COMPRESSION AND DATA REDUCTION I

Stability-preserving Lossy Compression for Large-scale Partial Differential Equations

Qian Gong[1], Mark Ainsworth[2], Jieyang Chen[3], Xin Liang[4], Liangji Zhu[5], Ethan Klasky[5], Tushar Athawale[1], Qing Liu[6], Anand Rangarajan[5], Sanjay Ranka[5], Scott Klasky[1]

[1]Oak Ridge National Lab
[2]Brown University
[3]U of Oregon
[4]U of Kentucky
[5]U of Florida
[6] NJIT

SC25 ST. LOUIS hpc ignites.

The Era of Data Deluge

- Data growth far outpaces storage and I/O bandwidth improvements → moving and storing data are now costlier than computing it
- Data reduction is essential for energy-efficient workflow and sustainable data management
  - Dramatically reduce data volume and accelerate processing while maintaining scientific fidelity and workflow stability

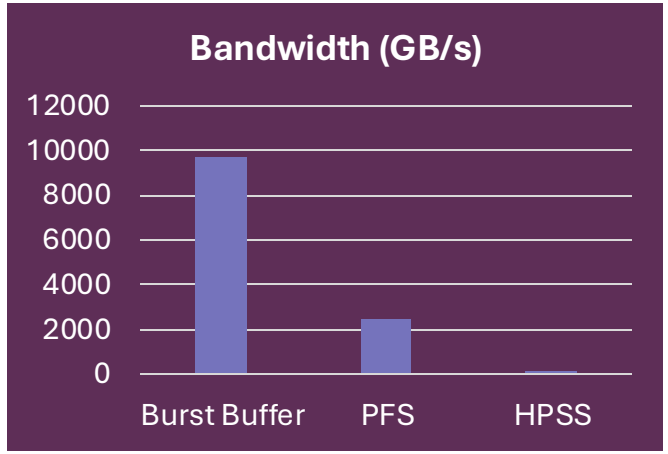

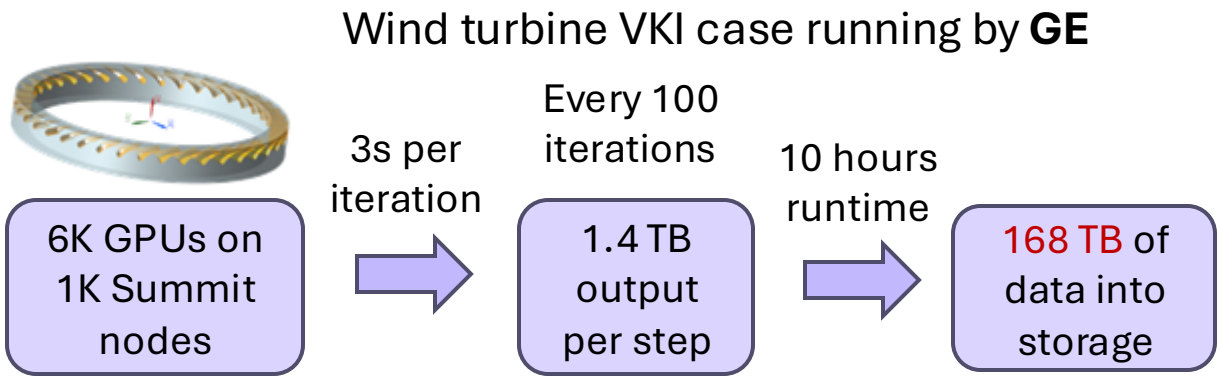

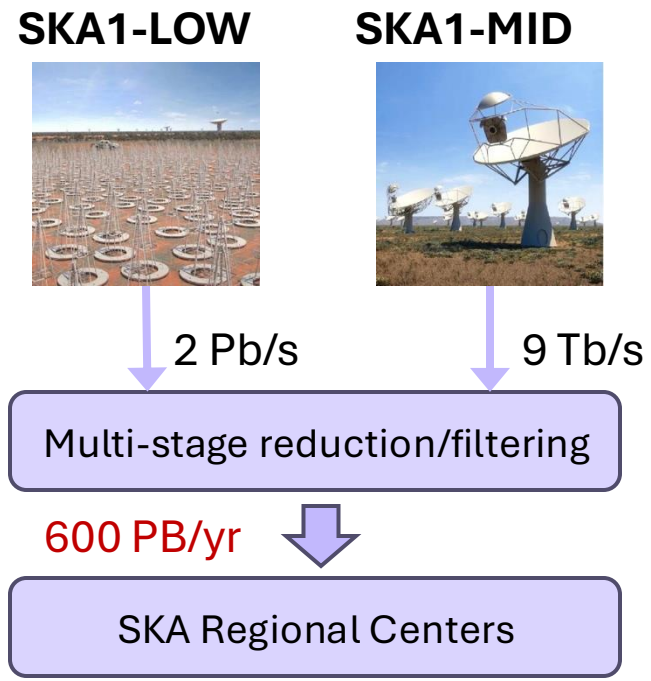

Lossy Compression

- Lossy compression enables 10-500x reduction rate on floating-point data with **controlled scientific impact**

- Research opportunities:
  - Quantified scientific impact: primary data, quantities-of-interest (QoI)
  - Structured and unstructured data: data decorrelation, dimensionality reduction
  - Co-design with I/O, streaming, simulation and data management workflow, data lifecycle
  - Hardware acceleration: GPUs, FPGA, SmartNICs
  - AI-based data reduction & compression for AI: QoI preservation, foundation data reduction models, AI workflow cost optimization

The Missing Piece: Impact of Compression on Long-term Simulation Stability

- Static errors vs dynamically evolving errors:
  - Lossy compression can be applied for checkpointing
  - Physical stability after lossy checkpoint-restart remains underexplored → essential for trustiness as instabilities compromise scientific reliability
  - Even small perturbations can amplify in PDE simulations

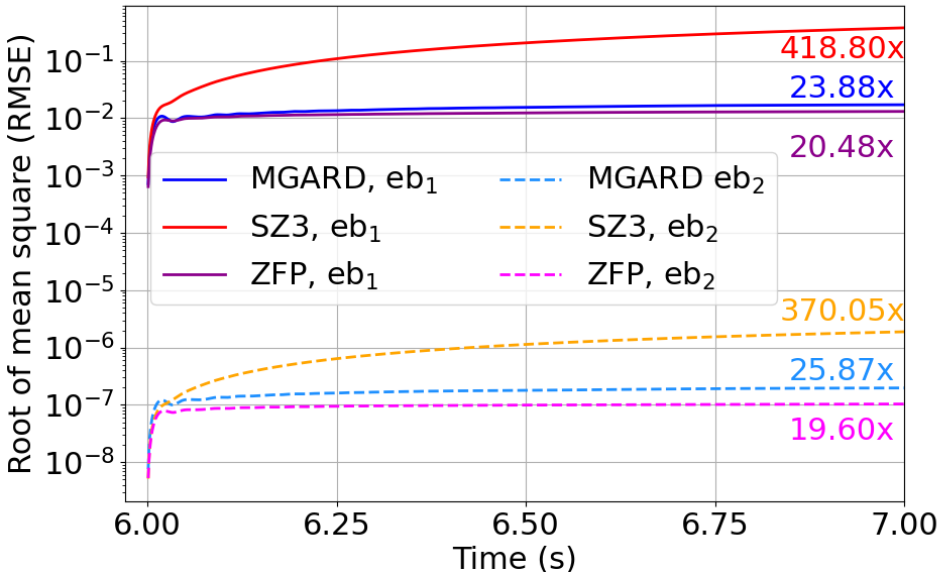

Error propagation in a wave equation simulating the motion of a 2D pond surface responding to raindrops

- Execute the simulation for 6 seconds ($\Delta t = 1x10^{-3}$), compressing the field using three lossy compressors with a relative $eb_1 = 10^{-3}$ (CR~29x) and $eb_2 = 10^{-8}$ (CR~5x)
- A spike in error presents immediately after restart from the lossy checkpoints, and continues to grow, regardless of the error bound and compressors being used
- Similar behaviors seen for other PDEs

Lossy Checkpoint/Restart (C/R) for PDEs

- **Why choose C/R?**
  - A vital technique for fault-tolerance, resilience, and spatiotemporal state replay
  - Instability from compression-incurred error
- **Why choose the wave equation?**
  - Parabolic PDEs: compression perturbation naturally dissipates over time
  - Second-order hyperbolic PDEs: energy conservation → unstable restarts occur when compression breaks physical invariants **vs** errors in standard norms (e.g., L^2, L-inf)
- **Governing equation:**

$$\frac{\partial^2 u}{\partial t^2} - c^2\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}\right) = S(x,y,t)$$

Periodic condition

  - 2nd order finite difference method (FEM):

$$\frac{u_{i,j}^{n+1} - 2u_{i,j}^{n} + u_{i,j}^{n-1}}{\Delta t^2} - c_{i,j}^2\left(\frac{u_{i+1,j}^{n} - 2u_{i,j}^{n} + u_{i-1,j}^{n}}{\Delta x^2} + \frac{u_{i,j+1}^{n} - 2u_{i,j}^{n} + u_{i,j-1}^{n}}{\Delta y^2}\right) = s_{i,j}^{n}$$

  - The recomposition of u^{n+1} requires both u^n and u^{n-1}

Violation of Energy Conservation during Lossy Checkpointing

- Energy conservation in wave equation:
  - In the absence of source, the total energy remains constant in time---system exchanges energy between kinetic (KE) and potential (PE) forms

$$PE(u^n) = \frac{1}{2}\sum_{i,j}\left[\left(\frac{u^n_{i+1,j} - u^n_{i,j}}{\Delta x}\right)^2 + \left(\frac{u^n_{i,j+1} - u^n_{i,j}}{\Delta y}\right)^2\right]\Delta x \Delta y$$

$$KE(u^n, u^{n-1}) = \frac{1}{2c^2}\sum_{i,j}\left(\frac{u^n_{i,j} - u^{n-1}_{i,j}}{\Delta t}\right)^2 \Delta x \Delta y$$

  - Compression injects artificial energy
  - Instability arises if imbalanced perturbation present in KE and PE

Adjusting the resulting energy errors by changing Δh, Δt: smaller energy perturbation leads to less error growth

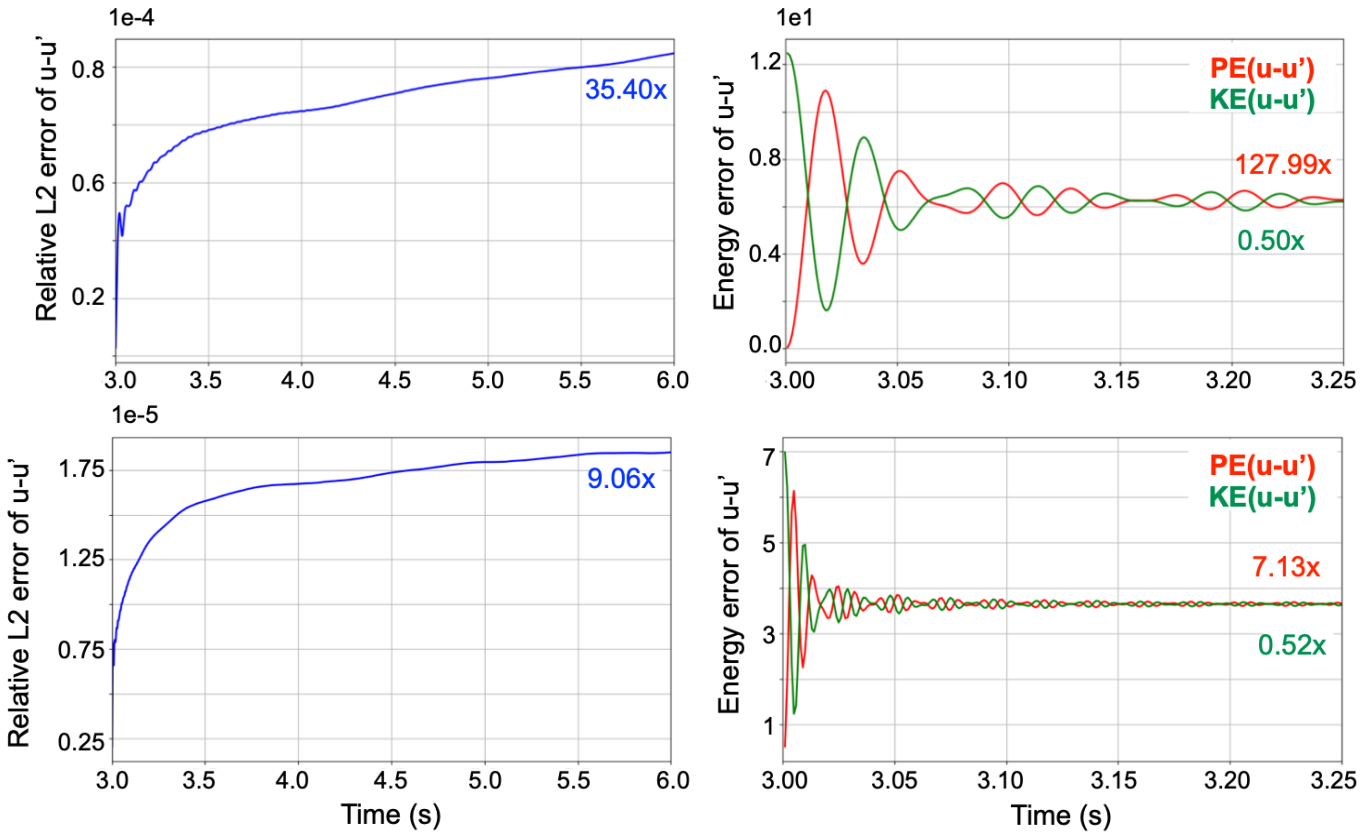

Energy-conserved Compression for Checkpointing

- Goal: reduce u^n and u^{n-1} while minimizing the perturbation to $KE(u^n, u^{n-1})$ and $PE(u^{n-\frac{1}{2}})$ → balancing $KE(e^n, e^{n-1})$ and $PE(e^{n-\frac{1}{2}})$

- Define: $u^D = \frac{1}{2}(u^n - u^{n-1})$, and $u^A = \frac{1}{2}(u^n + u^{n-1})$
  - Total energy becomes: $KE(u^n, u^{n-1}) + PE(u^{n-\frac{1}{2}}) = KE(u^D) + PE(u^A)$
  - Now the goal becomes choosing the compressed version of $\tilde{u}^D$ and $\tilde{u}^A$, which minimize the perturbation and balance in energy measured by $KE(u^D - \tilde{u}^D) + PE(u^A - \tilde{u}^A)$,

- The reconstruction can be obtained through:

$$\tilde{u}^n = \tilde{u}^D + \tilde{u}^A \text{ and } \tilde{u}^{n-1} = \tilde{u}^A - \tilde{u}^D$$

Error Control via S-norm

- We will compress u^D and u^A associating with KE and PE using separate norms:

$$\epsilon_{PE} = \frac{1}{h^2}\|u^A - \tilde{u}^A\|_1 \text{ and } \epsilon_{KE} = \frac{1}{c^2\Delta t^2}\|u^D - \tilde{u}^D\|_0$$

- MGARD defines a series of s-norms:
  - If we know how we'll use the reduced data, we can compress using appropriate norms
  - The choice of s (i.e., norm) affects how data is compressed → s=0 minimizing L^2 error (good for KE)and s=1 equivalent to the norm of PE
- Balancing perturbations in KE and PE yields:
  - $\tau_{KE} : \tau_{PE} = c^2\Delta t^2 C_{PE} : h^2$
  - $\|u^D - \tilde{u}^D\|_0 \leq \tau_{KE}$ and $\|u^A - \tilde{u}^A\|_1 \leq C_{PE}\,\tau_{PE}$

Error preservation using MGARD s=0 and s=1

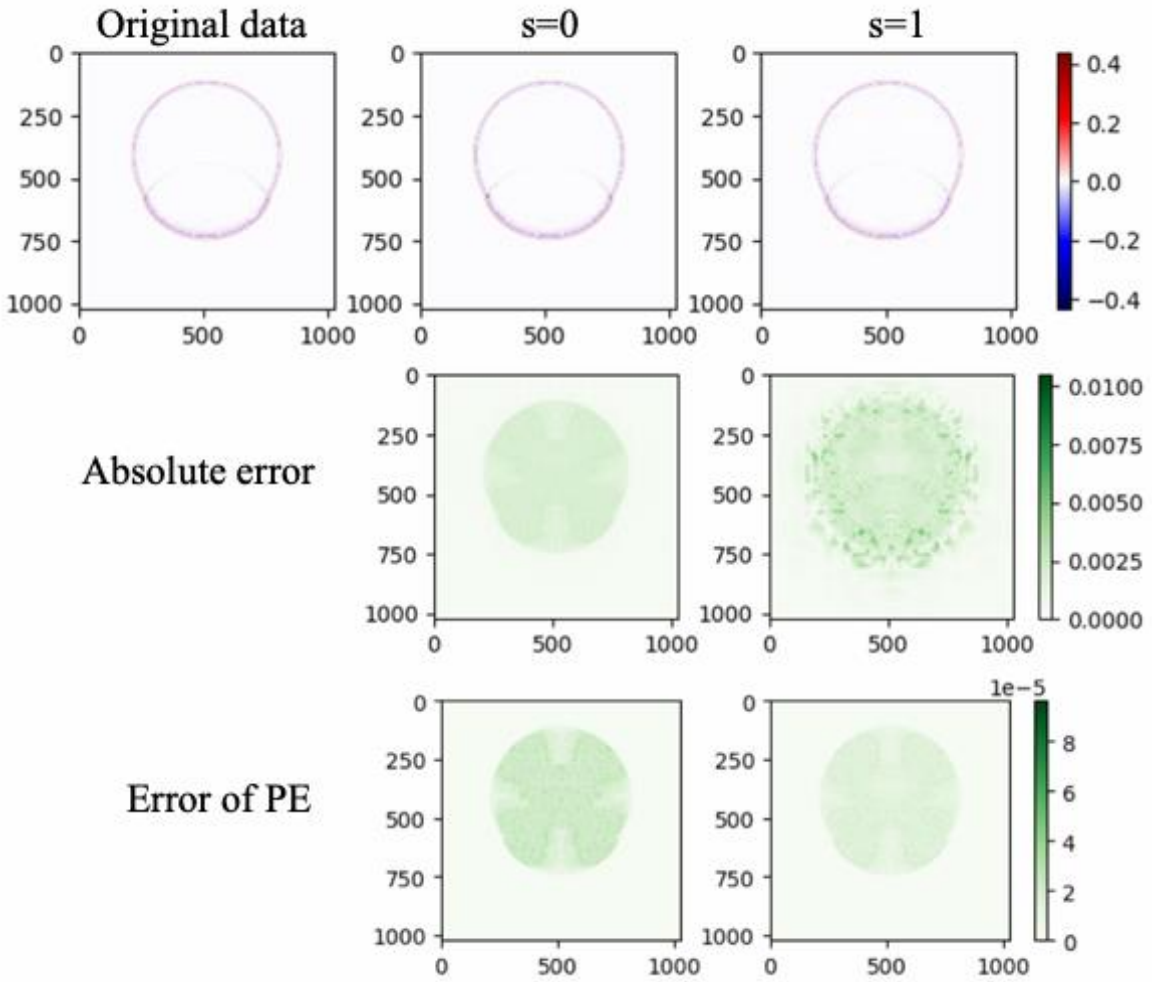

Design Overview

- Stability-controlled lossy checkpointing pipeline

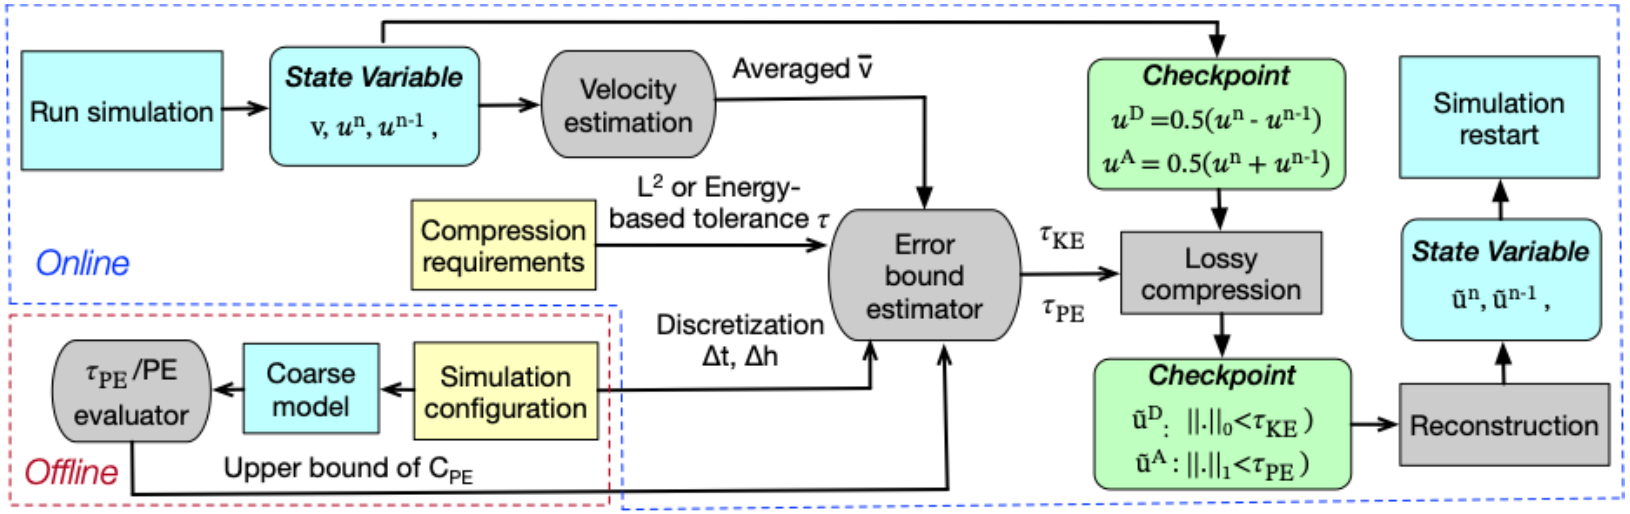

- Implication of L^2 errors in u^n
  - $\|u^n - \tilde{u}^n\|_{l^2} = \|(u^D - \tilde{u}^D) + (u^A - \tilde{u}^A)\|_{l^2} \leq \tau_{kE} + \|u^A - \tilde{u}^A\|_{l^2}$
  - We further prove $\|u^A - \tilde{u}^A\|_{l^2} \leq \gamma \|u^A - \tilde{u}^A\|_1$, where $\gamma = \frac{M^2}{\pi^2}$

Detailed proof in the paper

Algorithm 1: Error tolerance assignment

Input: Requested eb_{TYPE} and τ_n, C_{PE}, grid points N, physical domain size M, averaged velocity $\bar{c}$

if $eb_{TYPE} \in L^2$ **then**
  $\tau_{\mathrm{PE}} \leftarrow \tau_n \Big/ \left(\frac{M^2}{\pi^2} + \frac{\bar{c}^2 \Delta t^2 C_{\mathrm{PE}}}{h^2} \right)$;
  $\tau_{\mathrm{KE}} \leftarrow \tau_{\mathrm{PE}} \frac{\bar{c}^2 \Delta t^2 C_{\mathrm{PE}}}{h^2}$;
else if $eb_{TYPE} \in \tau_{KE}$ **then**
  $\tau_{\mathrm{PE}} \leftarrow \tau_{\mathrm{n}} \frac{h^2}{\bar{c}^2 \Delta t^2 C_{\mathrm{PE}}}$;
  $\tau_{\mathrm{L2}} \leftarrow \frac{M^2}{\pi^2} \tau_{PE} + \tau_{\mathrm{n}}$;
else if $eb_{TYPE} \in \tau_{PE}$ **then**
  $\tau_{\mathrm{KE}} \leftarrow \tau_{\mathrm{n}} \frac{\bar{c}^2 \Delta t^2 C_{\mathrm{PE}}}{h^2}$;
  $\tau_{\mathrm{L2}} \leftarrow \frac{M^2}{\pi^2} \tau_n + \tau_{\mathrm{KE}}$;
else
  Input error type not supported ;
return τ_{PE}, τ_{KE}, τ_{L2};

Experiment Setup

- 2D and 3D wave propagation simulation, w/ velocity models downloaded from the OPENFWI dataset and an explosive source:

$$s(t) = -2\alpha(t - t_0)e^{-\alpha(t-t_0)^2}$$

  - Run for 4.5 sec/2 sec for 2D (4096x4096) and 3D (900x900x900) cases, lossy checkpointing, then resuming for 2.5 sec/3 sec ($\Delta t = 5 \times 10^{-4}$ and 1×10^{-3} for 2D and 3D case)
  - Evaluating post-restart error propagation: L^{∞} (SZ3, ZFP), L^2 (MGARD-l2), energy-conserving (MGARD-energy)

- Error metrics:
  - Primary data error: RMSE$(u^n - \tilde{u}^n)$
  - Kinetic energy error: $KE(u^n - \tilde{u}^n)$
  - Potential energy error: $PE(u^n - \tilde{u}^n)$

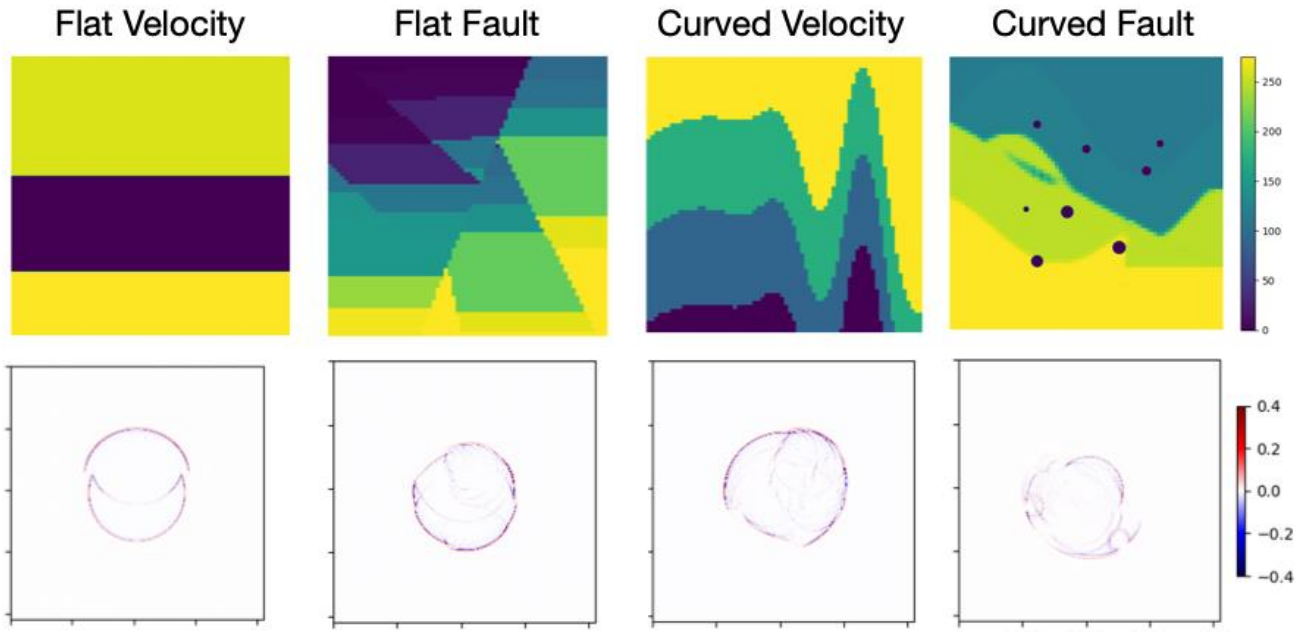

Different initial condition (Velocity)

Evaluation: Stability Preservation (L2- vs energy-based error control)

- Evaluating across different error tolerance and PDE parameters

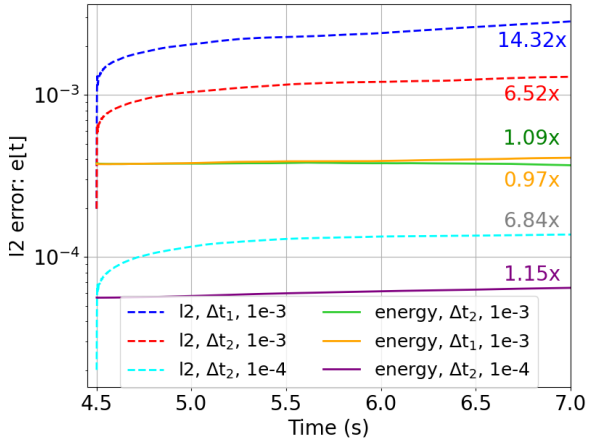

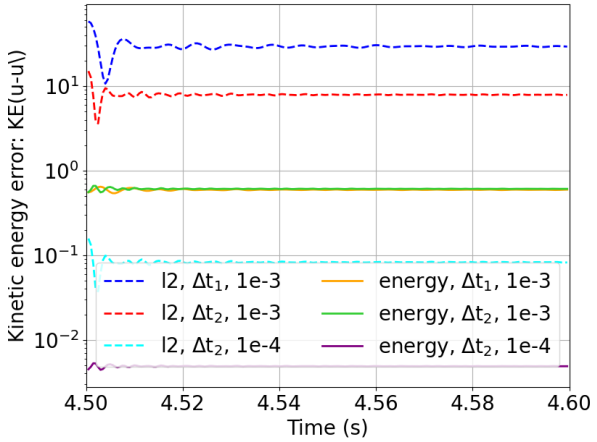

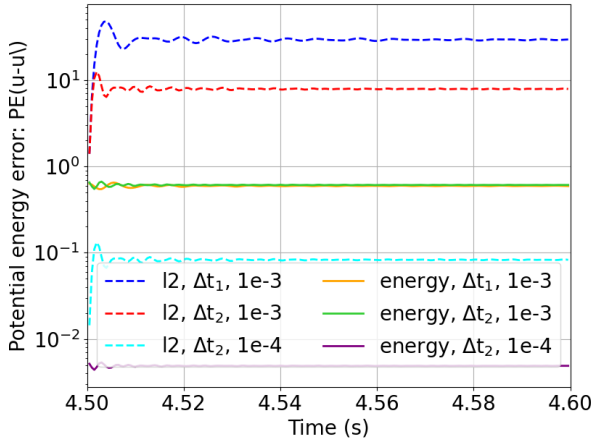

- Error propagation after a single C/R

fixing compression ratios (~70x) across compressors

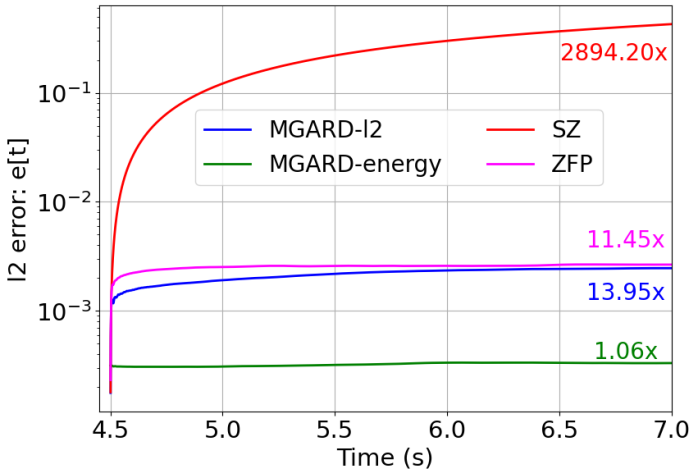

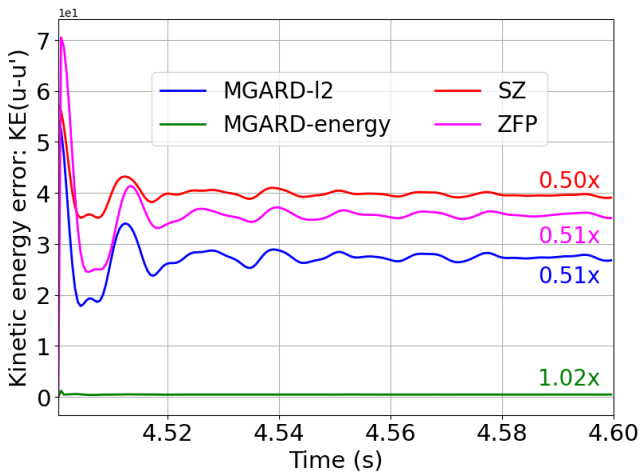

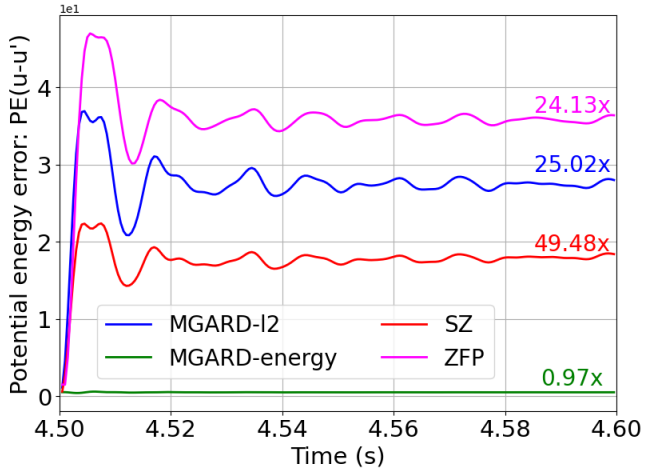

Evaluation: Stability Preservation

- Error Magnification (EM) after single C/R --- more materials **EM = final ϵ_T / ϵ_0 at restart**

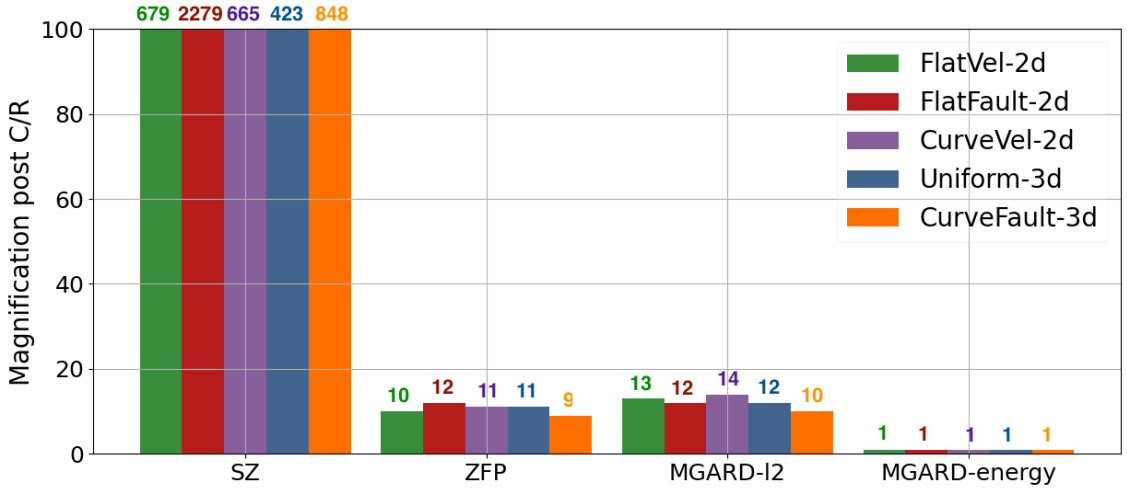

RMSE

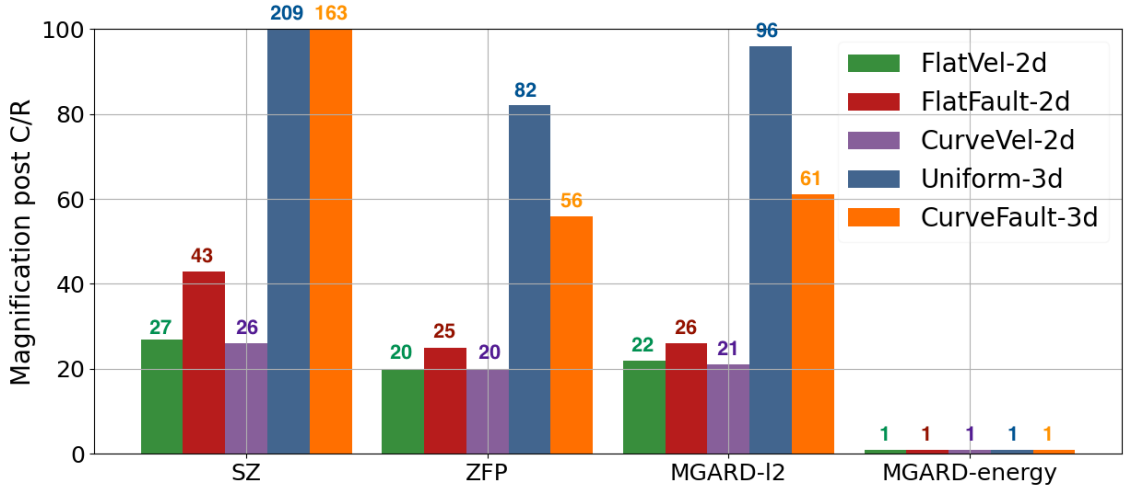

PE

- Error propagation after multiple C/R cycles:
  - With appropriate stability control, $\epsilon_p \leq p\tau$, where ϵ_p is the final error after p C/R cycles, and τ is the prescribed error tolerance used at each compression

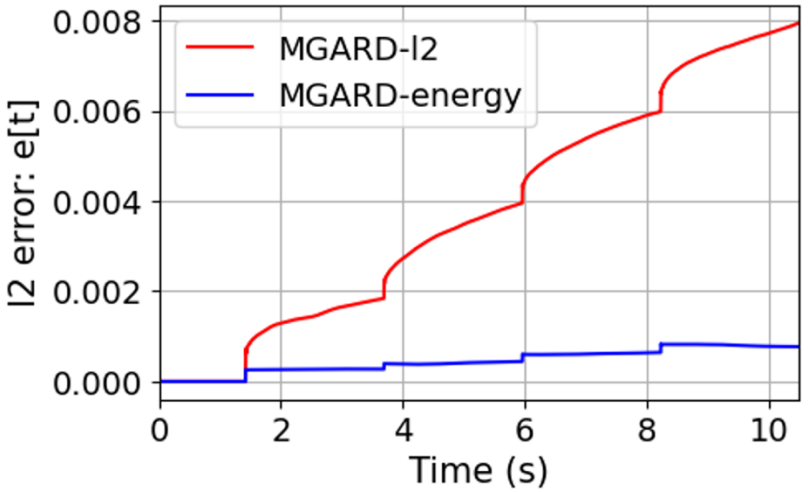

Evaluation: Checkpointing I/O Acceleration

- The I/O acceleration is achieved from the tighter error control ($EM \approx 1$)
  - When target at the same final error (ϵ_T), the stability preserving compression technique allows to prescribe larger error bounds (smaller compressed data), thus achieving larger compressibility than compression based on L^2 or L^∞ norm
  - Evaluation dataset: 4 GB and 17 GB per GPU
  - Targeting relative RMSE: 2×10^{-6}
  - Compression and decompression are performance than uncompressed (4-6x in reads, 3-8x in writes) and checkpoints compressed using GPU-accelerated MGARD (up to 1.8-2.2x)

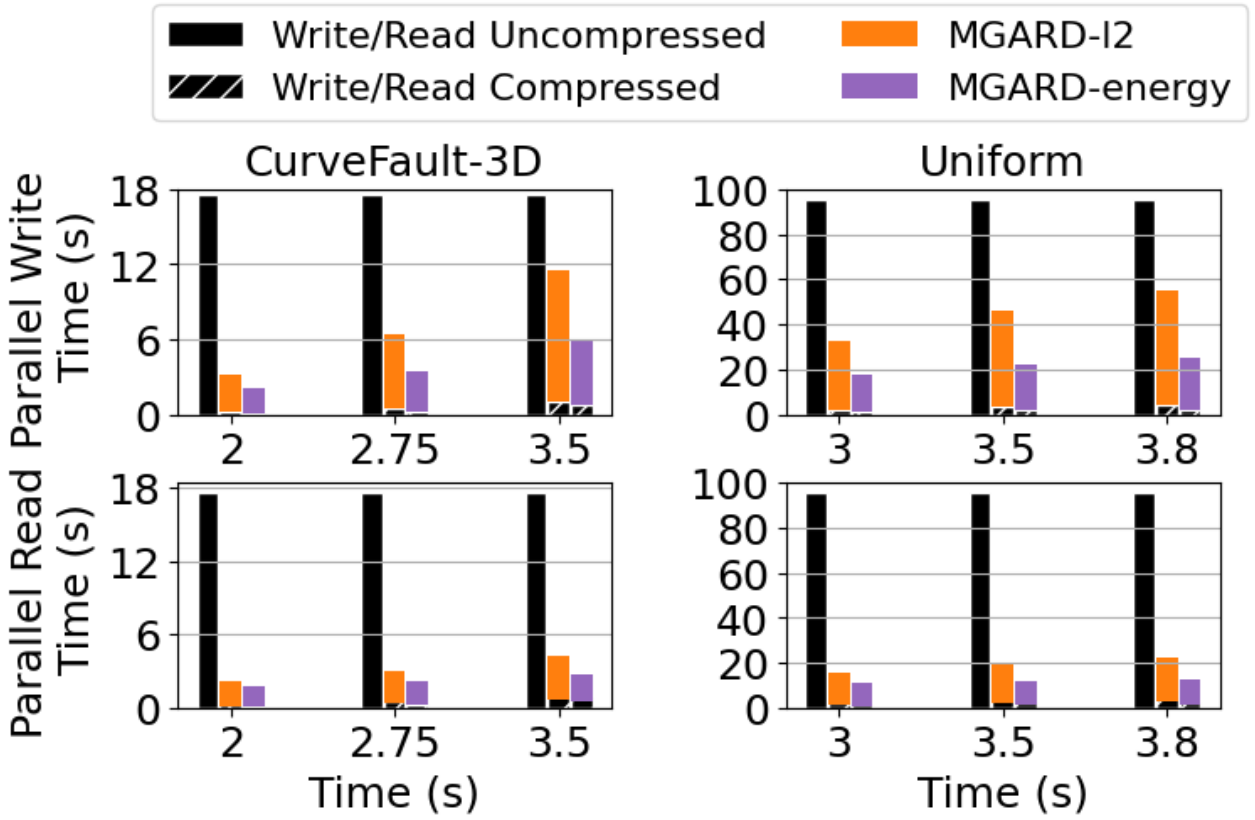

1,024 Frontier nodes

Conclusion & Future Work

- Summary:
  - We showed that checkpoint compression must obey physical conservation or else the post-restarted CFD becomes unstable
  - We incorporate the physics of second-order wave equation during checkpoint compression by separating two energy components and compressing each with distinct error controls
  - The proposed compression technique significantly reduced I/O and storage costs for checkpointing in large-scale PDE simulations while maintaining numerical stability.

- Looking ahead:
  - Extend the methodology beyond the second-order linear wave equation
  - Explore different compression strategies to ensure long-term stability and accuracy
  - Integration with modern checkpointing frameworks for deployment in production environment