OAK RIDGE National Laboratory

Understanding Performance-Quality Trade-offs in Scientific Visualization Workflows with Lossy Compression

Jieyang Chen, David Pugmire, Matthew Wolf, Nicholas Thompson, Jeremy Logan， Kshitij Mehta, Lipeng Wan, Jong Youl Choi, Ben Whitney, Scott Klasky

Scientific Data Group

Computer Science and Mathematics Division

Oak Ridge National Laboratory

ORNL is managed by UT-Battelle, LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

1

Motivation: The widening compute-data gap

- I/O bandwidth falls behind computation power
  - Needs fewer bytes/operation to sustain same computing efficiency
- Scientific visualization becomes I/O bound
  - Larger fraction of time is spent in I/O as computing currency and/or data size increase
- Current solutions
  - In situ methods: avoid doing I/O

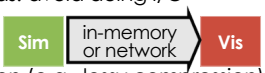

  - Data reduction (e.g., lossy compression):
    - Reduce data size by discard less important information
    - Substitute computing cycle for I/O

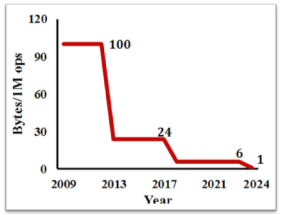

fewer bytes/operation

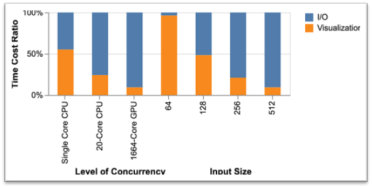

I/O cost in visualization

OAK RIDGE National Laboratory

2

2

Motivation

- Challenges with current solutions: a priori knowledge
  - In situ processing only has access to data while running
  - Reduced data has a pre-set accuracy
- Scientist generally don't know all operations they want to perform before running
  - pre-set accuracy can be problematic for unanticipated post-processing
- Goal in this work: Understand the impact of visualization on reduced data
  - How much I/O reduction can be achieved?
  - How would it impact visualization quality?
  - What is the impact on features?

OAK RIDGE National Laboratory

3

3

Background: popular lossy compression software

- MGARD[1]
  - Based on multigrid method
  - Offers a high degree of flexibility and guaranteed provision on loss incurred by the reduction
  - Allows preservation of derived quantities
- SZ[2]
  - Based on Lorenzo predictor and linear regression
  - Offer three modes of error bounds: ABS, REL, and PWR
- ZFP[3]
  - Based block discrete cosine transform
  - Supports fine-grained read and write access

1: M. Ainsworth et al. Multilevel Techniques For Compression And Reduction Of Scientific Data-quantitative Control Of Accuracy In Derived Quantities

2: X. Liang et al. Error-Controlled Lossy Compression Optimized for High Compression Ratios of Scientific Datasets

3: P. Lindstrom et al Fixed-Rate Compressed Floating-Point Arrays

OAK RIDGE National Laboratory

4

4

Designing visualization workflow for evaluation

$$\frac{\partial U}{\partial t} = r_u \nabla^2 U - UV^2 + f(1-U)$$
$$\frac{\partial V}{\partial t} = r_v \nabla^2 V + UV^2 - (f+k)V$$

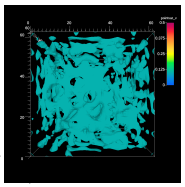

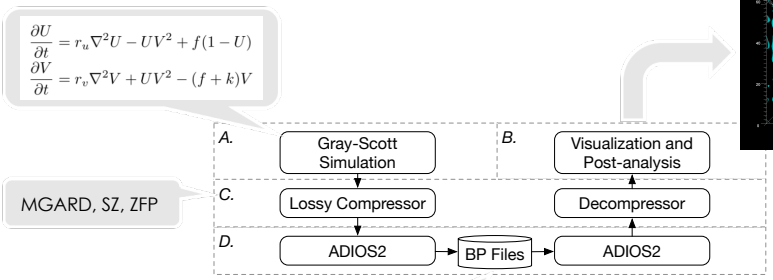

MGARD, SZ, ZFP

- BP self-describing file format
- publisher-subscriber framework for data exchange

5

Evaluation setup

- Simulation configuration
  - Largest possible grid size: 512^3
  - Total 10K iterations and dumping data for every 1k iterations
- Visualization
  - Using Marching Cubes to generate iso-surface (value = 0.1) based on variable V. (On both VTK-m and VisIt)
- Compressor error bound
  - From almost 0% to at least 30% relative error occurs

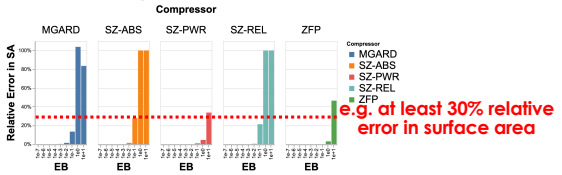

6

Evaluation metrics

	Metrics	Description/Definition	Tool Used
Performance	Compression Ratio	Original size/Compressed size	Our post-analyzer
	Read I/O cost	Including time cost for loading data visualization	
Quality	PSNR	$10\log_{10}(\frac{Max\ fluctuation^2}{Mean\ Square\ Error})$	Z-Checker
	Relative L^∞ errpr	$\frac{Max(absolute\ error)}{Max\ (input)}$	Our post-analyzer
	Relative error of iso-surface area (**SA**)	Visualization-based features	VisIt
	Relative error of number of connected components (**NCC**)	Visualization-based features	

7

Evaluation platform

- Software
  - **Compilers**: GCC v7.4.0
  - **Simulation code**: Gray-Scott*
  - **Data management**: ADIOS2
  - **Visualization**: VTK-m v1.4 and VisIt v3.0.1
  - **Lossy compressors**: MGARD v0.0.0.2, SZ v2.1.6, ZFP v0.5.5
- Hardware: a workstation
  - **CPU**: Intel 20 core Xeon
  - **Memory**: 32GB
  - **GPU**: Nvidia Quadro 4000M * 2

*https://github.com/pnorbert/adiosvm/tree/master/Tutorial/gray-scott

8

Evaluation goal

How mush of information/feature can be preserved with given CR?

CR

I/O Cost

General Error Metrics

PSNR

Rel. L^∞ Error

How would general error metrics reflect feature errors in visualization?

Rel. Error SA

Rel. Error NCC

Error in Vis. Features

How would the I/O cost-quality tradeoff be like?

SA: Surface Area
NCC: # of connected components

OAK RIDGE National Laboratory

9

9

Results – Compression ratio vs. Quality

SA: Surface Area
NCC: # of connected components

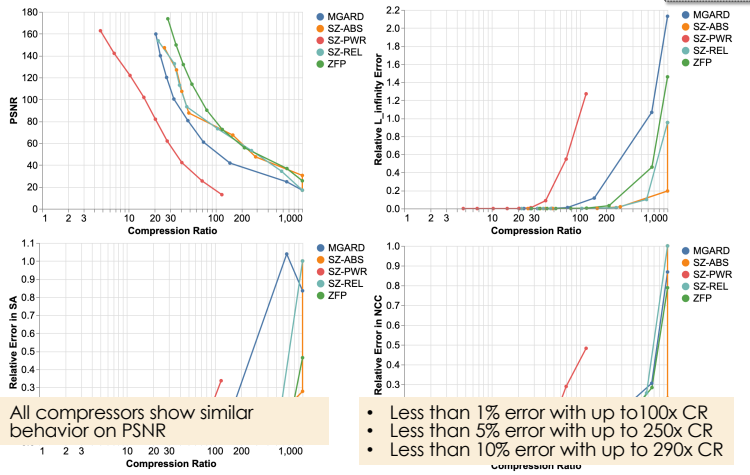

- All compressors show similar behavior on PSNR
- Less than 1% error with up to100x CR
- Less than 5% error with up to 250x CR
- Less than 10% error with up to 290x CR

OAK RIDGE National Laboratory

10

10

General Error Metrics vs. Feature Error

SA: Surface Area
NCC: # of connected components

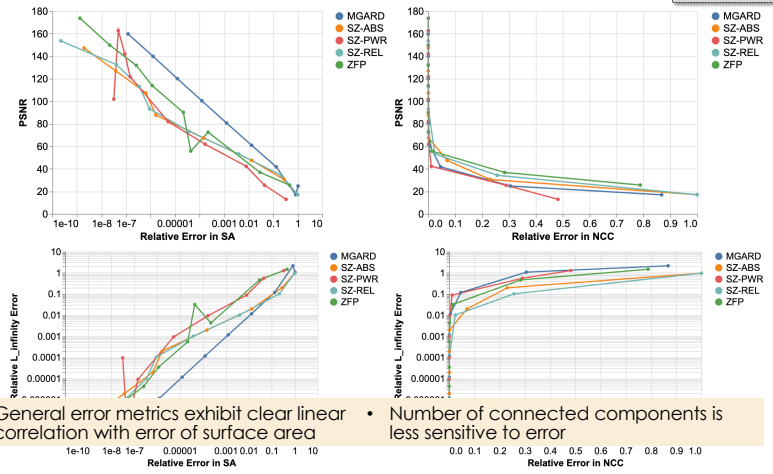

- General error metrics exhibit clear linear correlation with error of surface area
- Number of connected components is less sensitive to error

OAK RIDGE National Laboratory

11

11

Results – Quality vs. File read reduction

SA: Surface Area
NCC: # of connected components

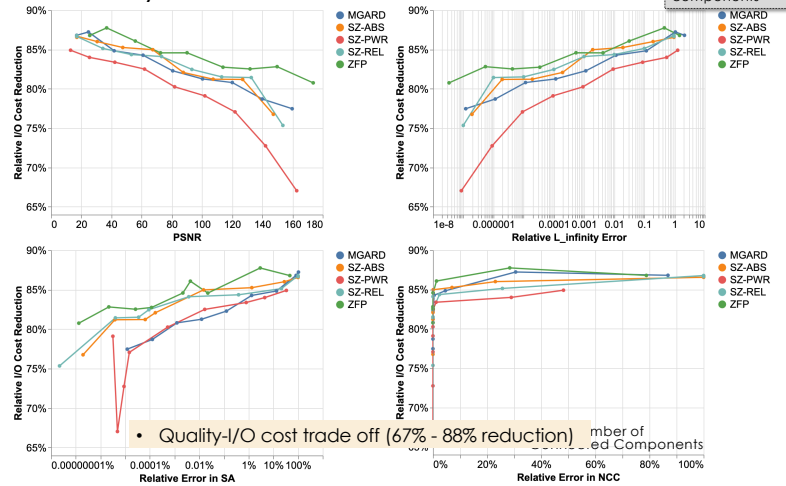

- Quality-I/O cost trade off (67% - 88% reduction)

OAK RIDGE National Laboratory

12

12

Evaluation summary

- About current lossy compression software
  - All tested lossy compression software provide decent compression ratio with adjustable parameters that can help preserve visualization quality
  - However, great care must be taken for parameters.
- About quantitative metrics
  - General error metrics provide good sensitivity to show the impact on some visualized features.
  - Number of connected component is less sensitive to information loss compared with surface area

Conclusion & Future works

- What we have learned?
  - Lossy compression can significantly reduce I/O cost in visualization tasks
  - Paving the road for studying impact on other data sharing methods (e.g. in situ method, staging, etc.)
  - Tuning the compressor parameters help us get an initial insight into the impact brings by lossy compression
- Future works
  - Extending to large scale: larger simulation runs with higher level of concurrency
  - Studying the impacts on preserving more complex features
  - Seeking different optimization strategies to reduce compression overhead

Thanks!

Q&A

Backup slides

Results – Storage space cost (compression ratio)

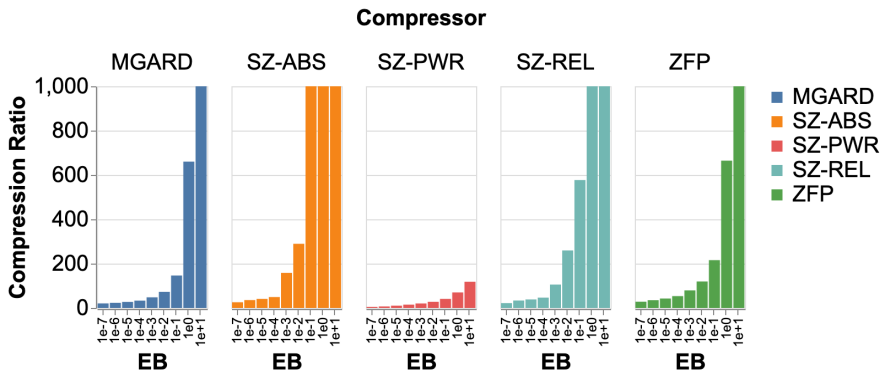

OAK RIDGE National Laboratory

Results – Error bound vs. Quality

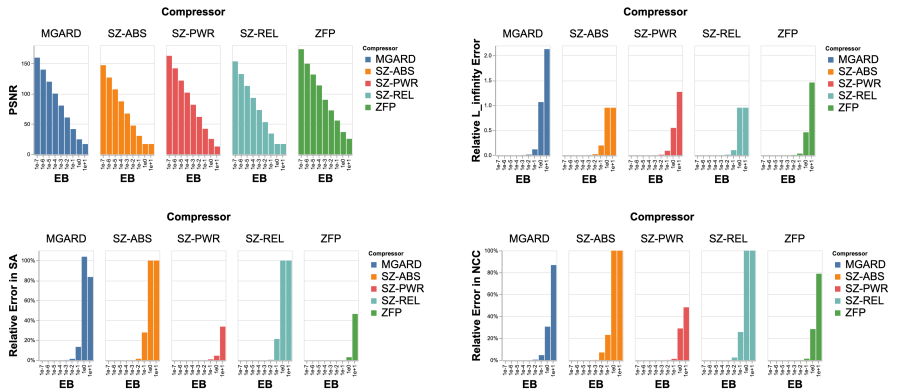

OAK RIDGE National Laboratory