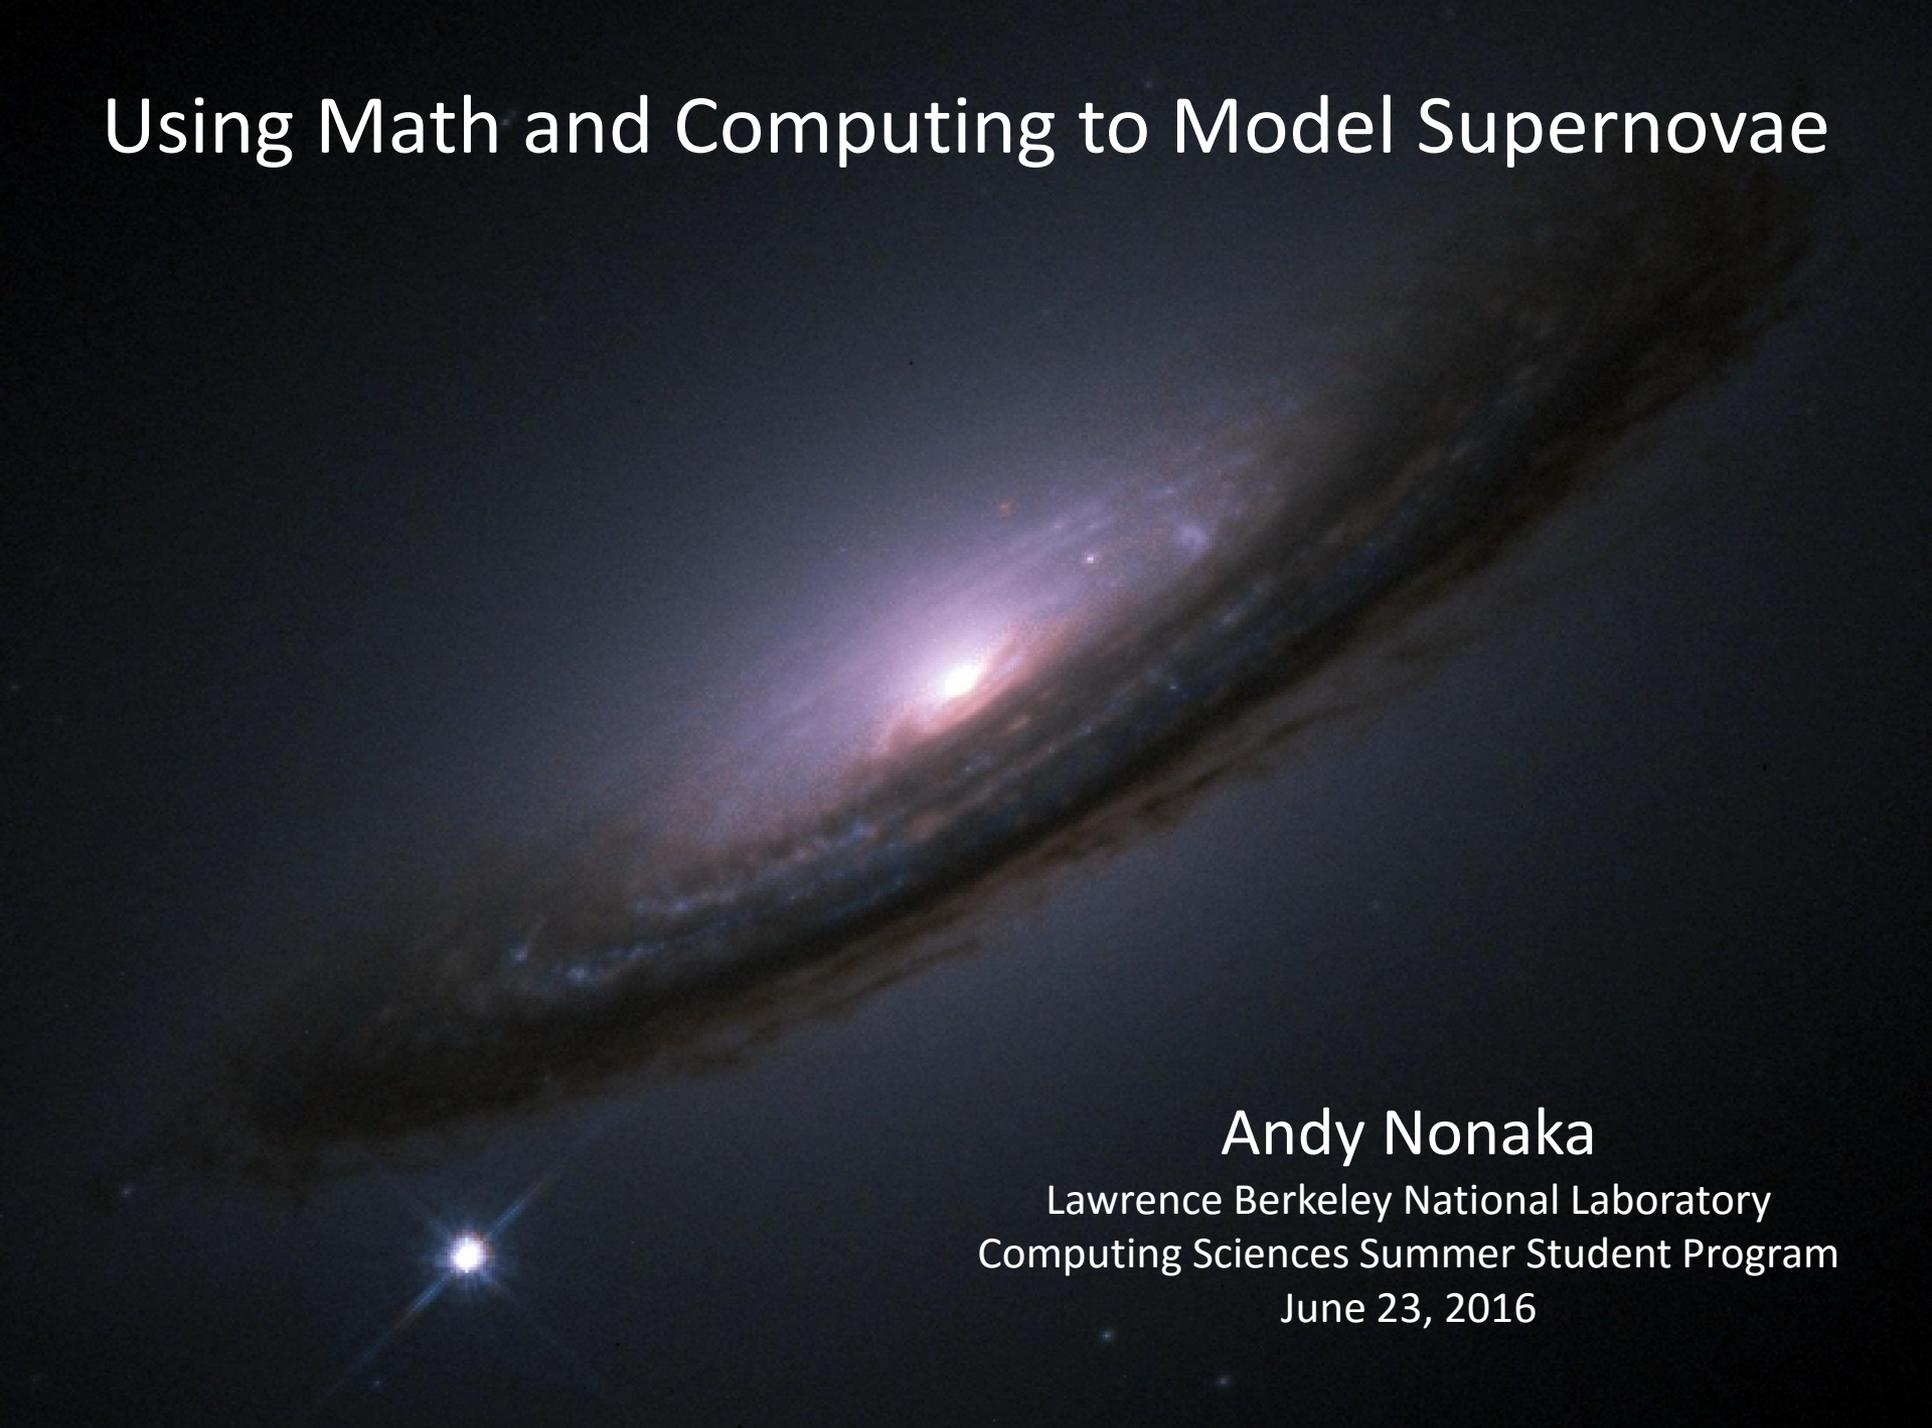


Using Math and Computing to Model Supernovae



Andy Nonaka

Lawrence Berkeley National Laboratory
Computing Sciences Summer Student Program
June 23, 2016

Galaxy NGC 4526 imaged by the Hubble Space Telescope (www.nasa.gov)

60 million light years away



SN1994D (Type Ia supernova)

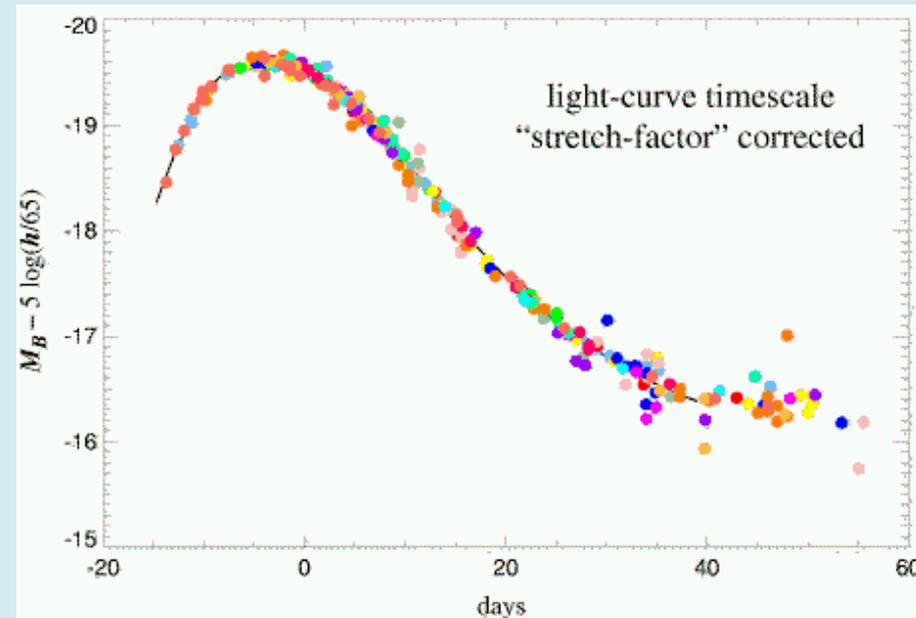
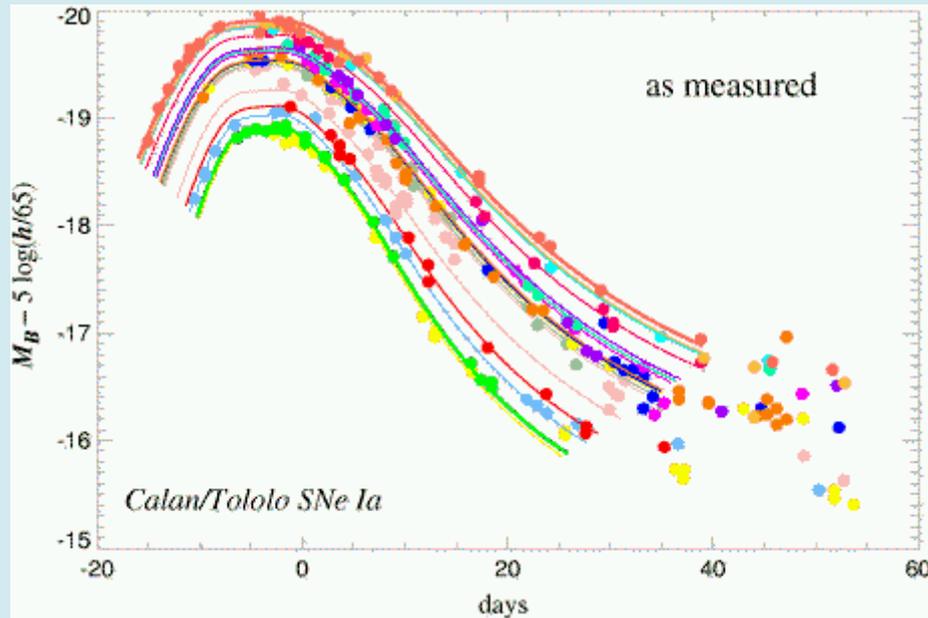
The supernova is as bright as the host galaxy!



- Why should we care?
- Using modern telescopes, Type Ia supernova light curves can now be observed several hundred times per year.
 - Spectra indicate that oxygen and calcium are present early, where as nickel, cobalt, and iron are present later.

Type Ia Supernovae are Distance Indicators

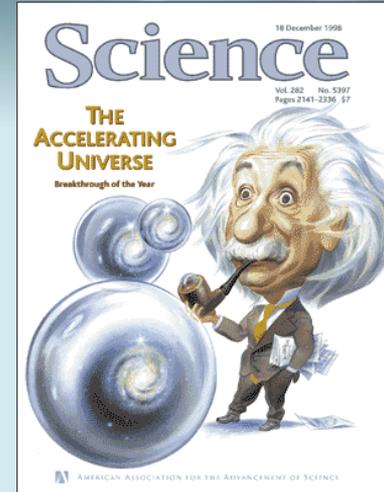
- By observing Type Ia supernovae at known, nearby distances, scientists have established a width-luminosity relationship; **wider = brighter**.



- 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

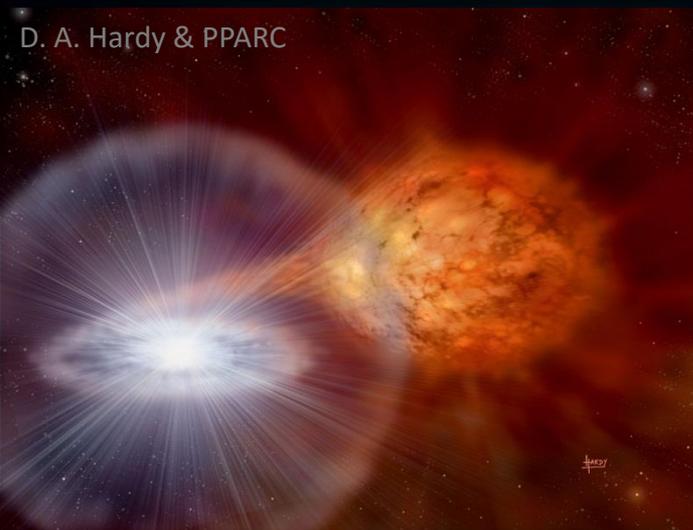
- Due to the observed redshift, we know the **speed** at which the host galaxy is moving away from us.
 - Led to discovery of the acceleration of the expansion of the universe in 1998
 - 2011 Physics Nobel Prize (Perlmutter, LBNL)
- Problem: We don't know how well the width-luminosity relationship holds 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...



Studying Type Ia Supernovae

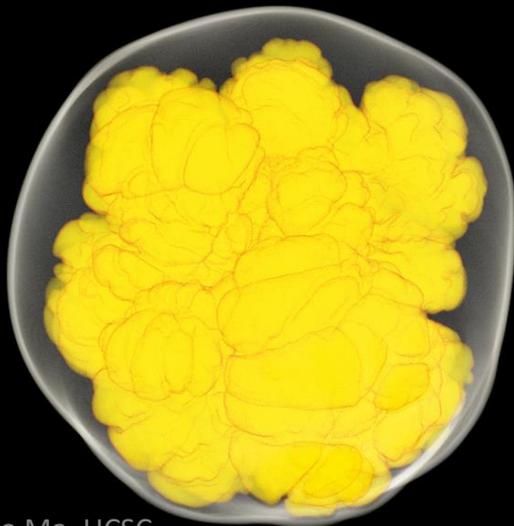
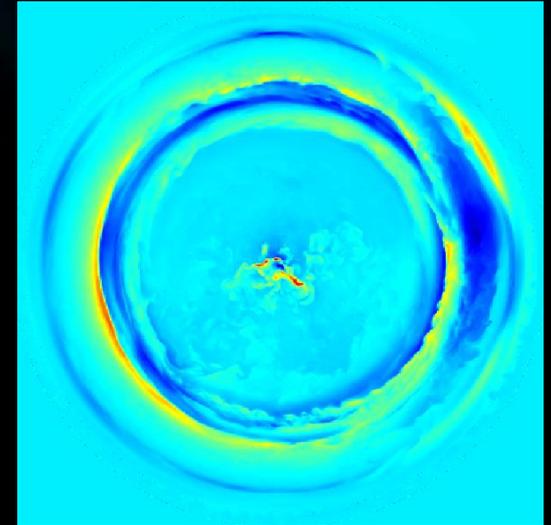
- We study this problem using math and computing
 - Develop mathematical models/equations describing stellar evolution and explosions
 - Develop numerical methods (algorithms) to solve these equations
 - Use supercomputers (10,000 – 100,000 CPUs) such as edison at NERSC.
- Requires expertise in applied math and computer science.
- Requires expertise in astrophysics (collaborate with experts in the field).

The Phases of Type Ia Supernovae: Single Degenerate Model



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**.



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

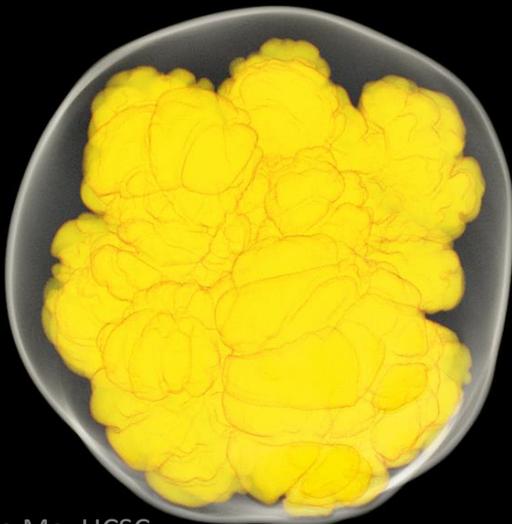
The resulting event is visible from Earth for **weeks to months**.



SN 1994D (High-Z SN Search team)

Computing the Explosion Phase

- Over the past decade, many have performed studies of the explosion phase using supercomputers.
 - Governed by well-understood (both theoretically and algorithmically) **fluid dynamics** equations.
 - A supercomputer can model this system in a few days or weeks, depending on spatial resolution.



Our CASTRO code is one of many publicly available codes capable of modeling such explosions.

Governing Equations

- Equations describing a compressible, reacting fluid/gas:

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

$$\frac{\partial(\rho \mathbf{u})}{\partial t} = -\nabla \cdot (\rho \mathbf{u} \mathbf{u}) - \nabla p + \rho \mathbf{g} \quad \text{conservation of momentum}$$

$$\frac{\partial(\rho E)}{\partial t} = -\nabla \cdot (\rho \mathbf{u} E + p \mathbf{u}) + \rho H + \rho \mathbf{u} \cdot \mathbf{g} \quad \text{conservation of energy}$$

ρ density

E total energy per unit mass

\mathbf{u} velocity

\mathbf{g} gravity

X_k mass fraction of species “k”

H energy release due to reactions

$\dot{\omega}_k$ reaction rate of species “k”

p pressure

Basic Solution Methodology

- Equations describing a compressible, reacting fluid/gas:

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

$$\frac{\partial(\rho \mathbf{u})}{\partial t} = -\nabla \cdot (\rho \mathbf{u} \mathbf{u}) - \nabla p + \rho \mathbf{g} \quad \text{conservation of momentum}$$

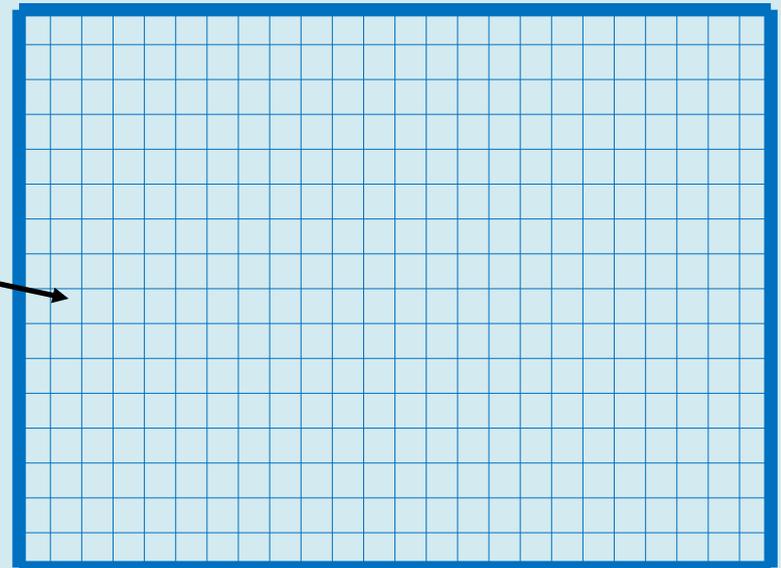
$$\frac{\partial(\rho E)}{\partial t} = -\nabla \cdot (\rho \mathbf{u} E + p \mathbf{u}) + \rho H + \rho \mathbf{u} \cdot \mathbf{g} \quad \text{conservation of energy}$$

- Finite volume approach.

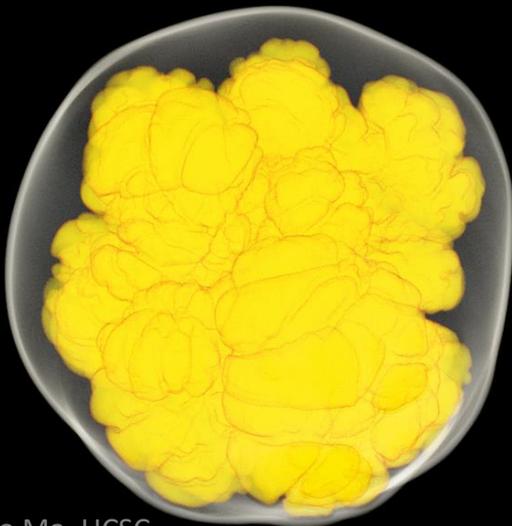
- Divide problem into grid cells

$\rho, \mathbf{U}, E, \text{etc.}$

- Advance solution incrementally over many time steps, Δt , until final time achieved



Computing the Explosion Phase



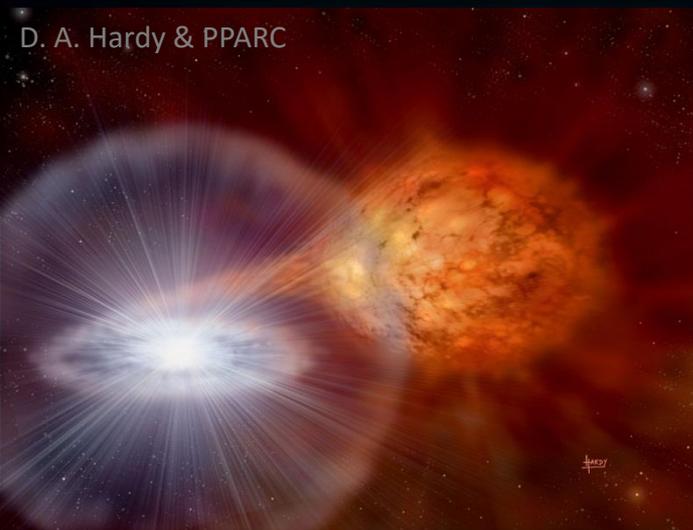
A major problem are the initial conditions, which have been based on “guesses”.

What is the initial state of the star?

Where are the first flames?

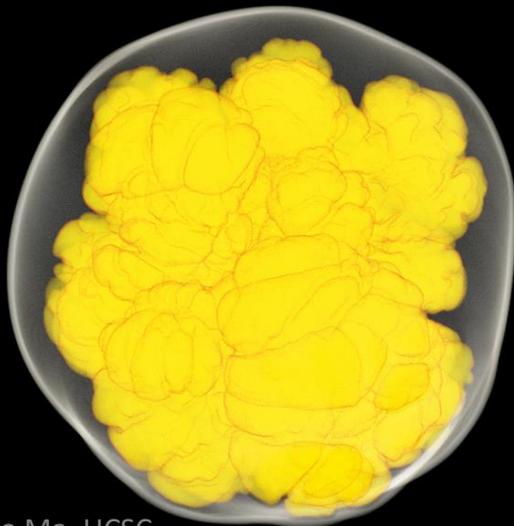
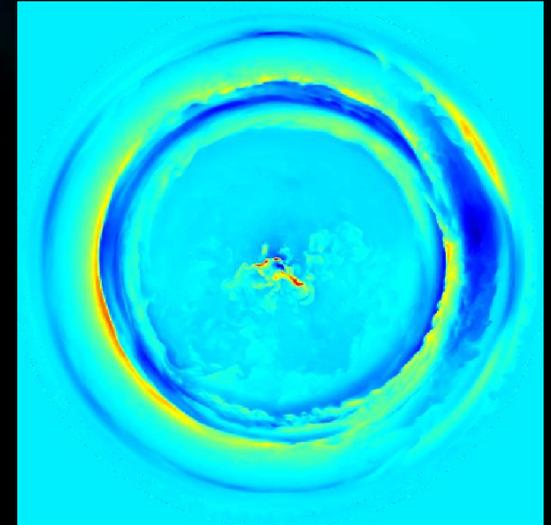
How many ignition points are there?

The Phases of Type Ia Supernovae: Single Degenerate Model



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**.



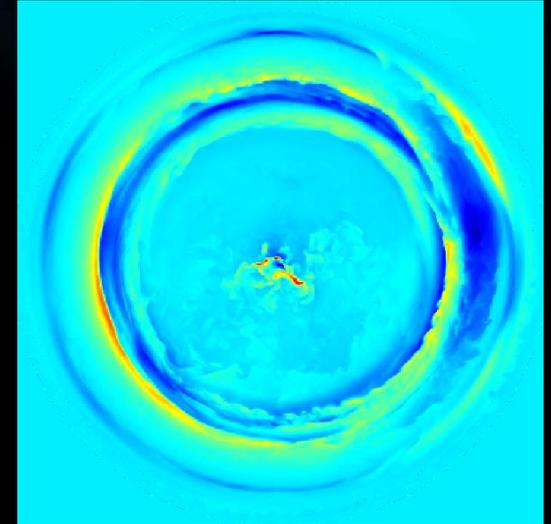
Flame (possibly) transitions to a detonation, causing the star to explode within **two seconds**.

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



SN 1994D (High-Z SN Search team)

Computing the Convective Phase



- We would like to simulate the last few hours of smoldering preceding the explosion to obtain initial conditions for CASTRO.
- Problem: It takes weeks on a supercomputer to simulate 2 seconds of real-time. How do we simulate hours?

Governing Equations

- Compressible, reacting fluid equations:

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

conservation of mass

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

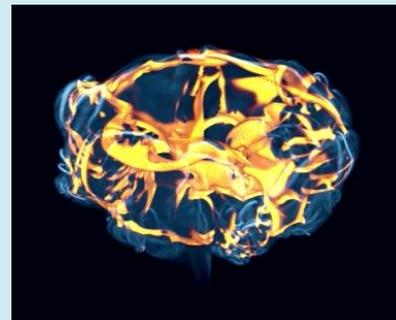
conservation of momentum

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

conservation of energy

- These equations describe 3 things:

- Motion of the fluid
- Nuclear reactions (burning)
- Sound waves



Smoldering Phase vs. Explosive Phase

- How is the smoldering phase different from the explosion phase?
 - “Low Mach Number” flow - fluid speed small compared to sound speed ($\sim 1\%$)
 - Sound waves carry little energy and have minimal impact on the overall solution
 - “Ignoring” them doesn’t significantly affect the solution.
- We have derived a new equation set that ignores the effect of sound waves, yet retains all the remaining physics, and is much more computationally efficient.

Low Mach Number Equation Set

- Derive new equations/model 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:

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

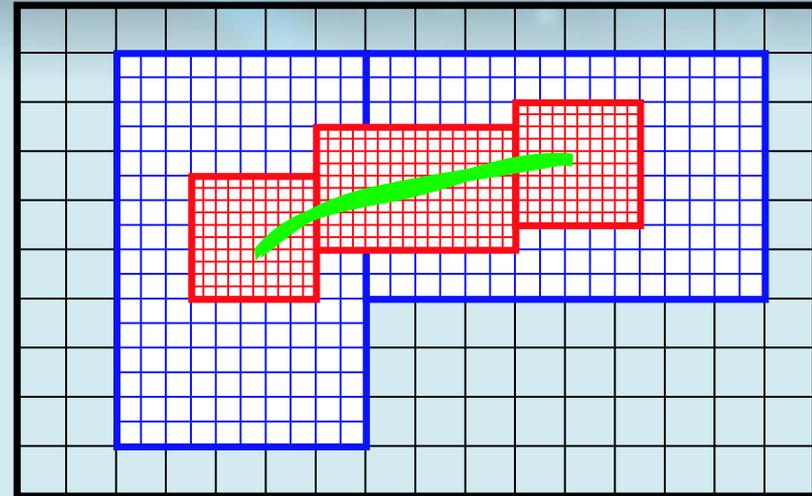
- Low Mach time step is a factor of $1/M$ larger than a compressible time step, enabling longer simulations!

Computational Efficiency

- In our white dwarf simulations, the peak Mach number varies from 0.01 – 0.05.
 - Net result: the low Mach number time step is a factor of 70 greater than a compressible time step
 - However, the low Mach number equation set is more complex and takes approximately 2.5 times longer advance a single time step.
 - Thus, to advance the solution to the final time, **MAESTRO** is a factor of $(70 / 2.5) \approx 28$ more efficient than a compressible algorithm, given the same number of computational resources for this problem.
 - Now we can simulate roughly 1 minute of the smoldering phase, but we are still looking to simulate several hours.

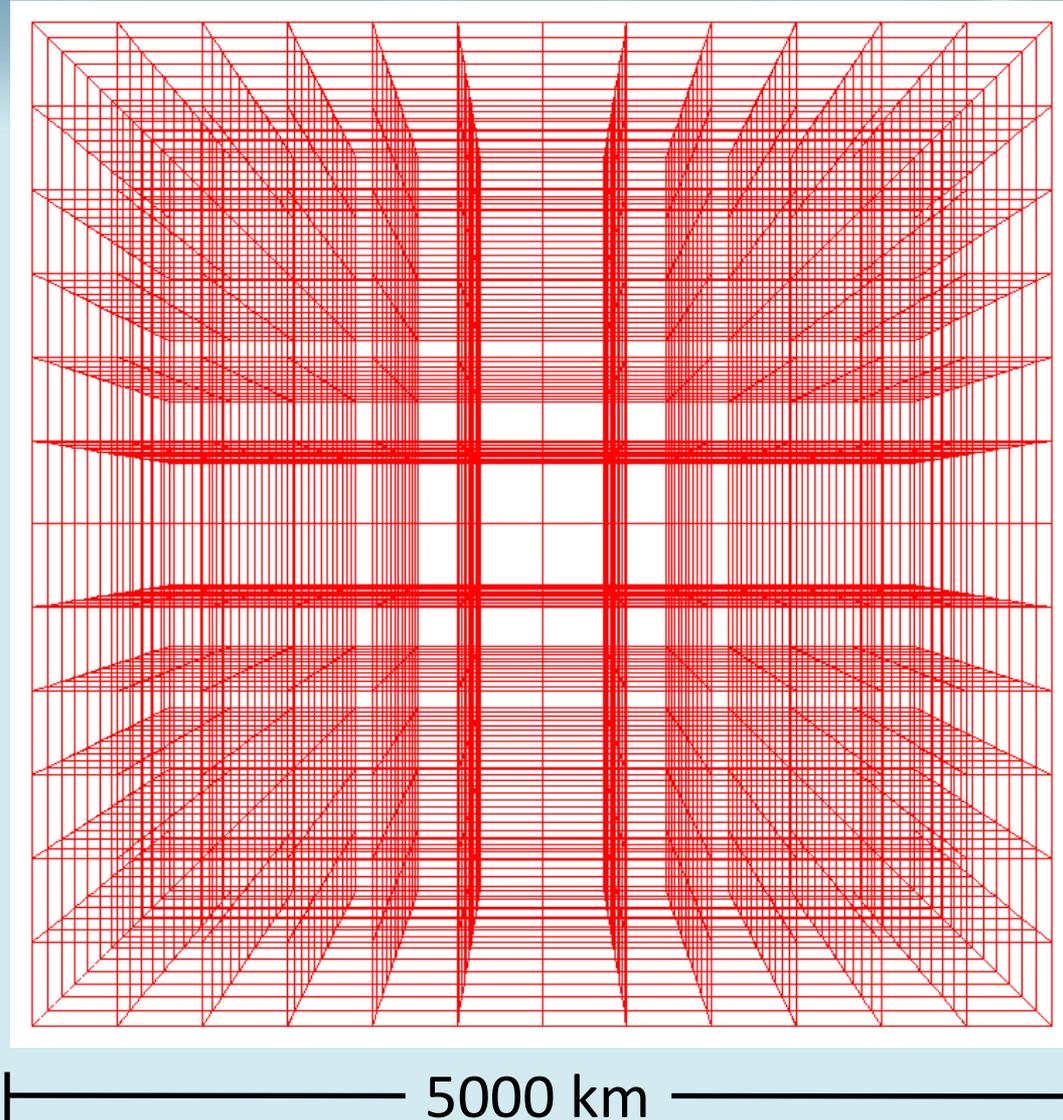
Adaptive Mesh Refinement

- Incorporate AMR using established techniques
 - Advance each level independently and synchronize solution between levels to maintain conservation
- For the full star problem, we need to consider our refinement 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



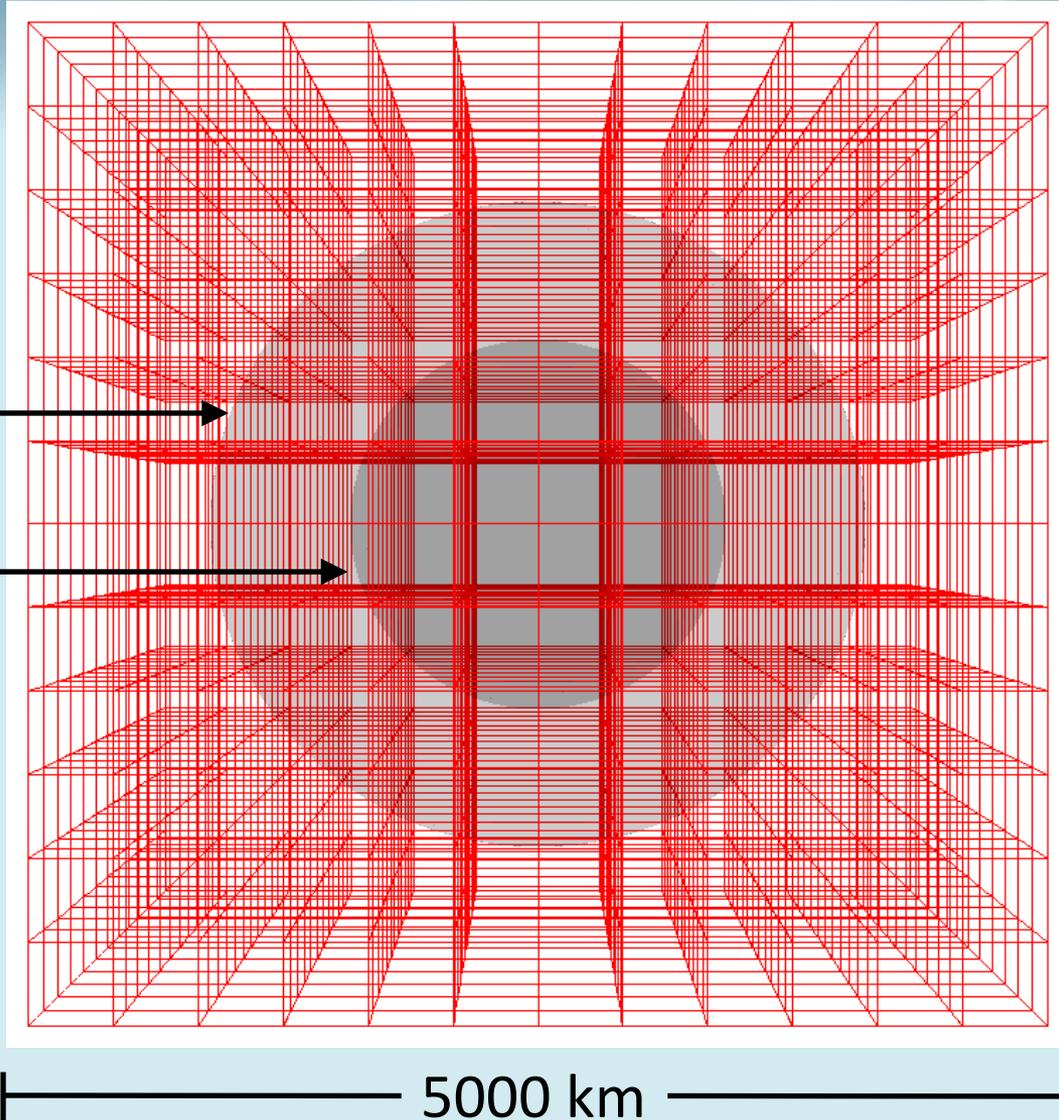
Adaptive Mesh Refinement

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

Edge of Star

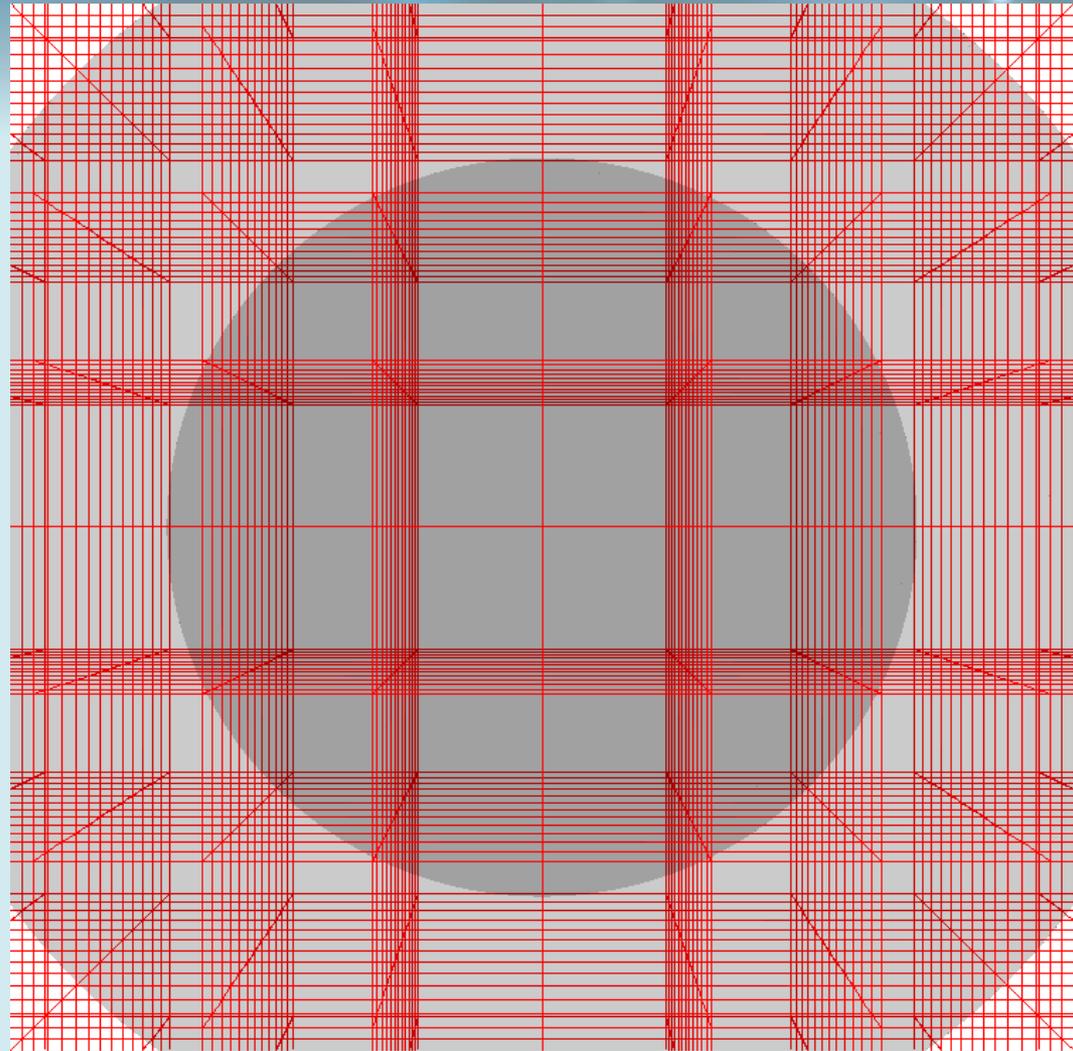


Convective Zone
Boundary



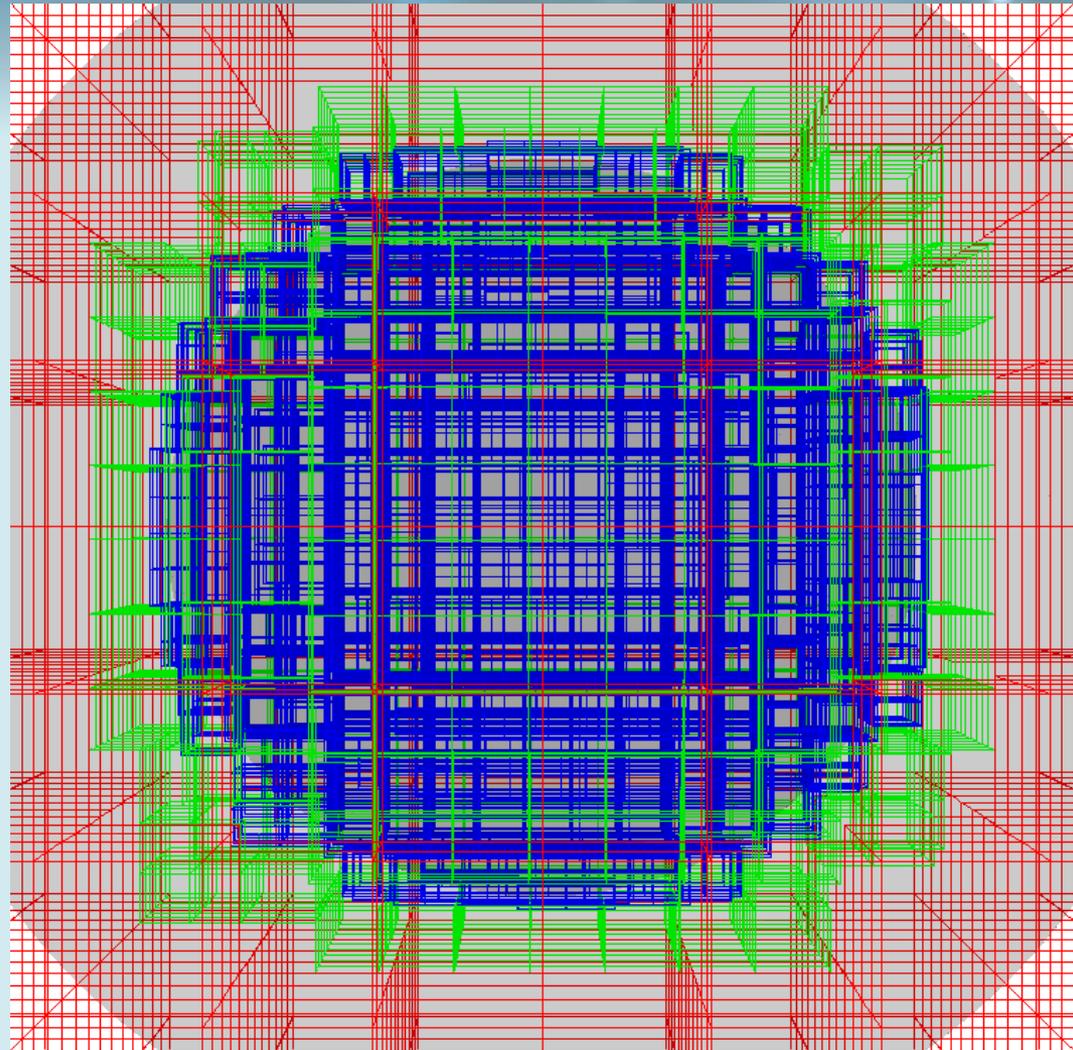
Adaptive Mesh Refinement

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



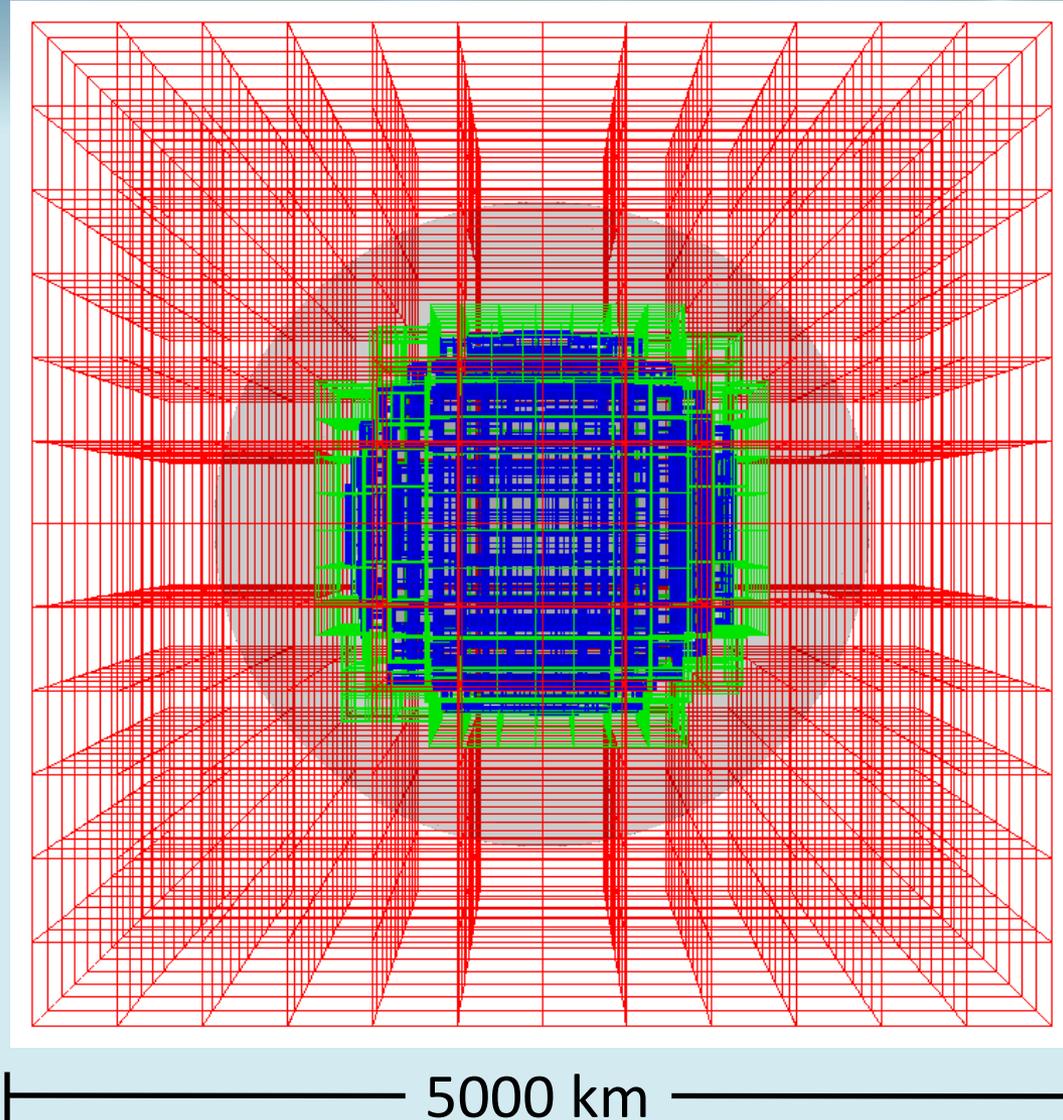
Adaptive Mesh Refinement

- 576^3 (8.7 km)
 - $1728 \cdot 48^3$ 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



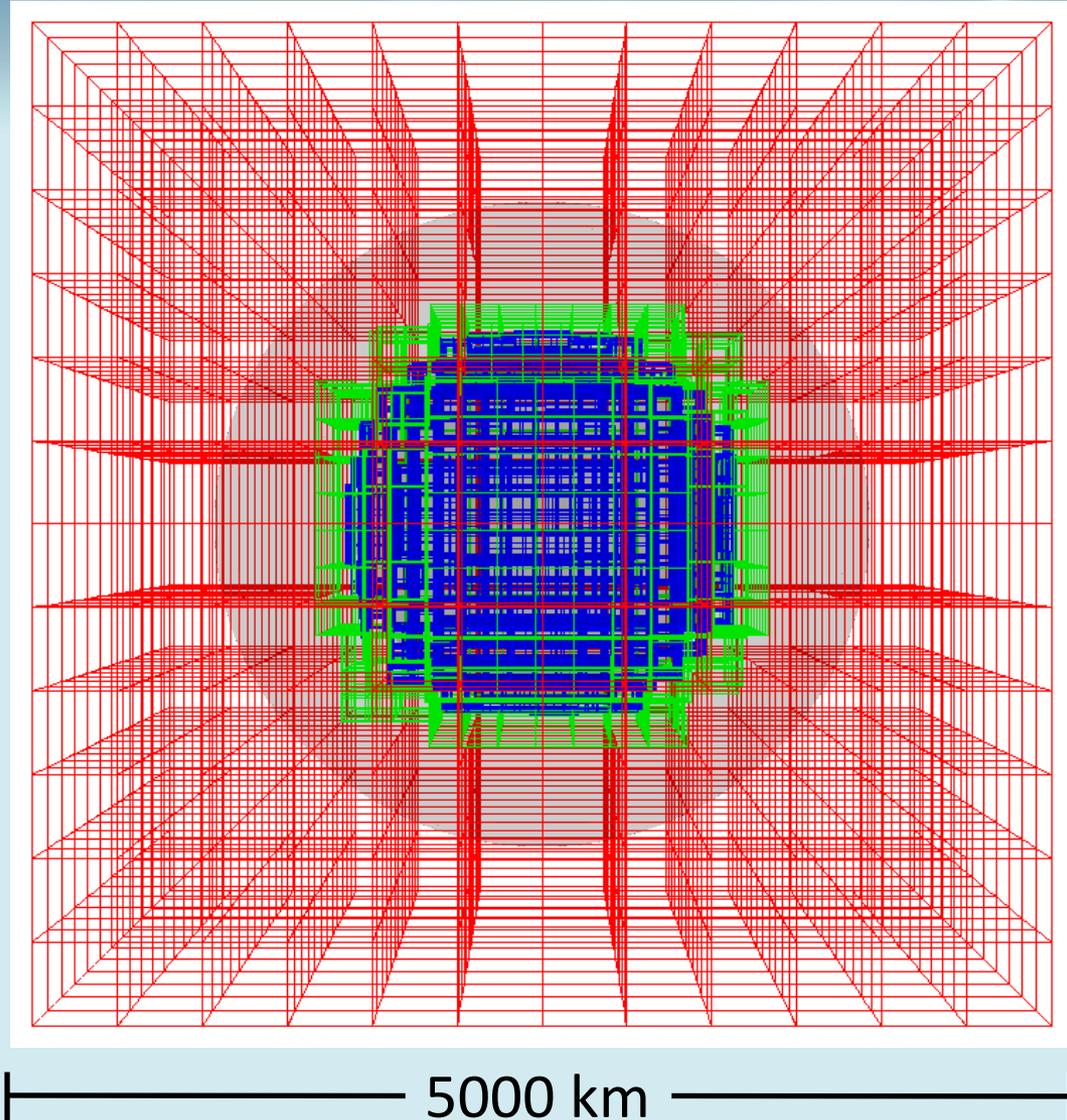
Adaptive Mesh Refinement

- 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**.



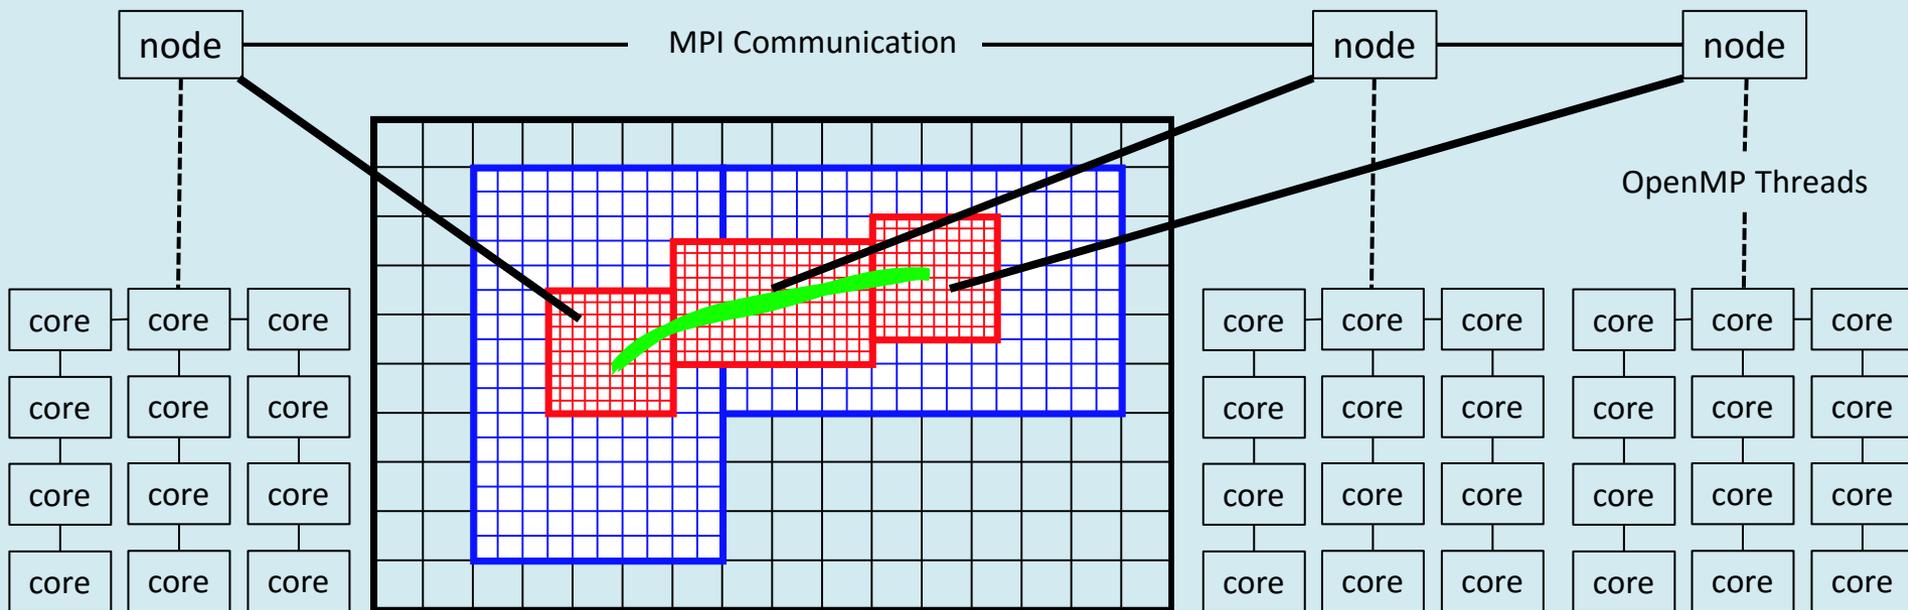
Adaptive Mesh Refinement

- In practice, we run most of the simulation using the coarsest resolution only and add AMR in the last few minutes as the star approaches ignition.
 - Allows us another **factor of 20 speedup**



Parallelization Strategy

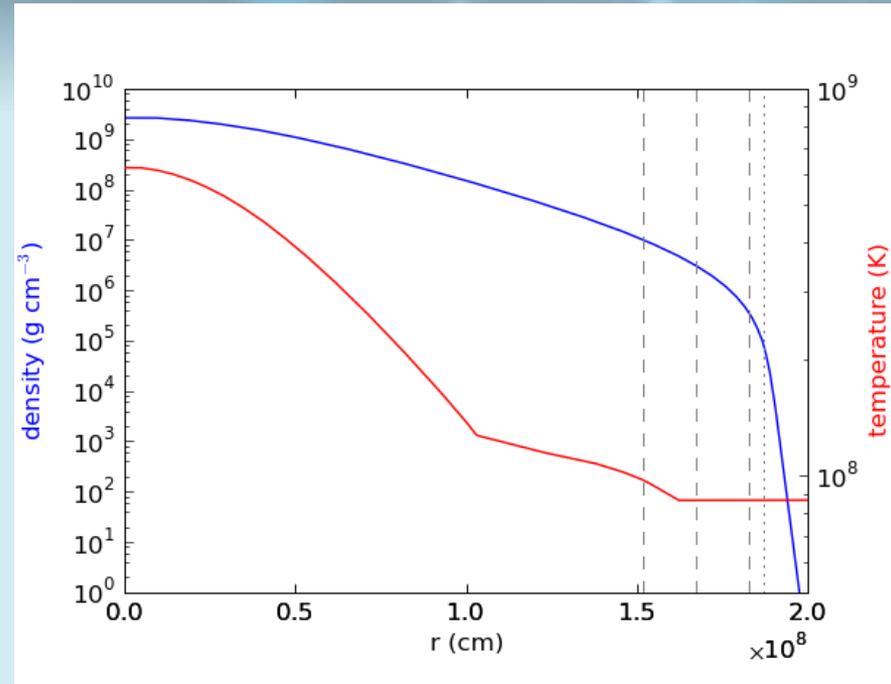
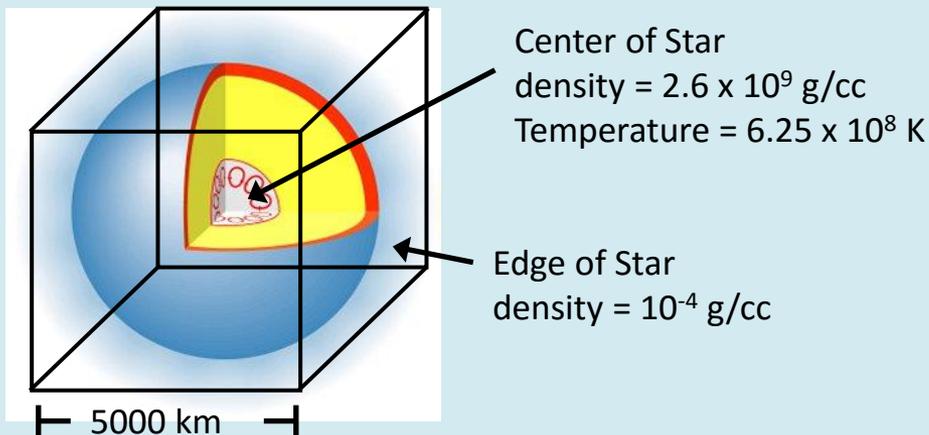
- Hybrid MPI/OpenMP approach to parallelization.
 - Nodes assigned to grids, threads spawned on cores to work on grids



- We are able to efficiently run our codes on 100,000+ processors using this approach.

White Dwarf Convection: Initial Conditions

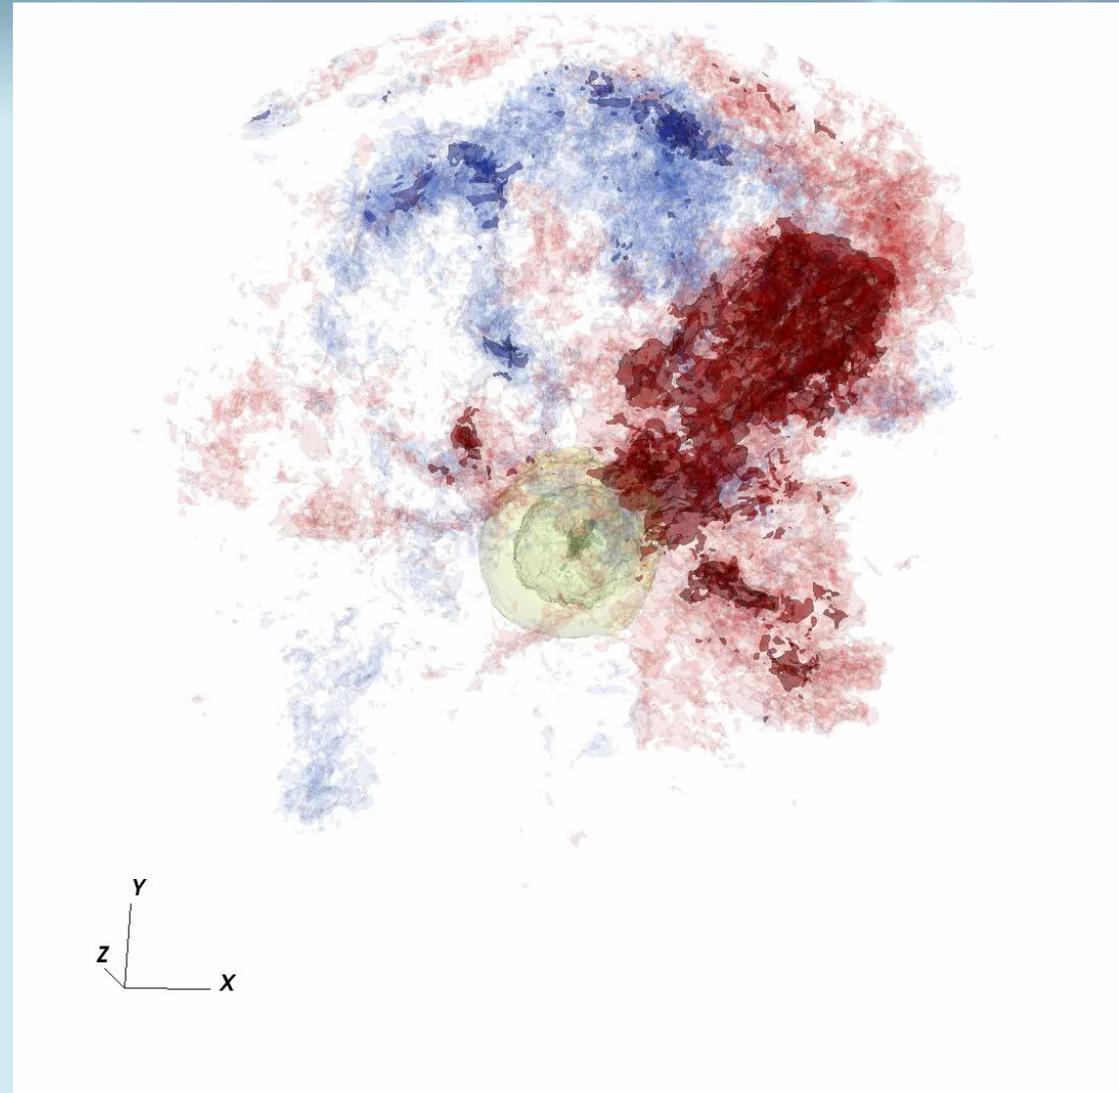
- Initial conditions
 - 1D model model mapped onto Cartesian grid
 - Random velocity perturbation added to prevent initial nuclear runaway



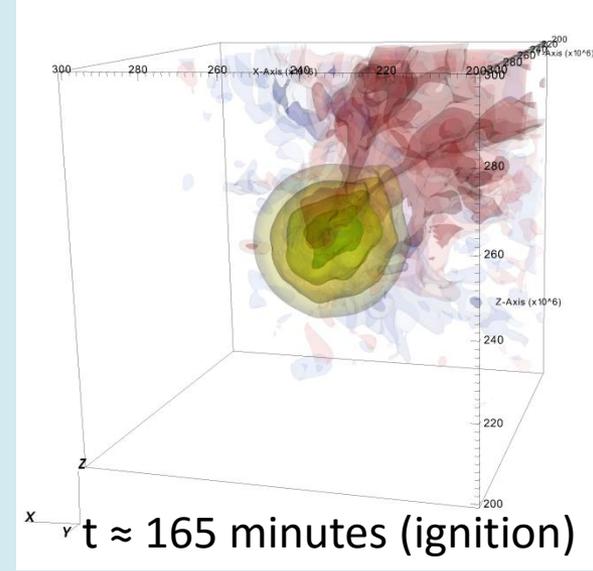
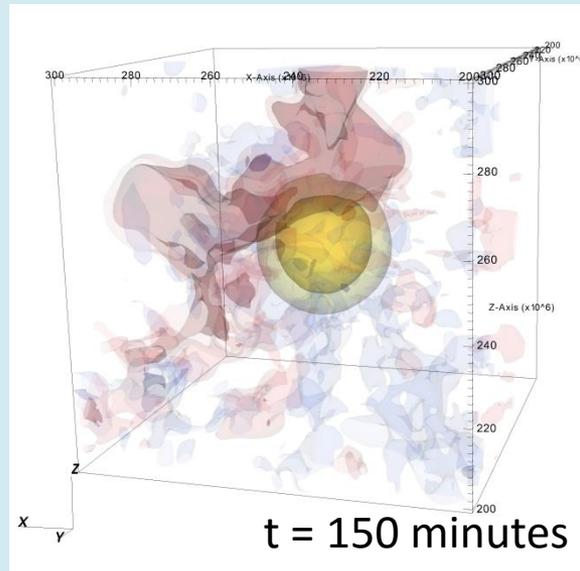
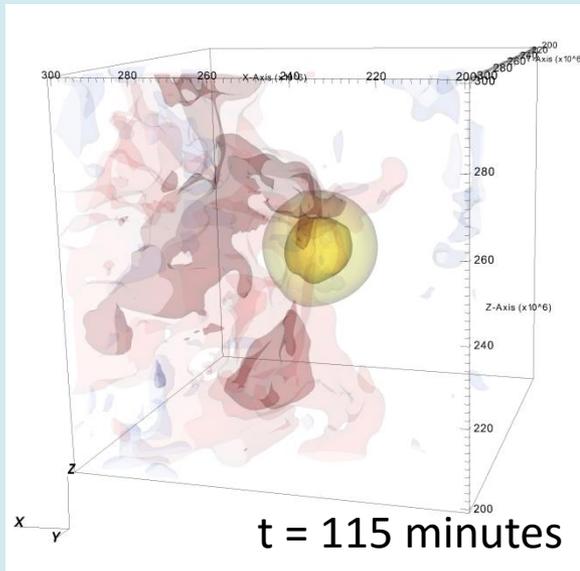
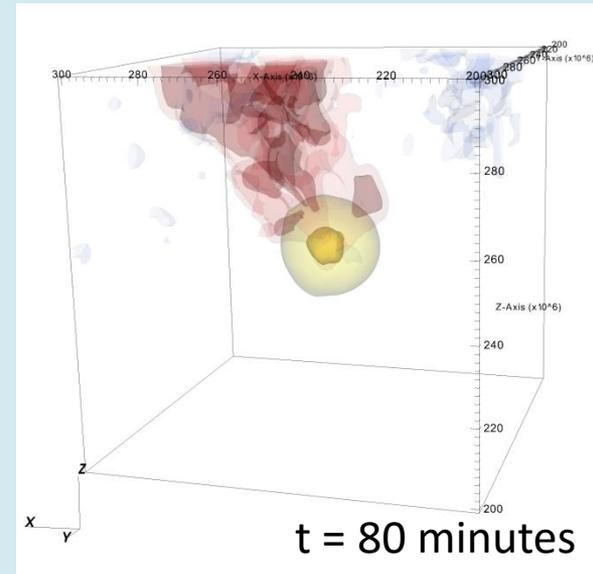
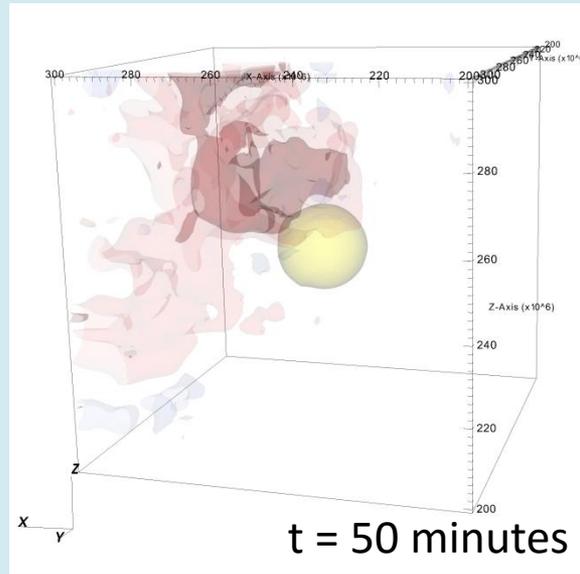
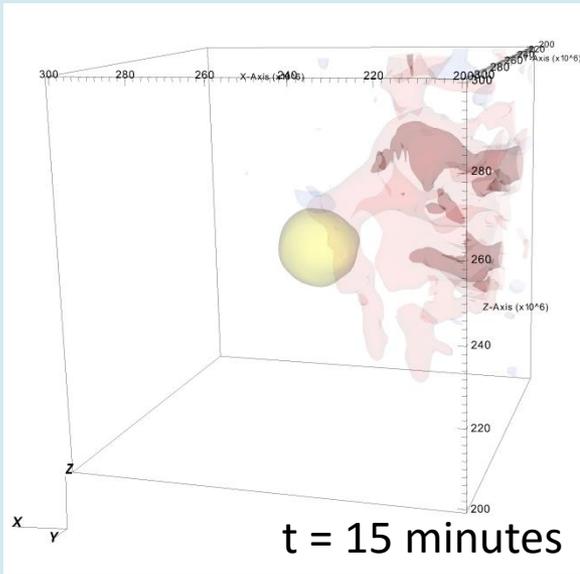
- Use 10K cores for 40 days (10 million CPU hours) to run effective 2304^3 resolution (2.2km zones) to ignition

White Dwarf 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

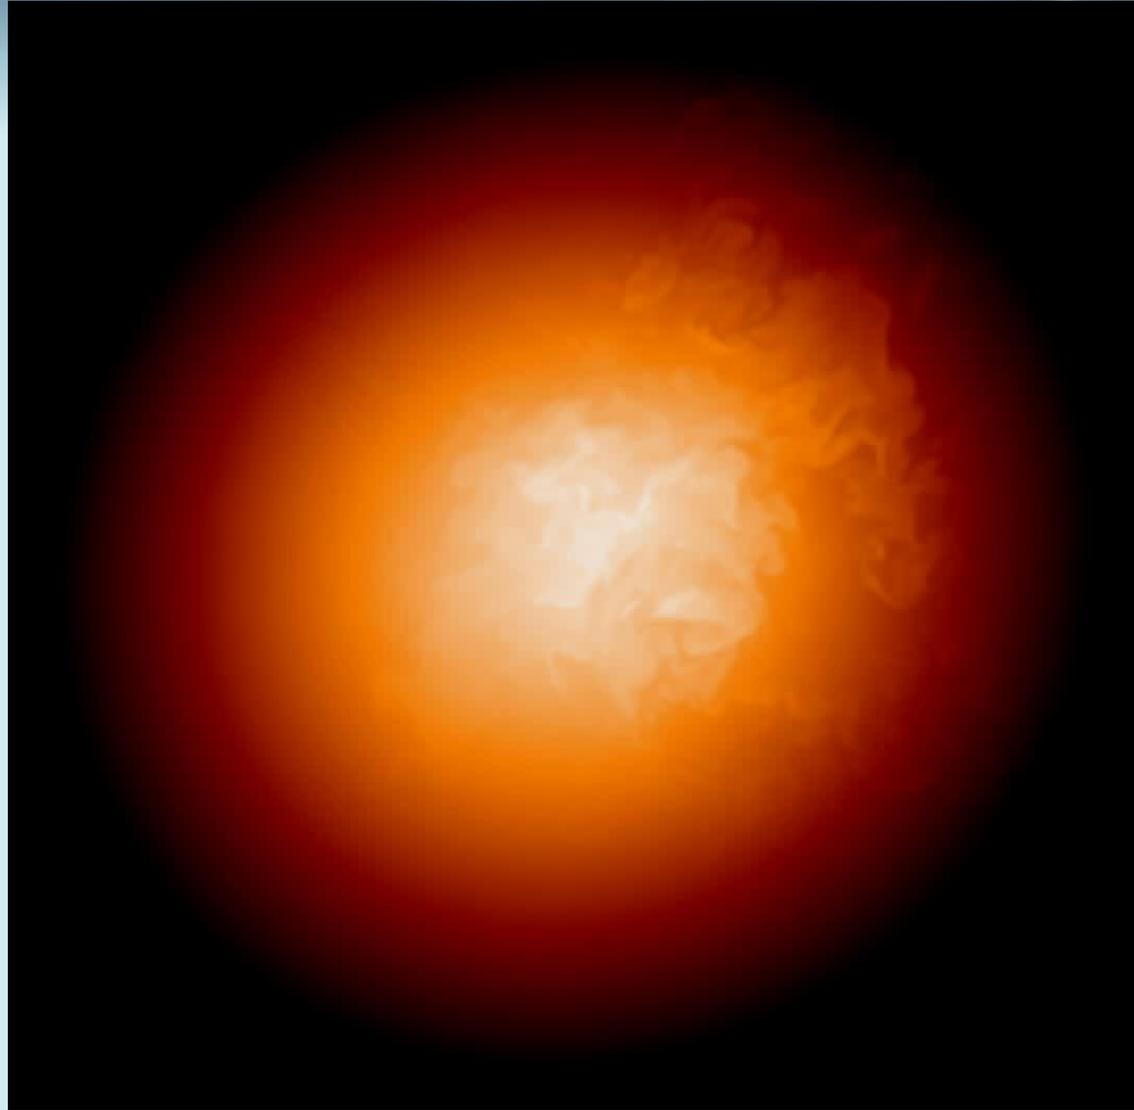


- Red / Blue = outward / inward radial velocity
- Yellow / Green = contours of increasing burning rate



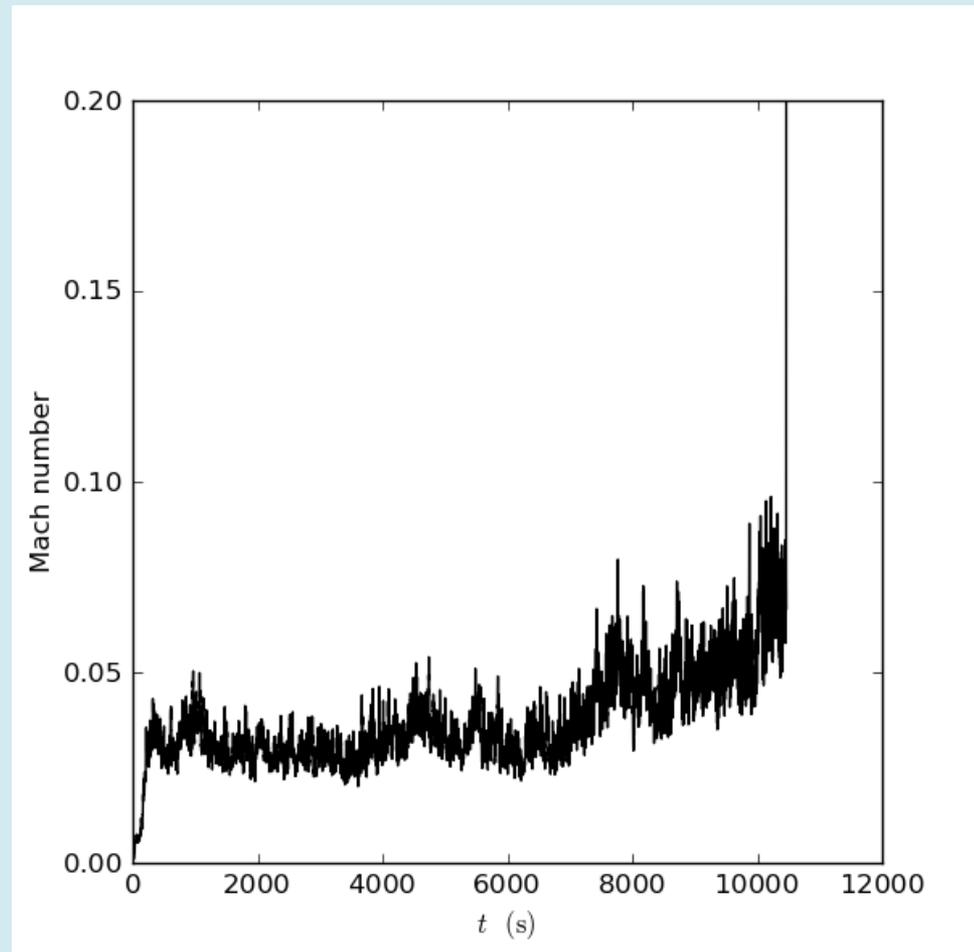
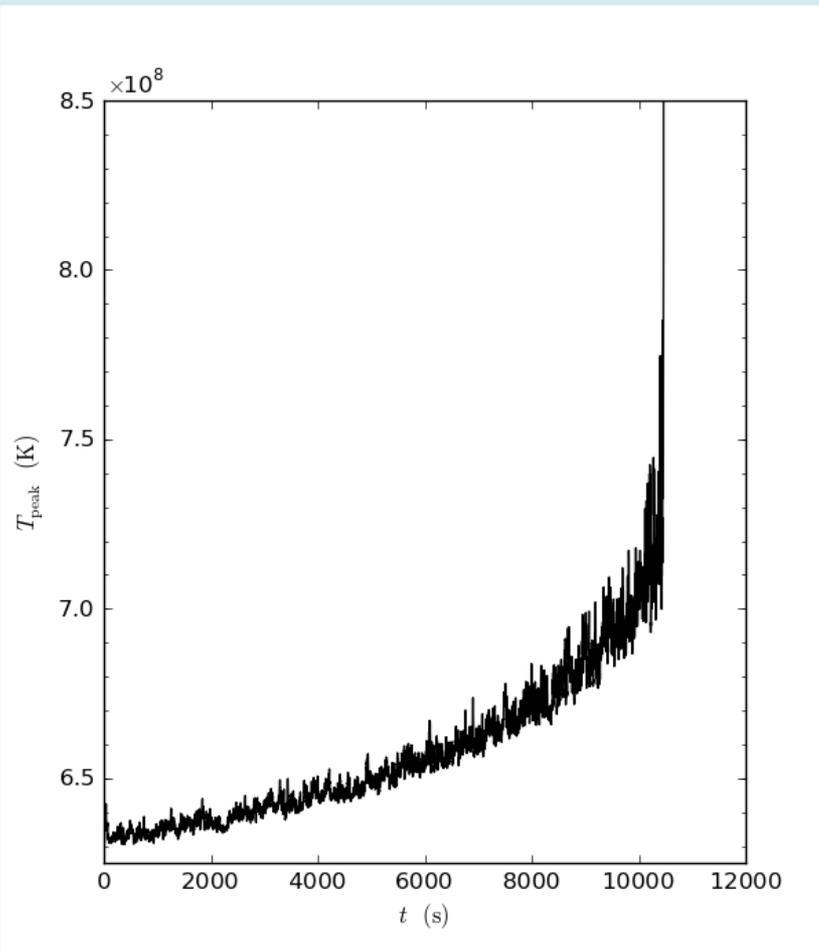
White Dwarf Convection: Ignition

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



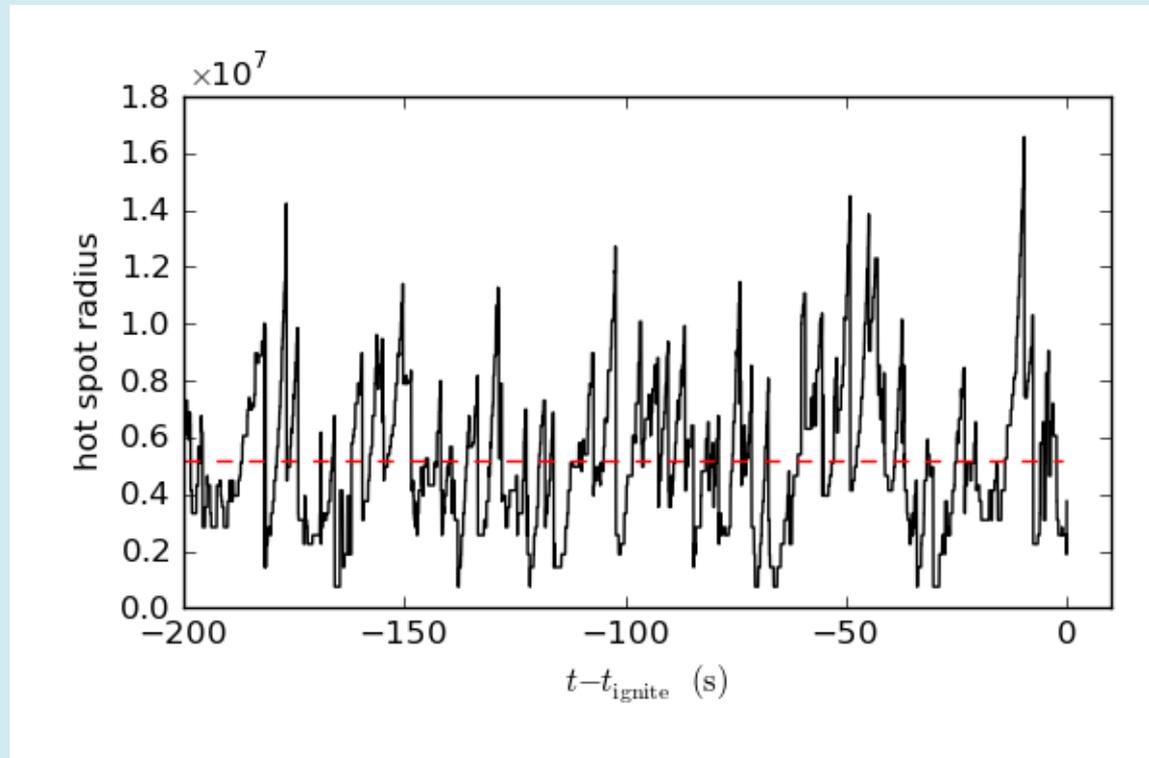
White Dwarf Convection: Long-Time Behavior

- Maximum temperature and Mach number vs. time



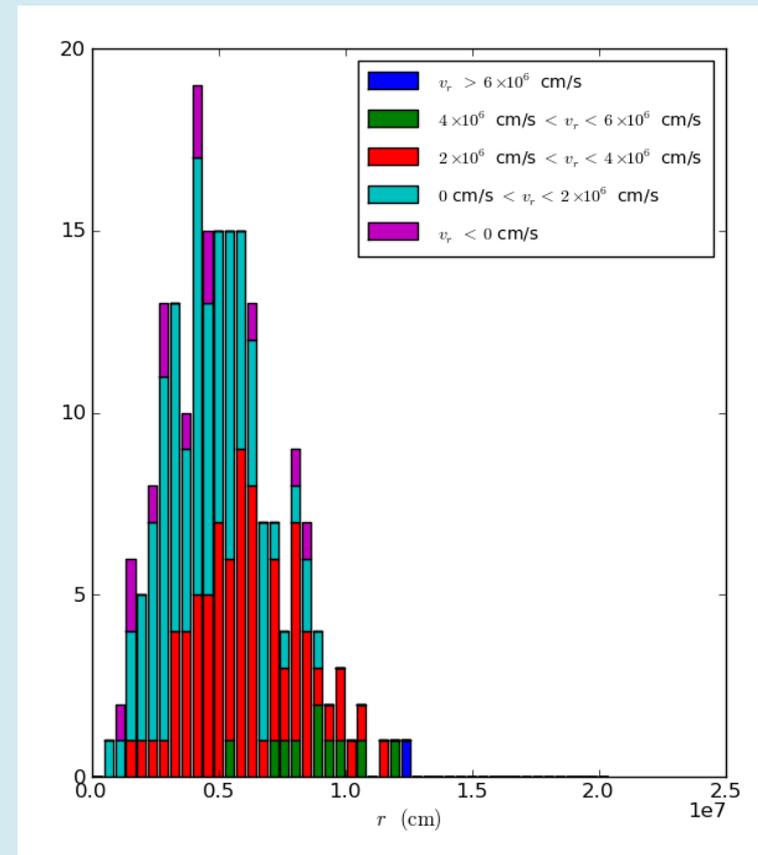
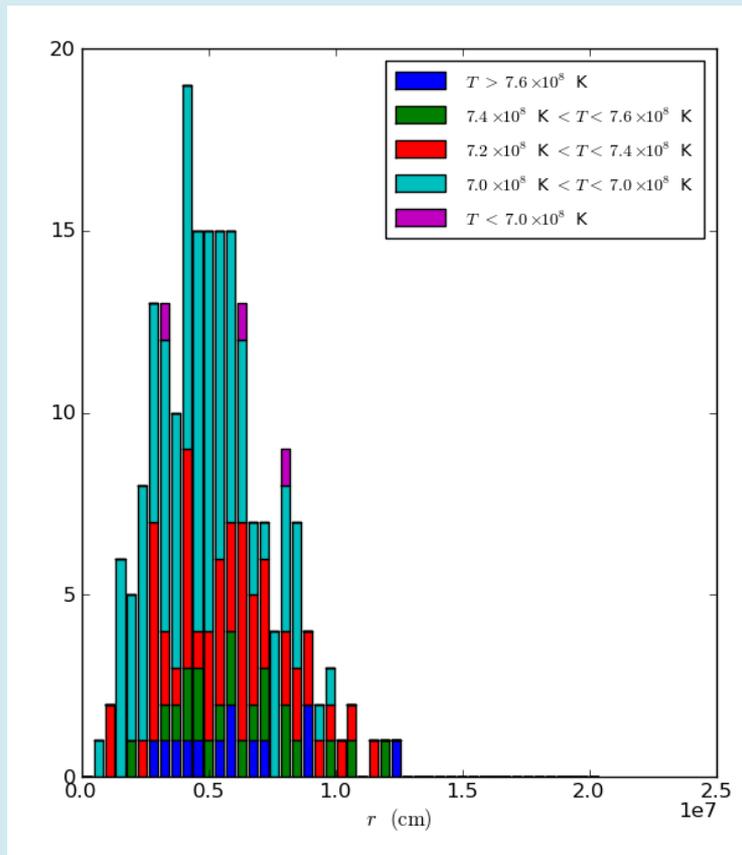
WD Convection: Ignition

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



White Dwarf 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



White Dwarf Convection Summary

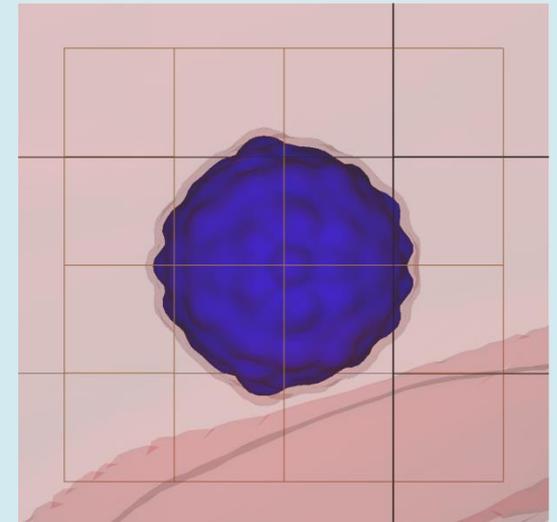
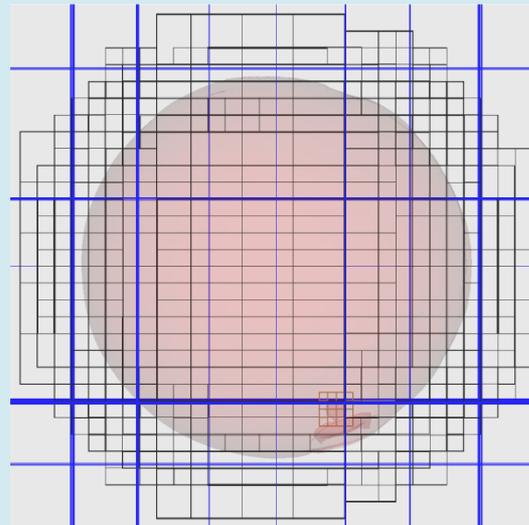
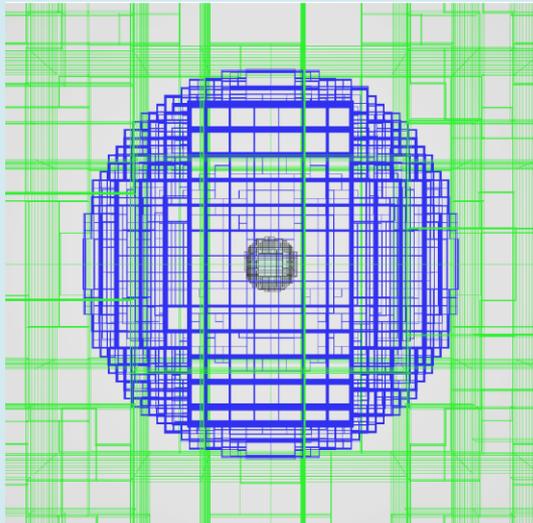
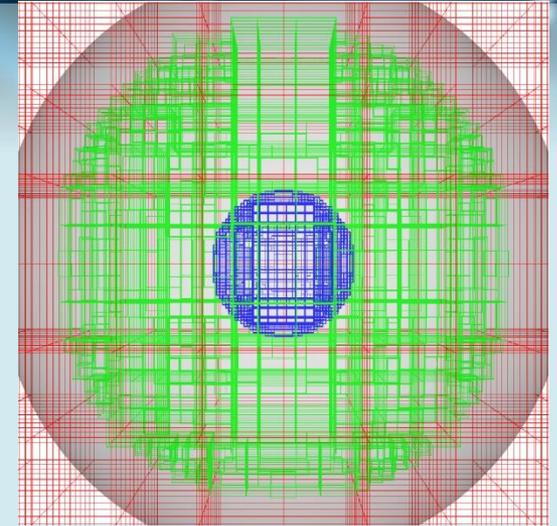
- We have performed the most detailed full-star calculations ever of convection up to the point of ignition in Type Ia supernovae
 - Low Mach number formulation
 - Adaptive mesh refinement
 - Factor of ~ 6000 speedup compared to traditional uniform resolution compressible approaches
 - Performing science at 10K-20K cores, scaling to 100K cores
- Main scientific conclusions:
 - Likely ignition radius of 50-70km
 - Single ignition point strongly favored
 - Characterization of full state of the star, including the background velocity field

Compressible Simulations with CASTRO

- Once the first flames have ignited, the fluid velocities become large compared to the sound speed, and the assumptions we used to derive the low Mach number equation set are no longer valid.
- We study post-ignition dynamics of early flames with the fully compressible code, CASTRO.
 - We can import the initial conditions directly from the MAESTRO simulation into the compressible code framework.

CASTRO Grid Configuration

- 5 levels of AMR
 - Divide 5000 [km] domain into 150 [m] zones
 - Effective $36,864^3$ zones would require 50 trillion grid points without AMR
 - With AMR we only use 1 billion zones

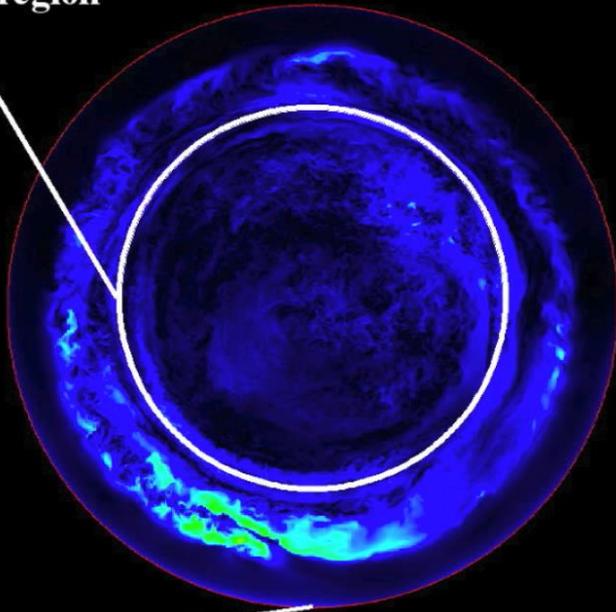


The Main Simulation

- We ran this simulation on 64,000 cores for 1 week.
 - Modeled first 0.5 seconds after ignition
- We included the background turbulent velocity field.
- We are interested in measuring properties of the spreading flame (size, rate of expansion) as well as the energy release and elemental production due to burning.

Slice through MAESTRO results of magnitude of velocity

Convective region



Stellar surface

Other Simulations

- We ran other simulations where we modified the ignition conditions and/or disabled the background velocity.

Simulation Name	Ignition Radius	Include Background Velocity?
AV	41km	Y
A0	41km	N
BV	10km	Y
B0	10km	N
CV	Center	Y
C0	Center	N

Comparison of simulations with different initial conditions

AV Simulation

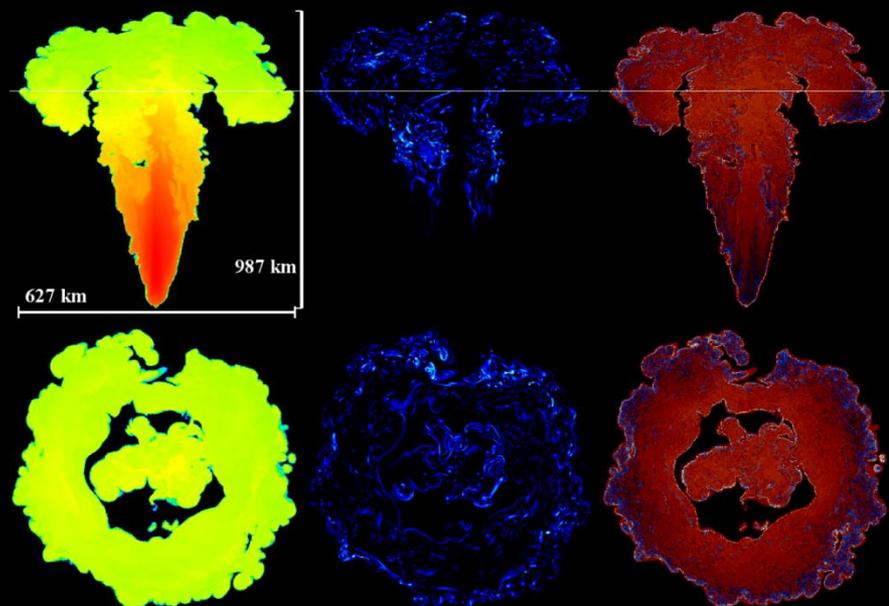
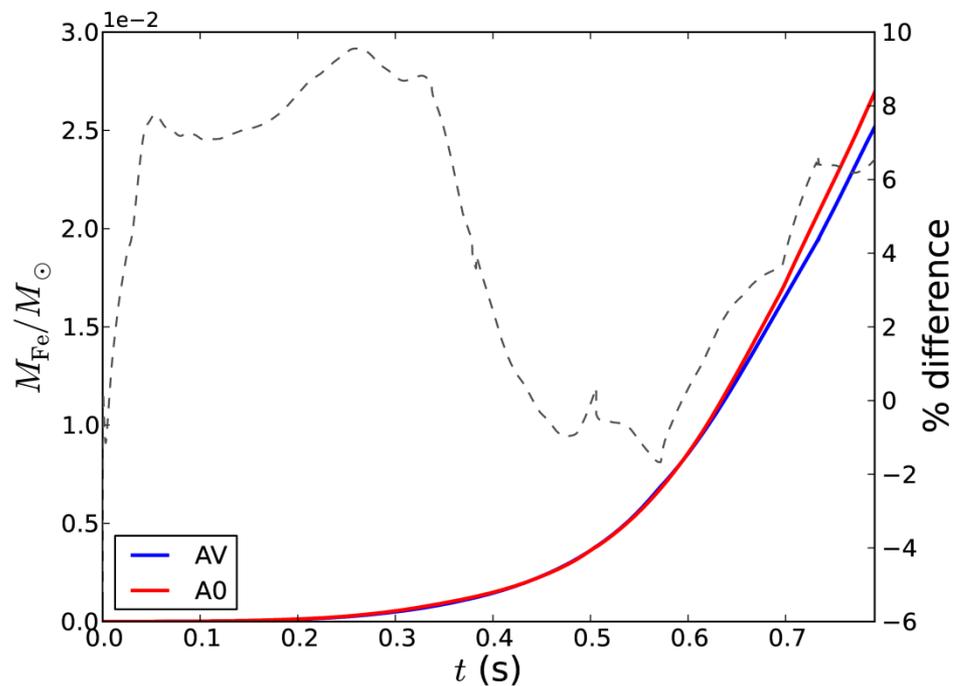
41km ignition point

Include background flow field

A0 Simulation

Same as above, but NO background flow field

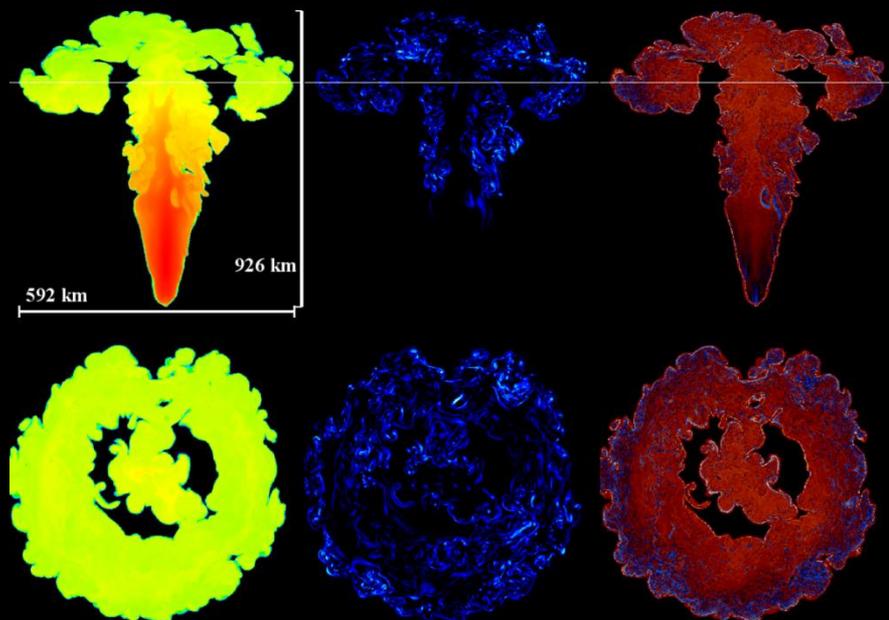
Iron Production



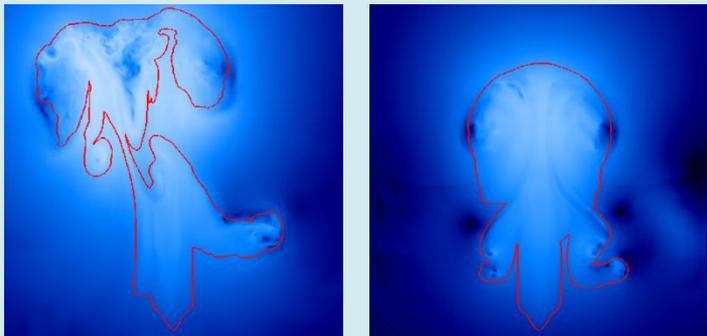
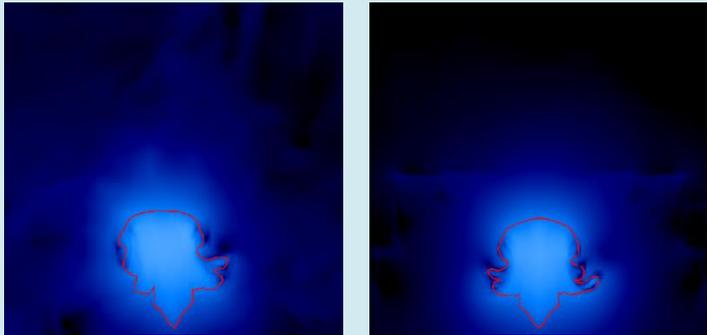
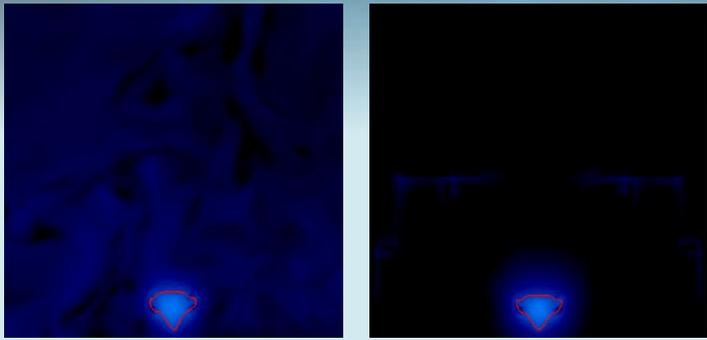
Temperature

Vorticity

Energy Release

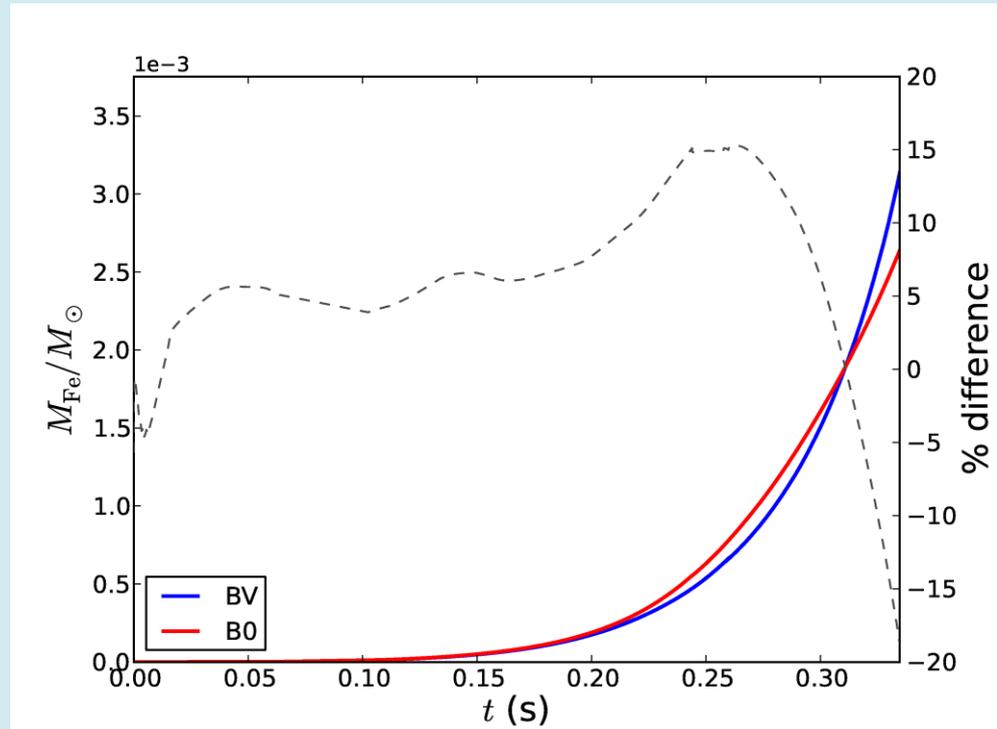


Effect of Velocity on Other Ignition Points

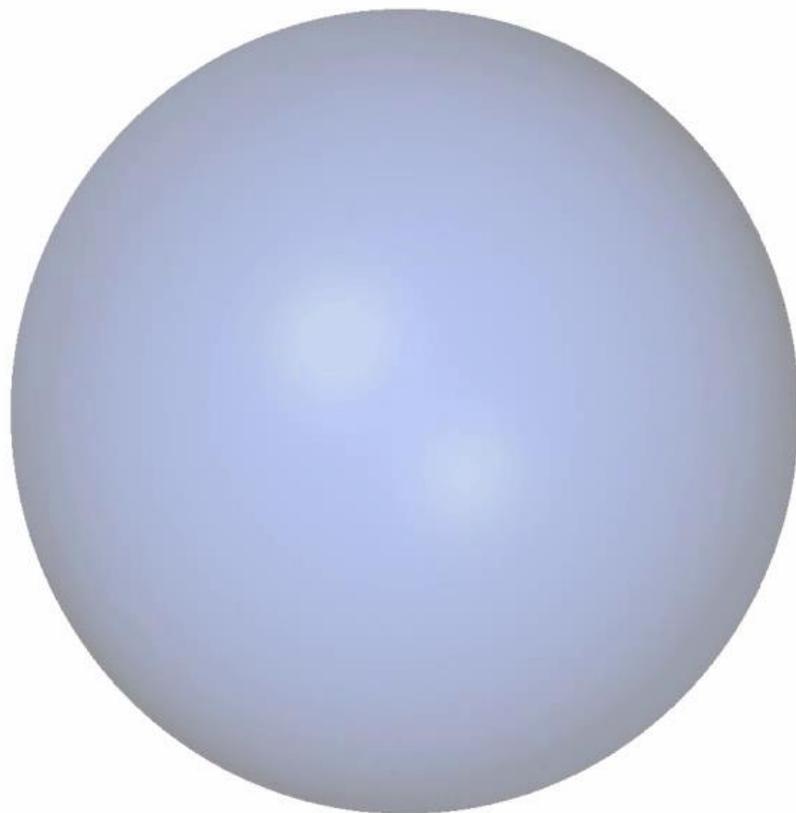


Comparing early flame evolution for artificial (10km) ignition with velocity field (left, “BV”) and without (right, “B0”).

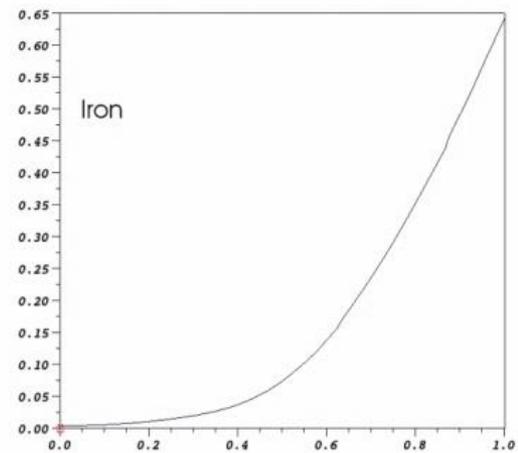
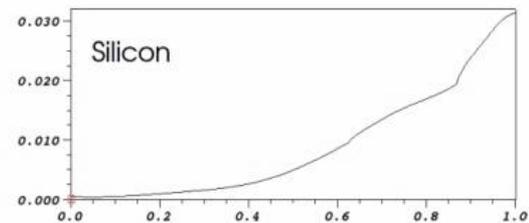
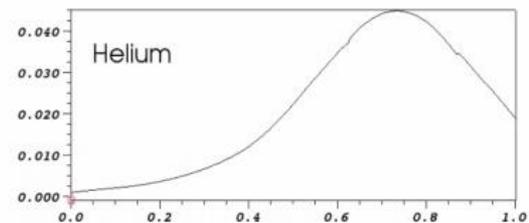
Iron Production



Star Surface



Amount produced
(in solar masses = $2e33g$)



Time (seconds)

Post Ignition Study Summary

- We have performed full-star simulations of early post ignition flame dynamics at unprecedented resolution
 - Compressible formulation
 - Adaptive mesh refinement
 - Performing science at 64K cores, scaling to 200K+ cores
- Main scientific conclusions:
 - Turbulent flow field has little effect on expected ignition conditions, but will have a stronger effect for more central ignition
 - Flame speeds prescribed by flame model have little effect since buoyant rise speeds dominate

Summary

- Recent advancements in mathematical modeling, numerical methods, and supercomputing allow us to gain new insight on complex phenomena such as Type Ia supernovae.
- In order to solve such problems, teams of interdisciplinary scientists, engineers, and mathematicians must closely work together.