Scaling HPC Applications Through Predictable and Reliable Data Reduction Methods

Sian Jin, Dingwen Tao (advisor) {sianjin, ditao}@iu.edu

INDIANA UNIVERSITY
BLOOMINGTON

Motivation

- **Large amount of data** from:
  - Large-scale scientific applications;
  - Large-scale machine learning (ML) applications.
- **High storage footprint and I/O bottleneck**:
  - Storage & computing power: not growing at the same pace.

Challenges And Goal

Adaptively balance between **compression ratio**, **compression throughput**, and **data quality** to **efficiently** utilize data reduction across a wide range of **HPC applications**.

Key Contributions

Deeply integrate the data reduction methods to applications to improve the overall performance while maintain the same level of end-result quality:

- **Compression ratio modeling**: predict compression ratio before compression happens.
- **Data quality modeling**: characterize the error introduced by lossy compressors and provide theoretical analysis for application-specific end-result quality.
- **In-situ compression**: build specifically designed frameworks to fully utilize the data reduction techniques.
- **I/O libraries**: accelerating the performance for general scientific data libraries, such as HDF5.

Ratio-Quality Model

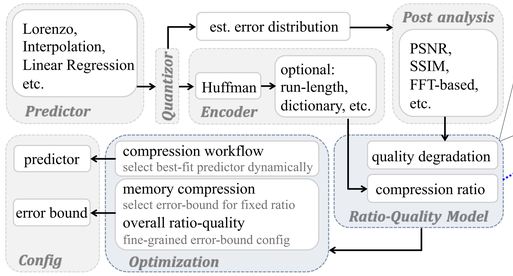

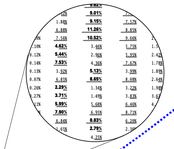

Evaluated on wide range of datasets

← Fig 1 An overview of ratio-quality modeling workflow for prediction-based lossy compression.

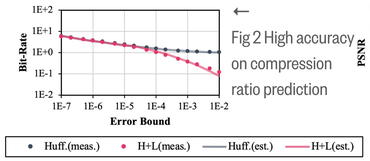

← Fig 2 High accuracy on compression ratio prediction

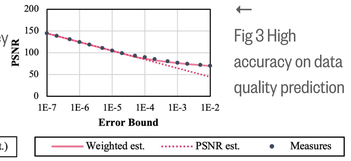

← Fig 3 High accuracy on data quality prediction

In-Situ Data Reduction for Scientific Applications

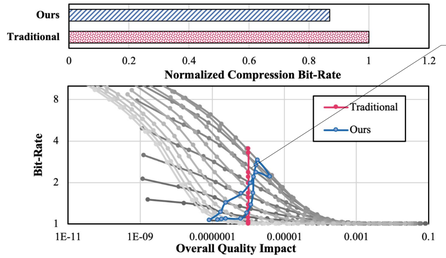

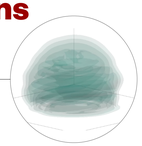

← Fig 4 RTM: In-situ data reduction for RTM dataset with multiple time steps in consideration for post-hoc analysis.

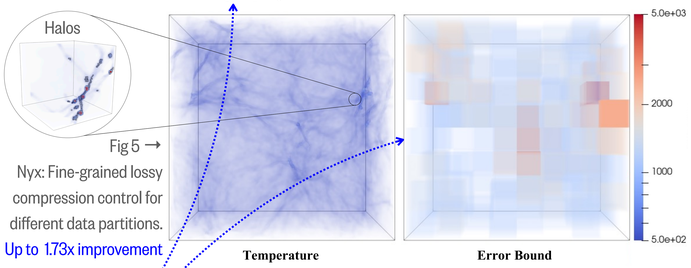

Fig 5 → Nyx: Fine-grained lossy compression control for different data partitions. Up to 1.73x improvement

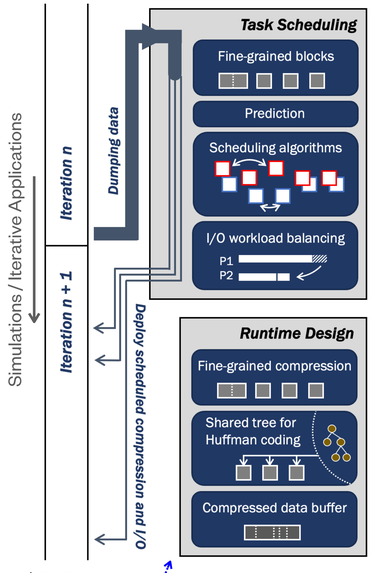

↑ Fig 9 An overview of Concealing Compression-Accelerated I/O for HPC Applications through In Situ Task Scheduling

Advanced Parallel Write With I/O Libraries

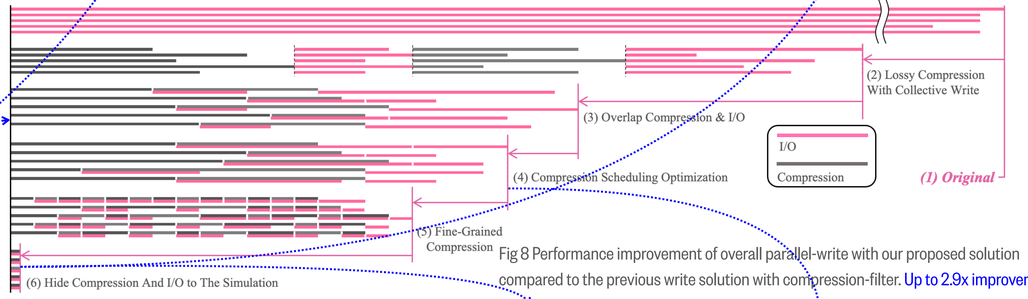

↑ Fig 6 Timeline of data aggregation with 5 processes and 2 data fields. Our solution overlaps compression with I/O with our general-purpose mechanism that fits most applications.

Fig 8 Performance improvement of overall parallel-write with our proposed solution compared to the previous write solution with compression-filter. Up to 2.9x improvement. ↓

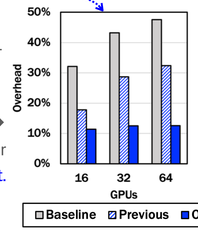

Fig 7 → Time overheads between the baseline, asynchronous I/O, and our solution across different scale. Up to 2.57x additional improvement.

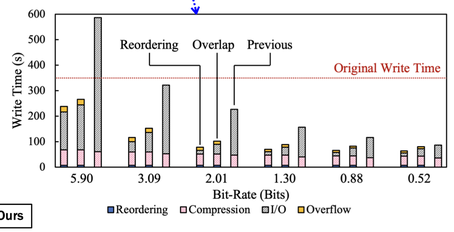

In-Situ Data Reduction for Machine Learning

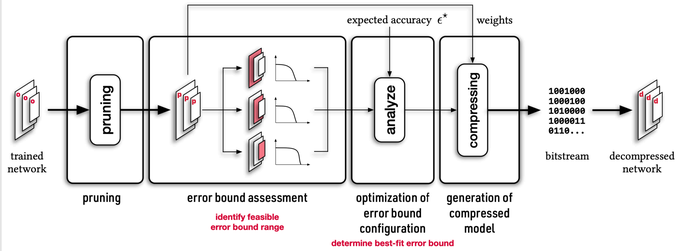

↑ Fig 10 Overview of DeepSZ framework for neural network compression. Reduce the model size by up to 116x with only < 0.3% accuracy degradation.

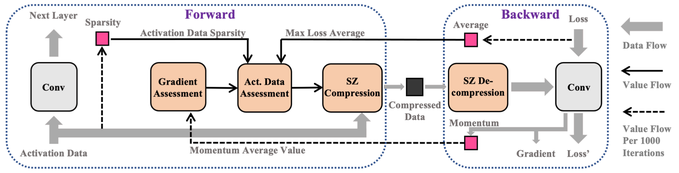

↑ Fig 11 Overview of our proposed memory-efficient CNN training framework COMET. Reduces the memory footprint during model training by up to 13.5x over baseline and 1.8x over previous.

Fig 12 → Comparison of validation accuracy between baseline training and COMET. COMET provides high memory reduction rate while maintaining the convergence speed of validation accuracy.

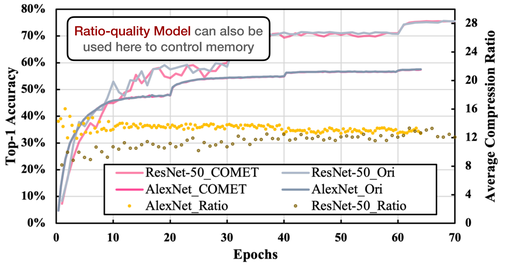

Primary Publications & Acknowledgments

EuroSys'24: Concealing Compression-accelerated I/O for HPC Applications through In Situ Task Scheduling
SC'22: Accelerating Parallel Write via Deeply Integrating Predictive Lossy Compression with HDF5
ICDE'22: Improving Prediction-Based Lossy Compression Dramatically via Ratio-Quality Modeling
VLDB'22: COMET: A Novel Memory-Efficient Deep Learning Training Framework by Using Error-Bounded Lossy Compression
HPDC'21: Adaptive Configuration of In Situ Lossy Compression for Cosmology Simulations via Fine-Grained Rate-Quality Modeling
IPDPS'20: Understanding GPU-Based Lossy Compression for Extreme-Scale Cosmological Simulations
HPDC'19: DeepSZ: A Novel Framework to Compress Deep Neural Networks by Using Error-Bounded Lossy Compression

Argonne NATIONAL LABORATORY
BERKELEY LAB
ECP EXASCALE COMPUTING PROJECT
Los Alamos NATIONAL LABORATORY EST. 1943
NSF
TACC
XSEDE
Publications