A General Framework for Error-controlled Unstructured Scientific Data Compression

Qian Gong[+], Zhe Wang[+], Viktor Reshniak[+], Xin Liang[*], Jieyang Chen[++], Qing Liu[^], Tushar Athawale[+], Yi Ju[=], Anand Rangarajan[#], Sanjay Ranka[#], Norbert Podhorszki[+], Scott Klasky[+]

Oak Ridge of National Laboratory[+], USA
University of Kentucky[*], USA
University of Florida[#], USA
University of Alabama[++], USA
Max Plank Computing and Data Facility[=], Germany
New Jersey Institute of Technology[^], USA

Data tsunami

Due to the constrained storage capacity and I/O bandwidth, scientific finds are increasingly constrained by the amount of data that can be outputted and stored

Wind turbine VKI case running by **GE**

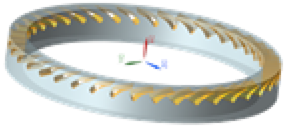

6K GPUs on 1K Summit nodes → 3s per iteration → 1.4 TB output per step → Every 100 iterations, 10 hours runtime → 168 TB of data into storage

XGC fusion code

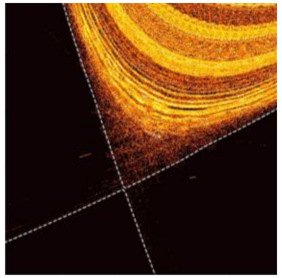

1 TB per step → 1024 Frontier nodes → 20 hours runtime → 2 PB of data into storage

SKA1-LOW **SKA1-MID**

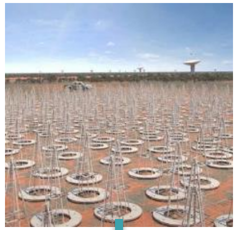

2 Pb/s 9 Tb/s

Multi-stage reduction/filtering

600 PB/yr

SKA Regional Centers

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

Data transmission & storage challenges

- Network bandwidth and storage capacity have been lagging behind
- It's critical to design techniques that can:
  - Effectively reduce the data sizes that go to tape and move around geographically distributed sites with quantifiable errors
  - Accelerate writing and retrieving of data without sacrificing important physics (e.g., time decimation risks of losing transitory information)

I/O throughput w/ ADIOS on Frontier

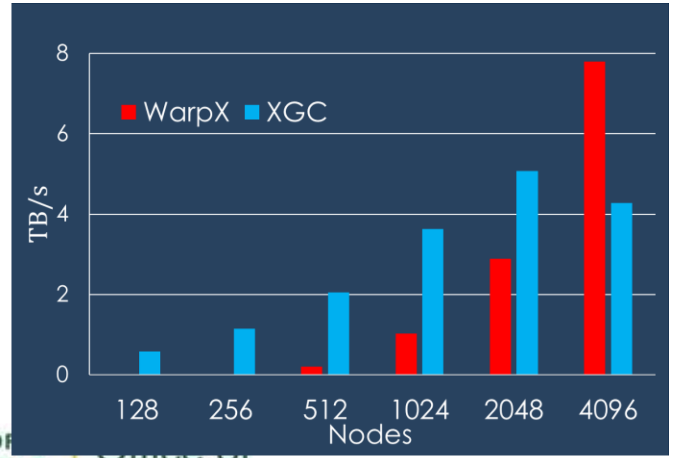

Short-term (e.g., parallel file system) vs long-term (e.g., tape) storage systems

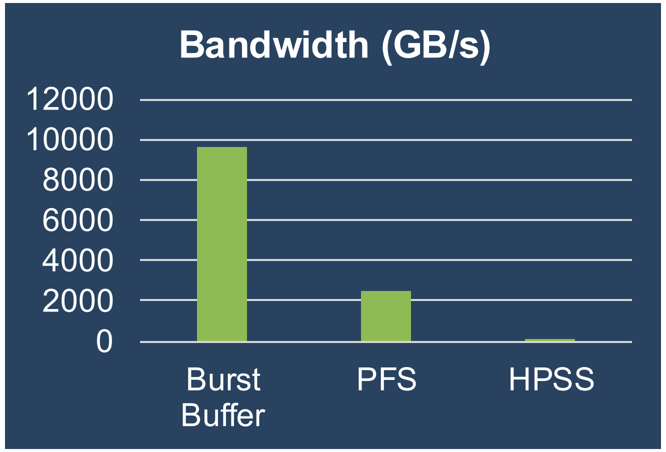

Lossy compression for science applications

- Conventional entropy-based lossless compression achieve limited compression ratios (CR) on scientific data
- Lossy compression provides
  - larger compression ratios allowing data to be written at more frequent rates and saved for longer time
    - ❑ MGARD compressed E3SM hourly data by 24x, leading to 4x reduction of storage and 10X more accurate cyclones prediction comparing to the 6-hourly output setting
    - ❑ MGARD + post-processing allowed to reduce the f-data from XGC by 300X while maintaining errors with 10^{-8} accuracy on 5 derived quantities
  - mathematically guaranteed compression-incurred errors
    - ❑ Primary data, quantities-of-interests (QoIs)

E3SM
(sea level pressure)

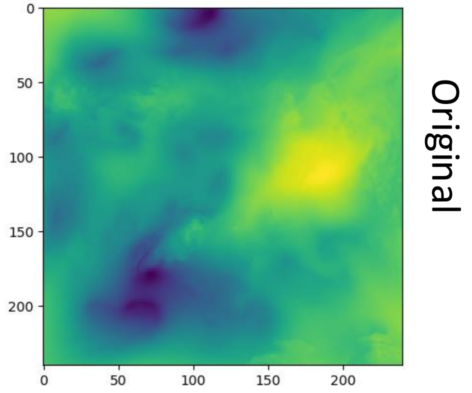

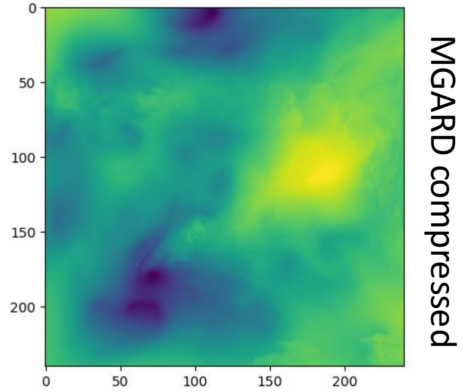

CR = 28.5x,
NRMSE = 5e-4,
compared to lossless with ZSTD: CR = 1.28x

Lossy compression for unstructured-meshed data -- limitation

- Unstructured meshes are widely utilized in finite element simulation to represent complex geometries
  - Time-series, multi-variate data on mesh vertices
- Most lossy compressors rely on block or stencil-based algorithms for data transformation, requiring
  - sterilization of mesh node data followed by compression on 1D arrays

    sub-optimal compression ratios due to overlooking of data correlation in the mesh topology
  - or value-based vertex traversal

    slowing down compression and requiring to storing the traversal graph

Data in GE VKI case reduced by MGARD

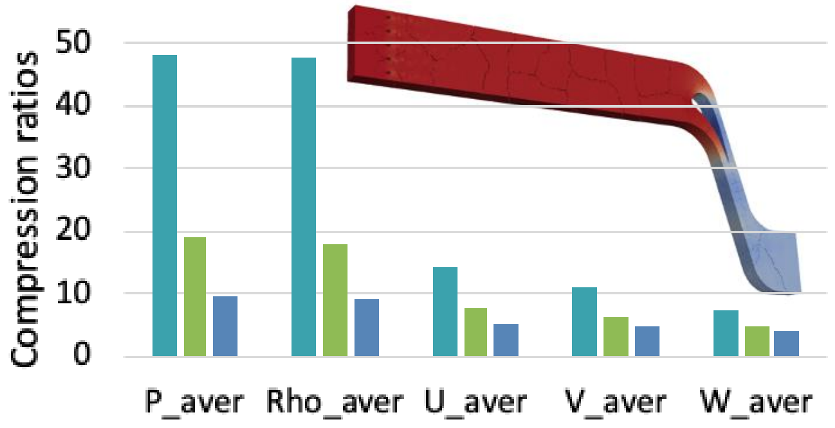

CR: 14.5x 8.2x 5.6x

Default order of mesh vertices

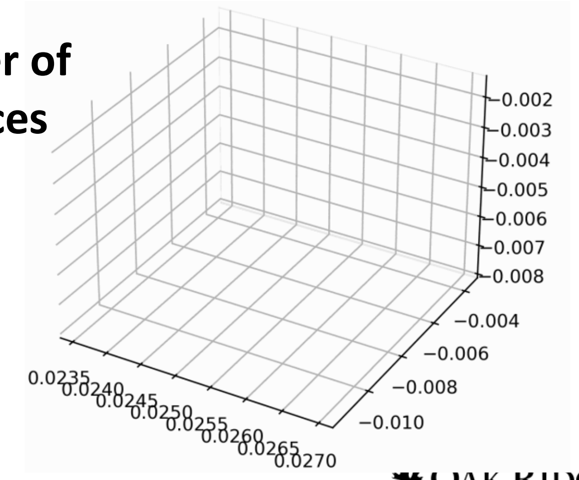

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

Motivation for multi-component unstructured compression

- Unstructured mesh data can be approximated by rectilinear grid data, with small errors
- The approximation errors and interpolated gridded data can be separately compressed in error-controlled fashion
  - Denote the interpolation function as $f(\bar{x}) = a^T\bar{x} = x_1$, where a^T is interpolation coefficient, $\bar{x}$ is associated node data, and x_1 is the resulted data on grid
  - Denote the back-interpolation function as g, computing the residuals on mesh nodes as $x_2 = x - g(x_1)$
- Can we obtain large compression ratios from both terms?
  - mesh residuals x_2 (small-values), and
  - interpolated data x_1 (higher correlation on structured grids)

Pressure produced by OpenFoam on a 2D mesh

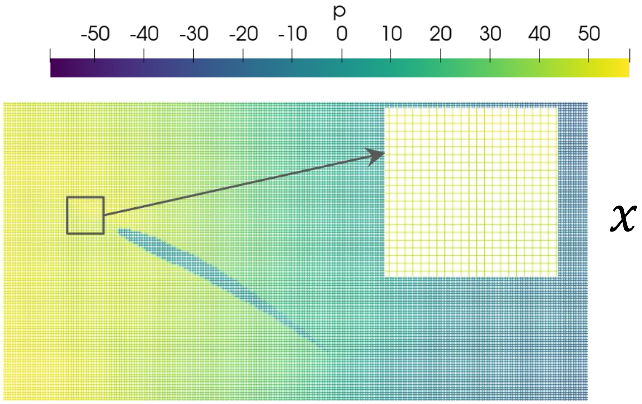

x

Interpolation on grids

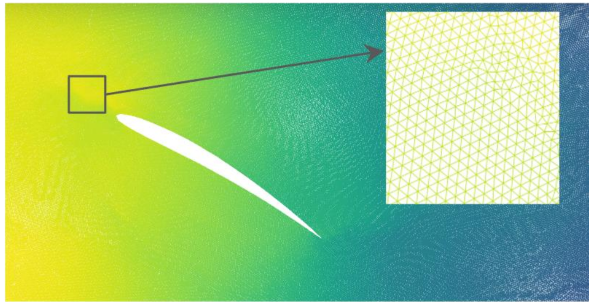

x_1

of x_1 equivalent to 39% of x_2 producing similar visualization → small approximation errors

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

Preliminary on mesh interpolation

- Interpolation function f (many-to-one), back-interpolation function g (one-to-many)
- Denote $c(.)$ as an error-controlled compressor used for compressing grid interpolation (x_1) and mesh residual (x_2), such that:
  - $|\epsilon_1| = max|x_1 - x'_1| \leq \tau_1$ and $|\epsilon_2| = max|x_2 - x'_2| \leq \tau_2$, where x'_1 and x'_2 are the reduced data
- The reconstruction of x using x'_1 and x'_2 can be captured by $x' = g(x'_1) + x'_2$
- To preserve the errors of x', we require a linear function g for back-interpolation, thus
  - the accumulated errors in x' can be bounded by $(\sum_j |a_j|)\ \tau_1 + \tau_2 \leq \tau$
  - a special case where g corresponds to the nearest neighbor values or piece-wise linear interpolation, the above formula can be reduced to $\tau_1 + \tau_2 \leq \tau$

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

Multi-component unstructured compression pipeline

- Initialization includes to construct a rectilinear grids and a mesh-to-grid mapping
  - reusable across timesteps and variables → small overhead in computational throughput
- Main compression pipeline consists of mesh-to-grid interpolation, residual calculation, and separate lossy compression – N-dimensional compression for grid interpolation and 1D compression for mesh residuals

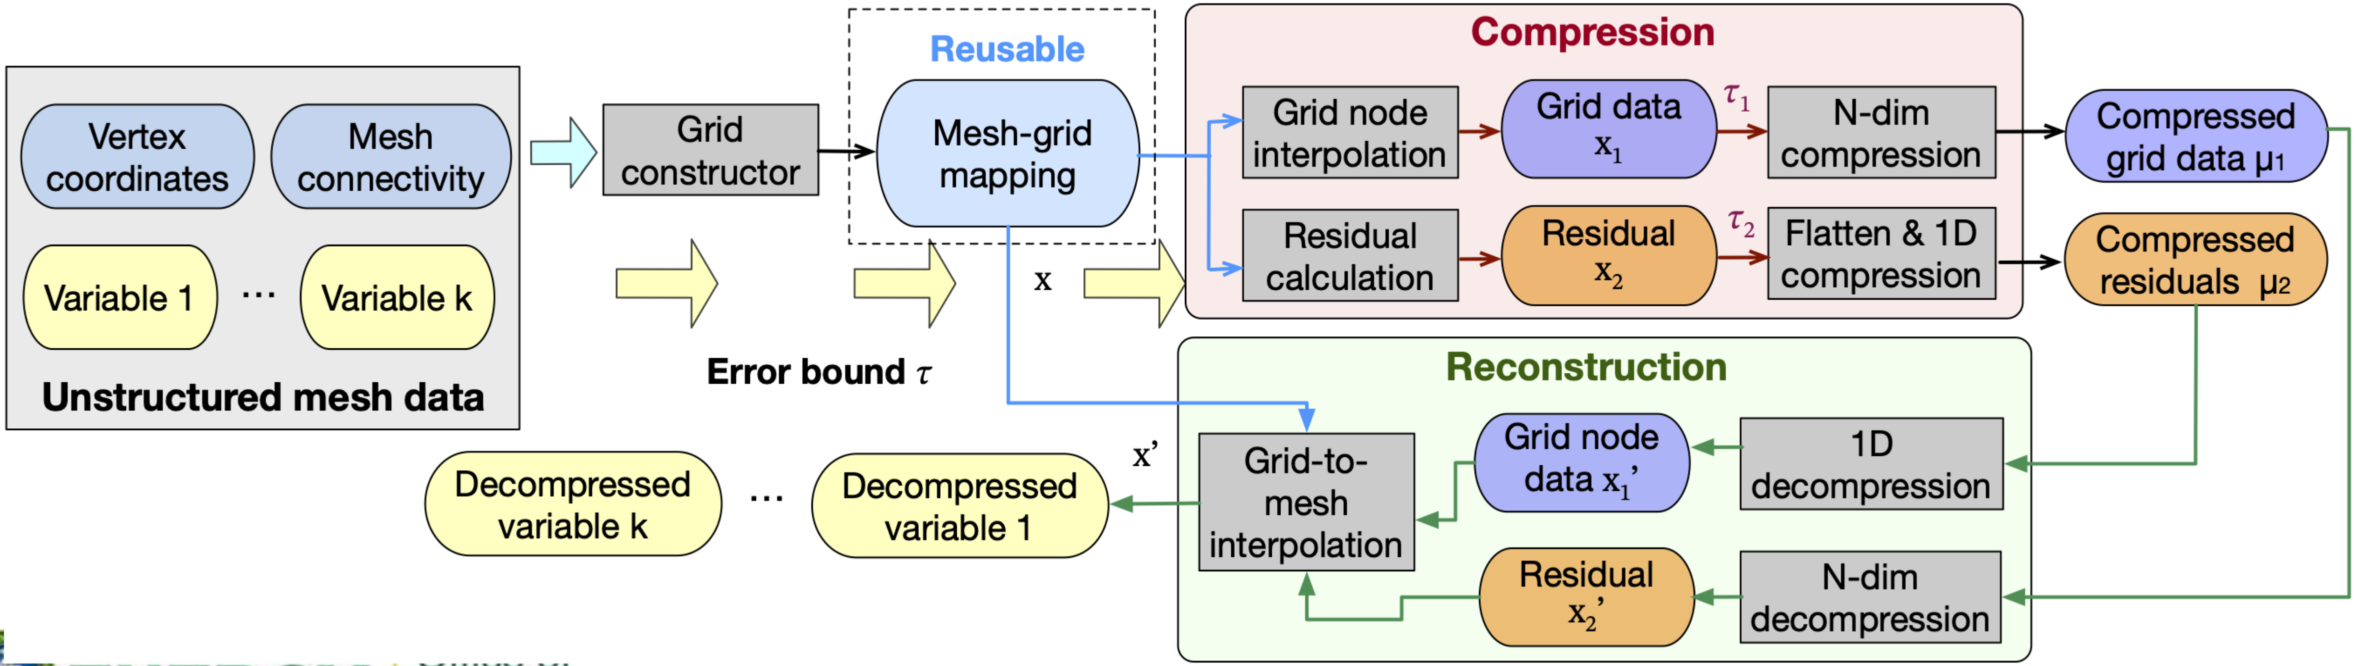

Hyperparameter optimization (data dependent)

- Spacing of the uniform grids:
  - # of grid points & values of mesh residuals

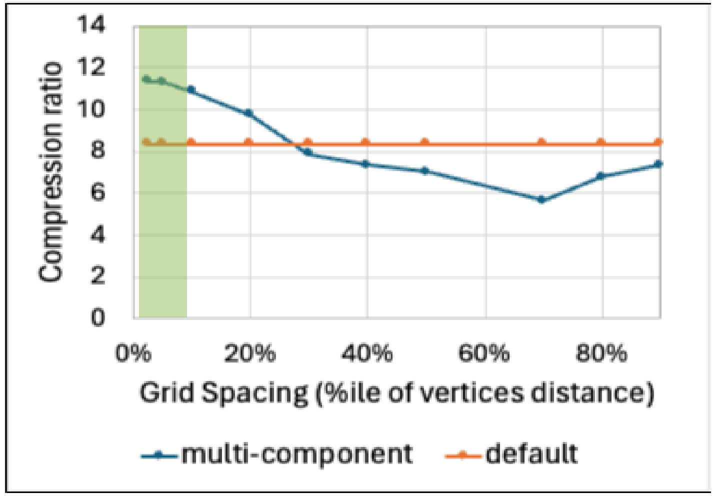

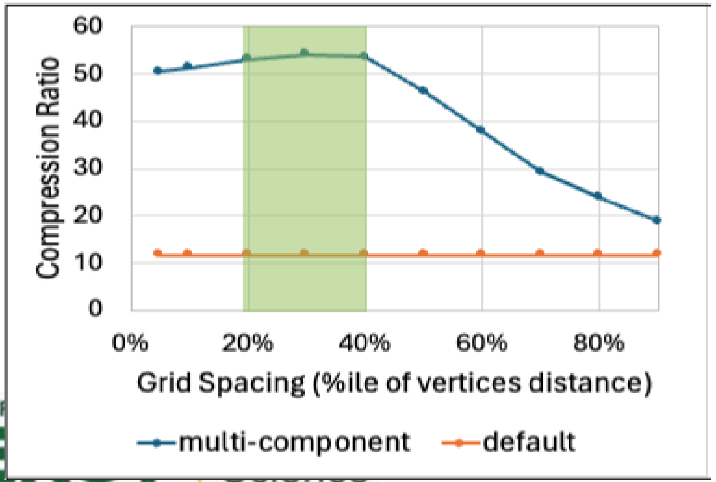

- Error bound for grid interpolation (τ_1) and residuals (τ_2): $\tau_1 + \tau_2 \leq \tau$

Nek5000 velocity (highly turbulent)

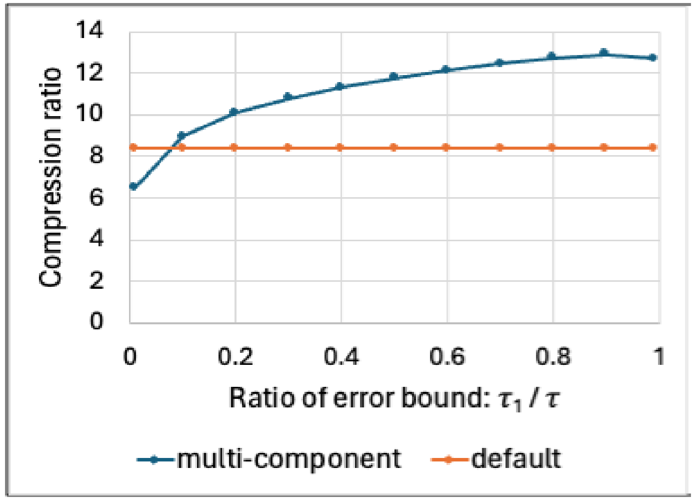

VKI pressure (relatively smooth)

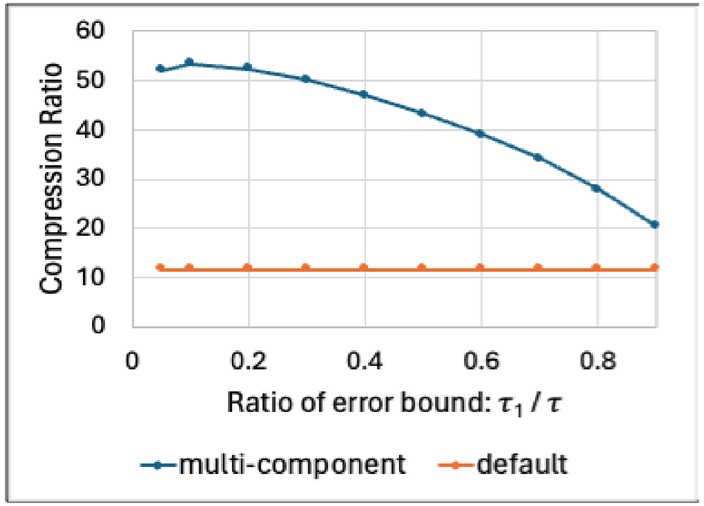

U.S. DEPAR ENE

OAK RIDGE National Laboratory

Experimental datasets

- Evaluations using four datasets on 2D, 3D unstructured meshes

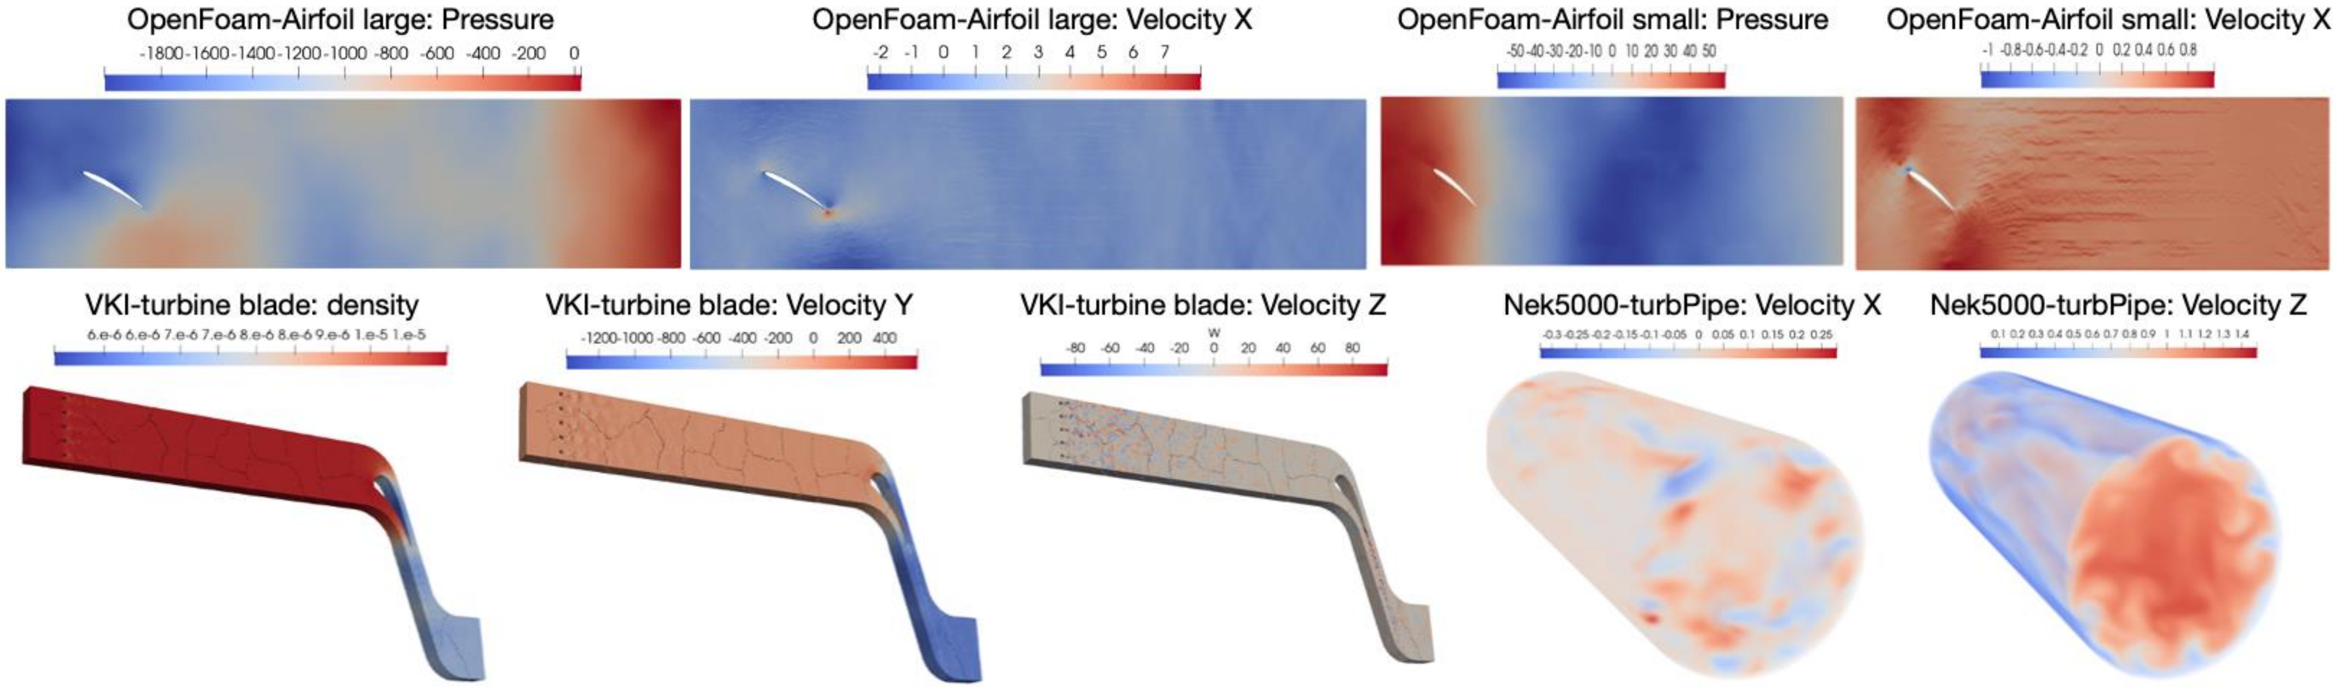

Dataset	# of attributes	# of vertices
Nek5000-turbPipe	3	1.5 million
VKI turbine blade	5	210 million
OpenFoam airfoil-small, large	3	0.5 million, 2 million

Improvements in compression ratios demo using MGARD

- Multi-component compression pipeline + MGARD for reducing the grid interpolation and mesh representation

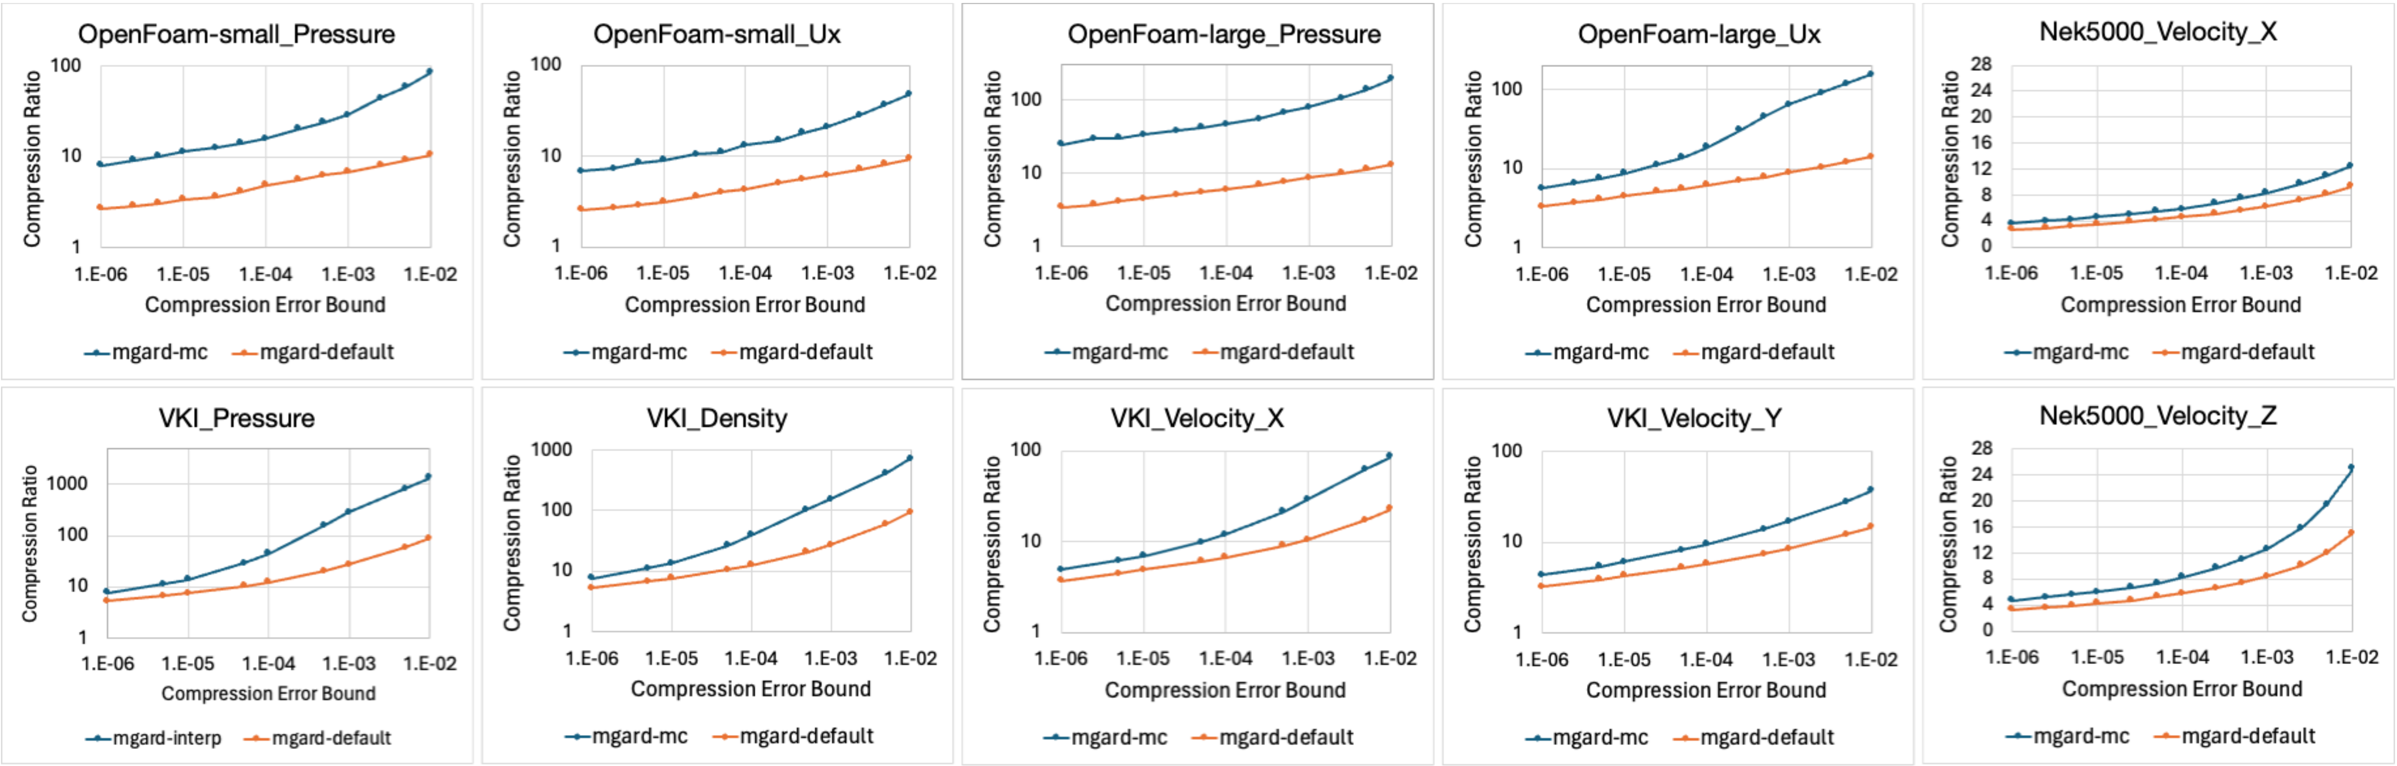

Larger improvements for high error bounds and smooth data

Demo the improvements for mores compressors

Compression ratio achieved by multi-component pipeline w/ associated compressors

- Improvement ratios = $\frac{CR_{mc}}{CR_0}$, optimizing the unstructured compression pipeline for MGARD, SZ, and zfp

Compression ratio achieved by the same compressors on flattened 1D data

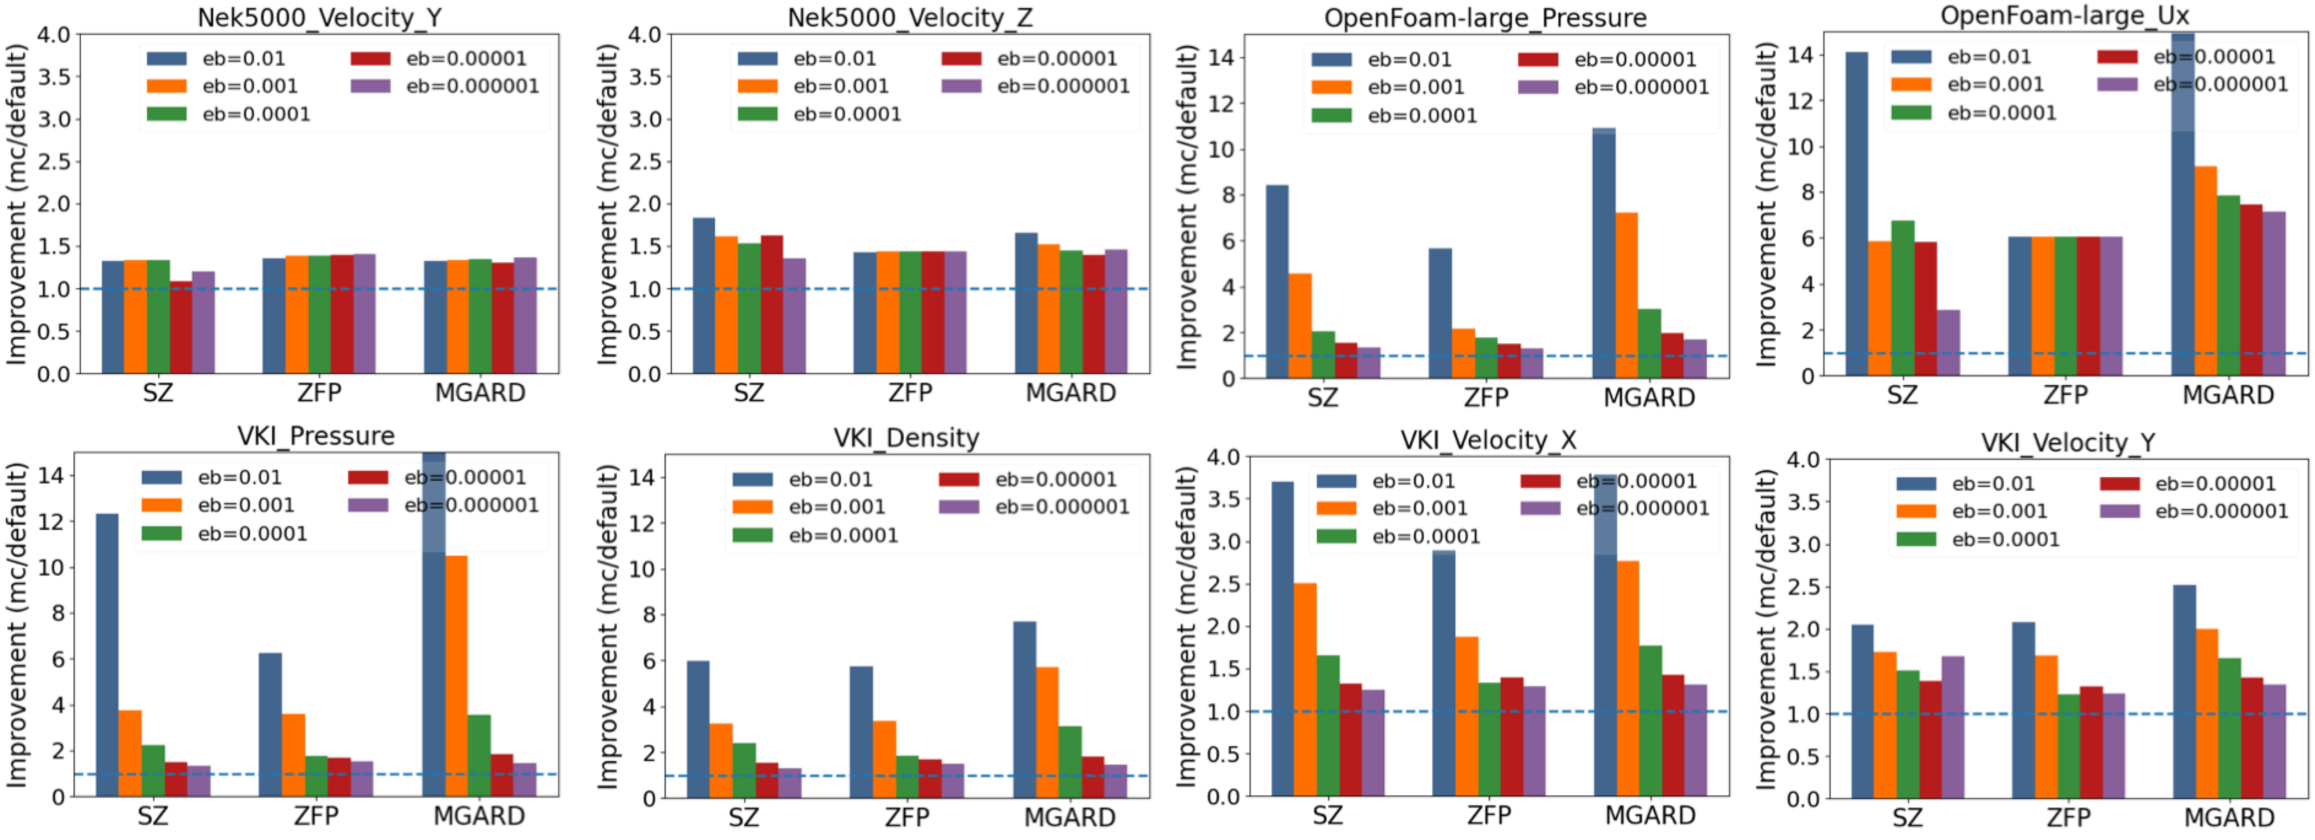

*In average **2.3-3.5x** improvement across the 3 compressors*

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

Throughput performance

- Overhead = $\frac{T_{mc} - T_0}{T_0}$

$T_{mc} - T_0$ → Extra time spent by multi-component pipeline w/ MGARD

T_0 → Time spent by MGARD on flattened 1D data

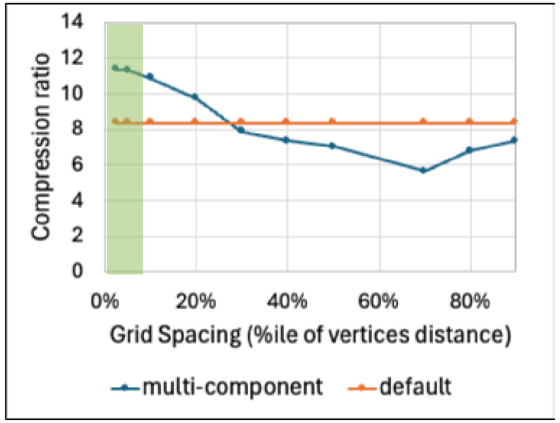

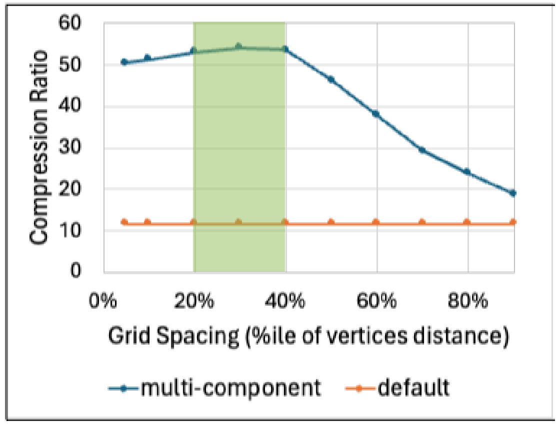

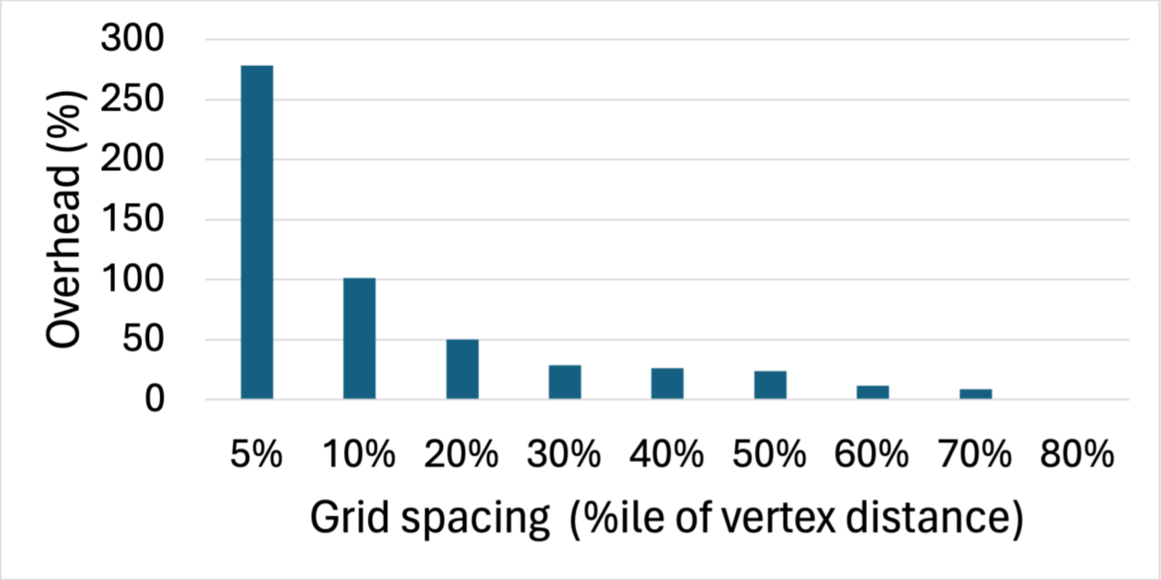

- Overhead comes from: 1) mesh-to-grid interpolation, and 2) compression of interpolated representation
- Percentage of overhead varies as a function of the grid granularity, and commonly range from 25%-100% with the optimal settings used by most datasets

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

Future explorations

- Our multi-component compression pipeline addresses the limitation of state-of-art lossy compressors on unstructured mesh data and only depends on mesh topologies
- Future optimization can be obtained by adaptively refining the grid and using variable order of functions for more accurate and compact interpolation
- Refactor the unstructured data, providing a series of representations with varied degree-of-free and accuracies
  - An alternative to single-error-bounded compression
  - Fine-grained error control and need-based retrieval