Ares – Simulating Type Ia Supernovae on Heterogeneous HPC Architectures

Landon Dyken [1], Alexander Holas [2, 3], Mark Ivan Ugalino [4], Nageeb Zaman [5]
(1) University of Illinois Chicago, Chicago, IL; (2) Heidelberg Institute for Theoretical Studies, Heidelberg, Germany;
(3) Zentrum für Astronomie der Universität Heidelberg, Institut für Theoretische Astrophysik, Heidelberg, Germany;
(4) University of Maryland, College Park, MD; (5) Louisiana State University, Baton Rouge, LA

Type Ia Supernovae are highly energetic thermonuclear explosions of white dwarfs,

which serve as standardizable distance markers that are essential for investigating the accelerating expansion of our Universe. The explosion physics that trigger these events are ***inherently multi-scale***, ranging from the usual diameter of a white dwarf at about 4×10^3 to 10^4 km to the carbon flame thickness $\sim 10^0$ cm, which poses a huge challenge in performing hydrodynamical simulations of these systems. To resolve the physical mechanism at every scale possible, we employ state-of-the-art **adaptive mesh refinement (AMR)** techniques within our hydro solvers.

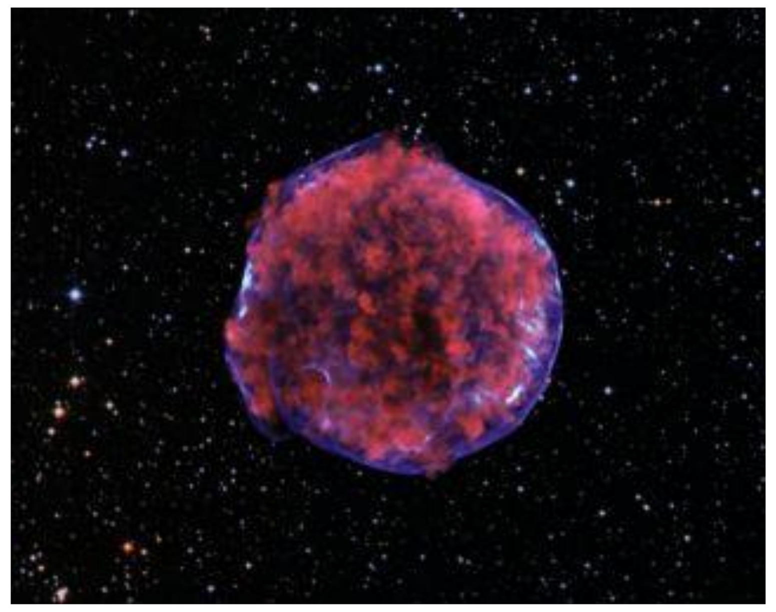

Tycho type Ia supernova remnant (SN 1572)
Credit: X-ray: NASA/CXC/Rutgers/K. Eriksen et al.; Optical: DSS

ARES

Multi-physics capabilities of Ares

Ares is built on the Parthenon framework for adaptive mesh refinement on distributed HPC clusters. We extend this framework by adding solvers for gravity, thermonuclear burning, and equation of state (EOS) necessary to simulate type Ia explosions.

Nuclear Network — To simulate the energy released through nuclear burning processes, we solve the equations of **nuclear statistical equilibrium** (NSE), where energies are high enough for strong interactions to equilibriate, for the mass fractions of nuclides.

Self-Gravity — To implement self-gravity in our simulations, we built a **monopole gravity solver**, creating a 1D gravity profile based on a shell-averaged density profile.

Equation of State — For the equation of state, we incorporated our extension of the existing **Singularity-EOS** developed at LANL **to include Helmholtz EOS** to support conditions in the degenerate gas interior of white dwarfs.

Adaptive Mesh Refinement in the Parthenon Framework

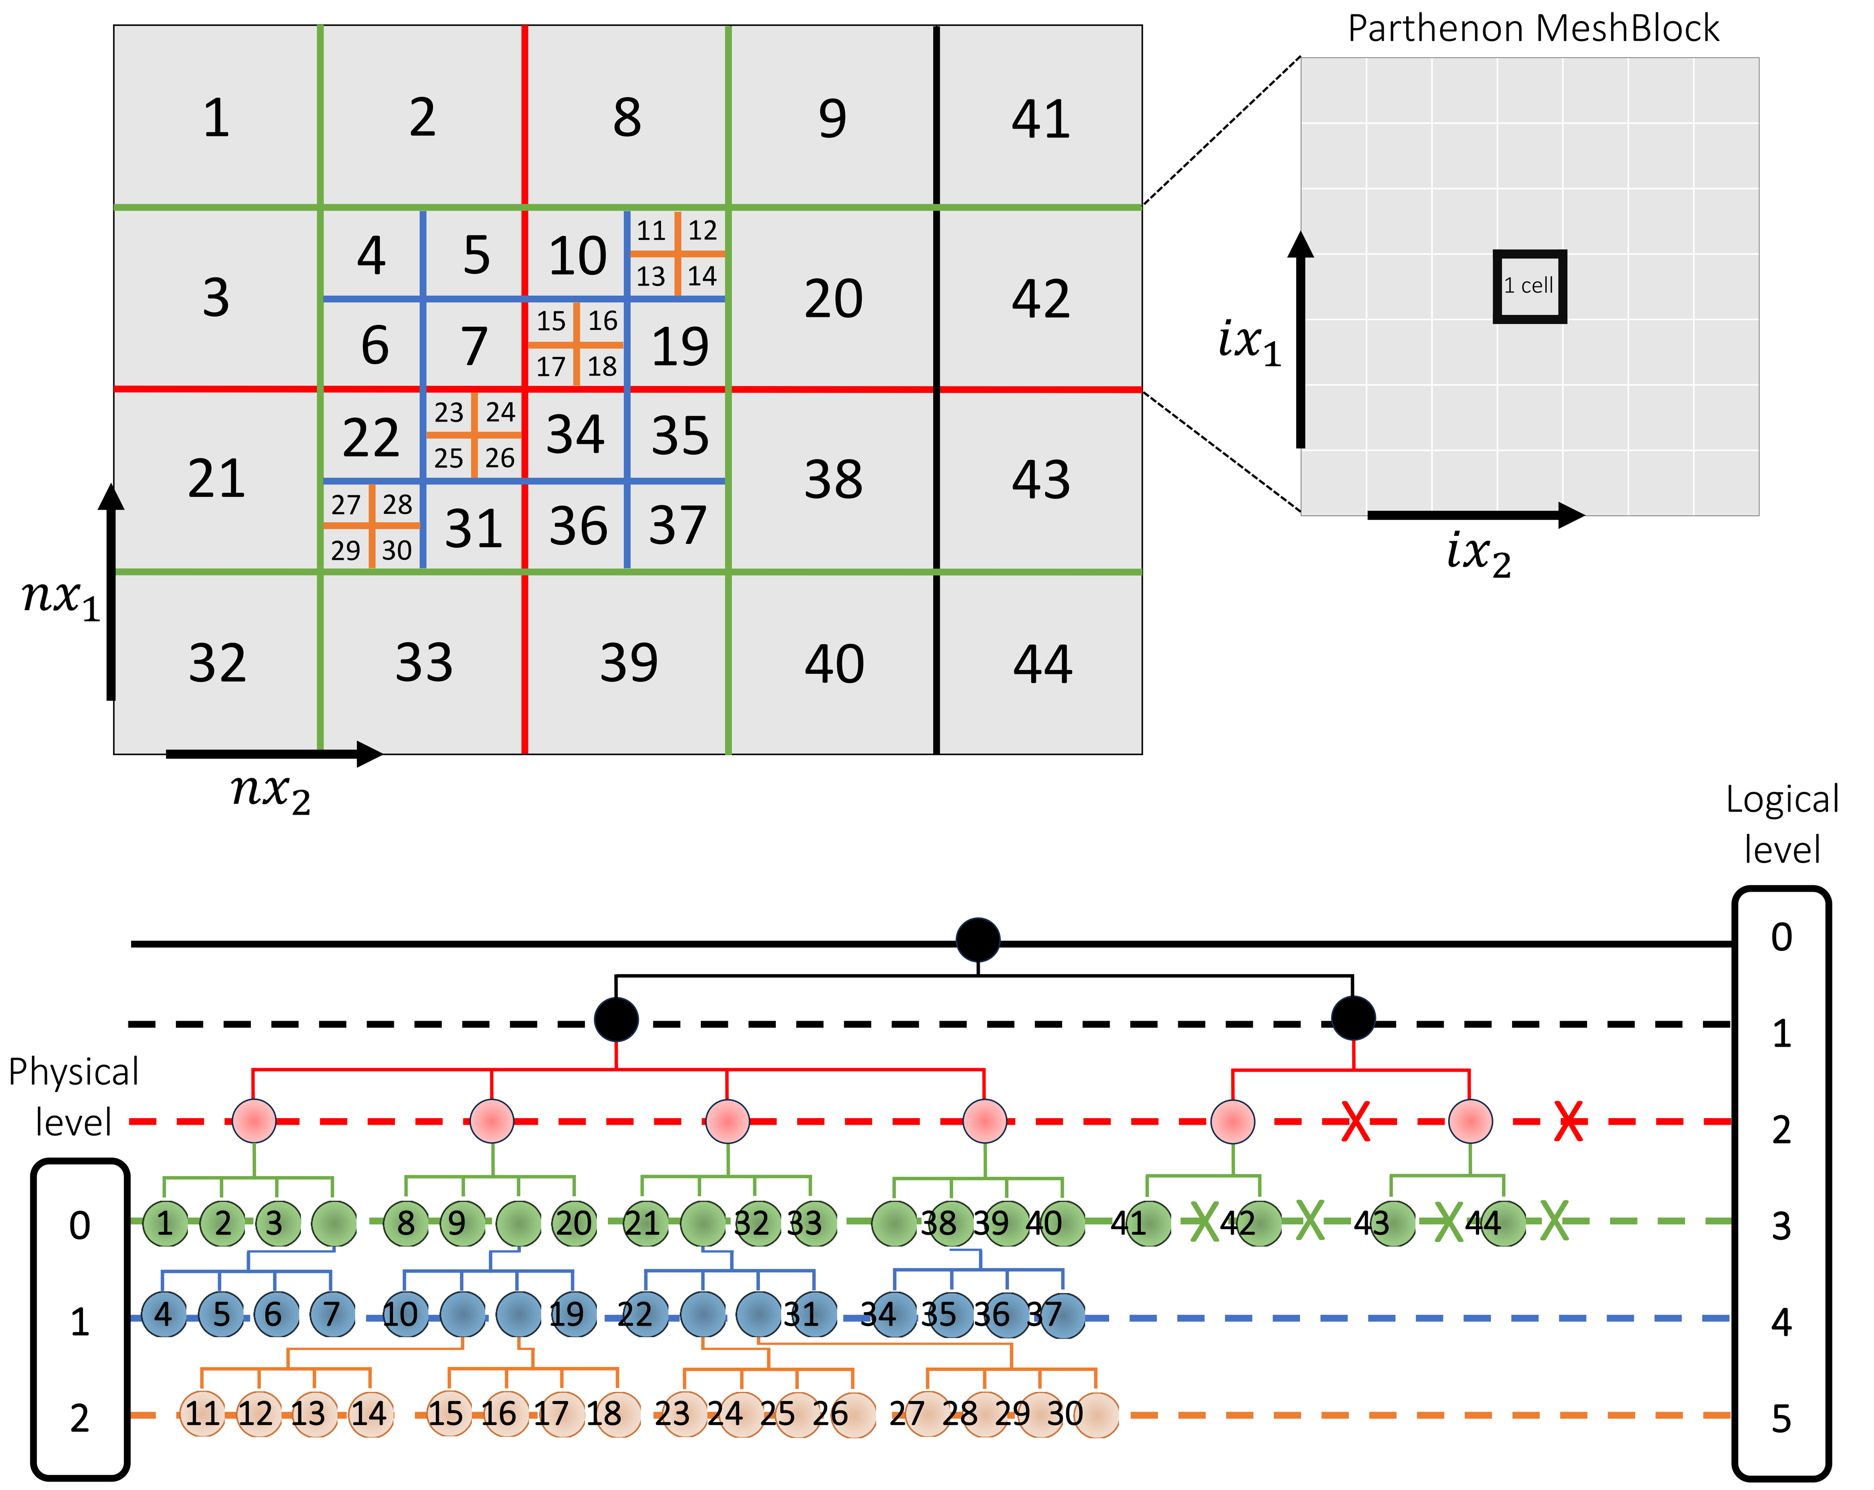

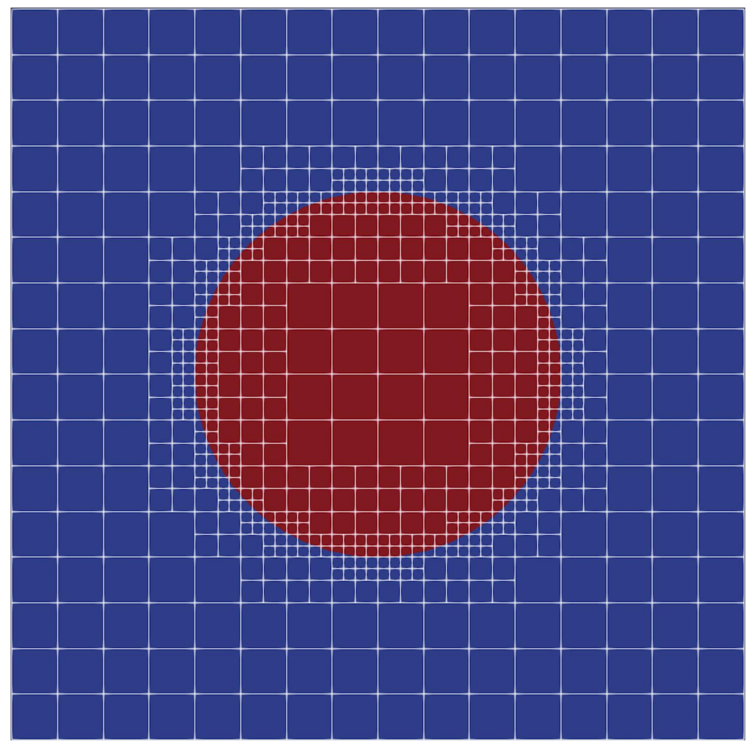

A collapsing gas sphere test problem for Ares with 3 levels of adaptive mesh refinement based on a density gradient criterion

However, these AMR-enabled simulations require *immense computational resources*.

Most existing codes are only designed to run on homogeneous CPU-only systems and are at risk of losing their competitiveness as there is a general shift towards heterogenous HPC architectures. There exist several efforts to enable these codes for GPUs, however, they are vendor specific. Solutions for performance portability like **Kokkos** facilitate new developments.

Inspired by this problem, we create the **first performance portable multi-physics massively-parallel hydrodynamics code Ares** based on the Parthenon AMR framework, which enables us to reach resolved scales that are out of reach for current state-of-the-art codes.

NSE Solver

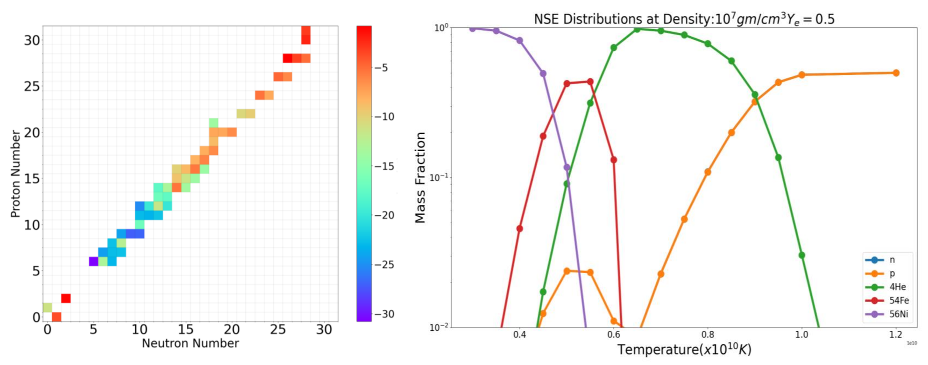

(left) Mass fractions of different isotopes relevant for thermonuclear burning for materials with density of 10^7 g/cm^3 and temperature of 6×10^9 Kelvin. (right) Mass fraction distribution for key isotopes within a range of temperatures relevant for thermonuclear burning.

Helmholtz Equation of State

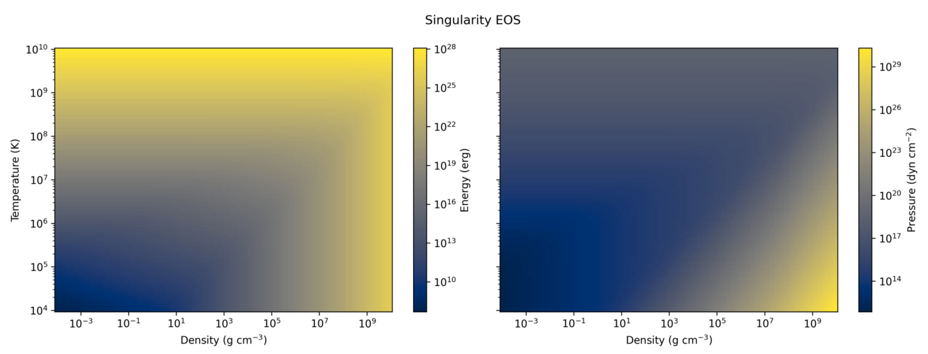

The internal energy (left) and pressure (right) calculated using the Helmholtz equation of state for an equal-mass fraction mixture of C-12 and O-16 in temperatures and densities relevant for degenerate stellar interiors. To our knowledge, this is the first implementation of the Helmholtz EoS on GPUs.

Self-Gravity Monopole Solver

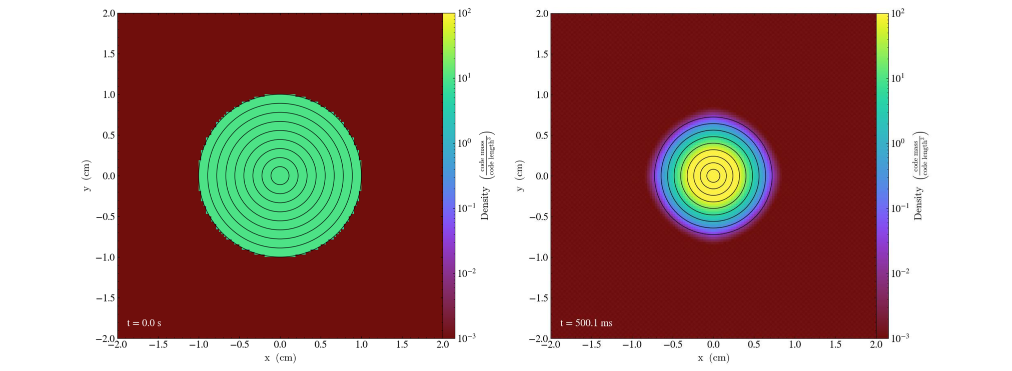

Instead of performing a full Poisson solve for self-gravity, we took advantage of the spherical symmetry of our system and implemented a monopole gravity solver. This solve requires mass binning shown above for two timesteps which we use to evaluate the gravitational acceleration for each zone in our computational domain.

Simulation Snapshot

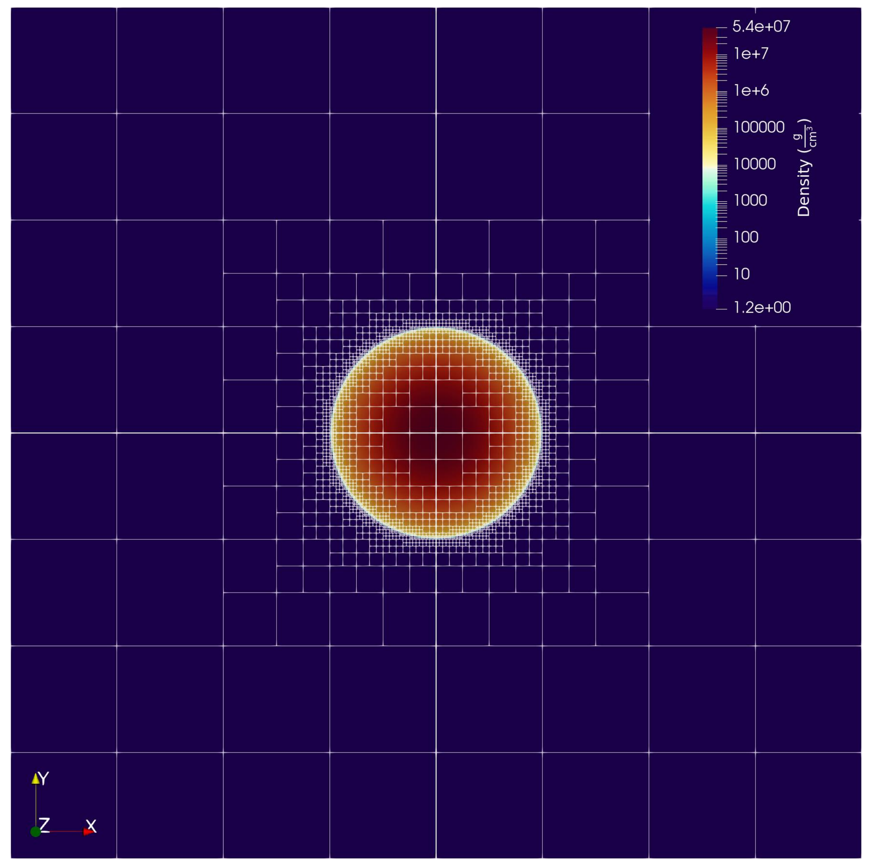

Slice plot of a white dwarf at initialization with 7 AMR levels using Ares, produced with Paraview 5.10.1

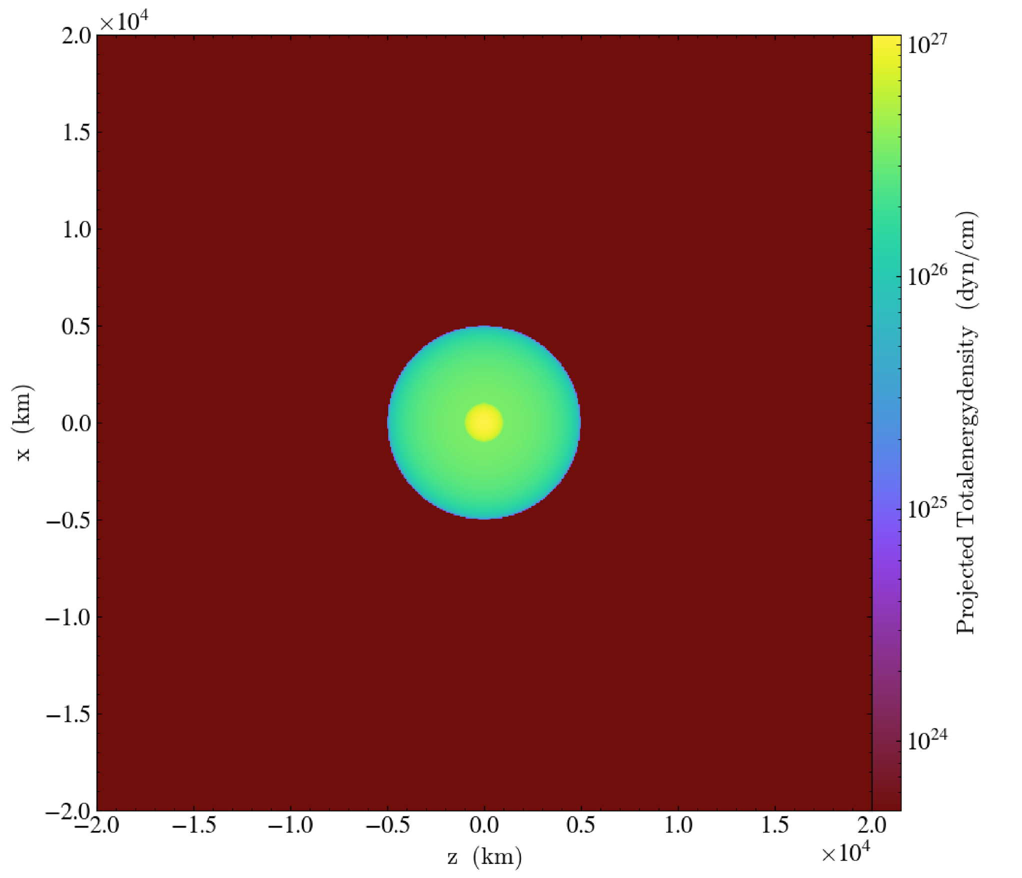

Slice plot of a white dwarf with an ignition hotspot of temperature set at 7 x 10^9 K, produced with Python package yt.

Computing Architecture on Chicoma

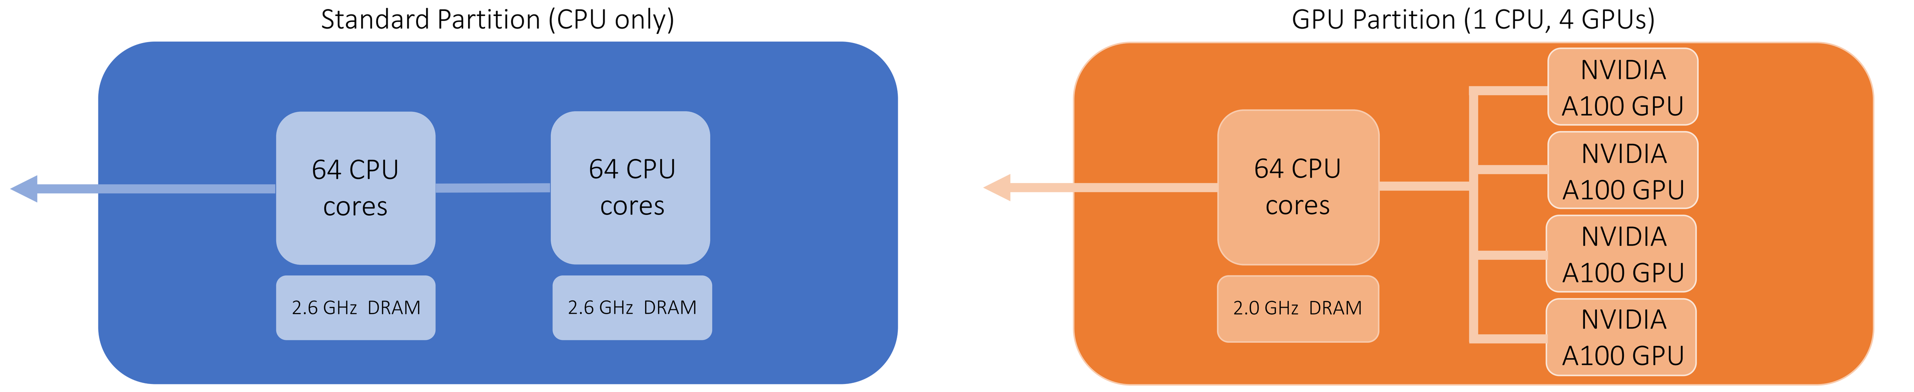

Software specification: 2 – Rome 64 core @ 2.6 GHz AMD EPYC 7H12 with 512 GB per node (16-32 GiB DIMM)

Software specification: 1 – Rome 64 core @ 2.0 GHz AMD EPYC 7713 with 256 GB per node (8-32 GiB DIMM)
4 – NVIDIA A100 Tensor Core GPUs 40 GB HBM2 / GPU

Scaling Study

We conducted scaling studies by evaluating Ares on a toy ideal gas sphere problem for 1000 cycles on LANL Chicoma. We varied MPI ranks from 1 to 64, with each rank using either 64 OpenMP-enabled CPU cores or 1 A100 GPU. To match the architecture, GPU runs used 4 MPI ranks per node while CPU runs used 1.

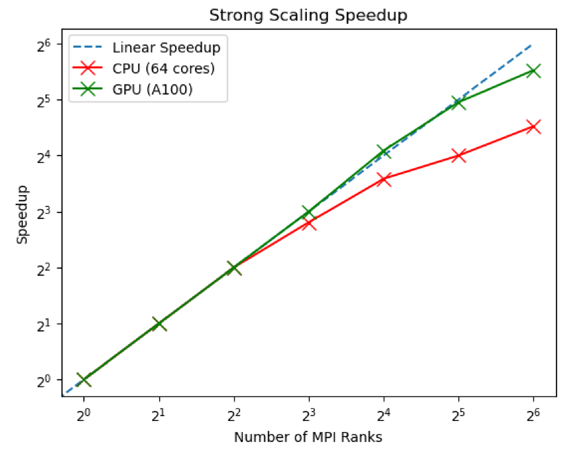

With adaptive mesh refinement enabled, strong scaling shows perfect speedup up to 32 GPUs, and near-linear scaling for both CPU and GPU ranks. **MPI ranks with one GPU were up to 50% faster than 64 OpenMP-enabled CPU cores.**

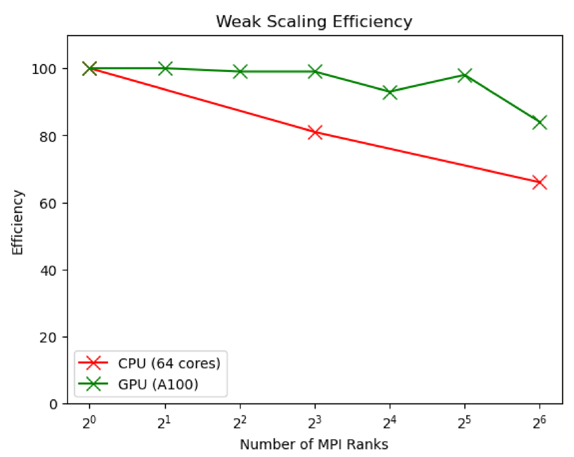

The weak scaling plot outlines that the CPU runs maintain **66% efficiency at 64 MPI ranks with OpenMP**. The weak scaling results for GPU runs show an **efficiency of 84% using 64 GPUs on 16 nodes**.

References

1. Parthenon: Grete et al. *Int J High Perform Comput Appl* (2022)
2. Kokkos: Christian R. Trott et al. *IEEE Trans. Parallel Distrib.* (2022) **33** 4
3. yt: Matthew J. Turk et al. (2011) *ApJS* **192** 9
4. Helmholtz EoS: F. X. Timmes and F. Douglas Swesty (2000) *ApJS* **126** 501
5. NSE solver: I. Seitenzahl et al. (2009) *At. Data Nucl. Data Tables* **95** 1

LA-UR-23-29212

Acknowledgements

We acknowledge the support and guidance of our mentors: Julien Loiseau, Hyun Lim, Samuel Jones, Forrest Glines, Philipp Edelmann, Jonah Miller, Brendan Krueger, Chris Mauney, Richard Berger, Andrew Reisner, Sumathi Lakshmiranganatha and Karen Tsai. We also thank Christoph Junghans, Kelley Thompson, Chris Fryer and Erika Maestas for their assistance throughout the summer school. We also acknowledge the financial support and funding from CNLS and ASC to enable this work.