Using a Developing MiniApp to Compare Platform Characteristics on Cray Systems

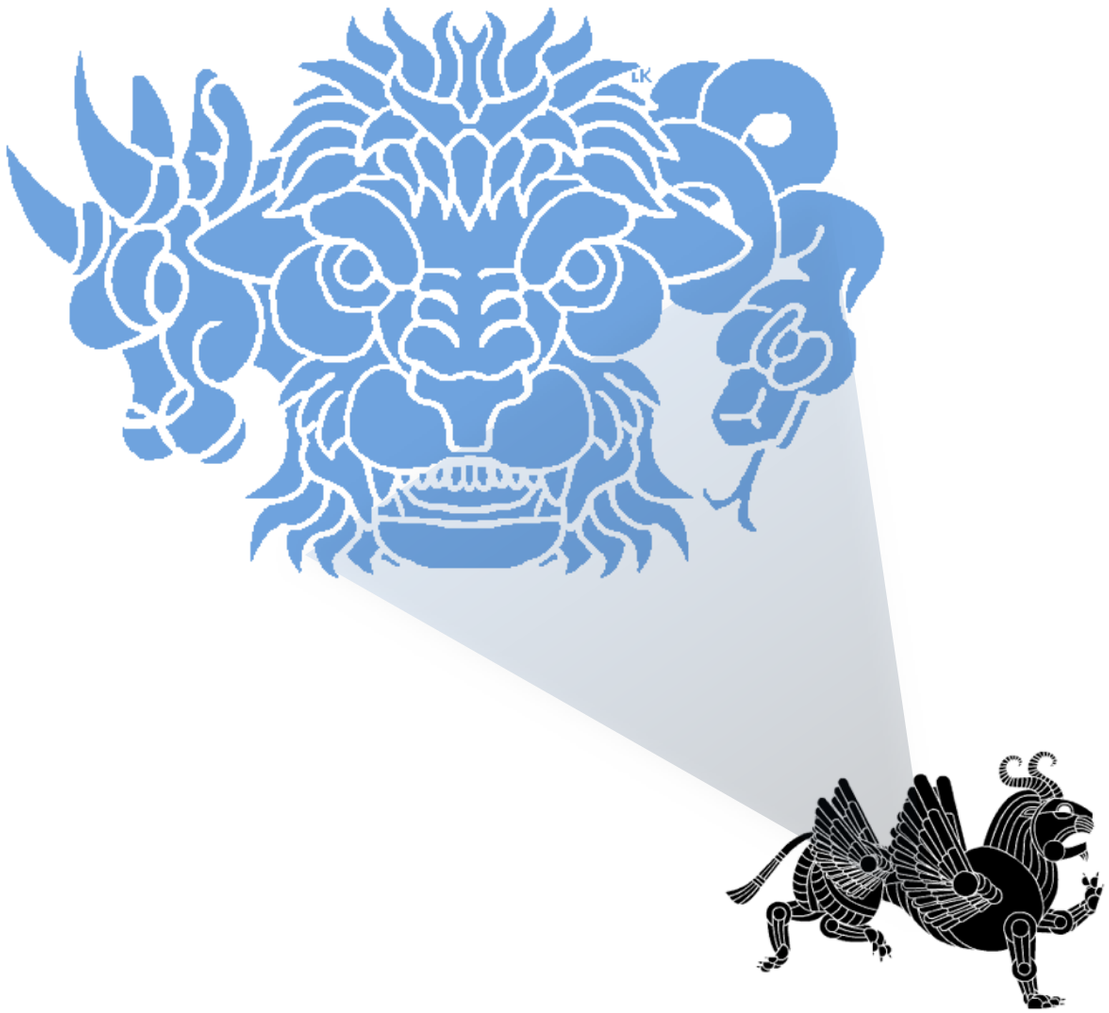

Bronson Messer

Scientific Computing & Theoretical Astrophysics Groups

Oak Ridge National Laboratory

ORNL is managed by UT-Battelle
for the US Department of Energy

OAK RIDGE National Laboratory

CHIMERA collaboration

- Steve Bruenn (Florida Atlantic University)
- John Blondin, Chris Mauney (NC State University)
- Eirik Endeve, Raph Hix, Austin Harris, Eric Lentz, Bronson Messer, Anthony Mezzacappa, Konstantin Yakunin (ORNL/UTK)

- Former Team Members
  - Reuben Budjiara, Austin Chertkow, Ted Lee

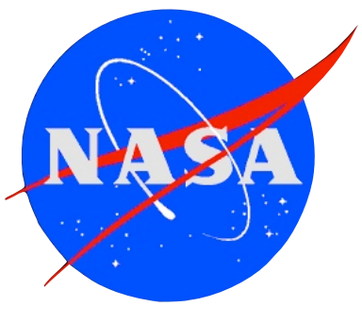

MiniApp work sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory

The research and activities described in this presentation were performed using the resources of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC0500OR22725.

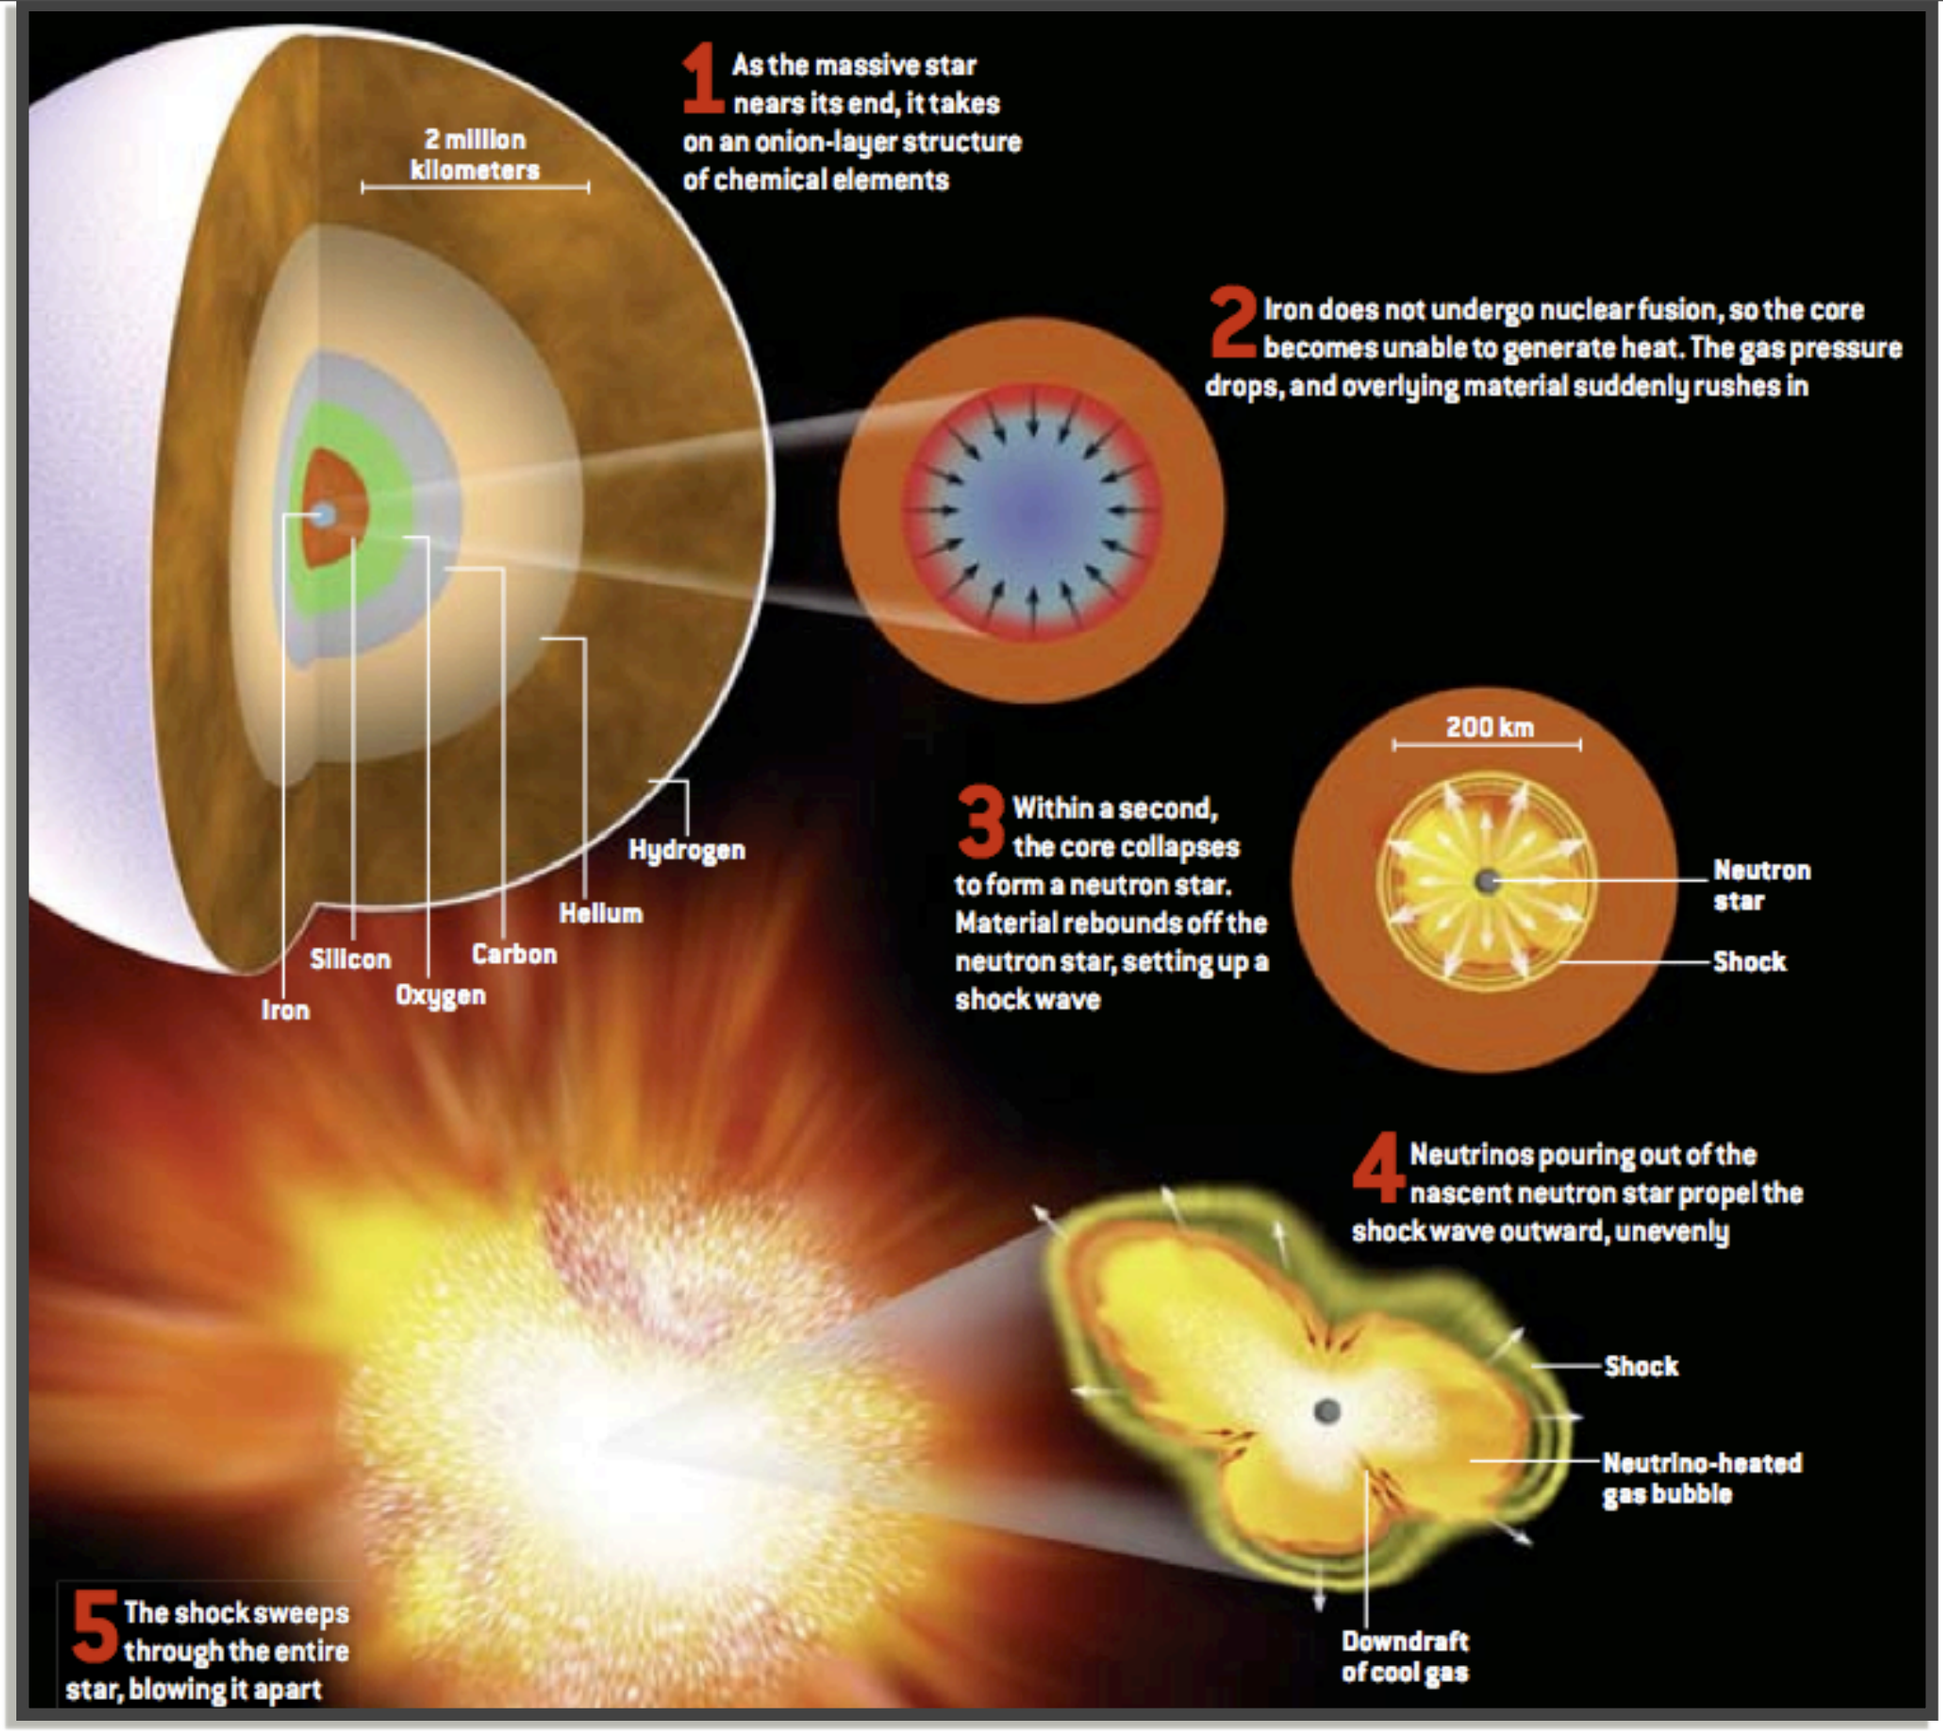

Hillebrandt & Janka 2006 (Sci Am)

How is the supernova shock revived?

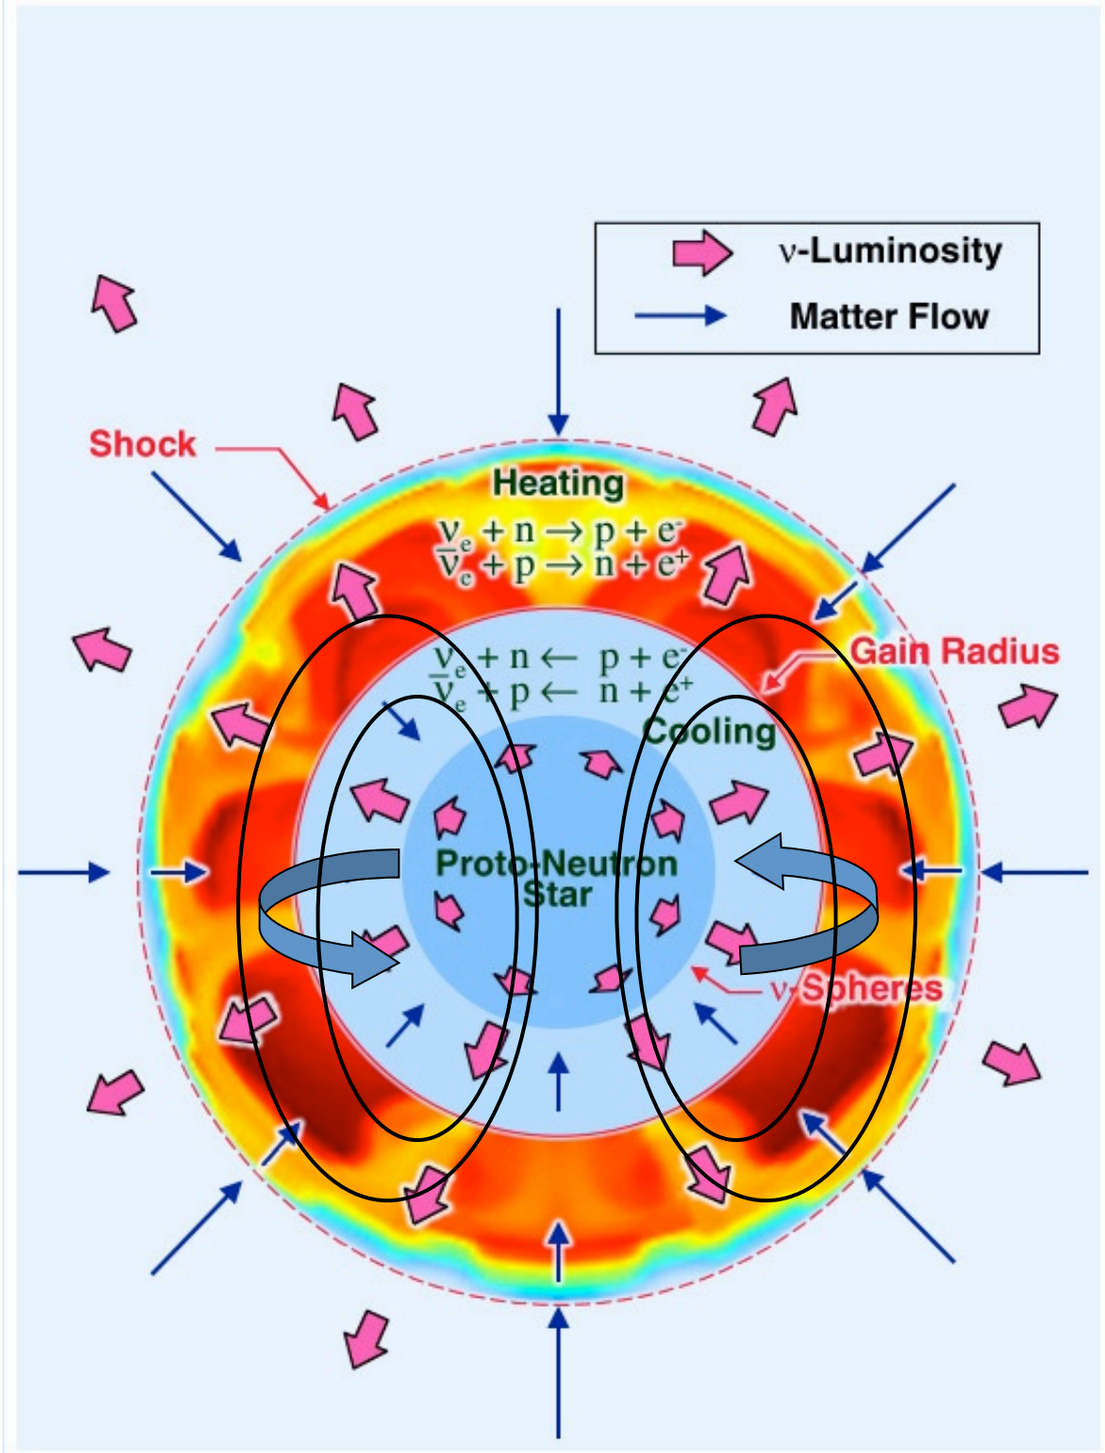

Known, Potentially Important Ingredients

- Gravity
- Neutrino Heating
- Convection
- Shock Instability (SASI)
- Nuclear Burning
- Rotation
- Magnetic Fields

Need 3D models with all of the above, treated with sufficient realism.

CHIMERA

- "Ray-by-ray-Plus" MGFLD Neutrino Transport
  - O(v/c), GR time dilation and redshift, GR aberration
- PPM Hydrodynamics (finite-volume)
  - GR time dilation, effective gravitational potential
  - adaptive radial grid
- Lattimer-Swesty EOS + low-density BCK EOS
  - K=220 MeV
  - low-density EOS (BCK+NSE solver) "bridges" LS to network
- Nuclear (Alpha) Network
  - 14 alpha nuclei between helium and zinc
- Effective Gravitational Potential
  - Marek et al. A&A, 445, 273 (2006)
- Neutrino Emissivities/Opacities
  - "Standard" + Elastic Scattering on Nucleons + Nucleon–Nucleon Bremsstrahlung

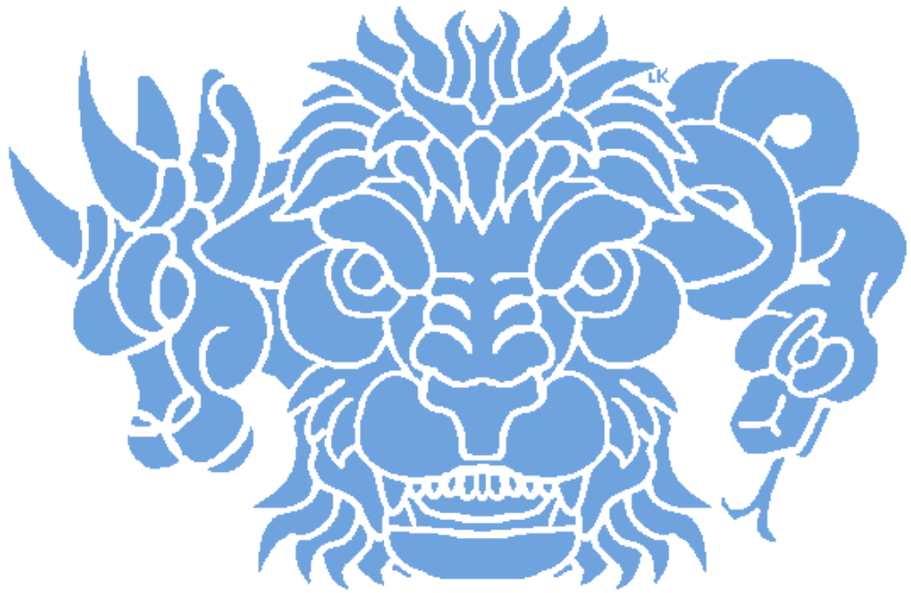

Average INCITE award of 50M cpu-hrs/year over the past 5 years ~ 10% of total INCITE available

15 solar mass 3D run

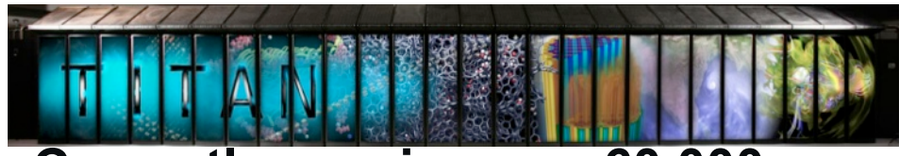

Currently running on ~60,000 cores

- 15 solar mass WH07 progenitor
- 540 radial zones covering inner 11000 km
- 180 phi zones (2 degree resolution)
- 180 theta zones in "constant mu" grid, from 2/3 degree at equator to one 8.5 degree zone at pole.
- Full opacities
- 0.1% density perturbations (10-30 km) applied at 1.3 ms after bounce in transition from 1D.

Entropy

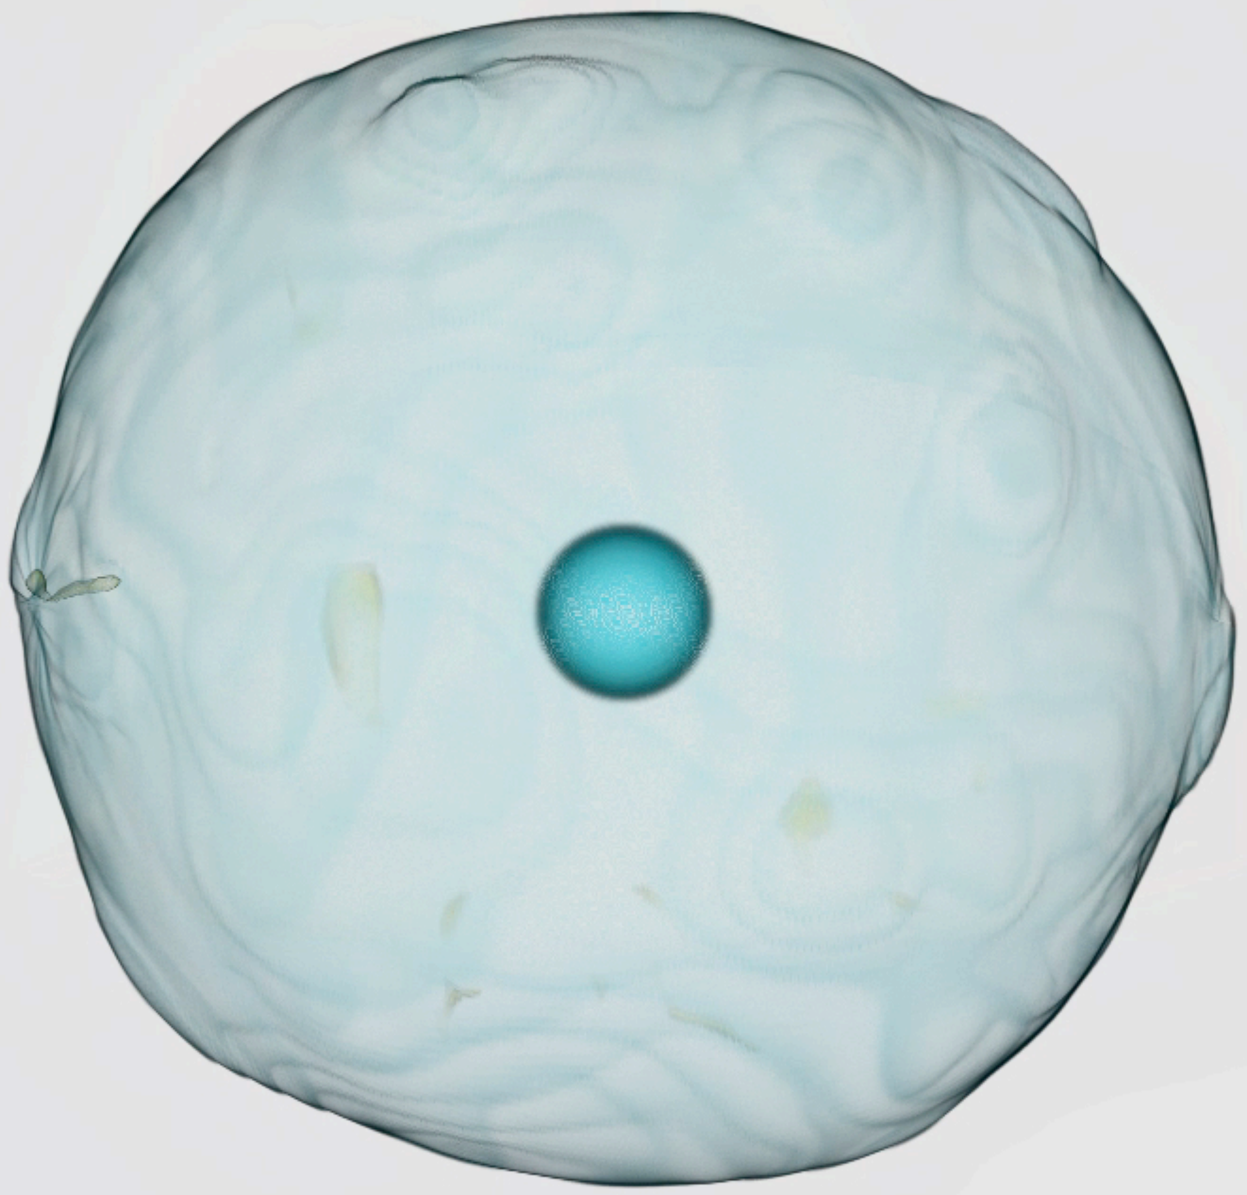

OLCF
Oak Ridge Leadership Computing Facility

CHIMERA program flow

- Directionally-split hydrodyamics is the “spine” of CHIMERA.
- Operator-split multiphysics is performed on strictly **node-local** data.

See Levesque’s VH-1 example from Monday’s Reveal tutorial

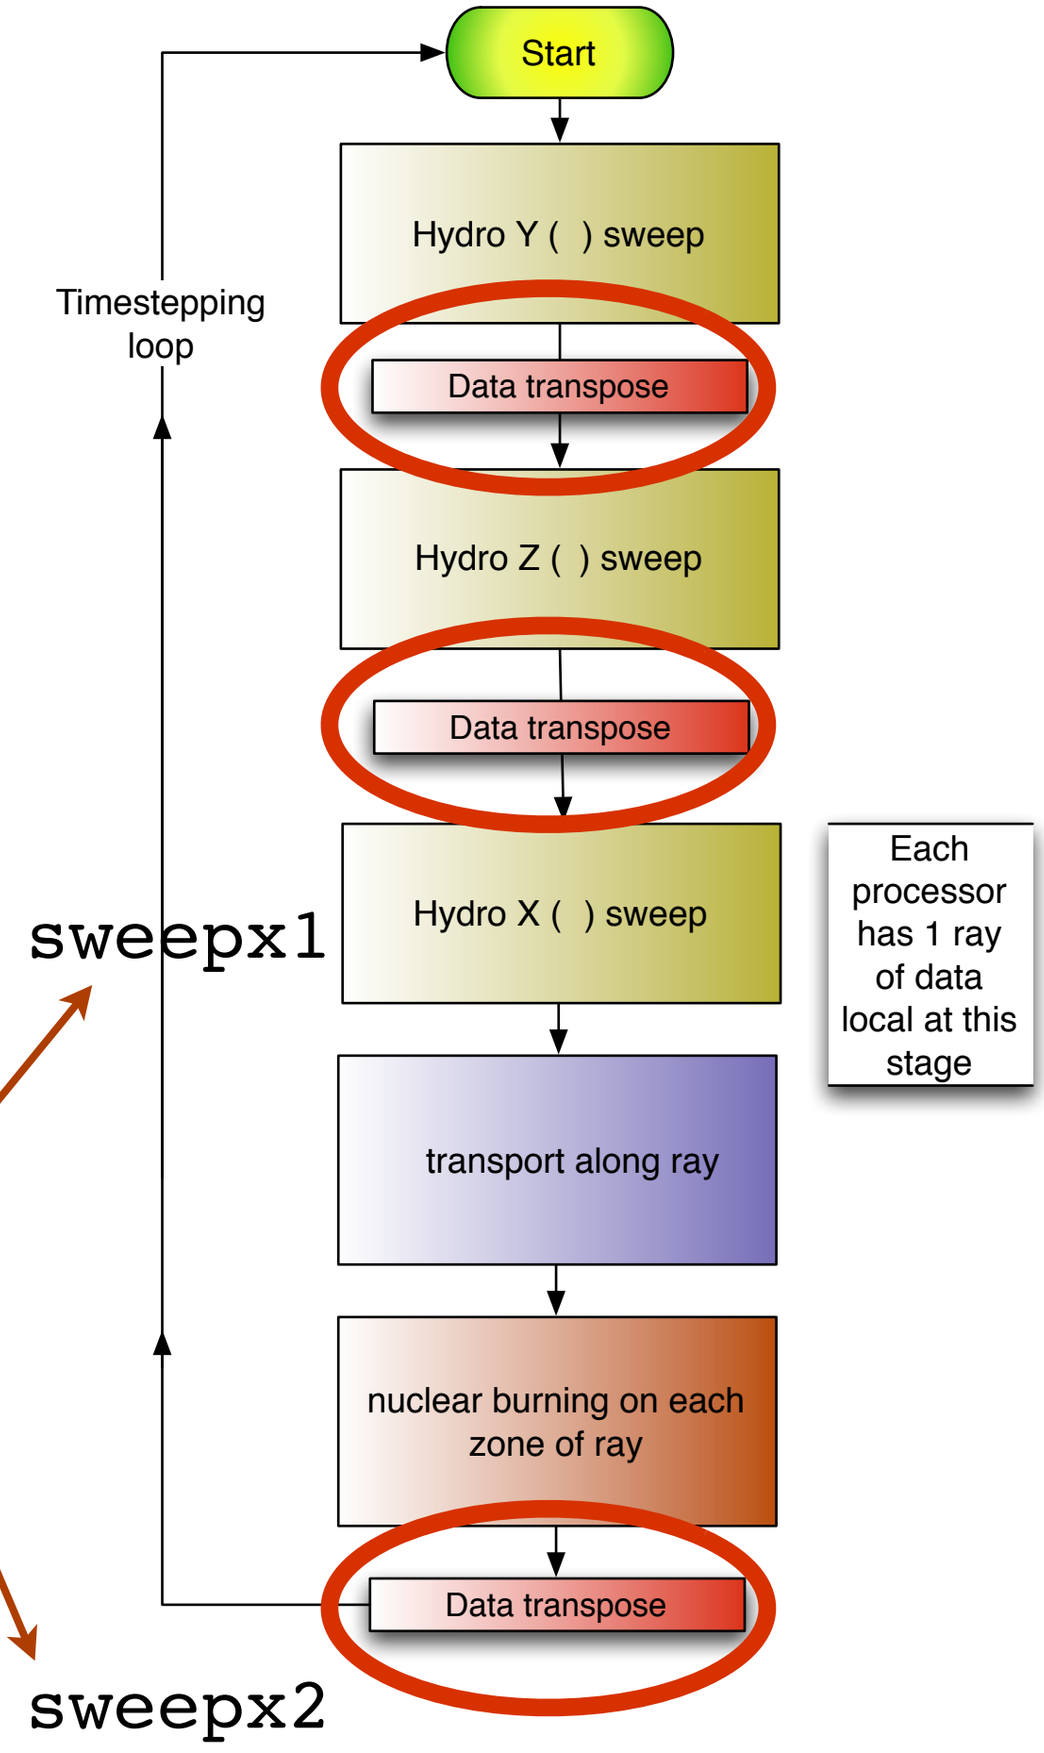

MPI domain decomposition

Using M*N processors; X data starts local to proc

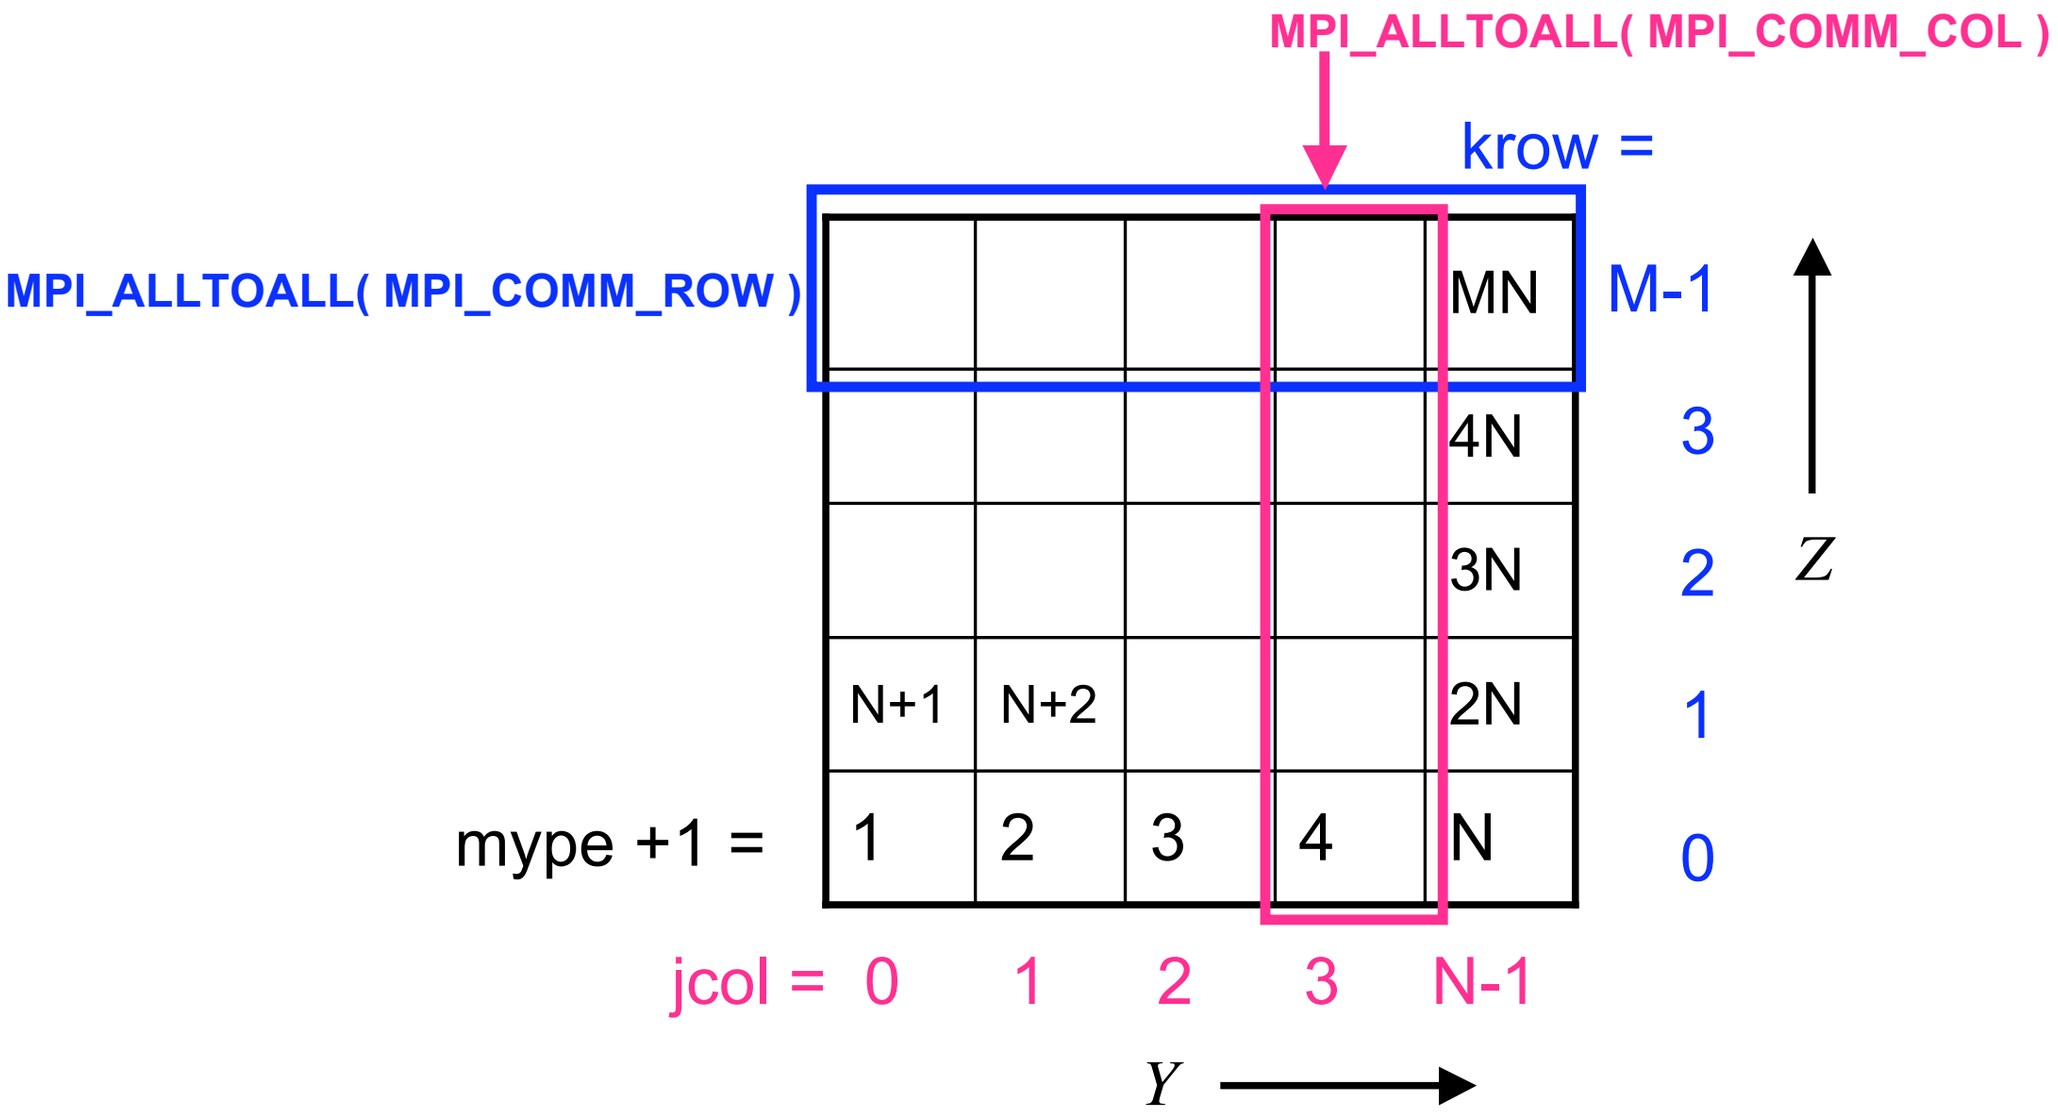

- Leads to particular modulo arithmetic for resolution vs. # of MPI ranks

What are MiniApps and why are they important?

- **MiniApps are reduced, proxy applications that encapsulate the salient performance characteristics of larger, full-sized applications**
- Current and near-future HPC architectures are becoming more and more complex
- HPC codebases are also complicated and unwieldy
  - often decades old
  - often very large (can be O(1M) lines of code)
  - have accreted complicated build systems and other architecture-specific features
- Trying to predict code performance on new architectures and understand the impact of algorithmic choices can be close to impossible.
- MiniApps provide a way forward for system architects, computer scientists, and applied mathematics researchers to help escape this conundrum.

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

Ziz*: A CHIMERA MiniAPP

- Designed to be modular (more modular than CHIMERA).
- Current beta version models the VH-1 (MVH3) ‘skeleton’ of CHIMERA - the stellar hydrodynamics
- CHIMERA: ~350K LOC ; Ziz ~ 2500 LOC
- Transport and nuclear kinetics (‘burning’) to be added
- Does Ziz mirror the performance of CHIMERA? Can it?

***A ziz is is a giant griffin-like bird in Hebrew mythology, often portrayed as something somewhat akin to a Greek chimera. The name is also easy to type.**

Weak scaling versus CHIMERA

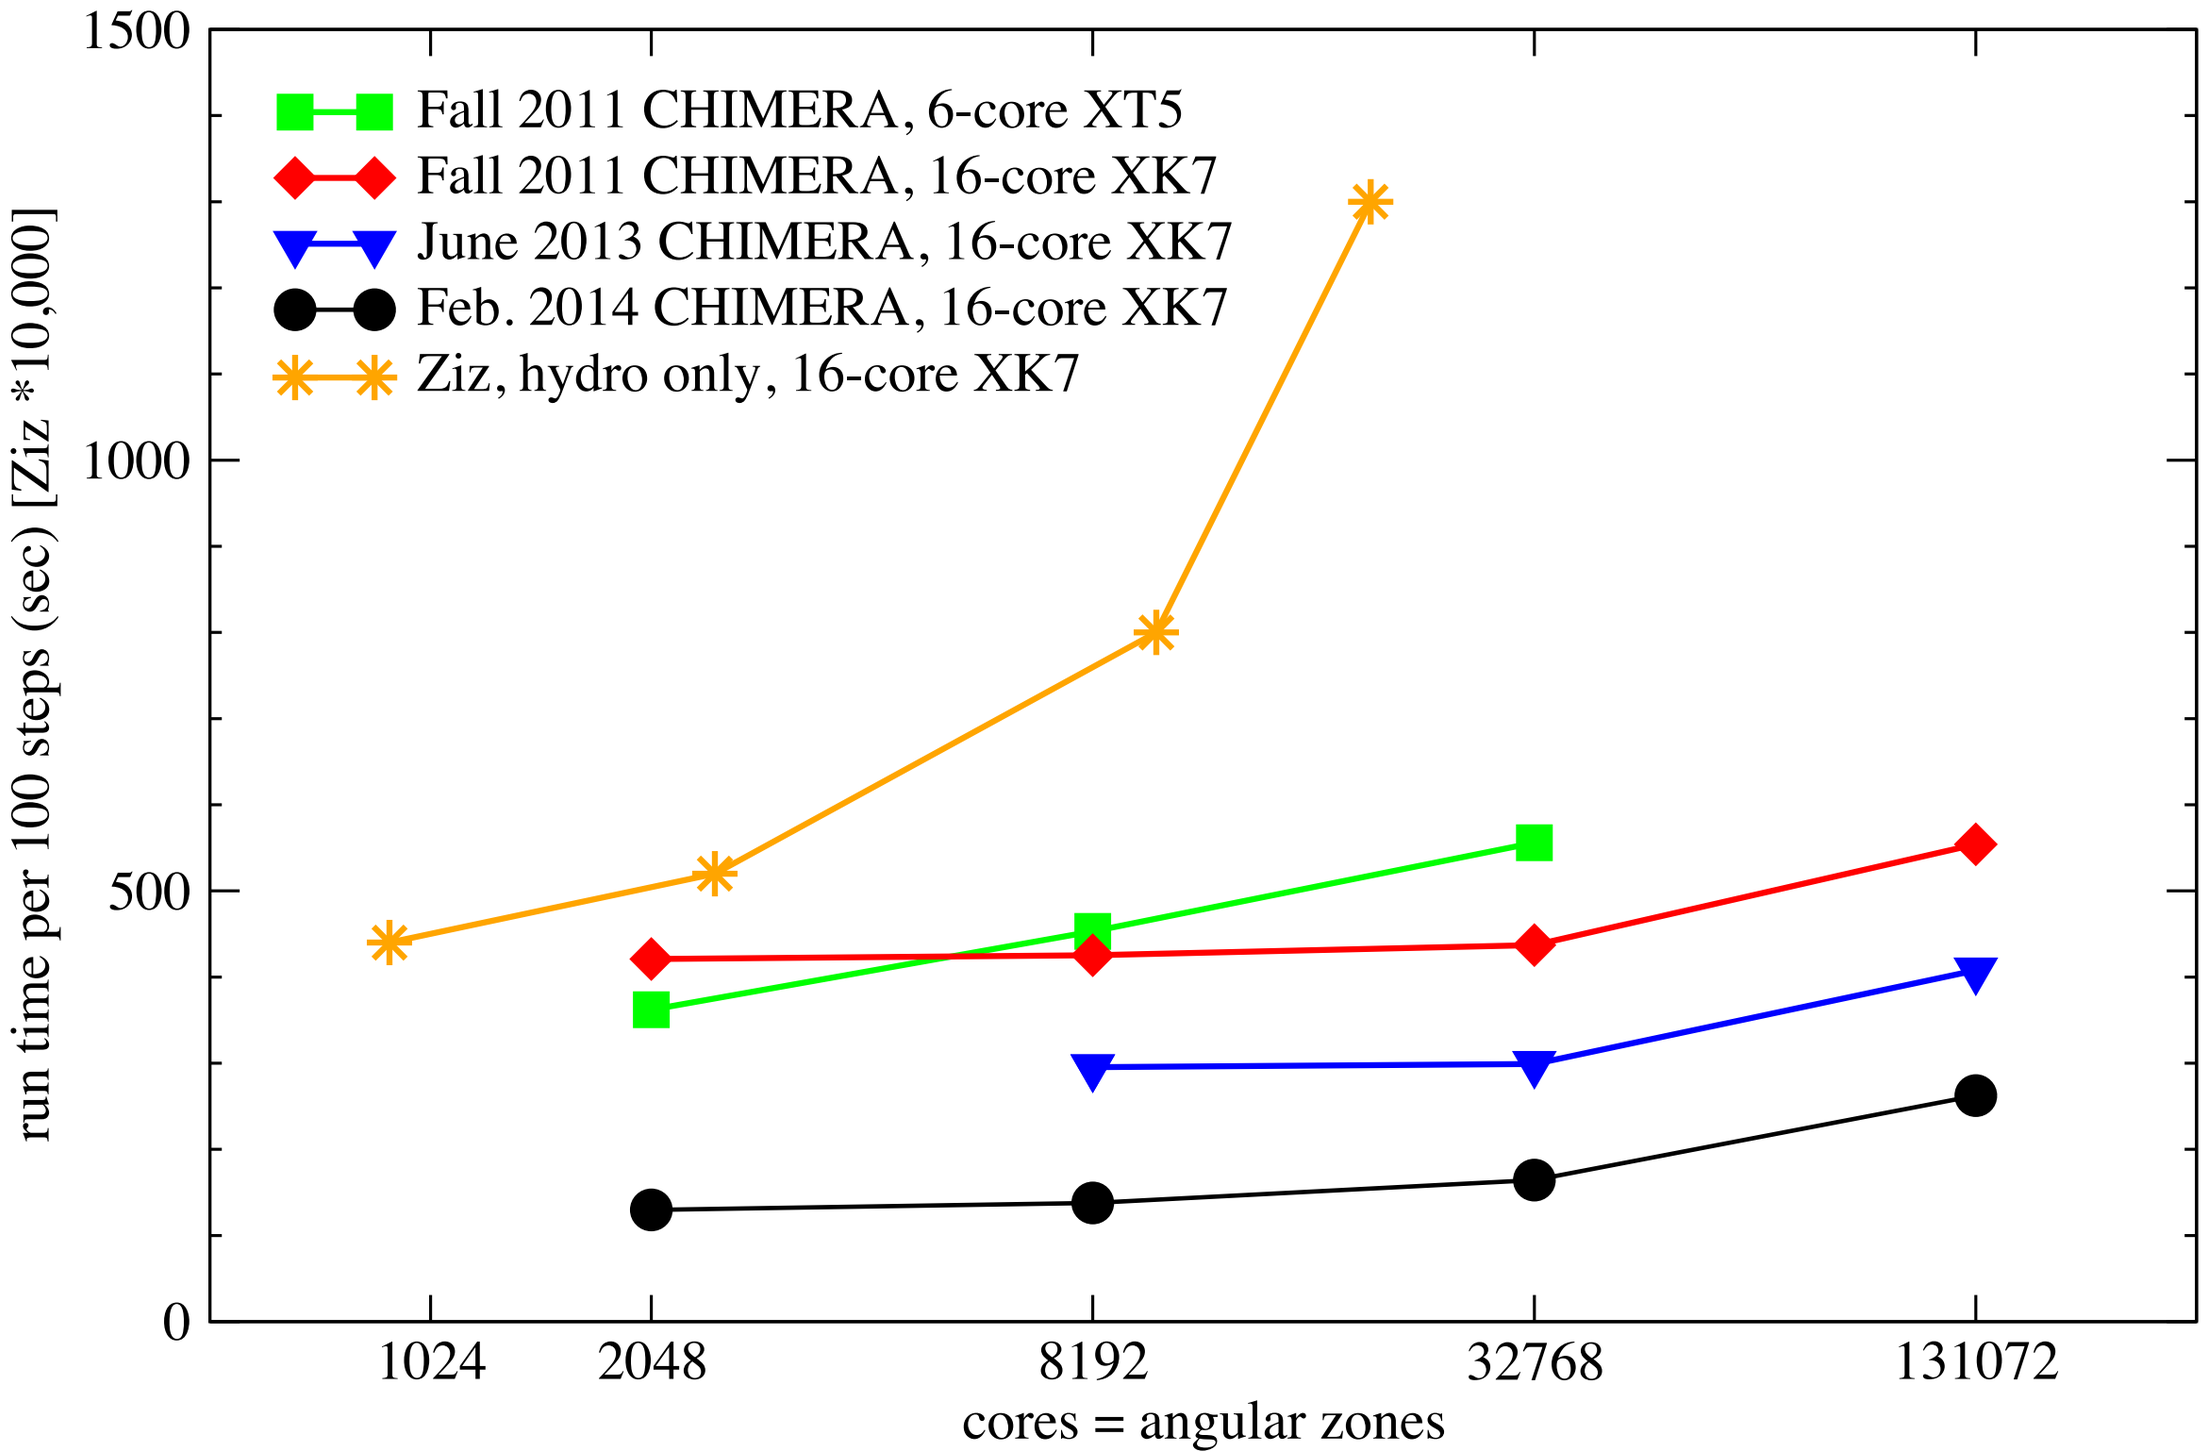

Weak scaling across platforms

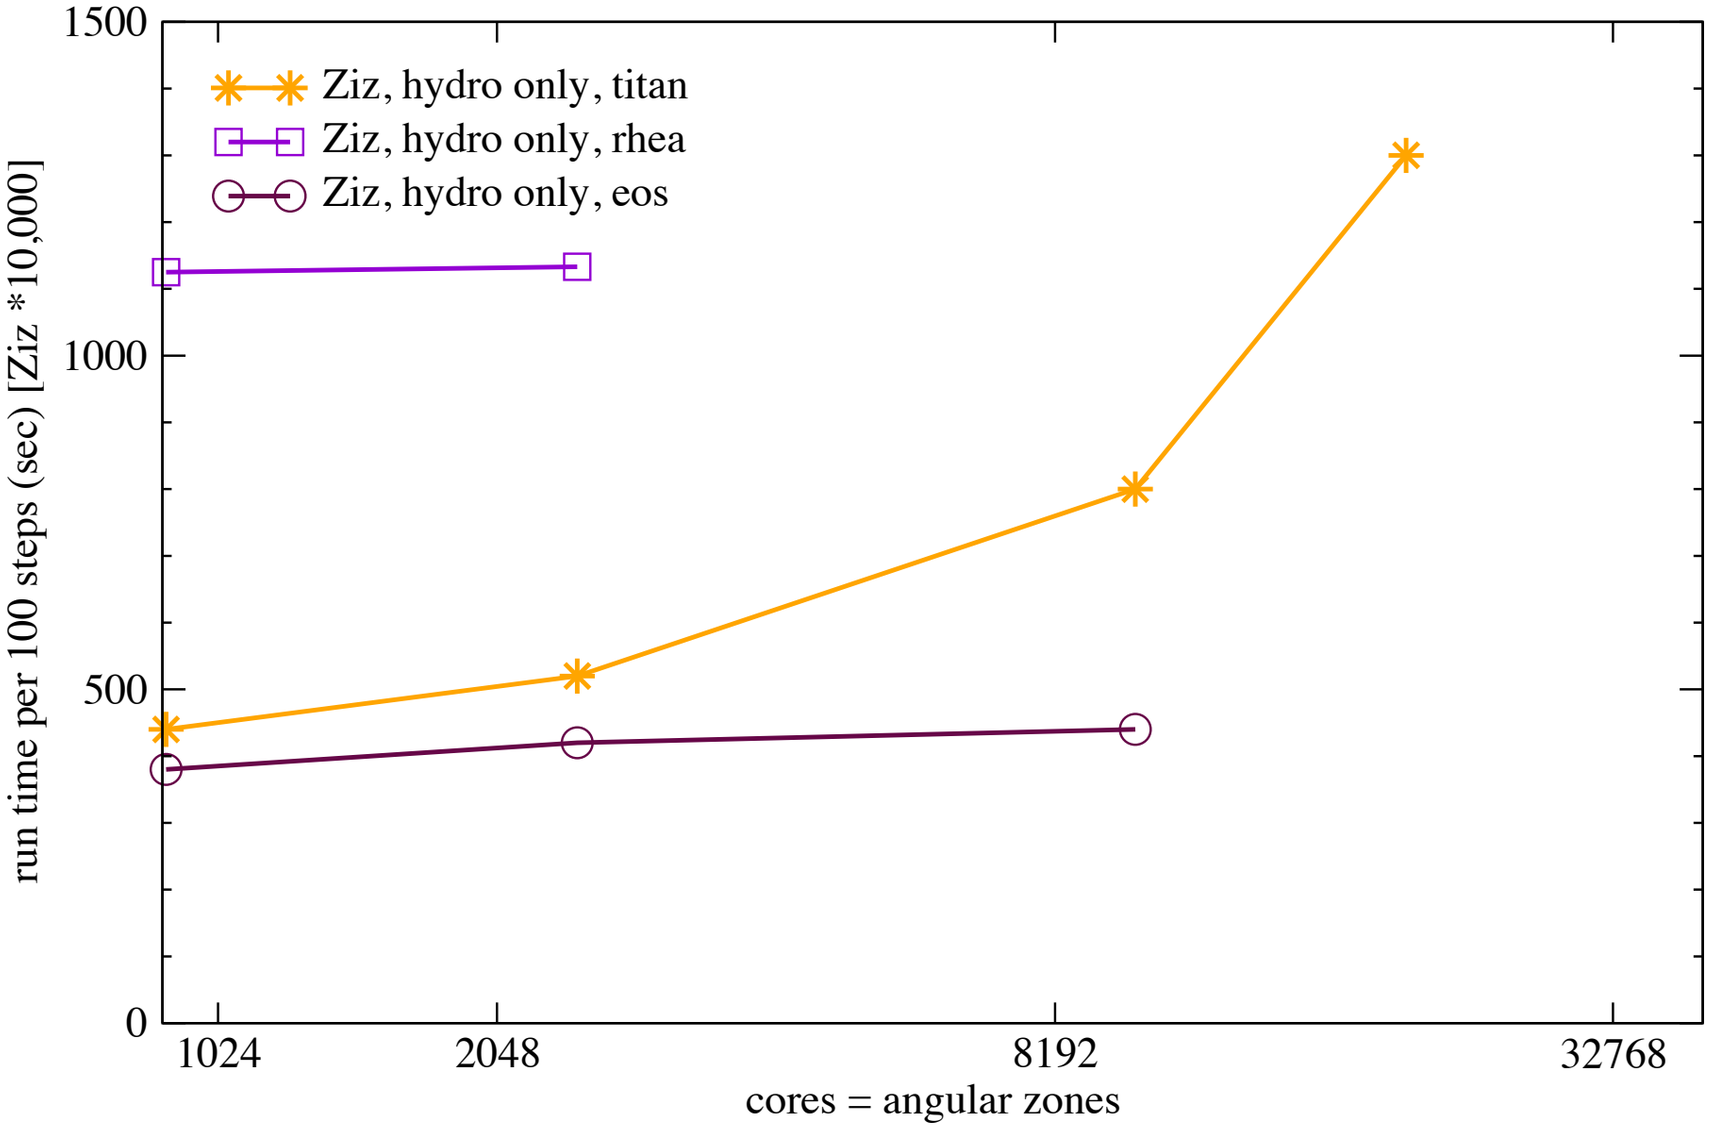

- Also unsurprising...
  - Aries interconnect better for sub-comm collectives
  - Rhea ‘whitebox’ cluster - low overall performance, but weak scaling is correspondingly not too bad

“Faking” computational intensity

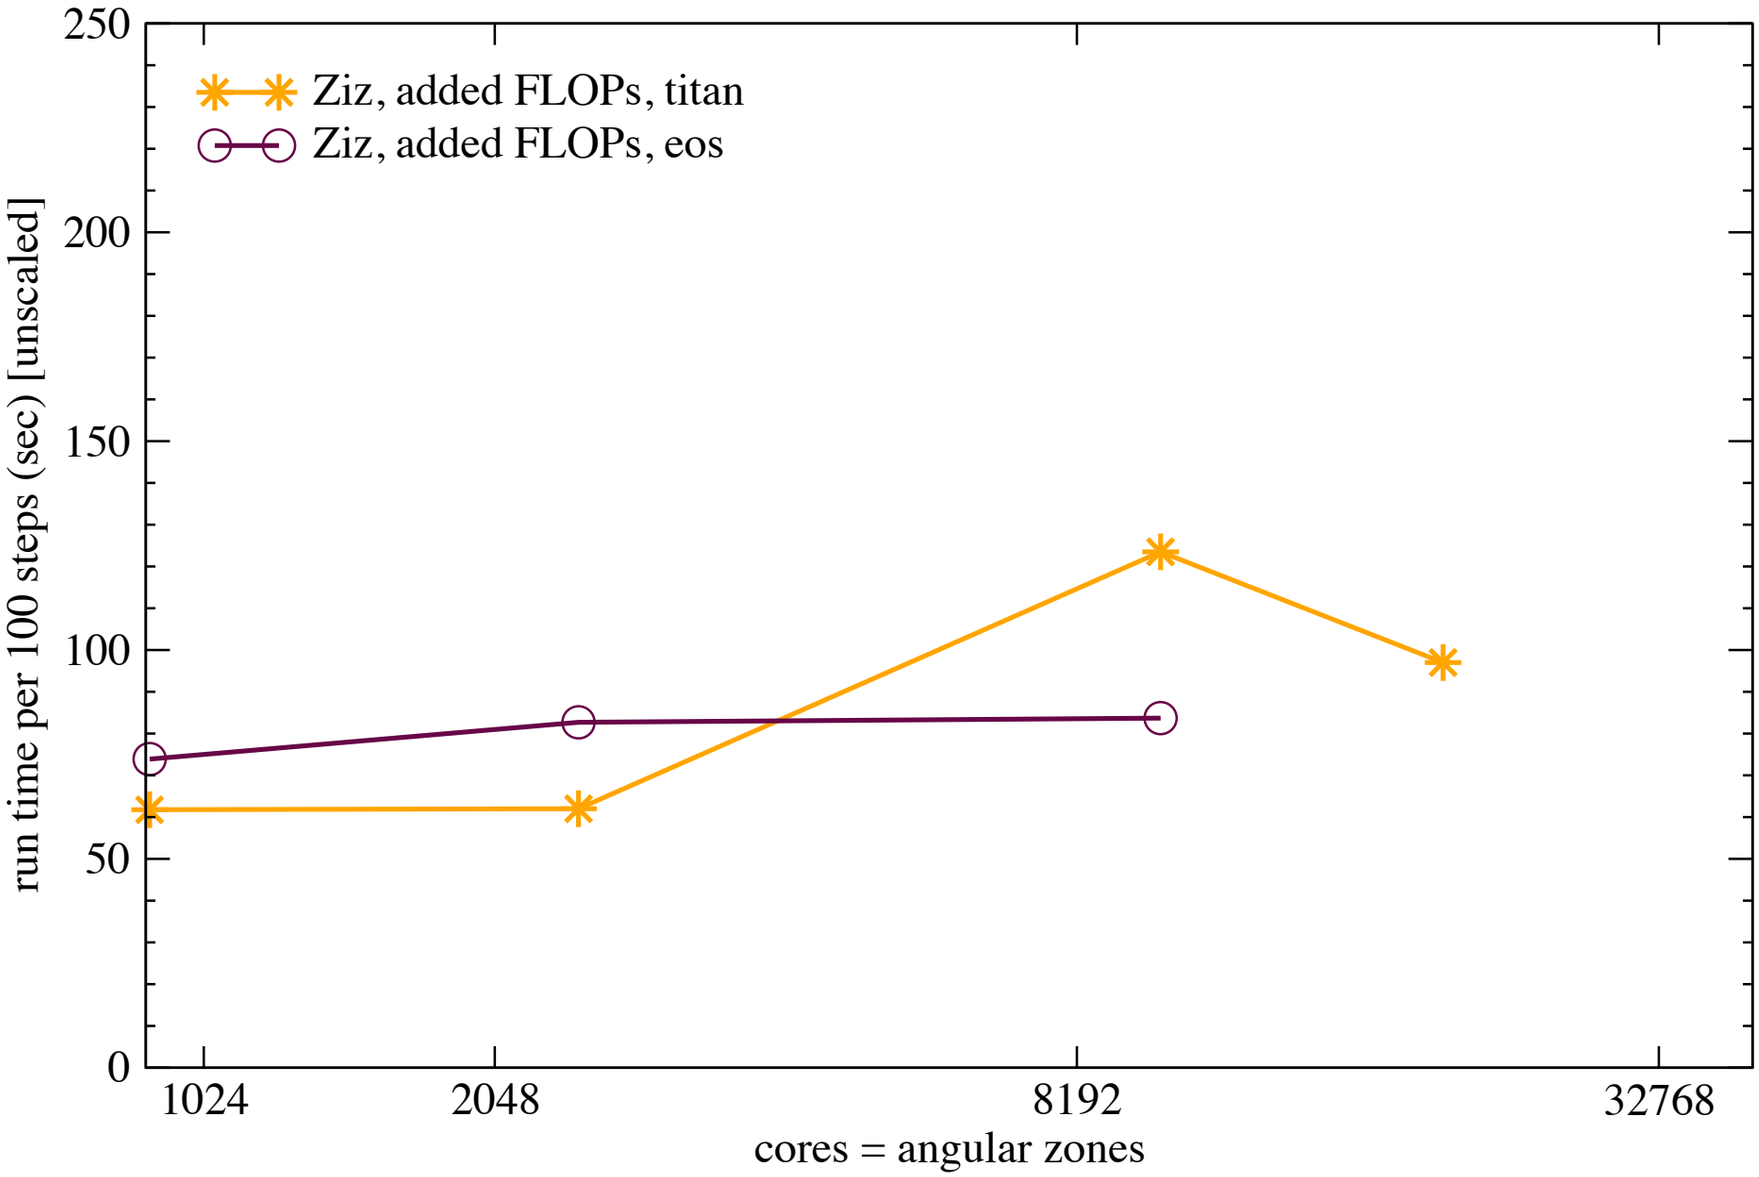

- Loops added after each of sweepx1 and sweepx2
- 70 MFLOPs/zone added - rough sum of transport and burning LU decomps
- 160 REAL*8’s/zone added to MPI payload

What happened?

- Unknown at present, but profiling (with CrayPat) provides a clue...

platform	hydro only	added FLOPs
Titan	`MPI_AllToAll` = 38%, `parabola` = 20%	`MPI_AllToAll` = 45%, `fakeflops` = 45%, `parabola` < 1%
Eos	`MPI_AllToAll` = 30%, `parabola` = 14%	`MPI_AllToAll` = 78%, `fakeflops` = 19%, `parabola` < 1%

- Increased FP performance and increased interconnect bandwidth (vs. Titan[Gemini]) on Eos seems to lead to load imbalance
  - Most (78%) of the `AllToAll` time on Eos @ 900 ranks is in sync; essentially none of that time is sync on Titan

Summary

- Ziz is a MiniApp being developed to model/mimic/suggest the performance of multiphysics codes like CHIMERA
- Interconnect and node-performance differences between platforms lead to expected performance differences in the hydro-only version of Ziz
- Adding superfluous FLOPs to this version leads to results than can seem counterintuitive at first.
  - This suggests that modules that do a better job of actually modeling the FLOPS that are executed in CHIMERA are required
- PS An OpenACC version of Ziz has also been tested (see, again, Levesque's tutorial from this conference). The node-to-node performance enhancement is ~194%.