Modeling Stellar Explosions with the AMReX Astrophysics Suite

Michael Zingale
(Stony Brook University)

in collaboration with

Ann Almgren, John Bell, Andy Nonaka, Weiqun Zhang (LBNL),
Chris Malone (LANL),
Adam Jacobs (MSU),
Max Katz (Nvidia),
Maria Barrios Sazo, Alan Calder, Doug Swesty, Don Willcox (Stony Brook),
Stan Woosley (UCSC)

Stony Brook
University

Support from DOE Office of Nuclear Physics, NSF Astronomy & Astrophysics.
Nvidia Titan X Pascal donated by Nvidia via GPU Grant Program
Computer time via DOE INCITE @ OLCF/ORNL and NERSC/LBNL

Astrophysical Interests

- Convective burning and explosions
  - Supernovae (both thermonuclear and gravitational)
  - X-ray bursts and novae (thermonuclear explosion of accreted material on a surface of compact object)
  - General stellar evolution, including post main-sequence evolution of massive stars
- Convection challenges
  - Often the convection is highly subsonic
  - Difficult for traditional astrophysical hydrodynamics codes
- Explosive challenges
  - Burning and hydrodynamics can decouple
- New algorithms are needed for efficient simulation of convective astrophysical flows

Challenges of Multiphysics

- Stars involve:
  - Hydrodynamics (including turbulence and instabilities)
  - Combustion
  - Self-gravity
  - Radiation / diffusion
  - Magnetic fields
- *Several different physical processes with different character* (hyperbolic, elliptic, parabolic) and timescales
  - Inefficient to just discretize in space and use a method-of-lines ODE integrator to advance the solution
    - Timestep restricted by stiffest system

Multiscale Challenges

- *Nature is 3-d*
  - Convection driven by nuclear energy release
  - Fluid instabilities / turbulence
  - Localized burning / runaway
  - Rotation
- Range of lengthscales can be enormous
- Solutions (?)
  - Adaptive mesh refinement
  - Subgrid scale models

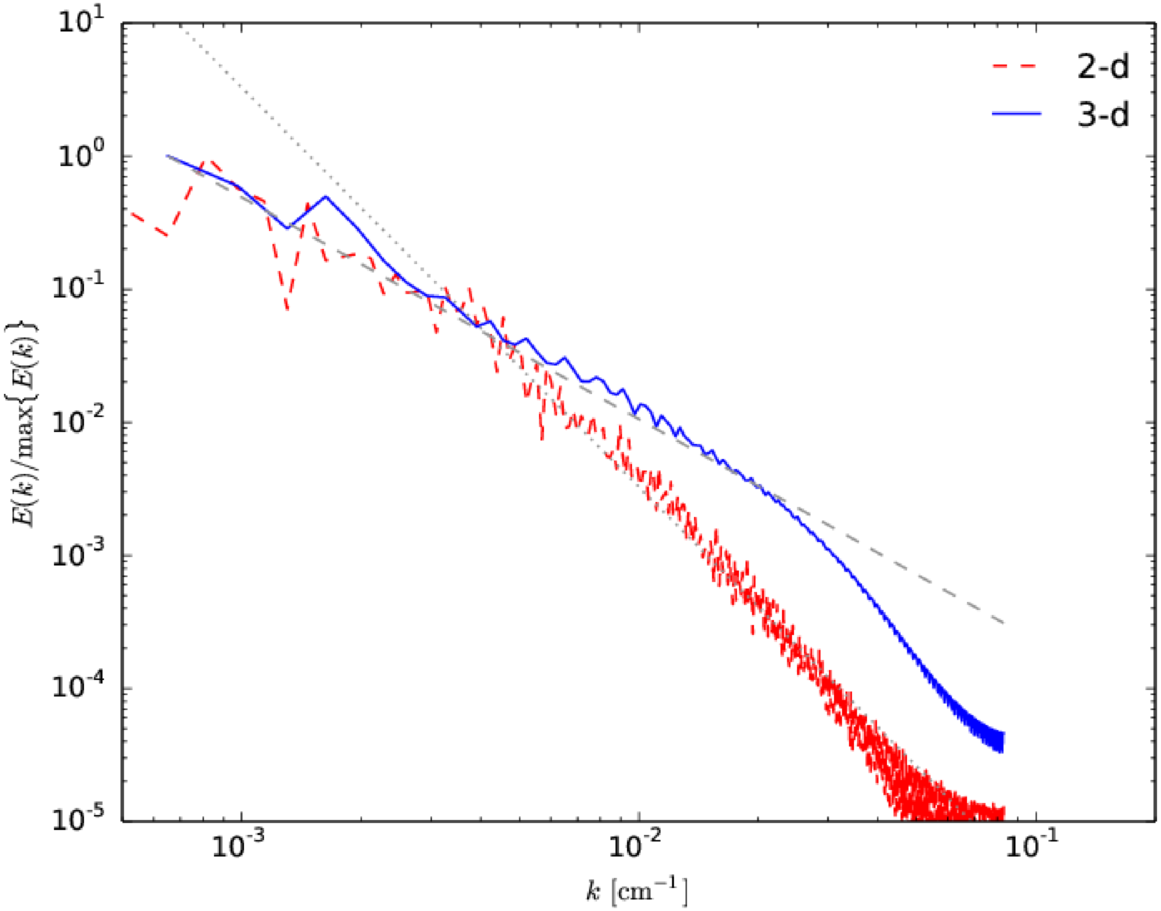

Temporal Challenges

- Many astrophysical explosions exhibit a range of relevant timescales
  - Stellar evolution up to point of explosion / remnant formation ~ millions to 10s of billions of years
  - Simmering convective phase ~ millenia to days/hours
  - Explosion ~ seconds to hours
  - Radiation transport ~ weeks to months
- *No single algorithm can model a star from start to finish*

Low Mach Hydrodynamics

- With explicit timestepping, information cannot propagate more than one zone per step

$$\Delta t = \min \left\{ \frac{\Delta x}{|u| + c} \right\}$$

- For M ≪ 1 :

$$\Delta t \approx \frac{\Delta x}{c}$$

- We want:

$$\Delta t \approx \frac{\Delta x}{|u|}$$

- For very low Mach number flows, *it takes ~ 1/M timesteps for a fluid element to move more than one zone—can't we do better?*

▶ A Mach 0.01 front moving to the right (a) initially, (b) after 1 step, (c) after 100 steps.

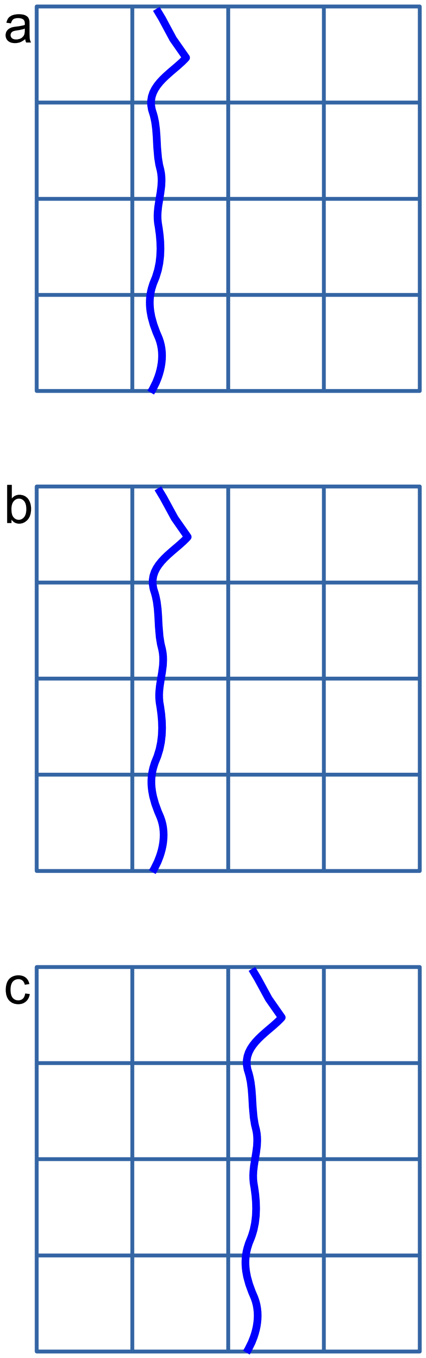

Low Speed Divergence Constraints

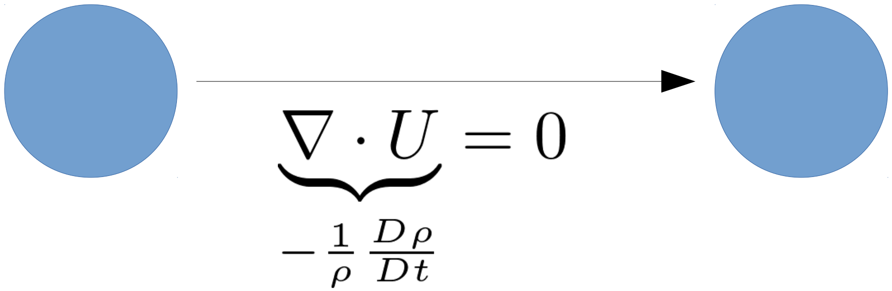

$$\underbrace{\nabla \cdot U}_{-\frac{1}{\rho}\frac{D\rho}{Dt}} = 0$$

Incompressible: density of a fluid element doesn’t change as it is advected

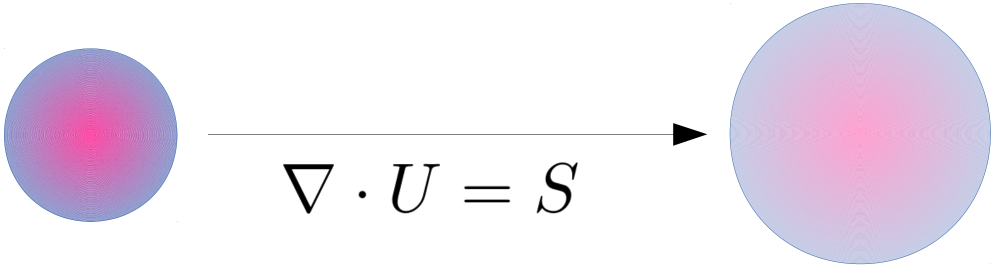

$$\nabla \cdot U = S$$

Low Mach combustion: heat release in fluid element is a source of divergence, fluid element expands

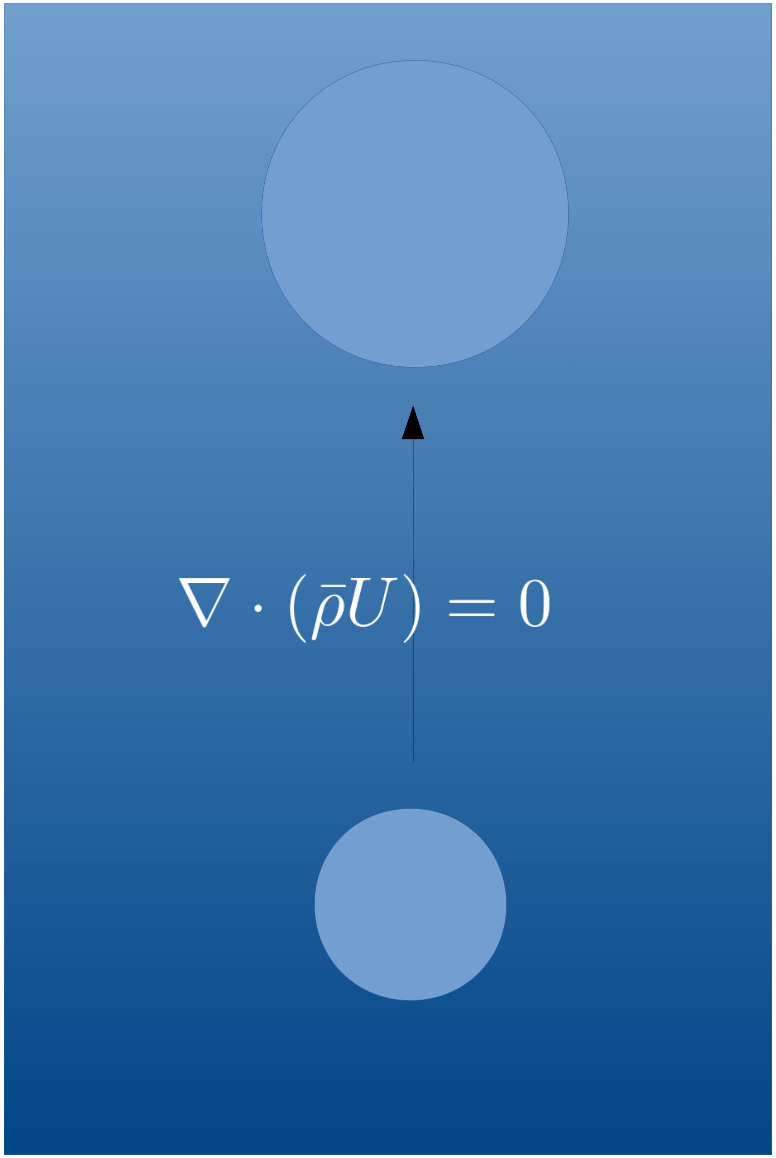

$$\nabla \cdot (\bar{\rho} U) = 0$$

Anelastic: fluid element adiabatically expands as it buoyantly rises

Maestro: Low Mach Hydro

- 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
  - 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
  - S represents 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

Castro

- Castro is the fully compressible counterpart to Maestro
  - 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

Refs:
Almgren et al. 2010
Zhang et al. 2011
Zhang et al. 2013

AMReX: Block-Structured AMR

- AMReX handles the grid management and parallel distribution.
  - F90 or C++/Fortran library

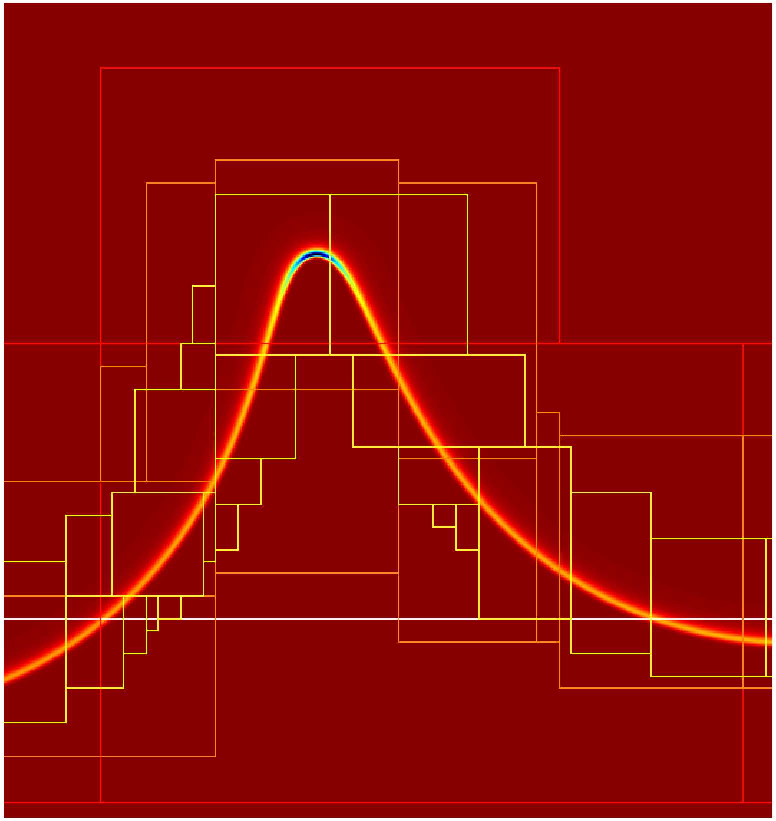

- Hybrid parallelism model based on MPI and OpenMP
  - Tiling to improve manycore performance
- Efficient cell-centered and node-centered geometric multigrid solvers
- Nightly regression test of all the codes
- Sidecar: separate processor group can do runtime analysis while the simulation is running
- Extensive performance and memory profiling tools built-in
- Highly portable

Scaling

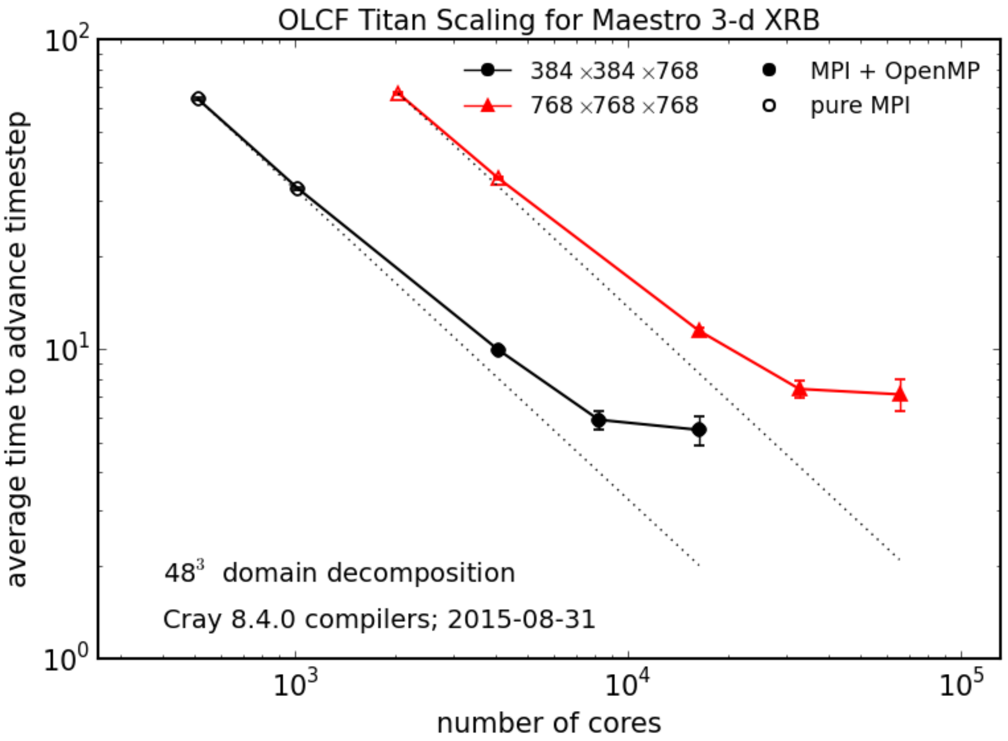

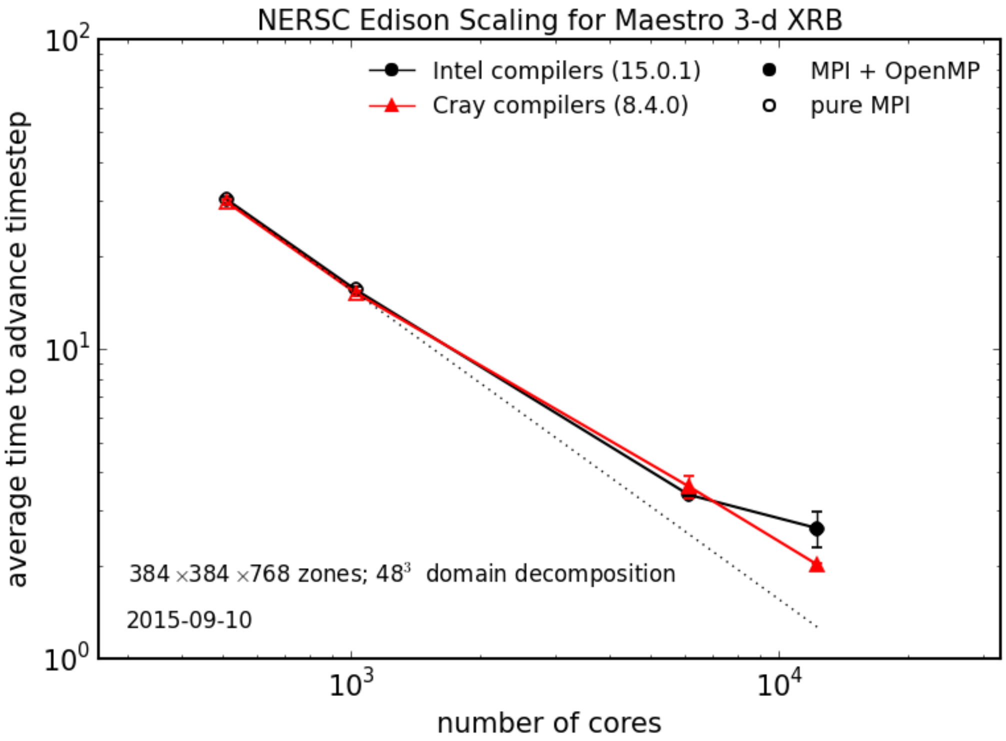

The upturn at the end of each curve is where we use a single MPI task per node, so the OpenMP is going across sockets / NUMA nodes

Open Science

- *Every line of code needed to rerun the simulations shown (SN Ia convection, sub-Ch convection, WD mergers, & XRB) is in our public github repos*
  - `http://github.com/AMReX-Astro`
  - Including inputs files, analysis scripts, submission scripts, etc…
  - repos: MAESTRO, Castro
  - These are our actual development repos
- All output files store the git hashes of the source, the machine name, compiler versions and flags, values of all runtime parameter, …
  - Most papers include the github hash of the repos used for simulations
  - Reproducibility is part of the scientific method

Type Ia Supernovae

- No H; strong Si, Ca, Fe lines
- Occur in old populations
- Bright as host galaxy, $L \sim 10^{43}$ erg s-1
- ^{56}Ni powers the lightcurve
- Act as standard candles
- *General consensus: thermonuclear explosion of a carbon/oxygen white dwarf*
  - What progenitor?

SN 1994D (High-Z SN Search team)

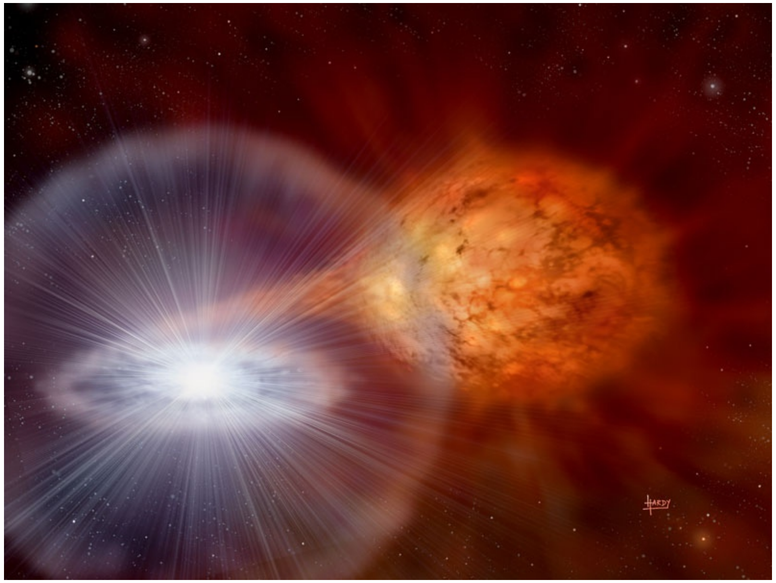

(David A. Hardy & PPARC)

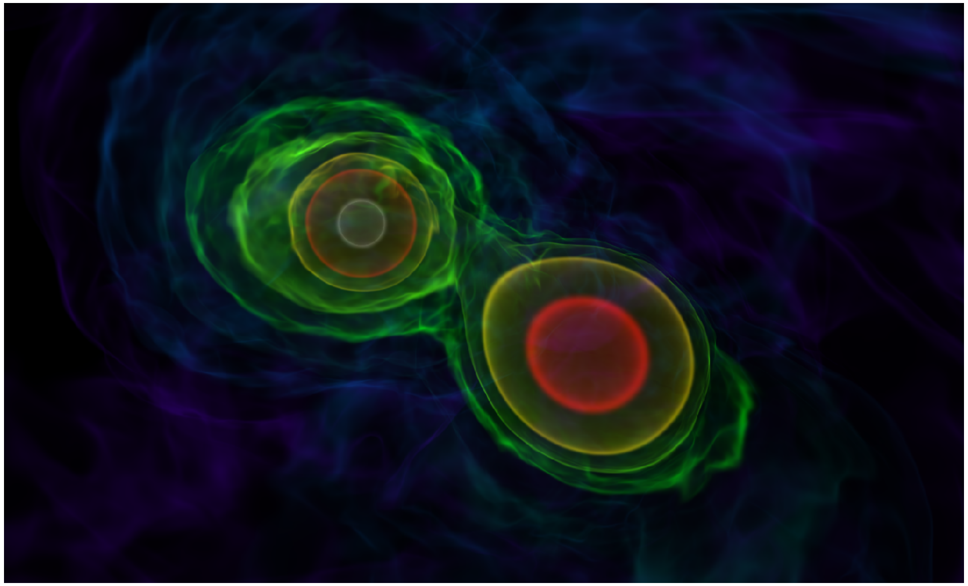

Variations in SNe Ia

- Chandra model:
  - Massive WD accretes from companion to Chandra mass; simmering of C begins at center, burning front ignition consumes star
  - Does nature make massive Wds?
  - Does the burning front remain subsonic?
- Double degenerates:
  - Two WDs inspiral, explosion either prompt or after (long term?) accretion
  - Can we avoid the accretion induced collapse?
  - Can we get an explosion that looks like a SNe Ia?
- Sub-Chandra model:
  - Double detonation: ignite in He layer on surface of WD, shock converges at center of underlying C/O WD and detonates inside out
  - Can we hide the He?
  - Can we make normal SNe Ia?
- What gives the variation in the peak brightness of events?

Convection Preceding Chandra Model

- Explosion in Chandra model for SN Ia preceded by centuries of simmering / convection
  - Sets explosion initial conditions
- Dipole / jet feature seen (as in previous calculations)
  - Asymmetry in radial velocity field
  - Direction changes rapidly
- Ignition is localized
  - Single point, off-center favored

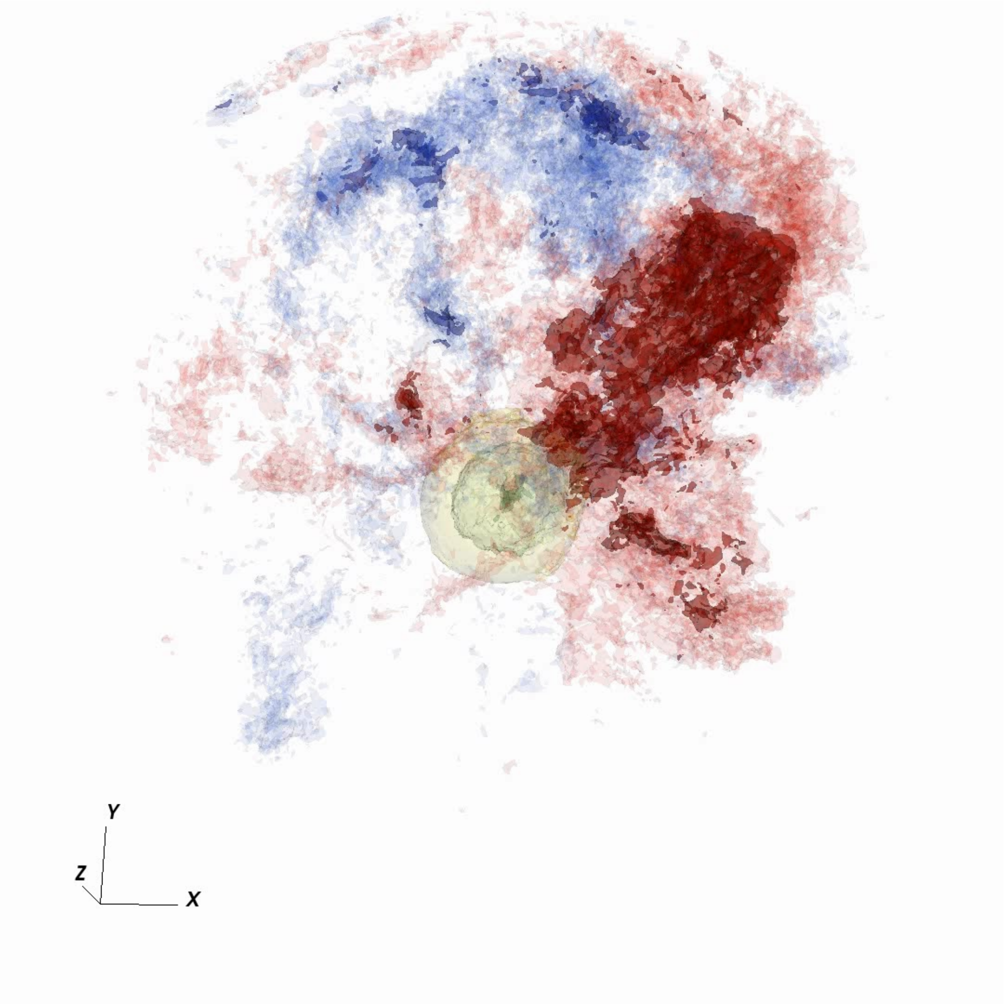

Radial velocity field (red = outflow; blue = inflow) in an 1152^3 non-rotating WD simulation.

Refs:
Zingale et al. 2009
Zingale et al. 2011
Nonaka et al. 2012

Ignition Radius Likelihood

- Distribution of likely ignition locations
  - Average hotspot radius over 1 s intervals
  - Consider final 200 s of evolution
- Vast majority of hotspots are moving outward from the center
- *Off-center ignition likely*

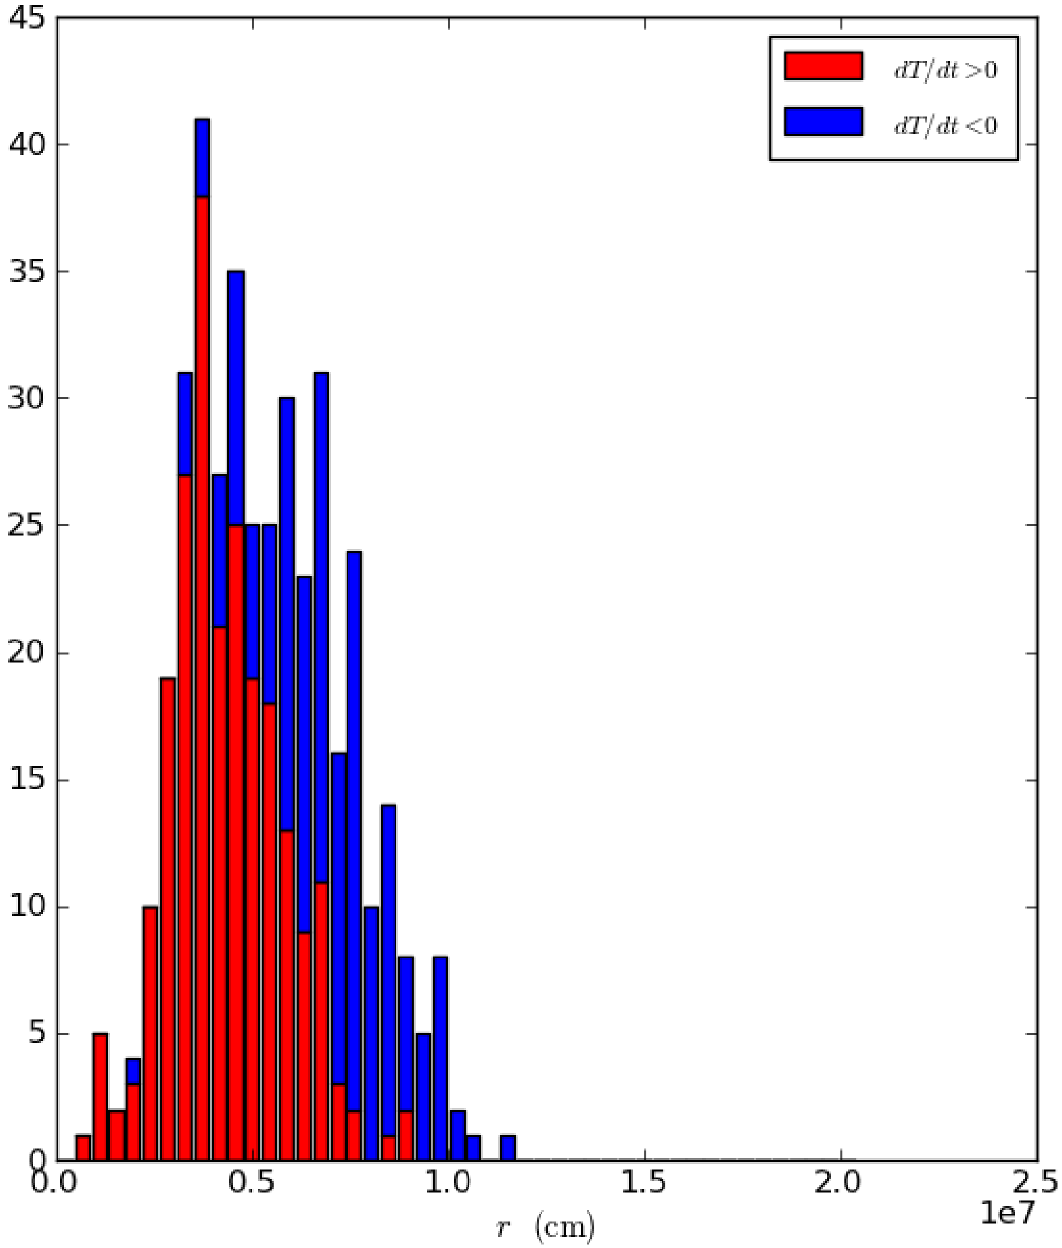

► Histogram of likely ignition radii from 576^3 non-rotating model. Hotspot radii are averaged into 1 s intervals and colored by sign of temperature change

On To Explosion...

- Mach number gets large (ignition): restart in our compressible code, Castro
  - Same underlying BoxLib discretization
  - Same microphysics
  - Solves the fully compressible Euler equations using an unsplit PPM method
- Basic findings:
  - Off-center ignition: background turbulence doesn't strongly affect flame propagation.
  - Central ignition: convective turbulence can push the flame off-center.
  - Single-degenerate model almost always produces an asymmetric explosion
  - Single spot = small amount of burned mass = less expansion = higher density when DDT occurs

Castro (including MGFLD radiation solver) is freely available at: https://github.com/BoxLib-Codes

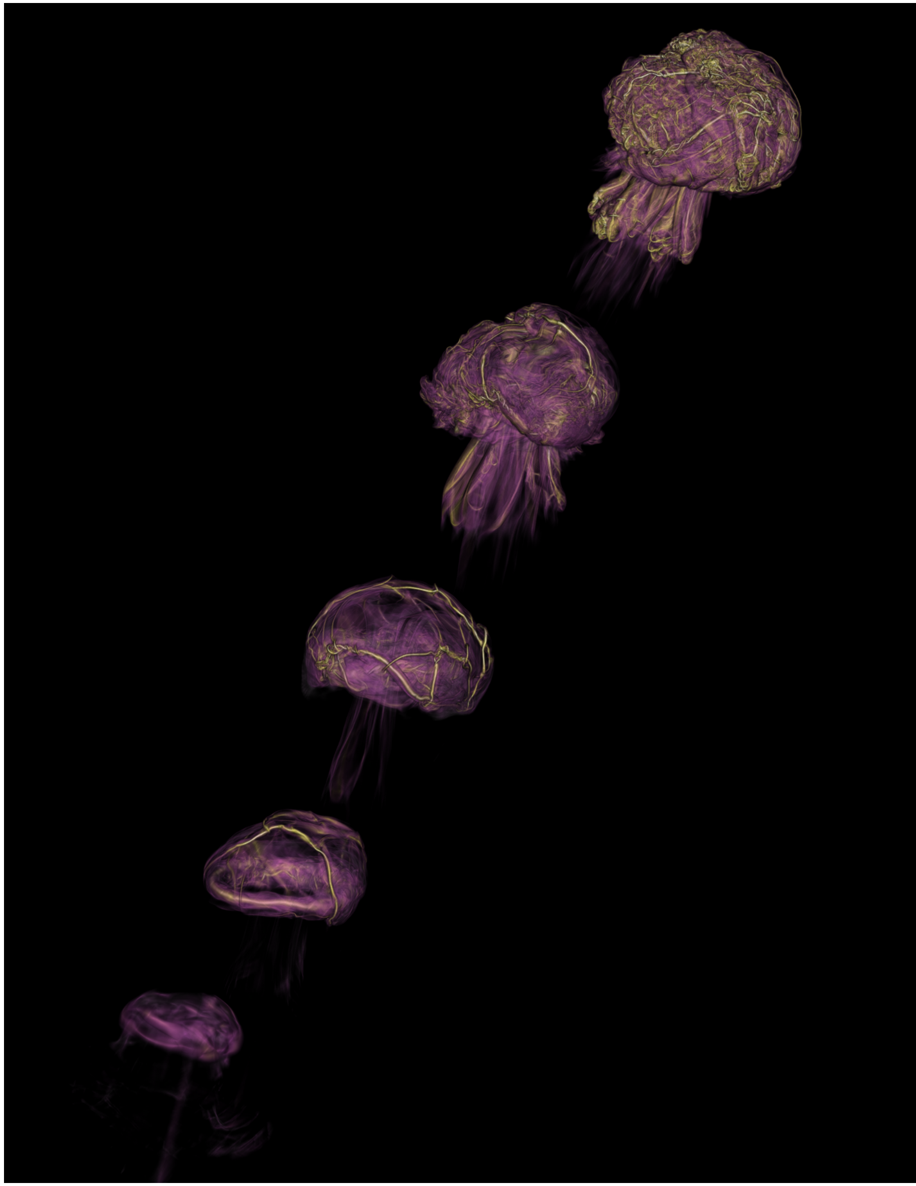

(Malone et al. 2014)

sub-Chandra SNe Ia Models

- Basic idea:
  - Burning begins in an accreted helium layer on the surface of a low(er) mass white dwarf
  - Detonation
- How does the burning transfer to the C/O core?
  - Edge lit: direct propagation of detonation across interface. May require ignition at altitude
  - Double detonation: compression wave converges at core, ignites second detonation at the center of the WD
- Main problem: how much surface He is too much?
- Potential progenitors: Iax class SNe (Foley et al. 2013)
  - Lower velocity, lower peak magnitude, hot photosphere
- Our focus:
  - What does the ignition in the He layer look like?
  - What variety of outcomes can we expect for different masses?

Runaway

- Runaway driven by 3-alpha and $^{12}C(\alpha,\gamma)^{16}O$
  - Next set of calculations will use a bigger network

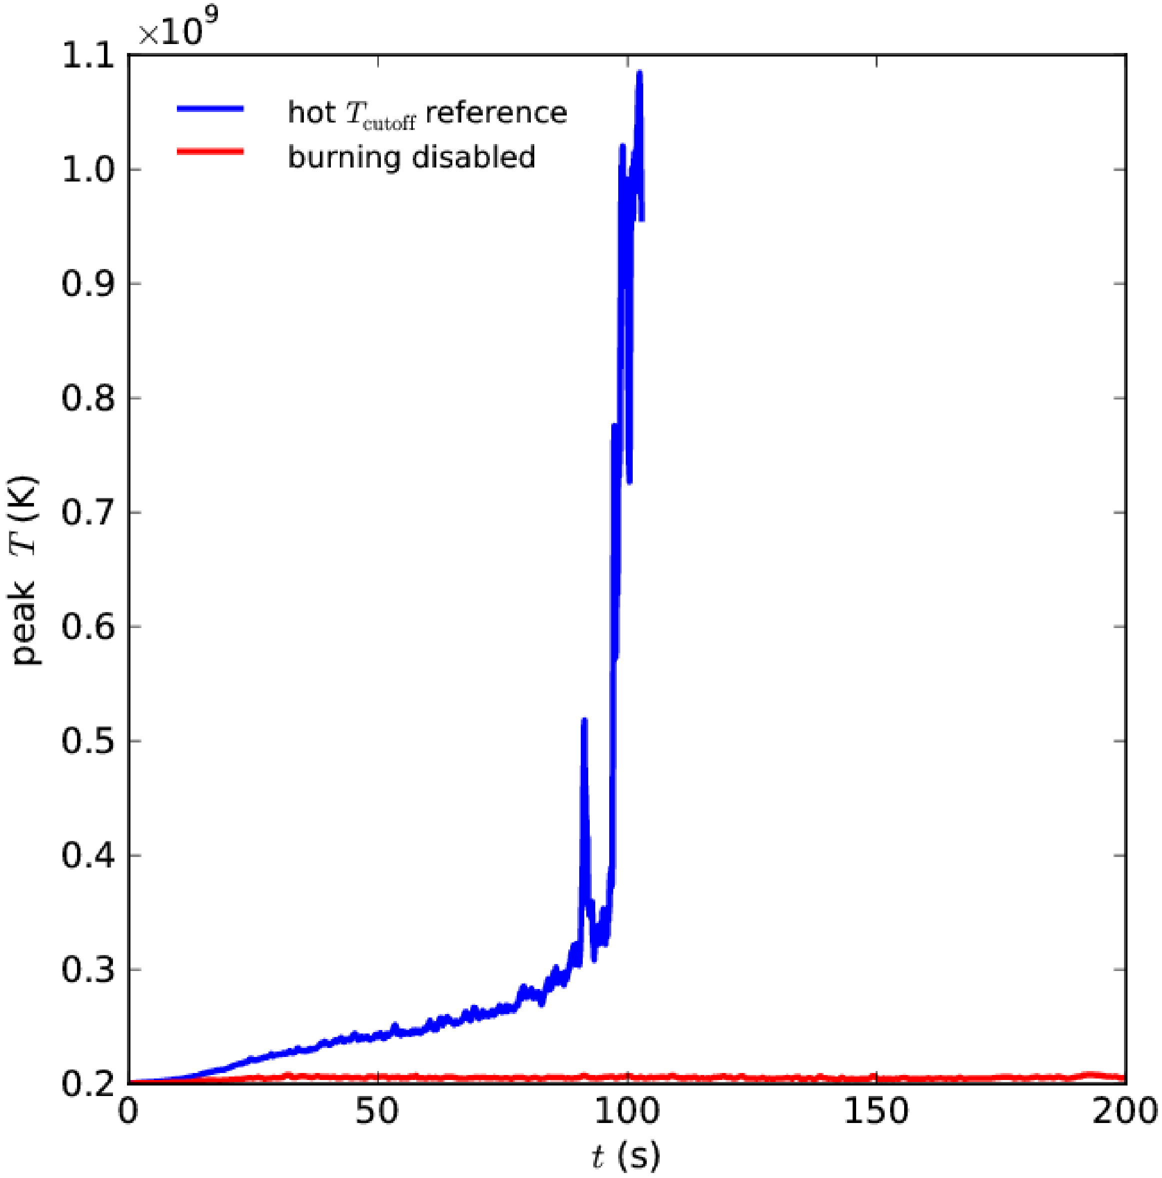

It's always a good idea to run your calculations without any burning—just to see how quiet the star is.

Sub-Chandra He Convection

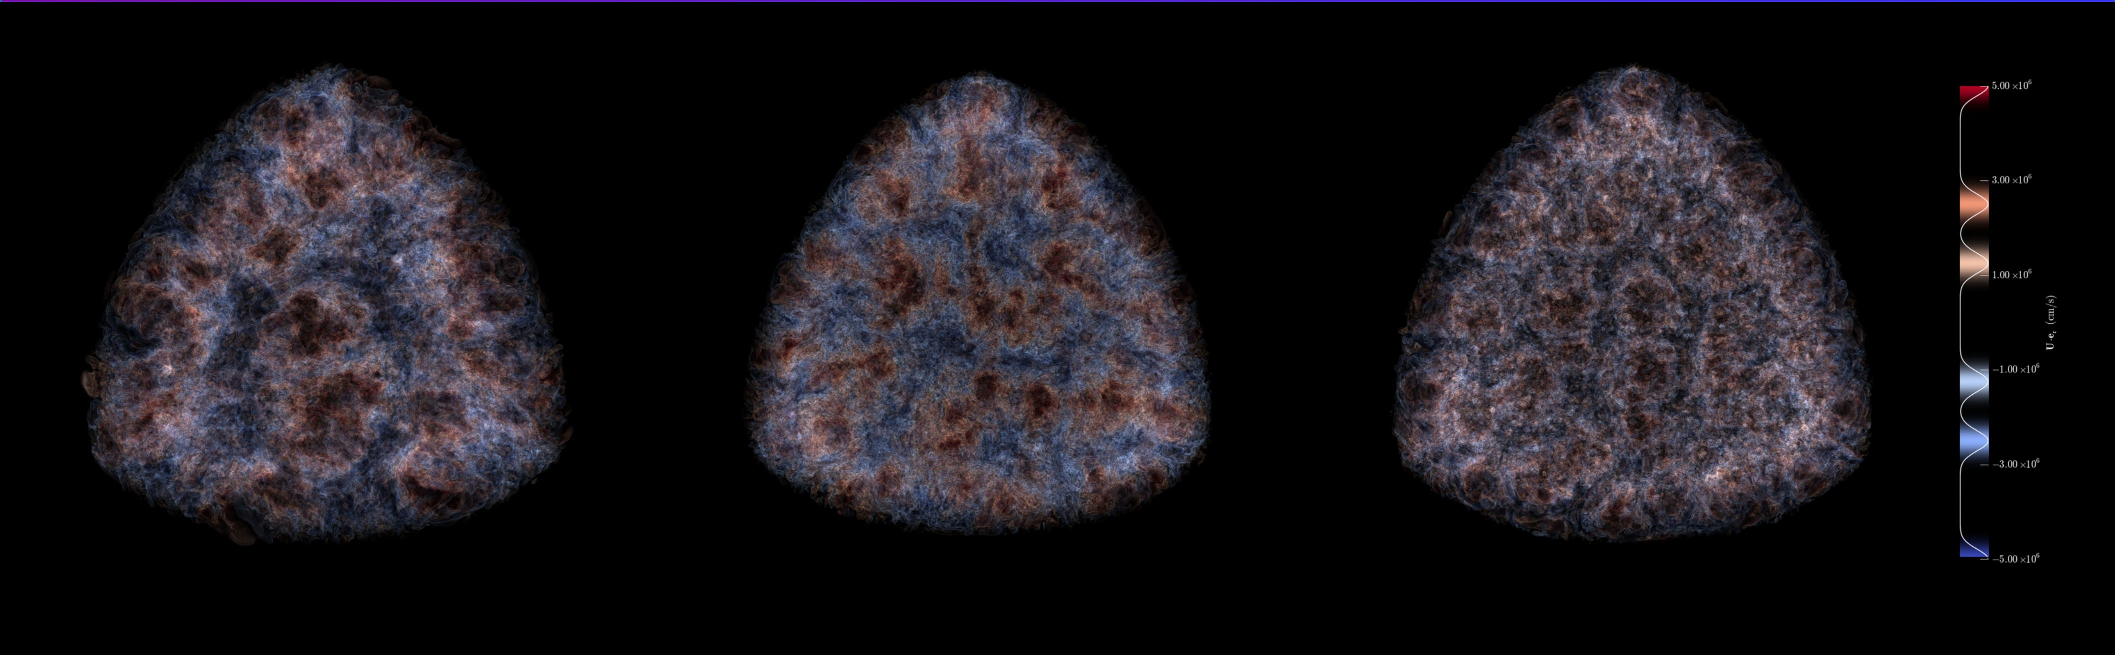

Visualization using yt

- Variations for different WD and He layer masses (Adam Jacobs's thesis)
- Cellular pattern forums
  - Length scale converged with resolution
  - Hot spots rise up and expand
- Potentially multiple hot spots simultaneously
- Three types of outcomes
  - Localize runaway on short timescale
  - Nova-like convective burning
  - Quasi-equilibrium (?)

Refs:
Zingale et al. 2013
Jacobs et al. 2016

Mergers

- We are also working on WD mergers (Max Katz thesis)
- Focus is comparison of high-resolution grid-based simulations to other studies to understand numerical challenges

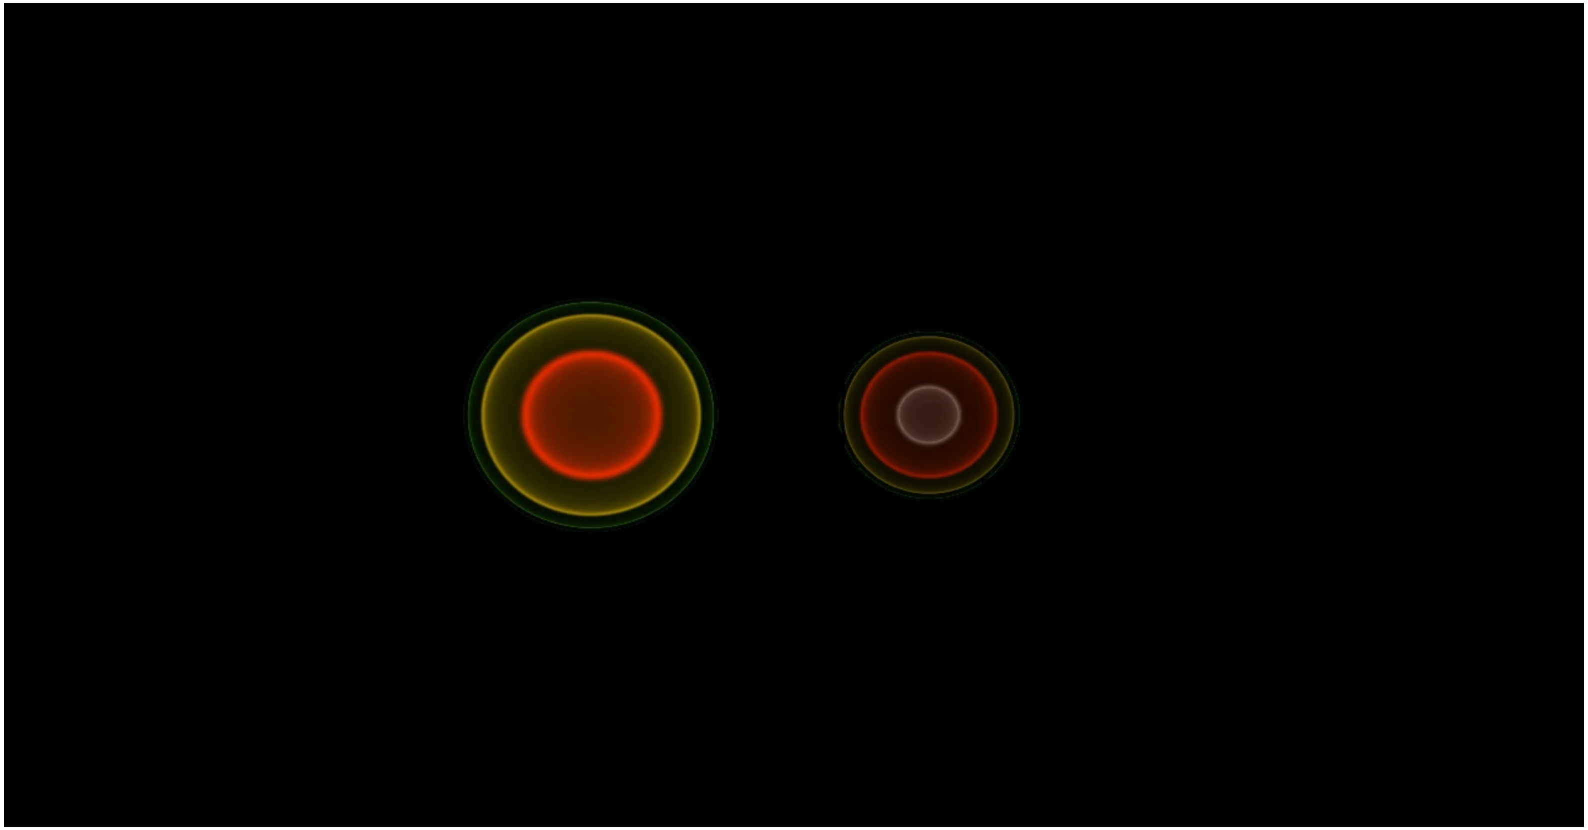

Visualization using yt

Refs:
Katz et al. 2016

X-ray Bursts

- Thermonuclear runaway in thin accreted H/He layer on surface of a neutron star
- Accretion timescale ~ hours to days
- Runaway timescale ~ seconds
- > 70 sources known, some with 10s or more individual bursts.
- Potential site for rp-process nucleosynthesis

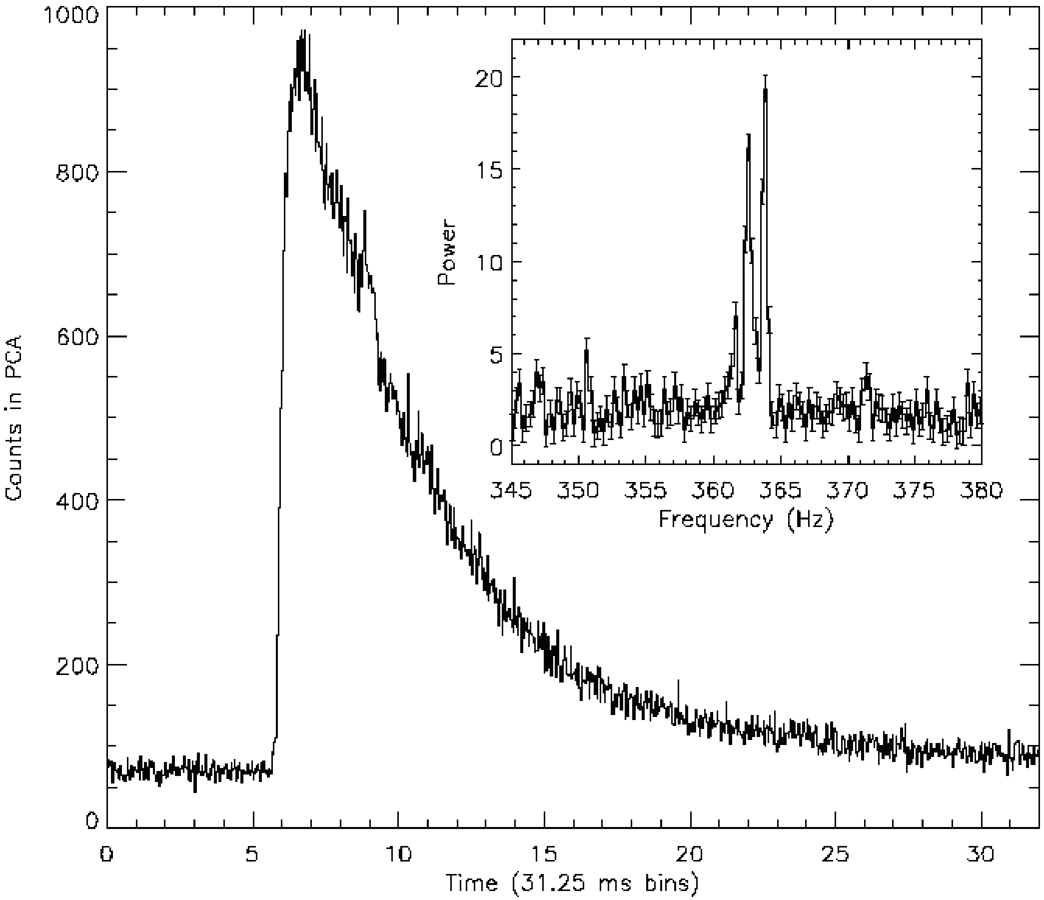

Strohmayer et al., 1996, ApJ, 469:L9

Outstanding XRB Questions

- How does the fuel spread over the surface?
- How does the ignition begin?
- Is the burning localized?
- Does convection modify the nucleosynthesis?
- What are the effects of rotation?
- Does convection bring ash to the surface?

These are all multi-dimensional effects

X-ray Bursts

- Current calculations:
  - 512 × 512 × 768 zones
  - 6 cm resolution
  - 11 nuclei network
    - Captures H burning (hot CNO), 3-α, rp-process breakout
  - T increase over 10^9 K, evolve for 0.02 s
- Next steps:
  - Bigger domains
  - Variety of initial models

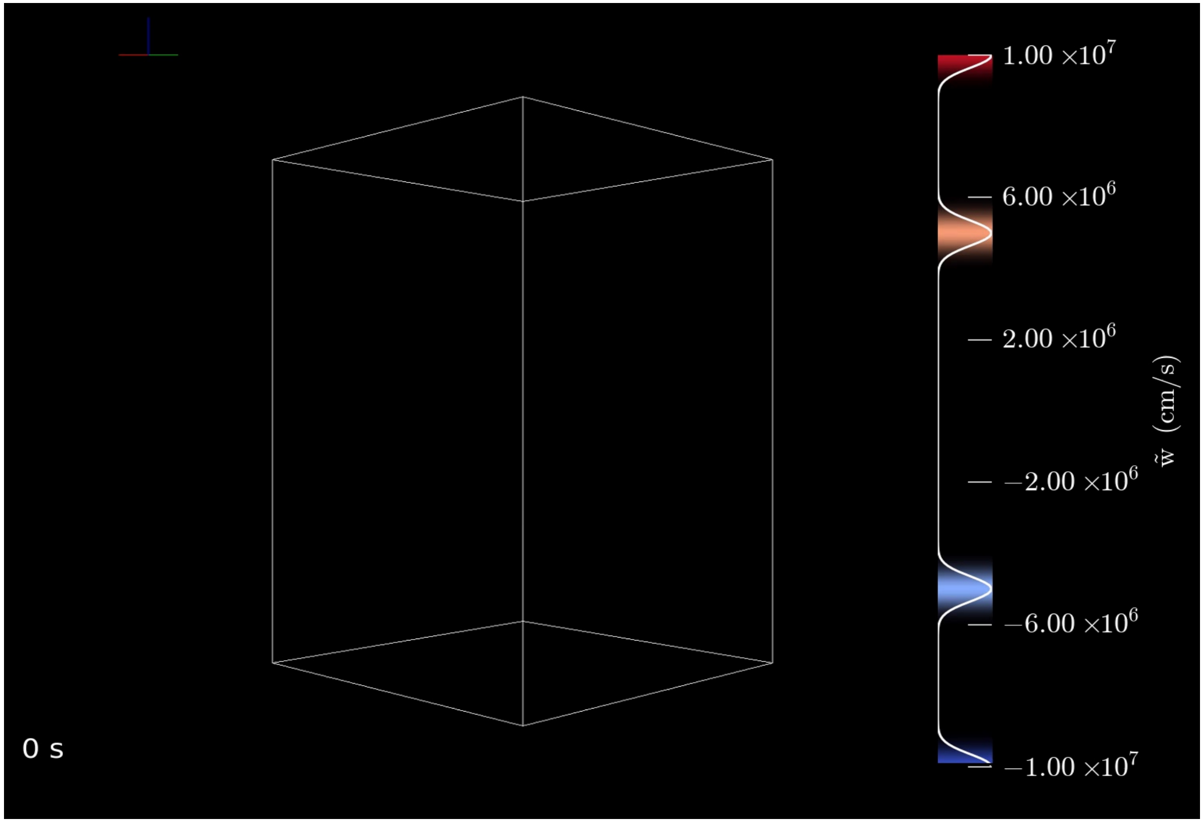

Visualization using yt

Refs:
Malone et al. 2011
Malone et al. 2014
Zingale et al. 2015

GPU Offloading

- Basic GPU strategy
  - Use unified memory in the backend to hide data motion complexity
  - Standard MFIter loops over boxes are unmodified
  - Fortran kernels require minimal changes
    - Standard “loop over zones in a box” approach is used unmodified
- Success requires that we keep data resident on device as much as possible
- Initial target: hydrodynamics
  - Goal: entire hydro update on GPUs
  - A simplified proxy-app of Castro’s hydrodynamics was made: StarLord
  - Performance on summitdev: 3-d hydrodynamics on 1 Pascal GPU is 70x faster than a single Power8 core
- Presently we rely on CUDA Fortran; future may use OpenACC, OpenMP

Accelerating Microphysics

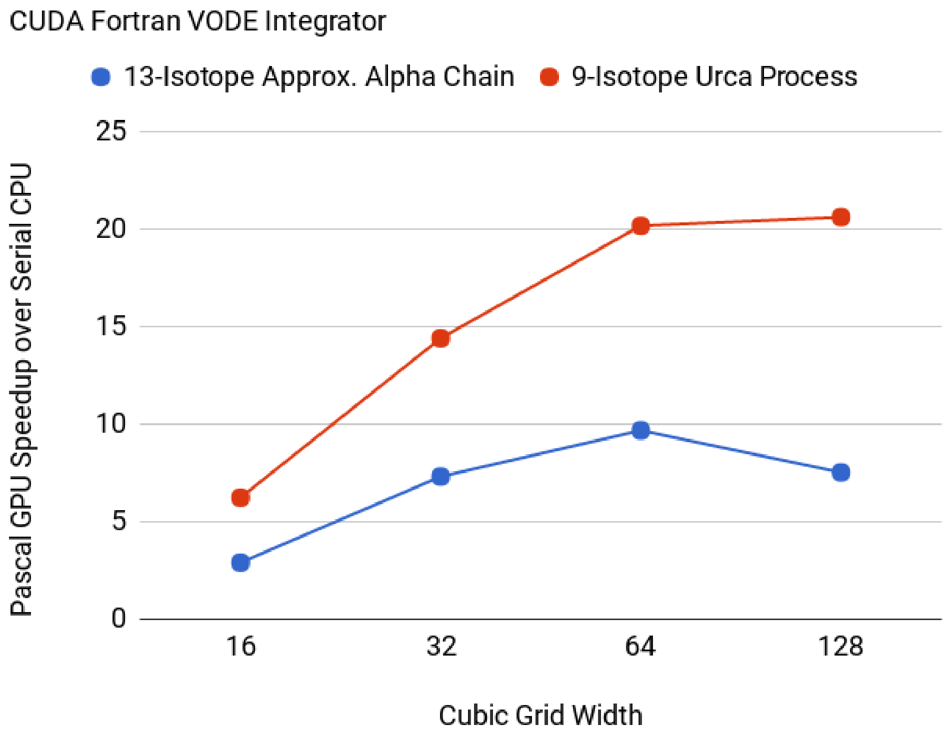

- Integrating reaction networks makes up a large fraction of runtime for some problems
- Reaction network is a system of ODEs (one for each nucleus)
  - Stiff system—needs implicit methods
  - Cost is in evaluation of RHS/Jacobian
- Strategy: do the entire ODE integration on GPUs
  - State data moved at start
  - ODE driver manages timestep (subcycling), calling RHS, Jacobian, …
  - Results moved to CPU at the end
- Different zones can have vastly different burn complexity
  - Some reactions scale as T^{40}—very stiff

Pyro: Hydro by Example

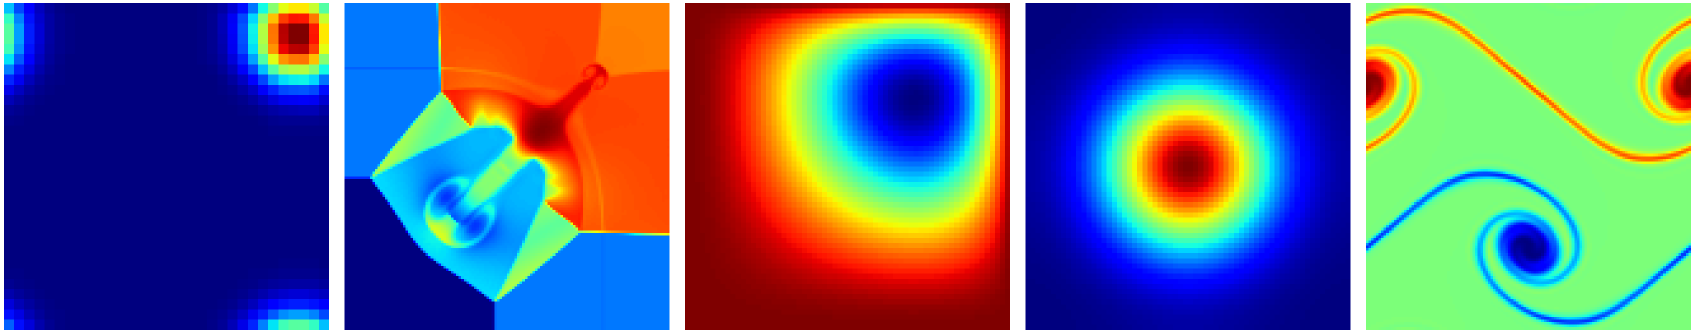

- A “python hydro” code designed with clarity in mind to teach students about simulation techniques
- 2-d solvers for:
  - Linear advection
  - Compressible hydrodynamics
  - Elliptic equations (via multigrid)
  - Implicit diffusion
  - Incompressible hydrodynamics
  - Low Mach number atmospheric flows
  - Coming soon: gray flux limited diffusion radiation hydrodynamics
- BSD-3 licensed, up on github: https://github.com/zingale/pyro2

Intro to Comp Hydro for Astro

- Current contents:
  - Simulation Overview
  - Classification of PDEs
  - Finite Volume Grids
  - Advection
  - Burgers’ Equation
  - Euler Equations: Theory
  - Euler Equations: Numerical Methods
  - Elliptic Equations and Multigrid
  - Diffusion
  - Multiphysics Applications
  - Reacting Flows
  - Planning a Simulation
  - Incompressible Flow and Projection Methods
  - Low Mach Number Methods

- Freely available under a CC license
- Code for every figure and method shown is provided
- Contributions welcomed via github issues and PRs

Summary/Future

- Astrophysical modeling requires the cooperation of many different domain scientists
- Chandra SNe Ia:
  - single-point, off-center ignition
  - convection calculations in the Chandra used for explosion
  - Urca process calculations are starting
- Sub-Ch SNe Ia: variety, likely single-point…
- WD mergers: infrastructure developed to allow for accurate, high-resolution modeling
- XRBs: well-developed convective field realized
  - Working on understanding flame propagation now
- Maestro development directions: rotation, higher-order, acoustics, MHD, rotation, ???
- Castro development: improved radiation, better coupling, GPU hydro
- Releasing simulation codes / problem files is part of scientific reproducibility