fixed-length codes are more efficient to decode. Our objective is to find the best trade-off between compression rate and encoding performance for in-memory search tree keys.

3.3 Compression Schemes

Based on the above dictionary encoding models, we next introduce six compression schemes implemented in HOPE, as shown in Figure 4, We select these schemes from the three viable categories (FIVC, VIFC, and VIVC). Each scheme makes different trade-offs between compression rate and encoding performance. We first describe them at a high level and then provide their implementation details in Section 4.

Single-Char is the FIVC compression algorithm used in Huffman encoding and arithmetic encoding. The fixed-length intervals have consecutive single characters as the boundaries (e.g., [a, b), [b, c)). The dictionary symbols are 8-bit ASCII characters, and the dictionary has a fixed 256 entries. The codes assigned to the symbols are Hu-Tucker codes. Hu-Tucker codes are optimal order-preserving prefix codes Figure 4a shows an example dictionary segment.

[3]Arithmetic encoding does not operate the same way as a typical dictionary encoder. But its underlying principle matches this category.

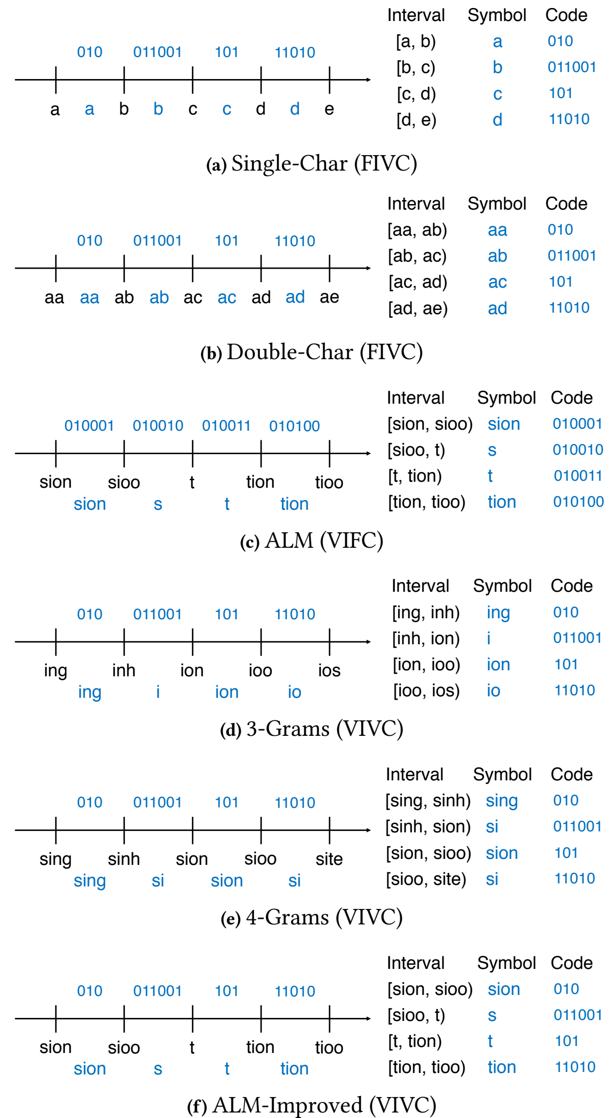

Figure 4: Compression Schemes – Example dictionary segments.

Double-Char is a FIVC compression algorithm that is similar to Double-Char, except that the interval boundaries are consecutive double characters (e.g., [aa, ab), [ab, ac)). To make the dictionary complete, we introduce a terminator character ∅ before all ASCII characters to fill the interval gaps between [a'\255', b) and [b'\0', b'\1'), for example, with interval [b∅, b'\0'). Figure 4b shows an example dictionary. This scheme should achieve better compression than Single-Char because it exploits the first-order entropy of the source strings instead of the zeroth-order entropy.

ALM is a state-of-the-art VIFC string compression algorithm. To determine the interval boundaries from a set of sample source strings (e.g., initial keys for an index), ALM first selects substring patterns that are long and frequent. Specifically, for a substring pattern s, it computes $len(s) \times freq(s)$, where $freq(s)$ represents the number of occurrence of s in the sample set. ALM includes s in its dictionary if the product is greater than a threshold W. It then creates one or more intervals between the adjacent selected symbols. The goal of the algorithm is to make the above product (i.e., length of common prefix × access frequency) for each interval as equal as possible. The detailed algorithm is described in [12].

ALM uses monotonically increasing fixed-length codes. Figure 4c shows an example dictionary segment. The dictionary size for ALM depends on the threshold W. One must binary search on W's to obtain a desired dictionary size.

3-Grams is a VIVC compression algorithm where the interval boundaries are 3-character strings. Given a set of sample source strings and a dictionary size limit n, the scheme first selects the top $n/2$ most frequent 3-character patterns and adds them to the dictionary. For each interval gap between the selected 3-character patterns, 3-Grams creates a dictionary entry to cover the gap. For example, in Figure 4d, "`ing`" and "`ion`" are selected frequent patterns from the first step. "`ing`" and "`ion`" represent intervals [`ing`, `inh`) and [`ion`, `ioo`) on the string axis. Their gap interval [`inh`, `ion`) is also included as a dictionary entry. 3-Grams uses Hu-Tucker codes.

4-Grams is a VIVC compression algorithm similar to 3-Grams with 4-character string boundaries. Figure 4e shows an example. Compared to 3-Grams, 4-Grams exploits higher-order entropy; but whether it provides a better compression rate over 3-Grams depends on the dictionary size.

ALM-Improved improves the ALM scheme in two ways. First, as shown in Figure 4f, we replace the fixed-length codes in ALM with Hu-Tucker codes because we observe access skew among the intervals despite ALM's "equalizing" algorithm. Second, the original ALM counts the frequency for every substring (of any length) in the sample set, which is slow and memory-consuming. In ALM-Improved, we simplify this by only collecting statistics for substrings that are suffixes of the sample source strings. Our evaluation in Section 6 shows that using Hu-Tucker codes improves ALM's compression rate while counting the frequencies of string suffixes reduces ALM's build time without compromising the compression rate.

4 HOPE

We now present the design and implementation of HOPE. There are two goals in HOPE's architecture. First, HOPE must minimize its performance overhead so that it does not negate the benefits of storing shorter keys. Second, HOPE must be extensible. From our discussion in Section 3, there are many choices in constructing an order-preserving dictionary encoding scheme. Although we support six representative schemes in the current version of HOPE, one could, for example, devise better heuristics in generating dictionary entries to achieve a higher compression rate, or invent more efficient dictionary data structures to further reduce encoding latency. HOPE can be easily extended to include such improvements through its modularized design.

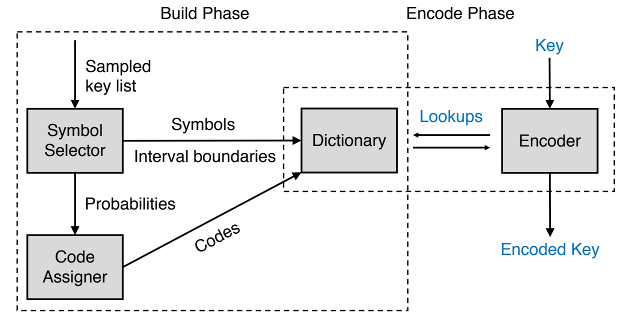

Figure 5: The HOPE Framework – An overview of HOPE's modules and their interactions with each other in the two phases.

4.1 Overview

As shown in Figure 5, HOPE executes in two phases (i.e., Build, Encode) and has four modules: **Symbol Selector**, **Code Assigner**, **Dictionary**, and **Encoder**. A DBMS provides HOPE with a list of sample keys from the search tree. HOPE then produces a Dictionary and an Encoder as its output. We note that the size and representativeness of the sampled key list only affect the compression rate. The correctness of HOPE's compression algorithm is guaranteed by the dictionary completeness and order-preserving properties discussed in Section 3.1. In other words, any HOPE dictionary can both encode arbitrary input keys and preserve the original key ordering.

In the first step of the build phase, the Symbol Selector counts the frequencies of the specified string patterns in the sampled key list and then divides the string axis into intervals based on the heuristics given by the target compression scheme. The Symbol Selector generates three outputs for each interval: (1) dictionary symbol (i.e., the common prefix of the interval), (2) interval boundaries, and (3) probability that a source string falls within that interval.

The framework then gives the symbols and interval boundaries to the Dictionary module. Meanwhile, it sends the probabilities to the Code Assigner to generate codes for the dictionary symbols. If the scheme uses fixed-length codes, the Code Assigner only considers the dictionary size. If the scheme uses variable-length codes, the Code Assigner examines the

<table>
<tr><th>Scheme</th><th>Symbol Selector</th><th>Code Assigner</th><th>Dictionary</th><th>Encoder</th></tr>
<tr><td>Single-Char</td><td>Single-Char</td><td rowspan="2">Hu-Tucker</td><td rowspan="2">Array</td><td rowspan="6">Fast Encoder</td></tr>
<tr><td>Double-Char</td><td>Double-Char</td></tr>
<tr><td>ALM</td><td>ALM</td><td>Fixed-Length</td><td>ART-Based</td></tr>
<tr><td>3-Grams</td><td>3-Grams</td><td rowspan="3">Hu-Tucker</td><td rowspan="2">Bitmap-Trie</td></tr>
<tr><td>4-Grams</td><td>4-Grams</td></tr>
<tr><td>ALM-Improved</td><td>ALM-Improved</td><td>ART-Based</td></tr>
</table>

Table 1: Module Implementations – The configuration of HOPE's six compression schemes.

probability distribution to generate optimal order-preserving prefix codes (i.e., Hu-Tucker codes).

When the Dictionary module receives the symbols, the interval boundaries, and the codes, it selects an appropriate and fast dictionary data structure to store the mappings. The string lengths of the interval boundaries inform the decision; available data structures range from fixed-length arrays to general-purpose tries. The dictionary size is a tunable parameter for VIFC and VIVC schemes. Using a larger dictionary trades performance for a better compression rate.

The encode phase uses only the Dictionary and Encoder modules. On receiving an uncompressed key, the Encoder performs multiple lookups in the dictionary. Each lookup translates a part of the original key to some code as described in Section 3.1. The Encoder then concatenates the codes in order and outputs the result. This encoding process is sequential for variable-length interval schemes (i.e., VIFC and VIVC) because the remaining source string to be encoded depends on the results of earlier dictionary lookups.

We next describe the implementation details for each module. Building a decoder is optional because our target workload for search trees does not require reconstructing the original keys.

4.2 Implementation

HOPE users can create new compression schemes by combining different module implementations. HOPE currently supports the six compression schemes described in Section 3.3. For Symbol Selector and Code Assigner, the goal is to generate a dictionary that leads to the maximum compression rate. We no longer need these two modules after the build phase. We spend extra effort optimizing the Dictionary and Encoder modules because they are on the critical path of every search tree query.

Symbol Selector: It first counts the occurrences of substring patterns in the sampled keys using a hash table. For example, 3-Grams considers all three-character substrings, while the ALM examines substrings of all lengths. For Single-Char and Double-Char, the interval boundaries are implied because they are fixed-length-interval schemes. For the remaining schemes, the Symbol Selectors divide intervals using the algorithms described in Section 3.3: first identify the most frequent symbols and then fill the gaps with new intervals.

The ALM and ALM-Improved Symbol Selectors require an extra *blending* step before their interval-division algorithms. This is because the selected variable-length substrings may not satisfy the *prefix property* (i.e., a substring can be a prefix of another substring). For example, "`sig`" and "`sigmod`" may both appear in the frequency list, but the interval-division algorithm cannot select both of them because the two intervals on the string axis are not disjoint: "`sigmod`" is a sub-interval of "`sig`". A blending algorithm redistributes the occurrence count of a prefix symbol to its longest extension in the frequency list [12]. We implement this blending algorithm in HOPE using a trie data structure.

After the Symbol Selector decides the intervals, it performs a test encoding of the sample keys using the intervals as if the code for each interval has been assigned. The purpose of this step is to obtain the probability that a source string (or its remaining suffix) falls into each interval so that the Code Assigner can generate codes based on those probabilities to maximize compression. For variable-length-interval schemes, the probabilities are weighted by the symbol lengths of the intervals.

Code Assigner: Assume that the Code Assigner receives N probabilities. To assign fixed-length codes to them, the Code Assigner outputs monotonically increasing integers $0, 1, 2, \cdots, N-1$, each using $\lceil \log_2 N \rceil$ bits. For variable-length codes, HOPE uses the Hu-Tucker algorithm to generate optimal order-preserving prefix codes. One could use an alternative method, such as *Range Encoding* [39] (i.e., the integer version of Arithmetic Encoding). Range Encoding, however, requires more bits than Hu-Tucker to ensure that codes are exactly on range boundaries to guarantee order-preserving.

The Hu-Tucker algorithm works in four steps [1]. First, it creates a leaf node for each probability and then lists the leaf nodes in interval order. Second, it recursively merges the two least-frequent nodes where there are no leaf nodes between them (unlike Huffman). This is where the algorithm guarantees order. After constructing this probability tree, it next computes the depth of each leaf node to derive the lengths of the codes. Finally, the algorithm constructs a tree by adding these leaf nodes level-by-level starting from the deepest and then connecting adjacent nodes at the same level in pairs. HOPE uses this Huffman-tree-like structure to extract the final codes. Our Hu-Tucker implementation in the Code Assigner uses an improved algorithm [49] that runs in $O(N^2)$ time instead of $O(N^3)$ as in the original paper [26].

Dictionary: A dictionary in HOPE maps an interval (and its symbol) to a code. Because the intervals are connected and disjoint, the dictionary needs to store only the left boundary of each interval as the key. A key lookup in the dictionary

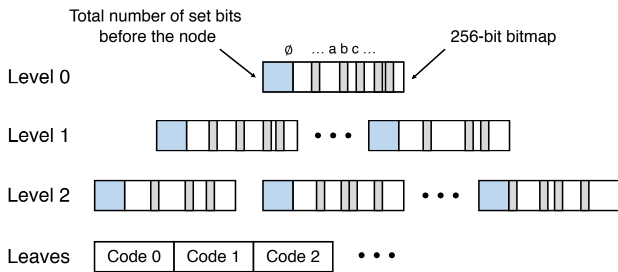

Figure 6: 3-Grams Bitmap-Trie Dictionary – Each node consists of a 256-bit bitmap and a counter. The former records the branches of the node and the latter represents the total number of set bits in the bitmaps of all the preceding nodes.

then is a "greater than or equal to" index query. For the values, we store only the codes along with the lengths of the symbols to determine the number of characters from the source string that we have consumed at each step.

We implemented three dictionary data structures in HOPE. The first is an *array* for the Single-Char and Double-Char schemes. Each dictionary entry includes an 8-bit integer to record the code length and a 32-bit integer to store the code. The dictionary symbols and the interval left boundaries are implied by the array offsets. For example, the 97th entry in Single-Char has the symbol a, while the 24770th entry in Double-Char corresponds to the symbol aa[4]. A lookup in an array-based dictionary is fast because it requires only a single memory access and the array fits in CPU cache.

The second dictionary data structure in HOPE is a *bitmap-trie* used by the 3-Grams and 4-Grams schemes. Figure 6 depicts the structure of a three-level bitmap-trie for 3-Grams. The nodes are stored in an array in breadth-first order. Each node consists of a 32-bit integer and a 256-bit bitmap. The bitmap records all the branches of the node. For example, if the node has a branch labeled a, the 97th bit in the bitmap is set. The integer at the front stores the total number of set bits in the bitmaps of all the preceding nodes. Since the stored interval boundaries can be shorter than three characters, the data structure borrows the most significant bit from the 32-bit integer to denote the termination character ∅.

Given a node $(n, bitmap)$ where n is the count of the preceding set bits, its child node pointed by label l at position $n + popcount(bitmap, l)$[5] in the node array. Our evaluation shows that looking up a bitmap-trie is 2.3× faster than binary-searching the dictionary entries because it requires fewer memory probes and has better cache performance.

Finally, we use an ART-based dictionary to serve the ALM and ALM-Improved schemes. ART is a radix tree that supports variable-length keys [33]. We modified three aspects

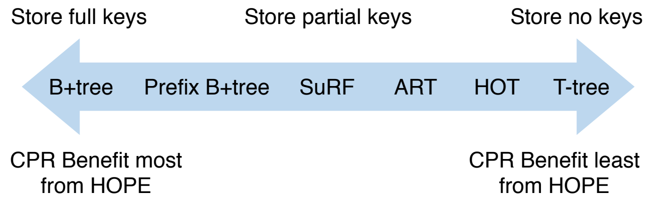

Figure 7: Search Tree on Key Storage – Search trees get decreasing benefits from HOPE, especially in terms of compression rate (CPR), as the completeness of key storage goes down.

of ART to make it more suitable as a dictionary. First, we added support for *prefix keys* in ART. This is necessary because both abc and abcd, for example, can be valid interval boundaries stored in a dictionary. We also disabled ART's *optimistic common prefix skipping* that compresses paths on single-branch nodes by storing only the first few bytes. If a corresponding segment of a query key matches the stored bytes during a lookup, ART assumes that the key segment also matches the rest of the common prefix (a final key verification happens against the full tuple). HOPE's ART-based dictionary, however, stores the full common prefix for each node since it cannot assume that there is a tuple with the original key. Lastly, we modified the ART's leaf nodes to store the dictionary entries instead of tuple pointers.

Encoder: HOPE looks up the source string in the dictionary to find an interval that contains the string. The dictionary returns the symbol length L and the code C. HOPE then concatenates C to the result buffer and removes the prefix of length L that matches the symbol from the source string. It repeats this process on the remaining string until it is empty.

To make the non-byte-aligned code concatenation fast, HOPE stores codes in 64-bit integer buffers. It adds a new code to the result buffer in three steps: (1) left-shift the result buffer to make room for the new code; (2) write the new code to the buffer using a bit-wise OR instruction; (3) split the new code if it spans two 64-bit integers. This procedure costs only a few CPU cycles per code concatenation.

When encoding a batch of sorted keys, the Encoder optimizes the algorithm by first dividing the batch into blocks, where each block contains a fixed number of keys. The Encoder then encodes the common prefix of the keys within a block only once, avoiding redundant work. When the batch size is two, we call this pair-encoding. Compared to encoding keys individually, pair-encoding reduces key compression overhead for the closed-range queries in a search tree. We evaluate batch encoding in [52] (Appx. B).

5 INTEGRATION

Integrating HOPE in a DBMS is a straightforward process because we designed it to be a standalone library that is independent of the target search tree data structure.

[4] $24770 = 96 \times (256 + 1) + 97 + 1$. The +1's are for the terminators ∅

[5] The POPCOUNT CPU instruction counts the set bits in a bit-vector. The function $popcount(bitmap, l)$ counts the set bits up to position l in $bitmap$.

When the DBMS creates a new search tree, HOPE samples the initial bulk-loaded keys and constructs the dictionary (i.e., the build phase). After that, every query for that tree, including the initial bulk-inserts, must go through the Encoder first to compress the keys. If the search tree is initially empty, HOPE samples keys as the DBMS inserts them into the tree. It then rebuilds the search tree using the compressed keys once it sees enough samples. We use a small sample size because it guarantees fast tree rebuild, and it does not compromise the compression rate (see [52] (Appx. A)).

We typically invoke the HOPE framework's Build Phase only once because switching dictionaries causes the search tree to rebuild, which is particularly expensive for large trees. Our assumption is that the value distribution in a database column is relatively stable, especially at the substring level. For example, "`@gmail.com`" is likely a common pattern for emails. Because HOPE exploits common patterns at relatively fine granularity, its dictionary remains effective in compressing keys over time. We evaluated HOPE under a dramatic key distribution change [52] (Appx. C) and observed a compression rate decreases as expected, with simpler schemes such as Single-Char less affected. We note that even if a dramatic change in the key distribution happens, HOPE is not required to rebuild immediately because it still guarantees query correctness. The system can schedule the reconstruction during maintenance to recover the compression rate.

We applied HOPE to five in-memory search trees:

- **SuRF**: The Succinct Range Filter [51] is a trie-based data structure that performs approximate membership tests for ranges. SuRF uses succinct data structures to achieve an extremely small memory footprint.
- **ART**: The Adaptive Radix Tree [33, 34] is the default index structure for HyPer [28]. ART adaptively selects variable-sized node layouts based on fanouts to save space and to improve cache performance.
- **HOT**: The Height Optimized Trie [17] is a fast and memory-efficient index structure. HOT guarantees high node fanouts by combining nodes across trie levels.
- **B+tree**: We use the cache-optimized TLX B+tree [10] (formerly known as STX). TLX B+tree stores variable-length strings outside the node using reference pointers. The default node size is 256 bytes, making a fanout of 16 (8-byte key pointer and 8-byte value pointer per slot).
- **Prefix B+tree**: A Prefix B+tree [15] optimizes a plain B+tree by applying prefix and suffix truncation to the nodes [24]. A B+tree node with prefix truncation stores the common prefix of its keys only once. During a leaf node split, suffix truncation allows the parent node to choose the shortest string qualified as a separator key. We implemented both techniques on a state-of-the-art B+tree [32] other than TLX B+tree for better experimental robustness.

HOPE provides the most benefit to search trees that store the full keys. Many tree indexes for in-memory DBMSs, such as ART and HOT, only store partial keys to help the DBMS find the record IDs. They then verify the results against the full keys after fetching the in-memory records. To understand HOPE's interaction with these different search trees, we arrange them in Figure 7 according to how large a part of the keys they store. The B+tree is at one extreme where the data structure stores full keys. At the other extreme sits the T-Tree [31] (or simply a sorted list of record IDs) where no keys appear in the data structure. Prefix B+tree, SuRF, ART, and HOT fall in the middle. HOPE is more effective towards the B+tree side, especially in terms of compression rate. The query latency improvement is the difference between the speedup due to shorter keys and the overhead of key compression. For the B+tree family, shorter keys means larger fanouts and faster string comparisons. Although tries only store partial keys, HOPE improves their performance by reducing the tree height. We next analyze the latency reduction of using HOPE on a trie.

Let l denote the average key length and cpr denote the compression rate (i.e., uncompressed length / compressed length). The average height of the original *trie* is h. We use t_{trie} to denote the time needed to walk one level (i.e., one character) down the *trie*, and t_{encode} to denote the time needed to compress one character in HOPE.

The average point query latency in the original *trie* is $h \times t_{trie}$, while this latency in the compressed *trie* is $l \times t_{encode} + \frac{h}{cpr} \times t_{trie}$, where $l \times t_{encode}$ represents the encoding overhead. Therefore, the percentage of latency reduction is:

$$\frac{h \times t_{trie} - (l \times t_{encode} + \frac{h}{cpr} \times t_{trie})}{h \times t_{trie}} = 1 - \frac{1}{cpr} - \frac{l \times t_{encode}}{h \times t_{trie}}$$

If the expression > 0, we improve performance. For example, when evaluating SuRF on the email workload in Section 7, $l = 21.2$ bytes, and $h = 18.2$ levels. The original SuRF has an average point query latency of $1.46\mu s$. t_{trie} is, thus, $\frac{1.46\mu s}{18.2} = 80.2ns$. Our evaluation in Section 6 shows that HOPE's Double-Char scheme achieves $cpr = 1.94$ and $t_{encode} = 6.9ns$ per character. Hence, we estimate that by using Double-Char on SuRF, we can reduce the point query latency by $1 - \frac{1}{1.94} - \frac{21.2 \times 6.9}{18.2 \times 80.2} = 38\%$. The real latency reduction is usually higher (41% in this case as shown in Section 7) because smaller tries also improve cache performance.

6 HOPE MICROBENCHMARKS

We evaluate HOPE in the next two sections. We first analyze the trade-offs between compression rate and compression overhead of different schemes in HOPE. These microbenchmarks help explain the end-to-end measurements on HOPE-integrated search trees in Section 7.

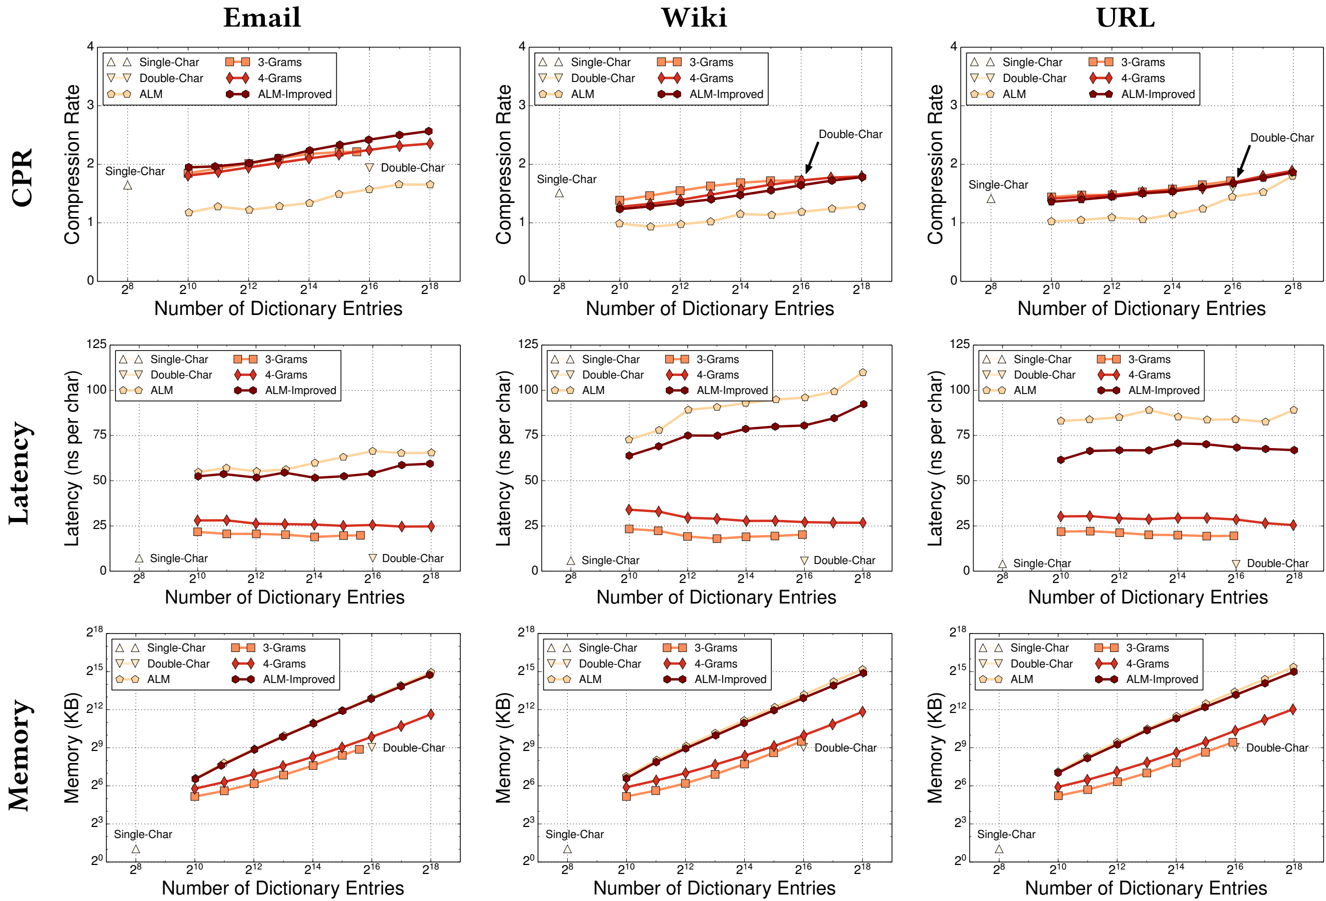

Figure 8: Compression Microbenchmarks – Measurements of HOPE's six schemes on the different datasets.

We use the following datasets for all our experiments:

- **Email:** 25 million email addresses (host reversed – e.g., "com.gmail@foo") with an average length of 22 bytes.
- **Wiki:** 14 million article titles from the English version of Wikipedia with an average length of 21 bytes [6].
- **URL:** 25 million URLs from a 2007 web crawl with an average length of 104 bytes [2].

We run our experiments using a machine equipped with two Intel® Xeon® E5-2630v4 CPUs (2.20GHz, 32 KB L1, 256 KB L2, 25.6 MB L3) and 8×16 GB DDR4 RAM.

In each experiment, we randomly shuffle the target dataset before each trial. We then select 1% of the entries from the shuffled dataset as the sampled keys for HOPE. Our sensitivity test in [52] (Appx. A) shows that 1% is large enough for all schemes to reach their maximum compression rates. We repeat each trial three times and report the average result.

6.1 Performance & Efficacy

We first evaluate the runtime performance and compression efficacy of HOPE's six built-in schemes listed in Table 1. HOPE compresses the keys one-at-a-time with a single thread. We vary the number of dictionary entries in each trial and measure three facets per scheme: (1) the compression rate, (2) the average encoding latency per character, and (3) the size of the dictionary. We compute the compression rate as the uncompressed dataset size divided by the compressed dataset size. We obtain the average encoding latency per character by dividing the execution time by the total number of bytes in the uncompressed dataset.

Figure 8 shows the experimental results. We vary the number of dictionary entries on the x-axis (log scaled). The Single-Char and Double-Char schemes have fixed dictionary sizes of 2^8 and 2^{16}, respectively. The 3-Grams dictionary cannot grow to 2^{18} because there are not enough unique three-character patterns in the sampled keys.

Compression Rate: The first row of Figure 8 shows that the VIVC schemes (3-Grams, 4-Grams, ALM-Improved) have better compression rates than the others. This is because VIVC schemes exploit the source strings' higher-order entropies to optimize both interval division and code assignment at the same time. In particular, ALM-Improved compresses the keys more than the original ALM because it uses a better pattern

extraction algorithm, and it uses the Hu-Tucker codes that leverage the remaining skew in the dictionary entries' access probabilities. ALM tries to equalize these weighted probabilities but our improved version has better efficacy. We also note that a larger dictionary produces a better compression rate for the variable-length interval schemes.

Encoding Latency: The latency results in the second row of Figure 8 demonstrate that the simpler schemes have lower encoding latency. This is expected because the latency depends largely on the dictionary data structures. Single-Char and Double-Char are the fastest because they use array dictionaries that are small enough to fit in the CPU's L2 cache. Our specialized bitmap-tries used by 3-Grams and 4-Grams are faster than the general ART-based dictionaries used by ALM and ALM-Improved because (1) the bitmap speeds up in-node label search; and (2) the succinct design (without pointers) improves cache performance.

The figures also show that latency is stable (and even decrease slightly) in all workloads for 3-Grams and 4-Grams as their dictionary sizes increase. This is interesting because the cost of performing a lookup in the dictionary increases as the dictionary grows in size. The larger dictionaries, however, achieve higher compression rates such that it reduces lookups: larger dictionaries have shorter intervals on the string axis, and shorter intervals usually have longer common prefixes (i.e., dictionary symbols). Thus, HOPE consumes more bytes from the source string at each lookup with larger dictionaries, counteracting the higher per-lookup cost.

Dictionary Memory: The third row of Figure 8 shows that the dictionary sizes for the variable-length schemes grow linearly as the number of dictionary entries increases. Even so, for most dictionaries, the total tree plus dictionary size is still much smaller than the size of the corresponding uncompressed search tree. These measurements also show that our bitmap-tries for 3-Grams and 4-Grams are up to an order of magnitude smaller than the ART-based dictionaries for all the datasets. The 3-Grams bitmap-trie is only 1.4× larger than Double-Char's fixed-length array of the same size.

Discussion: Schemes that compress more are slower, except that the original ALM is strictly worse than the other schemes in both dimensions. The latency gaps between schemes are generally larger than the compression rate gaps. We evaluate this trade-off in Section 7 by applying the HOPE schemes to in-memory search trees.

6.2 Dictionary Build Time

We next measure how long HOPE takes to construct the dictionary using each of the six compression schemes. We record the time HOPE spends in the modules from Section 4.2 when building a dictionary: (1) Symbol Selector, (2) Code Assigner, and (3) Dictionary. The last step is the time required to populate the dictionary from the key samples. We present only the Email dataset for this experiment; the results for the other datasets produce similar results and thus we omit them. For the variable-length-interval schemes, we perform the experiments using two dictionary sizes (2^{12}, 2^{16}).

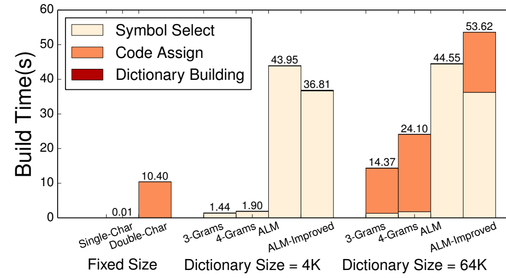

Figure 9: Dictionary Build Time – A breakdown of the time it takes for HOPE to build dictionaries on a 1% sample of email keys.

Figure 9 shows the time breakdown of building the dictionary in each scheme. First, the Symbol Selector dominants the cost for ALM and ALM-Improved because these schemes collect statistics for substrings of all lengths, which has a super-linear cost relative to the number of keys. For the other schemes, the Symbol Selector's time grows linearly with the number of keys. Second, the time used by the Code Assigner rises dramatically as the dictionary size increases because the Hu-Tucker algorithm has quadratic time complexity. Finally, the Dictionary build time is negligible compared to the Symbol Selector and Code Assigner modules.

7 SEARCH TREE EVALUATION

To experimentally evaluate the benefits and trade-offs of applying HOPE to in-memory search trees, we integrated HOPE into five data structures: SuRF, ART, HOT, B+tree, and Prefix B+tree (as described in Section 5). Based on the microbenchmark results in Section 6, we evaluate six HOPE configurations for each search tree: (1) Single-Char, (2) Double-Char, (3) 3-Grams with 64K (2^{16}) dictionary entries, (4) 4-Grams with 64K dictionary entries, (5) ALM-Improved with 4K (2^{12}) dictionary entries, and (6) ALM-Improved with 64K dictionary entries. We include the original search trees as baselines (labeled as "Uncompressed"). We choose 64K for 3-Grams, 4-Grams, and ALM-Improved so that they have the same dictionary size as Double-Char. We evaluate an additional ALM-Improved configuration with 4K dictionary size because it has a similar dictionary memory as Double-Char, 3-Grams (64K), and 4-Grams (64K). We exclude the original ALM scheme because it is always worse than the others.

7.1 Workload

We use the YCSB-based [21] index-benchmark framework proposed in the Hybrid Index [50] and later used by, for

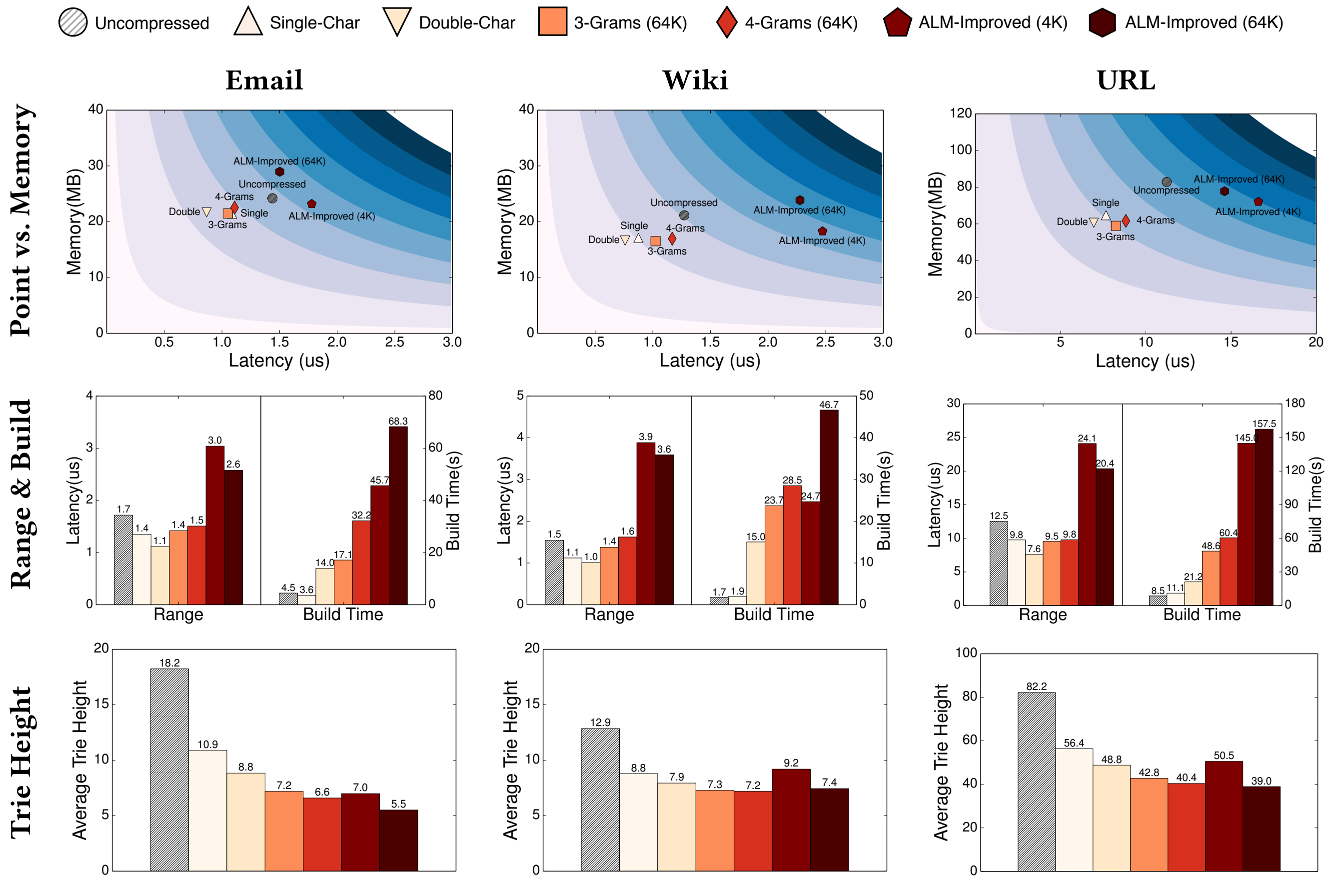

Figure 10: SuRF YCSB Evaluation – Runtime measurements for executing YCSB workloads on HOPE-optimized SuRF with three datasets. The trie height is the average height of each leaf node after loading all keys.

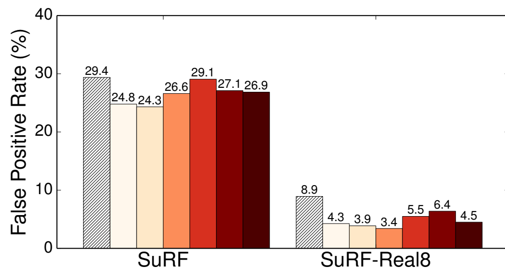

Figure 11: SuRF False Positive Rate – Point queries on email keys. SuRF-Real8 means it uses 8-bit real suffixes.

example, HOT [17] and SuRF [51]. We use the YCSB workloads **C** and **E** with a Zipf distribution to generate **point** and **range** queries. Point queries are the same for all trees. Each range query for ART, HOT, B+tree, and Prefix B+tree is a start key followed by a scan length. Because SuRF is a filter, its range query is a start key and end key pair, where the end key is a copy of the start key with the last character increased by one (e.g., ["`com.gmail@foo`", "`com.gmail@fop`"]). We replace the original YCSB keys with the keys in our email, wiki and URL datasets. We create one-to-one mappings between the YCSB keys and our keys during the replacement to preserve the Zipf distribution.[6]

7.2 YCSB Evaluation

We start each experiment with the building phase using the first 1% of the dataset's keys. Next, in the loading phase, we insert the keys one-by-one into the tree (except for SuRF because it only supports batch loading). Finally, we execute 10M queries on the compressed keys with a single thread using a combination of point and range queries according to the workload. We obtain the point, range, and insert query latencies by dividing the corresponding execution time by the number of queries. We measure memory consumption (HOPE size included) after the loading phase.

Figures 10 to 12 show the benchmark results. Range query and insert results for ART, HOT, B+tree, and Prefix B+tree are in [52] (Appx. D) because the performance results are

[6] We omit the results for other query distributions (e.g., uniform) because they demonstrate similar performance gains/losses as in the Zipf case.

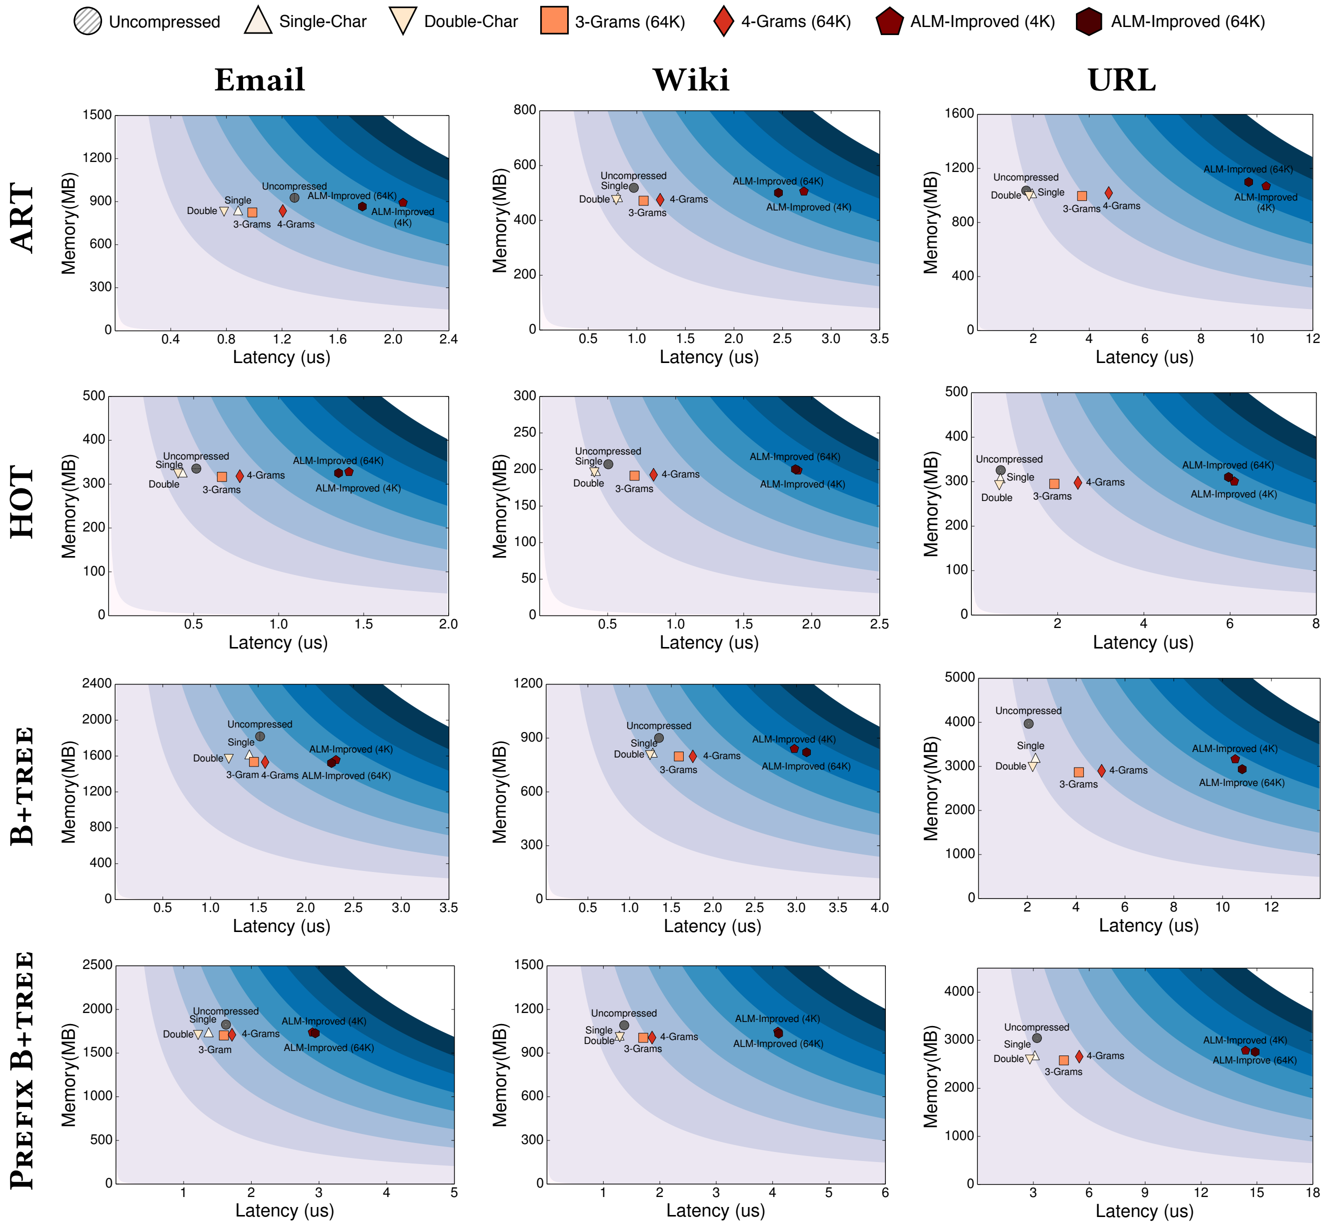

Figure 12: YCSB Point Query Evaluation on ART, HOT, B+tree, and Prefix B+tree – Measurements for executing YCSB point query workloads on HOPE-optimized indexes.

similar to the corresponding point query cases for similar reasons. We first summarize the high-level observations and then discuss the results in more detail for each tree.

High-Level Observations: First, in most cases, multiple schemes in HOPE provide a *Pareto improvement* to the search tree's performance and memory-efficiency. Second, the simpler FIVC schemes, especially Double-Char, stand out to provide the best trade-off between query latency and memory-efficiency for the search trees. Third, more sophisticated VIVC schemes produce the lowest search tree memory in some cases. Compared to Double-Char, however, their small additional memory reduction does not justify the significant performance loss in general.

SuRF: The heatmaps in the first row of Figure 10 show the point query latency vs. memory trade-offs made by SuRFs with different HOPE configurations. We define a cost function $C = L \times M$, where L represents latency, and

M represents memory. This cost function assumes a balanced performance-memory trade-off. We draw the equi-cost curves (as heatmaps) where points on the same curve have the same cost.

HOPE reduces SuRF's query latencies by up to 41% in all workloads with non-ALM encoders. This is because compressed keys generate shorter tries, as shown in the third row of Figure 10. According to our analysis in Section 5, the performance gained by fewer levels in the trie outweighs the key encoding overhead. Although SuRF with ALM-Improved (64K) has the lowest trie height, it suffers high query latency because encoding is slow for ALM-Improved schemes.

Although the six HOPE schemes under test achieve compression rates of 1.5–2.5× in the microbenchmarks, they only provide ∼30% memory savings to SuRF. The reason is that compressing keys only reduces the number of internal nodes in a trie (i.e., shorter paths to the leaf nodes). The number of leaf nodes, which is often the majority of the storage cost, stays the same. SuRF with ALM-Improved (64K) consumes more memory than others because of its larger dictionary.

Section 6.1 showed that ALM-Improved (4K) achieves a better compression rate than Double-Char for email keys with a similar-sized dictionary. When we apply this scheme to SuRF, however, the memory saving is smaller than Double-Char even though it produces a shorter trie. This is because ALM-Improved favors long symbols in the dictionary. Encoding long symbols one-at-a-time can prevent prefix sharing. As an example, ALM-Improved may treat the keys "`com.gmail@c`" and "`com.gmail@s`" as two separate symbols and thus have completely different codes.

All schemes, except for Single-Char, add computational overhead in building SuRF. The dictionary build time grows quadratically with the number of entries because of the Hu-Tucker algorithm. One can reduce this overhead by shrinking the dictionary size, but this diminishes performance and memory-efficiency gains.

Finally, the HOPE-optimized SuRF achieves lower false positive rate under the same suffix-bit configurations, as shown in Figure 11. This is because each bit in the compressed keys carries more information and is, thus, more distinguishable than a bit in the uncompressed keys.

ART, HOT: Figure 12 shows that HOPE improves ART and HOT's performance and memory-efficiency for similar reasons as for SuRF because they are also trie-based data structures. Compared to SuRF, however, the amount of improvement for ART and HOT is less. This is for two reasons. First, ART and HOT include a 64-bit value pointer for each key, which dilutes the memory savings from the key compression. More importantly, as described in Section 4.2 and Section 5, ART and HOT only store partial keys using *optimistic common prefix skipping* (OCPS). HOT is more optimistic than ART as it only stores the *branching points* in a trie (i.e., the minimum-length partial keys needed to uniquely map a key to a value). Although OCPS can incur false positives, the DBMS will verify the match when it retrieves the tuple. Therefore, since ART and HOT store partial keys, they do not take full advantage of key compression. The portion of the URL keys skipped is large because they share long prefixes. Nevertheless, our results show that HOPE still provides some benefit and thus are worth applying to both data structures.

B+tree, Prefix B+tree: The results in Figure 12 show that HOPE is beneficial to search trees beyond tries. Because the B+trees use reference pointers to store variable-length string keys outside of each node, compressing the keys does not change the tree structure. In addition to memory savings, the more lightweight HOPE schemes (Single-Char and Double-Char) also improve the B+tree's query performance because of faster string comparisons and better cache locality. We validate this by running the point-query workload on email keys and measuring cache misses using `cachegrind` [5]. We found that Double-Char on TLX B+tree reduces the L1 and last-level cache misses by 34% and 41%, respectively.

Compared to plain B+trees, we observe smaller memory saving percentages when using HOPE on Prefix B+trees. This is because prefix compression reduces the storage size for the keys, and thus making the structural components of the B+tree (e.g., pointers) relatively larger. Although HOPE provides similar compression rates when applied to a Prefix B+tree and a plain B+tree, the percentages of space reduction brought by HOPE-compressed keys in a Prefix B+tree is smaller with respect to the entire data structure size.

As a final remark, HOPE still improves the performance and memory for highly-compressed trees such as SuRF. It shows that HOPE is orthogonal to many other compression techniques and can benefit a wide range of data structures.

8 CONCLUSIONS

We introduced HOPE, a dictionary-based entropy-encoding compressor. HOPE compresses keys for in-memory search trees in an order-preserving way with low performance and memory overhead. To help understand the solution space for key compression, we developed a theoretical model and then show that one can implement multiple compression schemes in HOPE based on this model. Our experimental results showed that using HOPE to compress the keys for five in-memory search trees improves both their runtime performance and memory-efficiency for many workloads.

Acknowledgements. This work was supported by the National Science Foundation under awards SDI-CSCS-1700521, IIS-1846158, and SPX-1822933, Google Research Grants, and the Alfred P. Sloan Research Fellowship program.

REFERENCES

[1] 2005. My explanation of the Hu-Tucker Algorithm. http://www-math.mit.edu/~shor/PAM/hu-tucker_algorithm.html.

[2] 2007. URL Dataset. http://law.di.unimi.it/webdata/uk-2007-05/uk-2007-05.urls.gz.

[3] 2017. Zlib. http://www.zlib.net/.

[4] 2018. LZ77 and LZ78. https://en.wikipedia.org/wiki/LZ77_and_LZ78.

[5] 2019. Cachegrind. http://valgrind.org/docs/manual/cg-manual.html.

[6] 2019. English Wikipedia Article Title. https://dumps.wikimedia.org/enwiki/20190701/enwiki-20190701-all-titles-in-ns0.gz.

[7] 2019. LZ4. https://lz4.github.io/lz4.

[8] 2019. Memory Prices. https://www.jcmit.net/memoryprice.htm.

[9] 2019. Snappy. https://github.com/google/snappy.

[10] 2019. TLX B+tree (formerly STX B+tree). https://github.com/tlx/tlx.

[11] Daniel Abadi, Samuel Madden, and Miguel Ferreira. 2006. Integrating compression and execution in column-oriented database systems. In *Proceedings of SIGMOD'06*. ACM, 671–682.

[12] Gennady Antoshenkov. 1997. Dictionary-based order-preserving string compression. *The VLDB Journal* 6, 1 (1997), 26–39.

[13] Gennady Antoshenkov, David Lomet, and James Murray. 1996. Order preserving string compression. In *ICDE'96*. IEEE, 655–663.

[14] Nikolas Askitis and Ranjan Sinha. 2007. HAT-trie: a cache-conscious trie-based data structure for strings. In *Proceedings of ACSC'07*. Australian Computer Society, Inc., 97–105.

[15] Rudolf Bayer and Karl Unterauer. 1977. Prefix B-trees. *ACM Transactions on Database Systems (TODS)* 2, 1 (1977), 11–26.

[16] Bishwaranjan Bhattacharjee et al. 2009. Efficient index compression in DB2 LUW. *Proceedings of VLDB'09* 2, 2 (2009), 1462–1473.

[17] Robert Binna, Eva Zangerle, Martin Pichl, Günther Specht, and Viktor Leis. 2018. HOT: A Height Optimized Trie Index for Main-Memory Database Systems. In *Proceedings of SIGMOD'18*. ACM, 521–534.

[18] Carsten Binnig, Stefan Hildenbrand, and Franz Färber. 2009. Dictionary-based order-preserving string compression for main memory column stores. In *Proceedings of SIGMOD'09*. ACM, 283–296.

[19] Matthias Boehm et al. 2011. Efficient in-memory indexing with generalized prefix trees. *BTW* (2011).

[20] Zhiyuan Chen, Johannes Gehrke, and Flip Korn. 2001. Query optimization in compressed database systems. *ACM SIGMOD Record* 30, 2 (2001), 271–282.

[21] Brian F Cooper et al. 2010. Benchmarking cloud serving systems with YCSB. In *Proceedings of SoCC'10*. ACM, 143–154.

[22] Siying Dong. 2018. personal communication. 2018-12-26.

[23] Edward Fredkin. 1960. Trie memory. *Commun. ACM* 3, 9 (1960).

[24] Goetz Graefe et al. 2011. Modern B-tree techniques. *Foundations and Trends® in Databases* 3, 4 (2011), 203–402.

[25] Steffen Heinz, Justin Zobel, and Hugh E Williams. 2002. Burst tries: a fast, efficient data structure for string keys. *ACM TOIS* 20, 2 (2002).

[26] Te C Hu and Alan C Tucker. 1971. Optimal computer search trees and variable-length alphabetical codes. *SIAM J. Appl. Math.* 21, 4 (1971).

[27] David A Huffman. 1952. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE* 40, 9 (1952), 1098–1101.

[28] Alfons Kemper and Thomas Neumann. 2011. HyPer: A hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. In *Proceedings of ICDE'11*. IEEE, 195–206.

[29] Thomas Kissinger, Benjamin Schlegel, Dirk Habich, and Wolfgang Lehner. 2012. KISS-Tree: smart latch-free in-memory indexing on modern architectures. In *Proceedings of DaMoN'12*. ACM, 16–23.

[30] Robert Lasch et al. 2019. Fast & Strong: The Case of Compressed String Dictionaries on Modern CPUs. In *DaMoN'19*. ACM.

[31] Tobin J Lehman and Michael J Carey. 1985. *A study of index structures for main memory database management systems.* Technical Report. University of Wisconsin-Madison Department of Computer Sciences.

[32] Viktor Leis, Michael Haubenschild, and Thomas Neumann. 2019. Optimistic Lock Coupling: A Scalable and Efficient General-Purpose Synchronization Method. *IEEE Data Eng. Bull.* 42, 1 (2019), 73–84.

[33] Viktor Leis, Alfons Kemper, and Thomas Neumann. 2013. The adaptive radix tree: ARTful indexing for main-memory databases. In *Proceedings of ICDE'13*. IEEE, 38–49.

[34] Viktor Leis, Florian Scheibner, Alfons Kemper, and Thomas Neumann. 2016. The ART of practical synchronization. In *DaMoN'16*. ACM.

[35] Justin J Levandoski, David B Lomet, and Sudipta Sengupta. 2013. The Bw-Tree: A B-tree for new hardware platforms. In *Proceedings of ICDE'13*. IEEE, 302–313.

[36] Yinan Li, Craig Chasseur, and Jignesh M Patel. 2015. A padded encoding scheme to accelerate scans by leveraging skew. In *Proceedings of SIGMOD'15*. ACM, 1509–1524.

[37] Chunwei Liu, McKade Umbenhower, Hao Jiang, Pranav Subramaniam, Jihong Ma, and Aaron J. Elmore. 2019. Mostly Order Preserving Dictionaries. In *Proceedings of ICDE'19*. IEEE.

[38] Yandong Mao, Eddie Kohler, and Robert Tappan Morris. 2012. Cache craftiness for fast multicore key-value storage. In *Proceedings of EuroSys'12*. ACM, 183–196.

[39] G Martin. 1979. Range encoding: an algorithm for removing redundancy from a digitised message. In *Video and Data Recording Conference, Southampton, 1979*. 24–27.

[40] Markus Mäsker, Tim Süß, Lars Nagel, Lingfang Zeng, and André Brinkmann. 2019. Hyperion: Building the Largest In-memory Search Tree. In *Proceedings of SIGMOD'19*. ACM, 1207–1222.

[41] Donald R Morrison. 1968. PATRICIA—practical algorithm to retrieve information coded in alphanumeric. *JACM* 15, 4 (1968), 514–534.

[42] Ingo Müller, Cornelius Ratsch, and Franz Faerber. 2014. Adaptive String Dictionary Compression in In-Memory Column-Store Database Systems. In *EDBT'14*. 283–294.

[43] Vijayshankar Raman et al. 2013. DB2 with BLU acceleration: So much more than just a column store. In *Proceedings of VLDB'13*, Vol. 6. VLDB Endowment, 1080–1091.

[44] Vijayshankar Raman and Garret Swart. 2006. How to wring a table dry: Entropy compression of relations and querying of compressed relations. In *Proceedings of VLDB'06*. VLDB Endowment, 858–869.

[45] Jun Rao and Kenneth A Ross. 2000. Making B+-trees cache conscious in main memory. In *ACM SIGMOD Record*, Vol. 29. ACM, 475–486.

[46] Michael Stonebraker, Samuel Madden, Daniel J Abadi, Stavros Harizopoulos, Nabil Hachem, and Pat Helland. 2007. The end of an architectural era:(it's time for a complete rewrite). In *Proceedings of VLDB'07*. VLDB Endowment, 1150–1160.

[47] Ziqi Wang, Andrew Pavlo, Hyeontaek Lim, Viktor Leis, Huanchen Zhang, Michael Kaminsky, and David G Andersen. 2018. Building a bw-tree takes more than just buzz words. In *Proceedings of SIGMOD'18*. ACM, 473–488.

[48] Ian H Witten, Radford M Neal, and John G Cleary. 1987. Arithmetic coding for data compression. *Commun. ACM* 30, 6 (1987), 520–540.

[49] JM Yohe. 1972. Hu-Tucker Minimum Redundancy Alphabetic Coding Method [Z] (Algorithm 428). *Commun. ACM* 15, 5 (1972), 360–362.

[50] Huanchen Zhang, David G Andersen, Andrew Pavlo, Michael Kaminsky, Lin Ma, and Rui Shen. 2016. Reducing the storage overhead of main-memory OLTP databases with hybrid indexes. In *Proceedings of SIGMOD'16*. ACM, 1567–1581.

[51] Huanchen Zhang, Hyeontaek Lim, Viktor Leis, David G Andersen, Michael Kaminsky, Kimberly Keeton, and Andrew Pavlo. 2018. SuRF: practical range query filtering with fast succinct tries. In *Proceedings of SIGMOD'18*. ACM, 323–336.

[52] Huanchen Zhang, Xiaoxuan Liu, David G. Andersen, Michael Kaminsky, Kimberly Keeton, and Andrew Pavlo. 2020. Order-Preserving Key Compression for In-Memory Search Trees. arXiv:cs.DB/2003.02391