

Fixed-PSNR Lossy Compression for Scientific Data

Dingwen Tao,^{*} Sheng Di,[†] Xin Liang,[‡] Zizhong Chen,[‡] and Franck Cappello^{†§}

^{*}The University of Alabama, AL, USA

[†]Argonne National Laboratory, IL, USA

[‡]University of California, Riverside, CA, USA

[§]University of Illinois at Urbana-Champaign, IL, USA

tao@cs.ua.edu, sdi1@anl.gov, xliang007@ucr.edu,

chen@cs.ucr.edu, cappello@mcs.anl.gov

Abstract—Error-controlled lossy compression has been studied for years because of extremely large volumes of data being produced by today’s scientific simulations. None of existing lossy compressors, however, allow users to fix the peak signal-to-noise ratio (PSNR) during compression, although PSNR has been considered as one of the most significant indicators to assess compression quality. In this paper, we propose a novel technique providing a fixed-PSNR lossy compression for scientific data sets. We implement our proposed method based on the SZ lossy compression framework and release the code as an open-source toolkit. We evaluate our fixed-PSNR compressor on three real-world high-performance computing data sets. Experiments show that our solution has a high accuracy in controlling PSNR, with an average deviation of 0.1 ~ 5.0 dB on the tested data sets.

I. INTRODUCTION

Today’s high-performance computing (HPC) applications are producing extremely large amounts of data, such that data storage, transmission, and postanalysis are becoming a serious problem for scientific research. The problem comes from the limited storage capacity and I/O bandwidth of parallel file systems in production facilities. For example, the cosmology Hardware/Hybrid Accelerated Cosmology Code (HACC) [1] is a typical example of parallel executions facing this issue. According to the HACC researchers, the number of particles to simulate can be up to 3.5 trillion in today’s simulations (and even more in the future), which leads to 60 PB of data to store; yet a system such as the Mira supercomputer [4] has only 26 PB of file system storage. For now, HACC researchers have to control the data size by a temporal decimation (i.e., dumping the snapshots every k time steps in the simulation), degrading the consecutiveness of simulation in time dimension and losing important information unexpectedly. To this end, lossy compression has been studied for years [2], [3].

Peak signal-to-noise ratio (PSNR) is one of the the most critical indicators used to assess the distortion of reconstructed data versus original data in lossy compression. Unlike other evaluation metrics such as maximum compression error, PSNR measures the overall data distortion, which is closely related to the visual quality [4]. Most of the existing lossy data compressors (such as SZ [5] and ZFP [6]), however, are executed based on either a specific absolute error bound or a pointwise relative error bound. To adjust lossy compression for a particular data set to reach a target PSNR (or mean squared error (MSE)), users have to run the lossy compressor multiple times each with different error-bound settings, a tedious and

time-consuming task especially when masses of fields need to be compressed (e.g., CESM simulation involves 100+ fields [7]). In this paper, we explore a novel technique providing fixed-PSNR lossy compression for scientific users.

The contributions of this paper are as follows.

- We explore an in-depth analysis to precisely predict the overall data distortion (such as MSE and PSNR) based on state-of-the-art lossy compressors.
- We propose a novel fixed-PSNR method that can allow users to fix PSNR during the lossy compression based on our accurate prediction of PSNR.
- We implement our proposed fixed-PSNR method based on the SZ lossy compression framework and release the code as an open-source tool under a BSD license.
- We evaluate our implemented fixed-PSNR method using three real-world HPC data sets. Experimental results show that our proposed solution can precisely control the overall data distortion (i.e., PSNR) accurately during the compression with negligible overhead. In absolute terms, the average deviation between the user-specified PSNRs (before the compression) and the actual PSNRs (after the decompression) can be limited within 0.1 ~ 5.0 dB on the tested data sets.

The rest of the paper is organized as follows. In Section II, we discuss state-of-the-art lossy compression for scientific data and traditional error-control strategies. In Section III, we discuss how to estimate the overall data distortion for l^2 -norm-preserving lossy compression. In section IV, we formulate equations to estimate the l^2 -norm-based data distortion and propose our new fixed-PSNR error-control approach. In Section V, we present our experimental evaluation results. In section VI, we summarize our conclusions and discuss the future work.

II. BACKGROUND

In this section, we first discuss state-of-the-art lossy compression methods for scientific data and then describe the traditional error-control strategies of existing lossy compressors.

A. Scientific Data Compression

Compression techniques for HPC scientific data have been studied for years. The data compressors can be split into two categories: lossless and lossy. Lossless compressors make sure that the reconstructed data set after the decompression is

exactly the same as the original data set. Such a constraint may significantly limit the compression ratio (up to 2 in general [8]) on the compression of scientific data. The reason is that scientific data are composed mainly of floating-point values and their tailing mantissa bits could be too random to compress effectively. State-of-the-art lossy compressors include SZ [5], [9], ZFP [6], ISABELA [10], FPZIP [11], SSEM [12], and NUMARCK [13]. Basically, they can be categorized into two models: prediction based and transform based. A *prediction-based compressor* predicts data values for each data point and encodes the difference between every predicted value and its corresponding real value based on a quantization method. Typical examples are SZ [5], [9], ISABELA [10], and FPZIP [11]. A *transform-based compressor* transforms the original data to another space where most of the generated data is close to zero, such that the data can be stored with a certain loss in terms of user-required error bounds. For instance, SSEM [12] and ZFP [6] adopt a discrete Wavelet transform (DWT) and a customized orthogonal transform, respectively.

SZ [5], [9] is a state-of-the-art error-bounded lossy compressor designed for HPC scientific data. As discussed above, it falls into the category of prediction-based model and comprises *three* main steps: (1) predicting the value for each data point by its preceding neighbors in the multidimensional space; (2) performing a customized Huffman coding [14] to shrink the data size significantly; and (3) applying the GZIP lossless compressor [15] on the encoded bytes to further improve the compression ratio. The current version (v 1.4) of SZ [5] adopts the Lorenzo predictor [16] as the default prediction method in Step (1).

B. Traditional Error-Control Methods of Lossy Compression

Existing state-of-the-art error-bounded lossy compressors provide a series of different error-control strategies to satisfy the demands of users. For example, ISABELA [10] can guarantee the relative error of each reconstructed data point fall into a user-set pointwise relative error bound. ZFP [6] has three compression modes: fixed-accuracy mode, fixed-rate mode, and fixed-precision mode. Specifically, the fixed-accuracy mode can ensure the absolute error of each reconstructed data point less than a user-set absolute error bound. The fixed-rate mode can guarantee that the data will be compressed to a fixed number of bits. The fixed-precision mode can ensure the number of retained bits per value, which can lead to a control of pointwise relative error. SZ [5] offers three methods to control different types of errors: absolute error, pointwise relative error, and value-range-based relative error. Note that unlike the pointwise relative error that is compared with the value of each data point, value-range-based relative error is compared with the value range. To the best of our knowledge, this work is the first attempt to design an error-control approach that can precisely control the overall data distortion (such as MSE and PSNR) for lossy compression of scientific data.

III. L^2 -NORM-PRESERVING LOSSY COMPRESSION

In this section, we first introduce the prediction-based lossy compression. Then we infer that the pointwise compression

error (i.e., the difference between the original value of each data point and its decompressed value) is equal to the error introduced by quantization [17] or embedded coding (EC) [18]. We present two theorems to demonstrate that the overall distortion-based l^2 norm of the reconstructed data can stay the same as the distortion in the second step.

In the compression phase of the prediction-based lossy compression, the first step is to predict the value of each data point and calculate the prediction errors. This step will generate a set of prediction errors (denoted by X_{pe}) based on the original data set (denoted by X), data point by data point, during the compression. The second step is to quantize or encode these prediction errors. The third step is optional for further reducing the data size by entropy encoding.

During the decompression, one needs to reconstruct the prediction errors based on quantization method or EC and then reconstruct the overall data set. Note that the procedure of constructing the decompressed data set (denoted \tilde{X}), data point by data point, is based on the reconstructed prediction errors (denoted \tilde{X}_{pe}) during the decompression.

In what follows, we derive that the following equation must hold for prediction-based lossy compression.

$$X - \tilde{X} = X_{pe} - \tilde{X}_{pe} \quad (1)$$

During the compression, the prediction method generally predicts the value of each data point based on the data points nearby in space because of the potential high consecutiveness of the data set. The Lorenzo predictor [16], for example, approximates each data point by the values of its preceding adjacent data points. Since the neighboring data values to be used to reconstruct each data point during the decompression are actually the decompressed values instead of the original values, in practice, one has to assure that the compression and decompression stage have exactly the same prediction procedure (including the data values used in the prediction method); otherwise, the data loss will be propagated during the decompression. Hence, the predicted values during the compression must be equal to the predicted values during the decompression. That is, we have $X_{pred} = \tilde{X}_{pred}$. Then, we can derive Equation (1) based on the following two equations: $X_{pe} = X - X_{pred}$ and $\tilde{X} = \tilde{X}_{pe} + \tilde{X}_{pred}$.

Based on Equation (1), we can easily derive the following theorem.

Theorem 1. *For prediction-based lossy compression, the overall l^2 -norm-based data distortion is the same as the distortion (introduced in the second step) of the prediction error.*

Theorem 1 can be extended to the transform-based lossy compression methods with orthogonal transforms (such as DWT), as shown below. Because of space limitations, we omit the details of the proof.

Theorem 2. *For orthogonal-transform-based lossy compression, the overall l^2 -norm-based data distortion is the same as the distortion (introduced in the second step) of the transformed data.*

This theorems indicates that *the overall distortion of the*

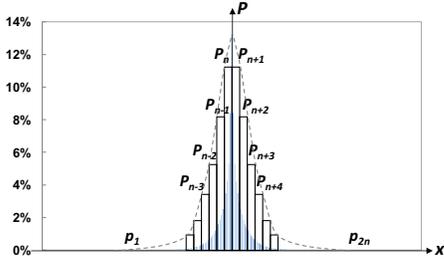


Fig. 1. Example of the distribution and uniform quantization of the prediction errors generated by the SZ lossy compressor on one ATM data field.

decompressed data can be estimated by the data distortion introduced in the second step for l^2 -norm-preserving lossy compression.

IV. DESIGN OF FIXED-PSNR LOSSY COMPRESSION

In this section, we propose our design of fixed-PSNR error-control method for lossy compression of scientific data.

As shown in Theorem 1 and 2, the l^2 -norm-based error, such as MSE, introduced by the second step of prediction-based and orthogonal-transform-based lossy compression stays unchanged after decompression. Therefore, we can estimate the l^2 -norm based compression error by estimating the error of the second step. For the following discussion, we focus mainly on *quantization* rather than EC for the second step.

Quantization converts the prediction errors or transformed data (denoted by X^t) to another set of integer values that are easier to compress. In quantization, the value range is split into multiple bins. Then, the compressor goes through all the prediction errors or transformed data to determine in which bins they are located, and it represents their values by the corresponding bin indexes, which are integer values. During the decompression, the midpoint of each quantization bin will be used to reconstruct the data located in the bin; it is called the quantized value in the following discussion. The effectiveness of the data reduction in quantization depends on the distribution of the prediction errors or transformed data.

We assume $P(x)$ to be the probability density function of X^t , that is, $X^t \sim P(x)$. The probability distribution of X^t is symmetric in a large majority of cases, based on our experience. The blue area in Figure 1 shows the typical probability distribution of the prediction errors generated by the SZ lossy compressor using a data field from a production climate simulation (called CSEM-ATM). Therefore, we assume $P(x)$ to be symmetric without loss of generality, and we focus only on n bins on one side. The number of quantization bins is represented by $2n$. We denote δ_i the length of the i th quantization bin, where $\delta_i = \delta_{(2n+1)-i}$ ($1 \leq i \leq n$) because of the symmetry property.

We denote \tilde{X}^t as the quantized values of X^t . The MSE between X^t and \tilde{X}^t can be calculated by

$$\begin{aligned} MSE(X^t, \tilde{X}^t) &= E_{X^t}[(X^t - \tilde{X}^t)^2] \\ &= 2 \int_0^{+\infty} (x - \tilde{x})^2 \cdot P(x) dx, \end{aligned} \quad (2)$$

where $E[\cdot]$ represents the expectation. Note that \tilde{x} is a step function, since the values in each bin are quantized to the same value. Lossy compressors such as NUMARACK [13], SSEM

[12], and SZ [5] often use the midpoint of the quantization bin to approximate the values located in it. Therefore, $\tilde{x} = \frac{s_i + s_{i+1}}{2} = m_i$ when $s_i \leq x < s_{i+1}$. We can further estimate the MSE based on the probability density function $P(x)$ and the step function \tilde{x} as follows.

$$\begin{aligned} MSE &= 2 \sum_{i=1}^n \int_{s_i}^{s_{i+1}} (x - \hat{x})^2 \cdot P(x) dx \\ &\approx 2 \sum_{i=1}^n (P(m_i) \cdot \int_{s_i}^{s_{i+1}} (x - m_i)^2 dx) \\ &= 2 \sum_{i=1}^n (P(m_i) \cdot \int_0^{\delta_i} (x - \frac{\delta_i}{2})^2 dx) = \frac{1}{6} \sum_{i=1}^n \delta_i^3 P(m_i) \end{aligned} \quad (3)$$

We then can calculate the normalized root mean squared error (NRMSE) and peak signal-to-noise ratio (PSNR) as follows:

$$NRMSE = \frac{\sqrt{MSE}}{vr} = \left(\sum_{i=1}^n \delta_i^3 P(m_i) \right)^{\frac{1}{2}} / \sqrt{6} vr \quad (4)$$

$$\begin{aligned} PSNR &= -20 \cdot \log_{10}(NRMSE) \\ &= -10 \cdot (\log_{10} \left(\sum_{i=1}^n \delta_i^3 P(m_i) \right) + 2 \cdot \log_{10} vr + \log_{10} 6), \end{aligned} \quad (5)$$

where vr represents the value range of the original data X . So far, we have established the estimation equation for l^2 -norm-based compression error with general quantization.

Uniform quantization is the most straightforward yet effective quantization approach, which is adopted by SZ lossy compressor. With this approach, all quantization bins have the same length, (i.e., $\delta_1 = \dots = \delta_{2n} = \delta$). On the other hand, the $2n$ quantization bin can cover all the prediction errors as long as the number of bins is large enough; hence, $\sum_{i=1}^n P(m_i) \approx P(x \geq 0) / \delta = 1/2\delta$. Equations (5) now can be further simplified as follows:

$$PSNR = 20 \cdot \log_{10}(vr/\delta) + 10 \cdot \log_{10} 12. \quad (6)$$

Equation (6) says that *the PSNR depends only on the unified quantization bin size regardless of the distribution of X^t* . For example, the SZ lossy compressor sets the bin size δ to twice the absolute error bound (i.e., eb_{abs}) to ensure the maximum pointwise compression error within eb_{abs} . Therefore, based on Equation (6), our PSNR estimation for SZ lossy compressor becomes

$$PSNR = 20 \cdot \log_{10}(vr/eb_{abs}) + 10 \cdot \log_{10} 3. \quad (7)$$

Note that eb_{abs}/vr is the value-range-based relative error bound (denoted by eb_{rel}) defined by SZ. Thus we can estimate SZ's PSNR precisely based on the user-set eb_{rel} .

Similar to the definitions of fixed-rate and fixed-accuracy, we can design a fixed-PSNR mode for the SZ lossy compression framework. The term "fixed" here means that the compression can be performed based on a specific given target (such as bit-rate or accuracy). More generally, the following theorem can be drawn from our analysis.

Theorem 3. *The prediction-based and orthogonal-transform-based lossy compression methods with uniform quantization in the second step are fixed-PSNR regardless of the prediction errors or transformed data. In other words, the l^2 -norm-based compression error can be estimated precisely by the unified*

TABLE I
DATA SETS USED IN EXPERIMENTAL EVALUATION

	Dim. Size	# of Fields	Data Size	Example Fields
NYX	2048 × 2048 × 2048	6	206 GB	baryon_density, temperature
ATM	1800 × 3600	79	1.5 TB	CLDHGH, CLDL0W
Hurricane	100 × 500 × 500	13	62.4 GB	QICE, PRECIP, U, V, W

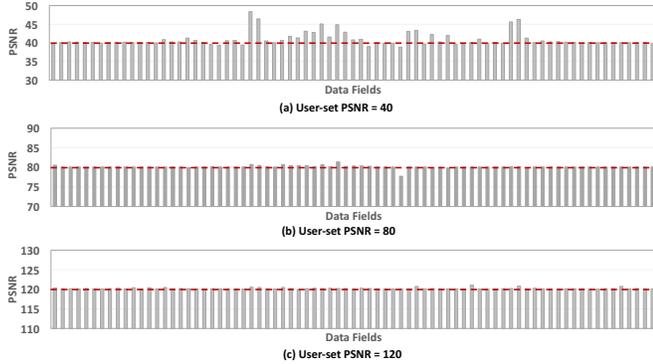


Fig. 2. Evaluation of proposed fixed-PSNR mode on all data fields in ATM.

quantization bin size and the value range of original data.

We can further simplify Equation (7) to the following equation.

$$eb_{rel} = \sqrt{3} \cdot 10^{-\frac{PSNR}{20}} \quad (8)$$

Based on Equation (8), we propose our fixed-PSNR error-control approach based on the SZ compression framework. It consists of the following three steps: (1) get user-required PSNR, (2) estimate value-range-based relative error bound based on Equation 7, and (3) perform original SZ lossy compression with the estimated eb_{rel} . Compared with the original compression workflow, the only computational overhead of our approach is the time to calculate the value-range-based relative error bound based on Equation (8) for each data field, which is negligible.

V. EXPERIMENTAL EVALUATION

In this section, we evaluate our proposed fixed-PSNR lossy compression on three real-world HPC simulation data sets.

We conduct our evaluation using the Bebop cluster [19] at Argonne National Laboratory. Each node is equipped with two Intel Xeon E5-2695 v4 processors and 128 GB of memory. We perform our evaluations on various single floating-point HPC application data sets including 2D CESM-ATM data sets from climate simulation [7], 3D Hurricane data sets from the simulation of the hurricane ISABELA [20], and 3D NYX data sets from cosmology simulation [21]. Each application involves many simulation snapshots (or time steps). The details of the data sets are described in Table I.

Figure 2 shows the evaluation results of our proposed fixed-PSNR lossy compression on all the fields in the ATM data sets with different user-set PSNRs. The three PSNR values—40 dB, 80 dB, and 120 dB—typically represent the low, medium, and high compression-quality scenarios, respectively. We plot the red dash lines showing the user-set PSNRs for ease comparison. Our fixed-PSNR lossy compressor can precisely

TABLE II
EVALUATION OF PROPOSED FIXED-PSNR MODE WITH SZ ON NYX, ATM, AND HURRICANE DATA SETS

User-set PSNR (dB)	NYX		ATM		Hurricane	
	AVG	STDEV	AVG	STDEV	AVG	STDEV
20	24.3	1.82	21.9	3.34	25.0	6.52
40	41.9	2.32	40.9	1.80	42.0	3.97
60	60.7	0.74	60.2	0.62	60.5	0.74
80	80.1	0.05	80.1	0.35	80.1	0.32
100	100.1	0.07	100.2	0.17	100.1	0.39
120	120.1	0.01	120.2	0.19	120.3	0.63

control the PSNRs and meet the PSNR’s demands for more than 90+% ATM data fields on the average. Note that here “meet” means that the actual PSNR after the decompression is equal or higher than the user-set PSNR before the compression.

Table II presents the experimental results of our proposed fixed-PSNR lossy compression on the NYX, ATM, and Hurricane data sets. The *AVG* and *STDEV* in the table represent the average value and standard deviation of the actual PSNRs of all the data fields in the set after the decompression. The table illustrates that our solution provides a very high accuracy in controlling PSNR on these HPC data sets. Specifically, in absolute terms, the deviation between the user-set PSNRs and the actual PSNRs can be limited within 0.1 ~ 5.0 dB on average. We note that Figure 2 and Table II both illustrate that the higher the PSNR of demand, the better our fixed-PSNR method performs. The reason can be that the accuracy of our estimation of the MSE (as shown in Equation 3) will decrease as the size of the quantization bin increases (i.e., higher error bound in the uniform quantization).

VI. CONCLUSION AND FUTURE WORK

In this paper, we propose a novel approach that can perform the lossy compression in terms of the user’s specific demand on overall distortion of data (PSNR). To the best of our knowledge, this is the first attempt to control the overall data distortion, which provides practical support to help efficiently tune the lossy compression to a target compression quality. Our experiments with three well-known real-world data sets show that the deviation between the user-specified PSNR before the compression and the actual PSNRs after the lossy compression can be limited within 0.1 ~ 5.0 dB on the average. In future work, we will explore more techniques to further improve the fixed-PSNR lossy compression, especially for the low compression-quality demands.

ACKNOWLEDGE

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, under contract DE-AC02-06CH11357, and supported by the National Science Foundation under Grant No. 1619253. We gratefully acknowledge the computing resources provided on Bebop, a high-performance computing cluster operated by the Laboratory Computing Resource Center at Argonne National Laboratory.

REFERENCES

- [1] S. Habib, V. Morozov, N. Frontiere, H. Finkel, A. Pope, K. Heitmann, K. Kumaran, V. Vishwanath, T. Peterka, J. Insley *et al.*, “Hacc: extreme scaling and performance across diverse architectures,” *Communications of the ACM*, vol. 60, no. 1, pp. 97–104, 2016.
- [2] D. Tao, S. Di, Z. Chen, and F. Cappello, “Exploration of pattern-matching techniques for lossy compression on cosmology simulation data sets,” in *High Performance Computing - ISC High Performance 2017 International Workshops, DRBSD, ExaComm, HCPM, HPC-IODC, IWOPH, IXPUG, P³MA, VHPC, Visualization at Scale, WOPSSS*, 2017, pp. 43–54.
- [3] D. Tao, S. Di, Z. Chen, and F. Cappello, “In-depth exploration of single-snapshot lossy compression techniques for n-body simulations,” in *2017 IEEE International Conference on Big Data*. IEEE, 2017, pp. 486–493.
- [4] S. Guthe and W. Straßer, “Real-time decompression and visualization of animated volume data,” in *Visualization, 2001. VIS’01. Proceedings*. IEEE, 2001, pp. 349–572.
- [5] 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*. IEEE, 2017, pp. 1129–1139.
- [6] P. Lindstrom, “Fixed-rate compressed floating-point arrays,” *IEEE transactions on visualization and computer graphics*, vol. 20, no. 12, pp. 2674–2683, 2014.
- [7] J. Kay, C. Deser, A. Phillips, A. Mai, C. Hannay, G. Strand, J. Arblaster, S. Bates, G. Danabasoglu, J. Edwards *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.
- [8] P. Ratanaworabhan, J. Ke, and M. Burtcher, “Fast lossless compression of scientific floating-point data,” in *Data Compression Conference, 2006. DCC 2006. Proceedings*, 2006, pp. 133–142.
- [9] S. Di and F. Cappello, “Fast error-bounded lossy hpc data compression with sz,” in *2016 IEEE International Parallel and Distributed Processing Symposium*. IEEE, 2016, pp. 730–739.
- [10] S. Lakshminarasimhan, N. Shah, S. Ethier, S.-H. Ku, C.-S. Chang, S. Klasky, R. Latham, R. Ross, and N. F. Samatova, “Isabela for effective in situ compression of scientific data,” *Concurrency and Computation: Practice and Experience*, vol. 25, no. 4, pp. 524–540, 2013.
- [11] P. Lindstrom and M. Isenburg, “Fast and efficient compression of floating-point data,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 1245–1250, 2006.
- [12] N. Sasaki, K. Sato, T. Endo, and S. Matsuoka, “Exploration of lossy compression for application-level checkpoint/restart,” in *Parallel and Distributed Processing Symposium (IPDPS), 2015 IEEE International*, 2015, pp. 914–922.
- [13] Z. Chen, S. W. Son, W. Hendrix, A. Agrawal, W.-k. Liao, and A. Choudhary, “Numarck: machine learning algorithm for resiliency and checkpointing,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Press, 2014, pp. 733–744.
- [14] D. A. Huffman, “A method for the construction of minimum-redundancy codes,” *Proceedings of the IRE*, vol. 40, no. 9, pp. 1098–1101, 1952.
- [15] L. P. Deutsch, “Gzip file format specification version 4.3,” 1996.
- [16] L. Ibarria, P. Lindstrom, J. Rossignac, and A. Szymczak, “Out-of-core compression and decompression of large n-dimensional scalar fields,” in *Computer Graphics Forum*, vol. 22, no. 3. Wiley Online Library, 2003, pp. 343–348.
- [17] A. Gersho and R. M. Gray, *Vector quantization and signal compression*. Springer Science & Business Media, 2012, vol. 159.
- [18] J. M. Shapiro, “Embedded image coding using zerotrees of wavelet coefficients,” *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3445–3462, 1993.
- [19] LCRC Bebop cluster, <https://www.lcrc.anl.gov/systems/resources/bebop>, online.
- [20] Hurricane ISABEL simulation data, <http://vis.computer.org/vis2004contest/data.html>, online.
- [21] NYX simulation, <https://amrex-astro.github.io/Nyx>, online.