

Clemson University

TigerPrints

All Theses

Theses

5-2024

Analyzing an In-line Compression Management System for Improved Performance in a High-Performance Computing Environment

Steven Platt
platt9@clemson.edu

Follow this and additional works at: https://tigerprints.clemson.edu/all_theses



Part of the [Data Storage Systems Commons](#)

Recommended Citation

Platt, Steven, "Analyzing an In-line Compression Management System for Improved Performance in a High-Performance Computing Environment" (2024). *All Theses*. 4241.

https://tigerprints.clemson.edu/all_theses/4241

This Thesis is brought to you for free and open access by the Theses at TigerPrints. It has been accepted for inclusion in All Theses by an authorized administrator of TigerPrints. For more information, please contact kokeefe@clemson.edu.

ANALYZING AN IN-LINE COMPRESSION MANAGEMENT SYSTEM FOR
IMPROVED PERFORMANCE IN A HIGH-PERFORMANCE COMPUTING
ENVIRONMENT

A Thesis
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science
Computer Engineering

by
Steven Platt
May 2024

Accepted by:
Dr. Jon C. Calhoun, Committee Chair
Dr. Melissa Smith
Dr. Tao Wei

Abstract

High-performance computing (HPC) has enabled advancements in computation speed and resource cost by utilizing all available server resources and using parallelization for speedup. This computation scheme encourages simulation model development, massive data collection, and AI computation models, all of which store and compute on massive amounts of data. Data compression has enhanced the performance of storing and transferring this HPC application data to enable acceleration, but the benefits of data compression can also be transferred to the active allocated memory used by the application. In-line compression is a compression method that keeps the application memory compressed in allocated memory, decompressing content as the data is needed by the algorithm. The actively decompressed data size in allocated memory can be limited by grouping and compressing the data into blocks informed by the application's computational kernel's data access patterns and a selected compressor. Several factors are considered to tune an in-line compression scheme for a kernel. This research explores factors such as block size and compressor choice on the runtime and memory usage of the Matrix Multiplication kernel (MM). Matrix multiplication is a fundamental algorithm in most HPC kernels that computes a linear transformation on a set of vectors. MM kernels provide a baseline for evaluating in-line compression due to MM's row-based data access patterns and the usage of several matrices to compute the resulting matrix. Like traditional data compression, trade-offs between memory size and the runtime necessary for compression necessitate tuned parameters for each kernel. The results of this research demonstrate essential parameter trends, trade-offs, and the importance of locality in the kernel's data access patterns.

Acknowledgments

I extend my heartfelt gratitude to Dr. Jon C. Calhoun for his invaluable guidance and insightful feedback throughout the journey of this thesis. I greatly appreciate his willingness to teach so much of his research, exploration, analysis, and presentation expertise in such a short amount of time. His encouragement and expertise made this research experience possible, and I am very grateful!

I am deeply thankful to Dr. Melissa Smith and Dr. Tao Wei for their constructive criticism and support through my time in the ECE department. From deep hallway design discussions to casual overviews, the entire ECE department has been a support for this endeavor. Special thanks to ECE Department Chair Dr. Hai Xiao, Graduate Program Coordinator Dr. Harlan Russell, and Dr. Hassan Raza for their feedback and openness.

The peers, teammates, and research partners I have had these past 2 years have been phenomenal. Special thanks to Emily Lattanzio, Tripp Herlong, Mac Shaughnessy, and the many others who became impromptu sounding boards and sources of inspiration.

I owe a debt of gratitude to my parents and family, whose unwavering love and encouragement have been critical in pursuing this goal. They have been a driving force behind this accomplishment. Special thanks to Grandpa Waggoner for inspiring me to pursue a Master's degree.

This material is based upon work supported by the National Science Foundation under Grant No. SHF-1910197 and SHF-1943114.

Table of Contents

Title Page	i
Abstract	ii
Acknowledgments	iii
List of Figures	v
1 Introduction	1
1.1 Purpose of this Research	2
2 Background and Related Work	4
2.1 Compression	4
2.2 Matrix Multiplication	8
2.3 Related Work	9
3 In-line Compression Architecture	10
3.1 Methods Summary	10
3.2 Compression Manager Analysis	15
3.3 In-line Compression Metrics	15
4 Experimental Results	18
4.1 Testing Environment	18
4.2 Memory Footprint Analysis	19
4.3 Timing	25
5 Conclusions and Discussion	34

List of Figures

2.1	Matrix Multiplication Implementation	8
3.1	Standard Compression Workflow	11
3.2	In-line Compression Workflow	11
4.1	Total allocated memory usage for matrix-matrix multiplication on a matrix of order $n = 50$	20
4.2	In-line compression ratio for various compressors for matrix-matrix multiplication with order $n = 100$	22
4.3	Full In-line Compression Ratio by Dimensions and Error Bound	24
4.4	Compressor Call Counts per Matrix with respect to Block Counts	26
4.5	B -Transpose Kernel Access Pattern on Decompression Counts	29
4.6	Runtime with All Compressors and Dimensions	31
4.7	Runtime with SZ-ABS and Block Count	32

Chapter 1

Introduction

High-performance Computing (HPC) is the practice of sharing computing resources and tasks among several nodes to enable streamlined data processing and increased computing power. Data scientists use HPC systems to process and generate terabytes and petabytes of data through simulations and data collection [32].

Memory storage sizes and transfer speeds are bottlenecks in HPC operations [23, 3, 19, 7]. To mitigate this bottleneck, HPC systems use data compression schemes to exchange extra computations to reduce the size of data to be stored or transferred.

Compression is a necessary tool for large-data HPC systems in two ways. First, it reduces the size of the data needed to be stored in memory storage. This enables efficient usage of the memory space to allow more data storage. Second, reduced data storage means that less data needs to be transferred to allocated memory, enabling faster transfer times than transferring the complete original data.

Processing any data requires loading the data into a node's allocated memory. Any compressed data must be decompressed for calculations to have any meaning. The standard approach is to load all of the required application data in the kernel's allocated memory space. This approach requires using large nodes with massive memory spaces, enabling fast transfers and calculations across the entire dataset. This approach would enable fast calculations, but enabling this approach has some drawbacks. The more memory a kernel needs, the less likely the resources will be available to perform that function. Decompressing and storing an entire dataset requires nodes with memory suitable for such large data. The process would require nodes to be selected that have the memory

capacity necessary for that process. Because there is a limited number of large nodes available, those that are available have a high demand. This makes enabling nodes with large memory expensive time-wise due to waiting for the necessary resources. In a cloud-based system, this time-cost becomes a monetary cost to access a high-demand resource.

Costly nodes encourage searching out a solution that could process large data with smaller, less powerful nodes. One could search out the smallest possible node with the memory space to analyze data the standard way, but that begs the question if those resources are being used effectively. HPC systems conserve their memory usage by having the data compressed during memory and transfer. Could the benefits of compression be transferred to the algorithmic steps in the kernel to reduce the size of a kernel's active allocated memory?

Standard HPC kernels store application data in its compressed form in memory storage. Accessing this memory requires decompression and compression operations between memory storage and allocated memory. The data is decompressed completely in allocated memory. To transfer the benefit of compression to allocated memory, application data must be stored in a compressed form in allocated memory. Accessing and replacing data would require decompression and compression operations to be performed as part of the continuous sequence of operations for that kernel, or in-line with the kernel.

1.1 Purpose of this Research

This research analyzes the properties of in-line compression to explore its benefits and limitations across input parameters and configurations. Changes to the setup of in-line compression impact the number of compression/decompression operations requested, the runtime of the system, and the compression ratio or memory benefit from utilizing in-line compression.

In-line compression builds on preexisting kernels to improve their allocated memory access, and any improvements provided by in-line compression are relative to the performance of the kernel. Matrix-multiplication is used as the general starting kernel due to its abundance and importance in HPC operations as well as its ability to demonstrate features like locality and multiple datasets that are present in many other kernels that could benefit from in-line compression techniques.

Compressor choice impacts the overall compression possible when using in-line compression. Each compressor is built with a specific purpose in mind. Lossless compressors decompress

the data to exactly the original dataset, prioritizing accuracy over compression ratio and runtime. Lossy compressors sacrifice some accuracy within an acceptable specified range to obtain benefits in compression ratio and runtime. The interactions of compressor choice with other properties of in-line compression such as blocking and spacial locality make this property interesting.

An interesting aspect of in-line compression is the use of splitting the kernel dataset into several sections, or blocks. This is enabled by blocking the data at compression time and compressing each block separately. Each block can be compressed or decompressed individually, allowing for the decompressed data usage section to be dynamic, only having the data decompressed when that data is needed. That is, while all data is compressed in active memory, only the specified smaller blocks of data that are needed for each operation in the kernel need to be decompressed at any one time. By adjusting the block size, the compressor overhead and the decompressed data size are affected.

The properties obtained by this research informs decisions for improving the performance of in-line compression in systems involving many in-line kernels. This thesis makes the following contributions to in-line compression development:

- Discusses the importance and potential use-cases for both full and partial in-line compression.
- Uses matrix-matrix multiplication as an exemplary kernel to analyze design trade-off decisions went adapting code to use both types of in-line compression.
- Experimental results on dense matrices representative of those in ML and scientific computing show in-line compression significantly reduces the memory footprint for a dense matrix-matrix multiply by more than $300\times$
- Determine that data layout transformations that leverage knowledge of both the data access pattern and the compressed array structure are required to obtain good performance.

The following chapters lay a framework for understanding and building upon in-line compression. Chapter 2 introduces key background concepts for the construction of in-line compression, such as compression itself, kernels, and the purpose and use cases for in-line compression. Chapter 3 presents the in-line compression architecture and describes how in-line compression can be done. This chapter also highlights key parameters and presents an in-line compression performance model. Chapter 4 examines key parameters experimentally by exploring the results of in-line compression performance on a Matrix-Multiplication kernel.

Chapter 2

Background and Related Work

In-line compression is based on several mainstays in HPC development. Existing compressors, with their differing design philosophies and architectures, can be incorporated into an in-line compression application. In-line compression applications incorporate these compression techniques in the same continuous operational sequence as a kernel performing a specified algorithm. The kernel that is selected determines several options for how compression can be implemented in that kernel, and each compression option impacts the overall runtime and memory performance of the in-line compression algorithm. In-line compression algorithms vary based on the use case of the kernel, how much memory is available to the kernel, and the memory access patterns of the kernel. Understanding the use cases for in-line compression is important to determine if a kernel's active memory usage can be improved by this system.

2.1 Compression

Data compression is an effective tool to reduce memory storage and transfer requirements in HPC systems. [6, 17, 27] Data compression is accomplished by algorithms called compressors and decompressors. When data is put through a compressor, the compressor uses an algorithm to represent the information and produce ideally smaller amounts of data that represents the original data. A decompressor takes this compressed file and expands it back to the size of the original data.

Compression techniques enable data to be represented using fewer bits than the original data. There is an art to selecting a compression algorithm for a given set of data, and compressors

can be chosen based on the properties of the data. Compression algorithms utilize data properties such as redundant data, repeating patterns, frequency of characters, and proximity to neighboring data to encode the data into a smaller representation. Decompression algorithms take the compressed file and decode it back to a representation of the original data.

2.1.1 Compression Ratio

There are several important factors to consider when compressing data. The compression ratio or rate of the compression quantifies the change of data size between the original size and the compressed size. The bigger the rate, the better the compression, and the smaller the compressed size. Another important consideration of compression is the compression speed, as the memory reduction comes at the cost of additional processor time. A final consideration is the impact of the compression on the integrity of the data. Some systems require the compression process to be able to exactly restore the original data, while others can sacrifice some acceptable level of accuracy for improvements in runtime and rate.

Compression ratio is a metric for relating the data's original memory size to its compressed size according to equation 2.1, with the compression ratio CR being the ratio between the original data size o and the compressed data size c . The compression ratio is greater than 1 if the compressed size is smaller than the original data size and indicates positive compression. Negative compression is revealed when the compression ratio is less than one, indicating a larger compressed file.

$$CR = \frac{o}{c} \tag{2.1}$$

The compressor's runtime is directly related to the operations that are performed in the compression and decompression algorithms. The time complexity is an indicator of whether a specific compressor is good for a specific application. For instance, a slower but more accurate compressor is great for large scientific data where accuracy is important, but it is not as useful for streaming applications that require fast decompression operations.

Another feature that can balance memory and time usage is the error bound. This parameter enables the user to select a reasonable amount of error that the kernel data can incur without becoming useless or misleading. The decompressed data would not be an exact copy of the original data, but rather an approximation within the specified error bound. Tightening an error bound

requires more memory and runtime, while loosening an error bound allows enables the compressor to select a representative value from a range of values. This representative value can reduce runtime and memory usage.

2.1.2 In-line Compression

In-line compression provides the possibility to reduce the active memory size by storing kernel data in its compressed form in allocated memory. This provides several benefits for HPC kernel operations.

First, reduced active memory size enables the kernel to be run on smaller memory nodes. These nodes are more plentiful and lower cost than their state-of-the-art counterparts, enabling more processes to be run than on a larger node. Even when performing a serial kernel, the use of a smaller node enables the larger nodes to be used for larger calculations.

Second, running kernels on smaller nodes promises improvements through parallelization. While in-line compression takes more runtime to process the data due to the extra compression operations, the ability to run on smaller kernels means one can utilize many of those kernels. One in-line compression kernel may take longer, but if the kernel can be parallelized, the multitude of smaller nodes can be utilized to enable faster computation than even the standard kernel operation.

2.1.3 Lossless and Lossy Compression

Lossless compression algorithms are designed to guarantee that the data generated by decompression perfectly matches the original uncompressed data. There is no loss in the data accuracy; the error is zero. This type of data is useful for applications such as measuring sensitive data, recording rare events, and fields where precision and accuracy are critical. This accuracy comes at a cost of increased runtime. Lossless compression ratios range around 2, so memory sacrifices are made to ensure the accuracy of the data. Examples of lossless compression schemes include the PNG image format and Huffman lossless compression.

Error-bounded lossy compression sacrifices an acceptable level of accuracy to gain increases in compression ratio and runtime. The acceptable error is a parameter of the compressor that can ensure a specific error bound range. State-of-the-art lossy compression ratios can enable compression ratios in the range of 20. Applications such as images and audio receive a large storage reduction

by utilizing lossy compression.

Another application for lossy compression for HPC environments is the compression of scientific floating-point data. Because the data is floating point, it is by nature an approximation. By utilizing a lossy floating-point compression scheme, floating-point data can take advantage of reduced memory usage while still maintaining enough accuracy for its use case.

There are two major error-bounded lossy floating-point compressors used in this analysis. SZ is an adaptive prediction-based error-bounded lossy compression framework developed at Argonne National Laboratory [6, 28, 16]. This compression is enabled by dividing the data into non-overlapping blocks that then use a prediction method in each block to estimate the next floating point value. If the next value lies within a given error bound or accuracy to several predictions, the correct prediction is encoded. Otherwise, the exact value is recorded. A lossless Huffman encoding is used to compress the data one last time. SZ provides three main error-bounding modes. The absolute error-bounding (SZ-ABS) limits the error to within a specified absolute error bound. The point-wise-relative error-bounding mode (SZ-PWR) limits the error to within the user-provided proportion of the dataset’s range. The Peak-Signal-to-Noise-Ratio error-bounding mode (SZ-PSNR) guarantees that the PSNR calculation of the data retains the minimum user-provided value.

ZFP is a block-wise transformation-based compressor developed at Lawrence Livermore National Laboratory. [17] ZFP’s compression process includes dividing the input data into 4^d blocks, with d being the dimensionality of the data. ZFP encodes the transformation coefficients produced after orthogonal block transformations of those aligned and normalized blocks.

Two ZFP error-bounding modes are utilized in this thesis. First, the fixed-accuracy mode (ZFP-ACC) limits the absolute error to the user set error bound, similar to SZ-ABS. Second, the rate-based ZFP error-bounding mode (ZFP-RATE) supports random access into the compressed data at the cost of compression ratio.

Many compressors are written to meet a variety of compression needs, such as floating-point data that can sacrifice some accuracy for better compression. To explore various compressors for a use case, several compressors need to be tested. Most compressors have their own APIs to access their settings and run compression operations. Testing various compressors would normally mean a complete redesign of software to configure and run using different compressors. LibPressio [30] is a software library designed to provide a common interface for various compression algorithms within HPC environments. It allows developers to integrate compression and decompression functionalities

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \times \begin{bmatrix} b_{11} & b_{12} & b_{13} \\ b_{21} & b_{22} & b_{23} \\ b_{31} & b_{32} & b_{33} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} & c_{13} \\ c_{21} & c_{22} & c_{23} \\ c_{31} & c_{32} & c_{33} \end{bmatrix}$$

Figure 2.1: Matrix Multiplication Implementation

into their applications more easily by abstracting away the complexities of individual compression libraries. This abstraction enables a unified interface for accessing different compression algorithms, allowing users to experiment with and compare various compression methods to find the most suitable ones for their specific use cases.

2.2 Matrix Multiplication

Matrix multiplication is a common linear transformation kernel that is common in many HPC use cases, such as scientific computing, graphics processing, machine learning model training, and cryptography. For this reason, matrix multiplication is the representative kernel used in testing in-line compression.

Given matrices A and B where the number of columns in A matches the number of rows in B , matrix C is calculated by Equation 2.2, where n is the size of the matrices. Figure 2.1 provides a visual guide for what the data access pattern for matrix multiplication is. For each of the output elements in C , the corresponding row in A and column in B are accessed and linearly transformed to generate that output element.

$$C_{ij} = \sum_{k=1}^n A_{ik} \times B_{kj} \tag{2.2}$$

The time complexity of the standard matrix multiplication kernel is $\mathcal{O}(n^3)$. While optimizations to matrix multiplication have been made to improve its runtime complexity, the standard matrix multiplication calculation is used as an example for in-line compression.

2.3 Related Work

Compression schemes that reduce memory use are the basis for memory savings in HPC. Much research is dedicated to use cases and improvements for data compression. [7] details use cases for floating-point lossy data, the type of data used in this research. Compressed plots and images keep the data size lower while still retaining acceptable visual fidelity for the image’s use case. These compressed files are smaller than uncompressed data, reducing the memory footprint of the data. Smaller files implies a reduced intensity of the data stream when transferring data from memory storage [26, 1, 8, 31, 15].

These software-level benefits can be used to reduce the I/O time of the system, accelerate checkpoint/restart systems, network speedup [11, 10, 12], and accelerate the overall kernel execution. Reduced checkpoint times [13, 4, 29] allows HPC simulation checkpoints to have a smaller impact on the overall application operations while still providing the memory backup.

Caching is enables speedup by keeping frequent or imminent resources close to where computation is performed. This can be performed through hardware and software caches that save a copy of frequent data to reduce access times for later operations. Software caches [22, 25, 9] are useful to improve I/O performance and to cache input for parallel tasks. Hardware caches [2, 5, 20, 14, 21] have been paired with compression to expand the size of main memory and hardware caches. In-line compression has been explored utilizing software caches, specifically how cache configuration impacts the performance of full in-line compression [24].

Chapter 3

In-line Compression Architecture

3.1 Methods Summary

In-line compression is a transformative data management system that uses compression and aggregation techniques to reduce the amount of data stored and transferred in allocated memory in high-performance computing environments. This section explores the differences in purpose and implementation between standard memory storage compression and in-line compression.

3.1.1 In-line Compression Overview

The purpose of in-line compression differs from compression for memory storage. Compression is commonly used to reduce the size of kernel data that is stored in memory storage such as a hard drive or cloud memory. Its purpose is to reduce both the storage size and the size of data needed to be transferred to a node for computation. The purpose of in-line compression is to reduce the amount of allocated memory in a node that is needed to access all of the data necessary for a kernel's computation.

Figure 3.1 illustrates the standard method of decompressing data to use with a kernel. The compression and decompression operations occur between the compressed data housed in memory storage and the decompressed data storage in allocated memory. The entire dataset is decompressed in main memory.

Figure 3.2 demonstrates the in-line compression implementation. The compressed data size




Figure 3.1: Standard Compression Workflow




Figure 3.2: In-line Compression Workflow

is copied into allocated memory. When data is needed, only what is needed for a specific operation is decompressed. This means that the compression and decompression operations retrieve and store data from the allocated memory alone. Note that if the entire compressed data gets decompressed, the allocated memory usage is larger than its standard counterpart. The compressed data needs to be arranged such that individual sections (or blocks) of that data can be decompressed as needed. Block selection is a feature of interest when exploring the performance of in-line compression.

Enabling in-line compression in a kernel involves three main steps: configuring input data in the format for that in-line compression configuration, storing compressed data in allocated memory, and moving compressor calls to data access time. First, the kernel data is arranged into blocks of data that are individually compressed. Any manager for in-line compression needs to keep track of which block each element of data is in so that the correct block can be randomly accessed when the data is needed at runtime. Second, the in-line configured data is copied over to allocated memory. This enables the compressor to quickly access the data blocks that are needed for decompression without relying on slower memory storage bus access. Third, instead of decompressing all blocks before the kernel's operation to keep the data close to the kernel, the decompression calls need to

be moved into the kernel's data access scheme. That is, when a new row or data from a specific block is needed, any old blocks are compressed back into allocated memory, and the correct blocks are decompressed to take their places.

Some potential options for in-line data compression's improvement for some kernels include software memory caching, kernel upgrade, and parallelization. If a block is used frequently, that block could be cached to eliminate the additional compression and decompression operations required to access it regularly. Upgrading the kernel to one with better time complexity improves performance, but that performance is gained through the kernel operations, not through any benefit in-line compression adds to the kernel. Additional threads and parallelizing in-line compression provides an opportunity to leverage the purpose of HPC systems and utilize smaller nodes in parallel.

3.1.2 Blocks

3.1.2.1 Block Types

Breaking a dataset into blocks of data enables each block to be compressed and decompressed individually, enabling random access to compressed data. How the data is blocked informs runtime and memory performance.

There are two main types of in-line compression blocking schemes: full and partial. The full in-line compression scheme involves compressing each sub-dataset individually. For the matrix multiplication kernel, each matrix is compressed individually. When a matrix is needed, that entire matrix is decompressed. This enables fast access times as the entire matrix is present in allocated memory, but the allocated memory must store the entire decompressed matrix.

The partial in-line compression scheme enables the user to keep less active allocated memory decompressed by limiting the size of each block of compressed data. This results in many smaller compressed data blocks. The user specifies a block size, the number of elements in each compression block. For all of the data in all of the sub-datasets, data is compressed in groups of that block size. Accessing smaller blocks involves decompressing the blocks needed, manipulating the data in those blocks, and re-compressing that data block. These additional compression calls increase the overhead of the compressor on runtime but limits the allocated memory usage. This enables less active memory usage as only a necessary fraction of the data is decompressed during each kernel operation. How these blocks are selected can play a role in memory and runtime performance of

in-line compression.

3.1.2.2 Block Tuning Factors

Block size is an integral component of how in-line compression works, and the block size parameter needs to be tuned for its use case. The extreme block sizes provide qualitative insight into how the block sizes affect memory and runtime performance. On one extreme, full in-line compression, where every matrix is its own block, provides access to an entire dataset close to memory. This means that the decompressed matrix is taking a relatively large space in allocated memory, but only one set of compression and decompression operations are needed to access this data. Full in-line compression causes a small runtime increase, but it enables an entire matrix to be decompressed with easy access to the kernel at the cost of large allocated memory usage. On the opposite extreme, a partial in-line compression scheme with each element being its own block (that is, a block size of 1) has the opposite performance. The allocated memory usage is very minimal, as only the elements needed for each operation are compressed at once, but this memory savings comes at a large runtime cost to enable the many compression and decompression calls to access those small blocks. Tuning these block sizes are essential to providing the memory improvements necessary for a kernel on user-specified hardware while keeping the runtime cost at a tolerable limit.

Block count also subtly affects the memory usage in the compressed data. Each compressed block has some level of overhead metadata that enables the decompression of that block. Having more blocks can decrease the impact of the decompressed data on the allocated memory size, but more blocks can increase the compressed data's impact on allocated memory through 2 main methods. First, the increase in blocks increases the amount of memory attributed to metadata. As block sizes decrease toward the minimum of 1 element per block, the metadata could even be larger than the encoded data depending on the chosen compressor. Second, less data in blocks means less locality for the compressor to take advantage of when compressing the data. Compressors like SZ and ZFP that rely on interpolating from neighboring data have better compression ratios when more data is compressed as compressions can be informed by more data [6, 18, 17]. These increases in compressed memory size as the number of blocks increases enables the memory improvements to eventually become saturates, limiting more allocated memory size improvements enabled by block size.

The context of data is important in selecting blocks for that dataset. Data rarely exists

in a vacuum, and patterns in types of data can be leveraged for better compression. These better compressions then reduce the compressed data size in memory, aiding the goal of reducing allocated memory usage. Patterns for some compressors include exploring the dimensionality of the data. For instance, a picture is 2-dimensions, and patterns can be inferred by exploring pixels around a point of interest.

Block dimensionality can be used to leverage the dimensionality used by compressors. 1-dimensional blocks group the data as if it is just a list of data. This would be the equivalent of only looking at a row of a picture's pixels to draw information from. 2 and 3 dimensional blocks could keep data blocked together if they are frequently decompressed together for compressors that use multiple dimensions. If a compressor makes decisions using multiple dimensions, then having blocks that consider these multiple dimensions may be useful to reduce the compressed data size in allocated memory.

3.1.2.3 Kernel Block Access

Decompressing data when the kernel needs it is the main means for in-line compression to reduce the allocated memory usage. HPC kernels access data differently depending on what operations they are performing. A kernel could access data row-by-row, sequentially, or even random access. When selecting the block sizes and locations, it is important to take the kernel's access patterns into consideration. Blocking choices that decompress all and only the required data for each operation and accomplish this with minimal decompression operations balance reducing the decompressed allocated memory usage with managing the runtime overhead for compression.

The matrix multiplication access pattern provides a useful surrogate to explore this concept of kernel access patterns affecting the data blocking. To calculate a single element in the output matrix, all of the elements in the corresponding row in the first matrix and the corresponding column in the second matrix must be accessed.

Block choice matters when tuning in-line compression for a kernel, and the matrix multiplication kernel is no exception. Because each kernel operation requires access to a 1-dimensional slice of multiple matrices, 1-dimensional blocking enables streamlined access to the specific required data. An entire row or column needs to be decompressed at once to enable each operation. Thus, selecting a block size that is related to the shared row and column length limits extra decompression operations.

3.2 Compression Manager Analysis

Managing the compression and decompression of memory blocks inside of the allocated data is essential to the implementation of in-line compression. In-line compression for this thesis was managed through an in-line compression manager (ICM). The ICM enables the selection of compressor, compressor options such as error bound, block size, and other parameters for in-line compression.

The ICM provides a good starting point for understanding in-line compression. It enables the blocking and compressed data management necessary to enable full and partial in-line compression. It provides options to explore matrix dimensionality, block sizing, compressor choice, error bound, and optimization choices. The main disadvantages of the ICM's implementation is the lack of certain options. Currently only SZ and ZFP are supported compressors, although the use of LibPressio makes adding more compressors trivial. Data blocking in the ICM is only 1-dimensional, which limits the context a compressor has to enable higher compression.

3.3 In-line Compression Metrics

The general performance of in-line compression algorithms can be modelled by understanding how in-line compression is set up. Input parameters such as kernel, compressor, specified compression ratio, original data size, and block size play a role in affecting the allocated memory usage and the runtime of the kernel.

3.3.1 Compression Ratio

Compression ratio (CR) is typically defined with respect to the uncompressed data size (U) and the compressed data size (C) as reflected in Equation 3.1. A simple ratio easily demonstrates whether positive compression occurred (if $CR < 1$).

$$CR = \frac{U}{C} \tag{3.1}$$

While this compression ratio information is useful, it does not demonstrate the size of the active allocated memory data. Since reducing the allocated memory usage is the main goal of in-line compression, having a memory measurement that takes the allocated data size into account is essen-

tial to accurately assess the memory implications of in-line compression. Equation 3.2 describes the compression ratio equation used for the following results. Because the allocated memory size is the compressed data size (C) and the decompressed block sizes (DB), the modified in-line compression ratio (iCR) is a compression ratio that takes the entire allocated memory into account, rather than just the compressed size of the data.

$$iCR = \frac{U}{C + DB} \quad (3.2)$$

This modified compression ratio reveals how block choice can affect the overall allocated memory usage. Take a matrix multiplication kernel that process 2 matrices (A, B) and generate a third matrix (C) where $AB=C$. First, one can note that reducing memory involves a balance between the compressed data size and the size of the decompressed blocks. Reducing allocated memory usage involves reducing or mitigating the growth of these two compression outputs.

Second, full in-line compression reveals a possibility for negative compression. If the kernel’s in-line compression is implemented with full in-line compression, note that all 3 matrices are always decompressed to access the data if the kernel is using full in-line compression. This means that the uncompressed data is the same as the decompressed block data; all of it is decompressed. No matter the size of the compressed data, the allocated memory size is larger in this configuration than if in-line compression had not been used at all. This case illustrates when full in-line compression is useful. If a different kernel was used, one that needed access to only a subset of the matrices, then only a subset of the data needs to be decompressed. The in-line compression ratio aids in describing allocated memory usage regardless of the matrix dimensions.

3.3.2 Compressor Counts

Each time that a block is compressed or decompressed, the specified compressor is called. While the compressor choice and size of the data do play a role in runtime, the number of compressor calls indicates whether there are inefficiencies in the blocking for the specified kernel. Because the time to compress and the time to decompress are different, it is important to distinguish which type of compressor operation is more prevalent for which matrix.

Two methods of approaching this runtime metric is to consider its response to block size and kernel access patterns. First, the more blocks there are, the more compression and decompression

operations needed to access the block. These compressor calls increase runtime through the overhead necessary to access the data. Second, efficient blocking dependent on the kernel’s access patterns can reduce the number of compressor calls. Depending on how the kernel accesses data and how the blocks are configured, compressor calls have extreme cases. If several blocks are accessed to decompress all of the elements needed for an operation, several compressor calls need to be made. A block choice that aligns with the kernel’s access patterns can reduce the number of compressor calls.

3.3.3 Runtime

In-line compression achieves its allocated memory usage reduction at the cost of runtime. Understanding what factors go into runtime is important to understand how runtime growth can be controlled as improvements are made in memory. Equation 3.3 identifies what factors influence the runtime of an in-line compression kernel.

$$Runtime = OriginalKernelTime + \sum(AllCompressionAndDecompressionTimes) \quad (3.3)$$

The runtime for an in-line compressed kernel is the sum of three main factors: the original kernel’s runtime, the time for all compression operations, and the time for all decompression operations. The number of compressor calls directly impacts the time for compressions and decompressions. The kernel choice itself is a constant if in-line compression is to be selected for that specific algorithm. The time complexities of the compressor accesses and the original kernel are important to predicting the runtime of the in-line compression kernel itself.

The compressor calls, with its time complexity, occur inside the kernel data access calls based on the original kernel, with its time complexity. At low matrix dimension sizes and low block counts, the compressor calls have a larger impact on the relative runtime. As the dimension size of the matrices increase, the time for compressor accesses may be amortized by the time taken by the original kernel. For example, if compression operations have a time complexity of $\mathcal{O}(n^3 = 2)$ while the kernel it is operating inside has a time complexity of $\mathcal{O}(n^3)$, the runtime at larger dimensions approaches the runtime of the original kernel itself. Exploring how block count and dimensions affect runtime is important to mitigate the runtime increase to enable allocated memory reduction.

Chapter 4

Experimental Results

4.1 Testing Environment

4.1.1 Hardware

Tests were run using the Clemson’s University’s Palmetto Cluster Phase 19b. Each compute node in this phase contains 2 Intel Xeon 6230R CPUs along with 372 GB of DRAM. Each test is run using a single core. This is done to profile and understand the performance of a single process. Future work will expand this work to investigate in-line compression for shared memory and distributed memory applications.

4.1.2 Software

The software utilized for these tests were built on GCC (GNU Compiler Collection) version 12.1.0. The compressors SZ version 2.1.8.1 and ZFP version 1.0.0 were used as floating-point error-bounded lossy compression benchmarks, and compressor management was maintained through LibPressio version 0.97.3. Python 3.11.6 served as the primary scripting language for orchestration and automation of computational tasks, with additional support from Seaborn version 0.13.2 and Matplotlib version 3.8.2 for data visualization and analysis.

4.1.3 Testing Format

The matrices used for matrix multiplication were square matrices of a user-specified dimension (n in plots). These matrices were populated with linear transformations of π .

The lossy floating-point compressors used in this research are error bounded. Unless otherwise noted, the error bounds utilized are $\epsilon = 0.01$ for SZ_ABS, SZ_PWR (point-wise relative), ZFP_RATE, and ZFP_ACC (accuracy). For SZ_PSNR, $\epsilon = 40db$.

4.2 Memory Footprint Analysis

4.2.1 Allocated Size

In-line compression reduces the amount of allocated memory that is being used by the application by storing data outside the working set in compressed form. Only when data is needed by the application is it decompressed just before it is operated on. One key parameter to control how much data is stored decompressed is the size of each block that the data is decomposed into prior to compression. We first experiment and determine the impact of block size on the allocated memory size, and it also illustrates how memory is organized inside of the allocated memory.

This experiment demonstrates how in-line compression is an effective means of reducing the allocated memory size. To understand the resulting allocated memory size, one must understand how in-line compression organizes the data in the allocated memory. Allocated memory contains 2 main sections: the compressed application data, and the decompressed data needed for a single operation — i.e., kernel. This experiment highlights how the sizes of those sections trade-off to result in overall data reduction.

Figure 4.1 illustrates the used memory size and the contents of that memory at various block counts. This experiment performs matrix multiplication on matrices of order $n = 50$. We use the SZ compressor with an absolute error bound of $\epsilon = 0.01$.

The height of each bar is the amount of allocated memory needed for each block size. Each bar is subdivided into the memory that is needed for this application. Matrices A, B , and C all need to be compressed, and some level of memory is needed to store their decompressed blocks. The lighter shades represent the compressed section of memory, while the darker shades represent the largest decompressed blocks for that block size.




Figure 4.1: Total allocated memory usage for matrix-matrix multiplication on a matrix of order $n = 50$.

The overall trend of this plot is that the allocated memory size decreases as the number of blocks increases. This test demonstrates that the allocated memory size is a trade-off between the compressed size of the entire data and the decompressed data sizes. More blocks means that less decompressed data needs to be stored for each block. However, more blocks requires more compressor metadata overhead and a potential for the smaller blocks to not be able to be compressed as well. Future in-line compression algorithms should be designed to limit or reuse meta-data.

The key takeaway of this experiment is that partial in-line compression is made possible due to the balance of the decompressed block sizes and the compressed data size. This trade off in block count needs to be tuned to enable the memory usage reduction that enables in-line compression's use case of being able to run kernels with large data on nodes with smaller allocated memory space. A second takeaway is that this memory balance has another input that was held constant in this experiment: the data size.

4.2.2 In-line Compression Ratio

Considering the full input data size while attempting to minimize the allocated data size requires a new metric to keep track of the allocated memory’s size relative to the original data size. Normally for most compressors, compression ratio is used as the metric, as it compares the uncompressed size to the compressed data size. In this case, a modified in-line compression ratio is used to relate the original data size compared to the overall allocated memory usage (see Section 3.3.1). The use of this modified compression ratio is necessary to contribute for the uncompressed blocks of data that are also present in allocated memory. Essential inputs toward the in-line compression ratio include the original data size and the number of blocks.

4.2.2.1 Block Count

Figure 4.2 compares the in-line compression ratio achieved by varying the number of blocks that we decompose our matrix of order $n = 100$ into before compressing with various versions of SZ and ZFP. The error bounds is $\epsilon = 0.01$ for SZ_ABS, SZ_PWR (point-wise relative), ZFP_RATE, and ZFP_ACC (accuracy). For SZ_PSNR, $\epsilon = 40db$.

The in-line compression ratio increases with block count, but it has an upper limit. The allocated memory size decreases by dividing each matrix into more blocks, thus having less data decompressed in main memory. The choice of underlying compressor also impacts the compression ratio, although only in the compressed data section of the allocated memory. We see ZFP performance the best due to its lower meta-data overhead than SZ. This becomes magnified for large block counts where the number of elements per block is small.

A compression ratio of $300\times$ is remarkable, but where is this benefit coming from? The main contributor to this is the decrease in the size of the allocated memory size due to less decompressed data. With more blocks per data array, less memory is accessed for each calculation — i.e., the decompressed working set is small. However, as the block size shrinks, more compression and decompression operations are needed for the computation. With the additional additional compression and decompression operations the overhead in runtime increase; thus longer application execution time (see Section 3.3.2).

The compression ratio of ZFP is particularly notable. For high block counts, its compression ratio outshines SZ. This is possible by having less data in each block for high block counts, but also

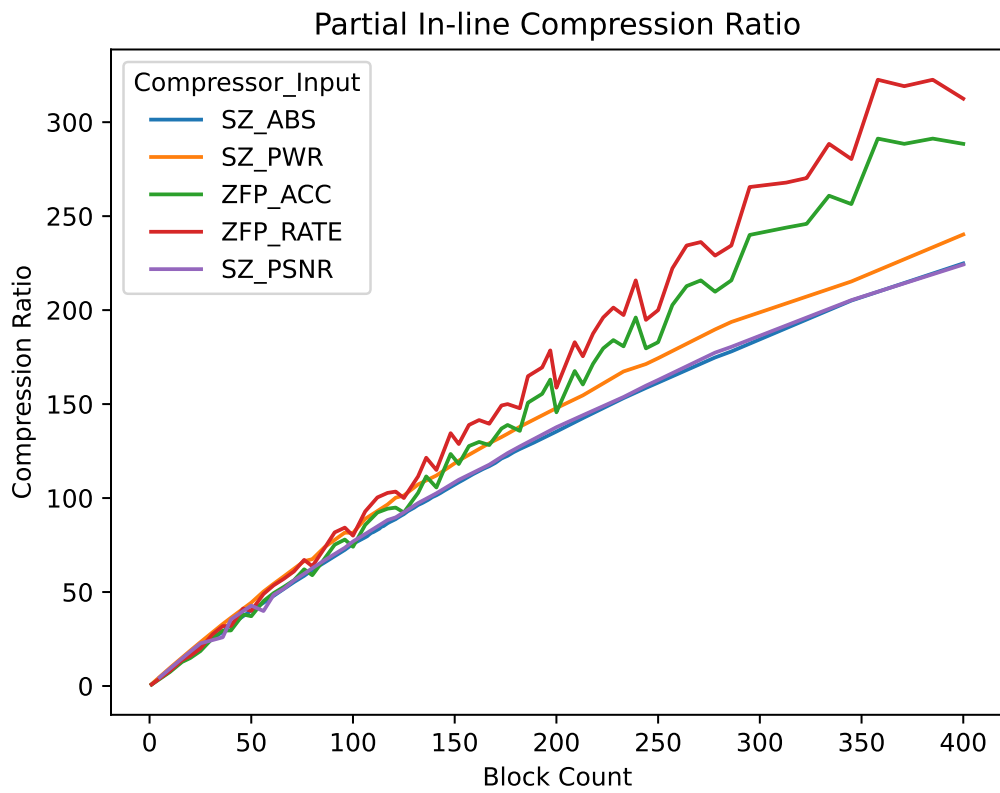


Figure 4.2: In-line compression ratio for various compressors for matrix-matrix multiplication with order $n = 100$.

due to the metadata of both compressors. ZFP has less metadata than SZ [18]. For a few larger blocks, this metadata difference is imperceptible. For several smaller blocks however, the metadata takes on a larger proportion of the memory usage. Thus, ZFP demonstrates a higher compression ratio than SZ for large block counts.

4.2.2.2 Data Size

This experiment compares the matrix dimensions to the compression ratio for a full in-line compressed matrix multiplication kernel. Full in-line compression was chosen to analyze a simple example with 1 block per matrix, including how data is impacted by the error bound.

An experiment on how data size impacts the in-line compression ratio is essential due to the difference from traditional compression ratios. The entire allocated memory size is compared to the uncompressed data size. Thus, it is important to consider how much of the data set is decompressed, even if using full in-line compression where an entire matrix is compressed together.

Similar to the previous in-line compression ratio experiment, this experiment compares the in-line compression ratios over square matrices orders ranging from $n = 5$ to $n = 1000$. The SZ compressor is used with the indicated error bounds to illustrate their affect on in-line compression, whether full or partial. Results for other compressors are similar and omitted for clarity.

Figure 4.3 illustrates how data size impacts compression ratio. Because compression is possible through locality and repetition, it is expected that smaller data sizes have a smaller compression ratio. The effects of the error bound on each block's compression are passed through to in-line compression, with the tightest error bound having the smallest compression ratio, and the loosest error bound can sacrifice accuracy to gain space.

One note is that this experiment is using full in-line compression, or a single block per matrix. This means that for the in-line compression ratio equation 3.2, if all of the matrices in the dataset are involved in each operation in the kernel, the uncompressed data size and the decompressed block sizes are the same. This situation is present with matrix-matrix multiplication and yields negative compression no matter what the data size.

This situation highlights that full compression is great for some use-cases but not for others. If a subset of the matrices are used per kernel operation, then only a subset of the matrices need to be decompressed for each operation. The uncompressed dataset is greater than the necessary decompressed matrices for each operation, so depending on the size of the compressed dataset,

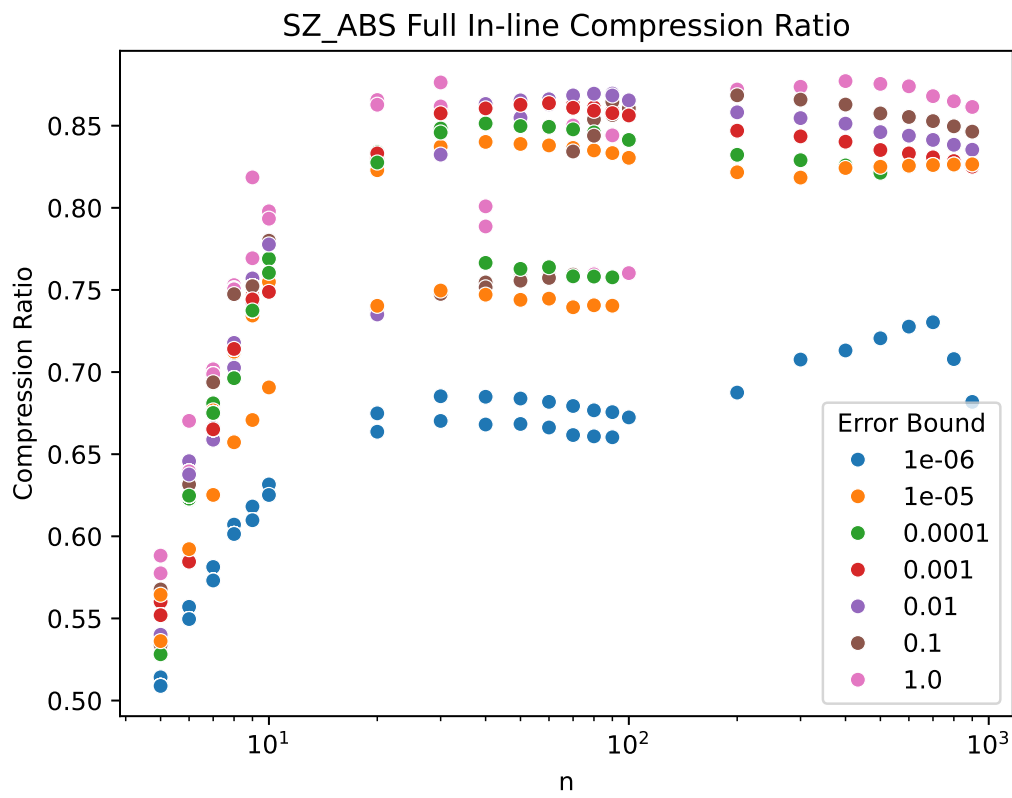


Figure 4.3: Full In-line Compression Ratio by Dimensions and Error Bound

positive compression is possible.

In-line compression can yield negative compression if the decompressed data size is larger than the difference between the uncompressed data size and the compressed data size. The error bound and data size impact the in-line compression performance in the same way they impact the compressor that is chosen.

4.3 Timing

Similar to how traditional compression improves the compression ratio by leveraging more time and computational resources, obtaining a decrease in allocated memory usage comes at a cost of runtime. The main source of additional time usage is in the compression and decompression operations. This section presents the general compressor call pattern as well as some techniques to reduce the number of compression calls and to keep the runtime manageable.

4.3.1 Compression/Decompression Calls

4.3.1.1 Block Counts

Runtime analysis requires determining what operations takes up time, and the main timing difference between a standard kernel and an in-line kernel is the number of compression and decompression counts. Understanding what causes these compressor count increases is essential to selecting proper settings to reduce the runtime increases.

Figure 4.4 identifies which matrices are compressed and decompressed for the in-line compressed matrix multiplication kernel. A portion of each matrix is required to compute the next output matrix. The order in which operations are done matters for runtime because the correct portions of the matrices need to be decompressed in allocated memory. Identifying the reasons behind which blocks are decompressed more frequently enables understanding for better fitting a kernel for in-line compression with respect to its runtime.

This experiment compares the number of compression and decompression calls with respect to the number of blocks per kernel. The dimensions of the matrices analyzed are 50x50, and the matrix dimension is divisible by 5 and 10. It is worth noting that all compressors with the predefined standard error bounds are present for this experiment. The B matrix is transposed to ensure that all data is visible (see Figure 4.5).

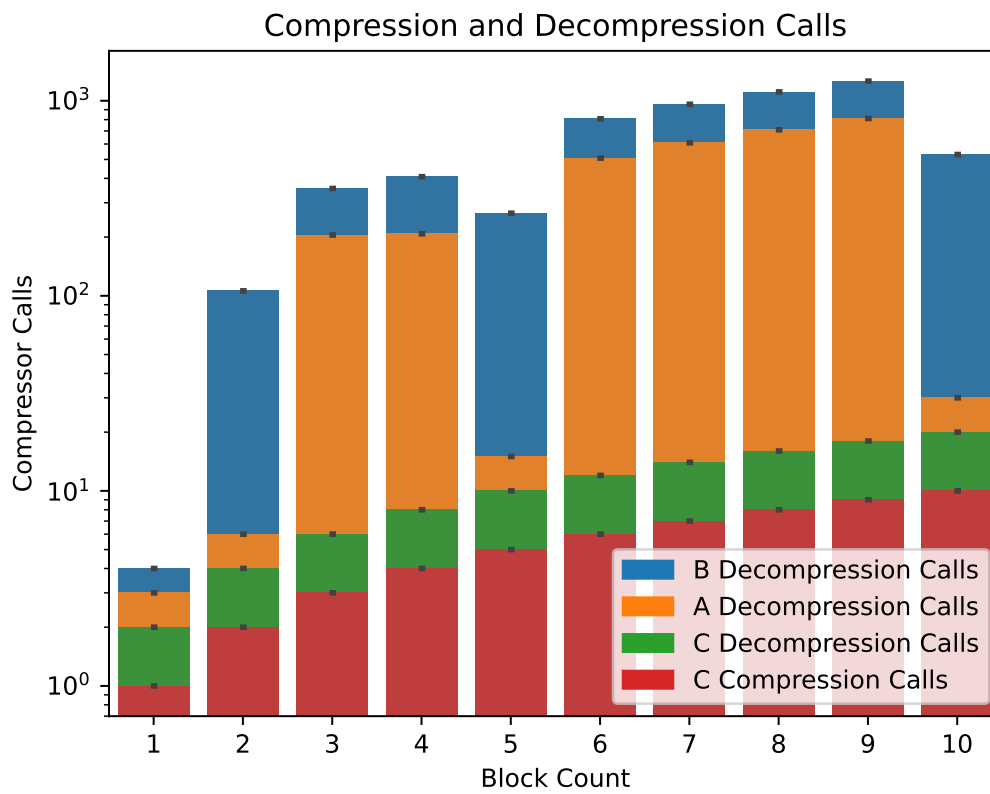


Figure 4.4: Compressor Call Counts per Matrix with respect to Block Counts

As the number of blocks increases, the number of compression and decompression operations also increases. Block counts of 1, 2, 5, and 10 are interesting due to having the least number of compression calls (Full) or actually having less compressor calls than neighboring block counts. The only matrix that is compressed is matrix C , as it is the one with the results that need to be stored. Note that the number of matrix C compressions is equal to the block size. This is because the order in which operations are done iterates through all C matrix entries, then through the blocks containing the entries in A and B to compute the C entry.

The decompressions provide the bulk of the compressor calls. All of the C blocks are decompressed once to access them to compute the resulting entries. This step could be removed for kernels where intermediate operations don't need to be stored, like for matrix multiplication. The number of A and B decompression counts provide factors of interest, particularly with how the selected decompressed blocks align with the kernel's access pattern.

For each entry in C , a single row from A and a single column for B need to be accessed to perform a standard matrix multiplication operation. How the blocks are selected impacts which ones are available to get data from. For block size 1 (full in-line), the number of decompressions is one for each matrix, as all of the rows and columns have been decompressed by the time they are needed for any of the matrix multiplication operations.

Block sizes 5 and 10 have different characteristics from their neighboring block counts. Not only are their overall compressor counts less than their neighbors, but the A matrix requires significantly less decompression counts. Because all dimensions tested are divisible by 2, 5 or 10, this illustrates an important kernel access pattern for matrix multiplication that can be used to reduce decompression calls. Because the blocks used for testing are 1-dimensional, each block can be a row or a part of a row. In this case, block sizes of 2, 5, and 10 are factors of the matrix row length. This feature limits future redundant operations for other data in that block (for example, if a block overflows into another row).

While block counts of 2, 5, and 10 have a significantly smaller number of compressor calls due to a decrease in A matrix decompressions, the B matrix decompressions look significantly larger. This is an artifact of the logarithmic scale used in this plot. A linear scale reveals that B matrix decompressions are directly related to block count. B is dependent on the kernel's access pattern.

This experiment also showcases what does not affect the number of compressor counts. Data dimensions, compressor choice, and error bound do not affect the compressor counts because they

are part of the compression process. They are not used to select how many counts, but rather how the data is compressed in those operations. What does affect compressor choice is kernel access pattern locality and the choices of how blocks are implemented.

4.3.1.2 Matrix Multiplication Informed Blocking

The previous experiment demonstrated how block choices can affect the compression and decompression operation counts. The choice in block dimensions and size details what kind of data is present in each block, and thus what is available during that step of the computation. The purpose of this experiment is to give an example of this property using the matrix multiplication kernel access pattern.

Each element in the output matrix C requires access to the corresponding row of matrix A and the corresponding column of matrix B . Because blocks with this system are 1-dimensional row-wise, one can access the partial or complete row with a limited number of decompression operations. However, since the data needed from the B matrix is by columns, the column data is contained in many row-based blocks. To compensate for this, an option for the kernel is to optimize the B matrix for the row-based blocking by transposing the B matrix. This enables the blocks to access the data needed by the kernel with less decompression calls.

For this experiment, the B Optimization flag for that kernel transposes the B matrix and updates the algorithm accordingly. The dimensions for matrices in this experiment were 100x100. All compressors with the predefined standard error bounds are present for this experiment.

Figure 4.5 indicates a large decompression count difference between using using the B -transpose optimization and keeping the standard operation. For the 100x100 matrix, since the block size is less than or equal to the number of elements in the row, accessing a column of B requires accessing 100 different blocks. When this operation is accounted for all 10,000 elements in the output matrix C , that setup requires 1,000,000 decompression operations for the B matrix alone. In all of these decompression steps, many unused elements are being decompressed along with the column data of interest. Transposing B allows the kernel access pattern to match the blocking pattern for in-line compression to limit the number of decompression counts and unnecessary data accesses.

The implications of this experiment go beyond just matrix multiplication. Whatever the base kernel is, an important means of limiting a runtime increase is to manage the block selection and kernel configuration to have blocks that meaningfully relate to the data needed for each operation

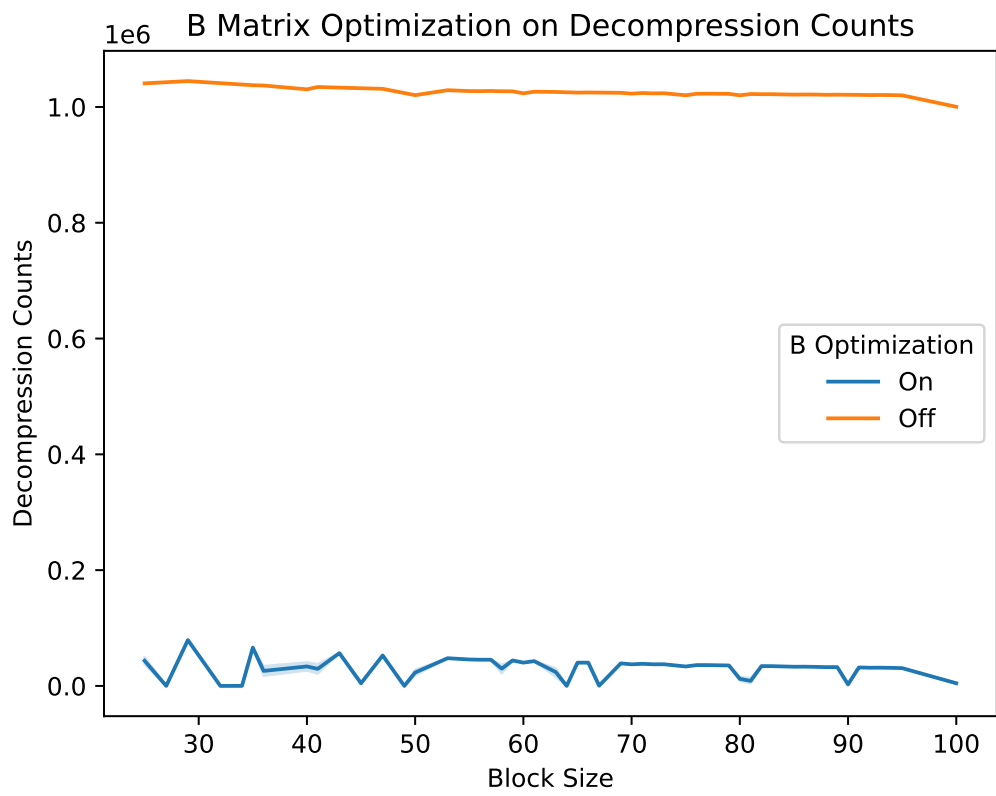


Figure 4.5: *B*-Transpose Kernel Access Pattern on Decompression Counts

in the kernel.

4.3.2 Runtime

After examining the sources of compression and decompression operations, how do these observations relate to runtime? The original kernel’s runtime complexity, the compressor, and the block count all affect the overall runtime of the in-line compression kernel. Figures 4.6 and 4.7 demonstrate how these choices impact the overall runtime of the in-line compression kernel.

These experiments relate runtime to the same parameters that memory was compared to: data size and block count. Additionally, the compressor choice is included to provide a better picture on all of the factors affecting runtime. These experiments compare the runtime to the matrix data dimensions, while also highlighting the compressor choice and the block counts. The runtime of the original kernel without in-line compression is also included for comparison.

The dimensions of the matrices analyzed range from 5x5 to 1000x1000. All compressors with the predefined standard error bounds are present for this experiment. *B* optimization is on for these experiments.

Adding any in-line compression to a kernel increases the runtime to compress and decompress, as Figure 4.6 demonstrates. The original matrix multiplication kernel has a $O(n^3)$ time complexity, and each of the in-line methods with various compressors follows suit. Each of the compressors has a similar relationship to what it had to memory in Figure 4.2: ZFP has the larger compression ratio and has the smaller runtime, while SZ has the smaller compression ratio and the larger runtime. An interesting point of note in Figure 4.6 is that around $n=700$, SZ-ABS gets a faster runtime briefly. This is peculiar that a compressed version would run faster than the kernel it is based on. An explanation for this could be the hardware cache optimizations, but this could be an area for future interest.

Another feature of interest with Figure 4.6 is that for all compressors with block size 1, the runtime of the compression gets amortized into the overall kernel runtime for large n . This indicates the possibility of in-line compression kernel runtimes in the same time complexity as the original kernel.

Figure 4.7 illustrates the direct relationship between runtime and the number of blocks. Each block adds additional compressor time to the original kernel time complexity. Not only is it additional time, but almost an order of magnitude more.




Figure 4.6: Runtime with All Compressors and Dimensions

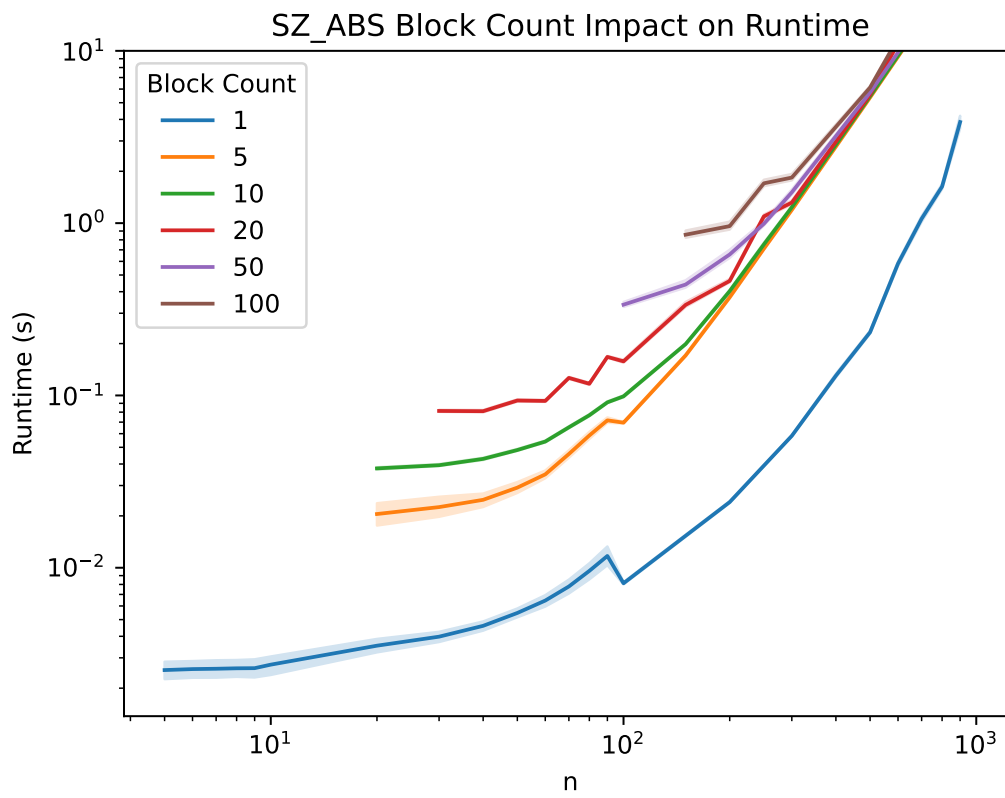


Figure 4.7: Runtime with SZ-ABS and Block Count

The in-line compression kernel has the time complexity of the original kernel with the linear increase of the number of compressor operations and the impact of the original compressor's runtime. Compared with Figure 4.6, Figure 4.7 reveals the possibility for tuning and improvement of the system itself. Offloading compressions to other threads and improved processor choice can mitigate the growth of runtime caused by additional blocks.

Chapter 5

Conclusions and Discussion

In-line compression provides a viable approach to reducing the allocated memory usage for HPC kernels with large kernel data. This thesis explored the runtime and memory performance of in-line compression on the matrix multiplication kernel. Matrix multiplication was selected as the exemplary kernel due to its prevalence in HPC systems and its ease of illustrating the properties of in-line compression.

The use cases for full and partial in-line compression were explored, revealing the trade-off between allocated memory usage and runtime. Larger block sizes provide a small allocated memory size reduction for a small runtime cost, while smaller block sizes enable greater reduction in allocated memory usage at the cost of runtime for compressing and decompressing those blocks. The spacial locality of those blocks impacts compressability, and block choices that leverage knowledge of both the data access pattern and the compressed array structure are required to obtain good performance.

Experimental results from the matrix multiplication in-line compression kernel were analyzed for key parameters' affects on runtime and allocated memory usage. Findings include insights into how in-line compression works and methods for memory and runtime optimization. In-line compression is possible through a balance between the compressed data size and the amount of data decompressed. The decompressed data size is dependent on the block size, while the compressed data is dependent on the location of blocks and the compressor choice. Compression and decompression calls can be minimized by blocking with respect to the kernel's access pattern. Improvements are possible which can enable the compression overhead time to be amortized by the original kernel's runtime.

In-line compression enables allocated memory reduction in large-data HPC kernels, and future work can be performed to improve this memory reduction method. Analyzing in-line compression on various kernels beyond standard matrix multiplication aids in identifying other features of interest for this method. Kernels like tiled matrix multiplication take advantage of 2-dimensional locality, so multi-dimensional blocks promise improvements in the compressed data size. The tuning of in-line compression parameters enables its impact on more kernels and applications than just matrix multiplication.

Bibliography

- [1] Bulent Abali, Hubertus Franke, Xiaowei Shen, Dan E. Poff, and T. Basil Smith. Performance of hardware compressed main memory. In *Proceedings of the Seventh IEEE International Symposium on High-Performance Computer Architecture, HPCA '01*, pages 73–81. IEEE, 2001.
- [2] Angelos Arelakis, Fredrik Dahlgren, and Per Stenstrom. Hycomp: A hybrid cache compression method for selection of data-type-specific compression methods. In *Proceedings of the 48th International Symposium on Microarchitecture*, pages 38–49, 2015.
- [3] Steve Ashby, Pete Beckman, Jackie Chen, Phil Colella, Bill Collins, Dona Crawford, Jack Dongarra, Doug Kothe, Rusty Lusk, and Paul Messina. The opportunities and challenges of exascale computing. Technical report, Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee, 2010.
- [4] Allison H. Baker, Haiying Xu, John M. Dennis, Michael N. Levy, Doug Nychka, Sheri A. Mickelson, Jim Edwards, Mariana Vertenstein, and Al Wegener. A methodology for evaluating the impact of data compression on climate simulation data. In *Proceedings of the 23rd International Symposium on High-performance Parallel and Distributed Computing, HPDC '14*, pages 203–214, New York, NY, USA, 2014. ACM.
- [5] Xi Chen, Lei Yang, Robert P. Dick, Li Shang, and Haris Lekatsas. C-pack: A high-performance microprocessor cache compression algorithm. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 18(8):1196–1208, 2009.
- [6] S. Di and F. Cappello. Fast error-bounded lossy hpc data compression with sz. In *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 730–739, May 2016.
- [7] Jack Dongarra, Bernard Tourancheau, Franck Cappello, Sheng Di, Sihuan Li, Xin Liang, Ali Murat Gok, Dingwen Tao, Chun Hong Yoon, Xin-Chuan Wu, Yuri Alexeev, and Fred-eric T Chong. Use cases of lossy compression for floating-point data in scientific data sets. *Int. J. High Perform. Comput. Appl.*, 33(6):1201–1220, nov 2019.
- [8] Fred Douglass. The compression cache: Using on-line compression to extend physical memory. In *Proceedings of 1993 Winter USENIX Conference*, pages 519–529, 1993.
- [9] Brian Van Essen, Henry Hsieh, Sasha Ames, Roger Pearce, and Maya Gokhale. Di-mmap – a scalable memory-map runtime for out-of-core data-intensive applications. *Cluster Computing*, 18:15–28, 2015.
- [10] Rosa Filgueira, Malcolm Atkinson, Alberto Nuñez, and Javier Fernández. An adaptive, scalable, and portable technique for speeding up mpi-based applications. In *Euro-Par 2012 Parallel Processing*, pages 729–740, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

- [11] Rosa Filgueira, David E. Singh, Jesús Carretero, Alejandro Calderón, and Félix García. Adaptivecompi: Enhancing mpi-based applications' performance and scalability by using adaptive compression. *The International Journal of High Performance Computing Applications*, 25(1):93–114, 2011.
- [12] Lisa Fischer, Sebastian Götschel, and Martin Weiser. Lossy data compression reduces communication time in hybrid time-parallel integrators. *Computing and Visualization in Science*, 19(1):19–30, Jun 2018.
- [13] Tanzima Zerín Islam, Kathryn Mohror, Saurabh Bagchi, Adam Moody, Bronis R. de Supinski, and Rudolf Eigenmann. Mcengine: A scalable checkpointing system using data-aware aggregation and compression. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC '12*, pages 17:1–17:11, Los Alamitos, CA, USA, 2012. IEEE Computer Society Press.
- [14] Animesh Jain, Parker Hill, Shih-Chieh Lin, Muneeb Khan, Md E. Haque, Michael A. Laurenzano, Scott Mahlke, Lingjia Tang, and Jason Mars. Concise loads and stores: The case for an asymmetric compute-memory architecture for approximation. In *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 1–13. IEEE, 2016.
- [15] Scott Levy, Kurt B. Ferreira, and Patrick G. Bridges. Improving application resilience to memory errors with lightweight compression. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '16*, pages 28:1–28:12, Piscataway, NJ, USA, 2016. IEEE Press.
- [16] X. Liang, S. Di, D. Tao, S. Li, S. Li, H. Guo, Z. Chen, and F. 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, December 2018.
- [17] P. Lindstrom. Fixed-rate compressed floating-point arrays. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2674–2683, December 2014.
- [18] Tong Liu, Jinzhen Wang, Qing Liu, Shakeel Alibhai, Tao Lu, and Xubin He. High-ratio lossy compression: Exploring the autoencoder to compress scientific data. *IEEE Transactions on Big Data*, PP, 03 2021.
- [19] Paul Messina. The u.s. d.o.e. exascale computing project - goals and challenges. Technical report, NIST, 2017.
- [20] Joshua San Miguel, Jorge Albericio, Andreas Moshovos, and Natalie Enright Jerger. Doppelgänger: A cache for approximate computing. In *Proceedings of the 48th International Symposium on Microarchitecture*, pages 50–61, 2015.
- [21] Gennady Pekhimenko, Vivek Seshadri, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry. Base-delta-immediate compression: Practical data compression for on-chip caches. In *Proceedings of the 21st International Conference on Parallel Architectures and Compilation Techniques*, pages 377–388, 2012.
- [22] Swann Perarnau, Judicael A. Zounmevo, Balazs Gerofi, Kamil Iskra, and Pete Beckman. Exploring data migration for future deep-memory many-core systems. In *2016 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 289–297. IEEE, 2016.
- [23] S. Ranjan. Performance modeling of inline compression with software caching for reducing the memory footprint in pysdc. Master's thesis, 2023.

- [24] S. Ranjan, D. Fulp, and J. C. Calhoun. Exploring the impacts of software cache configuration for in-line compressed arrays. In *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7, Waltham, MA, USA, 2022.
- [25] Ali Shafiee, Meysam Taassori, Rajeev Balasubramonian, and Al Davis. Memzip: Exploring unconventional benefits from memory compression. In *2014 IEEE 20th International Symposium on High Performance Computer Architecture (HPCA)*, pages 638–649. IEEE, 2014.
- [26] T. Basil Smith, Bulent Abali, Dan E. Poff, and R. B. Tremaine. Memory expansion technology (mxt): Competitive impact. *IBM Journal of Research and Development*, 45(2):303–309, 3 2001.
- [27] 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), 2014.
- [28] D. Tao, S. Di, Z. Chen, and F. Cappello. Significantly improving lossy compression for scientific data sets based on multidimensional prediction and error-controlled quantization. In *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 1129–1139, May 2017.
- [29] Dingwen Tao, Sheng Di, Xin Liang, Zizhong Chen, and Franck Cappello. Improving performance of iterative methods by lossy checkpointing. In *Proceedings of the 27th International Symposium on High-Performance Parallel and Distributed Computing, HPDC '18*, pages 52–65, New York, NY, USA, 2018. ACM.
- [30] R. Underwood, V. Malvoso, Jon C Calhoun, Sheng Di, and Franck Cappello. Productive and performant generic lossy data compression with libpressio. In *2021 7th International Workshop on Data Analysis and Reduction for Big Scientific Data (DRBSD-7)*, pages 1–10. IEEE, 2021.
- [31] Paul R. Wilson, Scott F. Kaplan, and Yannis Smaragdakis. The case for compressed caching in virtual memory systems. In *Proceedings of the Annual Conference on USENIX Annual Technical Conference, ATEC '99*, pages 8–8, Berkeley, CA, USA, 1999. USENIX Association.
- [32] P. Xenopoulos, J. Daniel, M. Matheson, and S. Sukumar. Big data analytics on hpc architectures: Performance and cost. In *2016 IEEE International Conference on Big Data (Big Data)*, pages 2286–2295, Washington, DC, USA, 2016.