

THE UNIVERSITY OF CHICAGO

FAST AND EFFECTIVE COMPRESSION FOR IOT SYSTEMS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF COMPUTER SCIENCE

BY

CHUNWEI LIU

CHICAGO, ILLINOIS

DECEMBER 2022

Copyright © 2022 by Chunwei Liu
All Rights Reserved

Table of Contents

LIST OF FIGURES	v
LIST OF TABLES	viii
ACKNOWLEDGMENTS	ix
ABSTRACT	x
1 THESIS STATEMENT	1
2 INTRODUCTION	2
3 MOSTLY ORDER PRESERVING DICTIONARIES	6
3.1 Background	9
3.1.1 Dictionary Encoding	9
3.1.2 Consensus Protocol	12
3.2 MOP Overview	13
3.2.1 MOP Definitions	13
3.2.2 Cascade-MOP	15
3.2.3 Dictionary Generation Assumptions	15
3.3 Methodology	16
3.3.1 Initialization	16
3.3.2 MOP Encoding	18
3.3.3 Finalization	20
3.4 MOP Queries	20
3.4.1 Equality Operator	20
3.4.2 Range Operator	21
3.4.3 Sort Operator	23
3.5 Experimental Setup	25
3.6 Evaluation	26
3.6.1 Range Filtering Evaluation	26
3.6.2 SORT Evaluation	31
3.6.3 MOP Generation	31
3.6.4 Overall Compression Performance	45
3.7 Conclusion	46
4 DECOMPOSED BOUNDED FLOATS FOR FAST COMPRESSION AND QUERIES	48
4.1 Background and Related Work	51
4.1.1 Numeric Data Representation	52
4.1.2 Compression for Decimal Numbers	54
4.1.3 Query-friendly Storage Layout	57
4.2 BUFF Overview	58
4.2.1 Bounded Float	58

4.2.2	Float Splitting and Compression	61
4.2.3	Handling Outliers	63
4.2.4	Query Execution	65
4.3	Evaluation	72
4.3.1	Datasets	72
4.3.2	Optimized Baselines	73
4.3.3	Performance Overview	75
4.3.4	Compression Performance	76
4.3.5	Query Performance	79
4.3.6	Benchmark Evaluation	86
4.4	Conclusions	87
5	ADAEDGE: A DYNAMIC COMPRESSION SELECTION FRAMEWORK FOR RESOURCE CONSTRAINED DEVICES	88
5.1	Background	91
5.1.1	Lossless and Lossy Compression	91
5.1.2	Compression selection	93
5.1.3	Multi-Armed Bandit Problem	96
5.2	AdaEdge Overview	98
5.2.1	AdaEdge Framework Constraints	98
5.2.2	AdaEdge Workflow	102
5.2.3	Optimization target	107
5.2.4	Recoding optimization	108
5.2.5	LRU-based compression	109
5.3	Evaluation	110
5.3.1	Machine learning accuracy with given model	111
5.3.2	Adaptive Compression Selection	112
5.4	Conclusion	121
6	CONCLUSION	122
	REFERENCES	124

List of Figures

2.1	Various categories of IoT sensors	3
3.1	Tension between three properties of OP dictionary	7
3.2	Key space for MOP dictionary	14
3.3	Key space for C-MOP dictionary	15
3.4	MOP generation example	17
3.5	Steps for range filtering with C-MOP	22
3.6	Sort with MOP	24
3.7	Parquet Columnar Store Format	26
3.8	Range filter overview on taxi dataset	27
3.9	Percent of records encoded by ordered keys.	27
3.10	Type-2 filter on taxi dataset	27
3.11	Range query on TPC-H dataset	29
3.12	TPC-H type 2 range filter with selectivity 0.00001	29
3.13	TPC-H Type-2 range filter with selectivity 0.27869	29
3.14	C-MOP range filter on taxi dataset with increasing the number of cascading levels.	30
3.15	SORT Operator Runtime	32
3.16	Average generation time.	33
3.17	Evaluating the effect of lookahead on MOP generation performance (Runtime and Ordered Percentage).	34
3.18	Evaluating the effect of pitch on MOP ordered percentages using uniform, zipf, and taxi workloads.	35
3.19	Worker configurations	36
3.20	Evaluating the effect of coordinator batch size on MOP generation runtime using uniform and zipf workloads. 24 workers are being used with worker batch size set to 1	37
3.21	Evaluating the effect of MOP layer ratio on MOP ordered percentages using uniform, zipf, and taxi workloads.	38
3.22	Evaluating the effect of dynamic MOP layer ratios when correcting for data distribution changes.	39
3.23	Evaluating the space saving benefit of C-MOP with regard to padding ratio.	42
3.24	Evaluating the benefit of C-MOPs with regard to padding ratio when using uniform and histogram insertion techniques for the second layer.	43
3.25	Evaluating the impact of sorted data and sampling strategy on a MOP's ordered ratio.	44
3.26	Encoded Column Size and Dict Padding	45
3.27	Scanning and decoding 180 million values.	46
4.1	Many datasets only span a fixed range of values with limited precision, which is a small subset of broad float number range and precision.	49
4.2	Compression examples for different approaches.	52
4.3	Extracting range bounded integer part and precision bounded fractional part for float number	61

4.4	Byte-oriented columnar layout	62
4.5	Sparse encoding condenses sub-column.	64
4.6	Datasets have bounded range and precision	72
4.7	Compression performance overview (greater is better, CR: compression ratio, Cthr: compression throughput). BUFF outperforms in query and compression throughput.	76
4.8	Compression ratio shows the BUFF are comparable to the best and always better than vanilla Gorilla (Sprintz and Scaled-slice fails on PMU).	77
4.9	Compression throughput shows BUFF performs best in most cases with user given range stats.	79
4.10	Low selective equality filter	79
4.11	Low selective range filter	79
4.12	High selective equal filter	79
4.13	Max throughput	82
4.14	Sum throughput	83
4.15	Materialization throughput	84
4.16	Materialization and aggregation with targeted precision boosts throughput by only reading the leading bytes.	85
4.17	Relative 1NN accuracy on 128 UCR datasets shows reduced precision has a limited impact on accuracy.	86
4.18	TSBS query runtime for float attributes.	87
5.1	AdaEdge overview	99
5.2	compression should be able to handle the given signal generation rate. An example shows a signal with 4 million data points generated per second.	100
5.3	The compressed file size should be within the network transmission capacity. Bars show the compression egress rate on a 4GHz signal with double type. Lines show the network transmission capacity per second (MB/s).	101
5.4	AdaEdge keeps evolving the ingested data with the given storage budget for offline mode devices.	102
5.5	AdaEdge keeps ingesting the latest data while applies different levels of compression on the early segments to free space for incoming data. The red rectangle denotes an arbitrary segment in AdaEdge, with the rectangle length showing its current size.	105
5.6	Decision tree model accuracy on UCI	112
5.7	Random forest model accuracy on UCI	112
5.8	KNN model accuracy on UCI	112
5.9	KMeans model accuracy on UCI	113
5.10	Decision tree model accuracy on UCR	113
5.11	Random forest model accuracy on UCR	113
5.12	KNN model accuracy on UCR	113
5.13	KMeans model accuracy on UCR	113
5.14	Machine learning model accuracy on target compression ratio	115
5.15	Sum aggregation query over target compression ratio	116
5.16	Max aggregation query over target compression ratio	116

5.17	Sum aggregation query and random forest complex optimization target over target compression ratio with weighted value $w_1 = 0.625$ and $w_2 = 0.375$	117
5.18	Compression speed and random forest complex optimization target over target compression ratio with weights $w_1 = 0.524$ and $w_2 = 0.476$	118
5.19	KMeans accuracy loss over ingestion time	120

List of Tables

3.1	Evaluating cardinality estimation techniques.	41
4.1	Bounded float still keeps decimal precision	59
4.2	Number of bits needed for targeted precision	60

ACKNOWLEDGMENTS

It has been a wonderful time at the University of Chicago. I would not have reached this point without the incredible environment of creativity, kindness, and enthusiasm created by my mentors, colleagues, friends, and family.

First and foremost, I would like to thank my advisor, Prof. Aaron J. Elmore, for being a great mentor in research and life. His professional expertise and insights helped me dive into cutting-edge research topics; his invaluable suggestions guided me through challenges in research projects; and his passion for research and life profoundly fostered my critical thinking, open mindset, collaborative attitude and communication. He is always willing to provide his best effort and support during every phase of my Ph.D. study. Without his continued support and patience in this challenging journey, I would not reach the destination.

I also wish to thank Prof. Sanjay Krishnan and Prof. Raul Castro Fernandez for serving on my doctoral dissertation committee, and Prof. Michael J. Franklin for being on my master's committee. Their invaluable suggestions and insights helped shape the research direction and improve my thesis. I am very fortunate to have them as mentors and role models during graduate school.

I would also like to thank my collaborators, Dr. John Paparrizos and Dr. Hao Jiang, for their enormous help and contribution to my research over the years. We had numerous enlightening discussions on my research projects. Their hard work and valuable suggestions are the key components to my achievement. Working with them is an enjoyable and rewarding experience.

Finally, I would like to thank my family for their unwavering and unconditional support. I thank my parents Wenzhong Liu and Hongbing Yue, and my grandparents Xinglin Liu and Yulan Ma, who have always been encouraging me.

ABSTRACT

The Internet of Things (IoT) enables connections of trillions of sensors and data collection for connectivity and analytics. The amount of IoT-generated data has exploded due to the rapid growth of interconnected IoT devices from a wide range of IoT applications. IoT systems need an effective solution to ingest, store, and analyze the exploding scale of IoT data to make the best of the limited resources in IoT systems, such as network, storage, energy, and computation power. Given that IoT data is endless, heterogeneous, and dynamic, the thesis proposes new compression techniques to handle those special data features. MOP handles the endlessness of IoT data by pre-allocating encoding space for the dynamic incoming data and enables in-situ queries in the encoded domain for fast query execution. BUFF addresses the heterogeneity of IoT data by applying decomposed but compact encoding space for the IoT numeric data with different statistics and achieves efficient query execution support according to the host's hardware. In addition, IoT's dynamic data imposes challenges to traditional compression selection strategies as there is no one-size-fits-all solution for IoT systems with varying data statistics, complex workloads, and constrained hardware resources. We introduce AdaEdge as a hardware-conscious encoding selection framework for resource-constrained devices.

CHAPTER 1

THESIS STATEMENT

Modern applications and systems are generating vast amounts of data every day. It is expected that the amount of data generated in the future will continue to grow and largely be driven by the Internet of Things (IoT) devices. Systems for storing and analyzing IoT data need an effective solution to ingest, store and analyze the exploding IoT data to make the best of the limited resources. To solve this problem, we propose new compression techniques that compress the endless and heterogeneous IoT signal data quickly and effectively while supporting efficient and accurate query execution according to the hardware specifics. These compression techniques, together with resource-aware compression selection strategy, can adaptively select the optimal compression approach to meet the system resource constraints and workload requirements. The thesis divides the problem into the following research tasks to achieve the goal.

- **Endless streaming data** Compression for IoT data should handle the endlessness of the signal streaming data and adapt to any data distribution change. The compression should provide efficient in-situ query execution in the encoded space.
- **Heterogeneous numeric data** An IoT-focused compression approach has to handle the heterogeneity of IoT data, including the most popular numeric data type with varying data statistics such as distribution, range, precision. The compression should also provide efficient query execution on simple and advanced hardware.
- **Compression selection for dynamic data** The compression selection strategy for IoT systems should be aware of changing data statistics, workload, and limited resources, including network bandwidth connectivity, local storage and computation power.

CHAPTER 2

INTRODUCTION

Internet of Things (IoT) systems generate massive amounts of heterogeneous data from diverse sensors every second. Interconnected smart IoT devices, including consumer electronics such as smartphones, smart home devices, autonomous vehicles, wearables, and connected healthcare, are widely used by IoT applications. Collecting a large amount of sensor-generated data series is a critical first step before any downstream analytics are performed. Recently, the International Data Corporation predicted that the global amount of data will reach 175 *ZB* by 2025 [94]. Billions of IoT devices will be a significant driver of this growth, which contributes more than 50% of the data volume. To keep up with the storage demands stemming from all this data creation, new effective data storage layout and compression techniques are necessary, instead of relying on increasing capacity supplied by the storage industry or costly public cloud resources.

Many IoT systems use the cloud as a data storage and query processing backend. Depending on the IoT system design, the sensor sends data to its hub node - an edge device or a router. The data is ingested and assembled into segments before sending to the cloud. Then the data processing steps, such as compression, indexing, archiving, and query, are typically processed on the cloud side. However, the cloud backend IoT framework has the following problem: data transmission to the cloud is costly, especially for use cases with unstable or poor network connections. Additionally, all raw data is uploaded to the cloud, including useless data which may never be touched in the future, which wastes network and storage resources.

As edge devices are getting more computational power and edge computing is becoming more popular, pushing down compression, query, and analysis to the edge is an appealing option. It is beneficial to many system by minimizing the data size needed to send, saving network bandwidth, saving the edge and cloud resources for storing data, and speeding up the query with the right encoding techniques. Such pushing down also allows query processing on




Figure 2.1: Various categories of IoT sensors

the edge directly, which could significantly decrease query latency and make the system more robust to the network disconnection. All these benefits contribute to the system’s overall cost saving on storage, network, and computation. However, the edge is usually resource-constrained compared with cloud capabilities in terms of network bandwidth, storage, and computation power. The limited resource brings challenges to compression design and query execution, which are not well addressed by traditional compression approaches. Therefore, more dedicated compression solutions are necessary for IoT systems.

The compression for IoT data should consider three unique features of real-time IoT data: IoT data is **heterogeneous**, **endless** and **dynamic**. IoT system generates heterogeneous data because of the variety of data acquisition devices [70], as is shown in Figure 2.1. The **heterogeneous** data here does not only mean the data with many different data types but also the data with varying data statistics, including range, precision, cardinality, and distributions. IoT data typically differs from this historical dataset. The traditional analysis

over static and fixed data that can provide whole data statistics for the data processing and compression task. In contrast, the IoT system generates **endless** streaming, and its statistics change over time **dynamically**.

To compress IoT data efficiently, we need to build dedicated compression approaches for IoT data and a compression selection framework to handle dynamic data features and different system constraints. This thesis proposes MOP [74] to address the endless IoT data by providing efficient single pass, best effort order-preserving dictionary encoding for streaming data. The thesis also introduces BUFF [73] to deal with the heterogeneous data from IoT systems by enabling range and precision aware decomposed float compression for IoT numeric data. With an extensive evaluation of many compression for IoT data, we find there is no one-size-fits-all solution for IoT data compression. An adaptive compression selection strategy is needed to choose the best compression fit for the incoming IoT data based on the data statistics and workload. This thesis proposes AdaEdge as a dynamic compression selection framework for resource constrained devices. We here give each system a high-level overview as follows.

MOP: Dictionary encoding, or domain encoding, is an important form of compression that uses a bijective mapping to replace attributes from a large domain (i.e. strings) with a finite domain (i.e. 32 bit integers). This encoding both reduces data storage and allows for more efficient query execution. Traditional dictionary encoding only supports efficient equality queries, while range queries require that encoded values are decoded for evaluating the predicates. An order-preserving dictionary allows for range queries without decoding by ensuring that encoded keys follow the same order as the values in the dictionary. While this approach enables efficient queries, it requires that the full set of values is known to create the mappings. In this work, we bridge this gap by introducing mostly ordered dictionaries that use a best-effort dictionary generation based on sampling the input dataset. Query evaluation on a mostly ordered dictionary avoids decoding when possible and gracefully degrades performance as the ratio of ordered values decreases.

BUFF: Modern data-intensive applications often generate large amounts of low precision float data with a limited range of values. Despite the prevalence of such data, there is a lack of an effective solution to ingest, store, and analyze bounded, low-precision, numeric data. To address this gap, we propose BUFF, a new compression technique that uses a decomposed columnar storage and encoding methods to provide effective compression, fast ingestion, and high-speed in-situ adaptive query operators with SIMD support. Compared to the state-of-the-art float compression techniques, BUFF achieves up to $50\times$ speedup for filtering and aggregations, and $1.5\times$ improvement for ingestion with a single thread, while offering comparable compression sizes to the state-of-the-art approaches.

AdaEdge: Different compression approaches may have varying accuracy impacts and efficiency on different workloads. There is no one-size-fits-all compression solution for IoT systems with varying data statistics, workloads, and hardware constraints. The thesis introduces AdaEdge as a hardware-conscious encoding selection framework for resource-constrained devices. AdaEdge is a best-effort compression selection framework to preserve application-critical information as much as possible within the system constraints. AdaEdge uses a LRU based compression policy to compress the data based on its staleness and importance. AdaEdge applies multi-armed bandit algorithm to assist the compression selection, optimizing workload targets such as compression ratio, compression throughput, workload accuracy, or their weighted combinations. It provides lossy and lossless compression selection depending on the hardware constraints. AdaEdge supports an online mode to handle the network constraints for edge node, and an offline mode with data evolving policy to conserve workload-specific information.

In the rest of this dissertation, we first present the aforementioned mostly dictionary encoding designed for endless streaming data in Chapter 3, and range and precision aware decomposed float compression in 4. We introduce AdaEdge as a hardware-conscious encoding selection framework for resource-constrained devices in Chapter 5. The thesis concludes in Chapter 6.

CHAPTER 3

MOSTLY ORDER PRESERVING DICTIONARIES

Database compression is critical for current data-intensive systems where the rate of data growth is outstripping growth of processing speeds, memory capacity, and I/O bandwidth. Columnar databases enable efficient compression in two primary ways. First, the ability to organize data by attributes reduces entropy and thereby improves compression effectiveness. Second, through use of columnar encoding [9] (i.e. run-length encodings, delta encoding, and dictionary encoding) the database supports efficient in-situ query processing, whereas prior use of byte-oriented compression (i.e. `gzip`, `snappy`) requires decompression as a blocking step before query execution [48], which can be CPU-intensive [40].

While all columnar encodings offer compression benefits over unencoded data, query benefits can be limited given the value domain and data distribution [9]. For non-numeric data types, such as dates and strings, dictionary encoding can translate a large and near infinite domain to a smaller finite and dense domain, often integers [98]. Dictionary encoding, or domain encoding, creates a bijective mapping from an arbitrary infinite source domain (*value domain*) to a fixed size target domain (*code domain*). The translated codes as well as the mapping are stored as the encoded result. Since the domain is smaller, it allows values further compression using methods such as bit-packed encoding (e.g. truncating unnecessary bits, such as using 3 bits to represent integers 0-7) [98, 56], as well as supporting fast query execution via efficient hardware instructions [56]. As a result, dictionary encoding is widely supported in many analytic platforms systems [102, 18, 98, 1, 76, 49], and offers significant benefits for string data types that are common in many domains, such as enterprise data [41] and open data initiatives [56]. Note that we limit our focus to global dictionaries that maintain a single mapping for an entire column, instead of a local per-block dictionary.

For query filtering with equality predicates, dictionary encoding allows the query to translate the predicate value(s) to the code domain and evaluate the predicates directly on the encoded data. Prior work studies variations of dictionary encoding algorithms on their

trade-off between compression ratios and decoding speed [79]. However, for query range predicates, the system must decode encoded values to evaluate the predicates. To avoid this expensive translation, the use of order-preserving dictionaries can allow range predicates to be evaluated without decoding [20]. Here, order-preserving dictionaries work by ensuring that encoded keys maintain the same order as values (e.g. for mappings $k_1 = v_1$ and $k_2 = v_2$, $k_1 < k_2$ iff $v_1 < v_2$). Note that we refer to order-preserving dictionaries as *OP* for short.

Generating the mapping to support OP dictionaries comes at a cost. If the code domain is a finite domain, such as 32-bit integers, then the entire set of values must be known and fixed before the dictionary encoding can occur, as preserving the order requires sorting the values first [20]. If the code domain is an infinite domain, such as double values, an order preserving scheme can be generated for a set of unknown values using a Dewey Decimal style coding, where new values can be inserted using increasing precision [104]. While this approach works when the domain is unknown or not-fixed, it suffers from reduced compression (e.g. larger code domain values and no bit-packing), less efficient CPU operations, and lacks the ability for dense SIMD operations that depend on small key values [56].




Figure 3.1: Tension between three properties of OP dictionary

Therefore, in Figure 3.1, we propose a conjecture for dictionary encodings that it is impossible to simultaneously provide more than two out of the three properties: the dictionary

is order preserving (i.e. OP), the encoded keys have a finite domain (i.e. integer codes), and if the encoding tolerates an unknown value domain (i.e. unknown values). For example, if we support an unknown value domain and OP, the integer code domain property can be violated when there is no more “code space” for new values, and a more complex code domain must be used to keep the OP property. Otherwise, if we require OP and integer codes, recoding (e.g. redoing the encoding with a known domain) is needed when there is no valid code space to maintain ordering.

While many applications know the domain of values to encode a priori, several important scenarios exist where the values are unknown when the encoding occurs. First, for data that is continuously being loaded into a data warehouse or lake, the entire set of values may not ever be known or known before analysis must begin – either of which could prevent use of an OP dictionary if integer codes are desired. Second, for massive static datasets, generating an OP dictionary requires taking a full pass on the dataset before encoding to learn the full value set. Given the potential costs of scanning and parsing the dataset, this extra pass may be prohibitive for scenarios that want to begin analysis quickly.

Observe that if a dictionary encoding supports unknown values and an integer code domain, we may achieve the OP goal by pre-allocating a huge code space for a dictionary, but at the cost of larger bit representation for each value. Despite this, one cannot guarantee there will be no code space conflicts in the future, especially in the presence of skewed data distributions. This inability to guarantee ordering though, leads to the key research questions of this paper which are *can a database system leverage a dictionary encoding that is partially ordered and given the ability to sample a dataset, what is the ideal pre-allocation key allocation?*

In this paper, we present a best-effort order preserving dictionary encoding: a mostly order preserving dictionary encoding (*MOP*). It pre-allocates a code space for an order preserving dictionary based on estimated statistics, such as record count and cardinality. A backup code space is reserved to handle potential code space conflicts, which we explore

both with unordered and cascading ordered variations. For distributed generation, MOP leverages a leader protocol to synchronize and update local dictionary views of its writers. For a MOP that has 90% of the keys ordered, we are able to reduce query filtering latency by up to 47% compared to decoding a standard dictionary, and 9% slower than a order preserving dictionary according to our experiments. In addition we propose a nested MOP, *Cascade-MOP* or *C-MOP*, that minimizes the amount of disordered data at the cost of more complex query evaluation.

To evaluate the effectiveness of MOP, we implement a prototype within the open-source columnar format Parquet [18]. With this prototype we consider how to construct a MOP without knowing the value domain a priori and demonstrate how to leverage MOP to accelerate range queries and sort operators. The rest of the paper is organized as follows. Section 4.1 discusses related work. Section 4.2 overviews MOP and key definitions. Sections 3.3 and 3.4.2 cover MOP generation and query execution respectively. Section 3.4.3 evaluates MOP for both range filtering and generation. The paper concludes in Section 3.7.

3.1 Background

While there is a large body of research on compression and encoding for columnar databases, we mainly focus on dictionary encoding and its query optimizing on encoded data. In this section, we review previous efforts on this topic. We first introduce dictionary encoding and review the state-of-the-art, then briefly discuss consensus protocols which our dictionary encoding is based on in a distributed environment.

3.1.1 Dictionary Encoding

Dictionary encoding creates a bijective mapping between values of variable length to compact integer codes, and replaces the original data entries with corresponding codes. The dictionary used in the encoding process is then attached to the encoded data, as part of the encoding

output. Dictionary encodings are information-lossless as described by Lempel and Ziv, and such encodings can achieve the theoretical lower bound of compression ratio defined by Shannon entropy [67, 115]. A dictionary can be global or local. For global dictionaries, a single dictionary is used for the entire column; for local dictionaries multiple dictionaries exist, and each are valid only for some partition of the data (i.e. block, page, or row-group). Dictionaries may be random such that keys are assigned with no order, or may be order-preserving such that keys preserve an order of the values (e.g. for mappings $k_1 = v_1$ and $k_2 = v_2$, $k_1 < k_2$ iff $v_1 < v_2$). Note that while order-preserving dictionaries can be local or global, they primarily are used for global dictionaries for query performance benefits [20].

Compression techniques help reduce I/O operations and data storage, and therefore are prevalent for analytic systems. Many byte-oriented or block-oriented compression techniques require data to be decompressed before querying [54]. These compression techniques, such as GZip and Lempel-Ziv [116], typically require expensive CPU cycles. Columnar encoding schemes, such as run-length encoding or dictionary encoding, trade-off CPU overhead for larger storage and the ability to directly query on the encoded data [9].

Research has explored optimizing dictionary encoding for database systems. Chen et al. [30] proposes a hierarchical dictionary encoding scheme for string attributes, supporting encoding at different granularities (attribute level, word level, prefix level, etc.). Paradies et al. [88] demonstrates an entropy-based approach that is adaptive to user query patterns. Column-oriented databases, such as MonetDB [118] and C-Store [102], use dictionary encoding to allow arbitrary types to be converted to consecutive integer codes; this enabling other encoding schemes, such as bit-packing (i.e. truncating unnecessary bits), run-length encoding, and delta encoding to be applied, further reducing storage size.

In addition to the encoded data storage, the dictionary itself can have substantial contribution to storage space. Muller et al. [79] make a thorough comparison between various dictionary compression algorithms regarding decoding speed and space consumption, and propose an empirical decision tree based approach to select a dictionary algorithm that is

either optimized for access speed or storage space on a given dataset. Zukowski et al. [119] propose PDICT compression scheme that allows infrequent values to be exceptions from the dictionary in order to reduce dictionary size on skewed frequency distributions, thus better compression performance.

In addition to compression, research explores how dictionary encodings affect query performance. Chen et al. [30] design a compression-aware optimizer to estimate overhead brought by accessing compressed data. Ray et al. [93] and Abadi et al. [9] show that by rewriting query predicates, one can efficiently skip the decoding step and execute queries directly on encoded data, which is beneficial to query performance due to reduced I/O requests and leverage in MOP. Jiang et al. [56] demonstrates how a SIMD-based algorithm can filter up to 18 billions entries per second on dictionary encoded data that is bit-packed.

While standard dictionary encodings only support querying directly on encoded data for equality predicates, an order-preserving dictionary [14] can support direct evaluation for range queries. For non-supported queries, either the encoded values must be decoded using the dictionary and evaluated against the predicates, or the predicates must be evaluated against the entire dictionary and passing keys are recorded. Antoshenkov et al. [15] propose an order-preserving data structure for DBMS and Binnig et al. [20] adapt the idea for in-memory analytic databases. Order preserving dictionaries require that the entire value space is known before encoding the dataset.

Dictionaries can be implemented using various data structures, such as an array, hash table, or trie [44] and different compression strategies, such as Huffman coding, Hu-Tucker coding [52], front coding [111], and RE-PAIR [66]. These compression strategies have full access to entire value corpus and thus can achieve better compression ratio, they all suffer from higher latency and low throughput [60], rendering them unsuitable for many database applications. In this paper, we focus our discussion on hash table-based dictionary with no compression strategy applied on the dictionary.

3.1.2 Consensus Protocol

When multiple distributed, shared-nothing nodes participate in the dictionary generation process in parallel, consensus protocols need to be employed to ensure a global unique key-value mapping. Consensus protocols are a family of algorithms for distributed systems that makes sure when multiple write attempts for a key occur, only one value is chosen for all participating nodes.

Many previous systems employ consensus protocols that provide eventual consistency to improve system throughput. Wu et al. [112] presents using distributed logs to maintain unified dictionary views on each node. Amazon’s Dynamo [34] designs a key-value store that employs a quorum and version based eventual-consistency algorithm to resolve value-conflict between nodes. These are closely related to our distributed-dictionary generation problem. However, in our scenario, each node encodes data stream and emits the results to output devices in real-time. Nodes only maintain local data copies of limited size, and are not allowed to re-encode data that has been written out. Eventual consistency does not fit in this case as there’s no time bound on when a change will be propagated to nodes, which may cause incorrect data to be emitted. In this paper, we limit our discussion to consensus protocols that are able to provide a stronger consistency guarantee. We focus on two families of consensus protocols: protocols with a distinguished leader, such as GFS [47], Raft [80], and RAMCloud [81], and decentralized protocols, such as Paxos [65].

Protocols with Distinguished Leader In this type of protocol, a unique node is designated([47, 81]) or elected([80]) to be leader, and all other nodes serve as followers. Followers who want to make write operations send requests to leader, who computes a final image based on the requests and broadcasts the results to all followers.

The distinguished leader protocol approach is easy to understand and simpler to implement [80]. In addition, the leader always maintains an up-to-date image of global dictionary, allowing easy access to the information. However, a single leader architecture is vulnerable to single node failure and requires a complicated process to recover. In addition, the leader

node can easily become performance bottleneck under high throughput.

Decentralized Protocols Two-phase commit (2PC) and Paxos are well-known decentralized protocols. These protocols allow any node in the system to propose a change. The nodes then vote to approve or reject the change. A unanimous agreement (2PC) or majority agreement (Paxos) is needed for the change to be approved. A decentralized protocol is more robust against single node failure, and allows the throughput to be evenly distributed among nodes in the cluster. But it may generate more network traffic when voting to reach an agreement. It can also suffer from competition of requests from different nodes and thus fail to make progress when trying to reach consensus. Bully algorithm [46] alleviates the blocking issue by always allowing the nodes with higher ID to move on.

In this paper, we implement both the distinguished leader protocol and decentralized protocol 2PC with bullying for distributed dictionary generation.

3.2 MOP Overview

Here we introduce the MOP organization, key space layout, necessary data structures, and a natural extension of MOP.

3.2.1 MOP Definitions

A Mostly Order Preserving Dictionary (MOP) is a bijective mapping of keys K and values V consisting of two sections. As shown in Figure 3.2 the first section is order preserving and the second is not. Key space is pre-allocated for the ordered section and the disordered section grows as needed, such that after initialization the MOP specifies at most m keys are ordered, while the dictionary can grow to n keys (with $n \geq m$).

The initialization phase looks at a sample (or head) of records from the dataset to be loaded by collecting statistics and a set of values (B) to bootstrap the ordered section. These values are distributed evenly throughout the ordered section. Our *allocation strategy*




Figure 3.2: Key space for MOP dictionary

determines how much space to use for the ordered section and how keys should be spaced in the section. During the generation phase, when inserting a value $v \in V$ into the MOP to get its corresponding key k , we distribute the incoming items into the ordered section. When code space conflicts happen (e.g., cannot insert the value without violating ordering), the unsettled items will be appended to the end of code space in the disordered section sequentially (e.g., v 's key k will be $m < k \leq n$). We refer to this as a *spillover*. Given a MOP dictionary, *ordered ratio* r is defined to indicate the proportion of ordered entries in the dictionary:

$$r = 1 - \frac{n - m + 1}{|V|}$$

A *padding ratio* p is defined to indicate the proportion of empty keys in the ordered space:

$$p = \frac{n - |V|}{n}$$

We use r and p for query and storage evaluation respectively. Therefore, the goal of MOP is to maximize the ordered ratio r as much as possible, while trying to avoid excessive “bloat” of the key space by minimizing p . In Section 3.3 we discuss how MOP allocates key space and assigns keys. Note that unless bit-packing is specified prior to encoding, n is not fixed and can grow. If bit-packing is used n is fixed and the dictionary can only hold 2^n keys, and if additional keys are needed the dictionary must be entirely recoded (e.g. regenerate). For the rest of the paper we focus on the case where n is dynamic, unless specified.

Given that it is impossible to guarantee that the ordered ratio of a MOP = 1, there will be some records that cannot fit into the ordered section. A natural extension to explore




Figure 3.3: Key space for C-MOP dictionary

is instead of a single ordered section, is cascading ordered sections before defaulting to a disordered section.

3.2.2 Cascade-MOP

With a *Cascade-MOP* (*C-MOP*), we nest multiple levels of order-preserving sections before a single disordered-section. By default we limit nesting to eight levels, but this is configurable. Here, when spilling a value to the disordered section we instead treat it as a new order-preserving section. Rather than appending to the end of the key space, we allocate a new nested ordered key space for those unsettled items. For this new zone we allow for refinement in our allocation strategy, in particular how much space to allocate and how to bias the initial placement of spilled values. We call each order-preserving section an OP zone. Figure 3.3 shows the layout of a C-MOP with k levels. The first zone runs from keys $(0, e_0)$, the next from $(e_0 + 1, e_1)$, and so on until the number of OP zones grows until it reaches the user-defined limit (8 by default). A final disordered section exists after the OP zones. In Section 3.3.2 we discuss determining the number of cascading levels.

3.2.3 Dictionary Generation Assumptions

For this paper we assume there is a single coordinator (thread, process, or node) that is responsible for assigning keys to values. There are n workers (threads, processes, or nodes) who are each reading a partition of a data file that contains records to write to a database. For attributes that are encoded using a dictionary, each worker maintains a local copy of the dictionary. When a value is to be encoded it first checks the local dictionary, and if

the value is not found it requests the key from the coordinator by providing the value. For performance reasons, the worker may batch requesting values together and the coordinator may batch requests from workers together – this is respectively referred to as worker batch size and coordinator batch size. When waiting for keys the workers block writing the current data records, and the leader broadcasts new keys to all workers, which is subsequently added to their local dictionaries when processed. Once the key for the value is found the worker writes out the encoded key for the record. We assume that once a worker writes an encoded value it cannot be changed, unless the process restarts (e.g. recoding). While we discuss the process for one attribute encoding for a data load, the process works for multiple attributes with separate metadata and attribute identifiers in the request.

3.3 Methodology

In this section we discuss the steps of MOP and C-MOP generation. This is broken into three main steps: initialization, encoding, and finalization. For C-MOP, we explain how spillover differs. For exposition we assume there is a coordinator managing the dictionary and multiple workers encoding values and requesting keys from the coordinator. We assume a shared-nothing architecture with workers starting with a distinct partition of the input dataset. Other approaches, such as threads with locks, are valid with minimal changes.

3.3.1 Initialization

When generating a new dictionary all workers notify a coordinator about the attribute to encode, which may come from a file containing multiple attributes. This message includes the worker’s coarse estimation about the data to encode (i.e. file size). The leader responds by requesting from each worker a sample of the dictionary to bootstrap the MOP. We refer to the percentage of records as *lookahead* (i.e. a lookahead of 0.10 is a 10% a sample). The worker preprocesses the sample of records before sending out a sample of the dictionary,




Figure 3.4: MOP generation example

which includes a set of records and several features (i.e. sortedness, estimated cardinality). These features help the coordinator understand the values to encode. By default, workers use the head of the file to sample unless otherwise specified. In rare cases, the leader will force workers to redo sampling using a uniform sampling strategy. This occurs if the features collected from workers indicate a high probability the attribute is sorted. To determine sortedness of an attribute, we use Kendall’s τ which generates a real number in $[-1, 1]$. An output of 1 denotes a fully sorted attribute and -1 denotes a fully inverse-sorted attribute.

Once the sample is collected, the number of unique values in the sample data is counted and divided by the lookahead proportion to determine the estimated cardinality of the unique values in the data. The size of the order preserving section’s key space is pre-allocated based on the estimated cardinality and a slack parameter which determines slack space between keys. The sample values are then uniformly distributed over the key space as shown in Figure 3.4. It is worth noting that this process could leverage more advanced sample-based distinct value estimators [29]; however, our experiments in Section 3.6.3 show these approaches add limited value to MOP performance. Given apriori distribution information, the sample values can be biased to the distributed key space, which we discuss in Section 3.6.3.

The space left between two ordered keys allows new input unique values to be ingested if encountered later in order to retain the overall ordering. We calculate this slack space

by dividing a controlled parameter *pitch* by the lookahead proportion. By adjusting *pitch*, the slack space can be scaled to allow for more or less key space available for new unique value batches. A higher *pitch* gives a larger initial dictionary with more space between the bootstrapped values. Too much slack may result in a dictionary to be sparse or padded (e.g. many empty cells), and too little slack results in the inability to order keys properly. As we later show, C-MOPs can achieve a high ordered ratio with a *pitch* of 1, and single layer MOPs can benefit from a *pitch* of 3, but this increases the key space. Once the initial bootstrapped dictionary is generated, the coordinator sends it to all workers, and the encoding process begins.

3.3.2 MOP Encoding

In the encoding stage, the remaining unique values get inserted into the MOP in groups of values, or batches. Each worker can batch a configured number of values (worker batch) then send them to the coordinator, which in turn can batch a set of requests (coordinator batch) to insert into the MOP. In addition to amortizing message overhead, batching multiple values for insertion can help with proper spacing of values and reduce spillover.

For every value in each batch, if there is space in the order preserving section, the value is inserted there uniformly between the two already inserted ordered values that the new value's ordering falls between. If multiple values in this batch fall in the same range, they will be evenly spaced. Figure 3.4 shows a sample key insertion for the first batch values.

If no space is left for inserting a value in the ordered section, the value will be spillover either to a disordered section for MOPs or to a cascading ordered section for C-MOPs.

MOP Spillover

With MOP, when values spillover they are added sequentially to the end of the dictionary in the disordered section. For example, in Figure 3.4 in the second batch, no space exists for *carrot* between *bean* and *cherry*. The *carrot* value then gets appended to the end of the

ordered key space in the disordered section. This batching process repeats until all data has been processed.

C-MOP Spillover

The C-MOP encoding process works in a similar fashion, with the difference being when there is no space left in an order preserving section for inserting new values, the spillover(s) may create another order preserving zone. Here, the process will add any pending requests to the current set of values to encode. All spilled-over values will be used to seed the next OP zone. Spillover can result in creating a final disordered section if we believe the encoding is nearing completion or we hit a prescribed maximum number of OP zones, which is the strategy we use.

When creating the next OP zone, we determine how large to make the zone based on the prior zone's size multiplied by a *MOP zone ratio*, which is set between 0.20 and 2.0. To determine the ratio we examine the estimated cardinality for the prior zone (scaled by estimated progress of the dataset), and the observed cardinality for the prior zone(s). If the observed cardinality is less than the estimated cardinality, we set the MOP zone ratio to 0.2 with the expectation that the spillover is due to the distribution being slightly deviated. If the observed cardinality is much higher than the estimated cardinality, we set the ratio closer to 2.0 as we assume our estimates were way off or that the distribution in the dataset has changed and we should account for larger capacity to avoid introducing an additional OP zone.

We explore biasing the placement of values into new OP zones using a histogram of the prior zone(s). Here, we create a histogram of values bucketed on the original sample values. We then look at all values currently inserted into the dictionary and count the number of values that either landed between the sample values or would have landed between the sample values if space permitted. Next, we partition the new OP zone into sections whose sizes are scaled by the histogram data. Values inserted into these OP zones will be distributed evenly

within their respective sections. As we later demonstrate, this approach works well under certain use cases, but we opt for the same MOP allocation strategy of uniformly placing values in the key space for simplicity.

3.3.3 Finalization

Once all values are encoded, the process is finalized. For each OP zone and the disordered section we create a zone map of min and max values for query evaluation and record the starting position of this zone (i.e. the first key). Figures 3.2 and 3.3 respectively show the final organization for MOP and C-MOP. The generated dictionary is written to the output file.

3.4 MOP Queries

In this section, we describe how MOP/C-MOP supports query operator execution efficiently. MOP is optimized for equality predicates, range predicates, and sorting. These operators can work directly on encoded data with little or no decoding overhead. For other operators, data needs to be decoded and materialized, just as when using a normal dictionary. The goal of our query evaluation is to rewrite any query q that filters or sorts on a MOP encoded attribute x , such that it minimizes decode operations and evaluates as much of the query directly on the encoded keys, as opposed to evaluating predicates against decoded values.

3.4.1 Equality Operator

An equality operator ($x = v_a$, $x \neq v_a$) can be performed efficiently on MOP just as on a normal dictionary. When taking $x = v_a$, for example, we first check if there exists a key k in MOP satisfying $MOP(k) = v_a$. A miss means no value in the column matches the predicate and the execution can return prematurely. Otherwise, we scan the encoded entries k_i , looking for $k_i = k$. This operation skips decoding operations on the encoded entries and

thus saves CPU overhead.

3.4.2 Range Operator

Without loss of generality, we discuss in this section how to execute an inclusive range operator ($x \in (v_a, v_b)$) on a MOP encoded column; the approach can be generalized to a compare operator, exclusive ranges, and `like` predicates with constant prefixes, as well as conjunctive and disjunctive combinations of these operators.

MOP and C-MOP can both be viewed as a combination of multiple order-preserving dictionaries (OP) and an unordered dictionary (DIS), where MOP contains only one OP, and C-MOP can have multiple OP sections. We first discuss the execution of a range operator on OP and DIS separately, then show how to combine the results to obtain the final result. An inclusive range operator for an OP can be easily extracted from a query q via query rewriting. To evaluate $x \in (v_a, v_b)$, we first find the corresponding key range $k_a = \min(\{k | OP(k) \geq v_a\})$ and $k_b = \max(\{k | OP(k) \leq v_b\})$, then perform a scan on the encoded column, looking for entries satisfying $k \in (k_a, k_b)$. This process involves no decoding operations.

When executing a range operator on DIS, we compare the zone map (v_{min}, v_{max}) of DIS against the query range and consider three cases.

Type 1 $(v_{min}, v_{max}) \cap (v_a, v_b) = \emptyset$. In this case, DIS does not contain any value within the query range, and can be safely skipped for the query.

Type 2 Here, the query range and zone map overlap. In this case, for the keys in DIS we perform a decode operation and compare the decoded result against the range. Let k_s be the starting key in DIS. For each encoded entry k , we check $k \geq k_s$ to make sure the key belongs to DIS, and, if so, we decode $v = DIS(k)$ and check if $v \in (v_a, v_b)$. This operation only involves decoding keys belonging to DIS, which is relatively small compared to decoding the entire column.

Type 3 $(v_{min}, v_{max}) \subseteq (v_a, v_b)$. In this case, all keys in DIS are included in the query range.

Let k_s be the starting key in DIS; we can rewrite the range query to be $k \geq k_s$ and execute it on the encoded column. This involves no decoding operations.

As a MOP contains two sections, OP and DIS, with disjoint key ranges, executing a range operator on MOP is equivalent to first applying the operator to OP and DIS separately then performing a disjunction of the results. As an example, we consider a query $x \in (apple, cherry)$ on the MOP shown in 3.4. The operator on OP is rewritten as $k \in (0, 4)$. As "cherry" is within the zone map of DIS, we also need to perform decoding for keys belonging to DIS, which yields $k \geq 14 \wedge decode(k) \in (apple, cherry)$. The result, as the disjunction of the two parts, is then

$$k \in (0, 4) \vee (k \geq 14 \wedge decode(k) \in (apple, cherry))$$

Similarly, query $x \in (apple, mango)$ will be translated to $k \in (0, 11)$ on OP, and, as "mango" is larger than v_{max} of DIS, it is translated on DIS as $k \geq 14$, and the result is

$$k \in (0, 11) \vee k \geq 14$$




Figure 3.5: Steps for range filtering with C-MOP

C-MOP is comprised of multiple OP sections and a DIS section, with their key ranges

non-overlapping, and thus can be processed in a similar manner as described above in MOP. More precisely, we first execute the operator on each section independently and perform a disjunction of the result. To further optimize query evaluation, we utilize a zone map for each ordered section, or zone, to minimize unnecessary zone look ups. To further optimize the result, we also perform range merging. If the final result contains range clauses that are adjacent to each other, e.g., $k \in (k_1, k_2] \vee k \in (k_2, k_3)$, they will be merged together as a larger range $k \in (k_1, k_3)$. This step reduces the number of ranges to be evaluated and potentially improves efficiency. This is shown in 3.5. Evaluating a range operator for a C-MOP works as follows:

Check the zone map for the disordered section to decide the query type: As with a MOP, step (1) is to check the type of query to determine if the disordered section is needed.

Check the zone map for each OP zone to see if skipping is possible: For each OP zone, step (2) checks its zone map before getting the appropriate keys in that zone. If the query range is disjoint with zone map, this zone is skipped.

Range query translation per OP zone: If skipping a zone is not feasible, for step (3) the range query is translated into the qualified key range for this OP zone (i.e. $k_a = \min(\{k | OP(k) \geq v_a\})$ and $k_b = \max(\{k | OP(k) \leq v_b\})$). After query translation, each OP zone outputs either skipped or a qualified key range that is added a disjunctive predicate.

Qualified range merging: A large number of qualified ranges adds more integer operations for verification on each record. In order to eliminate unnecessary compare operations, we merge key ranges into a larger one if they are adjacent in step (4).

3.4.3 Sort Operator

A sort operator for a MOP shares similar logic with range filtering, by leveraging the orderedness of the MOP dictionary. To avoid decoding the disordered section, we temporarily expand the key space to sort disordered keys in relation to the sorted section. The MOP sort operator pre-scans the dictionary and temporarily assigns dictionary entries in the dis-

ordered section with a float key that follows the orderedness in the ordered section and new float keys. We then build a mapping from the old integer code to the new float code for entries in the disordered section. We apply a sorting algorithm on encoded integer values and translated float values, without decoding any values.




Figure 3.6: Sort with MOP

To perform sort operation on a MOP encoded column, we create a bijection $sort : \mathcal{N} \rightarrow \mathcal{R}$ which maps all keys in $MOP(OP, DIS)$ to a real number space, satisfying that

1. For $(k, v) \in MOP$, the dictionary $(sort(k), v)$ is fully order-preserving.
2. For $k \in OP$, $sort(k) = k$.

With property 1, we can apply $sort$ to the encoded entries $\{k_i\}$, and do a merge sort on the sequence $\{sort(k_i)\}$. In addition, by property 2, the dictionary $sort$ uses contains much less entries than MOP (it only contains keys in DIS), making the sorting operation more efficient than first decoding the entries and then perform sorting.

Now we explain how to define the bijection $sort$ for $MOP(OP, DIS)$. We construct a sorted list L_{sort} , and a dictionary D_{key} . For each entry $(k_o, v_o) \in OP$, we insert v_o into L_{sort} . Then for each entry (k_d, v_d) in DIS , we perform a search in L_{sort} to find two adjacent entries (v_1, k_1) and (v_2, k_2) satisfying $v_1 < v_d < v_2$, insert v_d into L_{sort} , and insert $(k_d, \frac{k_1+k_2}{2})$ into D_{key} . The $sort$ bijection is then defined as

$$sort(k) = \begin{cases} k & k \in OP \\ D_{key}(k) & k \notin OP \end{cases}$$

The goal of our experimental evaluation is twofold. The first section evaluates MOP’s and C-MOP’s ability to improve query performance for range filtering and sorting. We compare these against order preserving and dictionaries that require decoding (e.g. non-order preserving). We evaluate across varied ordered ratios and selectivity ratios with a synthetic and real-world dataset. The second section evaluates MOP’s and C-MOP’s ability to generate highly ordered dictionaries and to understand what are the critical factors in doing so.

3.5 Experimental Setup

All experiments were performed on servers with 2 Intel(R) Xeon(R) CPUs E5-2670 v3 @ 2.30GHz, 128GB memory, 250GB HDD, Gigabit Ethernet, and Ubuntu 14.04. Experiments use a real-world dataset of `taxi` rides from New York City [7], the `lineitem` table from TPC-H, and two synthetic datasets derived from an English word dictionary based on `uniform` and `zipf` distributions. Unless otherwise stated we use scale factor 30 for TPC-H and 15 million rides from the taxi dataset. The default MOP configuration has a lookahead of 0.1, pitch size of 1, and a worker batch size of 20. The default C-MOP layer ratio is 0.2.

For these experiments we use the Parquet file format as our query performance testbed. Parquet is an open-source column-oriented file format for distributed analytic frameworks, such as Spark and Impala [114, 64]. It provides efficient data compression and encoding schemes with the ability to handle complex nested data types (e.g. lists and maps) [76]. In the Parquet file format, values from each column are logically organized to be adjacent and physically stored in contiguous memory locations for improved compression and query I/O. Parquet supports distributed writers by storing metadata at the end of the file. As Parquet only has native support for local dictionary encoding, we add a global dictionary encoding feature into Parquet. We write a stand-alone query engine in Scala that filters and joins on attributes with equality and range predicates. Depending on the query and encoding predicate evaluation can occur directly on the dictionary key (i.e. query rewriting) or by




Figure 3.7: Parquet Columnar Store Format

decoding the encoded keys.

As shown in Figure 3.7, a Parquet file is made up of several row groups, which are indexed by block metadata saved in a file footer. A row group consists of several column chunks with metadata also in the file footer. In each column chunk, there are data pages and a dictionary page if dictionary encoding enabled. The dictionary is stored in a dictionary page per column chunk. Columns are aligned in row group level, which means all data for a given row is organized in the same row group. In the file footer metadata is organized about column chunk metadata, zone maps (e.g. min, max and number of nulls), encoding, and compression information. In order to evaluate query performance of both local dictionary and global dictionary with the same system, we implement global and OP dictionaries in Parquet.

3.6 Evaluation

We firstly show query performance on MOPs encoded datasets to prove their abilities of improving query performance for range filtering and sorting operator. Then we show MOPs generation experiment results to understand the critical factors of generating MOPs.

3.6.1 Range Filtering Evaluation

In this section, we encode the given dataset under different MOP or C-MOP configurations in order to generate various ordered ratios. Order preserving and standard dictionary encoded datasets are generated for comparison. All encoded datasets are stored in Parquet’s file




Figure 3.8: Range filter overview on taxi dataset




Figure 3.9: Percent of records encoded by ordered keys.




Figure 3.10: Type-2 filter on taxi dataset

format.

Using a stand-alone Java query execution framework, we evaluate the three types of range queries discussed in Section 3.4.2 by decoding dictionary keys (decoded) and directly evaluating queries on keys via query rewriting (direct). In the following experiments, we use *MOP Decoded* for cases where we use the MOP organization but fully decode every value for checking the predicate, and we use *MOP Direct* to indicate best-effort filtering on encoded keys when possible (i.e. types 1, 2, and 3). *Dict Decoded* represents filtering on a dense dictionary (e.g. no padding) when decoding every value for a predicate evaluation. *OP Direct* indicates a dense sorted dictionary with predicate evaluation without any decoding on an OP encoded dataset. For all filtering evaluation, we do not materialize the output values and instead evaluate as a filter for late materialization that creates a bitmap indicating the records that satisfy the predicates [102]. Unless otherwise stated, we run each query 15 times and report the average running time.

Taxi Dataset Filtering

We evaluate query performance on the taxi ride dataset. We encode column *pickup_latitude* from the *trip_data* table with OP and MOP dictionary encoding respectively and use Parquet’s default encoding for other attributes [7]. With MOP’s default configuration, we can achieve a 94.2% ordering. We manually tune MOP parameters to get dictionaries with different ordered ratios.

Figure 3.8 shows the range query performance on the default MOP encoded dataset compared with the OP encoded dataset. MOP Direct performs similarly to OP Direct in terms of range filters for Type-1 and Type-3. The two counterparts have similar performance as neither decode the value during query processing. However, MOP Decoded is about 10% slower than Dict Decoded for all types of range queries as there are more entries in the MOP dictionary due to dictionary padding. For Type-2 range filters, MOP Direct is slightly slower than OP Direct but is still far more efficient than decoded filtering. We have a detailed analysis on Type-2 range queries with varying ordered proportions in following experiments.

Figure 3.9 shows the percentage of records encoded by ordered key versus MOP ordered ratio. Unlike the linear trend for TPC-H, which has a uniform distribution, the percentage of records encoded by ordered key grows rapidly on the Taxi dataset. This is typical for skewed distributions as it is more likely to capture the frequent records early. Therefore, Type-2 range queries can still perform efficiently with small MOP sorted ratios (i.e. 30%).

As shown in Figure 3.10, Type-2 direct filtering on MOP encoded datasets takes more time as the ordered ratio decreases, as more values need to be decoded. MOP Decoded performs better at the same time due to fewer entries being present in the MOP dictionary. Regardless, MOP Direct always outperforms the best decoded filter, even on datasets with low ordered ratios.

TPC-H Dataset Filtering

Using a TPC-H dataset of scale 30, we encode the table *lineitem* into a Parquet file with different dictionary encodings for the *shipdate* column and Parquet’s default encoding for other columns. We apply a range filter on the *shipdate* column with varying selectivity.

We achieve a 100% ordered ratio on the *shipdate* attribute with the MOP default space allocation strategy. Therefore, we manually adjusted our space allocation strategy to get MOP encoded files with different ordered ratios. In the following experiments, we fix slack to 5 and change the lookahead.




Figure 3.11: Range query on TPC-H dataset

Figure 3.12: TPC-H type 2 range filter with selectivity 0.00001

Figure 3.13: TPC-H Type-2 range filter with selectivity 0.27869

In Figure 3.11, we show the performance of three types of range queries on datasets with different dictionary encoding variations. We generate a MOP encoded file with an ordered ratio of 90.1% by manually setting the lookahead to 0.00001. These results show that direct query on MOP performs 1.5 to 2.2 times better than the regular decoded query version but 7% to 14% slower than a direct query on OP encoded datasets. The two direct queries perform quite similar to each other on Type-1 range queries, while they have a relatively greater performance difference on Type-2 and Type-3 queries. MOP direct Type-2 queries are slower as they decode records with keys from the disordered section before verifying the records. It is also slower than OP direct query on Type-3 queries as one more integer operation is needed to verify the records with disordered key. Even though disordered section decoding is not eliminated for Type-2 range queries, the MOP Direct query is still quite efficient because decoding is avoided on more than 90% of records. MOP direct query performance for Type-2 range queries varies as the proportion of records encoded by the ordered section changes, which is further dependent on the MOP ordered ratio. Again, Figure 3.9 shows the percentage of records encoded by the MOP ordered section increases uniformly as the MOP ordered ratio increases. This trend is typical for most uniformly distributed workloads as every value has relatively similar frequencies. Therefore, the MOP direct query performance on Type-2 range queries should be proportional to the MOP ordered ratio.

As we decrease the lookahead from 0.0001 to 0.000001, the MOP ordered ratio decreases




Figure 3.14: C-MOP range filter on taxi dataset with increasing the number of cascading levels.

from 99.8% to 25.9%. Figures 3.12 and 3.13 show the query performance for Type-2 queries on a dataset with different MOP ordered ratios. The dashed lines show the baseline query times for the regular decoded query and direct query on an OP encoded file, and the bars show the query time on a MOP encoded file with corresponding ordered ratios. Overall, the decoded query becomes more efficient as the ordered ratio decreases as less key space is used for the MOP dictionary, resulting in more efficient value decoding. However, it does not offset the overhead caused by an increased decoding workload in the MOP direct query. As more values need to be decoded for a direct query, the query time increases proportionally as the ordered ratio decreases. Please note that the datasets we are using are tuned to show the Type-2 query performance on different ordered ratios. However, we can achieve 100% ordered ratio using MOP default settings, where MOP’s performance is almost identical to that of OP.

C-MOP Filtering

Figure 3.14 shows range filter performance for a C-MOP with an increasing number of cascading levels for the Taxi dataset; TPC-H results are not included as it does not benefit from cascading due to its uniform distribution and low cardinality. Three sub-figures correspond to types of range filters respectively. In this experiment set, there are 1 qualified key range and 2 qualified key ranges on C-MOP datasets for queries of Type-1 and Type-3 respectively after range merging. For Type-2 queries, the number of qualified key ranges is

always equal to the number of OP Zones after range merging. With deeper cascading levels, there will be more OP Zones thus more qualified key ranges need to be checked, resulting in more integer operations for each record. However, the overhead from checking multiple qualified key ranges is alleviated by qualified range merging and offset by fewer disordered keys needing to be decoded. Therefore, the increasing cascading level shows minimal impact on filtering performance overall and in the next section we demonstrate space savings gained by cascading.

3.6.2 *SORT Evaluation*

In this experiment, we implement a MOP sort operator based on a recursive quick sort algorithm as described in Section 3.4.3. Figure 3.15 shows the sort performance on the Taxi dataset. The bars indicate end-to-end sorting time for a MOP encoded dataset with different ordered ratios, and the dashed lines show the baseline query time for regular decoded sorting and direct sorting on an OP encoded file. Overall, MOP sorting is slightly worse than OP direct sorting, but outperforms a decoded sorting. In spite of some inevitable translation overhead for entries from the disordered section, sorting operations primarily on integers and some floats are far more efficient than that of strings. MOP sorting takes more time as the ordered ratio decreases as there is more encoded value translation needed from entries in the disordered section. However, MOP sorting is still superior to decoded, even with a relatively low ordered ratio.

3.6.3 *MOP Generation*

In this section, we analyze the effects of changing each of the MOP generation parameters: lookahead, pitch, number of layers, and MOP layer ratio. We also look at the MOP's ability to dynamically adjust the number of keys allocated and the way values are distributed based on incoming data. We measure MOP generation performance by looking at the runtime of the generation and the ordered percentage of the resulting dictionary when using different




Figure 3.15: SORT Operator Runtime

datasets. Given the results of the prior section, for many workloads an ordered ratio of at least 70% gives the largest filtering performance improvement, which is trivial to achieve for most cases. As the ordered ratio nears 90%, filtering performance nears that of order preserving dictionaries. For many of these experiments, we highlight a trade-off between space and orderedness to allow the reader to understand how the impact of generation parameters. Based on these observations we believe a pitch size of 1, lookahead of 10%, worker batch size of 20, and at least 3 cascading layers with an auto layer-ratio set to 0.2 results in a good compromise of orderedness and performance. Therefore, unless otherwise stated, the experiments use this configuration. The default setup uses one server for the coordinator and one for all workers, with experiments including network latency in runtimes.

As shown in Figure 3.16, we evaluate MOP, C-MOP and OP generation performance with scaling the number of workers. We also show a local non-ordered dictionary as a performance baseline. For MOPs, we use a configuration with lookahead of 10%, worker batch size of 100. For C-MOP, we use a cascading level of 4. In Figure 3.16(a), the dictionaries are generated on the *shipdate* of the *lineitem* table with TPC-H scale 30. In Figure 3.16(b), the dictionaries are generated on the pickup latitude of the full taxi dataset, which is roughly 30





Figure 3.16: Average generation time.


GB and 125,000 values spread over 173 million records. This experiment parses the input CSV file and writes out the entire Parquet file. The leader and workers are co-located on a single machine that has 16 HDDs striped via RAID 0. Note that for a single OP worker, one process reads the file, sorts the keys, and then writes the file. For multiple workers, each worker reads their segment and sends values to the leader, who after sends the updated dictionary for workers to encode with. MOP is faster than OP in most cases and performs close to local dictionary encoding in TPC-H due to the relatively low cardinality. When the cardinality for the target column is relatively high, as with taxi, MOP has overhead for frequent coordination with the leader. MOP outperforms OP in most cases due to not learning the domain first (i.e., OP 1st scan), except for occasional cases with few workers where one worker blocks more than others when waiting for keys (i.e. 2 workers on Taxi). C-MOP performs quite close to MOP in terms of generation time as C-MOP only applies more key allocation for the spillover values, whose number is usually small.

Lookahead

For these MOP generation experiments, we test MOP performance when changing lookahead (sampling rate of the first X% of the file). Here, all datasets use approximately 173 million records. Figure 3.17(a) shows resulting ordered percentages and speeds of generating MOPs



((a)) Ordered Percentage



((b)) Runtime (ms)

Figure 3.17: Evaluating the effect of lookahead on MOP generation performance (Runtime and Ordered Percentage).

with different data skews. Skewed datasets have more repeated values, so key allocation costs are reduced leading to lower runtimes. Skewed workloads are also more predictable, leading to higher ordered percentages.

The higher variation in the uniform distribution workloads resulted in less predictable data that also lead to lower ordered percentages. Generating a MOP using the taxi dataset also resulted in higher ordered percentages than either of the synthetic workloads. This is because MOP predicted that the cardinality of the data was larger than it actually was, resulting in extra slack space being allocated. This extra slack space allowed a higher percentage of batch values to fit in the ordered section at the cost of a more inflated key space.

In Figure 3.17(b), we also show the costs of different lookahead percentages. As expected there is a linear growth for the datasets due to sampling from the head of the file, and the Taxi dataset is slower due to parsing 14 attributes to encode a single attribute. The Zipf and Uniform workload each only has one attribute per record, reducing the CPU intensive task of parsing and validating [40]. This result shows that high lookahead percentages can come with a high cost.




Figure 3.18: Evaluating the effect of pitch on MOP ordered percentages using uniform, zipf, and taxi workloads.


Pitch

Here we discuss the impact of scaling the slack space between sampled values with the pitch parameter. The slack space, or the empty space left between sample values, is calculated by dividing pitch by lookahead, which for these experiments is set to 10%. In Figure 3.18, we show how adding more space in between sample values affects the MOP’s ability to make ordered datasets. Overall, all workloads see similar benefits from an increase in pitch with respect to the ordered percentage.

As pitch scales the initial space between sample values, it changes the amount of room left for encoding values in the ordered section of the dictionary, which results in fewer values in the disordered section. By changing the pitch, we are able to adjust the ordered percentage of the resulting MOP, but to increase the ordered percentage we incur the costs of having a larger key space, which can result in a larger file (Sec. 3.6.4) and worse query performance [56].

Worker Configuration

For these experiments we measure how worker batch size and the number of workers affects MOP generation. In these experiments, we do not materialize the encoded file, the leader process and worker processes are run on different machines to account for network latency, and the generation is run on files with 100,000 values. Figure 3.19(a) shows how changing the



((a)) Evaluating the effect of worker batch size on MOP generation runtime using uniform and zipf workloads ((b)) Evaluating the effect of worker number on MOP generation runtime using uniform and zipf workloads

Figure 3.19: Worker configurations

the worker batch size affects the runtime of the MOP generation. By increasing the amount of batching being done by the workers, we increase the number of new values the worker will discover at once. As all values new to the worker need to be sent to the coordinator as a proposal, increasing the amount of items sent at once will decrease the total number of messages that need to be sent. This will in turn reduce message passing overhead and subsequently runtime. In Figure 3.19(b), we show how parallelism, through adding more workers, affects the speed at which a MOP can be generated. As expected, we see that the sharp decrease in the runtime early on indicates that the MOP generation is able to effectively leverage parallelism, but with adding many workers the benefit gains diminish as the coordinator is not able to allocate keys faster than the workers propose them. Partitioning the keyspace for multiple leaders could enable higher parallelism if needed.

Coordinator Batch size

This experiment shows how changing the coordinator batch size, or the number of worker proposals being received at once before processing, affects the MOP generation. The coordinator batch size was tested for values between 1 and 22.




Figure 3.20: Evaluating the effect of coordinator batch size on MOP generation runtime using uniform and zipf workloads. 24 workers are being used with worker batch size set to 1

In Figure 3.20, we show that, for uniform workloads, by increasing the coordinator batch size, up to a point, the key allocation process runs faster as the coordinator can process groups of values more efficiently than a single value at a time. However, by making the coordinator do more work before sending keys back to the workers, we incur costs on the worker side because the workers must wait for a response before they continue processing values. This results in the eventual increase in runtime that happens with large coordinator batch sizes. Unlike the uniform workload, the zipf workload does not see much change in runtime as workers will not have to send as many messages to the coordinator. This means leader does not need to process many values, and the workers do not need to wait for many responses.

Impact of C-MOP Layers' Size and Spacing

These experiments show how additional C-MOP layers affect ordered percentages, where all values within any C-MOP layer are considered ordered and all values that fail to fit into any level to be included in the disordered section. The tests in Figure 3.21 show that increasing the layer count increases ordered percentage as each successive layer creates more space for encoded values to be inserted into. Figure 3.21 also shows that increasing the MOP layer




Figure 3.21: Evaluating the effect of MOP layer ratio on MOP ordered percentages using uniform, zipf, and taxi workloads.

ratio parameter increases the ordered percentage of the dictionary. The first layer of the C-MOP is calculated in the same fashion as in the normal, single-layer MOP. The successive layer sizes are set by multiplying the previous layer size by the MOP layer ratio. For example, if the MOP layer ratio is 0.20 and the first layer's key space size is 1000, the second and third layer key space sizes would be 200 and 40 respectively. In this experiment, we tested ratios between 0.10 and 2.00. Larger ratios result in larger key spaces of the new layers, so more space exists in the ordered part of the dictionary for new values. Therefore, by increasing this ratio, we can increase the orderedness of the MOP, but the size of the key space being allocated will grow. Given these results and the minimal query overhead, we believe at least

three layers should be used. In the following experiments we further analyze the impact of sizing the layers.

Handling Distribution Changes

In this section, we will discuss the C-MOP’s ability to correct for a worst-case scenario of large changes in incoming data distributions. By recalculating the estimated cardinality when new layers are being created, the C-MOP can dynamically grow the MOP layer ratio if the number of distinct incoming values was underestimated initially. This ratio will be set between 0.20 and 2.0 as previous experimental data showed that a too small ratio would not sufficiently increase ordered percentage and a too large ratio would over-inflate the key space.




Figure 3.22: Evaluating the effect of dynamic MOP layer ratios when correcting for data distribution changes.

To demonstrate how C-MOP can correct for large distribution changes, we ran an experiment on a synthetic file where the first 20% of the file had a zipf distributed workload and the remaining 80% was uniform. When varying lookahead from 10% to 50%, Figure 3.22 shows how 1, 2, 4, 8, and 16 layer C-MOPs using both a static MOP layer ratio percentage

of 0.20 and the dynamic MOP layer ratio handles distribution changes based on the amount of data collected on the new distribution.

For lookaheads less than or equal to 0.20, both single layer MOPs and static ratio C-MOPs performed poorly as they had no information of the distribution change. However, dynamic ratio C-MOPs expect to need more space after seeing the distribution change, and work to correct initial estimated cardinality mispredictions. Furthermore, as each successive layer will have more time to learn, adding additional layers will greatly improve ordered percentages. This shows that adaptively changing the MOP layer ratio allows for robustness in the presence of adverse datasets.

Cardinality Estimation

Table 3.1 compares our simple cardinality estimator that divides the number of distinct values in the sample by the lookahead and a more advanced adaptive estimator [29] that separately estimates high frequency and low frequency cardinalities based on the sample. As the simple estimator generally overestimates the actual cardinality, it generates MOPs/C-MOPs with more padding and a higher ordered ratio than the adaptive estimator which generally has lower and more accurate cardinality predictions. For good MOP generation performance, we look to have a high ordered ratio and a low padding ratio to improve query performance and decrease file size. However, there is a tradeoff between the ordered and padding ratios. Increasing the number of keys allocated to the ordered section of the dictionary will increase the ordered ratio at the cost of extra padding and vice versa. As the cardinality estimation only affects the size of the key space, it does not improve general MOP performance. On a given dataset, for a fixed pitch, an estimator could have good ordered ratio performance, good padding ratio performance, or be somewhere in the middle. Scaling the key space through pitch allows the MOP to be easily tuned along the ordered ratio-padding ratio performance spectrum for any given cardinality estimate—so long as the estimate is not completely off base. As our simple estimator and adaptive estimator produce reasonable estimates, we chose to

use the naive estimator as it is less computationally expensive. However, as described in the next experiment, the ordered ratio-padding ratio trade-off results in no padding ratio benefit when using one estimator over another.

Method	Pitch	Zipf		Uniform		Taxi	
		Ordered Ratio	Padding Ratio	Ordered Ratio	Padding Ratio	Ordered Ratio	Padding Ratio
MOP Simple	1	0.711	0.699	0.662	0.457	0.910	0.850
	2	0.818	0.840	0.802	0.688	0.950	0.924
	4	0.891	0.918	0.896	0.837	0.975	0.962
	8	0.945	0.956	0.950	0.959	0.988	0.981
MOP Adaptive	1	0.470	0.265	0.686	0.501	0.755	0.356
	2	0.582	0.481	0.813	0.711	0.796	0.540
	4	0.698	0.677	0.902	0.848	0.868	0.755
	8	0.806	0.823	0.953	0.923	0.919	0.868
C-MOP Simple	1	0.960	0.740	0.743	0.532	0.979	0.879
	2	1.000	0.868	0.890	0.741	1.000	0.939
	4	1.000	0.933	0.973	0.867	1.000	0.970
	8	1.000	0.966	1.000	0.933	1.000	0.984
C-MOP Adaptive	1	0.621	0.302	0.771	0.573	0.821	0.445
	2	0.779	0.529	0.903	0.761	0.865	0.615
	4	0.934	0.720	0.975	0.877	0.932	0.800
	8	1.000	0.853	1.000	0.938	0.985	0.893

Table 3.1: Evaluating cardinality estimation techniques.

Key Space

Here we compare key space sizes between single-layered MOPs and multi-layered C-MOPs with both the simple and adaptive estimators. For each experiment, we adjusted the pitch until dictionaries were 70%, 80%, and 90% ordered. The average key space needed to achieve these ordered percentages are shown in Figure 3.23. All MOPs were generated using the same sets of zipf, uniform, and taxi data. C-MOP generation can better leverage key space in several, smaller order preserving sections than one large one as new layers can correct for distribution changes and help correct mispredictions made while sampling. By adding a second layer, we can generally see a reduction in the key space needed to produce a dictionary when compared to single-layered MOPs. The MOP generation is better able to leverage the allocated space in several, smaller order preserving sections than one large one. If the initial

sample is not fully representative of the distribution of data, there will be values that the MOP did not account for. These values will then overflow into the successive layers and influence the distributions of the new layers. The layering in C-MOPs can then work to correct mispredictions made in the initial sampling process.




Figure 3.23: Evaluating the space saving benefit of C-MOP with regard to padding ratio.

The choice of estimator is also shown to have no effect on the padding ratio as reaching a target ordered ratio requires adjusting pitch until a specific slack value is reached. For example, to achieve an ordered ratio of 0.70, there needs to be enough slack space between sample values for 70% of the data to fit into the ordered section(s), regardless of the initial cardinality estimation.

MOP generation performance with regard to padding is dependent on the distribution of values, so we do not see any padding benefit when using different estimators. We therefore use a simple estimator as it is less computationally expensive.

Figure 3.24 also shows how different techniques for allocating keys to successive layers affects the key space used. When the distribution stays constant throughout the file, as in the uniform workload, a significant benefit can be had by using a histogram of the previous data to influence the distribution in the successive layers. The histogram insertion technique is able to make corrections based on all data seen before the first spillover into a new layer, so it is able to notice congested segments in the dictionary and take care of values that spill over

in these regions effectively. However, in cases where distributions do not stay constant, the uniform insertion technique has small benefits over the histogram method. As the uniform method does not bias key spaces based on any previous information, it is not penalized for predicting wrong, which may happen when changes in incoming values cause spillovers in unexpected ranges, like seen in the zipf and taxi workloads. Therefore, while a histogram based approach can benefit certain cases, we opt for simplicity and stay with a uniform allocation strategy.




Figure 3.24: Evaluating the benefit of C-MOPs with regard to padding ratio when using uniform and histogram insertion techniques for the second layer.

Column Sortedness

As described in Section 3.3.1, MOP draws the sample from the file head by default and falls back to a uniform sampling strategy if the file appears sorted (i.e. Kendall’s Tau [61] ≥ 0.8 or ≤ -0.8). In Figure 3.25 we observe generating a MOP on sorted columns to justify this approach. Experiments are run on both a sorted and randomized version of three different columns, one with a zipf distribution, one with a uniform distribution, and one from the taxi dataset. The MOPs for the randomized columns were generated using a head sample, and the MOPs for the sorted columns were generated using both a head sample and a uniform sample. Each sampling strategy for each column was then used to generate both a 1-layer

MOP and an 8-layer C-MOP.




Figure 3.25: Evaluating the impact of sorted data and sampling strategy on a MOP’s ordered ratio.

When using head sampling on sorted columns, MOPs have poor ordered percentages. All batch values being inserted after the sample will be greater than the largest already inserted value, so the only available room in the ordered section is in the remaining space after the last sample value.

When forced to resample uniformly, the sample values will reflect the overall column distribution, so ordered percentages increase. However, generation then incurs the costs of taking a uniform sample. To read and parse the file sequentially, the entire column may have to be read before the batching process. In certain cases, sampling may be cheap, such as a large file stored on a distributed block store, and uniform sampling is then preferred. For this work we target head-based samples, as we assume that random access is expensive and reading random records is expensive due to string escape characters (e.g. line breaks occurring outside of record delimiters).

3.6.4 Overall Compression Performance

In this experiment set, we mainly focus on dictionary and bit-packing hybrid encoding compression performance. With dictionary encoding only, there is no storage overhead difference for MOP regardless of the key space used, as each record always takes exactly 4 bytes, which is not the case when bit-packing is introduced. Here, we use bit-packing encoding to further encode the targeted attributes in the previous experiment and report the column size. We use bit-packing locally in each partition of a Parquet file (called row groups) that truncates the keys to use the fewest bits to represent the largest key in the partition (i.e. 3 bits needed for keys 0–7).




Figure 3.26: Encoded Column Size and Dict Padding

As is shown in Figure 3.26(a), the dashed line indicates the column size for the order preserving dictionary encoded dataset and the bars show the column size for MOP encoded datasets. The blue line corresponding to the right vertical coordinate represents the padding ratio in the MOP dictionary. For certain ordered ratios there is the same storage cost compared with OP. Even though the key space used for MOP increases, the number of bits to represent the max value does not change. For TPC-H one more bit is added on each record for the dataset with ordered ratio from 72.0% to 90.1%, and two more bits are needed if we want achieve MOP ordered ratio 99.8%. Three more bits are needed to get MOP ordered

ratio $> 90\%$ on targeted attribute of the Taxi dataset. According to our experiment, 9.9% and 11.3% extra storage (compared with OP encoded column size) are required respectively for targeted attribute in the TPC-H and Taxi dataset to achieve a MOP ordered ratio $> 90\%$.




Figure 3.27: Scanning and decoding 180 million values.

Figure 3.27 shows the impact of padding ratios on scanning and decoding. Here, we take TPC-H lineitem shipdate and generate various MOPs with varied padding ratios (via ordered ratios). We then do a full scan of the column and decode every encoded key sequentially. These results show that a larger dictionary negatively impacts decoding performance.

3.7 Conclusion

In this paper we introduce mostly order preserving dictionaries (MOP) for supporting efficient range queries on encoded datasets. We present a technique for generating a MOP with a limited sample of the input dataset while minimizing the size of the dictionary. In addition, we introduce a variation that uses cascading MOPs (C-MOP) that has multiple levels of ordered keys. We present query rewriting rules to minimize decoding of keys to minimize predicate evaluation latency. We implement MOP and C-MOP in the open-source columnar framework, Parquet, and evaluate query and generation performance. Our results

demonstrate that MOPs are able to accelerate range filtering and sorting, and achieve high order ratios with small samples.

CHAPTER 4

DECOMPOSED BOUNDED FLOATS FOR FAST COMPRESSION AND QUERIES

Modern applications and systems are generating massive amounts of low precision floating-point data. This includes server monitoring, smart cities, smart farming, autonomous vehicles, and IoT devices. For example, consider a clinical thermometer that records values between 80.0 to 120.9, a GPS device between -180.0000 to 180.0000, or an index fund between 0.0001 to 9,999.9999. The International Data Corporation predicts that the global amount of data will reach $175ZB$ by 2025 [94], and sensors and automated processes will be a significant driver of this growth. To address this growth, data systems require new methods to efficiently store and query this low-precision floating-point data, especially as data growth is outpacing the growth of storage and computation [86].

Several popular formats exist for storing numeric data with varied precision. A *fixed-point* representation allows for a fixed number of bits to be allocated for the data to the right of the radix (i.e., decimal point) but is not commonly used due to the more popular floating-point representation. *Floats* (or floating-point) allows for a variable amount of precision by allocating bits before and after the radix point (hence the floating radix point), within a fixed total number of bytes (32 or 64). For floats, the IEEE float standard [51] is widely supported both by modern processors and programming languages, and is ubiquitous in today's applications.

However, two main reasons result in floats not being ideal for many modern applications: (i) an overly high-precision and broader range that wastes storage and query efficiency; and (ii) not being amenable for effective compression with efficient in-situ filtering operations. For the former reasons, many databases offer custom precision format, typically referred to as a *numeric* data type. For example, consider Figure 4.1 that shows how precision varies for the IEEE float standard and sample application requirements on GPS and temperature datasets.




Figure 4.1: Many datasets only span a fixed range of values with limited precision, which is a small subset of broad float number range and precision.

The numeric approach allows for “just enough” precision but is not well optimized for efficient storage and filtering. Several methods have been explored for the latter reason, but due to the standard format, compression opportunities are limited for compression effectiveness, throughput, and in-situ query execution. This paper proposes a new storage format that extends the ideas of a numeric data type that supports custom precision but is optimized for fast and effective encoding while allowing to work directly over the encoded data.

To motivate our design, we first describe short-comings with popular compression techniques for bounded-precision floating data. Specifically, the recently proposed Gorilla method [89] is a delta-like compression approach that calculates the XOR for adjacent values and only saves the difference. Gorilla achieves compression benefits by replacing the leading and trailing zeros with counts. Gorilla is a state-of-the-art compression approach for floats, but it does not work well on low-precision datasets as low-precision does not impact float’s representation similarity. In addition, Gorilla’s encoding and decoding steps are slow because of its complex variable coding mechanism. Alternative float compression approaches leverage

integer compression techniques by scaling the float point value into integer [21]. Despite their simplicity, these approaches rely on multiplication and division operations that are usually expensive and, importantly, cause overflowing problems when the input value and the quantization (i.e., scaling) factor are too large.

General-purpose byte-oriented compression approaches, such as Gzip [35], Snappy [109] and Zlib [45], can also be applied to compress floats. The input float values are serialized into binary representation before applying byte-oriented compression. These compression approaches are usually slow because of multiple scans looking for commonly repeated byte sub-sequences. Furthermore, full decompression is needed before any query evaluation. Dictionary encoding [98, 74, 15] is applicable for float data, but it is not ideal since the cardinality of input float data is usually high, resulting in an expensive dictionary operation overhead. Note that we are only considering lossless formats or formats with bounded loss (e.g., configurable precision). Lossy methods (e.g., spectral decompositions [38] and representation learning [83]) may compress the data more aggressively but often at the cost of losing accuracy.

Decades of database systems demonstrate the benefit of a format that allows for a value domain that includes a configurable amount of precision. We take this approach, and to address the above concerns, integrate ideas from columnar systems and data compression for our proposed method, **BoUnded Fast Floats** compression (BUFF). BUFF provides fast ingestion from float-based inputs and a compressible decomposed columnar format for fast queries. BUFF ingestion avoids expensive conditional logic and floating-point mathematics. Our storage format relies on a fixed-size representation for fast data skipping and random access, and incorporates encoding techniques, such as bit-packing, delta-encoding, and sparse formats to provide good compression ratios¹ and fast adaptive query operators. When defining an attribute that uses BUFF, the user defines the precision and optionally the minimum and maximum values (e.g., 90.000-119.999). Without the defined min and max

1. Compression ratio is defined as $\frac{\text{compressed_size}}{\text{uncompressed_size}}$.

values, our approach infers these through the observed range, but at a decreased ingestion performance. Our experiments show the superiority of BUFF over current state-of-the-art approaches. Compared to the state-of-the-art, BUFF achieves up to $35\times$ speedup for low selective filtering, up to $50\times$ speedup for aggregations and high selective filtering, and $1.5\times$ speedup for ingestion with a single thread, while offering comparable compression sizes to the state-of-the-art.

We start with a review of the relevant background (Section 4.1), which covers numeric representations (Section 4.1.1) and a wide range of compression methods (Section 4.1.2). In Section 4.2, we present BUFF compression and query execution with four contributions:

- We introduce a “*just-enough*” bit representation for many real-world datasets based on the observation of their limited precision and range (Section 4.2.1).
- We apply an aligned decomposed *byte-oriented columnar storage* layout for floats to enable fast encoding with progressive, in-situ query execution on compressed data (Section 4.2.2).
- We propose *sparse encoding* to handle outliers during compression and query execution (Section 4.2.3).
- We devise an *adaptive filtering* to automatically choose between SIMD and scalar filtering (Section 4.2.4).

We evaluate BUFF with its competitors in terms of compression ratio and query performance in Section 4.3, and conclude in Section 4.4.

4.1 Background and Related Work

In this section, we first introduce several numeric data representations, including the most popular float format. Then, we review several float compression approaches with a set of

Value	Fixed<2,20> Integer bits Fractional bits Sign bit	32-bits Float Exponent bits Mantissa bits	Gorilla Control bits: 11 if current meaningful bits does not fit bit range of previous meaningful bits Control bits: 0 if xor=0							
0.66	00.10101000111101011100	00111111001010001111010111000011	Step 1 XOR with previous value Step 2 write control bits, length and difference bits							
1.41	01.01101000111101011100	0011111101101000111101011100001	Leading zeros Meaningful bits Trailing zeros 0011111000111101011100001010010							
1.41	01.01101000111101011100	0011111101101000111101011100001	00000000100111001000111100100010 [b11,d8,d22] 10011100100011110010001							
1.50	01.10000000000000000000	001111111000000000000000000000	00000000000000000000000000000000 # leading zeros (decimal) # meaningful bits (decimal) [b0]							
2.72	10.10111000010100011111	0100000001011100001010001111011	0000000011101000111101011100001 [b11,d9,d22] 11101000111101011100001							
3.14	11.00100011110101110001	01000000010010001111010111000011	011111111011100001010001111011 [b11,d1,d31] 111111111011100001010001111011							
size(bits)	132	Integer bits Fractional bits 192	Control bits: 10 if use previous meaningful bits range 176							
Sprintz										
Step 1 Quantize×100	Step 2 prediction Lookup table	Predicted value	Step 3 delta	Step 4 bit-packing	Step 1 quantize	Step 2 mostly16	Step 1 dictionary encoding	encoded	Step 2 bit-packing	Step 1 LZ77; Step 2 Huffman coding
66	66	66	66	1000010	66	0x0042	0.66 0	0	000	3f28f5c3
141	66	75	141	1001011	141	0x008D	1.41 1	1	001	3fb47ae1
141	141	0	141	0000000	141	0x008D	1.50 2	1	001	[8,8]
150	141	9	150	0001001	150	0x0096	2.72 3	2	010	3fc00000
272	150	121	272	1111001	272	0x0110	3.14 4	3	011	402e147b
314	272	42	314	0101010	314	0x013A		4	100	404[40,5]
size(bits)				42		96			210	328*

*For Gzip, we skip showing step 2 Huffman coding for space, and report the final encoded size.

Figure 4.2: Compression examples for different approaches.

examples. Finally, we include popular query-friendly data structures that have influenced our compression format.

4.1.1 Numeric Data Representation

Numeric data representation is the internal representation of numeric values in any hardware or software of digital systems, such as programmable computers and calculators [100]. Many data representations are developed according to different system and application requirements. The most popular implementations include fixed-point (including numeric data type) and float-point.

Fixed-point uses a tuple $\langle sign, integer, fractional \rangle$ to represent a real number R :

$$R = (-1)^{sign} * integer.fractional$$

Fixed-point partitions its bits into two fixed parts: signed/unsigned integer section and fractional section. For a given bits budget N , the fixed-point allocate a single bit for sign, I bits for the integer part, and F bits for the fractional part (where $N = 1 + I + F$). Fixed-point always has a specific number of bits for the integer part and fractional part. Figure 4.2 shows examples of fixed-point; as we can see from the “Fixed” column, the radix point’s location

is always fixed, no matter how large or small the corresponding real number is. Fixed-point arithmetic is just scaled integer arithmetic, so most hardware implementations treat fixed-point numbers as integer numbers with logically decimal point partition between integer and fractional part. Fixed-point usually has a limited range ($-2^I \sim 2^I$) and precision ($2^{(-F)}$), thus it can encounter overflow or underflow issues. A more advanced dynamic fixed-point representations allow moving the point to achieve the trade-off between range and precision. However, this is more complicated as it involves extra control bits indicating the location of the point. Fixed-point is rarely used outside of old or low-cost embedded microprocessors where a floating-point unit (FPU) is unavailable. Fixed-point is currently supported by few computer languages as floating-point representations are usually simpler to use and accurate enough.

Floating-point uses a tuple $\langle sign, exponent, mantissa \rangle$ to represent a real number R :

$$R = (-1)^{sign} * mantissa * \beta^{exponent}$$

Instead of using a fixed number of bits for the integer part and fractional part, Float point reserves a certain number of bits for the exponent part and mantissa, respectively. The base β is an implicit constant given by the number representation system, while in our floating-point representation, β equals to 2. The number of significant digits does not depend on the position of the decimal point. For a given bit budget N , the floating-point allocates a single bit for sign, E bits for the exponent part, and M bits for the mantissa part (where $N = 1 + E + M$). The IEEE standard defines the format for single and double floating-point numbers using 32 and 64 bits, respectively. The exponent is encoded using an offset-binary representation with the zero offset by exponent bias:

$$R = (-1)^{sign} * 1.mantissa * 2^{exponent}$$

The IEEE standard specifies a 32-bits float as having: Sign bit: 1 bit, Exponent width: 8

bits (offset by 127), and Significand precision: 24 bits (23 explicitly stored). Figure 4.2 shows several examples of 32-bits float format; as we can see from the ‘32-bits float’ column, the underlined bits correspond to the fractional bits, and the radix point is floating base on its represented value.

Numeric data type allocates enough bits for a given real number depending on the precision and scale of the input. Numeric data type can store numbers with a very large number of digits and perform calculations exactly. To declare a column of type numeric, we use the syntax: `NUMERIC(precision, scale)`. The scale is the count of decimal digits in the fractional part, to the right of the decimal point. The precision is the total count of significant digits in the whole number, that is, the number of digits to both sides of the decimal point (e.g., 65.4321 has a precision of 6 and a scale of 4).

The Numeric data type is used in most database systems, e.g., MySQL, PostgreSQL, Redshift, and DB2, when exactness is required. However, the flexibility and high precision come at the cost of more control bits associated with each value. The storage cost of the Numeric type varies depending on the system implementation. PostgreSQL uses two bytes for each group of four decimal digits, plus 3 ~ 8 bytes overhead for each value. In MySQL, Numeric values are represented in a binary format that packs nine decimal digits into four bytes. This makes the Numeric type inefficient for space. Besides, the query on Numeric data type is either by materializing to float point for approximate calculation or using its own exact arithmetic which is very expensive compared to the integer types or the floating-point types described earlier.

4.1.2 Compression for Decimal Numbers

Due to the popularity of float-point data, several compression techniques exist, including methods for float-point data, general-purpose methods, and migrated approaches from integer compression. We limit our discussion to lossless compression techniques.

Gorilla [89] is an in-memory time-series database developed by Facebook. It introduces

two encodings to improve delta encoding: delta-of-delta (delta encoding on delta encoding, e.g., [101,102,103, 104] are encoded as [101,1,0,0]) for timestamps, which is usually an increasing integer sequence, and XOR-based encoding for value domain, which is a float type. In the XOR-based float encoding, successive float values are XORed together, and only the different bits (delta) are saved. The delta is then stored using control bits to indicate how many leading and trailing zeroes are in the XOR value. Similar compression techniques are also used in numerical simulation [72, 53, 71] and scientific computing [92, 27]. Figure 4.2 shows an example of Gorilla encoding with the first step calculating XOR, and the second step writing the control bits, number of leading zeros, number of meaningful bits, and actual meaningful bits. We use 32-bits to save space even though Gorilla was originally designed for 64-bits format. Gorilla is the state-of-the-art approach for float compression, which is widely used in many time series database systems, such as InfluxDB [4] and TimescaleDB [6].

However, Gorilla compression uses variable-length encoding, which means records are not aligned with their encoded bits. The encoded bits have to be decoded sequentially to reach the target records. Each value depends on its previous record; thus, for a target value, all previous values must be decompressed. Those features impede effective record skipping and random data access.

Sprintz [21] was initially designed for integer time series compression. Sprintz employs a forecast model to predict each value based on previous records via a lookup table, and then encodes the delta between the predicted value and the actual value. Those delta values are usually closer to zero than the actual value, making it smaller when encoded with bit-packed encoding. It is also possible to apply Sprintz to floats by first quantizing the float into an integer. As is shown in Figure 4.2, we first multiply input values by 100 and get a series of integers. Then we predict the next value based on the prediction model. In the third step, we calculate the delta between the prediction and actual value, and compress the delta with bit-packing encoding in the last step. Sprintz uses fixed-length for encoded values, which

is easy to locate the target value. However, to decode the target value, we need to fully decode prior values, since previous values are still needed to predict the current value and then decode the value with compressed delta. Furthermore, the multiplication of decimal numbers is costly and potentially raises integer overflow issues during quantization.

Mostly encoding encodes the data column where its data type is more extensive than most of its stored values required. Mostly encoding is used for integer and numeric data type compression in the Amazon Redshift data warehouse [49]. For numeric values, Mostly encoding quantifies floats into an integer before applying compression. Most of the column values are encoded to smaller data representation with mostly encoding, while the remaining values that cannot be compressed are stored in their original form. Mostly encoding can be less effective than no compression when a high portion of the column's values cannot be compressed. Figure 4.2 shows an example of Mostly encoding. It first quantifies the float values and then applies MOSTLY16 (2-byte) for the given data as it needs 9 bits for each quantized value (we use hex to save space). This representation is less effective than bit-packing encoding that stores input value using as few bits as possible.

General-purpose byte-oriented compression encodes the input data stream at the byte level. Popular techniques, such as Gzip [35], Snappy [109] and Zlib [45] derive from the LZ77 family [117] that looks for repetitive sub-sequence within a sliding window on the input byte stream and encodes the recurring sub-sequence as a reference to its previous occurrence (first step in Figure 4.2). For a better compression ratio, Gzip applies Huffman encoding on the reference data stream (second step in Figure 4.2). Snappy only applies LZ77 but skips Huffman encoding for higher throughput. Byte-oriented compression treats the input values as a byte stream and encodes them sequentially. The data block needs to be fully decompressed before any original value can be accessed.

4.1.3 Query-friendly Storage Layout

Compression should not only keep the data size small but also support fast query execution. Prior work speeds up query performance by introducing a hierarchical data representation layout. Some works [68, 107, 91, 78] are orthogonal to data compression but are applicable to BUFF to speed up the query performance.

DAQ [91] uses a bit-sliced index representation and computes the most significant bits of the result first then less significant ones. With bit-sliced index representation, it performs efficient approximation and provides efficient algorithms for evaluating predicates and aggregations for unsigned integers type. Column Sketch [50] introduces a new class of indexing scheme by applying lossy compression on a value-by-value basis, mapping original data to a representation of smaller fixed width codes. Queries are evaluated affirmatively or negatively for the vast majority of values using the compressed data, and check the original data for the remaining values only if needed. PIDS [58] identifies common patterns in string attributes from relational databases and uses the discovered pattern to split each attribute into sub-attributes and further compress it. The sub-attributes layout enables a predicate push-down to each attribute. And the intermediate result from the prior column could be projected to the following attributes efficiently with the aid of fast data skipping and progressive filtering.

BitWeaving [68] provides fast scans in a main-memory database by exploiting the parallelism available at the bit level in modern processors. It organizes the input codes horizontally or vertically into a processor word, allowing early pruning techniques to avoid accesses on unnecessary data at the bit level and speed up the scan performance. ByteSlice [42] is another main memory storage layout that supports both highly efficient scans and lookups. Similar to BitWeaving, ByteSlice assumes encoded binary as input. But instead of stripping the input code in bit units, ByteSlice decomposes the input code at the byte level and pads the trailing bits into a byte to achieve better read/write speed and be compatible with SIMD instructions. MLWeaving [107] is an in-memory data storage layout that allows efficient materialization of quantized data at any level of precision. Its memory layout is based on

BitWeaving, where the first bit in a batch of input values is saved as a word, followed by the second bits of the same batch, and so on. Data at any level of precision can be retrieved by following a different access pattern over the same data structure. Using lower precision input data is of special interest, given their overall higher efficiency.

These techniques speed up query performance by progressively filtering out disqualified records and narrowing down the records that need to be parsed. They mainly process a query in their defined unit (either sub-attributes or bits) on their columnar data layout. This fine-grained columnar data layout is the foundation for BUFF.

4.2 BUFF Overview

Many applications generate data that varies within a specific range with limited precision. Based on our discussion of state-of-the-art, none can leverage this feature to support efficient compression and query execution. To alleviate those deficiencies, we propose a novel compression method for floats. The compression works at two levels: eliminating the less significant bits based on a given precision (Section 4.2.1), and splitting float into the integer part and fractional part, then compressing those bits separately with two splitting strategies (Section 4.2.2). We then introduce sparse encoding to handle outliers to avoid compression performance deterioration (Section 4.2.3). In the end, we describe query execution on our compressed data format (Section 4.2.4). With these techniques, we achieve both good compression ratio and outstanding query performance.

4.2.1 Bounded Float

Computer systems can only operate on numeric values with finite precision and range. Using floating-point values as real numbers does not clearly identify the precision with which each value must be represented. Too small precision yields inaccurate results, and too big wastes computational and storage resources [36]. We target applications that need limited precision

and data that falls within a finite range. Therefore, our proposed compression takes full advantage of the range and precision features of a given dataset.

Bounded Precision

Float point data-type uses most bits for *mantissa*, which is not necessarily needed for many sensors or applications since many sensors usually provide limited decimal precision (e.g., the clinical thermometer has one place of precision 0.1 °F). However, the decimal precision is not aligned with those *mantissa* bits in a float due to the floating radix point in the format. Float format makes the most of its *mantissa* bits to get the best approximation for the given number. With 0.1 for example, all 52 *mantissa* bits are used to get its approximation (0.10000000000000000555) in IEEE double format. This mechanism is intended to support high precision for float format, but it impedes efficient float compression in most real-world cases. Those *mantissa* trailing bits provide far more extra precision than required, which leads to a huge bits change with a tiny number fluctuation. For example, Gorilla achieves good compression by first applying XOR to the current number with the previous one, then removing the leading and trailing zeros. The problem here is that *mantissa* bits are too sensitive to the number changes, making it difficult to leverage the trailing zero benefits. Suppose we know the precision of the incoming numbers. In that case, we can provide just

Table 4.1: Bounded float still keeps decimal precision

Real number		Sign	Exp(8 bits)	Mantissa(23 bits)	Actual value
3.14	binary		010000000	10010001111010111000011	3.1400001
	bounded-2		010000000	10010001100000000000000	3.13671875
					≈ 3.14

enough precision support by eliminating the less significant *mantissa* bits. At the same time, we are still able to provide a guaranteed precision float for downstream analytic. Table 4.1 shows float representation of a given value 3.14. As is shown in the binary row, the float format saves the closest approximation as 3.1400001. If we only care about two decimal positions after the decimal point, as is shown with *bounded-2* row, we can remove some

trailing bits, while its value is still possible to approximate to the original value with two places of precision. The removal of insignificant bits saves a lot of bits when representing a number with a given precision. This is one motivation for BUFF. To efficiently determine the number of fractional bits need for different required precision, we run brutal-force verification on all possible numbers under each precision requirement. We get the maximal fractional bits needed for its corresponding precision, as shown in Table 4.2, which becomes the lookup table for bounded precision. We only provide up to 10 digit places precision here as it is good enough for most datasets. We can extend to higher precision, but it is less meaningful as commonly used float only has 52 bits for mantissa.

Table 4.2: Number of bits needed for targeted precision

Precision	1	2	3	4	5	6	7	8	9	10
Bits needed	5	8	11	15	18	21	25	28	31	35

Bounded Range

Another observation for most application scenarios is bounded range, The measured value is delimited by a lower and an upper measuring bound limit that defines the measuring span. The bounded range could result from either the measuring range of the instrument (e.g., clinical thermometer measuring range from $35^{\circ}C$ to $42^{\circ}C$), the physical definition of measurand (e.g., CPU usage is defined between 0% to 100%), or based on other domain knowledge. The bounded range is another feature that is helpful when compressing the data since we only focus on the given range, and we can encode the number in a more concise but accurate manner. BUFF can work with no provided ranges, but as our experiments show, it hinders ingestion performance as the compression and allocation needs to find the minimal and maximal observed value when encoding a set of values. With bounded precision and range, we can get rid of the less useful and redundant bits initially designed for higher precision and broader range. With those two techniques, we can get a condensed representation of the input data. In addition to the compression benefits, BUFF also adopts float splitting

to organize our compressed data better.

4.2.2 Float Splitting and Compression

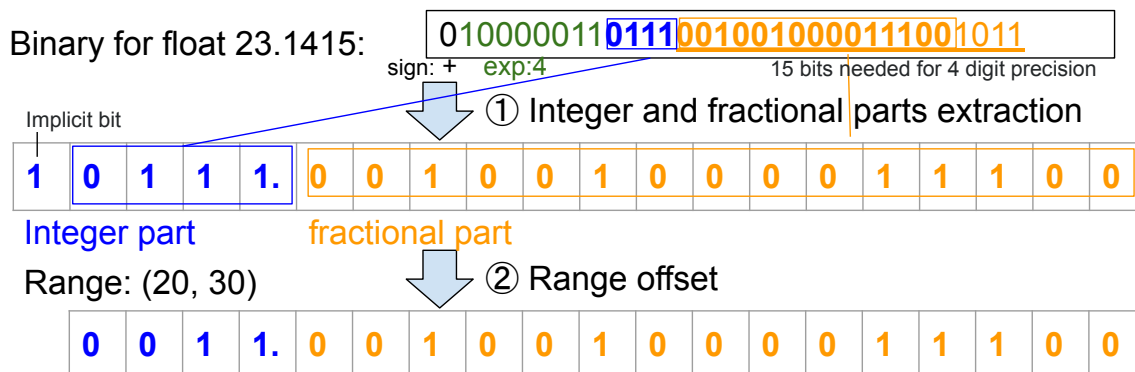


Figure 4.3: Extracting range bounded integer part and precision bounded fractional part for float number

The basic idea of float spitting is decomposing the input number into integer and fractional parts. This division separates the given number by radix point in a logical way and organizes each component into a different physical location. We can then apply efficient compression to those two parts, respectively. Bit splitting on fixed-point numbers is intuitively easy to do as its fixed position for the decimal point. However, the splitting is not easily applicable to float data directly, since float numbers are not aligned by the radix point. We extract both parts of a given float number with bit operations. First, we extract the exponent with a bit-mask to locate the decimal point. With the exponent value, we then bit-shift the float number and extract all left-side bits (including one implicit bit) as the integer part, and right-side bits as the fractional part. Figure 4.3 shows how this process works for a float value 23.1415 with scale 4 represented in 32-bit IEEE format. The exponent (marked with green font) is decoded as 4, which means the integer part has 5 bits (with a single leading 1 given implicitly by IEEE standard). We then ① extract 4 leading bits from the mantissa (marked with red) and add a leading 1 to get the integer part. Based on Table 4.2, we extract 15 following bits as fractional part for required precision 4. ② With the range information either given by user or detected by our encoder, we offset the integer

align them on byte boundaries. Assume we have m bits to read/write, this could be finished by $\lfloor \frac{m}{8} \rfloor$ times byte read/write, and at most one more bit read/write for remaining $m - \lfloor \frac{m}{8} \rfloor$ bits. Whereas float splitting divides m bits into x bits and y bits (where $x + y = m$), and this leads to $\lfloor \frac{x}{8} \rfloor + \lfloor \frac{y}{8} \rfloor$ times byte read/write, and at most twice bit read/write for remaining $x - \lfloor \frac{x}{8} \rfloor$ bits and $y - \lfloor \frac{y}{8} \rfloor$ bits. Additionally, the integer part usually takes less than one byte for many datasets because of the narrow value range for many measurands. To avoid slowness of bits read/write, we can always pad int bits with several following fractional bits to extend it into a byte. This keeps the fine-grained float splitting, but good read/write performance as well. If the goal is to save m bits anyway, we should split the bits in a more efficient way. This is the motivation of our advanced splitting version: byte-oriented float splitting.

Byte-oriented float splitting compression works similarly to the float splitting discussed previously. We first extract integer bits and just enough fractional bits that meet the precision requirement, then offset the intermediate value by minimal value before writing those bits in the byte unit separately. Each byte unit is treated as a *sub-column* and stored together. We store the remaining trailing bits of each record as the last sub-column, as is shown in Figure 4.4. We save the range and precision information as metadata along with compressed data to guarantee the decompression capability. With byte-oriented float splitting, compression and decompression are improved because of the fast byte reads/writes. We can also support custom variable precision materialization and achieve more efficient query execution because of a byte-oriented splitting layout.

4.2.3 Handling Outliers

In addition to variable precision materialization support brought by byte-oriented splitting, we can also handle outliers efficiently.

Outliers in the input data sequence are the data points that differ significantly from others. Outliers enlarge the value range significantly and further inflates the code space,




Figure 4.5: Sparse encoding condenses sub-column.

which requires more bits to encode the input data sequence. For example, 99% of the input value can be represented by 1 byte, while 1% needs 4 bytes. Encoding all numbers in the same length inflates the code space to 4 bytes because of the 1% outliers. The code space inflating, if not handled well, will deteriorate the compression effectiveness and query efficiency. Outliers can be too large or too small that are far away from the common range. While in either case, the outliers usually only inflate the higher bits either with padding zeros or padding ones. When there are both large and small outliers, the padding bits can be any bit sequence between all zeros and all ones. This requires us to detect outliers and distinguish the most frequent value in the higher encoded byte. We first head sample the dataset and then apply Boyer–Moore majority voting algorithm [24] in each compressed byte unit to find the majority item (occurring more than 50%). If the majority item has a frequency greater than 90%, we get most frequent value and re-organize the corresponding byte column to achieve better compression by *sparse encoding* as shown in Figure 4.5: we keep the frequent value at the beginning and use a compressed bit-vector to indicate the outlier records, followed by outlier values in the current byte column. The outlier handling enables better compression performance and faster query execution with the aid of bit vectors. Higher byte columns can be skipped for most queries that target on some normal values only.

Sparse encoding requires a number of bits to be allocated per frequent value, to indicate

if a value uses a frequent value or is an outlier. Compared with pure byte representation, sparse encoding further reduces the compressed data size if only there is a frequent item with a frequency greater than $1/8$, assuming 1 bit overhead per record in the bit-vector. However, the outlier scheme comes with costs that special logic is required to handle outliers during compression and query execution. So we apply factor analysis on the frequency of frequent item through micro-benchmark with controlled outlier ratio. According to our experiments, sparse encoding achieves a significant query boost compared with pure byte representation when the frequency of frequent item is greater than 90%, which is the current default threshold in BUFF.

4.2.4 Query Execution

In addition to the compression benefits, byte-oriented float splitting improves the query execution on the target float columns in the following aspects: byte-oriented data arrangement enables fast reading/writing and value parsing. Fixed-length coding helps effective data skipping to save CPU resources. The sub-column layout enables predicate push-down and progressive query execution. For efficient query execution directly on the encoded sub-columns, we apply query rewriting to decompose original query operators into a combination of independent query operators on sub-columns. Those translated query operators could be executed independently and combined to obtain the final results. If we evaluate each query operator on the sub-column one at a time, it is possible to skip some records that are determined to be qualified or disqualified already by previous sub-column query operators. Therefore, only a small amount of records need to be parsed and further checked.

Progressive Filtering

With byte-oriented float splitting compression, we can progressively evaluate the translated filter on each sub-column sequentially, and obtain the final results when all filters are finished. For a given predicate $x \text{ } OPC$ where x is the attribute stored using BUFF, OP is the operator,

and C is the operand of this predicate. Instead of decoding each record and matching the predicate, we rewrite the filtering predicate on the target column into a combination of filtering predicates on its compressed sub-columns. Assume there are k sub-columns $x_1, x_2, x_3, \dots, x_k$, we can get translated operands on those sub-columns as $C_1, C_2, C_3, \dots, C_k$ and then apply filtering on each sub-column to obtain the final results [58].

To save computation in the filtering, we can avoid evaluating some records based on our current observations. During the progressive filtering process, all the intermediate results on each sub-column are shipped by a bit-vector, similar to late materialization [101]. Two bit-vectors *RESULTS* and *TO_CHECK* are maintained to indicate the qualified rows and suspecting rows, respectively. The following filter only needs to further parse and check the suspecting rows in *TO_CHECK* bit-vector, then add qualified row numbers into *RESULTS* and suspensive row into *TO_CHECK* for the filter of the following sub-column. When the filtering level goes deeper, there will be fewer records that need to be checked. We can even early prune the filter execution in some special cases if there is no item in the *TO_CHECK* bit-vector. Even for the cases where early pruning is not happening, progressive filtering can avoid a lot of value parsing and check, especially valuable for the trailing bits sub-column where value parsing is more expensive. Progressive filtering is eligible for most filtering query executions. We describe its detailed implementation for *EQUAL* and *GREATER* predicate here. Other predicates such as *GREATER_EQUAL* or *LESS_EQUAL* can be achieved by logically combine existing predicates, *NOT_EQUAL* can be achieved by negating the *EQUAL* results.

Equality filtering predicate $x = C$ can be decomposed as $x_1 = C_1 \wedge x_2 = C_2 \wedge x_3 = C_3 \wedge \dots \wedge x_k = C_k$ by query rewriting, which intuitively indicates the equality predicate could hold if and only if equality holds on all sub-columns. This can be formulated as:

$$x = C \iff \bigwedge_{i=1}^k (x_i = C_i)$$

The predicate push-down enables efficient progressive filtering on the sub-columns where we can execute the most selective predicate to maximize the filtering benefits - to filter out disqualified records as much as possible in the early phase. For filtering on the first sub-column, qualified rows are added into *TO_CHECK* bit-vector for filtering on the following sub-columns. Data skipping is used in the following evaluation, and only rows in *TO_CHECK* are parsed and further evaluated. Disqualified rows are removed from *TO_CHECK* bit-vector. In the last sub-column, all qualified rows are added into *RESULTS* bit-vector as a final result for the filtering predicate $x = C$.

With the previous example in Section 4.2.2, if given an equality filtering predicate $x = 23.1415$, we can extract the bits of bounded range and precision 0011.001001000011100 with the aid of range and precision in the metadata. Then we can rewrite the predicate as $x_1 = 00110010 \wedge x_2 = 01000011 \wedge x_3 = 100$. On the encoded file shown in Figure 4.4, progressive filtering starts with evaluating all values in the first sub-column x_1 , and row 1 and $n - 1$ are qualified with the first predicate, so they are added to *TO_CHECK*. On the second sub-column x_2 , only row 1 and $n - 1$ need to be checked, and row $n - 1$ is disqualified, thus removed from *TO_CHECK*. On the third column x_3 , only row 1 is evaluated and qualified, thus row 1 is added into *RESULTS* bit-vector returned as a final result.

Range filtering predicate $x > C$ can be decomposed as a combination of $x_i = C_i$ and $x_i > C_i$ after query rewriting. For example, $x_1 > C_1$ is sufficient to induce $x > C$. While $x_1 = C_1$ indicates further evaluation is needed for x_2 . Then we proceed to sub-column 2 predicate evaluation. Similarly, if $x_2 > C_2$, we have $x > C$. Otherwise if $x_2 = C_2$, we proceed to sub-column 3. The evaluation is repeated until all sub-columns are processed. It can be formulated as:

$$x > C \iff \bigvee_{i=1}^k [\bigwedge_{j=1}^{i-1} (x_j = C_j) \wedge (x_i > C_i)]$$

For the first step, all records are parsed and evaluated. Only rows in *TO_CHECK* bit-

vector from the previous step are decoded and evaluated in the following evaluation step. There are two categories of records generated for each step: qualified rows (where $x_i > C_i$) are added to *RESULTS* bit-vector and unsettled rows (where $x_i = C_i$) are added to a new *TO_CHECK* bit-vector. On the last sub-column, all qualified rows are added to *RESULTS* bit-vector, which is returned as a final result for the filtering predicate $x > C$. During any filtering step, an early stop is activated when the *TO_CHECK* bit-vector is empty and all following bits are skipped. *NULL* is returned for equality filtering, and the intermediate *RESULTS* bit-vector is returned for range filtering in this case. We introduced cold start filtering cases where all records are potentially qualified. However, our filtering operators can be applied to the cases where intermediate query results are provided in the form of bit-vector, by using given bit-vector as *TO_CHECK* for start sub-column.

Progressive Aggregation

Byte-oriented float splitting compression also facilitates some aggregation queries with the aid of techniques mentioned above [91]. Query execution could be improved either by data skipping for Min/Max queries or by transforming double operations into integer ones for Sum/Average queries.

Min/Max information is stored as metadata in the compressed byte array for each compressed segment. Therefore, efficient metadata lookup is sufficient to answer those queries. However, metadata lookup is not applicable for custom query scope predefined by a filter. Our compression approach also supports fast custom scope query execution, as encoded values are independently and neatly arranged in each sub-column level. Min/Max aggregation queries perform similarly with the equality query discussed previously. Starting with the first sub-column, the query executor finds the greatest/smallest value and puts all corresponding rows into *TO_CHECK* bit-vector. The query executor then proceeds to the next sub-column, evaluates all records corresponds to rows belonging to the previous *TO_CHECK* bit-vector, and generates a new *TO_CHECK* bit-vector if applicable. The min/max value

is assembled in the end and returned as final results after all sub-columns are evaluated. A *RESULTS* bit-vector indicating the rows with min/max value is also returned along with query results. Min/Max queries are efficient with BUFF because only very limited bits are evaluated during the query execution with the aid of progressive filtering and data skipping. In most cases, only a single row is qualified and added into *TO_CHECK* bit-vector for the first sub-column. Thus, the following sub-columns evaluation is merely skipping to the target row and assembling the specific single record, which is more efficient than decoding all bits and comparing the decoded values.

Sum/Average aggregation queries are very efficient with our compression approach with no full decoding needed. The byte-oriented splitting layout enables fast access to each byte/bits component corresponding to different precision. Rather than decoding every record into float value before applying the summation operation, value is accumulated on each sub-column in the form of small integer until the final summation of the intermediate sum results from each sub-column adjusted by its corresponding exponent and base value. The final summation is returned as *Sum* result or further processed to get the *Average* result. For example, on the encoded file shown in Figure 4.4, we first sum records in each sub-column, and get $sum(x_1)$, $sum(x_2)$ and $sum(x_3)$ respectively. Then we scale the summation with its corresponding basis. In this example, basis for column x_1 is 2^{-4} , x_2 is 2^{-12} and x_3 is 2^{-15} . Then the final sum result is $sum(x) = sum(x_1)*2^{-4} + sum(x_2)*2^{-12} + sum(x_3)*2^{-15} + B \times N$ where B is the base value (min value) of this encoding block and N is the number of entries. We can further divide by N to get the average. The value is not fully decompressed to the original double value during the query execution, and all the execution is achieved with integer arithmetic, which improves the query performance.

Variable Precision Materialization and Aggregation

Reduced precision has emerged as a promising approach to improve power and performance trade-offs [31]. Customized precision originates from the fact that many applications can tol-

erate some loss in quality during computation, as in media processing and machine learning. The byte-oriented encoding layout also enables variable precision materialization in addition to full precision required by downstream analytic. For example, consider a system-agnostic query like `printf("%.2f", latitude) where AVG(temp)>90°F`. We can skip reading the trailing bytes/bits for `temp` column if we get $AVG(temp) > 90^\circ$ with several leading bytes, and only several leading bytes of `latitude` are needed for two digits precision printing. For any reduced precision value read, BUFF aligns the requested precision with the least byte boundary covering the required precision. For example, for a given compressed dataset with 4 bits for the integer and 18 bits for the fractional part (supports 5 digit precision at most), BUFF materializes 2 bytes ($\lceil(4 + 11)/8\rceil$) to support 3-digit precision query.

BUFF also provides deterministic approximate aggregation query processing on data with reduced precision as suggested by DAQ [91]. Recall previous example in Section 4.2.2, on the encoded file shown in Figure 4.4, we can get deterministic approximate aggregation result by only reading the partial data. Similar to progressive aggregation execution in Section 4.2.4, we can get an sum approximation by reading the first sub-column as $sum(x) = sum(x_1) * 2^{-4} + B \times N$ with a deterministic lower bound $= sum(x_1) * 2^{-4} + B \times N$ where all trailing bits are zeros, and an upper bound $= sum(x_1) * 2^{-4} + (B + 2^{-4}) \times N$ where trailing bits are ones. Similarly for `avg` query, we get a deterministic bound by reading first sub-column as $[sum(x_1) * 2^{-4}/N + B, sum(x_1) * 2^{-4}/N + (B + 2^{-4})]$. As we read and process more data, we can provide a more accurate result.

Data Skipping

Data skipping is another technique used to enhance query performance. After scanning a sub-column, we know some records in the remaining columns do not satisfy the predicate, and we can skip them. The byte-oriented encoding layout enables efficient data skipping in the following aspects: byte-oriented layout guarantees efficient data reading and skipping without handling any cross-boundary issue. The fixed coding length for each component

enables efficient skipping with simple bits length calculations. The trailing bits in the last component are mostly skipped based on the previous evaluation. In some cases, we can early prune the query execution and totally skip checking the following bytes and trailing bits if no records need to be further checked. Besides, data skipping enables fast records access for custom row-level scope queries, such as queries on a given time interval for time series data. Data can still be skipped efficiently without parsing the irrelevant entries.

Adaptive filtering

Different from prior works that use SIMD to accelerate query execution [57, 110, 90, 42], BUFF uses a workload-adaptive execution strategy. When SIMD is available, BUFF will use SIMD as much as possible. However, SIMD is not always the best when compared with the scalar progressive filtering. SIMD is usually good on low selective filtering, where a large percentage of records are evaluated. While progressive filtering is extremely efficient on high selective filtering tasks. For each sub-column, BUFF will choose either the filtering strategy based on density ratio $r = \#candidates/\#records$. By default, SIMD execution is used for the first sub-column unless a sparse bit-vector is provided as input. For the following sub-columns, BUFF uses SIMD execution if the density ratio is high and progressive filtering otherwise.

We perform a sensitivity analysis to find the crossover point. Progressive filtering takes more time as the density ratio increases and crosses over the constant SIMD cost at a density ratio 0.06, which is the current SIMD threshold used in BUFF's adaptive filtering. The threshold number 0.06 means progressive filtering covers many query predicates since the average item frequency for sub-column is $1/2^8 \approx 0.0039$, assuming a uniform distribution. And for other highly skewed cases, sparse encoding helps in query execution. Please be noted that, the SIMD threshold changes on different hardware.

4.3 Evaluation

Compression approaches should keep the compressed data size smaller and support fast query execution, including fast filtering and aggregation queries. This section applies the compression techniques mentioned above on various datasets covering common applications with bounded float and reports both compression and query performance. In addition, we also include end-to-end complex query evaluation on Time Series Benchmark Suite (TSBS) [5].

All experiments were performed on servers with 2 Intel(R) Xeon(R) CPUs E5-2670 v3 @ 2.30GHz, 128GB memory, 250GB HDD, and Ubuntu 18.04. All our implementations and experiments are done in Rust 1.49.0 and we use AVX2_m256i for SIMD. To use ByteSlice for floats we re-implement ByteSlice in Rust, also implement Gorilla and Sprintz in Rust. For Gzip and Snappy, we use rust flate2 and parity-snappy libraries respectively. We apply Gzip level-9 in our experiments. All our implemented baseline achieve their claimed throughput (e.g., Snappy: $250MB/s$ Gzip: $20-50MB/s$, Sprintz: $200MB/s$).

4.3.1 Datasets

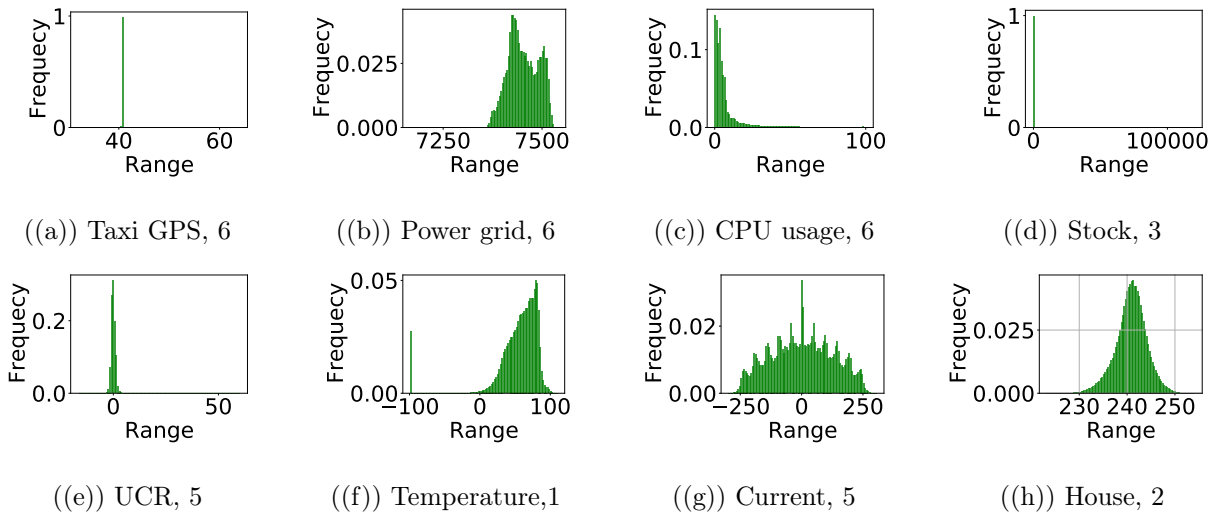


Figure 4.6: Datasets have bounded range and precision

Our dataset covers various applications: stock data [2], taxi data from NYC (GPS) [7],

DevOps monitoring (CPU) [96], city temperature (TEMP) [63], humidity [25, 26], power grid (PMU) [99], magnet motor (Current) [8], household electric power consumption (House) [3] and machine learning [82]. As most datasets are shared using a *txt* or *csv* format, we can easily extract the precision by counting the fractional digits. As we target a compression technique that would work with application code or a database, we evaluate an ingestion process that sends data in the IEEE 64-bit float format which is used by many applications and compression evaluations [108, 72, 31, 27, 33, 89]. Figure 4.6 shows the data range, distribution and scale (precision) for representative measurements of those datasets. The x-axis shows data range including all outliers if any. For GPS we use the latitude attribute. The data distribution varies on different measurements. Those datasets cover data with different precision with/without outliers. Many datasets have a very skewed distribution, where most records are clustered around a tiny area. As an example, most data points in NYC’s GPS dataset fall into the city area with coordinate between $(40.702541, -74.007347)$ and $(40.793961, -73.883324)$. The high decimal position with scale 6 guarantees sub-meter level precision of the data points. Similarly, all these measurements have limited range and bounded precision. In addition to those read world datasets, we also generate datasets from TSBS. Here, we use the IoT dataset with a scale 1000, which includes 9.1 billion data points about tracking information for trucks.

4.3.2 *Optimized Baselines*

During our evaluation, we found that the idea of bounded precision can improve Gorilla. Therefore, we provide an optimized version of Gorilla for comparison. In addition, we include fixed-point with bounded precision and ByteSlice variations for float as baselines.

Optimizing Gorilla

When we apply Gorilla on our datasets, we get a poor compression performance for most datasets. For datasets, such as CPU and Temperature, the compressed size is close or even

larger than the original size. After detailed analysis, we notice that many data points in these datasets fluctuate over value 0 or ± 2 . There are two reasons why Gorilla is sub-optimal on datasets covering those values. First, float numbers fluctuating around 0 flips the sign bit (the leading bit). Also, when floats fluctuate around an absolute value 2, they flip the leading bit of exponent bits as exponent value 1 is encoded as 10000000 in a 32-bits float. Second, high precision floating numbers make the mantissa bits sensitive to the value change, especially for those less significant bits. For these two reasons, Gorilla fails to perform efficient trailing zero compression and performs poorly on these datasets.

To improve the compression performance of vanilla Gorilla, we have two different directions to work on:

- Leverage the leading zeros by offsetting the input numbers to avoid the "pitfall" value points.
- Use bounded precision float for the input numbers to eliminate the less significant bits in the mantissa part.

We verified those two solutions on our datasets. We found the former solution improves compression ratio performance slightly but at the cost of slowing down the compression throughput significantly, since we have to offset the value before we can apply gorilla encoding. However, the latter solution significantly improves compression performance. In the following evaluation section, we include the latter Gorilla version (GorillaBD). Nevertheless, GorillaBD only improves the compression performance and can do nothing to speed up the query execution because of the complex variable encoding/decoding mechanism inherited from Gorilla. This motivates us to devise a new compression for floats.

Bounded Fixed-point

As discussed in Section 4.1.1, fixed-point is not precision aware, since it uses every bit to approximate the value no matter what precision is needed. Additionally, fixed-point is not

optimized for custom ranges, since it represents a range $-2^I \sim 2^I$. In order to include fixed-point into our evaluation, we implement bounded fixed-point (as *FIXED*) as a baseline with bounded precision applied to achieve just-enough precision and delta encoding used to offset the custom range. In *FIXED*, we scan the file and encodes the datasets with bounded fixed point representation.

ByteSlice variations for float

ByteSlice is designed for an input code of a bit-packed integer, which is not directly applicable to float numbers. In order to apply ByteSlice on floats, a conversion from the integer code is needed to meet the ByteSlice input requirement. We can either borrow from *BUFF* to extract the aligned bit representation of float (**BUFF-SLICE**) or scaling float into integer (**SCALED-SLICE**), and then bit-pack those code into ByteSlice. We compare against both ByteSlice variations.

4.3.3 Performance Overview

The radar chart of Figure 4.7 shows the overall performance of all float compression approaches above. We normalized the performance results into a range between 1 and 10 on each dimension uniformly, and then take the average of all workable datasets (*Sprintz* and *SCALED-SLICE* fail on some dataset which are not included here). ByteSlice variations do not provide a specialized aggregation solution, so we materialize before max and sum operation. For all metrics here, higher is better. As we can see, the state-of-the-art *Gorilla* approach does not work well on those datasets. *Sprintz* and *FIXED* perform well in terms of compression ratio and materialization performance on those datasets while performing poorly on other performance dimensions. ByteSlice variations are good on materialization performance but sacrifice compression ratio because of bit padding. *BUFF* performs best in terms of query and compression throughput. In addition, it handles outliers data quite well and achieves a better compression ratio on outlier-included datasets (e.g., *GPS*, *Stock*




Figure 4.7: Compression performance overview (greater is better, CR: compression ratio, Cthr: compression throughput). BUFF outperforms in query and compression throughput.

dataset). For some query types, such as min/max and high selective equality filtering, BUFF is up to $50\times$ faster than other compression baselines. General-purpose approaches (e.g., Gzip and Snappy) perform poorly in these dimensions.

4.3.4 Compression Performance

We evaluate both the compression ratio and compression throughput on our datasets and benchmarks. A good compression must be fast and effective.

Compression Ratio

Figure 4.8 shows the compression ratio for each compression approaches. Gorilla compression performs poorly on most of our datasets as there are very limited repeated or similar adjacent values. GorillaBD works better than Gorilla since many trailing bits are formatted




Figure 4.8: Compression ratio shows the BUFF are comparable to the best and always better than vanilla Gorilla (Sprintz and Scaled-slice fails on PMU).

by bounded precision float, resulting in a better compression ratio. BUFF performs similarly with Sprintz and FIXED since they leverage the bounded precision and range to map the numbers into almost the same finite encoding space. Value range and value precision determine the number of bits needed for the integer and the fractional parts. The slight compression ratio difference between Sprintz and BUFF comes from the dividing of bits, which wastes some code space when the bits are not fully used. However, this brings a huge query benefits, which we will show in Section 4.3.5. In addition, as mentioned in Section 4.2.3, BUFF can detect outliers and enables sparse encoding to compress the leading byte column further, resulting in a better compression ratio on GPS and Stock dataset compared with Sprintz and FIXED. ByteSlice variations pads the trailing bits for better query performance, so it is less effective than BUFF, FIXED, and Sprintz on most datasets. The Byte-oriented compression approaches compress data by searching common sequences at the byte level, so Byte-oriented compression performs efficiently on temperature datasets because of the low cardinality brought by low data precision but performs poorly on datasets with less repeated values.

Compression throughput

In addition to the compression ratio, we also evaluate compression throughput. We count the time, including loading data from memory, applying compression, and writing back to memory. Figure 4.9 shows compression throughput for all float compression approaches. The leftmost bar corresponds to BUFF with/without user-provided range stats. With ranges given by a user, our method can skip the range checking step in the float decomposing step and achieve higher compression throughput. We refer to it as “Fast-Buff” as is indicated by the star marker bar with the light green bar as a base. Overall, BUFF outperforms all its competitors on most datasets. Gorilla relies on *XOR* between two adjacent values, then dynamically decides the number of leading and trailing zeros, and writes the middle residue bits. Gorilla’s complicated encoding mechanism and variable encoded length limit its encoding throughput.

BUFF is faster than FIXED and Sprintz because of its fast byte-oriented writing. According to profiling, BUFF byte-oriented writing is 6% ~ 12% faster than writing bits. Sprintz for float also relies on arithmetic multiplications to quantify input float values into integer values, which causes overflow issues such that Sprintz and SCALED-SLICE fail on the PMU dataset. Overall byte-oriented compression is slower than BUFF and Sprintz. Snappy achieves the highest compression throughput on the CPU dataset, but Snappy performs poorly on the CPU dataset in the previous compression ratio experiment, where no compression is achieved at all. According to previous experiments, Gzip compression provides a higher compression ratio but uses more CPU resources than Snappy, thus lowering compression throughput. Compared with ByteSlice variations, BUFF writes fewer bits but pays more overhead on writing the last trailing bits. Therefore, we can see very similar compression throughput on most datasets for both approaches, but BUFF is faster on GPS and Stock datasets where more space is saved by sparse encoding.




Figure 4.9: Compression throughput shows BUFF performs best in most cases with user given range stats.

4.3.5 Query Performance




Figure 4.10: Low selective equality filter




Figure 4.11: Low selective range filter




Figure 4.12: High selective equal filter

We also evaluate decompression and query performance. We measure execution time, including reading the data block from memory, parsing data, and evaluating the query. For all queries evaluated, we leverage each compression approach’s characteristics to speed up its query as much as possible. Techniques, such as query rewriting, early pruning, and progressive filtering, are used for each approach whenever feasible, as discussed in each experiment section. We use the most frequent value (lower selectivity) in the corresponding dataset as a predicate operand to conduct a fair comparison for all equality filtering and range filtering. Progressive filtering is less efficient with low selectivity predicate as fewer

values are skipped during query execution. We also run high selectivity queries to show BUFF’s potential. SCALED-SLICE and Sprintz bars for PMU dataset are missing as they fail on PMU dataset because of an overflow issue.

Filtering

We include filtering performance evaluation with query rewriting, early pruning, progressive filtering, data skipping and SIMD execution enabled when possible. For BUFF, instead of materializing all records before applying the filtering predicate, we rewrite the query into query predicates combination on all sub-columns, then apply the translated query predicate on each sub-column sequentially. We also use adaptive filtering and data skipping with the aid of intermediate bit-vector results. For Sprintz and FIXED, we apply query rewriting techniques to avoid decompression overhead. For ByteSlice variations, we adopt their SIMD solution for filtering queries.

Our first experiment set shows low selective equality filtering performance. In Figure 4.10, BUFF achieves outstanding performance compared to all the other competitors. This speedup is mainly from adaptive filtering and data skipping. As we execute the predicate on the first sub-column, only a tiny part of the input records are qualified and need to be further checked in the subsequent sub-column predicates. For most datasets, less than 5% of entries are qualified and need to be further checked after filtering on the first sub-column with BUFF. However, for GPS and Stock datasets, more than 90% entries are still qualified because of code space inflated by outliers. As we hit the frequent value in this case, we can skip decoding the first sub-column by directly using a bit-vector from sparse encoding, which speeds up the query. Sprintz and FIXED filtering use query rewriting to avoid record decompression as well. Compressed bits are extracted and compared with the translated to the target directly by using integer arithmetic. ByteSlice variations are always the second to the best approaches due to the SIMD speedup and early stopping. BUFF has a better data and code locality, as it works on the same predicate for the current sub-column. While ByteSlice needs to load

all sub-columns and jump to different sub-column with corresponding predicate evaluation in each iteration. Additionally, always using SIMD can slow down ByteSlice’s throughput as a single match, resulting in SIMD loading and evaluating for all adjacent records within the same SIMD word length. In BUFF, data skipping can efficiently jump to the target record and avoid unnecessary data loading. Except for these aforementioned approaches, decompression is required before predicate evaluation on float numbers. General-purpose byte-oriented compression approaches are sequential algorithms and need to decompress the last bits before the original data can be accessed. Gorilla encodes value with different bits and variable bits representation, which impedes query translation and data skipping. These bring a high latency for filtering the compressed data with those methods.

Figure 4.11 shows the low selective *greater than* range filtering experiments with the same operand. As we can see, the range filtering performs similarly to equality filtering but with smaller throughput overall. The selectivity of the range query is much lower than previous equality filtering. There are more qualified records in this case, which triggers the *IF* branch that edits the bitmap more frequently. Thus filtering throughput is deteriorated compared with previous equality filtering experiments. Similarly, ByteSlice variations is slower than equality filtering because of more comparison and less early stop. However, BUFF still performs best. Similar to equality filtering, all other approaches require full decompression. As we can see from the figure, their filtering performance is proportional to their materialization performance, which we will discuss later in Section 4.3.5. Figure 4.12 shows high selective equality queries that filter out majority of records. BUFF is more efficient on the high selective query as more bits are skipped by progressive filtering. Overall, BUFF achieves 20× for average query speedup compared with the second-best float compression approach. BUFF has similar trends on high selective range queries, and it achieves 16× on average query speedup compared to the second-best float compression method (we omit the figure).

Aggregation

An aggregation query is another frequent query for numeric data analysis. We evaluate the query performance of *min/max* and *sum* for all compression approaches.




Figure 4.13: Max throughput

Figure 4.13 shows *min/max* query throughput. The query time includes bit parsing and query execution. Even though we can quickly extract *min/max* value from data statistics in the metadata block for BUFF, we force aggregation query execution on compressed data for a fair comparison, and this is necessary for aggregation with a filter. In addition to this setting, we enable all the techniques used in previous experiments. For *min/max* aggregation, we can use progressive filtering and data skipping for BUFF. For FIXED and SPRINTZ, we find the *min/max* integer code directly then convert this value into float numbers. For other approaches, float numbers are decompressed for evaluation. The result shows that BUFF achieves 22× on average query speedup than the second-best float compression approaches SPRINTZ or FIXED. The *min/max* value is usually with a low frequency such that only a tiny portion of entries are added to the bit-vector and need to be checked in the following sub-columns. This makes BUFF extremely efficient for *min/max* aggregation queries since




Figure 4.14: Sum throughput

a huge amount of entries (bits) are skipped during the evaluation step.

The sum query performs similarly to materialization as all compressed data has to be decompressed, and all techniques mentioned previously are ineligible in this scenario. Figure 4.14 shows the *sum* query throughput on the same dataset. BUFF outperforms all its competitors. For SPRINTZ, FIXED and BUFF, we use lazy materialization to avoid the overhead of full decompression for each entry. The sum intermediate result is calculated with an integer representation for SPRINTZ and FIXED, then converted back to float numbers at the very end. This is similar to sum operator with BUFF, where the sum is calculated on each sub-column with its integer representation respectively and concatenated into a final float-point result in the end. The performance difference between those approaches comes from the more efficient reading and parsing of BUFF and no data dependency between adjacent numbers. All other approaches perform poorly since full decompression is inevitable.




Figure 4.15: Materialization throughput

Full Materialization

Materialization performance varies a lot for each compression approach. Overall, materialization overhead is comprised of bit parsing time and assemble time. Fixed-length encoding helps for bit parsing. As is shown in Figure 4.15, SPRINTZ and FIXED perform better than BUFF as they read single fixed-length data chunk during materialization, which is faster than Gorilla that reads variable-length bits for each record, and BUFF where multiple data chunks need to be fetched as requested, then parsed and concatenated into a float number. However, BUFF is still comparable with SPRINTZ and FIXED thanks to the efficient reading and parsing of BUFF. ByteSlice variations are faster than BUFF because of its fast byte reading and decoding benefits. Since BUFF bit extraction is more efficient than the scaled version on decompression, BUFF-SLICE is always faster than SCALED-SLICE. Gzip is the slowest among all candidates since it has to decode Huffman encoding before restoring the recurring sub-sequences. Snappy achieves a better performance than Gzip because it skips Huffman encoding for higher throughput.

Materialization and Aggregation with Requested Precision




Figure 4.16: Materialization and aggregation with targeted precision boosts throughput by only reading the leading bytes.

We also show the versatility of BUFF on variable precision materialization and aggregation support, which are especially useful when users only care about limited precision in formatted output and fast estimation of aggregation queries. Figure 4.16 shows the materialization and aggregation with requested precision on the UCR dataset. The red line-based right y-axis shows the bits needed to support the requested precision. The query performance linearly increases as bytes read decreases. A single expensive bit read can significantly slow down the query performance, as we can see from the first set of results on full data with 3 bytes and 1 bits read.

1-Nearest Neighbor

One nearest neighbor (1NN) is a popular machine learning tasks where an entry X is assigned to the nearest neighbor's class. In addition to SQL tasks, We also include 1NN experiments to explore the impacts of BUFF on 1NN accuracy.

We run 1NN task on the UCR dataset collection with different BUFF precision for the input float numbers, and report the 1NN accuracy result for all 128 datasets. For better

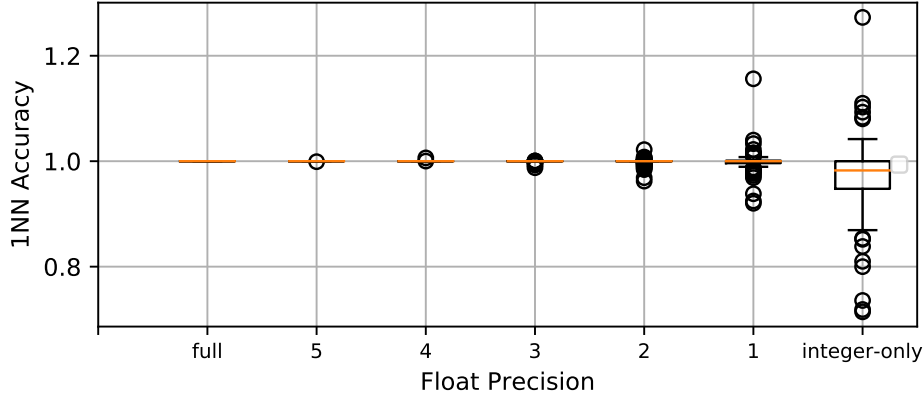


Figure 4.17: Relative 1NN accuracy on 128 UCR datasets shows reduced precision has a limited impact on accuracy.

representation, the accuracy number for each dataset is normalized by its corresponding accuracy of full precision input. As is shown in Figure 4.17, relative accuracy decreases gracefully as the BUFF precision goes down. Nevertheless, we still achieve outstanding performance even with reduced precision greater than one. This experiment shows BUFF has a good potential of saving storage and network transmission cost for some machine learning tasks without loss of precision.

4.3.6 Benchmark Evaluation

In this set of experiments, we show complex query performance with TSBS. On TSBS, we generate “last-loc”, “low-fuel” and “high-load” queries, which respectively include projection on longitude and latitude attributes (project), filtering on a fuel_state attribute containing single sub-column with BUFF encoding (range-single), and filtering on current load attribute (range). In addition to the float-related operations, those end-to-end queries also include any filter, join, sorting on timestamp, string, or integer attributes. The runtime for these columns is independent of float compression approaches, so we take do not show them and only focus analysis on query runtime of float attributes. For Gorilla, the float-associated cost makes up 54% to 76% of the total runtime. However, BUFF can diminish those costs to 1 or 2 orders of magnitude smaller on those queries. Figure 4.18 highlights the query runtime




Figure 4.18: TSBS query runtime for float attributes.

on float attributes, BUFF always performs the best over all other baseline and is faster than ByteSlice variations on TSBS range filter queries since adaptive filtering chooses progressive filtering execution for the highly selective filter.

4.4 Conclusions

This paper proposes BUFF, a novel decomposed float compression for low precision floating-point data generated everywhere and every day. BUFF uses “just-enough” precision for a given dataset, and leverages the data distribution feature (range, frequent value) to simplify its encoding mechanism and enhance the compression performance. The compression supports efficient in-situ adaptive query on encoded data directly. In addition, BUFF provides fast aggregation and materialization for different precision levels. BUFF’s precision bounded technique is also applicable to the state-of-the-art Gorilla method and improves Gorilla’s compression performance significantly. Our evaluation shows BUFF achieves fast compression and query execution, while keeps comparable compression ratio to the best approaches. Future work includes studying the performance of Buff for common downstream complex analytical tasks, such as classification [87], clustering [84, 85], and anomaly detection [23].

CHAPTER 5

ADAEDGE: A DYNAMIC COMPRESSION SELECTION FRAMEWORK FOR RESOURCE CONSTRAINED DEVICES

The Internet of Things (IoT) refers to the billions of physical devices worldwide connected to the internet, all collecting and sharing data. IoT devices generate vast amounts of data from various sensors every day. Data systems typically employ different compression approaches to reduce on-disk or in-memory data sizes and improve bandwidth and storage utilization on the centralized cloud. As edge devices gain more computational power and popularity, supporting compression, query, and analysis on the edge side is a good option to reduce the IoT system's cost. Pushing compression down to the edge makes ingested data pre-compressed, which benefits many system aspects. It minimizes the data that needs to transfer, saves network bandwidth, saves the edge or cloud resources for storing data, and typically improves query performance. However, the edge is usually resource-constrained compared with cloud servers in terms of network bandwidth, storage, and computation power. Those limited resources bring new challenges and constraints to compression design and selection.

The IoT data from a signal sensor can statistically change over time, so it is difficult for a compression approach to always have efficient and effective compression over the data stream. Modern data systems support a variety of lossless and lossy compression approaches to process the data for mixed workloads. Different compression approaches have diverse compression performance and even impact the workload accuracy differently in the cases of lossy compression. Because of dynamic data features and various workloads across extended IoT applications, it is impossible to provide a one-size-fits-all compression solution for all kinds of data and tasks. Therefore, an adaptive lossless and lossy compression selection strategy is necessary to select the optimal compression for the incoming IoT data based on the data statistics, query workloads, and limited resources available on the host.

Compression selection is an optimization task accounting for the data features, hardware

specifics, and workload information. Compression selection can optimize for different performance targets, including space usage, compression throughput, and query performance. Abadi decision tree [10] is one of the first projects to explore compression selection for in-situ query execution on compressed data. It introduces a rule-based encoding selection approach that relies on the global knowledge of the dataset features. LEA [28] is a compression selection framework that builds difference regression models to predict the compression and query performance, then selects the optimal compression according to the optimization target. CodecDB [59] builds a neural network model to predict optimal compression based on the collected data features. It also provides specialized operators which can operate on encoded columns directly. Hyrise [22] uses separate models to predict compressed size and query operator performance, then use linear programming or greedy heuristics to get the best encoding for each data chunk. Beside these projects, Open-source columnar formats such as Apache Parquet [18], ORC [17] and Arrow Feather [16], and some commercial data systems such as Vertica [77] and MySQL [37] choose to hard-coded compression selection based on the column data type. However, those conventional compression selection focuses mainly on lossless compression approaches and optimizes for either storage or SQL query performance. Besides, they do not consider the more limitations and constraints for edge devices in the IoT systems, such as the cases where an offline edge device needs to ingest and hold more data than its storage capacity temporarily.

Edge devices usually have more constraints and limitations from the hardware and system environment, including limited storage, bandwidth, and computation power. Traditional lossless compression approaches come with a compression limitation defined by the entropy of the data. So it is impossible to always keep 100% accurate data representation with limited system storage resources. Lossy compression approaches are widely used by many time-series databases to control storage consumption. Some systems directly remove data ingested earlier than a given time or exceeding a storage threshold. RRDtool [105] and InfluxDB [4] remove old data exceeding a certain period to take back space for new data.

ModelarDB [55] trades the query accuracy for storage space by selecting the best linear model to approximate the data with user-defined error bounds according to the storage budget. SummaryStore [11] reclaims the storage space by replacing the data with aggregate summaries with a specific compression ratio. TVStore [13] uses a time-varying compression framework to bound storage. These projects mainly rely on a single lossy compression approach to compress the data according to the storage budget. A selection strategy among different lossy compression approaches is not discussed. Furthermore, previous works mainly focused on the lossy impacts for the accuracy of dashboard queries and some SQL operations, while widely used machine learning tasks are not well studied as an optimization target.

Different lossy compression approaches may impact the accuracy and efficiency of different workloads. There is no one-size-fits-all compression solution for IoT systems with varying data statistics, workloads, and hardware constraints. We, therefore, propose AdaEdge as a hardware-conscious encoding selection framework for resource-constrained devices. AdaEdge is a best effort compression selection framework with the following contribution:

- Multi-armed bandit assisted lossy and lossless compression selection depending on the hardware constraints and workload target.
- Online and offline mode compression selection solutions for resource-constrained devices to conserve mission-critical information.
- Multiple optimization targets supported, including compression ratio, throughput, aggregation queries and machine learning task accuracy.
- LRU-based compression policy to compress the data based on its importance.

5.1 Background

In this section, we first introduce several lossy and lossless compression approaches. We then discuss the compression selection problem in IoT system setting and frame the compression selection problem as a multi-armed bandit problem. Finally, we introduce Multi-Armed Bandit Problem, its variations, and applications.

5.1.1 Lossless and Lossy Compression

There are two types of compression: lossless compression where the original data can be fully recovered from compressed data, and lossy compression where the original data cannot be recovered. We discussed many lossless compression approaches in Section 4.1.2, including byte-compression techniques, such as Gzip and Snappy, and lightweight encodings, such as dictionary encoding, Gorilla [89], Sprintz [21], and BUFF. Lossless compression is widely used to persist historical data and unstructured media data. However, lossless compression capabilities are always bounded by the entropy of the input data. The entropy conveyed by a piece of information defines the minimal bits required to represent the corresponding information losslessly. This implies a lower bound for a 100% accurate representation of the given information. When the system does not provide enough storage budget, all the existing data compression selection frameworks fail to get an encoding solution. However, these scenarios are quite typical in IoT systems. The dynamic resource and hardware constraints in IoT systems may require extremely aggressive compression to meet the limited constraints. In this case, if we still want to keep data smaller, we need to pay the cost of partial information loss while we want to keep the minimal impact on the task accuracy.

In order to handle such scenarios, we include some basic lossy compression in AdaEdge. Lossy compression can achieve a much higher compression ratio than lossless compression by giving up partial information. To better serve our use case, the lossy compression used by AdaEdge should be easily configurable with any given compression ratio and easy for

recoding without much effort for decompression. Our AdaEdge framework currently supports the following lossy compression approaches.

BUFF is designed to be a lossless compression approach, but it can be used as lossy compression if we compromise on the precision of the float value. In AdaEdge, we include the BUFF lossy version as a lossy compression option, as noted by BUFF-lossy. BUFF-lossy achieves more aggressive compression by ignoring insignificant bits. This is the least significant change to the given input value, which is beneficial to the tree-based machine learning tasks, which are sensitive to the feature value change.

Piecewise Linear Approximation (PLA) [97] is a linear generalization method that enables linear approximation for the values under each data segments, and the resource budget bounds the number of segments. Largest Triangle Three Buckets (LTTB) is a variant of the widely accepted and used Visvalingam-Whyatt (VW) algorithm [106]. LTTB is used as the line generalization method of TVStore [12] and TimeScaleDB [6]. As compared to other line generalization algorithms, LTTB has much lower complexity. It can compress data in almost a single pass while preserving visually important points like its counterparts. Simplicity and data preservation are two key features of LTTB. The linear generalization method can keep the visual similarity of the signal, so it is good for dashboard queries.

Piecewise Aggregate Approximation (PAA) [113, 62] is designed to reduce the dimensionality for time series data mining by representing time series as mean values of segments. While the simple dimensionality reduction can be used as a form of lossy compression. PAA saves space by keeping the mean values of an equal-sized window, and it's configurable to achieve different levels of lossy approximation by adjusting the window size. Since PAA keeps the mean values for each window, it is usually good at preserving Sum and Avg aggregation query accuracy.

Fourier transform [39] converts a signal from its original space or time domain to a representation in the frequency domain and vice versa. Fast Fourier Transform (FFT) is an algorithm that computes the discrete Fourier transform of a sequence, or its inverse. Once

the original sequence is transformed into its composite frequency parts, the less important frequencies could be ignored from the data without distorting the original sequence a lot, thus reducing file sizes. FFT achieves more aggressive compression by ignoring more components with high frequency. It is usually good for the distance evaluation of high-dimensional data.

RRD-sample. We also include a random sampling method to simulate the RRDTTool compression logic which bounds storage by deleting data when the storage quota is reached. So instead of deleting the old segment, AdaEdge keeps one random value from this segment for late queries and duplicates the value to a whole segment for the relevant workload when necessary.

Overall, all the lossy compression in our framework are easily configurable to achieve the desired compression ratio. And easy to adjust the level of compression for the given data. Given storage budget, AdaEdge is able to choose from Lossy compression sets including BUFF-lossy, PAA, FFT, PLA, RRD-sample for aggregation queries and machine learning tasks such as Decision tree, Random forest, KNN, K-means.

5.1.2 *Compression selection*

Data compression is crucial for many data systems and applications. Compression helps save network and storage resources, which is extremely important for some resource-constrained devices and systems. Depending on the system hardware, the compression has to work under a series of resource limitations while guaranteeing the accuracy of tasks. Ideally, lossless compression is qualified for the given hardware limitations, and the data can be fully recovered before sending to the downstream applications. At the same time, lossy compression approaches are needed when there are more harsh storage constraints. The compression should be more aggressive to save system resources while minimally affecting task accuracy. The lossless compression selection will mainly focus on the compressed data size and compression throughput, while the lossy compression selection will be more focus on the accuracy loss of the workload.

The resource constraints derive from many aspects of the system, including data ingested rate, limited memory, storage and network bandwidth. The constraints can also come from the requirements of the downstream application and workload: the compression should help speed up the workload and diminish the workload accuracy loss. Those constraints impose requirements on the compression ratio, compression throughput, query throughput, and task accuracy.

We formulate compression selection as an optimization problem for the resource-constrained system. The system constraints are denoted as I_{bgt} for given signal ingestion rate (assuming no control on the signal generating rate), B_{bgt} for network bandwidth, and S_{bgt} for local storage capacity. The data sequence is organized in the form of segments with a fixed number of consecutive data sequences (original size denoted as U). There is only one compression scheme is selected for each segment. $v_{ij} = 0, 1$ indicates whether compression c_j is selected for segment x_i , then we have $\sum_{j=1}^k v_{ij} = 1$ for any given segment i assuming k compression approaches in the compression candidate set C . In terms of the compression performance, we also use r_{ij} to indicate the estimated compression ratio of compression c_j on segment x_i . t_{ij} indicates the estimated compression throughput of compression c_j on segment x_i . e_{ij} indicates the estimated query error of lossy compression c_j on segment x_i . For n segments, we can get the total compressed size:

$$S = U \times \sum_i^n \sum_j^k (r_{ij} * v_{ij})$$

And we set the average egress rate as:

$$B = \frac{I_{bgt}}{n} \times \sum_i^n \sum_j^k (r_{ij} * v_{ij})$$

So when we optimize for the given network bandwidth B_{bgt} , we formulate the lossless com-

pression selection as

$$\arg \min_V S(V) \text{ s.t. } I_{comp}(V) \geq I_{bgt} \text{ and } B(V) \leq B_{bgt}$$

where V is the compression assignment for each segment comprised by v_{ij} and $I_{comp}(V)$ indicates the compression throughput with given V . The lossless compression selection should handle the given ingested signals and generate compressed data with egress rate under the network capacity. If no solution for this optimization problem is feasible, a lossy compression selection problem is needed to meet the harsh constraints. To optimize for the given network bandwidth B_{bgt} , we formulate the lossy compression selection as

$$\arg \min_V E_T(V) \text{ s.t. } I_{comp}(V) \geq I_{bgt} \text{ and } B(V) \leq B_{bgt}$$

where $E_T(V)$ is the accuracy loss for the given task T with the given compression assignment V . The goal switches to minimizing the task accuracy loss within the system constraints.

For the systems without a network to constantly egress the data out, the compression selection target switches to keep ingesting the new data while minimizing the task accuracy loss when shedding the data. We formulate the compression selection problem as

$$\arg \min_V E_T(V) \text{ s.t. } I_{comp}(V) \geq I_{bgt} \text{ and } S(V) \leq S_{bgt}$$

where we keep evolving the data to meet the system's limited storage budget while ensuring the selected compression approaches are capable of handling the given signal.

In AdaEdge, the input signals are organized into data segments with a fixed number of data points. The compression thread keeps compressing the ingested segments to meet the system constraints. As different approaches compress the segments, we have a better estimation of each approach, which in turn serves the system with better compression suggestions. Therefore, we formulate the compression selection for IoT data as a multi-armed

bandit (MAB) problem. The compression selection component should be able to automatically select the optimal compression approach, given the workload, data arrival rates, and resource capacity. The compression selection requires a lightweight framework to efficiently summarize the past compression performance and provide compression suggestions based on the latest observation from system feedback. So it is a perfect fit for the multi-armed bandit (MAB) problem.

5.1.3 Multi-Armed Bandit Problem

The Multi-Armed Bandit Problem (MAB), also called K-armed bandit problem [103, 75], is a problem where an agent repeatedly makes a choice among a fixed number of different options, or actions which return agent rewards. Each action has a different probability distribution that determines how much reward is given, and these distributions are unknown to the agent beforehand. The objective is to maximize the expected total reward over some time period, which is called total expected return.

Assuming we have a bandit with k arms, each of the k arms is associated with an expected reward $q(a)$ given that arm action a is selected. We denote the arm selected at time step t as A_t , its corresponding reward as R_t . If action a is selected at time step t , we have the expected reward

$$q(a) = \mathbb{E}[R_t | A_t = a]$$

Since we do not have the reward value for each action beforehand, we need a good estimation of it. We denote the reward estimation for action a at time step t as $Q_t(a)$. The closer $Q_t(a)$ is to $q(a)$, the better action bandit can make. In the ideal case where estimation $Q_t(a)$ is 100% accurate, we would always select the action with the highest value to maximize the total return. The action chosen at time step t can be formulated as:

$$A_t = \arg \max_a Q_t(a)$$

This is the intuition of greedy action, where the agent always selects the arm action with the highest reward estimation. When the agent follows the greedy suggestion, it exploits the bandit system's current knowledge. While purely exploiting will only bias the bandit system on the former knowledge, some proper exploring action is needed. The exploring action is basically trying on the non-greedy actions, which helps to improve the reward estimation of the non-greedy actions. Exploitation action is the right thing to do to maximize the expected reward in one step, but exploration may produce a greater total reward in the long run.

There are many sophisticated variations for balancing exploration and exploitation for particular mathematical formulations of the multi-armed bandit problem. ϵ -greedy is a basic solution of balance exploration and exploitation for MAB problem, with ϵ possibility to choose the non-greedy action. The Optimistic ϵ -Greedy Algorithm is a simple modification by setting the initial action estimates to high values in a regular greedy algorithm to push it explores the whole candidate action searching for optimal action. Rather than performing exploration by simply selecting an arbitrary non-greedy action, chosen with a probability ϵ that remains constant, the Upper Confidence Bound (UCB) algorithm changes its exploration-exploitation balance as it gathers more knowledge of the environment. It moves from being primarily focused on exploration, when actions that have been tried the least are preferred, to instead concentrating on exploitation, selecting the action with the highest estimated reward. There are many other variations for bandit algorithms, such as Gradient Bandit [43] and Contextual Bandits [32], which are orthogonal to this work.

5.2 AdaEdge Overview

AdaEdge is an adaptive hardware-conscious compression selection framework for resource-constrained devices. AdaEdge can select the optimal compression according to system resource limitations, incoming data statistics, and query workload.

5.2.1 *AdaEdge Framework Constraints*

We first introduce the basic constraints of the AdaEdge framework. As we can see from Figure 5.1, There are several limitations and constraints in AdaEdge. The first one is the signal generation rate from the sensor, which ranges from zero to millions of data points per second. In AdaEdge, we assume no back pressure control to the sensors regarding the data generation rate, so the ingestion rate is a hard constraint for our framework. AdaEdge should be able to ingest all incoming data and process with proper compression, or downsampling when necessary. In our compression selection step, all selected compression must be able to handle the given ingestion rate from the sensors.

The network bandwidth is another hard constraint for AdaEdge. While broadband networks are common, the low bandwidth networks are still popular and they play critical roles in the modern IoT ecosystem. The network bandwidth varies depending on the network connection technology. In practice, the bandwidth changes for a cellular network from 0.01 Mbps to 200 Mbps. In addition, the network bandwidth varies significantly in cases of network congestion or edge device movement. Network disconnection is typical for IoT edge devices in some environments, such as the agriculture, aerospace exploration, and mining industry. AdaEdge is able to handle all different network connection cases and provide proper data ingestion and compression solutions accordingly.

An edge node has a very limited storage space compared to an advanced cloud server. Storage devices are carefully accessed on the edge device. For the sake of device lifetime and power consumption, some industrial systems are designed to avoid using storage devices for

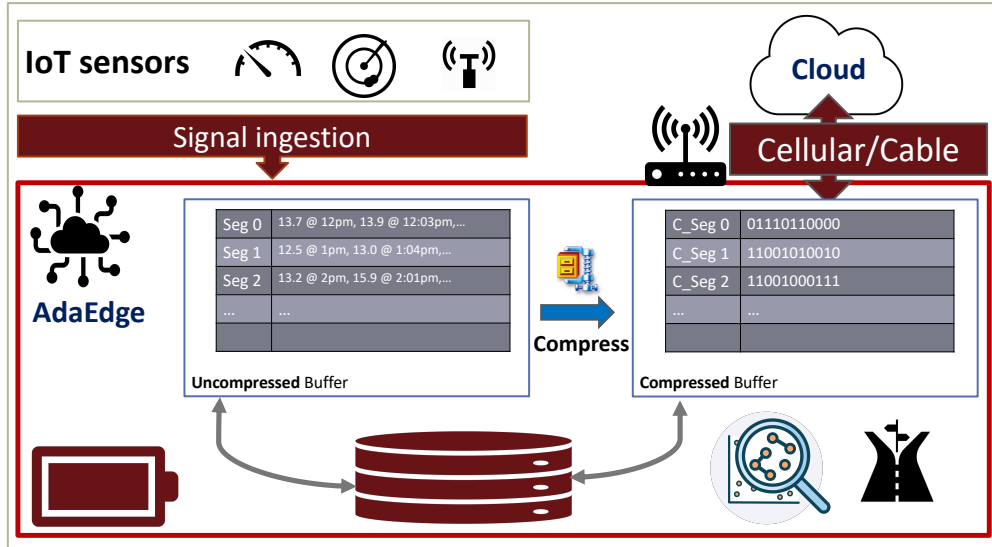


Figure 5.1: AdaEdge overview

daily operations but only save the storage devices for a particular use during offline cases.

The last limited resource on edge is power. Many edge devices are designed for demanding environments and deployed in remote areas without a sustainable power supply. When the battery becomes the only power supply for a given edge node, any task running on the edge should consider the power consumption. In AdaEdge, we mainly focus on the network and storage constraints, and we leave the power constraints as a future work.

Based on those constraints, we classify the IoT use case into two modes. In terms of network connection, we introduce online and offline modes for AdaEdge.

Online mode

The online mode views the edge node as a data hub, and the goal is to transfer the signal data as much as possible within the given constraints, including signal rate and network bandwidth. For the online mode, AdaEdge keeps all data ingested and compressed in memory before sending those segments out through network protocol. The selected compression should be able to handle the given signal generation rate, and the compressed file size should be within the network transmission capacity.

Figure 5.2 show an example with 4 million data points generated per second. The line shows the ingestion rate (the size of data generated per second), and the bars are the in-




Figure 5.2: compression should be able to handle the given signal generation rate. An example shows a signal with 4 million data points generated per second.

gestion rate (the size of data compressed per second at full speed) for different compression approaches. We can see from the figure that most compression methods are qualified to compress the signal example except for Gzip and Kernel methods, which are usually slow in compressing the data. In addition to the ingestion rate, we also investigate the compression egress performance. As we can see from Figure 5.3, the first bar shows the egress rate of the original 4GHz signal without compression. The other bars show the compression egress rate with different compression applied. A lower egress rate on the same signal means better compression performance. Thus less network bandwidth is needed to transmit the compressed data. Without compression before data transmission, it is impossible to transfer the ingested data to the cloud with the given network setting. But assuming we use 4G here, then many lossless compression such as Sprintz, BUFF, dictionary encoding, and lossy compression can send out the compressed data under the 4G network. However, if with the 3G network, there is no qualified lossless compression. So depending on the query workload, we need to choose corresponding lossy compression to minimize the query error rate.

For the online mode, we keep all data ingested and compressed in memory before we send




Figure 5.3: The compressed file size should be within the network transmission capacity. Bars show the compression egress rate on a 4GHz signal with double type. Lines show the network transmission capacity per second (MB/s).

out those segments through a network protocol. We omit saving data to the local storage to save storage space for future offline cases, save power and prolong the lifespan of the storage device. We choose the best lossless compression by default. In the particular case of no qualified lossless compression, we chose lossy approaches optimized for the target workload.

Offline mode

The edge node serves as a local computation and storage node for scenarios with poor network or no connection. The goal switches to keep ingested data as much as possible for data offloading if a future network connection is expected. If information loss is inevitable because of a tight storage budget, then it should use the lossy compression, which has the smallest negative impact on the workload. We define the lossy representation with a minimal negative impact on the workload as mission-critical information for a target workload.

As is shown in Figure 5.4, the offline mode performs differently from the online mode, which always optimizes for the compressed size. Instead, load shedding is needed when local storage runs out of space. For all ingestion data, we use the local disk as temporary storage. When there is not enough space, we do more aggressive compression (even lossy compression) on less valuable or less informative segments instead of removing some segments.

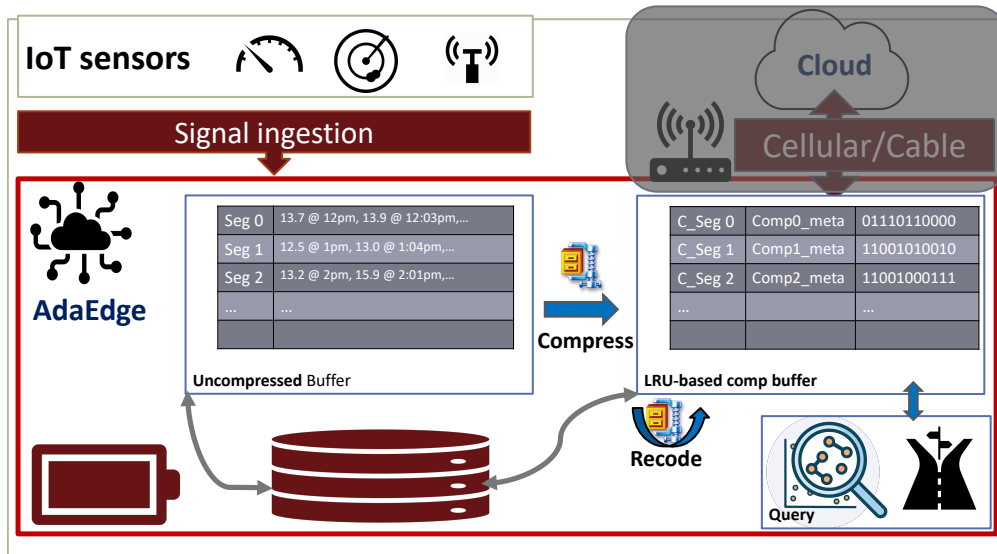


Figure 5.4: AdaEdge keeps evolving the ingested data with the given storage budget for offline mode devices.

The informativeness can be measured by the query usage of the given segment. A query counter for each segment is a straightforward solution for measuring the informativeness of segments. However, the informativeness should also reflect the contribution of each segment to the query. For example, segment A with 1% qualified entries should be less informative than segment B with 99%. The ratio of qualified entries for a segment is defined by the number of entries qualified for a given query divided by the total number of entries in the current segment. Since there are different definitions of informativeness for a segment, AdaEdge builds a dedicated segment management component with standard *GET* and *PUT* APIs for different policies. AdaEdge uses a LRU based compression policy by default, which is discussed in Section 5.2.5.

5.2.2 AdaEdge Workflow

AdaEdge allows the collection and aggregation of data from multiple device clients. The framework can be configured to accept multiple signals corresponding to a metric that the sensor produces, such as temperature and humidity. AdaEdge ingests data points generated by the remote sensor clients, caches them into fixed-size arrays, attaches a timestamp

for the segment, and pushes it to the uncompressed buffer. As the uncompressed buffer is being filled, compression threads offload data from the uncompressed buffer and process them. The compression selection threads are adaptively configured to use different compression approaches based on storage capacity, network bandwidth, ingestion rate, and specified analytical tasks. If the uncompressed buffer exceeds its capacity, which may happen when the ingestion rate exceeds the compression speed, the data is flushed to the disk. The compression threads push the compressed data into a compressed buffer pool, which can also flush to disk. AdaEdge can execute queries or analyses (e.g., aggregation queries, clustering) over the compressed data or the raw time-series segments in the uncompressed buffer. Each segment in the framework is associated with metadata describing its compression configurations, which can be used for the codec to do further compression or decompression for the downstream workload.

AdaEdge currently supports byte-compression techniques, such as Gzip and Snappy, and lightweight encodings, such as dictionary encoding, Gorilla [89], Sprintz [21], and BUFF [73] for numeric data. AdaEdge also supports specialized lossy time-series representation methods (i.e., lossy compression), such as Piecewise Aggregate Approximation (PAA) [113, 62], Fourier transform [39], Piecewise Linear Approximation (PLA) [97], and BUFF-lossy. These approaches differ in terms of compression ratio, throughput, query efficiency, and workload accuracy. There is no one size fits all approach for any time series or task. With system resource limitations defined by edge devices, AdaEdge is able to switch the compression approach according to the optimization target adaptively. AdaEdge currently supports optimization targets, including space (data compression ratio), time (compression runtime), and accuracy (machine learning model and aggregation query accuracy).

The compression selection component in AdaEdge is able to automatically select the compression approach, given the workload, data arrival rates, and resource capacity. AdaEdge compression selection requires a lightweight framework to efficiently summarize the past compression performance and provide compression suggestions based on the latest obser-

vation from system feedback. The compression selection component of AdaEdge is built based on the multi-armed bandit (MAB) problem. As the signal is ingested into the system, AdaEdge compresses the uncompressed segments and puts them into the compressed buffer. Once compression is finished for segments, AdaEdge runs an evaluation process to monitor the compression ratio, compression throughput, target workload accuracy, Etc. This information is used for the MAB component to estimate the potential reward by applying each compression candidate. Recall the MAB greedy algorithm logic, the compression selection for AdaEdge is selecting the compression arm with the highest estimated value of the optimization target. So the $Q_t(a)$ from section 5.1.3 will be the optimization target estimation for compression action a at time step t .

In AdaEdge, there could be more than one MAB instances for the AdaEdge compression selection component, depending on the use case. For AdaEdge online mode, a target compression ratio can be derived from the existing system constraints. For example, an edge node ingests the given signal containing double-precision floating-point number (8 bytes per value) at ingestion rate I points per second and network bandwidth B bps. In order to send the ingested signal out through the given hardware, the derived target compression ratio is $R = B/(8 * I)$ if we ignore the network packet header overhead for each network layer. With this target compression ratio R , AdaEdge applies lossless compression on ingested segments by default and falls back to lossy compression if R is not achievable by lossless compression. For the lossless encoding selection, there is no impact on the task accuracy loss, so the MAB optimization target is the compressed size only. When finished compressing the segment, the compressed segment size is updated to the corresponding arm in the MAB component, which is used to improve the reward estimation of MAB. The compression selection decision is made based on those estimations. When lossless compression doesn't work for the target compression ratio R , another MAB instance is initialized for the lossy compression selection. Lossy compression means the decompressed data is not 100% restoration of the original data; thus, it negatively impacts the workload accuracy. All lossy compression approaches cur-

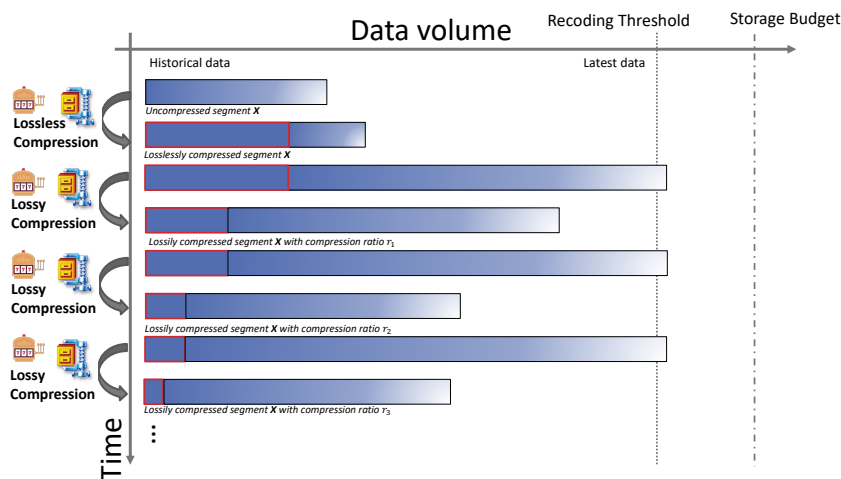


Figure 5.5: AdaEdge keeps ingesting the latest data while applies different levels of compression on the early segments to free space for incoming data. The red rectangle denotes an arbitrary segment in AdaEdge, with the rectangle length showing its current size.

rently supported by AdaEdge are easily configurable for any given target compression ratio, so the compressed segment from different lossy compression has the same size, which meets the target compression ratio. Therefore, the compressed file size is no longer critical for lossy compression, while task accuracy becomes the most crucial factor. The workload accuracy could be the MAB optimization target. AdaEdge periodically checks the optimization target for all lossy compressed segments and updates the latest reward value for the corresponding arm to support better future action suggestions.

The offline mode for AdaEdge is more complicated than the online mode, as there is no clear target compression ratio derived from the system constraints. An actual use case could be an offline edge device without any network connection, so there is no way to egress the ingested data. The data must keep evolving within the limited storage budget given by the system. The goal of AdaEdge compression selection in the offline mode is to try to keep mission-critical information as much as possible to minimize the potential task accuracy loss under the given storage budget. The data evolving here is achieved by applying more aggressive compression on the data with less critical information. Figure 5.5 shows the cascade compression for the offline mode. For a given segment in the uncompressed buffer,

AdaEdge applies lossless compression by default to save space for future segments. When a pre-defined threshold is reached, the recoding process is waked up to further compressed the segments in the compressed buffer. As the system keeps ingesting more and more data points, there will be more space needed, thus, more aggressive compression will be applied to the old data segments. The compression sequence of the segments depends on the cache maintenance policy, which will be discussed in Section 5.2.5.

Similar to the online mode, AdaEdge creates multiple MAB instances for different compression levels for offline mode. Each MAB instance corresponds to a specific compression ratio range. This design is based on the observation that the optimization target changes significantly across different compression ratio ranges, and a single MAB instance for lossy compression selection is hard to reflect the compression ratio impact on the optimization target. For a given uncompressed segment, the lossless compression selection MAB instance suggests the optimal lossless compression and compresses the segment to save space. So the compression size is the optimization target for the lossless compression selection MAB instance. The compressed size is then updated to the MAB instance for future estimation. AdaEdge keeps ingest the new data points until the predefined threshold θ is reached. When more space is needed for the new ingested segments, AdaEdge initializes the lossy compression on the compressed segments, by default shrinking the size to half of the input size where we can derive a target compression ratio for the current segment. So the lossy compression selection component needs to refer to the MAB instance corresponding to the target compression ratio range, and the suggested action is the optimal compression for the optimization target within that specific compression ratio range. AdaEdge periodically evaluates the target task for each MAB instance on the compressed segment belonging to its corresponding compression range. AdaEdge keeps compressing and recoding the segment within the given storage budget until the egress option is available.

5.2.3 Optimization target

AdaEdge currently supports single workload targets, including aggregation query (min, max, sum, average) accuracy, machine learning task accuracy, and compression throughput. It is easy to deploy other workload targets with a clear quantitative evaluation rule in AdaEdge.

In AdaEdge, the machine learning task accuracy evaluation on a lossy compressed segment is defined as follows: with lossy compression, we can get its compressed representation R_l . Before feeding it into the machine learning task, we need to decompress it into X_l . Given the task T and a pre-trained model M , we get the prediction results $Y_l = M(X_l)$. For the machine learning task evaluation on lossy compressed data, we define the accuracy as the ratio of matched prediction labels compared with predicted labels on the original data:

$$ACC_{ml} = \frac{|\{x|x \in X \text{ AND } M(x_l) = M(x)\}|}{|X|}$$

Here, we assume the task model M is the pre-trained model on the original dataset X and provided to the compression task as a given input model. Model M could be either trained on the edge device or pre-downloaded from the cloud. AdaEdge has a dedicated serialization and deserialization module to manage the machine learning model instance. AdaEdge loads and deserializes the corresponding binary file into a machine learning instance for performance evaluation when a particular machine learning task is defined as the optimization target.

In addition to the machine learning tasks, AdaEdge defines relative loss for aggregation queries as other work does for approximate query processing evaluation [69]:

$$Acc_{agg} = 1 - |V_{true} - V_{lossy}|/|V_{true}$$

Where V_{true} is the real value we get from the aggregation operator on the original data, while V_{lossy} is the estimated value on the lossy decompressed data. The compression throughput

is defined as

$$C_{thr} = S_o/T_c$$

where S_o is the original file size and T_c corresponds to the compression runtime. A fast compression usually means less number of instruction for the compression codec, which consumes less power [19]. Therefore, a fast compression may correspond to a power efficient compression approach, which is very importance for IoT applications.

In addition to the single optimization target, AdaEdge also supports a hybrid weighted target, which we call a complex optimization target. We leave the user with the flexibility of assigning different weights for each kind of task to support custom optimization targets for the system. The complex target is defined as

$$target_c = w_1 \times \overline{ACC}_{agg} + w_2 \times \overline{ACC}_{ML} + w_3 \times \overline{C}_{thr}$$

where \overline{ACC} s and \overline{C}_{thr} are normalized, and $w_1 + w_2 + w_3 = 1$.

5.2.4 Recoding optimization

AdaEdge may apply multiple times of recoding for the same segment. Traditional compression requires full decompression before recoding, which may slow down the recoding process. Similar to "virtual decompression" in TVStore [13], AdaEdge tries to avoid any unnecessary decoding as much as possible. All the lossy compression approaches supported in AdaEdge are easy to apply recoding without requiring any decompression beforehand. When more aggressive compression is required, AdaEdge checks the original and target compression methods and applies direct recoding on the compressed data without decompression if possible.

AdaEdge currently supports "virtual decompression" for the recoding tasks with identical source and destination compression approaches. For example, we truncate the BUFF compressed data by removing the trailing bits to achieve more aggressive compression. We

also apply PAA compression on the PAA compressed data to achieve PAA compression with greater window size. We eliminate more insignificant frequency channels to compress the FFT encoded segments further. We apply PLA compression on the PLA encoded segment to get a smaller number of key data points to save more space. More similar work can be done by enabling direct transcoding between different compression approaches, which need specific compression optimization for each compression pair, so we keep that as future work.

5.2.5 *LRU-based compression*

The compression and recoding sequence on the segments is important for minimizing the negative impacts on the target workload. For example, a newly ingested fresh data are usually more important than the old data in the system. More aggressive compression should be applied to the old data when necessary, and latest data should be kept lossless if possible. Prior work, RRDTool [105] removes the old data in a round-robin way. TVstore [13] and SummaryStore [11] apply more aggressive compression on the segment with an early timestamp.

AdaEdge supports LRU-based compression policy, where the least recently used segments are compressed first, and the frequently used data segments are less likely to be lossily compressed. The LRU-based segment management component is implemented based on a list. The segment management component returns to the compression or recoding threads with the least recently used segments at the head of the list, and appends the compressed segment to the end of the list. Meanwhile, all segments touched by the running queries will be moved to the end of the list to ensure their least chance of being lossy compressed. The LRU-based compression policy works well for AdaEdge where ingesting the signal and some aggregation queries are the main tasks. The segment management component is easy to adapt to other compression sequence policies. More investigation on different policies will be interesting, which we leave as future work.

5.3 Evaluation

In this section, we focus on the evaluation of AdaEdge for aggregation queries and machine learning tasks with varying compression approaches. We implement Gzip, Snappy, Gorilla, Sprintz, and BUFF for lossless compression, and PAA, PLA, FFT, BUFF-lossy and random sample for lossy methods in AdaEdge. The MAB-assisted compression selection component’s action space for lossless compression selection includes Gzip, Snappy, Gorilla, BUFF, and Sprintz. We use the default setting for Gzip, Snappy, and Gorilla compression. For BUFF and Sprintz, we keep 4 digit precision on the CBF dataset, 5 digit precision on the UCR dataset, and 6 digit precision on the UCI dataset according to the dataset statistics. The action space for lossy compression selection includes lossy compression approaches such as PAA, PLA, FFT, BUFF-lossy, and RRD-sample. For any given compression ratio target, AdaEdge tunes the lossy compression parameters (such as the window size, number of data points budget, number of frequency channels, and number of bits) accordingly to achieve the compression ratio target.

We use optimistic ϵ -greedy for all the compression selection to encourage the exploration of the whole candidate action for optimal action. For the offline mode experiments, we set ϵ to 0.1 for MAB to better explore the knowledge of sub-optimal actions. For online mode experiments, we set ϵ to 0.01 for better exploiting the optimal MAB action. We show the compression selection performance of AdaEdge in both online and offline modes with single or complex optimization targets.

All experiments were performed on servers with 2 Intel(R) Xeon(R) CPUs E5-2670 v3 @ 2.30GHz, 128GB memory, 250GB HDD, and Ubuntu 18.04. All our implementations and experiments are done in Rust 1.49.0.

5.3.1 Machine learning accuracy with given model

In our first set of experiments, we apply lossy compression (BUFF, PAA, FFT, PLA) on UCR and UCI datasets and then apply different machine learning tasks to those lossy representations. UCI and UCR dataset are two popular machine learning dataset, including ~ 250 datasets from different fields.

Figure 5.6 to 5.13 show the *ACC* defined in Section 5.2.3 and compression ratio on the UCR and UCI datasets with machine learning models, including decision tree, random forest, KNN, and KMeans, respectively. We train the model on the original dataset and then apply those models to the lossy compressed data to evaluate the predicted label changes. The horizontal axis is the compression ratio achieved by corresponding compression approaches, and the vertical axis shows the corresponding relative accuracy defined in Section 5.2.3. As we can see from Figures 5.6 and 5.7, the tree-based task is sensitive to lossy compression, since slight variations of the key data features might result in a completely different branch, thus a different prediction in the inference phase. Overall, BUFF-lossy works best in most compression ratio configurations as it has a minimal distortion (bit level difference) from the original values. But we do notice BUFF-lossy performance deteriorates at a low compression ratio. As is shown in Figures 5.10 and 5.11, BUFF-lossy is worse than FFT, PAA and PLA around compression ratio 0.12 in the decision tree and random forest evaluation on the UCR dataset. In addition, BUFF-lossy fails to support our dataset for compression ratios lower than 0.11. With Figures 5.12 and 5.13, we can see that both BUFF-lossy and FFT benefits high dimensional distance based task. PLA performs the worst overall for the machine learning tasks on our dataset.

With those experiments, we can see that the machine learning accuracy performance varies for different lossy compression approaches on different target compression ratio targets. It demonstrates that a simple rule-based encoding selection solution is not good enough to handle the changing data features and varying workloads. A dynamic compression selection model is necessary to provide suggestions of optimal compression based on the incoming




Figure 5.6: Decision tree model accuracy on UCI




Figure 5.7: Random forest model accuracy on UCI

data statistics, host system requirements, and downstream workload.

5.3.2 Adaptive Compression Selection

We have seen the overall compression performance on historical data collections. We then evaluate the AdaEdge compression selection on the steaming data with the running ingestion process. In this experiment, AdaEdge ingests the data from a dummy client generating data points from the CBF dataset [95], which is a simulated dataset derived from a library of distribution bases. The received data is then ingested as uncompressed segments. AdaEdge applies compression selection on the ingested data segments according to the system constraints. For comparison, we also include baselines with fixed compression combinations.

Online mode




Figure 5.8: KNN model accuracy on UCI




Figure 5.9: KMeans model accuracy on UCI




Figure 5.10: Decision tree model accuracy on UCR




Figure 5.11: Random forest model accuracy on UCR




Figure 5.12: KNN model accuracy on UCR




Figure 5.13: KMeans model accuracy on UCR

We start with AdaEdge online mode, where we have a clear target compression ratio derived from the ingestion data rate and network bandwidth. AdaEdge adaptively chooses the optimal compression for a given workload target. The compression process fetches the segments from the uncompressed buffer and applies compression based on the MAB component estimation. Figure 5.14 shows the machine learning accuracy loss across different target compression ratio. The horizontal axis is the target compression ratio derived by system constraints, and the vertical axis shows the corresponding relative accuracy loss. The smaller the task accuracy loss is, the better compression it is on this given dataset. The MAB line is the compression selection component used by AdaEdge. We also show the lossless compression performance with solid lines and lossy baselines with dash lines. For lossy compression baselines, PAA, PLA, and FFT can handle any target compression ratio between 1.0 inclusive and 0 exclusive, while BUFF-lossy does not support a compression ratio below 0.125 on the CBF dataset. Lossless compression has no accuracy loss for any workload, so they always get 0 for accuracy loss. We show negative values for different lossless compression for an easily distinguishable layout in the figure. AdaEdge uses optimal lossless compression when possible and falls back to lossy compression otherwise. From Figures 5.14, MAB based compression selection always selects the optimal compression for each target compression ratio. It selects BUFF-lossy for compression range above 0.125 and choose either PAA or FFT for other compression ratio range. The variation on the MAB line comes from the exploration cost of MAB, in which sub-optimal action is taken, thus inflating the accuracy loss.

We also test AdaEdge encoding selection performance on the aggregation queries. Figures 5.15 and 5.16 show the relative aggregation accuracy loss for each compression approaches. Since the accuracy loss is very low, we use a log scale here for the vertical axis. Similarly to the previous experiment, we denote all the lossless baselines with a value smaller than $1.00e - 18$ for better layout. Please note that there is no accuracy loss for all the lossless compression. The AdaEdge compression selection works well on those two workloads, as it

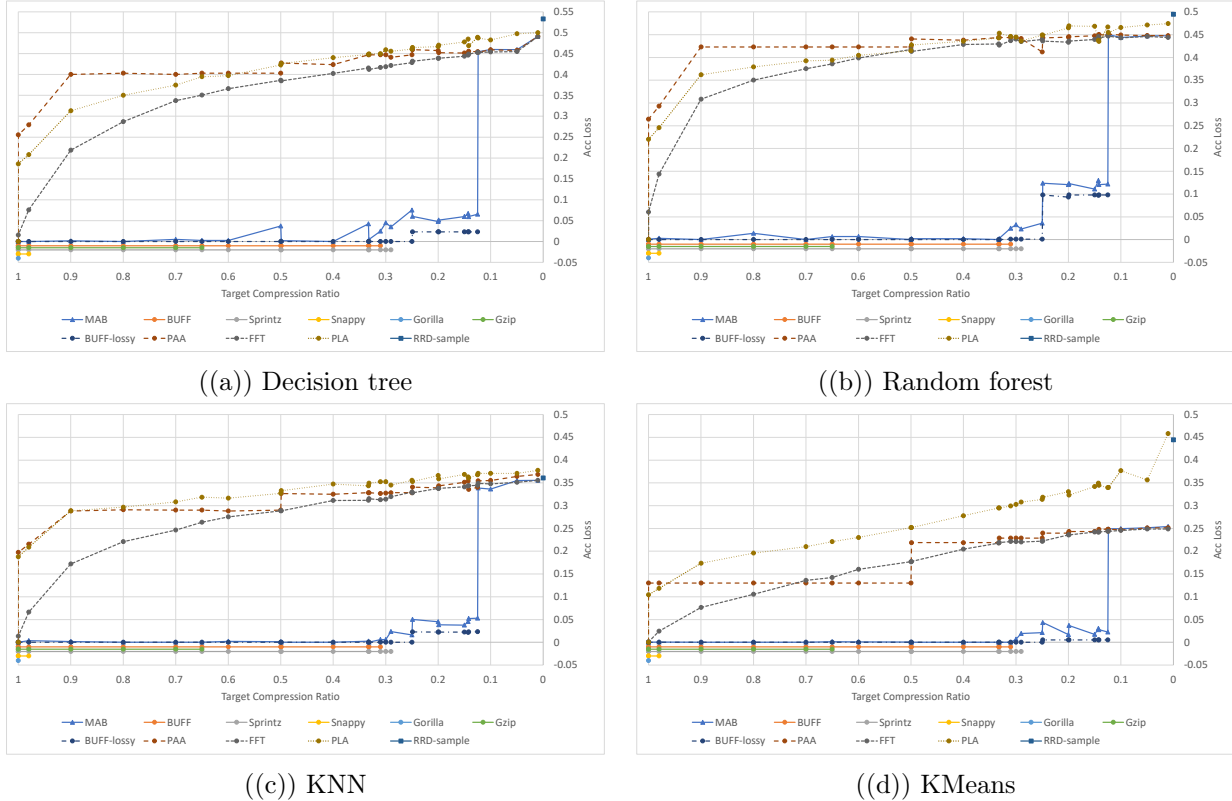


Figure 5.14: Machine learning model accuracy on target compression ratio

picks PAA or FFT as optimal compression, which is the ground truth optimal for the sum aggregations. The spikes on the MAB lines are also from the MAB exploration step. It's prominent because of the log scale, which amplifies the error from sub-optimal actions on a tiny amount of segments. Similarly, AdaEdge chooses PLA for all the compression ranges because of the outstanding performance of PLA for max aggregation queries.

It is straightforward for AdaEdge to optimize for a single target. In the following experiment, we run compression selection experiments optimizing for complex workload targets. For time series data systems, aggregation query and machine learning task accuracy are sometimes both important. AdaEdge also supports a complex optimization target considering both aggregation query and machine learning accuracy. A knob is provided to balance two of them.

$$A = \arg \max_a (w_1 \times Acc_{agg}(a) + w_2 \times Acc_{ML}(a))$$

Figure 5.17 shows the sum aggregation query and random forest complex optimization target

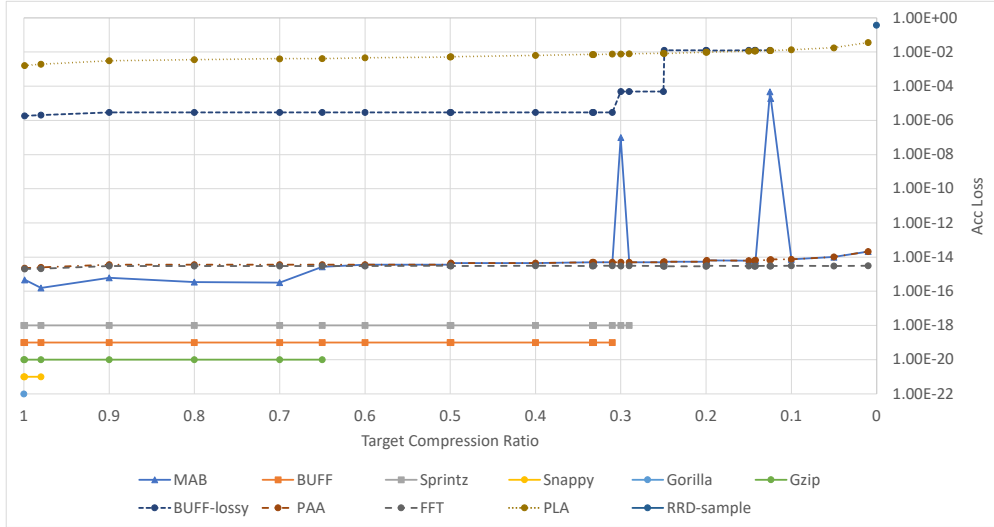


Figure 5.15: Sum aggregation query over target compression ratio

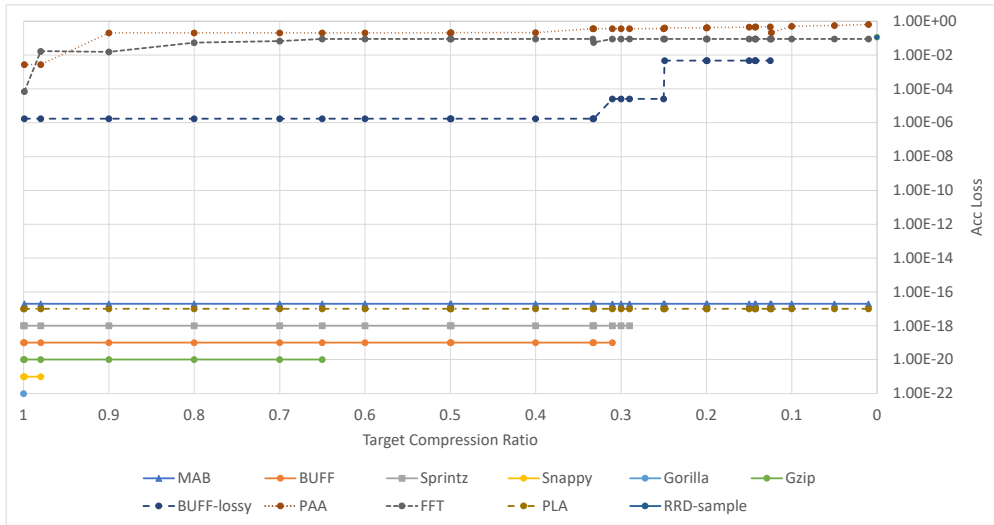


Figure 5.16: Max aggregation query over target compression ratio

for the online mode where a clear target compression ratio can be derived from the system constraints. According to the baselines, there are two crossover points for the complex optimization target: the first is a compression ratio of around 0.8, and the second is around 0.25. FFT is the ground truth optimal compression in the range 1 to 0.8, then BUFF-lossy is the optimal compression in the range 0.8 to 0.25, and then FFT becomes the optimal again in the range 0.25 to 0. According to Figure 5.17, AdaEdge can adapt to the best compression in most range. The outlier around compression ratio 0.15 is because of the variation from the sum aggregation query.

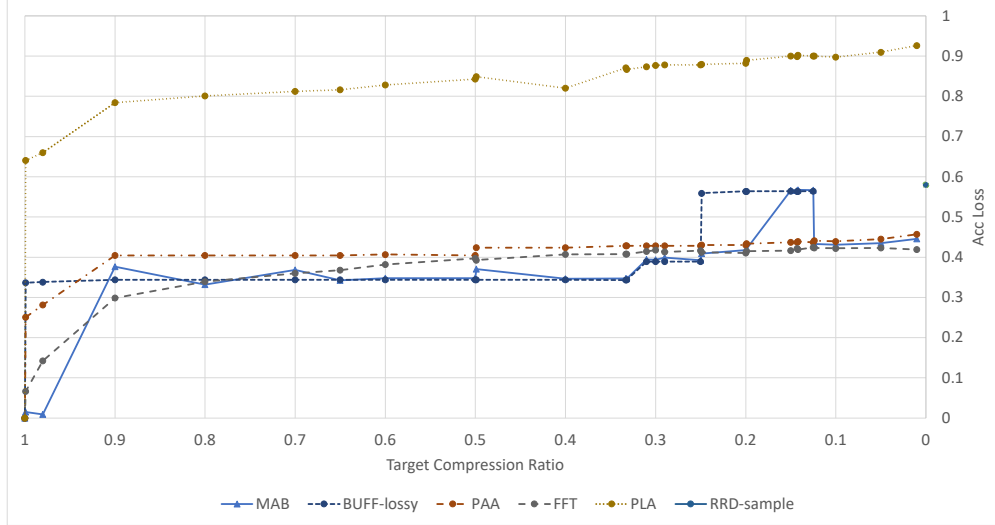


Figure 5.17: Sum aggregation query and random forest complex optimization target over target compression ratio with weighted value $w_1 = 0.625$ and $w_2 = 0.375$.

Accuracy is not the only dimension for system performance evaluation. Data systems care about compression speed in addition to task accuracy. AdaEdge is also able to support the balanced optimization target between those two. Figure 5.18 shows the compression selection performance with the optimization target as the combination of compression speed and task accuracy.

$$A = \arg \max_a (w_1 \times C_{thr}(a) + w_2 \times Acc_{ML}(a))$$

As we can see from Figure 5.18, the vertical axis is the complex target for speed and accuracy combination, so greater is better in this figure. Overall, the MAB is capable of selecting the best compression. We can see a crossover point at 0.25 between PAA and BUFF-lossy, and AdaEdge can handle those cases well.

Offline mode

AdaEdge offline mode aims to preserve workload-critical information as much as possible within the system storage budget. So the compression process keeps fetching uncompressed segments from the uncompressed buffer and applies a different level of compression based on the segment importance. In this experiment, we reserve 10MB of space for AdaEdge to





Figure 5.18: Compression speed and random forest complex optimization target over target compression ratio with weights $w_1 = 0.524$ and $w_2 = 0.476$.

ingest 80MB of data points with a recoding threshold 0.8, meaning recoding is needed to free up space when the space usage reaches 0.8. The horizontal axis in Figure 5.19 shows the ingestion timestamp. The two vertical axes correspond to space usage and machine learning task accuracy loss. In this ingestion experiment, there are basically two types of encoding selection: lossless compression selection that optimizes for minimizing compression ratio to save space as much as possible, and lossy compression selection that optimizes for machine learning tasks accuracy. Here we also include the baselines with all possible compression pairs in the form "Lossless, Lossy". There are, in total, 25 baselines covering all compression pair combinations. We only show several top compression pair representatives to keep the figure clear and readable. "MAB, MAB" is our solution with MAB-based encoding selection strategies. During the ingestion experiment, we will periodically check the space usage shown with the red lines in the figures.


In this experiment, we use 0.8 as a threshold for initializing a recoding process. As we can see from Figure 5.19, all the compression pairs are capable of bounding the space usage to a safe level defined by the recoding threshold, as we can see from the red lines. The increasing slope of MAB space usage is slighter than Gzip, Snappy or Gorilla baselines shown in subfigures 5.19(g), 5.19(h) and 5.19(i) respectively, since MAB lossless compression selects

Sprintz as the optimal lossless compression with the smallest compressed file size. The blue line to the bottom shows the KMeans task accuracy loss. The task accuracy loss starts to increase when the space usage line reaches the recoding threshold, where the lossy recoding process is initialized. The slower the task accuracy loss line increases, the better compression it is. As we can see, MAB-based lossy compression selection is always the best to choose the optimal compression to minimize task accuracy loss.


In this experiment, MAB-based lossy compression selection chooses BUFF-lossy as optimal lossy compression and then uses PAA as optimal lossy compression when BUFF-lossy fails to compress the segment further. We can also see from the baseline subfigure 5.19(b), the accuracy loss stays low when the recoding first starts, and increases quickly when more data is ingested, requiring more aggressive compression for the old data. BUFF-lossy fails and falls back to RRD-sample to further compress the segments in the late phase of the ingestion experiment. In addition to KMeans, we also test AdaEdge on other machine learning tasks. The experiments show that AdaEdge could also adapt to the best solution on other machine learning tasks, including KNN, random forest, and decision tree.




((a) MAB, MAB




((b) Sprintz, BUFF-lossy




((c) Sprintz, PAA




((d) Sprintz, FFT




((e) Sprintz, PLA




((f) Sprintz, RRD



((g) Gzip, BUFF-lossy



((h) Snappy, BUFF-lossy



((i) Gorilla, BUFF-lossy

Figure 5.19: KMeans accuracy loss over ingestion time

5.4 Conclusion

AdaEdge provides an efficient and effective compression selection framework for data systems with a clear target compression ratio or a clear storage budget. AdaEdge supports single and complex optimization goals for the whole system, allowing users to choose the custom optimization target for their data systems. It applies LRU based compression policy to compress the data depending on its usage. It supports lossy and lossless compression selection depending on the hardware constraints and user-defined workload target. AdaEdge deploys a multi-armed bandit algorithm to assist the compression selection, and provides online and offline modes to conserve workload-specific information as much as possible within the system constraints.

CHAPTER 6

CONCLUSION

Data systems employ compression to reduce on-disk or in-memory data sizes, improve bandwidth and storage utilization, and save costs. Conventional compression approaches have traditionally focused on minimizing file sizes. This strict focus on size alone, however, overlooks considerations of critical importance to maximizing query execution performance. Improperly-applied compression may hurt query performance, as many compression methods require complete decompression before query execution may proceed. Such compression approaches waste computation resources and further optimization opportunities. Instead, a good compression should be able to trade off compressed size with the ability to query directly on the compressed data and speed up the query according to the host system’s hardware. When designing a new compression, we should treat the compression as an optimization opportunity for the workloads on the given data. The thesis applies this rule when creating MOP, BUFF, and AdaEdge for IoT systems.

In Chapter 3, this thesis proposes MOP to encode a stream data in a best-effort order-preserving way without full knowledge of the incoming data. MOP pre-allocates encoding space for the incoming new entries, and optimizes the orderedness of the dictionary for better query performance. The thesis presents a technique for generating a MOP with a limited sample of the input stream while minimizing the size of the dictionary. In addition, the thesis introduces a variation that uses cascading MOPs with multiple levels of ordered keys. The query on MOPs performs similarly to order-preserving dictionaries with minimal overhead for decoding. MOPs present query rewriting rules to minimize decoding of keys to minimize predicate evaluation latency. The evaluation on Parquet demonstrates that MOPs are able to accelerate range filtering and sorting, and achieve high order ratios with small samples.

In Chapter 4, this thesis proposes BUFF to handle the heterogeneous IoT numeric data with varying precision and ranges. BUFF addresses the heterogeneity of IoT data by applying decomposed but compact encoding space for the IoT numeric data with different

statistics. BUFF achieves efficient query execution support according to different hardware. BUFF supports efficient in-situ adaptive query execution on encoded data directly. In addition, BUFF provides fast aggregation and materialization for different precision levels. The evaluation shows BUFF achieves fast compression and query execution while keeping a comparable compression ratio to the best approaches.

Similar rules also hold for compression selection. Instead of focusing on the compressed size only, compression selection also takes the query performance and task accuracy (when lossy compression approaches are included) into consideration. In Chapter 5, this thesis proposes AdaEdge as a hardware-conscious compression selection framework for resource-constrained devices. AdaEdge supports single or multiple optimization objectives for the system to enable more robust compression selection focusing on not only size but also speed and accuracy. AdaEdge introduces a LRU-based compression policy to compress the data based on its importance and staleness. AdaEdge uses Multi-Armed Bandit assisted lossy and lossless compression selection to adapt to the system hardware constraints and optimize for the workload targets. AdaEdge supports multiple optimization targets, including compression ratio, throughput, aggregation queries and machine learning task accuracy. In addition, AdaEdge provides online and offline mode compression solutions for resource-constrained devices to conserve mission-critical information. The evaluation shows that AdaEdge is capable of choosing optimal compression for customized optimization targets defined by users while still complying with the system constraints.

Compression is not an isolated system design decision anymore. To better serve the overall system performance goals, the design of data compression should adjust to the data access patterns. A good compression should not only compress the data effectively but also support efficient compression and query execution. With this rule, we can design fast and effective compression solutions to boost the system's overall performance goals: make data smaller, and query faster and more accurate simultaneously.

References

- [1] Apache CarbonData. <https://carbondata.apache.org/>.
- [2] Huge stock market dataset. <https://www.kaggle.com/borismarjanovic/price-volume-data-for-all-us-stocks-etfs>.
- [3] Individual household electric power consumption data set. <https://archive.ics.uci.edu/ml/datasets/individual+household+electric+power+consumption>.
- [4] Influxdb - open source time series, metrics, and analytics database. <https://www.influxdata.com/>.
- [5] Time series benchmark suite (tsbs). <https://github.com/timescale/tsbs>.
- [6] Time-series data simplified. <https://www.timescale.com>.
- [7] Tlc trip record data. www.nyc.gov/html/tlc/html/about/triprecorddata.shtml.
- [8] Torque characteristics of a permanent magnet motor. <https://www.kaggle.com/graxlmaxl/identifying-the-physics-behind-an-electric-motor>.
- [9] Daniel Abadi, Samuel Madden, and Miguel Ferreira. Integrating Compression and Execution in Column-oriented Database Systems. In *SIGMOD*, pages 671–682, 2006.
- [10] Daniel Abadi, Samuel Madden, and Miguel Ferreira. Integrating compression and execution in column-oriented database systems. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 671–682, 2006.
- [11] Nitin Agrawal and Ashish Vulimiri. Low-latency analytics on colossal data streams with summarystore. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 647–664, 2017.
- [12] Yanzhe An, Yue Su, Yuqing Zhu, and Jianmin Wang. Tvstore: Automatically bounding time series storage via time-varying compression.
- [13] Yanzhe An, Yue Su, Yuqing Zhu, and Jianmin Wang. TVStore: Automatically bounding time series storage via Time-Varying compression. In *20th USENIX Conference on File and Storage Technologies (FAST 22)*, pages 83–100, Santa Clara, CA, February 2022. USENIX Association.
- [14] G. Antoshenkov, D. Lomet, and J. Murray. Order preserving string compression. In *ICDE*, pages 655–663, Feb 1996.
- [15] Gennady Antoshenkov. Dictionary-based order-preserving string compression. *VLDB Journal*, 6(1):26–39, 1997.
- [16] Apache Foundation. Apache Arrow Feather. <https://arrow.apache.org/docs/python/feather.html>.

- [17] Apache Foundation. Apache ORC. <https://orc.apache.org>.
- [18] Apache Foundation. Apache Parquet. <https://parquet.apache.org/>.
- [19] Kenneth C Barr and Krste Asanović. Energy-aware lossless data compression. *ACM Transactions on Computer Systems (TOCS)*, 24(3):250–291, 2006.
- [20] Carsten Binnig, Stefan Hildenbrand, and Franz Färber. Dictionary-based order-preserving string compression for main memory column stores. In *ACM SIGMOD*, pages 283–296. ACM, 2009.
- [21] Davis Blalock, Samuel Madden, and John Gutttag. Sprintz: Time series compression for the internet of things. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–23, 2018.
- [22] Martin Boissier. Robust and budget-constrained encoding configurations for in-memory database systems. *Proceedings of the VLDB Endowment*, 15(4):780–793, 2021.
- [23] Paul Boniol, John Paparrizos, Themis Palpanas, and Michael J Franklin. Sand: streaming subsequence anomaly detection, 2021.
- [24] Robert S Boyer and J Strother Moore. Mjrtj—a fast majority vote algorithm. In *Automated Reasoning*, pages 105–117. Springer, 1991.
- [25] Javier Burgués, Juan Manuel Jiménez-Soto, and Santiago Marco. Estimation of the limit of detection in semiconductor gas sensors through linearized calibration models. *Analytica chimica acta*, 1013:13–25, 2018.
- [26] Javier Burgués and Santiago Marco. Multivariate estimation of the limit of detection by orthogonal partial least squares in temperature-modulated mox sensors. *Analytica chimica acta*, 1019:49–64, 2018.
- [27] Martin Burtscher and Paruj Ratanaworabhan. Fpc: A high-speed compressor for double-precision floating-point data. *IEEE Transactions on Computers*, 58(1):18–31, 2008.
- [28] Lujing Cen, Andreas Kipf, Ryan Marcus, and Tim Kraska. Lea: A learned encoding advisor for column stores. In *Fourth Workshop in Exploiting AI Techniques for Data Management*, pages 32–35, 2021.
- [29] Moses Charikar, Surajit Chaudhuri, Rajeev Motwani, and Vivek Narasayya. Towards estimation error guarantees for distinct values. In *PODS*, pages 268–279, 2000.
- [30] Zhiyuan Chen, Johannes Gehrke, and Flip Korn. Query optimization in compressed database systems. *SIGMOD Record*, 30(2):271–282, 2001.
- [31] Stefano Cherubin and Giovanni Agosta. Tools for reduced precision computation: A survey. *ACM Computing Surveys (CSUR)*, 53(2):1–35, 2020.

- [32] Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings, 2011.
- [33] Florent De Dinechin, Jérémie Detrey, Octavian Cret, and Radu Tudoran. When fpgas are better at floating-point than microprocessors. In *FPGA*, volume 8, page 260, 2008.
- [34] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Voshall, and Werner Vogels. Dynamo: Amazon’s highly available key-value store. In *Proceedings of Twenty-first ACM SIGOPS Symposium on Operating Systems Principles*, SOSP ’07, pages 205–220, New York, NY, USA, 2007. ACM.
- [35] Peter Deutsch et al. Gzip file format specification version 4.3. Technical report, RFC 1952, May, 1996.
- [36] Hank Dietz, Bill Dieter, Randy Fisher, and Kungyen Chang. Floating-point computation with just enough accuracy. In *International Conference on Computational Science*, pages 226–233. Springer, 2006.
- [37] Paul DuBois. *MySQL*. Pearson Education, 2008.
- [38] Adam Dziedziec, John Paparrizos, Sanjay Krishnan, Aaron Elmore, and Michael Franklin. Band-limited training and inference for convolutional neural networks. In *International Conference on Machine Learning*, pages 1745–1754. PMLR, 2019.
- [39] Christos Faloutsos, M. Ranganathan, and Yannis Manolopoulos. Fast subsequence matching in time-series databases. In *SIGMOD*, pages 419–429, 1994.
- [40] Yuanwei Fang, Chen Zou, Aaron J. Elmore, and Andrew A. Chien. UDP: A Programmable Accelerator for Extract-transform-load Workloads and More. In *IEEE/ACM Micro*, pages 55–68, 2017.
- [41] Franz Färber, Sang Kyun Cha, Jürgen Primsch, Christof Bornhövd, Stefan Sigg, and Wolfgang Lehner. Sap hana database: data management for modern business applications. *ACM Sigmod Record*, 40(4):45–51, 2012.
- [42] Ziqiang Feng, Eric Lo, Ben Kao, and Wenjian Xu. Byteslice: Pushing the envelop of main memory data processing with a new storage layout. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 31–46, 2015.
- [43] Abraham D Flaxman, Adam Tauman Kalai, and H Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. *arXiv preprint cs/0408007*, 2004.
- [44] Edward Fredkin. Trie memory. *Commun. ACM*, 3.9:490–499, 1960.
- [45] Jean-loup Gailly and Mark Adler. Zlib compression library. 2004.

- [46] Hector Garcia-Molina. Elections in a distributed computing system. *IEEE transactions on Computers*, (1):48–59, 1982.
- [47] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. The google file system. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles*, pages 20–43, Bolton Landing, NY, 2003.
- [48] G. Graefe and L. D. Shapiro. Data compression and database performance. In *Symposium on Applied Computing*, pages 22–27, Apr 1991.
- [49] Anurag Gupta, Deepak Agarwal, Derek Tan, Jakub Kulesza, Rahul Pathak, Stefano Stefani, and Vidhya Srinivasan. Amazon redshift and the case for simpler data warehouses. In *ACM SIGMOD*, pages 1917–1923. ACM, 2015.
- [50] Brian Hentschel, Michael S Kester, and Stratos Idreos. Column sketches: A scan accelerator for rapid and robust predicate evaluation. In *Proceedings of the 2018 International Conference on Management of Data*, pages 857–872, 2018.
- [51] David Hough. Applications of the proposed ieee 754 standard for floating-point arithmetic. *Computer*, (3):70–74, 1981.
- [52] Te C Hu and Alan C Tucker. Optimal computer search trees and variable-length alphabetical codes. *SIAM Journal on Applied Mathematics*, 21(4):514–532, 1971.
- [53] Martin Isenburg, Peter Lindstrom, and Jack Snoeyink. Lossless compression of predicted floating-point geometry. *Computer-Aided Design*, 37(8):869–877, 2005.
- [54] Balakrishna R. Iyer and David Wilhite. Data Compression Support in Databases. In *VLDB*, pages 695–704, 1994.
- [55] Søren Kejser Jensen, Torben Bach Pedersen, and Christian Thomsen. Scalable model-based management of correlated dimensional time series in modelardb+. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, pages 1380–1391. IEEE, 2021.
- [56] Hao Jiang and Aaron J Elmore. Boosting data filtering on columnar encoding with simd. In *DAMON*, page 6. ACM, 2018.
- [57] Hao Jiang and Aaron J Elmore. Boosting data filtering on columnar encoding with simd. In *Proceedings of the 14th International Workshop on Data Management on New Hardware*, pages 1–10, 2018.
- [58] Hao Jiang, Chunwei Liu, Qi Jin, John Paparrizos, and Aaron J Elmore. Pids: attribute decomposition for improved compression and query performance in columnar storage. *Proceedings of the VLDB Endowment*, 13(6):925–938, 2020.
- [59] Hao Jiang, Chunwei Liu, John Paparrizos, Andrew A Chien, Jihong Ma, and Aaron J Elmore. Good to the last bit: Data-driven encoding with codecdb. In *Proceedings of the 2021 International Conference on Management of Data*, pages 843–856, 2021.

- [60] Shunsuke Kanda, Kazuhiro Morita, and Masao Fuketa. Practical string dictionary compression using string dictionary encoding. In *Big Data Innovations and Applications*, pages 1–8. IEEE, 2017.
- [61] M. G. Kendall. A New Measure Of Rank Correlation. *Biometrika*, 30(1-2):81, 1938.
- [62] Eamonn Keogh, Kaushik Chakrabarti, Michael Pazzani, and Sharad Mehrotra. Dimensionality reduction for fast similarity search in large time series databases. *Knowledge and information Systems*, 3(3):263–286, 2001.
- [63] Kelly Kissock. University of dayton kettering labs window energy use analysis. 2007.
- [64] Marcel Kornacker, Victor Bittorf, Taras Bobrovitsky, Casey Ching Alan Choi Justin Erickson, Martin Grund Daniel Hecht, Matthew Jacobs Ishaan Joshi Lenni Kuff, Dileep Kumar Alex Leblang, Nong Li Ippokratis Pandis Henry Robinson, David Rorke Silvius Rus, John Russell Dimitris Tsirogiannis Skye Wanderman, and Milne Michael Yoder. Impala: A modern, open-source sql engine for hadoop. In *Proceedings of the 7th Biennial Conference on Innovative Data Systems Research*, 2015.
- [65] Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, May 1998.
- [66] N Jesper Larsson and Alistair Moffat. Off-line dictionary-based compression. *Proceedings of the IEEE*, 88(11):1722–1732, 2000.
- [67] A. Lempel and J. Ziv. On the complexity of finite sequences. *IEEE Trans. Inf. Theor.*, 22(1):75–81, September 1976.
- [68] Yinan Li and Jignesh M Patel. Bitweaving: fast scans for main memory data processing. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pages 289–300, 2013.
- [69] Xi Liang, Zechao Shang, Sanjay Krishnan, Aaron J Elmore, and Michael J Franklin. Fast and reliable missing data contingency analysis with predicate-constraints. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 285–295, 2020.
- [70] Libelium. 50 Sensor Applications for a Smarter World. http://www.libelium.com/50_sensor_applications/, 2019. [Online; accessed 19-January-2022].
- [71] Peter Lindstrom. Fixed-rate compressed floating-point arrays. *IEEE transactions on visualization and computer graphics*, 20(12):2674–2683, 2014.
- [72] Peter Lindstrom and Martin Isenburg. Fast and efficient compression of floating-point data. *IEEE transactions on visualization and computer graphics*, 12(5):1245–1250, 2006.
- [73] Chunwei Liu, Hao Jiang, John Paparrizos, and Aaron J Elmore. Decomposed bounded floats for fast compression and queries. *Proceedings of the VLDB Endowment*, 14(11):2586–2598, 2021.

- [74] Chunwei Liu, McKade Umbenhowe, Hao Jiang, Pranav Subramaniam, Jihong Ma, and Aaron J Elmore. Mostly order preserving dictionaries. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 1214–1225. IEEE, 2019.
- [75] Aditya Mahajan and Demosthenis Teneketzis. Multi-armed bandit problems. In *Foundations and applications of sensor management*, pages 121–151. Springer, 2008.
- [76] Sergey Melnik, Andrey Gubarev, Jing Jing Long, Geoffrey Romer, Shiva Shivakumar, Matt Tolton, and Theo Vassilakis. Dremel: interactive analysis of web-scale datasets. *PVLDB*, 3(1-2):330–339, 2010.
- [77] Micro Focus International plc. Vertica Encoding Types. <https://www.vertica.com/docs/9.2.x/HTML/Content/Authoring/SQLReferenceManual/Statements/encoding-type.htm>.
- [78] Ingo Müller, Andrea Arteaga, Torsten Hoefler, and Gustavo Alonso. Reproducible floating-point aggregation in rdbms. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*, pages 1049–1060. IEEE, 2018.
- [79] Ingo Müller, Cornelius Ratsch, Franz Faerber, et al. Adaptive string dictionary compression in in-memory column-store database systems.
- [80] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, pages 305–319, Philadelphia, PA, 2014. USENIX Association.
- [81] John Ousterhout, Arjun Gopalan, Ashish Gupta, Ankita Kejriwal, Collin Lee, Behnam Montazeri, Diego Ongaro, Seo Jin Park, Henry Qin, Mendel Rosenblum, Stephen Rumble, Ryan Stutsman, and Stephen Yang. The ramcloud storage system. *ACM Trans. Comput. Syst.*, 33(3):7:1–7:55, August 2015.
- [82] John Paparrizos. 2018 ucr time-series archive: Backward compatibility, missing values, and varying lengths, January 2019. <https://github.com/johnpaparrizos/UCRArchiveFixes>.
- [83] John Paparrizos and Michael J Franklin. Grail: efficient time-series representation learning. *Proceedings of the VLDB Endowment*, 12(11):1762–1777, 2019.
- [84] John Paparrizos and Luis Gravano. k-shape: Efficient and accurate clustering of time series. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 1855–1870, 2015.
- [85] John Paparrizos and Luis Gravano. Fast and accurate time-series clustering. *ACM Transactions on Database Systems (TODS)*, 42(2):1–49, 2017.
- [86] John Paparrizos, Chunwei Liu, Bruno Barbarioli, Johnny Hwang, Ikradya Edian, Aaron J Elmore, Michael J Franklin, and Sanjay Krishnan. Vergedb: A database for iot analytics on edge devices. In *CIDR*, 2021.

- [87] John Paparrizos, Chunwei Liu, Aaron J Elmore, and Michael J Franklin. Debunking four long-standing misconceptions of time-series distance measures. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 1887–1905, 2020.
- [88] Marcus Paradies, Christian Lemke, Hasso Plattner, Wolfgang Lehner, Kai-Uwe Sattler, Alexander Zeier, and Jens Krueger. How to juggle columns: an entropy-based approach for table compression. In *IDEAS*, pages 205–215. ACM, 2010.
- [89] Tuomas Pelkonen, Scott Franklin, Justin Teller, Paul Cavallaro, Qi Huang, Justin Meza, and Kaushik Veeraraghavan. Gorilla: A fast, scalable, in-memory time series database. *Proceedings of the VLDB Endowment*, 8(12):1816–1827, 2015.
- [90] Orestis Polychroniou, Arun Raghavan, and Kenneth A Ross. Rethinking simd vectorization for in-memory databases. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 1493–1508, 2015.
- [91] Navneet Potti and Jignesh M Patel. Daq: a new paradigm for approximate query processing. *Proceedings of the VLDB Endowment*, 8(9):898–909, 2015.
- [92] Paruj Ratanaworabhan, Jian Ke, and Martin Burtscher. Fast lossless compression of scientific floating-point data. In *Data Compression Conference (DCC'06)*, pages 133–142. IEEE, 2006.
- [93] Gautam Ray, Jayant R. Haritsa, and S Seshadri. Database Compression: A Performance Enhancement Tool. 09 2004.
- [94] David Reinsel-John Gantz-John Rydning. The digitization of the world from edge to core. *Framingham: International Data Corporation*, 2018.
- [95] Naoki Saito. *Local feature extraction and its applications using a library of bases*. Yale University, 1994.
- [96] Mohammad Shahrads, Rodrigo Fonseca, Íñigo Goiri, Gohar Chaudhry, Paul Batum, Jason Cooke, Eduardo Laureano, Colby Tresness, Mark Russinovich, and Ricardo Bianchini. Serverless in the wild: Characterizing and optimizing the serverless workload at a large cloud provider. *arXiv preprint arXiv:2003.03423*, 2020.
- [97] Hagit Shatkay and Stanley B Zdonik. Approximate queries and representations for large data sequences. In *ICDE*, pages 536–545. IEEE, 1996.
- [98] Vishal Sikka, Franz Färber, Wolfgang Lehner, Sang Kyun Cha, Thomas Peh, and Christof Bornhövd. Efficient transaction processing in SAP HANA database: the end of a column store myth. In *ACM SIGMOD*, pages 731–742. ACM, 2012.
- [99] Emma Stewart, Anna Liao, and Ciaran Roberts. Open μ pmu: A real world reference distribution micro-phasor measurement unit data set for research and application development. 2016.

- [100] Jon Stokes. *Inside the machine: an illustrated introduction to microprocessors and computer architecture*. No Starch Press, 2007.
- [101] Mike Stonebraker, Daniel J Abadi, Adam Batkin, Xuedong Chen, Mitch Cherniack, Miguel Ferreira, Edmond Lau, Amerson Lin, Sam Madden, Elizabeth O’Neil, et al. C-store: a column-oriented dbms. In *Making Databases Work: the Pragmatic Wisdom of Michael Stonebraker*, pages 491–518. 2018.
- [102] Mike Stonebraker, Daniel J. Abadi, Adam Batkin, Xuedong Chen, Mitch Cherniack, Miguel Ferreira, Edmond Lau, Amerson Lin, Sam Madden, Elizabeth O’Neil, Pat O’Neil, Alex Rasin, Nga Tran, and Stan Zdonik. C-store: A Column-oriented DBMS. In *Very Large Data Bases*, pages 553–564. VLDB Endowment, 2005.
- [103] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [104] Igor Tatarinov, Stratis D Viglas, Kevin Beyer, Jayavel Shanmugasundaram, Eugene Shekita, and Chun Zhang. Storing and querying ordered XML using a relational database system. In *ACM SIGMOD*, pages 204–215. ACM, 2002.
- [105] Tobi Oetiker. Rrdtool: round robin database tool. <http://oss.oetiker.ch/rrdtool/>.
- [106] Maheswari Visvalingam and James D Whyatt. Line generalisation by repeated elimination of points. *The cartographic journal*, 30(1):46–51, 1993.
- [107] Zeke Wang, Kaan Kara, Hantian Zhang, Gustavo Alonso, Onur Mutlu, and Ce Zhang. Accelerating generalized linear models with mlweaving: A one-size-fits-all system for any-precision learning. *Proceedings of the VLDB Endowment*, 12(7):807–821, 2019.
- [108] Nathan Whitehead and Alex Fit-Florea. Precision & performance: Floating point and iee 754 compliance for nvidia gpus. *rn (A + B)*, 21(1):18749–19424, 2011.
- [109] Wikipedia contributors. Snappy (compression) — Wikipedia, the free encyclopedia. [https://en.wikipedia.org/w/index.php?title=Snappy_\(compression\)&oldid=977673115](https://en.wikipedia.org/w/index.php?title=Snappy_(compression)&oldid=977673115), 2020. [Online; accessed 14-September-2020].
- [110] Thomas Willhalm, Nicolae Popovici, Yazan Boshmaf, Hasso Plattner, Alexander Zeier, and Jan Schaffner. Simd-scan: ultra fast in-memory table scan using on-chip vector processing units. *Proceedings of the VLDB Endowment*, 2(1):385–394, 2009.
- [111] I. H. Witten, A. Moffat, and T. C. Bell. Managing gigabytes: Compressing and indexing documents and images. *IEEE Transactions on Information Theory*, 41(6):2101–, Nov 1995.
- [112] Gene T.J. Wu and Arthur J. Bernstein. Efficient solutions to the replicated log and dictionary problems. In *Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing*, PODC ’84, pages 233–242, New York, NY, USA, 1984. ACM.

- [113] Byoung-Kee Yi and Christos Faloutsos. Fast time sequence indexing for arbitrary lp norms. *VLDB*, 2000.
- [114] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. Spark: Cluster computing with working sets. *HotCloud*, 10(10-10):95, 2010.
- [115] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Trans. Inf. Theor.*, (3):337–343, September 1977.
- [116] J. Ziv and A. Lempel. Compression of individual sequences via variable-rate coding. *IEEE Trans. Inf. Theor.*, (5):530–536, September 1978.
- [117] Jacob Ziv and Abraham Lempel. A universal algorithm for sequential data compression. *IEEE Transactions on information theory*, 23(3):337–343, 1977.
- [118] Marcin Zukowski, Peter A. Boncz, Niels Nes, and Sándor Héman. MonetDB/X100 - A DBMS In The CPU Cache. *IEEE Data Eng. Bull.*, pages 17–22, 2005.
- [119] Marcin Zukowski, Sandor Heman, Niels Nes, and Peter Boncz. Super-scalar ram-cpu cache compression. In *ICDE*, pages 59–59, 2006.