

An Optimized Error-controlled MPI Collective Framework Integrated with Lossy Compression

Jiajun Huang,^{*} Sheng Di,[†] Xiaodong Yu,[‡] Yujia Zhai,^{*} Zhaorui Zhang,[§] Jinyang Liu,^{*} Xiaoyi Lu,[¶] Ken Raffenetti[†] Hui Zhou[†] Kai Zhao^{||} Zizhong Chen^{*} Franck Cappello[†] Yanfei Guo[†] Rajeev Thakur[†]

^{*}University of California, Riverside

[†]Argonne National Laboratory

[‡]Stevens Institute of Technology

[§]The Hong Kong Polytechnic University

[¶]University of California, Merced

^{||}Florida State University

Abstract—With the ever-increasing computing power of supercomputers and the growing scale of scientific applications, the efficiency of MPI collective communications turns out to be a critical bottleneck in large-scale distributed and parallel processing. The large message size in MPI collectives is particularly concerning because it can significantly degrade the overall parallel performance. To address this issue, prior research simply applies the off-the-shelf fix-rate lossy compressors in the MPI collectives, leading to suboptimal performance, limited generalizability, and unbounded errors. In this paper, we propose a novel solution, called *C-Coll*, which leverages error-bounded lossy compression to significantly reduce the message size, resulting in a substantial reduction in communication cost. The key contributions are three-fold. (1) We develop two general, optimized lossy-compression-based frameworks for both types of MPI collectives (collective data movement as well as collective computation), based on their particular characteristics. Our framework not only reduces communication cost but also preserves data accuracy. (2) We customize SZx, an ultra-fast error-bounded lossy compressor, to meet the specific needs of collective communication. (3) We integrate *C-Coll* into multiple collectives, such as MPI_Allreduce, MPI_Scatter, and MPI_Bcast, and perform a comprehensive evaluation based on real-world scientific datasets. Experiments show that our solution outperforms the original MPI collectives as well as multiple baselines and related efforts by 1.8–2.7×.

Index Terms—Lossy Compression, MPI Collective, Distributed Systems, Scientific Datasets

I. INTRODUCTION

MPI collectives provide high-performance collective communications in distributed systems, making a significant impact on various research fields such as scientific applications, distributed machine learning, and others [1], [2], [3], [4], [5], [6]. With the advent of exascale computing and deep learning applications, the demand for large-message MPI collectives has increased. For example, in image classification tasks, VGG19 [7] and ResNet-50 [8] have 143 million and 25 million parameters, respectively, with communication overheads of 83% and 72% [6]. Therefore, optimizing MPI collectives for large messages has become essential [9], [10], [11].

MPI collectives consist of both internode communication and intranode communication, and the former is often the major concern. The overall collective performance is usually limited by the efficiency of internode communications because

of limited network bandwidth. Therefore, optimizing internode collective communication is critical to improving the overall performance of MPI collectives. This topic has been a focus of research for decades, with state-of-the-art algorithms achieving notable improvements. However, with the increasing demand for large-message MPI collectives, further optimization remains necessary [12], [13], [11]. Lossy compression [14], [15], [16], [17] (rather than lossless compression [18], [19], [20]) is a promising solution to mitigate this MPI collective performance issue because of its ability to significantly reduce the message size.

Although lossy compression has been widely used to resolve many other scalability issues in high-performance computing, such as reducing memory footprint [21], reducing storage space [22], [23], and avoiding duplicated computation [24], only a few studies have explored its use in this direction, and all expose certain limitations. To elaborate, Zhou et al. [25] proposed GPU-compression enhanced point-to-point communication by integrating MPC [26] and 1D fixed-rate ZFP [17] into MVAPICH2 [27]. Their approach, referred to as *CPR-P2P*, simply involved compressing the messages before transmission and decompressing them after reception, leading to significant performance overhead due to the non-negligible time required for compression and decompression. Meanwhile, Zhou et al. [28], [29], [30] proposed several additional approaches to improve multiple MPI collectives using 1D fixed-rate ZFP [17] on GPUs. Their methods, however, focus on fixed-rate compression¹, introducing two major limitations: (1) compression errors cannot be bounded, leading to an uncontrolled accuracy, and (2) the compression quality is considerably lower compared to the fixed-accuracy mode² in ZFP, as demonstrated by prior research [31].

The aforementioned limitations of compression-enabled MPI collective algorithms motivate us to develop a new efficient MPI collective framework which leverages lossy compression technique to significantly improve the MPI col-

¹Fixed-rate compression means that the lossy compression would be performed based on a user-specified fixed compression ratio.

²Fixed-accuracy, also known as error-bounded lossy compression, compresses data based on a user-specified error bound.

lective performance. However, this brings in three direct technical challenges. **A** Devising a general framework that can effectively hide the communication cost and choose an appropriate timing to call lossy compression is non-trivial. **B** The data loss nature of the lossy compression brings up a critical concern on the accuracy of collective operations. **C** Existing lossy compressors are not designed for the collective context, leading to suboptimal collective performance because of unnecessary overheads when they are directly applied in MPI collectives [14], [15], [32], [33], [17].

Our developed framework is named as compression-facilitated MPI collective framework (*C-Coll*), which can address the aforementioned limitations and challenges. To the best of our knowledge, this is the first-ever framework that provides a general high-performance solution for compression-integrated MPI collectives. Moreover, this is the first accuracy-aware design, which ensures the accelerated collective performance with error-bounded lossy compression does not compromise data quality. To be more specific, our contributions include:

- To address challenge **A**, we introduce two efficient frameworks, which can significantly accelerate both types of MPI collectives. Specifically, the first framework notably diminishes the compression overhead in the collective data movement operations (e.g., Scatter, Bcast and All-gather), thereby achieving substantial performance improvement. The second one hides communication inside of compression in collective computation (e.g., Reduce-scatter), which in turn enhances the performance of collective computation. Moreover, these two frameworks can be combined together to speed up more advanced MPI collective operations such as All-reduce.
- To address challenge **B**, we devise several strategies to effectively control the error propagation in the *C-Coll* framework. On the one hand, we carefully select the most suitable error-bounded lossy compressor based on a comprehensive analysis of various state-of-the-art lossy compression methods, considering multiple aspects of their performance on MPI collectives. On the other hand, we carefully regulate the number of compression operations and also develop an efficient method to resolve the imbalance in collective communication introduced by the error-bounded lossy compression.
- To address challenge **C**, we customize SZx to cater to the specific needs of collective communication. Specifically, we redesign the compression workflow of SZx and implement a pipelined version of SZx to overlap compression and communication. As a result, we significantly decrease the communication cost of our *C-Coll* framework by up to $4.9\times$.
- To prove the error-bounded nature of the *C-Coll* framework, we perform an in-depth mathematical analysis to derive the limited impact of lossy compression on error propagation. Theoretical analysis demonstrates that the final error of collective data movement framework

of *C-Coll* can be well-bounded within the compression error bound set by the user. For collective computation framework, we can derive the final aggregated error is bounded well with a high probability. For example, if there are 100 nodes, the aggregated error is bounded in the range of $[-\frac{20}{3}\hat{e}, \frac{20}{3}\hat{e}]$ with a probability of 95.44%, where \hat{e} is the compression error bound.

- To demonstrate the generality of our design, we integrate *C-Coll* into multiple collectives, such as MPI_Allreduce, MPI_Scatter, and MPI_Bcast, and evaluate the performance carefully with various scientific datasets generated by real-world applications (such as Seismic Imaging [34], Hurricane Simulation [35], and Climate Simulation [36], [37]). Experiments with 128 Xeon Broadwell compute nodes from a supercomputer show that other related efforts or baselines exhibit undesired performance degradation on MPI_Allreduce because of the significant compression overhead. In comparison, our solution—the lossy-compression-based Allreduce (we call it C-Allreduce)—outperforms the original MPI_Allreduce by $2.1\times$. Moreover, our compression-based scatter and Bcast operations (namely C-Scatter and C-Bcast) achieve up to $1.8\times$ and $2.7\times$ performance improvements compared with the original MPI_Scatter and MPI_Bcast. We also use a real-world use-case (image stacking analysis) to validate the practical effectiveness of C-Allreduce, which shows up to $1.5\times$ performance gain over MPI_Allreduce, while preserving a high data integrity/accuracy during the collective operations.

The rest of the paper is organized as follows: we introduce background and related work in Section II and detail our design and optimization in Section III. Evaluation results are presented in Section IV followed by conclusion and future work in Section V.

II. BACKGROUND AND RELATED WORK

In this section, we discuss the background and related work. We first introduce MPI collective communication, followed by a discussion on high-speed lossy compressors and their integration with MPI implementations. The focus of our study is on lossy compression. This emphasis is due to the significantly lower compression ratios observed with lossless methods when applied to scientific datasets [14], [15].

A. MPI Collective Communication

There are many types of MPI collective operations, which can be divided into two sub-categories — collective data movement and collective computation according to their communication patterns.

1) *Collective data movement*: Collective data movement includes gather, allgather, scatter, all-to-all, and so on. The gather operation collects the data from different processes and stores the collected data into the root process. In comparison, allgather stores the collected data to every participated process. As an opposite of gather, the scatter operation divides the data in the root process and sends the split data to all the

processes. The all-to-all operation acts as the “allscatter”, which collectively scatters data on each process to each other.

2) *Collective computation*: Allreduce/reduce are two popular collective computation operations. We use MPI_SUM as an example to explain the working principle as it is frequently used. The reduce routine will sum up all the data entries from all the processes in the same communicator and store the sum into the root process. The Allreduce does the same thing but keeps a copy of the sum on every process in that communicator. Another widely-used operation is the reduce-scatter, which acts like a combination of the reduce and scatter: the reduced sum is scattered into all the processes.

B. High-speed Lossy Compressors

Compression developers and scientific researchers have shown great interest in high-speed lossy compression due to its ability to achieve a very high compression ratio. SZ [14], [15], [32] is an example of a fast, error-bounded lossy compressor that offers performance similar to other compressors, including FPZIP [38] and SZauto [16]. Another compressor that is known for its relatively high compression ratios and even faster compression speed than SZ is ZFP [17]. However, neither of these compressors can match the speed of SZx [33], which achieves a compression throughput of 700-900 MB/s on CPUs [33]. Additionally, our comprehensive experiments and analysis, as discussed in Section III-C, demonstrate that SZx is the most suitable lossy compressor for MPI collective operations. As such, we developed our customized compressor for MPI collectives based on SZx, which will be detailed in Section III-C.

C. Lossy Compression-enabled MPI Implementations

Researchers have shown interest in using lossy compression to improve MPI communication performance for years. Zhou et al. proposed GPU-compression enhanced point-to-point communication [25], and several optimized MPI collective operations [28], [29], [30] using 1D fixed-rate ZFP [17] on GPUs, but their solutions are either showcasing limited overlapping between compression and communication or subject to the fixed-rate mode compression, leading to the substandard compression quality, as will be demonstrated later in Section IV-E. Conversely, our general framework can optimize the collective performance of all MPI collectives while maintaining controlled errors.

III. C-COLL DESIGN AND OPTIMIZATION

Figure 1 presents the overall design architecture of *C-Coll*. We highlight the newly designed modules as green boxes. The primary contributions lie in the performance optimization layer and the middleware layer. We carefully characterize the performance of multiple state-of-the-art error-bounded lossy compressors in the context of MPI collectives and select the best-qualified compressor – SZx, which will be detailed in Section III-C. We also propose a series of performance optimization strategies specifically for both of the two collective types (data movement and collective computation),

as indicated in the figure. The corresponding details will be discussed in Sections III-A1, III-A2, and III-E2.

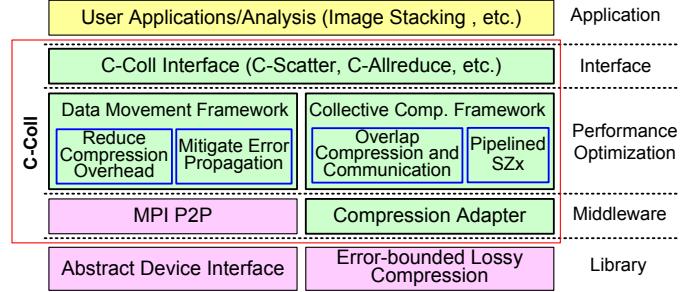


Fig. 1: Design architecture (yellow box: applications; green box: new contributed modules; purple box: third-party).

A. Two Proposed Novel Frameworks for Compression-enhanced Collectives

To integrate lossy compression into MPI collective communications, at least two important aspects must be considered: performance and accuracy. In general, MPI collective operations can be divided into two groups: collective data movements, and collective computation. Instead of directly using the CPR-P2P method, we propose two frameworks to implement collective communications for each group, which can maximize the collective operation performance.

1) *Collective data movement framework*: In this subsection, we detail our strategies for addressing the issues of communication imbalance and compression overhead in our optimized collective data movement framework. For collective data movement operations, each process in the same communicator needs to communicate with each other to exchange data. If we directly use the CPR-P2P method, the sender needs to compress the data every time before sending it, and the receiver is required to decompress the data upon the data arrival. Most of the compression and decompression overheads in CPR-P2P, actually, can be avoided by carefully setting the timing of compression operations, as the original data have not been modified during the intensive communications. Besides, the CPR-P2P can cause unbalanced communications in that input data on different processes have various compressed data sizes. Such unbalanced communication will slow down the overall collective performance, resulting in a sub-optimal performance. With our framework, we can balance the communications with a fixed pipeline size as the compressed data sizes are decided at the beginning of the intensive communications. In the following text, we illustrate our idea with two examples: a ring-based allgather algorithm and a binomial tree broadcast algorithm. The same philosophy can be easily extended to other collective algorithms in this category.

Figure 2 shows the high-level comparison of our proposed framework versus the CPR-P2P in the ring-based allgather algorithm. In this algorithm, $N-1$ rounds are required to get the gathered results on every process, where N is the number of processes in the communicator. In order to use lossy compression to reduce communication cost in the ring-based

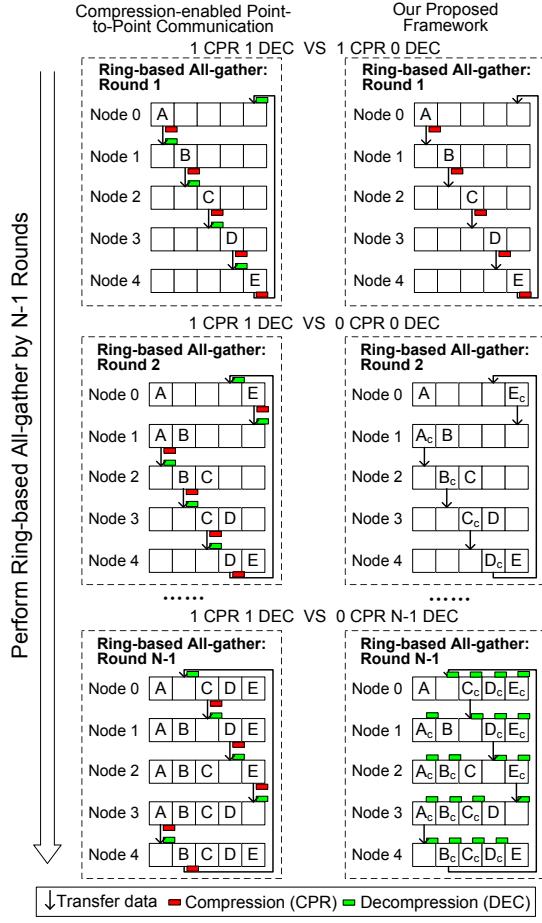


Fig. 2: High-level design of our collective data movement framework in the ring-based allgather algorithm to mitigate compression error propagation. A means the original data and A_c means the compressed data. This rule applies to other data chunks as well. This algorithm completes in $N-1$ rounds, where N is the number of processes.

algorithm, the straight-forward idea is performing compression and decompression at each round. Instead, our design does not decompress the received data until the last round, thus significantly decreasing the compression overhead from $(N-1) \cdot T_{chunk}$ to T_{chunk} , where T_{chunk} is the compression cost of one chunk. When N is large, our novel framework could have nearly $N \times$ better performance compared with CPR-P2P in terms of compression. Note that the decompression cost remains the same.

Figure 3 presents a similar comparison in the binomial tree broadcast algorithm. There are $\log_2 N$ rounds before the completion, where N refers to the total process count. We notice that our proposed framework could reduce the compression and decompression costs from $\log_2 N \cdot (T_{comp} + T_{decom})$ to $T_{comp} + T_{decom}$ and the performance improvement is $\log_2 N \times$, where T_{comp} and T_{decom} are the compression time and decompression time of the data at the root process, respectively.

Besides the compression cost, the CPR-P2P method can lead to an undesirable error propagation issue during collective sends and receives, as the same data chunk undergoes multiple

rounds of compression and decompression. Our framework also solves this issue by compressing the same data chunk for only one time. Similar to the analysis of compression overhead, for the absolute error bounded compression, our proposed framework can decrease the worst case accuracy loss by $(N-1) \times$ and $(\log_2 N) \times$ in the ring-based allgather and binomial tree broadcast algorithm, respectively.

2) **Collective computation framework:** For collective computation routines, the data entries from all processes in the same communicator need to collectively compute with each other. Unlike in the case of collective data movement, the data transferred in this communication pattern can be updated. As a result, the previous framework cannot be utilized here thus we need to propose a new framework. Despite the updated transferred data precluding us from diminishing compression, we find an opportunity to hide communication inside the compression and decompression. To clearly elaborate our proposed design, we use the ring-based reduce_scatter algorithm as an instance. Note that this framework can be easily extended to other collective computation operations.

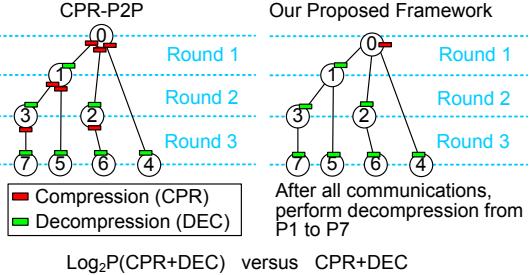


Fig. 3: High-level design of our collective data movement framework in the binomial tree broadcast algorithm. It completes in $\log_2 N$ rounds, where N is the number of processes.

Our proposed framework for collective computation is depicted in Figure 4. It employs a ring-based reduce_scatter algorithm, where each process is required to exchange message chunks with its neighboring processes. For large datasets, these chunk sizes become substantial as they are determined by dividing the size of the input data by the number of processes. In the initial CPR-P2P model, compression and decompression occur before and after any communication, respectively. Consequently, a single round incurs three types of overhead: compression/decompression for one message chunk, send/receive operations for the compressed message chunk, and reduction operation for one message chunk. Typically, the compression-related overhead is more significant than the send/receive overhead, as compressed data sizes are considerably smaller than their original counterparts. In our redesigned approach, we significantly mitigate the send/receive overhead by actively pulling communication progress within the compression and decompression phases. This substantially reduces the overall communication time.

B. Theoretical Analysis of Error Propagation in C-Coll

In this section, we prove the error-bounding nature of our C-Coll framework mathematically. In the following analysis,

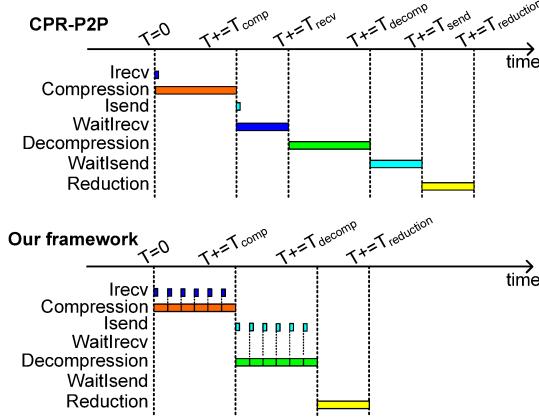


Fig. 4: High-level design of our collective computation framework in the ring-based reduce-scatter algorithm.

we assume the lossy compression error e for data x follows a normal distribution, without loss of generality. Specifically, the normal distribution is represented as $e \sim N(\mu, \sigma^2)$ within the range of $[x - \hat{e}, x + \hat{e}]$, where \hat{e} is the compression error bound. This is exemplified in Figure 5, in which we compress climate, weather, seismic wave datasets by SZ3 and ZFP. It clearly shows the normal distribution curve generated by Maximum Likelihood Estimation (MLE) fits the measured compression error values very well for different application datasets.

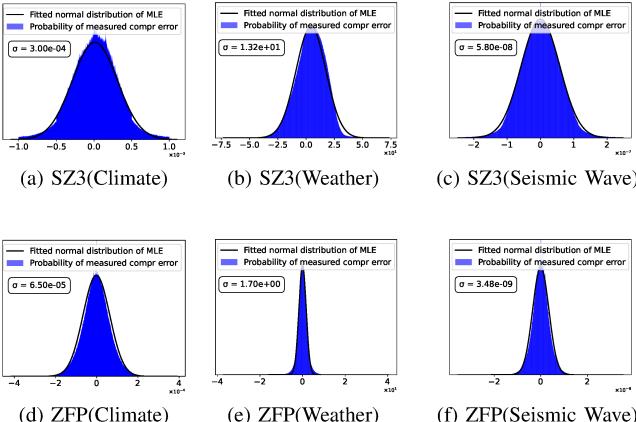


Fig. 5: Exemplifying the normal distribution property of compression errors.

For the collective communication primitive in MPI, the data will be aggregated gradually from each node during the communication process, which is shown in Fig. 2, Fig. 3 and Fig. 4. With lossy compression integrated, the compression error will also be aggregated in the data aggregation stage. In the collective data movement framework of *C-Coll*, each data chunk is compressed only once. Thus, the final error for each data point is within \hat{e} . Unlike the data movement framework, the compression error is aggregated in the collective computation framework of *C-Coll*, and the aggregation function often involves the *Sum*, *Average*, *Max*, *Min* operations for the floating point data. We further illustrate the error propagation for

these operations. For the collective computation framework, we assume there are n data that are collected from n nodes, and the compression error for each of them is e_i , which follows the normal distribution $e_i \sim N(\mu_i, \sigma_i^2)$.

Theorem 1. *Based on the above analysis, the final aggregated error for Sum operation falls into the interval $[-2\sqrt{n}\sigma, 2\sqrt{n}\sigma]$ with the probability of 95.44%, where n is the number of computing nodes in MPI and σ is the variance of the error bound of the lossy compressor.*

Proof. The linear combination (e.g., *Sum* in MPI) of normally distributed random variables also follows a normal distribution that is shown in the formula (1), where a_i denotes various constants for n data.

$$\sum_{i=0}^n a_i e_i \sim N\left(\sum_{i=0}^n a_i \mu_i, \sum_{i=0}^n a_i^2 \sigma_i^2\right) \quad (1)$$

For the *Sum* operation in the collective computation framework, the compression error will be involved gradually with the aggregation chain, which is shown in the formula (2):

$$\begin{aligned} x_{sum} &= (((x_1 + e_1) + x_2) + \dots + e_n) \\ &= (x_1 + e_1) + (x_2 + e_2) + \dots + (x_n + e_n) \end{aligned} \quad (2)$$

Where x_i indicates the data collected from node i and x_{sum} represents the final aggregated data. We further demonstrate that the compression error e_2 remains to follow the normal distribution after compression through Figure 6, which illustrates the probability of the measured compression error (e_2) with the fitted MLE curve. The same property also holds for subsequent compression errors in the aggregation chain such as e_3, e_4 , and so on.

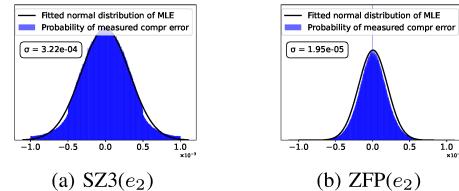


Fig. 6: Exemplifying the normal distribution property of compression error e_2 .

From the formula (2), we can calculate the aggregated compression error as $\tilde{e}_{sum} = \sum_{i=0}^n e_i$. The aggregated error \tilde{e}_{sum} follows the normal distribution as formula (3):

$$\tilde{e}_{sum} \sim N\left(\sum_{i=0}^n \mu_i, \sum_{i=0}^n \sigma_i^2\right) \quad (3)$$

The formula (3) indicates that the variance σ^2 of the aggregated error \tilde{e}_{sum} will be restricted well. When we utilize the same compression error bound across various nodes, the aggregated error conforms to a normal distribution represented as $\tilde{e}_{sum} \sim N(0, n\sigma^2)$. Therefore, the final aggregated error falls within the interval $[-2\sqrt{n}\sigma, 2\sqrt{n}\sigma]$ with the probability of 95.44% according to the properties of normal distribution. \square

Since the compression error is bounded by the error bound \hat{e}_i and the error follows the norm distribution, we can assume that $\hat{e}_i \approx 3\sigma_i$ (\hat{e}_i bounded to $3\sigma_i$ with probability of 99.74%).

Corollary 1. *Based on the above assumption, the final aggregated error falls within the interval $[-\frac{2}{3}\sqrt{n}\hat{e}, \frac{2}{3}\sqrt{n}\hat{e}]$ with the probability of 95.44%. For example, if there are 100 nodes, the final aggregated error will be bounded within the range $[-\frac{20}{3}\hat{e}, \frac{20}{3}\hat{e}]$ with a probability of 95.44%.*

Corollary 2. *Being similar to the Sum operation, the final aggregated error for the Average operation in collective computation framework follows the normal distribution $\hat{e}_{avg} \sim N(0, \frac{\sigma^2}{n})$. The final aggregated error will be reduced extremely compared to the original error by n times.*

Theorem 2. *For the Max, Min operations, the final error follows the normal distribution $\hat{e}_{max,min} \sim N(0, (2 - \frac{n+2}{2^n})\sigma^2)$.*

Proof. Since we need to compare the data from the neighbored node gradually, there is a $\frac{1}{2}$ probability that we can choose the non-compressed data. Otherwise, the selected data will contain an error within the error bound \hat{e} . Therefore, the variance of the final aggregated error can be calculated as following formula (4):

$$\frac{1}{2^n}n\sigma^2 + \frac{1}{2^{n-1}}(n-1)\sigma^2 + \dots + \frac{1}{2}\sigma^2 = (2 - \frac{n+2}{2^n})\sigma^2 \quad (4)$$

Moreover, the FXR mode cannot control the error bound, which may cause fairly high compression errors on some data points unexpectedly. Compared with ZFP, SZx is an ultra-fast error-bounded compressor driven by the fixed-accuracy mode.

TABLE I: OVERALL COMPRESSION/DECOMPRESSION THROUGHPUT (MB/S)

Datasets	RTM			Hurricane			CESM-ATM		
	ABS	Com	Decom	Com	Decom	Com	Decom	Com	Decom
SZx	1E-2	1742	3309	1687	3640	666	1251		
	1E-3	1479	2723	895	1659	533	918		
	1E-4	1288	2215	644	1168	526	822		
ZFP(ABS)	ABS	Com	Decom	Com	Decom	Com	Decom		
	1E-2	1383	1444	492	634	240	273		
	1E-3	1082	1141	307	397	170	191		
	1E-4	783	811	170	209	128	136		
ZFP(FXR)	FXR	Com	Decom	Com	Decom	Com	Decom		
	4	610	601	251	319	335	397		
	8	438	413	129	142	200	203		
	16	324	311	82	81	119	112		

TABLE II: COMPRESSION RATIOS (ORIGINAL DATA SIZE / COMPRESSED DATA SIZE)

Datasets	RTM			Hurricane			CESM-ATM			
	ABS	min	avg	max	min	avg	max	min	avg	max
SZx	1E-2	88	116.3	124.1	121	123.1	124.1	4.9	8.5	22.8
	1E-3	26.1	49.4	115.1	14.9	17.4	29.1	3.3	5.1	13.1
	1E-4	9.5	30.4	111.3	6.9	7.2	8.8	2.4	3.4	8
ZFP(ABS)	ABS	min	avg	max	min	avg	max	min	avg	max
	1E-2	67.7	87.6	125.2	18.1	18.3	18.6	4.7	8.1	23.1
	1E-3	30.6	58.4	123.2	10.5	10.7	11	3.4	5.6	15
	1E-4	13.4	38	120.1	5.4	5.5	5.8	2.4	3.8	9.7
ZFP(FXR)	FXR	min	avg	max	min	avg	max	min	avg	max
	4	8	8	8	8	8	8	8	8	8
	8	4	4	4	4	4	4	4	4	4
	16	2	2	2	2	2	2	2	2	2

TABLE III: COMPRESSION QUALITIES (PSNR)

Datasets	RTM			Hurricane			CESM-ATM			
	ABS	min	avg	max	min	avg	max	min	avg	max
SZx	1E-2	31.2	37.3	61.5	25.2	26.2	26.8	46	50.6	55.9
	1E-3	40.6	51.4	80.9	40.3	42	42.8	64.2	70.3	73.4
	1E-4	60.2	72	102.6	62.4	63.6	64.3	82.8	93.5	97.2
ZFP(ABS)	ABS	min	avg	max	min	avg	max	min	avg	max
	1E-2	36.6	45.7	71.8	33	34	34.3	52.6	56.8	61
	1E-3	49.2	59.5	88.2	47.8	49	50	69.1	73.2	77.6
	1E-4	69.3	80.1	111	68.5	69.7	70.4	92.2	96.6	100.6
ZFP(FXR)	FXR	min	avg	max	min	avg	max	min	avg	max
	4	41.9	46.4	56.1	30.9	31.6	32.3	30.3	34	39.9
	8	66.7	71.6	80.1	56	57.8	64.9	58.8	60.5	62.9
	16	113.8	118.5	127.3	103.7	105.5	113	106.4	108.1	110.6

In addition to the above analysis, we also compare SZx, ZFP(ABS), and ZFP(FXR) regarding the compression throughput, ratio, and quality, using different datasets with various error bounds or rates, as shown in Table I, II, and III. To ensure a fair evaluation, all compression and decompression processes were executed using a single thread on an Intel Xeon E5-2695v4 CPU. We adopt the 1D compression mode in that the dimensional information will have to be skipped due to the 1D chunk-wise design in most of the MPI collectives. The information about the datasets we used here is detailed in Table IV, which will be found in Section IV-A. Because of the space limit, we use the QVAPORF and CLOUD fields in the Hurricane and CESM-ATM dataset, respectively, in our experiments, and other fields exhibit very similar results. We observe that SZx is much faster than ZFP(ABS) by up to $4.1\times$ in compression and $5.7\times$ in decompression with similar compression ratios and qualities. Additionally, with similar compression qualities, ZFP(FXR) has the lowest compression speed and compression ratio compared to both ZFP(ABS) and SZx. Therefore, we develop our customized compressor based on SZx in terms of the context of MPI collectives. For

the purpose of comparison, we also implement compression-enabled point-to-point communication-based collectives based on both ZFP(FXR) and ZFP(ABS), which serve as baselines.

D. Characterization of Performance Bottlenecks

We integrate SZx into point-to-point communication of the ring-based allreduce algorithm, in order to understand the key bottlenecks of the collective performance, which will be a fundamental work to guide our optimization strategies. In this characterization, we use ring-based allreduce which serves as a very good example, because it consists of both collective data movement (allgather) and collective computation (reduce_scatter). Figure 7 presents the detailed performance breakdown for direct integration of SZx under a series of data sizes. Our analysis reveals that in the original ring-based allreduce algorithm, the all-gather operation accounts for approximately 60% of the overall execution time, as the communications are not overlapped at all with each other, unlike the reduce-scatter stage. The Wait operation is the second most time-consuming one, which waits for the completion of non-blocking send and receives before conducting the reduction operation. The remaining operations in the original allreduce algorithm include Memcpy (local data copies), Reduction (Reduce operations), and Others (data allocations and other calculations), which together constitute about 20% of the overall execution time. After integrating the SZx, the bottleneck turns into the ComDecom (compression and decompression) as the data needs to be compressed before being sent out and needs to be decompressed every time the receiver receives them. We also observe a reduction in Allgather and Wait times, suggesting a decrease in transferred data, and both of the MPI-related time can be further optimized. However, the Others part also takes a significant amount, specifically 23% in the 278MB case. This is because the SZx requires users to free compression-generated buffers after the compressor is called, resulting in a significant overhead.

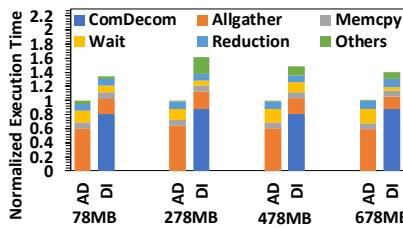


Fig. 7: Compare the performance of AD(Original MPI_Allreduce) and the DI(Direct Integration) of SZx from 78MB to 678MB with a 200MB step.

E. Step-wise Optimizations

In this subsection, we detail our step-wise optimization strategies, a principal contribution of this paper. To facilitate explanation, we present our implementation of a ring-based allreduce algorithm that incorporates lossy compression. These optimization strategies are applicable to other collective operations as well. Notably, in the ring-based allreduce, the data transfer required for each process is just $\frac{2(N-1)}{N} \cdot D_{input}$,

where D_{input} represents the input data size and N is the number of processes. Thus, this design is highly efficient for processing long messages. Furthermore, our integration of lossy compression significantly enhances the efficiency of collective communication involving long messages. We have termed our lossy-compression-enabled MPI_Allreduce as C-Allreduce, with ‘C’ denoting compression. Subsequently, we will discuss the design and implementation specifics of C-Allreduce in a structured, step-by-step approach.

1) Utilize our collective data movement framework:

To reduce the compression overhead and balance communications, we utilize the data movement framework that we presented in Section III-A1. At the beginning, every process compresses its local data and stores the compressed data size. Then, every process synchronizes with each other to collect the compressed data sizes in a local integer array *compressed_sizes*. As the compressed data size only has four bytes, this step is very fast. After that, all processes get the sum of all the compressed data sizes, noted as *total_count*. Then, each process communicates with each other with a fixed pipeline size until every process has sent *to_send* = *total_count* - *compressed_sizes*[*send_rank*] and received *to_recv* = *total_count* - *compressed_sizes*[*self_rank*] from other processes in a ring communication pattern. After all communications end, every process starts to decompress all the received compressed data and store the decompressed data in the receive buffer. Note that they do not need to decompress the data that are compressed by themselves. After this step, we can significantly decrease the time spent by compression and allgather communications compared with the direct integration of SZx. Besides, our solution can also preserve the quality/accuracy of the data very well because of the error-bounding feature, which will be demonstrated later in Section IV-E.

2) Customize SZx to reduce communication overhead with our collective computation framework:

In order to use our collective computation framework in the reduce-scatter stage of the ring-based allreduce algorithm, we need to redesign the compression workflow of SZx so that we can consistently poll the progress of the Isend and Irecv inside of the compression and decompression. Therefore, we design and implement the PIPE-SZx (pipelined SZx) based on the original SZx. Instead of compressing the original data as a whole, we divide the compression process into small chunks, each of which handles 5120 data points. Between the compression of two adjacent chunks, we actively poll the communication progress of the non-blocking receive. However, the compressed data of each chunk cannot be simply combined together, otherwise the compressed data cannot be correctly decompressed because each compressed chunk is of variable uncertain length. To solve this problem, we decide to store the compressed data of all chunks in the same output buffer and pre-allocate enough memory space (four bytes per chunk, small memory consumption) at the front of the buffer for storing the compressed data sizes of those chunks together (essentially a kind of index), instead of storing them along with the compressed data chunks.

Such a design is more cache-friendly, thus having lower overhead. During the decompression, we maintain a chunk-starting-location pointer based on the recorded compressed chunk sizes to tell the algorithm where the decompression operation should start for each chunk. We repeat this process chunk by chunk and poll the progress of the non-blocking send between decompression chunks. Through this optimization, we can hide the communication in the reduce-scatter stage inside of compression, which further improves the performance of our C-Allreduce design.

IV. EXPERIMENTAL EVALUATION

In this section, we present and discuss the evaluation results.

A. Experimental Setup

Since inter-node communication is the major bottleneck for collectives as discussed previously, we utilized a 128-node cluster with one process per node in our experiments. Each node is equipped with two Intel Xeon E5-2695v4 Broadwell processors. Furthermore, each NUMA node contains 64 GB of DDR4 memory, resulting in a total of 128 GB of memory per node. The nodes are interconnected via Intel Omni-Path Architecture (OPA), providing a maximum message rate of 97 million per second and a bandwidth of 100 Gbps.

TABLE IV: Information of The Scientific Datasets

Applications	# files	Dimensions	Descriptions
RTM[34]	70	849×849×235	Seismic Wave
Hurricane[35]	48×13	100×500×500	Weather Simulation
CESM-ATM[40]	26×33	1800×3600	Climate Simulation

MPI collectives are common operations used in the simulation analysis. For instance, generating stacking images in RTM (essentially an allreduce sum operation) is a typical real-world example [41], which will be demonstrated at the end of this section. We also utilize various datasets from a series of applications, including the reverse time migration (RTM) dataset, Hurricane dataset, and CESM-ATM dataset, to evaluate our solution. Table IV shows the detailed specifications of these datasets. Our baselines consist of original MPI collectives in MPICH 4.1.1, MPI collectives implemented by compression-enabled point-to-point communications with SZx, fix-accuracy ZFP, and fix-rate ZFP (version 0.5.5). For our experiments, we adopt a two-stage approach, including a warm-up stage and an execution stage. We conduct 10 runs for each stage and report the average results to present the general performance.

B. Step-wise Optimizations to C-Allreduce with Performance Analysis

In this section, we carry out optimizations to our C-Allreduce (*C-Coll* enhanced Allreduce) integrated with SZx step by step and show the performance on 16 Broadwell nodes. The ring-based allreduce that we implement contains a reduce-scatter stage and an allgather stage. Thus, we breakdown the total execution time into several parts: ComDecom (i.e., the time needed to compress and decompress data), Allgather (i.e., the time required to transfer data in the allgather stage),

TABLE V: Step-wise Optimizations of C-Allreduce

Method (Abbr.)	Description about the Allreduce implementation
Original MPI_Allreduce (AD)	No compression
Direct Integration (DI)	Implemented with CPRP2P
Novel Design (ND)	Optimized by our collective data movement framework
Overlapped Optimization (Overlap)	Optimized by our collective computation framework

Memcpy (i.e., the time spent on coping data in the reduce-scatter stage), Wait (i.e., the non-overlapped time spent on transferring data in the reduce-scatter stage), Reduction (i.e., the time required to do reduction operation), and Others (i.e., the time needed to do other data allocation and calculations). We adopt the RTM dataset to demonstrate our step-wise optimizations, which is the largest one among all three representative scientific datasets we evaluate. All the variants benchmarked in the following contexts are summarized in Table V.

1) Evaluating our collective data movement framework:

Figure 8 shows the performance improvement achieved by our novel design in the allgather stage for data sizes ranging from 28MB to 678MB. This Novel Design (ND) leverages our collective data movement framework, resulting in a considerable reduction in both compression and decompression time. Notably, at the data size of 128MB, the ND achieves a speed-up of up to 1.48× when compared to the Direct Integration (DI) approach discussed in Section III-D. Additionally, our balanced communication in Allgather using the ND is up to 7.1× faster than the unbalanced communication in DI at 628MB. We analyze the key reason why our solution can obtain a significant performance improvement as follows. In fact, to overcome the compression bottleneck and balance MPI communications, we utilize our collective data movement framework that pre-compresses the data before transmission and decompresses it after all communications, rather than using expensive compression-enabled point-to-point communication (CPR-P2P) in collective routines. This novel design can significantly reduce the amount of compression required during collective communication. Using CPR-P2P also brings unbalanced communications as the compressed data sizes may vary, but we can balance the communications with a fixed pipeline size in our new design because we do not need to compress the data every time before we send it.

Note that using CPR-P2P may accumulate errors during intensive collective communication, such as ring-based communication, as the same data is repeatedly passed from one process to another. Therefore, we have utilized our new framework to ensure that errors in the final results are bounded, which we discuss in the application evaluation section IV-E.

2) Evaluating reduced communication overhead with our collective computation framework:

We demonstrate the effectiveness of our collective computation framework in this section. From Figure 9, it is evident that the carefully overlapped version (Overlap) leads to significantly less Wait time compared with the previous version (ND), resulting in a performance boost of up to 4.9× for the data size of 678MB. The rationale for this performance improvement is shown in the following text. To utilize our collective computation

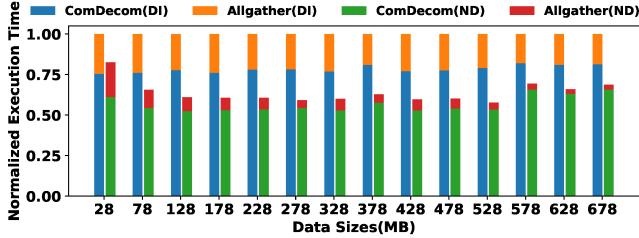


Fig. 8: Compare the performance of DI (Direct Integration) and our ND (Novel Design) from 28MB to 678MB.

framework for hiding communication during compression, we design and implement PIPE-SZx (pipelined SZx), which could break the compression process into small chunks and allow us to overlap the compression with communication in a fine-grained pipelined manner. As a result, we can significantly reduce MPI_Wait time in the reduce-scatter stage, which uses non-blocking communications.

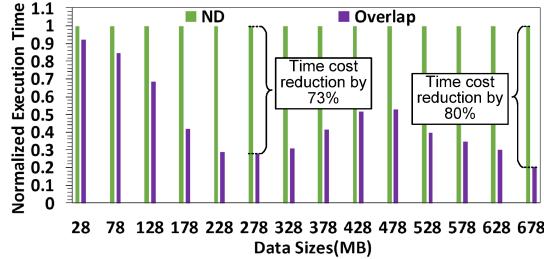


Fig. 9: Compare the wait time of ND (Novel Design) and Overlap (Overlap Optimization) from 28MB to 678MB with a 50MB step.

3) Evaluating overall running time of different optimizations: In Figure 10, we compare the end-to-end execution time of our compression-enabled Allreduce with step-wise optimizations against the Original MPI_Allreduce (AD). We can see that the end-to-end performance of our compression-enabled Allreduce gradually increase with these optimizations and it consistently beats the AD across various message sizes with our Overlapped Optimization (Overlap). We call the fully optimized variant as C-Allreduce and will thoroughly evaluate both the performance and accuracy of it in the following sections.

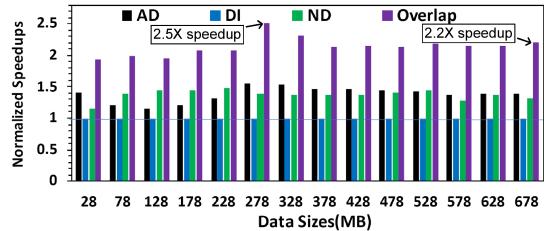


Fig. 10: Compare the overall running time of different optimizations from 28MB to 678MB with a 50MB step.

C. End-to-end Comparisons of C-Allreduce with Baselines

In this section, we compare the performance of our C-Allreduce with four different baselines on various data

sizes, node numbers, and datasets. The baselines include the original MPI_Allreduce without compression (Allreduce), MPI_Allreduce implemented by compression-enabled point-to-point communication (CPR-P2P) with fixed-rate ZFP (ZFP(FXR)), fixed-accuracy ZFP (ZFP(ABS)), and SZx. The SZx and ZFP(ABS) have a fixed accuracy of 1E-3, and the ZFP(FXR) has a rate of 4, if not specifically mentioned.

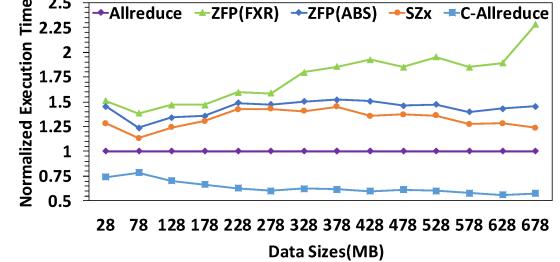


Fig. 11: Compare the performance of our C-Allreduce and multiple baselines from 28MB to 678MB with a 50MB step on 128 nodes.

1) Evaluating with different data sizes on 128 nodes: In this section, we present the performance of our C-Allreduce, along with related baselines, using data sizes ranging from 28MB to 678MB on 128 Broadwell nodes. The RTM dataset is adopted in these experiments. From Figure 11, we can observe that the SZx-integrated baseline performs the best among the three compression-integrated baselines, and ZFP(FXR) has the worst performance. This is because SZx has a faster compression speed compared to ZFP(FXR) and ZFP(ABS). Additionally, ZFP(ABS) and SZx have better compression ratios than ZFP(FXR) as we previously mentioned in III-C. However, none of the three CPR-P2P baselines can outperform the original Allreduce. By employing our novel framework and step-wise optimizations, our C-Allreduce achieves significantly higher performance than the original Allreduce with up to 1.8 \times performance improvement.

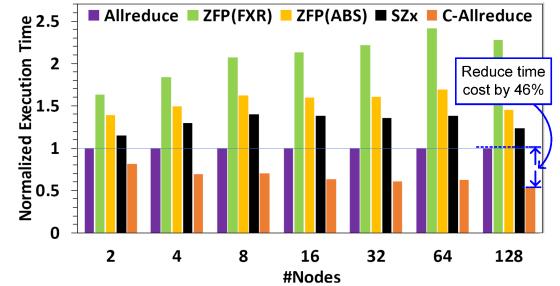


Fig. 12: Compare the performance of our C-Allreduce and multiple baselines from 2 to 128 nodes with 678MB message.

2) Evaluating with different node numbers with 678MB data: To demonstrate the scalability of our approach, we compare the normalized execution time of our C-Allreduce and four different baselines using a fixed data size of 678MB across 2 to 128 nodes. The RTM dataset is adopted in these experiments. As shown in Figure 12, our C-Allreduce outperforms all the baselines across various node numbers. It can reach a performance boost of up to 1.8 \times compared to

the original Allreduce. Similar to our observation in Section IV-C1, we found that all compressor-integrated baselines exhibit performance degradation compared to Allreduce, with the SZx-integrated baseline performing the fastest among them.

TABLE VI: Compression Ratios

Datasets	Hurricane		CESM-ATM
Fields	PRECIPf	QGRAUPf	CLOUDf
CPR	33.8	58.3	39.9

3) *Evaluating with different application datasets*: In this section, we evaluate the performance of our C-Allreduce and related baselines on different datasets using the error bound of 1E-4. Due to the page limitation, we only show the SZx baseline here as it has the best performance among compression-enabled baselines. The PRECIPf, QGRAUPf, and CLOUDf fields are from the Hurricane application dataset and the Q field is originated from the CESM-ATM application dataset. They have different data distributions and averaged compression ratios as shown in Table VI. From Figure 13, we can notice that our C-Allreduce has the best performance among all the implementations. Specifically, $1.74 \times$ and $1.58 \times$ speedups can be achieved in QGRAUPf and PRECIPf fields, respectively. Such consistent high performance of C-Allreduce across varied data fields highlights its efficiency and adaptability. On the contrary, the fastest compression-enabled baseline (SZx) still has performance degradation compared with the original Allreduce due to the limited performance of the CPR-P2P in all cases.

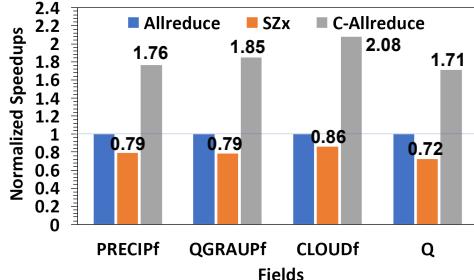


Fig. 13: Compare the performance of our C-Allreduce and multiple baselines in different application datasets.

Apart from the performance, we also evaluate the accuracy using visualization and numerical metrics including the widely-used peak signal-to-noise ratio (PSNR) [42] and normalized root mean squared error (NRMSE) [43]. Figure 14 represents the visual and numerical evaluation of our C-Allreduce with the Hurricane application dataset. The excellent image quality, along with the great PSNR(60.04) and NRMSE(1E-3) demonstrate that our C-Allreduce delivers a well-controlled accuracy. The same phenomenon can be witnessed in the Figure 15, where we evaluate the accuracy of C-Allreduce with the CESM-ATM application dataset. Thus, we can conclude that the selected error bound is suitable for other applications like Hurricane and CESM-ATM.

D. Generalizability Demonstration on Other MPI Collectives

We have demonstrated the high performance of our C-Allreduce, consisting of C-Allgather and C-Reduce-scatter. To show-

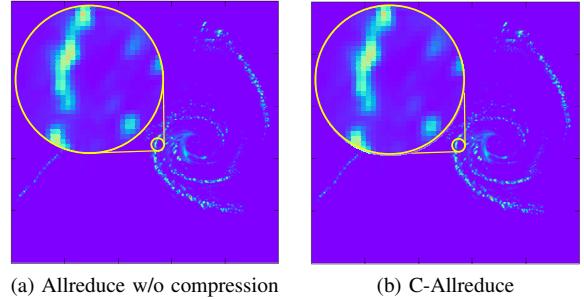


Fig. 14: The visualization and numerical evaluation of the accuracy of our C-Allreduce using Hurricane application dataset. Here the PSNR and NRMSE of the sub-figure(b) are 60.04 and 1E-3, respectively.

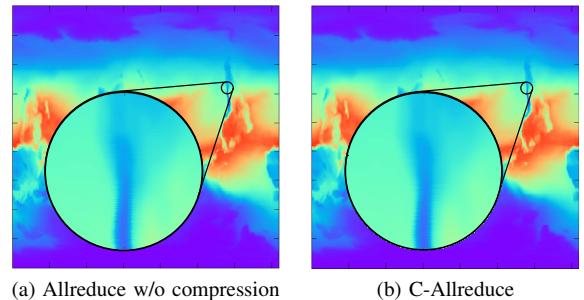


Fig. 15: The visualization and numerical evaluation of the accuracy of our C-Allreduce using CESM-ATM application dataset. Here the PSNR and NRMSE of the sub-figure(b) are 59.19 and 1E-3, respectively.

case the generalizability of our frameworks and optimizations, we also present C-Bcast and C-Scatter, which utilize the ubiquitous binomial tree algorithm adopted by MPICH. We conduct experiments on the RTM dataset ranging from 28MB to 678MB using 16 Broadwell nodes. In Figure 16, we present the speedups of our C-Scatter and C-Bcast, normalized against the original MPI_Scatter and MPI_Bcast (Baseline), respectively. We also compare our C-Scatter and C-Bcast with the SZx-integrated CPR-P2P baselines (SZx). Our C-Scatter is $1.8 \times$ faster than Baseline, while our C-Bcast has up to $2.7 \times$ speedup compared to Baseline. These performance improvements originated from the decreased data transfer volume and diminished compression overheads from our proposed frameworks. The speedups are even more significant compared to our C-Allreduce because collective data movement benefits more from our frameworks than collective computation. However, the SZx-integrated CPR-P2P baselines are much slower than the baseline due to significant compression overheads.

E. Evaluation of Image Stacking Performance and Accuracy

Image stacking is a well-known technique used in various scientific domains, including climate simulation and geology, to generate high-quality images by combining different images. For example, MPI is used by researchers to sum the individual images into final images [41]. In this experiment,

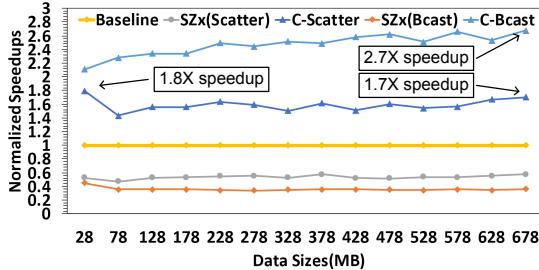


Fig. 16: Generalizability demonstration of our proposed framework and optimizations with MPI_Scatter and MPI_Bcast from 28MB to 678MB with a 50MB step.

we conduct image stacking of the RTM dataset on 16 nodes. As each snapshot has different value ranges, we use the fixed-accuracy (ABS) mode to compress the data so that each snapshot contributes a similar amount of errors rather than letting the snapshots with large value ranges dominate the errors. Three absolute error bounds (i.e., 1E-2, 1E-3, and 1E-4) are selected to demonstrate the flexibility between the accuracy and performance of our C-Allreduce. The same error bounds are used for the baseline SZx and ZFP(ABS). For ZFP(FXR), three fixed rates (4, 8, and 16) are selected and the related compression ratios are 8, 4, and 2.

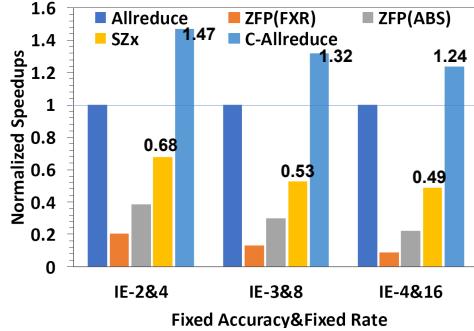


Fig. 17: Compare the image stacking performance of our C-Allreduce and multiple baselines in different error bounds for SZx and ZFP(ABS) and rates for ZFP(FXR).

We show the performance results and validate the high quality of the stacked images generated under our compression-integrated MPI collective framework, as shown in Figures 17 and 18, respectively. For our C-Allreduce, we could see that, with the increase of error bounds, the performance drops but the quality of the reconstructed image rises. The highest speedup can be witnessed in the 1E-2 case, where our C-Allreduce has 1.5 \times higher performance compared to the original Allreduce. Nevertheless, the three compression-integrated baselines (i.e., SZx, ZFP(ABS), and ZFP(FXR)) all result in performance degradation compared with the original Allreduce, regardless of the absolute error bounds or rates. With an error bound of 1E-4, our C-Allreduce shows an excellent reconstructed image quality, whereas the ZFP(ABS) integrated baseline with the same error bound cannot achieve the same quality. This is because that our proposed frameworks can significantly decrease the error propagation of the

reconstructed data, while the CPR-P2P cannot. Due to the page limitation, we do not show the reconstructed image of the SZx integrated baseline as it has even worse quality compared with ZFP(ABS). Besides, the ZFP(FXR) baseline with a rate of 4 has the worst reconstructed quality and the stacked image is completely different from the original one, as the fixed-rate mode cannot ensure a bounded accuracy.

In addition to the visualization evaluation, we complement our findings with numerical metrics to further demonstrate the exceptional accuracy of our approach. For an error bound of 1E-2, our C-Allreduce method achieves a Peak Signal-to-Noise Ratio (PSNR) [42] of 42.86 and a Normalized Root Mean Square Error (NRMSE) [43] of 7E-3, which illustrate a suboptimal data quality. However, when the error bound is tightened to 1E-3, we observe a notable increase in accuracy, with the PSNR soaring to 57.97 and the NRMSE dropping to 1E-3, which represent a great data quality. Consistent with the trends observed in image quality, the highest accuracy is attained at the 1E-4 error bound. At this threshold, our approach reaches a PSNR of 79.57 and maintains an NRMSE of 1E-4. These numerical metrics further validate the well-controlled accuracy of our method. In a nutshell, our C-Allreduce integrated with *C-Coll* framework can remarkably increase the performance of the original Allreduce and also preserves the quality of the original datasets very well at the error bound of 1E-3 and 1E-4 (see Figure 18 (e) and (f)).

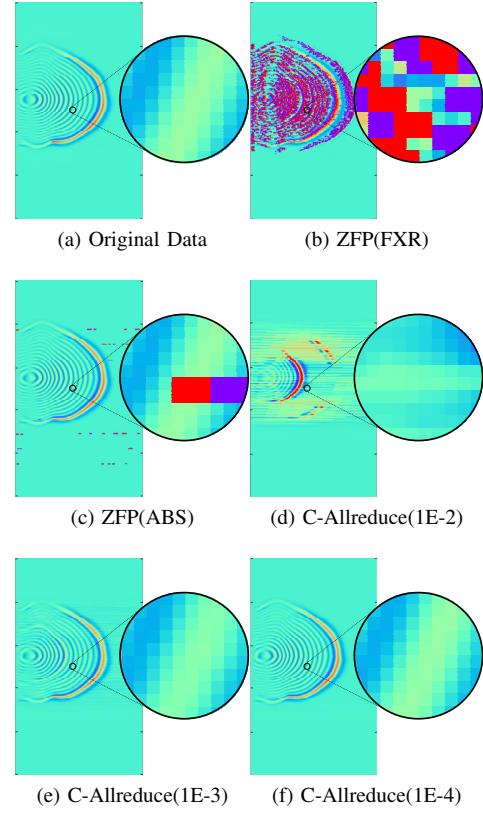


Fig. 18: Compare the reconstructed image qualities of our C-Allreduce with different ZFP-integrated baselines.

V. CONCLUSION AND FUTURE WORK

In this paper, we introduce *C-Coll*, a novel design for lossy-compression-integrated MPI collectives that significantly improves performance with bounded errors. Our two proposed high-performance frameworks for compression-integrated MPI collectives, together with customized pipe-lined SZx, enable us to implement C-Allreduce, which outperforms the original Allreduce by up to **2.1×** while preserving high data quality. We demonstrate the generalizability of our approaches through C-Scatter and C-Bcast, which outperform the original MPI_Scatter and MPI_Bcast by up to **1.8×** and **2.7×**, respectively. In summary, our research has addressed the issues of sub-optimal performance, lack of generality, and unbounded errors in lossy-compression-integrated MPI collectives, laying the foundation for future research in this area. Moving forward, we plan to expand our research by implementing more *C-Coll* based collectives and deploying our design on other hardware, such as GPUs and AI accelerators.

ACKNOWLEDGMENT

This research was supported by the Exascale Computing Project (ECP), Project Number: 17-SC-20-SC, a collaborative effort of two DOE organizations – the Office of Science and the National Nuclear Security Administration, responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering and early testbed platforms, to support the nation’s exascale computing imperative. The material was supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research (ASCR), under contract DE-AC02-06CH11357, and supported by the National Science Foundation under Grant OAC-2003709, OAC-2104023, and OAC-2311875. The experimental resource for this paper was provided by the Laboratory Computing Resource Center on the Bebop cluster at Argonne National Laboratory.

REFERENCES

- [1] A. A. Awan, K. Hamidouche, J. M. Hashmi, and D. K. Panda, “S-caffé: Co-designing mpi runtimes and caffe for scalable deep learning on modern gpu clusters,” in *Proceedings of the 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, 2017, pp. 193–205.
- [2] Y. Wang and M. Borland, “Pelegant: A parallel accelerator simulation code for electron generation and tracking,” in *AIP Conference Proceedings*, vol. 877, no. 1. American Institute of Physics, 2006, pp. 241–247.
- [3] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard et al., “Tensorflow: A system for large-scale machine learning,” in *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, 2016, pp. 265–283.
- [4] A. Ayala, S. Tomov, X. Luo, H. Shaeik, A. Haidar, G. Bosilca, and J. Dongarra, “Impacts of multi-gpu mpi collective communications on large fft computation,” in *2019 IEEE/ACM Workshop on Exascale MPI (ExaMPI)*. IEEE, 2019, pp. 12–18.
- [5] A. Jain, A. A. Awan, H. Subramoni, and D. K. Panda, “Scaling tensorflow, pytorch, and mxnet using mvapich2 for high-performance deep learning on frontera,” in *2019 IEEE/ACM Third Workshop on Deep Learning on Supercomputers (DLS)*. IEEE, 2019, pp. 76–83.
- [6] A. M. Abdelmoniem, A. Elzanaty, M.-S. Alouini, and M. Canini, “An efficient statistical-based gradient compression technique for distributed training systems,” 2021.
- [7] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2015.
- [8] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- [9] S. Chunduri, S. Parker, P. Balaji, K. Harms, and K. Kumaran, “Characterization of mpi usage on a production supercomputer,” in *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2018, pp. 386–400.
- [10] M. Bayatpour and M. A. Hashmi, “Salar: Scalable and adaptive designs for large message reduction collectives,” in *2018 IEEE International Conference on Cluster Computing (CLUSTER)*, 2018, pp. 1–10.
- [11] P. Patarasuk and X. Yuan, “Bandwidth optimal all-reduce algorithms for clusters of workstations,” *Journal of Parallel and Distributed Computing*, vol. 69, no. 2, pp. 117–124, 2009.
- [12] G. Almási, P. Heidelberger, C. J. Archer, X. Martorell, C. C. Erway, J. E. Moreira, B. Steinmacher-Burow, and Y. Zheng, “Optimization of mpi collective communication on bluegene/l systems,” in *Proceedings of the 19th Annual International Conference on Supercomputing*, ser. ICS ’05. New York, NY, USA: Association for Computing Machinery, 2005, p. 253–262. [Online]. Available: <https://doi.org/10.1145/1088149.1088183>
- [13] R. Thakur, R. Rabenseifner, and W. Gropp, “Optimization of collective communication operations in mpich,” *The International Journal of High Performance Computing Applications*, vol. 19, no. 1, pp. 49–66, 2005.
- [14] S. Di and F. Cappello, “Fast error-bounded lossy hpc data compression with sz,” in *2016 ieee international parallel and distributed processing symposium (ipdps)*. IEEE, 2016, pp. 730–739.
- [15] D. Tao, S. Di, and F. Cappello, “Significantly improving lossy compression for scientific data sets based on multidimensional prediction and error-controlled quantization,” 06 2017.
- [16] K. Zhao, S. Di, X. Liang, S. Li, D. Tao, Z. Chen, and F. Cappello, “Significantly improving lossy compression for hpc datasets with second-order prediction and parameter optimization,” in *Proceedings of the 29th International Symposium on High-Performance Parallel and Distributed Computing*, 2020, pp. 89–100.
- [17] P. Lindstrom, “Fixed-rate compressed floating-point arrays,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, pp. 2674–2683, 2014.
- [18] L. P. Deutsch, “Gzip file format specification version 4.3,” <https://datatracker.ietf.org/doc/draft-deutsch-gzip-spec/03/>, 1996.
- [19] J. loup Gailly and M. Adler, “zlib,” <https://www.zlib.net/>.
- [20] Y. Collet, “Zstandard – real-time data compression algorithm,” <http://facebook.github.io/zstd/>, 2015.
- [21] X.-C. Wu, S. Di, E. M. Dasgupta, F. Cappello, H. Finkel, Y. Alexeev, and F. T. Chong, “Full-state quantum circuit simulation by using data compression,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC ’19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: <https://doi.org/10.1145/3295500.3356155>
- [22] S. Li, S. Di, X. Liang, Z. Chen, and F. Cappello, “Optimizing lossy compression with adjacent snapshots for n-body simulation data,” in *2018 IEEE International Conference on Big Data (Big Data)*, 2018, pp. 428–437.
- [23] K. Zhao, S. Di, D. Perez, X. Liang, Z. Chen, and F. Cappello, “Mdz: An efficient error-bounded lossy compressor for molecular dynamics,” in *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, 2022, pp. 27–40.
- [24] A. M. Gok, S. Di, Y. Alexeev, D. Tao, V. Mironov, X. Liang, and F. Cappello, “Pastrí: Error-bounded lossy compression for two-electron integrals in quantum chemistry,” in *2018 IEEE International Conference on Cluster Computing (CLUSTER)*, 2018, pp. 1–11.
- [25] Q. Zhou, C. Chu, N. S. Kumar, P. Kousha, S. M. Ghazimirsaeed, H. Subramoni, and D. K. Panda, “Designing high-performance mpi libraries with on-the-fly compression for modern gpu clusters,” in *2021 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2021, pp. 444–453.
- [26] A. Yang, H. Mukka, F. Hesaaraki, and M. Burtscher, “Mpc: A massively parallel compression algorithm for scientific data,” 2015.
- [27] M. TEAM. (2020) MVAPICH2-X 2.3 User Guide. [Online]. Available: <https://mvapich.cse.ohio-state.edu/static/media/mvapich/mvapich2-x-userguide.pdf>
- [28] Q. Zhou, P. Kousha, Q. Anthony, K. Shafie Khorassani, A. Shafi, H. Subramoni, and D. K. Panda, “Accelerating mpi all-to-all communication with online compression on modern gpu clusters,”

in High Performance Computing: 37th International Conference, ISC High Performance 2022, Hamburg, Germany, May 29 – June 2, 2022, Proceedings. Berlin, Heidelberg: Springer-Verlag, 2022, p. 3–25. [Online]. Available: https://doi.org/10.1007/978-3-031-07312-0_1

[29] Q. Zhou, Q. Anthony, L. Xu, A. Shafi, M. Abduljabbar, H. Subramoni, and D. K. D. Panda, “Accelerating distributed deep learning training with compression assisted allgather and reduce-scatter communication,” in 2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2023, pp. 134–144.

[30] Q. Zhou, Q. Anthony, A. Shafi, H. Subramoni, and D. K. D. Panda, “Accelerating broadcast communication with gpu compression for deep learning workloads,” in 2022 IEEE 29th International Conference on High Performance Computing, Data, and Analytics (HiPC), 2022, pp. 22–31.

[31] R. Underwood, S. Di, J. C. Calhoun, and F. Cappello, “Fraz: A generic high-fidelity fixed-ratio lossy compression framework for scientific floating-point data,” in 2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS). Los Alamitos, CA, USA: IEEE Computer Society, may 2020, pp. 567–577. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/IPDPS47924.2020.00065>

[32] X. Liang, S. Di, D. Tao, S. Li, S. Li, H. Guo, and F. Cappello, “Error-controlled lossy compression optimized for high compression ratios of scientific datasets,” 12 2018, pp. 438–447.

[33] X. Yu, S. Di, K. Zhao, J. Tian, D. Tao, X. Liang, and F. Cappello, “Ultrafast error-bounded lossy compression for scientific datasets,” in Proceedings of the 31st International Symposium on High-Performance Parallel and Distributed Computing, ser. HPDC ’22. New York, NY, USA: Association for Computing Machinery, 2022, p. 159–171. [Online]. Available: <https://doi.org/10.1145/3502181.3531473>

[34] S. Kayum et al., “GeoDRIVE – a high performance computing flexible platform for seismic applications,” First Break, vol. 38, no. 2, pp. 97–100, 2020.

[35] Hurricane ISABEL simulation data, <http://vis.computer.org/vis2004contest/data.html>, 2004, online.

[36] J. E. Kay and et al., “The Community Earth System Model (CESM) large ensemble project: A community resource for studying climate change in the presence of internal climate variability,” Bulletin of the American Meteorological Society, vol. 96, no. 8, pp. 1333–1349, 2015.

[37] J. W. Hurrell, M. M. Holland, P. R. Gent, S. Ghan, J. E. Kay, P. J. Kushner, J.-F. Lamarque, W. G. Large, D. Lawrence, K. Lindsay et al., “The community earth system model: a framework for collaborative research,” Bulletin of the American Meteorological Society, vol. 94, no. 9, pp. 1339–1360, 2013.

[38] P. Lindstrom and M. Isenburg, “Fast and efficient compression of floating-point data,” IEEE Transactions on Visualization and Computer Graphics, vol. 12, no. 5, p. 1245–1250, sep 2006. [Online]. Available: <https://doi.org/10.1109/TVCG.2006.143>

[39] K. Zhao, S. Di, M. Dmitriev, T.-L. D. Tonello, Z. Chen, and F. Cappello, “Optimizing error-bounded lossy compression for scientific data by dynamic spline interpolation,” in 2021 IEEE 37th International Conference on Data Engineering (ICDE), 2021, pp. 1643–1654.

[40] K. Zhao, S. Di, X. Lian, S. Li, D. Tao, J. Bessac, Z. Chen, and F. Cappello, “SDRBench: Scientific data reduction benchmark for lossy compressors,” in 2020 IEEE International Conference on Big Data (Big Data), 2020, pp. 2716–2724.

[41] J. Gurhem, H. Calandra, and S. G. Petitot, “Parallel and distributed task-based kirchhoff seismic pre-stack depth migration application,” in 2021 20th International Symposium on Parallel and Distributed Computing (ISPDC), 2021, pp. 65–72.

[42] S. S., R. Ganesan, and R. Isaac, “Experimental analysis of compacted satellite image quality using different compression methods,” Advanced Science, vol. 7, 03 2015.

[43] M. V. Shcherbakov, A. Brebels, N. L. Shcherbakova, A. P. Tyukov, T. A. Janovsky, V. A. Kamaev et al., “A survey of forecast error measures,” World applied sciences journal, vol. 24, no. 24, pp. 171–176, 2013.