Petascale Simulations of Type Ia Supernovae

Andy Nonaka
Lawrence Berkeley National Laboratory
May 8, 2011
Bay Area Scientific Computing Day


Galaxy NGC 4526 imaged by the Hubble Space Telescope (www.nasa.gov)
60 million light years away; located in the Virgo constellation
SN1994D (Type Ia supernova)
The supernova is as bright as the host galaxy!

Outline

- Motivation: Type Ia Supernovae
- MAESTRO: Low Mach Number Astrophysical Solver
  - Mathematical formulation to exploit time scales of interest
  - Adaptive Mesh Refinement (AMR)
  - Massively parallel (100,000 cores) implementation
- Primary Collaborators
  - Ann Almgren, John Bell, Mike Lijewski, Candace Gilet: LBNL
  - Mike Zingale, Chris Malone: Stony Brook University

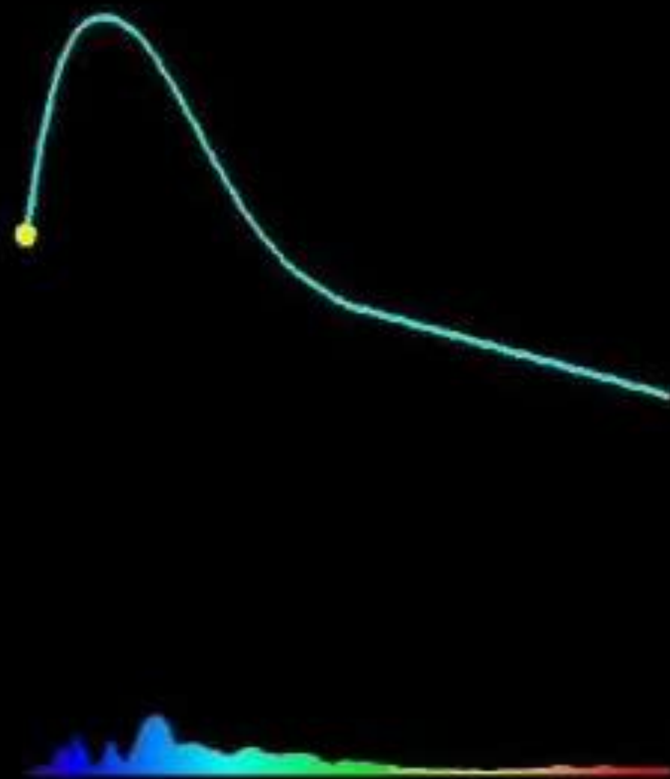

- Using modern telescopes, Type Ia supernova light curves can now be observed several hundred times per year:
  - Spectra contains silicon, lacks hydrogen
  - Peak powered by radioactive decay of nickel

Type Ia Supernovae are Distance Indicators

- By observing Type Ia supernovae Ia at known, nearby distances, scientists have established a width-luminosity relationship; wider = brighter.

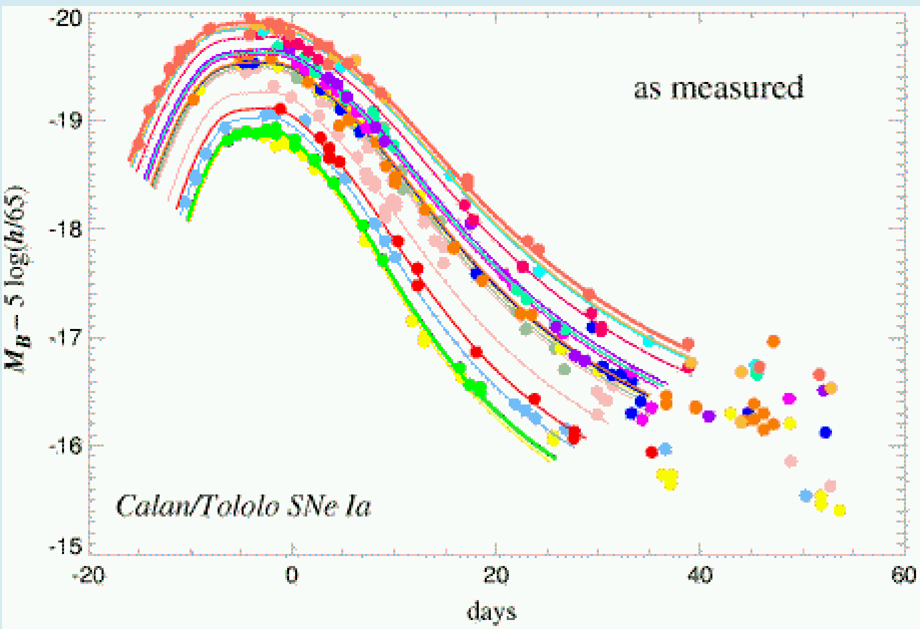

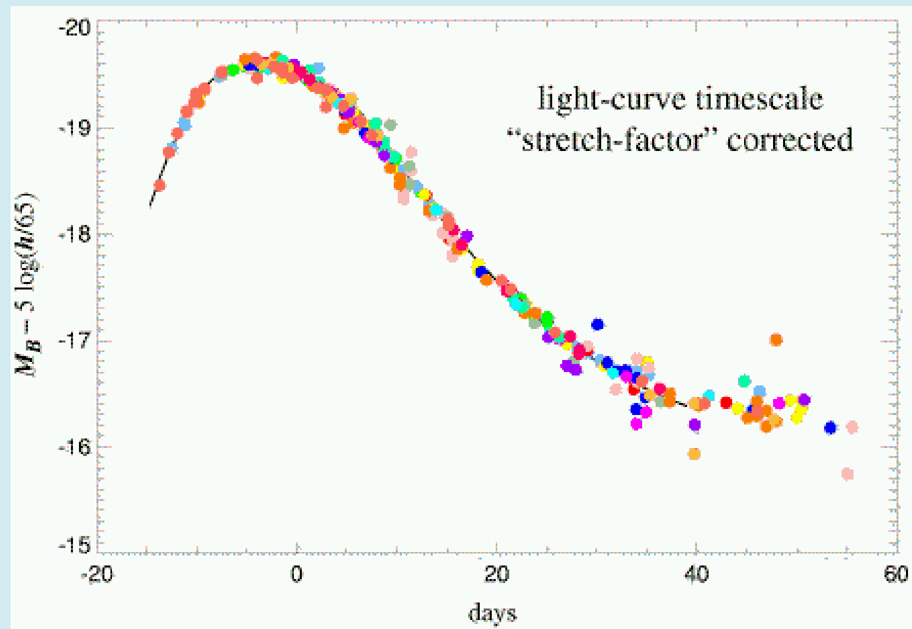

- Theory: by observing the peak luminosity and decay rate, we can determine the distance to a host galaxy.
  - Particularly useful for mapping distant galaxies since they are so bright!

Type Ia Supernovae are Speed Indicators

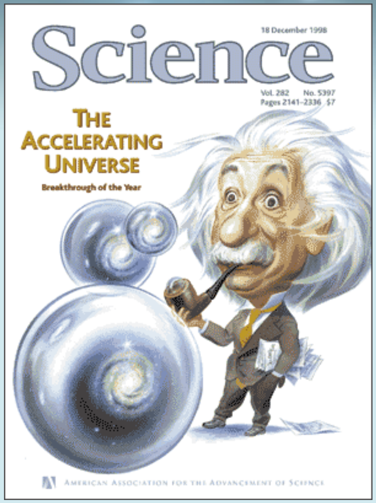

- Due to the observed redshift, we know the speed at which the host galaxy is moving away from us.
  - Led to discovery of the acceleration of the expansion of the universe (1998)
- Problem: We don't know how well the width-luminosity relationship holds for distant Type Ia supernovae
  - Farther away = earlier in the life of the universe
  - Composition of stars was different back then…
  - Not even sure if accepted models properly describe nearby events…
- How will we study this problem using computers?

The Phases of Type Ia Supernovae: Single Degenerate Model

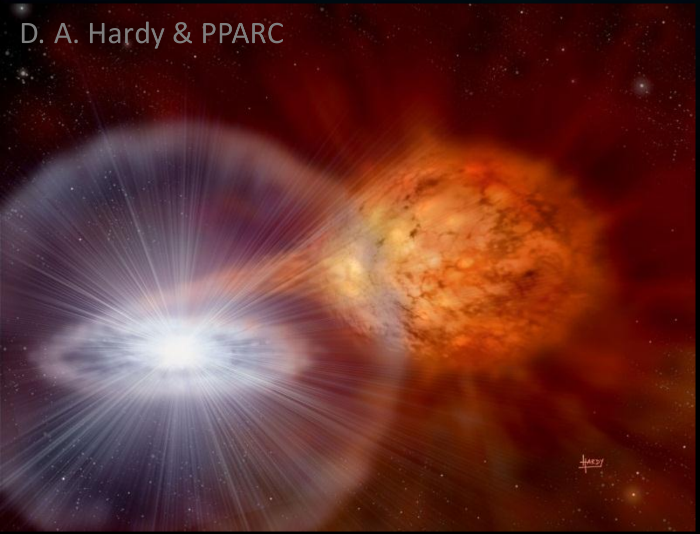

D. A. Hardy & PPARC

A white dwarf accretes matter from a binary companion over millions of years.

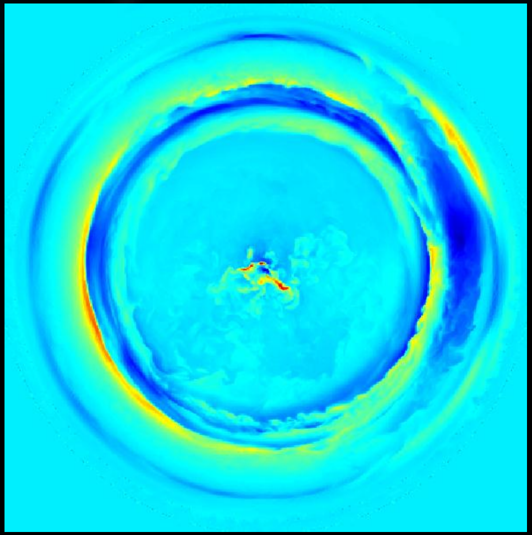

Smoldering phase characterized by subsonic convection and gradual temperature rise lasts hundreds of years.

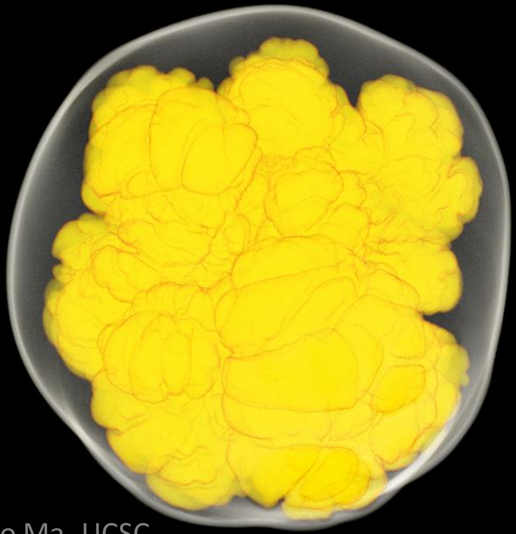

Haitao Ma, UCSC

Flame (possibly) transitions to a detonation, causing the star to explode within two seconds.

SN 1994D (High-Z SN Search team)

The resulting event is visible from Earth for weeks to months.

Each Phase has Different Computational Requirements

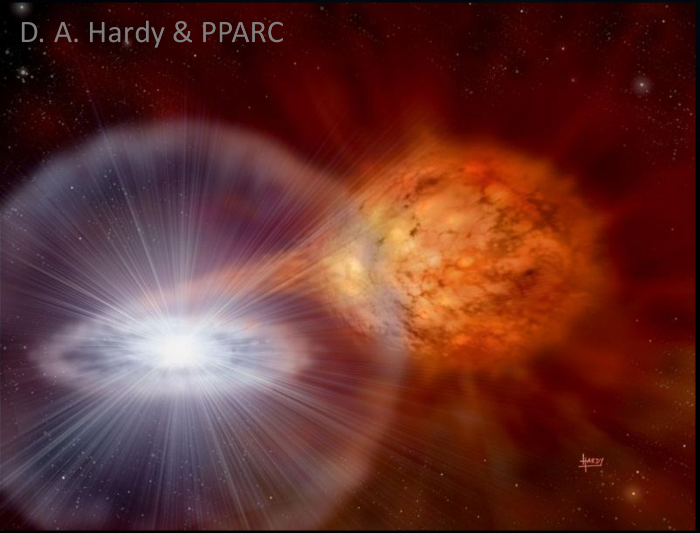
D. A. Hardy & PPARC

KEPLER
(Woosley, UCSC)

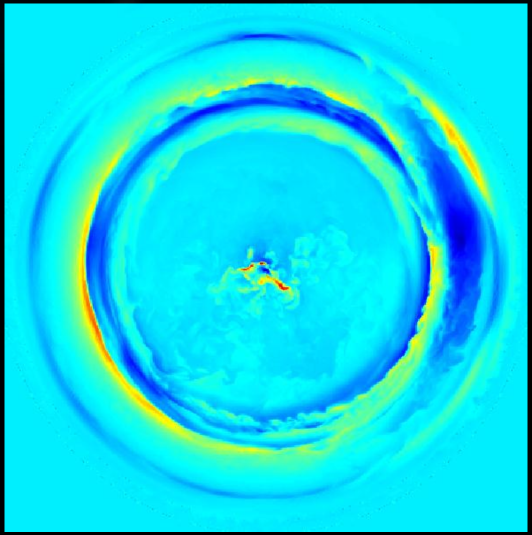
MAESTRO
(LBNL)

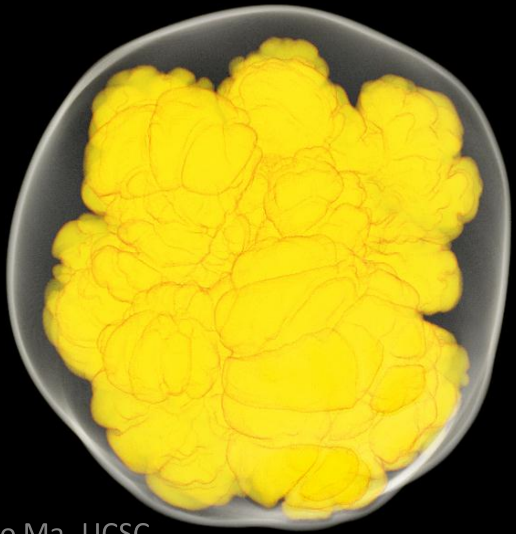
CASTRO (LBNL)

Haitao Ma, UCSC


SEDONA
(Kasen, Nugent, Thomas, LBNL)

SN 1994D (High-Z SN Search team)

Computing the Convective Phase

MAESTRO (LBNL)

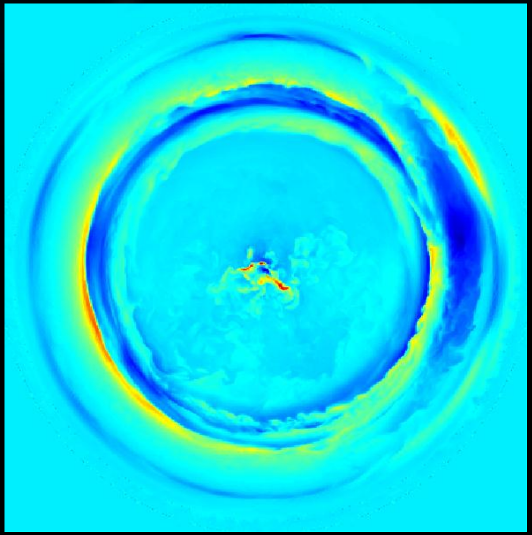

- In this talk we will focus on the last few hours of convection preceding ignition.
  - We wish to use MAESTRO to determine the initial conditions for the detonation / explosion phase for CASTRO
    - Previous studies have artificially seeded hot ignition points into their initial conditions
  - Low Mach number regime; $M = U/c$ is $O(10^{-2})$
  - Long-time integration infeasible using fully compressible approach

MAESTRO: Low Mach Number Astrophysics

- Algorithmic Details

What is MAESTRO?

- Key Theme
  - Having access to hundreds of thousands of CPUs is essential, but still not enough to solve this problem. We must also utilize a special mathematical formulation as well as AMR technology to address this problem.

What is MAESTRO?

- MAESTRO is a massively parallel, finite volume, AMR code for low Mach number astrophysical flows
  - Massively Parallel: Scales to 100,000 cores
  - Finite Volume: Solution in each Cartesian cell represents the average over the cell
  - AMR: Block-structured approach with logically rectangular grids

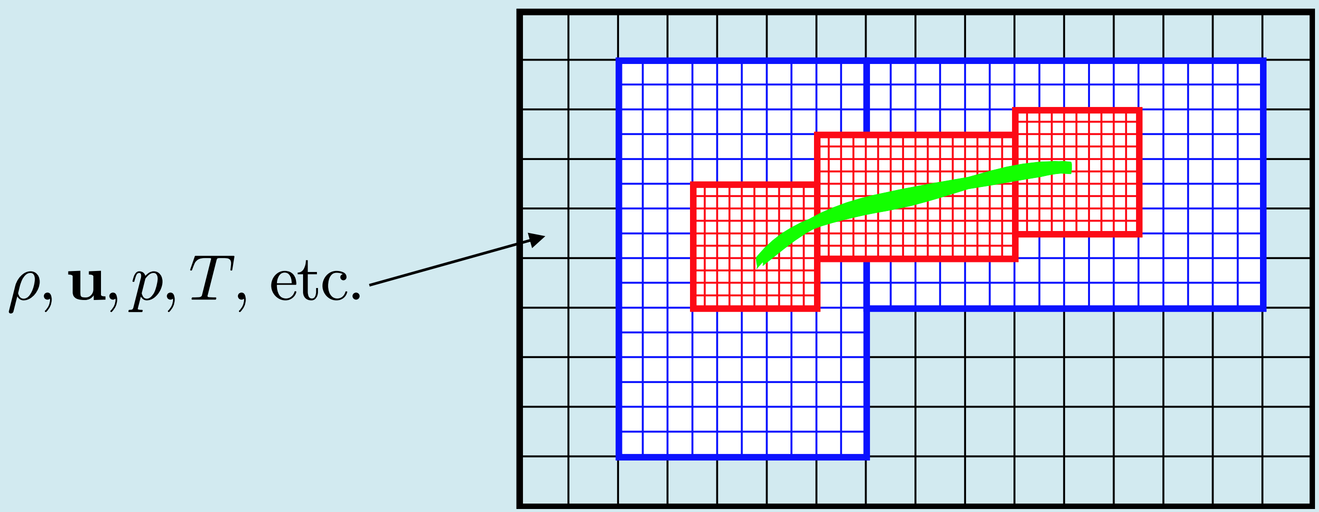

  - Low Mach Number: Fluid speed is small compared to the speed of sound
  - Astrophysical Flows: Modular equation of state and reaction networks

Low Mach Number Equation Set

- Equation set derived using low Mach number asymptotics
  - Mach number: M = U/c
  - Looks similar to the standard equations of compressible flow, but sound waves have been analytically removed
    - Enables time steps constrained by the fluid velocity CFL, not the sound speed CFL:

$$\Delta t_{\mathrm{compressible}} < \frac{\Delta x}{|u| + c} \qquad \Delta t_{\mathrm{lowMach}} < \frac{\Delta x}{|u|}$$

    - Low Mach time step is a factor of 1/M larger than a compressible time step, enabling long-time integration!

Low Mach Number Equation Set

- Derived from fully compressible equation set

$$\frac{\partial(\rho X_k)}{\partial t} = -\nabla \cdot (\rho X_k \mathbf{u}) + \rho \dot{\omega}_k$$

conservation of mass

$$\frac{\partial(\rho \mathbf{u})}{\partial t} = -\nabla \cdot (\rho \mathbf{u}\mathbf{u}) - \nabla \pi + \rho \mathbf{g}$$

conservation of momentum

$$\frac{\partial(\rho h)}{\partial t} = -\nabla \cdot (\rho h \mathbf{u}) + \rho H$$

conservation of energy

ρ	density	h	specific enthalpy
$\mathbf{u}$	velocity	$\mathbf{g}$	gravity
X_k	mass fraction of species “k”	H	reaction heating
$\dot{\omega}_k$	reaction rate of species “k”	π	deviation from ambient pressure

Low Mach Number Equation Set

- Our system is closed with an equation of state, which keeps system in thermodynamic equilibrium.
  - Differentiate equation of state along particle paths to represent as a divergence constraint:

$$\nabla \cdot (\beta_0 \mathbf{u}) = \beta_0 S$$

$\beta_0 \rightarrow$ captures expansion/contraction of fluid due to changes in altitude

$S \rightarrow$ captures local compressibility effects due to reactions and thermal diffusion

- Numerical enforcement of divergence constraint analogous to solution methodology for incompressible flow
  - Pressure-projection method involving a variable-coefficient Poisson solve

Numerical Methodology

- Strang splitting couples advection/reaction/diffusion
  - Advection using Godunov approach
  - Reactions using stiff ODE solver
  - Diffusion semi-implicit (multigrid)
  - Divergence-constraint requires elliptic solve (multigrid)

$$
\begin{aligned}
\frac{\partial(\rho X_k)}{\partial t} &= -\underline{\nabla \cdot (\rho X_k \mathbf{u})} + \underline{\rho \dot{\omega}_k} \\
\frac{\partial(\rho \mathbf{u})}{\partial t} &= -\underline{\nabla \cdot (\rho \mathbf{u}\mathbf{u})} - \underline{\nabla \pi} + \underline{\rho \mathbf{g}} \\
\frac{\partial(\rho h)}{\partial t} &= -\underline{\nabla \cdot (\rho h \mathbf{u})} + \underline{\nabla \cdot \kappa \nabla T} + \underline{\rho H}
\end{aligned}
$$

$$
\underline{\nabla \cdot (\beta_0 \mathbf{u}) = \beta_0 S}
$$

Computational Efficiency

- In our white dwarf simulations, the Mach number is approximately $M \approx 0.05$
  - The low Mach number time step is a factor of **<u>70</u>** greater than a compressible time step
  - However, a low Mach number time step takes approximately 2.5 times longer to compute, mostly due to the linear solvers (multigrid)
- Thus, to advance the solution to the final time, MAESTRO is a factor of (70 / 2.5) ≈ 28 more efficient than a compressible algorithm, given the same number of computational resources for this problem

Adaptive Mesh Refinement

- Incorporate AMR using established techniques
  - Advance each level independently and synchronize fluxes, velocities, and pressure at coarse-fine interfaces

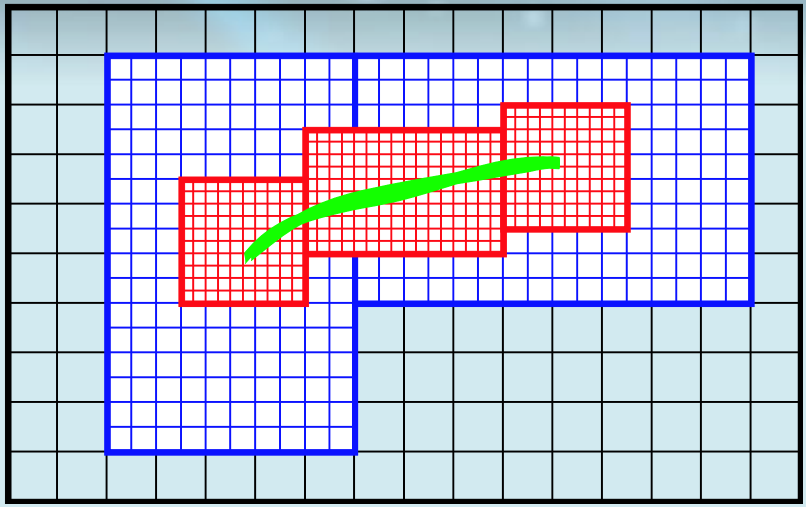

- For the full star problem, we need to consider our tagging criteria
  - Burning occurs near core, driving flow in the inner-convective region of the star.
  - We expect ignition point(s) to be near the center of the star

Adaptive Mesh Refinement

- 576^3 (8.7 km)
  - $1728 \cdot 48^3$ grids
  - 191 Million Cells

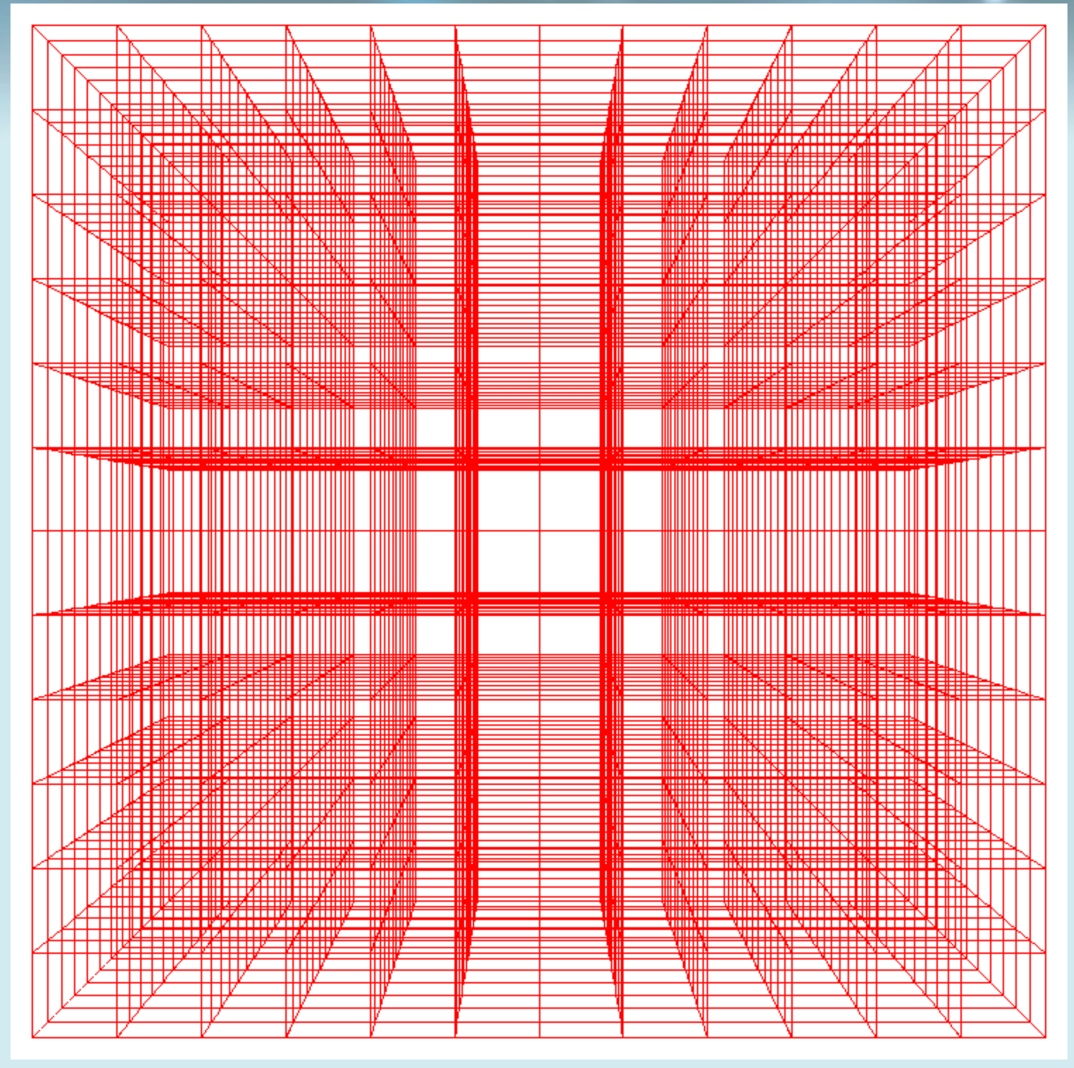

Adaptive Mesh Refinement

- 576^3 (8.7 km)
  - $1728 \cdot 48^3$ grids
  - 191 Million Cells

Edge of Star

Convective Zone Boundary

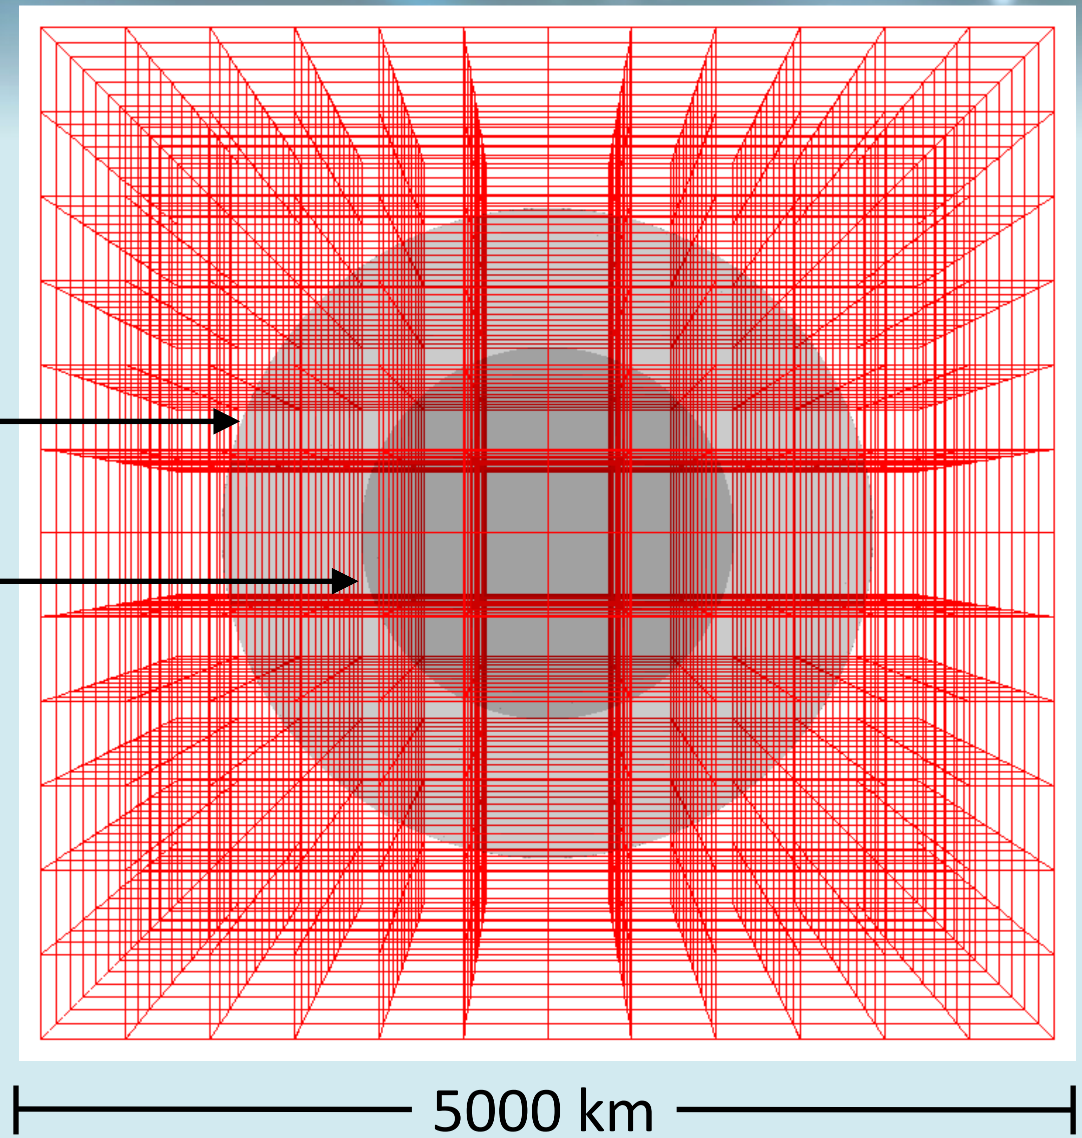

Adaptive Mesh Refinement

- 576^3 (8.7 km)
  - $1728 \cdot 48^3$ grids
  - 191 Million Cells

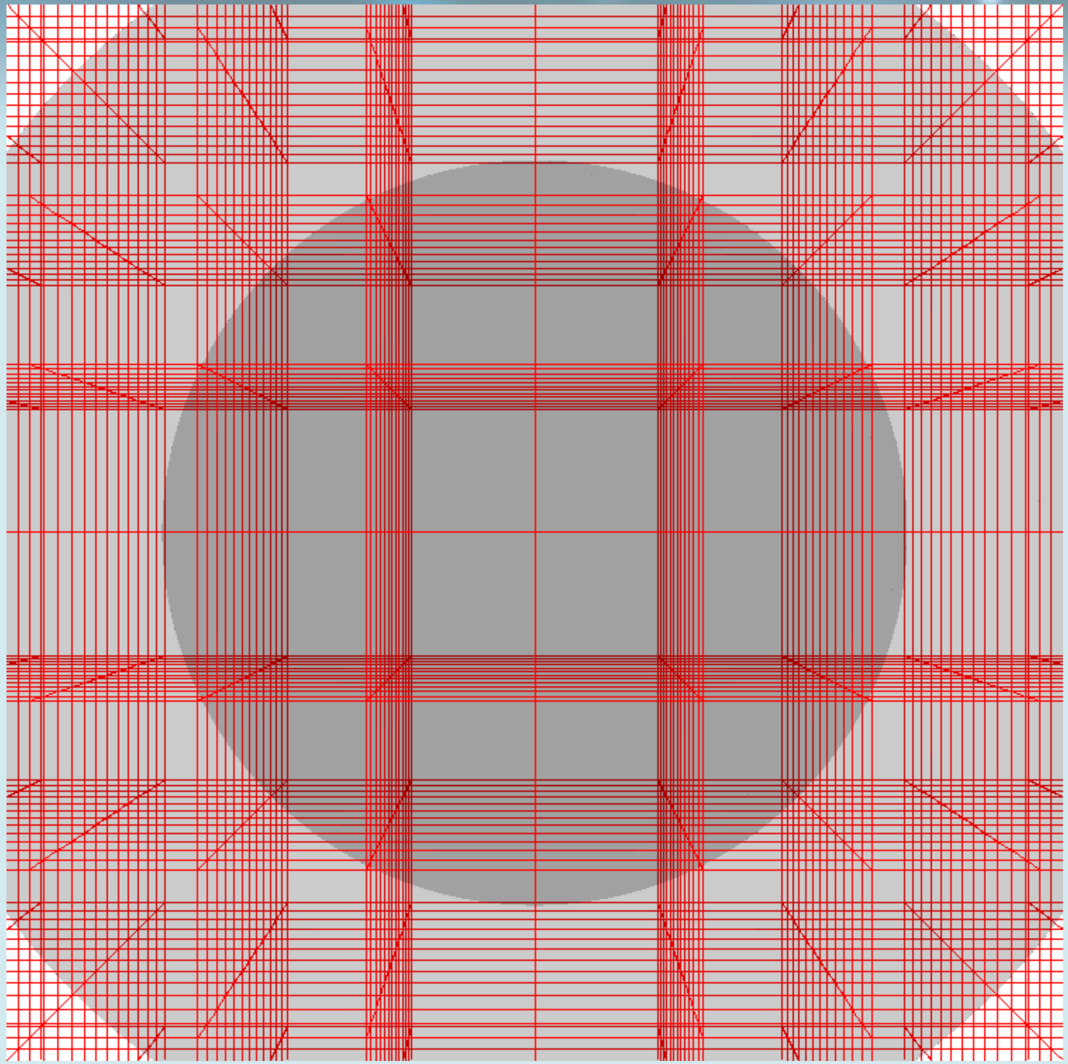

Adaptive Mesh Refinement

- 576^3 (8.7 km)
  - 1728 · 48^3 grids
  - 191 million cells
- 1152^3 (4.3 km)
  - 1684 grids
  - 148 million cells
  - 9.7% of domain
- 2304^3 (2.2 km)
  - 3604 grids
  - 664 million cells
  - 5.4% of domain

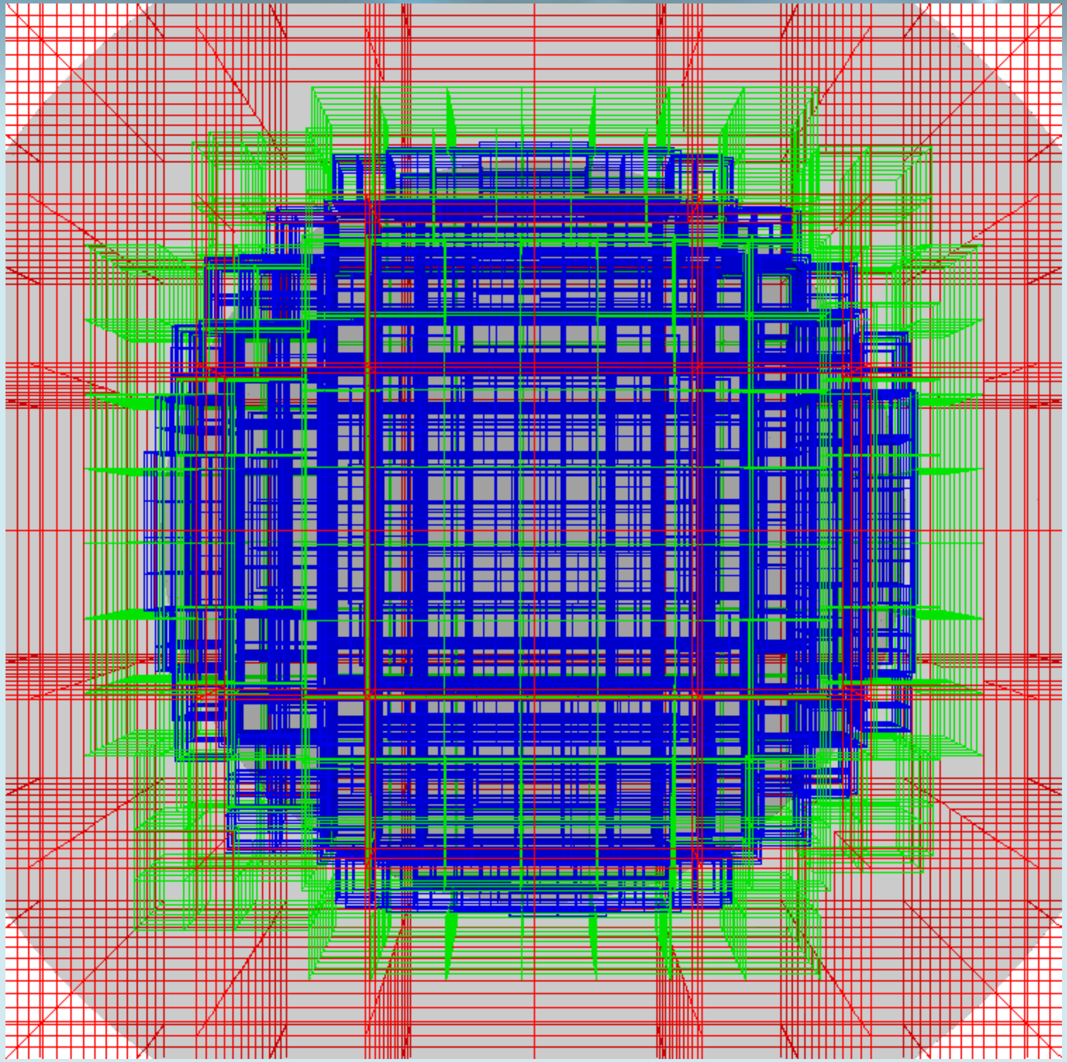

Adaptive Mesh Refinement

- A 2304^3 simulation with no AMR would contain 12.2 billion cells.
- Our simulation contains a total of 1.0 billion cells, requiring a factor of 12 less work.
  - Excluding AMR overhead, which is several percent

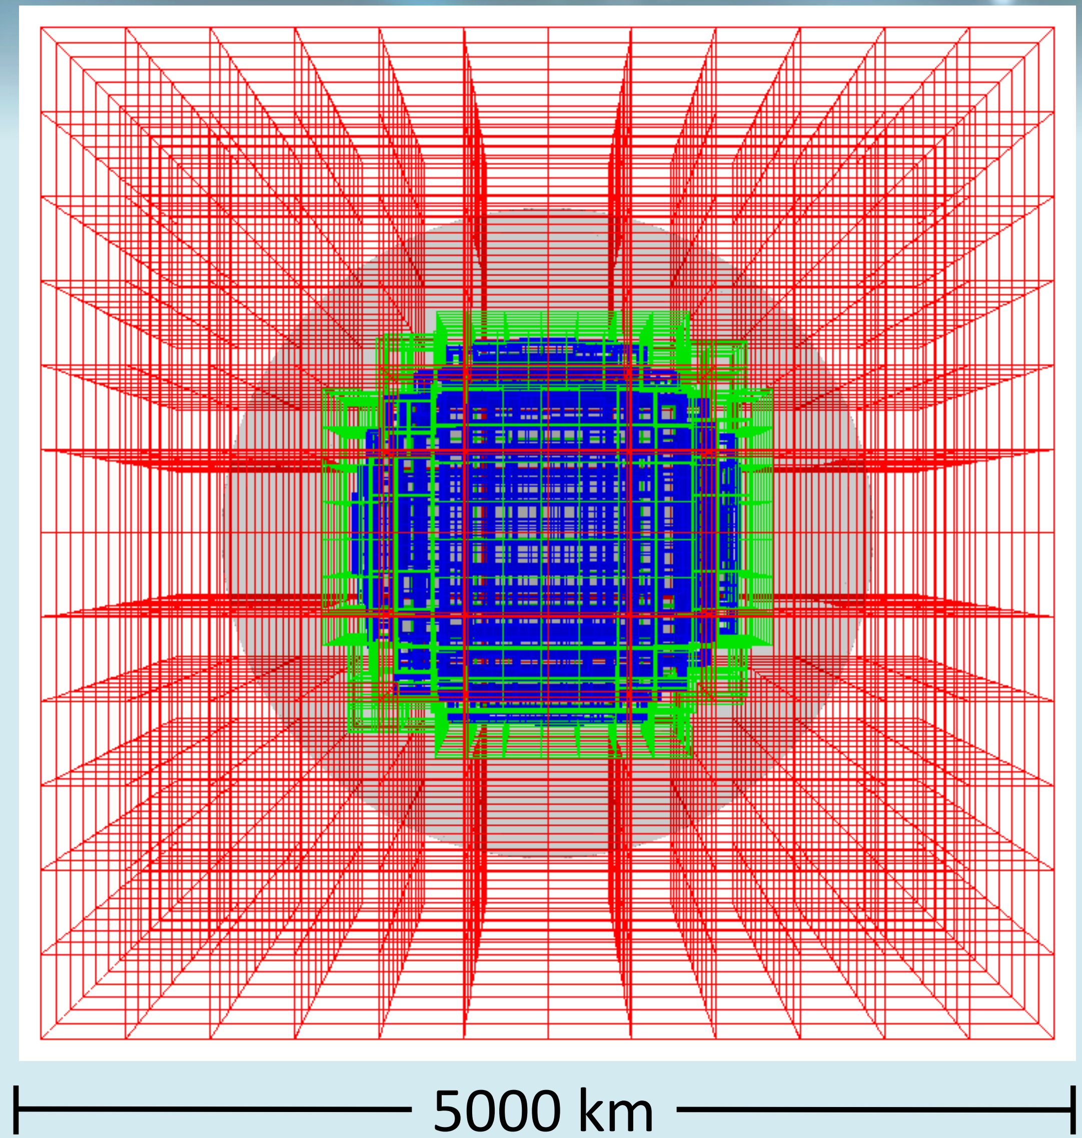

Parallelization Strategy

- Hybrid MPI/OpenMP approach to parallelization.
  - Nodes assigned to grids, threads spawned on cores to work on grids

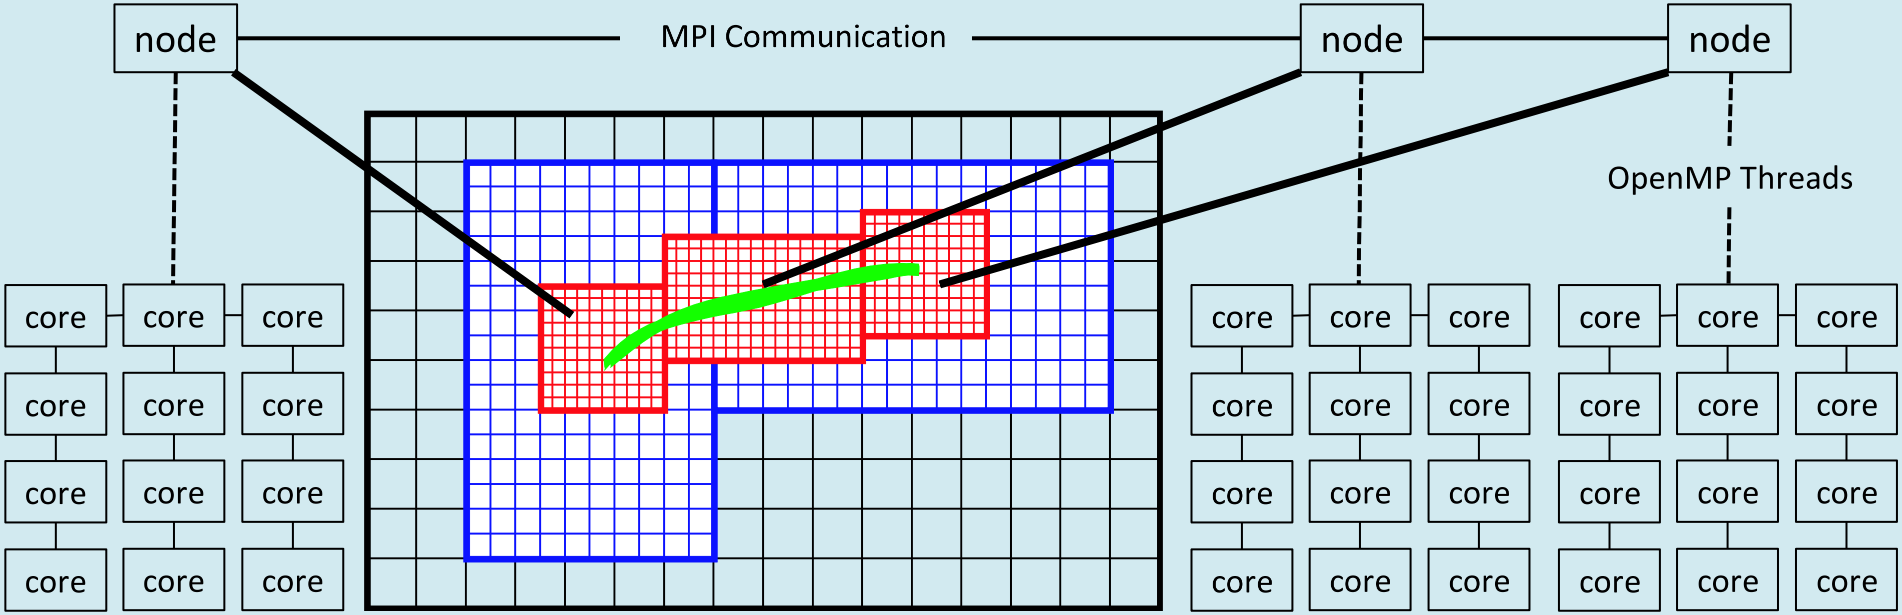

- Allows scaling to a factor of *nthreads* greater cores than pure MPI due to reduced communication time

- Weak scaling results for a 2-level Type Ia supernova simulation
  - Performed on jaguar at OLCF; 12 threads per MPI process
  - Each MPI process was assigned to a single 128^3 grid at each level
    - 768 processor simulation uses effective 1024^3 resolution
    - 96,000 processor simulation has effective 5120^3 resolution

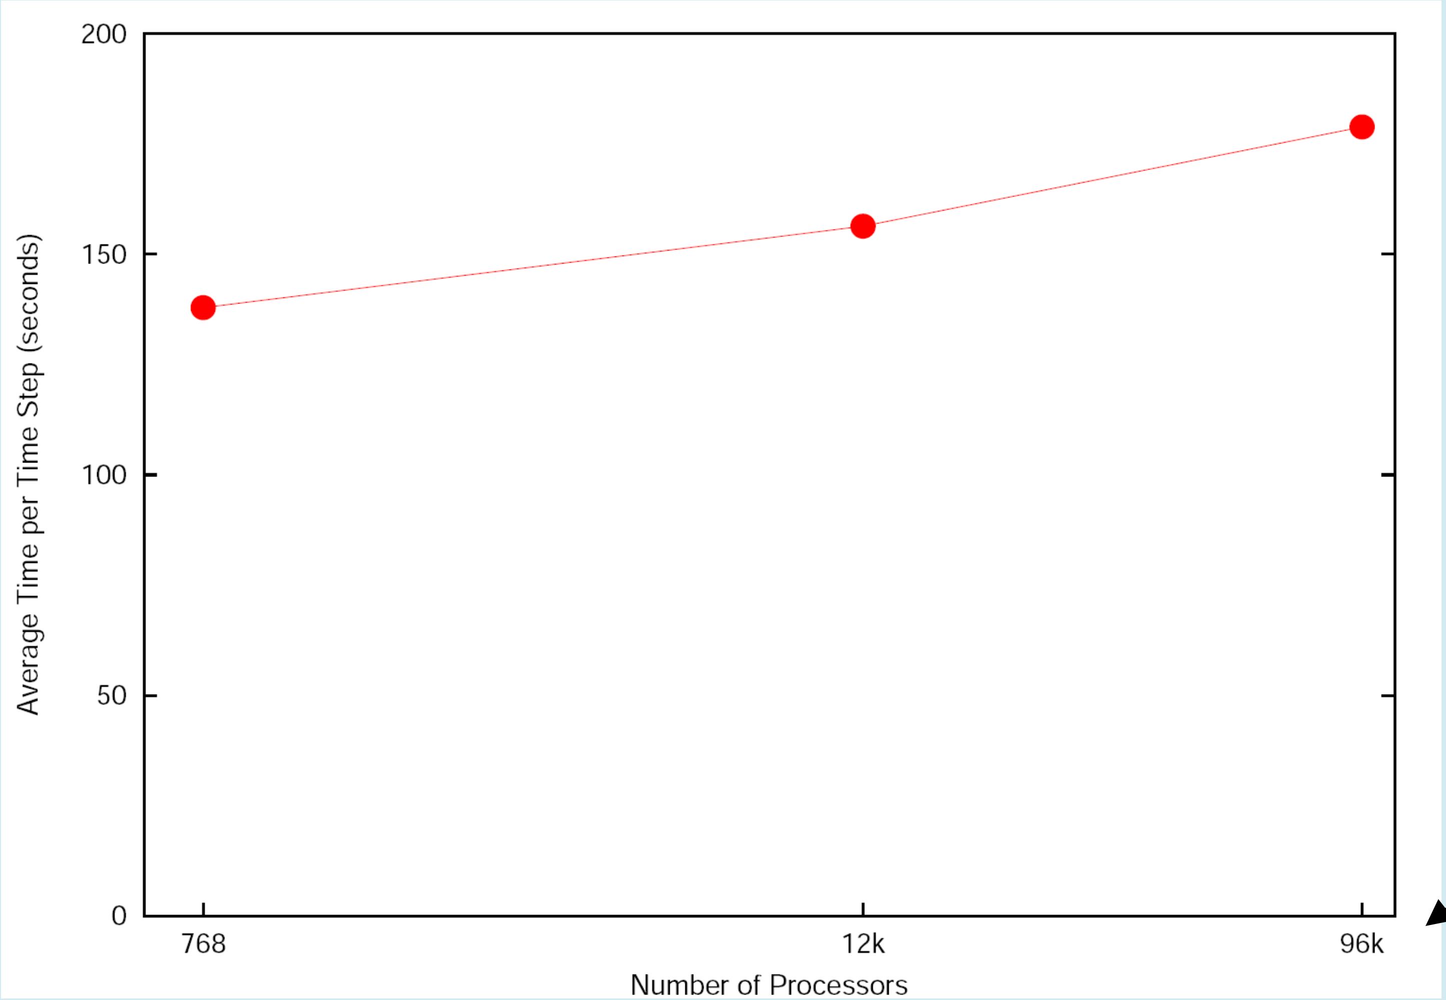

Weak scaling results using a different number of threads per core would scale the numbers on the x-axis

MAESTRO: Low Mach Number Astrophysics

- Scientific Results

White Dwarf Convection: Initial Conditions

- Initial conditions
  - 1D KEPLER model mapped onto Cartesian grid
  - Random velocity perturbation added to prevent initial nuclear runaway

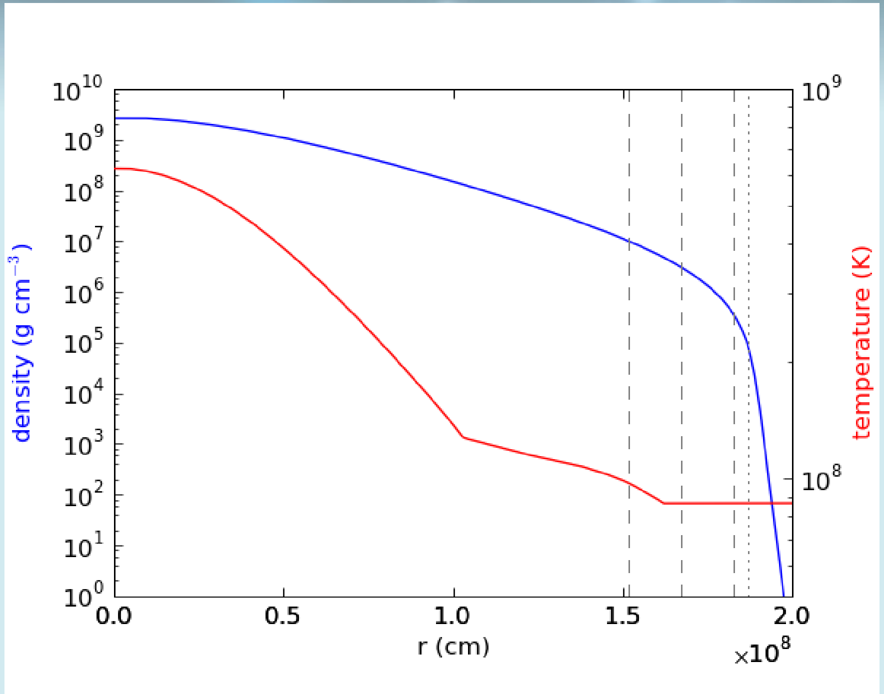

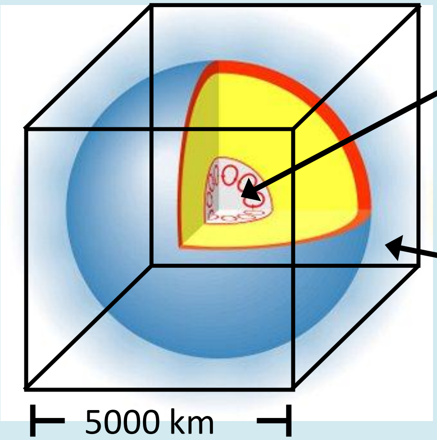

Center of Star
density = 2.6×10^{9} g/cc
Temperature = 6.25×10^{8} K

Edge of Star
density = 10^{-4} g/cc

- Use 10K cores for 40 days (10 million CPU hours) to run effective 1152^3 resolution (4.3km zones) to ignition

White Dwarf Convection: Long-Time Behavior

- Maximum temperature and Mach number vs. time

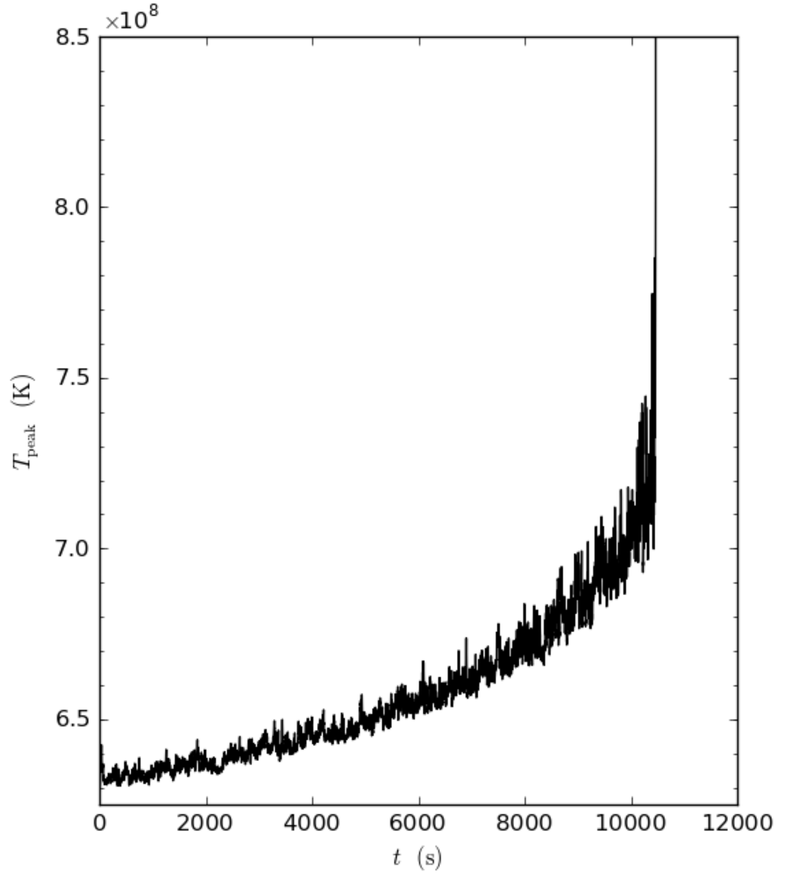

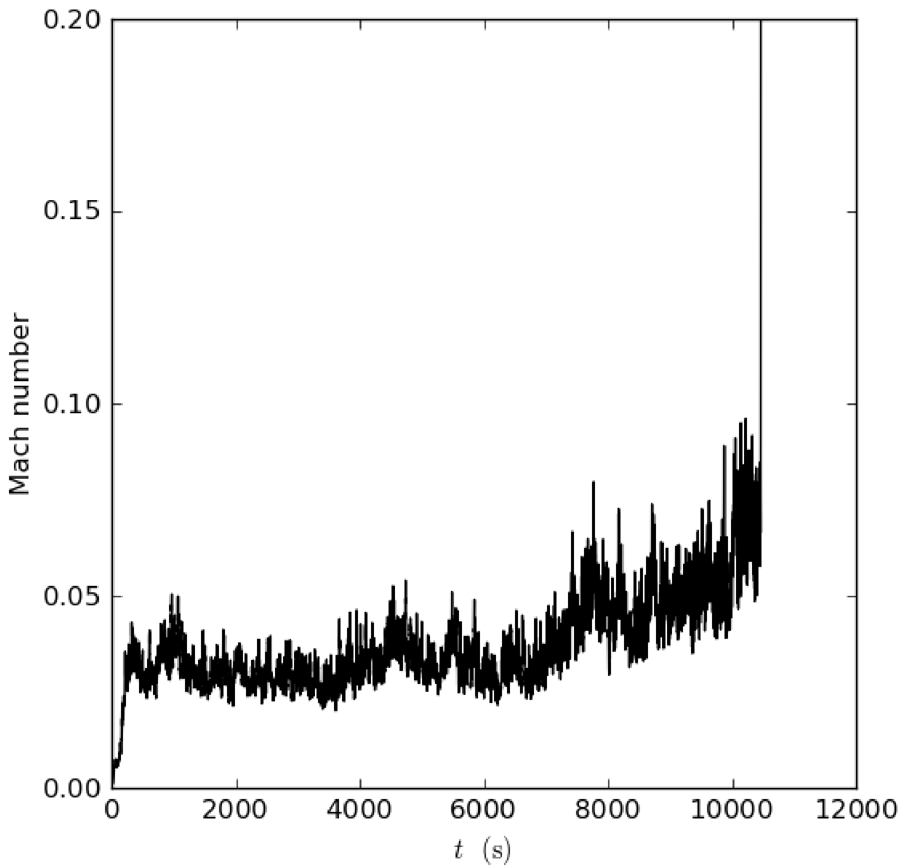

- Red / Blue = outward / inward radial velocity
- Yellow / Green = contours of increasing burning rate

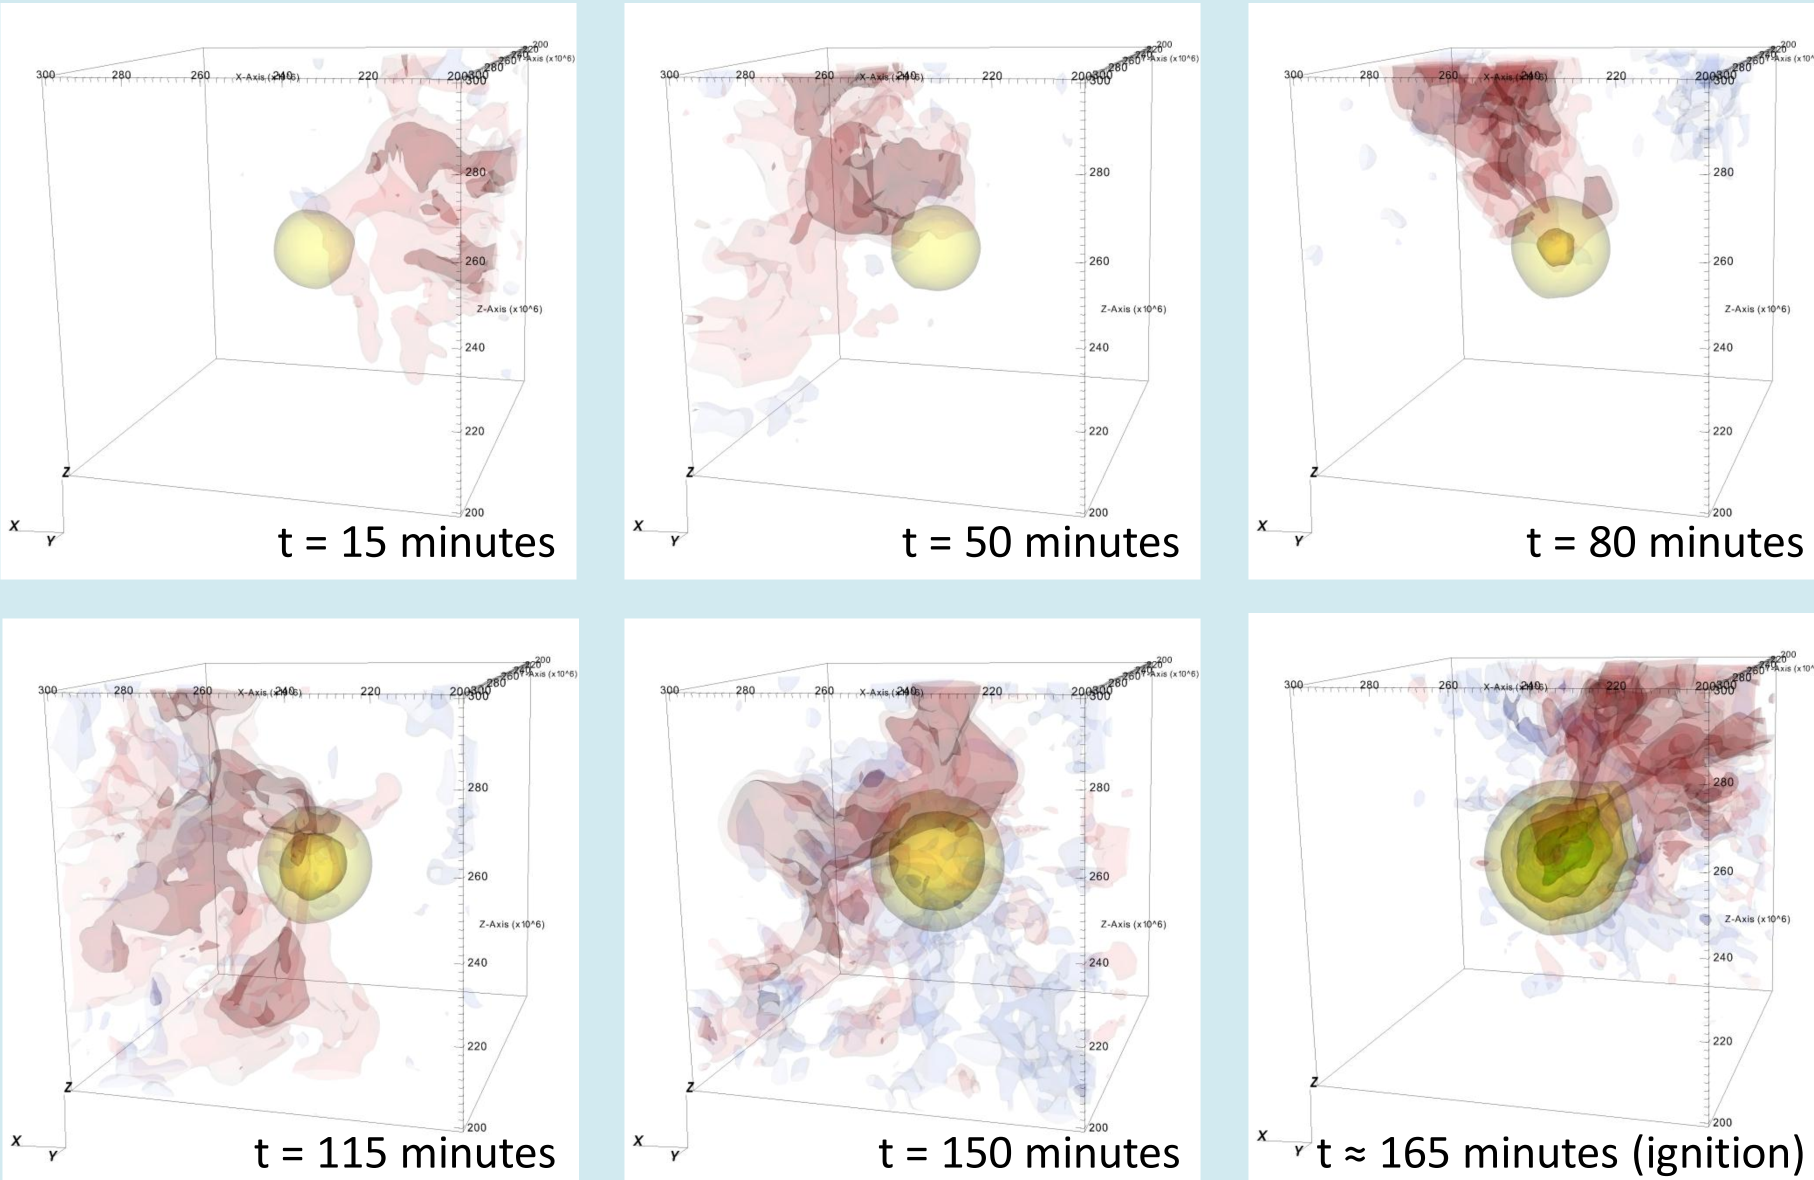

WD Convection: Ignition

- Convective flow pattern a few minutes preceding ignition
  - Inner 1000 km^3 of star
  - Effective 2304^3 resolution (2.2km) with 3 total levels of refinement
  - Red / Blue = outward / inward radial velocity
  - Yellow / Green = contours of increasing burning rate

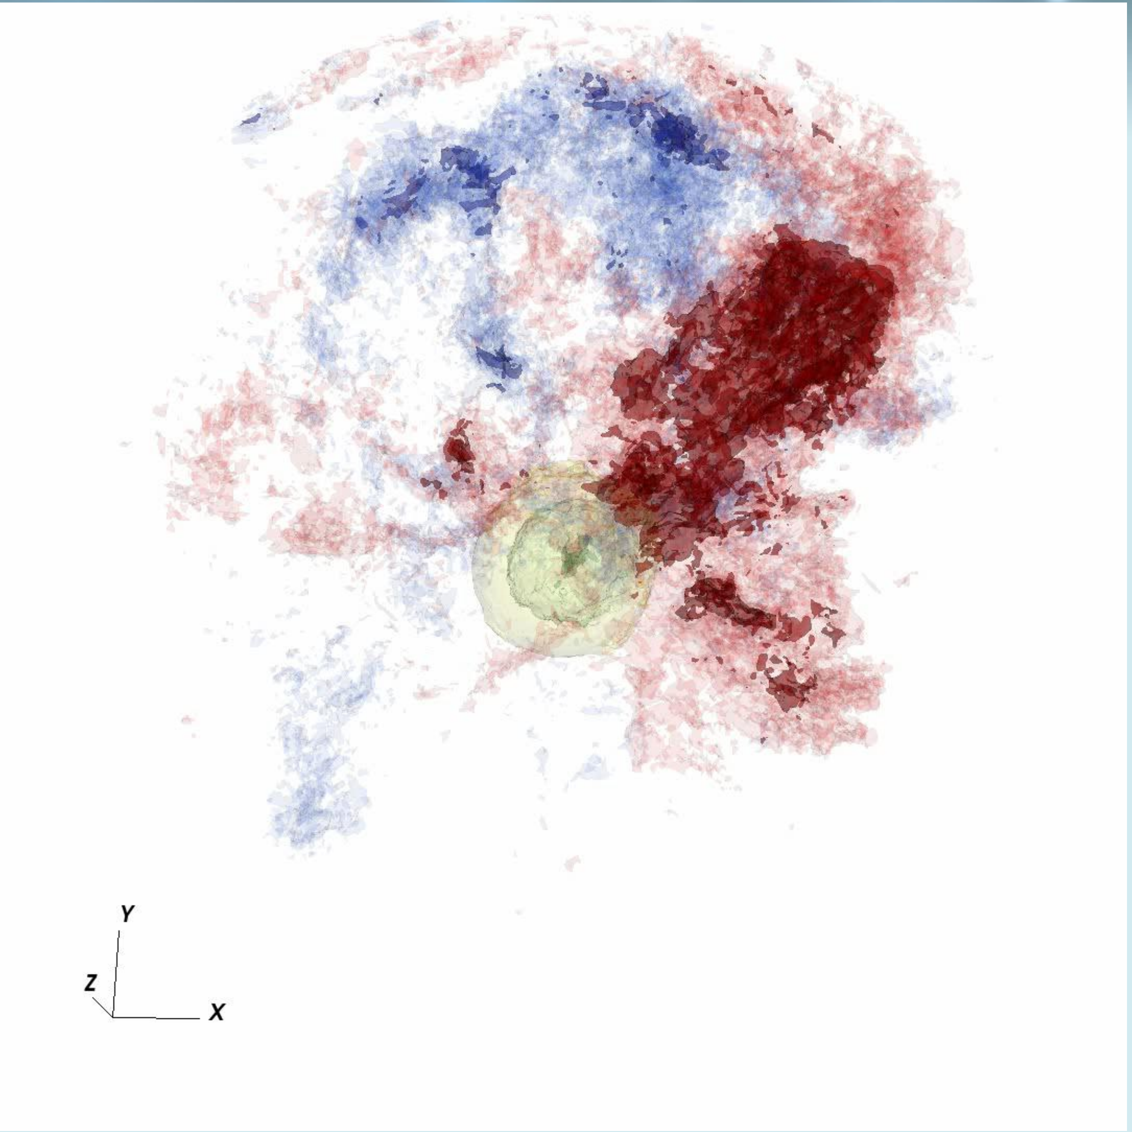

WD Convection: Ignition

- Same data from the previous simulation
- 2D slice of temperature profile a few minutes preceding ignition

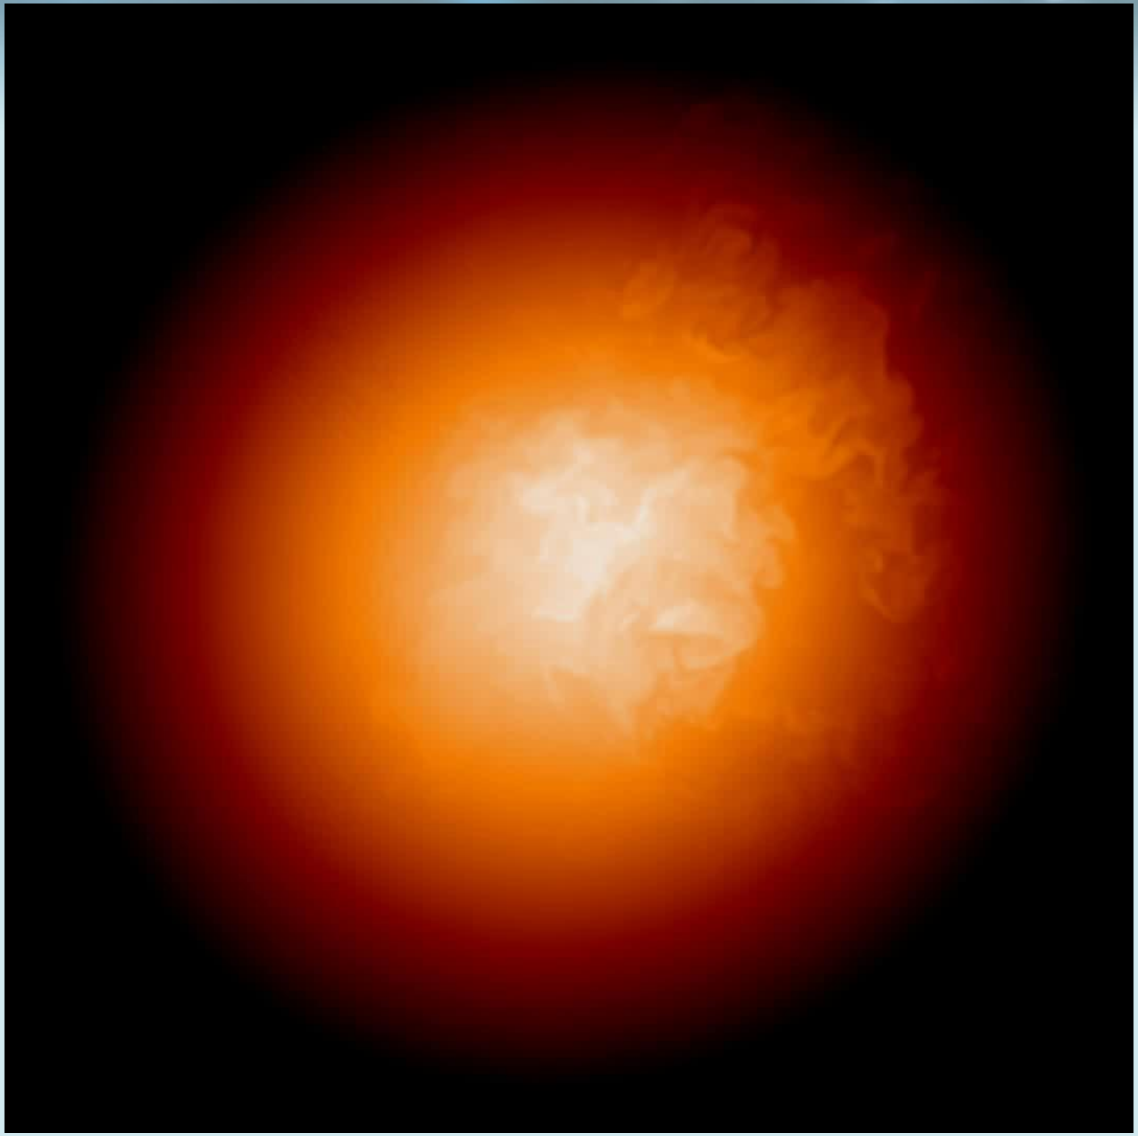

WD Convection: Ignition

- Examining the radius of the hot spot over the last few minutes indicates ignition radius of 50-70 km off-center is favored.

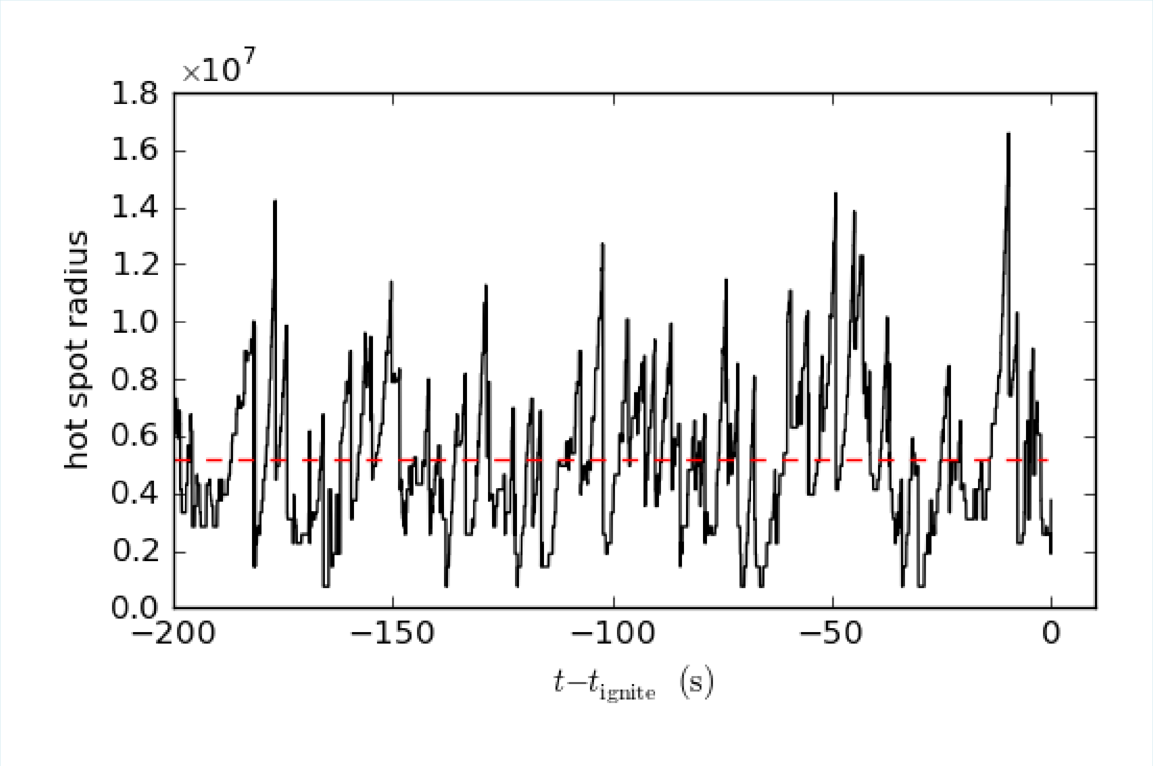

WD Convection: Ignition

- Histograms of ignition conditions over the final 200 seconds
  - (Left) Temperature and location of peak hot spot
  - (Right) Radial velocity and location of peak hot spot

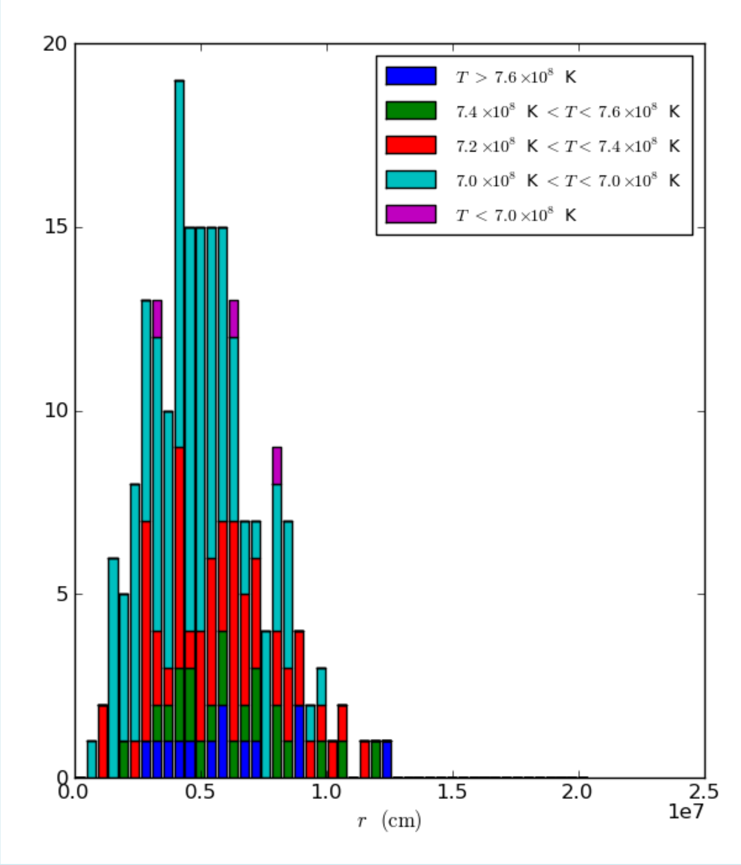

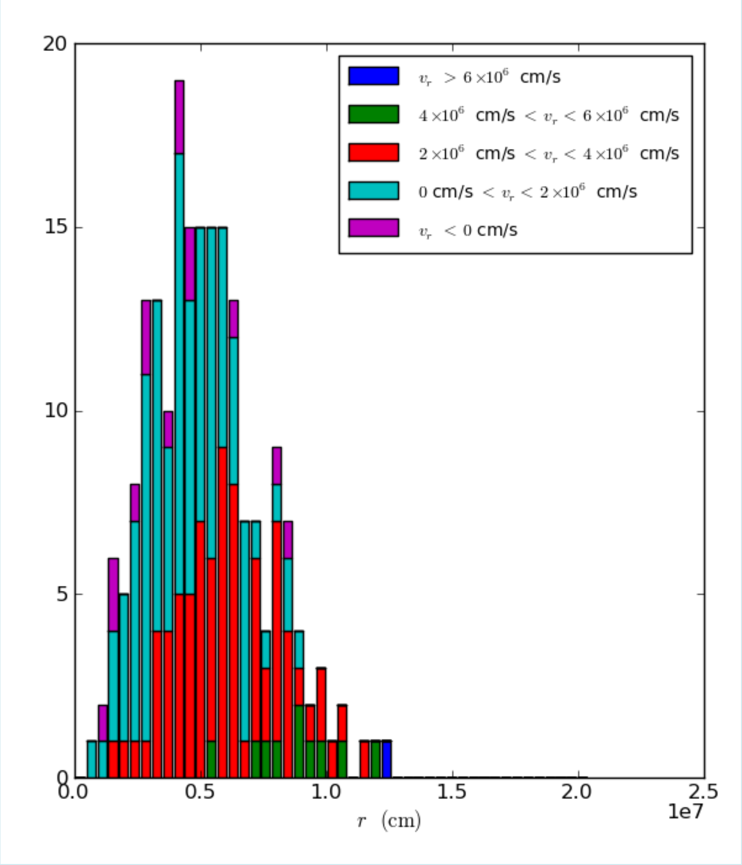

Summary / Future Work

- We have performed the first-ever full star simulations of convection preceding ignition in Type Ia supernovae
  - Low Mach number formulation
  - Adaptive mesh refinement
  - Performing science at 10K-20K cores, scaling to 100K cores
- What's next?
  - Examine the distribution of hot spots from our newest high-resolution studies.
  - Tracer particles to further understand development of hot spots
  - Examine role of turbulence and its effects on the first flames
  - Perform simulations in our compressible framework, CASTRO, using MAESTRO data as initial conditions