Convection in X-ray Bursts

Michael Zingale
Stony Brook University

in collaboration with
Ann Almgren, John Bell, Andy Nonaka (LBL); Chris Malone (LANL),
Stan Woosley (UCSC)

Supported by DOE/Office of Nuclear Physics, DE-FG02-06ER41448 and DE-FG02-87ER40317 to Stony Brook, and NSF award AST-1211563.
Computer time: National Energy Research Scientific Computing Center (Office of Science, DOE DE-AC02-05CH11231)

Convection in Astrophysics

- Evolution of many stellar systems dominated by convective transport of energy
  - 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
- Often the convection is highly subsonic
  - Challenging for traditional astrophysical hydrodynamics codes
- *New algorithms are needed for efficient simulation of convective astrophysical flows*

Multidimensional Simulations

- Nature is 3-d
  - Convection driven by nuclear energy release
  - Fluid instabilities / turbulence
  - Localized burning/runaway
  - Rotation
- Challenging simulations
  - Large computing / memory requirements
  - Making sense of enormous amounts of data
- SNe Ia and XRBs begin with periods of low speed convection driven by nuclear energy release
  - Requires ability to model the domain for long timescales
- Requires a different algorithmic approach than those traditionally used in astrophysics

Simulating Low Mach Phenomena

- 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.

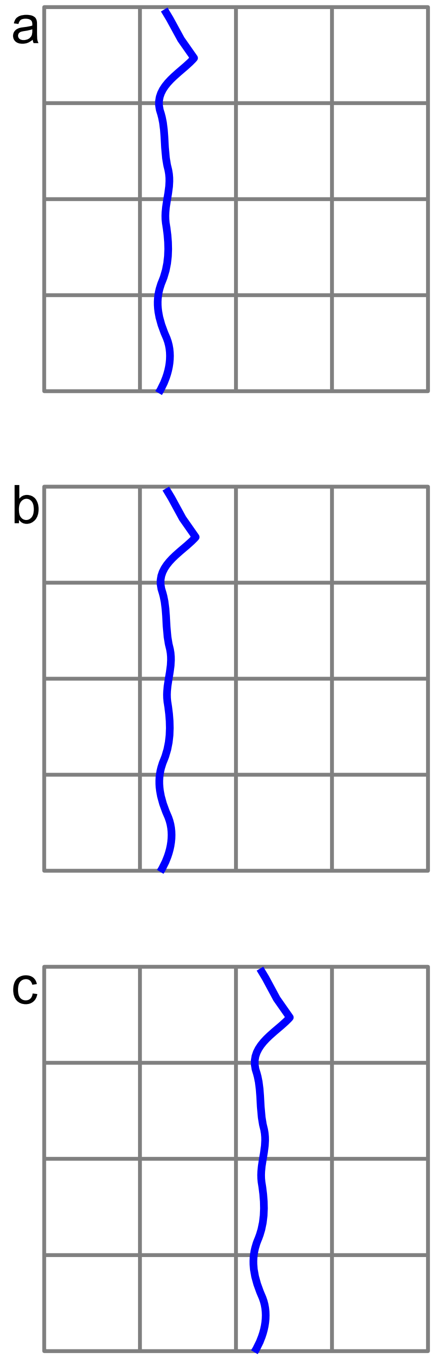

Maestro: Low Mach Number Hydro

- 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
  - Analytically enforce hydrostatic equilibrium through base state:

$$\nabla p_0 = \rho_0 g$$

- 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

Maestro: Low Mach Number Hydro

- Solved via a fraction step method:
  - Advection terms handled with an unsplit Godunov method
  - Diffusion (if used) via an implicit solve with multigrid
  - Projection enforces constraint, solved via multigrid (two solvers: cell-centered and nodal)
  - Reactions via Strang-splitting (local implicit ODE integration)
  - Overall second-order in space and time
- Supports a general equation of state
  - Includes some recent ideas on energy conservation in low Mach systems with general equations of state
- Supports arbitrary reaction networks
  - Multiple species advected
  - New coupling mode (SDC) underway
- Lagrangian tracer particles
- Weak scaling to $O(10^5)$ processors
  - MPI + OpenMP hybrid approach to parallelism via BoxLib

Maestro Applications

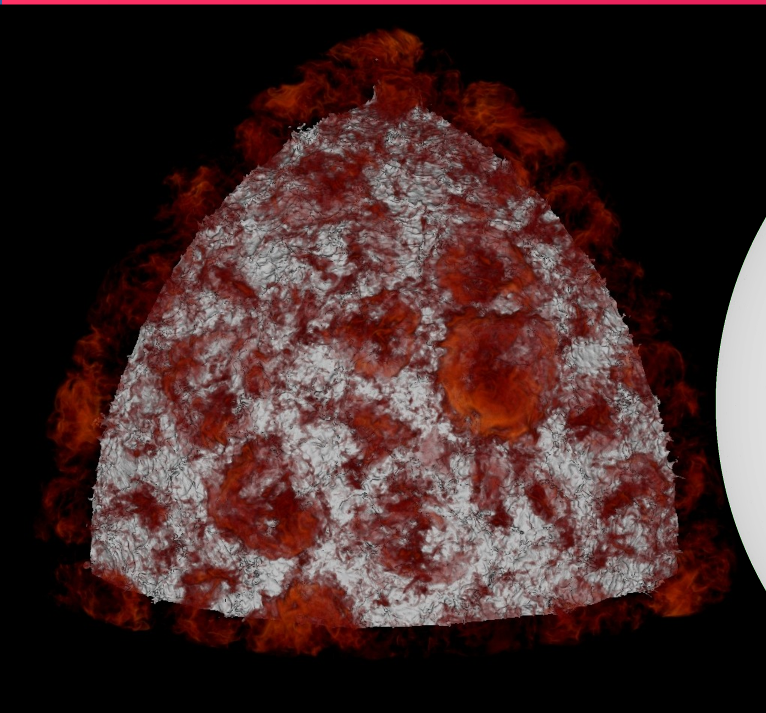

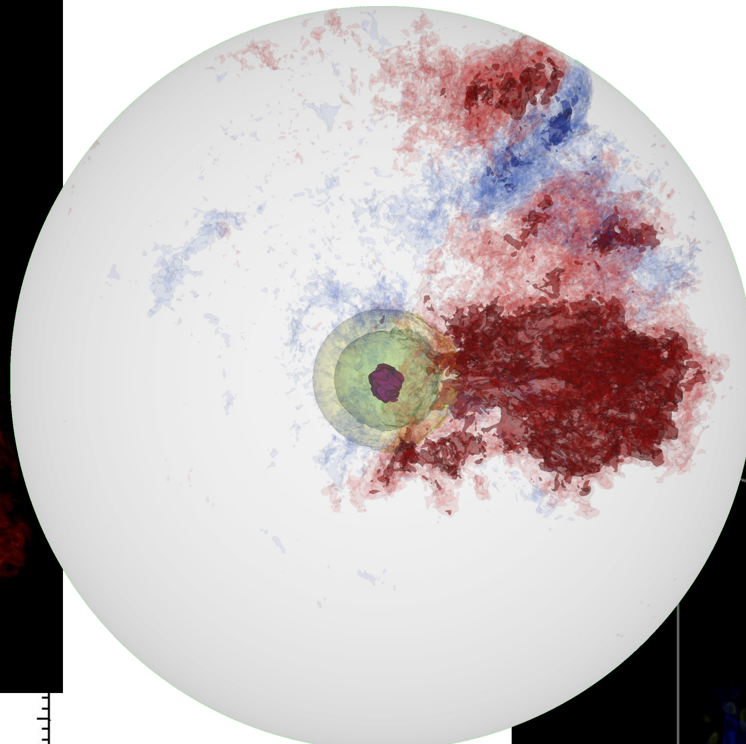

Clockwise from far left: He shell convection in a sub-Chandra SNe Ia model, convection in the Chandra-mass SNe Ia model, convection in a mixed H/He XRB, H core convection in massive stars

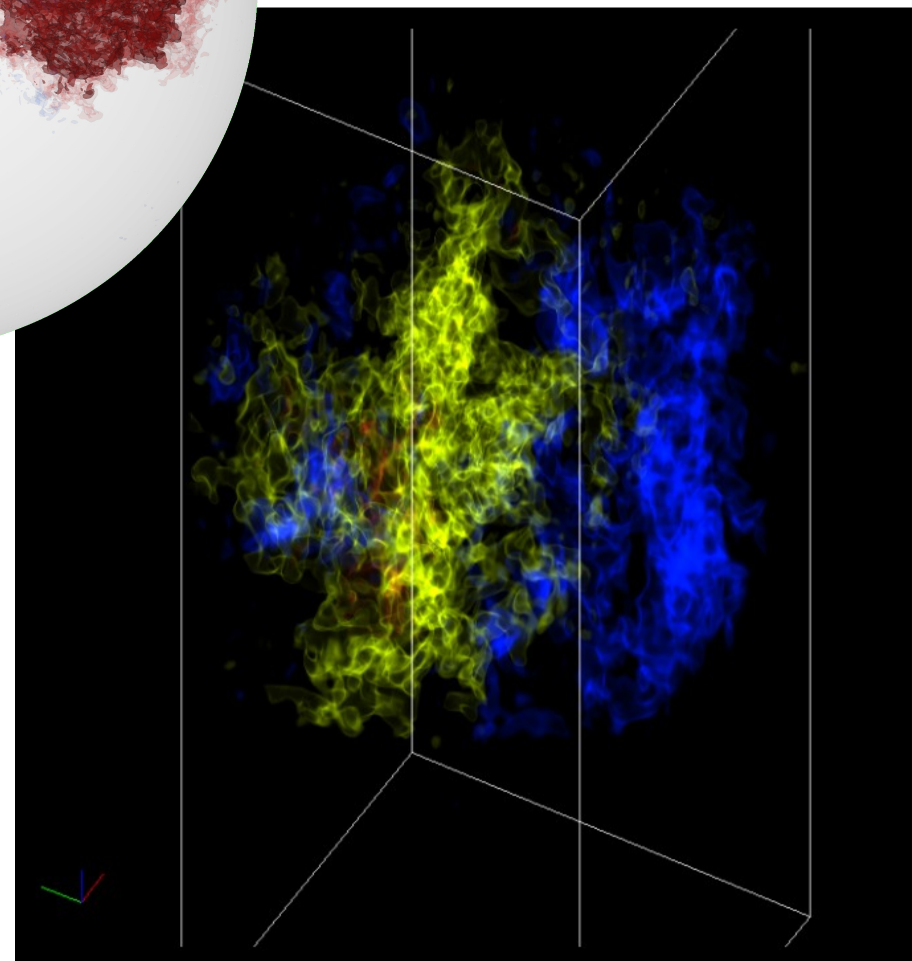

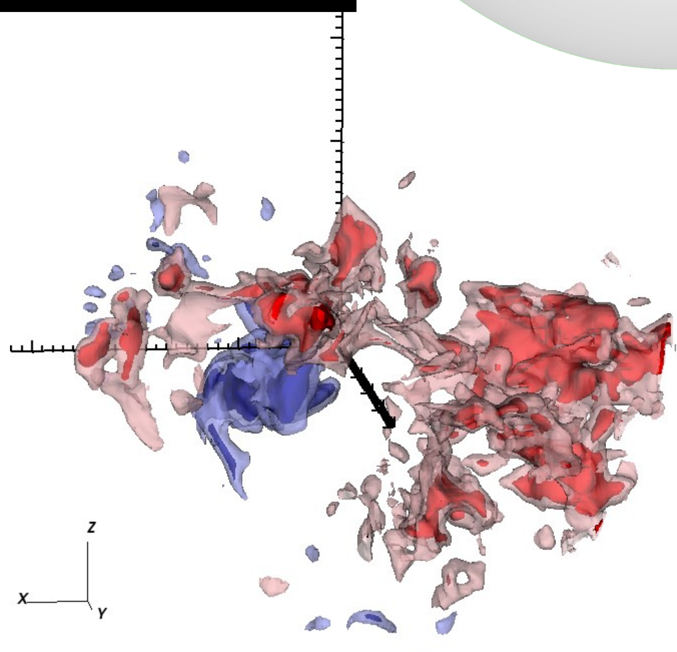

Community Support

- Maestro is a large code
  - Publicly available at: http://bender.astro.sunysb.edu/Maestro/
  - > 250 page users guide
  - E-mail support list
- Engage community members to run new applications. We provide the support
- Potential applications in astrophysics include:
  - Classical novae
  - URCA process in white dwarfs
  - Proto-neutron star cooling
  - Core convection in massive stars
  - Shell burning in evolved stars
  - Multidimensional core-collapse SNe progenitor models (recently started)
  - Convection in exoplanetary interiors

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

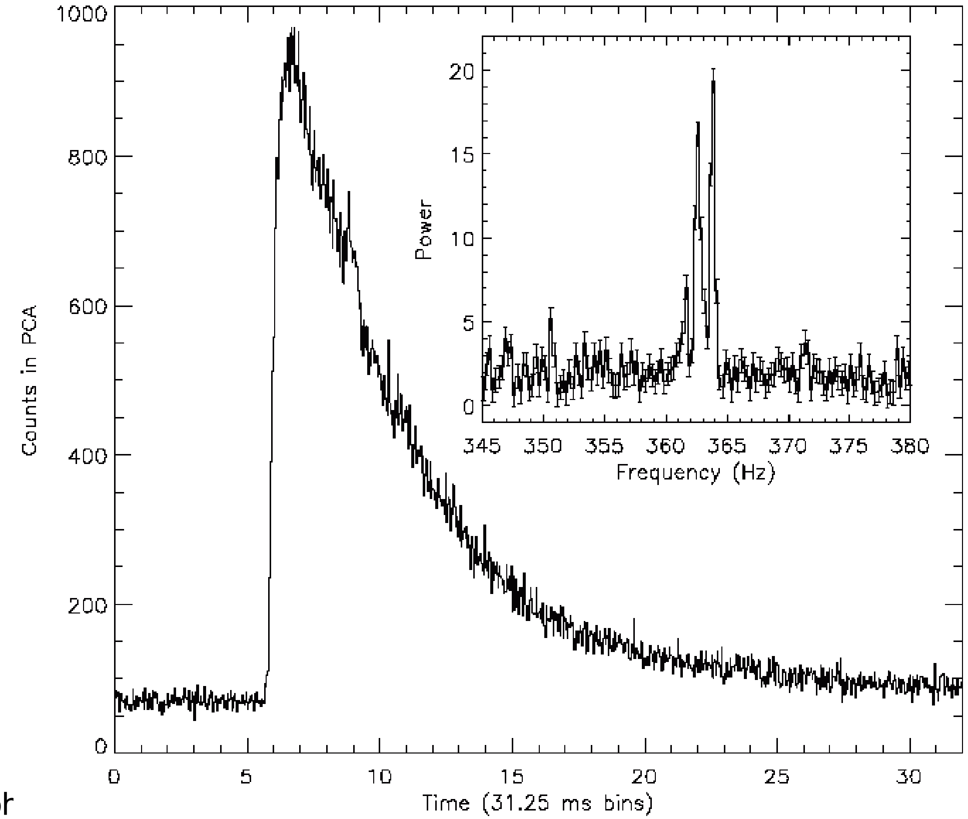

Strol

Outstanding Questions for XRBs

- 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?
- How do we use observations to infer properties of the nuclear equation of state?

These are all multi-dimensional effects

Modeling Convective Burning in XRBs

- Suite of 2-d models of mixed H/He bursts
  - Relaxed resolution requirements compared to He burning
  - 11 nuclei rp network
  - Large temperature increase, but nonlinear runaway not yet seen
- 3-d simulations running at NERSC right now.

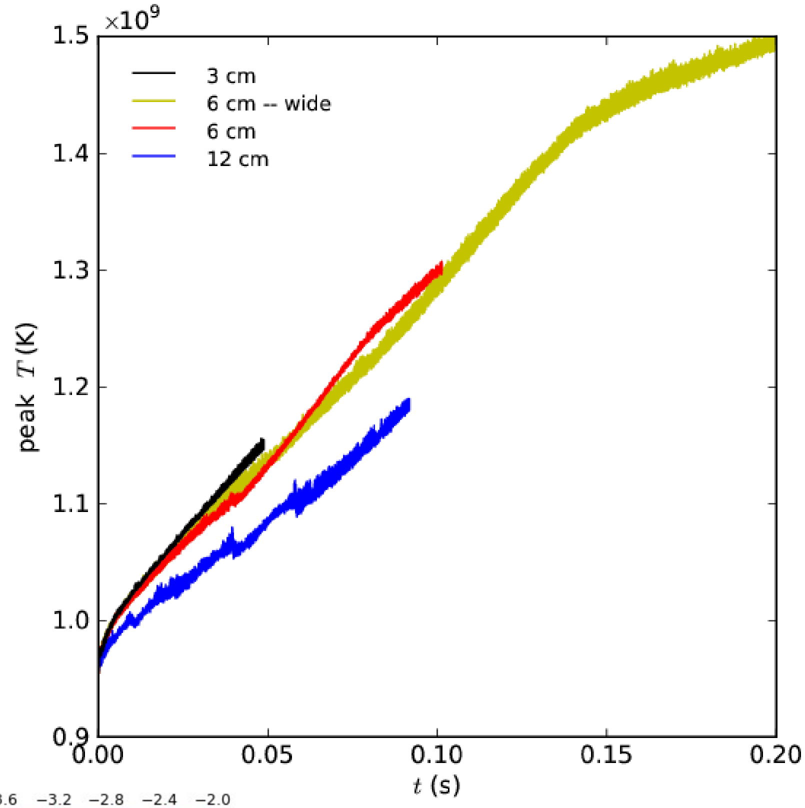

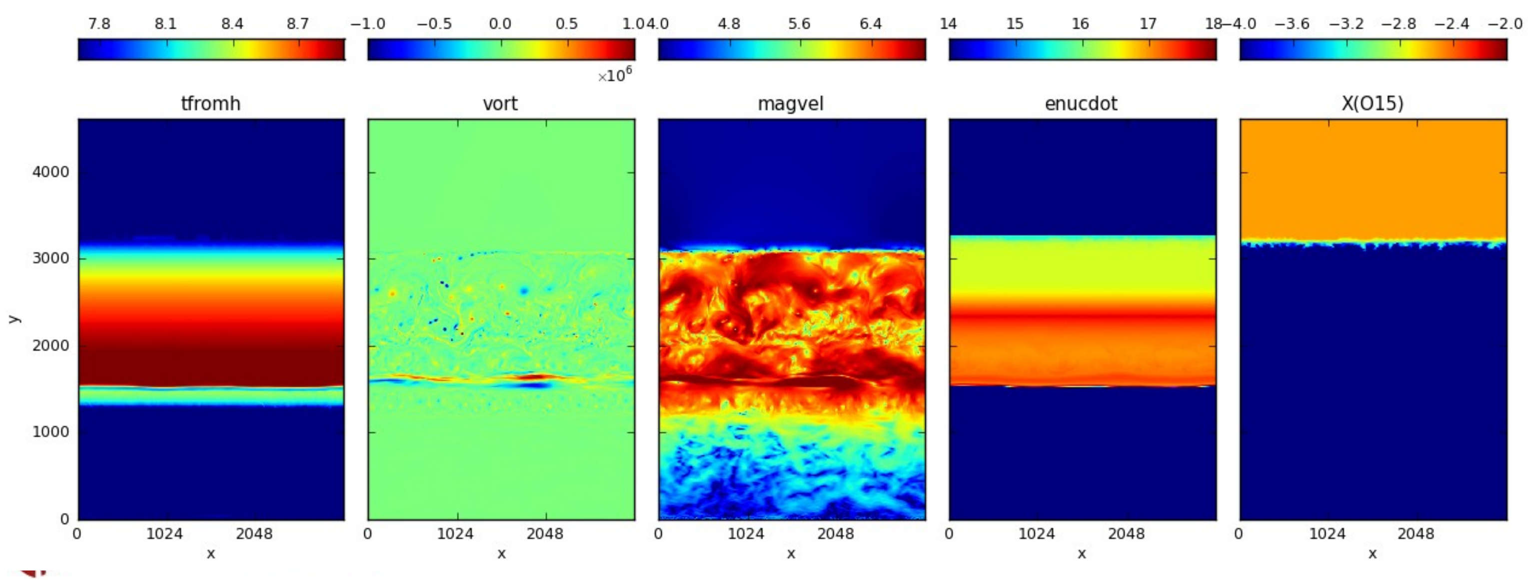

Modeling Convective Burning in XRBs

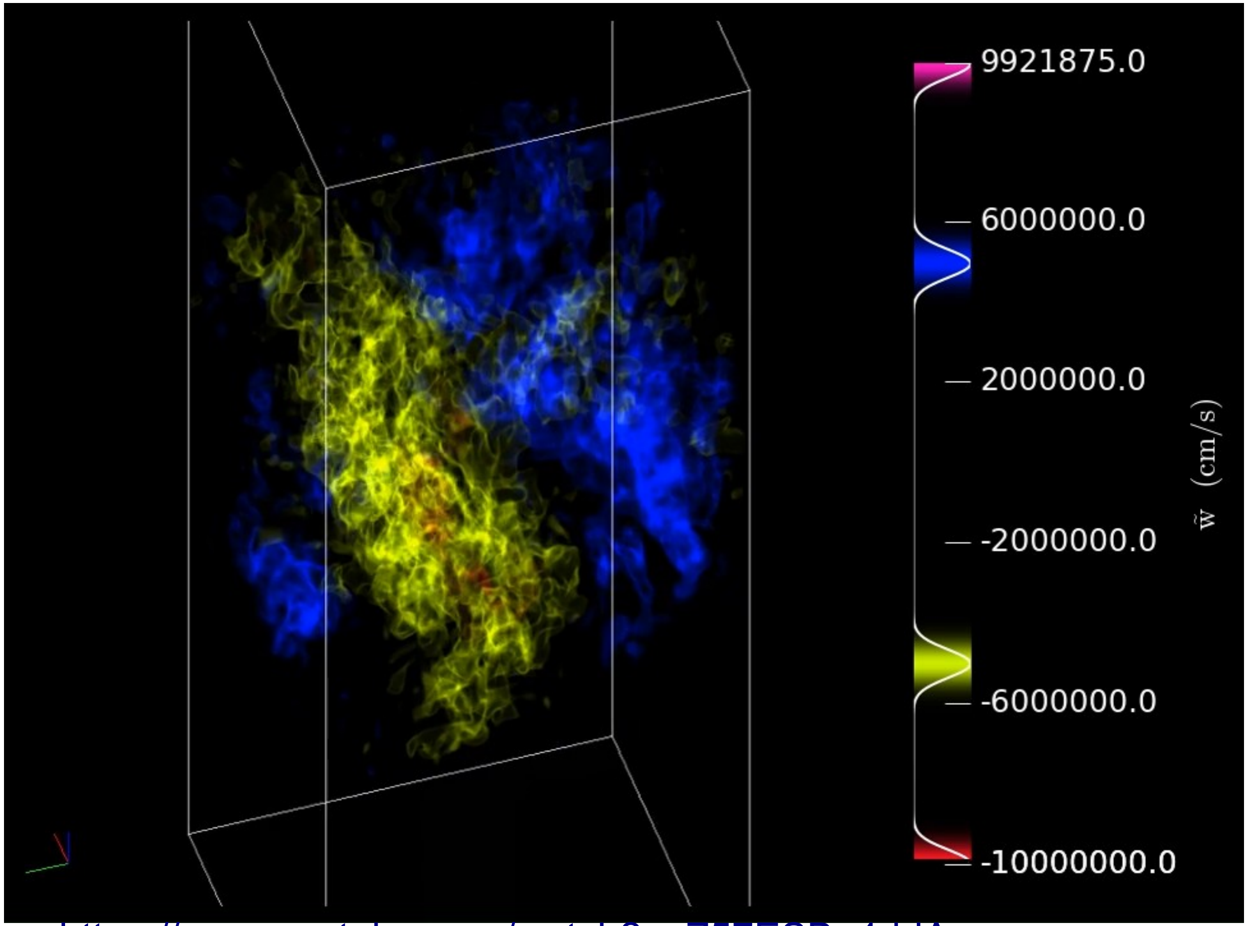

https://www.youtube.com/watch?v=Z57ESBa4ddA

- Initial 3-d simulation are exploring the computational requirements, domain size, etc.

Computational Needs

- Current 3-d calcs (running right now on Edison):
  - 6 cm resolution, but small domain (confines convection)
  - 11-isotope network (approximations create some artificial convective behavior)
  - Running to 100,000 steps will take 800,000 CPU hours on Edison (not including any charging factors)
    - Core sizes: domain decomposition is 192 64^3 grids: 192 MPI tasks $\times$ 6 OpenMP threads / MPI task
      - 64^3 grids seems to give the best performance, so the number of cores is then dictated by the domain size and resolution we desire.
    - I/O performance: (on Edison) 17 GB/s (each plotfile is 23 GB; 10 TB / simulation).
      - We spend less time to write a plotfile than to advance the solution 1 step. We write out every few hundred steps.
  - This job fits in about 1/5th of our allocation
    - That's where we like to run, so that we can do multiple jobs to understand sensitivity to the problem assumptions

*Note: information concerns the m1938 "OT" repo—not all the repos for our collaboration

Workflow

- ~10 TB output per simulation
  - Some checkpoints saved to HPSS
  - All plotfiles saved to HPSS
- Data kept at NERSC
  - Connections too slow to bring it back to NY
- Meta-data stored for reproducibility
  - Code git hashes, build machine/date/dir, runtime parameter values, compiler versions and flags, …
- Analysis
  - Global quantities output every step
  - Post-processing on plotfiles: VisIt (previously) and yt (preferred)

XRB Scaling @ NERSC

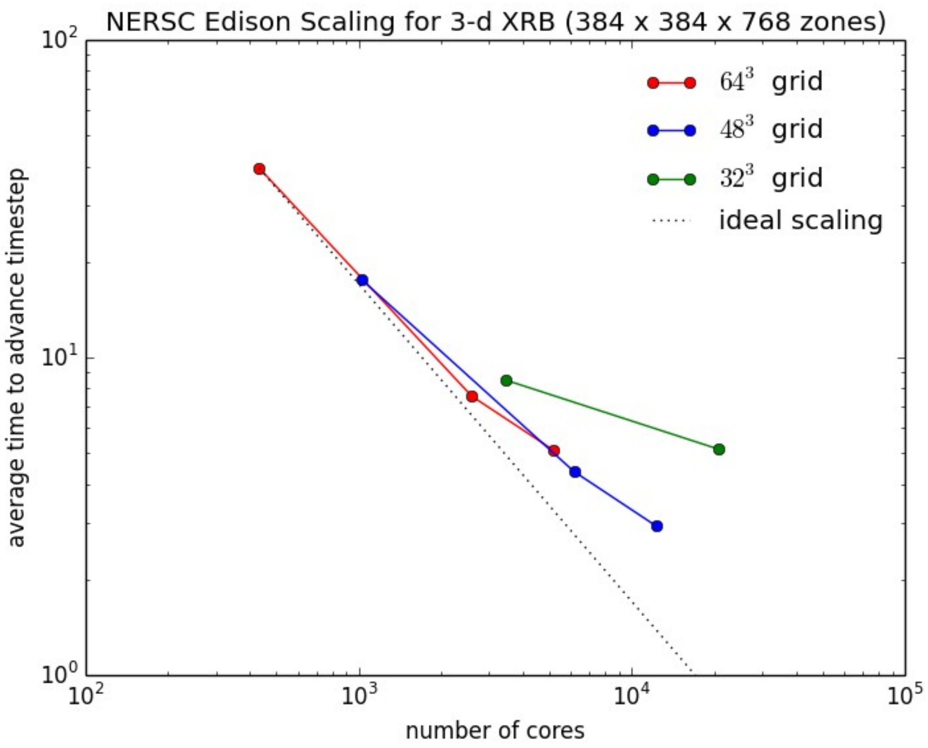

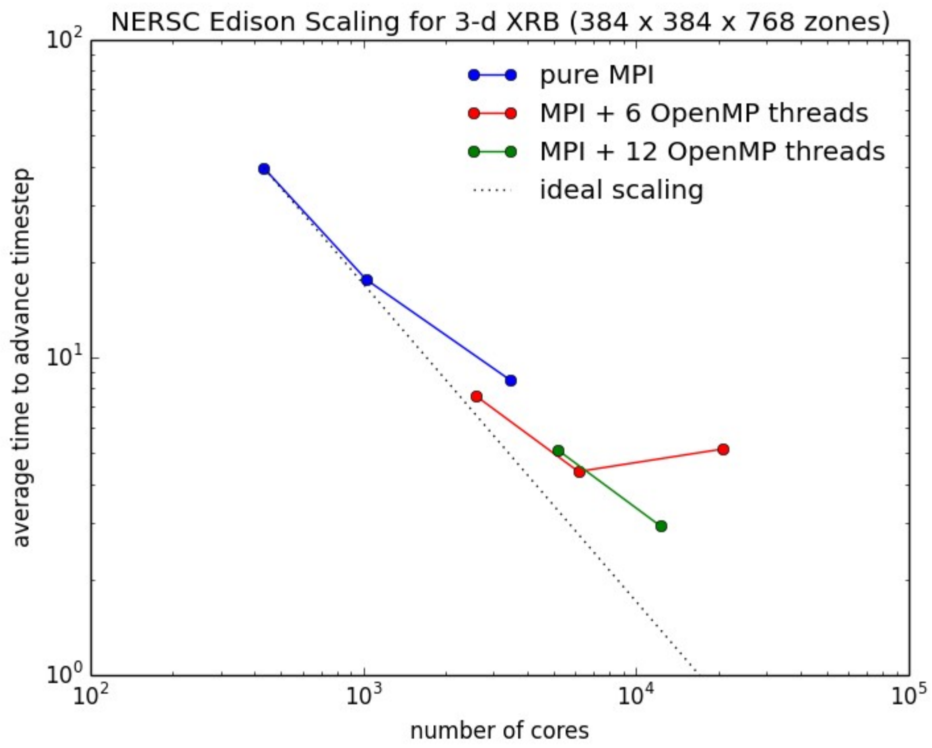

- 3-d XRB strong scaling (running on Edison):
  - $384^2 \times 768$ grid, 3 different domain decompositions (smaller than our target size)
  - MPI only or MPI + 6/12 threads. Always use all 24 cores on an Edison node.
  - Advection and reactions scale very well
  - Multigrid (especially the nodal solve) is the current bottleneck

XRB Scaling @ NERSC

64^3 grids

MPI	threads	advection	MAC	nodal	reactions	misc	total
432	1	14.031400	3.520500	2.002190	23.230700	8.548250	39.490900
432	6	2.011130	1.022320	1.180420	3.315770	1.341420	7.558950
432	12	1.253060	0.756383	1.145770	1.807880	0.862014	5.099680

48^3 grids

MPI	threads	advection	MAC	nodal	reactions	misc	total
1024	1	7.353820	1.760450	1.209890	10.139900	6.860430	17.607400
1024	6	1.211910	0.808226	1.011530	1.640570	1.219000	4.375010
1024	12	0.717778	0.582928	0.896226	0.844047	0.669709	2.947060

32^3 grids

MPI	threads	advection	MAC	nodal	reactions	misc	total
3456	1	1.922680	1.576740	2.900350	2.743110	1.973430	8.489440
3456	6	0.444619	1.338010	2.907420	0.525098	0.427345	5.148430

Where Do We Want To Go?

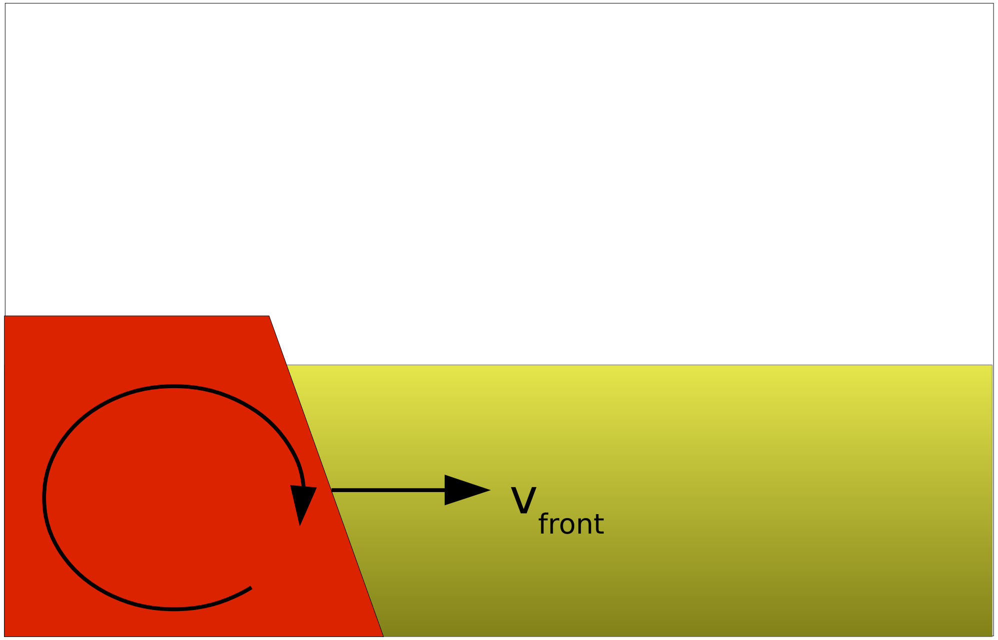

- Lateral flame propagation with resolved nuclear physics
  - Low Mach methods cannot (currently) describe two different scale heights (fuel and ash)
  - Lengthscale for Coriolis force to balance pressure gradient (Rossby length): $L = \sqrt{gH_0}/f$ ~ few km
    - Much bigger domain that we currently use

Future Computational Needs

- Next steps:
  - Wider domain (2$\times$ in each lateral dimension): 4$\times$ more expensive
  - Bigger network using vectorized ODE integrator on GPUs
  - Subgrid burning model: allows lower resolution, bigger domains, longer simulation times
  - Improved threading
- Future algorithmic developments:
  - Lateral gradients in the low Mach formulation
  - Magnetic fields
- NERSC help:
  - Training (webinars work ok—some of us teach and cannot easily travel)
  - Performance analysis ("hotspots") consulting
- Concerns:
  - We don't want to write custom code for each platform

Summary/Future

- Maestro is a mature simulation tool for exploring convective reacting flows in hydrostatic environments
  - Available: http://bender.astro.sunysb.edu/Maestro/
  - Lots of potential applications we need help with
- Smallscale 3-d simulations are running at NERSC now
  - Bottlenecks are identified, some plans to improve them are underway
- The future plans are always bigger, bigger, bigger
  - Both new physics and new algorithmic improvements are needed