Maestro and Castro:

Simulation Codes for Astrophysical Reacting Flows

Michael Zingale[1], Ann Almgren[2], Maria G. Barrios Sazo[1], Vince Beckner[2], John Bell[2], Doreen Fan[2], Brian Friesen[2], Adam Jacobs[3], Max Katz[4], Chris Malone[5], Andy Nonaka[2], Donald E. Willcox[1], Weiqun Zhang[2]

[1]Stony Brook University; [2]Lawrence Berkeley National Laboratory; [3]Michigan State University; [4]NVIDIA Corp; [5]Los Alamos National Laboratory

Stellar explosions are multiphysics problems—modeling them requires the coordinated input of gravity solvers, reaction networks, transport, and hydrodynamics together with microphysics recipes to describe the physics of matter under extreme conditions. Furthermore, these models involve following a wide range of spatial and temporal scales, which puts tough demands on simulation codes. We developed the codes Maestro and Castro to meet the computational challenges of these problems. Maestro uses a low Mach number formulation of the hydrodynamics to efficiently model convection. Castro solves the fully compressible radiation hydrodynamics equations to capture the explosive phases of stellar phenomena. Both codes share the same microphysics and use the AMReX library to provide adaptive mesh refinement and manage the parallelism.

AMReX

(Rendleman et al. 2000, Zhang et al. 2016)

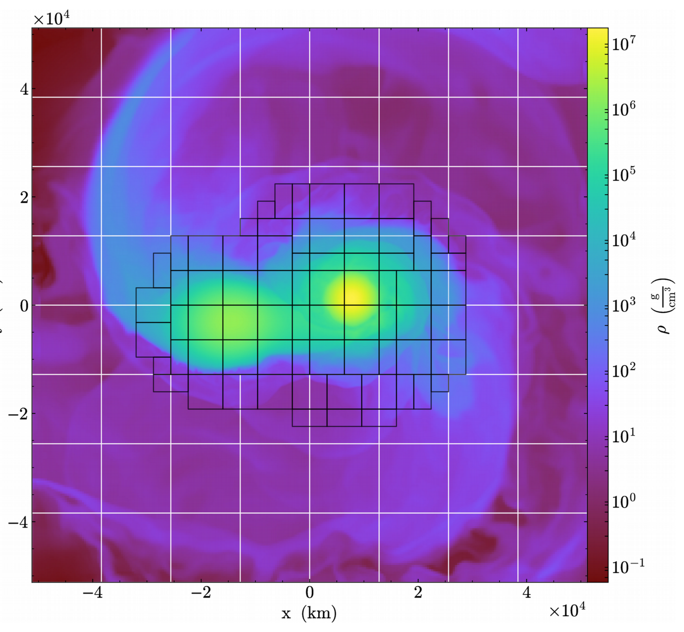

- Block structured AMR
  - Domain divided into varying sized rectangular patches
  - AMReX managed the grid and parallel distribution.
  - Domain codes write kernels that operate on a patch.
  - User-defined tagging criteria for refinement
- C++/Fortran library
- Hybrid parallelism model based on MPI and OpenMP
  - Coarse-grained distribution of grid patches to nodes using MPI
  - OpenMP threads work within grids
    - Either threading of loops over zones or logical tiling of grids
- Efficient cell-centered and node-centered geometric multigrid solvers provided
- Nightly regression test of all the codes
- Extensive performance and memory profiling tools built-in
- Portability: AMReX-based codes (including Maestro and Castro) run on anything from a laptop to the latest supercomputers
- GPU offloading in development
- Visualization with yt and VisIt

Open development model

- Open development repositories—all codes are freely available on github
- Planning / review through github issues and pull requests
- Extensive User's Guides and mailing lists are provided to help new users

Maestro: low Mach number hydro

(Almgren et al. 2006a,b, 2008; Zingale et al. 2009, Nonaka et al. 2010)

- Reformulation of compressible Euler equations
  - Retain compressibility effects due to heating and stratification
  - Asymptotic expansion in Mach number decomposes pressure into thermodynamic and dynamic parts
  - General equation of state supported
  - Hydrostatic equilibrium analytically enforced:

$$\nabla p_0 = \rho_0 g$$

- Elliptic constraint on velocity field:

$$\nabla \cdot (\beta_0 \mathbf{U}) = \beta_0 \left(S - \frac{1}{\overline{\Gamma}_1 p_0} \frac{\partial p_0}{\partial t} \right)$$

  - β_0 is a density-like variable that captures atmospheric stratification
  - S represents reactions and external heating sources
- Self-consistent evolution of base state
- Timestep based on bulk fluid velocity, not sound speed
- Brings ideas from the atmospheric, combustion, and applied math communities to nuclear astrophysics
- Solution methodology
  - Unsplit PPM advection
  - Implicit thermal diffusion (solved via multigrid)
  - Approximate projection enforces velocity constraint (using cell-centered and node-centered multigrid)
  - Strang-splitting for reactions

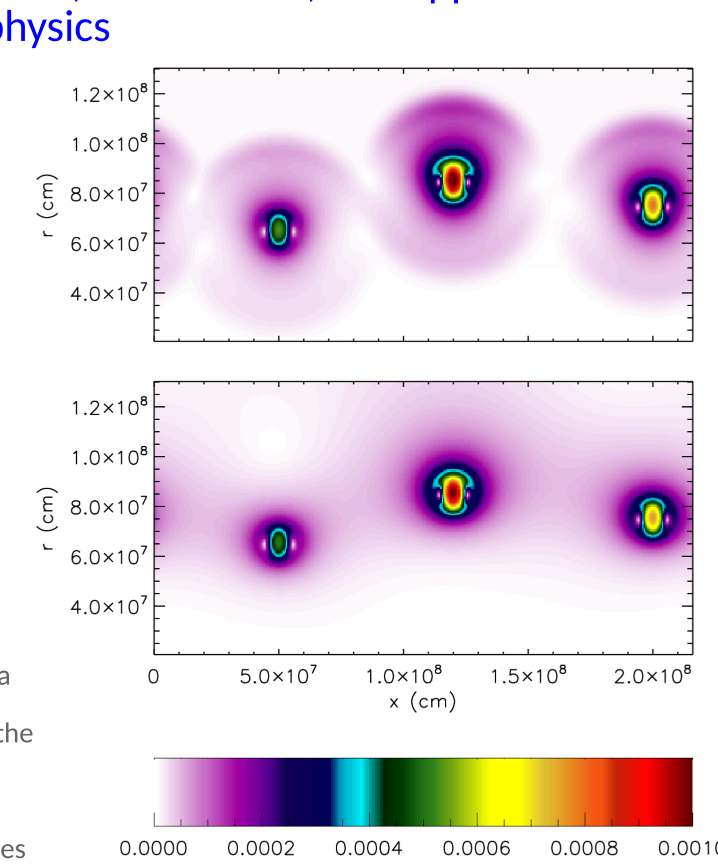

Mach number for 3 burning rising bubbles from a compressible code (top) and Maestro (low Mach number hydrodynamics method) (bottom). For the compressible code, we see the sound waves propagating outward from the disturbances. In Maestro, the velocity field adjusted instantaneously. Note that the flow in the bubbles looks identical between the two codes.

Castro: compressible radiation hydro

(Almgren et al. 2010, Zhang et al. 2011, 2013)

- 1-, 2-, and 3-dimensional unsplit, 2nd-order hydrodynamics
- Multigroup flux-limited diffusion radiation hydrodynamics, including terms to O(v/c)
- Adaptive mesh refinement with subcycling in time; jumps of 2x and 4x between levels
- Arbitrary equation of state
- General nuclear reaction networks
- Explicit thermal diffusion
- Full Poisson gravity (with isolated boundary conditions), conservative flux formulation
- Rotation (in the co-rotating frame) in 2-d axisymmetric and 3-d
- Ability to restart from a Maestro calculation to bring it into the compressible regime

Parallel performance

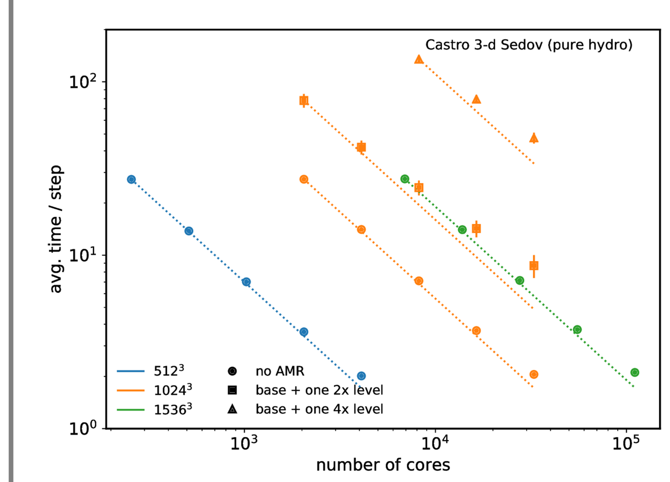

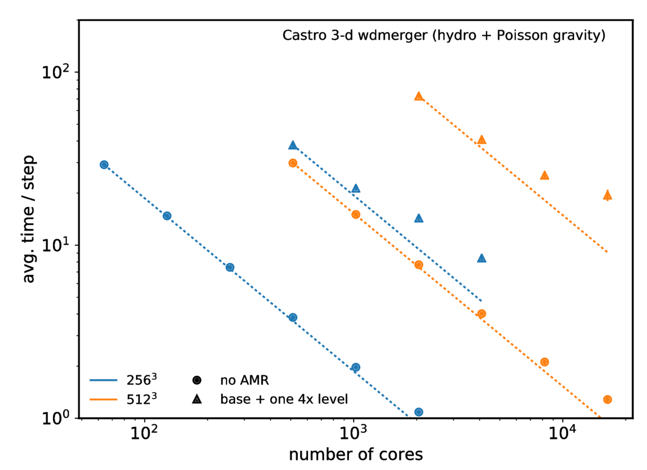

(left) Castro strong scaling on OLCF Titan for pure hydro (Sedov w/ real EOS). The different colors represent different base resolutions. Runs are done with jumps of 2x and 4x in refinement and show excellent strong scaling across all processor counts. (right) Castro strong scaling on OLCF Titan for a hydro + Poisson gravity (wdmerger) problem. This uses a multipole solver to determine Dirichlet boundary conditions and then geometric multigrid to solve for the potential in the interior. Two coarse grid sizes are shown with and without a 4x jump in refinement and show excellent scaling.

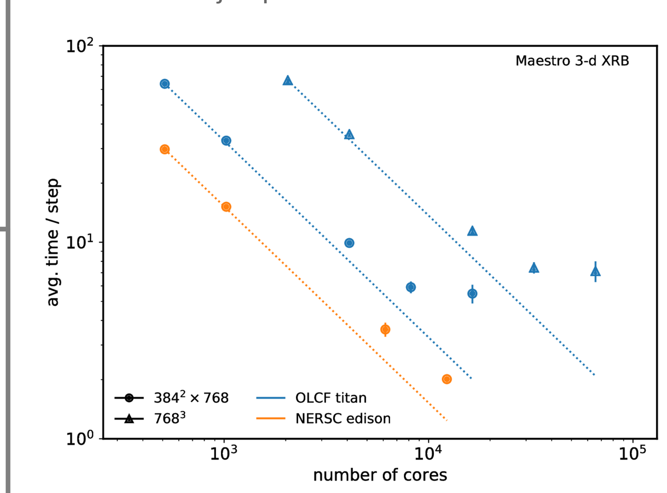

Maestro strong scaling on NERSC Edison and OLCF Titan. Two different problem sizes are shown, demonstrating good scaling to high core counts. The turn off at the end at OLCF is from crossing NUMA nodes.

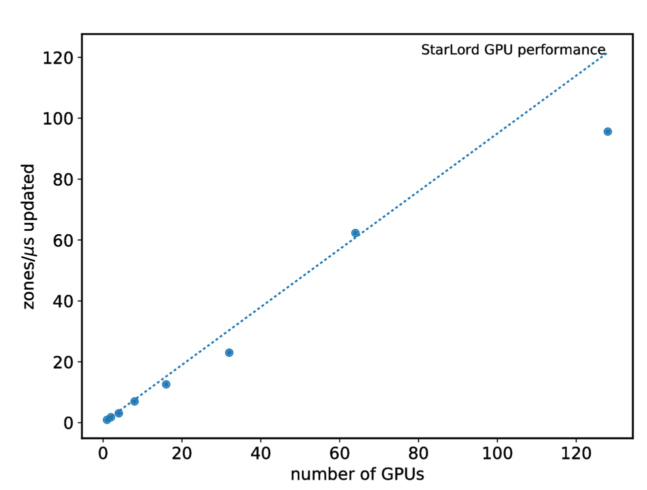

Castro proxy code StarLord GPU scaling on the OLCF Summitdev machine. We see good scaling behavior to 64 GPUs in this test of a pure hydro (Sedov + real EOS).

Shared Microphysics

https://github.com/StarKiller-astro/Microphysics/

- Common EOS and reaction networks appropriate for stellar hydrodynamics
- Reaction righthand side decoupled from integration method
- JINA ReacLib interface via pynucastro (http://pynucastro.github.io/pynucastro/)

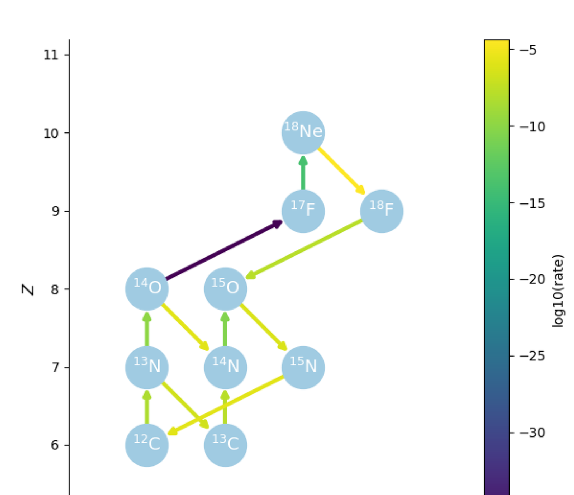

▶ Example pynucastro network for the CNO cycle (with breakout), colored according to the reaction rate. This network can be exported to the StarKiller Microphysics and used directly in Maestro and Castro

Applications

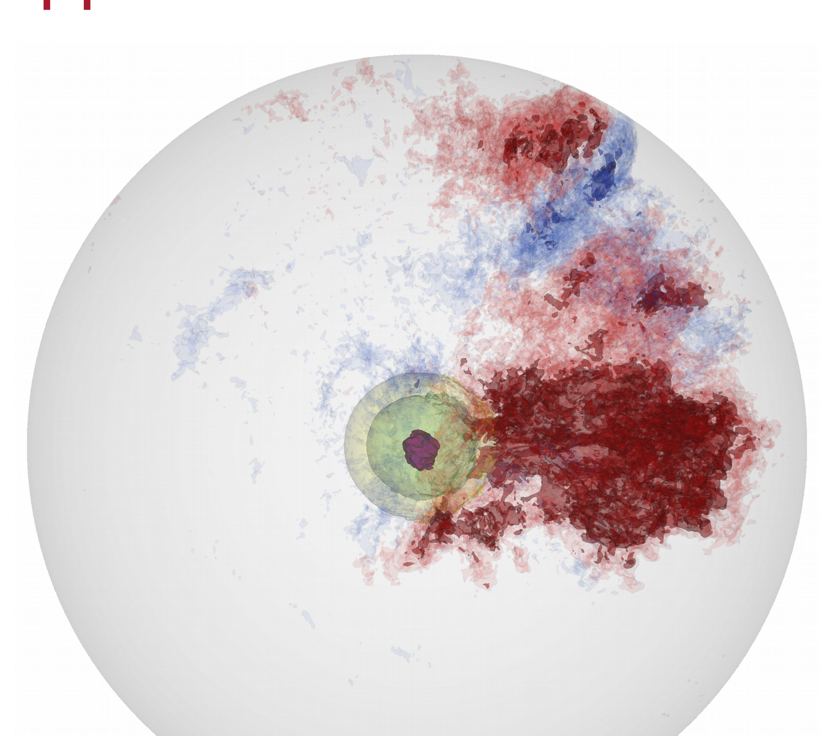

▲ Maestro calculation of the convection preceding ignition in the Chandrasekhar mass model of a Type Ia supernova. The radial velocity is shown (red: outflow, blue: inflow) in the inner 1000 km of the star, computed with a resolution of 2 km. We see a strong jet-like feature in the outflow which rapidly changes direction with time. (Zingale et al. 2011, Nonaka et al. 2012, used OLCF Titan)

▼ Castro simulation of the inspiral and merger of two white dwarfs, as a model of Type Ia supernova. Full self-gravity with isolated-boundary conditions is used in a conservative fashion to accurately model the merger event. A series of calculations with different mass white dwarfs will be used to understand the conditions when the merger takes place. (Katz et al. 2016, used OLCF Titan)

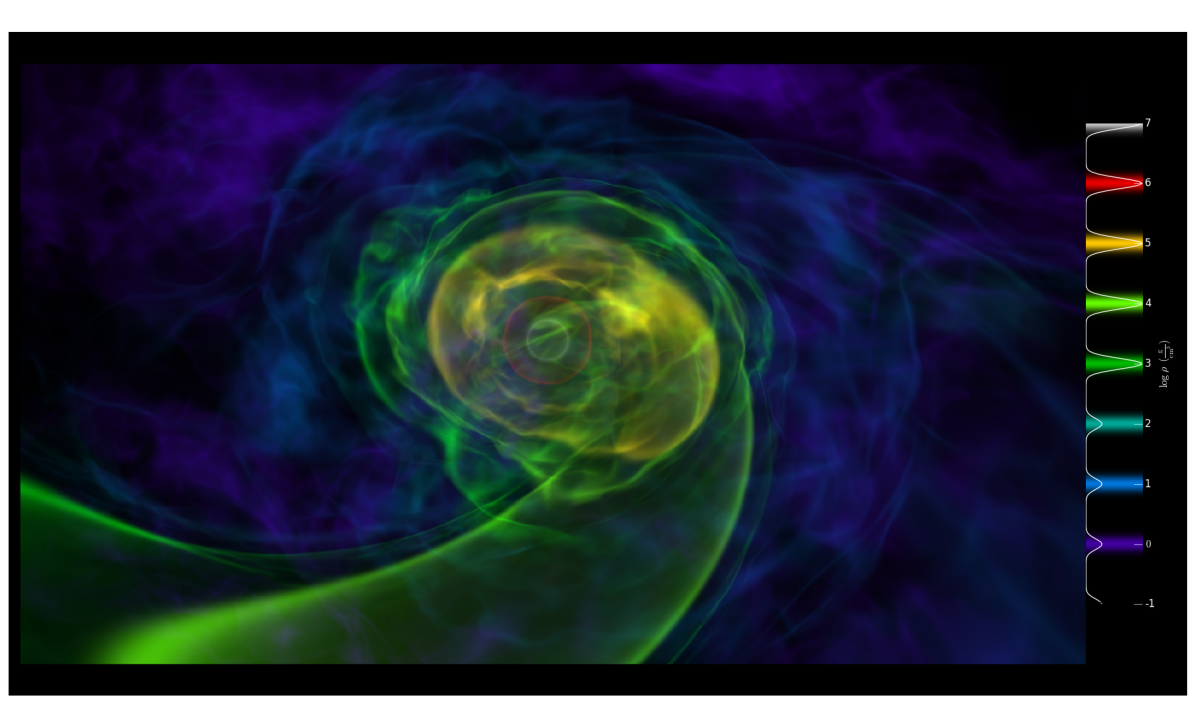

▶ Maestro simulation of convection in the accreted helium layer on a white dwarf, showing the radial velocity (red = outflow; blue = inflow). Three different models are shown: (top) a 0.8 $M_\odot$ white dwarf with a 0.12 $M_\odot$ He layer; (middle) a 1.0 $M_\odot$ white dwarf with a 0.04 $M_\odot$ He layer; and (bottom) a 1.2 $M_\odot$ white dwarf with a 0.02 $M_\odot$ He layer. We see that the convection breaks into cells, with the size of the cells roughly equal to the thickness of the He layer. This convection models the early stages of the sub-Chandra model for Type Ia supernovae. (Zingale et al. 2013, Jacobs et al. 2016, used OLCF Titan)

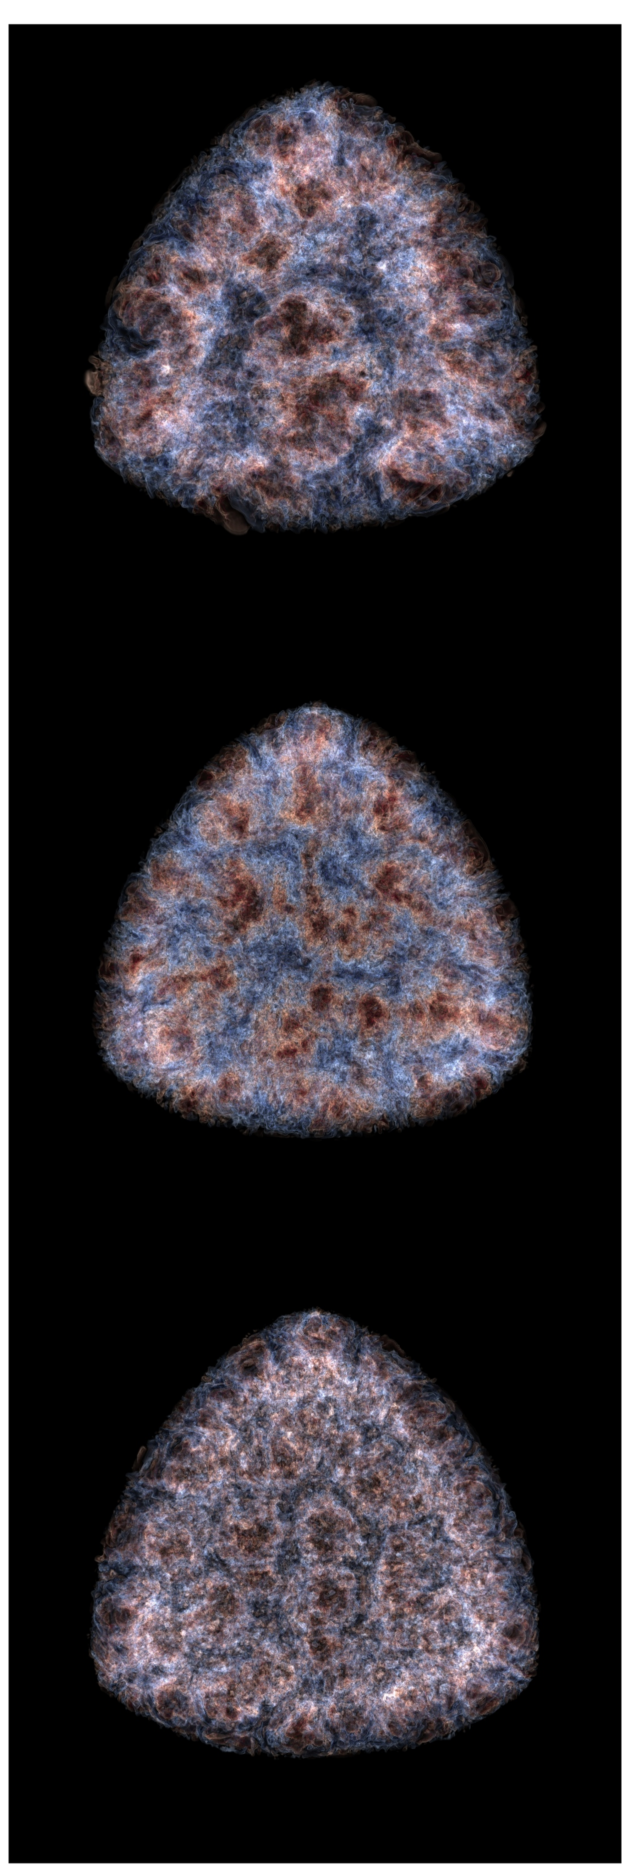

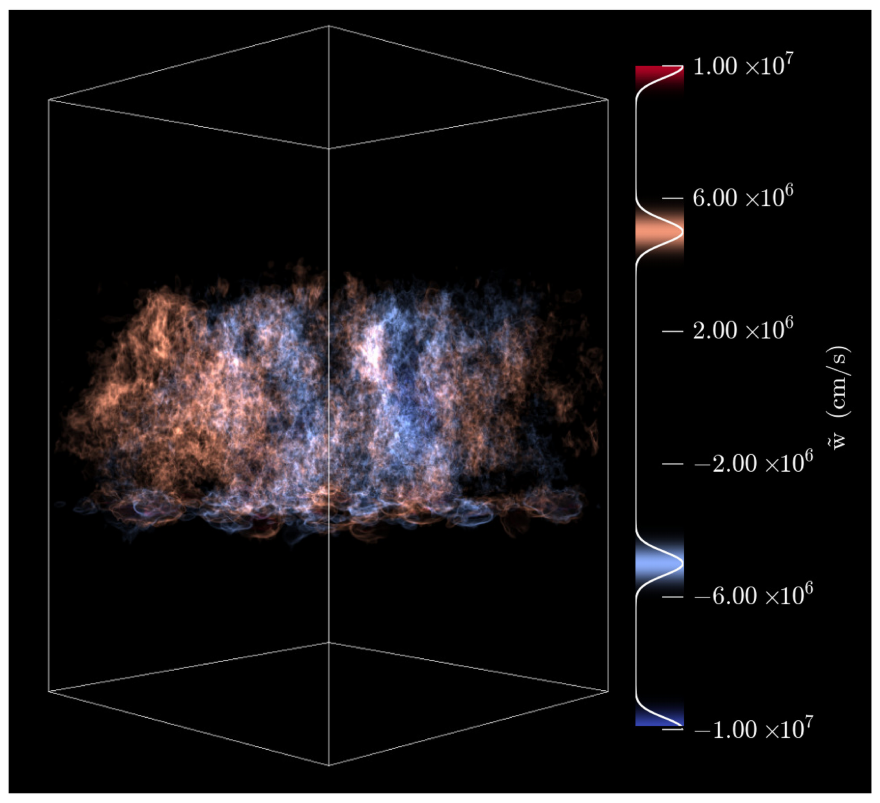

▶ Castro simulation of a buoyant thermonuclear flame during a Type Ia supernova, shortly after ignition. The bright rope-like structures trace strong vortex tubes, which are initially confined to a torus and then break apart to spread to both the flame surface and the ashes of the burning, as seen here. As turbulence sets in, the vortex tubes are stretched to thinner scales and fully permeate the flame, altering it's morphology and evolution. (Malone et al. 2014, used OLCF Titan; NCSA Blue Waters)

◀ Maestro model of convection in the accreted H/He layer on the surface of a neutron star as a model of an X-ray burst. The coloring shows the vertical velocity (red = upflows; blue = downflows), highlighting the convective cells. (Zingale et al. 2015, used OLCF Titan)

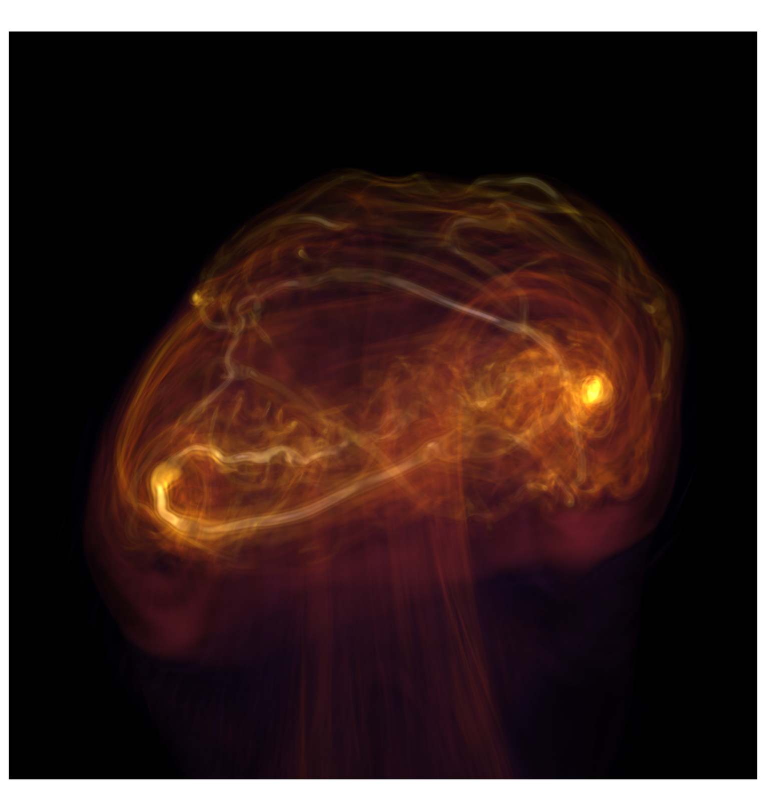

Future developments

- Maestro:
  - Port to C++ framework
  - Rapid rotation
  - Higher-order accuracy
- Castro:
  - New hydrodynamics solvers / MHD
  - Improved coupling of reactions, hydrodynamics, and gravity
- Microphysics infrastructure:
  - GPU acceleration of EOS and reactions

The work at Stony Brook was supported by a DOE/Office of Nuclear Physics grant No. DE-FG02-87ER40317. The work at LBNL was supported by the Applied Mathematics Program of the DOE Office of Advance Scientific Computing Research under U.S. Department of Energy under contract No. DE-AC02-05CH11231.

An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. This research used resources of the Oak Ridge Leadership Computing Facility located in the Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract DE-AC05-00OR22725. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Visualizations were performed using yt, VisIt and matplotlib. We thanks NVIDIA Corp for the grant of a Titan Xp GPU board for our development.

Freely available:
https://github.com/amrex-astro/