Present and Future Requirements for

m461:Stellar Explosions in Three Dimensions

Tomek Plewa (Florida State University)

+ 3 graduate students, Artur Gawryszczak (Warsaw), Konstantinos Kifonidis (Munich), Andrzej Odrzywolek (Cracow), Ju Zhang (FIT), Andrey Zhiglo (Kharkov)

1. m461: Stellar Explosions in Three Dimensions

- Summarize your projects and expected scientific objectives through 2014
  - **Modeling and simulations of transient phenomena in stellar astrophysics driven by either radiation or thermonuclear processes**
  - **Numerical solution of a coupled system of PDEs and ODEs**
  - **Tame nonlinearity!**

- Our goal is to …
  - **Explain observed properties of exploding stellar objects**

- Present focus is …
  - **Neutrino-driven core-collapse supernova explosions**

- In the next 3 years we expect to …
  - **Link models to observations**

Astronomy, Astrophysics, HEDP

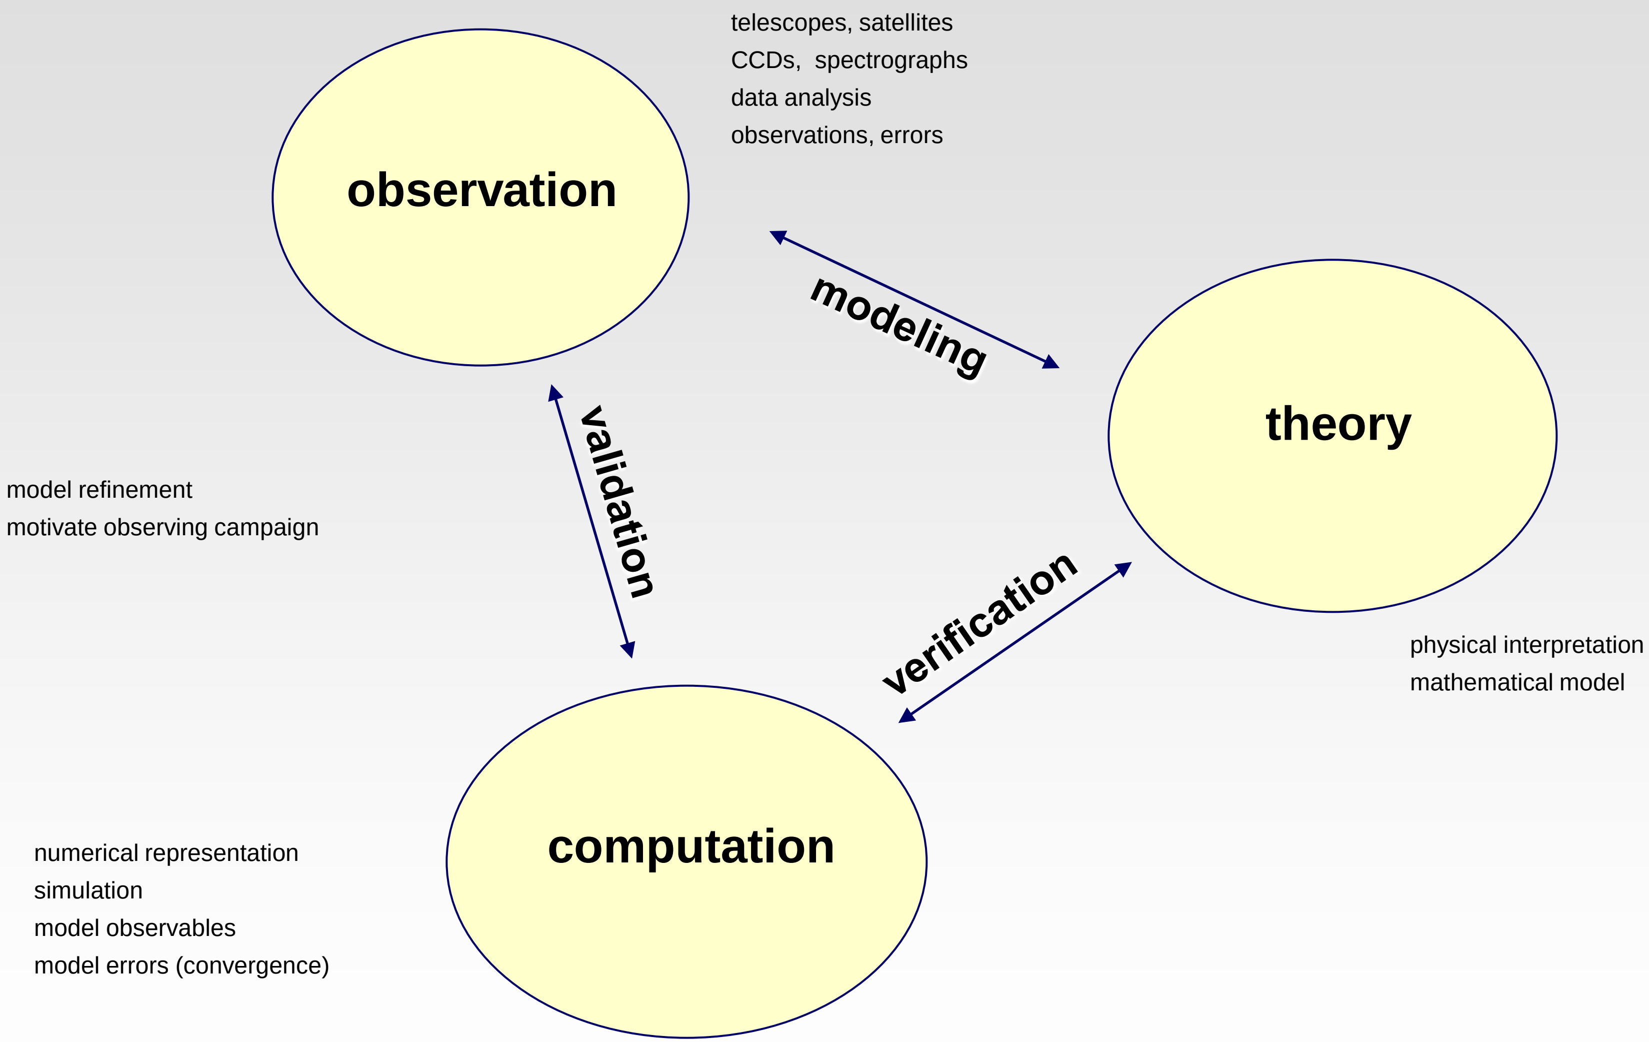

Objects of Interest

- AGB stars
- Classical Novae
- X-ray Bursts
- Thermonuclear Supernovae (Type Ia)
- Core-collapse Supernovae (Type II & Ib/Ic)
- Hypernovae, collapsars, gamma-ray bursts

Astronomer's View

Type II
Massive
Single
H-rich

Neutron Stars
M > 1.4 M
X-ray bursts

White Dwarfs
m < 1.4 Msun
(classical novae)

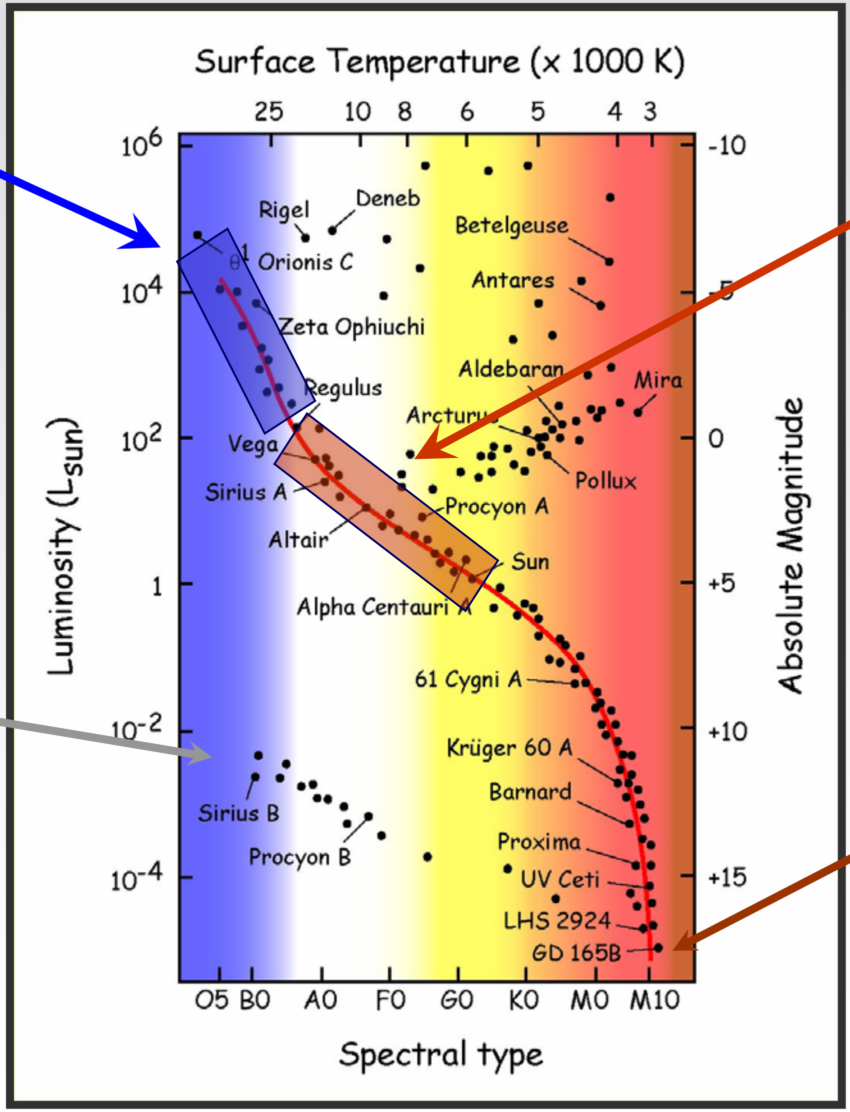

Type Ia
Medium mass
Binary
H/He-free

Brown Dwarfs
m > 0.075 Msun

Astrophysicist's Natural Reaction

2. Current HPC Methods

(see slide notes)

- Algorithms used
  - **Adaptive mesh discretization ($4{,}096^3$ effective meshes)**
  - **Finite volume compressible hydrodynamics**
  - **Multigrid**
  - **Particle tracing (10^6 particles)**

- Codes
  - **FLASH (MPI)**
  - **HOTB (OPenMP)**
  - **Nucnet (MPI or OpenMP)**

- Quantities that affect the problem size or scale of the simulations (grid? particles? basis sets? Other?)
  - **Grid**
  - **Time evolution**

Modeler's View

- A set of PDEs and ODEs

$$\partial_t \mathbf{U} + \nabla \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U})$$

$$\nabla^2 \Phi = 4\pi G \rho$$

- PDEs of every possible type
- ODEs frequently stiff
- Complex equation of state (first closure relation)
- Multidimensional (4D...7D, more closure relations)
- Various discretization methods (finite volume solvers, multigrid, particles, subgrid, front tracking)
- adaptive in space and time
- prone to produce demonstration runs
- ”unlimited” computing resources (“tree barking”)

Transport

- A set of PDEs and ODEs

$$\partial_t \mathbf{U} + \nabla \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U})$$

$$\nabla^2 \Phi = 4\pi G\rho$$

- PDEs of every possible type
- ODEs frequently stiff
- complex equation of state (first closure relation)
- multidimensional (4D…7D, more closure relations)
- various discretization methods (finite volume solvers, multigrid, particles, subgrid, front tracking)
- adaptive in space and time
- prone to produce demonstration runs
- ”unlimited” computing resources (“tree barking”)

Self-Gravity

- A set of PDEs and ODEs

$$\partial_t \mathbf{U} + \nabla \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U})$$

$$\nabla^2 \Phi = 4\pi G \rho$$

- PDEs of every possible type
- ODEs frequently stiff
- complex equation of state (first closure relation)
- multidimensional (4D...7D, more closure relations)
- various discretization methods (finite volume solvers, multigrid, particles, subgrid, front tracking)
- adaptive in space and time
- prone to produce demonstration runs
- ”unlimited” computing resources (“tree barking”)

Nuclear Physics

- A set of PDEs and ODEs

$$\partial_t \mathbf{U} + \nabla \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U})$$

$$\nabla^2 \Phi = 4\pi G \rho$$

- PDEs of every possible type
- ODEs frequently stiff
- complex equation of state (first closure relation)
- multidimensional (4D...7D, more closure relations)
- various discretization methods (finite volume solvers, multigrid, particles, subgrid, front tracking)
- adaptive in space and time
- prone to produce demonstration runs
- ”unlimited” computing resources (“tree barking”)

2. Current HPC Requirements

(see slide notes)

- Architectures currently used
  - **Distributed memory (FLASH) or SMP (HOTB) CPU clusters**
- Compute/memory load
  - **1,000 cores**
  - **1 GB per core**
  - **50,000 CPU hours per run (model)**
  - **10^6 CPU hours per year**
  - **10-20 models**
- Data read/written
  - **0.5/5 TB per model**
  - **20 GB checkpoints**
  - **0.5 TB per model moved out of NERSC, little moved in**
- Necessary software, services or infrastructure
  - **F90, C++, MPI, OpenMP, Python, VisIt, svn/git**
- Known limitations/obstacles/bottlenecks
  - **load imbalance**
  - **memory bandwidth**
  - **non-scalable data structures**
- Hours requested/allocated/used in 2010
  - **1.0/2.5/2.0 million**
- Additional info from the templates

Example #1: Binary WD Merger (DD)

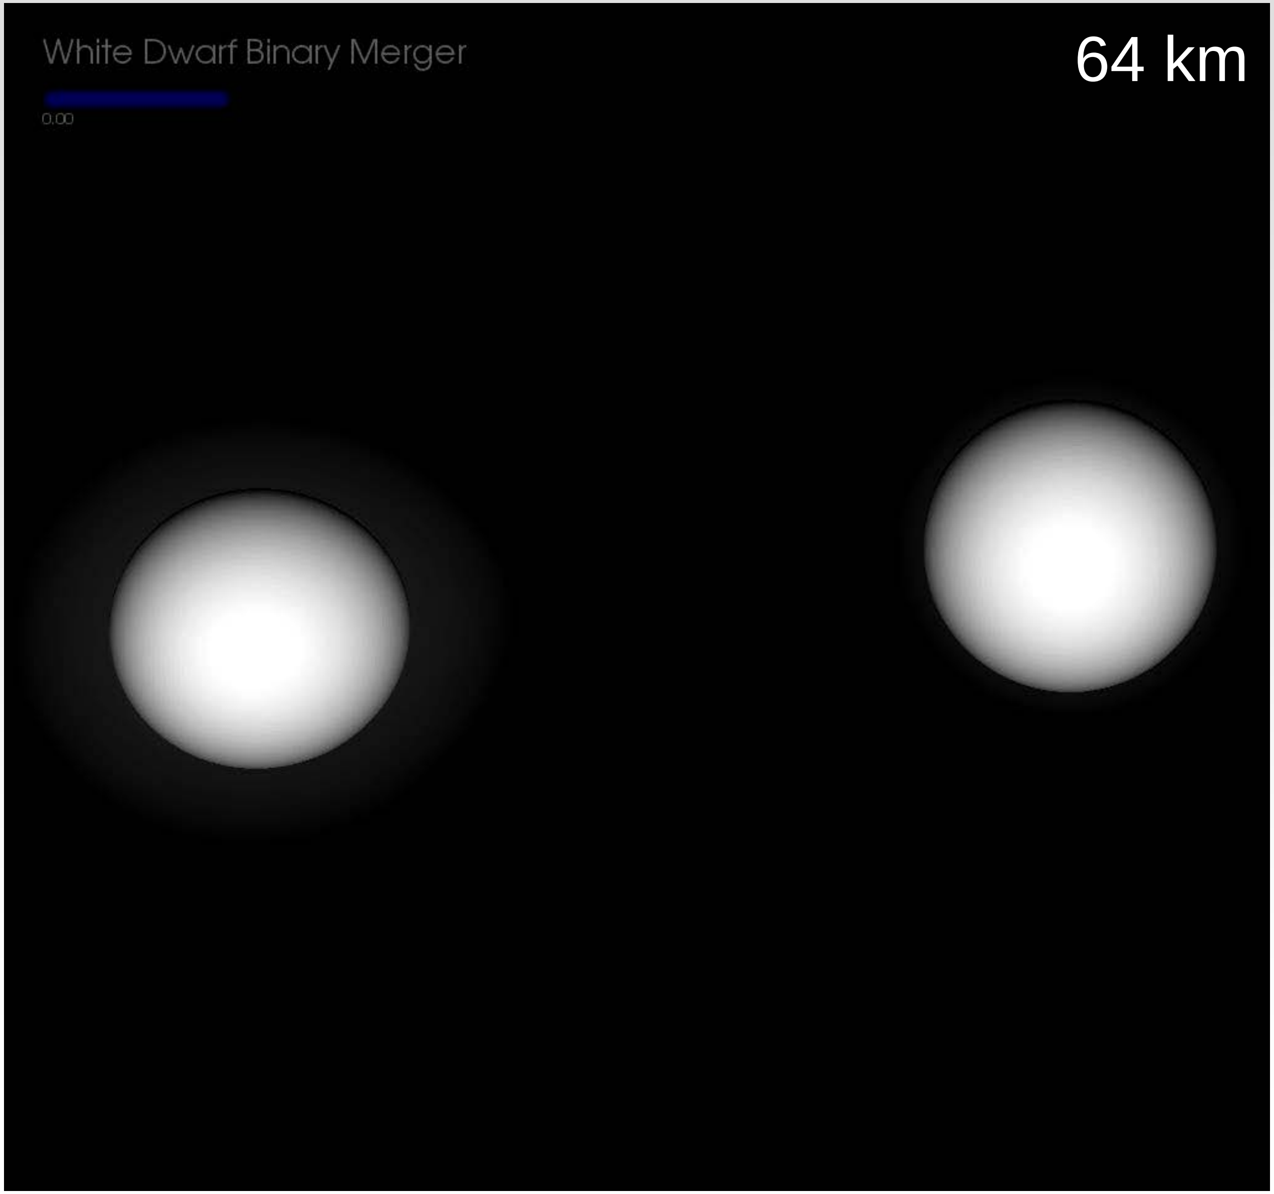

Work done in collaboration with A. Gawryszczak (Copernicus Center, Warsaw). Computing cycles: DOE NERSC.

FLASH/WDM Parallel Performance

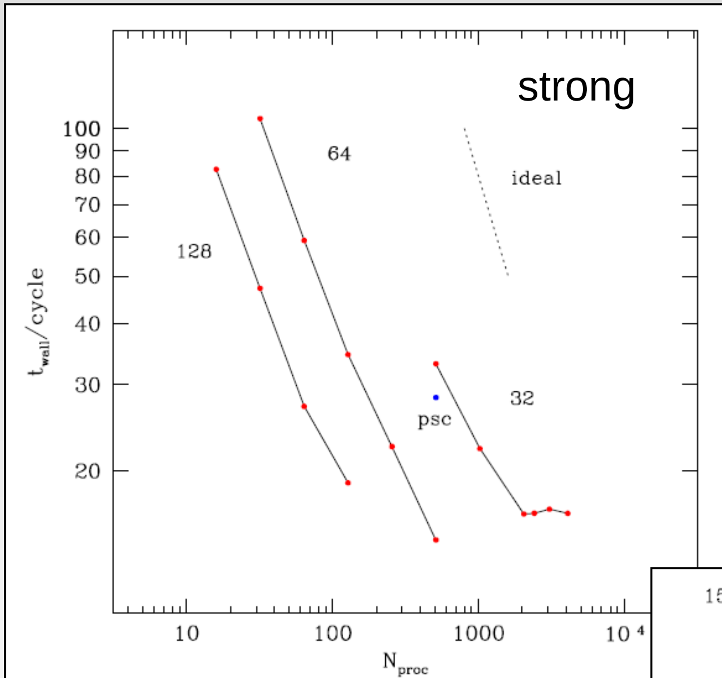

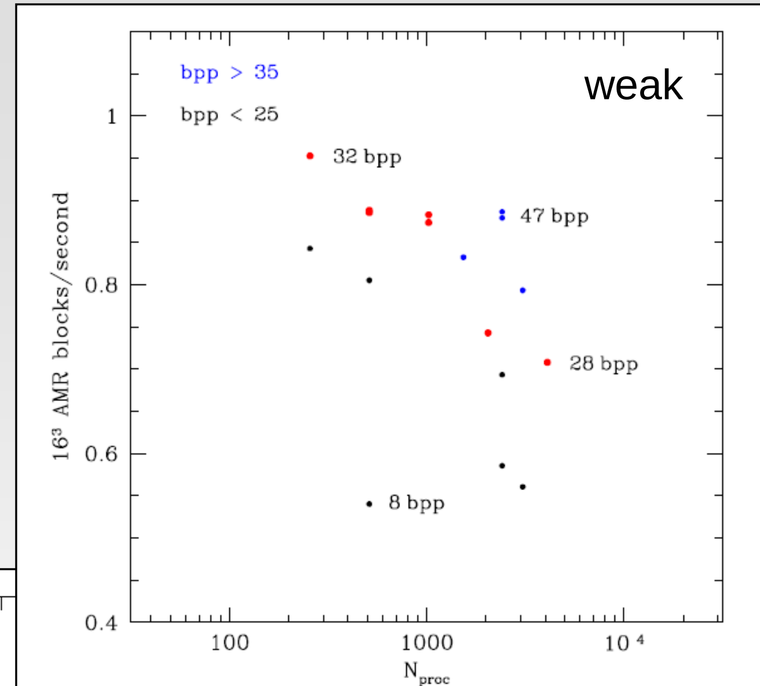

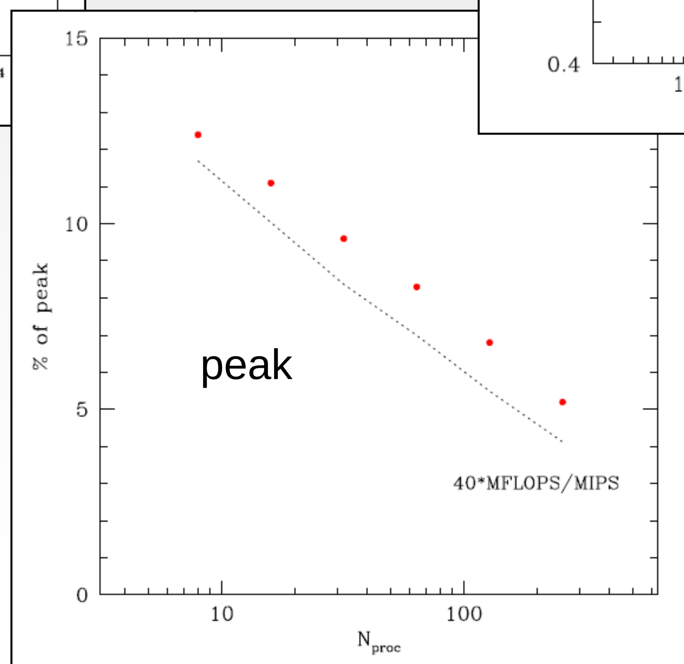

Example #2: Core-Collapse SN Explosions

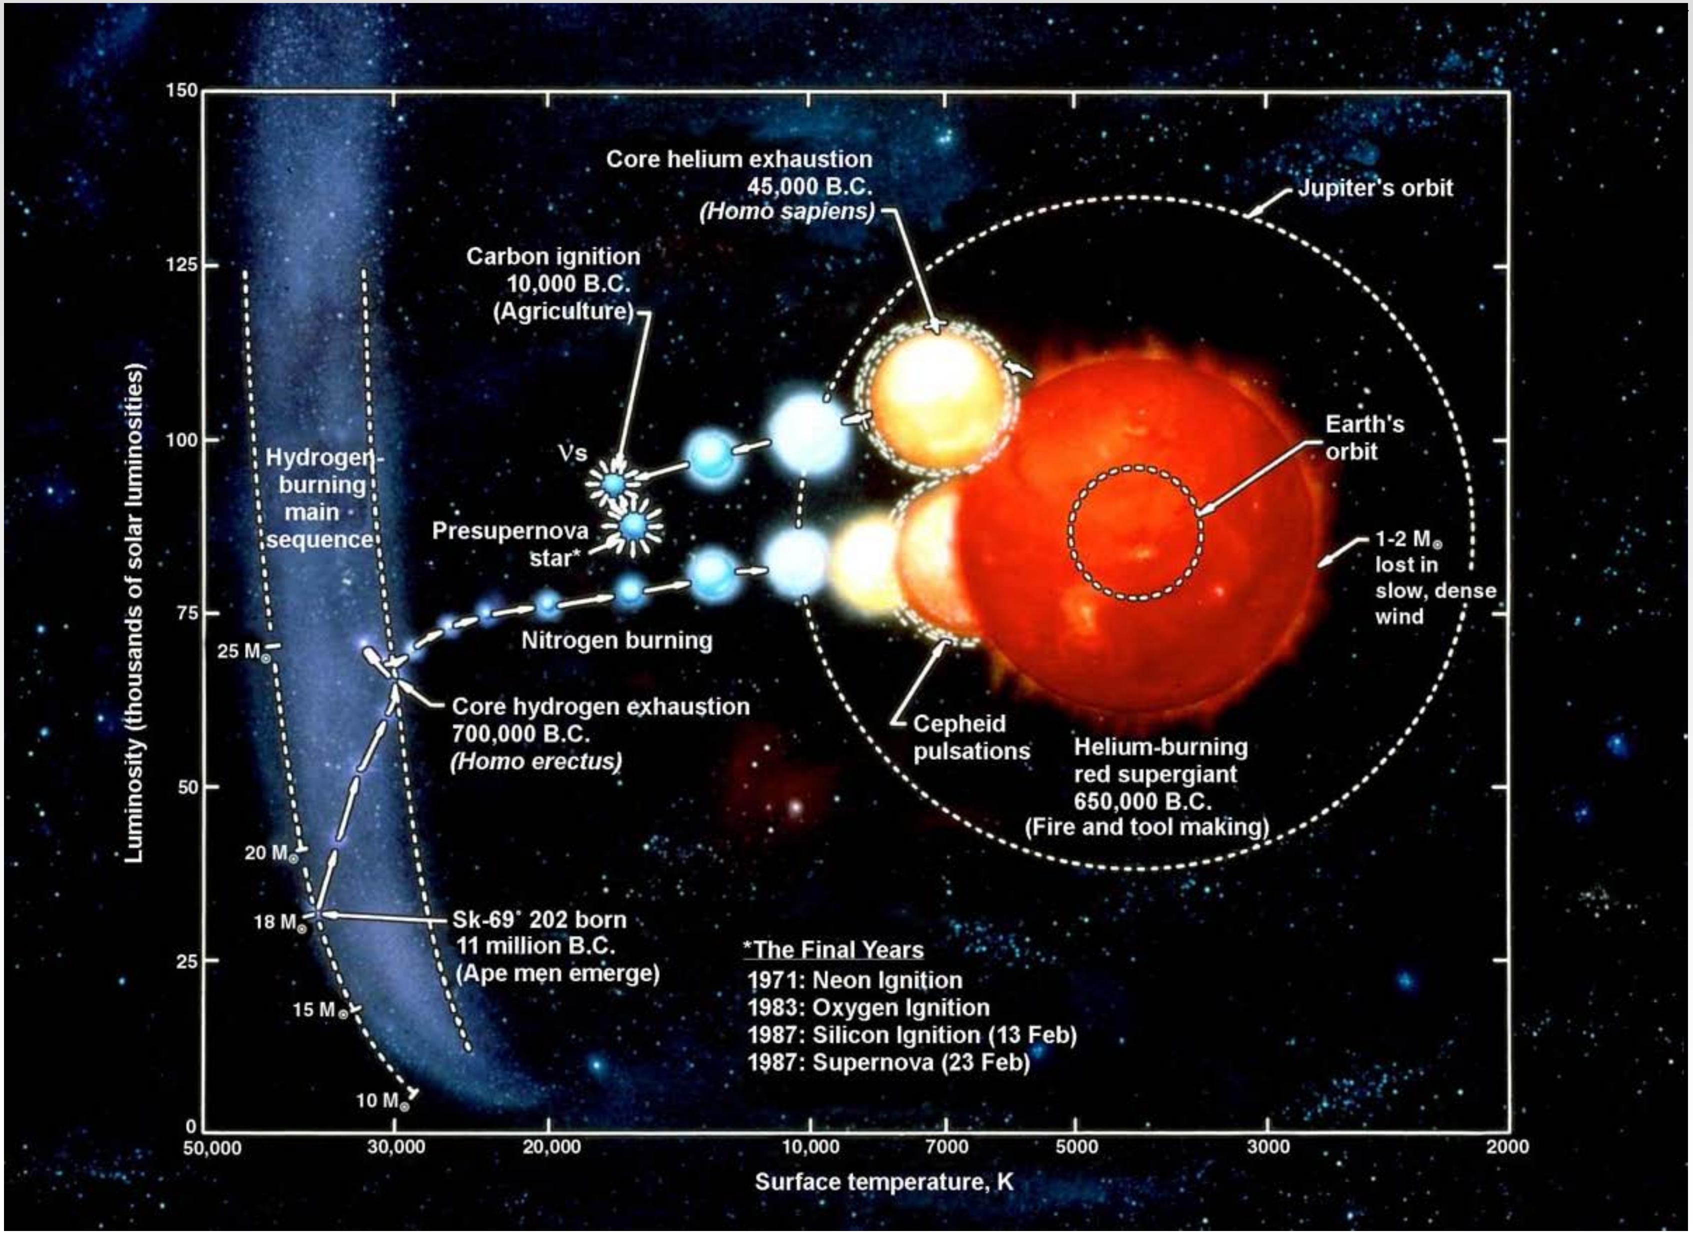

Core-Collapse Shock Revival

- Massive stars
- Gravity bombs with energy extracted by neutrinos
- Accretion shock originally too weak
- Revived by neutrino heating of the post-shock matter, a.k.a. Standing Accretion Shock Instability

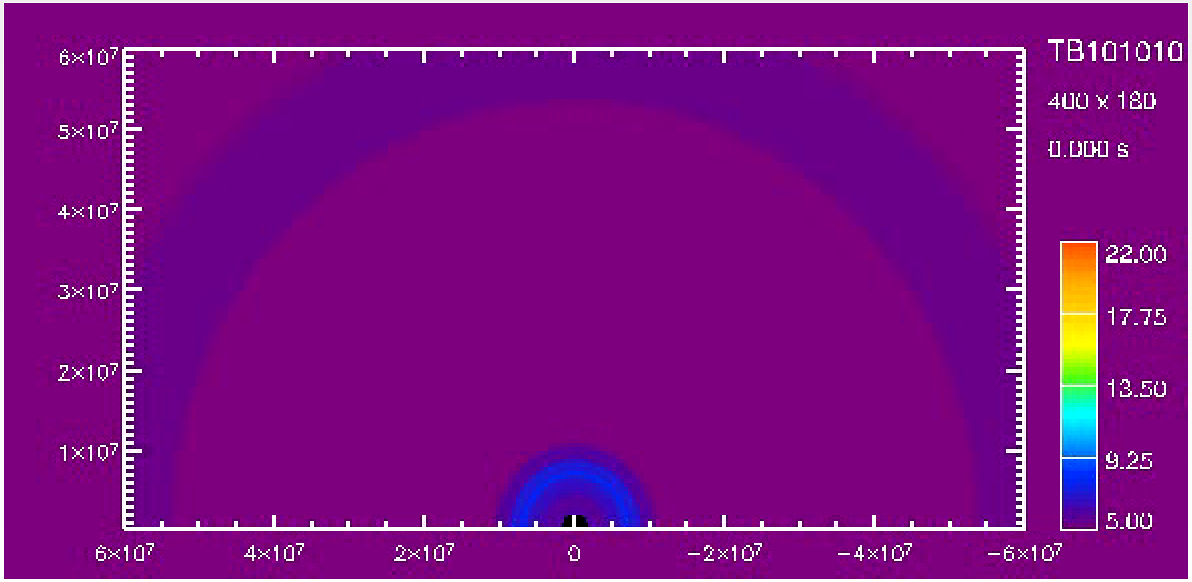

SASI in 3D

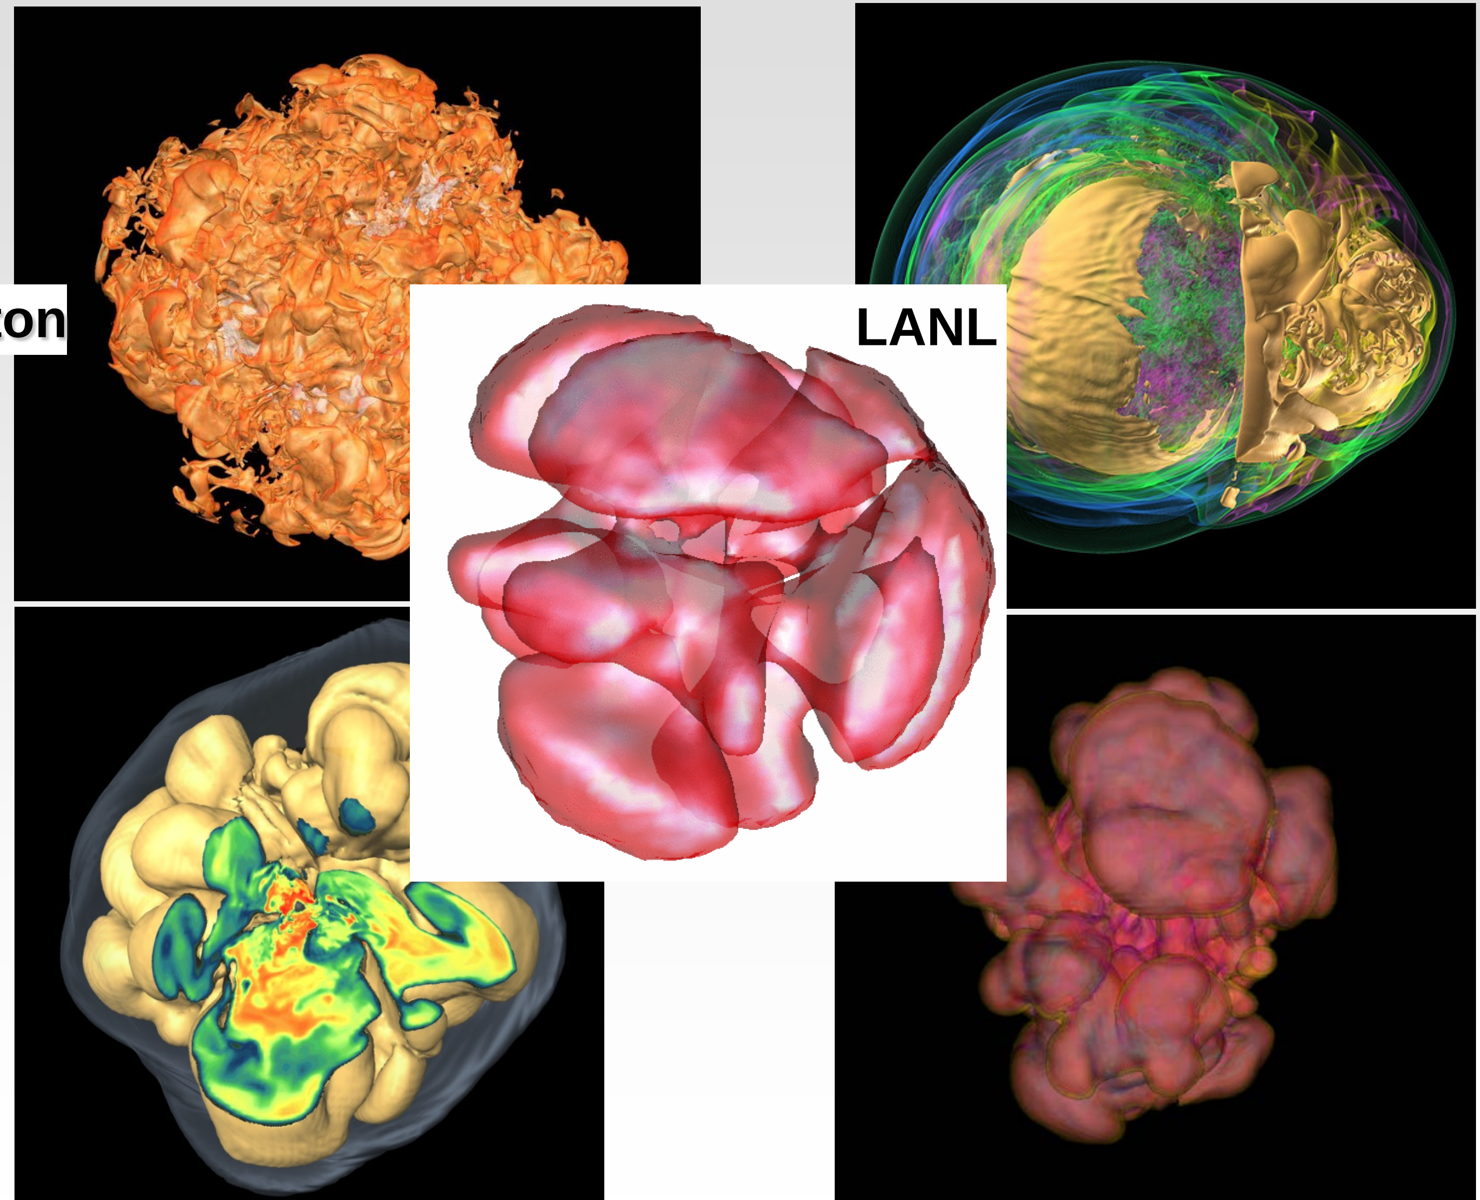

Spring 2010: Nordhaus & Burrows

- 3-D aids explosions compared to 2-D
- This is in essence an extension of what has been found by Janka & Mueller re 1-D to 2-D

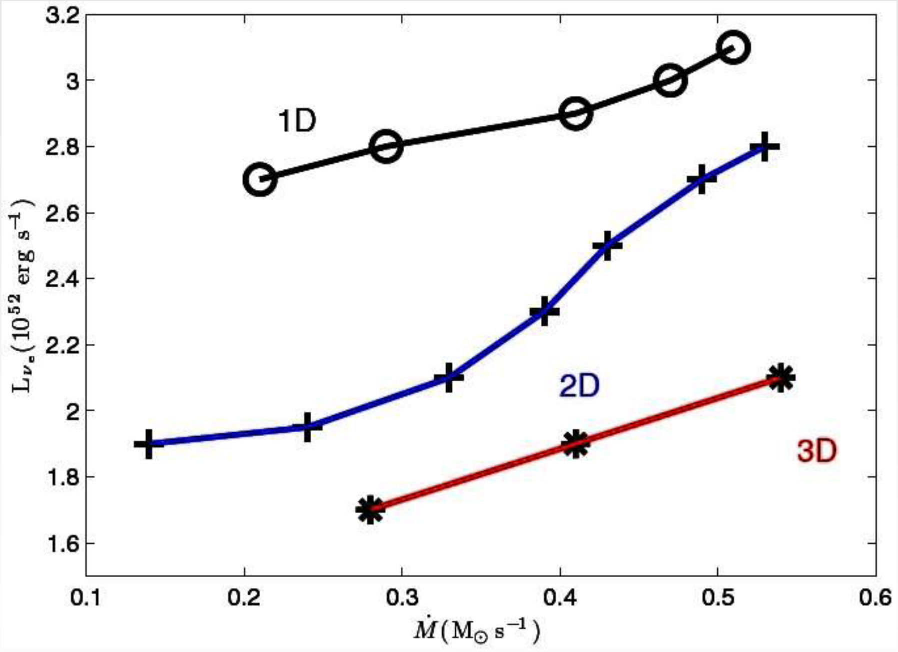

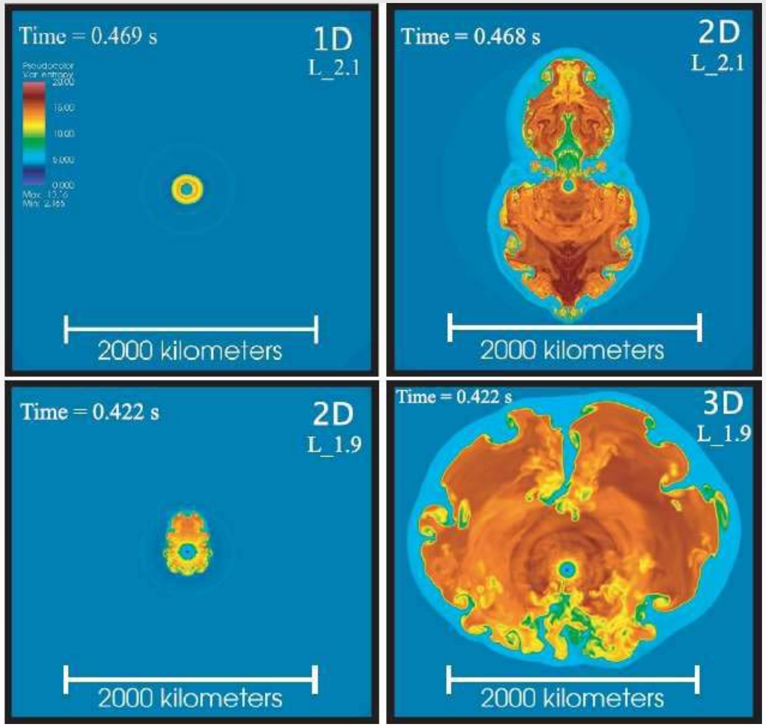

Fall 2010: Janka & MPA group

- Janka et al. fail to confirm Nordhaus & Burrows result

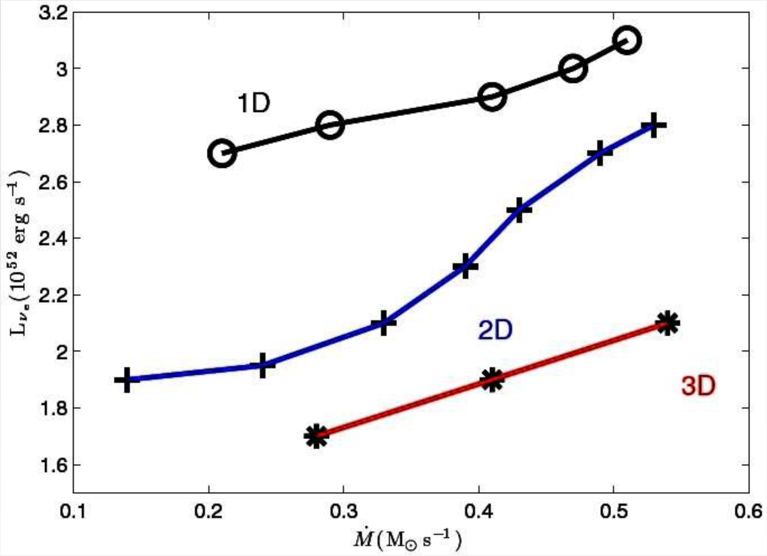

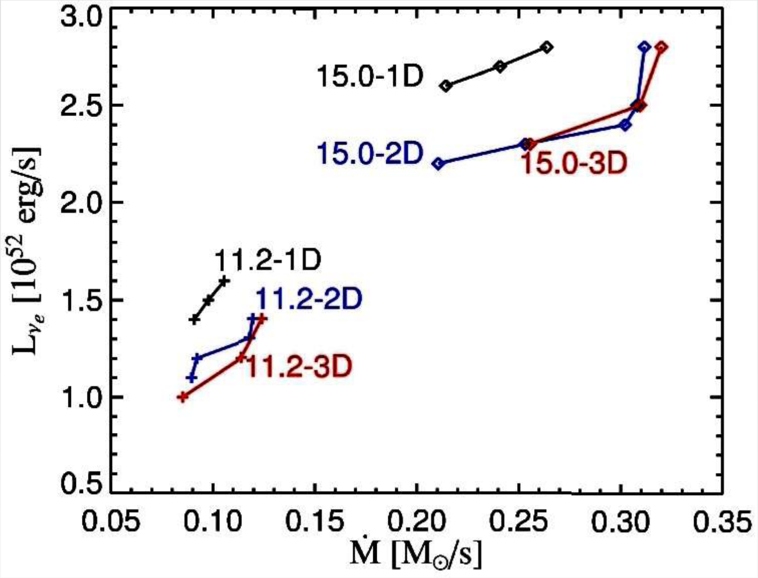

- This does not increase our confidence in numerical modeling!
- Increasing numerical resolution will NOT resolve the above discrepancy

The Challenge Continues

- Theory incomplete ☹
- Simulations unsuccessful ☹
- Experiments limited ☹

Example #3: Postexplosion Mixing

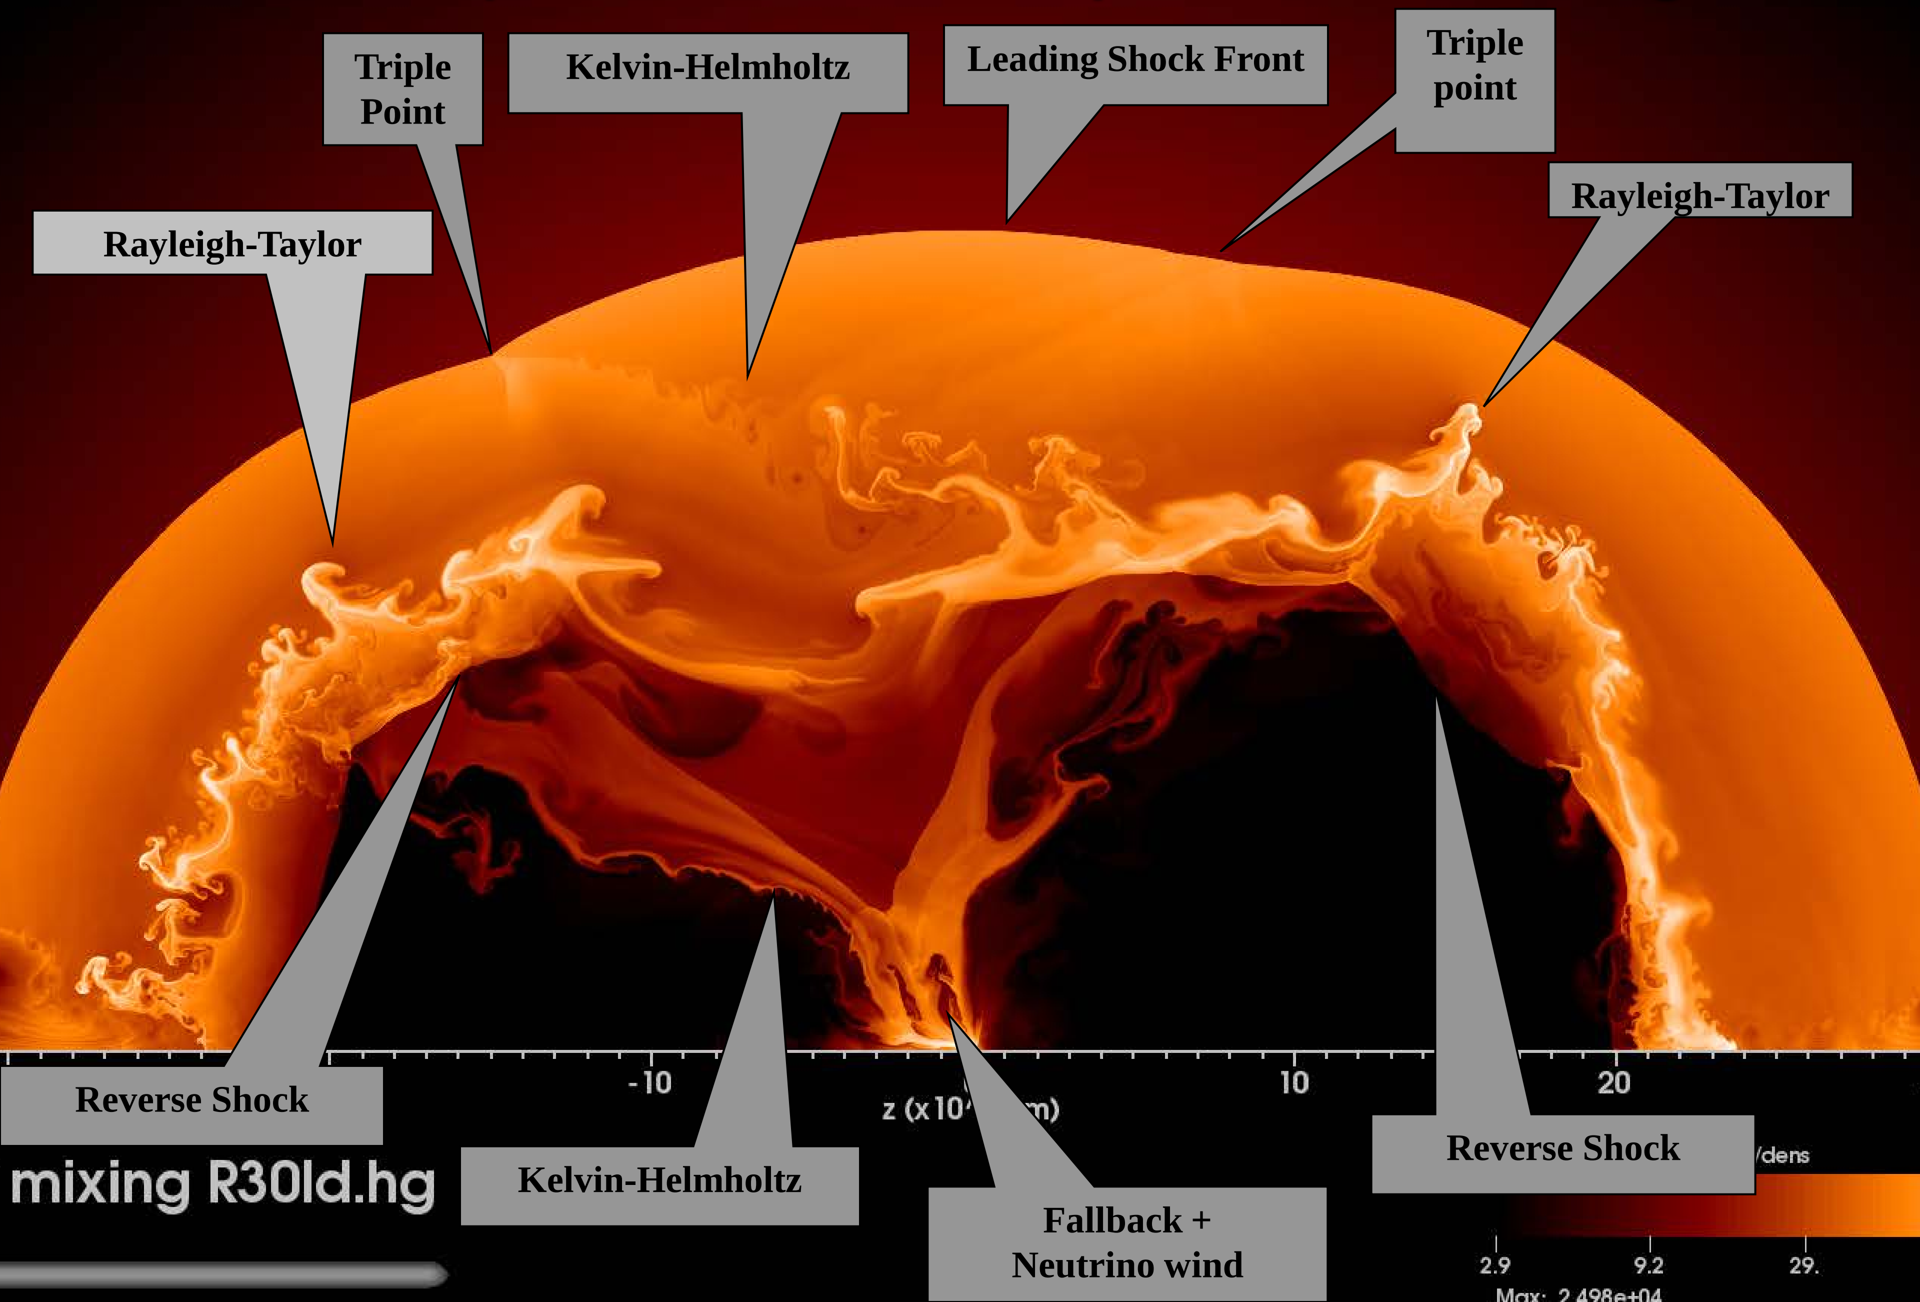

Non-SN NP Examples

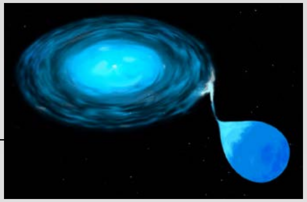

Thermonuclear runaways, but now in thin layers of material accreted by a degenerate star from a non-degenerate companion star (yes, most stars are binaries!)

- Degenerate star is a neutron star: X-ray burst
  - strong gravity
  - strongly degenerate matter
  - ignition/propagation unknown
- Degenerate star is a white dwarf: Classical Nova
  - weaker gravity
  - moderately degenerate
  - source of mixing (dredge up) unknown

NP Example: Classical Nova Runaway

^{14}O

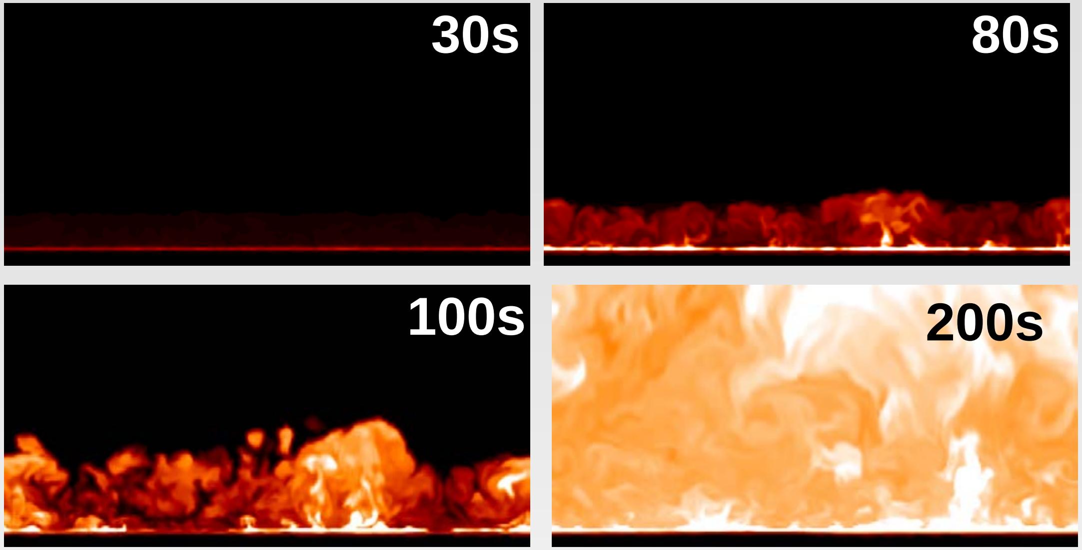

Kercek et al. (1999)

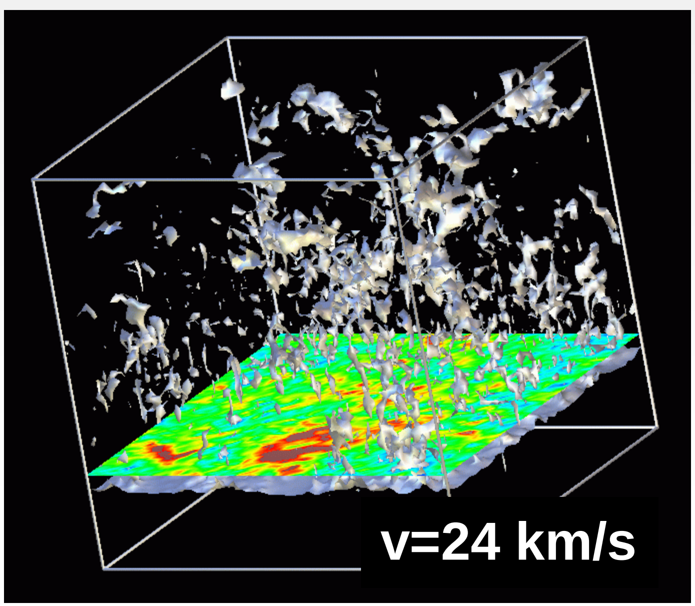

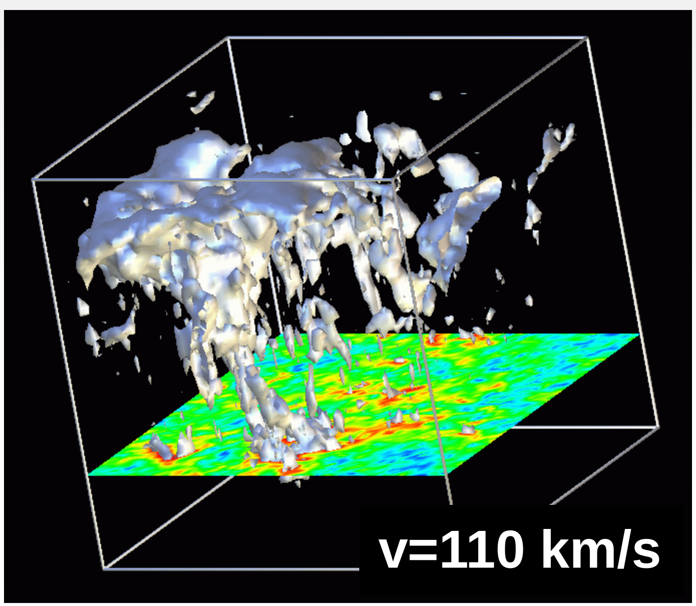

NP Challenge: Hot CNO cycle

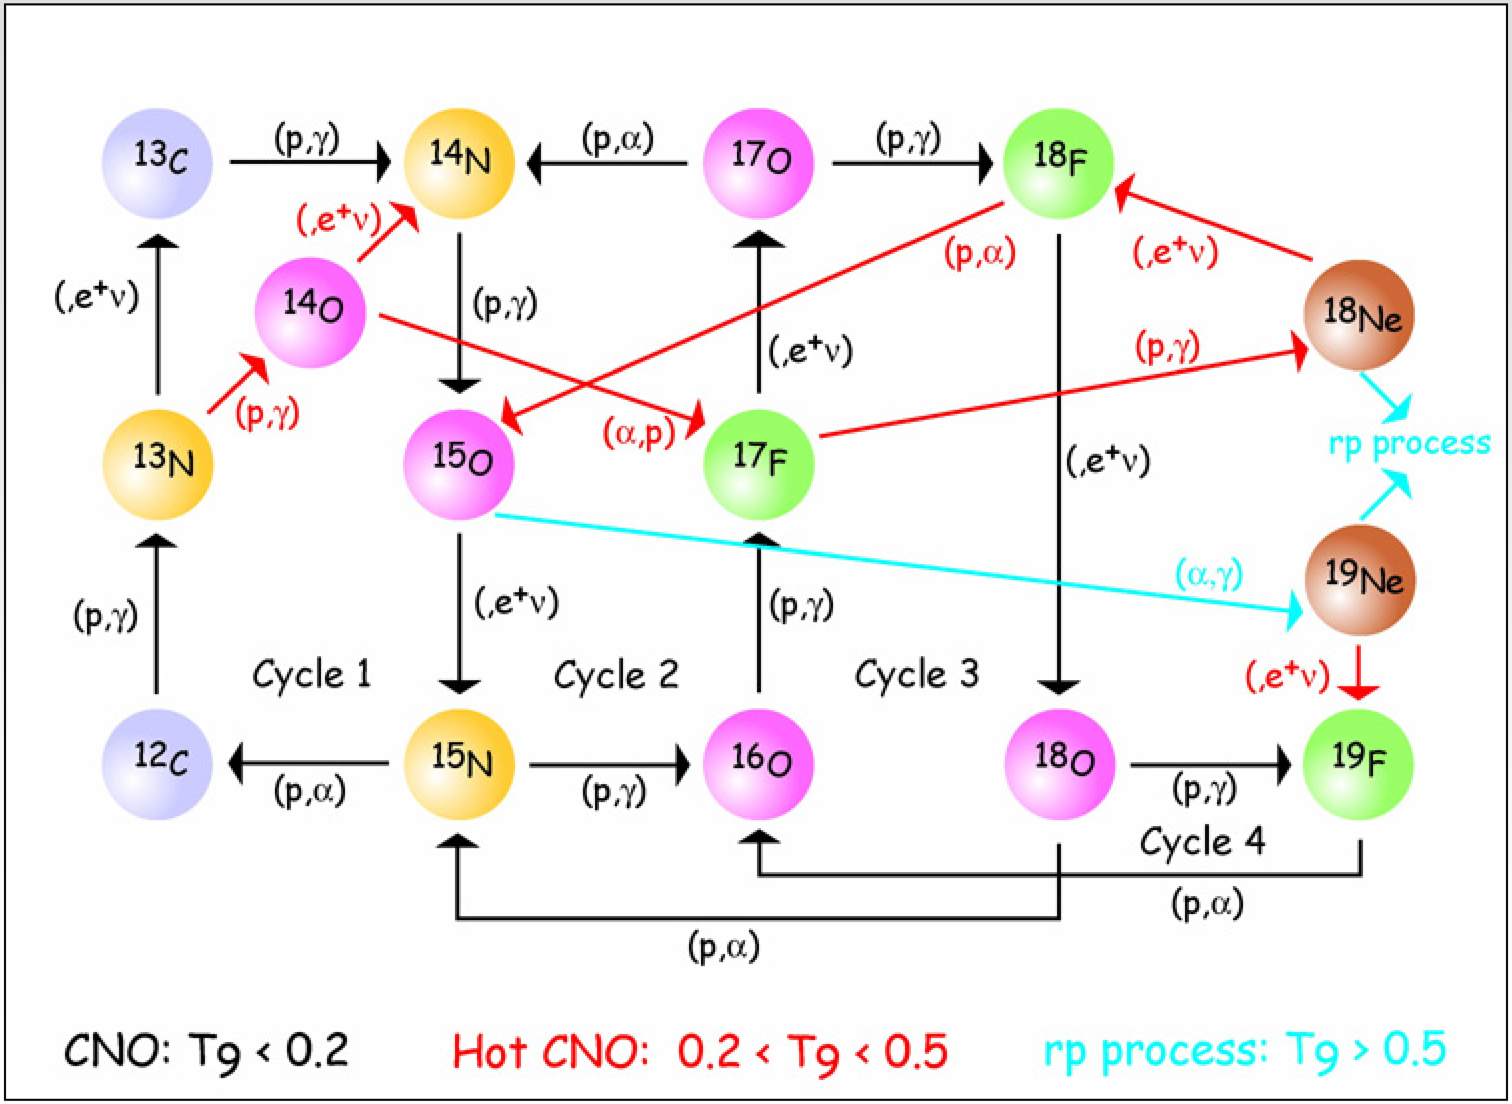

Wallace & Woosley (1981)
graph: F. X. Timmes

Next 5 Years

1. **Hot CNO cycle in classical novae**
   - T~300x10^{6}K (or ~27 keV)
   - rho ~150 g/cc
   - 30% uncertainty for some of the reaction rates

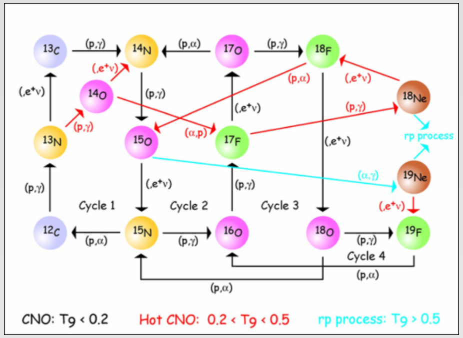

2. **^{12}C+^{12}C ignition in binary WD (DD SN Ia)**
   - T~2x10^{9}K (~180 keV)
   - rho ~1x10^6 g/cc
   - admixture of ^{4}He

3. **Parametrized core-collapse SN explosions**
   - energetics
   - mixing and asymmetries
   - observable imprints of the SN engine

Next 10 Years

1. **Paths to explosions**
   - coupled physics (turbulence, mixing, diffusion)
   - long-term evolution
   - subgrid scale models

2. **Paths in explosions**
   - coupled physics (radiation-matter interactions)
   - nucleosynthesis in the neutrino-mediated plasmas
   - radiation sources
   - Connecting scales (DNS -> LES)

3. **Paths past explosions**
   - long-term evolution
   - non-LTE physics
   - from discovery to predictions
   - sensitivities? adjoints?

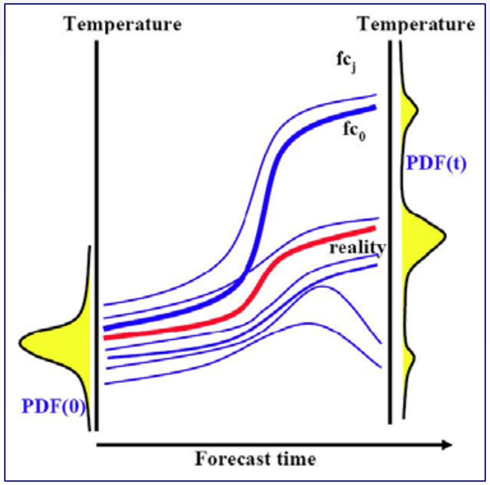

Steinheimer et al. (2010)

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

(see slide notes)

- Upcoming changes to codes/methods/approaches to satisfy science goals
  - **Approximate neutrino transport**
  - **Scalable nonlinear multigrid**
- Changes to Compute/memory load
  - **Compute x20**
  - **Memory x10**
- Changes to Data read/written
  - **Data x10**
- Changes to necessary software, services or infrastructure
  - **GPGPU tools?**
- Anticipated limitations/obstacles/bottlenecks on 10K-1000K PE system.
  - **Load imbalance (both due to discretization and physics)**
  - **Data locality**
  - **Scalability of global data structures**
- Key point is to directly link upcoming NERSC requirements to science goals

Strategy for New Architectures

- How are you dealing with, or planning to deal with, many-core systems that have dozens or hundreds of computational cores per node?
  - **MPI (OpenMP proved inefficient/less general)**
- How are you dealing with, or planning to deal with, systems that have a traditional processor augmented by some sort of accelerator such as a GPU or FPGA or similar?
  - **Planned**
    - **Preprocessor directives**
    - **Kernel extraction and analysis**
    - **Consider SC (CASTRO) or NSF (Athena) codes**

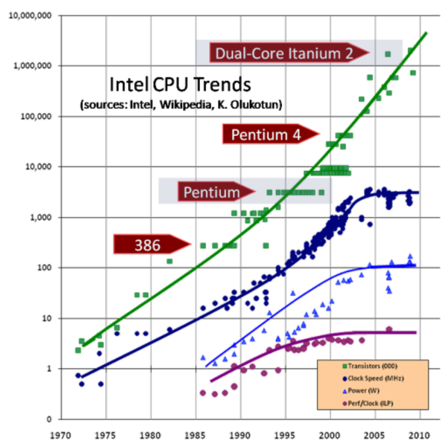

Sutter (2009)

4. Summary

- What new science results might be afforded by improvements in NERSC computing hardware, software and services?
  - **Physics-enhanced models (i.e. MHD)**
  - **Physics-coupling studies (in time, in space)**
  - **Model sensitivities**
- Recommendations on NERSC architecture, system configuration and the associated service requirements needed for your science
  - **Continue providing access to large memory per core systems**
  - **Actively support transition to limited memory per core architectures**
  - **Increase shared data space (NGFS, /project)**
  - **Provide users with control over scratch purging process**
- NERSC generally acquires systems with roughly 10X performance every three years. What significant scientific progress could you achieve over the next 3 years with access to 50X NERSC resources?
  - **Systematic studies of supernova explosion dependence on the progenitor structure**
  - **=> Progenitor structure studies!**
  - **Model databases for sensitivity studies**
- What "expanded HPC resources" are important for your project?
  - **Data storage (both runtime and archival)**
  - **Data analysis and visualization**
  - **Remote connectivity (both for raw data transfers and interactive/X applications)**
- General discussion