Lossy Scientific Data Compression With SPERR

Samuel Li (Presenter)
and John Clyne

May 25th, 2022

This material is based upon work supported by the National Center for Atmospheric Research, which is a major facility sponsored by the National Science Foundation under Cooperative Agreement No. 1852977.

Overview

- Lossy means that a small amount of error is introduced to every data point.
  - Very often achieves 10X data reduction
  - In contrast to lossless compression, above 2X is very rare.

- **SPERR** applies to floating-point volumetric data
  - Numeric simulation output.
  - 2D slices or 3D volumes in nature.

- **SPERR** is built on wavelet transforms, and is unique in providing a fixed point-wise-error compression ability.

- Compared to popular lossy compressors, **SPERR** has the best rate-distortion curve that we know of at a cost of more computationally intensive.

- **SPERR** is currently under active development with more enhancements coming, and is open-source: https://shaomeng.github.io/SPERR/

How Does SPERR Work

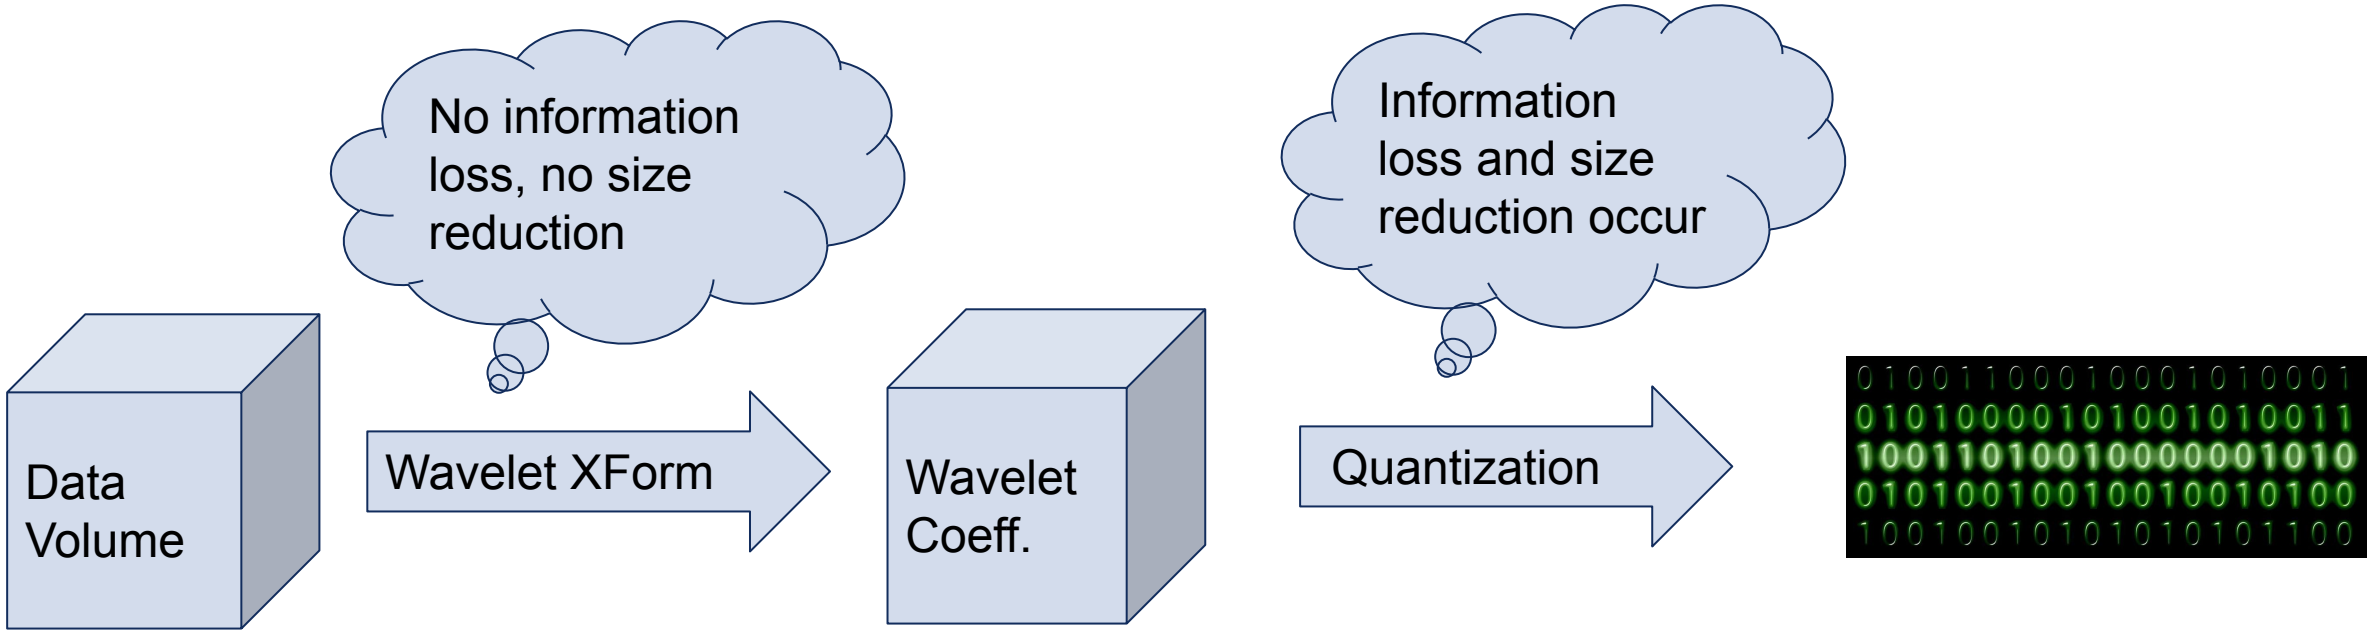

- Wavelet transform compacts much information into a small number of **significant coefficients**.
  - Sometimes this ability is called energy concentration.

- Quantization saves approximations of significant coefficients using various number of bits.
-
- Quantization finishes when a storage budget is met. I.e., **fixed-size compression**.

Error Distribution

- SPERR is designed to minimize average error of reconstruction.
- It results in normal error distributions.
  - The maximum point-wise error is about one order of magnitude bigger than the mean square error

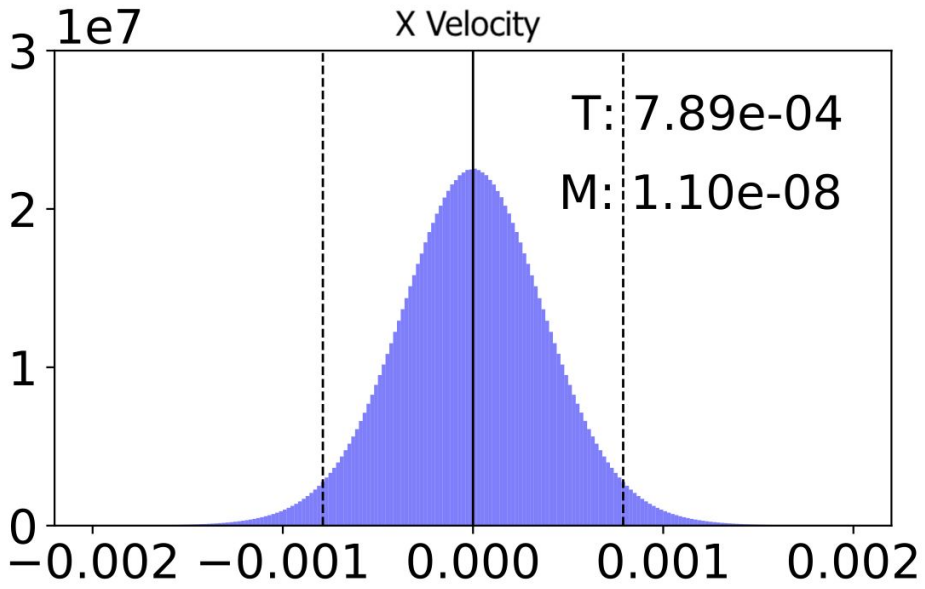

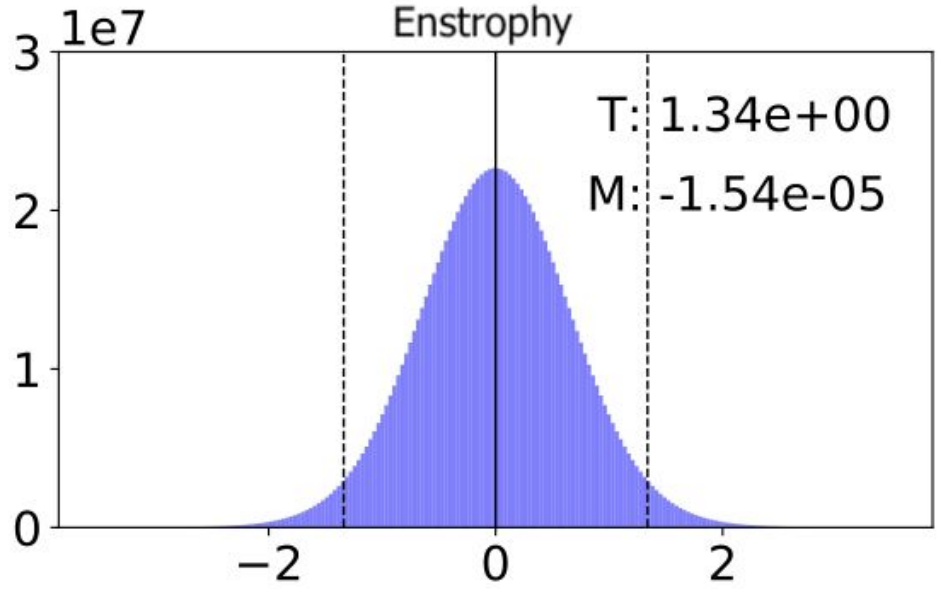

How Does SPERR Work – Fixed Point-Wise Error

Outliers: data points with error exceeding a user-defined threshold.
→ They live on the tails of those normal error distributions.

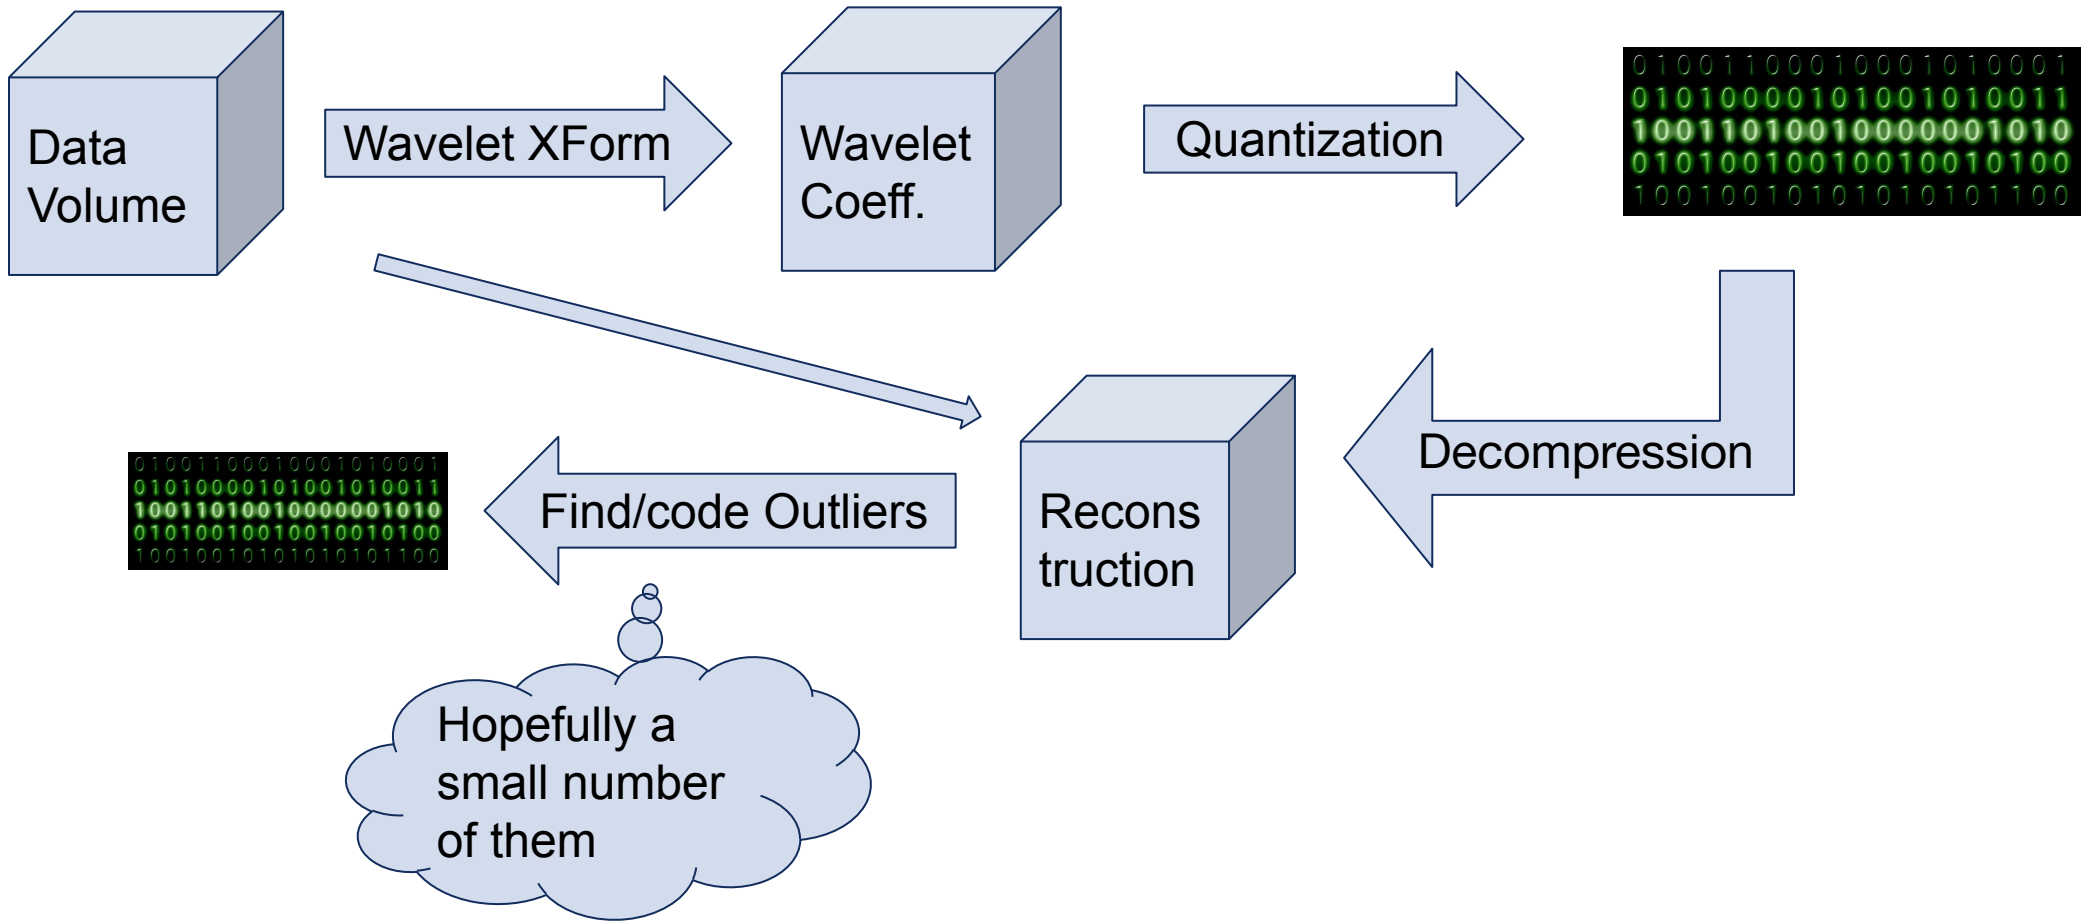

Error Distribution After Outlier Correction

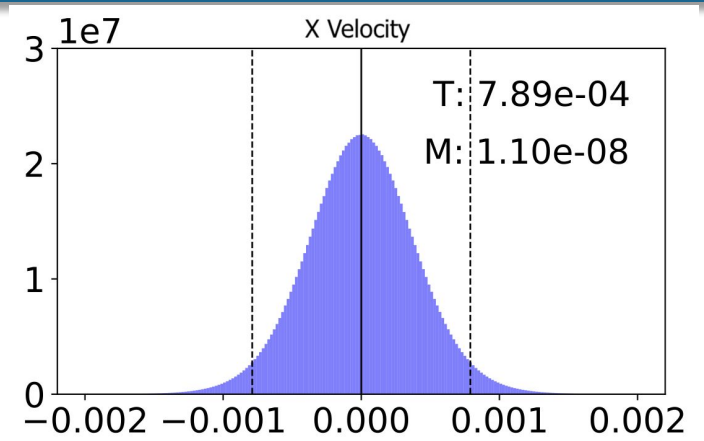

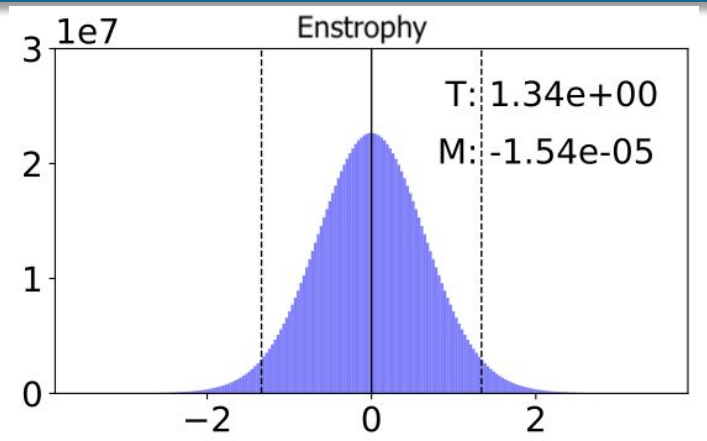

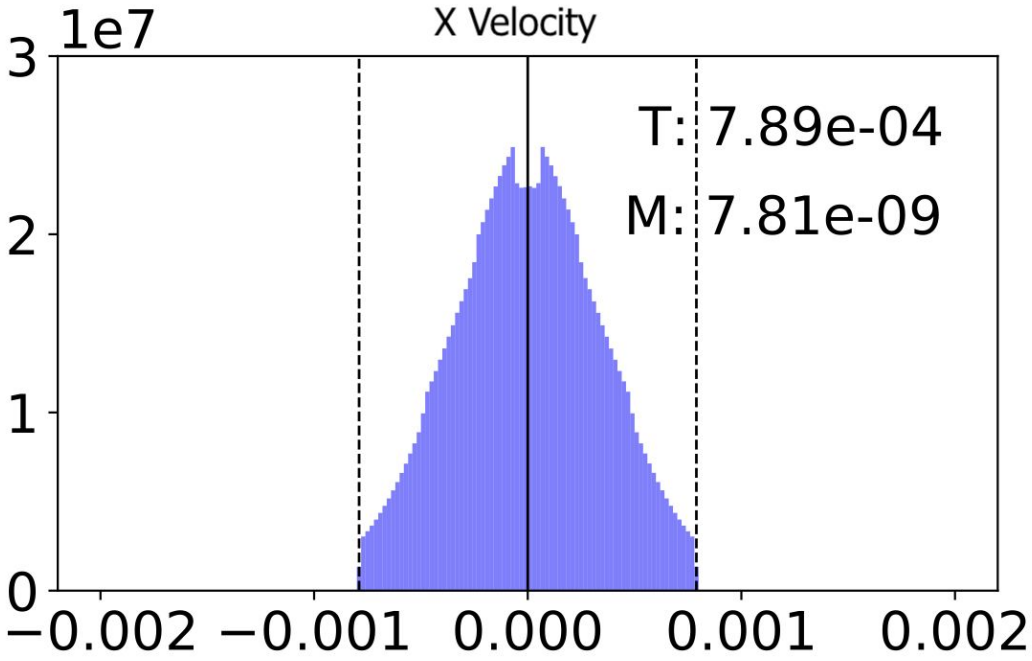

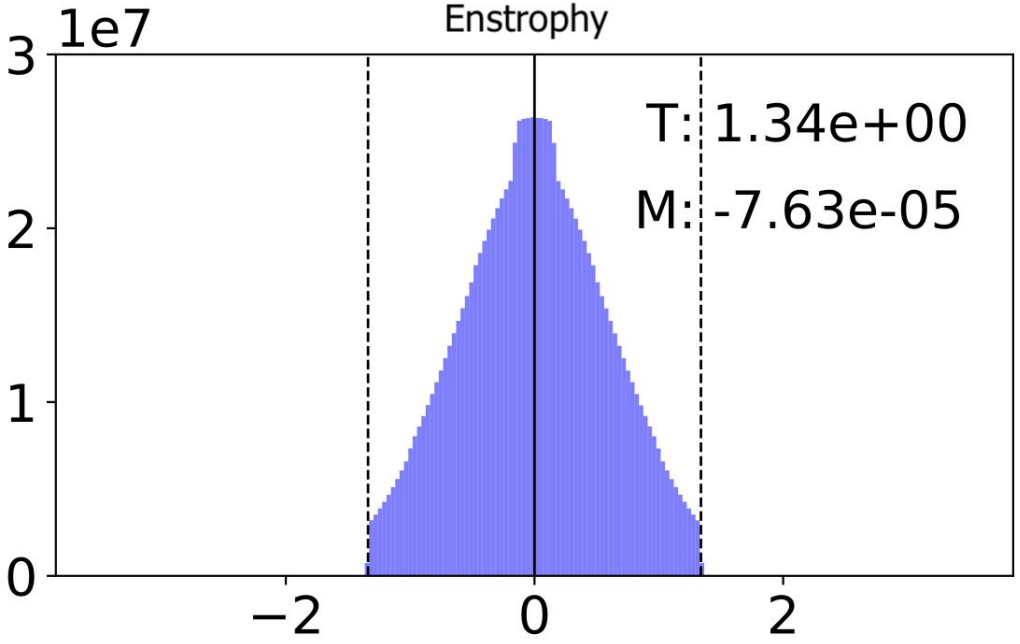

How Does SPERR Compare

Comparators: ZFP and SZ
Tool: rate-distortion curves

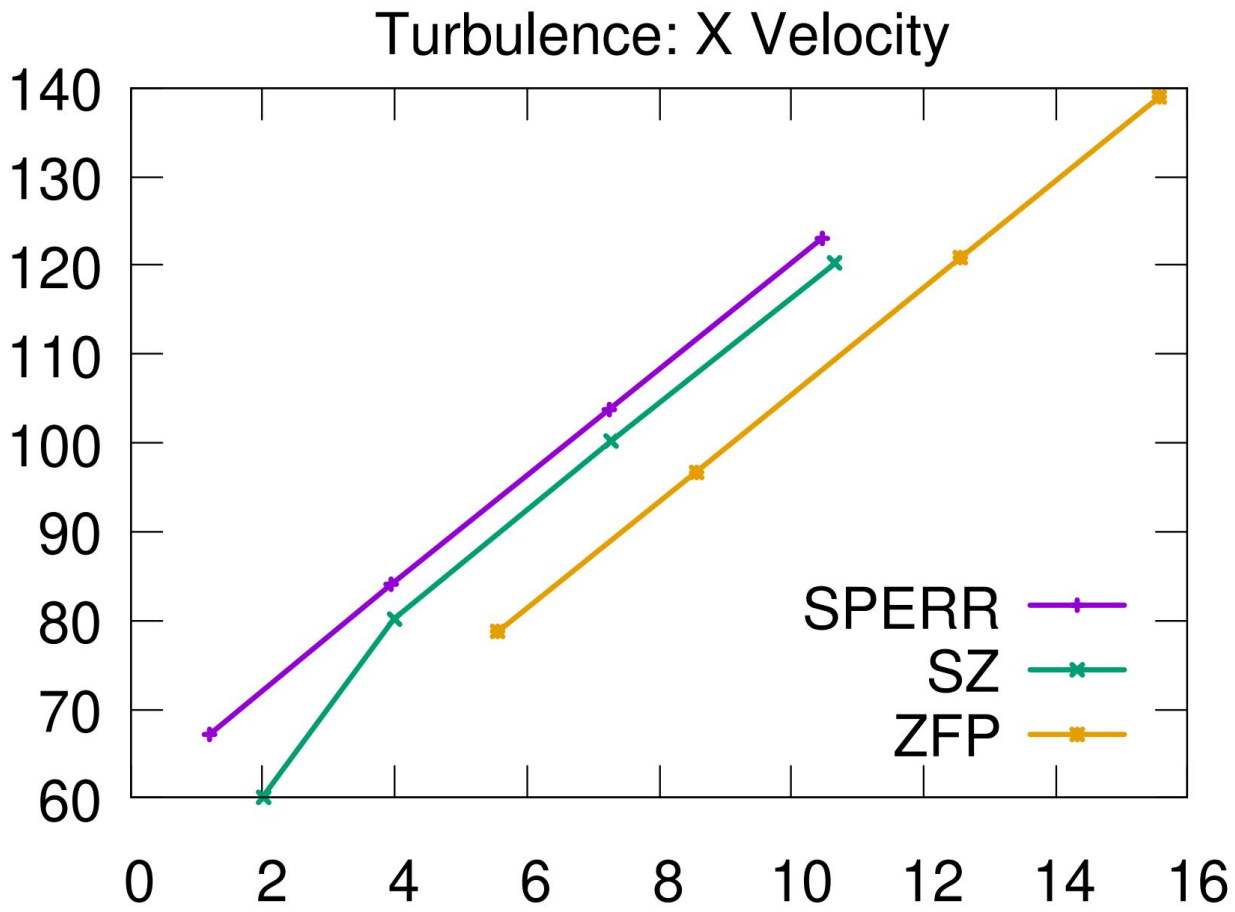

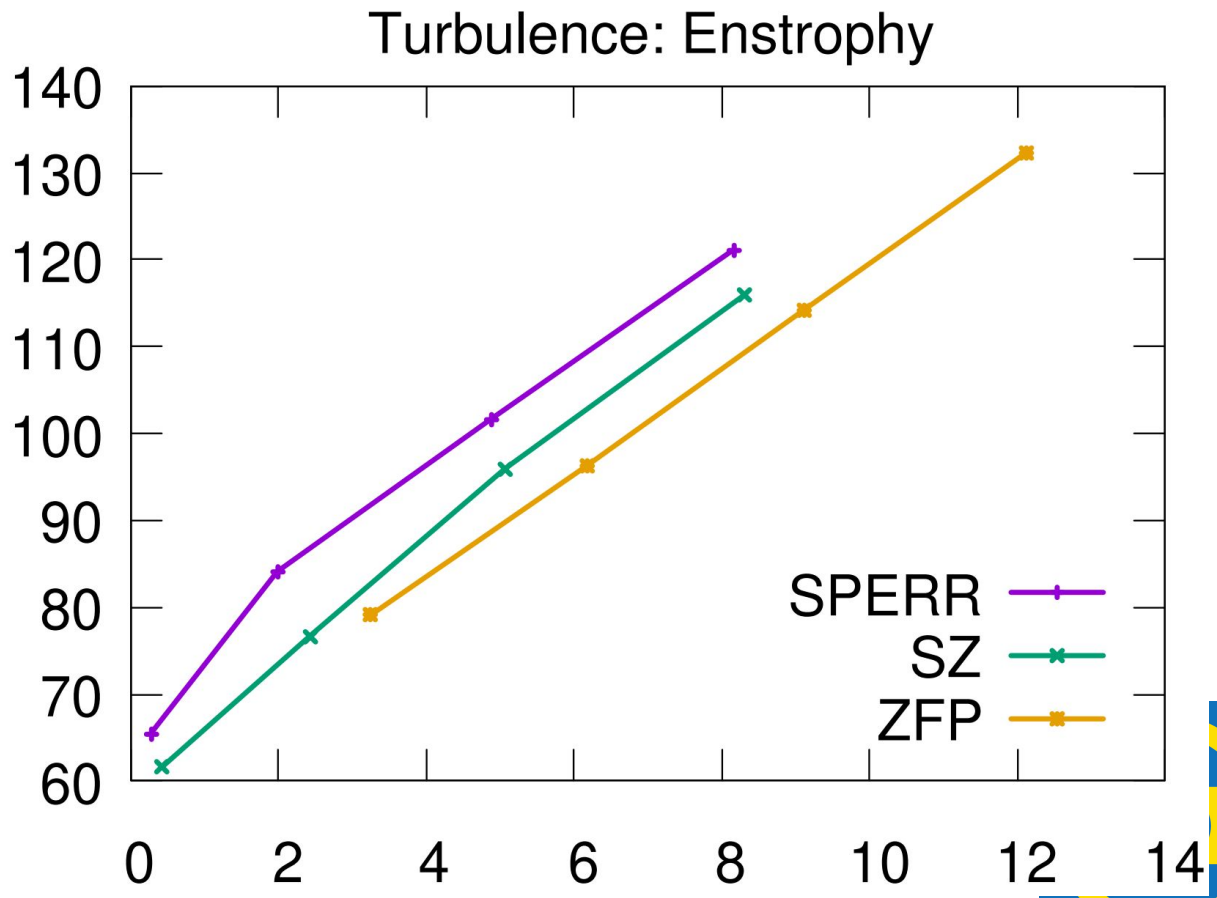

How Does SPERR Compare – cont.

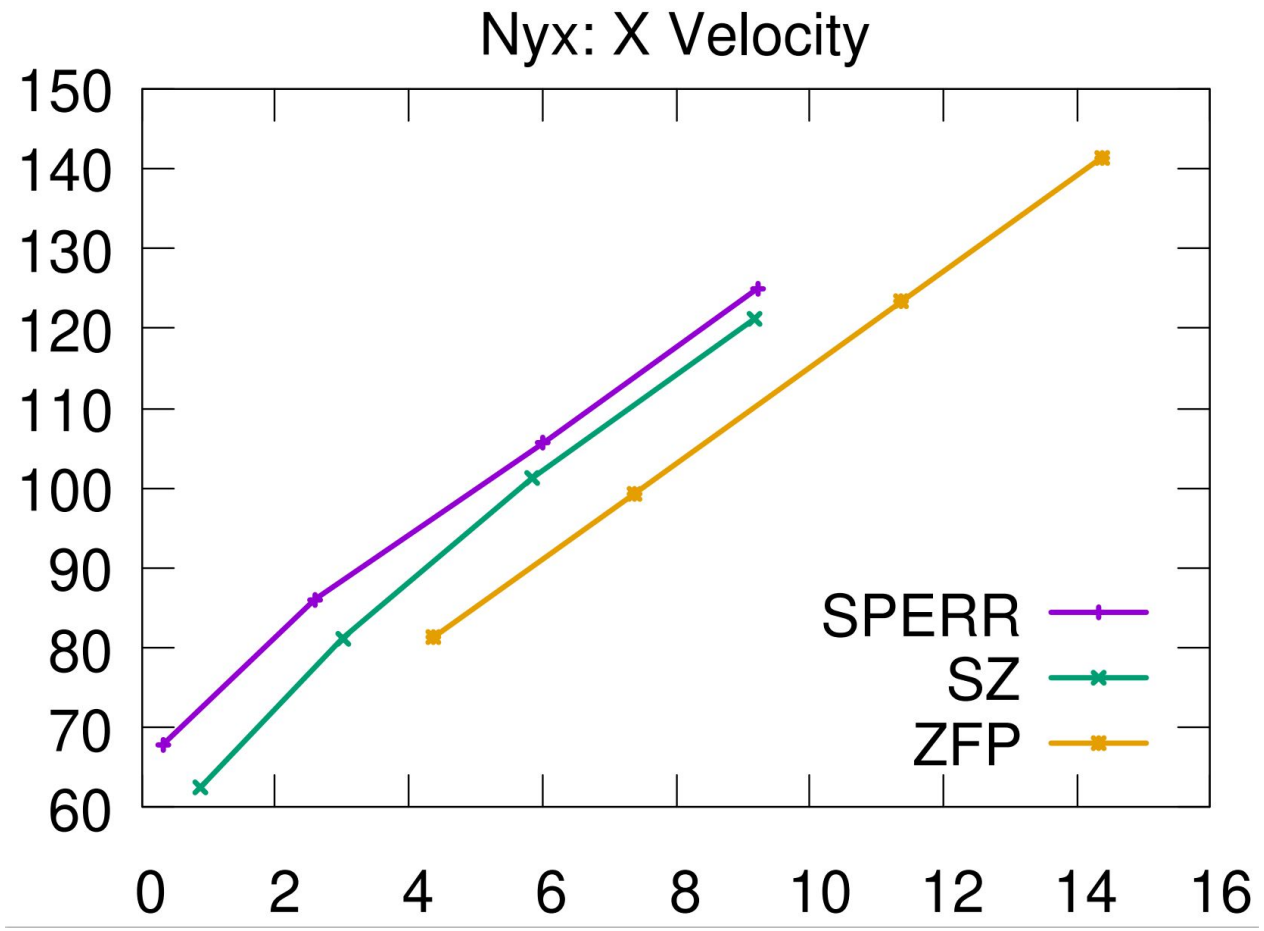

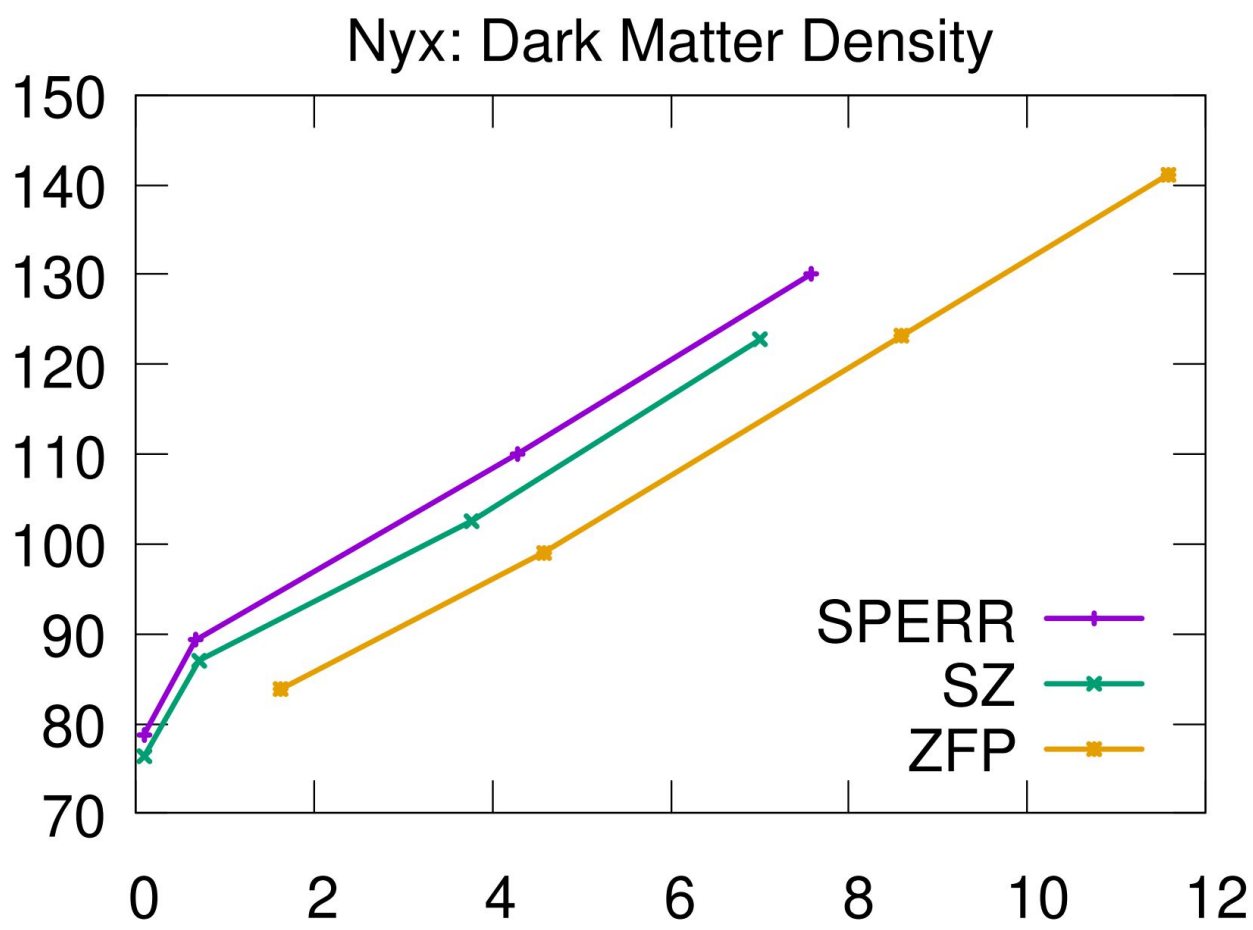

How Does SPERR Compare – cont.

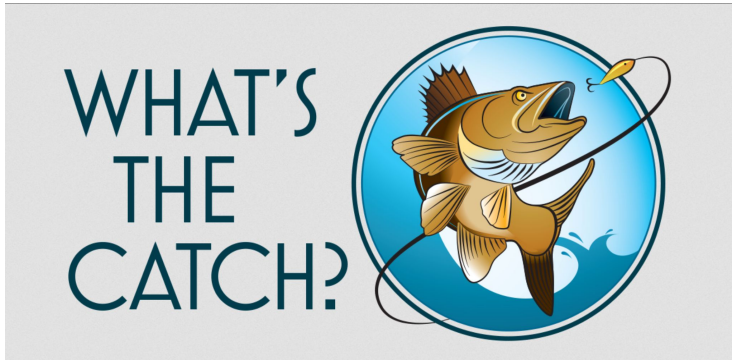

The biggest downside of SPERR is that it's slow.

- Our observation is that it's often a few times (2X–10X) slower than SZ and ZFP, depending on compression configurations.

- Parallelization:
  - It parallelizes on multi-core CPUs using domain decomposition (default: 128^3 domain size).
  - It isn't obvious how to implement SPERR efficiently on GPUs.

Current Development

- Performance improvement
- Multi-resolution support
  - E.g., when reconstructing a 512^3 volume, it produces approximations of sizes 256^3, 128^3, 64^3, etc. for explorative analysis.
  - No extra computational or storage cost; high quality approximations.
  - Enabled by wavelet transforms
- Regions of interest (ROI) decompression support
  - Enabled by the domain decomposition parallelization strategy.
- Fixed average error (measured by PSNR) compression mode.
- https://shaomeng.github.io/SPERR/

Lossy Scientific Data Compression With SPERR

Samuel Li (Presenter)
and John Clyne

May 25th, 2022

This material is based upon work supported by the National Center for Atmospheric Research, which is a major facility sponsored by the National Science Foundation under Cooperative Agreement No. 1852977.