Modeling Core Collapse Supernovae on the NERSC IBM-SP

Eric Myra & Doug Swesty
State Univ. of NY at Stony Brook

Collaborators:

Dennis Smolarski (Santa Clara Univ.)

Polly Baker (Indiana Univ.)

Ed Bachta (Indiana Univ.)

Ground to be covered...

- Motivation (why is the problem worth working on)
- A glimpse of the supernova scenario
- A glimpse of how we turn physics into a mathematical model
- A glimpse at parallel solvers and their performance
- Computational bottlenecks and challenges of the next stages

Motivation: The National Acad. Science Committee on Physics in the Universe's Eleven Big Questions

1. What is dark matter?
2. What is the nature of dark energy?
3. How did the universe begin?
4. Did Einstein have the last word on gravity?
5. What are the masses of neutrinos and how have they shaped the evolution of the universe?
6. How do cosmic accelerators work, and what are they accelerating?
7. Are protons unstable?
8. What are the new states of matter at exceedingly high density and temperature?
9. Are there additional space-time dimensions?
10. How were the elements from iron to uranium made?
11. Is a new theory of matter and light needed at the highest energies?

Motivation: Physics

Supernovae serve as "laboratories" for fundamental physics.

- ***Strong Interaction:***
  - serve as constraints on nuclear free-energy parameters through neutron stars mass ranges
    - incompressibility, symmetry parameters

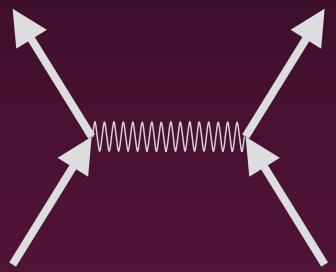

- ***Weak Interaction:***
  - reaction rates
  - neutrino masses and mixing

- ***Gravitational Interaction:***
  - expected to be visible to gravitational wave detectors

Motivation: Astronomy

Supernovae are needed to explain the universe's structure and make-up.

- We are star stuff!
  - Supernovae are the primary source of heavy elements.
- Supernovae are among the most energetic explosions in the cosmos.
- Affect the ISM (heating and compression) out to 100 light years.
- SN1987A commenced the science of extra-solar-system neutrino astronomy.
  - some day: to be repeated by not-too-close observation
  - no neutrino-radiation poisoning, please!
- Supernovae are a primary prospective source for gravitational-wave astronomy.

Motivation: Computation and Modelling

Supernova modelling is a challenging problem algorithmically:

- Diverse branches of physics:
  - relativistic hydrodynamics; non-equilibrium radiation
  - nuclear and weak-interaction physics
- Diverse coupling between branches:
  - matter-radiation coupling
  - multi-phase statistical mechanics
- Diverse timescales (strong, weak, hydrodynamic, diffusive, etc.)
- Diverse physical regimes (*e.g.,* densities can vary by 15 orders)

- A 1% problem!

$$\frac{\text{SN explosion energy}}{\text{proto-neutron star B.E.}} \approx \frac{10^{51}\text{erg}}{10^{53}\text{erg}} \approx 1\%$$

Motivation: Computation and Modeling

Supernova modelling is a challenging problem computationally.

Accurate modelling requires:

- multi-dimensional computations
- careful energy bookkeeping
- many nested iterative loops
- scalable pre-conditioners and sparse linear solvers for large parallel computers
- a tera- or petaFLOP-scale solution

The best methods (for the physics and the computation) require the fastest computers available.

The Storyline of the Project:

How to turn this…

…into this!

Stellar Core Collapse

-Massive stars evolve by burning lighter Elements, by thermonuclear fusion, into heavier elements

-High mass stars will produce elements Up to Silicon & Sulfur , which then burn into Iron

-Star has an Onionskin-like structure With layers of successively heavier elements

-Each burning stage progresses More rapidly

-When the Iron core mass becomes about 1.2-1.4 solar masses the core can no longer sustain itself against the pull of gravity & it collapses

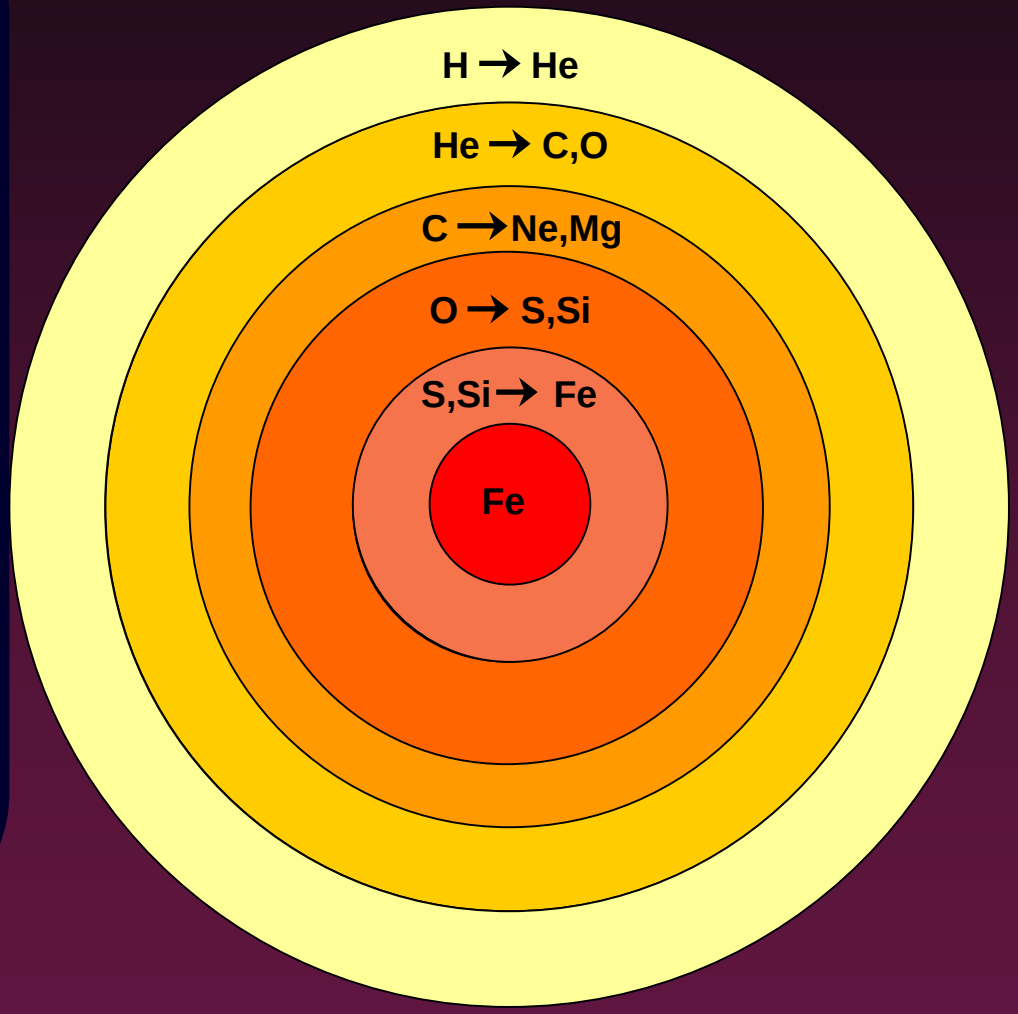

The strong force causes the core to bounce

- Nuclear matter goes through a series of phase transitions at about 10^{14} g/cm^3
  - pasta phases, bubble phases, etc.
- Above nuclear density, strong force is repulsive and dominant
- Collapse stops at center
  - soon after, stops everywhere else in sonic region
- A shock wave forms on the boundary of subsonic and supersonic material

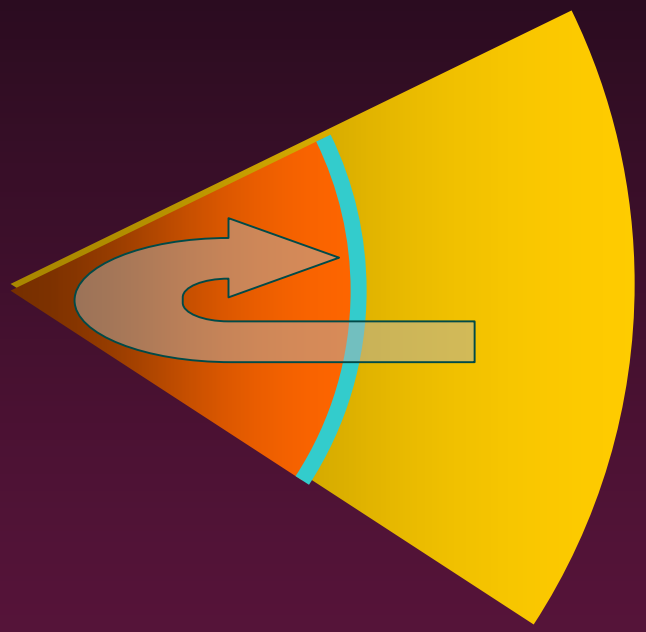

The shock wave starts to advance...

- The force of the core bounce starts the shock advancing.
- Supersonic material in the outer core...
  - falls through the shock.
  - is dissociated (costing 8 MeV/baryon ~ 2 x 10^{51} erg/0.1M_{sun}).
- Once the shock is beyond the neutrinosphere, μ's freely stream.
- Dissociation and μ emission deplete the shock of energy.
- Need to deliver about 10^{51} ergs to the mantle for an explosion.
- This does not happen in realistic calculations.

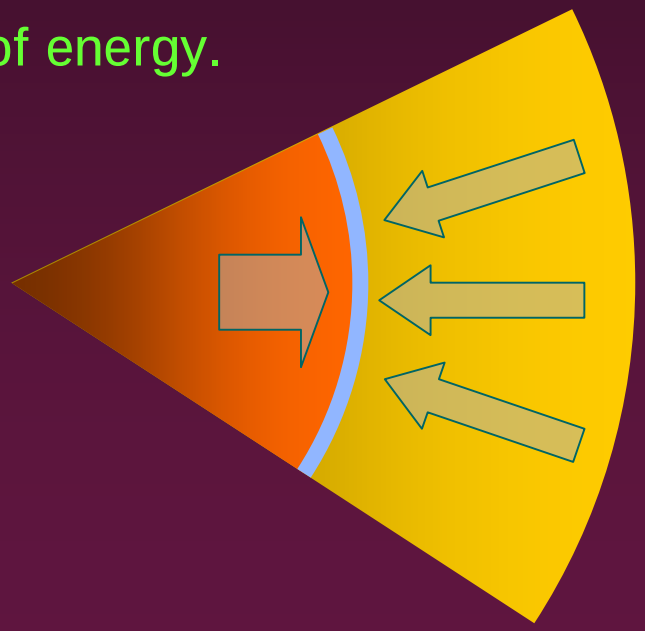

"Prompt"shocks eventually stall!

Supernova Explosion

-Shock wave heats matter as it moves outward

-Hot matter produces neutrino/anti-neutrino pairs of three types $e^+ + e^- \rightarrow \nu_x + \nu_x$

-Shock wave stalls before it reaches outer edge of iron core

-Hot (high entropy) matter starts rising upward (convection)

-A combination of convection and neutrino heating of the matter somehow revitalizes the stalled shock wave and causes the observed explosion

-Explosion leaves either a neutron star or a black hole behind

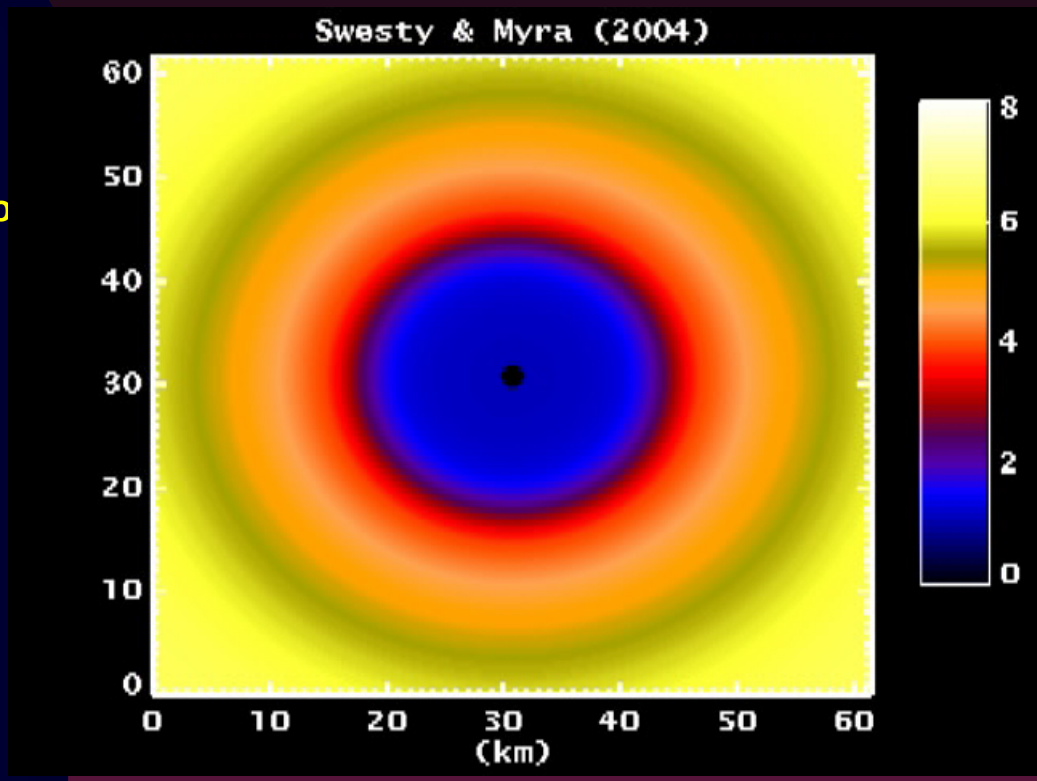

Microphysics

T = Temperature
ρ = Density
Y_e = Electron fraction
E_ε = Neutrino Energy Density
$\overline{E}_\varepsilon$ = Anti-neutrino Energy Density

EOS is obtained from table via Thermodynamically Consistent Interpolation Scheme; also maintains continuity in deriviatives of EOS w.r.t. T & ρ

Bi-quintic Hermite Interpolation in T & ρ, Linear Interpolation in Y_e

Neutrino Microphysics coefficients from tri-linear interpolation scheme

Equation of State

$$E = E(T, \rho, Y_e)$$

$$P = P(T, \rho, Y_e)$$

Emissivity

$$S_\varepsilon = S_\varepsilon(T, \rho, Y_e)$$

Absorption Opacity

$$\kappa_\varepsilon^a = \kappa_\varepsilon^a(T, \rho, Y_e)$$

Diffusion Coefficient

$$D_\varepsilon = D_\varepsilon(T, \rho, Y_e)$$

Scattering Coefficients

$$\kappa^{in}(\varepsilon, \varepsilon') = \kappa^{in}(\varepsilon, \varepsilon', T, \rho, Y_e)$$

$$\kappa^{out}(\varepsilon, \varepsilon') = \kappa^{out}(\varepsilon, \varepsilon', T, \rho, Y_e)$$

Pair Emissivity

$$G(\varepsilon, \varepsilon') = G(\varepsilon, \varepsilon', T, \rho, Y_e)$$

O(v/c) Multigroup Neutrino Energy Equation

One pair of neutrino/anti-neutrino energy eqns. for each of three neutrino flavors: e,μ, τ

$$\begin{cases} \dfrac{\partial E_\epsilon}{\partial t} + \nabla\cdot(E_\epsilon \mathbf{v}) + \nabla\cdot\mathbf{F}_\epsilon - \epsilon\dfrac{\partial}{\partial\epsilon}(\mathsf{P}_\epsilon) : \nabla\mathbf{v} = \mathcal{S}_\epsilon \\[2ex] \dfrac{\partial \bar{E}_\epsilon}{\partial t} + \nabla\cdot(\bar{E}_\epsilon \mathbf{v}) + \nabla\cdot\bar{\mathbf{F}}_\epsilon - \epsilon\dfrac{\partial}{\partial\epsilon}\left(\bar{\mathsf{P}}_\epsilon\right) : \nabla\mathbf{v} = \bar{\mathcal{S}}_\epsilon \end{cases}$$

Energy density $= E_\epsilon \equiv \frac{1}{c}\int I_\epsilon d\Omega$ Pressure tensor $= \mathsf{P}_\epsilon \equiv \frac{1}{c}\int I_\epsilon \Omega^2 d\Omega$

Flux $= \mathbf{F}_\epsilon \equiv \int I_\epsilon \Omega d\Omega$

Flux-Limited Diffusion (FLD) Approximation

Diffusion coefficient

$$\mathbf{F}_\epsilon = -D_\epsilon \nabla E_\epsilon \qquad \bar{\mathbf{F}}_\epsilon = -\bar{D}_\epsilon \nabla \bar{E}_\epsilon$$

$$\mathsf{P}_\epsilon = \chi_\epsilon E_\epsilon \qquad \bar{\mathsf{P}}_\epsilon = \bar{\chi}_\epsilon \bar{E}_\epsilon$$

Tensor Eddington factor

Optically thin limit	Optically thick limit
$D_\epsilon \rightarrow \dfrac{-cE_\epsilon}{\lvert\nabla E_\epsilon\rvert}$	$D_\epsilon \rightarrow \dfrac{c}{3\kappa_\epsilon}$
$\chi_\epsilon \rightarrow 1$	$\chi_\epsilon \rightarrow \dfrac{1}{3}$

Parallel Implementation

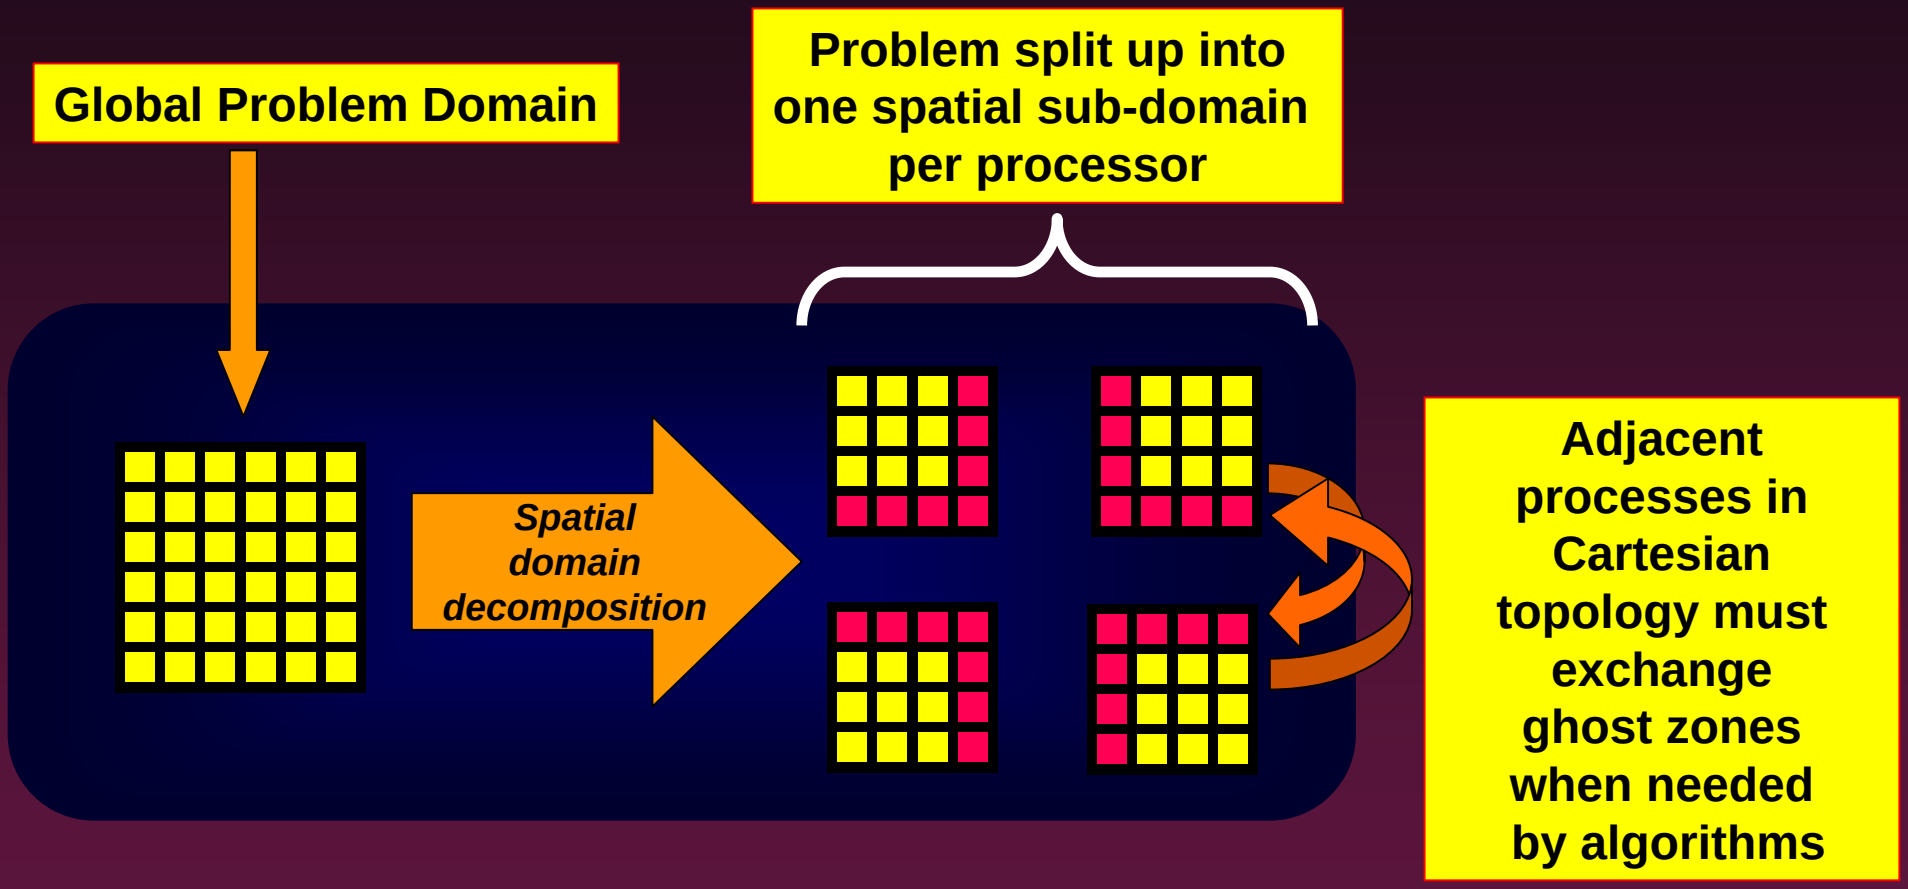

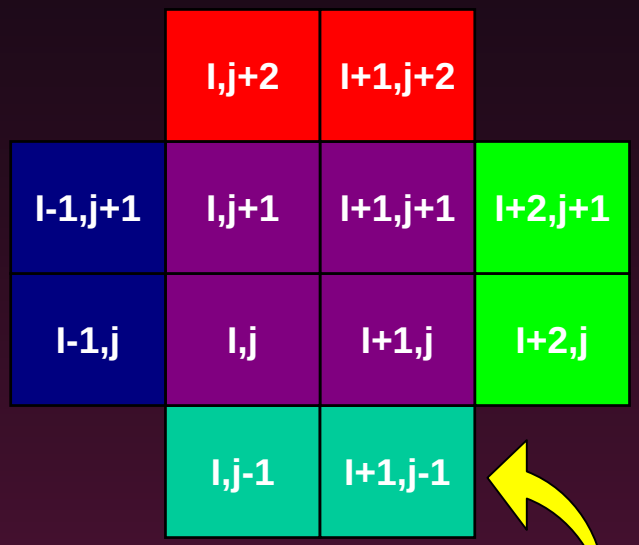

Matrices arising from local spatial operators require only nearest-neighbor communication to implement matrix-vector multiplies

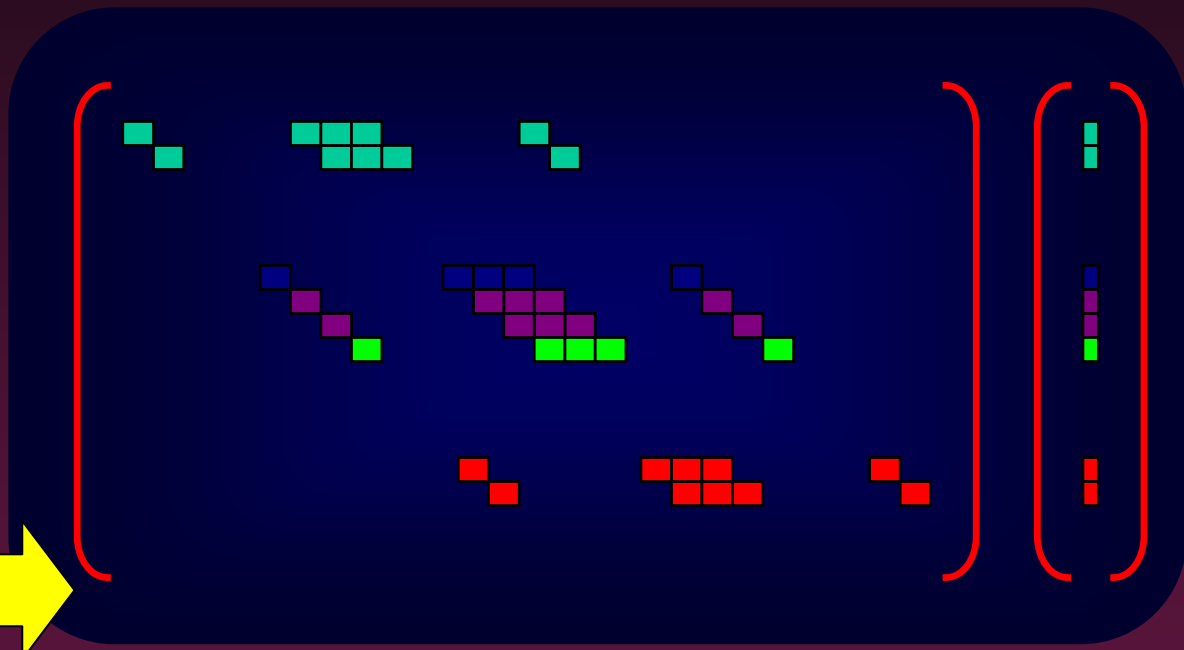

Colors indicate the ownership of zones by different processes

Colors also indicate how the block banded matrix and the vectors arising from difference eqns. are distributed among processors

Implicit Transport Evolution

-The speed of light sets the CFL stability limit for the transport

- Explicit transport timesteps will be very short:
  < 1/10th the hydrodynamic CFL timestep

-Weak interaction timescales are even shorter which create stiff coupling between matter & neutrinos

-Therefore we must use implicit methods to solve the transport equations

-This implies we will have to solve large sparse sets of linear or non-linear equations at every timestep.

-With lexicographic ordering linear & non-linear systems will have a block & band structure

1-D Multigroup Diffusion Matrix

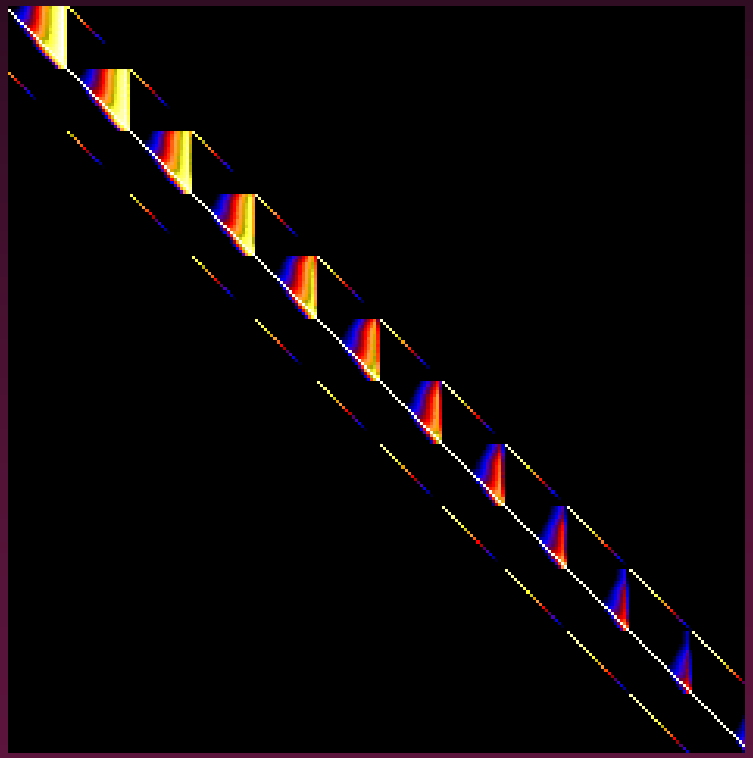

The Matrices Are Sparse

DENSE BLOCKS OF COUPLED ENERGY GROUPS DOWN MAIN DIAGONAL

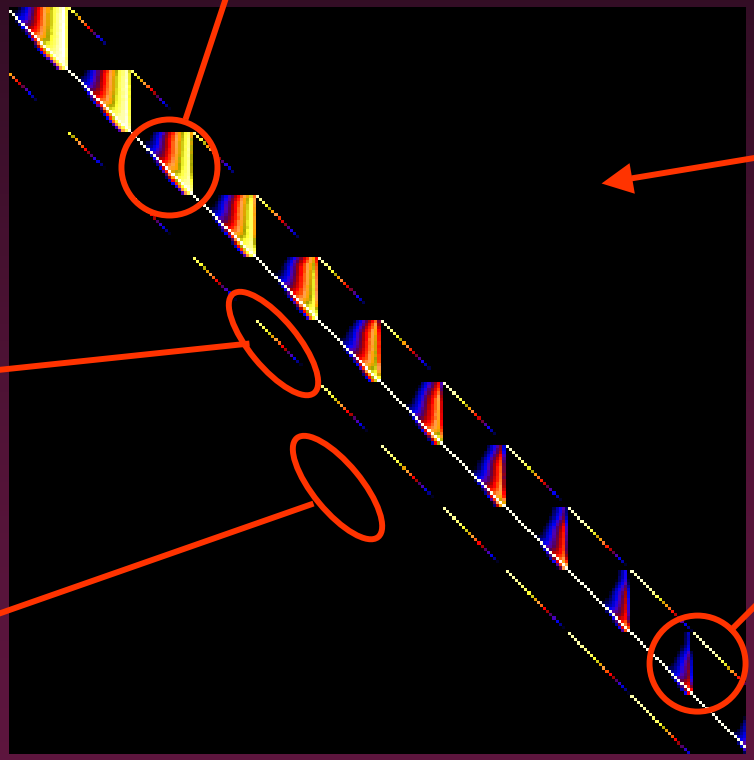

BLACK AREAS ARE WHERE THE MATRIX HAS ENTRIES OF ZERO.

OUTLIERS CORRESPOND TO SPATIAL COUPLING

MORE BANDS OF OUTLIERS IN HIGHER SPATIAL DIMENSIONS

ENERGY COUPLING DECREASES FURTHER OUT IN THE STAR (LOWER DENSITY OF SCATTERING CENTERS)

BiCGSTAB Algorithm

for Ax=b (Van der Vorst 1994)

For details consult:

Iterative Methods for Solving Linear Systems (A. Greenbaum 1997)

Iterative Methods for Sparse Linear Systems (Y. Saad 2003)

Templates for the Solution of Linear Systems (Barrett, et al. 1993)

Compute $r^{(0)} = b - A x^{(0)}$
Choose $q = r^{(0)}$
FOR i=1,2,...
  $\rho_{i-1} = (q^T r^{(i-1)})$
  IF $r_{i-1} = 0$ exit
  IF i = 1
    $p^{(i)} = r^{(i-1)}$
  ELSE
    $\beta_{i-1} = (\rho_{i-1} / \rho_{i-2})(\alpha_{i-1} / \omega_{i-1})$
    $p^{(i)} = r^{(i-1)} + \beta_{i-1}(p^{(i-1)} - \omega_{i-1} v^{(i-1)})$
  ENDIF
  $h = M^{-1} p^{(i)}$
  $v^{(i)} = Ah$
  $\alpha_i = \rho_{i-1} / (q^T v^{(i)})$
  $s = r^{(i-1)} - \alpha_i v^{(i)}$
  IF $|s|$ is small, set $x^{(i)} = x^{(i-1)} + \alpha_i h$ and exit
  $g = M^{-1} s$
  $t = Ag$
  $\omega_i = (t^T s) / (t^T t)$
  $x^{(i)} = x^{(i-1)} + \alpha_i h + \omega_i g$
  $r^{(i)} = s - \omega_i t$
  IF $|r^{(i)}| < \varepsilon |b|$ exit
ENDFOR

BiCGSTAB Algorithm

for Ax=b (Van der Vorst 1994)

```
Compute r^(0) = b-Ax^(0)
Choose q = r^(0)
FOR i=1,2,...
    ρ_(i-1) = (q^T r^(i-1))
    IF  r_(i-1) = 0  exit
    IF  i = 1
        p^(i) = r^(i-1)
    ELSE
        β_(i-1) = (ρ_(i-1) / ρ_(i-2)) (α_(i-1) / ω_(i-1))
        p^(i) = r^(i-1) + β_(i-1) (p^(i-1) - ω_(i-1) v^(i-1))
    ENDIF
    h = M^-1 p^(i)
    v^(i) = Ah
    α_i = ρ_(i-1) / (q^T v^(i))
    s = r^(i-1) - α_i v^(i)
    IF |s| is small, set x^(i) = x^(i-1) + α_i h  and exit
    g = M^-1 s
    t = Ag
    ω_i = (t^T s) / (t^T t)
    x^(i) = x^(i-1) + α_i h + ω_i g
    r^(i) = s - ω_i t
    IF |r^(i)| < ε |b|  exit
ENDFOR
```

nearest neighbor communication for MATVECs

BiCGSTAB Algorithm

for Ax=b (Van der Vorst 1994)

Compute $r^{(0)} = b - Ax^{(0)}$
Choose $q = r^{(0)}$
FOR i=1,2,...
  $\rho_{i-1} = (q^T r^{(i-1)})$
  IF $r_{i-1} = 0$ exit
  IF i = 1
    $p^{(i)} = r^{(i-1)}$
  ELSE
    $\beta_{i-1} = (\rho_{i-1} / \rho_{i-2})(\alpha_{i-1} / \omega_{i-1})$
    $p^{(i)} = r^{(i-1)} + \beta_{i-1}(p^{(i-1)} - \omega_{i-1} v^{(i-1)})$
  ENDIF
  $h = M^{-1} p^{(i)}$
  $v^{(i)} = Ah$
  $\alpha_i = \rho_{i-1} / (q^T v^{(i)})$
  $s = r^{(i-1)} - \alpha_i v^{(i)}$
  IF $|s|$ is small, set $x^{(i)} = x^{(i-1)} + \alpha_i h$ and exit
  $g = M^{-1} s$
  $t = Ag$
  $\omega_i = (t^T s) / (t^T t)$
  $x^{(i)} = x^{(i-1)} + \alpha_i h + \omega_i g$
  $r^{(i)} = s - \omega_i t$
  IF $|r^{(i)}| < \varepsilon |b|$ exit
ENDFOR

nearest neighbor communication for preconditioner

nearest neighbor communication for MATVECs

BiCGSTAB Algorithm

for Ax=b (Van der Vorst 1994)

Global reduction operation for inner products

nearest neighbor communication for preconditioner

nearest neighbor communication for MATVECs

Compute $r^{(0)} = b - Ax^{(0)}$
Choose $q = r^{(0)}$
FOR i=1,2,...
  $\rho_{i-1} = (q^T r^{(i-1)})$
  IF $r_{i-1} = 0$ exit
  IF i = 1
    $p^{(i)} = r^{(i-1)}$
  ELSE
    $\beta_{i-1} = (\rho_{i-1} / \rho_{i-2})(\alpha_{i-1} / \omega_{i-1})$
    $p^{(i)} = r^{(i-1)} + \beta_{i-1}(p^{(i-1)} - \omega_{i-1} v^{(i-1)})$
  ENDIF
  $h = M^{-1} p^{(i)}$
  $v^{(i)} = Ah$
  $\alpha_i = \rho_{i-1} / (q^T v^{(i)})$
  $s = r^{(i-1)} - \alpha_i v^{(i)}$
  IF $|s|$ is small, set $x^{(i)} = x^{(i-1)} + \alpha_i h$ and exit
  $g = M^{-1} s$
  $t = Ag$
  $\omega_i = (t^T s) / (t^T t)$
  $x^{(i)} = x^{(i-1)} + \alpha_i h + \omega_i g$
  $r^{(i)} = s - \omega_i t$
  IF $|r^{(i)}| < \varepsilon |b|$ exit
ENDFOR

Newton-Krylov Solves

N-K solution via Newton-BiCGSTAB

Currently using sparse parallel approximate inverse preconditioning (physics based) -- highly parallel

Transport Courant # ~ 10-20 means we can predict where to place elements of approximate inverse

Typically 2-3 Newton iterations

Typically 10-20 Bicgstab iterations per Newton iteration

Implicit solvers have to be ultra-fast -- 10^5 - 10^6 timesteps to complete simulation

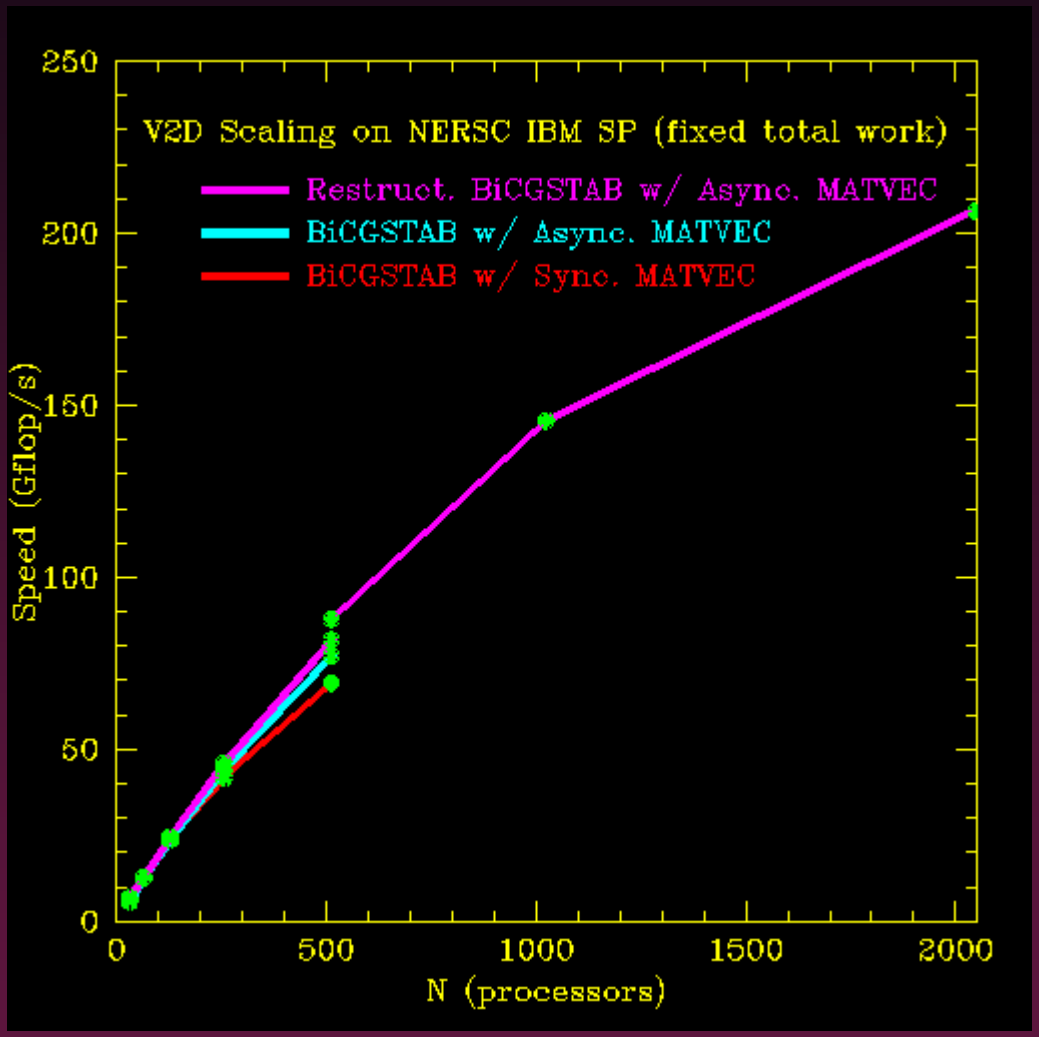

Convection Occurs....

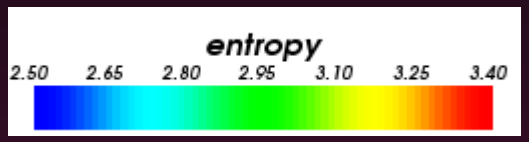

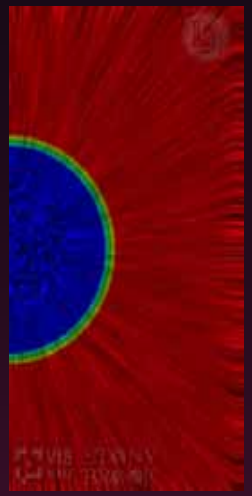

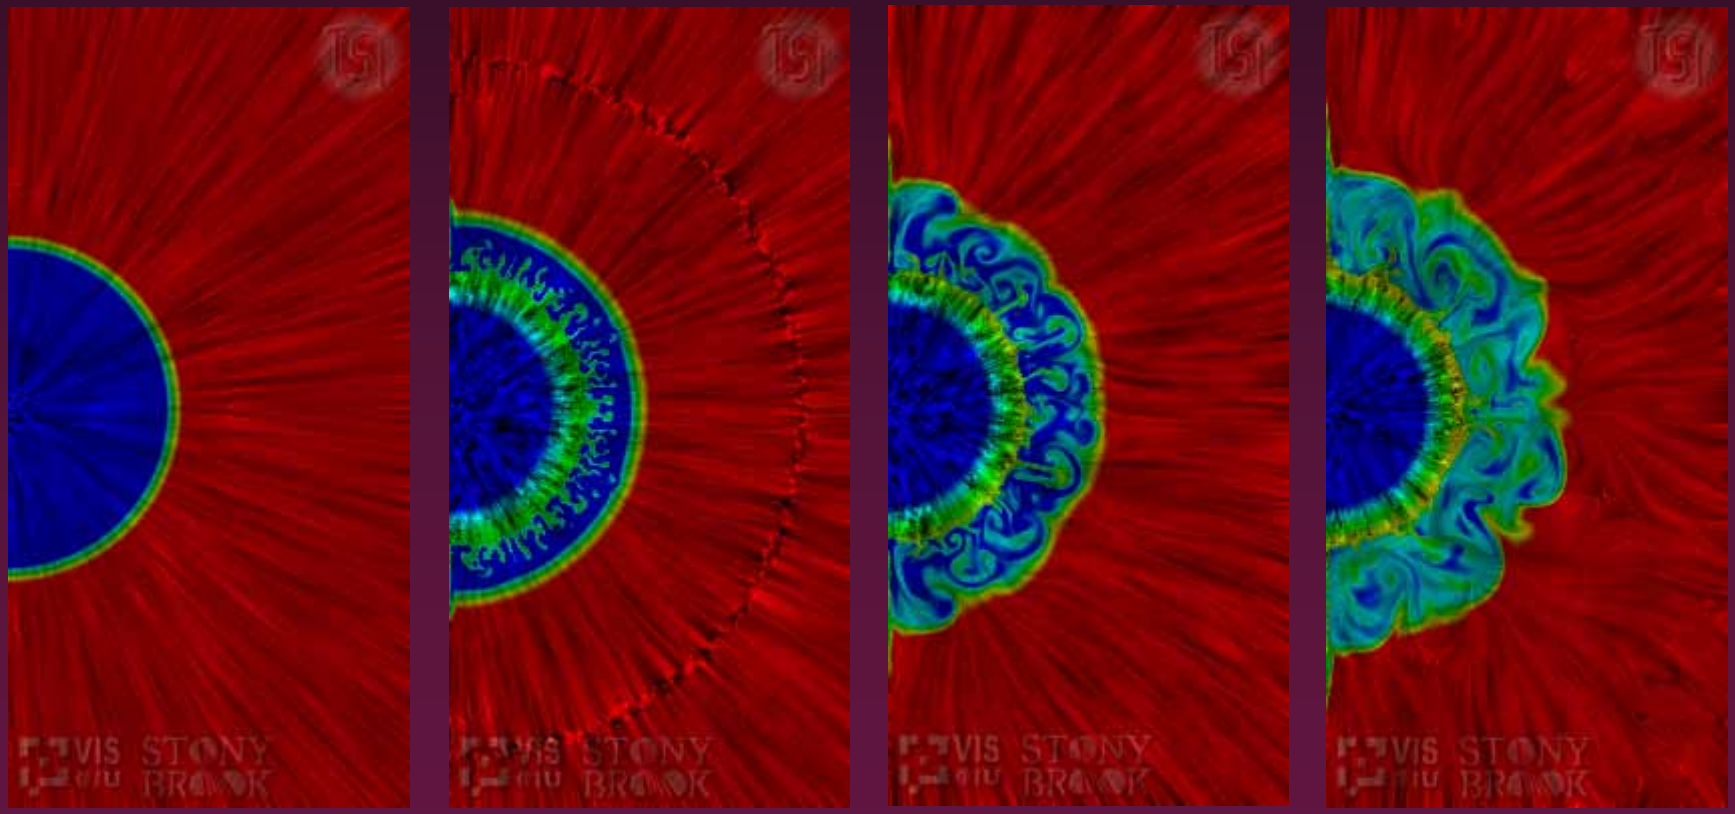

New Result: Radiative & Advective Fluxes Can Oppose Each Other:

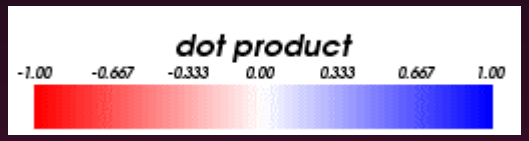

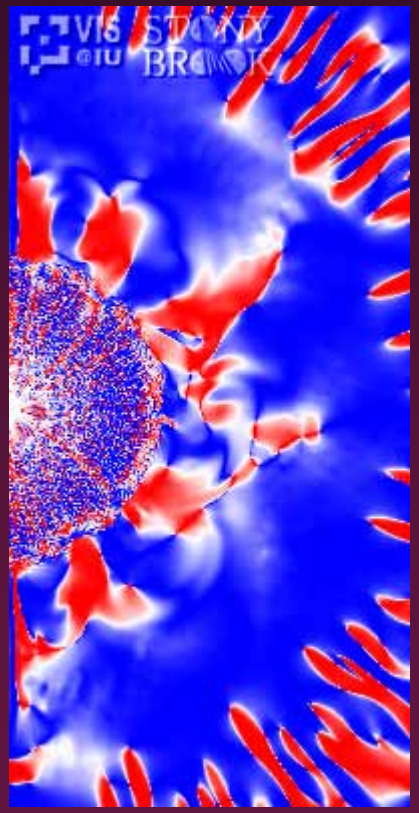

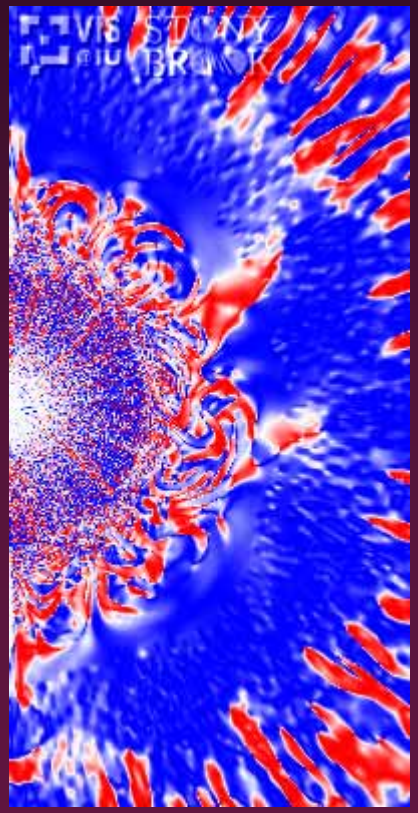

Conclusions

- We are able to carry out highly scalable implicit radiation-hydrodynamic models on the NERSC IBM SP
- There is a successful parallel preconditioning strategy that speeds up iterative performance of our Newton-Krylov solvers without compromising scalability
- The availability of Seaborg has allowed us to carry out new computational science at very large scales
- The advective neutrino fluxes play a very important role in core collapse supernova models
- 3-D models coming soon!