Radiation Hydrodynamics at Scale with FleCSI

[1] U Texas at Austin
[2] Michigan State University
[3] California Institute of Technology
[4] U Tennessee, Knoxville
[5] University of Stuttgart
[6] Florida State University

CO-DESIGN SUMMER SCHOOL 2024, CCS-7

Mammadbaghir Baghirzade[1], Shihab Shahriar Khan[2], Yoonsoo Kim[3], Sudarshan Neopane[4], Alexander Strack[5], Farhana Taiyebah[6]

Physics

RADIATION HYDRODYNAMICS

- Co-moving frame formulation

$$\partial_t \begin{bmatrix} \rho \\ \rho\mathbf{v} \\ e \\ E \end{bmatrix} + \nabla\cdot \begin{bmatrix} \rho\mathbf{v} \\ \rho\mathbf{v}\mathbf{v} + p\mathbf{I} \\ (e+p)\mathbf{v} \\ E\mathbf{v} \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{f}_r \;\; \text{radiation force} \\ \text{heating/cooling} \;\; \dot{q} + \mathbf{f}_r\cdot\mathbf{v} \\ -\nabla\cdot\mathbf{F} - \underbrace{\mathbf{P}\cdot\nabla\mathbf{v}}_{\text{photon-tiring}} - \dot{q} \end{bmatrix}$$

ρ : mass density / $\vec{v}$: fluid velocity / e : fluid energy density
E : radiation energy density / F : radiation energy flux / P : radiation pressure tensor

FLUX-LIMITED DIFFUSION

- Employs "bridge law" to emulate optically thin regime beyond the diffusion limit
- Computationally efficient proxy of realistic microphysics

- Radiation energy flux :

$$F^i \approx -\lambda\nabla E$$

'flux limiter' for causal constraint

- Radiation pressure tensor :

$$P^{ij} \rightarrow \left\{ \begin{array}{cc} E(\delta_{ij}/3) & \text{optically thick} \\ (\nabla E)^i(\nabla E)^j/(\nabla E)^2 & \text{optically thin} \end{array} \right\}$$

Implementation

NUMERICAL ALGORITHMS

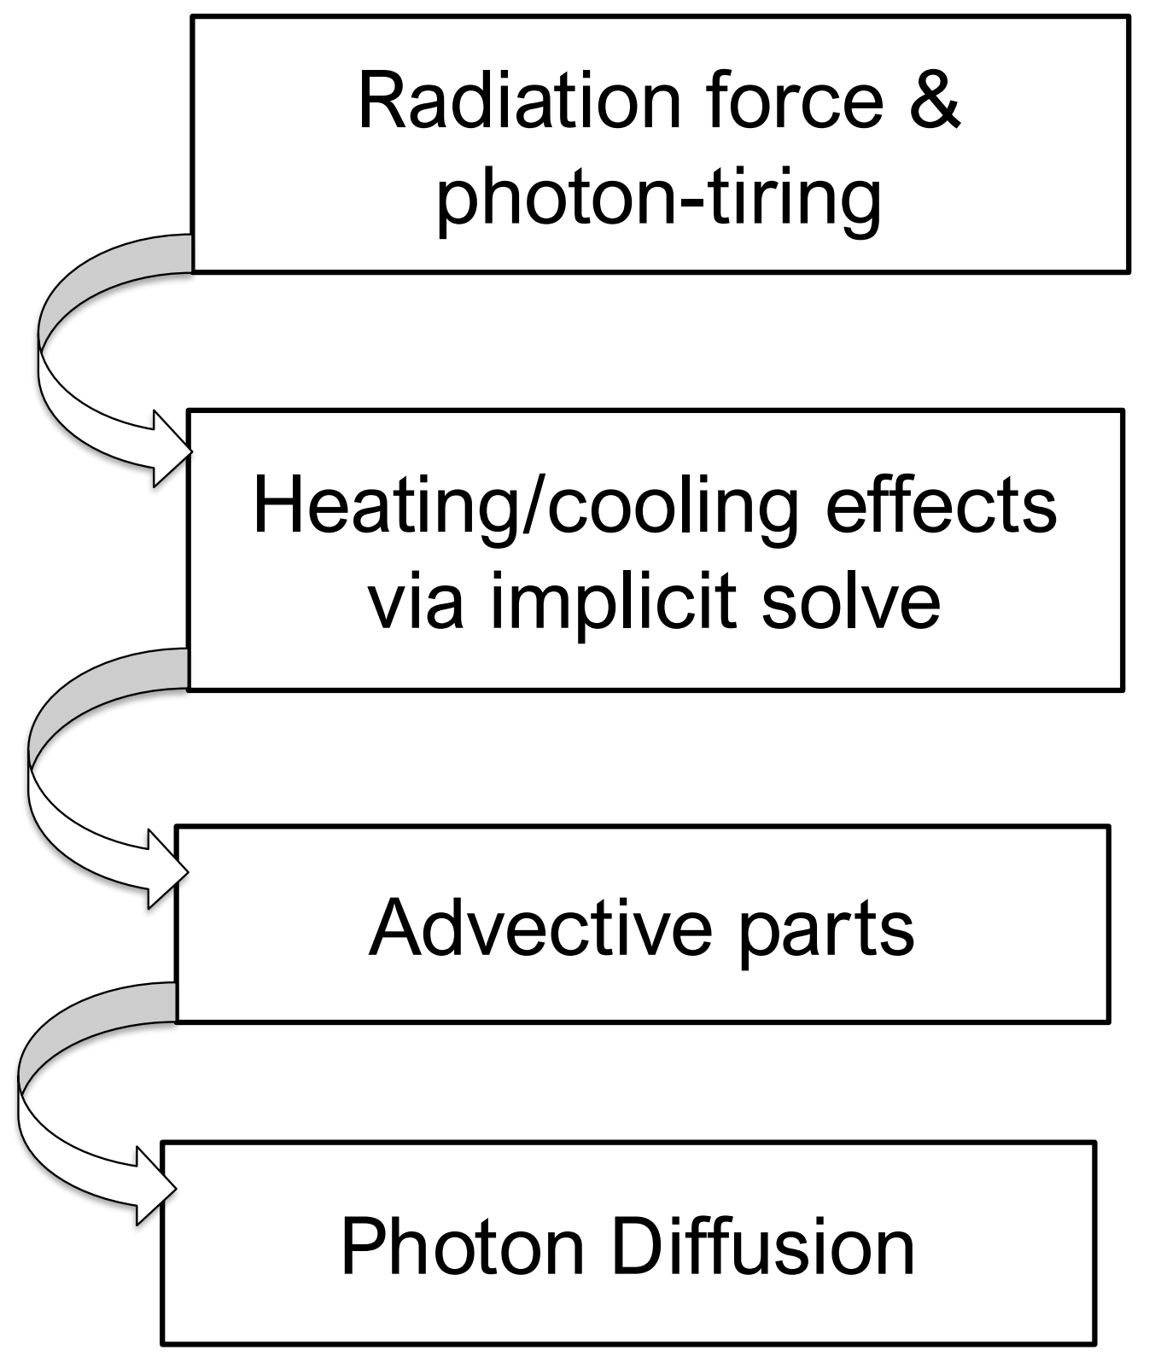
Radiation force & photon-tiring
Heating/cooling effects via implicit solve
Advective parts
Photon Diffusion

MULTIGRID SOLVER

- To beat down different frequency residuals.

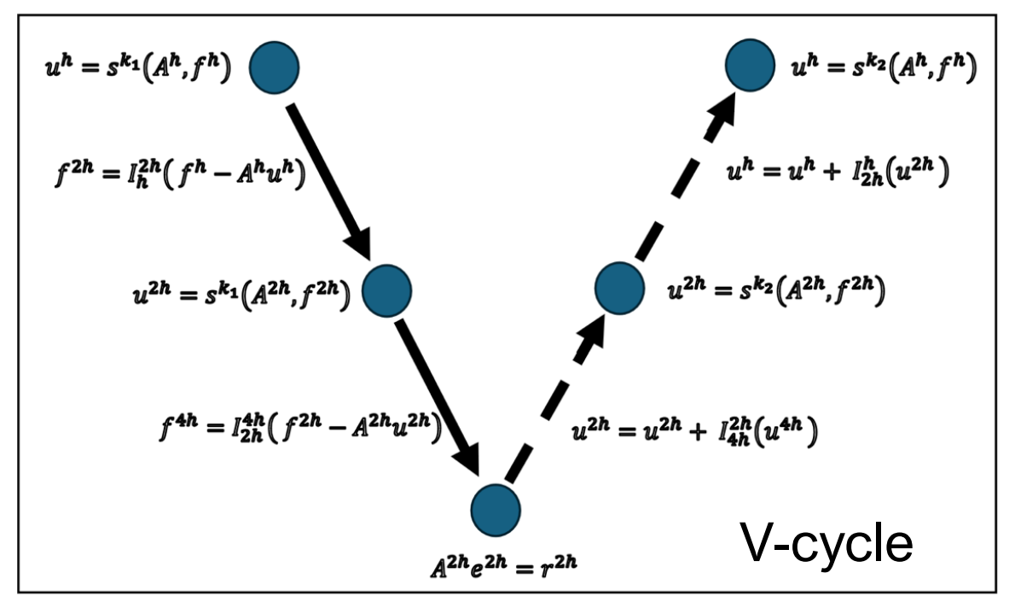

Red-Black Gauss-Seidel Smoother

- For better performance on parallel architecture.

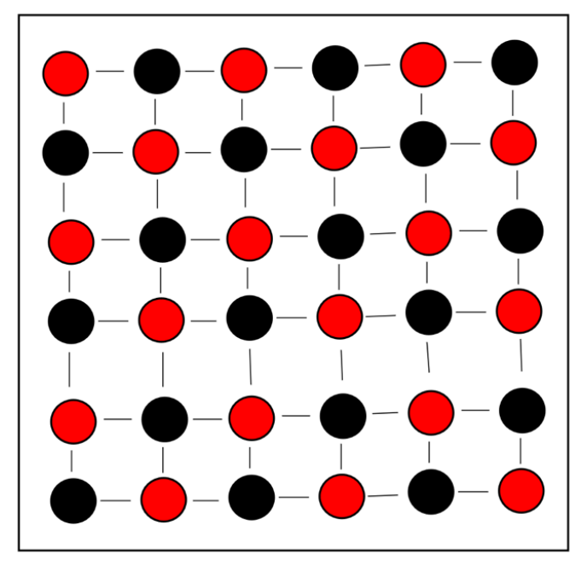

PARALLELIZATION & ACCELERATION

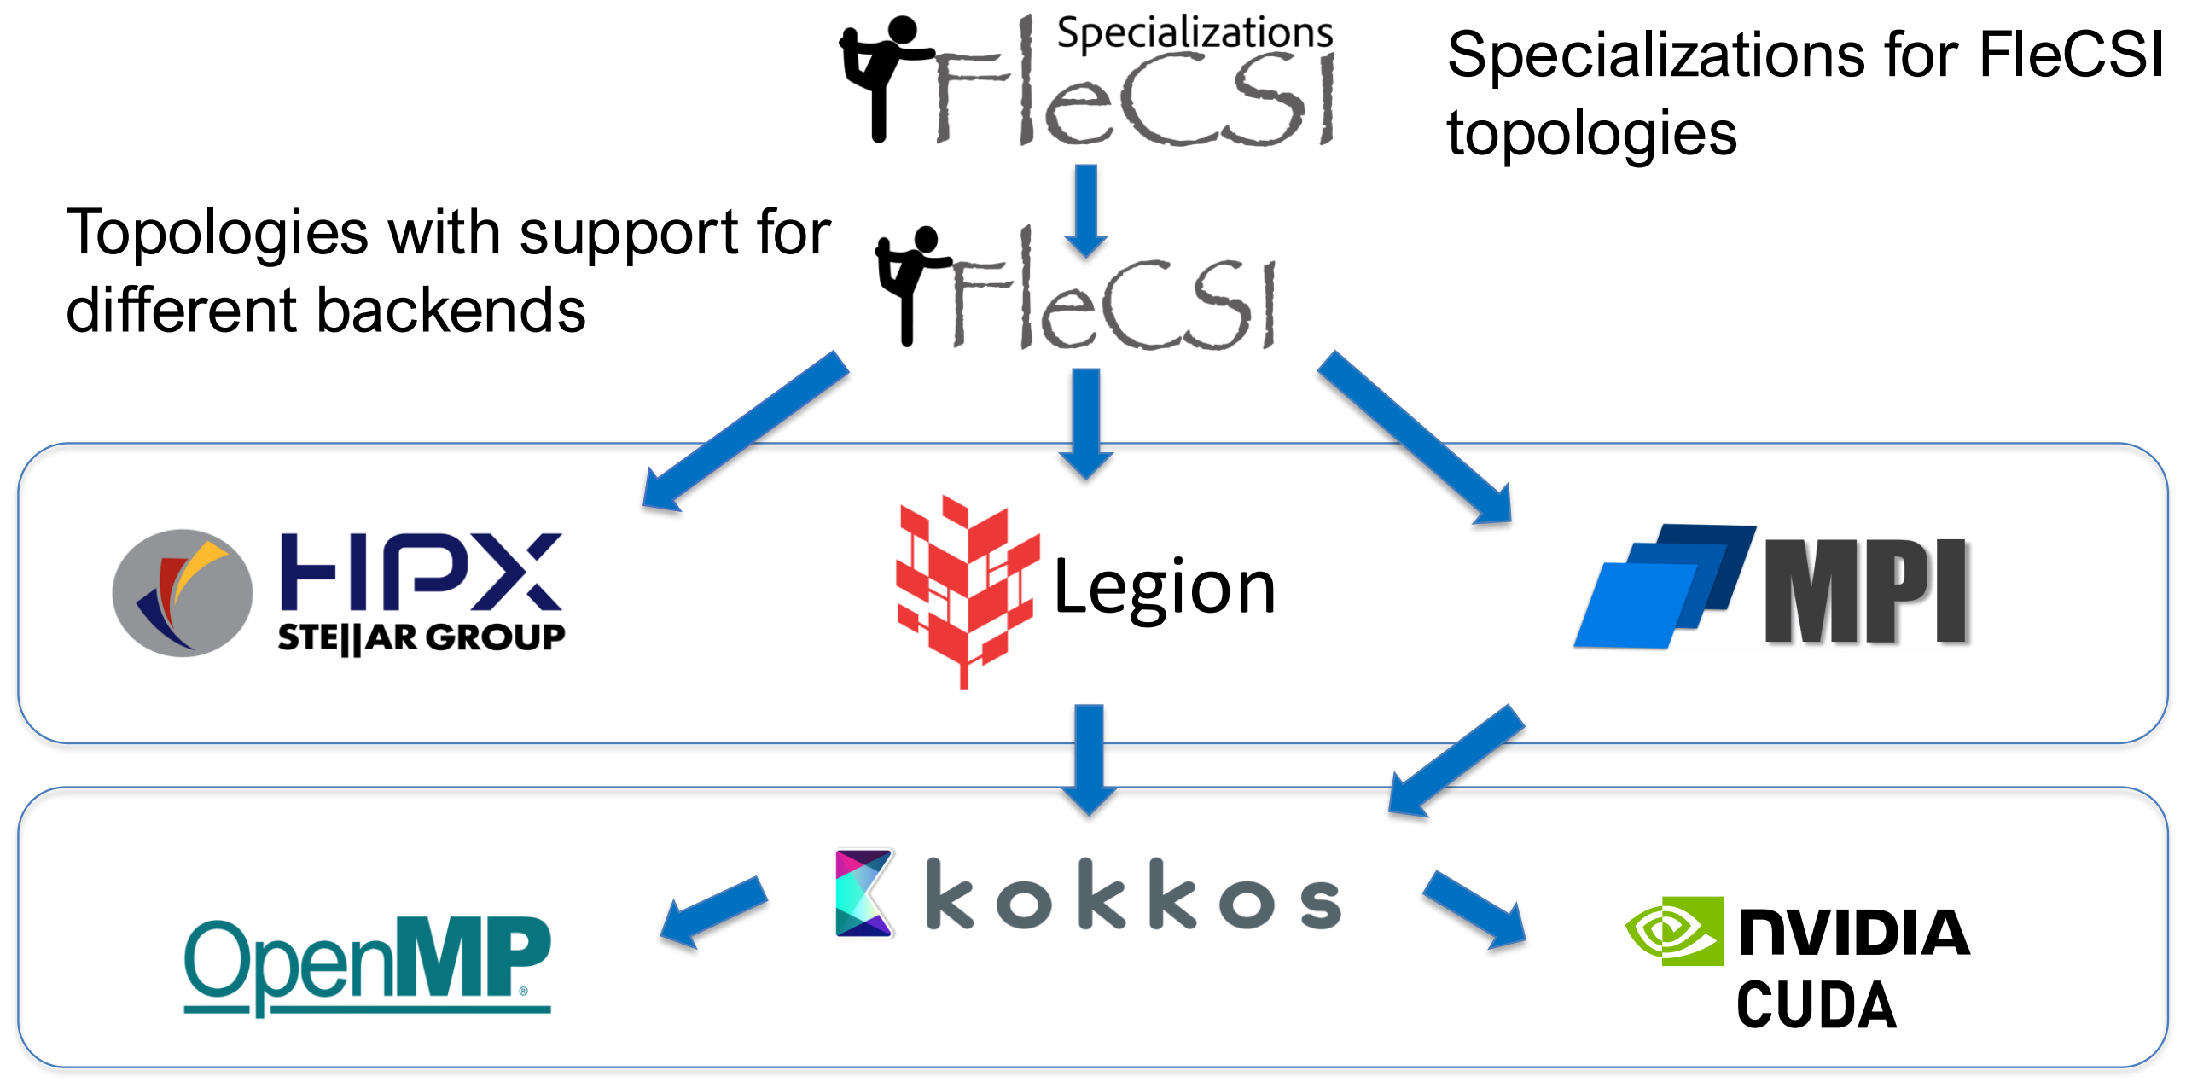

Specializations for FleCSI topologies

Topologies with support for different backends

Abstract

- ✓ **Hydrodynamic And Radiative Diffusion (HARD) code** is the first publicly available code coupling radiation and hydrodynamics that is highly scalable and portable to heterogenous HPC architectures.
- ✓ Radiation hydrodynamics plays a crucial role in high energy density physics (e.g. inertial confinement fusion, core collapse supernova, neutron transport).
- ✓ FleCSI adds support for multiple backends and makes it easy to target different architectures. It also facilitates the addition of new complex physics.

Scalable & Portable Radiation Hydrodynamics code

RADIATIVE RAYLEIGH-TAYLOR INSTABILITY

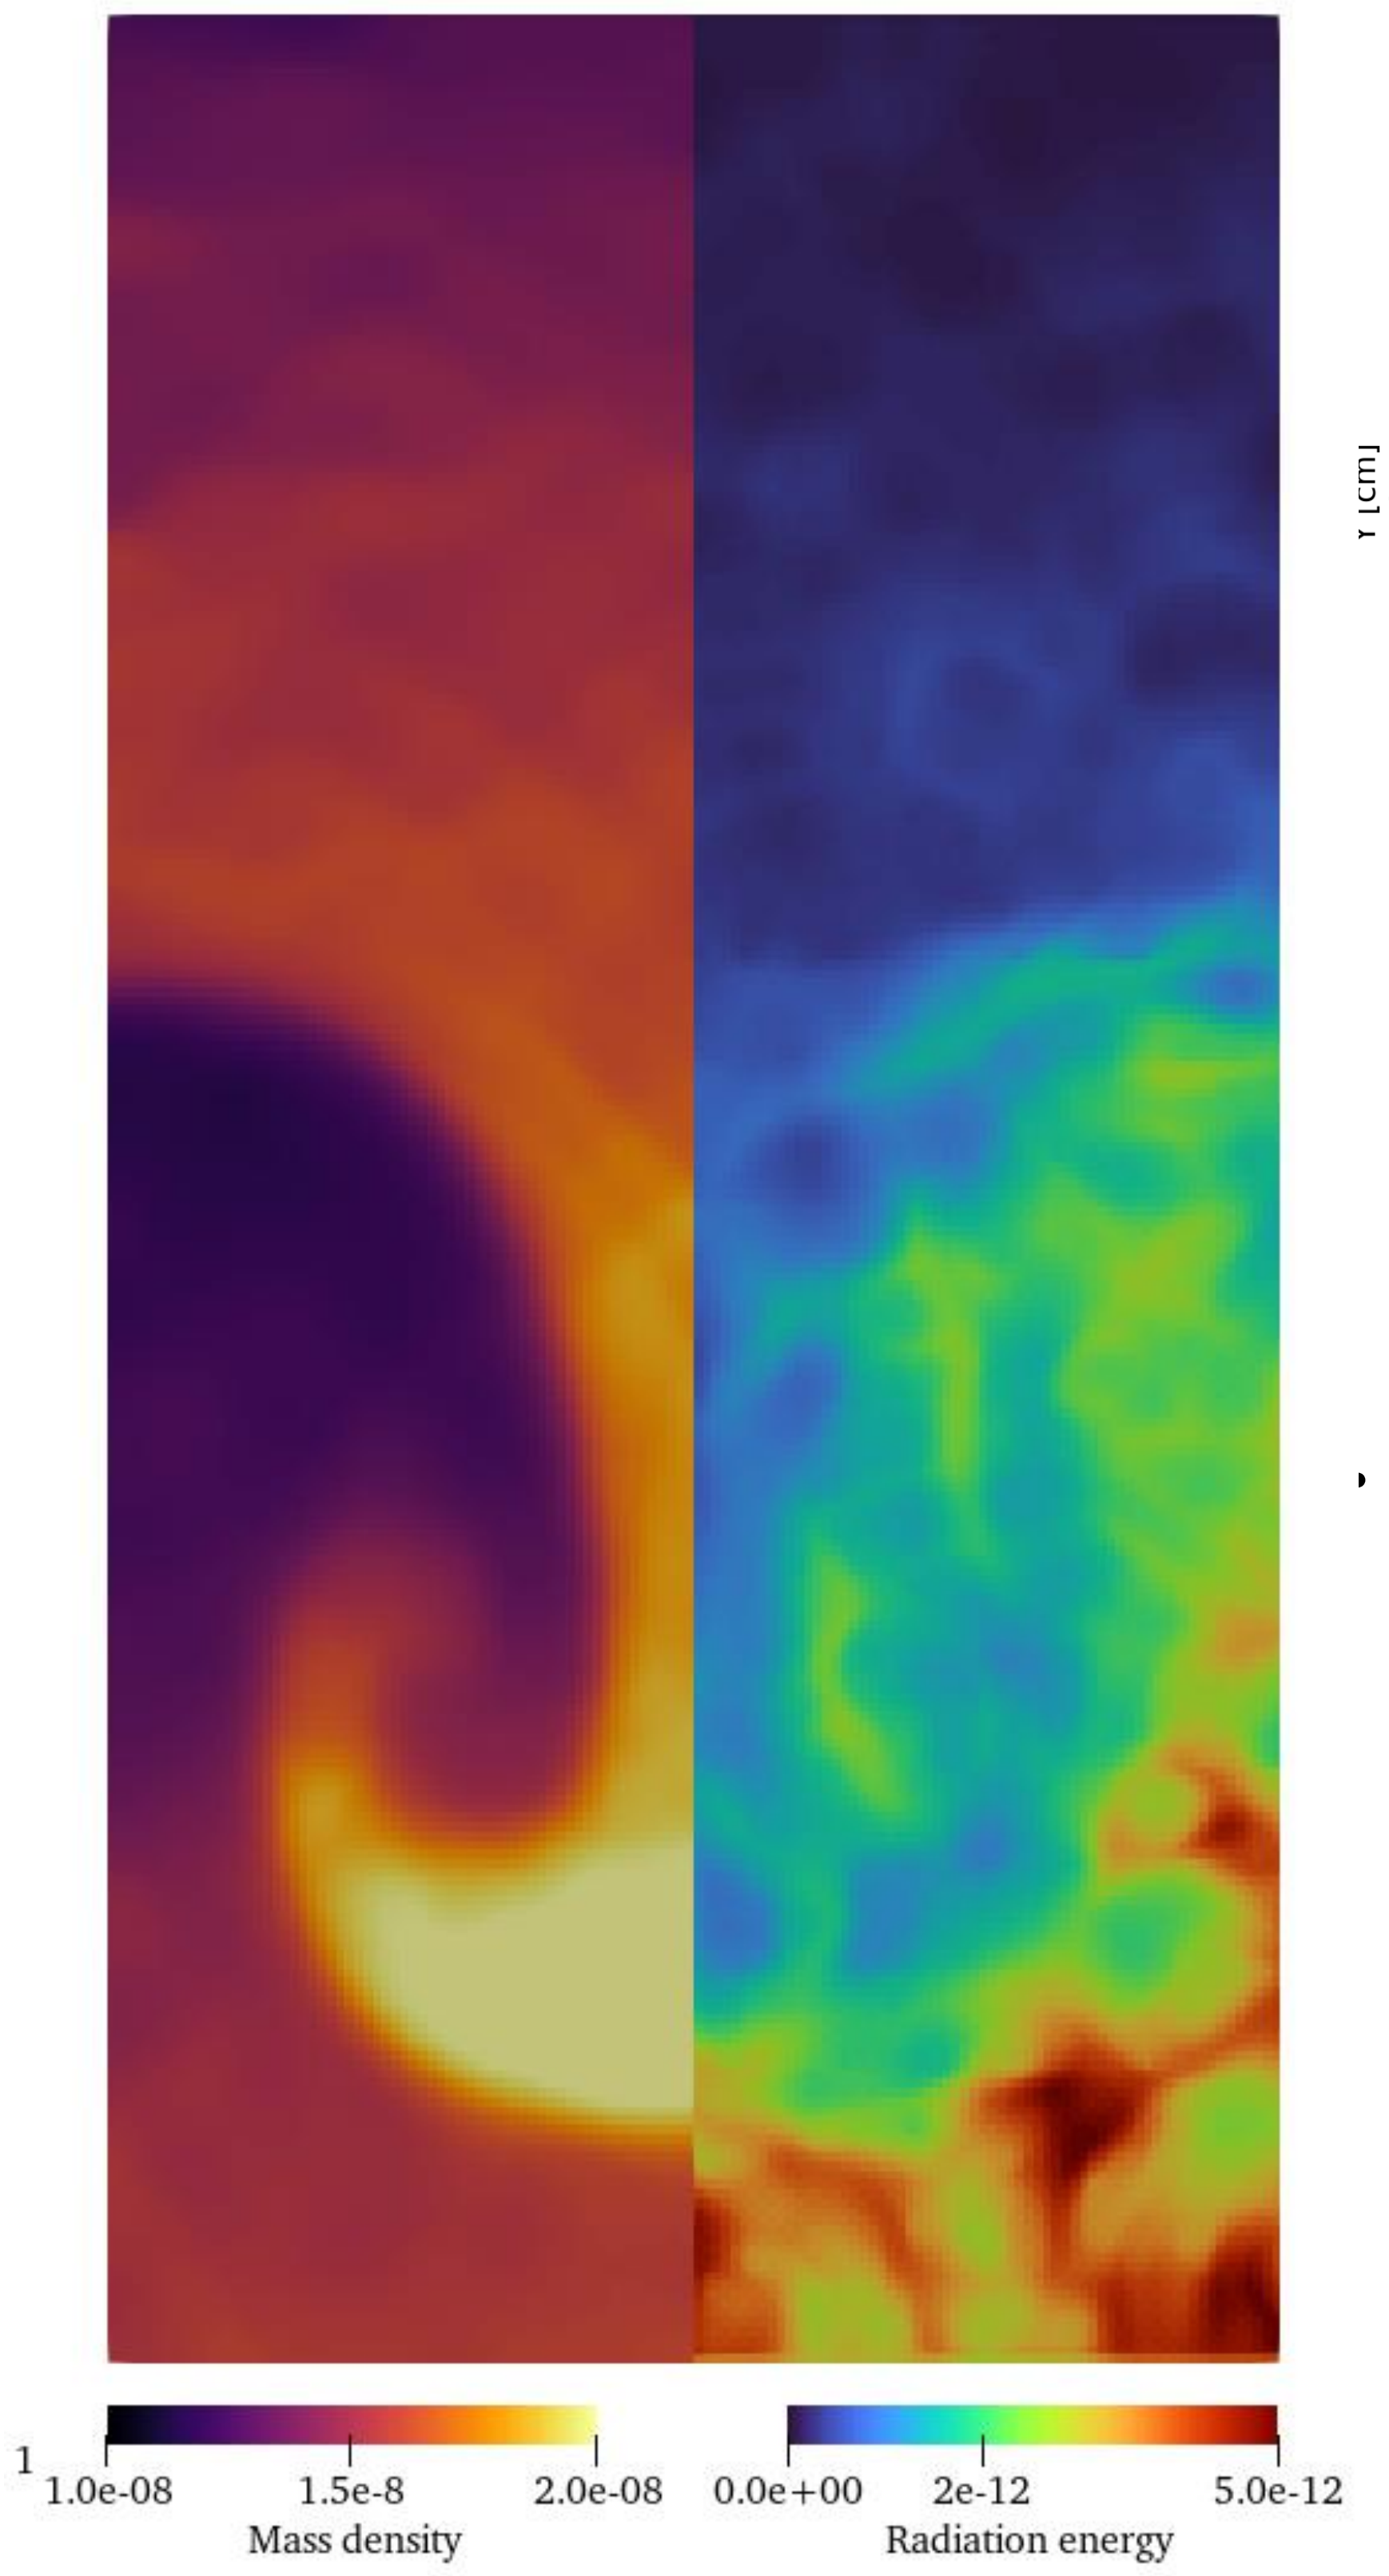

RADIATION PRESSURE-DRIVEN EXPANSION

- Radiation energy density was set to be decreasing toward the edge of each shapes

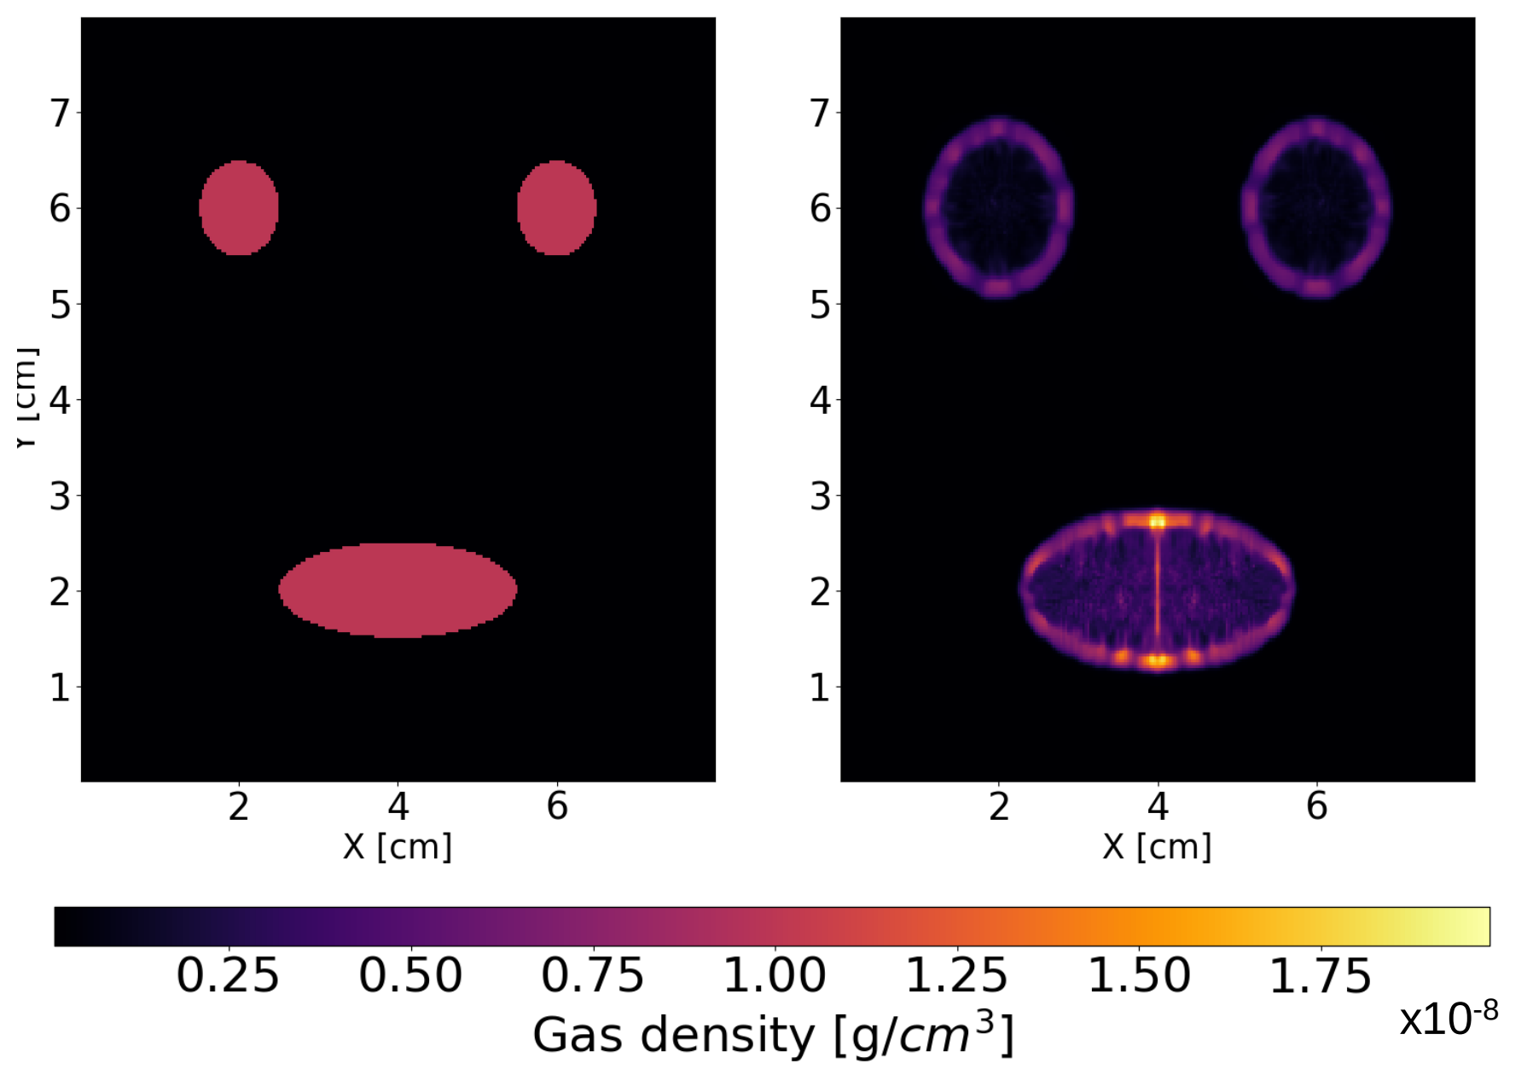

Initial mass density distribution (left) and at t = 4 μs (right)

SHADOW TEST

- Formation of shadow behind optically thick cylinder

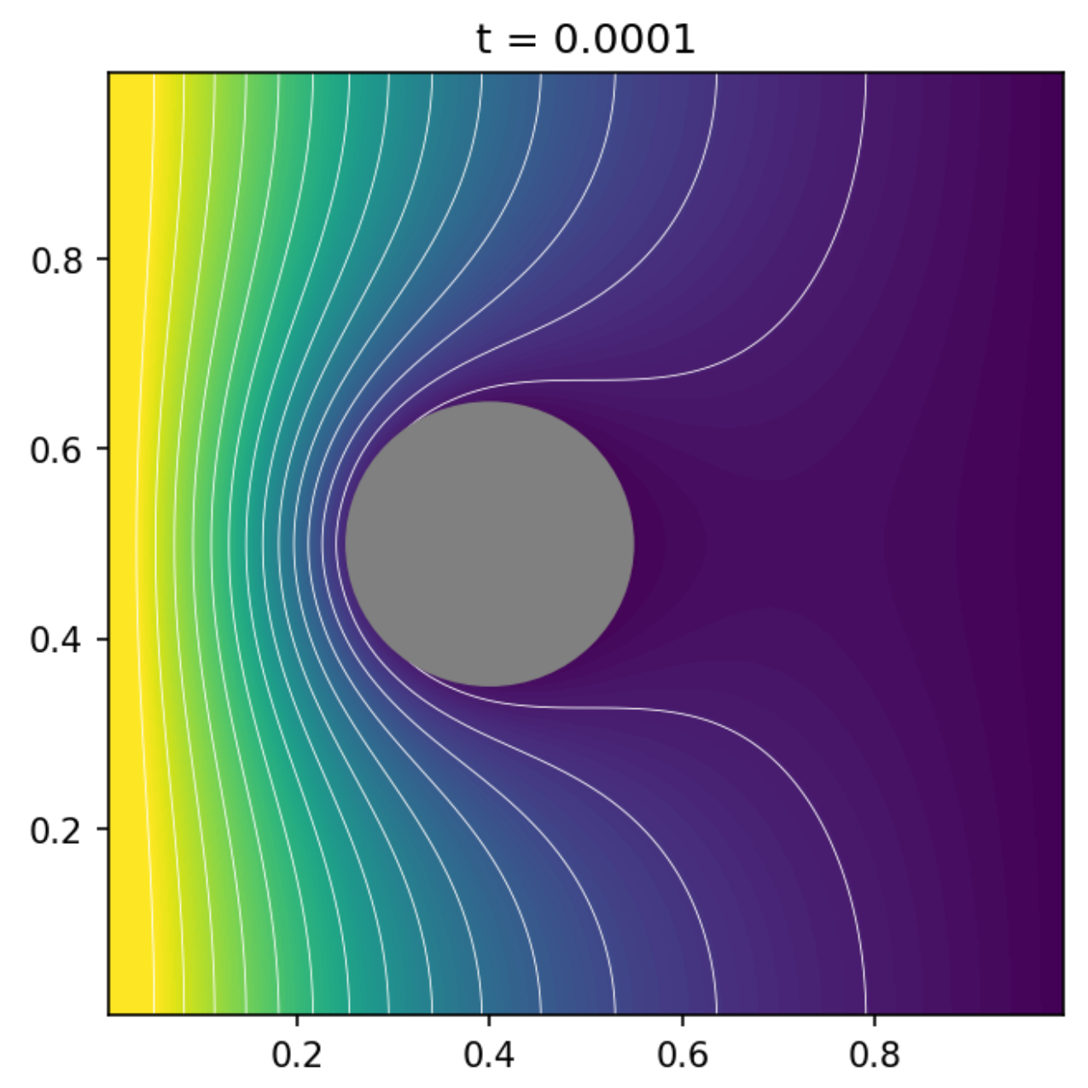

Performance

STRONG SCALING

Strong scaling runtimes on up to 16 nodes on the Darwin cluster for a 1D Rankine-Hugoniot problem using 2^{27} cells. Each node contains a dual-socket Intel Xeon E5-2695 configuration with a total of 36 physical cores.

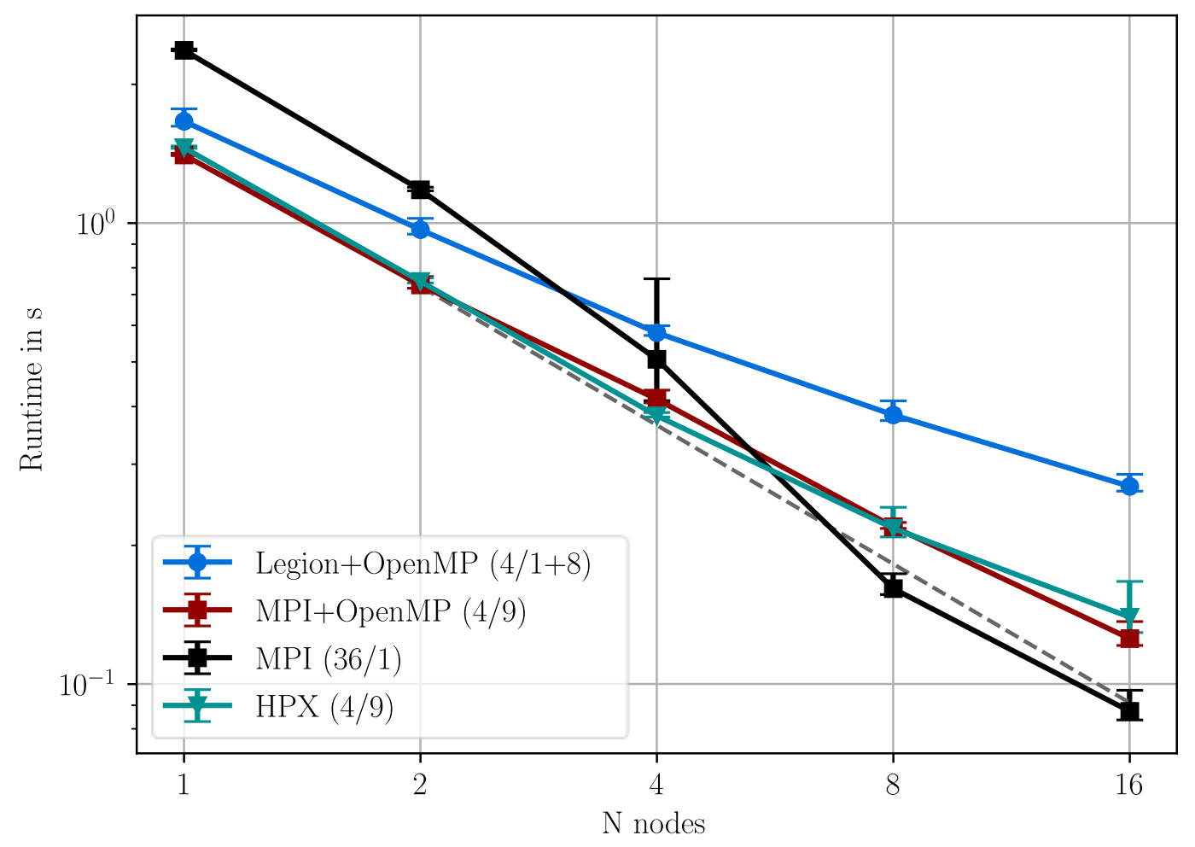

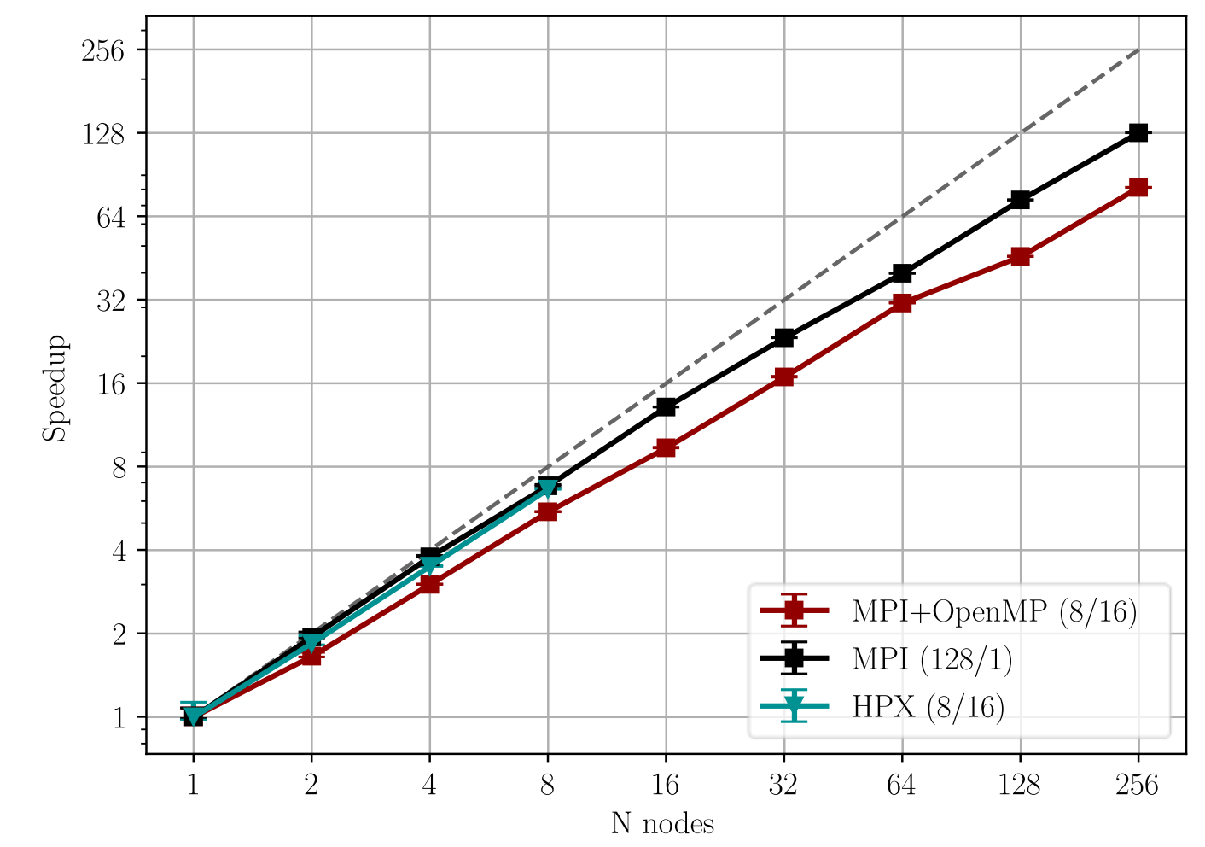

Strong scaling speedup on up to 256 nodes on the Chicoma supercomputer for a 3D Rankine-Hugoniot problem using 1024^3 cells. Each node contains a dual-socket AMD EPYC 7H12 configuration with a total of 128 physical cores.

GPU & PORTABILITY

- The hydrodynamic functionalities have been ported to GPU using FleCSI kernel task, using MPI+Kokkos backend.
- This shows significant on-node performance improvement across all mesh sizes. However, smaller GPU memory limits the maximum mesh size we can experiment per-node.
- For a mesh size of 256^3, the runtime on Darwin, using a V100 GPU was 9x faster, using a A100 GPU was 25x faster compared to other FleCSI backends.

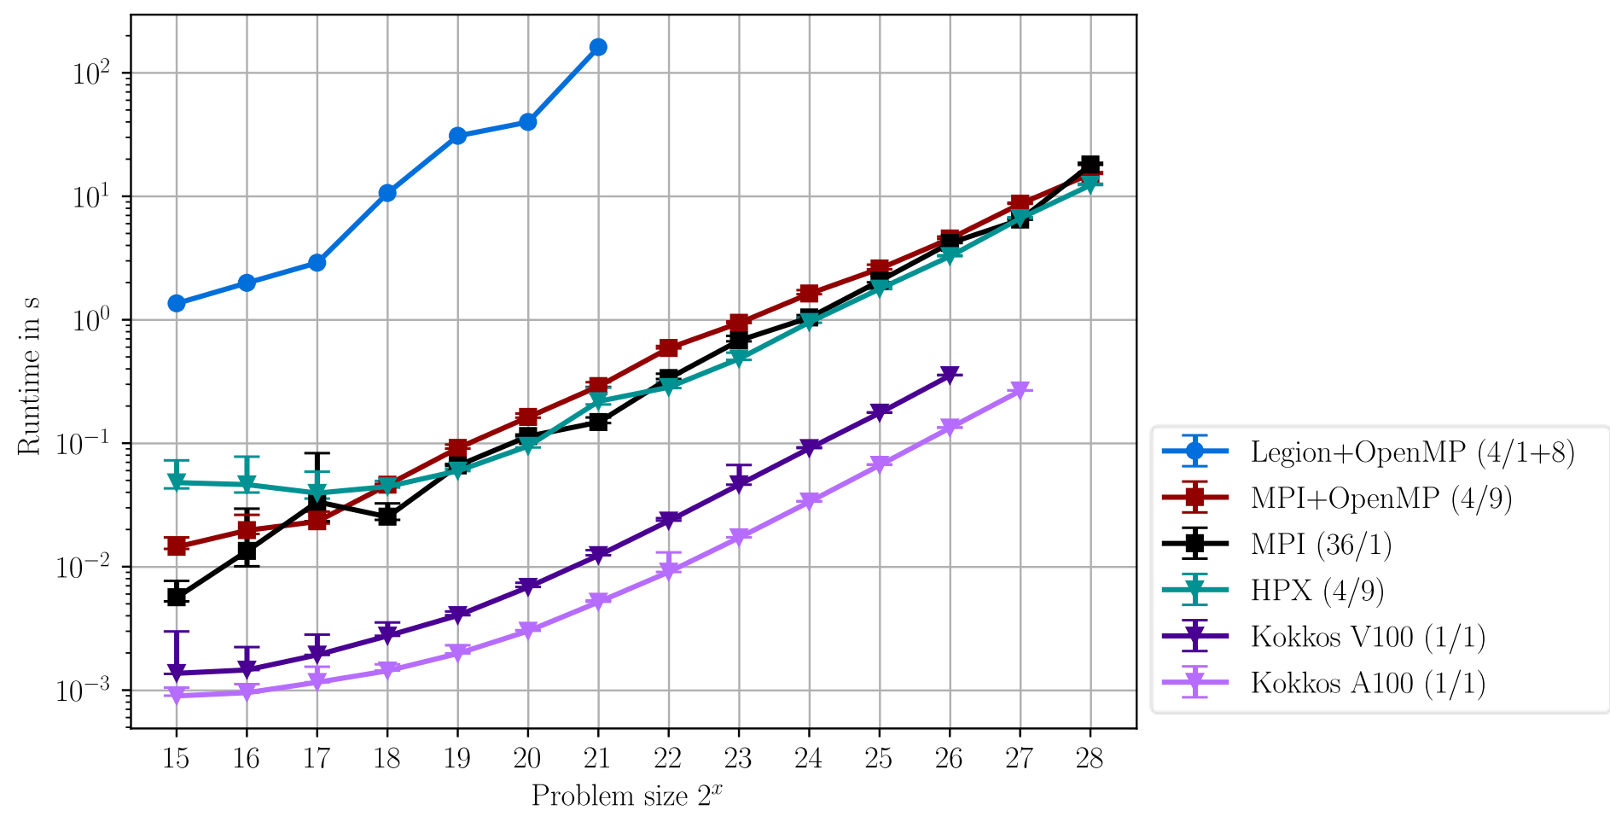

On-Node Size Scaling: GPU runtime (in purple) compared to several CPU backends.
(A100 results where obtained on Chicoma)

Summary and Future Work

- HARD code accurately captures matter-radiation coupling.
- HARD achieved near linear scalability up to 16 nodes and for multiple backends.
- Incorporate more complex schemes for matter-radiation coupling.
- GPU support for radiation-coupled hydrodynamics and multi-GPU setting.

Acknowledgment

- We acknowledge Darwin and Institutional Computing for computing resources.
- Big thanks to CCS-7 group office, Erika Maestas, Kelly Thompson, Christoph Junghans, and thanks to our mentors, Ben Bergen, Richard Berger, Brendan Krueger, Hyun Lim, Julien Loiseau, Andrés Yagüe López, Maxim Moraru, Nirmal Prajapati, Thomas Vogel, and Andrew Reisner.

LA-UR-24-27564