



**3D Hydrodynamic Simulations of
Rotating, Massive, Main Sequence Stars:
*What Texascale Days Deliver for us***


Paul Woodward, Huaqing Mao, Univ. of Minnesota

Falk Herwig, University of Victoria

8/3/23

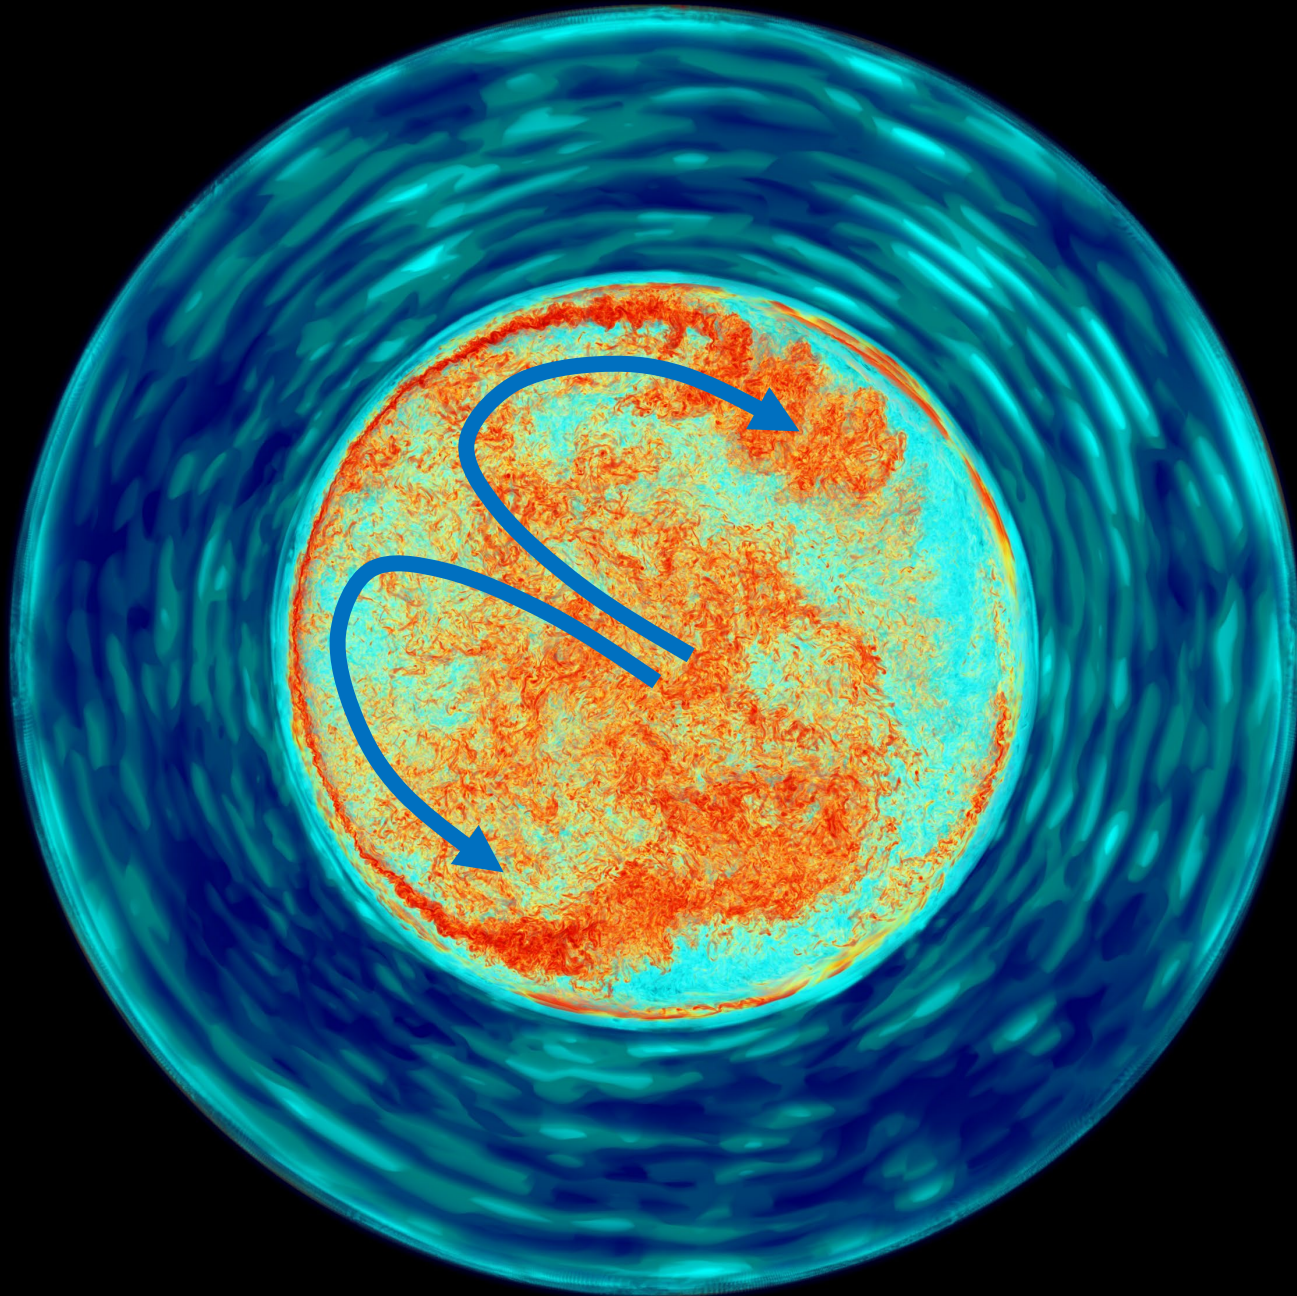
- We are studying stars.
- We have known what they are for 100 years, but there are still issues.
- The general evolution of stars was sorted out in the 1970s, but still issues.
- In the 1920s and 1930s, one could study stars observationally, apply the new theory of the atom and atomic spectra, and make analytic models.
- Nuclear physics provided the tools to understand stellar evolution, but only after the invention of the computer.
- Today, computation is the only way we know to fully understand stars.
- Together with knowledge of the microphysics, 1-D simulations assuming spherical symmetry could be carried out on computers, beginning with the CDC 6600 and 7600 in the late 1960s and 1970s.
- But stars are 3D objects.
- 1-D simulations represent 3-D phenomena via simplified models.
- 3-D simulations in recent years have shown those models to be wrong.
- Frontera is powerful enough to permit 3-D simulations to correct the models

- I will begin with our run M284, a year ago, on a grid of 2688^3 cells.
- $25 M_{\odot}$ star near beginning of the main sequence stage.
- Convective core out to 1517 Mm.
- Luminosity enhancement factor = 1000
Makes the velocities 10 times larger.
Makes the growth of the convection region in radius 1000 times faster.
Brings the average Mach numbers up to about 0.025, still very small.
Leaves the character of the convection flow unchanged.
Like using a small-scale model and wind tunnel to design an airplane.
- Characteristic convective velocity, when non-rotating, $v_{\text{conv}} = 0.01375$
- Mach numbers up to 0.06, so slow flow, but explicit hydro is justified.
- Non-rotating star (run M284) develops global dipole circulation pattern.
- Rotating case influenced by Coriolis and centrifugal forces.




- 25 M_{\odot} star near beginning of the main sequence stage, non-rotating.
- M284, dump 684, thin slice.
- Horizontal component of velocity, perpendicular to radius.
- Convective core to 1517 Mm.
- Characteristic convective velocity, $v_{\text{conv}} = 0.01375$
- Grid of 2688^3 cells
- Global dipole circulation.

In the non-rotating case, a global dipole circulation develops in which the flow strikes the convective boundary and is deflected in all directions. When the star rotates, the portion deflected in the direction of rotation has the greater centrifugal acceleration, and it tends to remain aloft, while the other portion descends.



- 25 M_{\odot} star near beginning of the main sequence stage, non-rotating.
- M284, dump 684, thin slice.
- Vorticity magnitude; smaller in the envelope by factor $< 1/10$
- Convective core to 1517 Mm.
- Characteristic convective velocity, $v_{\text{conv}} = 0.01375$
- Grid of 2688^3 cells
- Global dipole circulation.

In the non-rotating case, a global dipole circulation develops in which the flow strikes the convective boundary and is deflected in all directions. When the star rotates, the portion deflected in the direction of rotation has the greater centrifugal acceleration, and it tends to remain aloft, while the other portion descends.



- 25 M_{\odot} star near beginning of the main sequence stage, non-rotating.
- M284, dump 684, far hemi
- Concentration of entrained gas from above the convective boundary.
- Convective core out to 1517 Mm.
- Characteristic convective velocity, $v_{\text{conv}} = 0.01375$
- Grid of 2688^3 cells

In the non-rotating case, a global dipole circulation develops in which the flow strikes the convective boundary and is deflected in all directions. When the star rotates, the portion deflected in the direction of rotation has the greater centrifugal acceleration, and it tends to remain aloft, while the other portion descends.

*Movie of non-rotating core convection, run M284
Concentration of ingested gas into convection
zone in far hemisphere
Movie by Paul R. Woodward*

- 25 M_{\odot} star near beginning of the main sequence stage, non-rotating.
- M284, dump 684, far hemi
- Concentration of entrained gas from above the convective boundary.
- Convective core out to 1517 Mm.
- Characteristic convective velocity, $v_{\text{conv}} = 0.01375$
- Grid of 2688^3 cells

In the non-rotating case, a global dipole circulation develops in which the flow strikes the convective boundary and is deflected in all directions. When the star rotates, the portion deflected in the direction of rotation has the greater centrifugal acceleration, and it tends to remain aloft, while the other portion descends.

Frontera is powerful enough to permit us to run our non-rotating simulation forward in time for about 3 years, with a luminosity and thermal conduction enhancement of 10,000 times. Velocities are 21 times greater, and the time evolution is sped up by 10,000 times. This is enough to show us what this star wants to do if we only give it enough time.

We will plot terms in the kinetic energy equation, on the right, and in the entropy equation, on the left.

The time derivative terms in both equations are essentially zero.

This allows us to solve for the implied kinetic energy dissipation rate from the kinetic energy equation.

Plugging that into the entropy equation, we can then solve for the radial derivative of the convective heat energy flux.

That in turn allows us to find the convective heat energy flux as a function of radius, and to evaluate it at the convective boundary. If it is not zero there, then this is not a flow in statistical equilibrium.

3D hydrodynamic simulations of massive main-sequence stars. III. The effect of radiation pressure and diffusion leading to a 1D equilibrium model


HUAQING MAO,^{1,2} PAUL WOODWARD,^{1,2} FALK HERWIG,^{3,2} PAVEL A. DENISSEKOV,^{3,2} SIMON BLOUIN,^{3,2} AND WILLIAM THOMPSON³

¹*LCSE and Department of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA*

²*Joint Institute for Nuclear Astrophysics - Center for the Evolution of the Elements, USA*

³*Department of Physics and Astronomy, University of Victoria, Victoria, BC, V8W 2Y2, Canada*

Our work on core convection in this non-rotating star is written up in this article submitted to the *Astrophysical Journal*, and it basically contains the bulk of graduate student Huaqing Mao's Ph.D. thesis work.



Plots of the radial dependence of the portions of the total energy flux arising from conduction (radiation diffusion) and convection. These add up to the total luminosity arising from nuclear reactions. Over the 3-year course of this very long duration simulation, M252, at modest resolution (896^3 cells), a small additional region of convective penetration is established beyond the classic theoretical extent of the convection zone, marked by the broad vertical line. This development has a significant impact on the evolution of the star.

$$\begin{aligned} \frac{\partial}{\partial t} \overline{\left(\frac{1}{2}\rho u^2 4\pi r^2\right)} &= \overline{(\rho_1 \mathbf{u} \cdot \mathbf{g} - (\mathbf{u} \cdot \nabla)p_1) 4\pi r^2} \\ &\quad - \frac{\partial}{\partial r} \overline{\left(\frac{1}{2}\rho u^2 u_r 4\pi r^2\right)} - 4\pi r^2 \overline{\Phi} \end{aligned} \quad (22)$$

Assume dynamical equilibrium, so that the time derivative on the left vanishes, and then use this kinetic energy equation to solve for the kinetic energy dissipation rate.

$$\begin{aligned} -4\pi r^2 \overline{\Phi} &= \frac{\partial}{\partial t} \overline{\left(\frac{1}{2}\rho u^2 4\pi r^2\right)} - \overline{(\rho_1 \mathbf{u} \cdot \mathbf{g} - \mathbf{u} \cdot \nabla p_1) 4\pi r^2} \\ &\quad + \frac{\partial}{\partial r} \overline{\left(\frac{1}{2}\rho u^2 u_r 4\pi r^2\right)}. \end{aligned} \quad (24)$$

$$\frac{\partial}{\partial t}(\overline{\rho S}4\pi r^2) = 4\pi r^2 \frac{\overline{\Phi}}{T} + \overline{\frac{1}{T} \nabla \cdot (\mathbf{\Gamma} - \mathbf{F})}4\pi r^2 - \frac{\partial(\overline{\rho S u_r}4\pi r^2)}{\partial r} \quad (23)$$

Assume thermal equilibrium, so that the time derivative on the left vanishes, and then use this entropy equation to solve for the convective entropy flux.

$$\frac{\partial(\overline{\rho S u_r} r^2)}{\partial r} 4\pi = \frac{\overline{\Phi}}{T} 4\pi r^2 + \frac{1}{T} \frac{\partial(F_{\text{nuc}} - F_{\text{rad}})}{\partial r}. \quad (28)$$




Figure 22. The gradient of energy flux from nuclear heating and radiative diffusion, gradient of the advective entropy flux multiplied by temperature, rate of change in entropy multiplied by temperature, dissipation of kinetic energy measured at dump 6000 (**17188 h**) of run M252. Note the small role played by turbulent kinetic energy dissipation relative to the other terms in the entropy equation.





Figure 18. Top: work by pressure gradient and gravity field per unit time per unit radial distance, rate of change in kinetic energy per unit radial distance, radial derivative of total kinetic energy flux, dissipation derived from the turbulent dissipation model, measure dissipation rate per unit radial distance, averaged over 399 dumps (~ 1145 h) centered at dump 6000 (~ 17188 h) of M252; bottom: time



The character of the convection flow changes significantly as both dynamical (left) and thermal (right) equilibrium are approached over this very long duration simulation. Our analysis allows us to predict the final equilibrium.

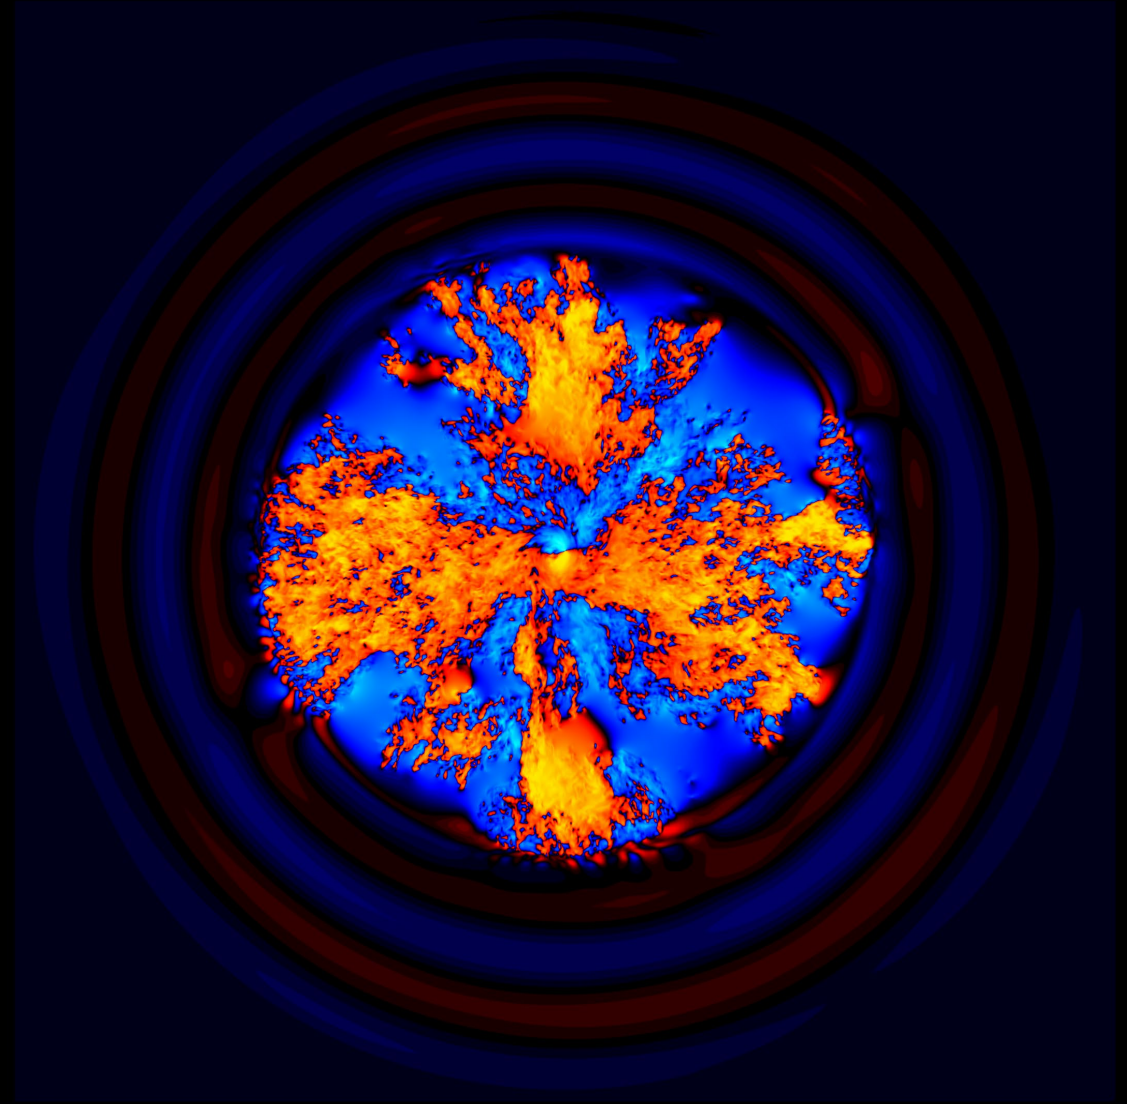

*Movie of non-rotating core convection, run M252
at very late times, run from dump 7000 with a
high dump frequency as run M253,
vorticity magnitude.
Movie by Paul R. Woodward*

- Note the separation of the turbulent convection region from the convective boundary.
- This separation is caused by a thin layer of gas that is continually heated by the thermal conduction flux as it decreases with radius to return to the level of the full luminosity from above.
- The added heat is constantly mixed away by the convection, which in this thin region carries heat inward rather than outward.

- I will now let this same star rotate initially as a solid body.
- Characteristic convective velocity, when non-rotating, $v_{\text{conv}} = 0.01375$
- Angular velocity of rotation = 2.86572×10^{-5} / sec.
- Rotation period = 2.54 days
- Equatorial rotation velocity at the convective boundary = $3.1623 \times v_{\text{conv}}$
- Non-rotating star (run M284) develops global dipole circulation pattern.
- Rotating case influenced by Coriolis and centrifugal forces.
- We do the calculation in a rotating frame, with the original, non-rotating, spherically symmetric star subtracted out of the nonlinear dynamical equations. We compute perturbations to this state.
- We spin up the non-rotating star gradually, to avoid exciting pulsations, before the convection gets going.
- The star rotates clockwise, in our view along the spin axis, and the coordinate system rotates with it.


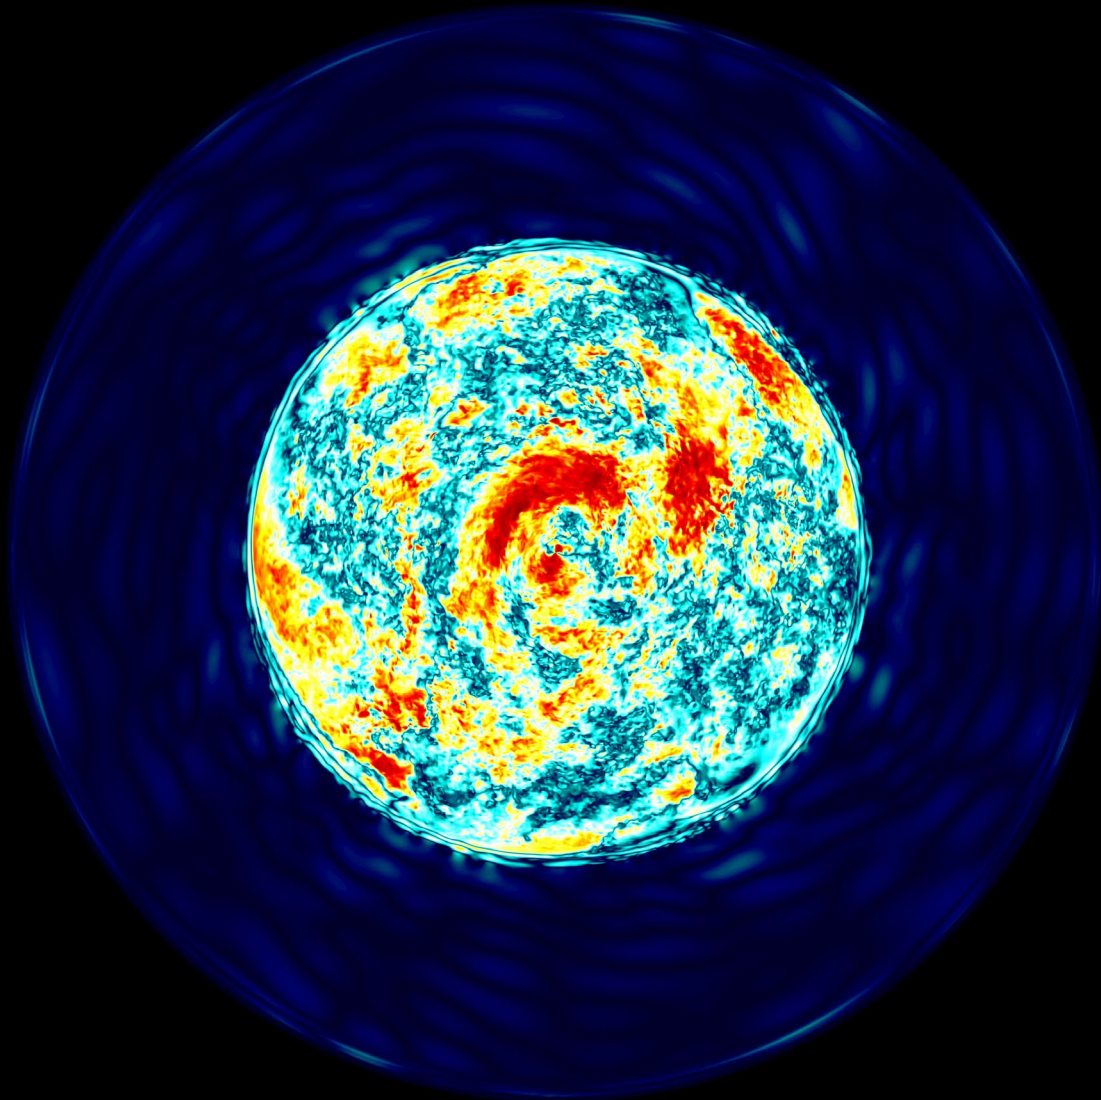
In the following slides, we see still images taken from movie animations. These give us time to absorb the relevant fluid behaviors we see in the movies. We can consider these slides of still images of 4 different variables as a training set that can help us to interpret the rapidly changing images in the movies. We first see matched pairs of images of the magnitude of the horizontal component of velocity (the component tangent to the sphere at that radius) and of the radial component of velocity. For the horizontal component of velocity, the magnitude increases as colors go from black to dark blue to aqua (cyan) to white to yellow and finally to red. For the radial velocity, red and yellow are outward and dark blue and aqua are inward, with black representing zero radial velocity. We then see matched pairs of images with vorticity magnitude in a 10% equatorial slice on the left and the log of the concentration, FV, of ingested envelope gas in the top hemisphere on the right. In the FV images the low concentrations in the center of the star are effectively transparent.

100



We begin with a Cartesian grid imprint, with gas rising along the 4 cardinal directions and sinking in the other 4.
Tangential velocity magnitude in equatorial slice; Radial velocity

120





The grid imprint is by now largely gone.

Tangential velocity magnitude in equatorial slice;

Radial velocity

140





A large outflow has developed striking the convective boundary at 3 o'clock. This generates flows along the boundary.

Tangential velocity magnitude in equatorial slice;

Radial velocity

160





The horizontal flows along the boundary are seen here in the red regions in the left-hand image.

Tangential velocity magnitude in equatorial slice;

Radial velocity



180



Here the opposing horizontal flows along the boundary are colliding, and the one that is prograde forces the other inward.

Tangential velocity magnitude in equatorial slice; Radial velocity

200


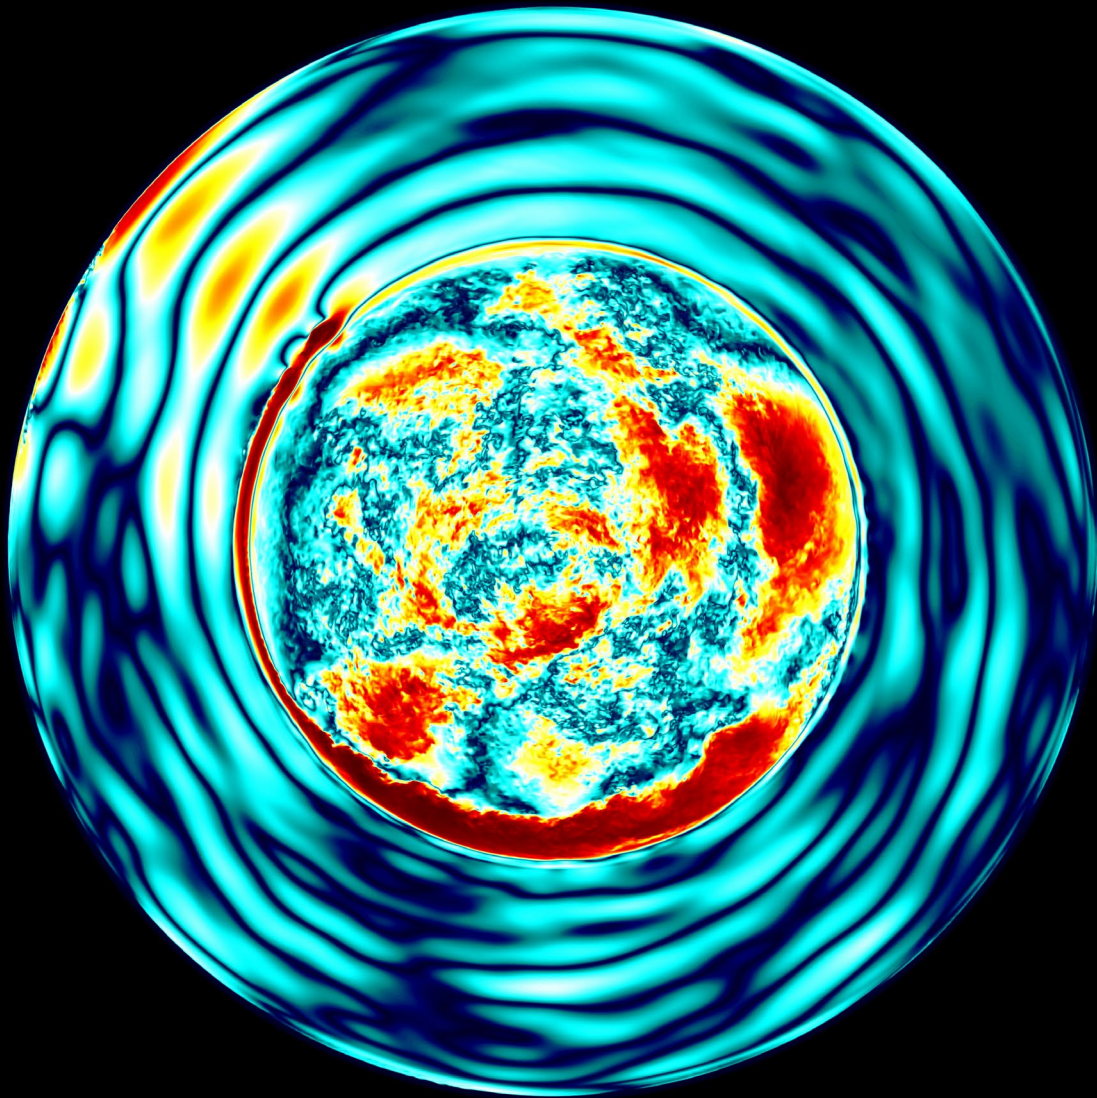


The internal gravity waves in the envelope gas have strengthened considerably. These can carry angular momentum.

Tangential velocity magnitude in equatorial slice;


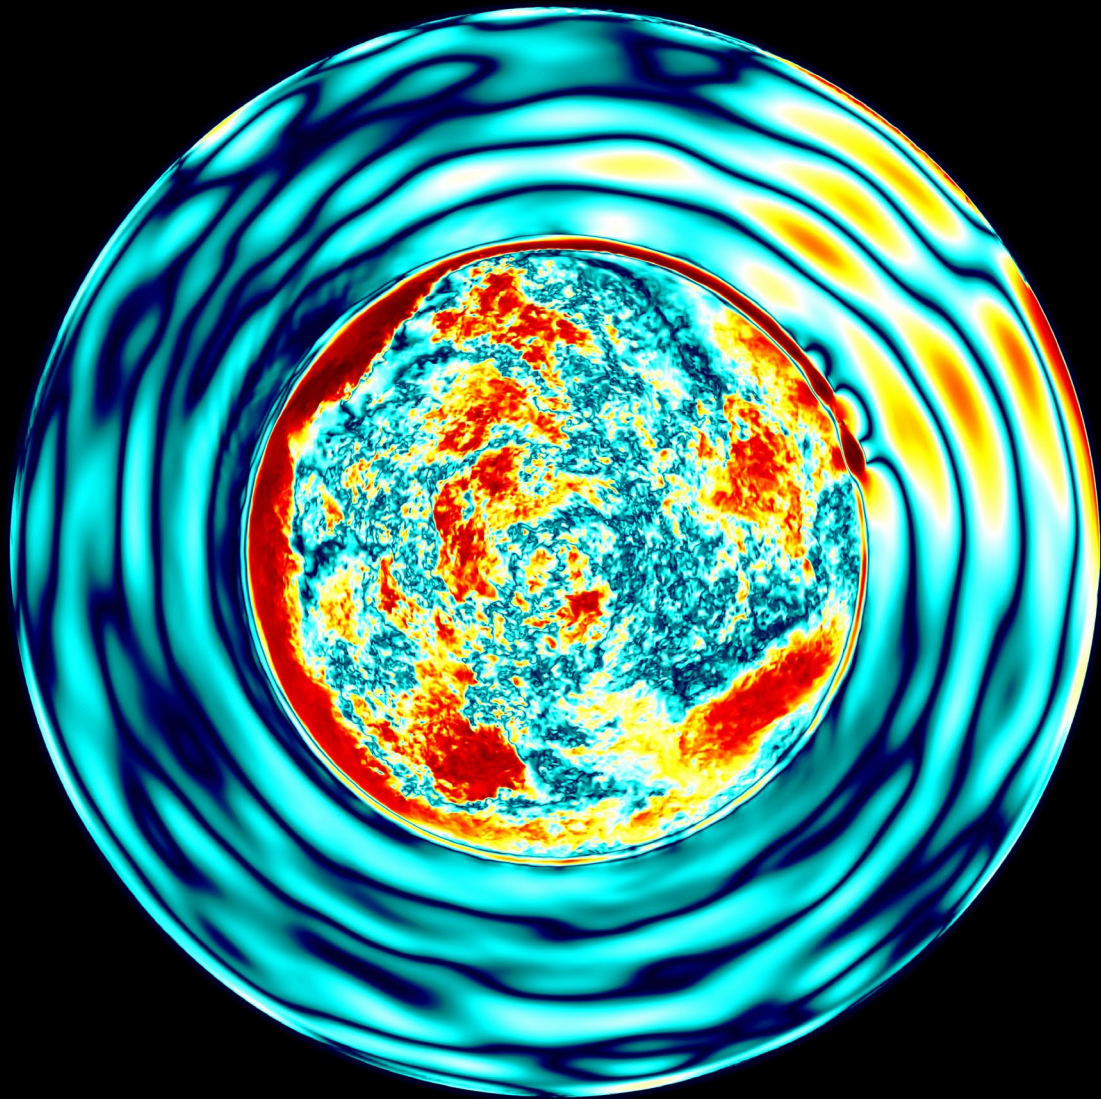
Radial velocity

220



The gravity waves are very strong now and are aligned with a tsunami-like interfacial gravity wave at 10 o'clock.
Tangential velocity magnitude in equatorial slice; Radial velocity

240





The strong gravity waves have a special signature in the radial velocity.

Tangential velocity magnitude in equatorial slice;

Radial velocity



260



Tangential velocity magnitude in equatorial slice;

Radial velocity

1020





The gravity waves have quieted down, and a general differential rotation has developed.

Tangential velocity magnitude in equatorial slice;

Radial velocity



100



Vorticity magnitude in equatorial slice;

Concentration of Entrained Gas

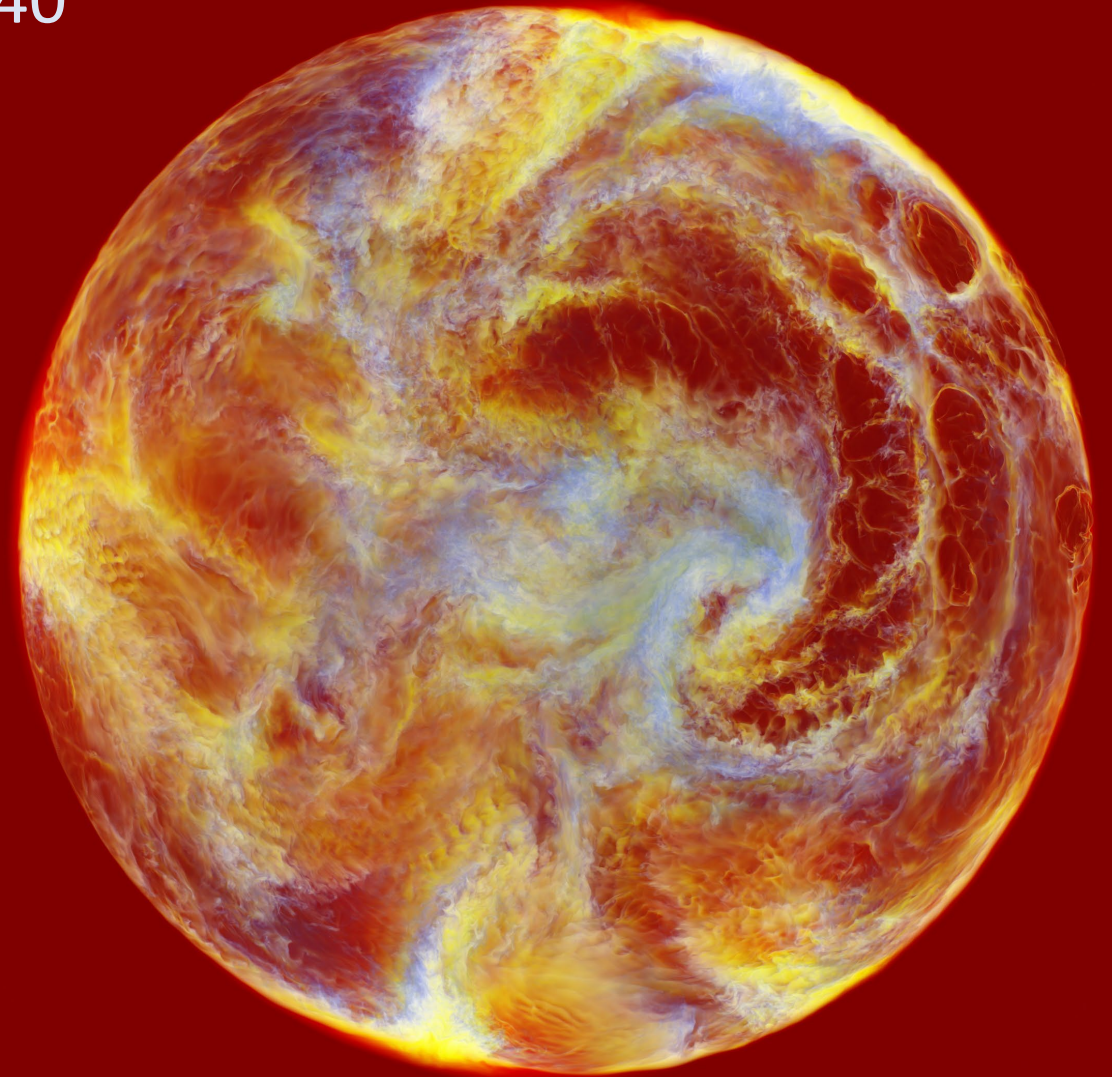

120



Vorticity magnitude in equatorial slice;

Concentration of Entrained Gas



140



Vorticity magnitude in equatorial slice;


Concentration of Entrained Gas

160




Vorticity magnitude in equatorial slice;

Concentration of Entrained Gas





180



Vorticity magnitude in equatorial slice;

Concentration of Entrained Gas



200



Vorticity magnitude in equatorial slice;

Concentration of Entrained Gas



220



Vorticity magnitude in equatorial slice;

Concentration of Entrained Gas



240



Vorticity magnitude in equatorial slice;


Concentration of Entrained Gas

260

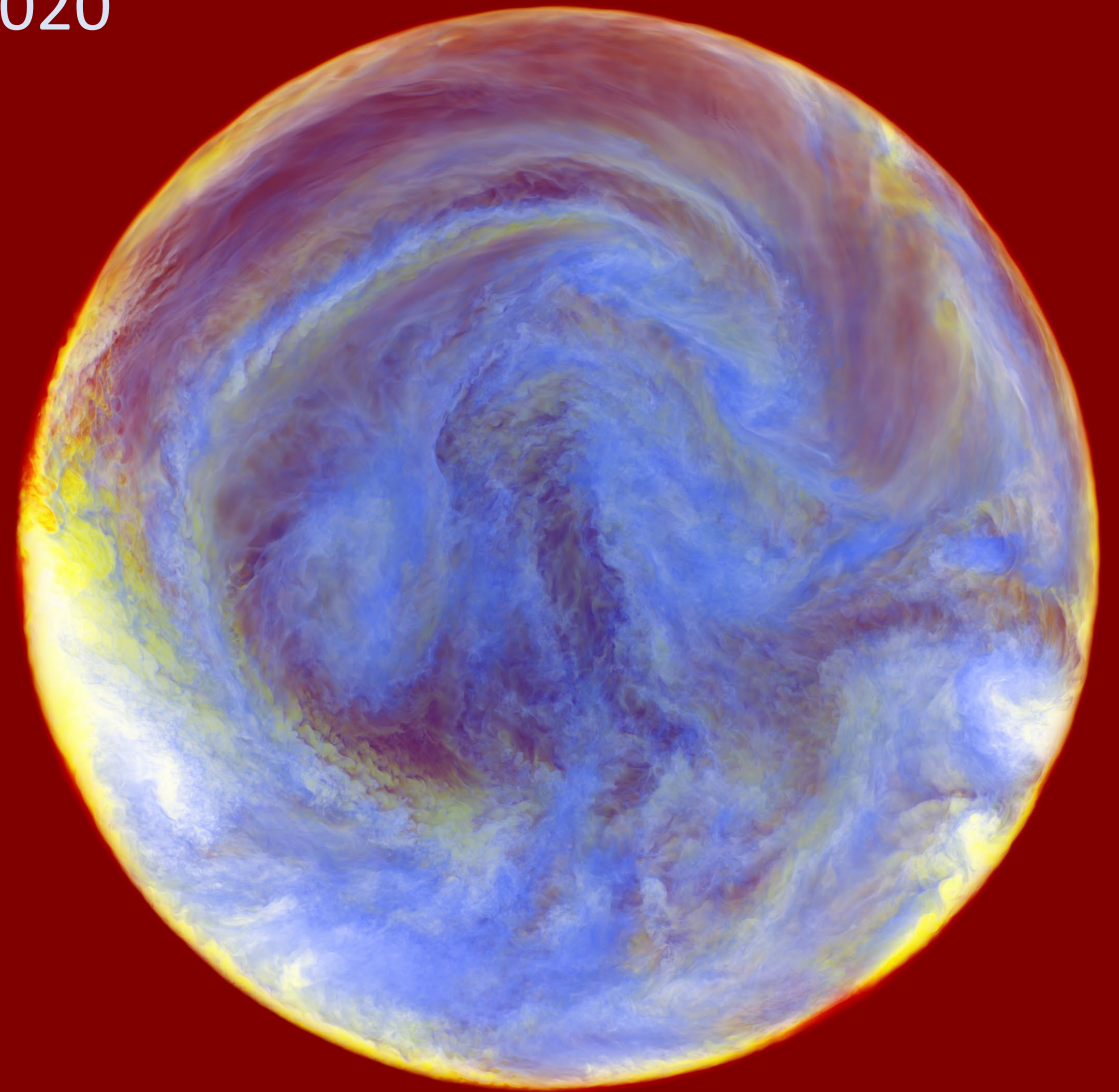


Vorticity magnitude in equatorial slice;

Concentration of Entrained Gas



1020



Vorticity magnitude in equatorial slice;


Concentration of Entrained Gas

*Movie of rotating core convection, run M163
FV far hemisphere close-up view.
Movie by Ben McDermott & Paul R. Woodward*

Run M163, volume rendering of the concentration of gas from above the convection zone that is ingested into that zone. As the log of this concentration increases from -4 to 0 , colors range from black, dark blue, lighter blue, aqua, white, yellow, and finally red. We look through very low concentrations near the center of the star to see the higher concentrations near the bowl-shaped convective boundary. Spin axis vertical.

*Movie of rotating core convection, run M163
FV top hemisphere close-up view.
Movie by Ben McDermott & Paul R. Woodward*

Run M163, volume rendering of the concentration of gas from above the convection zone that is ingested into that zone. As the log of this concentration increases from -4 to 0 , colors range from black, dark blue, lighter blue, aqua, white, yellow, and finally red. We look through low concentrations near the center of the star to see the higher concentrations near the bowl-shaped convective boundary. Spin axis into screen.




Run M163, dump 800

Rotating at 3.16 times the average velocity in the convection zone, this simulation produces only a suggestion of the Taylor columns we might expect.

Here we see a volume rendering of the axial component of the velocity, with red and yellow indicating upward motion and blue and aqua indicating downward motion.

Motions in the envelope gas seem relatively unaffected.

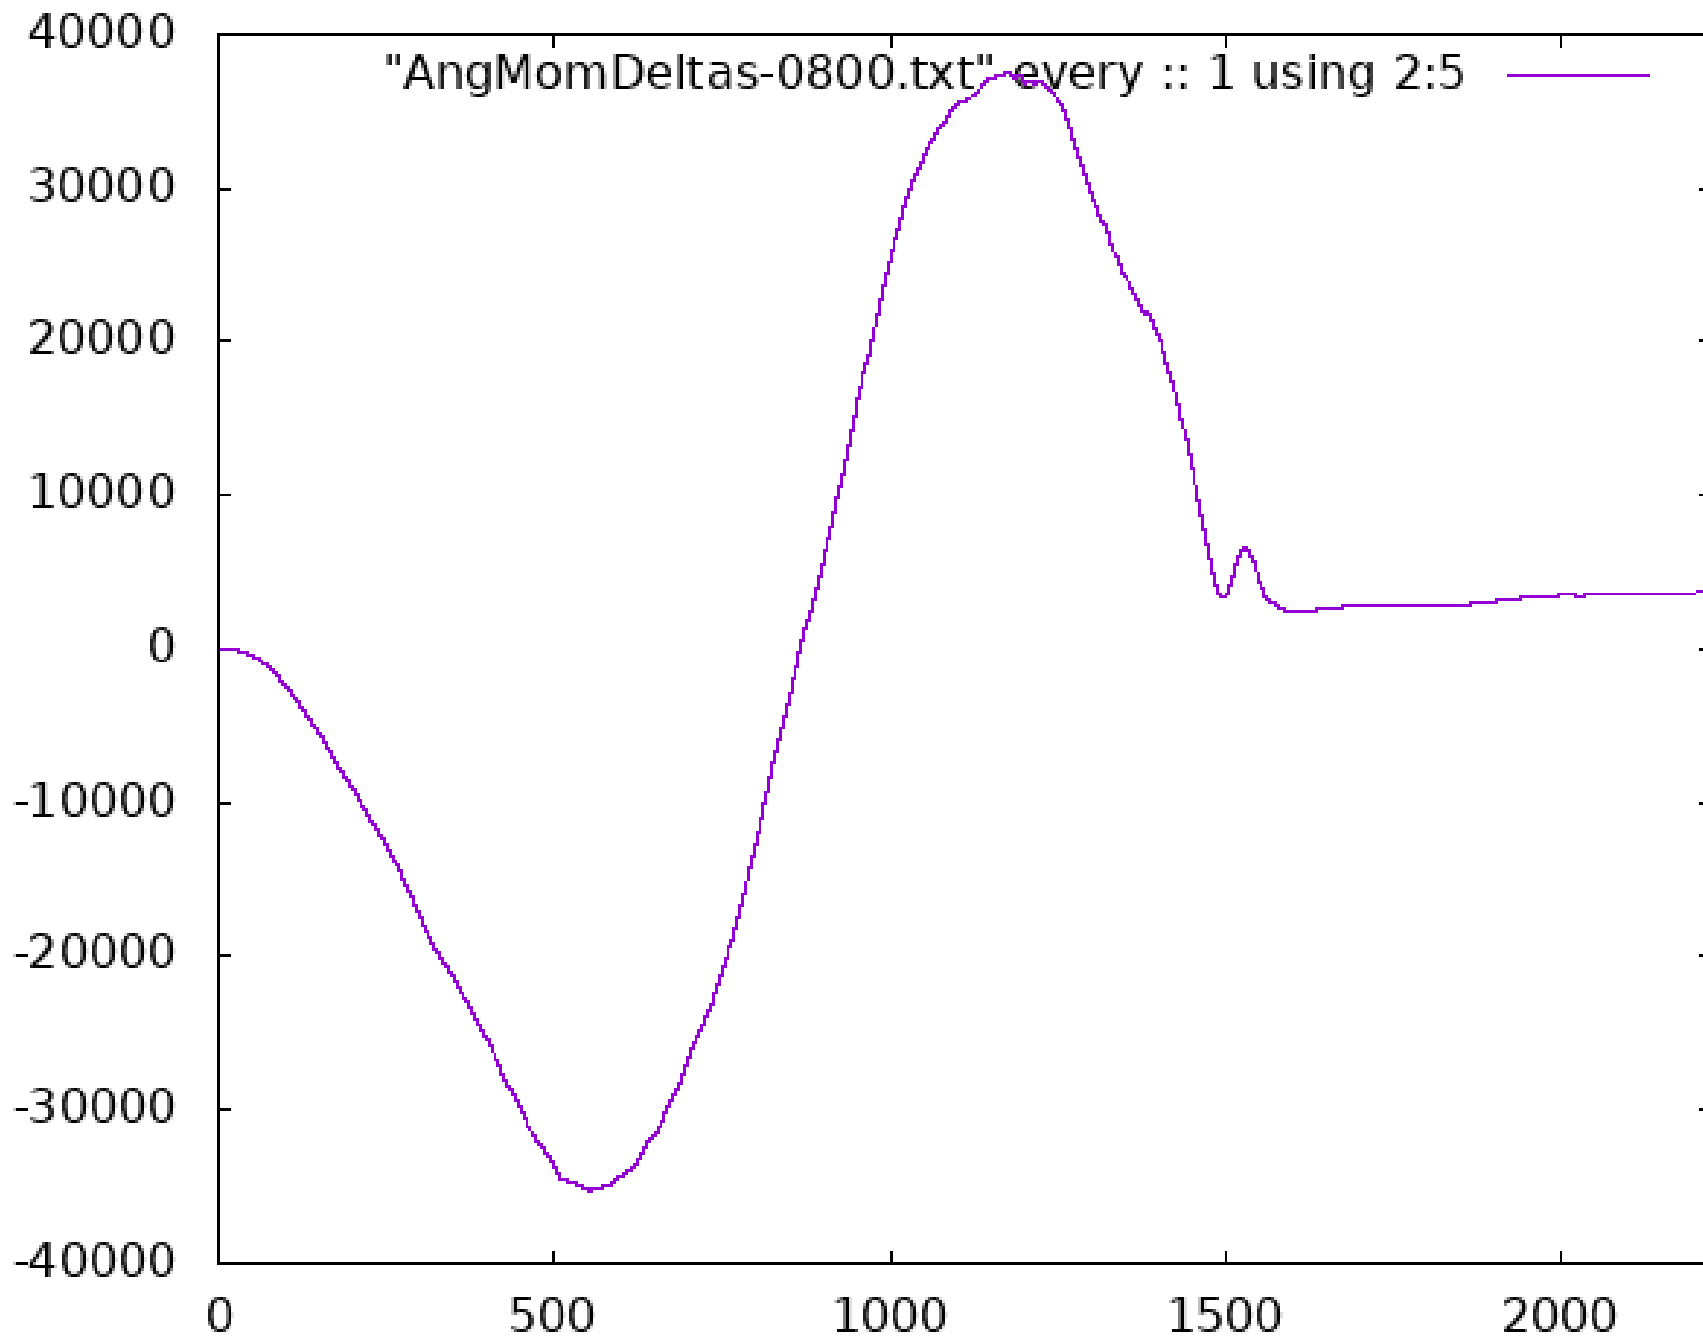


Run M137, dump 935

Rotating at 10 times the average velocity in the convection zone, this simulation produces the Taylor columns we would expect.

Entrainment of gas from above the convection zone occurs principally near the poles.


Overall, gas entrainment is reduced from the non-rotating case by about a factor of 3.



Here we see the angular momentum changes in the material of the entire star, in a mass coordinate, plotted against the original cylindrical radial location for run M163.

There is very significant transfer of angular momentum inside the convection zone, and also a small amount of outward transfer of angular momentum in the envelope.

delAngMom vs. R dump 500 to 1900 (AngMomDeltas-*.txt) Mass Coord



Here we see the angular momentum changes in the material of the entire star, in a mass coordinate, plotted against the original radial location for run M137, which is rotating 3.1623 times faster than M163. The behavior is very similar.

PPMstar Code Performance on Frontera:

John Cazes asked me to discuss how we made PPMstar run so well on Frontera.

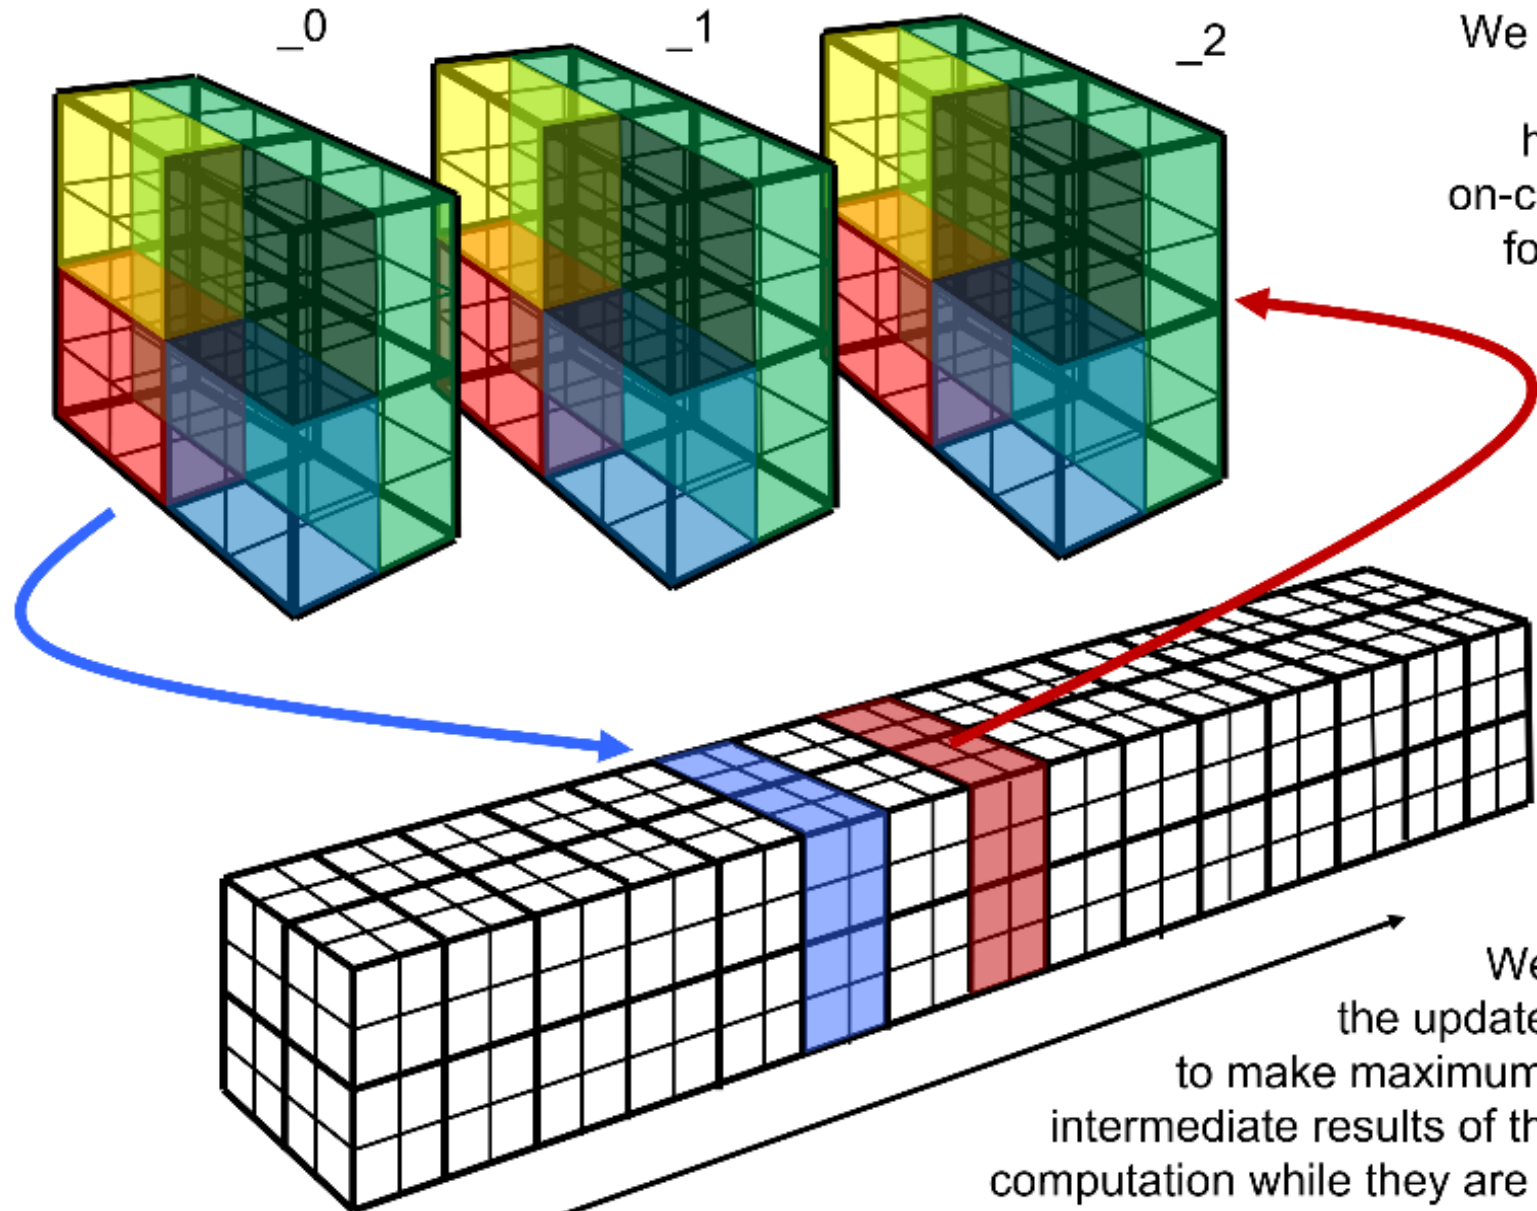
- **32-bit precision** throughout, enabled by code's perturbation form and PPB moment-conserving advection.
- **Overlapping MPI message passing** completely with computation.
Small number of large messages, sent early and received late.
Update half the star while messages for the other half travel.
- **Briquette data structure.**
All variables for blocks of 4^3 cells stored together in a *briquette record*.
- **Pipelined update of half-briquettes.**
Loop consisting of about 2300 SIMD-32 statements reads half-briquette in and writes out an updated half-briquette.
- **Explicit prefetching of 32-word vectors.** Important to avoid memcopies.
- **Explicit in-lining of subroutines with consolidated stack.**
Produces minimal data footprint on the chip. No pointless traffic to off-chip memory.
- **PPMstar is not memory-bandwidth-limited on any hardware.**
- **We zip all large data files in memory before writing them out to disk.**
This eliminates crashes upon restart dumps and keeps our disk usage down.

A Bit More on I/O in PPMstar – a Set of Techniques Anyone can Implement:

John Cazes asked me to discuss how we made PPMstar run so well on Frontera.

- We studied storage and I/O (and vis) in my lab, the LCSE, at Minnesota for 20 years.
- I went to Minnesota in 1985 with my code writing compressed dumps.
I compressed every saved number to 16 bits, passing it through a nonlinear map first.
- In 2005, I gave up on compressed dumps and instead wrote out vis dumps, with a single byte per variable per grid cell, plus doubly compressed “moms” dumps, with 32 variables saved with 16 bits each representing a briquette of 4×4×4 cells.
- Restart dumps are much less frequent, and in principle one needs only 2 of these per run
- Beginning with my code for Roadrunner at Los Alamos, I built into PPMstar an entire workflow, including the kind of I/O that a decade of working with disk arrays taught me.
- As soon as I initialize MPI, each rank calls a routine called “findmyself.”
- Based on the MPI rank, the process discovers if it is a worker, a team leader, or a team timekeeper. A team is usually 130 MPI ranks, with 128 workers and 2 administrators.
- The time keepers do global reductions as the workers work uninterruptedly.
- The team leaders collect reports and write them to disk.
- Reports are pictures of a worker’s bricks and radial profiles of many, many quantities.
- The team leaders are not allowed to hold up the workers in their work.

REFERENCE slide




We work with at most 3 half-bqs in on-chip cache for any one variable.

A single task for one thread of control is structured so that access to off-chip memory is minimized. A thread reads in a complete, densely packed data record for 2 planes of 16 cells, computes everything, and writes the new data record back. All the hundreds of intermediate results are held in cache on the CPU chip. No problem on any CPU. Big problem on any GPU. Context is 389 vectors for the simulation in the movie. Register file limited to 255. In future, we will need still more than our 389 today.

We pipeline the update process to make maximum reuse of intermediate results of the update computation while they are in cache.

This is how we now accommodate 32-wide SIMD processing. Do 2 planes at once.

REFERENCE slide



We work with at most 3 half-bqs in on-chip cache for any one variable.

A single thread works out of the L1 and L2 caches. It processes a grid pencil, which takes 32 trips through an update loop of some 2300 SIMD-32 f-ops. Each 1-D pass does everything, all the hydro, multifluid advection, radiation diffusion, and, potentially, nuclear burn. For the run shown, each MPI process updates 4 grid bricks of 16x16 grid pencils, and there are 8 such processes per node. Would prefer 2 MPI processes per node with each having 2 grid bricks of 32x32 grid pencils. Best would be smaller problem on same hardware

We pipeline the update process to make maximum reuse of intermediate results of the update computation while they are in cache.

This is how we now accommodate 32-wide SIMD processing. Do 2 planes at once.

Traditional Vectorization: Each thread of control updates an entire grid pencil. Vector lengths are moderate, ~ 128 , do each operation – EOS, interpolation, Riemann solves, conservation laws – over whole pencil of 16×128 cells.

Entire calculation fits into L2 cache, or at least L2+L3. Requires data transpose

Pipelined Sequence of Transverse-Vectorized Operations: Each thread of control updates an entire grid pencil, but it does so by entirely updating just one half-briquette of 32 cells at a time. Vector length is 32. Entire on-chip data context is ≤ 64 KB. *Does not fit into L1 cache, but does fit into L2.*

Briquette data records read optimally & transposed entirely within L1 cache.

Parallel Pipelined Sequences of Transverse-Vectorized Operations: Wow!

We are developing this for exascale hardware.

Teams of 4 or 8 threads work cooperatively out of an on-chip cache to jointly update a single grid pencil.

Entire on-chip data workspace is between 64 and 128 KB.

In principle, cost is one extra half-briquette record update equivalent/pencil

Traditional Vectorization: Works fabulously if you have time-traveled to 1992.

Pipelined Sequence of Transverse-Vectorized Operations: Works fabulously.

We developed this mode for LANL's Roadrunner and the CELL processor.

It works extremely well on all CPUs we have ever used.

For complicated grid cell updates, 500 to 1000 vector temporaries are required to reside in cache, with minimum lengths of 16 or 32 words on today's equipment.

On a CPU, there is enough room. On a GPU, there is not.

Parallel Pipelined Sequences of Transverse-Vectorized Operations: Wow!

We are developing this for exascale hardware.

Teams of 4 to 8 threads can cooperatively perform a single grid pencil update, working out of a small, shared, on-chip cache of 128 KB or more.

All present GPU hardware has such shared on-chip caches, and there is enough space in them to enable teams of 4 or 8 simultaneous threads.

Makes work for more cores without forcing the problem size to expand.

Prefetch # 1.0
Boundary #1.0
EOS - FVsCsEs
Prefetch # 2.0
Interpolate
Prefetch # 3.0
Boundary #2.0
Riemann
Prefetch # 4.0
Boundary #3.0
Update FVs
Prefetch # 5.0
Update Mom En
Write Back





Prefetch # 1.0
13 *32 flops
275 *32 flops
Prefetch # 2.0
387 *32 flops
Prefetch # 3.0
28 *32 flops
118 * 32 flops
Prefetch # 4.0
29 *32 flops
441 *32 flops
Prefetch # 5.0
360 *32 flops
Write Back

So, what operations are we talking about?

- **EOS, uses data from this cell only, but at 2 different stages of the pipeline. $275 \times \text{SIMD-32}$**
- **Interpolation, uses data from 3 stages of the pipeline for its 5-cell stencils. $387 \times \text{SIMD-32}$**
- **Riemann Problems, uses data from 2 pipeline stages; gives interface velocities. $118 \times \text{SIMD-32}$**
- **Multifluid fractional volume update, uses data from 2 pipeline stages. $441 \times \text{SIMD-32}$**
- **Conservation law updates with new flux computations from half-time face vels. $360 \times \text{SIMD-32}$**
- **Interspersed prefetches of of 1 KB each & boundary condition ops. $70 \times \text{SIMD-32}$**

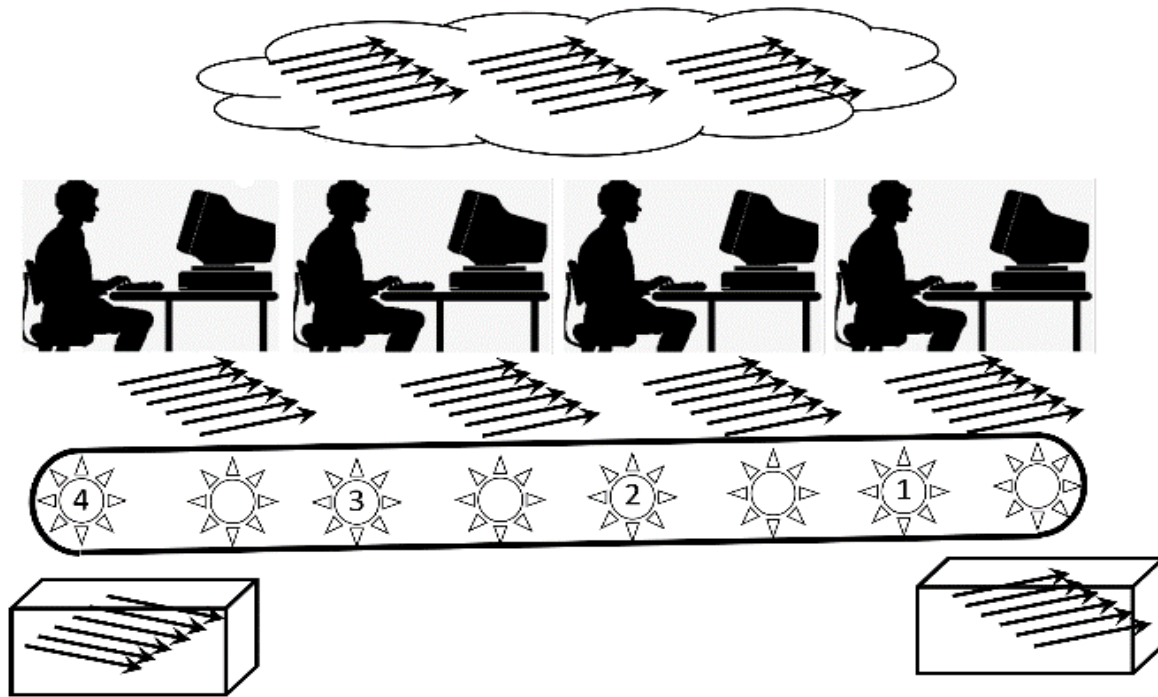
Breaks naturally into a sequence of operations of limited extent along the grid pencil (small number of pipeline stages) that begin at the point where data first appears in cache to the place where it is written back.

Analogy:



A Production-Line Analogy for our Pipelined Half-Briquette Update.

- An implementation with 4 production line stations – 4 threads of control.
- Each station works simultaneously on a different segment of the pipelined half-briquette update.
- Each does not begin work until the previous one has finished.



- The operations along the pipelined sequence naturally fall into large, discrete tasks.
- Each such task, generally, requires data from more than one pipeline stage – this means that the data comes from 1 or 2 previous iterations of the pipelined work.
- Each worker, or thread of control, works on a different iteration of the pipelined work.
- Upon the completion of each macro-task, data products are written over previous ones.
- Data inputs for each macro-task need be stored for only a single pipeline iteration.
- The overall data context on the chip is minimized: it is much, much less than $4\times$
- We get $4\times$ performance with much less than $4\times$ the required on-chip data context.

Here we show the sequence of pipeline iterations with the delays required for 4 threads of control to execute simultaneously. 5 macro-tasks are shown, of unequal length in flops.

Constraints:

No macro-task may begin until that task is complete for the previous pipeline iteration.


No pipeline iteration may begin until the iteration 4 times ago is complete.

These 2 constraints enforce a constant delay that remains

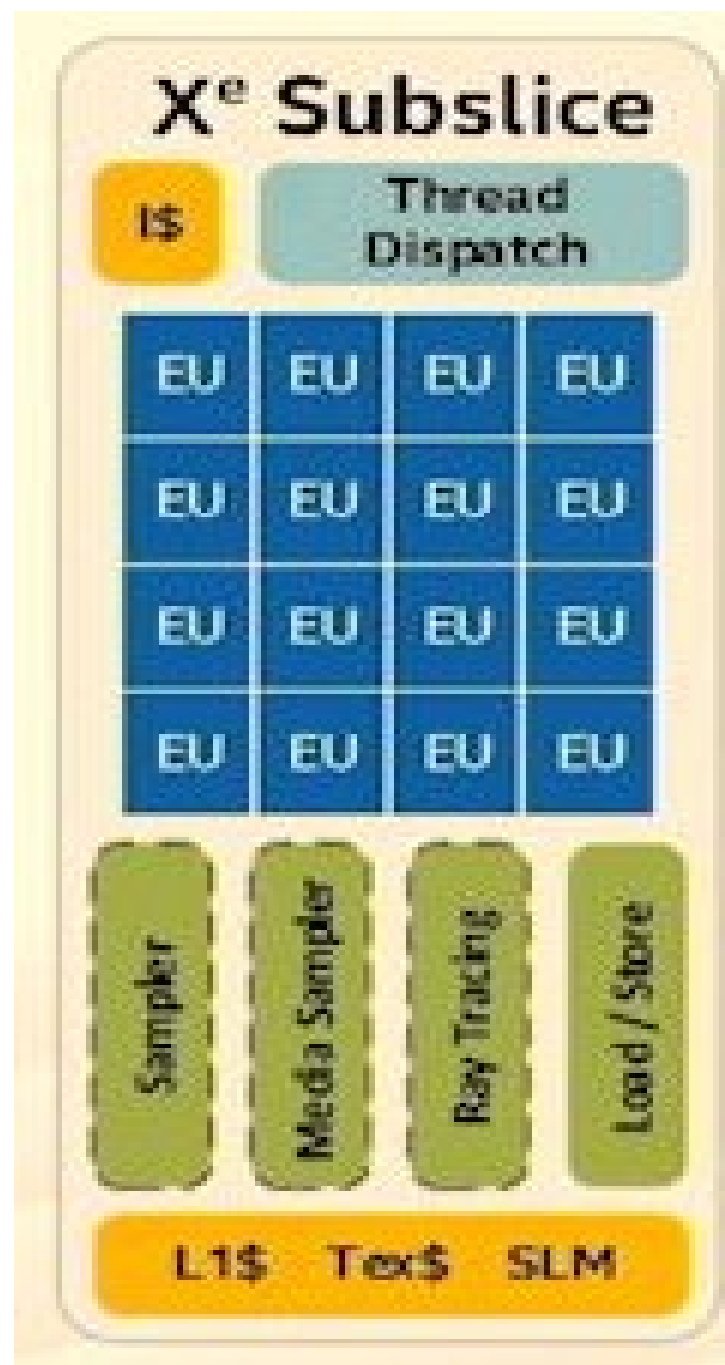
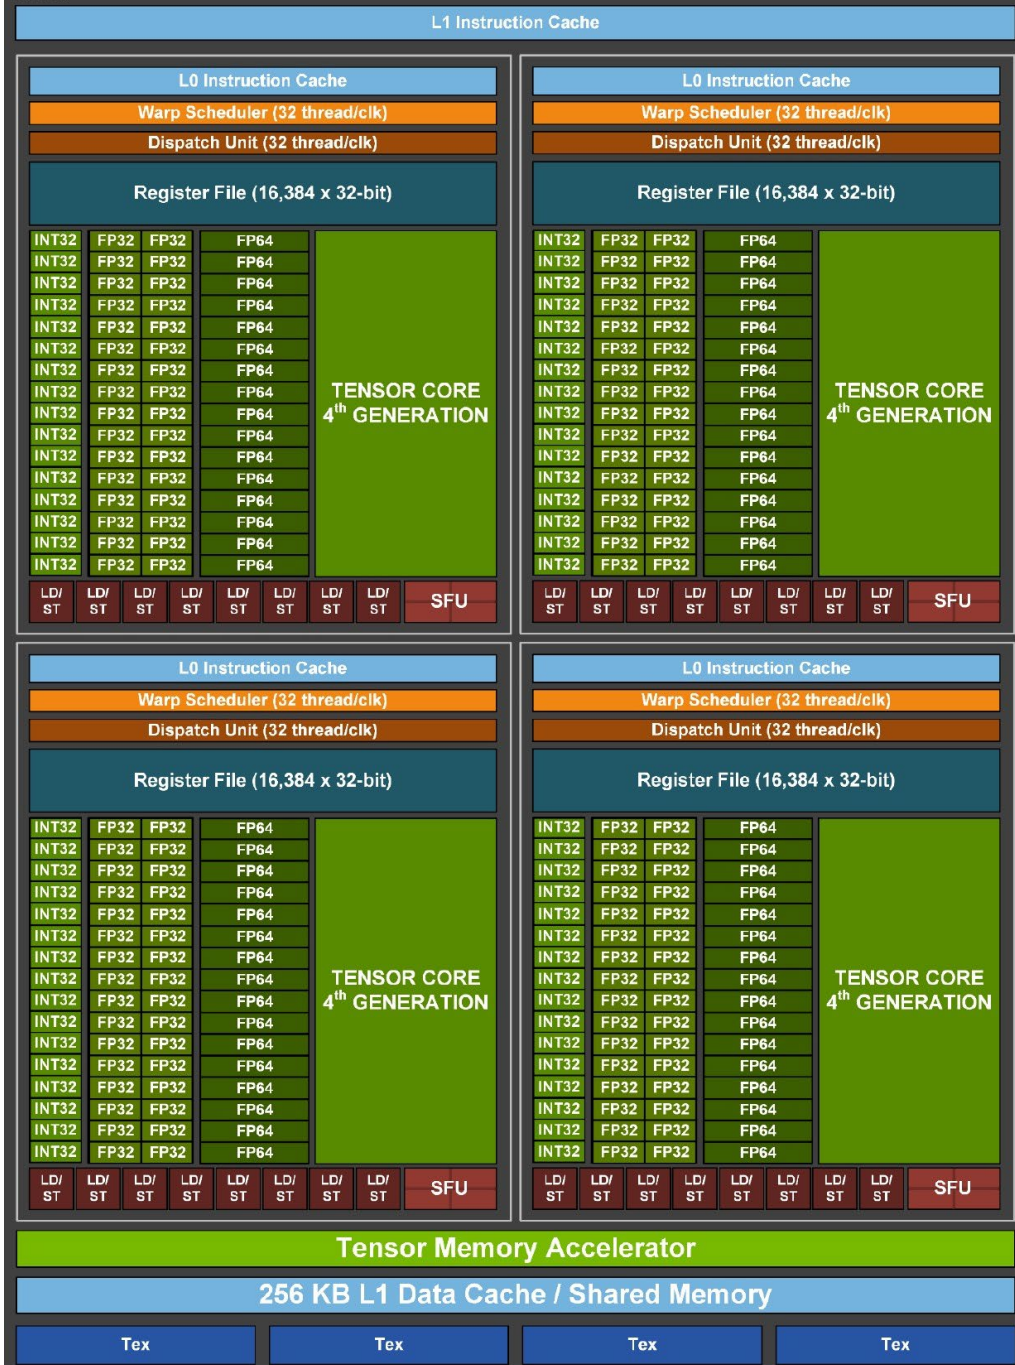
unchanged over

many, many iterations.

There is a start-up cost and an ending cost.



Each strip of operations shown here, representing one trip through the half-bq update loop, represents a highly optimized 32-way SIMD computation, that has the minimum vector length consistent with top speed and the minimum on-chip data context consistent with no unnecessary off-chip data access.



- A thread of control executes on one Nvidia SM quadrant or on a single pair of Intel EUs.
- The size of the instruction word allows only 255 registers per thread of control.
- The entire update cannot therefore be fit into the register file for a single thread of control.
- There is some extra space in the SLM/L1/SHM, but not much space.
- 8 threads of control must execute on each Nvidia SM or Intel Xe subslice.

*Movie of rotating core convection, run M163
vorticity magnitude,
10% equatorial slice close-up view.
Movie by Paul R. Woodward & Ben McDermott*

*Movie of rotating core convection, run M137
vorticity magnitude,
2% equatorial slice close-up view.
Movie by Paul R. Woodward*

The research reported here was supported by NSF through CDS&E grant 1814181, travel grant 2032010, and through grants of access to the Frontera computing system at TACC in Austin, Texas, where the simulations shown and discussed here were carried out and image rendering performed.

Partial support was also provided by NSF through the JINA-CEE physics frontier center, award PHY-1430152.

Herwig acknowledges funding through an NSERC Discovery Grant and a grant of access to the Compute Canada Niagara supercomputer operated by SciNet at the University of Toronto.

Herwig also acknowledges support for data analysis on the Astrohub online virtual research environment (<https://astrohub.uvic.ca>) developed and operated by the Computational Stellar Astrophysics group (<https://csa.phys.uvic.ca>) at the University of Victoria and hosted on the Compute Canada Arbutus Cloud at the University of Victoria.

Image rendering was performed using Ted Wetherbee's `srend` volume rendering package (written in Fortran) and included as part of the PPMstar code.

The movie images were rendered in part on Frontera and in part on an LCSE server at the University of Minnesota.

All the large runs shown and discussed here, runs M284, M163, and M137, were performed on Frontera during Texascale Days events using the PPMstar code running on 3510 nodes with grids of 2688^3 cells, at a sustained rate of 2.8 Pflop/s in 32-bit precision for M163 in 4/2023.