Astrophysical Modeling

- Cosmology - Mike Norman
- Type Ia Supernovae - Stan Woosley and John Bell
- Core-collapse Supernovae - (Adam Burrows),
  Stan Woosley, and John Bell
- General Relativistic Applications - ?

Physics
Turbulence/resolution
Radiation transport

COMPUTATIONAL ASTROPHYSICS CONSORTIUM

Purpose:

- Improve our understanding of supernovae of all types through the use of large scale computing.
- Design codes for the efficient study of hydrodynamics and radiation transport on the largest, fastest available machines.
- Train postdocs and graduate students in computational physics.
- Optimize and enhance the scientific return from astronomical missions - including JDEM, LSST, SWIFT,and ground-based supernova searches.

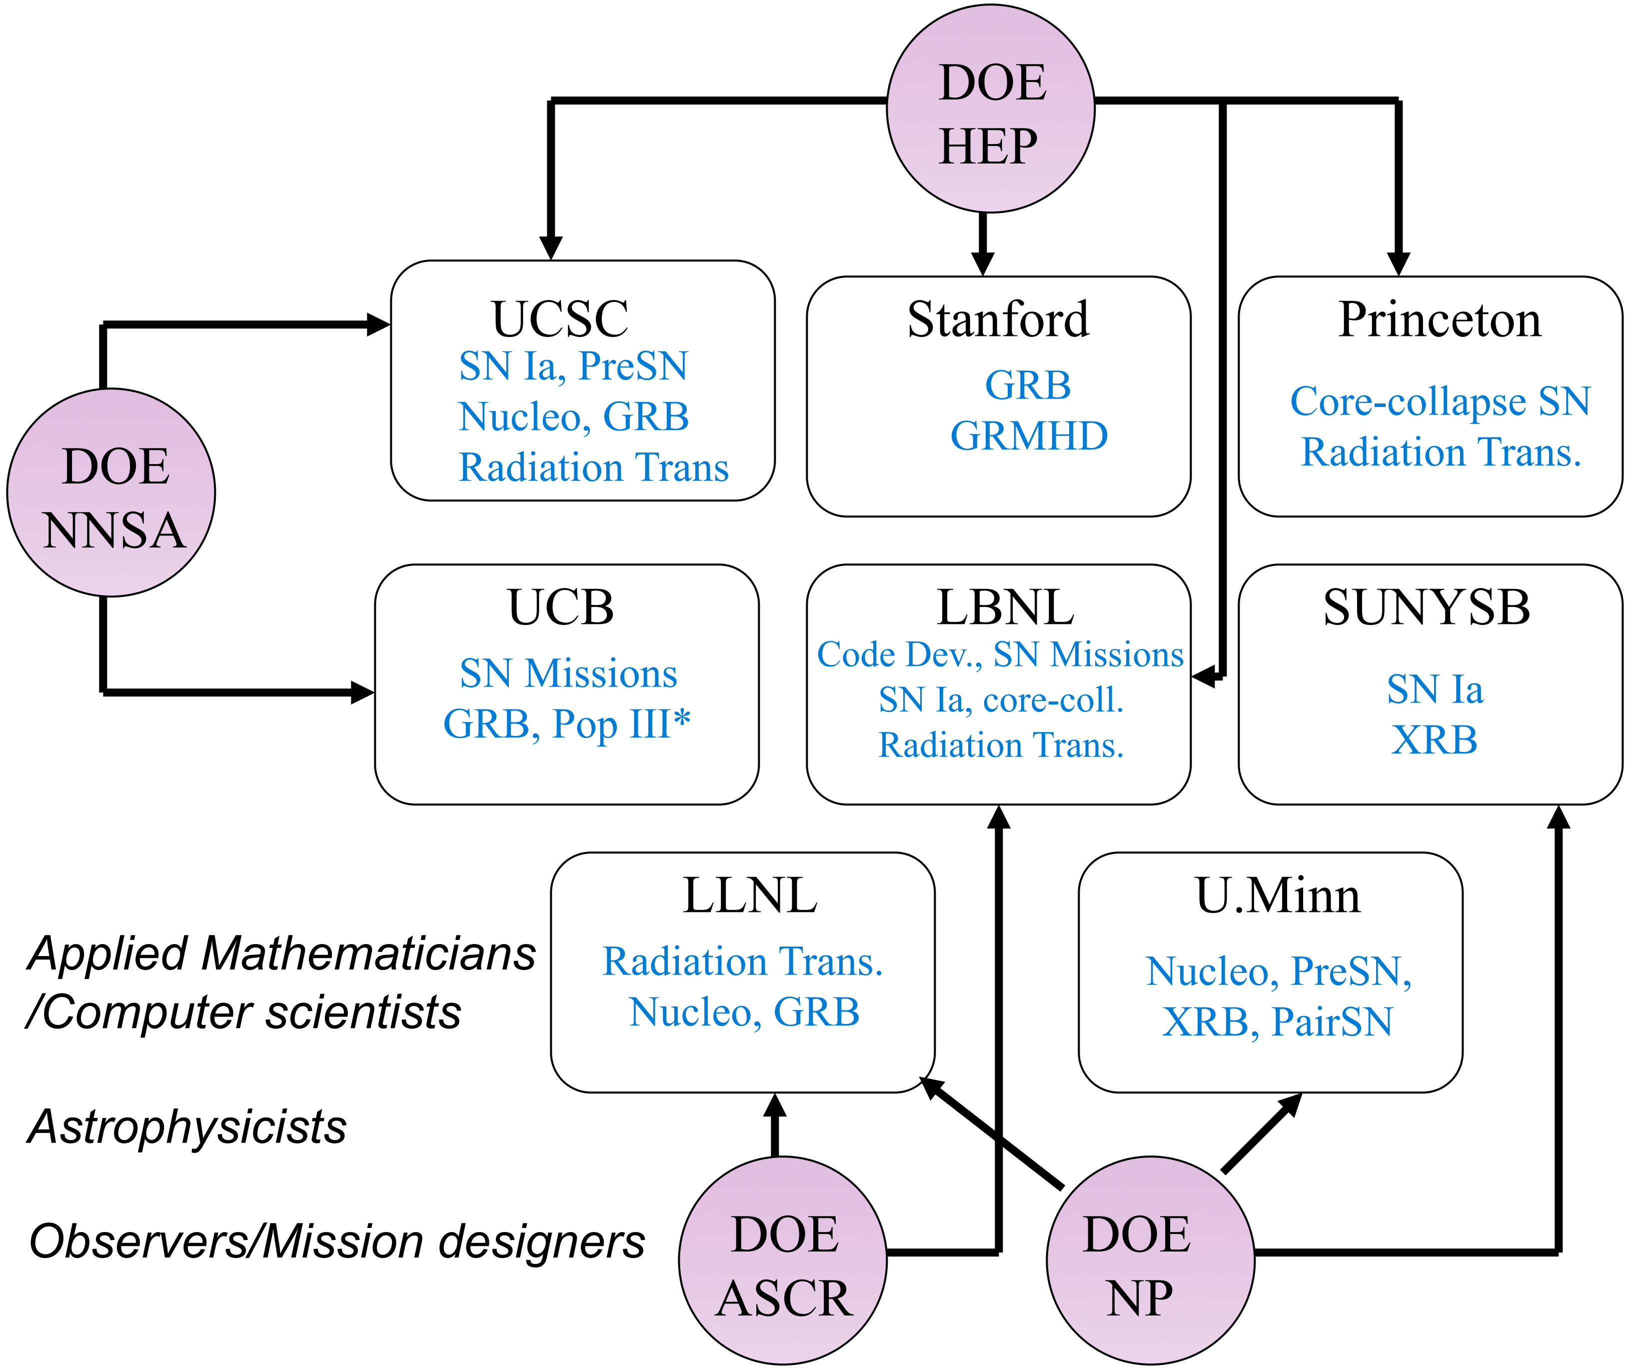
DOE HEP
UCSC
SN Ia, PreSN
Nucleo, GRB
Radiation Trans
Stanford
GRB
GRMHD
Princeton
Core-collapse SN
Radiation Trans.
DOE NNSA
UCB
SN Missions
GRB, Pop III*
LBNL
Code Dev., SN Missions
SN Ia, core-coll.
Radiation Trans.
SUNYSB
SN Ia
XRB
LLNL
Radiation Trans.
Nucleo, GRB
U.Minn
Nucleo, PreSN,
XRB, PairSN
Applied Mathematicians
/Computer scientists
Astrophysicists
Observers/Mission designers
DOE ASCR
DOE NP

Three years of preparation, code writing, and smaller scale simulations, and we are ready to compute on a larger scale…

Type Ia Supernovae

LBNL - John Bell, Ann Almgren, Andy Aspden

UCSC - Stan Woosley, Dan Kasen, Haitao Ma

SUNYSB - Mike Zingale, Chris Malone

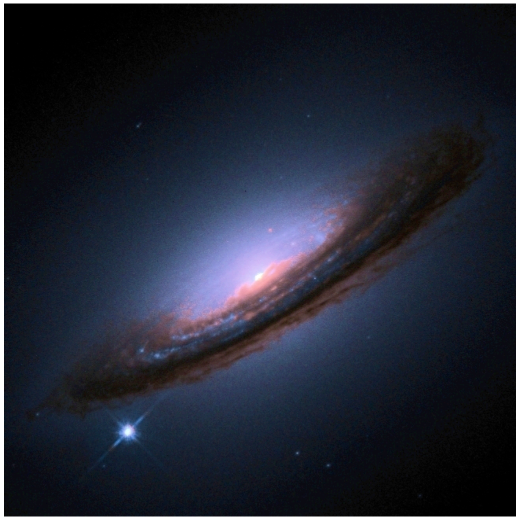

SN 1994D

SN 1998dh

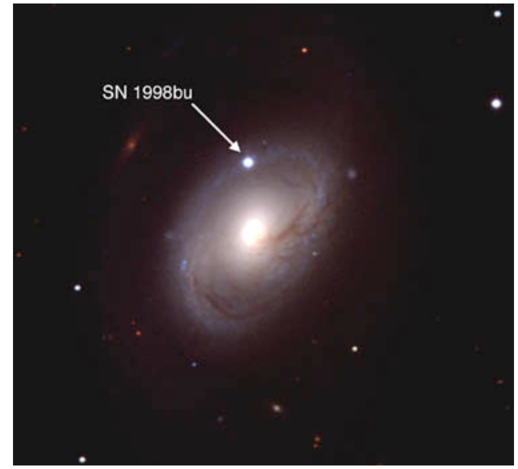

SN 1998bu

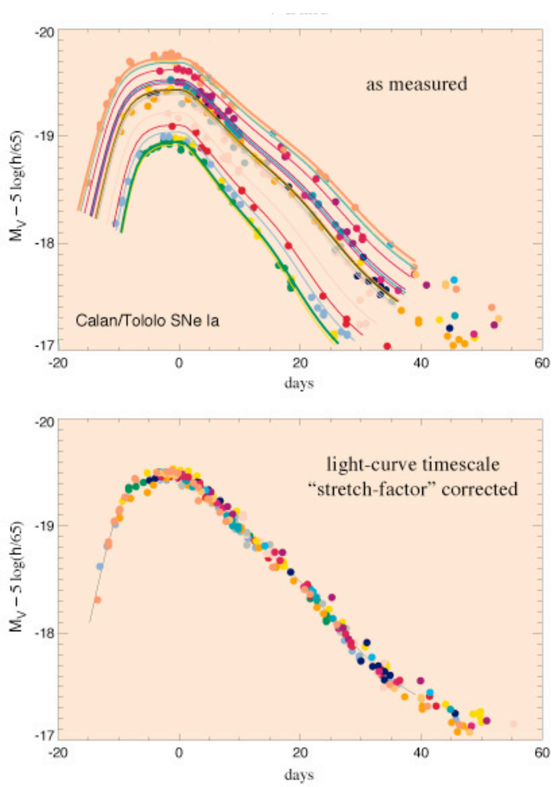

HST

Understanding Type Ia supernovae is an important both because they are a long standing problem in astrophysics and because of their application to (precision) cosmology.

Supernova Discovery History

Asiago Catalog (all supernova types)

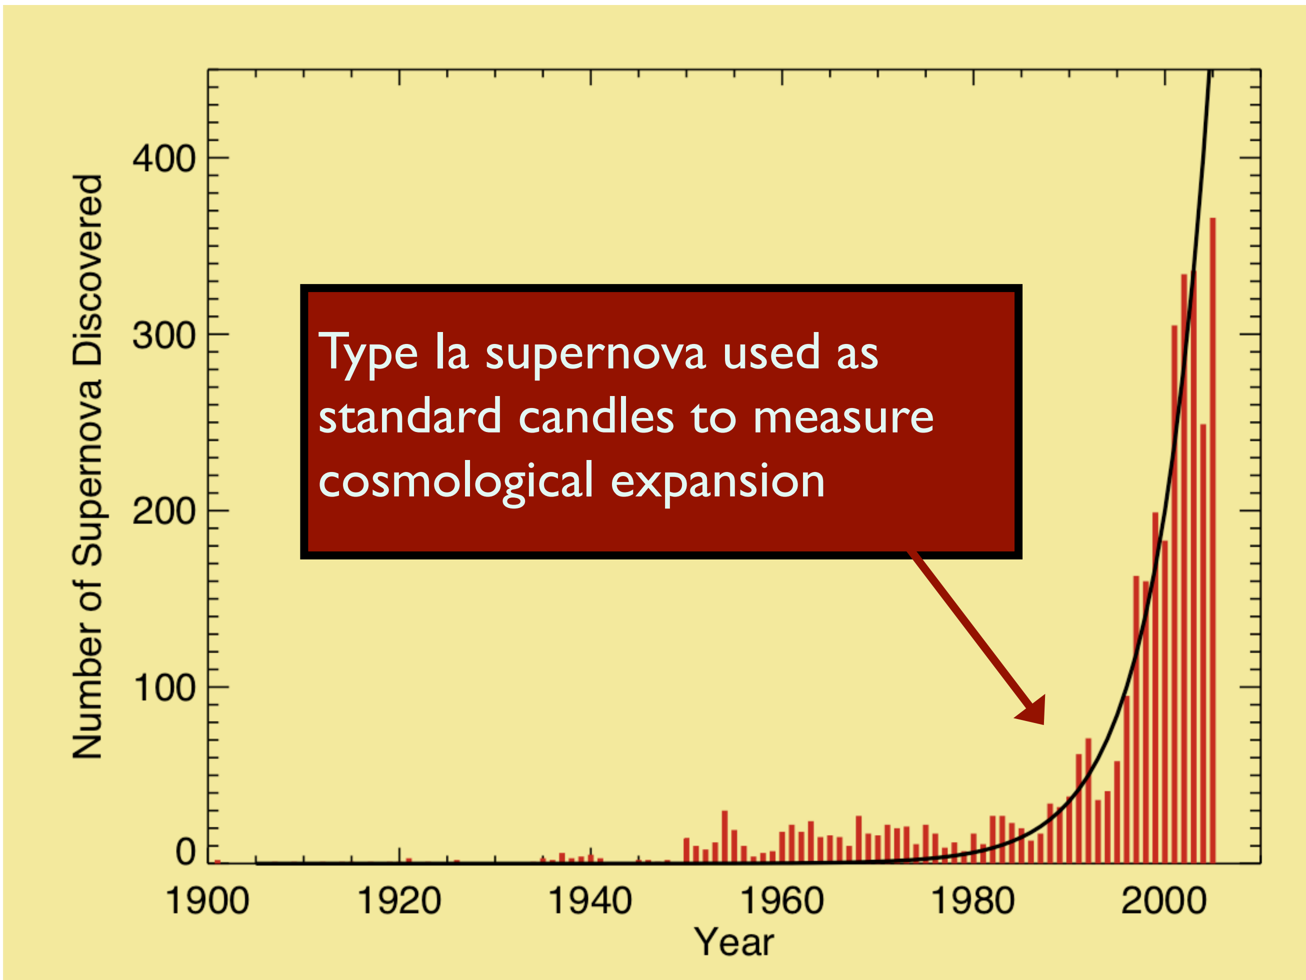

Discovery - Future

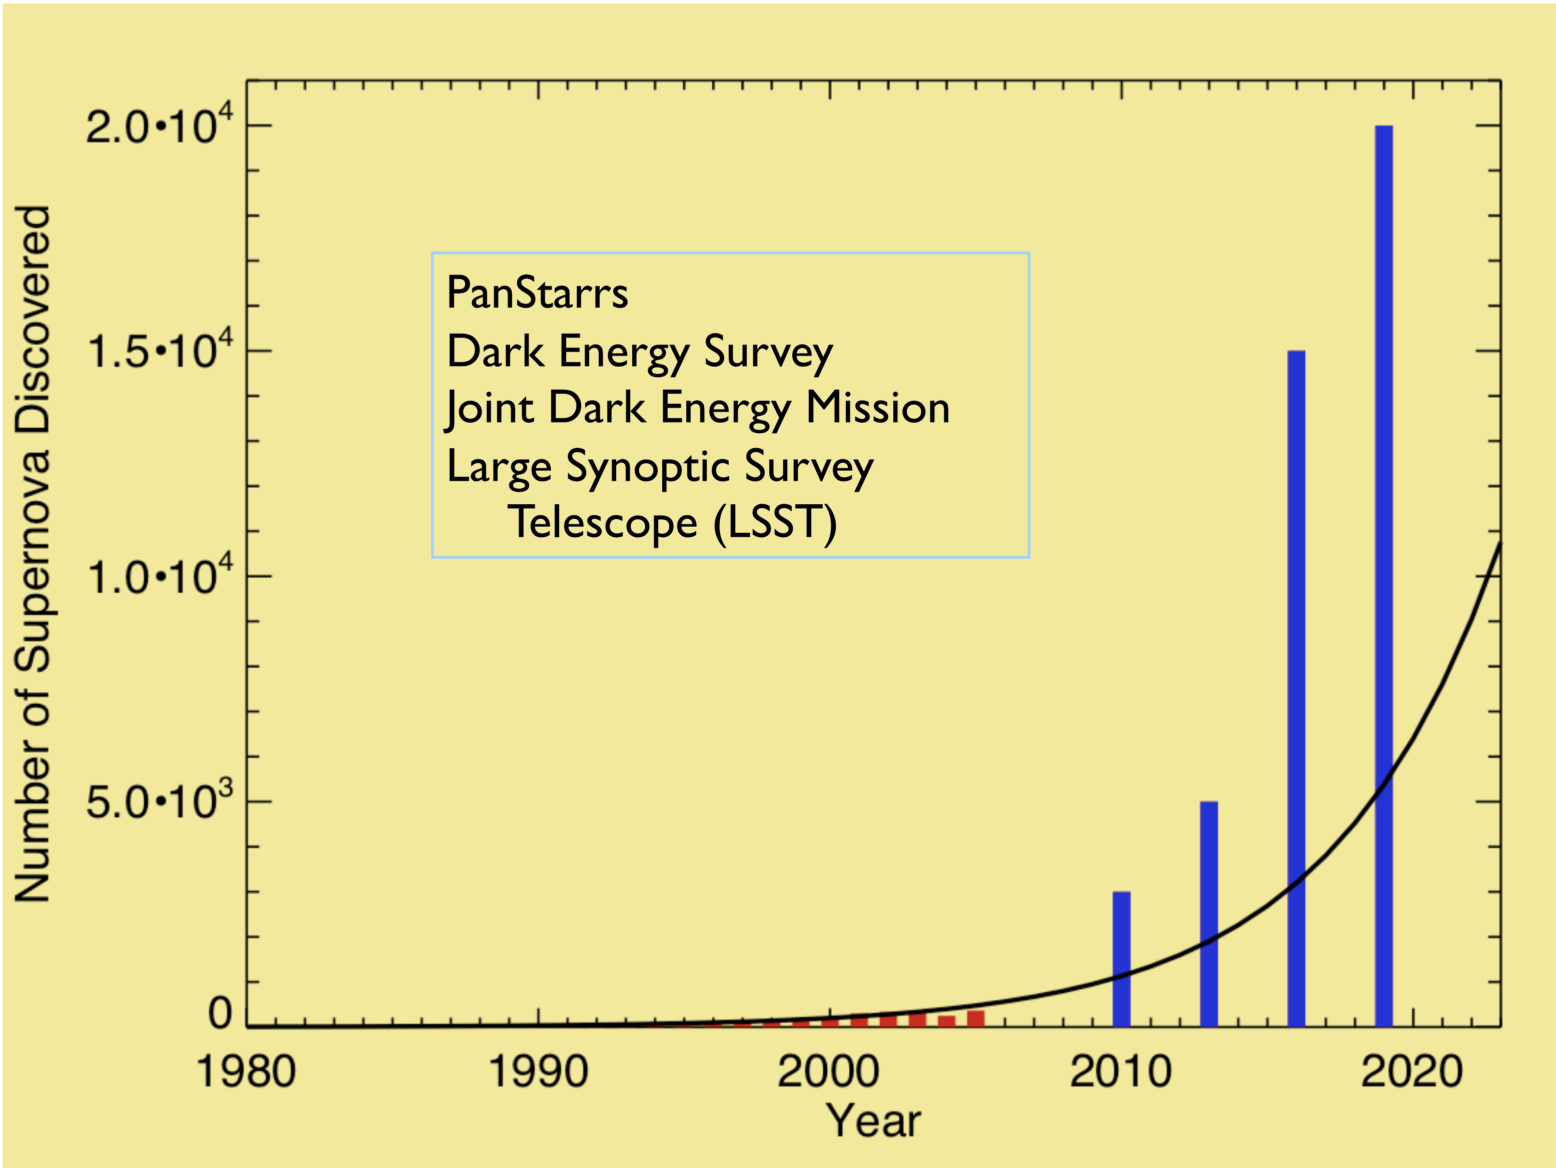

Type Ia Supernova - Basic Model

The best model for SN Ia is the thermonuclear explosion of a Chandrasekhar mass carbon-oxygen white dwarf (1.38 solar masses).

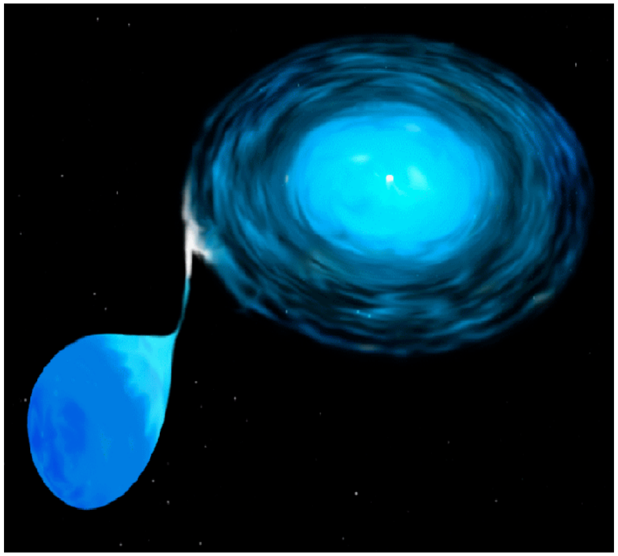

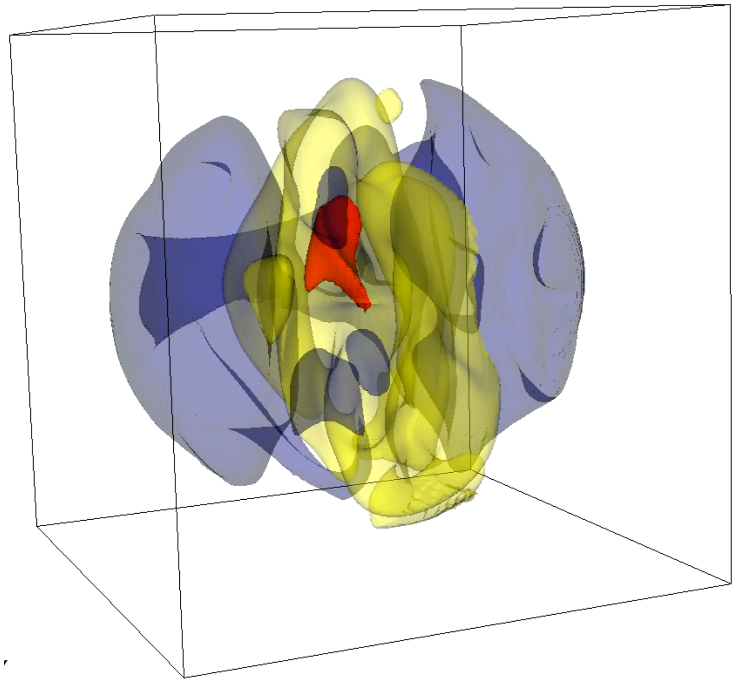

As material piles on the star, the central temperature increases, and carbon fusion reactions begin driving convection.

Eventually, the cooling cannot keep up with the energy release from reactions, and a burning front is born.

There are four areas where plan to make major progress:

- Determining the ignition conditions for the explosion. Where and when does the initial flame front form?
- How the flame moves through the star in response to the instabilities and turbulence that the burning itself creates
- How the subsonic deflagration makes a transition to detonation at late times. Can a spontaneous transition occur?
- What the final explosion looks like. What is its spectrum and brightness at all wavelengths, at all angles for all times?

QuickTime™ and a
decompressor
are needed to see this picture.

xyz slices of a calculation of SN Ia ignition
Central T = 6 x 10^8, last 100 seconds before explosion. Ignition occurs off center on one side

Reynolds Number implies barely turbulent.

The Explosion - Burning and Propagation

QuickTime™ and a
decompressor
are needed to see this picture.

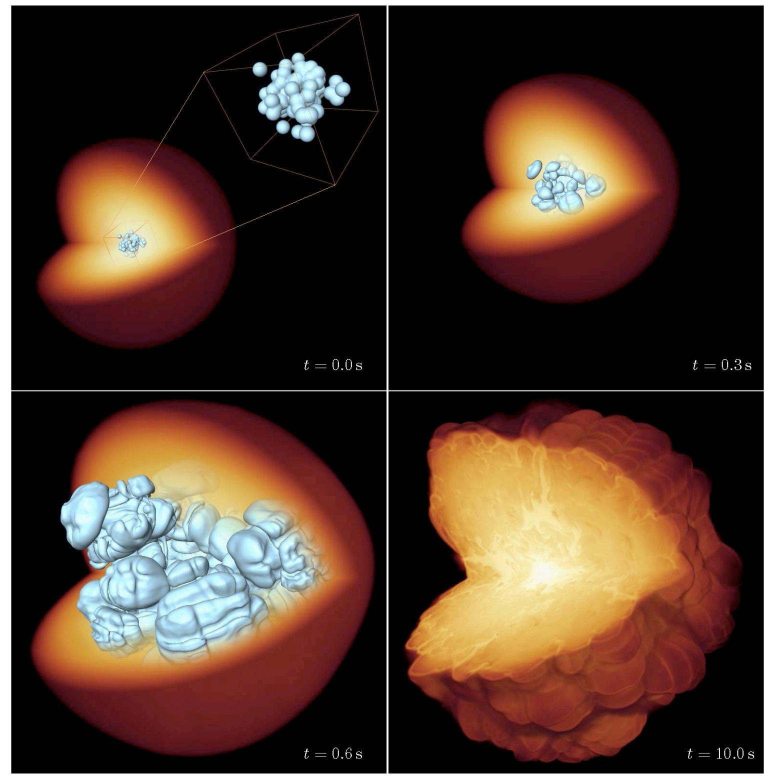

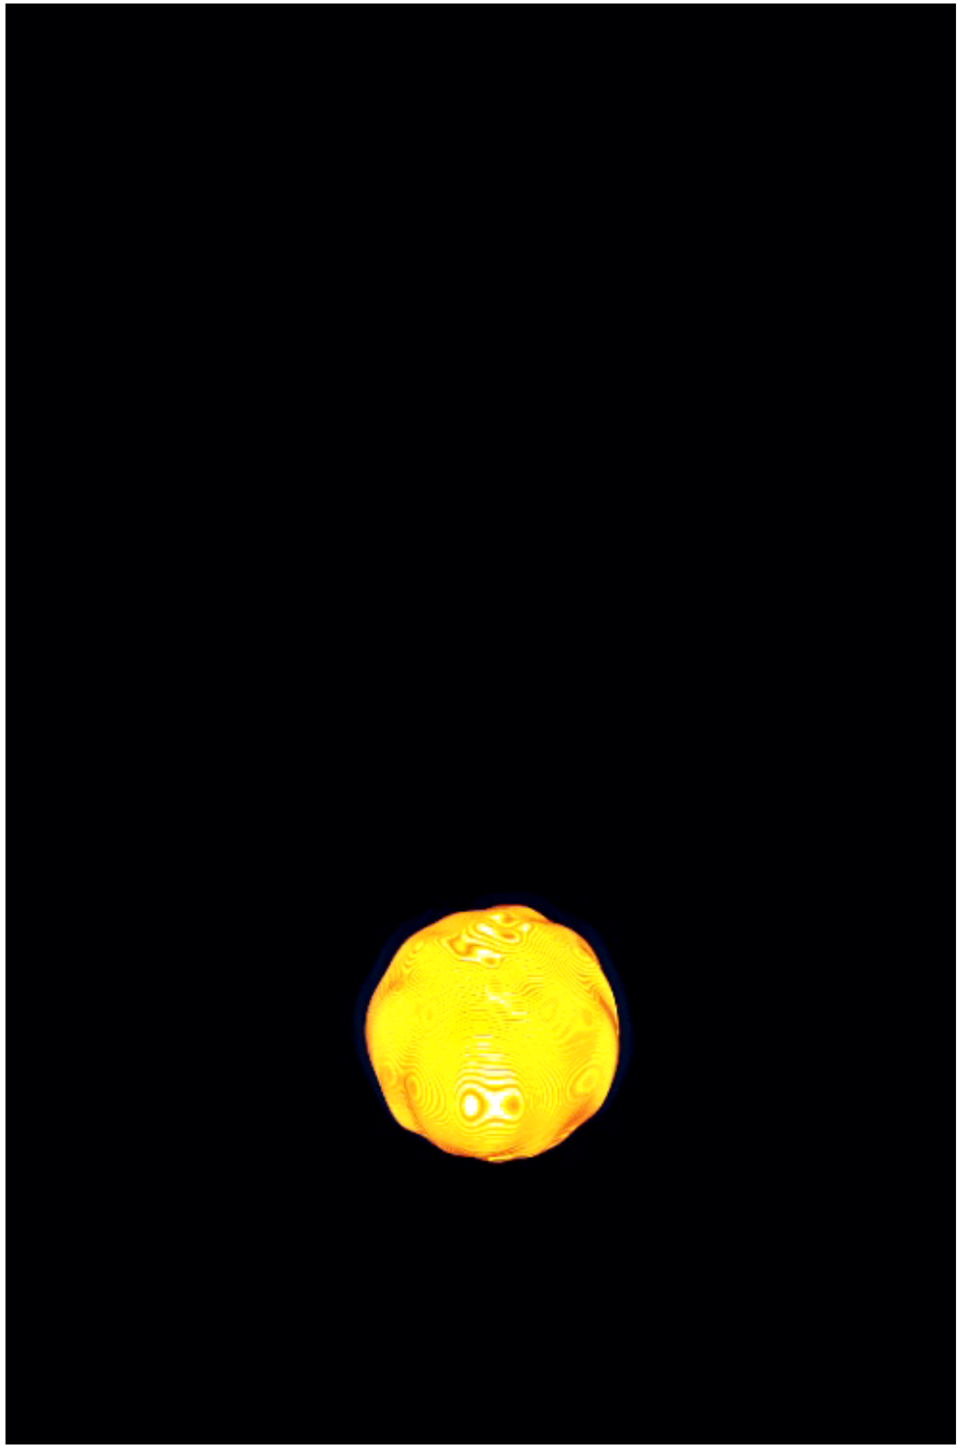

Burning Floating Bubble

500x500x2048
1 MCPUHr
so far at ATLAS

Aspden et al (2009)

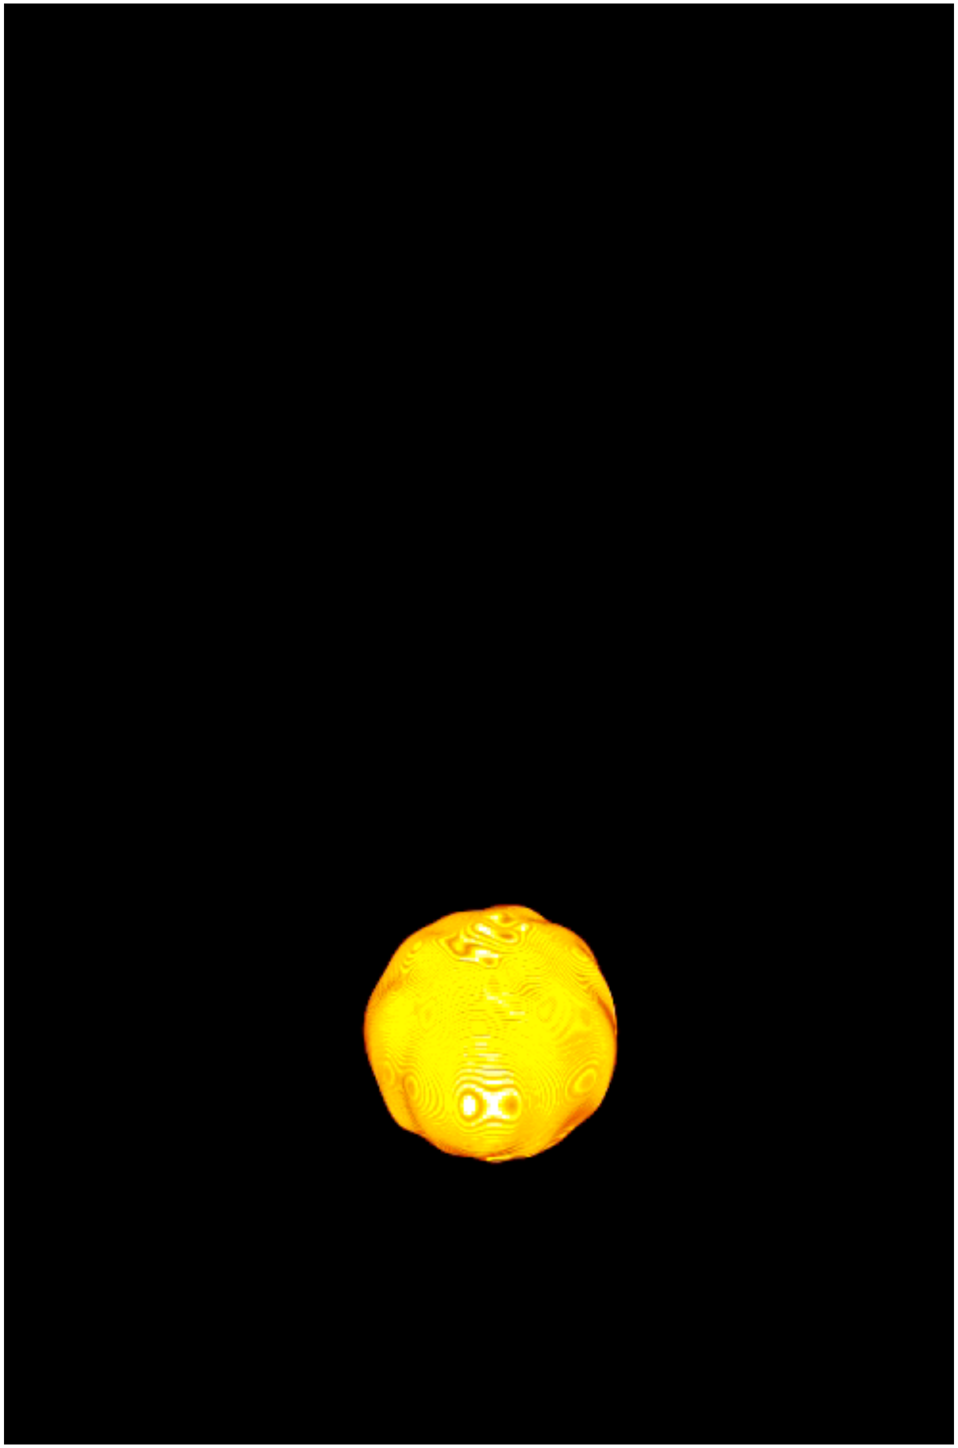

ENERGY
GENERATION
ONLY

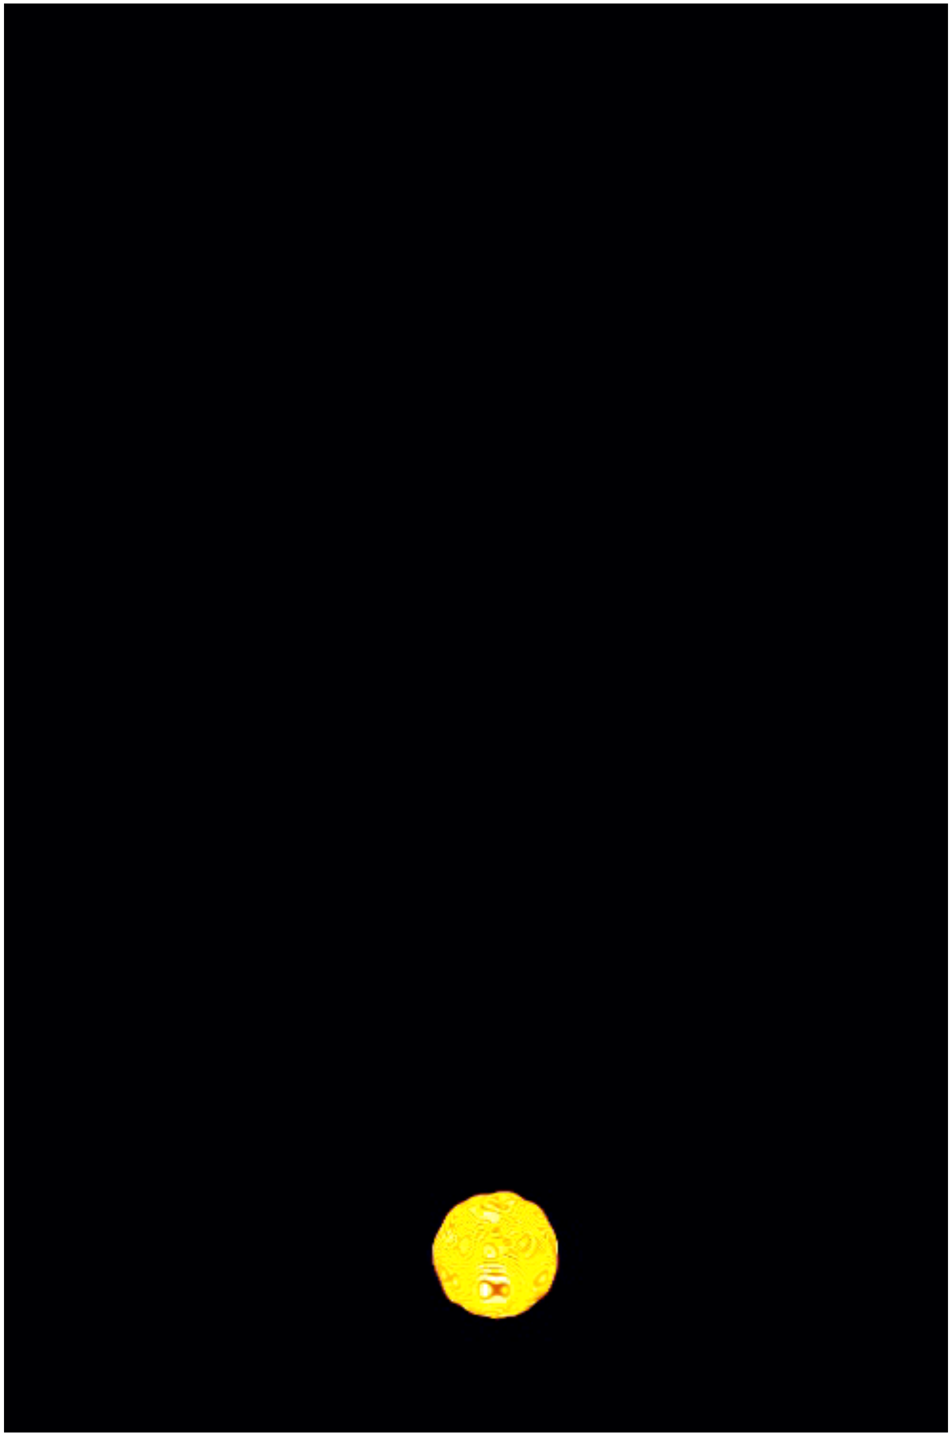

LOWER
RESOLUTION

QuickTime™ and a
decompressor
are needed to see this picture.

1024^3 zone calculation in Munich. Barely resolves integral scale for the turbulence. ~ 1 M CPU hr

SN Ia with CASTRO

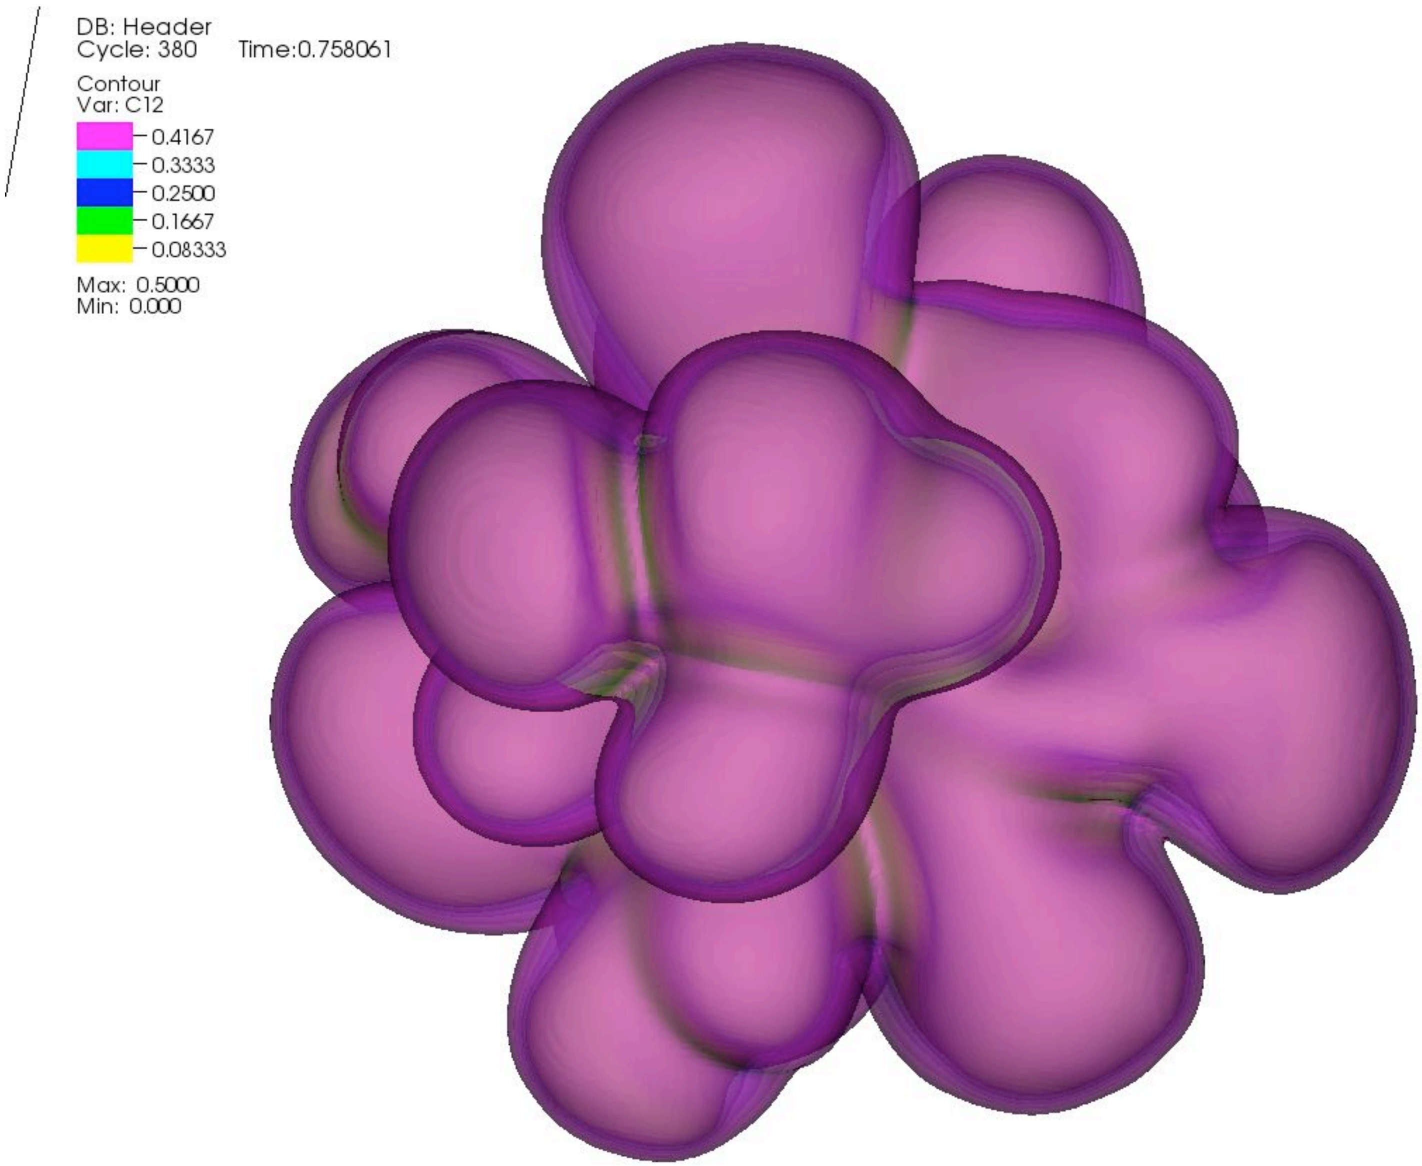

128^3 with two levels of AMR.

Equivalent to 512^3 for the finest zones (most of star at late times)

256 to 1024 CPU

5 hrs into run

CkPt = 23 GB

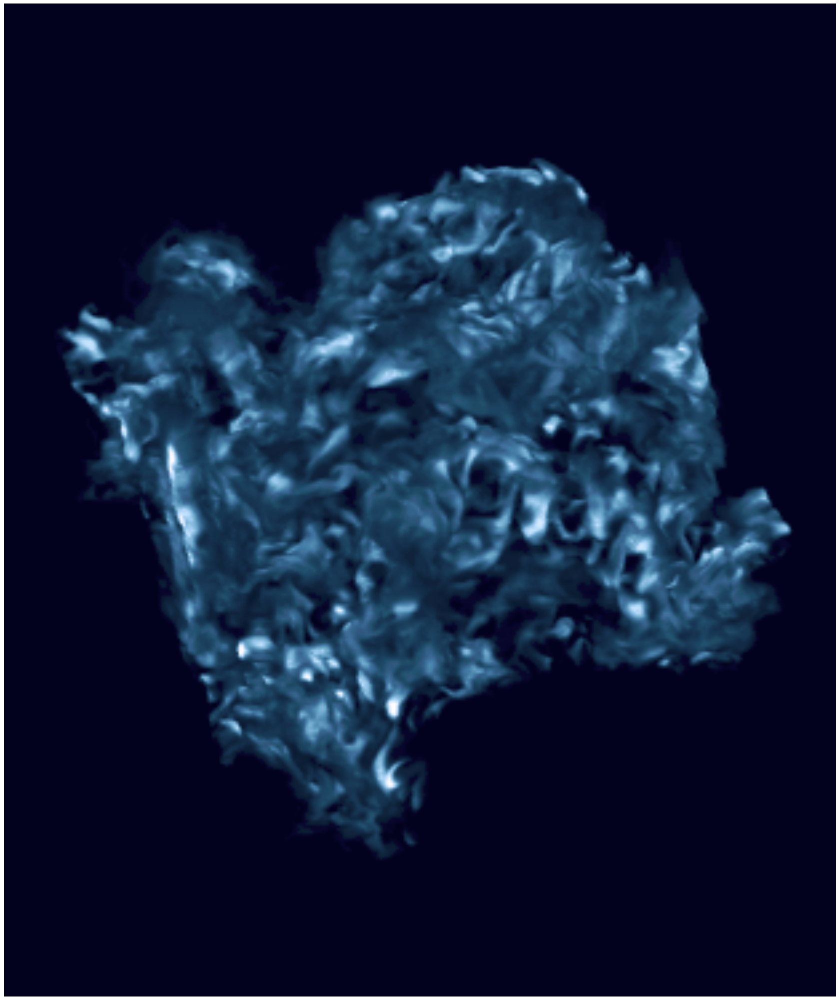

[illegible]= 2 x 10^7 g cm^{-3}

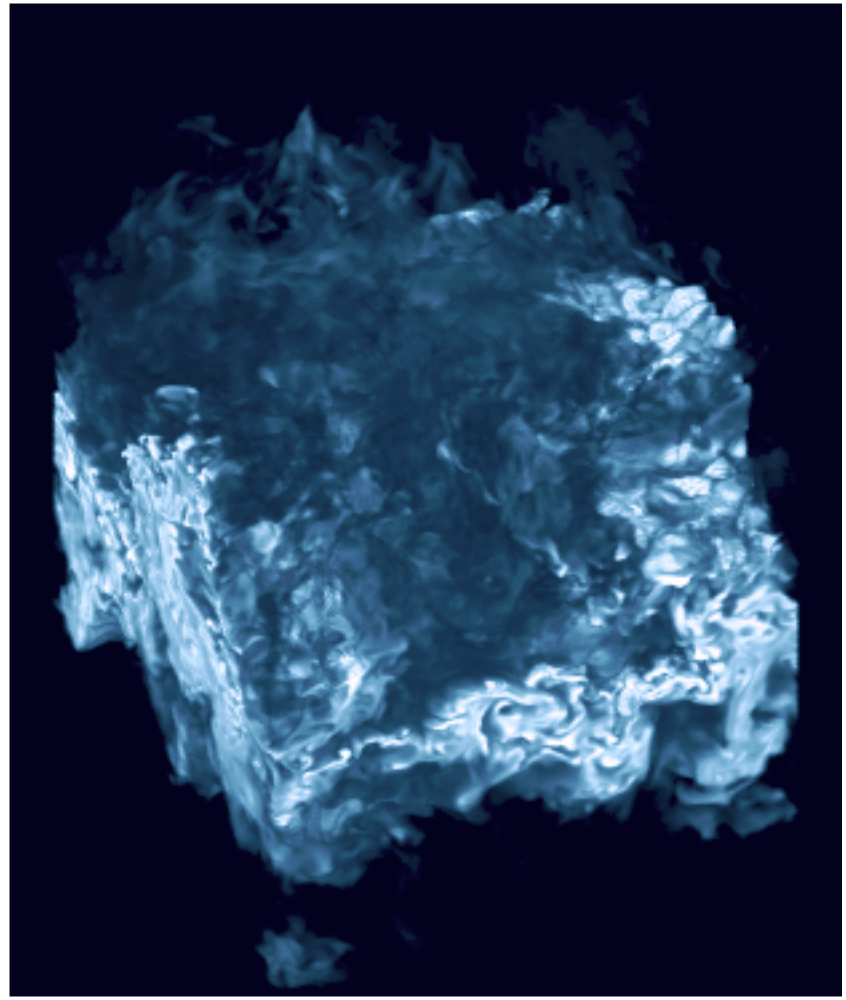

[illegible]= 1 x 10^7 g cm^{-3}

As the density declines below a critical value turbulence tears the flame leading to mixing and a transition to detonation. (> 1 MCPUhr; box is ~ 1 m)

WORK WITH KASEN AND RÖPKE
ON 2D MODELS FOR SN - NERSC and ORNL

- A grid of 44 2D SN Ia explosion models with “realistic” ignition and detonation conditions with multi-D light curves and spectra for each. Good agreement with the observed width-luminosity relation.

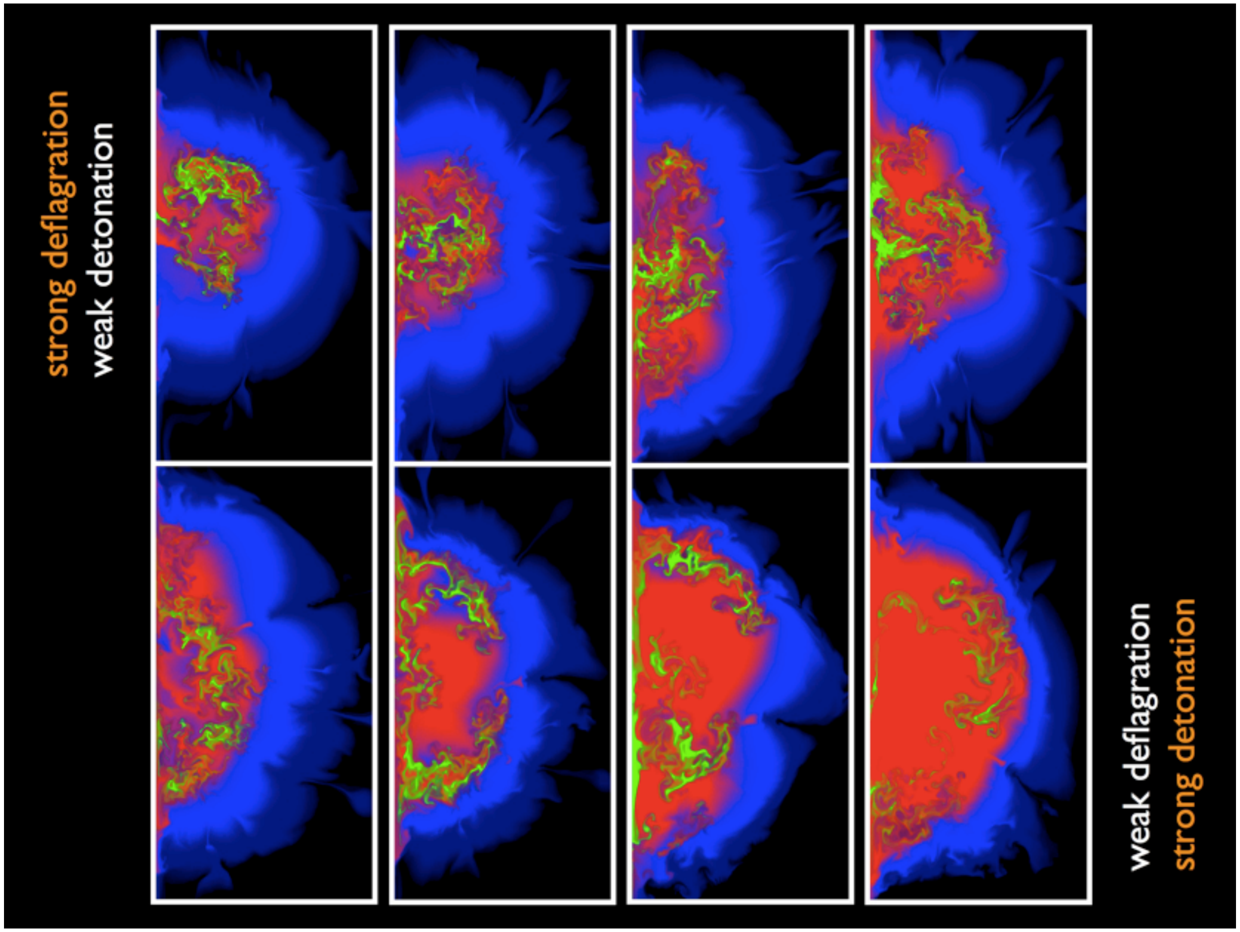

Kasen, Röpke, and Woosley, Nature, 2009.

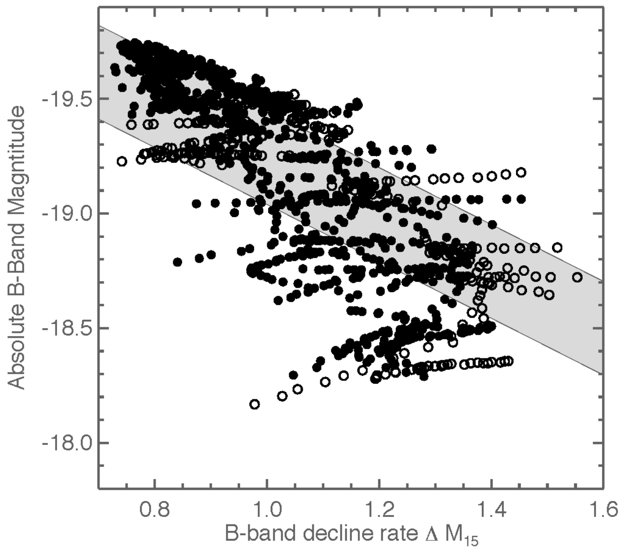

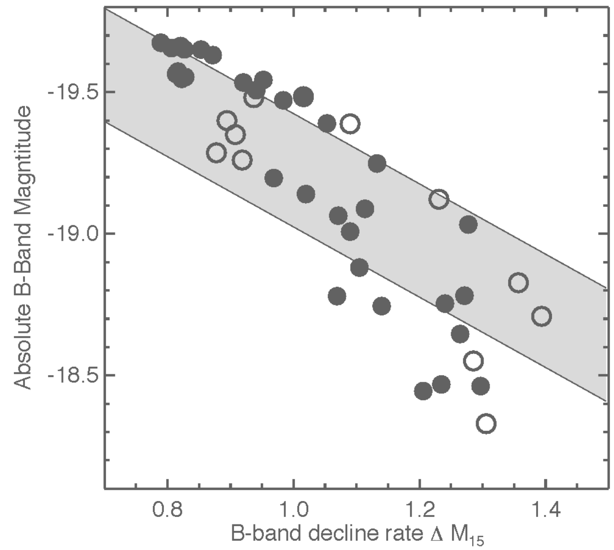

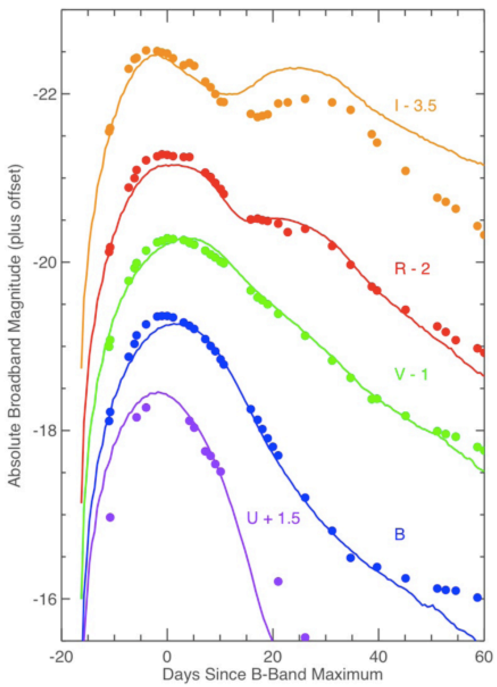

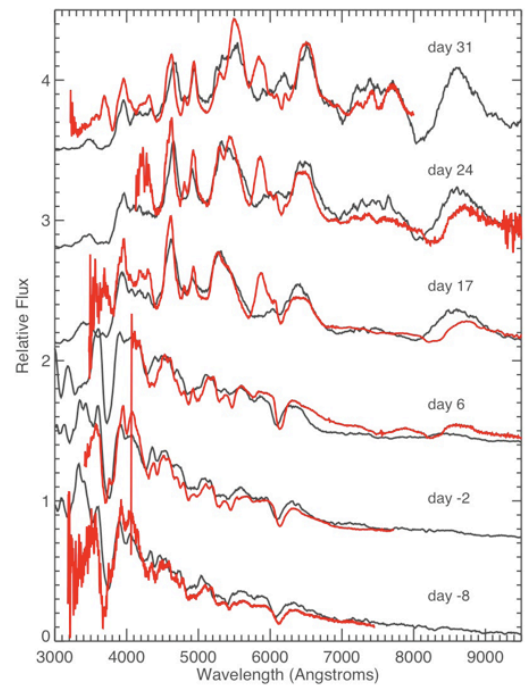

SN 2003 du
vs
Model

3-D supernova spectrum calculation

pure deflagration model from Roepke et al, 2007

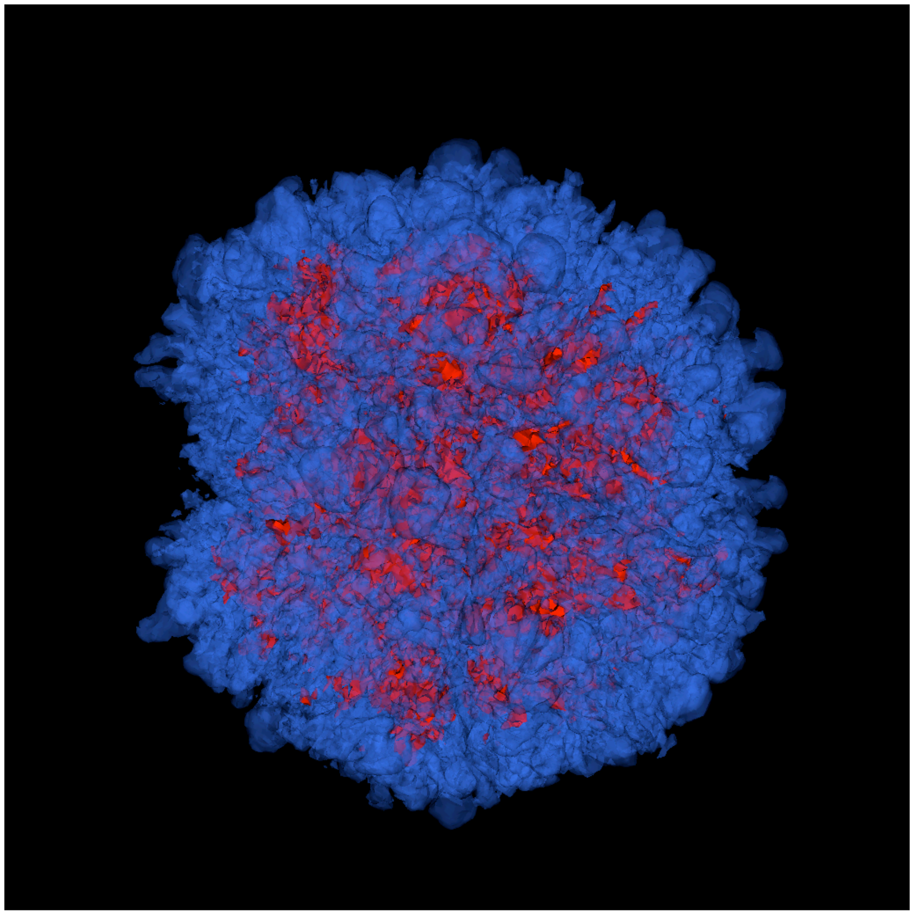

Calculation on 1000 cores
Cray XT4 Jaguar @ ORNL

spatial resolution
150 x 150 x 150

wavelength resolution
1 angstrom (2×10^4 points)

memory usage
1 TB

execution time 10,000 CPU-hours - one spectrum

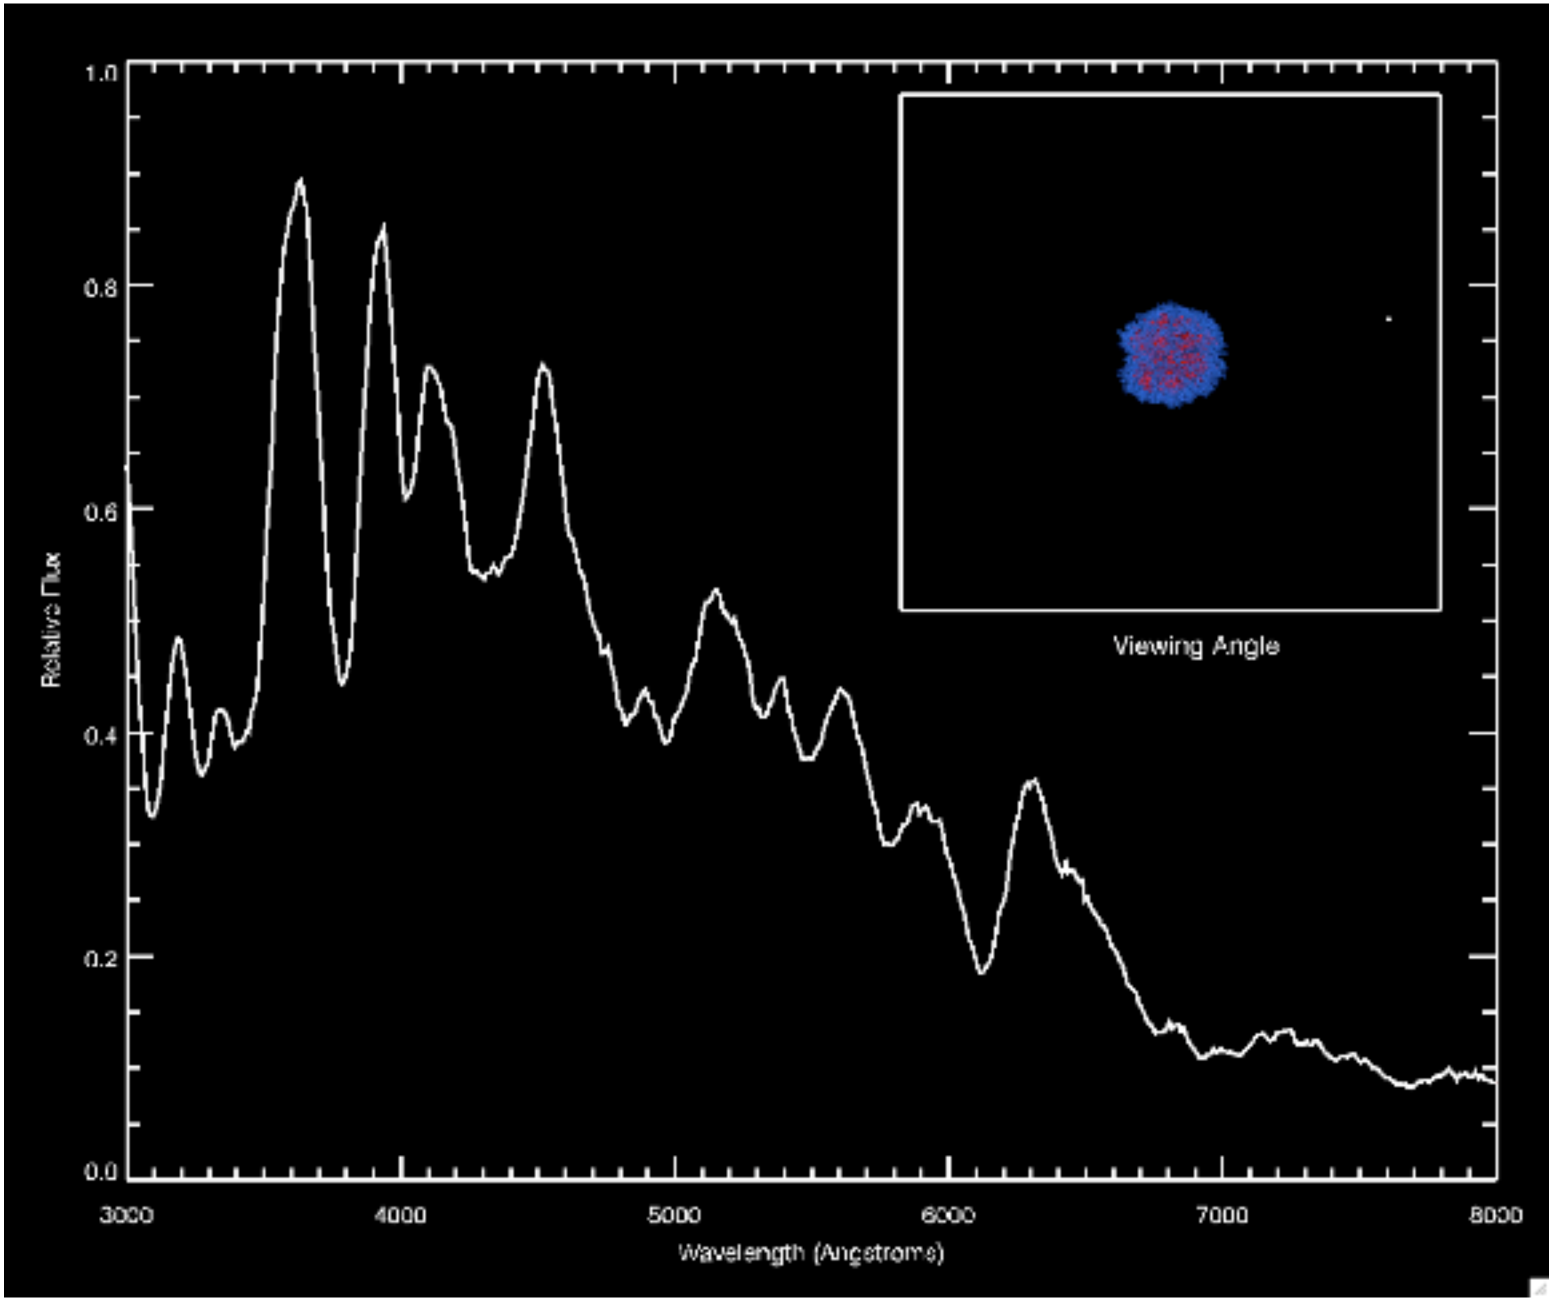

1.0
0.8
0.6
0.4
0.2
0.0
Relative Flux
3000
4000
5000
6000
7000
8000
Wavelength (Angstroms)
Viewing Angle

2. Current HPC Requirements

- Necessary software, services or infrastructure
  Compilers and visualization tools are presently adequate but we will need improvements in the future, Currently use F90, C++, OPENMP, MPI, htar, VisIT

- Current primary codes and their methods or algorithms
  MAESTRO - low Mach number code. Background hydrostatic equilibrium. Sound waves are filtered out of the system. To enforce the thermodynamics, an elliptic constraint on the velocity field is enforced

  CASTRO - Eulerian compressible radiation - hydrodynamics code, unsplit PPM, adaptive mesh, multiple time steps, spherical, Cartesian and cylindrical coordinates, general EOS, self gravity, reaction networks

  SEDONA - multi-dimensional Monte Carlo code. Currently assumes thermal level populations for all ions (LTE). Used at all wavelengths. Highly scalable.

MAESTRO parallel performance to 13,000 CPU

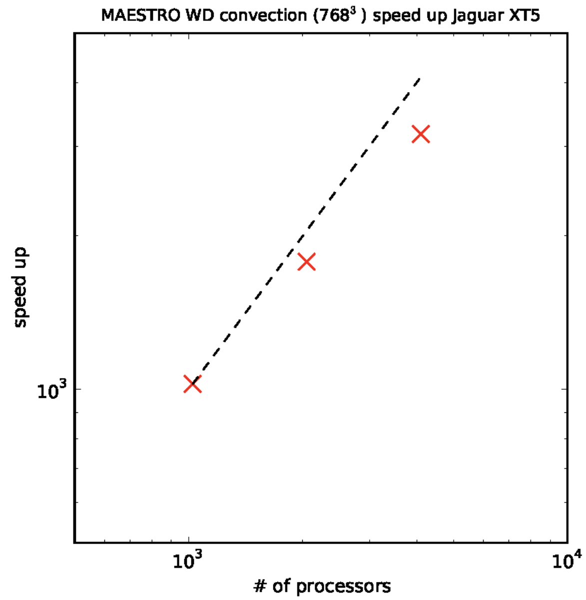

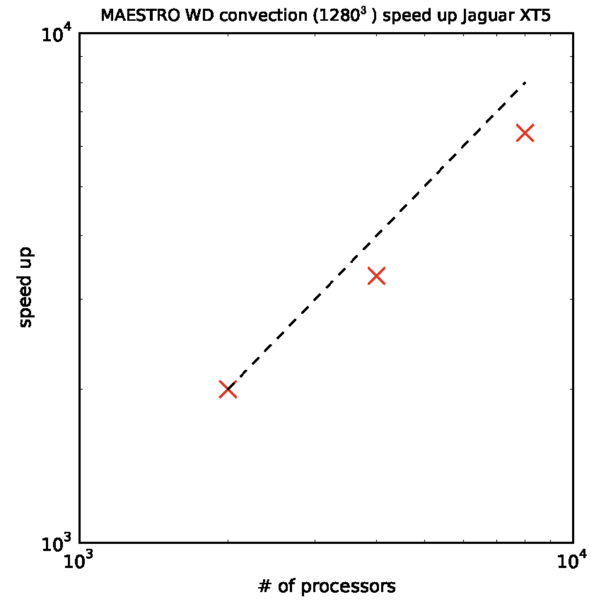

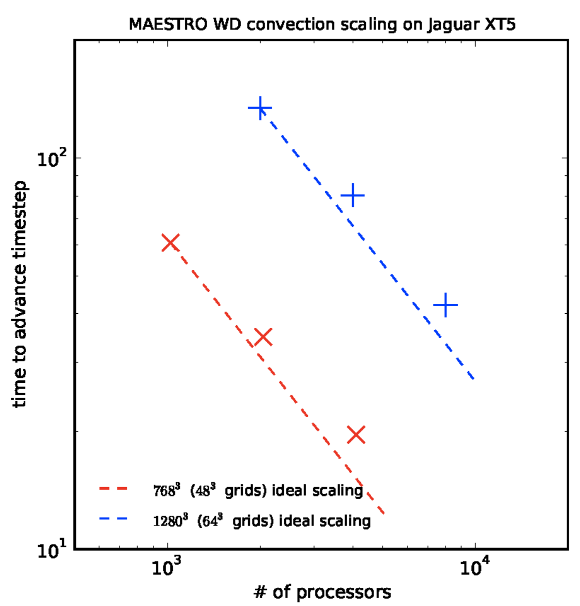

Figure 1: Maestro strong scaling results for the white dwarf convection problem proposed here. We note that this is a run of exactly the problem we intend to run, using the machine targeted in this proposal (the Jaguar XT5 machine at ORNL). Two different problem sizes were used, 768^3 and 1280^3. In each case, the number of grids in the domain decomposition was fixed as we increase the processor count. The top plots show the speed up curves for each resolution. The plot on the left shows the average time to advance a single timestep as a function of processor count. In all plots, ideal scaling is represented by a dashed line.

CASTRO parallel performance to 60,000 CPU without radiation

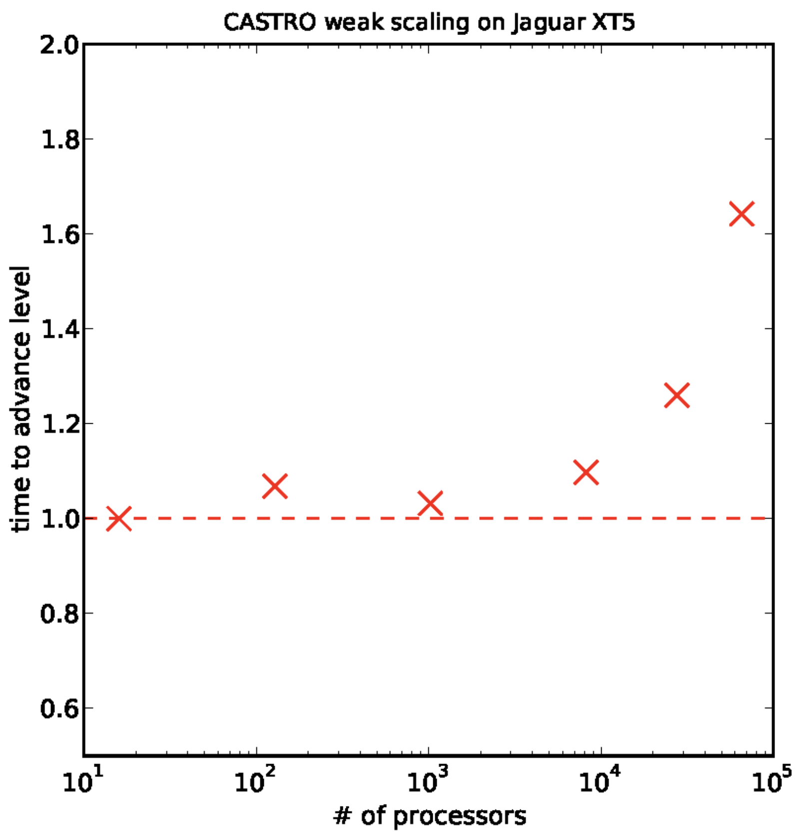

Figure 2: CASTRO weak scaling of an adaptive time step on the Jaguar XT5 machine. This problem put a white dwarf on the grid with the full stellar equation of state and monopole gravity. This represents the base physics needed for the full star explosion problems.

SEDONA parallel performance To >> 10,000

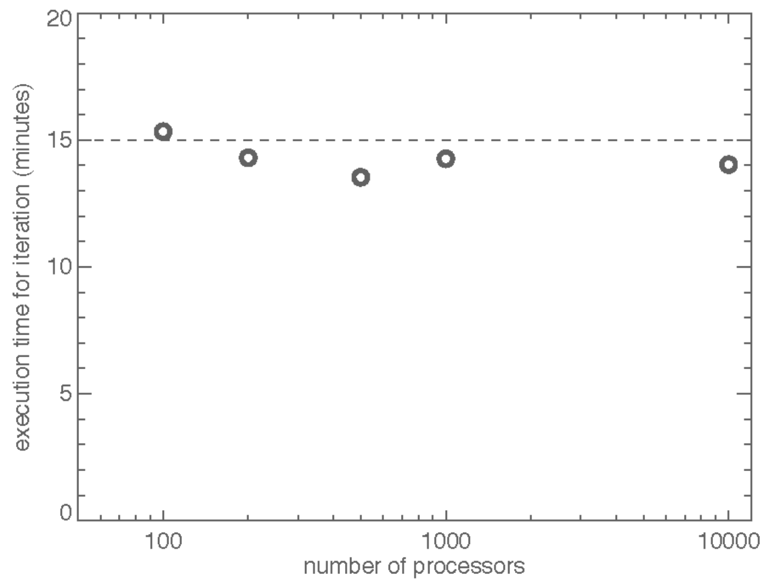

Figure 3: Weak scaling of the SEDONA code for the case of full replication parallelism, run on Jaguar at ORNL. This test problem was one iteration of a 2D SNe Ia light curve calculation. The total number of Monte Carlo photon packets was increased so as to keep the total number of packets per processor a constant.

	Ignition studies with MAESTRO	Full star explosion models with CASTRO	Studies of transition to detonation with CASTRO	Light curve and spectra calculations with SEDONA
Current year				
Number of runs	5 at 384^3	10 at 1024^3 deflagration stage	6	10
Mega-CPU hr - all runs	3	2	2	1
Number processors	1728	1024	1024	4096
Size check point file (GB)	10	20	20	0.01
Total output (TB)/run	6	6	2	0.01
On-line storage (TB)	3	3	1	0.01
Off-line storage (TB)	30	20	10	0.2
Next Year (2010)				
Number of runs	10 at 768^3	10 at 1024^3; 10 at 4096^3 include detonation	10	30 3D-LTE
Mega-CPU hr – all runs	60	40	10	12
Number processors	5K-10K	4K-16K	4K – 16 K	20 K
Size check point file (GB)	80	50	50	10
Total output (TB)/run	30	50	15	0.1
On-line storage (TB)	10	10	3	0.3
Off-line storage (TB)	300	500	150	3

Three years from now per year				
Number of runs	10 (higher resolution)	20 at 4096^3 include detonation	10	30 3D – LTE 2 x 3D -NLTE
Mega-CPU hr – all runs	75	> 50	15	32
Number processors	10K-15K	10K-40K	10 K – 40 K	20 K
Size check point file (GB)	100	150	100	10-100
Total output (TB)/run	40	100	40	0.1
On-line storage (TB)	10	20	10	0.3
Off-line storage (TB)	400	1000	300	4

3. HPC Usage and Methods for the Next 3-5 Years

- Upcoming changes to codes/methods/approaches

  Soon, CASTRO will be adapted to include level set tracking of burning fronts and a subgrid model for turbulence.

  Initial conditions for CASTRO runs will be taken from MAESTRO rather than being parameterized

  Include detonation physics and nucleosynthesis post-processing in CASTRO

  Modify SEDONA to be non-LTE in multi-dimensions

- Changes to Compute/memory load

  We anticipate at least 1.5 orders of magnitude increase in CPU, memory and storage use because the codes have now reached a state of readiness for production runs and we are moving to 3D.

3. HPC Usage and Methods for the Next 3-5 Years

- Changes to Data read/written

  Again we anticipate a greater than one order of magnitude increase - see table

- Changes to necessary software, services or infrastructure

  Improved programming models to support hieracrhical parallel approaches

  Tools for automatic program tuning

  Tools to facilitate rapid archiving and accessing of data

- Anticipated limitations/obstacles/bottlenecks on 10K-1000K PE system.

  Elliptic solves in MAESTRO

3. HPC Usage and Methods for the Next 3-5 Years

- Strategy for dealing with multi-core/many-core architectures

  We are currently using OPENMP as the model for loop-level parallelization. Preliminary results suggest that this will be an effective strategy, at least up to a modest number of cores per node. It would potentially be helpful to have a more “light-weight” approach with less overhead to starting threads than OPENMP. Our codes could potentially use GPUs or other accelerators effectively but the system would need to be configured to move data between into and out of the accelerator quickly; a huge latency in getting data into a GPU, for example, could make it difficult to use effectively.

4. Summary

- Recommendations on NERSC architecture, system configuration and associated service requirements needed for your science:

  Need Viz/analysis hardware at NERSC

- What significant scientific progress could you achieve over the next 5 years with access to ~50X NERSC resources?

  Actually we already need 50 x our current NERSC resources to accomplish next years goals. We are thus applying at a variety of facilities. The progress we plan to make includes moving to well-resolved 3D studies of ignition and full star 3D models of the explosion with low and moderate resolution of the flame

Core Collapse Supernovae

Princeton - Adam Burrows, Jason Nordhaus

LLNL - Louis Howell

LBNL - John Bell, Ann Almgren

UCSC - Stan Woosley, Candace Church, Luke Roberts

U. Minnesota - Alex Heger

Why is the problem interesting?

- The death of massive stars produces the most energetic explosions in the universe and is responsible for the production of most of the elements heavier than helium

- A laboratory in which novel particle physics and high density physics is important and can be tested

- A classic problem in astrophysics (60 years) and in computational astrophysics (40 years)

- Definitely a problem for the largest machines available

Core-Collapse Supernovae

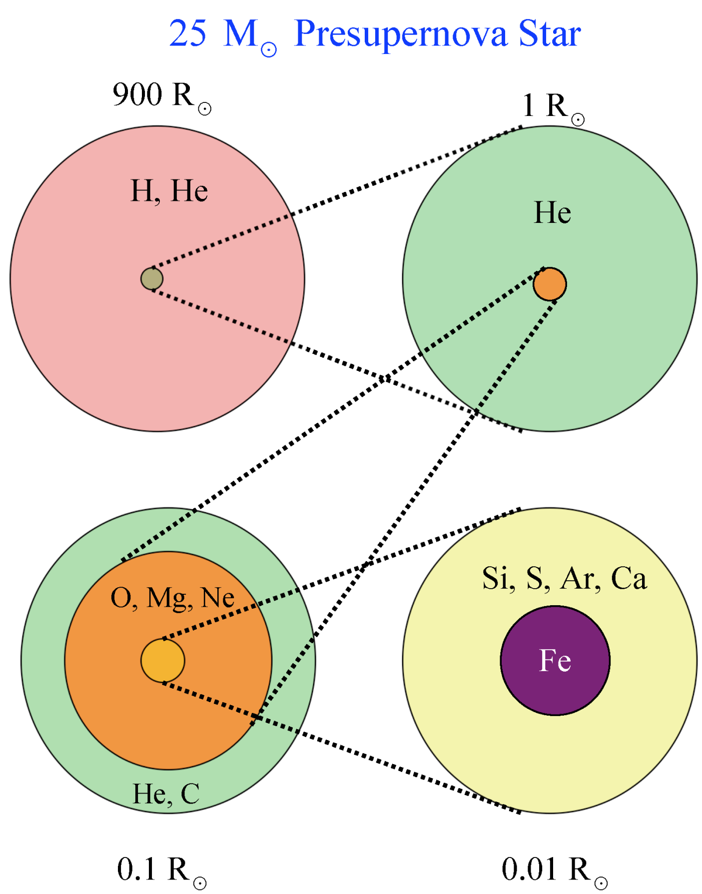

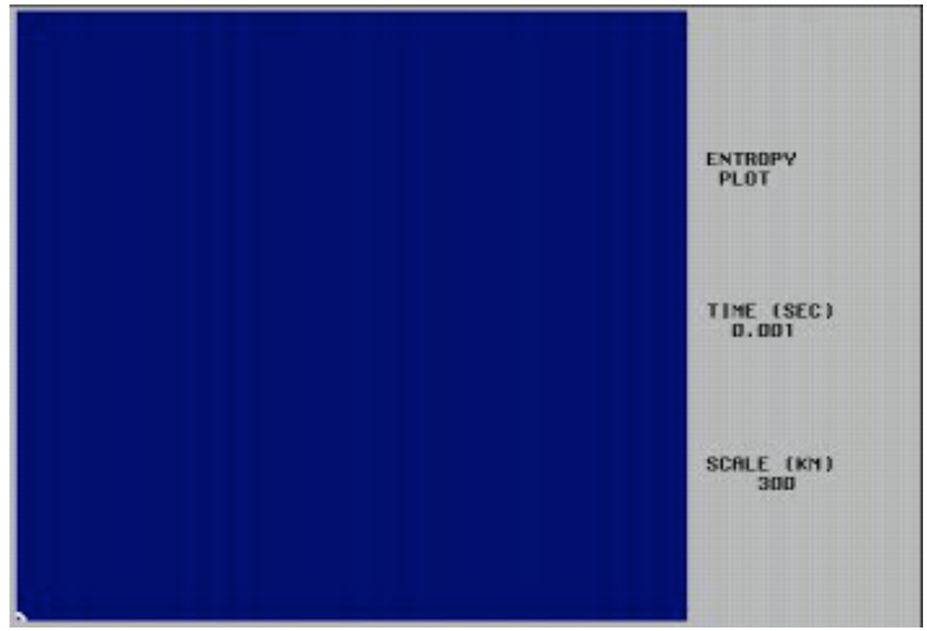

Why is this a problem for the biggest machines?

- Necessarily a 3D problem since turbulent convection isinvolved and the convection affects the efficiency of the energy deposited by the neutrinos. Six kinds of neutrinos non-thermal energy distribution.

- Radiation adds momentum space to this, making full radiation hydrodynamics problems essentially 3 + 3 = 6 dimensional problems. As each dimension is added, the computational burden increases mutliplicatively.

  Hence, what is 1000 X 1000 X 1000 (space) times 100,000 (timesteps) = 10^{14} generalized grid points in spacetime is 1000 X 1000 X 1000 (space) times 100,000 (time steps) TIMES 20 X 20 (angles) TIMES 20 (energy groups) [for full multi-angle treatment]

= 8 X 10^{17}, or ~10000 times more challenging.

- One can decrease this computational requirement by approximating the solution in full phase space by employing two-moment closures or flux limiting, but this only entails a factor of ~100 savings. Hence, there is a penalty of at least a factor of 100 in including radiation transport with 3D hydro, and this is assuming excellent scaling.

- Most work so far has been in 2D. A 3 dimensional hydro simulations with scientific merit and no radiation transport now requires ~1 million hours on Franklin. Hence, a corresponding flux-limited calculation would require ~100 Mhours.

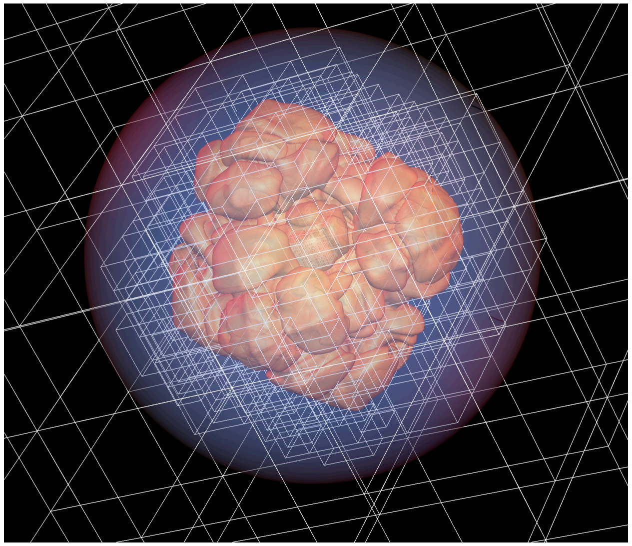

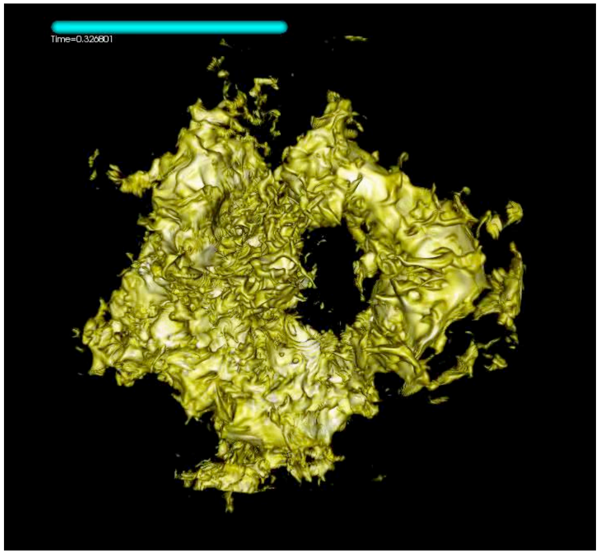

2. Current HPC Requirements

- **Architectures**

  Without neutrino transport, CASTRO has demonstrated good scaling to 60,000 CPU on Jaguar at ORNL and is running well on a variety of other architectures including FRANKLIN at NERSC and ATLAS at LLNL. With transport, architecture may be more of an issue. Currently we are restricted in runs doing MGFLD neutrino transport to 1000 CPU. In the next year we expect that to increase to 5000 CPU

- **Compute/memory load**

  CASTRO works well with 2 GB/CPU.

- **Data read/written**

  Checkpoint files for current runs are approximately 10 GB for 2D 128^2 runs; 200 GB for 3D hydro only runs. 3D with radiation will be 10 x bigger. Total output per 2D run is ~200 GB, for 3D hydro only, 5 TB. On line storage is ~5 TB. Archival storage about 20 TB.

2. Current HPC Requirements

- **Necessary software, services or infrastructure**

  Currently use F90, C++, OPENMP, MPI, htar. **hypre,** VisIT

- **Current primary codes and their methods or algorithms**

  CASTRO - Eulerian compressible radiation - hydrodynamics code, unsplit PPM, adaptive mesh, multi-group flux-limited diffusion, multiple time steps, spherical, Cartesian and cylindrical coordinates, general EOS including nuclear EOS, self gravity, reaction networks

Current Usage

Facilities Used or Using	NERSC OLCF ACLF NSF Centers Other:
Architectures Used	Cray XT IBM Power BlueGene Linux Cluster Other:
Total Computational Hours Used per Year	3.0 M Core-Hours
NERSC Hours Used in 2009	2.7 M Core-Hours
Number of Cores Used in Typical Production Run	2K – 16K
Wallclock Hours of Single Typical Production Run	150
Total Memory Used per Run	3000 - 25000 GB
Minimum Memory Required per Core	2 GB
Total Data Read & Written per Run	2000 GB
Size of Checkpoint File(s)	20-200 GB
Amount of Data Moved In/Out of NERSC	0.5 GB per year
On-Line File Storage Required (For I/O from a Running Job)	1000 GB and 2000 Files
Off-Line Archival Storage Required	50000 GB and 2500 Files

2. Current HPC Requirements

- **Known limitations/obstacles/bottlenecks**

  The radiation transport currently scales to only 1000 CPU even though the hydro scales much further. Howell expects to improve this to 5000 CPU next year. Further improvements may require improvements in solvers.

What could you do with 2 dex more resources - say 200 MCPU hr - that you can't do now?

- This would enable a few 3D rad/hydro simulation with multi-group flux-limited diffusion for a physically interesting number (~10^5 – 10^6) of time steps.

- We will soon have the code for this, but will need the resources to use them productively.

3. HPC Usage and Methods for the Next 3-5 Years

- **Upcoming changes to codes/methods/approaches**

  Improve scaling of radiation transport to much greater than 5000 CPU

  Explore alternate schemes for radiation transport including Monte Carlo

  Implement magneto-hydrodynamics

  Implement at least first order post-Newtonian gravity, red-shift corrections, etc.

  Apply MAESTRO to problems in presupernova evolution (convection) and perhaps either CASTRO or MAESTRO to MHD neutron star formation

- **Changes to Compute/memory load**

  We anticipate a 2 order of magnitude increase in CPU, memory and storage use because the code has now reached a state of readiness for production runs and we are moving to 3D. We are working on MPI/OPENMP approaches that will substantially reduce our memory/CPU requirements (currently 1 - 2 GB/core)

3. HPC Usage and Methods for the Next 3-5 Years

- **Changes to Data read/written**

  Again we anticipate a two order of magnitude increase. Low resolution full star 3D runs with neutrino transport will have restart dumps ~2 TB and total I/O of order 50 TB

- **Changes to necessary software, services or infrastructure**

  Improved programming models to support hieracrhical parallel approaches

  Tools for automatic program tuning

  Tools to facilitate rapid archiving and accessing of data

- **Anticipated limitations/obstacles/bottlenecks on 10K-1000K PE system.**

  No advanced radiation transport packages currently work on this many CPU

3. HPC Usage and Methods for the Next 3-5 Years

- **Strategy for dealing with multi-core/many-core architectures**
  We are currently using OPENMP as the model for loop-level parallelization. Preliminary results suggest that this will be an effective strategy, at least up to a modest number of cores per node. It would potentially be helpful to have a more "light-weight" approach with less overhead to starting threads than OPENMP. Our codes could potentially use GPUs or other accelerators effectively but the system would need to be configured to move data between into and out of the accelerator quickly; a huge latency in getting data into a GPU, for example, could make it difficult to use effectively.

Three Years

Computational Hours Required per Year	>150 MCPU hr
Anticipated Number of Cores to be Used in a Typical Production Run	20K-100K depending on development
Anticipated Wallclock to be Used in a Typical Production Run Using the Number of Cores Given Above	500-1000 Hr.
Anticipated Total Memory Used per Run	8-50 TB
Anticipated Minimum Memory Required per Core	0.5 GB
Anticipated total data read & written per run	4-400 TB

Anticipated size of checkpoint file(s)	2-10 TB
Anticipated On-Line File Storage Required (For I/O from a Running Job)	40 TB and 200-2000 Files Inconsistent with above
Anticipated Amount of Data Moved In/Out of NERSC	0.5 GB per year
Anticipated Off-Line Archival Storage Required	1 PB and 10000 Files

Ultimately MHD must be included.

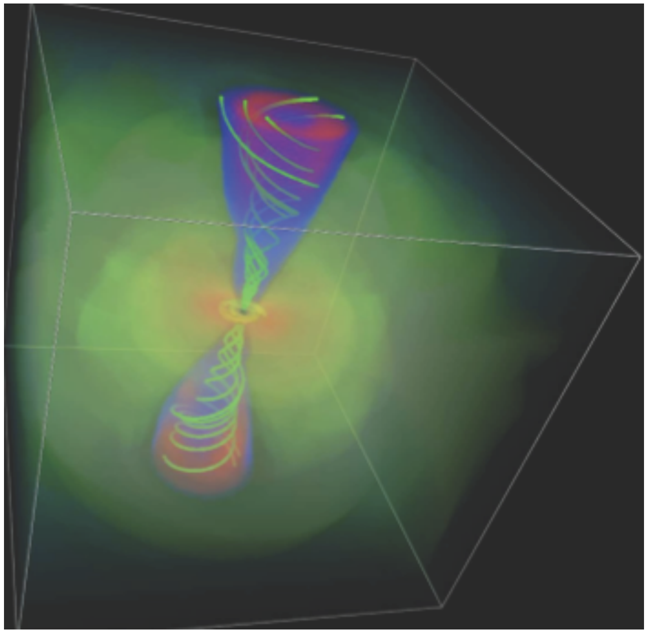

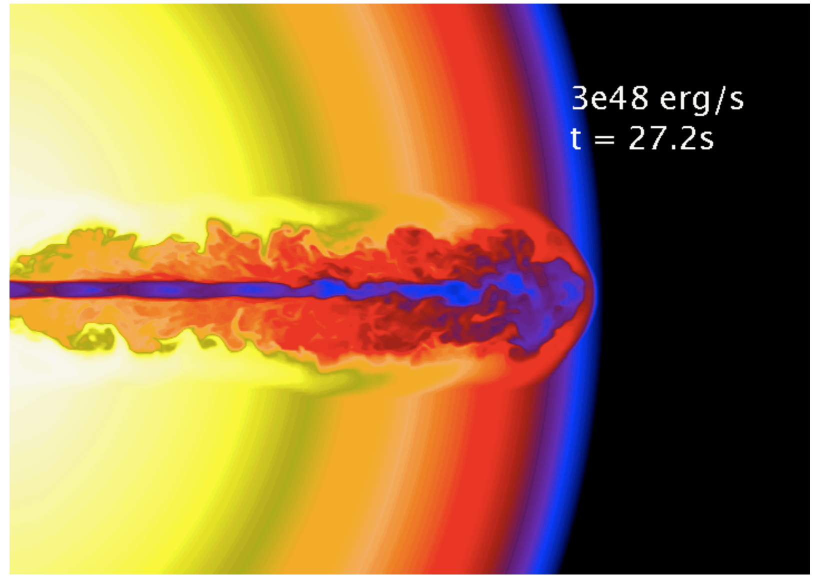

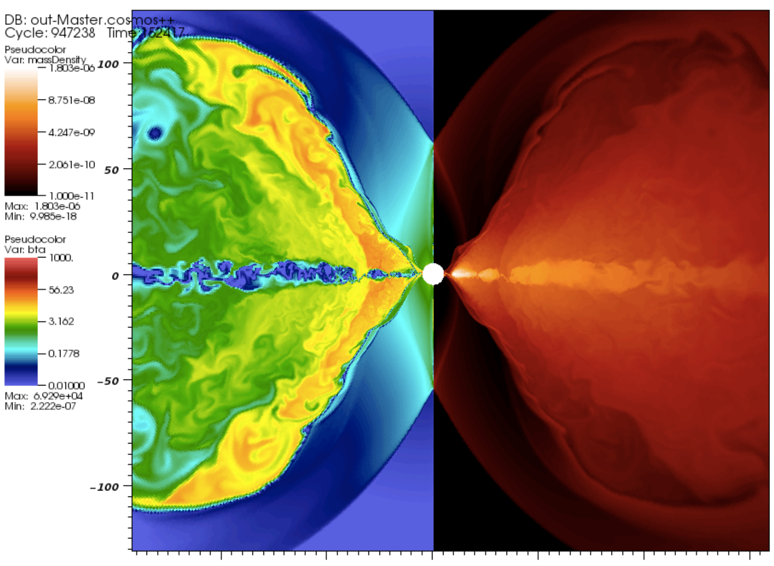

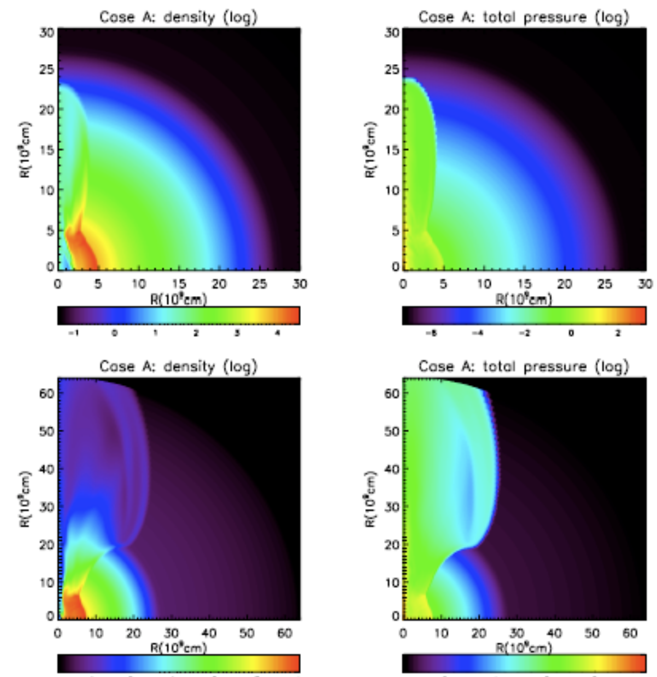

Other anticipated developments in next few years

- 3D studies of electron-positron pair instability supernovae
- Special relativity added to CASTRO. Treat relativistic jet propagation in gamma-ray bursts
- Nucleosynthetic postprocessing to obtain full yields both for studies of element production and spectra
- Light curves and spectra of all models using SEDONA
- Studies of radiation-hydrodynamics in collisions with circumstellar shells - Pulsational-pair-instability supernovae
- Studies of how the explosion and nucleosynthesis are affected by the assumed neutrino properties and mixing
- Studies of shock break-out using CASTRO and SEDONA
- Neutron star kicks, gravitational radiation, neutrino signal, etc.