

MILC: Inverted List Compression in Memory

Jianguo Wang Chunbin Lin Ruining He Moojin Chae
Yannis Papakonstantinou Steven Swanson

Department of Computer Science and Engineering
University of California, San Diego
{csjgwang, chunbinlin, r4he, mochae, yannis, swanson}@cs.ucsd.edu

ABSTRACT

Inverted list compression is a topic that has been studied for 50 years due to its fundamental importance in numerous applications including information retrieval, databases, and graph analytics. Typically, an inverted list compression algorithm is evaluated on its space overhead and query processing time. Earlier list compression designs mainly focused on minimizing the space overhead to reduce expensive disk I/O time in disk-oriented systems. But the recent trend is shifted towards reducing query processing time because the underlying systems tend to be memory-resident. Although there are many highly optimized compression approaches in main memory, there is still a considerable performance gap between query processing over compressed lists and uncompressed lists, which motivates this work.

In this work, we set out to bridge this performance gap for the first time by proposing a new compression scheme, namely, MILC (memory inverted list compression). MILC relies on a series of techniques including offset-oriented fixed-bit encoding, dynamic partitioning, in-block compression, cache-aware optimization, and SIMD acceleration. We conduct experiments on three real-world datasets in information retrieval, databases, and graph analytics to demonstrate the high performance and low space overhead of MILC. We compare MILC with 12 recent compression algorithms and experimentally show that MILC improves the query performance by up to 13.2 \times and reduces the space overhead by up to 4.7 \times .

1. INTRODUCTION

An inverted list is a sorted list of integers. Although simple, it is the standard structure in a wide range of applications. For instance, search engines usually rely on inverted lists to find relevant documents. Databases also heavily need inverted lists to accelerate SQL processing [3].

Inverted list compression is a topic that has been studied for 50 years due to its benefits in disk-oriented systems as

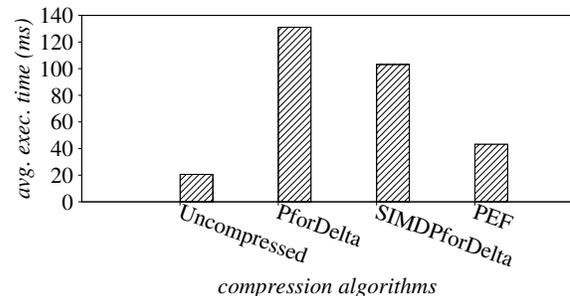


Figure 1: Executing queries over compressed (PforDelta [43], SIMDPforDelta [18], and PEF [25]) and uncompressed lists

well as recent memory-oriented systems. In disk-centric systems, compression can reduce expensive I/O time by shortening lists' sizes. Thus, list compression algorithms designed for disks (e.g., VB [32], Rice [28], Elias gamma [12]) mainly focus on reducing space overhead. The CPU decompression overhead is negligible compared to the saved I/O time due to the giant performance gap between disk and CPU. In recent memory-oriented systems, compression is also beneficial because it makes the system accommodate much more data than the physical memory capacity. For example, 100GB's raw data can be pushed to a server with 32GB DRAM. This reduces the total cost of ownership (TCO) since main memory is still an expensive resource. As a result, many compression algorithms have been developed for in-memory inverted lists, e.g., PforDelta [43], SIMDPforDelta [18], and PEF [25].

Motivation. We observe that there is still a considerable performance gap for query processing over compressed lists (with state-of-the-art compression algorithms) versus uncompressed lists in memory. For example, Figure 1 shows the (average) execution time of running 1000 real-world search engine queries over 300GB data.¹ It shows that the performance gap is 2.1 \times to 6.4 \times .

This raises an interesting question: *Is it possible to bridge this performance gap when operating on compressed data while still keeping low space overhead?* This work gives a positive answer to this question by proposing a new compression algorithm, namely, MILC (memory inverted list compression). Compared with uncompressed lists, MILC achieves a compression ratio up to 4.7 \times and executes queries nearly

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

Proceedings of the VLDB Endowment, Vol. 10, No. 8
Copyright 2017 VLDB Endowment 2150-8097/17/04.

¹We use the Web data described in Section 9.

as fast as that on uncompressed lists according to our experiments. Before diving into the technical descriptions, we define the problem first.

Problem statement. Given a sorted list L of n positive integers, the problem of inverted list compression is to compress L with as few as possible bits (smaller than the original list) while supporting query processing on compressed data as fast as possible. We mainly focus on supporting efficient membership testing – checking whether an element appears in a compressed list – because it is the core of many operations, e.g., intersection, union, difference, selection, join, successor finding, and top-k query processing.

Limitations of existing compression solutions. Existing compression algorithms for inverted lists, e.g., VB [32], Simple8b [2], GroupVB [9], and PforDelta [43], usually follow a golden rule that is to compute the differences (called *d-gaps*) between two consecutive integers (since all integers are sorted), and only encode the small d-gaps using fewer bits to save space. For example, let $L = \{8, 15, 20, 25, 35, 40, 52, 60, 65, 78, 90\}$, then existing solutions usually convert L to $L' = \{8, 7, 5, 5, 10, 5, 12, 8, 5, 13, 12\}$, where $L'[0] = L[0]$ and $L'[i] = L[i] - L[i - 1]$ ($i \geq 1$). But this is exactly why existing approaches cannot support membership testing efficiently: They have to decompress the entire list. Even with skip pointers as suggested in [24], still, they need to decompress at least one block of data on the fly. Moreover, the decompression overhead is high because they need to traverse the data at least twice in order to recover the original values: (1) decode each individual d-gap; (2) calculate prefix sums. Some compression algorithms may need more rounds, e.g., PforDelta requires another round of traversal to recover the exception values. Another important drawback of d-gap-based compression algorithms is that they are unfriendly to SIMD (single instruction multiple data) due to the inherent data dependencies in computing prefix sums [18]. Those compression algorithms that do not explicitly rely on d-gaps (such as EF [11, 34], PEF [25]) also have problems in dealing with membership testing efficiently as we explain more in related work.

Challenges. It is challenging to design a new compression approach to achieve similar query performance as uncompressed lists while keeping low space overhead, considering the problem has been studied for many years. This particularly holds today because the underlying hardware (e.g., big memory, new processors with large CPU caches and wide SIMDs) has changed significantly. The compression format should also be compliant with CPU cache lines and SIMD instructions such that membership testing can be executed even more efficiently. To solve the problem, we need to break the traditional rule by abandoning d-gaps. This will increase the space overhead naturally. Therefore, we need to design new techniques to reduce space overhead while maintaining high query performance.

Technical overview. To address the above challenges, we develop a novel compression scheme MILC (memory inverted list compression) that achieves similar membership testing performance with uncompressed lists. The basic idea of MILC is that it partitions an input list into different blocks, then it applies offset-based (instead of delta-based) encoding and uses the same number of bits to encode all the elements within the block (Section 4). This is crucial to the success of

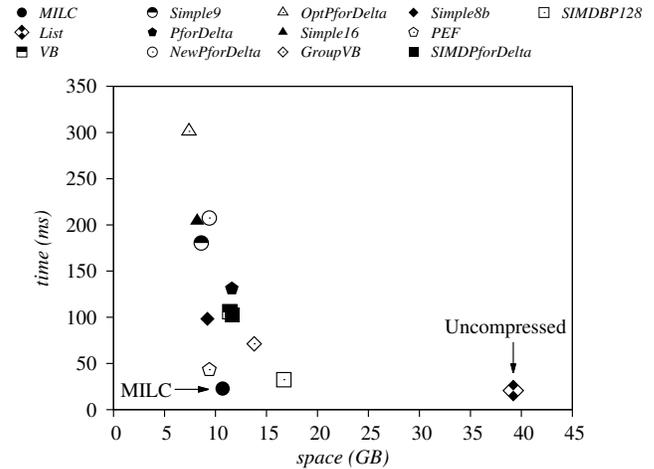


Figure 2: Experiments overview of MILC vs existing compression approaches on 300GB Web data in executing 1000 queries

MILC because it enables MILC to support membership testing directly on compressed data as we show in Section 4.

To further reduce space overhead and improve query performance, MILC employs four optimizations: (1) *Dynamic partitioning* (Section 5). It partitions a list into variable-sized blocks based on dynamic programming to minimize exception values in each block, i.e., elements in a block have low variance. This effectively reduces the space overhead because exceptions need more bits to represent, and all the other elements end up using the same high number of bits. (2) *In-block compression* (Section 6). MILC further splits every block into sub-blocks by smartly plugging in lightweight skip pointers to reduce the space overhead. (3) *Cache-aware optimization* (Section 7). MILC reorganizes data in a way that considers CPU cache line alignment. It can improve the performance of membership testing because CPU cache misses are reduced. (4) *SIMD acceleration* (Section 8). MILC also leverages SIMD for fast query processing.

Contribution. The main contribution of this work is a new compression scheme MILC that achieves similar membership testing performance with uncompressed lists while keeping low space overhead. MILC is tailored for modern computing hardware including big memory, fast CPU caches, and wide SIMD processing capabilities. To the best of our knowledge, this is the first inverted list compression algorithm that has such high query performance with low space overhead.

Experimental overview. We conduct experiments on datasets from information retrieval, databases, and graph analytics to demonstrate the advantages of MILC with a spectrum of 12 compression algorithms in terms of query performance and space overhead. Figure 2 shows a preview on 300GB Web data in answering 1000 user queries,² and Section 9 describes more details. It shows that MILC runs faster than existing compression approaches and consumes low space also. Thus, MILC represents the best tradeoff for inverted list compression in main memory in terms of time and space.

²We report the membership testing time to measure the effectiveness of MILC as described in Section 9.

```

SELECT d_year, s_nation, p_category,
       sum(lo_revenue - lo_supplycost) as profit
FROM   date, customer, supplier, part, lineorder
WHERE  lo_custkey = c_custkey
       AND lo_suppkey = s_suppkey
       AND lo_partkey = p_partkey
       AND lo_orderdate = d_datekey
       AND c_region = 'AMERICA'
       AND s_region = 'AMERICA'
       AND d_year = 1997
       AND p_mfgr = 'MFGR#1'

```

(a) star schema join

```

SELECT d_year, s_nation, p_category,
       sum(lo_revenue - lo_supplycost) as profit
FROM   lineorderfull
WHERE  c_region = 'AMERICA'
       AND s_region = 'AMERICA'
       AND d_year = 1997
       AND p_mfgr = 'MFGR#1'

```

(b) intersection

Figure 3: An example of star schema join and intersection

Paper organization. The rest of the paper is organized as follows. Section 2 presents the applications that require inverted list compression. Section 3 describes the existing work on inverted list compression. Section 4 presents the basic idea of MILC. Section 5 to Section 8 develop four optimizations used in MILC. Section 9 evaluates MILC experimentally. Section 10 concludes the work.

2. APPLICATIONS

In this section, we provide motivating applications that rely on inverted lists for efficient query processing. This means that a large range of applications can benefit from this work on inverted list compression.

2.1 Information retrieval

Information retrieval (IR) is a killer application of inverted lists to answer user queries with multiple terms [23]. IR systems store an inverted list for each term all the documents that contain the term. Taking the intersection or union of the lists for a set of query terms identifies those documents that contain all or at least one of the terms.

2.2 Database query processing

Inverted lists are also helpful in SQL databases, especially if there is logically one huge table and the query involves many predicates, e.g., a conjunction of predicates as shown in Figure 3b. In this case, most databases would precompute a list of matching row IDs for each predicate to facilitate the conjunction (or intersection) query [26].

Besides that, many *star schema joins* can also be framed as conjunctive queries as suggested in prior works [3, 26]. For instance, the star schema join in Figure 3a can be converted to the intersection query in Figure 3b where `lineorderfull` is a logical huge table that is created offline as follows [3, 26]: (1) `lineorderfull = lineorder` \bowtie `date` \bowtie `customer` \bowtie `supplier` \bowtie `part`; (2) add an additional *id* column to `lineorderfull`. Note that `lineorderfull` and `lineorder`

have the same number of tuples due to many-to-one mapping.³ Then the star schema join in Figure 3a can be reduced to the intersection query in Figure 3b as $L_1 \cap L_2 \cap L_3 \cap L_4$ where:

$L_1 = \pi_{id} \sigma_{c_region='AMERICA'}(\text{lineorderfull})$

$L_2 = \pi_{id} \sigma_{s_region='AMERICA'}(\text{lineorderfull})$

$L_3 = \pi_{id} \sigma_{d_year=1997}(\text{lineorderfull})$

$L_4 = \pi_{id} \sigma_{p_mfgr='MFGR\#1'}(\text{lineorderfull})$

Since all the lists are precomputed and stored in an index structure such as B-tree, then the query plan can be executed as follows: L_1 and L_2 are intersected first, then the results of $L_1 \cap L_2$ are intersected with L_3 , finally the results of $L_1 \cap L_2 \cap L_3$ are intersected with L_4 .

2.3 Graph analytics

Graph databases represent another family of advocates of inverted lists. There are usually two types of inverted lists in graph databases: adjacency lists and association lists. An adjacency list is dedicated for a vertex to maintain all neighborhood vertices connected with it. An association list is dedicated for an object (e.g., a Facebook page) to keep all relevant associations where an association is specified by a source object, destination object, and association type (e.g., `tagged-in`, `likers`) [33]. Many queries over these graphs can be answered efficiently using inverted lists. For example, finding “Restaurants in San Francisco liked by Mike’s friends” reduces to finding the intersection of the adjacency list of “Mike” and the association lists of “Restaurants” and “San Francisco”; discovering common friends among a group of people transforms to computing the intersection of several adjacency lists.

2.4 More applications

In addition, there are many other applications that heavily use inverted lists for fast query processing. For example, data integration systems build inverted lists for *q*-grams to find the most similar strings [15]. Data mining systems deploy inverted lists for fast data cube operations such as slicing, dicing, rolling up and drilling down [19, 21]. XML databases depend on inverted lists to find twig patterns efficiently [4]. Key-value stores also organize data elements falling into the same bucket (hash collision) with a chained list, which is essentially an inverted list [10].

3. RELATED WORK

In this section, we take a retrospective look at the major inverted list compression algorithms developed so far. Figure 4 shows a brief history.

As mentioned in Section 1, the common wisdom of a decent inverted list compression algorithm is to compute the deltas (a.k.a *d-gaps*) between two consecutive integers first and only encode the *d-gaps* to save space.⁴ To prevent from decompressing the entire list during query processing, it organizes those *d-gaps* into blocks (of say 128 elements per

³The many-to-one mapping is from a foreign key in the fact table to the primary key in the dimension table.

⁴Early compression algorithms (before 1990) do not follow this rule and encode each element of a list individually, e.g., Rice [28] and Elias gamma [12]. However, they are far worse than today’s compression algorithms, e.g., PforDelta, in terms of both query execution time and space overhead. Thus, we ignore them in this work.

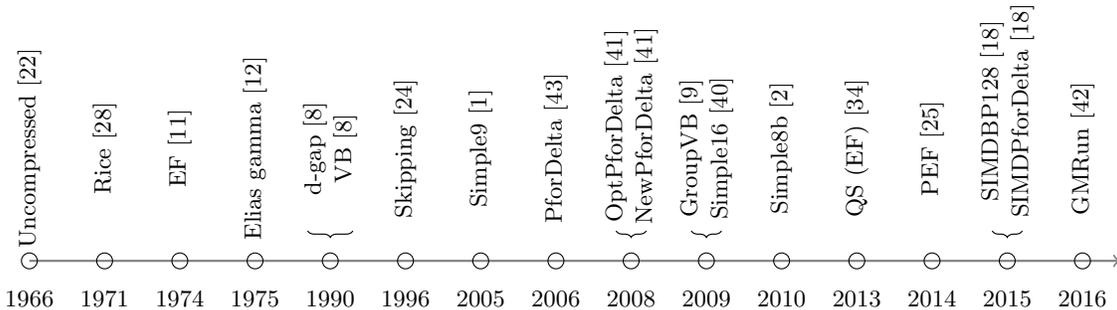


Figure 4: A brief history of representative inverted list compression approaches

block⁵) and builds a skip pointer per block such that only a block of data needs to be decompressed. Today, most excellent compression methods exactly follow this convention, including PforDelta [43] (and its descendants such as NewPforDelta [41] and OptPforDelta [40]), VB [8], GroupVB [9], Simple9 [1], Simple16 [40], and Simple8b [2].

Among them, PforDelta is a mature algorithm that is commonly used because it has a good tradeoff between query execution time (or decompression speed) and space overhead [40, 41]. The basic idea is that it compresses a block of 128 d-gaps by choosing the smallest b in the block such that a majority of elements (say 90%) can be encoded in b bits (called *regular values*). It then encodes the 128 values by allocating 128 b -bit slots, plus some extra space at the end to store the values that cannot be represented in b bits (called *exceptions*). Each exception takes 32 bits while each regular value takes b bits. In order to indicate which slots are exceptions, it uses the unused b -bit slots from the pre-allocated 128 b -bit slots to construct a linked list, such that the b -bit slot of one exception stores the offset to the next exception. In the case where two exceptions are more than 2^b slots apart, it adds additional forced exceptions between the two slots. Besides PforDelta, there are many variations, e.g., NewPforDelta [41] and OptPforDelta [41]. For example, OptPforDelta was designed to reduce the space overhead of PforDelta but at the expense of more decompression overhead.

However, PforDelta (as well as its variations) still takes considerable time to decompress a block of data, because it usually takes three phases for decompression: (1) It needs to copy the 128 b -bit values from the slots into an integer array via bit manipulations; (2) It then walks through the linked list of exceptions and copies their values into the corresponding array slots; (3) It also goes through the integer array again to perform prefix sums to recover the original values.

Recently, there is a resurgence of EF encoding [34] which is not directly based on d-gaps. Actually, EF encoding was originally proposed in 1974 [11], but it did not attract too much attention until 2013 when Vigna discovered that EF encoding can be competitive with PforDelta [34]. It encodes a sequence of integers using a low-bit array and a high-bit array. The low-bit array stores the lower $b = \log \frac{U}{n}$ bits of each element contiguously where U is the maximum possible element and n is the number of elements in the list.

⁵The block size represents a tradeoff between space and time and several existing works suggest 128 as the block size [2, 41].

The high-bit array then stores the remaining higher bits of each element as a sequence of unary-coded d-gaps. Later on, Giuseppe and Rossano improved it by leveraging the clustering property of a list, making it outperform PforDelta for some intersection queries but not union queries [25]. We call it PEF (Partitioned Elias Fano) in this paper. The fundamental problem of EF encoding (and its descendants including PEF) is that query processing is still not as efficient as it can be due to two reasons: (1) It needs to *sequentially* go through every *bit* in the high-bit array until a match is found, which requires many bit manipulations; (2) After that, it also needs to *sequentially* examine 2^b possible ties in the lower-bit array which can be slow if b is large.

In the literature, there were also proposals about reordering document IDs for better compression ratio, e.g., [40, 42]. This is orthogonal to this work and we do not consider them in this work. Besides that, this work focuses on data compression for inverted lists, which is also orthogonal to data compression in databases [20].

Currently, there is also a trend of leveraging SIMD to accelerate the decompression speed of existing compression methods, such as SIMDPforDelta [18]. The main idea is to reorganize data elements in a way such that a single SIMD operation processes multiple elements. However, for d-gap based compression approaches, computing prefix sums usually cannot leverage SIMD efficiently because of the intrinsic data dependencies [18].

Finally, we comment on the compression approaches used in Vertica [17] and Brighthouse [31] that are also relevant to this work. The main idea of both approaches is to partition a list into blocks and maintain metadata for each block to support query processing. However, MILC is significantly different from them in (1) how to compress the elements within a block; (2) how to partition the list into different blocks; and (3) how to store the metadata information.

4. BASIC COMPRESSION STRUCTURE

In this section, we present the basic compression structure as a starting point of MILC.

Storage structure. MILC’s basic structure follows the PforDelta compression algorithm in partitioning the list L into blocks but is different in compressing the data elements within a block. It splits L into $\lceil \frac{n}{m+1} \rceil$ partitions where $(m+1)$ is the size of each partition except the last partition if n is not divisible by $(m+1)$. The choice of m will be discussed later on. The first element (i.e., the minimal value) of each block serves as a skip pointer and all the skip

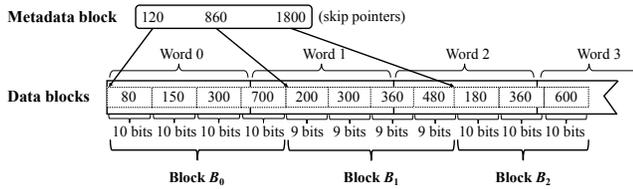


Figure 5: An example of storage format for $L = \{120, 200, 270, 420, 820, 860, 1060, 1160, 1220, 1340, 1800, 1980, 2160, 2400\}$ and $m = 4$

pointers are stored in a *metadata block*. Thus, each partition except the last one contains exactly m elements, called a *data block*. The metadata block contains $\lceil \frac{n}{m+1} \rceil$ elements (skip pointers); each element points to a data block.

MILC stores a data block as follows. Suppose the block contains the following m elements: $\{a_0, a_1, \dots, a_{m-1}\}$ and β is its skip pointer. MILC stores each element a_i as the difference between a_i and the skip pointer, i.e., $a_i - \beta$, instead of $a_i - a_{i-1}$ as in PforDelta [43]. We call it offset-based encoding instead of delta-based encoding. So the maximum difference is $(a_{m-1} - \beta)$, which can be encoded in $b = \lceil \log(a_{m-1} - \beta + 1) \rceil$ bits. Indeed, every element in the same data block is represented in b bits – unlike PforDelta, MILC does not use exceptions. Different blocks may use different number of bits to represent their values. To save space, MILC fully utilizes the 32 bits of a word by packing as many values as possible and padding the residual bits of the word (if any) with the next value if possible.

MILC stores the metadata block in the same format as PforDelta. Each entry in the metadata block contains the metadata information of a data block including the start value (32 bits), offset (32 bits), and the number of bits b (8 bits) to encode the data block.

Example. As an example, Figure 5 depicts the structure and storage format of $L = \{120, 200, 270, 420, 820, 860, 1060, 1160, 1220, 1340, 1800, 1980, 2160, 2400\}$ consisting of 14 elements and $m = 4$. It stores the list as follows: (1) It divides L into $\lceil \frac{14}{4+1} \rceil = 3$ partitions where each partition (except the last one) has 5 elements: $\{120, 200, 270, 420, 820\}$; $\{860, 1060, 1160, 1220, 1340\}$; $\{1800, 1980, 2160, 2400\}$. (2) It extracts the first element from each partition and puts it to the metadata block: $\{120, 860, 1800\}$. As a result, the data blocks are: $\{200, 270, 420, 820\}$ (the skip pointer is 120), $\{1060, 1160, 1220, 1340\}$ (the skip pointer is 860), and $\{1980, 2160, 2400\}$ (the skip pointer is 1800). (3) It subtracts the skip pointer from each data block. For example, for the first data block (B_0), since its skip pointer is 120, then it is stored a sequence of values by subtracting 120, i.e., $\{80, 150, 300, 700\}$. (4) It determines the smallest b in each block such that all the elements can be encoded in b bits, e.g., for block B_0 , the maximum number 700 can be encoded in 10 bits, thus, it uses 10 bits to represent every element in B_0 . (5) It serializes each data block as compact as possible (Figure 5). For example, B_0 has four 10-bit elements, but only the first three elements can be entirely packed into a 32-bit word. The fourth 10-bit element needs to span two words: the lower 2 bits are stored in the current word and the higher 8 bits are stored in a new word. Then B_1 is stored immediately after B_0 by sharing the last word in B_0 without wasting a single bit as is shown in Figure 5.

Next, we discuss the choice of m . If m is large, then it needs more bits to encode the data blocks because each data block spans a wide range, thus the overall space tends to be high. On the other hand, if m is small, then there will be more elements in the metadata block, which incurs high space overhead. Following the convention of PforDelta, we set m to 128 but other values are also possible. Later on in Section 5, we discuss the choice of m dynamically to minimize the overall space.

Supporting membership testing. MILC’s storage structure supports membership testing over a compressed list directly without decompressing a whole block, because MILC uses fixed-bit encoding to represent each element in the block using the same number of bits while preserving the order. Let e be a search key, then it performs binary search in the metadata block and jumps to the potential data block and runs another binary search but using a new key $(e - \beta)$ where β is the skip pointer of the data block.

Next, we explain how to implement binary search within a data block (as it is trivial to perform binary search in the metadata block as it is uncompressed). The problem requires bit manipulations because each element takes b bits, which are not necessarily 8 bits – **byte** type, 16 bits – **short** type, or 32 bits – **int** type that are natively accessible by a programming language. Observe that the core of binary search is obtaining the k -th value because binary search needs to consistently compare the search key with the middle value within a search range. Conventionally on the integer array, it is $A[k]$ to access the k -th value of an array A . But on the bit array, it requires a few bit manipulations to convert a b -bit value to a 32-bit value. For example in Figure 5, assume $b = 10$ and A be the compressed data blocks, then the first four values are:

- 1st: $(A[0] \& 0X03FF)$
- 2nd: $(A[0] \gg 10) \& 0X03FF$
- 3rd: $(A[0] \gg 20) \& 0X03FF$
- 4th: $(A[0] \gg 30) \mid ((A[1] \& 0X00FF) \ll 2)$

Space overhead analysis. It is evident that the space overhead of the storage format is high compared with PforDelta. Let us roughly analyze how high it is by assuming the elements in a list are equally apart to facilitate the analysis. Let θ be the gap between two consecutive elements in a block, m be the block size (e.g., $m = 128$), p be the exception ratio (e.g., $p = 10\%$ [40]), then PforDelta requires the following b bits to represent an element:

$$b = \lceil \log(\theta + 1) \rceil + 32 \times p \approx \log \theta + 3.2$$

Then for the basic compression structure, the gap now becomes $m \times \theta$. Thus, it requires the following b' bits to represent an element in the block:

$$b' = \lceil \log(m \times \theta + 1) \rceil \approx \log(128 \times \theta) = \log \theta + 7$$

That means the basic compression incurs $7 - 3.2 = 3.8$ more bits per element compared to PforDelta (but with much higher performance). Thus, in next sections, we present techniques to reduce the space overhead while keeping fast query performance.

Remark. It is worth noting that the basic structure of MILC presented in this section is based on PforDelta but with two important modifications. (1) Instead of computing deltas, MILC stores the offset values. (2) Instead of setting b such that a *majority* of elements are within 2^b , MILC deter-

mines b such that *all* elements are within 2^b , i.e., fixed-bit encoding. We also note that the basic structure of MILC is different from FOR [13], which also applies fixed-bit encoding. FOR was designed for storing a page of non-sorted elements. Thus, it needs to decompress the entire page during query processing. However, MILC can support membership testing directly over compressed data without even decompressing a whole block.

5. DYNAMIC PARTITIONING

In this section, we present a technique of dynamic partitioning to reduce the space overhead while keeping high query performance.

Why dynamic partitioning? The reason why the basic compression structure in Section 4 consumes much space is that it *evenly* partitions an input list into blocks (we call it *static partitioning*). So, if there are some exceptions⁶ in the block, then all the elements within the block have to use the same high number of bits to represent. In other words, static partitioning is vulnerable to data skew. As an example, if a data block is $\{3, 8, 10, 15, 150\}$, then it requires 8 bits just because of 150 (an exception) while the other values actually only need 4 bits to represent. Thus, it could save a lot of space if we can *dynamically* split a list in a way that similar (or close) elements are stored together to minimize exceptions.⁷

Thus, the problem is: Given a sorted list L of integers, how to split L into blocks such that the overall space overhead is minimized? The representation of each individual block still follows the fixed-bit encoding (Section 4) in order to support membership testing efficiently.

Dynamic partitioning. We propose a partitioning scheme by converting the problem to a dynamic programming problem for minimizing the overall space overhead. Dynamic programming can model the space overhead of partitioning a list at different positions such that it picks up a partitioning strategy with the lowest space overhead. Let E_i be the space overhead of representing $L[0 : i]$, then it splits $L[0 : i]$ at the j -th ($j < i$) position: $L[0 : j]$ and $L[j + 1 : i]$. Therefore, the space overhead of $L[0 : i]$ is the summation of the space overhead of $L[0 : j]$ and $L[j + 1 : i]$. Let $c(j, i)$ ($j \leq i$) be the space overhead of representing $L[j : i]$ and ℓ be the maximal size of a block, then,

$$E_i = \min_{j=\max\{0, i-\ell\}}^{i-1} (E_j + c(j + 1, i)) \quad (1)$$

Next, we analyze $c(j, i)$ used in Equation 1. Since the first element of the partition (i.e., $L[j]$) is stored in the metadata block as a skip pointer and the remaining values $L[j + 1 : i]$ are stored in a data block, then we compute the overhead of the two parts separately.

First of all, we analyze the space overhead of the skipping information (metadata block), which requires the following information per data block: (1) start value (32 bits), i.e., $L[j]$; (2) offset (32 bits) indicating where the data block starts from; (3) number of elements in the block (8 bits);

⁶A value is called an *exception* value if it is obviously larger than most other values in the block.

⁷Note that PforDelta does not have this issue because it uses different number of bits to represent regular elements and exceptions, but PforDelta cannot support membership testing directly on compressed data.

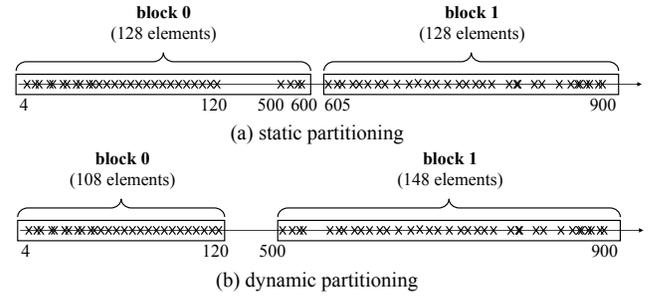


Figure 6: An example of dynamic partitioning

(4) number of bits to encode the block (8 bits). Thus, the skipping information per data block needs $32+32+8+8 = 80$ bits.

Second, we consider the space overhead of the data block $L[j + 1 : i]$. Recall that each element in the block is stored as the difference between it and $L[j]$. Among them, the maximal gap is $L[i] - L[j]$, which requires $\lceil \log(L_i - L_j + 1) \rceil$ bits. And there are $(i - j)$ elements in the block, thus, it requires $\lceil \log(L_i - L_j + 1) \rceil \times (i - j)$ bits in total. Therefore, $c(j, i)$ can be computed as follows:

$$c(j, i) = \lceil \log(L_i - L_j + 1) \rceil \times (i - j) + 80 \quad (2)$$

Example. Figure 6 shows an example where L contains 256 elements and $L = \{4, \dots, 120, 500, \dots, 600, 605, \dots, 900\}$. Using the fixed-length partitioning (or static partitioning) with the block size being 128 (Figure 6a), then L is partitioned into two blocks and the last element in the first block is 600. For the first block, each element takes $\lceil \log(600 - 4 + 1) \rceil = 10$ bits. While the dynamic partitioning (Figure 6b) can determine that the first 108 elements are similar and thus group them together. As a result, each element in the first block requires only $\lceil \log(120 - 4 + 1) \rceil = 7$ bits. For each element in the second block, it takes 9 bits for both static and dynamic partitioning. As a result, static partitioning takes $10 \times 128 + 9 \times 128 = 2432$ bits while dynamic partitioning takes $7 \times 108 + 9 \times 148 = 2088$ bits.

Determining the maximal block size ℓ . The maximal group size ℓ is very important in MILC's dynamic partitioning scheme: if it is too small (say $\ell = 1$), then the optimal partitioning can be missed; if it is too large (say $\ell = |L|$), it takes too much time to find the optimal partitioning. In Theorem 1, we show that the maximal block size after dynamic partitioning is less than 2λ , where λ is the number of bits to maintain the skipping information per block (i.e., $\lambda = 80$). As a result, we set $\ell = 160$ in Equation 1. Note that Theorem 1 is also very useful in Section 6, in determining lightweight skip pointers.

THEOREM 1. *A data block has at most 2λ elements after dynamic partitioning, where λ is the number of bits needed to store the skipping information per block.*

PROOF. We show that after partitioning, if a block still has more than 2λ elements, then we can always find a lower space cost by splitting the block into two parts, which contradicts with the optimality property achieved by dynamic programming. Without loss of generality, suppose a block contains m elements (Figure 7): a_0, a_1, \dots, a_{m-1} and β is the skip pointer of the block. We assume $m \geq 2\lambda$, next, we show

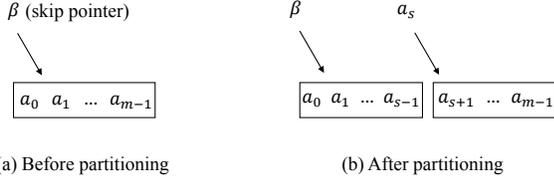


Figure 7: An example of illustrating the maximal size of a data block

that there always exists a lower cost by splitting the block into two.

Before partitioning, the total number of bits X required is (Figure 7a):

$$X = \lceil \log(a_{m-1} - \beta + 1) \rceil \times m$$

Then we split the block into two parts by picking up the middle value $a[s]$ (where $s = \lfloor \frac{m}{2} \rfloor$) as a skip pointer. Therefore the partitions are $[0 : s - 1]$ and $[s + 1 : m - 1]$. Then the total number of bits X' is (Figure 7b):

$$X' = \underbrace{\lceil \log(a_{s-1} - \beta + 1) \rceil}_{\text{1st block}} \times s + \underbrace{\lceil \log(a_{m-1} - a_s + 1) \rceil \times (m - 1 - s)}_{\text{2nd block}} + \underbrace{\lambda}_{\text{skip pointer } a_s}$$

Next we show $X' \leq X$ if $m \geq 2\lambda$ by introducing an intermediate variable Y :

$$Y = \lceil \log(a_{s-1} - \beta + 1) \rceil \times \frac{m}{2} + \lceil \log(a_{m-1} - a_s + 1) \rceil \times \frac{m}{2} + \lambda$$

Since $X' < Y$ no matter whether m is even or odd. Next, we show $Y \leq X$ if $m \geq 2\lambda$. $Y \leq X$ is equivalent to:

$$m \times (2\lceil \log(a_{m-1} - \beta + 1) \rceil - \lceil \log(a_{m-1} - a_s + 1) \rceil - \lceil \log(a_{s-1} - \beta + 1) \rceil) \geq 2\lambda$$

So, the remaining proof is to show $f = 2\lceil \log(a_{m-1} - \beta + 1) \rceil - \lceil \log(a_{m-1} - a_s + 1) \rceil - \lceil \log(a_{s-1} - \beta + 1) \rceil \geq 1$ since $m \geq 2\lambda$. We prove it by contradiction. Thus, we assume that $f = 0$ (as $f \geq 0$ and f is an integer).

$$\begin{aligned} \lceil \log(a_{m-1} - \beta + 1) \rceil - \lceil \log(a_{m-1} - a_s + 1) \rceil &= 0 \\ \lceil \log(a_{m-1} - \beta + 1) \rceil - \lceil \log(a_{s-1} - \beta + 1) \rceil &= 0 \end{aligned}$$

In order to hold:

$$\begin{aligned} \log(a_{m-1} - \beta + 1) - \log(a_{m-1} - a_s + 1) &< 1 = \log 2 \\ \log(a_{m-1} - \beta + 1) - \log(a_{s-1} - \beta + 1) &< 1 = \log 2 \end{aligned}$$

That is,

$$\begin{aligned} \log \frac{a_{m-1} - \beta + 1}{a_{m-1} - a_s + 1} &< \log 2 \\ \log \frac{a_{m-1} - \beta + 1}{a_{s-1} - \beta + 1} &< \log 2 \end{aligned}$$

That is,

$$\begin{aligned} a_{m-1} - \beta + 1 &< 2(a_{m-1} - a_s + 1) \\ a_{m-1} - \beta + 1 &< 2(a_{s-1} - \beta + 1) \end{aligned}$$

Summing up the left sides gives $2(a_{m-1} - \beta + 1) < 2(a_{m-1} - a_s + 1 + a_{s-1} - \beta + 1)$, i.e., $a_s - a_{s-1} < 1$, which is a contradiction since any two consecutive numbers differ at least 1. \square

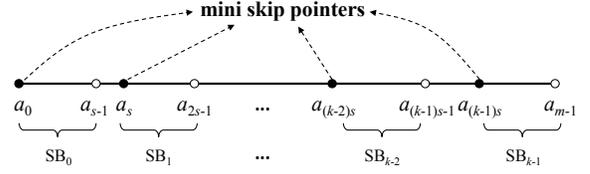


Figure 8: Split a data block into sub-blocks

Time complexity. Let n be the list size, then the time complexity of finding the optimal partitioning is $O(\ell n)$, which can be regarded as $O(n)$ since ℓ is a small constant ($\ell \leq 160$), as shown by Theorem 1.

Supporting membership testing. With dynamic partitioning, the structure supports membership testing efficiently in the same way as presented in Section 4, because a data block is still represented using fixed-bit encoding.

Remark. MILC is different from VSEncoding [30] because VSEncoding applies dynamic partitioning for PforDelta, i.e., it applies dynamic partitioning to a list of delta values. However, as we have shown in Section 4, this is exactly why membership testing is not supported efficiently due to the inevitable high decompression overhead. Besides that, VSEncoding chooses the maximal block size ℓ arbitrarily, which can either miss the optimal result or incur more preprocessing time. However, MILC solves the problem in an elegant way by proving in Theorem 1 that the maximal block will not exceed 2λ (which is 160).

6. IN-BLOCK COMPRESSION

In this section, we further reduce the space overhead of the compression from another angle while keeping fast query performance.

Why in-block compression? The dynamic partitioning groups similar elements to the same data block. However, all the elements in the same block have to use the number of bits based on the *maximal* element (i.e., the rightmost element) in the block in order to support fast search. But this on the other hand wastes some space for smaller elements. As an example in Figure 6b, after dynamic partitioning, the first block needs 7 bits to encode every element because the maximal value is 120. However, many smaller elements such as 10 and 20 do not necessarily need 7 bits. Therefore, in-block compression aims to use fewer bits to encode each element within a block to reduce the overall space overhead.

In-block compression structure. The main idea of in-block compression is to treat the elements in a data block as a micro inverted list and compress them using the approaches described in previous sections (with modifications) by splitting a block into sub-blocks. Before presenting the partitioning details, we answer the following question first: If partitioning a block into sub-blocks can reduce the overall space overhead, why previous dynamic partitioning (Section 5) – supposed to return a partitioning scheme with the *lowest* space cost – fails to capture such partitioning? That is because the overhead of maintaining a skip pointer within the block is *much smaller* than that outside the block. For example, it needs 80 bits to maintain a skip pointer outside the block as described in Section 5, but it only needs b (say 10) bits to maintain a skip pointer within the block (called a *mini* or *lightweight* skip pointer) as we explain below.

In particular, in-block compression applies the static partitioning method presented in Section 4 to *evenly* (except the last sub-block) split the elements into sub-blocks. Note that MILC does not apply the dynamic partitioning approach (Section 5) for in-block compression because that will incur more space as we explain later in the discussion part at the end of this section. Formally, suppose a block contains m elements (Figure 8): $\{a_0, a_1, \dots, a_{m-1}\}$, and let k be the number of sub-blocks, then the in-block compression partitions the block into the following k sub-blocks: $\{a_0, a_1, \dots, a_{s-1}\}, \{a_s, a_{s+1}, \dots, a_{2s-1}\}, \dots, \{a_{(k-2)s}, a_{(k-2)s+1}, \dots, a_{(k-1)s-1}\}, \{a_{(k-1)s}, a_{(k-1)s+1}, \dots, a_{m-1}\}$, where $s = \lfloor \frac{m}{k} \rfloor$. The first element of each sub-block serves as a mini skip pointer and all the mini skip pointers are stored together. Then, for every mini skip pointer in the block, it uses $\lceil \log(a_{m-1} - \beta + 1) \rceil$ bits where β is the skip pointer of the block. For every other element in the block, it uses the following number of bits b to encode:

$$b = \max\{\lceil \log(a_{s-1} - a_0 + 1) \rceil, \dots, \lceil \log(a_{m-1} - a_{(k-1)s} + 1) \rceil\} \quad (3)$$

Note that without in-block compression, each element originally takes $b' = \lceil \log(a_{m-1} - \beta + 1) \rceil$ bits and $b' \geq b$.

Besides that, in-block compression needs to maintain an extra 16-bit global information for all the sub-blocks: number of bits for encoding the sub-blocks (8 bits) and number of sub-blocks k (8 bits).

Example. Figure 9 illustrates an example of a list L with two data blocks B_0 and B_1 . Thus there are two skip pointers (stored in the format explained in Section 4) in the metadata block. Within each data block, it is further partitioned into sub-blocks. For example, the block B_0 consists of two sub-blocks and the block B_1 contains three sub-blocks. For all the sub-blocks within B_0 , it uses the same b_0 bits to encode each element, which it originally requires b'_0 bits ($b'_0 \geq b_0$). For all the sub-blocks within B_1 , it instead uses b_1 bits to represent each element. Figure 9 also highlights the 16-bit global information for each block B_0 and B_1 .

Determining the optimal number of skip pointers. The next question is: *How many mini skip pointers to add for a data block?* We solve the problem by analyzing the relationship of the overall space overhead T_k with k in order to find the optimal k .

$$T_k = \max\{\lceil \log(a_{s-1} - a_0 + 1) \rceil, \lceil \log(a_{2s-1} - a_s + 1) \rceil, \dots, \underbrace{\lceil \log(a_{m-1} - a_{(k-1)s} + 1) \rceil}_{\text{sub-blocks}}\} \times (m - k) + \underbrace{\lceil \log(a_{m-1} - \beta + 1) \rceil \times k}_{\text{mini skip pointers}} + \underbrace{16}_{\text{global information}} \quad (4)$$

To find the optimal number k^* , we can enumerate all possible solutions to find which value leads to the minimal space overhead. Since we do not want a sub-block contain too few elements, say it should contain at least 4 elements. Then, we can search k from 2 to $m/4$. Thus,

$$k^* = \arg \min_{k=2}^{m/4} T_k \quad (5)$$

Time complexity. The time complexity of finding the optimal partitioning (off-line) is $\sum_{k=2}^{m/4} k = O(\frac{m^2}{32}) = O(800) = O(1)$ since $m \leq 160$ from Theorem 1.

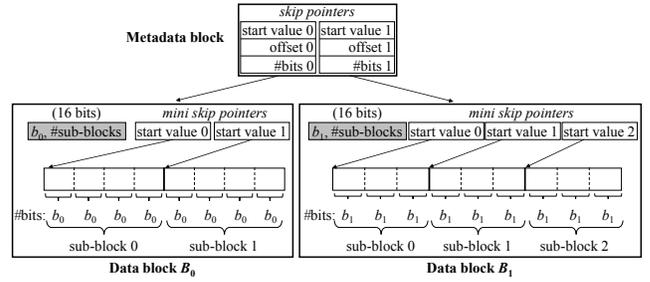


Figure 9: In-block compression

Supporting membership testing. It is a three-level structure where each level supports membership testing by using a revised key within a data block or a sub-block.

Discussion. We close this section by discussing two more questions: (1) Why not using dynamic partitioning to partition a block? (2) Can we further reduce the space overhead by partitioning a sub-block into sub-sub-blocks?

For the first question, it needs to maintain more skipping information to dynamically partition a data block into sub-blocks. The skipping information should at least contain: start value ($\lceil \log(a_{m-1} - \beta + 1) \rceil$ bits where β is the skip pointer of the block), number of elements (8 bits), offset (16 bits), and number of bits used to encode a sub-block (8 bits). Thus, a skip pointer needs $(\lceil \log(a_{m-1} - \beta + 1) \rceil + 32)$ bits, which is much higher since the current solution only needs $\lceil \log(a_{m-1} - \beta + 1) \rceil$ bits. On the other hand, it may not save too much space overhead in the sub-blocks because all data elements in a data block are very similar.

For the second question, it may not reduce the overall space anymore. That is because there is an extra space overhead associated with each split, i.e., 16 bits as highlighted in Figure 9. However, when there are few elements and each element uses very few bits, it is extremely difficult to save 16 bits anymore with further partitioning because in-block compression works if and only if $(m - k)(b' - b) \geq 16$. For example, suppose a block contains 8 elements: $\{10, 20, 30, 40, 50, 60, 70, 80\}$. Without partitioning, it takes $7 \times 8 = 56$ bits. With two partitions, i.e., 10 and 50 are promoted as mini skip pointers (taking $7 \times 2 = 14$ bits). The two sub-blocks become (after subtracting the skip pointer): $\{10, 20, 30\}$ and $\{10, 20, 30\}$. They take $3 \times 5 + 3 \times 5 = 30$ bits. Together with the mini skip pointers, the overall space overhead is $30 + 14 = 44$ bits, saving $56 - 44 = 12$ bits, which is less than 16 bits. Thus, we do not recommend further partitioning anymore.

7. CACHE-CONSCIOUS COMPRESSION

In this section, we further improve the layout of MILC such that it is more friendly to CPU cache lines for minimizing cache misses during membership testing.

What is cache-aware and why? Modern CPUs dedicate several layers of very fast caches (L1/L2/L3 cache) to alleviate the growing disparity between CPU clock speed and memory latency (a.k.a *memory wall*). Whenever a CPU instruction encounters a memory access, it first checks whether the accessed data resides in the caches. If yes, it accesses the data from the caches directly. Otherwise, a *cache line* (typically 64 bytes) of data is loaded from main memory to the caches. This will potentially evict other cache lines that are

in the caches. Thus, the goal of the cache-conscious design is to reduce cache misses by ensuring that a cache line brought from memory is fully utilized before it is being evicted.

Cache-aware design. We explain how to turn the compression structure presented in the previous section (Section 6) into a cache-aware structure. We classify the membership testing into two categories: within a metadata block (storing uncompressed skip pointers) and within a data block (storing compressed data).

For the membership testing within a metadata block, it is essentially the conventional binary search over an array. Previous studies have investigated it [16,27]. The main idea is to organize the elements into a B-tree structure with the node size being a CPU cache line (64 bytes). Note that the B-tree is materialized as an array using a level-order traversal manner without explicit storing any tree pointers, for saving space overhead. Thus, search can be efficiently executed by traversing the B-tree. However, there are two unique challenges in incorporating them into a fully functional compression structure: (1) The number of elements (i.e., skip pointers) may not form a perfect tree⁸ but most previous studies made such an assumption to save space overhead by not explicitly storing the tree pointers. We observe that only a collection of $17^h - 1$ elements can be converted to a h -level perfect tree. That is because a cache line contains $64/4 = 16$ elements (i.e., 17 children), then the total number of elements in a h -level perfect tree is:

$$\underbrace{16}_{\text{level 1}} + \underbrace{16 \times 17}_{\text{level 2}} + \underbrace{16 \times 17^2}_{\text{level 3}} + \dots + \underbrace{16 \times 17^{h-1}}_{\text{level } h} = 17^h - 1$$

However, there are many inverted lists and each inverted list has a different number of skip pointers that may not be $17^h - 1$. (2) Another unique challenge is how to find the corresponding data block after a skip pointer is located in the metadata block. This was not an issue for the non-cache-aware structure because the skip pointers and data blocks are stored in the same order. That is, a skip pointer and its data block have the same index number. However, if the skip pointers are stored in a cache-aware manner, the index numbers become completely different.

To solve the first challenge, we convert an array of sorted elements (i.e., skip pointers, non-cache-aware) to a complete tree instead of a perfect tree. A h -level complete tree [5] ensures that (1) only the last level is not full and all the elements in the last level are stored from left to right; (2) if the last level is removed, then it becomes a $(h - 1)$ -level perfect tree. We are not aware of any previous cache-aware designs having solved the problem, probably because they simply assumed the number of elements can form a perfect tree. But Schlegel et al. presented a solution in the SIMD area [29] that can be extended to cache-aware designs. The main idea is to determine for any element from the old non-cache-aware array the position in the new cache-aware array by developing a one-to-one mapping. Formally, let n be the number of elements (skip pointers), k be the number of elements that a cache line can accommodate ($k = 16$), H be the number of levels ($H = \lceil \log_k(n + 1) \rceil$), i be the element position in the old non-cache-aware array, and $g_n(i)$ be the

⁸A perfect tree (i.e., balanced and full) has three requirements [5]: (1) every node has precisely k entries where k is the fanout; (2) every intermediate node has exactly $k + 1$ children nodes; (3) every leaf node has the same depth.

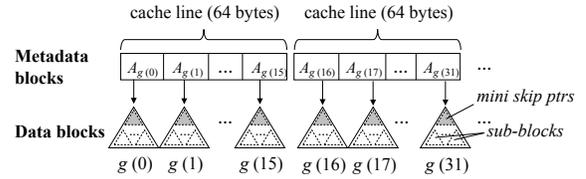


Figure 10: Cache-aware layout

position in the new cache-aware array. Then,

$$g_n(i) = \begin{cases} f_H(i) & \text{if } i \leq f_H^*(n) \\ f_{H-1}(i - o_H^*(n) - 1) & \text{otherwise} \end{cases} \quad (6)$$

We omit the explanations and proofs of these equations due to space constraints and refer interested readers to [29].

Next, we discuss how to tackle the second challenge. A simple solution is to compute a reverse mapping from the (new) position in the cache-aware array to the (old) position in the non-cache-aware array. However, that will take considerable time as the mapping has to be computed on the fly. MILC's solution is to change the storage of the data blocks such that they have the same order with their corresponding skip pointers. Figure 10 illustrates the design. The data blocks (as well as the skip pointers) are stored in the order of $g(0)$, $g(1)$, $g(2)$ and so on.

Next, we comment on the second type of membership testing that happens within a data block. It turns out that the existing design presented in the previous section is actually cache-aware. That is because each data block is organized as a two-level tree structure with the mini skip pointers being stored as the root node while the sub-blocks being stored as children nodes.

Supporting membership testing. The membership testing is executed efficiently by traversing an array of cache-aware skip pointers in the metadata block. Then it goes to the right data block to continue membership testing by using a revised key.

8. SIMD ACCELERATION

In this section, we discuss how to further improve the performance of MILC by leveraging the SIMD instructions.

What is SIMD-aware and why? A SIMD instruction operates on a s -bit register where s depends on different processors. Typically, s is 128, but more recent processors also support 256-bit or 512-bit SIMD operations. In this work, we use 128-bit SIMD instructions for a fair comparison with existing works [18]. The benefit of using SIMD instructions is to improve performance by processing multiple elements at a time.

Note that although compilers can automatically optimize some simple code with SIMD instructions, the optimizations are very limited, e.g., only for simple loops [14]. Thus, to fully exploit the SIMD instructions, we need to explicitly design a new storage structure that are amenable to SIMD.

SIMD-aware design. Recall in the previous discussions, the framework of MILC has two categories of blocks: metadata block and data blocks. The data elements in the metadata block are uncompressed while the data elements in the data blocks are compressed. Next, we discuss how to organize the data elements in both types of blocks into a SIMD-efficient structure.

For the elements (skip pointers) in the metadata block, they are stored as a contiguous sequence of cache lines. We store the data elements within the same cache line following the k-ary approach [29]. In this way, a SIMD operation can be applied to quickly find out which child node to visit [29].

For the elements in the data blocks, MILC basically follows the design in the previous section to support efficient membership testing while keeping a low space overhead. That is because a SIMD operation interacts with a s -bit SIMD register as a vector of *banks*, where a bank is a continuous section of b bits. For example, in SSE (Streaming SIMD Extensions) and AVX (Advanced Vector Extensions), b is 8 (byte type), 16 (short type), or 32 (int type). However, for inverted list compression, each element can be encoded in arbitrary number of bits b (not necessarily 8, 16, or 32 bits). Note that, there are other options if the space overhead is not a problem, e.g., by rounding up b to 8, 16, and 32.

Supporting membership testing. Given a search key, the first level of uncompressed data (skip pointers) will be accessed in a SIMD-aware manner following [29] to find out which block that potentially contains the key; then perform membership testing within the block.

9. EXPERIMENTS

In this section, we experimentally compare MILC against state-of-the-art inverted list compression schemes in terms of membership testing time and space overhead.

9.1 Experimental setting

Experimental platform. We conduct experiments on a commodity machine (Intel i7-4770 quad-core 3.40 GHz CPU, 64GB DRAM) with Ubuntu 14.04 installed. The CPU’s L1, L2, and L3 cache sizes are 32KB, 256KB, and 8MB. The CPU is based on Haswell microarchitecture which supports AVX2 instruction set. We use `mavx2` optimization flag for the SIMD optimization. We implement MILC in C++ and compile the code using GCC 4.4.7 with O3 enabled.

Datasets. In this work, we use the datasets from information retrieval, databases, and graph analytics.

(1) *Web data.* It is a collection of 41 million Web documents (around 300GB) crawled in 2012.⁹ It is a standard benchmark in the information retrieval community. We parse the documents and build inverted lists for each term. The query log contains 1000 real queries randomly selected from the TREC¹⁰ 2005 and 2006 (efficiency track).

(2) *DB data.* It is a star schema benchmark (SSB),¹¹ which includes one fact table (LINEORDER) and four dimension tables (CUSTOMER, SUPPLIER, PART, and DATE). We set the scale factor as 10 so the number of rows in the fact table is around 60 million. We use the query described in Section 2.2 for evaluation. Each list is allocated for a predicate on a logical huge table as described in Section 2.2. The list sizes are 11916634, 12028431, 9098421, and 11997098. Note that MILC is also applicable to other SSB queries. We use one single query for evaluation due to space limitation.

(3) *Graph data.* It is the twitter dataset crawled in 2009, which consists of 52,579,682 vertices and 1,963,263,821 edges. The data is widely used in graph analytics. Each list is

⁹<http://www.lemurproject.org/clueweb12.php>

¹⁰<http://trec.nist.gov/>

¹¹<http://www.cs.umb.edu/~poneil/StarSchemaB.PDF>

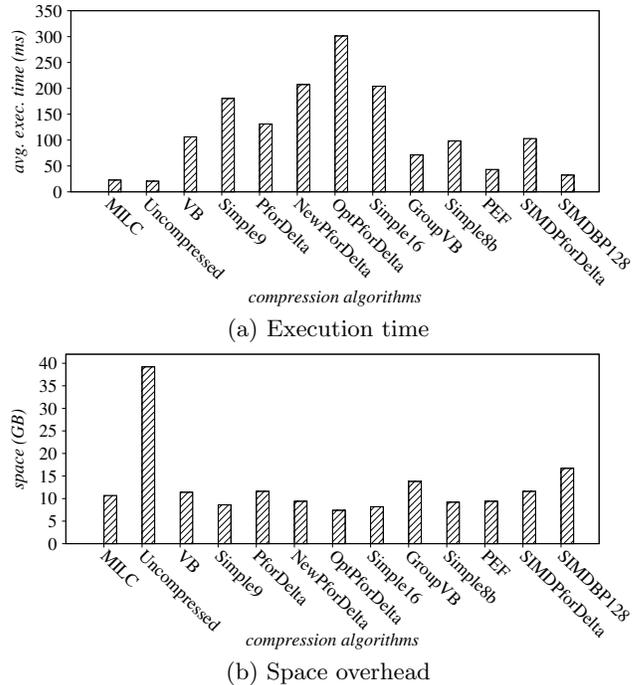


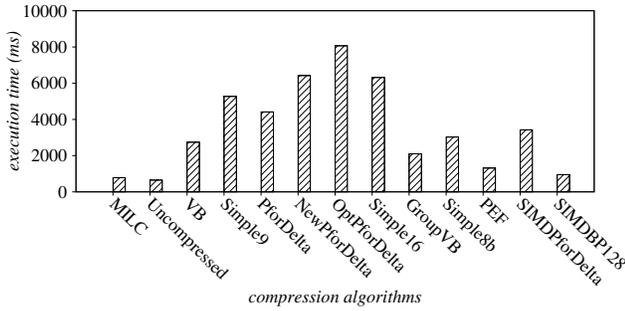
Figure 11: Comparing against existing compression approaches on Web data

an adjacency list of a vertex. We follow the methodology in [7] to evaluate the following query: “find out the common friends between a group of people”. Note that other queries could also be applied. The list sizes are 423640, 507777, 526292, and 779957, respectively.

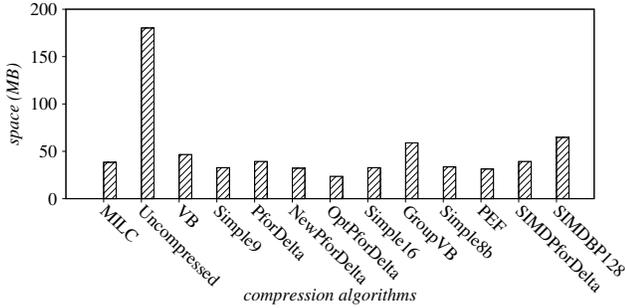
Competitors. We compare MILC with a wide range of recent compression approaches: Uncompressed, VB [32], PforDelta [43], OptPforDelta [41], NewPforDelta [41], Simple8b [2], Simple9 [1], Simple16 [40], GroupVB (a.k.a VarintGB) [9], SIMDBP128 [18], SIMDPforDelta [18], PEF [25]. We implement PEF from scratch and implement the other above-mentioned compression algorithms based on the source codes provided from [18]. For the uncompressed lists, we use conventional binary search. Note that we do not compare with some obviously low-performance encodings, such as Golomb, Rice, and Elias gamma. We also ignore those general purpose encoding schemes such as Snappy, LZ, LZ4, LZ0, or gzip, because they are much slower than PforDelta [18].

Evaluation methodology. In this work, we mainly use the following measurements for evaluation.

(1) *Execution time.* For each compression algorithm, we measure how fast it supports membership testing. In particular, we report the intersection time of each query since list intersection is essentially a series of membership testing operations to consistently find whether an element appears in a list. We use an SvS [6] that has been widely used in practice including Lucene. Assuming there are k lists L_1, L_2, \dots, L_k ($|L_1| \leq |L_2| \leq \dots |L_k|$) that are compressed. SvS decompresses the shortest list L_1 first. Then for each element $e \in L_1$, SvS checks whether e appears in L_2 (i.e., membership testing). Note that SvS does not need to decompress the entire L_2 due to skip pointers and it only needs to decompress a block of data that potentially contains e for membership testing. Then the results of L_1 and L_2 will be



(a) Execution time



(b) Space overhead

Figure 12: Comparing against existing compression approaches on DB data

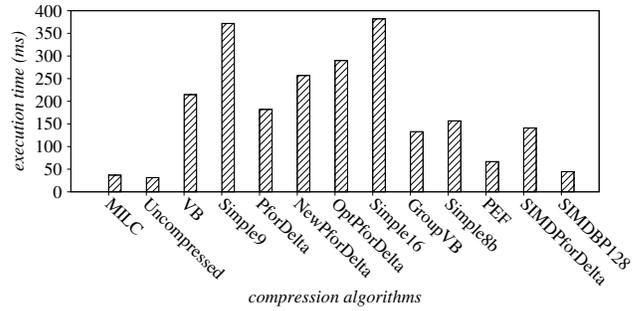
intersected with L_3 and the process continues until L_k .

(2) *Space overhead*. We also measure the space overhead that a compression algorithm takes.

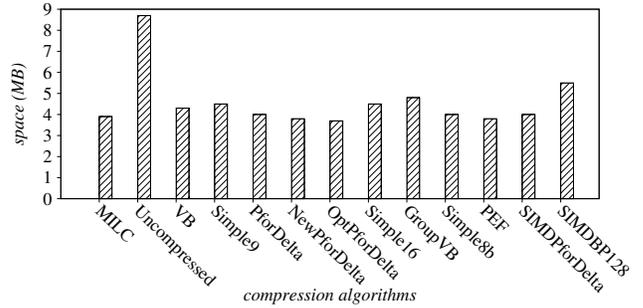
9.2 Comparing against existing compression approaches

In this experiment, we compare MILC with existing compression approaches on the two datasets. Note that MILC incorporates all the optimizations presented in Section 4, Section 5, Section 6, Section 7, and Section 8.

Figure 11 compares the average execution time and space overhead of on Web data. The execution time is measured by the average time (ms) of running those queries. Figure 11 shows that, (1) Compared with uncompressed lists, MILC achieves almost the same query performance but with a $3.7\times$ lower space overhead. The high performance is because MILC relies on fixed-bit encoding (instead of d-gaps as in most other compression algorithms) to support membership testing directly over compressed lists without decompressing even a whole block. MILC also relies on efficient architectural-aware data organizations such as cache-aware and SIMD-aware optimizations to achieve high query performance. The small space overhead is that MILC applies dynamic partitioning to store similar elements together. MILC also leverages the novel in-block compression technique to further reduce the space overhead. Figure 11 shows that, query processing on compressed lists can be (sometimes) executed as fast as that on uncompressed lists while keeping a low space overhead at the same time. (2) Compared with PforDelta, MILC is $5.75\times$ faster in execution time and 7.8% less in space overhead. The execution time saving is because PforDelta needs to decompress a whole block of data during membership testing because it is based on d-gaps but MILC does not need to do so. The space overhead saving is because



(a) Execution time



(b) Space overhead

Figure 13: Comparing against existing compression approaches on Graph data

PforDelta partitions a list statically while MILC partitions a list dynamically. And also, MILC applies in-block compression to further reduce the space overhead. (3) Compared with the variants of PforDelta, e.g., OptPforDelta and NewPforDelta, MILC has many advantages too. For example, MILC is $13.2\times$ faster (in execution time) than OptPforDelta while only incurring 44% more space. MILC is $9.1\times$ faster than NewPforDelta while only taking 13.8% more space. (4) Compared with Simple9 [1], Simple16 [40], and Simple8b [2], MILC is $7.9\times$, $9.0\times$, and $4.3\times$ faster but only consumes 24.4% , 30.4% , and 16.3% more space. (5) Compared with PEF [25], MILC is $1.9\times$ faster in execution time, while only consuming 13.8% more space. (5) Compared with the other compression algorithms, e.g., VB, GroupVB, SIMDBP128, and SIMDPforDelta, MILC runs $1.4\times$ to $4.6\times$ faster in query processing and takes 6.5% to 56.1% less space also.

Figure 12 shows the results of evaluating MILC on DB data. It shows that, (1) Compared with uncompressed lists, MILC needs similar execution time but only requires $3.7\times$ lower space overhead. (2) Compared with PforDelta, MILC is $5.65\times$ faster in execution time but MILC takes less space overhead. (3) Compared with OptPforDelta, MILC is $10.3\times$ faster in execution time but OptPforDelta only takes 38% less space overhead. (4) Compared with NewPforDelta, MILC is $8.2\times$ faster in execution time but NewPforDelta only incurs 16% less space overhead. (5) Compared with Simple9 [1], Simple16 [40], and Simple8b [2], MILC is $6.7\times$, $8.1\times$, and $3.9\times$ faster but only consumes 17.1% , 14.6% , and 14.3% more space. (6) Compared with PEF [25], MILC is $1.7\times$ faster in execution time, while only consuming 22.7% more space. (7) Compared with the other compression algorithms, e.g., VB, GroupVB, SIMDBP128, and SIMDPforDelta, MILC is $1.22\times$ to $3.5\times$ faster in query processing and takes 2.6% to 68.5% less space also.

Figure 13 shows the results of evaluating MILC on Graph data. We omit the descriptions of the results since they are largely similar to Figure 11 and Figure 12.

Overall, MILC represents the best tradeoff for inverted list compression especially in main memory databases compared among a spectrum of 12 existing compression algorithms.

10. CONCLUSION

In this work, we proposed a new compression approach MILC for encoding inverted lists in main memory. MILC is the first compression scheme that achieves similar query performance compared with uncompressed lists. Also, MILC is significantly faster than existing compression algorithms while keeping low space overhead. In the future, we plan to extend MILC to other storage devices including non-volatile main memory (NVMM) [39], SSDs [35,36], and HDDs. Another direction is to tailor MILC for supporting in-storage computing [37,38]. Besides that, it is also interesting to extend MILC to support more operations and queries, e.g., intersection, union, top-k query processing.

11. REFERENCES

- [1] V. N. Anh and A. Moffat. Inverted index compression using word-aligned binary codes. *IR*, 8(1):151–166, 2005.
- [2] V. N. Anh and A. Moffat. Index compression using 64-bit words. *SPE*, 40(2):131–147, 2010.
- [3] T. A. Björklund, N. Grimsmo, J. Gehrke, and O. Torbjørnsen. Inverted indexes vs. bitmap indexes in decision support systems. In *CIKM*, pages 1509–1512, 2009.
- [4] N. Bruno, N. Koudas, and D. Srivastava. Holistic twig joins: Optimal xml pattern matching. In *SIGMOD*, pages 310–321, 2002.
- [5] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, 3rd edition, 2009.
- [6] J. S. Culpepper and A. Moffat. Efficient set intersection for inverted indexing. *TOIS*, 29(1):1–25, 2010.
- [7] M. Curtiss, I. Becker, T. Bosman, S. Doroshenko, L. Grijincu, T. Jackson, S. Kunnatur, S. Lassen, P. Pronin, S. Sankar, G. Shen, G. Woss, C. Yang, and N. Zhang. Unicorn: A system for searching the social graph. *PVLDB*, 6(11):1150–1161, 2013.
- [8] D. R. Cutting and J. O. Pedersen. Optimizations for dynamic inverted index maintenance. In *SIGIR*, pages 405–411, 1990.
- [9] J. Dean. Challenges in building large-scale information retrieval systems: Invited talk. In *WSDM*, page 1, 2009.
- [10] B. Debnath, S. Sengupta, and J. Li. Skimpystash: Ram space skimpy key-value store on flash-based storage. In *SIGMOD*, pages 25–36, 2011.
- [11] P. Elias. Efficient storage and retrieval by content and address of static files. *JACM*, 21(2):246–260, 1974.
- [12] P. Elias. Universal codeword sets and representations of the integers. *TOIT*, 21(2):194–203, 1975.
- [13] J. Goldstein, R. Ramakrishnan, and U. Shaft. Compressing relations and indexes. In *ICDE*, pages 370–379, 1998.
- [14] Intel Corporation. *Intel® 64 and IA-32 Architectures Optimization Reference Manual*. 2012.
- [15] Y. Jiang, G. Li, J. Feng, and W.-S. Li. String similarity joins: An experimental evaluation. *PVLDB*, 7(8):625–636, 2014.
- [16] C. Kim, J. Chhugani, N. Satish, E. Sedlar, A. D. Nguyen, T. Kaldewey, V. W. Lee, S. A. Brandt, and P. Dubey. Fast: Fast architecture sensitive tree search on modern cpus and gpus. In *SIGMOD*, pages 339–350, 2010.
- [17] A. Lamb, M. Fuller, R. Varadarajan, N. Tran, B. Vandiver, L. Doshi, and C. Bear. The vertica analytic database: C-store 7 years later. *PVLDB*, 5(12):1790–1801, 2012.
- [18] D. Lemire and L. Boytsov. Decoding billions of integers per second through vectorization. *SPE*, 45(1):1–29, 2015.
- [19] X. Li, J. Han, and H. Gonzalez. High-dimensional olap: A minimal cubing approach. In *VLDB*, pages 528–539, 2004.
- [20] C. Lin, J. Wang, and Y. Papakonstantinou. Data compression for analytics over large-scale in-memory column databases (summary paper). *CoRR*, abs/1606.09315, 2016.
- [21] E. Lo, B. Kao, W.-S. Ho, S. D. Lee, C. K. Chui, and D. W. Cheung. Olap on sequence data. In *SIGMOD*, pages 649–660, 2008.
- [22] T. C. Lowe. Design principles for an on-line information retrieval system. Technical report, University of Pennsylvania, 1966.
- [23] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- [24] A. Moffat and J. Zobel. Self-indexing inverted files for fast text retrieval. *TOIS*, 14(4):349–379, 1996.
- [25] G. Ottaviano and R. Venturini. Partitioned elias-fano indexes. In *SIGIR*, pages 273–282, 2014.
- [26] V. Raman, L. Qiao, W. Han, I. Narang, Y.-L. Chen, K.-H. Yang, and F.-L. Ling. Lazy, adaptive rid-list intersection, and its application to index anding. In *SIGMOD*, pages 773–784, 2007.
- [27] J. Rao and K. A. Ross. Cache conscious indexing for decision-support in main memory. In *VLDB*, pages 78–89, 1999.
- [28] R. Rice and J. Plaunt. Adaptive variable-length coding for efficient compression of spacecraft television data. *TOCT*, 19(6):889–897, 1971.
- [29] B. Schlegel, R. Gemulla, and W. Lehner. K-ary search on modern processors. In *DaMoN*, pages 52–60, 2009.
- [30] F. Silvestri and R. Venturini. Vsencoding: Efficient coding and fast decoding of integer lists via dynamic programming. In *CIKM*, pages 1219–1228, 2010.
- [31] D. Ślęzak, J. Wróblewski, V. Eastwood, and P. Synak. Brighthouse: An analytic data warehouse for ad-hoc queries. *PVLDB*, 1(2):1337–1345, 2008.
- [32] L. Thiel and H. Heaps. Program design for retrospective searches on large data bases. *IPM*, 8(1):1 – 20, 1972.
- [33] V. Venkataramani, Z. Amsden, N. Bronson, G. Cabrera III, P. Chakka, P. Dimov, H. Ding, J. Ferris, A. Giardullo, J. Hoon, S. Kulkarni, N. Lawrence, M. Marchukov, D. Petrov, and L. Puzar. Tao: How facebook serves the social graph. In *SIGMOD*, pages 791–792, 2012.
- [34] S. Vigna. Quasi-succinct indices. In *WSDM*, pages 83–92, 2013.
- [35] J. Wang, E. Lo, M. L. Yiu, J. Tong, G. Wang, and X. Liu. The impact of solid state drive on search engine cache management. In *SIGIR*, pages 693–702, 2013.
- [36] J. Wang, E. Lo, M. L. Yiu, J. Tong, G. Wang, and X. Liu. Cache design of ssd-based search engine architectures: An experimental study. *TOIS*, 32(4):1–26, 2014.
- [37] J. Wang, D. Park, Y.-S. Kee, Y. Papakonstantinou, and S. Swanson. Ssd in-storage computing for list intersection. In *DaMoN*, pages 1–7, 2016.
- [38] J. Wang, D. Park, Y. Papakonstantinou, and S. Swanson. Ssd in-storage computing for search engines. *TC*, 2016.
- [39] J. Xu and S. Swanson. NOVA: A log-structured file system for hybrid volatile/non-volatile main memories. In *FAST*, pages 323–338, 2016.
- [40] H. Yan, S. Ding, and T. Suel. Inverted index compression and query processing with optimized document ordering. In *WWW*, pages 401–410, 2009.
- [41] J. Zhang, X. Long, and T. Suel. Performance of compressed inverted list caching in search engines. In *WWW*, pages 387–396, 2008.
- [42] Z. Zhang, J. Tong, H. Huang, J. Liang, T. Li, R. J. Stones, G. Wang, and X. Liu. Leveraging context-free grammar for efficient inverted index compression. In *SIGIR*, pages 275–284, 2016.
- [43] M. Zukowski, S. Heman, N. Nes, and P. Boncz. Super-scalar ram-cpu cache compression. In *ICDE*, 2006.