THE UNIVERSITY OF CHICAGO

HIERARCHICAL RESIDUAL ENCODING FOR MULTIRESOLUTION

SPATIAL-TEMPORAL DATA COMPRESSION

A DISSERTATION PROPOSAL SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCE

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

DEPARTMENT OF COMPUTER SCIENCE

BY

BRUNO BARBARIOLI

CHICAGO, ILLINOIS

OCTOBER 2022

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

Data compression is a key technique for reducing the cost of data transfer from storage to compute nodes. Increasingly, modern data scales necessitate lossy compression techniques, where exactness is sacrificed for a smaller compressed representation. One challenge in lossy compression is that different applications may have different accuracy demands. Today's compression techniques struggle in this setting either forcing the user to compress to the strictest accuracy demand, or to re-encode the data at multiple resolutions. This work proposes a simple, but effective multi-resolution compression algorithm for time series data, where a single encoding can be effectively decompressed at multiple output resolutions. There are a number of benefits over current state-of-the-art techniques. (1) The storage footprint of this encoding is smaller than re-encoding the data at multiple resolutions. (2) Similarly, the compression time is generally smaller than re-encoding at multiple resolutions. (3) Finally, the decompression latency with our encoding is significantly faster than single encodings at the strictest accuracy demand.

We further propose an extension to the time series case, where we also take into account the spatial dependence of certain data sets. We generalize our model to a high dimensional setting with applications ranging from image segmentation to overlapping satellite images.

CHAPTER 1

INTRODUCTION

In today's data analytics infrastructure, it is common for data storage to be separated from computational resources Angel et al. [2020]. For example, very large datasets can be stored in a block storage system like Amazon S3, and only transferred to compute nodes for analysis. Similarly, in edge computing, data might reside on edge nodes for privacy and cost reasons, and only be transferred to a central location when needed Paparrizos et al. [2021], Zeuch et al. [2020]. In such on-demand retrieval architectures, the time needed to transfer data between storage and computation nodes is an important bottleneck, and data compression is one of the main techniques for controlling the cost of data movement.

Traditionally, data compression approaches are divided into two categories: lossless and lossy. In lossless compression, the data are compressed and decompressed without loss of any information. Examples of such approaches include dictionary coding for string compression Färber et al. [2012], Antoshenkov [1997], GZip/BZip for byte-sequence compression Seward [1996], and turbo-coding for integer compression Hanzo et al. [2002]. In contrast, lossy compression allows for minor errors in the reconstruction that would not affect a downstream application. By nature, lossy compression is mostly aimed at high-dimensional quantitative data. Examples include JPEG for images Wallace [1992], H.264 for video Wiegand et al. [2003], and a variety of techniques for scientific data Cappello et al. [2019], Krasowska et al. [2021], Son et al. [2014]. Lossy compression techniques sacrifice accuracy (the degree of errors in the reconstruction) for storage size (the size of the compressed data).

To cope with ever growing datasets, lossy compression has been increasingly adopted in edge-computing and sensing systems Paparrizos et al. [2021], Zeuch et al. [2020], Pelkonen et al. [2015b], Shahid [2019], Paparrizos et al. [2020]. Data compression can be pushed towards the point of data collection to save on downstream data transfer, storage, and computation. For example, a visual dashboard monitoring a sensor only needs the data to

be accurate up to the screen pixel resolutions. On the other hand, machine learning models that consume the data are often robust to small amounts of imprecision in the input features. While such "compression pushdown" can be extremely effective, it is most useful when there is only a single downstream application consuming the data.

Multiple downstream applications can have differing accuracy demands (e.g., a visual dashboard requires a `1e-3` maximum error in all values, but the anomaly detection framework only requires `1e-1` precision). In such cases, the available pushdown strategies essentially are: (Strict Encoding) encode the data once at the strictest accuracy demand and (Multiple Encoding) re-encode the data at all of the different accuracy demands. The *strict encoding* strategy has the obvious drawback that it forces every application to pay the same data transfer cost as the one with the strictest accuracy requirement. On the other hand, the *multiple encoding* strategy allows different applications to selectively retrieve data encoded at their particular accuracy demands. However, the multiple encoding strategy has a steep cost in terms of compression throughput (i.e., linear in the number of applications), and local storage (i.e., stores one encoding per application).

This paper navigates this impasse and studies efficient methods for supporting multiple downstream consumers of lossy data. Ideally, it should be possible to store a *single encoding that can be selectively decompressed at all the different resolutions* and thus mitigate the downsides of both strategies described above. We call this problem the *multiresolution compression problem*, where the objective is to construct a single encoding of a dataset that can be decompressed at different cell-level error tolerances (hereafter called L_∞ errors). While related problems have been proposed in a number of adjacent areas such as in reduced-precision ML Wang et al. [2019] and approximate query processing Lazaridis and Mehrotra [2001], Buccafurri et al. [2011]; to the best of our knowledge, multiresolution compression has not been extensively evaluated in the data compression literature. This problem setting subtly changes the typical metrics of interest for data compression. An effective multireso-

lution compression algorithm: (Compression ratio) should require significantly less storage than a separate encoding at each error tolerance; (Compression latency) should be significantly faster to construct the encoding than a separate encoding at each error tolerance; (Decompression latency) should be faster to decompress data for higher error tolerances. This means that a multiresolution compression algorithm might not be the most effective at any single error tolerance, but in aggregate across multiple tolerances, it outperforms the strict and multiple encoding strategies.

This paper proposes an effective multiresolution compression algorithm, called Hierarchical Residual Encoding (HIRE), for uni-variate and multivariate time series data. HIRE assumes a mini-batch data acquisition model where data are streamed to the system in small blocks, and these blocks represent contiguous time-segments of collected data. HIRE constructs a synopsis data structure through a recursion of partitioning and residualizing steps. The partitioning step approximates a time series with a piecewise approximation, and the residualizing step calculates a signal that represents the approximation error. This error signal can further be approximated at increasingly finer granularities. We find that these residual signals are often highly compressible, since many values may lie under the error threshold of the strictest application. To retrieve data at a particular error threshold, there is a simple summation of the preceding layers. This result is not surprising as such residualizing operations are effective in time series forecasting Chatfield [1978], and in differential equation solvers Hutchinson and Raithby [1986].

We evaluate HIRE against a key set of baseline compression algorithms that: (1) can run efficiently in the online setting, (2) provide L_∞ error guarantees for point queries like HIRE, and (3) do not require dataset-specific modeling for compression. Compared to these baselines, we find that HIRE has the following desirable properties:

- Compression ratio (or similarly storage size) is much smaller than independently compressing at each resolution using the other compression methods.

3

- Compression time is smaller than independently compressing at each resolution using the other compression methods

- In most cases, decompression time for an error threshold is lower than decompression time of any other method for the same error threshold.

CHAPTER 2

BACKGROUND

First, we will motivate the general problem statement, describe current compression approaches, and give context to our contributions in multiresolution compression. Through this paper, we will consider the following running example.

[Edge Data Storage and Retrieval] Data transfer over a network is one of the most expensive (in terms of cost and energy) tasks in any distributed sensing application. Many recent sensing architectures argue for lazy data retrieval Paparrizos et al. [2021], where data are persisted on the edge for some period of time and only centralized if/when needed. Let's consider a simplified two server version of this architecture. Consider a sensor deployed at a climate research observatory that collects a time series of numerical data. The data are compressed online at an "edge server" during data collection (located at the observatory), and stored locally for a maximum of 10 days. The compressed versions can be retrieved from the edge server, transferred to a remote server (e.g., at a research university), and decompressed at the remote server.

2.1 Time Series Compression Basics

Let us consider a time series with observations:

$$X = [x_0, x_1, \ldots, x_{T-1}].$$

For now, let us assume that each x_i is a single floating point number (i.e., a univariate time series). We will relax this assumption later, but it will be easier to understand the baselines in the univariate case. A compression algorithm consists of an encoder and decoder pair

$$C_X = \mathsf{enc}(X) \qquad X' = \mathsf{dec}(C_X)$$

that produces a compressed representation of X called C_X and reconstructs an estimate of the original time series X' from C_X. Without loss of generality, we can consider C_X to be a vector as well. There are several important metrics of interest that describe the performance of such a compression algorithm:

- **Compression Ratio.** Let $H(\cdot)$ denote the size in bits of a vector. The compression ratio is defined as $\frac{H(C_X)}{H(X)}$. A lower compression ratio indicates better performance in terms of storage. *In our running example, the compression ratio directly affects how much data are transferred from the edge to the remote server.*

- **Compression Latency.** The compression latency of an algorithm is the time needed to produce C_X from X. *In our running example, the compression latency affects the data collection throughput of the edge server.*

- **Decompression Latency.** The decompression latency of an algorithm is the time needed to produce X' from C_X. *In our running example, the decompression latency affects how much extra time beyond data transfer is spent on the remote server.*

- **Reconstruction Error.** The difference between X and X' is called the reconstruction error, for some measure of dissimilarity. *In our running example, the reconstruction error measures how different the data processed on the remote server is compared to the edge server.*

While there are many such compression algorithms, we primarily focus ones with deterministic L_∞ reconstruction error guarantees. That is, we bound the maximum allowable error of the reconstructed time series $X' = \mathsf{dec}(\mathsf{enc}(X))$ with respect to the original time series X by specifying an error threshold parameter ε. The L_∞ reconstruction error $||X-X'||_\infty$ is defined as the maximum absolute disagreement between X and X' at any time-step i:

$$||X - X'||_\infty = \max_i |x_i - x_i'|$$

6

We then enforce an error guarantee by ensuring that the L_∞ error is within the pre-specified threshold ε, which means $||X - X'||_\infty \leq \varepsilon$. For example, ε may represent the maximum error that the observatory is willing to tolerate in the reconstructed time series to avoid a negative impact on subsequent weather forecasting tasks. We choose such guarantees because they are the strictest and most compatible with a wide variety of applications. The main trade-off in most techniques is when ε is increased (i.e., more error), the compression ratio decreases.

2.2 Compression with Bounded L_∞ Error

There is an extensive body of literature on time series compression techniques (refer to survey Keogh et al. [2004]). However, not all compression algorithms can provide L_∞ error guarantees. For example, a spectral approach like PCA, or an FFT, can only control the average error but not the worst-case error for any given observation. For techniques that do provide L_∞ error guarantees, they generally follow the same design pattern composing three main components: (1) Temporal Decorrelation, (2) Quantization, and (3) Byte Encoding.

2.2.1 Quantization

The most basic technique for compressing numerical data with an error guarantee is quantization. Quantization is the process of rounding a floating point number to nearest valid value of fixed precision. Even simple data quantization can be very effective and is employed in Amazon Redshift Gupta et al. [2015]. Proceeding to the technical formulation of quantization, let us suppose that x^-, x^+ are the minimum and maximum of X respectively. Define R_ϵ to be the *relative L_∞ error*, i.e., a desired error tolerance relative to the range of X:

$$R_\epsilon = \frac{\varepsilon}{x^+ - x^-}$$

A uniform quantization scheme cuts the range $[x^-, x^+]$ into $\frac{1}{R_\epsilon}$ equal-sized buckets [1]. To encode a dataset using this scheme, one sets each numerical value to its resident integer bucket (note that the floor function discretizes the formerly continuous input):

$$\mathsf{enc}(x_i) = \left\lfloor \frac{(x_i - x^-)}{(x^+ - x^-)} \cdot \frac{1}{R_\epsilon} \right\rfloor \tag{2.1}$$

We examine the quantization formula, we notice that the main term $\lfloor \frac{(x_i - x^-)}{(x^+ - x^-)} \cdot \frac{1}{R_\epsilon} \rfloor$ is integer-valued and ranges from 0 to $\lceil \frac{1}{R_\epsilon} \rceil$. In its most basic implementation, we can assign a fixed-length binary code to each of the integer values. This would result in storage size of $\log_2 \lceil \frac{1}{R_\epsilon} \rceil$ per values.

To decode, one simply reverses the transformation:

$$x_i' = \mathsf{dec}(\mathsf{enc}(x_i)) = \mathsf{enc}(x_i) \cdot R_\epsilon \cdot (x^+ - x^-) + x^- \tag{2.2}$$

As long as we also store x^+, x^-, R_ϵ, we can translate the stored integer into its corresponding floating point quantum. Unfortunately, regardless of whether quantization is applied to a column sampled from a normal distribution or to a time series with high autocorrelation, it yields the same mapping.

2.2.2   Temporal Decorrelation

For this reason, either before or after quantization most time series compression algorithms attempt to factor out all of the "predictable" terms, thereby only leaving uncorrelated errors to be compressed. The simplest approach to accomplish this is delta-encoding, which transforms a time series so that every value is represented as a successive difference from the previous value:

$$x_i^\delta = x_i - x_{i-1}$$

---

1. In the degenerate case of $\epsilon = 0$ one just keeps each element of X in its own bucket

This transformation is completely reversible with a left-to-right cumulative sum, so no information is lost. Since time series often have highly similar values along the time axis, the range of delta values is smaller in magnitude and variation than the original values. Thus, the deltas can often be effectively stored with reduced precision.

We can think of delta-encoding as a simple form of "predictive" compression. The previous value x_{i-1} can be interpreted as a simple model that predicts the value of its neighboring element Blalock et al. [2018]. The same trick would work for any function f of the previous j lagged elements where $0 < j < i$:

$$x_i^\delta = x_i - f(x_{i-1}, x_{i-2}, \ldots, x_{i-j})$$

Given the dependence of the current value on the previous (lagged) values, the structure of delta encoding is suitable for autoregressive models of any flavor. It is worth noting that the better the predictive model, the more skewed the X_δ values will be towards zero. If the effective domain of the numbers can be greatly reduced, then fewer bits can be used to represent X_δ.

In general, any time series modeling technique can be used to decorrelate the data. For example, there are a number of piecewise approximations for time series that exploit trend structure in a typical time series. Piecewise approaches decompose a time series into segments and use the segments to approximate the time series such as the: Piecewise Aggregate Approximation (PAA) Keogh and Pazzani [2000] and Piecewise Linear Approximation (PLA) Keogh et al. [2001]. Like delta encoding, we can think of piecewise approximation as a simple model that predicts the next value. One only needs to store the model and the compressed residual. If this model is accurate, the residual error is likely very sparse and highly compressible.

2.2.3 Byte-Encoding

Finally, after quantization and decorrelation, general-purpose compression algorithms can be used to simply translate the remaining data into a series of bytes and reduce redundancy. Most of these approaches are based on run-length encoding or the LZ77 algorithm that look for byte-level repetitions within sliding windows of data Seward [1996]. Generally speaking, byte-oriented techniques are lossless—meaning that the L_∞ error is always 0—so they can provide a trivial error guarantee. On numerical data the obvious limitation is that while two floating point numbers are close to each other numerically, there might not be a very strong similarity between their binary representations leading to poor compression with an LZ77 technique. This is why quantization and decorrelation are generally applied first to exploit the numerical structure.

2.3 The Multiresolution Problem

For any combination of quantization and decorrelation described above, the encoded representation C_X is tied to a specific target error. Let us consider how this can cause an unnecessary performance bottleneck in our running example.

[Two Applications With Different Error Tolerances] Consider two remote applications consuming the climate data collected at the observatory. The first application is a nightly report that is generated over all of the collected data. This report requires that every value processed is no more than $\varepsilon = 1e - 4$ of the true value. The second application is a minute-by-minute anomaly detector to detect whether the observatory has abnormal readings. This application is far more tolerant to error and requires that each value processed is only within $\varepsilon = 1e - 1$ of the true value.

If we use the state of the art, there are three clear solutions—all of which have significant drawbacks.

- **Compress at the Strictest Tolerance.** The most obvious approach would be to compress the data at the strictest error tolerance. Unfortunately, this would mean that every retrieval would pay a data transfer cost of the strictest threshold. *In our running example, the frequent anomaly detection tests would have to repeatedly transfer data encoded with a target of $\varepsilon = 1e - 4$ due to the relatively infrequent nightly reporting.*

- **Compress at All Tolerances.** Another approach would be to encode the data at all relevant error tolerances. While this approach allows each application to only transfer data encoded at an appropriate error tolerance, it shifts the burden towards encoding. Each additional error target would reduce the effective throughput available for data collection since we are compressing the data twice. *In our running example, we would cut our ingestion capacity by roughly a factor of two.*

- **Lazy Re-Encoding.** A hybrid approach would be to first encode the data at the strictest error tolerance, then at retrieval time decode the data, and re-encode it at the error target of each retrieving application. The core challenge is that decompression is often significantly slower than compression in many popular algorithms, and this hybrid approach incurs these costs at the edge. *In our running example, the frequent anomaly detection tests would trigger expensive re-encoding processes that would burden the edge server.*

These drawbacks suggest the need for a new type of time series compression algorithm aimed at supporting multiple downstream applications. Ideally, there should be a single encoding that can be selectively decompressed at different target resolutions.

[Multiresolution Compression] A multiresolution compression algorithm produces a single encoding C_X that can be selectively decomposed into sub-encodings $C_X^{\epsilon_1}, \ldots, C_X^{\epsilon_l}$ with corresponding error thresholds $\epsilon_1, \ldots, \epsilon_l$:

$$C_X = C_X^{\epsilon_1} \bigoplus C_X^{\epsilon_2} \bigoplus \cdots \bigoplus C_X^{\epsilon_l}$$

11

where $\oplus$ denotes some combination operation of the sub-encodings.

Decomposable encodings allow one to selectively transfer data for any target resolution. The multiple encoding strategy described above can be thought of as a trivial case where $\oplus$ is simply the concatenation operation over independent encodings of the data. The key issue with this strategy is that no work is shared between any of the encodings, and effectively sharing work will be the main premise of this paper. We will show that an additive decomposition where $\oplus$ is a linear combination operation is a simple but effective solution. Consequently, the goal of this paper is to investigate such algorithms, understand how to evaluate them, and how to optimize for different performance objectives in the multiresolution setting.

CHAPTER 3

MATHEMATICAL INTUITION

Now that we have motivated multiresolution compression, we will provide some basic technical intuition how such an algorithm can be constructed.

3.1 Residualization

From the previous section, let X be a time series and X' be an approximation of it (e.g., a decompression of a lossy encoding). The *residual* series is also a time series and is defined as:

$$R = X - X'$$

which is the difference between the original series and its reconstruction. It clearly also follows from this definition that if one knows the approximation and the residual, one can fully reconstruct the original series $X' + R = X$. Such a decomposition of terms is called an additive decomposition and is well-studied in modeling time-dependent phenomena outside of data compression Chatfield [1978]. For example, it is common to decompose time series into a trend component (which represents the general trend of the series) and noise (which represents variation along the trend). The trend component would be our X' and the noise would be our R.

Additive decompositions are recursive in nature since the residual R is itself another time series and can be further decomposed. Now, let us suppose that we have an approximation to R denoted R'. We can similarly define a residual series $S = R - R'$. Through substitution, we can see that our additive decomposition now looks like this:

$$X = X' + R' + S$$
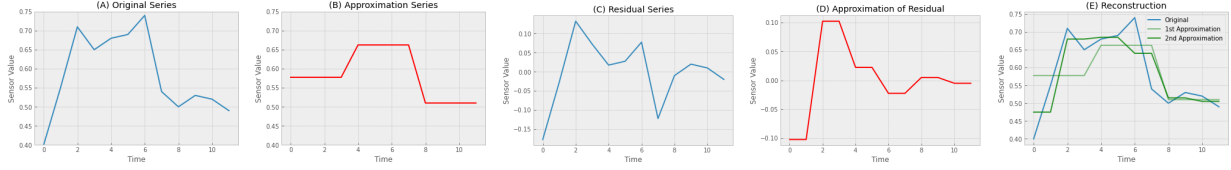
Figure 3.1: (A) Illustrates an example time series, (B) illustrates a coarse approximation of the time series, (C) illustrates the residual series, (D) illustrates a finer approximation of the residual series, and (E) shows how the summations of the residuals lead to progressively better approximations

This process can be repeated, further and further decomposing the residual, which yields a natural recurrence equation:

$$R_0 = X \tag{3.1}$$

$$A_i = \mathsf{approx}(R_i)$$

$$R_i = R_{i-1} - A_{i-1}$$

where $\mathsf{approx}(\cdot)$ is some approximation function. We can clearly see from this equation that for k such recursions:

$$X = \sum_{i=0}^{k-1} A_i \;\; + \;\; R_k$$

We leverage this basic intuition to construct an effective multiresolution compression algorithm. Written another way the equation above is simply $\sum_{i=0}^{k-1} A_i \approx X$. We need to construct a sequence of successive approximations that reduces the size of the residual signal until the residual is less than our desired error threshold.

3.2 Relevance to Time Series Compression

The next section will show how to leverage a series of successive refinements of the residual series to represent the original series. At each step, we compress the residual vector from the previous step, progressively increasing the fidelity of the compression. As more of the variation in the data is explained by each subsequent step, the residual vector becomes

14

smaller in magnitude and sparser.

Let us consider a concrete example. In Figure 4.1(A), we have plotted an example time series. One can coarsely approximate this time series with a piecewise constant approximation of 3 segments (Figure 4.1(B)). Between this approximation and the original series, there is a residual (Figure 4.1(C)). This residual can be captured by another piecewise constant approximation with 6 segments (Figure 4.1(D)). As long as each subsequent approximation includes additional information (i.e., increasing the number of segments), the remaining residual series reduces in magnitude. A summation of these approximations gets closer to the original value (Figure 4.1(E)). While the intuition is simple, realizing this idea turns out to be more difficult. The next section describes how we can enforce an error guarantee in this process of successive refinement and how to implement such an approach efficiently.

3.3   Novelty

Our new compression method for recursively approximating a time series has some similarities with other classical mathematical methods for approximating functions such as Fourier and Taylor methods. For instance, a Taylor series can in fact provide an L_∞ (upper bound) error for each distinct partial sum, and the error decreases as the size of the partial sum grows. However, there are key differences in our formulation because of how we refine our approximation and measure the resulting errors/bounds. In particular, in Taylor and Fourier series approximation error is only used to decide when to terminate the series (i.e., a stopping condition), which is entirely different from how we recursively approximate the error values themselves. We leverage the intuition that residuals often contain a lower entropy signal than the function values and by recursively applying approximation to the residuals at increasing granularity, the entropy of each successive level may decrease, leading to better compressibility. By setting the midrank function as the pool function, we show that the L_∞ error is strictly non-increasing with respect to the level of the hierarchy, a similar theoreti-

15

cal guarantee to the Taylor series; but unlike the Taylor series, HIRE operates on residual vectors and does not require that the function is analytic.

The novelty of our submission is therefore: i) a new problem definition highlighting the emerging need for multiresolution (de) compression systems, ii) hierarchical recursive approximation of residual vectors in the domain of time series compression with L_∞ guarantees—which has not been proposed before, iii) the use of pool function properties along with vector theory to propose linear time compression and decompression algorithms computing the errors on the fly, and iv) practical implementation of the theoretical algorithms making use of vectorization and parallel computing.

CHAPTER 4

HIRE: HIERARCHICAL RESIDUAL ENCODING

Based on our intuition from the previous section, we design an algorithm that constructs a progressively refined set of residual signals. This algorithm is called Hierarchical Residual Encoding (HIRE). HIRE can be thought of as a special temporal decorrelation operation much like delta coding or piecewise aggregate approximation. However, its special hierarchical structure allows for selective decoding at multiple error thresholds. We describe our method considering a univariate time series. We can also handle multivariate time series by running independent encodings.

4.1 Algorithm Basics

Before we introduce the algorithm, it would be informative to define a few general building blocks. As before, let $X = [x_0, x_1, \ldots, x_{T-1}]$ be a univariate time series represented as a vector of data, where $x_i \in \Re$ and $X \in \Re^T$. We further assume that the user provides us with a ε^* or the strictest error threshold that the encoding must guarantee.

4.1.1 Quantized Pooling

HIRE relies on a piecewise approximation to approximate each residual series: over disjoint windows, the value of the window is approximated by a single scalar aggregate. To use machine learning terminology, this operation is called (temporal) pooling. Pooling reduces the dimensionality of a data series along the time axis over a series of fixed-size windows. The pooling operation is defined by two parameters, an aggregation function f and a time series X. In particular let $p_f(X) : \Re^n \to \Re$, $p_f(X) = f(x_0, \ldots, x_{n-1})$, for any positive integer n. In other words, a pool function provides a concise estimate of a given time window with a single value.

The choice of pooling function is a hyper-parameter and different pooling functions have unique properties. We can think of them as different ways of approximating the values in a time window. For a window $X = [x_0, \ldots, x_{n-1}]$, we define:

- Mean: $f(X) = \frac{1}{n} \sum_i x_i$. The mean value is a natural pooling function that minimizes the squared error with respect to the window.

- Midrank: $f(X) = \frac{1}{2} (\max_i\{x_i\} - \min_i\{x_i\})$. We can show that midrank is actually the optimum pool function to minimize the L_∞ error of the residual vector in each node of the hierarchy, as described in the next subsection.

- Median: $f(X)$ returns the median value of vector X.

- Random: $f(X) = x_i$ with probability $1/n$. We can show that a random pooling function is robust to seasonal variation within windows of a time series.

To efficiently store the results of a pooling operation, we further quantize the aggregate value to a particular error threshold (using the process described in Section 2). In this way all pool values are integers, so more efficient compression methods can be used to compress/decompress them. During the decompression we derive the pool values P and round them to the corresponding real values before we take their sum.

4.1.2 Spline interpolation

Without quantization, pooling is a lossy operation for $n > 1$ and is only lossless for $n = 1$ (i.e., window size of 1). Thus, inverting this operation will only give us an approximation of the original series. To do so, we require some function that estimates the original time series from the pooled values. We define a spline function $s_w(p_f(X)) \approx X$, where $w = |X|$. The user can choose the spline function that they prefer. Our default option is a simple interpolation function that duplicates the value $p_f(X)$, w times. In particular, $s_w(p_f(X)) =$

$[p_f(X), \ldots \times w]$. This duplication-oriented spline can be thought of as constructing a step function interval with length w and value $p_f(X)$. Specifically, the default spline computes a step function where each interval captures a pooled value that covers w time steps. In doing so, we map from a lower dimensional summary to an approximation of the input series of a certain coarseness. With an increasingly smaller w, the step approximation of X improves proportionally to w itself.

4.1.3   Residual Vectors

The spline function returns an approximation of the original time series. Let $R_X = X - s_w(p_f(X))$ be the residual vector representing the error of the spline function with respect to the original vector. The key insight of our work is that residual vectors are generally more compressible than the original time series. It is clearly true that $X = R_X + s_w(p_f(X))$, and we can plug in this approximation into our recurrence equation of the previous section (Equation 3.1), where $\mathsf{approx}(R_i) = s_w(p_f(R_i))$.

4.2   Algorithm Description

The main idea is to run the procedure described above recursively following a hierarchical manner following the recurrence described in Equation 3.1. For simplicity we consider that size of the time series is a power of 2, so $T = 2^k$. First, we check if the L_∞ norm of vector X is at most *. If yes, then we can stop the recursion setting the pool value to 0. Otherwise, we compute the pool function $p_f(X)$ over the entire vector X. Then we set $R_X = X - s_w(p_f(X))$ as described above. In the next level of the recursion we set $X \leftarrow R_X[0, \ldots, T/2 - 1]$ and $X \leftarrow R_X[T/2, \ldots, T - 1]$, and we continue with the same way in both branches. Let $P_i = [p_{(i,0)}, \ldots, p_{(i,2^i-1)}]$, for $i = 0, \ldots, k$, be the vector of the pool values in the i-th level of the hierarchy sorted from left to right. Let $P = \bigcup_{i=0,\ldots,k} P_i$. We stop the recursion when the error is at most * or after having $|X| = 1$. In other words, as

19

we traverse down the hierarchy, there is a successively more accurate approximation of the residual at each level. After running this algorithm, the quantized pool values can be stored with any byte-encoding format.

We describe the pseudocode in Algorithm 1 (for simplicity, we describe the pseudocode considering that P is a global set of variables over the recursion) and provide a visual in Figure 4.1.

Algorithm 1: HIERARCHICAL

    Input   : $X, i, j,^*$
    Output: P

1 **if** $||X||_\infty \leq^*$ **then**
2   $p_{(i,j)} = 0$;
3   **return**;
4 $T = |X|$;
5 $p_{(i,j)} = p_f(X)$;
6 $R_X = X - s_T\left(p_f(X)\right)$;
7 **if** $T > 1$ **then**
8   HIERARCHICAL$(R_X[0, \ldots, T/2 - 1], i + 1, 2 \cdot j,^*)$;
9   HIERARCHICAL$(R_X[T/2, \ldots, T - 1], i + 1, 2 \cdot j + 1,^*)$;

Intuitively, we can think of a binary tree structure where the original time series lies in the root and its residual vector is split into two equal length sub-vectors creating two children in the tree structure. In each node of the tree, we store the corresponding singular pool value. Let $\mathcal{T}$ be the tree structure representing the hierarchical compression. For a node u in $\mathcal{T}$, let p_u be the pool value in this node. For example, if u is the j-th node in the i-th level, we have $p_u = p_{(i,j)}$. Let also R_u be the residual vector that is found from our algorithm at node u.

We note that it is not necessary to start the hierarchical compression from the zero level—considering the entire time series X. Instead, we can start the hierarchical compression from any level L, splitting the original time series into 2^L parts, $X[0, \ldots, T/2^L - 1], X[T/2^L, \ldots, 2 \cdot T/2^L - 1], \ldots, X[(2^L - 1) \cdot T/2^L, \ldots, T - 1]$ and run the Hierarchical al-
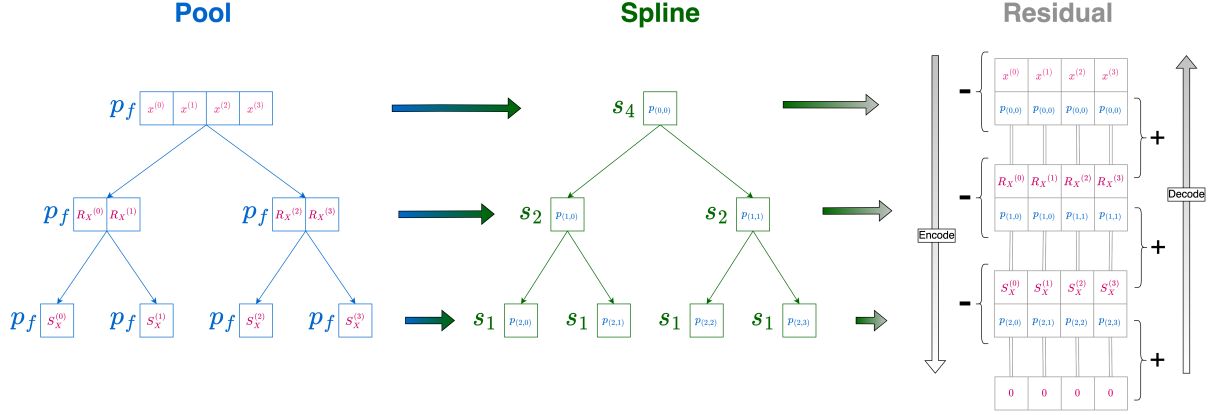
Figure 4.1: The main algorithmic workflow. Compression flows from left to right and top to bottom just like reading text. Decompression flows from bottom to top and inverts residualization with a summation.

gorithm for each of them independently. In the experiments, we mostly run the hierarchical compression for the last 10 or 20 levels, i.e., we set $L = k - 10$ or $L = k - 20$.

Finally, we note that it is not necessary to keep the pool values in different variables $p_{(i,j)}$ or p_u storing the indexes (i, j) or u. We only use this notation to make the description of the algorithm easier. Algorithm 1 can put all pool values in a single table P following the ordering: $P[h_1] = p_{(i_1, j_1)}, P[h_2] = p_{(i_2, j_2)}$ for $h_1 < h_2$ iff $i_1 <_2$ or $i_1 = i_2$ and $j_1 < j_2$. When the decompression algorithm needs to access $p_{(i,j)}$, it corresponds to the element $P[2^i + j - 1]$, so we have direct access in $O(1)$ time. Furthermore, as we will see in the next subsection, sometimes we might need to have access to the pool value of the parent or a the left (right) child of a node u. If u corresponds to the j-th node in the i-th level then $p_{(i-1, j \mod 2)}, p_{(i+1, 2 \cdot j)}, p_{(i+1, 2 \cdot j+1)}$ is the pool value of the parent, left child, and right child, respectively.

4.3   Decompression

Before we describe how we can decompress P to derive an approximation with L_∞ reconstruction of X (lossy) or the exact X (lossless) we show an interesting property of the residual

vectors over the nodes of $\mathcal{T}$.

Let u_m be the j-th node of the i-th level in the hierarchy. Let $\bar{X} = X[j \cdot w, \ldots, (j+1) \cdot w - 1]$, where $w = T/2^i$, be the corresponding part of the original time-series. Let $u_0 \to \ldots \to u_m$ be the path from the root of $\mathcal{T}$ to u_m. From their definitions, $R_{u_m} = \bar{X} - \sum_{\ell \le m} s_w(p_{u_\ell})$. Hence, it follows that $\max_h\{|R_{u_m}[h]|\} = ||\bar{X} - \sum_{\ell \le m} s_w(p_{u_\ell})||_\infty$, i.e., the L_∞ error of the sum of the spline vectors from the root node to the current node with respect to the original $\bar{X}$ vector, is the L_∞ error of the residual vector R_{u_m}.

We extend the previous observation to each level of $\mathcal{T}$. For each level $i \le k$ let E_i be the maximum L_∞ error of all residual vectors found in level i. Let $E = \bigcup_{i \le k} E_i$. Our system compresses both P, E using any known compression method. If M_1, M_2 are the compression methods used for P, E, respectively, the overall compression ratio of our method is $\frac{H(M_1(P)) + H(M_2(E))}{H(X)}$, where again $H(\cdot)$ denotes the size in bits.

Multiresolution Decompression   Let us see how our algorithm works with our running example. Recall that we have two different applications that require error thresholds of $\varepsilon = 1e - 4$ and $\varepsilon = 1e - 1$ after retrieving data from an edge server. Firstly, using HIRE, we construct a compressed residual encoding with $\varepsilon^* = 1e - 4$. Along with every retrieval request, a desired error threshold is sent . The edge server first finds the maximum error E_m such that $E_m \le$. Then we transfer all compressed pool values P_i for $i \le m$, from the edge server to the remote machine that made the request. Finally, the decompression procedure runs on the remote machine, finding $X' = \sum_{i \le m} s_{T/2^i}(P_i)$. Using the observations above, we have the guarantee that $||X - X'||_\infty \le$.

CHAPTER 5

OPTIMIZATIONS

In this chapter we describe multiple algorithmic and implementation optimizations that improve the running time of our algorithms.

5.1 Algorithmic Optimizations

5.1.1 Compression

We first note that Algorithm 1 runs in $O(kT)$ time. There are at most k levels of recursion and in each level we construct the residual vectors of all nodes. Hence, in each level we spend $O(T)$ time. We show how we can run Algorithm 1 in only linear, $O(T)$ time.

The main observation is that we do not need to construct the residual vectors explicitly. These vectors are helping to run the recursion algorithm and at the same time show what is the maximum L_∞ error in each node. We claim that we can still run the same algorithm without constructing the residual vectors. Let u be the j-th node in the i-th level and let $v_0 \to \ldots \to v_{i-1} = v_i = u$ be the path of the nodes from the root to node u. Let also $\bar{X} = X[j \cdot w, \ldots, (j+1) \cdot w - 1]$, for $w = T/2^i$, be the subset of time-series X that corresponds to the part of node u. It is straightforward to see that: $p_{(i,j)} = p_u = f\left(\bar{X} - \sum_{\ell=0}^{i-1} s_w(p_{v_\ell})\right)$. This is a very important observation because it shows that using only the original time series along with the previously computed pool values we can get the new pool value we need to store and compress.

Next, we show an efficient way to compute $f\left(\bar{X} - \sum_{\ell=0}^{i-1} s_w(p_{v_\ell})\right)$. Of course, the actual algorithm depends on the function f. Hence we check all of the main functions we used. This method can be extended to a large family of functions. For all functions we are using, we have the next observation: $f\left(\bar{X} - \sum_{\ell=0}^{i-1} s_w(p_{v_\ell})\right) = f(\bar{X}) - \sum_{\ell=0}^{i-1} p_{v_\ell}$. It is easy to maintain the second term $\sum_{\ell=0}^{i-1} p_{v_\ell}$ during the execution of the compression algorithm. Let S_v be the

sum of all pool values from the root to the node v. We can update $S_{child(v)} = S_v + p_{child(v)}$ in constant time. Hence, we only focus on how to compute $f(\bar{X})$ efficiently. In particular, we aim for constructing a data structure $\mathcal{D}$ in $O(T)$ time such that given a query range $[a, b]$, compute $f(X[a, \ldots, b])$ in $O(1)$ time.

We start considering the midrank function f. We pre-process X and we build a range MAX/MIN data structure $\mathcal{D}$ using the LCA technique Fischer and Heun [2007]. It is known that $\mathcal{D}$ can be computed in $O(T)$ time, it has $O(T)$ space, and can answer a range MAX or MIN query in $O(1)$ time. Hence we can run Algorithm 1 in $O(T)$ time.

Next, we consider the mean function f. Again, we need a data structure to find the mean of X in a query range. We compute and store the prefix sums of X: $\mathcal{D}[h] = \sum_{z=0}^{h} X[z]$. Overall, we construct a data structure $\mathcal{D}$ in $O(T)$ time such that given a range $[a, b]$ we return $f(X[a, b]) = \frac{\mathcal{D}[b] - \mathcal{D}[a-1]}{b-a+1}$ in $O(1)$ time. Again, we run Algorithm 1 for the mean function in $O(T)$ time.

It is straightforward how to get a random item in $X[a, \ldots, b]$ efficiently for the random function f. We just get a random number $t \in [a, b]$ and we return $X[t]$. So we can also run Algorithm 1 in $O(T)$ time for the random function. Next, we note that it is also straightforward to run the compression algorithm for the pool functions with quantization. The only difference is that when we find $f(X[a, b])$ we compute in $O(1)$ time the integer bucket it belongs to. Unfortunately, to the best of our knowledge, there is not any known data structure to compute the median function in a query range in $O(1)$ time, using at most linear pre-processsing time.

Finally, the data structure we used for the midrank pool function is actually needed in all pool functions to measure the L_∞ error in each node/level in the hierarchy. Previously, having a residual vector we could find the L_∞ error by checking the absolute values of its elements. In the optimum algorithm we do not construct explicitly the residual vectors, so we cannot do the same procedure. Instead, we argue as before. From the definitions

it follows that $R_u = \bar{X} - \sum_{\ell=0}^{i} s_w(p_{v_\ell})$. Hence, $e_u = ||R_u||_\infty = \max_h\{|\bar{X}[h] - S_u|\} = \max\{|\max_h\{\bar{X}[h]\} - S_u|, |\min_h\{\bar{X}[h]\} - S_u|\}$. As we explained above, we can calculate S_u (from the parent node) in $O(1)$ time. Using the same data structure $\mathcal{D}$ we used for the midrank function Fischer and Heun [2007], we can find the MAX and the MIN value of the original time series X in a query range, in $O(1)$ time. Hence, we can compute the L_∞ error in a node u of $\mathcal{T}$ in $O(1)$ time. The overall running time of Algorithm 1 for every pool function f we use (except the median), including the error calculation, is $O(T)$.

5.1.2 Decompression

The decompression algorithm we described in the previous section runs in $O(kT)$ time since we take the sum of the residual vectors to retrieve the original time-series or an approximation of it. We show how we can execute the decompression algorithm in only linear $O(T)$ time. More specifically, the algorithm can be executed in time linear to the number of nodes in the hierarchy that we need to retrieve to run the decompression procedure. While the algorithm is more tedious to describe than the compression algorithm, it is independent of the pool function f we used in the compression phase.

For simplicity, assume that the decompression method needs to take the sum over all the pool values (using spline interpolation) up to level $L \leq k$. The algorithm can be extended in case that we need to take the sum of vectors from different levels. We describe an algorithm doing it without computing the spline interpolation vectors explicitly. The main idea is the following: We run a sweep-line algorithm starting from left to right maintaining the total sum of all the corresponding pool values in the hierarchy. For example, imagine that we are currently considering index $h \leq T$ in the leaf node u. Let S_h be the total sum of the pool values from the root to u. We observe that the decompressed value is $X'[h] = S_h$. Hence, the goal is to maintain the correct values S_h over all indexes h from left to right. In order to derive the value S_h from S_{h-1}, we subtract the pool values that correspond to the nodes

that the index $h - 1$ belongs to and add the pool values that correspond to the nodes that the index h belongs to. The pseudocode can be seen in Algorithm 2.

Algorithm 2: FASTDECOMPRESSION

   Input : P
   Output: X'
1 $S_{-1} = 0$;
2 **for** $h = 0$ *to* $T - 1$ **do**
3    $i = L$;
4    $S_h = S_{h-1}$;
5   **while** $h \mod \frac{T}{2^i} == 0$ *AND* $i \geq 0$ **do**
6      $j = \frac{h}{T/2^i}$;
7      $S_h = S_h + p_{(i,j)} - p_{(i,j-1)}$;
8      $i = i - 1$;
9    $X'[h] = S_h$
10 **return** X';

The while condition in line 5 checks if there is a change in the pool values from index $h - 1$ to h at level i. Then variable j stores the node at level i that the index h belongs to. We update the value S_h subtracting the pool value $p_{(i,j-1)}$ (i.e., pool value in $(j-1)$-th node at level i) and adding the pool value $p_{(i,j)}$ (i.e., pool value in j-th node at level i). We recall that the notation $p_{(i,j)}$ is only used to simplify the description of the algorithm. We can always access the pool value in the j-th node in the i-th level taking the value $P[2^i + j - 1]$ in constant time. Algorithm 2 visits each compressed value in P two times, one to add it to the sum and one to subtract it from the sum. Hence the running time is $O(T)$.

5.2   Implementation Optimizations

Moving now from theory to practice, here we highlight a few systems best practices geared towards implementing a scalable encoder-decoder pair—that is, a pair with a low latency and runtime memory footprint.

5.2.1 Compression

The first step in optimizing the encoder (Algorithm 1) on a CPU is to convert the recursive formulation to an iterative formulation—thereby avoiding the allocation of unnecessary stack space. At each level, we apply the pool and spline operations in succession. A naive implementation of pooling might map the function of choice to each individual window. However, there is a large amount of extra work in that we might perform an unnecessary memory allocation operation to rearrange the array into windows of size n and then redundantly compute w small sums—only to eventually divide each sum by the exact same value. Instead, we can compute a prefix sum over the entire input with a single optimized function call and then in constant time deposit each pooled value into a pre-allocated array. This technique exploits the vectorization and the instruction-level parallelism present in conventional superscalar processors.

Now suppose that we have a multivariate time series $Y \in \Re^{T \times p}$ that consists of p univariate columns. We can apply the encoding algorithm to each column in parallel and thus achieve a latency speed up. Further suppose that we have b blocks in each univariate column. We can also apply the encoding algorithm to individual blocks—or groups of blocks, for that matter—which introduces a more granular form of parallelism.

5.2.2 Decompression

As an alternative approach to the linear time technique expressed in Algorithm 2, we can also in principle exploit threaded parallelism within the decoder. The HIRE decoding algorithm must calculate a linear combination over a large number of recomputed spline arrays. This summation need not be done in a sequential order due to the fact that the pooled values are already in memory. Visually, we can partition the hierarchy along the depth axis of the tree such that each individual spline reconstruction and summation operation (Figure 4.1) is assigned to a single thread. As a concrete example, if ten residual subtraction operations

are performed during encoding, ten addition operations must be performed during decoding. If we have two cores available, then we can assign five operation pairs (reconstruction and summation) to one thread and the remaining five to the other thread. We then perform a meta summation over the vectors returned by each thread which therefore yields the reconstructed time series X'.

5.3 Extensions

We briefly discuss some extensions of our main method. In particular, we extend our method to work with other error functions, and we show how we can split a time series to optimize the hierarchical compression algorithm.

We extend our algorithms to alternative error guarantees. In particular, we show how our compression method can bound the error of any L_p norm, extending the previous results for the L_∞ norm. For simplicity we focus on L_1, and L_2 norms. The main observation is that the residual vector R_X explicitly computed by Algorithm 1, or implicitly computed by the optimized algorithm, contains the absolute differences from the original vector. Hence, we argue that the L_p error of a node u is the L_p norm of vector R_X in node u. More specifically, in line 1 of Algorithm 1, we check whether $||X||_\infty \leq^*$. For any L_p norm we can use the condition $||X||_p \leq^*$ to check the L_p error in the current node of the hierarchical compression. If we want to measure the overall L_p error of the compressed time series we take the sum of the errors over the nodes within the same level. We can show that the optimum pool function to minimize the L_1 error is the median function, while the optimum pool function for the L_2 error is the mean function. For all L_p error functions, the compression algorithm runs in $O(kT)$ time, as we had with the L_∞ error (recall that k is the number of levels in the hierarchical compression). Furthermore, the linear time optimized algorithm can be applied for the L_2 error. The L_2 norm can be computed without constructing the residual vector explicitly by constructing prefix sums for both the values of the original time series and their

squared values. The mean function can be computed also in constant time in each node as we described in Section 5.1.1. On the other hand, the median function and the L_1 norm cannot be computed in $O(1)$ time (doing linear preprocessing) so the optimized algorithm cannot be used for the L_1 error. Finally, notice that the linear time optimized decompression algorithm is independent of the error function.
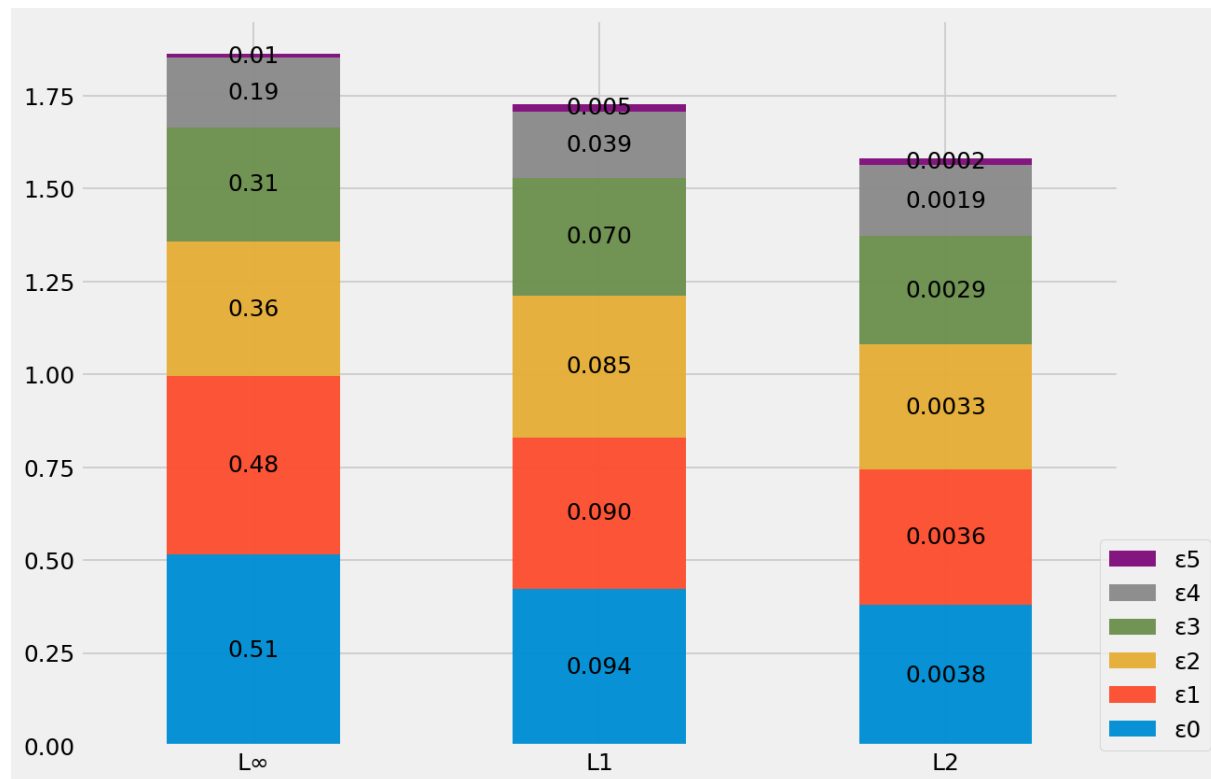

Figure 5.1: L_∞, normalized L_1, and normalized L_2 error values at the final 6 levels of $\mathcal{T}$ on 1024 samples from a IHEPC time series. The root level is blue, and we traverse upwards to the purple leaves. Note that the L_1 and L_2 bars are multiplied by constants c_1, c_2 for scale, with actual values shown in black.

So far, the compression algorithm always splits the time-series at the middle. While this method has many advantages, it might be possible that a different split may decrease the overall error in a higher rate. Hence, we explore different ways to split a time series during the compression algorithm. For example, using the optimum midrank function for bounding the L_∞ error, the best option is to split the time series such that the maxi-

mum difference of elements in each sub-time series is minimized. For example, given a time series X with N elements we want to find the element j such that $\max\{|\max_{i \leq j} X[i] - \min_{\ell \leq j} X[\ell]|, |\max_{i > j} X[i] - \min_{\ell > j} X[\ell]|\}$ is minimized. In order not to define a different optimization problem for each error function and each pool function, we consider the following splitting function to split a time series in any scenario: We split at the element j such that the sum of squared errors of the two sub-time series is minimized. In particular, the sum of squared errors splitting a time series X on j is defined as $\max\{\sum_{i \leq j}(X[i] - \bar{X}_{\leq j})^2, \sum_{\ell > j}(X[\ell] - \bar{X}_{>j})^2\}$, or $\sum_{i \leq j}(X[i] - \bar{X}_{\leq j})^2 + \sum_{\ell > j}(X[\ell] - \bar{X}_{>j})^2$, where $\bar{X}_{\leq j}$ is the mean of $\{X[1], \ldots, X[j]\}$ and $\bar{X}_{>j}$ is the mean of $\{X[j+1], \ldots, X[N]\}$. Overall, the sum of squared errors is defined as the sum of variances of the two sub-time series multiplied by their cardinalities. Intuitively, it captures how homogeneous each sub-time series is. Ideally, we would like to create homogeneous time series so that by applying a pool function we minimize its error. We show that both functions, maximum difference and the sum of squared errors, are increasing with respect to the number of elements in a time series. Hence, we run a binary search on the elements of X and for each element j we check, we evaluate the splitting function in both ranges $[1, .., j]$ and $[j + 1, .., N]$. Based on their comparison we continue the binary search either for larger or lower values of j. By constructing a data structure in linear time in the pre-processing phase, we can evaluate both splitting functions in a query range in $O(1)$ time, as described above. The binary search takes $O(\log T)$ steps, so the overall compression algorithm takes $O(T + k \log T) = O(T)$ time.

While, non-trivial splitting functions can help reducing the error faster, we notice that we should explicitly store the splitting index j (in contrast to splitting always at the middle where we do not store additional information). Storing the splitting indexes, the optimized decompression algorithm can still work in linear time. On the other hand, by splitting always at the middle, we do not have to store additional information.

CHAPTER 6

RELATED WORK

Of course, there has been substantial work in lossy numerical compression. Beyond the earlier discussion and the baselines used in our evaluation, there has been substantial work in lossy compression for scientific data Cappello et al. [2019], Krasowska et al. [2021], Son et al. [2014]. Like our study, most of these techniques focus on spatio-temporal data, where the data are organized on some continuous axis (such as space, time, or both). Example of such techniques include SZ family of compression algorithms Di and Cappello [2016], Liang et al. [2018], Zhao et al. [2021b] and the ZFP algorithm Diffenderfer et al. [2019]. These algorithms follow a familiar structure to those described in our work, and offer L_∞ error guarantees. They generally pre-process/transform the data, quantize it, and then, apply a byte-level encoding algorithm. We omit an extensive comparison because the problem settings are quite different. Scientific data compression algorithms generally focus on maximizing compression for data at rest, and the applicability of these techniques in an online or mini-batch setting is more limited. Furthermore, to the best our our knowledge, multiresolution extensions to these algorithms have not been developed.

There has also been significant interest in multiresolution problems in adjacent areas. In approximate query processing, the DAQ project Potti and Patel [2015] uses a vertical layout of floating point bits to construct incrementally more accurate query results with error guarantees. This approach does no compression (i.e., it does not save on storage), but it does reduce the query latency for aggregate queries. The MLWeaving project has a similar approach to achieve machine learning training at different levels of precision Wang et al. [2019]. Similarly, multiresolution trees have been widely applied in approximate query processing where data are aggregated at hierarchy predicates Lazaridis and Mehrotra [2001], Liang et al. [2021], Buccafurri et al. [2011]. Similar "multiresolution" results for aggregate queries are seen in wavelet techniques for AQP Chakrabarti et al. [2001], online aggregation Hellerstein

et al. [1997], and sketching Cormode et al. [2011]. This work inspires our approach in HIRE, but is unfortunately only restricted to answering aggregate queries and not point-lookups. Wavelet techniques beyond the scope of traditional linear algebra decompositions have also been explored for multiresolution matrix compression Kondor et al. [2014].

Thus, we focus our study on a key set of baseline compression algorithms that: (1) can run efficiently in the online setting with rapid incoming data, (2) provide L_∞ error guarantees for point queries, and (3) do not require dataset-specific modeling for compression. It is worth mentioning recent data compression work that has been excluded from this study. The DeepSqueeze project Ilkhechi et al. [2020] uses an auto-encoder to learn a low dimensional set of features that can represent the original dataset. In our experiments, we found that the "encoder" portion of the auto-encoder was very large in size (often the same order of magnitude as the data), and since it is dataset-specific, it has to be included in the compression ratio measurement. The encoder is required to compress any new data that arrives. Similarly, we build a simplified version of Squish that works assuming column-independence Gao and Parameswaran [2016a].

CHAPTER 7

EXPERIMENTS

We conducted our experiments on a Macbook Air M1 with 8GB RAM and 256GB SSD running macOS 12.4. All implementations were done in Python 3.9. Our technique, in addition to each baseline, was applied to 7 different multivariate time series data sets from the UCI Dua and Graff [2017] repository. We use TRC as our downstream compressor.

7.1 Datasets

We selected 7 datasets that would represent possible real world applications of our technique, where compressed multivariate time series could require multiple resolutions in order to fulfill latency requirements. Our datasets are from four different projects within the UCI repository: Heterogeneity activity recognition data set Stisen et al. [2015], from which we used four different data sets, phones accelerometer (PA), phones gyroscope (PG), watch accelerometer (WA), watch gyroscope (WG); Sensors for home activity monitoring (SHAM) data set Huerta et al. [2016]; Individual household electric power consumption data set (IHEPC) ; Bitcoin heist ransomware address data set (BC) Akcora et al. [2019]. We removed all non-numerical columns, since numerical values are the focus of our present research. We also removed missing data when present. As mentioned earlier, we assume a mini-batch model for data arrival. Since HIRE assumes mini-batches that are sized as powers of 2, we simply cut the different datasets to a multiple of our block sizes. This modification does not change our experiments and is simply done for consistency. The resulting dimensions for each dataset are displayed in Table 7.1.

Data set(s)	Columns	Rows
PA, PG, WA	3	262,144
SHAM	8	524,288
IHEPC	7	524,288
BC	7	1,048,576

Table 7.1: Dataset descriptions

7.2 Baselines

We compare HIRE with 9 different baselines ranging from state of the art methods for time series compression to canonical methods for tabular data. Our baselines feature both lossless and lossy techniques. Below is a brief description of each baseline:

- **Identity Gzip (IdG)**: Lossless compression baseline; we apply Gzip to an array of numbers represented as floating point values.

- **Quantize (Q)**: We convert each floating point number to an integer according to a user-defined error threshold, thereby saving exponent and mantissa bits (see 2.2.1 for more details). The numbers are stored as integers with bitpacking.

- **Quantize Gzip (QGZ)**: This method consists of a quantization step and Gzip as the downstream compressor.

- **Quantize TRC (QTRC)**: This method consists of a quantization step and the Turbo Range Coder (TRC) as the downstream compressor. TRC uses a Burrows–Wheeler transform (BWT) Burrows and Wheeler [1994] to rearrange blocks of values into runs of the same symbol (i.e. integer), and then applies an arithmetic encoder Witten et al. [1987] during the entropy encoding step.

- **Sprintz (Spz)**: Firstly, it predicts the current sample based on the previous sample and encodes its difference (see 2.2.2 for more details). Secondly, it bitpacks the errors

and stores metadata to allow for unpacking. Thirdly, it uses run length encoding on blocks of all zero errors. Lastly, it Huffman encodes the headers and payload Blalock et al. [2018].

- **AdaptivePiecewiseConstant (APC)** Piecewise approaches decompose a time series into segments and use the segments to approximate the time series. Our version of the algorithm adaptively sizes segments to enforce an error bound Chakrabarti et al. [2002]. The downstream data are compressed with GZip.

- **Gorilla (Grl)**: This technique employs a scheme that consists of a bitwise XOR between pairs of consecutive values. It then produces a lossless encoding for each pair based on the number of leading or trailing zeros and the meaningful bits present Pelkonen et al. [2015a]. The downstream data are compressed with GZip.

- **LFZip**: This method employs the prediction-quantization-entropy lossy framework. It first uses a model to predict the next value in the sequence based on the previous ones. Subsequently, the difference between the prediction and the actual value is obtained and quantized with a user determined Linf error. Finally, a version of BTW is applied to the quantized data.Chandak et al. [2020]

- **Buff**: This technique was designed to compress float point numbers. It works at two levels: eliminating less significant bits by adjusting for a certain precision and compressing the float and integer parts independently.. [1] Liu et al. [2021]
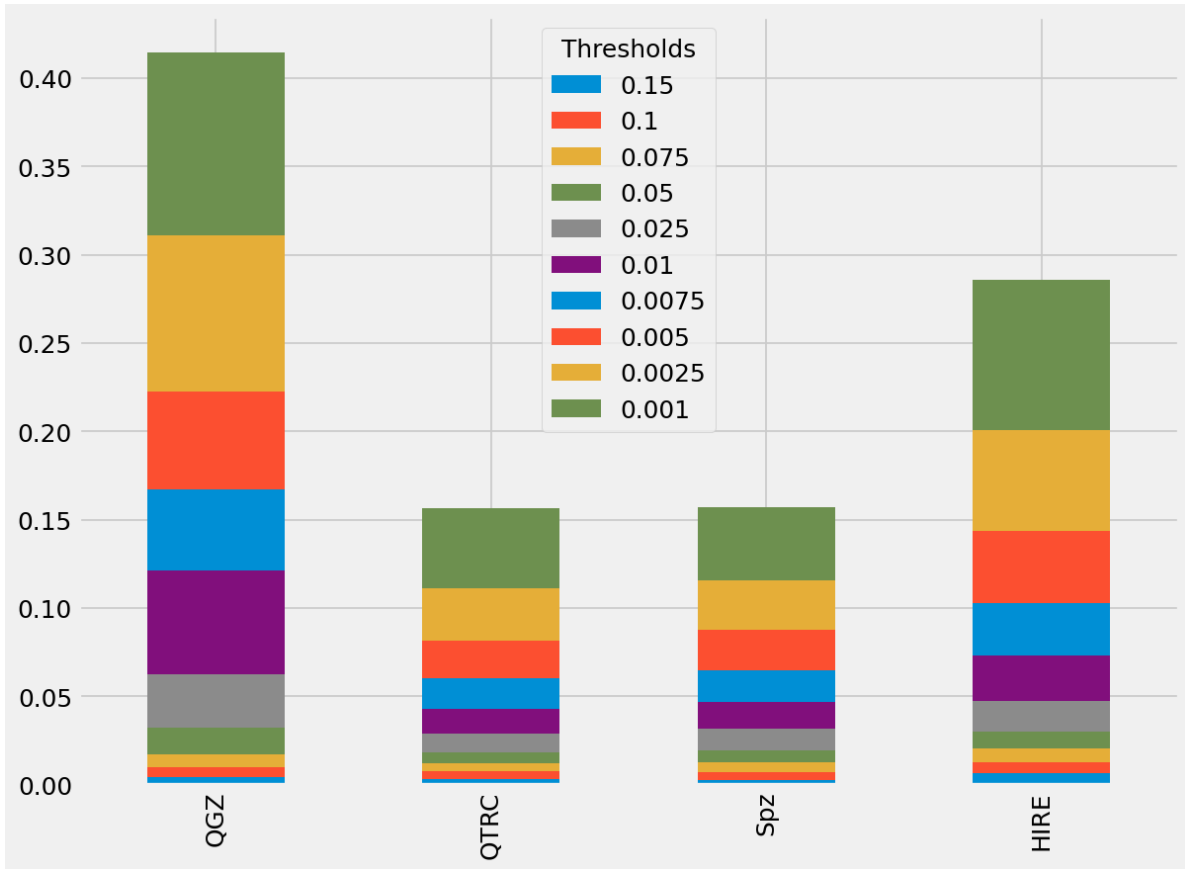
Figure 7.1: Phone Gyroscope

7.3   Performance Overview

Figure 7.1 exemplifies the main argument behind the benefits of using HIRE over competing methods. We display the compression ratio of compressed data at different error thresholds, corresponding to the four best approaches among those that we analyzed. The breakdown of compression ratio shows that at very low thresholds, we can always choose an error threshold using HIRE that will yield a more compressed representation than the competing methods. Hence, we need to encode the other methods at all possible error thresholds. On the other hand, because of how our method is constructed, we only need to store the lowest threshold and are able to retrieve all of the intermediate thresholds by traversing the tree structure until we reach the desired resolution.

Overall, we perform better than all of the competing methods when it comes to the *combined compression ratio* of all resolutions considered. For the experiments that we performed, 10 different resolutions were chosen that are inside the scope of real world usage: $\varepsilon^* \in \{0.15, 0.1, 0.075, 0.05, 0.025, 0.01, 0.0075, 0.005, 0.0025, 0.001\}$. HIRE requires half of the space to store all of the resolutions when compared to other baselines. Furthermore, the compression throughput of HIRE is significantly better than methods with low compression ratios, again due to the fact that we only need to run the encoder once to produce multiple resolutions.

7.4   Compression Ratio

We evaluated the compression performance of all of the methods on each of the seven datasets. The resulting compression ratios comprised of all of the resolutions are displayed in Table 7.2. Gorilla has the worst performance among all the baselines evaluated, as there

---

1. Buff does not directly allow for compression within a certain L_∞ error guarantee. It allows for a lossy compression by adjusting the float precision. Therefore we are only evaluating its performance as a lossless compression method.

are very few similar adjacent values on the real world data that we used. As we see Figure in 7.1, the best methods for a *single error threshold scenario* are Sprintz (Spz) and Quantized Turbo Range Coder (QTRC), but when summing up all of the different resolutions, their performance is on average two times worse than HIRE. One exceptional case is the SHAM data set, where the performance was really close to ours. On this specific dataset, the sample coefficient of variation (CV), the standard deviation over the mean, is extremely low due to the relative stability of the data. This impacts HIRE's smoothness throughout the various levels, which consequently results in a worse performance.

	PA	PG	WG	WA	SHAM	HIEPC	BC
IdGZ	5.554	4.023	4.498	4.643	7.698	2.172	2.306
Q	1.000	1.000	1.000	1.000	1.000	1.000	1.000
QGZ	0.526	0.415	0.351	0.431	0.172	0.436	0.175
QTRC	0.258	0.156	0.119	0.170	0.029	0.183	0.116
Spz	0.275	0.157	0.119	0.179	0.023	0.222	0.164
APC	1.844	1.004	0.986	1.093	0.207	0.551	0.296
Grl	8.184	5.920	6.389	8.067	7.657	2.956	8.067
LFZip	1.242	0.484	0.413	0.413	0.185	1.347	3.988
Buff	4.219	3.906	3.906	4.219	4.688	4.531	*
HIRE	0.116	0.085	0.070	0.085	0.021	0.091	0.061

Table 7.2: Compression ratio: sum of all resolutions

7.5 Compression and Decompression Throughputs

The evaluation of compression throughput includes data reading, the compression algorithm, and writing the compressed data to disk. HIRE outperforms all of the other low compression ratio baselines on average besides quantization alone (Q) as displayed in Table 7.3.[2] However, (Q) has more than a 10 times higher compression ratio than HIRE as we observe

---

2. We report the throughputs *without bitpacking* for the relevant baselines that use the routine prior to the execution of the downstream compressor. We found that in our experiments, bitpacking (implemented with https://pypi.org/project/bitarray/), seemed to introduce a latency bottleneck that skewed some of the results in favor of HIRE. As such, we decided to remove that step from the pipeline even though it inflates the results of those baselines.
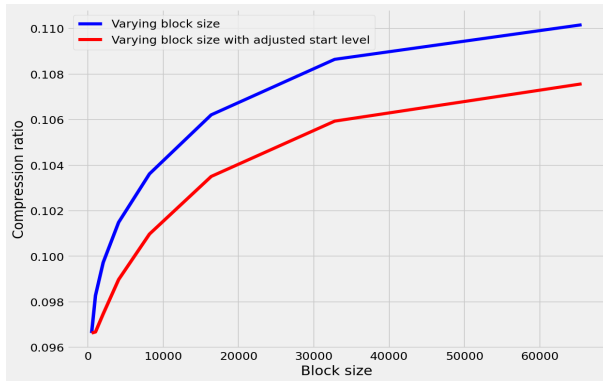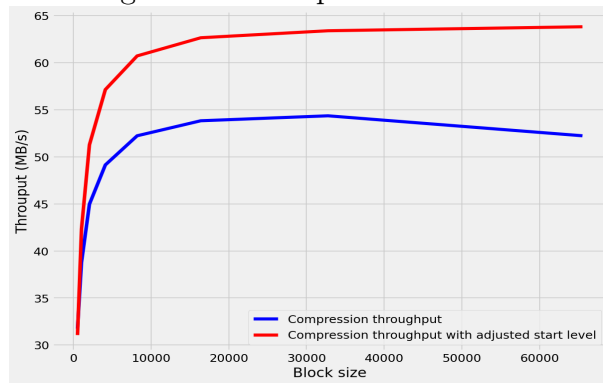
Figure 7.2: Compression ratio

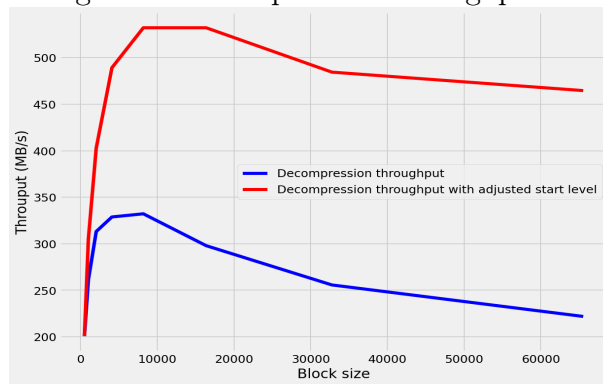Figure 7.3: Compression throughput

Figure 7.4: Decompression throughput

Figure 7.5: Breakdown of performance at different block sizes

in Table 7.2, making (Q) not practical in edge-computing applications. The main driver of HIRE's significant performance advantage is the fact that we require only a single call to our compression routine in order to produce multiple resolutions. This directly contrasts with the baseline methods, each of which compresses the data once per error threshold to produce 10 separable encodings. Gorilla once again performs considerably worse in this scenario due to its complex recursive encoding, where adjacent values are compared and bitwise operations executed [3].

	PA	PG	WG	WA	SHAM	HIEPC	BC
IdGZ	82	27	29	29	2	2	1.1
Q	1702	3855	2628	3648	454	399	419.2
QGZ	1949	102	111	68	36	3	1.2
QTRC	134	11	11	10	1	1	1.5
Spz	15	6	6	6	0.8	0.8	0.8
APC	27	50	51	46	223	52	138.9
Glz	3.1	3	3.1	3.1	0.3	0.5	3.1
LFZip	7.0	6.0	7.0	31.0	3.0	1.0	1.0
Buff	24.0	24.0	22.0	24.0	22.0	13.0	*
HIRE	1682.0	265.0	214.0	251.0	45.0	15.0	26.0

Table 7.3: Compression throughput in MB/s: sum of all resolutions

When it comes to decompression throughput, we report the average value for all of the different resolutions in Table 7.4. While HIRE performs better than all of the low compression ratio methods besides Quantize on the first four data sets, it is significantly outpaced by the quantization family and APC on the last three. Once again, this is likely due to the characteristics of the datasets. Applying HIRE to the first four datasets generates very smooth residuals, which allows for a shorter path of traversal along the tree until the desired resolution is reached. On data sets with extreme CVs, the traversal may not terminate until very low levels, which is overall detrimental to decompression throughput. One exceptional case is Gorilla, which has an extremely low decompression throughput when compared to the

---

3. We do note the caveat that Python is a high-level language, so the implementation that we used for Gorilla likely does not perform bitwise operations efficiently. This observation may explain *some* of the poor results for throughput, however the conclusion concerning compression ratio still stands fully in its own right.

other baselines. Once again, Gorilla's performance is impacted by both the datasets used in the analysis and its Python implementation.

	PA	PG	WG	WA	SHAM	HIEPC	BC
IdGZ	282.4	1013	971.4	906.0	96.52	136.0	146.6
Q	24811	24235	20942	21598	1118	1381	1566
QGZ	609	25827	15050	14482	963	1282	1320
QTRC	105.2	1451	1656	1514	480.5	197.0	285.4
Spz	64.6	163.0	160.47	160.9	21.5	21.5	21.4
APC	272.87	501.57	510.57	460.53	2234	523.7	1388
Glz	0.02	85.02	0.02	0.02	0.00	0.00	0.02
LFZip	74.10	71.77	73.15	68.34	5.68	7.26	7.18
Buff	272.01	272.01	346.58	297.31	336.28	153.11	*
HIRE	284.5	1927	1923	1925	247.3	72.7	217.1

Table 7.4: Decompression throughput in MB/s: average of all resolutions

7.6   Micro-benchmarks

We ran several different experiments in order to understand our method's most important hyperparameters: block size and start level. Additionally, we tested different pooling functions and their effects on the relevant metrics. All experiments in this subsection were performed on the Phones Accelerometer (PA) data set.

7.6.1   Block Size

A block represents a subset of the entire data set meant to be compressed. It allows for an online/streaming application of the method, since one can wait until a pre-determined block size is buffered before applying HIRE. However, it also affects the performance of the algorithm. We evaluate two different scenarios in which we vary the block size: with and without adjusting the starting level.

We notice in Figure 7.2 that as we increase the block size, the compression ratio decreases. This behavior can be attributed to two distinct reasons. Firstly, the larger the block size, the greater the range of values within each block. This phenomenon could possibly impact
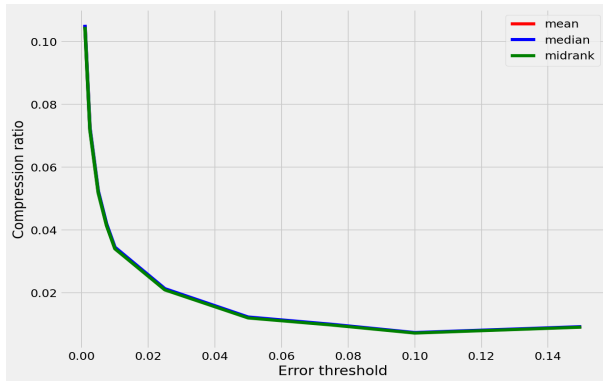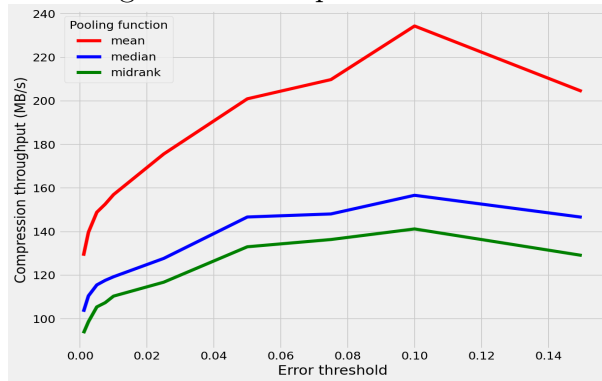
Figure 7.6: Compression ratio
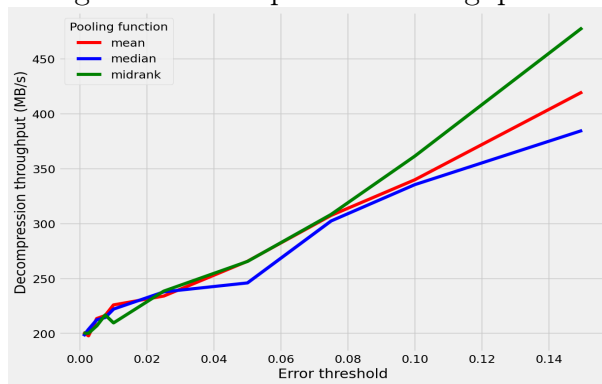

Figure 7.7: Compression throughput

Figure 7.8: Decompression throughput

Figure 7.9: Breakdown of performance for three different pooling functions

the variance of the residuals at each level, leading to less redundancy for the downstream compressor to exploit. Second, the presence of outliers can be mitigated at smaller block sizes, since their impact will be contained to a smaller subset of the data. Starting at a higher level in the tree structure also reduces the amount of data being stored, which leads to a positive impact on the compression ratio at each block size.

On the other hand, the increase in block size has a positive effect on compression throughput, especially when going from a very small block (512 time steps) to 8192 as we see in Figure 7.3. After that, the compression throughput plateaus on the adjusted level scenario and slightly decreases on the regular one. We attribute this fact to hardware limitations that impede further scalability due to the fixed number of CPU cores. Finally, when analyzing the decompression throughput shown on Figure 7.4, we can clearly see that once again it increases with the block size until it reaches a max value at roughly the same point as in compression—8192. This observation is also due to hardware constraints. The level-adjusted scenario performs significantly better than the regular one while being less affected by the hardware constraints.
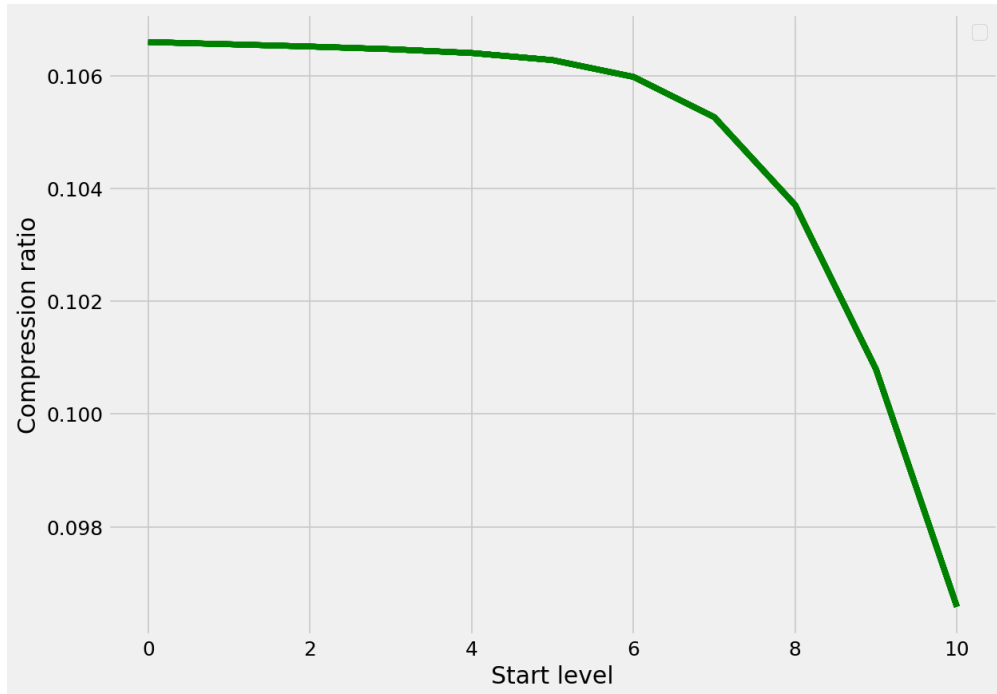
Figure 7.10: Compression ratio for different starting levels

7.6.2 Starting level

The starting level corresponds to the initial number of segments into which we break the original time series X and apply the pooling function; or equivalently, the initial number of nodes at level L of the binary tree $\mathcal{T}$. That is, we do not need to start the recursion at the first level, or even at the first few levels for that matter. Concretely, there is often little to no value in pooling large segments, particularly when the block size is purportedly large. We can therefore adjust the starting level in order to improve both compression ratio and throughput, albeit we lose the resolutions that correspond to the levels that are skipped. Such a trade-off is important in certain cases, as we may want to maintain a sufficient number of levels to allow for a specific number of resolutions.

Figure 7.10 shows that as we increase the starting level, the compression ratio decreases. Furthermore, the inverse relationship observed here is exponential in nature, since the number of values stored after pooling increases exponentially—specifically by a factor of 2—until

we reach the upper bound of T values at the leaf nodes of the hierarchy.

We observe an identical trend in compression throughput, which we will omit for the sake of brevity. Starting at a later level means that there are less spline interpolations to execute, and those which are executed have considerably less cost due—again to the exponential nature of the segments as we traverse down the tree.

7.6.3   Pooling function

The pooling function plays an important role in how the data are summarized and consequently the resulting residuals. We described its role in detail in Section 4.1.1. In Figure 7.6, we display the compression ratios at various error thresholds for the three pooling functions: mean, median, and midrank.

In Figures 7.7 and 7.8 respectively, we compare the compression and decompression throughputs achieved by the three pooling functions. The compression throughput of the mean is markedly higher at all of the thresholds we tested, which is consistent with the highly optimized implementation style that we chose (see 5.2). We note that the decompression throughput seems to stay somewhat similar for smaller error thresholds but then proceeds to diverge at larger error thresholds.

CHAPTER 8

CONCLUSION

We presented HIRE, a novel system for multiresolution compression that uses hierarchical residual encoding for time series data. We showed that strict and multiple encoding suffer from a high transfer cost and compression ratio in edge storage and retrieval applications. We proposed an efficient technique to handle multiresolution compression that alleviates the limitations of the previous methods. Our experiments validate that our system performs better than the baselines at multiresolution compression for edge computing applications.

HIRE can also be extended to the multivariate case. As mentioned, one simple way is to encode each time series independently. A more involved and efficient way is to extend our univariate hierarchical method. Instead of splitting the residual vector into two separate time series, we partition the values into four multivariate time series. Applying the residualizing step in a recursive manner, our compression and decompression algorithms could expand to handle the multiresolution problem in the multivariate case.

CHAPTER 9

SPATIAL-TEMPORAL EXTENSION

Spatial-temporal datasets are becoming prevalent in a wide range of applications, ranging from self-driving cars Caesar et al. [2020], security video Oh et al. [2011] to satellite images Lyons et al. [2022]. Furthermore, advances in hardware, more specifically GPUs, have allowed for the proliferation of scientific simulations that can generate petabytes of data Habib et al. [2013], creating the necessity to store significant amounts of data. The usefulness of compression techniques naturally arises as an effective way to manage such massive storage requirements.

We propose an extension to HIRE that leverages the spacial dependence of datasets in order to create an efficient multiresolution compression for high dimensional data. Contrary to state of the art methods for scientific data compression Zhao et al. [2021a], we allow for a single compression representation with multiple decompression resolutions, once again giving the user the flexibility necessary for different applications that could arise while requiring significantly less storage. Moreover, our method offers state of the art multiresolution compression speeds and competitive decompression throughput.

9.1   Research Plan

- **Oct/22 − Dec/22** : Framework implementation

- **Jan/23** : Experiments

- **Feb/23** : Paper write-up

- **Mar/01/2023** : VLDB23 Submission

REFERENCES

Cuneyt Gurcan Akcora, Yitao Li, Yulia R Gel, and Murat Kantarcioglu. Bitcoinheist: Topological data analysis for ransomware detection on the bitcoin blockchain. *arXiv preprint [Web Link]*, 2019.

Sebastian Angel, Mihir Nanavati, and Siddhartha Sen. Disaggregation and the application. In *12th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 20)*, 2020.

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

Davis W. Blalock, Samuel Madden, and John V. Guttag. Sprintz: Time series compression for the internet of things. *CoRR*, abs/1808.02515, 2018. URL `http://arxiv.org/abs/1808.02515`.

Francesco Buccafurri, Filippo Furfaro, Giuseppe M Mazzeo, and Domenico Saccà. A quadtree based multiresolution approach for two-dimensional summary data. *Information Systems*, 36(7):1082–1103, 2011.

M. Burrows and D. J. Wheeler. A block-sorting lossless data compression algorithm. Technical report, 1994.

Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020.

Franck Cappello, Sheng Di, Sihuan Li, Xin Liang, Ali Murat Gok, Dingwen Tao, Chun Hong Yoon, Xin-Chuan Wu, Yuri Alexeev, and Frederic T Chong. Use cases of lossy compression for floating-point data in scientific data sets. *The International Journal of High Performance Computing Applications*, 33(6):1201–1220, 2019.

Kaushik Chakrabarti, Minos Garofalakis, Rajeev Rastogi, and Kyuseok Shim. Approximate query processing using wavelets. *The VLDB Journal*, 10(2):199–223, 2001.

Kaushik Chakrabarti, Eamonn Keogh, Sharad Mehrotra, and Michael Pazzani. Locally adaptive dimensionality reduction for indexing large time series databases. *ACM Trans. Database Syst.*, 27(2):188–228, jun 2002. ISSN 0362-5915. doi:10.1145/568518.568520. URL `https://doi.org/10.1145/568518.568520`.

Shubham Chandak, Kedar Tatwawadi, Chengtao Wen, Lingyun Wang, Juan Aparicio Ojea, and Tsachy Weissman. Lfzip: Lossy compression of multivariate floating-point time series data via improved prediction. In *2020 Data Compression Conference (DCC)*, pages 342–351, 2020. doi:10.1109/DCC47342.2020.00042.

Chris Chatfield. The holt-winters forecasting procedure. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 27(3):264–279, 1978.

Graham Cormode, Minos Garofalakis, Peter J Haas, Chris Jermaine, et al. Synopses for massive data: Samples, histograms, wavelets, sketches. *Foundations and Trends® in Databases*, 4(1–3):1–294, 2011.

Sheng Di and Franck Cappello. Fast error-bounded lossy hpc data compression with sz. In *2016 ieee international parallel and distributed processing symposium (ipdps)*, pages 730–739. IEEE, 2016.

James Diffenderfer, Alyson L Fox, Jeffrey A Hittinger, Geoffrey Sanders, and Peter G Lindstrom. Error analysis of zfp compression for floating-point data. *SIAM Journal on Scientific Computing*, 41(3):A1867–A1898, 2019.

Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL `http://archive.ics.uci.edu/ml`.

Franz Färber, Norman May, Wolfgang Lehner, Philipp Große, Ingo Müller, Hannes Rauhe, and Jonathan Dees. The sap hana database–an architecture overview. *IEEE Data Eng. Bull.*, 35(1):28–33, 2012.

Johannes Fischer and Volker Heun. A new succinct representation of rmq-information and improvements in the enhanced suffix array. In *International Symposium on Combinatorics, Algorithms, Probabilistic and Experimental Methodologies*, pages 459–470. Springer, 2007.

Yihan Gao and Aditya Parameswaran. Squish: Near-optimal compression for archival of relational datasets. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1575–1584, 2016a.

Yihan Gao and Aditya G. Parameswaran. Squish: Near-optimal compression for archival of relational datasets. *CoRR*, abs/1602.04256, 2016b. URL `http://arxiv.org/abs/1602.04256`.

Dimitrios Giouroukis, Alexander Dadiani, Jonas Traub, Steffen Zeuch, and Volker Markl. A survey of adaptive sampling and filtering algorithms for the internet of things. In *Proceedings of the 14th ACM International Conference on Distributed and Event-based Systems*, pages 27–38, 2020.

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 *Proceedings of the 2015 ACM SIGMOD international conference on management of data*, pages 1917–1923, 2015.

Salman Habib, Vitali Morozov, Nicholas Frontiere, Hal Finkel, Adrian Pope, and Katrin Heitmann. Hacc: Extreme scaling and performance across diverse architectures. In *SC '13: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, pages 1–10, 2013. doi:10.1145/2503210.2504566.

Lajos Hanzo, Tong Hooi Liew, and Bee Leong Yeap. *Turbo coding, turbo equalisation and space-time coding.* John Wiley & Sons, 2002.

Joseph M Hellerstein, Peter J Haas, and Helen J Wang. Online aggregation. In *Proceedings of the 1997 ACM SIGMOD international conference on Management of data*, pages 171–182, 1997.

Ramón Huerta, Thiago Mosqueiro, Jordi Fonollosa, Nikolai Rulkov, and Irene Rodriguez-Lujan. Online decorrelation of humidity and temperature in chemical sensors for continuous monitoring. *Chemometrics and Intelligent Laboratory Systems*, 157, 07 2016. doi:10.1016/j.chemolab.2016.07.004.

BR Hutchinson and GD Raithby. A multigrid method based on the additive correction strategy. *Numerical Heat Transfer, Part A: Applications*, 9(5):511–537, 1986.

Amir Ilkhechi, Andrew Crotty, Alex Galakatos, Yicong Mao, Grace Fan, Xiran Shi, and Ugur Cetintemel. Deepsqueeze: deep semantic compression for tabular data. In *Proceedings of the 2020 ACM SIGMOD international conference on management of data*, pages 1733–1746, 2020.

Eamonn Keogh, Selina Chu, David Hart, and Michael Pazzani. An online algorithm for segmenting time series. In *Proceedings 2001 IEEE international conference on data mining*, pages 289–296. IEEE, 2001.

Eamonn Keogh, Selina Chu, David Hart, and Michael Pazzani. Segmenting time series: A survey and novel approach. In *Data mining in time series databases*, pages 1–21. World Scientific, 2004.

Eamonn J Keogh and Michael J Pazzani. Scaling up dynamic time warping for datamining applications. In *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 285–289, 2000.

Risi Kondor, Nedelina Teneva, and Vikas Garg. Multiresolution matrix factorization. In *International Conference on Machine Learning*, pages 1620–1628. PMLR, 2014.

David Krasowska, Julie Bessac, Robert Underwood, Jon C Calhoun, Sheng Di, and Franck Cappello. Exploring lossy compressibility through statistical correlations of scientific datasets. In *2021 7th International Workshop on Data Analysis and Reduction for Big Scientific Data (DRBSD-7)*, pages 47–53. IEEE, 2021.

Iosif Lazaridis and Sharad Mehrotra. Progressive approximate aggregate queries with a multi-resolution tree structure. *Acm sigmod record*, 30(2):401–412, 2001.

Xi Liang, Stavros Sintos, Zechao Shang, and Sanjay Krishnan. Combining aggregation and sampling (nearly) optimally for approximate query processing. In *Proceedings of the 2021 International Conference on Management of Data*, pages 1129–1141, 2021.

Xin Liang, Sheng Di, Dingwen Tao, Sihuan Li, Shaomeng Li, Hanqi Guo, Zizhong Chen, and Franck Cappello. Error-controlled lossy compression optimized for high compression ratios of scientific datasets. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 438–447. IEEE, 2018.

Peter Lindstrom. Fixed-rate compressed floating-point arrays. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2674–2683, 2014. doi:10.1109/TVCG.2014.2346458.

Chunwei Liu, Hao Jiang, John Paparrizos, and Aaron J. Elmore. Decomposed bounded floats for fast compression and queries. *Proc. VLDB Endow.*, 14(11):2586–2598, jul 2021. ISSN 2150-8097. doi:10.14778/3476249.3476305. URL `https://doi.org/10.14778/3476249.3476305`.

Mitchell Lyons, Kirk Larsen, and Matt Skone. CoralMapping/AllenCoralAtlas: DOI for paper at v1.3, June 2022. URL `https://doi.org/10.5281/zenodo.6622015`.

Alistair Moffat. Huffman coding. *ACM Computing Surveys (CSUR)*, 52(4):1–35, 2019.

Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, J. K. Aggarwal, Hyungtae Lee, Larry Davis, Eran Swears, Xiaoyang Wang, Qiang Ji, Kishore Reddy, Mubarak Shah, Carl Vondrick, Hamed Pirsiavash, Deva Ramanan, Jenny Yuen, Antonio Torralba, Bi Song, Anesco Fong, Amit Roy-Chowdhury, and Mita Desai. Avss 2011 demo session: A large-scale benchmark dataset for event recognition in surveillance video. In *2011 8th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 527–528, 2011. doi:10.1109/AVSS.2011.6027400.

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.

John Paparrizos, Chunwei Liu, Bruno Barbarioli, Johnny Hwang, Ikraduya Edian, Aaron J Elmore, Michael J Franklin, and Sanjay Krishnan. Vergedb: A database for iot analytics on edge devices. In *CIDR*, 2021.

Tuomas Pelkonen, Scott Franklin, Justin Teller, Paul Cavallaro, Qi Huang, Justin Meza, and Kaushik Veeraraghavan. Gorilla: A fast, scalable, in-memory time series database. *Proc. VLDB Endow.*, 8(12):1816–1827, aug 2015a. ISSN 2150-8097. doi:10.14778/2824032.2824078. URL `https://doi.org/10.14778/2824032.2824078`.

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, 2015b.

Navneet Potti and Jignesh M Patel. Daq: a new paradigm for approximate query processing. *Proceedings of the VLDB Endowment*, 8(9):898–909, 2015.

Julian Seward. bzip2 and libbzip2. *avaliable at http://www. bzip. org*, 1996.

John Shahid. Influxdb documentation, 2019.

Seung Woo Son, Zhengzhang Chen, William Hendrix, Ankit Agrawal, Wei-keng Liao, and Alok Choudhary. Data compression for the exascale computing era-survey. *Supercomputing frontiers and innovations*, 1(2):76–88, 2014.

Allan Stisen, Henrik Blunck, Sourav Bhattacharya, Thor Siiger Prentow, Mikkel Baun Kjærgaard, Anind Dey, Tobias Sonne, and Mads Møller Jensen. Smart devices are different: Assessing and mitigatingmobile sensing heterogeneities for activity recognition. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, SenSys '15, page 127–140, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450336314. doi:10.1145/2809695.2809718. URL `https://doi.org/10.1145/2809695.2809718`.

Gregory K Wallace. The jpeg still picture compression standard. *IEEE transactions on consumer electronics*, 38(1):xviii–xxxiv, 1992.

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.

Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the h. 264/avc video coding standard. *IEEE Transactions on circuits and systems for video technology*, 13(7):560–576, 2003.

Ian H. Witten, Radford M. Neal, and John G. Cleary. Arithmetic coding for data compression. *Commun. ACM*, 30(6):520–540, jun 1987. ISSN 0001-0782. doi:10.1145/214762.214771. URL `https://doi.org/10.1145/214762.214771`.

Steffen Zeuch, Eleni Tzirita Zacharatou, Shuhao Zhang, Xenofon Chatziliadis, Ankit Chaudhary, Bonaventura Del Monte, Dimitrios Giouroukis, Philipp M Grulich, Ariane Ziehn, and Volker Mark. Nebulastream: Complex analytics beyond the cloud. *Open Journal of Internet Of Things (OJIOT)*, 6(1):66–81, 2020.

Kai Zhao, Sheng Di, Maxim Dmitriev, Thierry-Laurent D. Tonellot, Zizhong Chen, and Franck Cappello. Optimizing error-bounded lossy compression for scientific data by dynamic spline interpolation. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, pages 1643–1654, 2021a. doi:10.1109/ICDE51399.2021.00145.

Kai Zhao, Sheng Di, Maxim Dmitriev, Thierry-Laurent D Tonellot, Zizhong Chen, and Franck Cappello. Optimizing error-bounded lossy compression for scientific data by dynamic spline interpolation. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, pages 1643–1654. IEEE, 2021b.