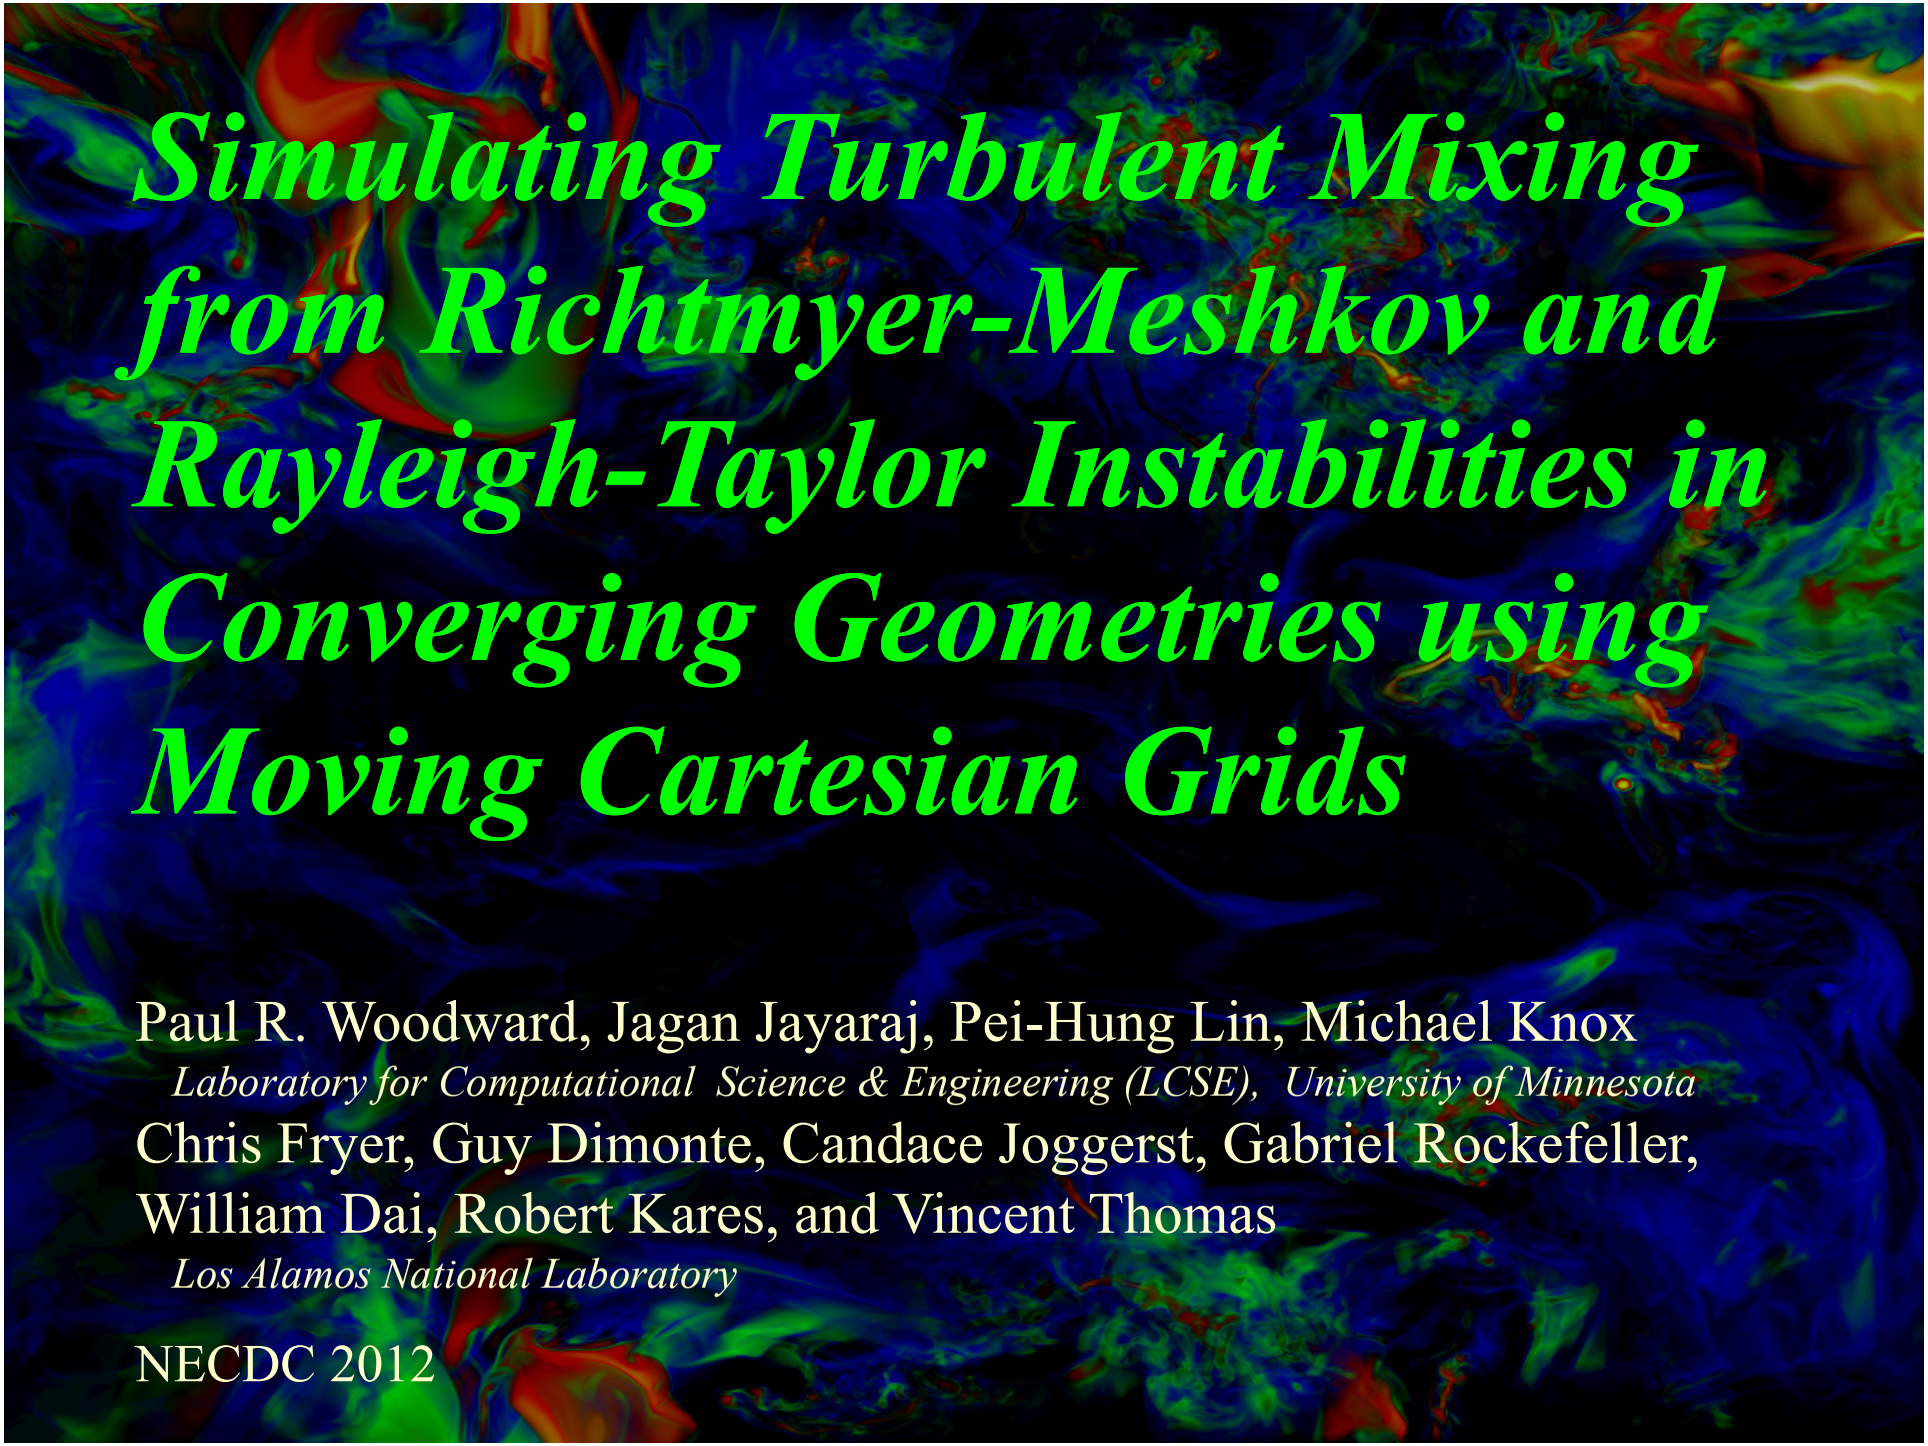

Simulating Turbulent Mixing from Richtmyer-Meshkov and Rayleigh-Taylor Instabilities in Converging Geometries using Moving Cartesian Grids

Paul R. Woodward, Jagan Jayaraj, Pei-Hung Lin, Michael Knox

Laboratory for Computational Science & Engineering (LCSE), University of Minnesota

Chris Fryer, Guy Dimonte, Candace Joggerst, Gabriel Rockefeller, William Dai, Robert Kares, and Vincent Thomas

Los Alamos National Laboratory

NECDC 2012

The Problem:

Simulate turbulent mixing in converging geometries with sufficient accuracy to validate subgrid-scale models of this phenomenon.

Such simulations should be feasible on today's large computing systems at a practical cost.

We would like to do such studies using LANL's codes from the Eulerian Application Project (EAP).

Methodology:

Explore the problem using the PPM code with PPB multifluid volume fraction advection.

This code is being built as a prototype to test strategies to take EAP codes onto advanced architectures.

Computing platforms include Roadrunner, and also large systems like Blue Waters, Sequoia, & TACC (MIC).

Questions to be Answered:

1) Can the code preserve the necessary symmetries of the problem to an adequate degree, and if so at what grid resolution?
2) Does the computed result converge in an appropriate statistical sense under grid refinement, and if so at what grid resolution and cost?
3) Is the converged result reasonable?
4) How does the required grid resolution scale with the size of the radial compression?
5) How does the required grid resolution scale with the amplitude of the initial perturbation?
6) Does AMR help, and if so, to what quantitative extent?
7) What does a converged run cost?
   - ✓ On Blue Waters? On Sequoia? At TACC?

Questions to be Answered:

1) What does this take (i.e. how do ***I*** do that)?
   - Does this require PPM hydro, and if so what about PPM hydro seems to be of critical importance?
   - Does this require PPB multifluid volume fraction advection, and if so, what about that is important?
   - Is "smart" dissipation for strong shocks needed?
   - Is grid jiggling needed?
   - Is there a dependence of the quality of the result upon the frequency/amplitude of the initial disturbance?
   - Could the results be improved by introducing a Navier-Stokes type of viscosity?
   - Could the results be improved by introducing an artificial surface tension?
   - ***How important is the implementation of the code for fast execution on advanced hardware?***

Statement of the original Youngs inerital confinement fusion test problem in 2008:

A simple spherical implosion (dimensionless units) [10]

Perturbation spectrum

$P(k) \sim \frac{1}{k^2}$, s.d.=0.0005

max. wavelength = 2.0

$\rho = 1.0 \quad p = 0.1$

Outer radius = 12

$\rho = 0.05, p = 0.1$

Outer radius= 10

Perfect gas equations of state $\gamma = 5/3$

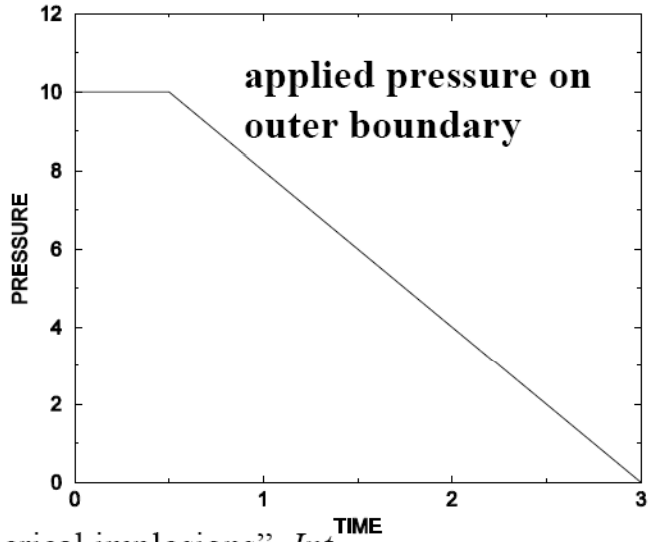

10. D.L.Youngs and R.J.R.Williams, "Turbulent mixing in spherical implosions", *Int. J. Numer. Meth. Fluids*. **56**, pp. 1597-1603 (2008)

http://www.cmla.ens-cachan.fr/fileadmin/Groupes/Cargese/2010/Cargese_2010-09-15_09-30.pdf

1D Lagrangian calculation

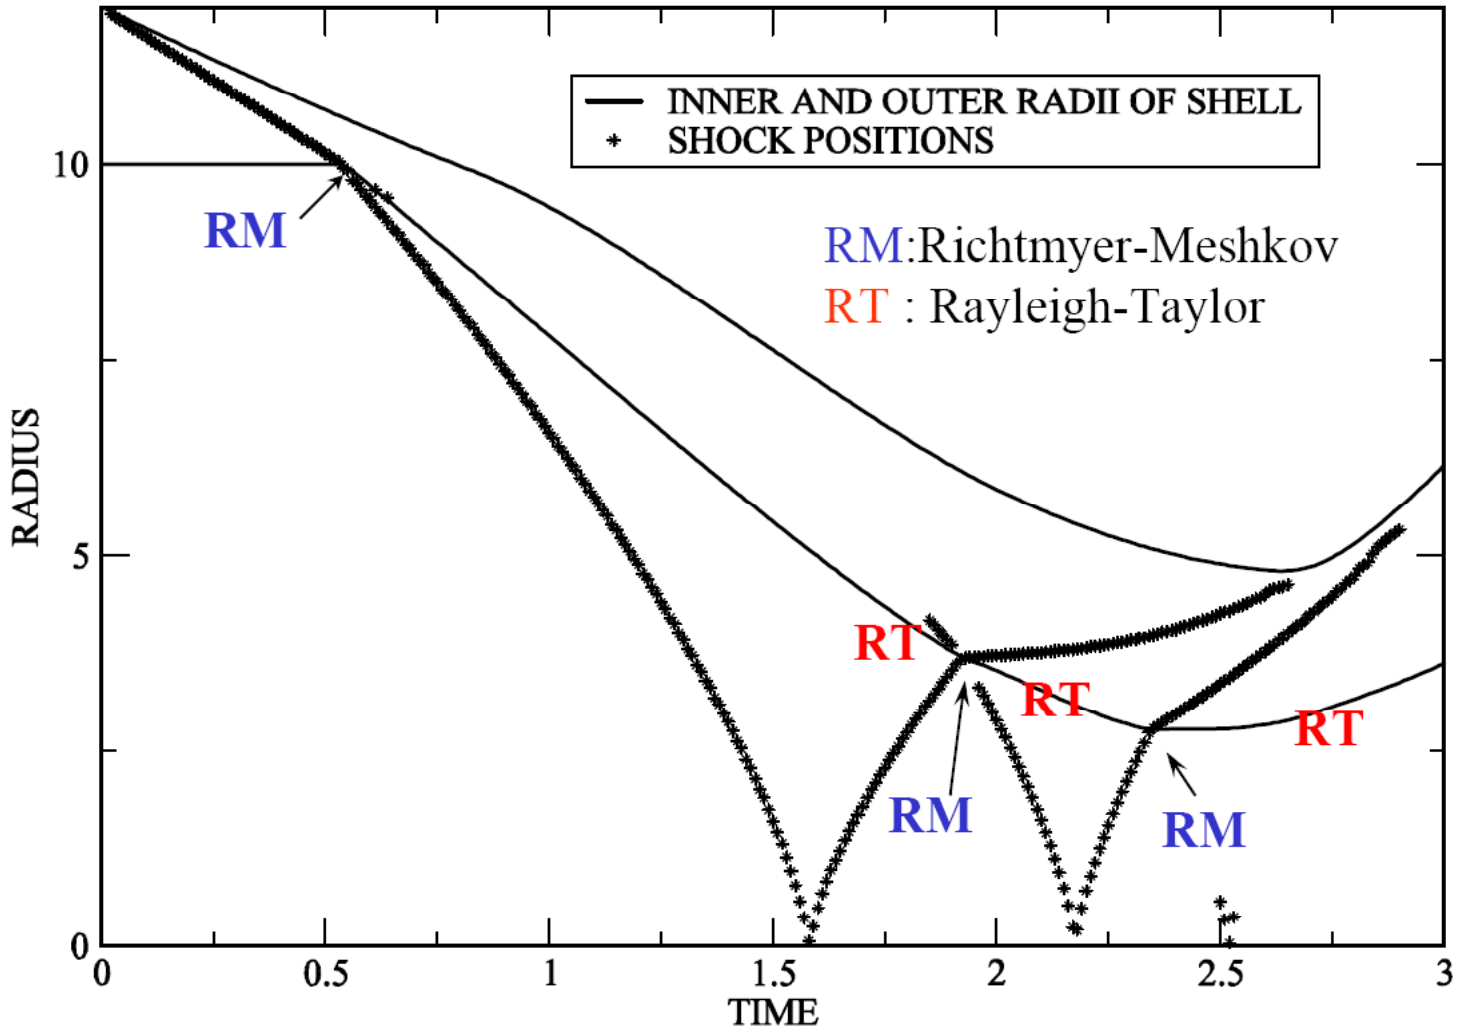

Note influence of initial conditions is complex: initial spectrum, amplification due to first shock and spherical convergence all set initial perturbations for late stage mixing

http://www.cmla.ens-cachan.fr/fileadmin/Groupes/Cargese/2010/Cargese_2010-09-15_09-30.pdf

3D SIMULATION THE SPHERICAL IMPLOSION

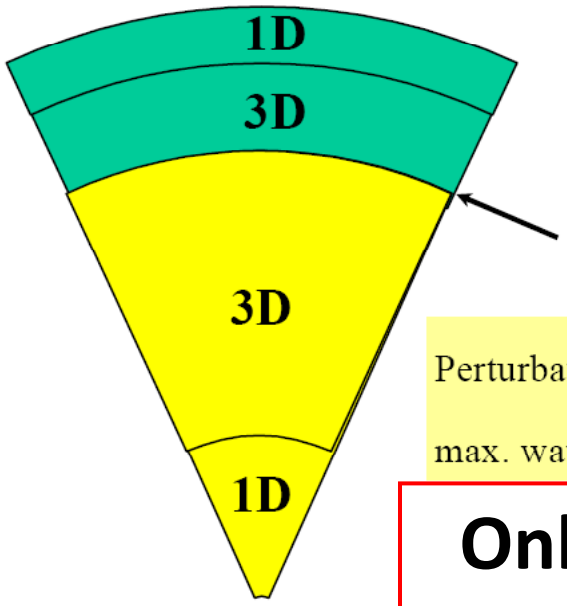

"square" patch perturbed by Fourier modes

Perturbation spectrum, $P(k) \sim \frac{1}{k^2}$, s.d.=0.0005

max. wavelength = shell thickness

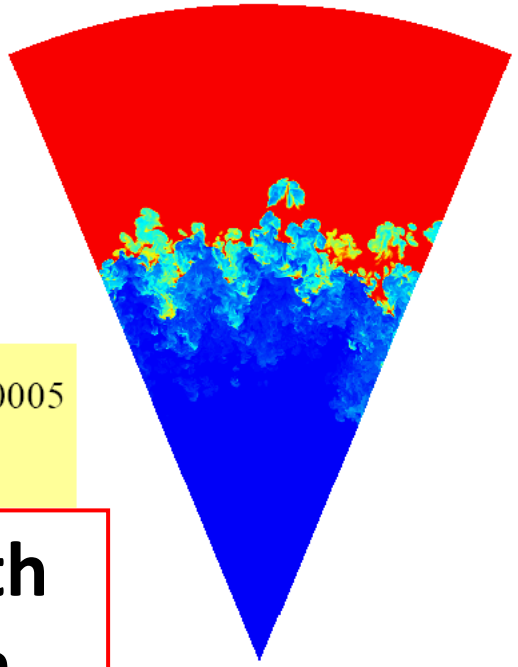

Only 45° sector, with 1-D region at origin.

Calculate sector $\frac{\pi}{2} - \frac{\pi}{8} < \theta, \phi < \frac{\pi}{2} + \frac{\pi}{8}$

Spherical polar mesh, Lagrangian in r-direction, 1D Lagrangian regions at origin and at outer boundary.

This is the original ICF test problem described by Youngs in 2008. We will improve upon it by computing the entire spherical region without the falsification of the boundary condition at the origin. We also modify it to produce a higher radial compression.

http://www.cmla.ens-cachan.fr/fileadmin/Groupes/Cargese/2010/Cargese_2010-09-15_09-30.pdf

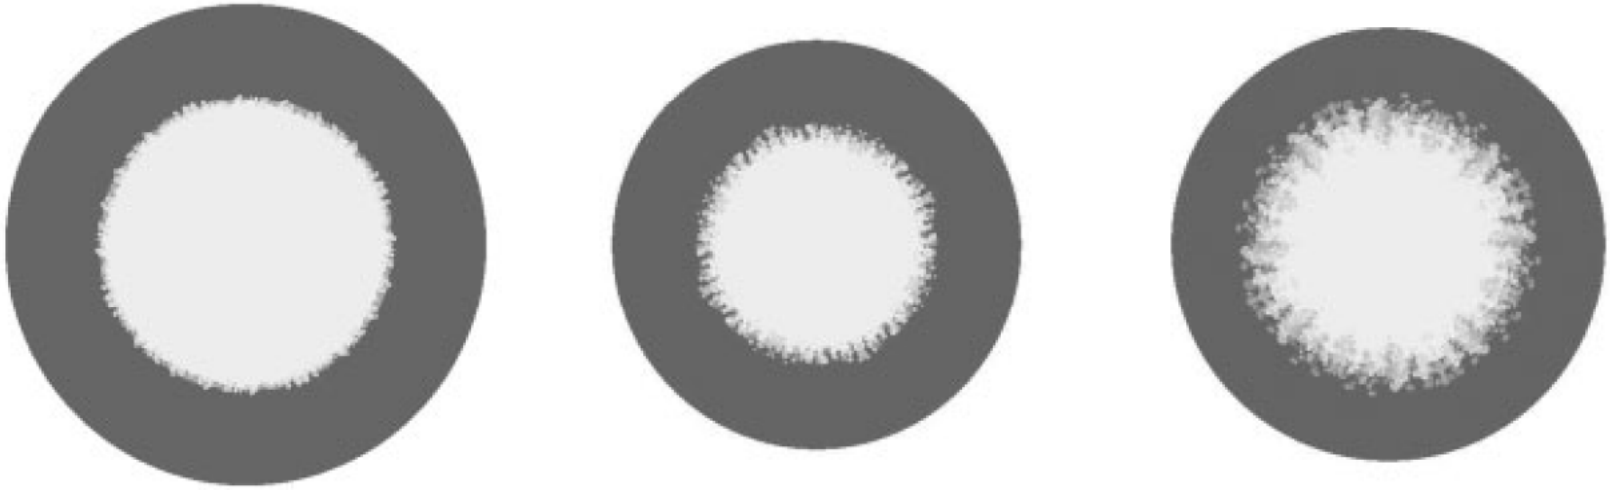

Figure 3. 2D sections through the 3D mesh for the spherical implosion, before, near and after maximum compression: (a) $t = 2.0$; (b) $t = 2.4$; and (c) $t = 2.8$.

Int. J. Numer. Meth. Fluids 2008; **56**:1597–1603
DOI: 10.1002/fld

This is the original ICF test problem described by Youngs in 2008. We will improve upon it by computing the entire spherical region without the falsification of the boundary condition at the origin. We also modify it to produce a higher radial compression.

Note the prominent mode 4 that has entered this simulation, arising from simulating only a 45° sector. It seems quite unlikely that this mode's emergence and growth to this dominant role in the solution could accurately represent the desired realistic behavior.

http://www.cmla.ens-cachan.fr/fileadmin/Groupes/Cargese/2010/Cargese_2010-09-15_09-30.pdf

Initial State for ICF Implosion by Radial Factor of 8 to 10:

Spherically symmetric base initial state.

For $r < 1.0$: $\rho = 0.01$, $p = 0.1$, $u = 0$, $\gamma = 5/3$.

For $1.0 < r < 1.2$: $\rho = 1.0$, $p = 0.1$, $u = 0$, $\gamma = 5/3$.

For $1.2 < r < 1.4$: $\rho = 0.1$, $p = 160$, $u = 0$, $\gamma = 5/3$.

For $1.4 < r < 1.5$: $\rho = 0.1$, $p = 160$, $u_r = -9.75\, r$, $\gamma = 5/3$.

Homologous radial grid contraction with $u_r(r,t) = u_r(r_0)$.

Perturb the inner surface of the dense shell with a radial displacement equal to $dr = A_{lm}\, Y_{lm}(\theta,\varphi)$,
with amplitude $A_{81\,39} = 0.04\,\pi / 39$
$A_{74\,36} = 0.02\,\pi / 39$

Smear transition regions with $f_{outer} = [1 + \sin(\pi(r-r_{trans})/\delta)]$
where r_0 is the radius of the transition, f_{outer} is the fraction of the outer (radial) state, and δ is the thickness.

Boundary Condition for ICF Implosion by Radial Factor 8-10:

Outside the radius $r_{bdry} = 1.4\,(1 - 9.75\,t)$

For $t < 0.05$

$\rho = 0.1$, $\quad p = 160$,
$u_r = -9.75\,r / (1 - 9.75\,t)$, $\quad \gamma = 5/3$.

For $t > 0.05$

$\rho = 0.1$,
$p = 160\ (0.1/160)^{10\,(t - 0.05)}$,
$u_r = -9.75\,r / (1 - 9.75\,t)$, $\quad \gamma = 5/3$.

Smear transition region with $f_{outer} = [1 + \sin(\pi(r - r_{bdry})/\delta)]$
where r_0 is the radius of the transition, f_{outer} is the fraction of the outer (radial) state, and δ, the transition thickness, is a constant number of grid cell widths, even as the grid cell width shrinks with the radially inward grid motion.

ICF Implosion by Radial Factor of 8-10:

Run to time t = 0.0525

Dump interval is 0.0025

Shock emerges from dense capsule around dump 6
and converges on the origin after dump 13,
around time 0.034

Reshock of Richtmyer-Meshkov fingers occurs shortly before
dump 16.

At dump 20 or 21, fuel is nearly at maximum compression.

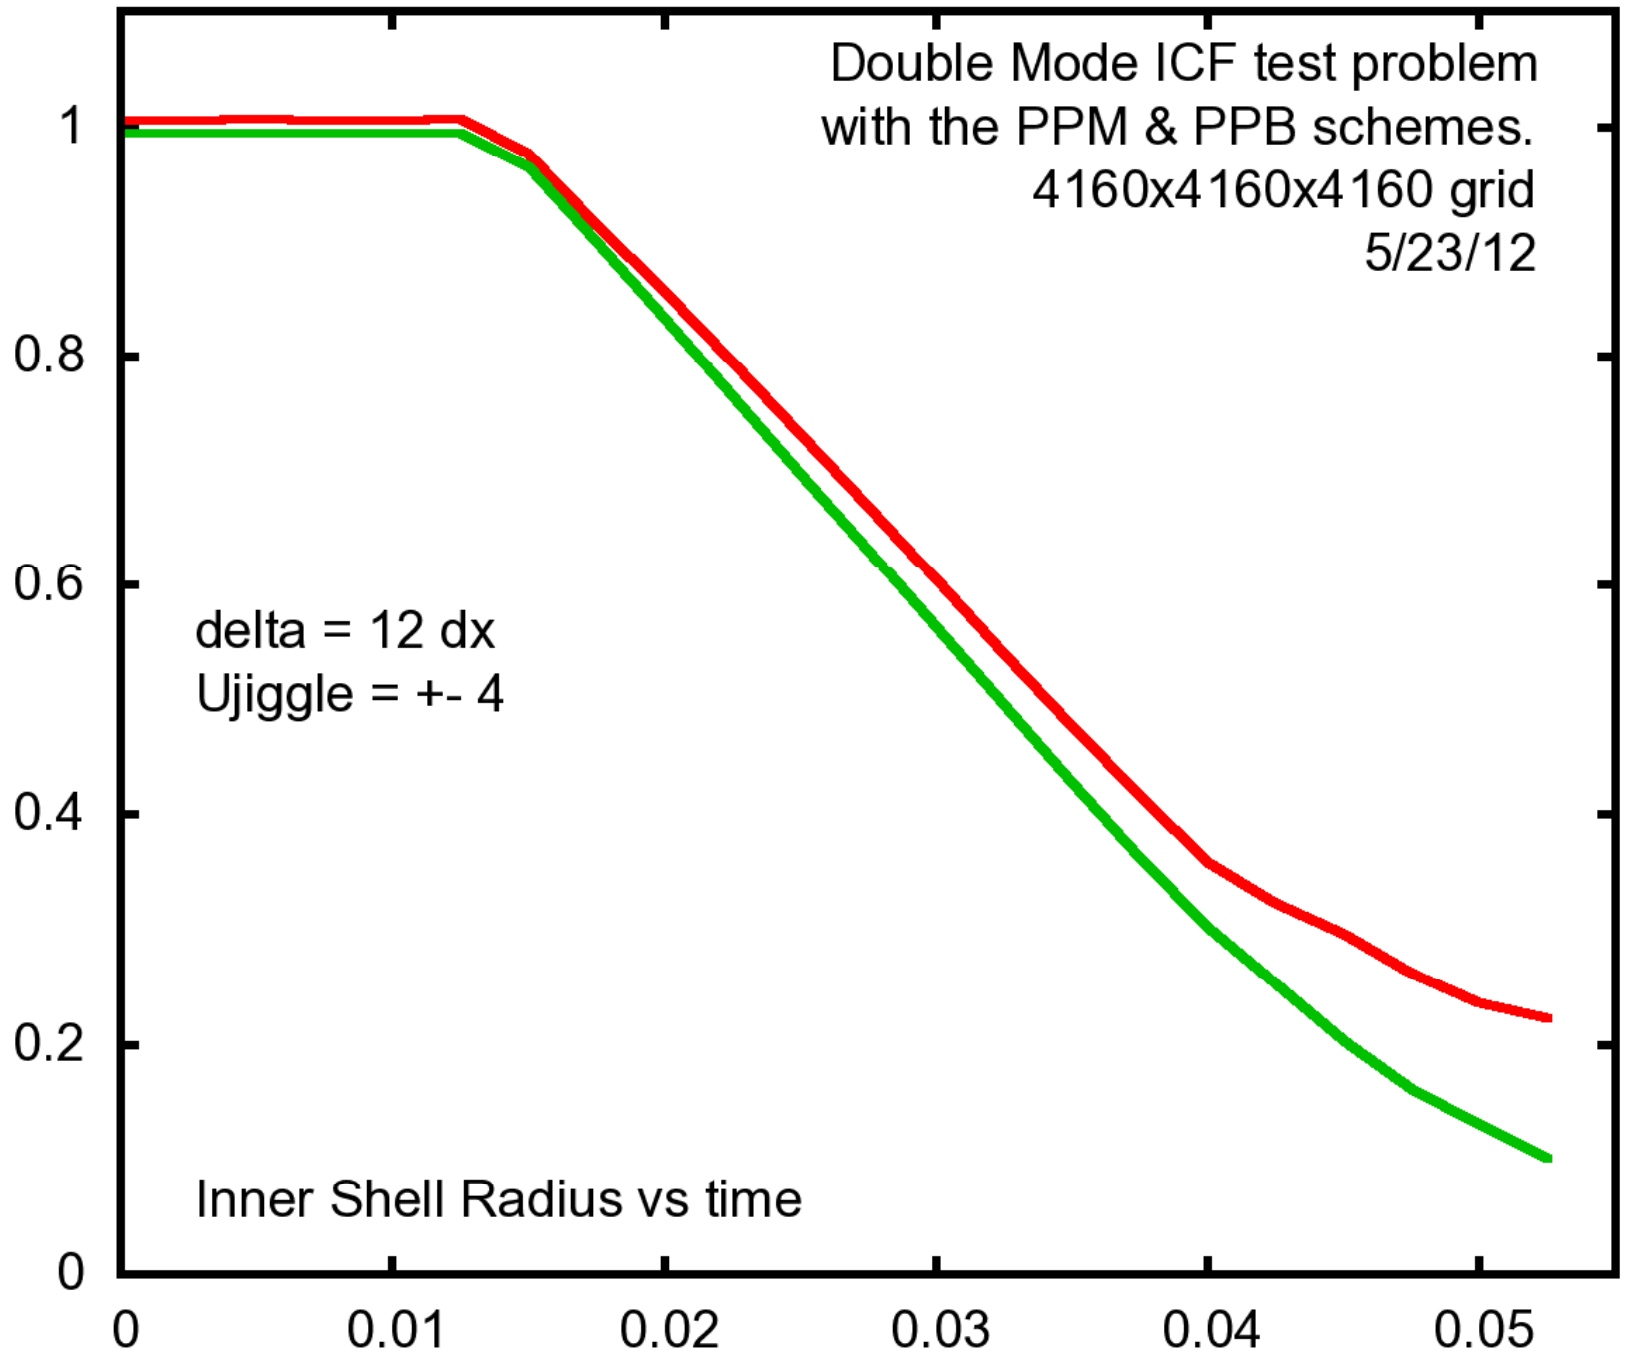
Double Mode ICF test problem
with the PPM & PPB schemes.
4160x4160x4160 grid
5/23/12
delta = 12 dx
Ujiggle = +- 4
Inner Shell Radius vs time
1
0.8
0.6
0.4
0.2
0
0
0.01
0.02
0.03
0.04
0.05

4160 x4160 x4160 grid,
Atwood # = 9.0476E-01, g = 0.0000E+00, Dynamic Viscosity = 0.00000E+00, CoefNS = 0.00000E+00
Lambda = 1.0000E+00, DisplaceAmpl = 3.22214E-03, Delta = 8.6538E-03, Ujiggle = +-4.00000E+00
Rho00 = 5.0000E-02, CEul0 = 1.82574E+00, Gamma = 1.66667E+00 Modes 7436 to 8139
Velocity Unit = 1.00000E+00 = sqrt(A*g*Wavelength / (1+A)) Power = 0.000E+00
Theta = Intgl{<fv*(1-fv)>}dy / Intgl{<fv>*(1-<fv>)}dy
PPM gas dynamics with PPB multifluid advection
#
#
Ndump Ncycle Time T Rescaled DT Courant # Bubble Ht Spike Ht Bubble UY Spike UY Theta

0	0	0.000000	0.000000	0.000004	0.500000	1.005617	0.994557	0.000000	0.000000	0.805600
1	540	0.002500	0.002500	0.000005	0.506072	1.005591	0.994625	0.165629	0.032163	0.781949
2	1052	0.005000	0.005000	0.000005	0.506124	1.006445	0.994718	0.005364	0.018167	0.762096
3	1584	0.007500	0.007500	0.000005	0.506053	1.005618	0.994716	-0.031614	0.004268	0.763996
4	2136	0.010000	0.010000	0.000004	0.505981	1.006287	0.994739	0.144815	0.013971	0.766733
5	2712	0.012500	0.012500	0.000004	0.505907	1.006342	0.994786	-6.058069	-5.864686	0.768681
6	3312	0.015000	0.015000	0.000004	0.505832	0.975997	0.965416	-18.333389	-18.890484	0.899603
7	3936	0.017500	0.017500	0.000004	0.505755	0.914674	0.900333	-24.044661	-26.520786	0.764109
8	4588	0.020000	0.020000	0.000004	0.505676	0.855773	0.832811	-23.831442	-26.987238	0.590905
9	5272	0.022500	0.022500	0.000004	0.505595	0.795518	0.765397	-24.747852	-26.906843	0.502433
10	5988	0.025000	0.025000	0.000003	0.505513	0.732034	0.698277	-25.348907	-26.900866	0.466842
11	6740	0.027500	0.027500	0.000003	0.505428	0.668773	0.630893	-25.373142	-26.982403	0.435551
12	7532	0.030000	0.030000	0.000003	0.505341	0.605169	0.563366	-25.630377	-26.971165	0.403180
13	8384	0.032500	0.032500	0.000003	0.549003	0.540621	0.496037	-25.538372	-26.921286	0.384529
14	15448	0.035000	0.035000	0.000001	0.443940	0.477476	0.428759	-24.855370	-26.673883	0.371277
15	17488	0.037500	0.037500	0.000001	0.476332	0.416344	0.362668	-23.682909	-25.329477	0.357527
16	19164	0.040000	0.040000	0.000002	0.485147	0.359062	0.302112	-18.478540	-21.794556	0.337281
17	20748	0.042500	0.042500	0.000002	0.509158	0.323952	0.253695	-12.552259	-19.743675	0.362481
18	22896	0.045000	0.045000	0.000001	0.477419	0.296301	0.203393	-12.363072	-18.337172	0.457712
19	25124	0.047500	0.047500	0.000001	0.492654	0.262136	0.162009	-11.948395	-14.418015	0.572075
20	27440	0.050000	0.050000	0.000001	0.492789	0.236559	0.131303	-7.919075	-11.984714	0.654028
21	30136	0.052500	0.052500	0.000001	0.498061	0.222541	0.102085	-5.607152	-11.687073	0.708574

t = 0.0025

The uniform Cartesian grid moves steadily inward. We render the same portion of the grid in each view here. so that instead of the grid moving inward, we see the fluid appear to move out-ward. The outer surface of the dense shell, however, moves at nearly the same velocity as the grid.

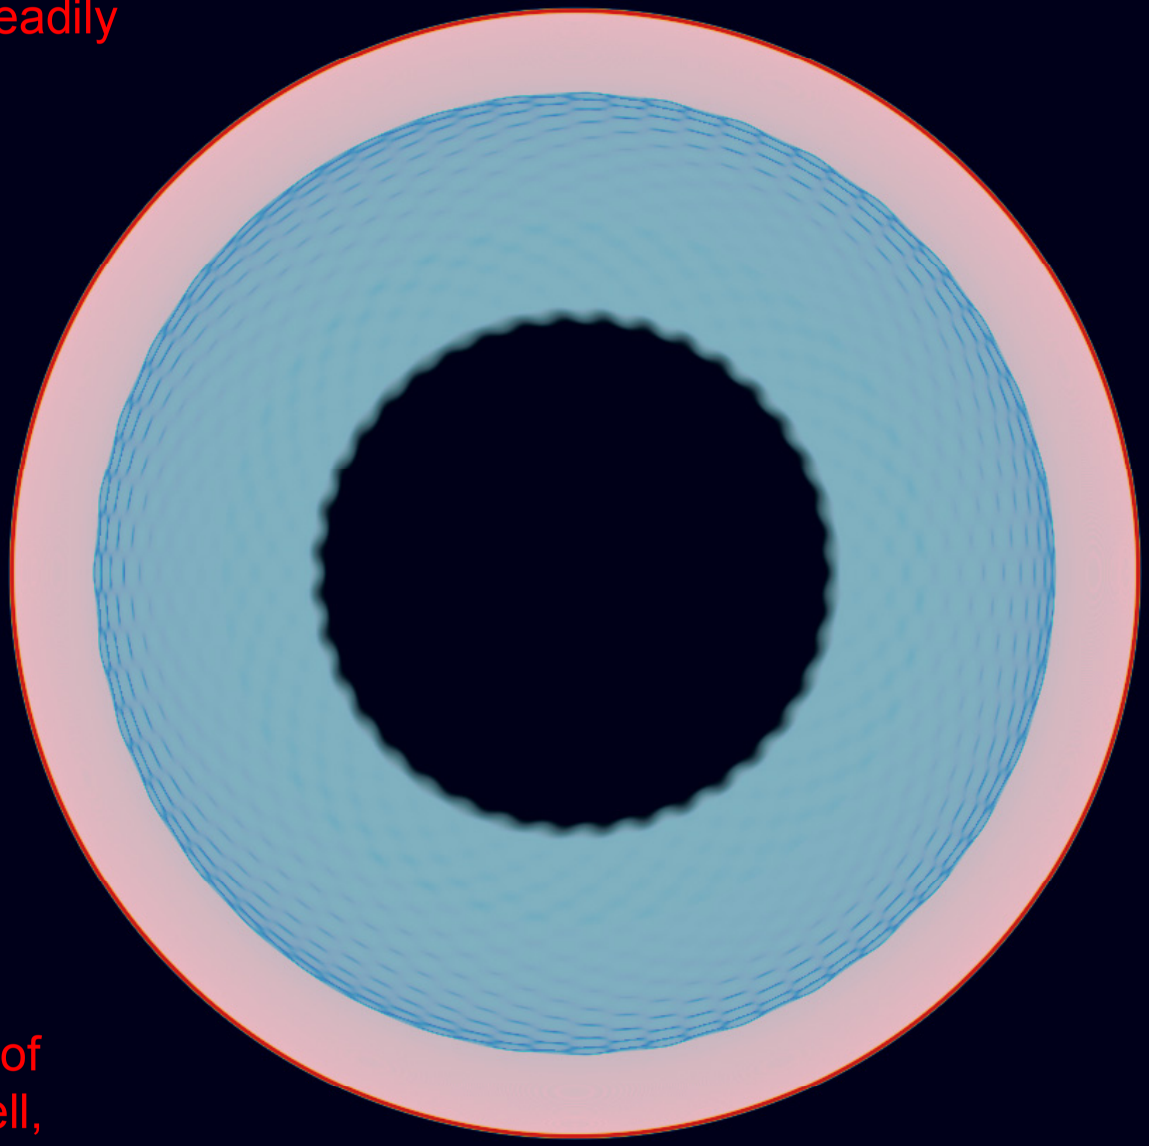

ρ

4160^3 grid

$4^{th} - 7^{th}$ 13^{ths} in Y

rholut 7

opacity 6

distance from midplane 1.6

t = 0.0050

Here we visualize only the 4^{th} – 7^{th} 13ths of the problem domain. This allowed us to work with only ¼ of the data, and it helps to give an impression of the 3-D geometry.

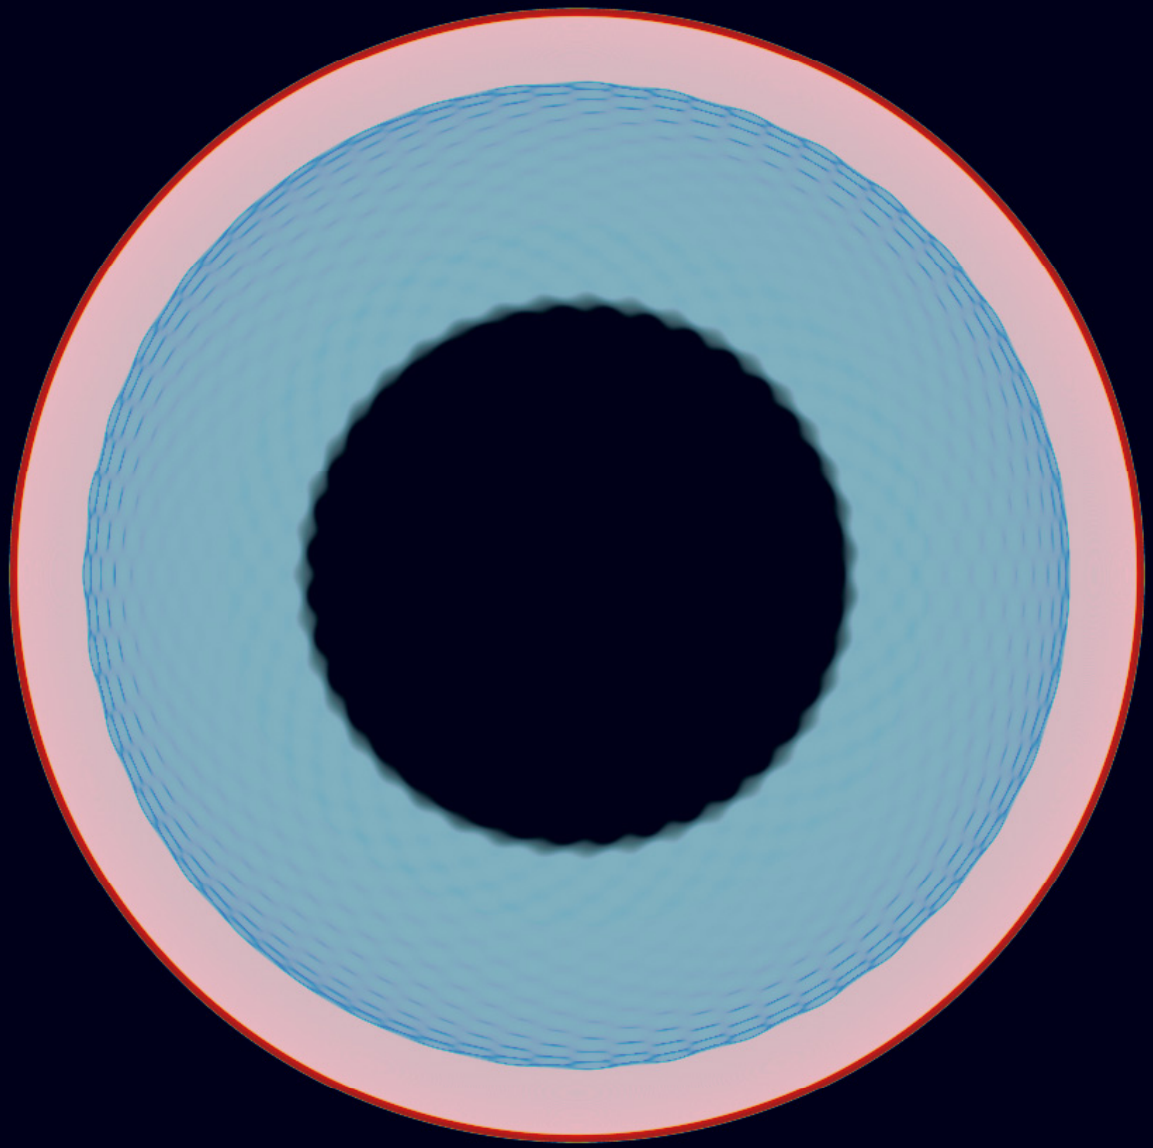

ρ

4160^3 grid

4^{th} – 7^{th} 13^{ths} in Y

rholut 7

opacity 6

distance from midplane 1.6

t = 0.0075

The compressed outer surface region here shows up as red. The initial dual spherical harmonic perturbation of the inner surface of the dense shell is clearly visible.

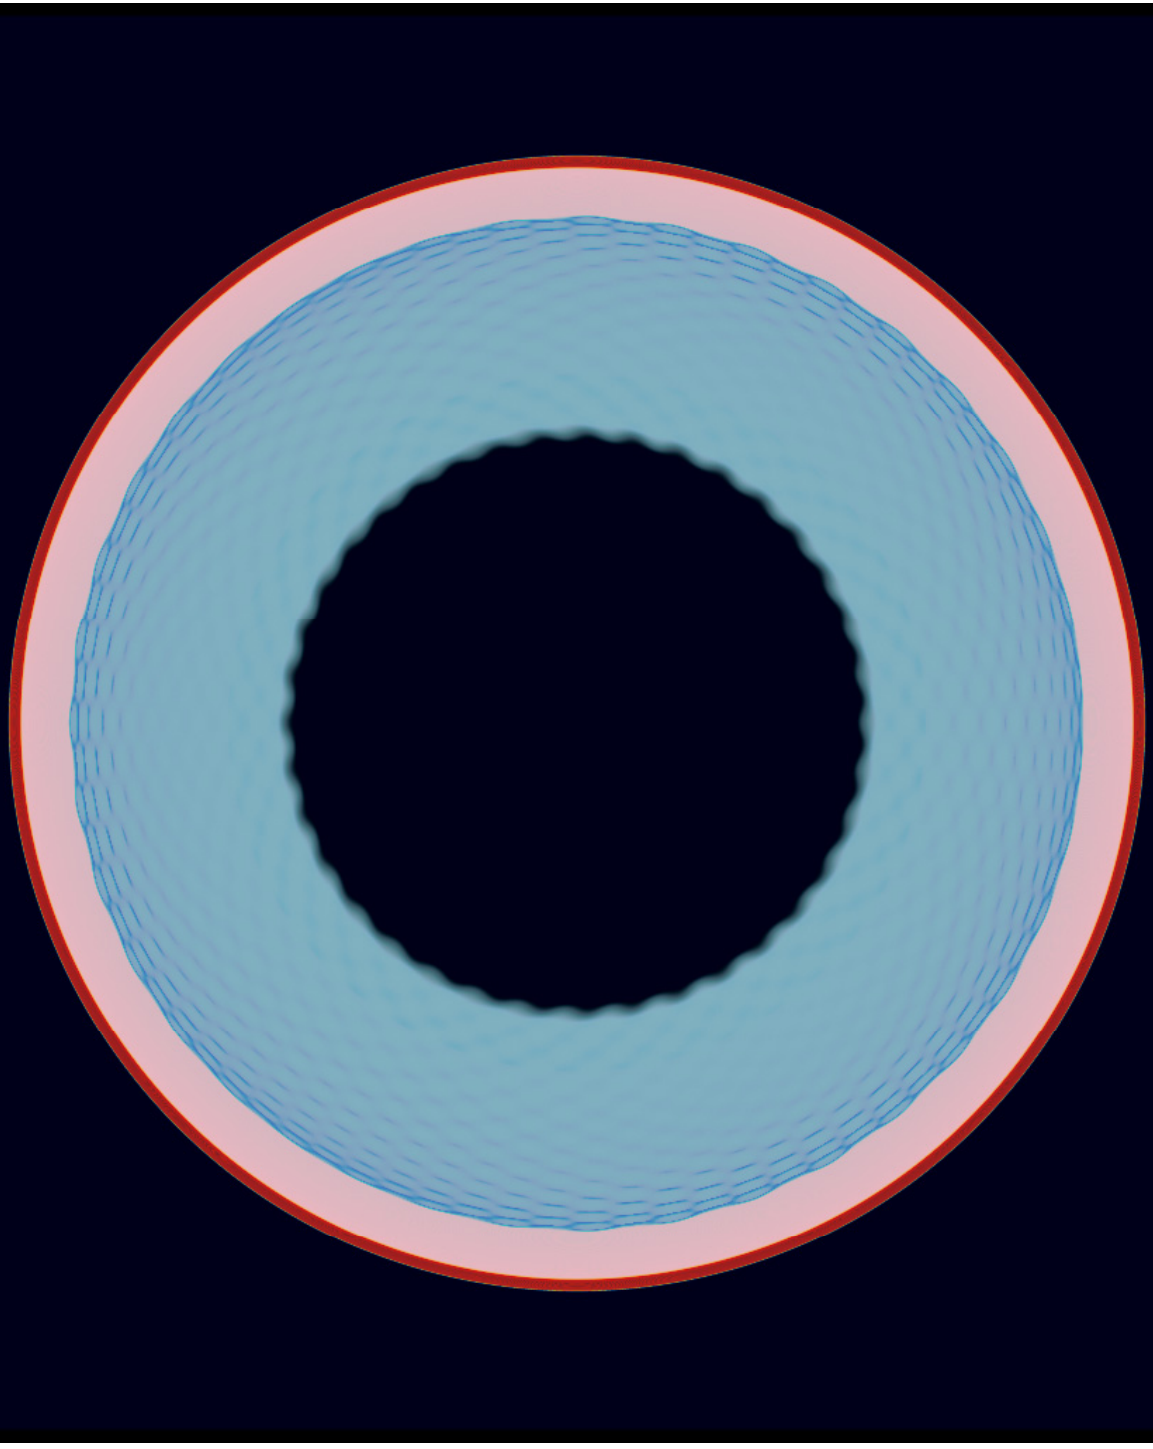

ρ

4160^3 grid

4^{th} – 7^{th} 13^{ths} in Y

rholut 7

opacity 6

distance from midplane 1.6

t = 0.0100

Although the initial disturbance is only 2% of the wavelength in amplitude, it shows up clearly, and we can even detect the modulation of $Y_{81\ 39}$ by $Y_{74\ 36}$, which produces a difference frequency of 3.

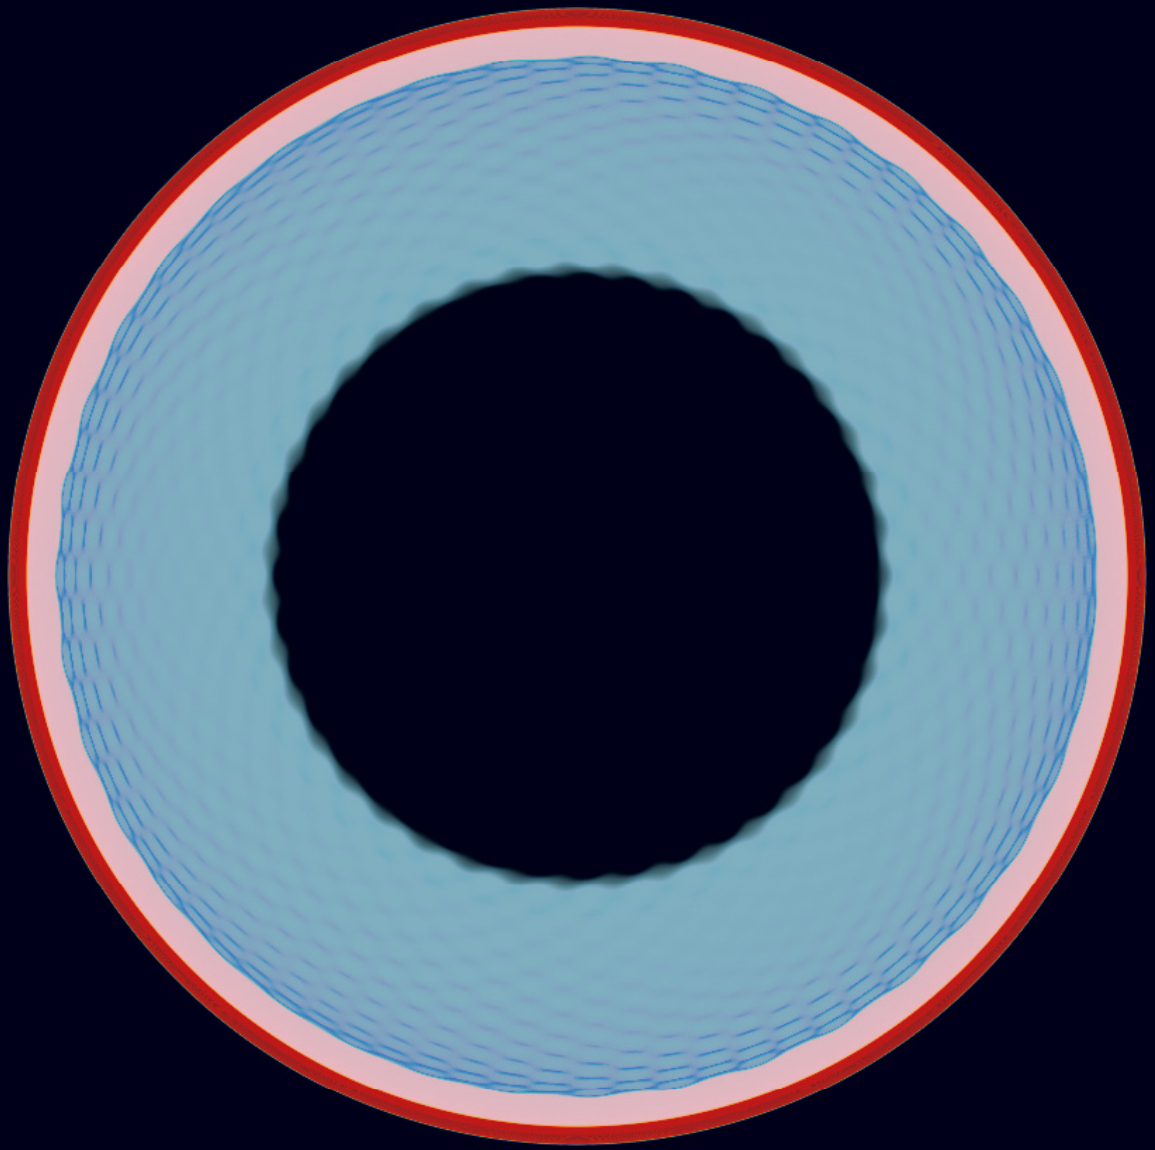

ρ

4160^3 grid

4^{th} – 7^{th} 13^{ths} in Y

rholut 7

opacity 6

distance from midplane 1.6

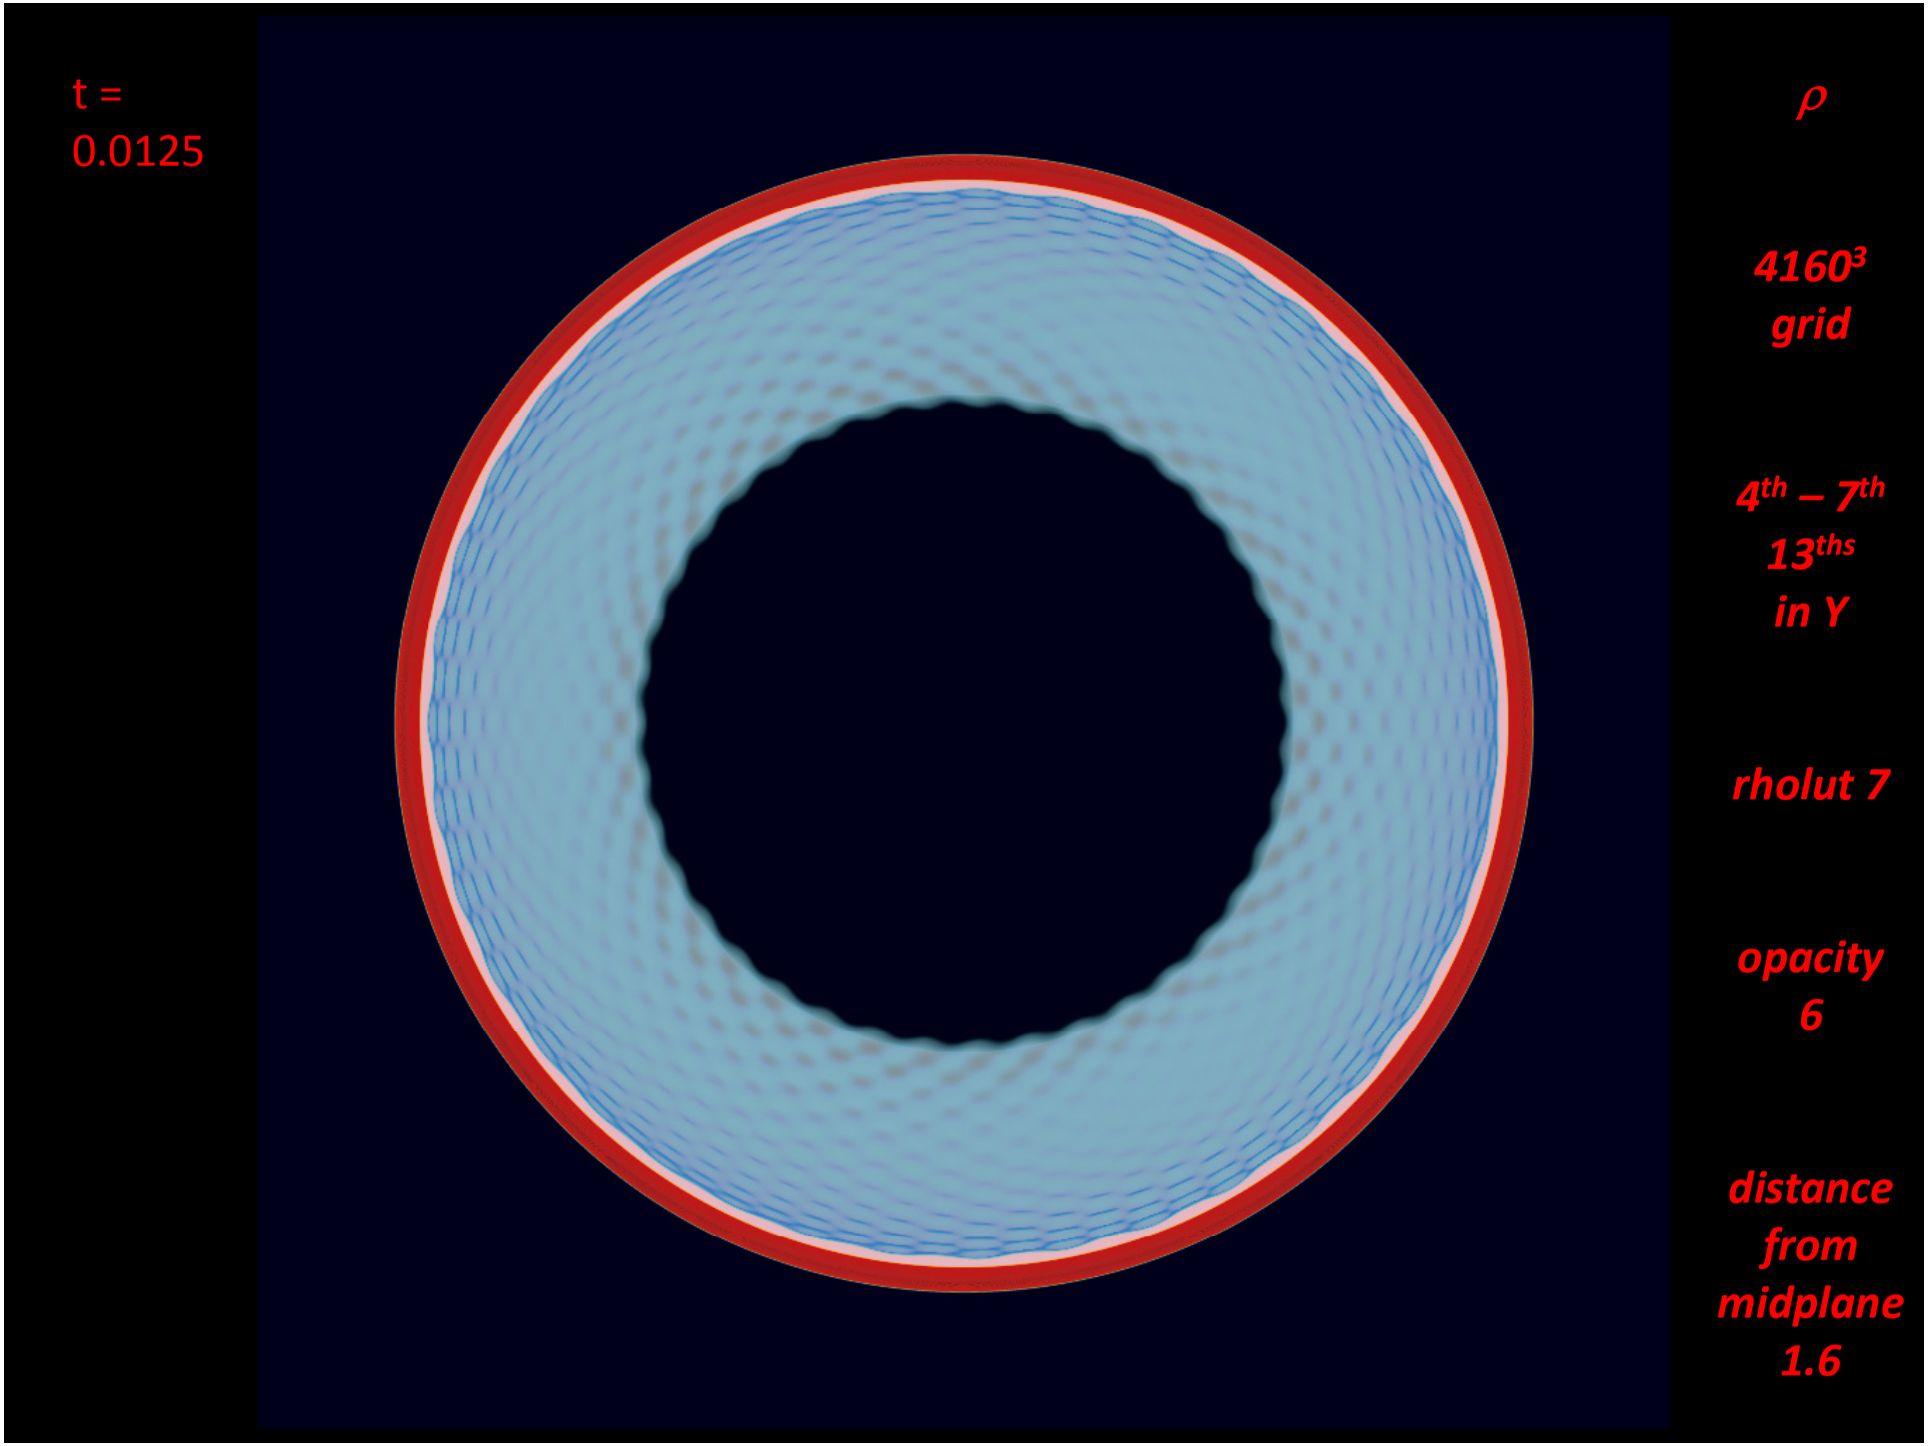
t =
0.0125
ρ
4160³
grid
4th – 7th
13ths
in Y
rholut 7
opacity
6
distance
from
midplane
1.6

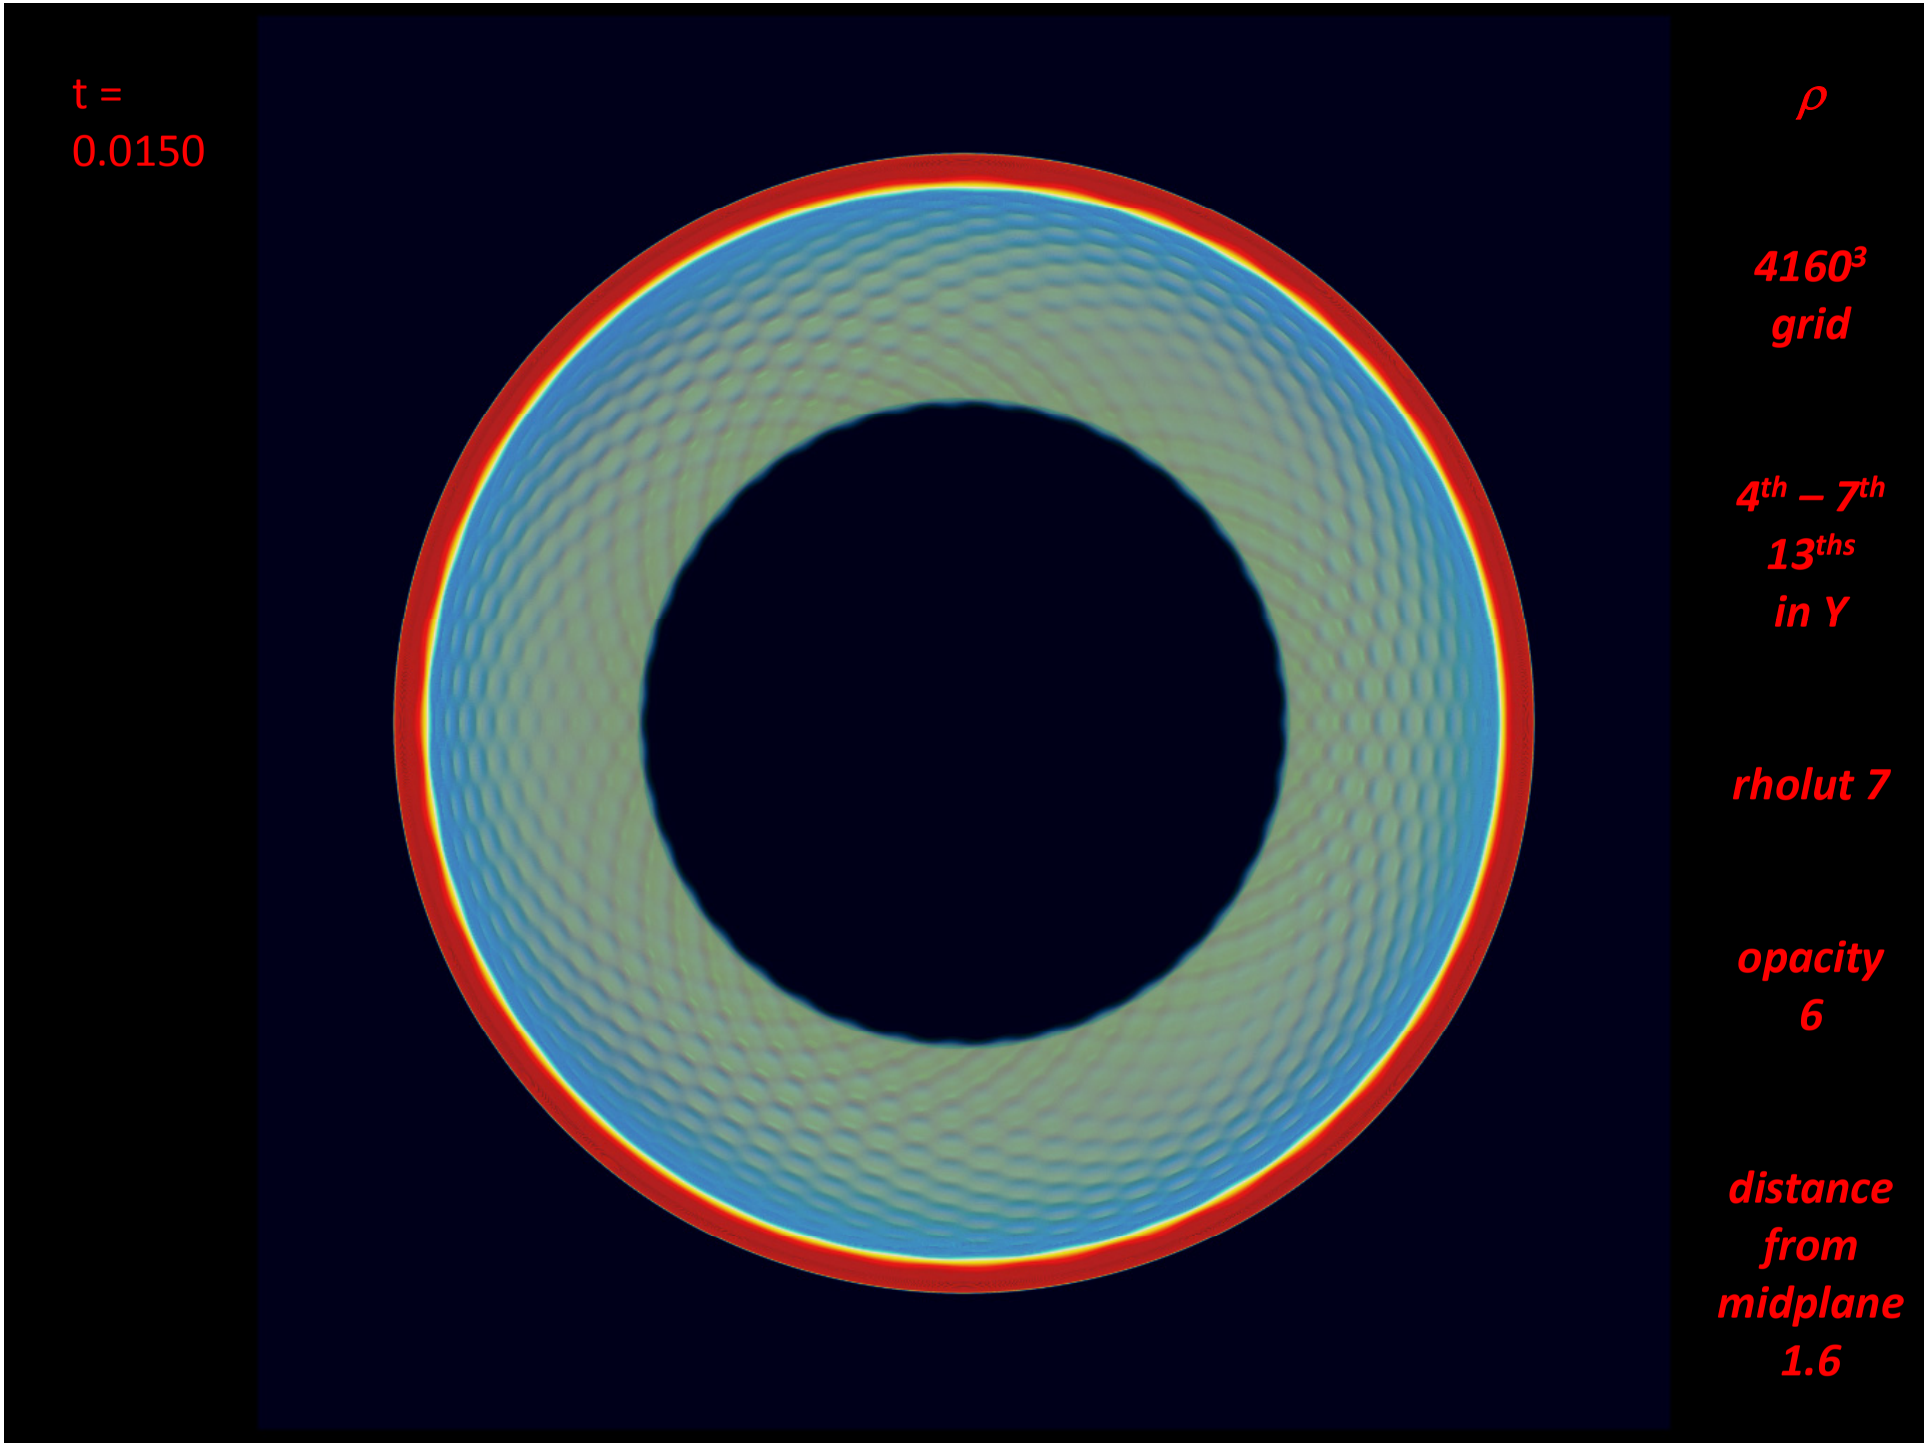
t =
0.0150
ρ
4160³
grid
4th – 7th
13ths
in Y
rholut 7
opacity
6
distance
from
midplane
1.6

t = 0.0175

The strong shock has left the dense shell and is moving through the light gas in the central region. The disturbance of the inner shell surface is amplified by the Richtmyer-Meshkov instability.

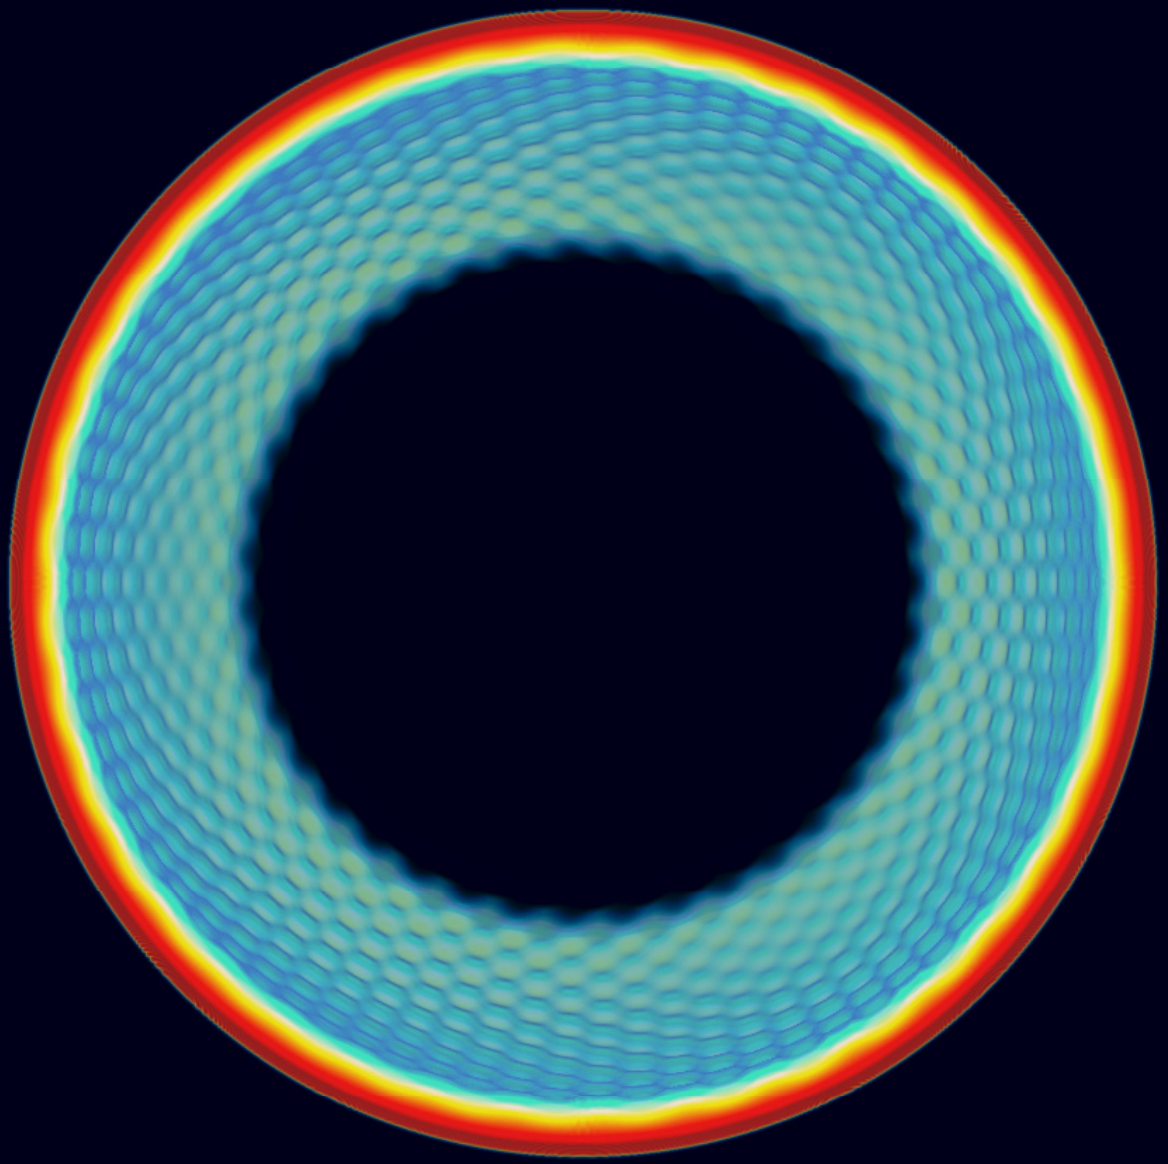

This dump was somehow overwritten with dump 8. We show the 2048^3 version.

ρ

2048^3 grid

3^{rd} – 4^{th} 8^{ths} in Y

rholut 7

opacity 3

distance from midplane 1.6

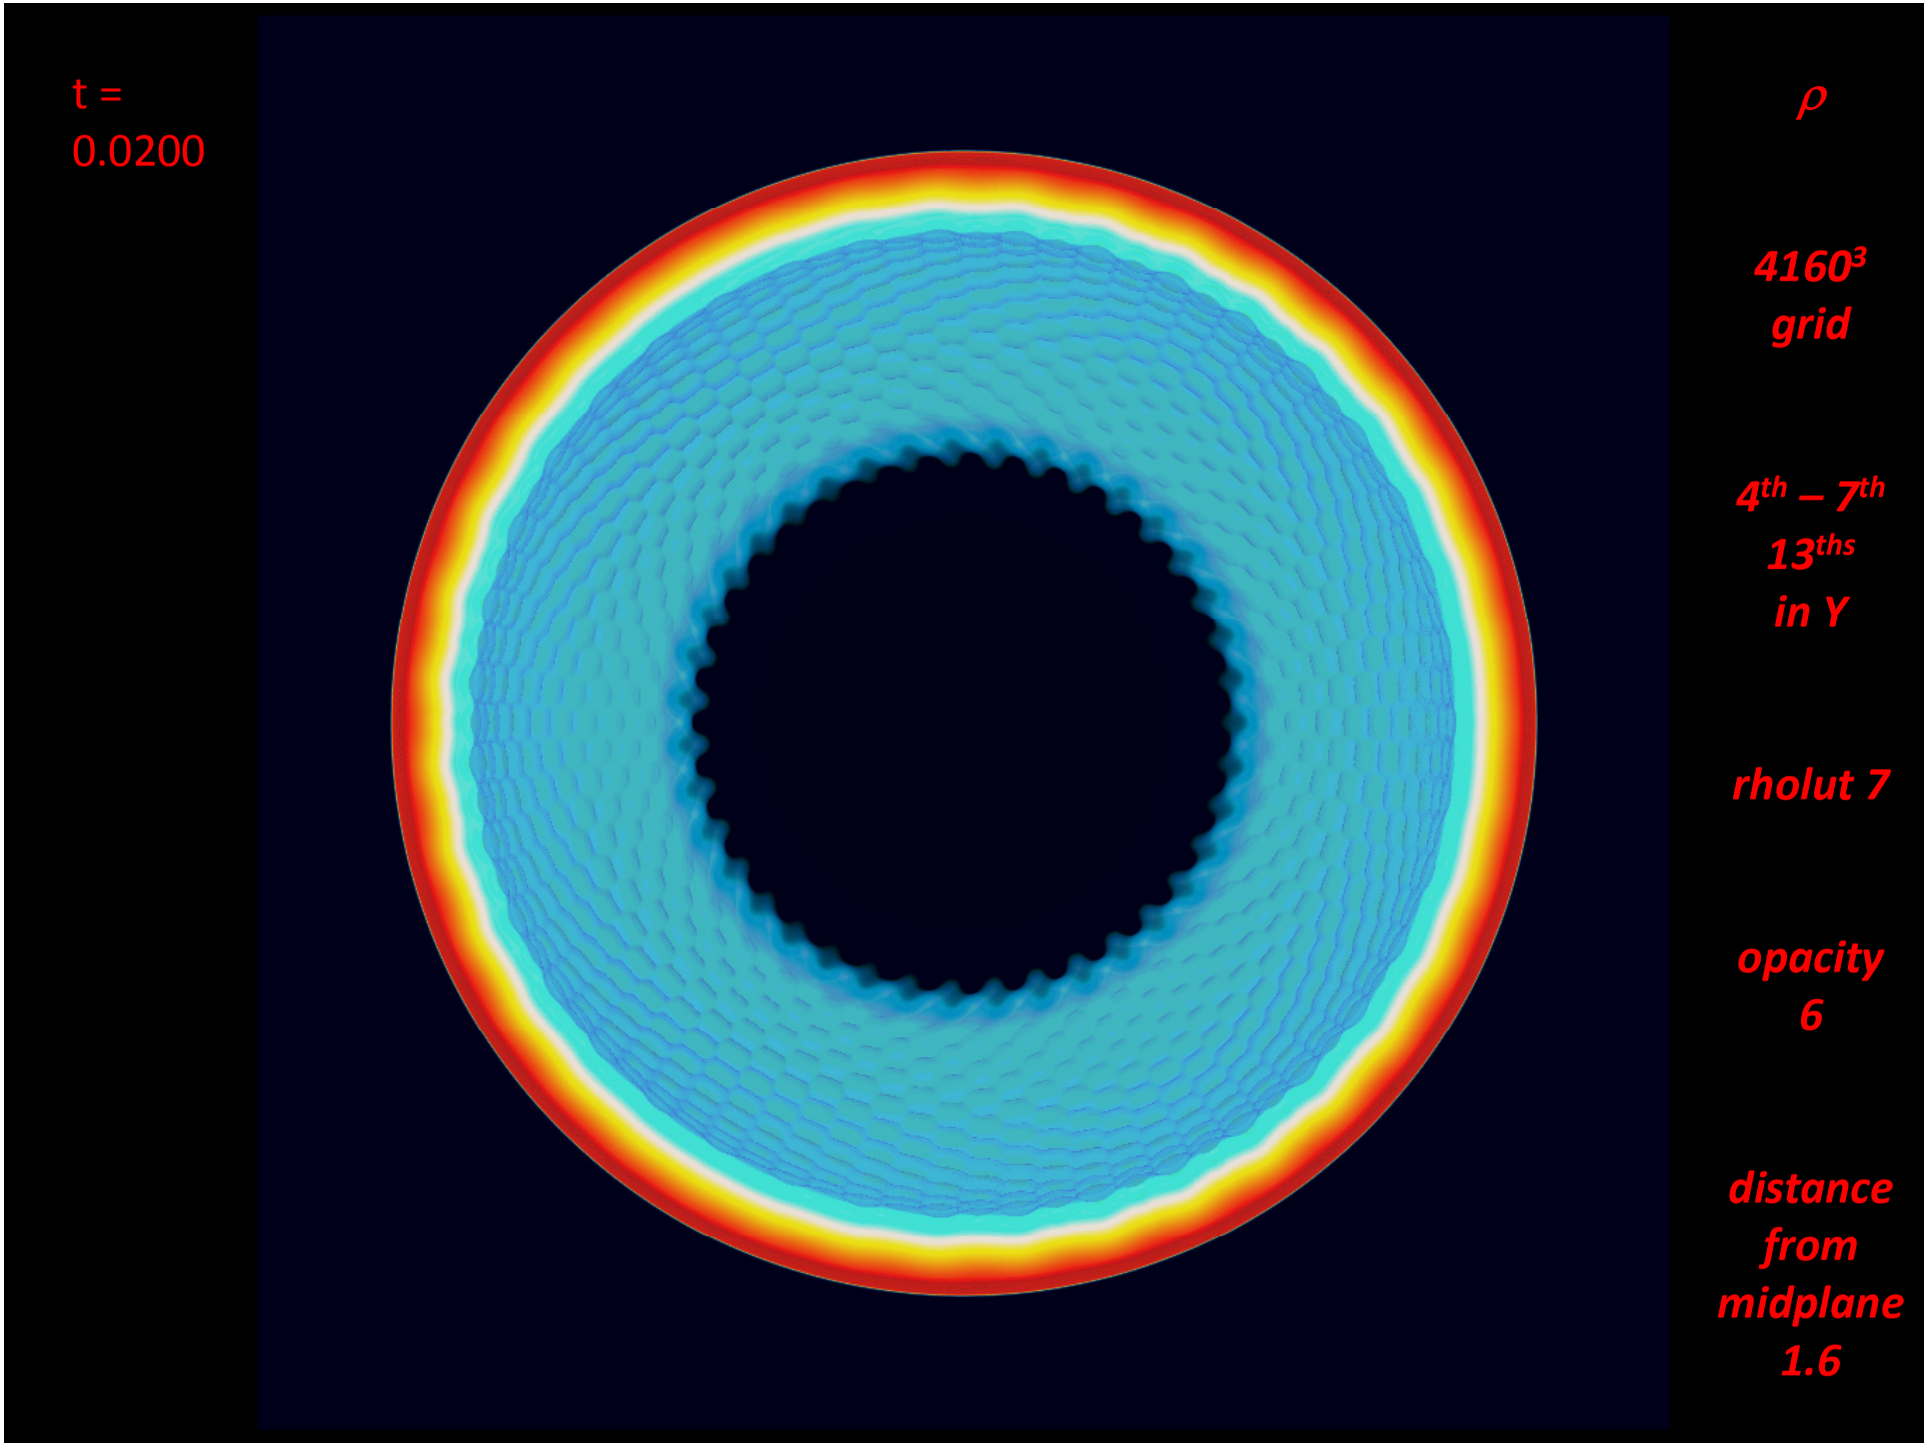
t =
0.0200
ρ
4160³
grid
4th – 7th
13ths
in Y
rholut 7
opacity
6
distance
from
midplane
1.6

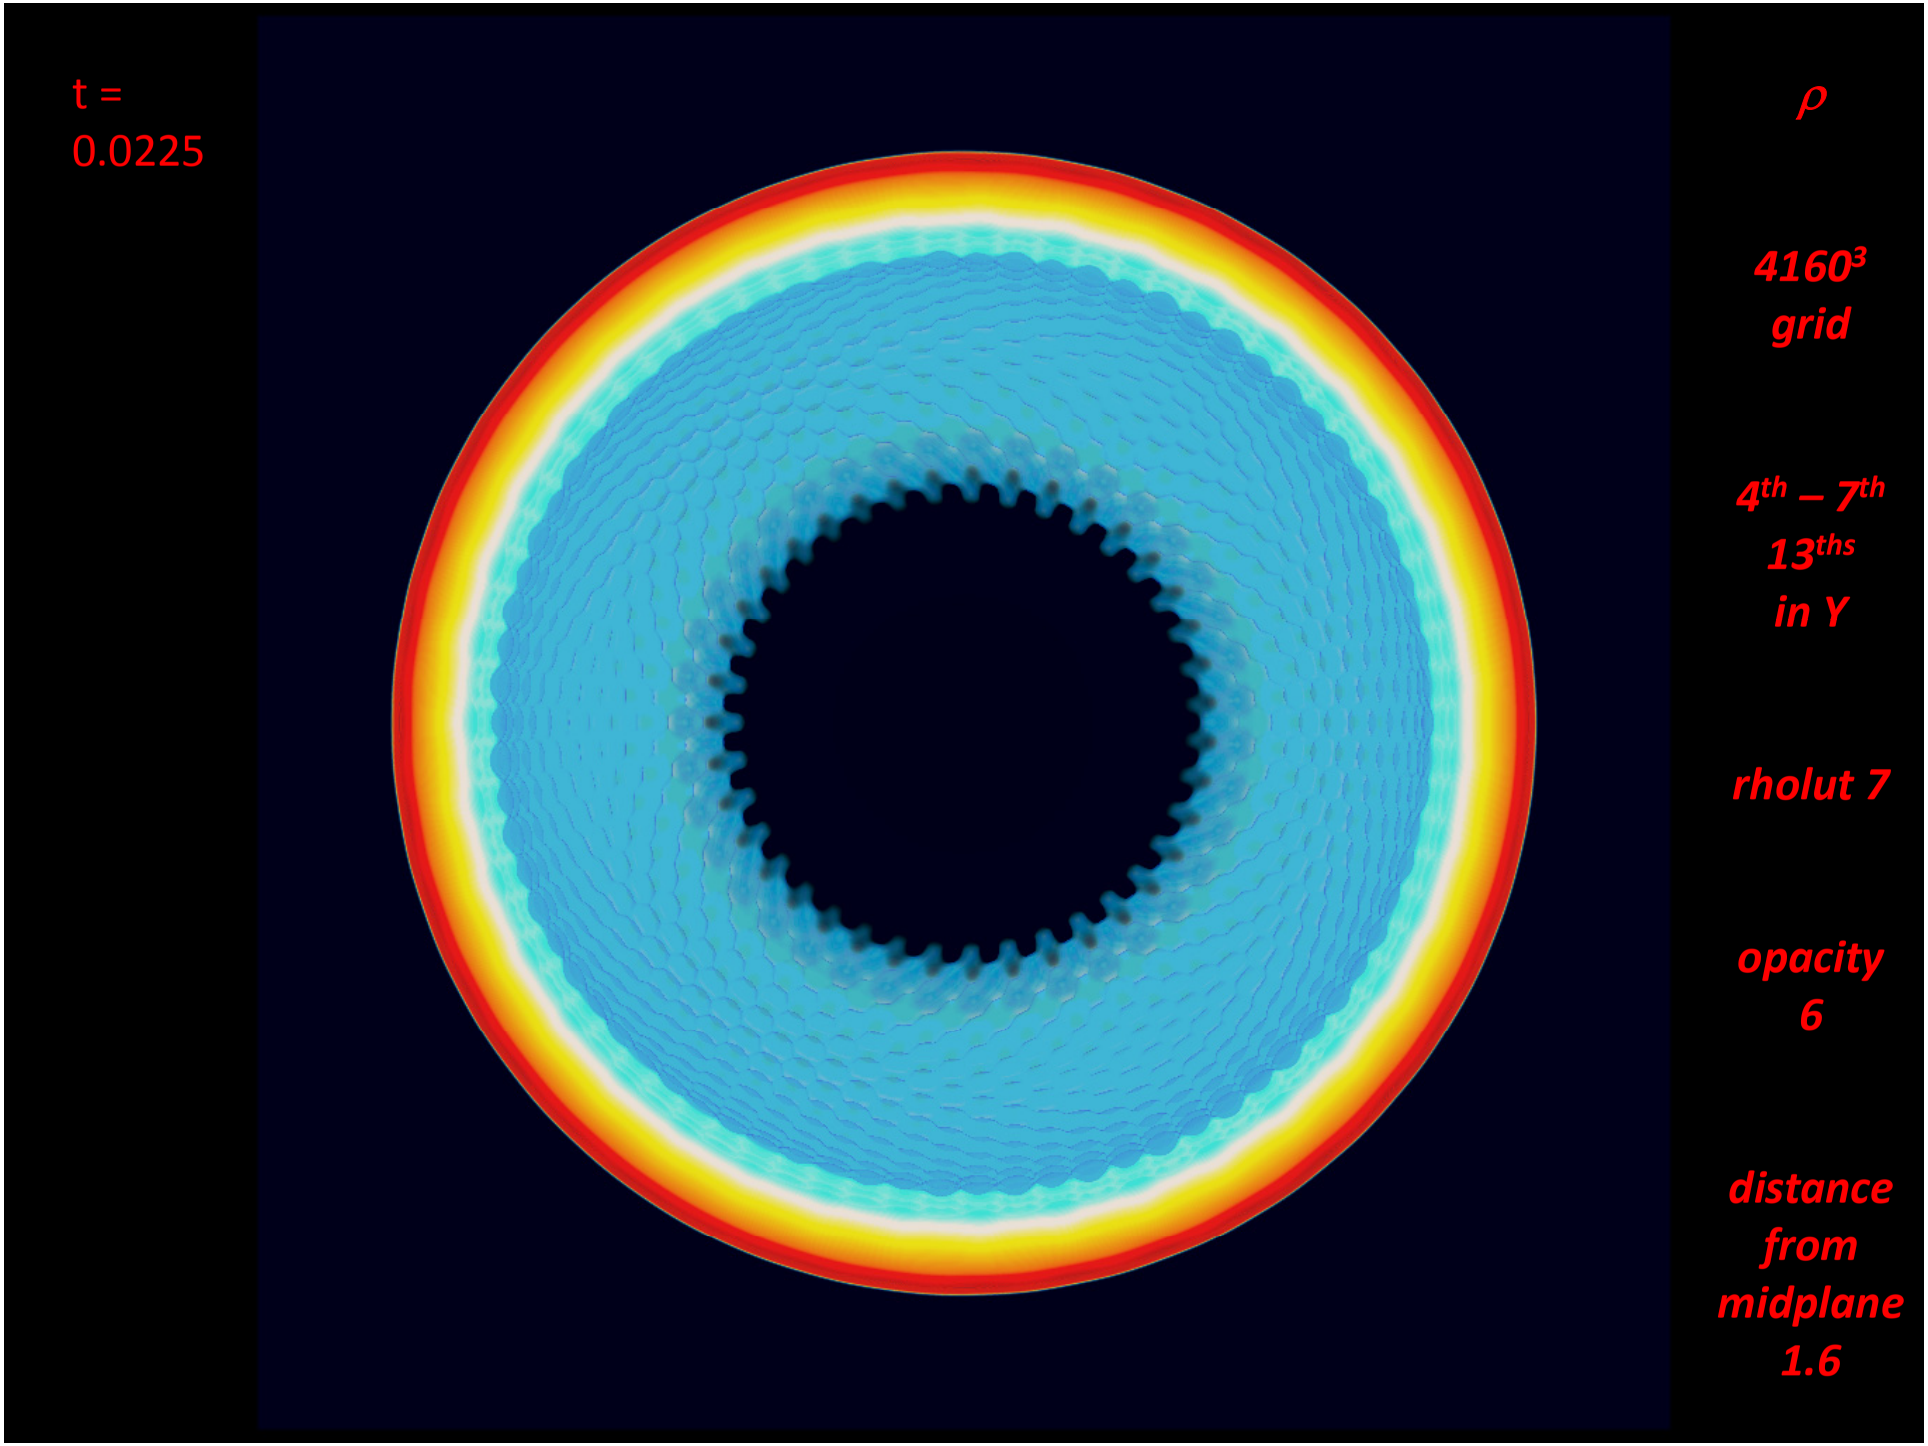
t =
0.0225
ρ
4160³
grid
4th – 7th
13ths
in Y
rholut 7
opacity
6
distance
from
midplane
1.6

t = 0.0250

At about this time, the strong rarefaction traveling outward through the dense shell has reached the surface, accelerating it more rapidly inward and setting off a strong Rayleigh-Taylor instability of this surface.

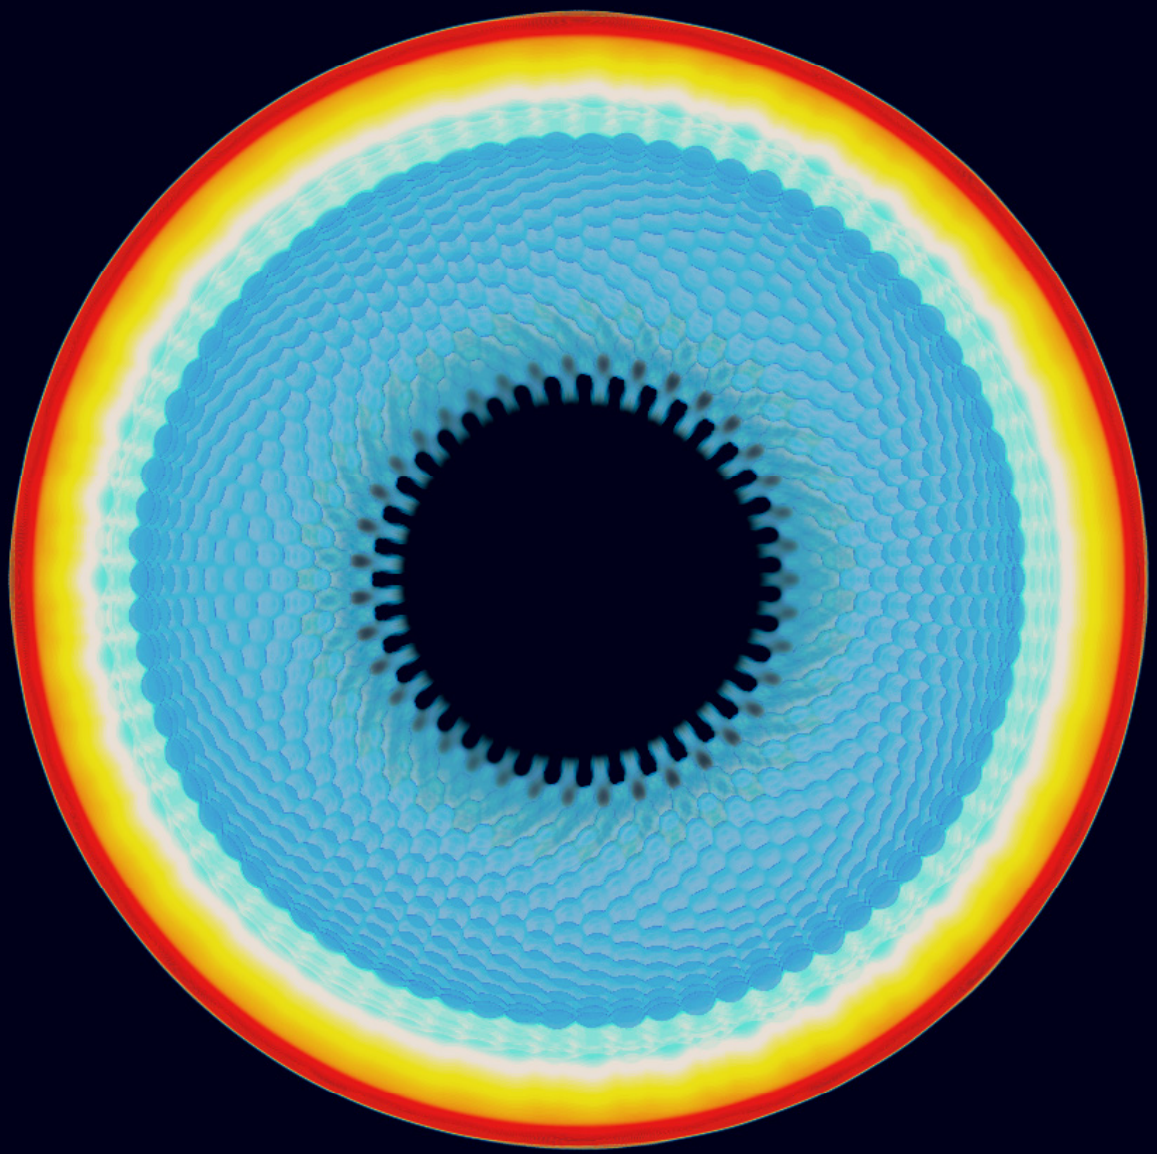

ρ

4160^3 grid

4^{th} – 7^{th} 13^{ths} in Y

rholut 7

opacity 6

distance from midplane 1.6

t = 0.0275

The imprint of the disturbance of the inner shell surface on the outer surface is now apparent, as is the growth of the Rayleigh-Taylor instability at that surface.

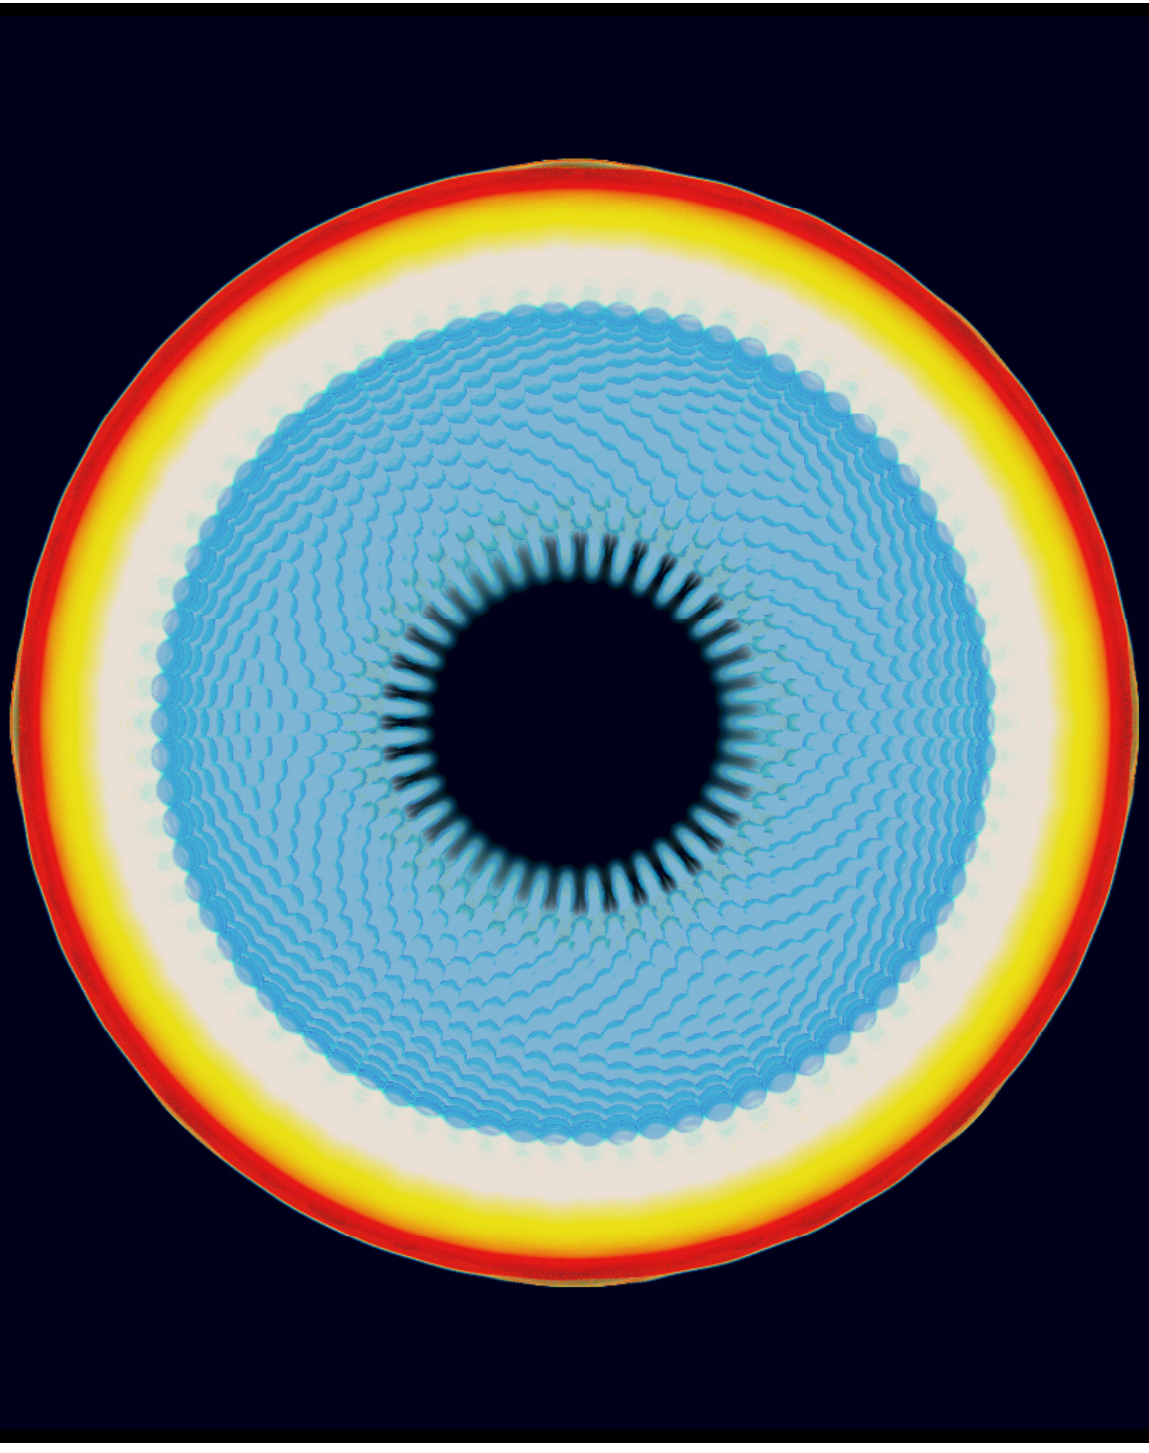

ρ

4160^3 grid

4^{th} – 7^{th} 13^{ths} in Y

rholut 7

opacity 6

distance from midplane 1.6

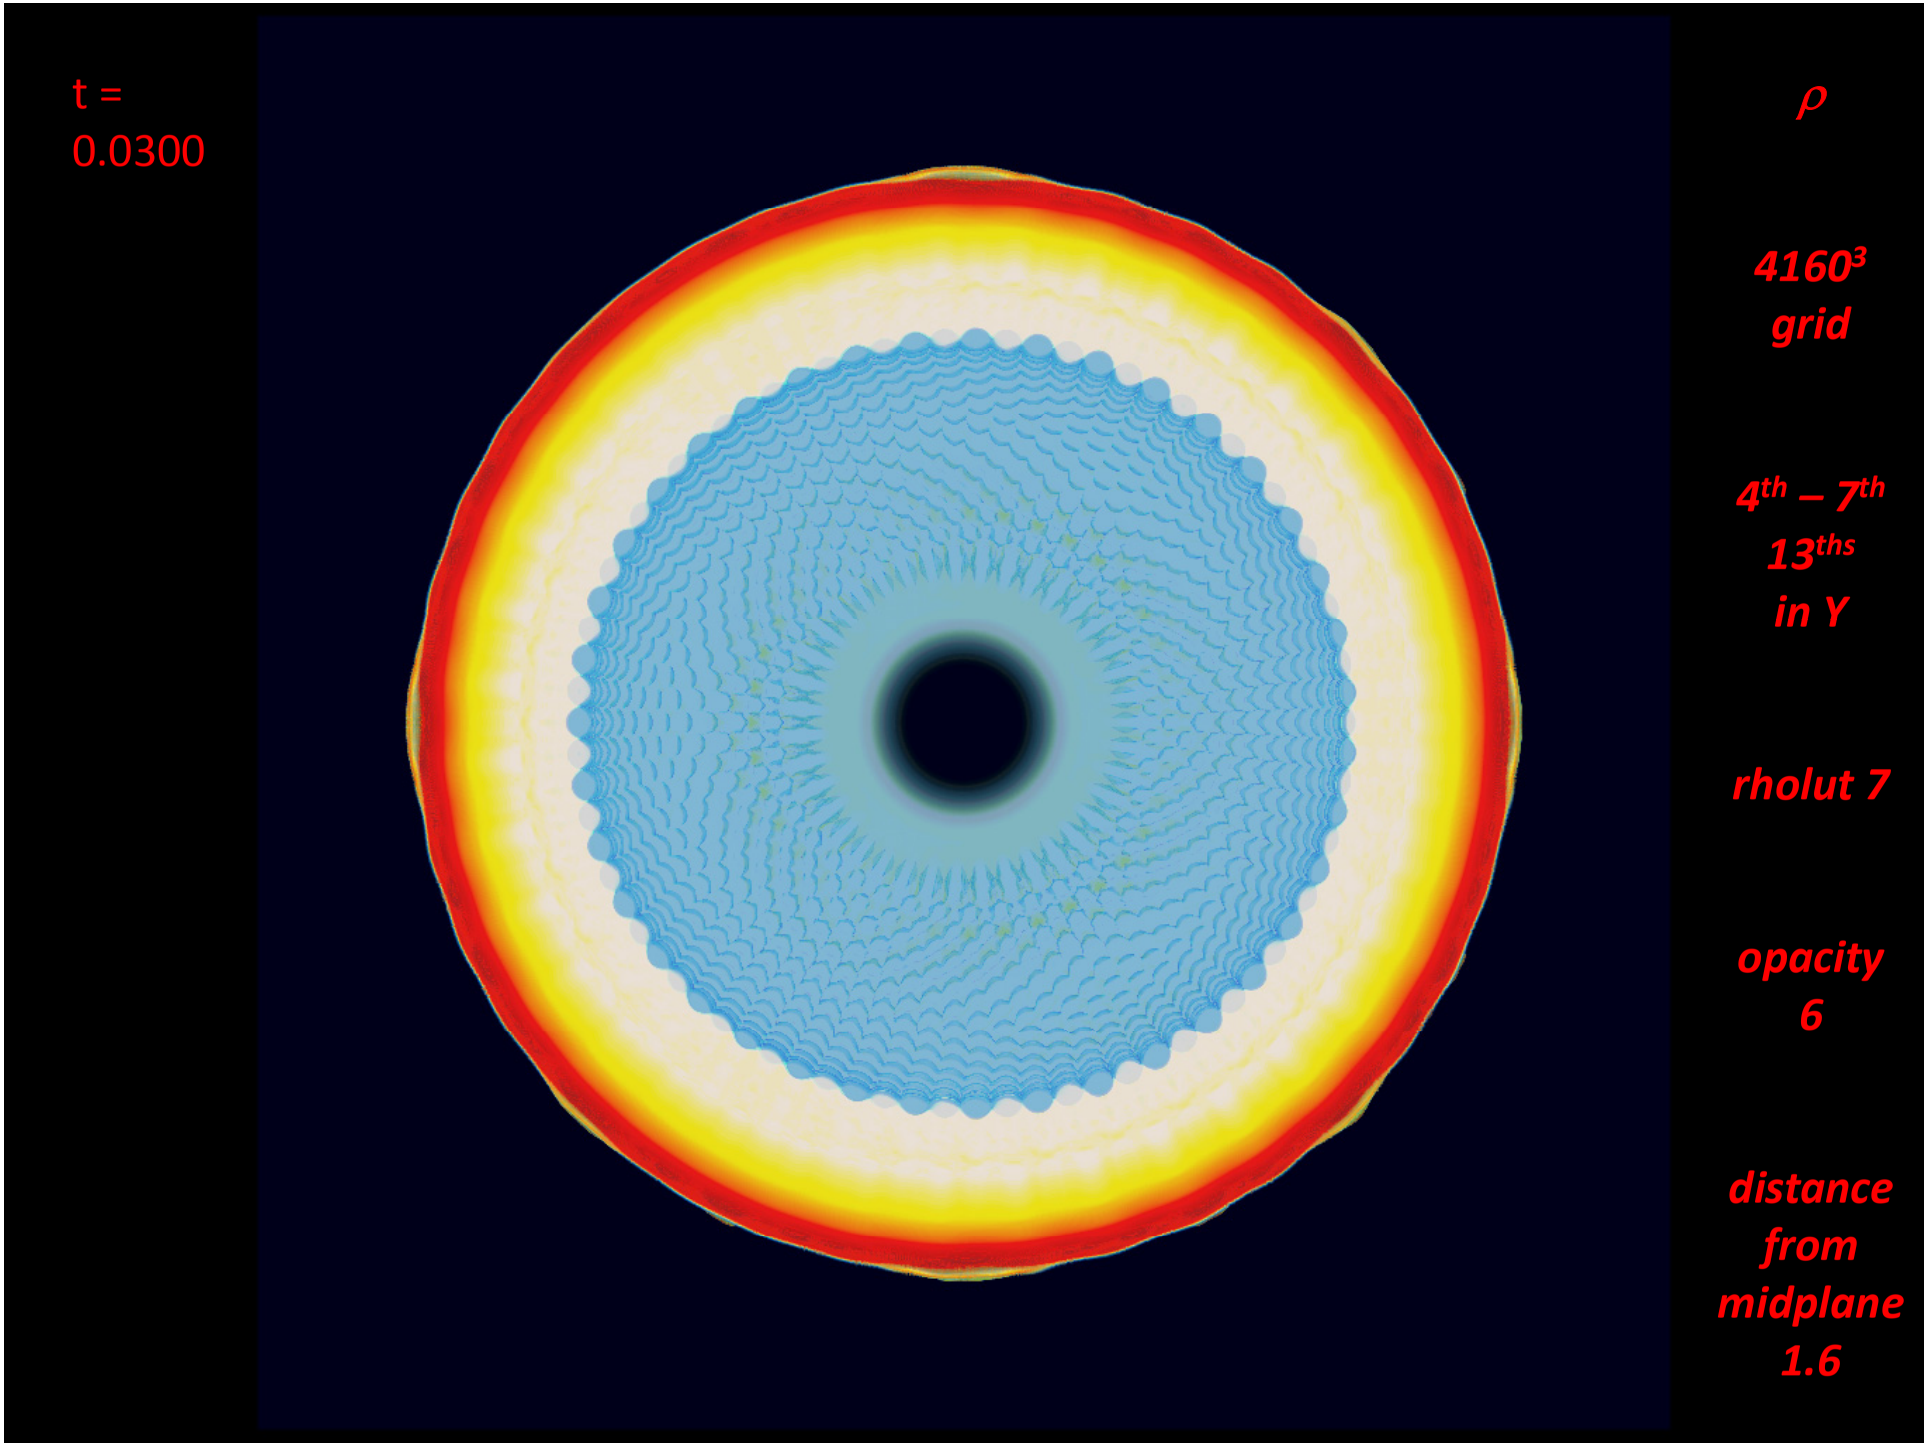
t =
0.0300
ρ
4160³
grid
4th – 7th
13ths
in Y
rholut 7
opacity
6
distance
from
midplane
1.6

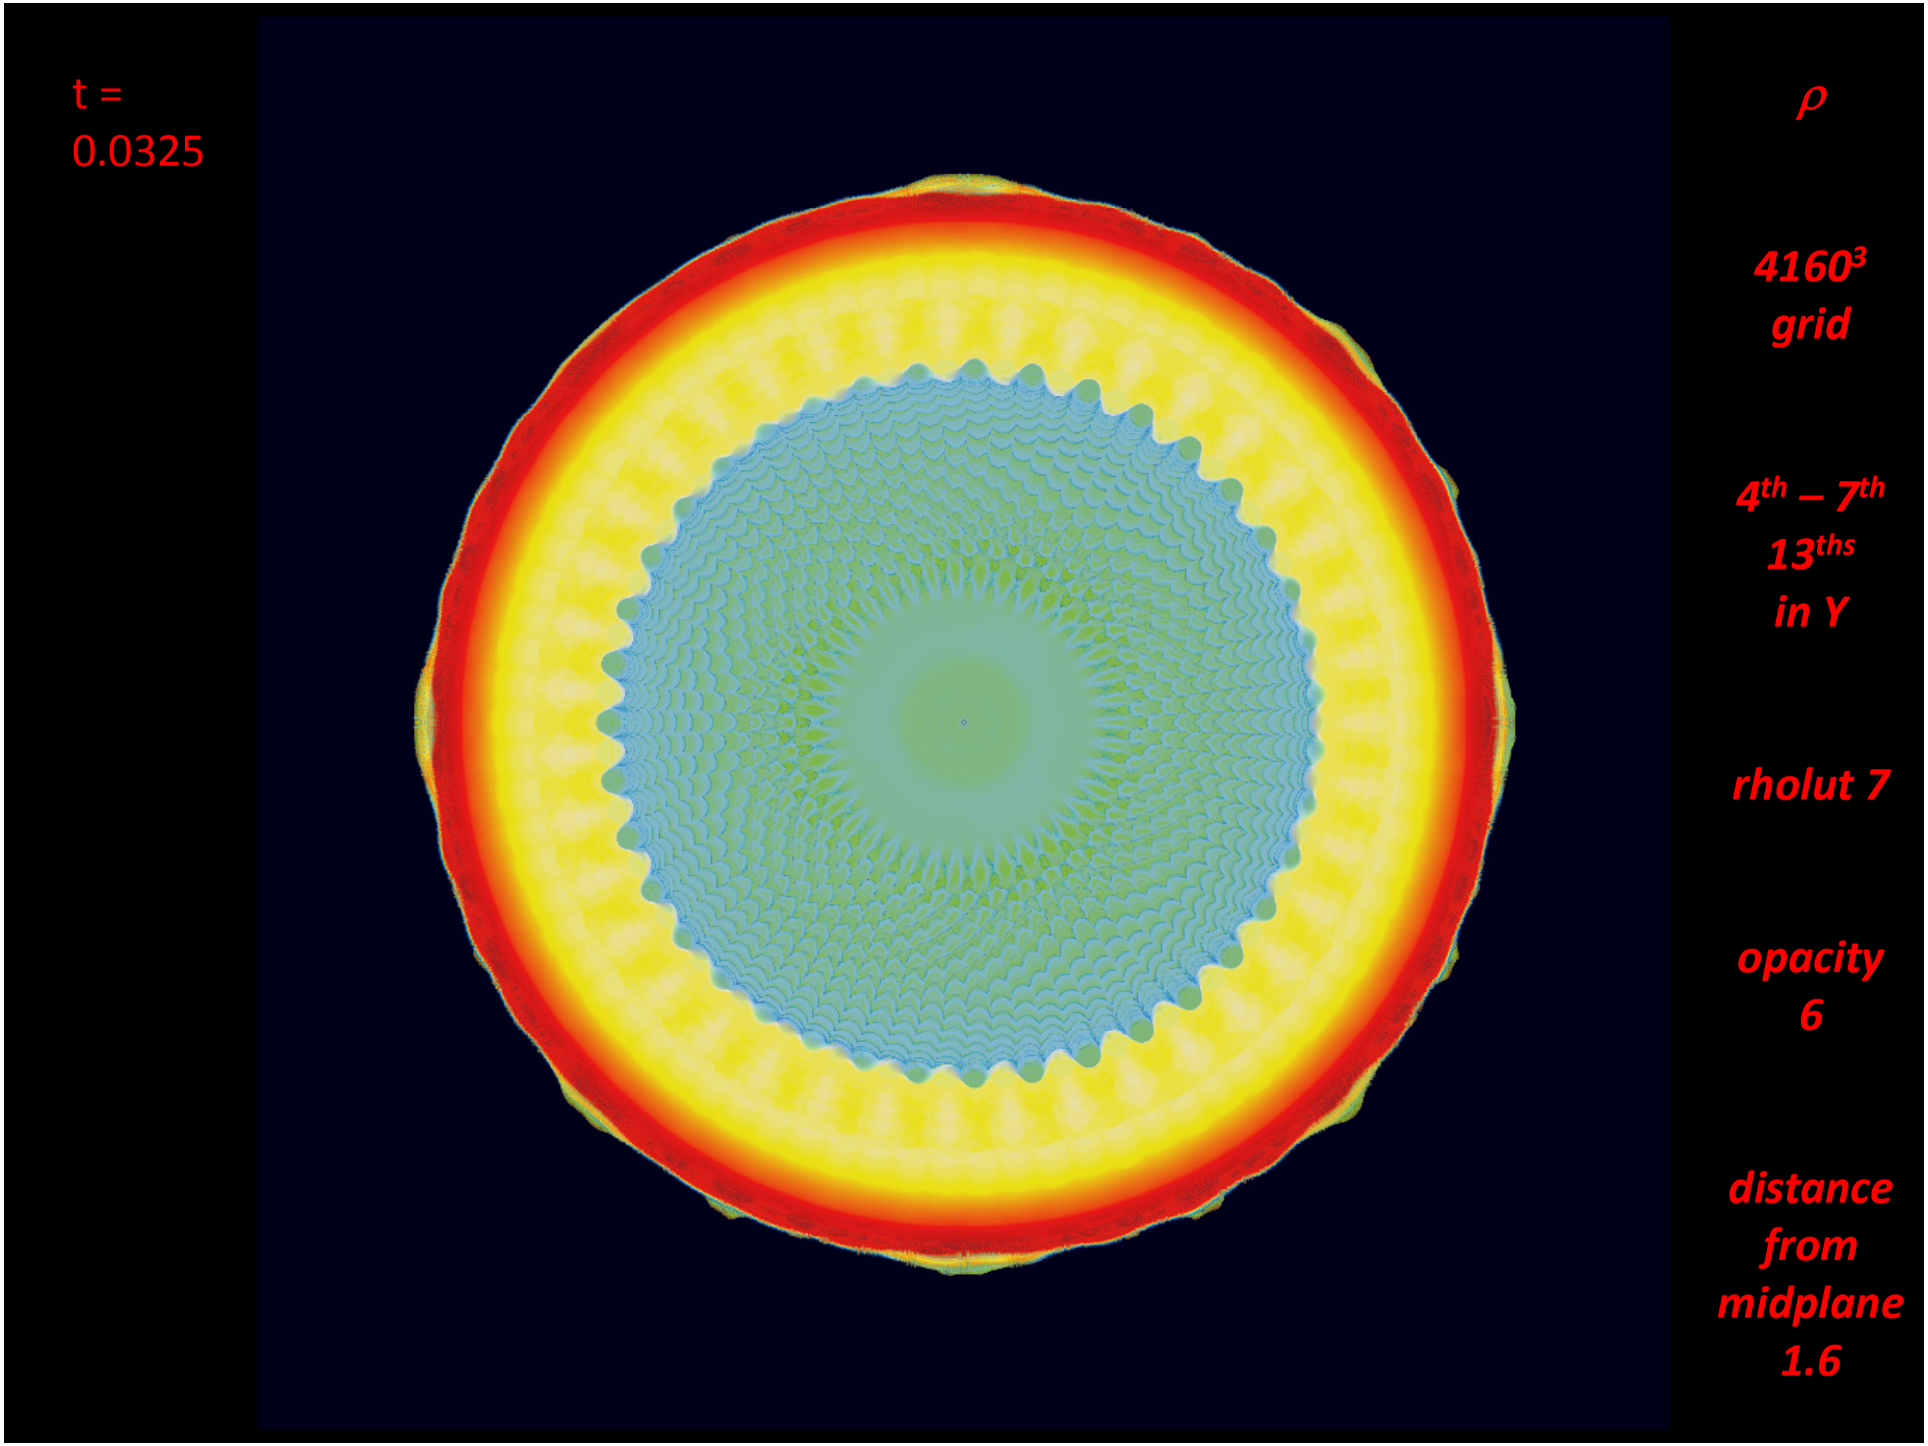
t =
0.0325
ρ
4160³
grid
4th – 7th
13ths
in Y
rholut 7
opacity
6
distance
from
midplane
1.6

t = 0.0350

Amplification of beat frequencies caused by grid plane crossings of the very thin outer surface of the shell by the Rayleigh-Taylor instability is especially evident now at the points of the compass.

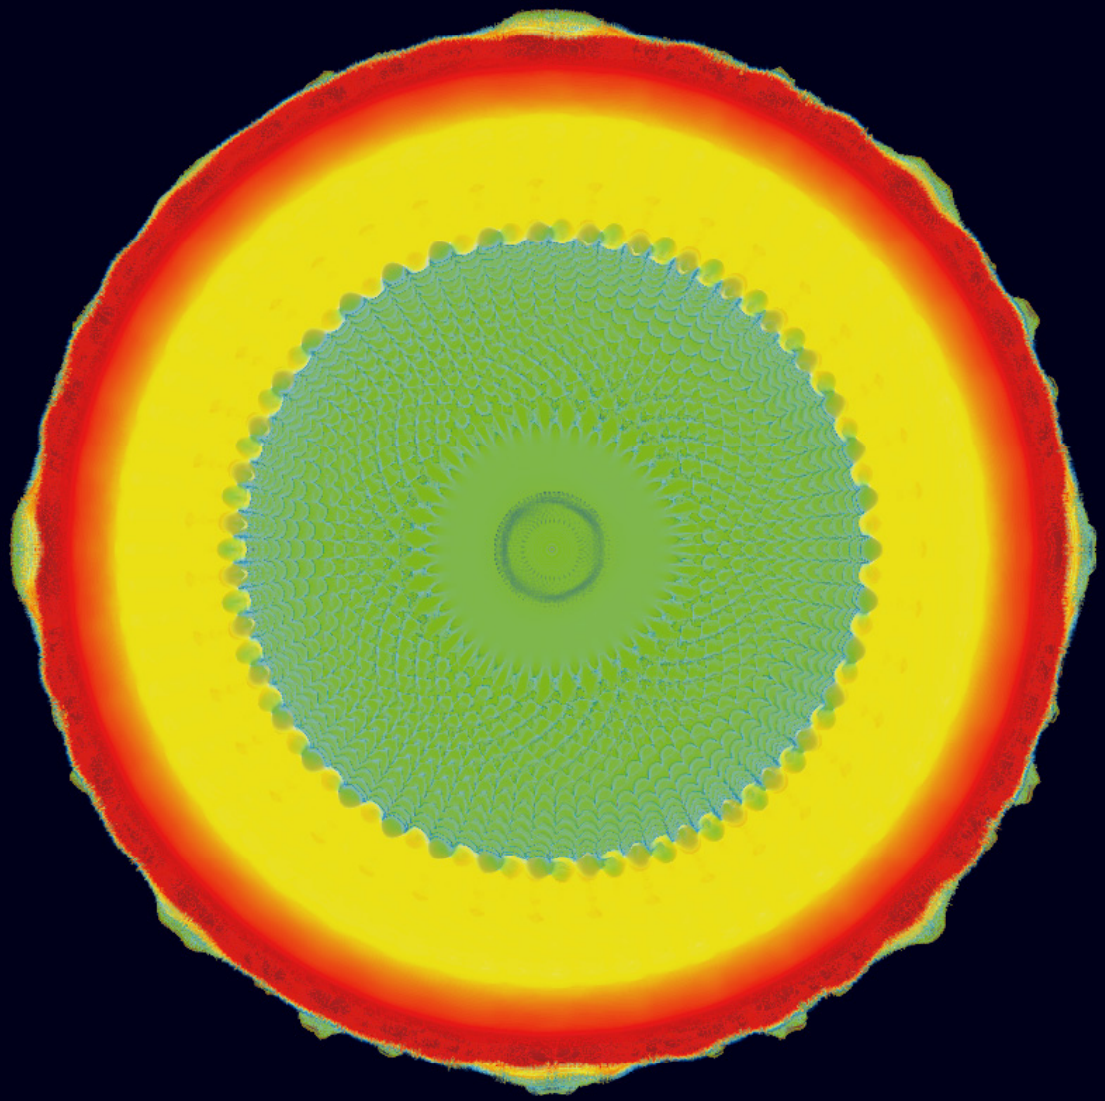

ρ

4160^3 grid

4^{th} – 7^{th} 13^{ths} in Y

rholut 7

opacity 6

distance from midplane 1.6

These amplified disturbances of numerical origin have only a visibly negligible effect upon the dynamics of the all-important inner surface of the shell.

t = 0.0375

The shock that has been reflected at the origin is now seen propagating back outward. It is clearly very round, as its very strong stability assures that it must be, even though on our Cartesian grid it cannot possibly have reflected correctly when right at the origin.

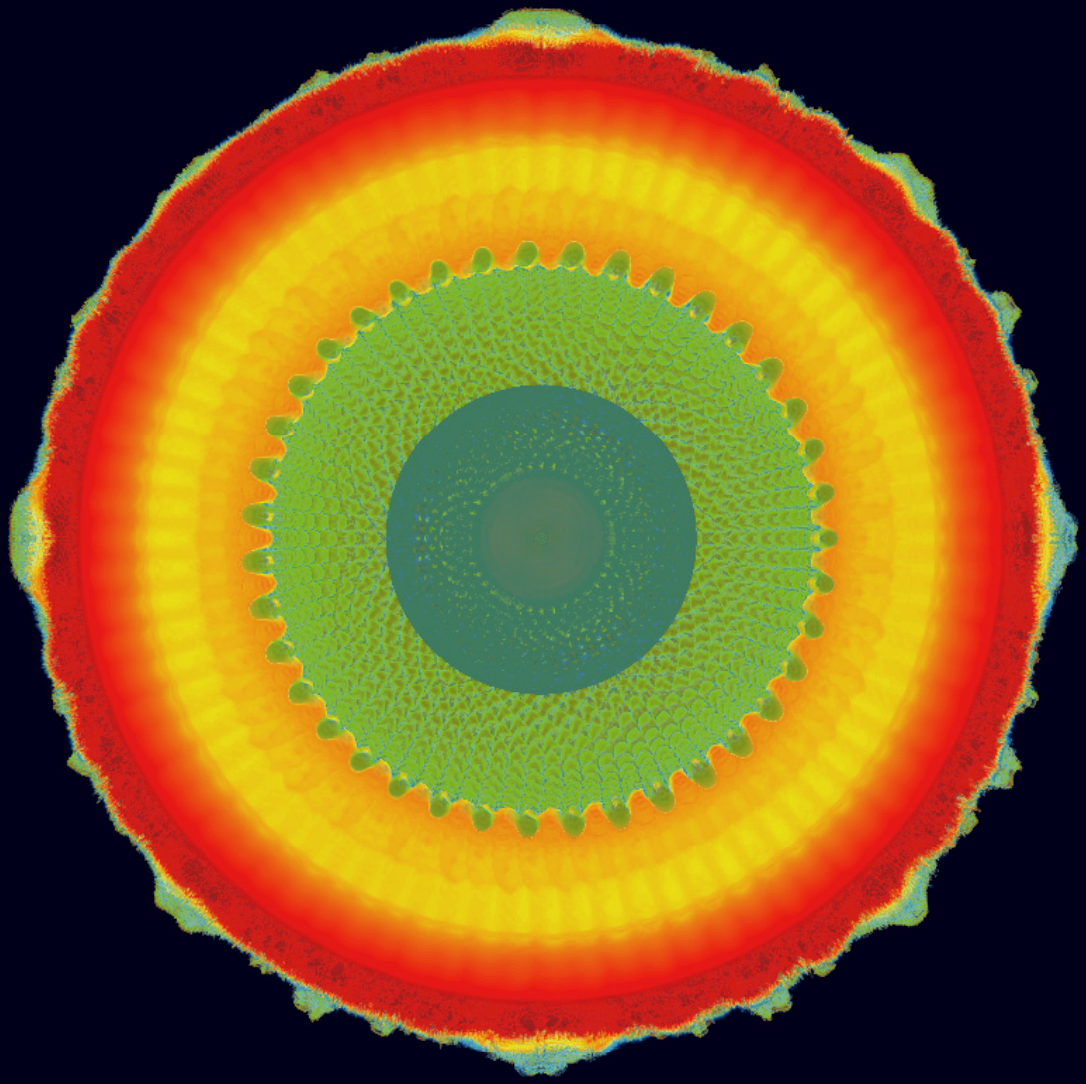

ρ

4160^3 grid

4^{th} – 7^{th} 13^{ths} in Y

rholut 7

opacity 6

distance from midplane 1.6

t = 0.0400

The outward propagating shock has just recently struck the dense shell, producing a strong deceleration and a strong Rayleigh-Taylor instability of that shell's inner surface. The Richtmyer-Meshkov fingers are now very strongly amplified. They were in the nonlinear regime of amplitude when struck, and a broad spectrum of small-scale modes results.

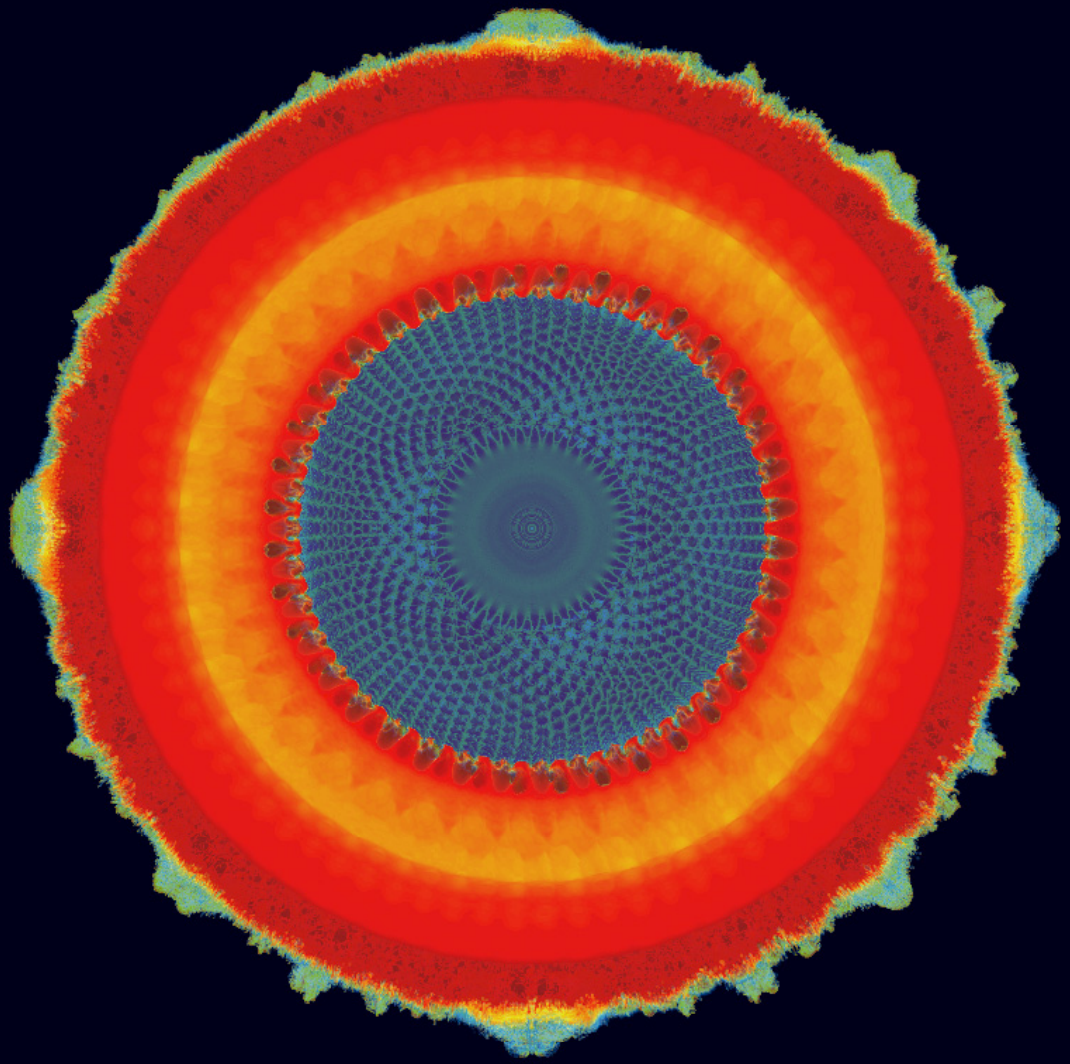

ρ

4160^3 grid

4^{th} – 7^{th} 13^{ths} in Y

rholut 7

opacity 6

distance from midplane 1.6

t =
0.0425

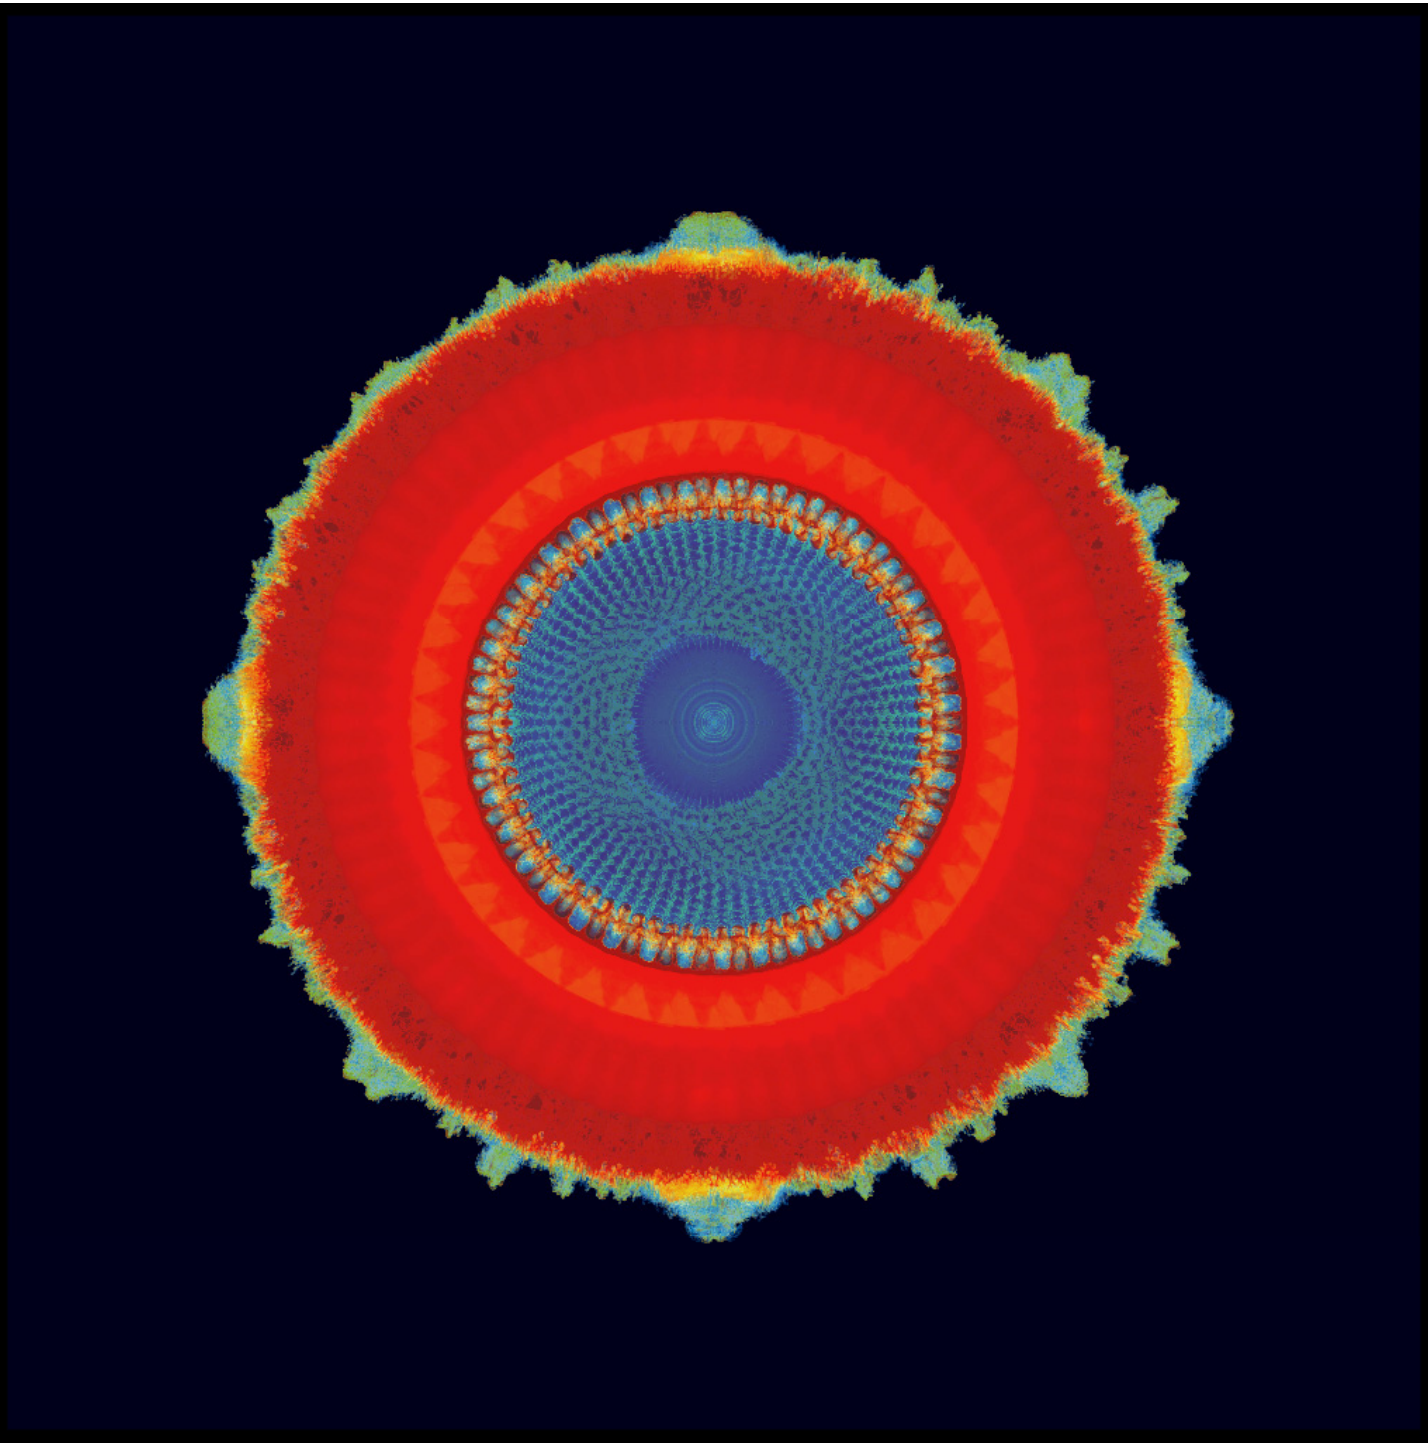

ρ

4160^3
grid

$4^{th} - 7^{th}$
13^{ths}
in Y

rholut 7

opacity
6

distance
from
midplane
1.6

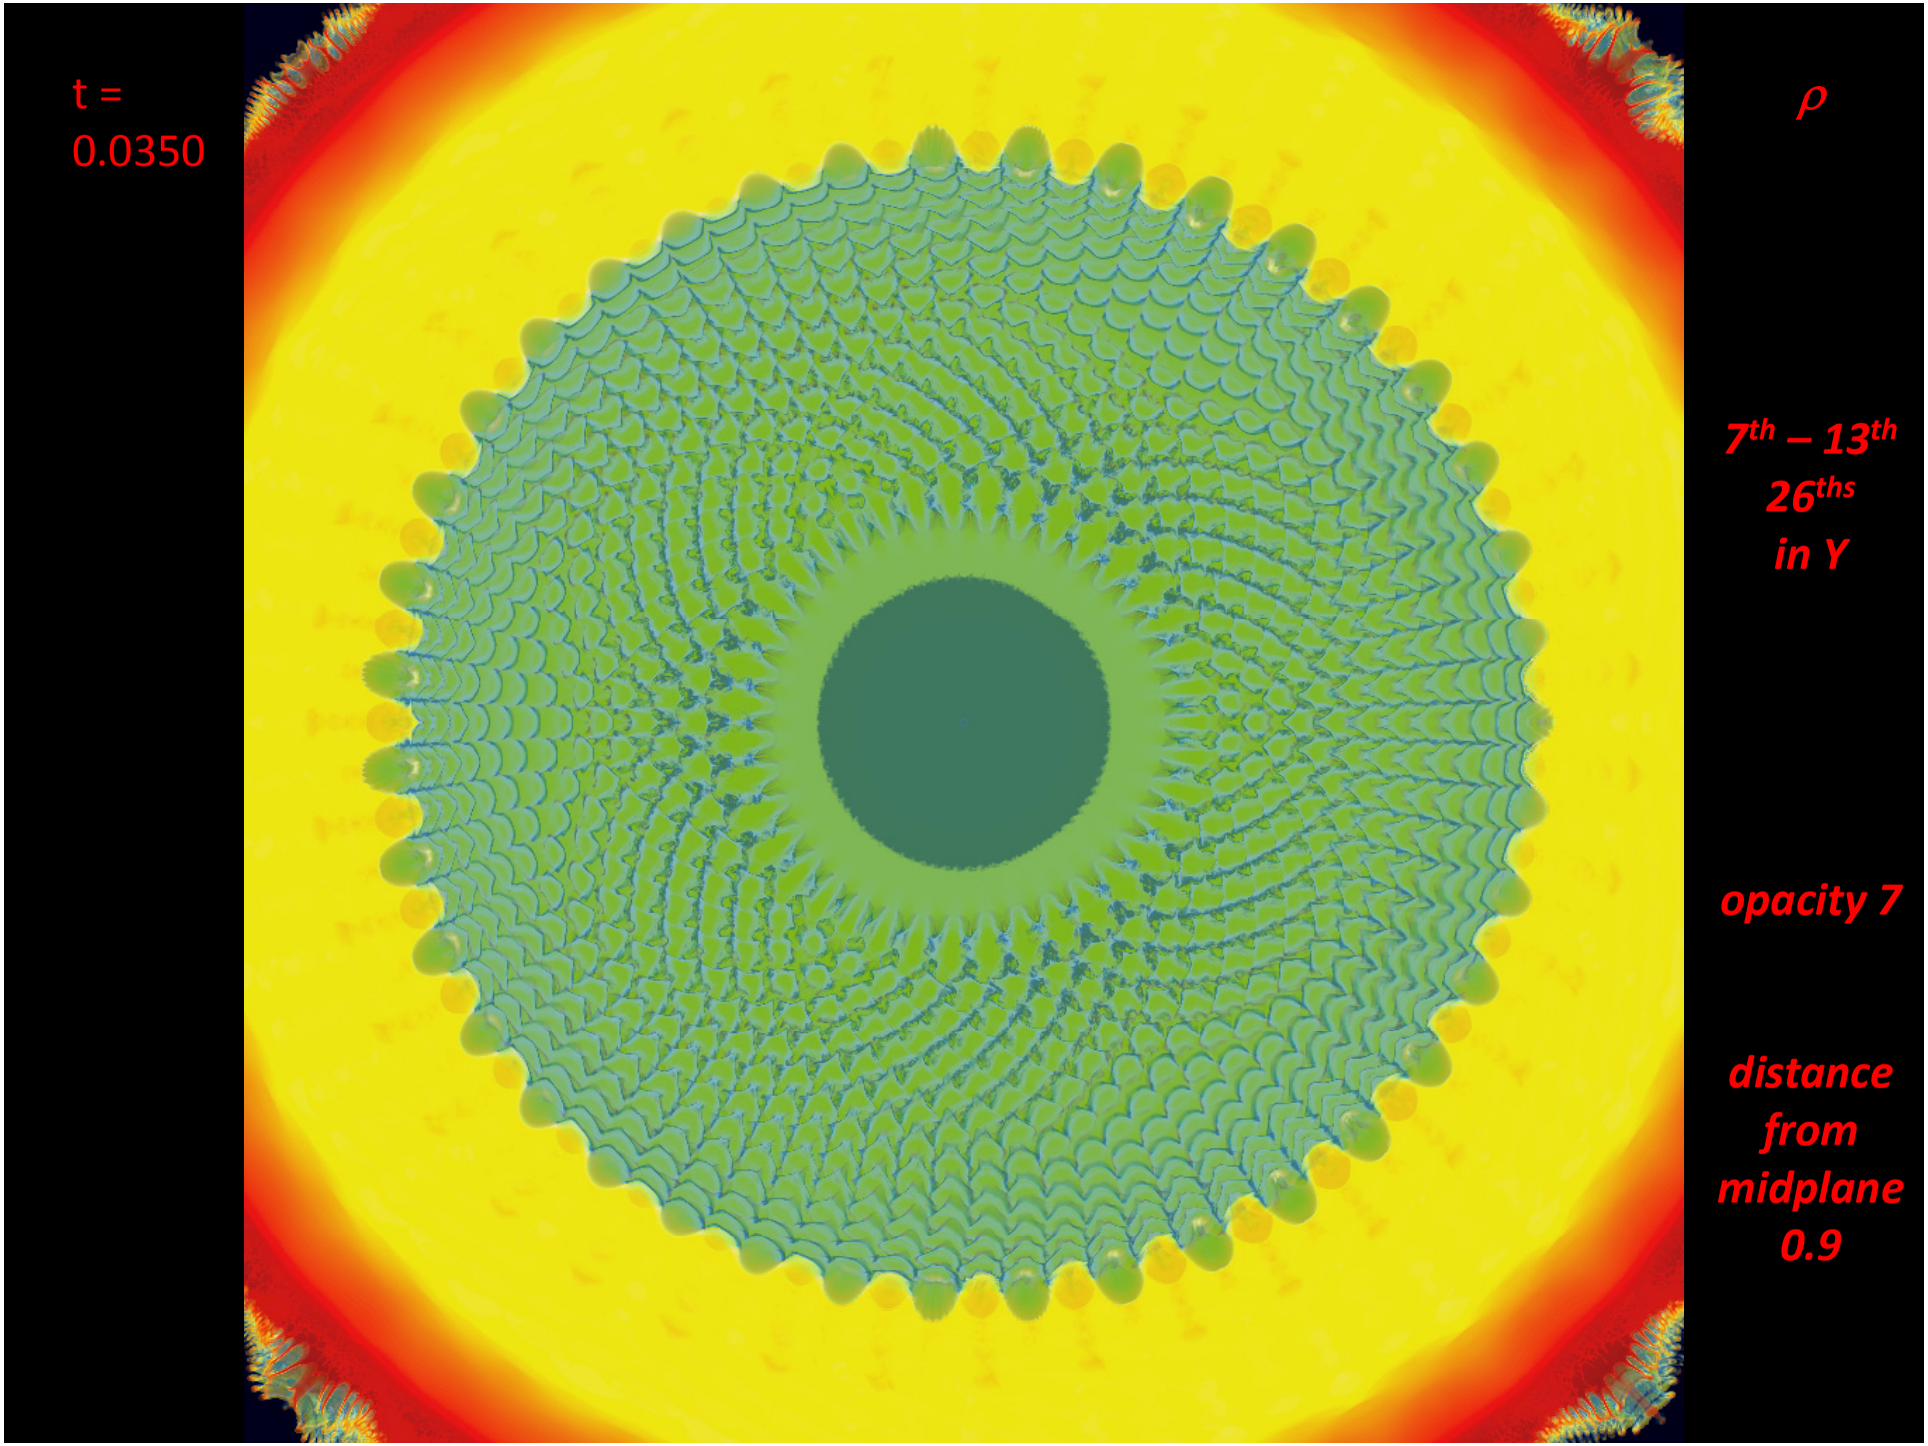
t =
0.0350
ρ
7th – 13th
26ths
in Y
opacity 7
distance
from
midplane
0.9

t = 0.0375

The out-ward propa-gating shock is just visible here about half way out to the inner surf-ace of the shell.

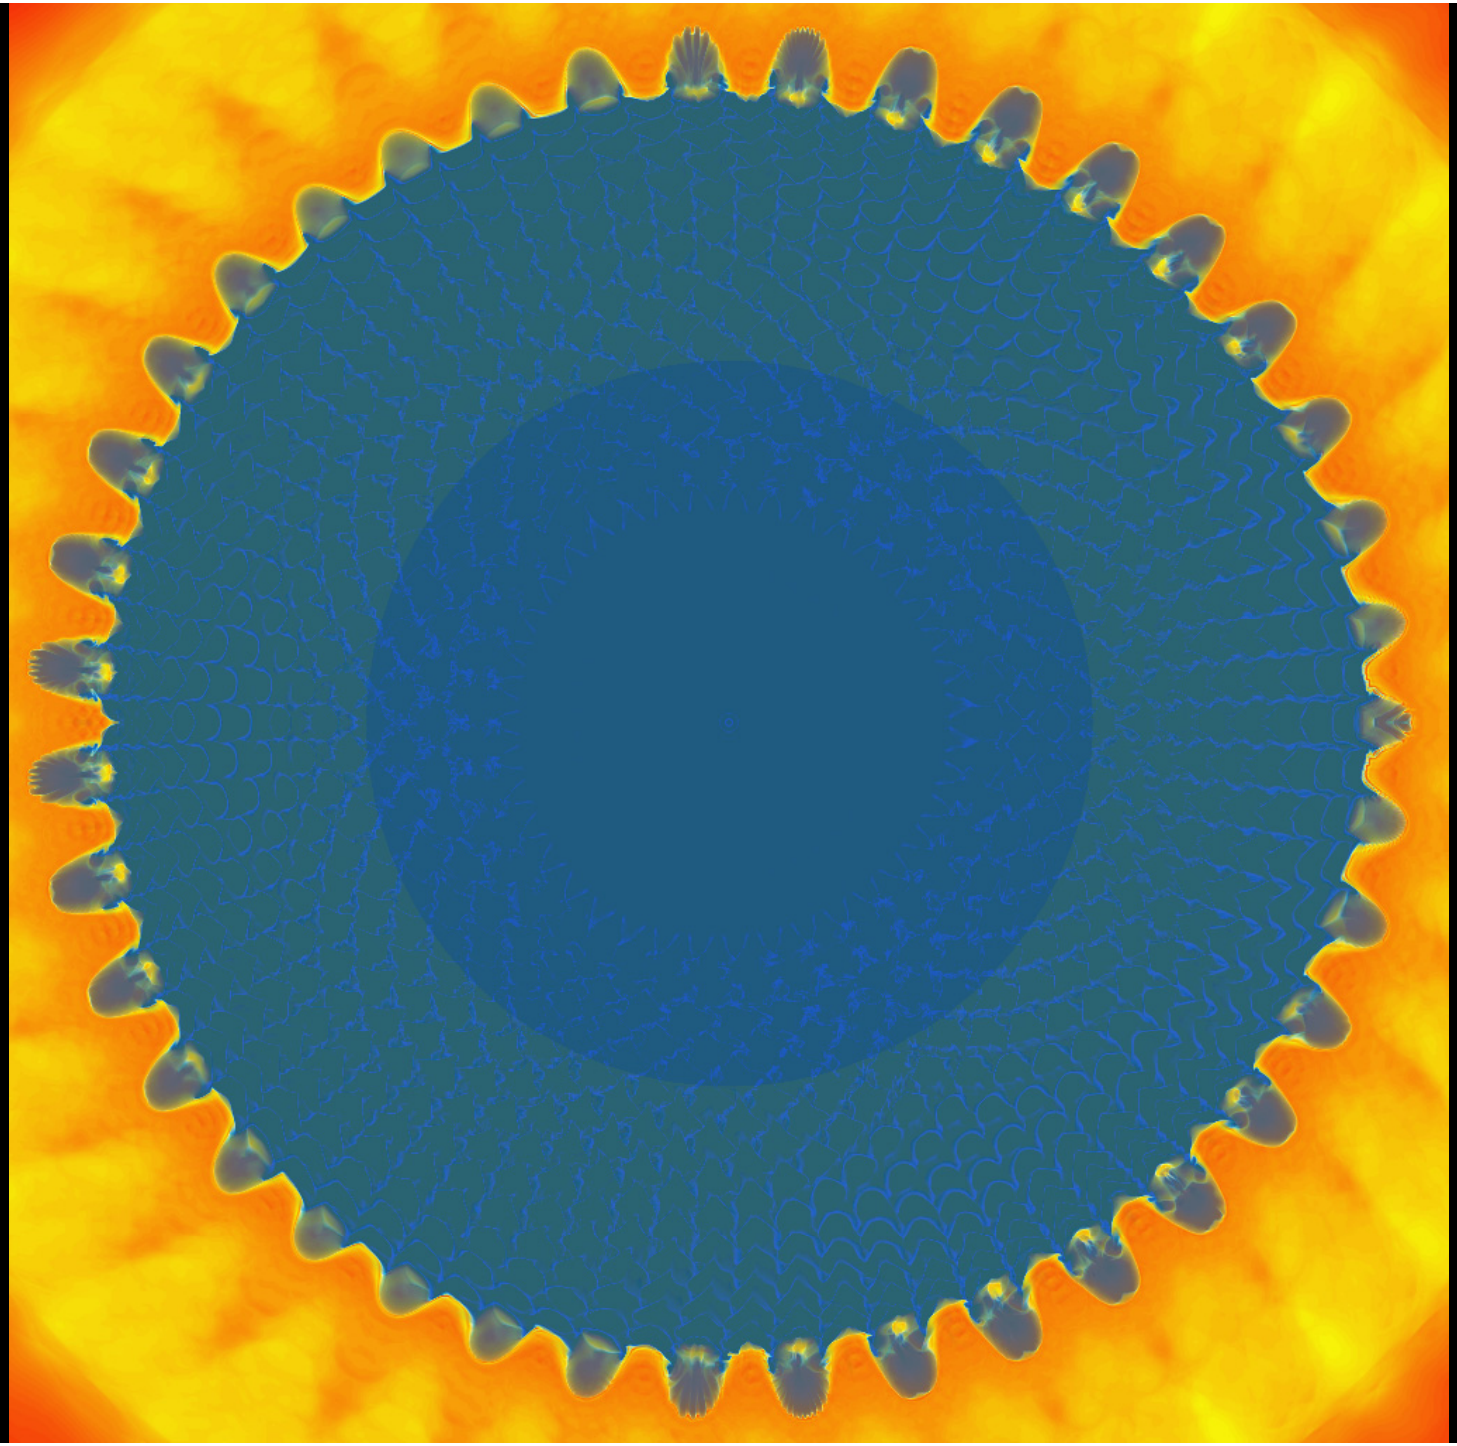

ρ

opacity 14

distance from midplane 0.7

t = 0.0400

The mode-mode modu-lation is clear here in the mode-3 pat-tern near the pole at the center of this view.

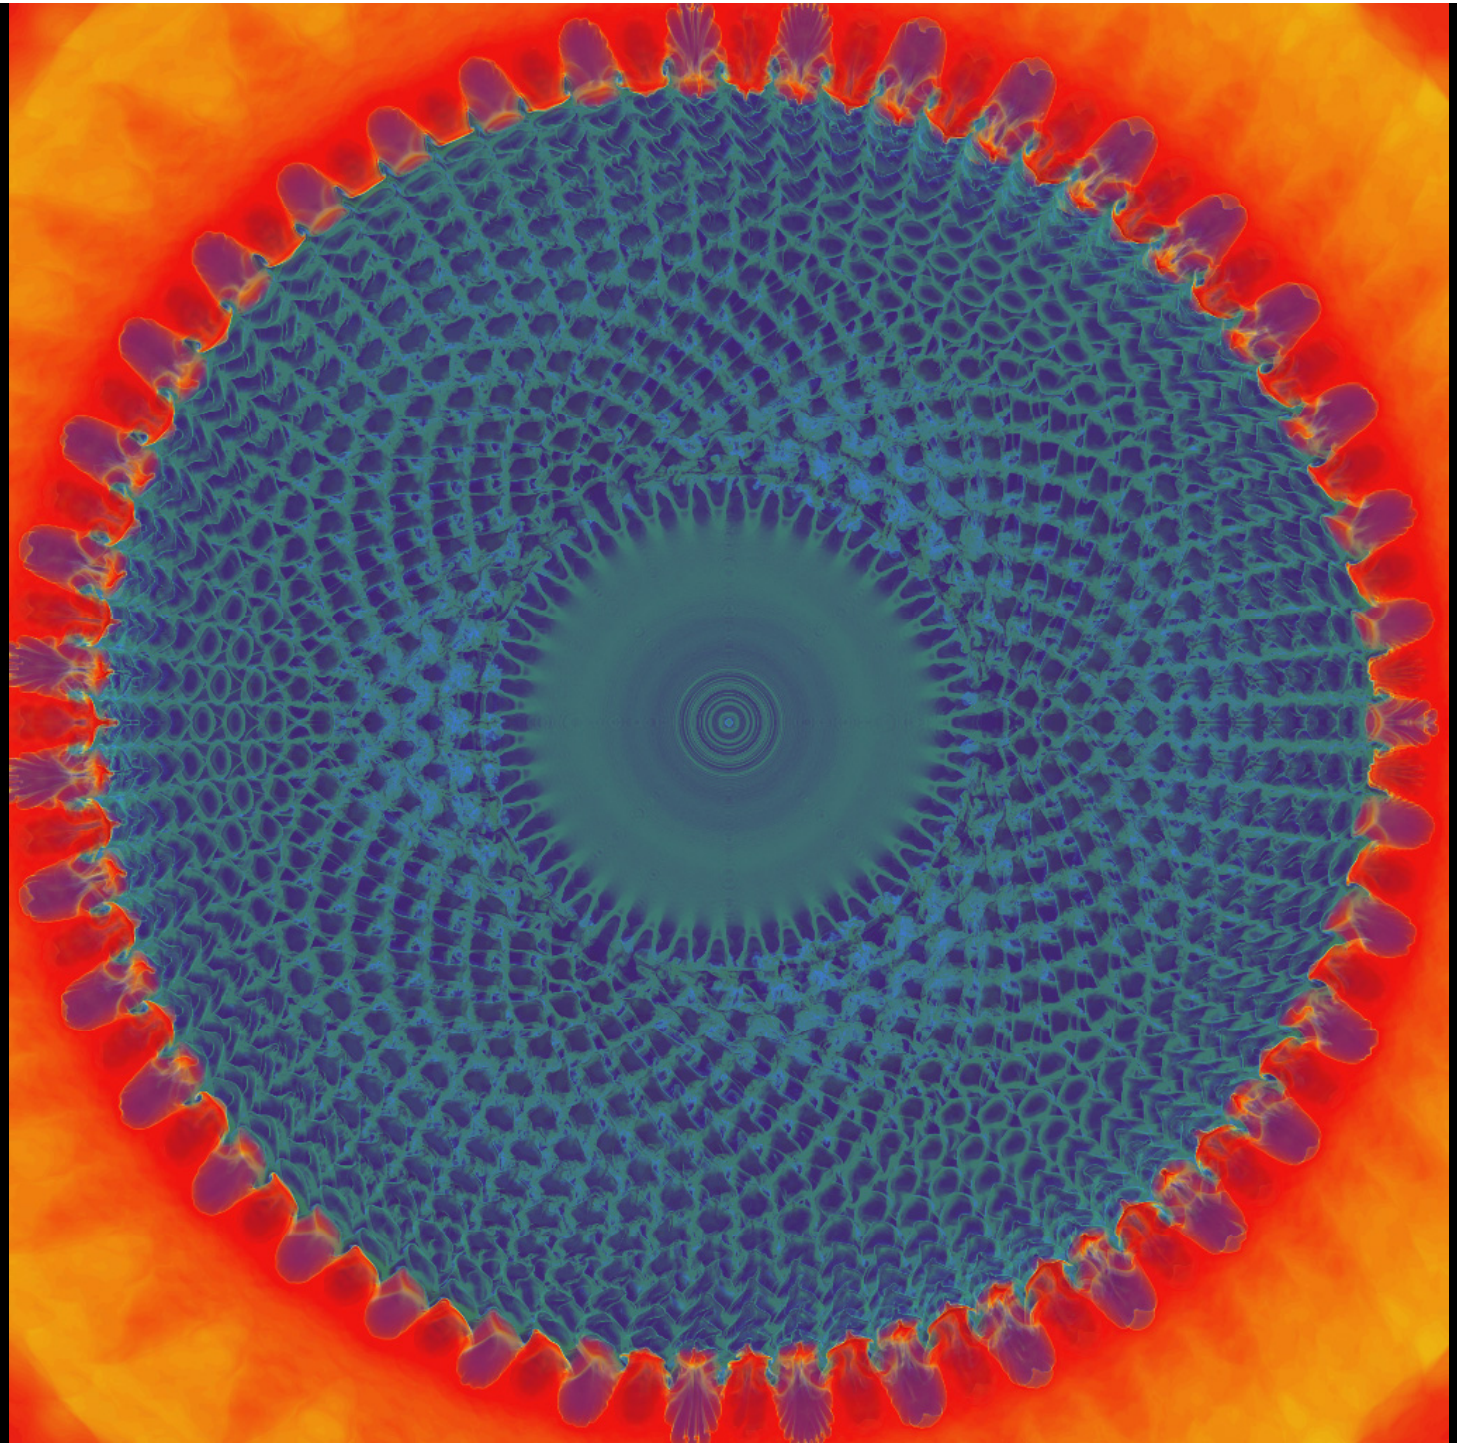

ρ

7^{th} – 13^{th} 26^{ths} in Y

opacity 7

distance from midplane 0.6

t = 0.0425

This high-opacity view makes clear that in the gas inside the shell there is a very com-plex density struct-ure.

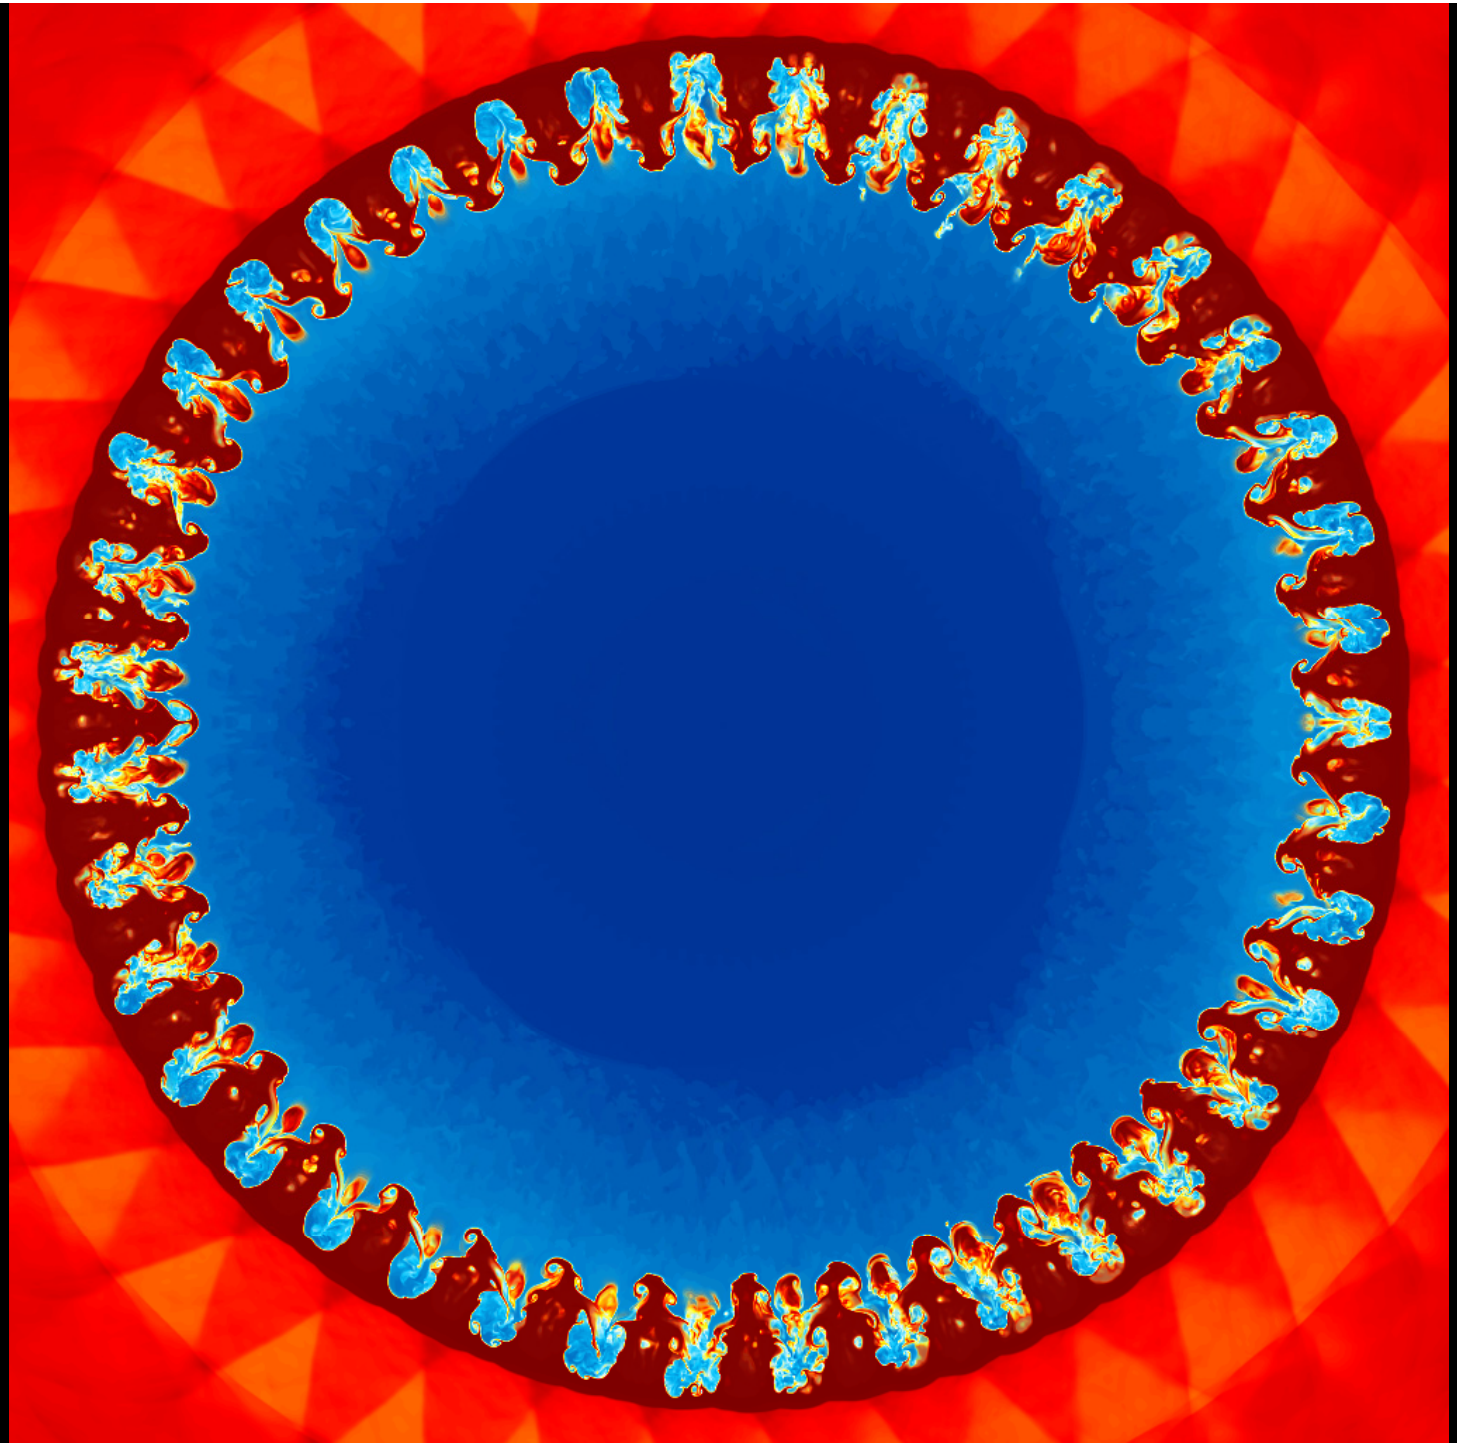

ρ

7^{th} – 13^{th} 26^{ths} in Y

thin midplane slice

opacity 112

distance from midplane 0.6

t = 0.0425

In this lower opacity view, we can see into the cavity. The mode-3 modu-lation is also clearly evident. The strong shock is very round.

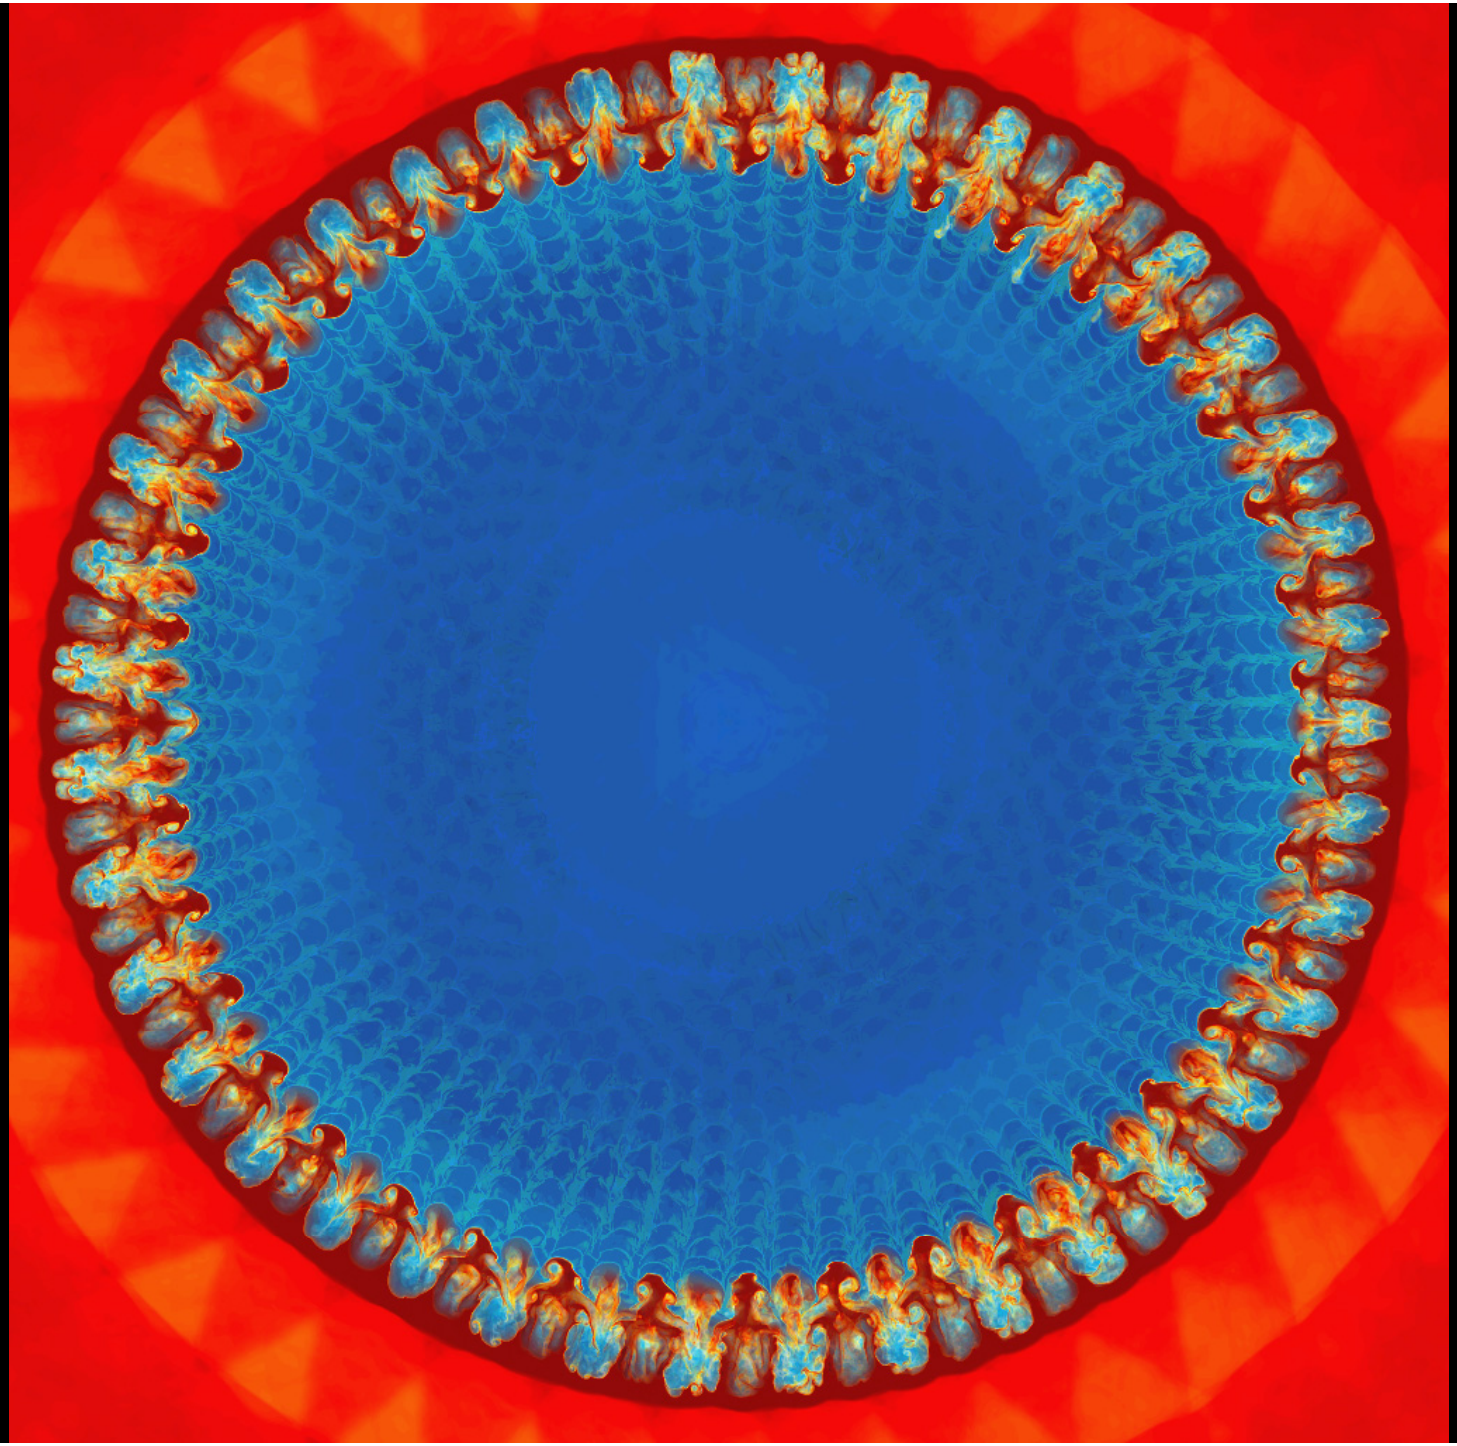

ρ

4160^3 grid

$7^{th} - 13^{th}$ 26^{ths} in Y

opacity 14

distance from midplane 0.6

t = 0.0425

Here we see the same view at the same time from a run at half the grid resolu-tion. The agree-ment is very close.

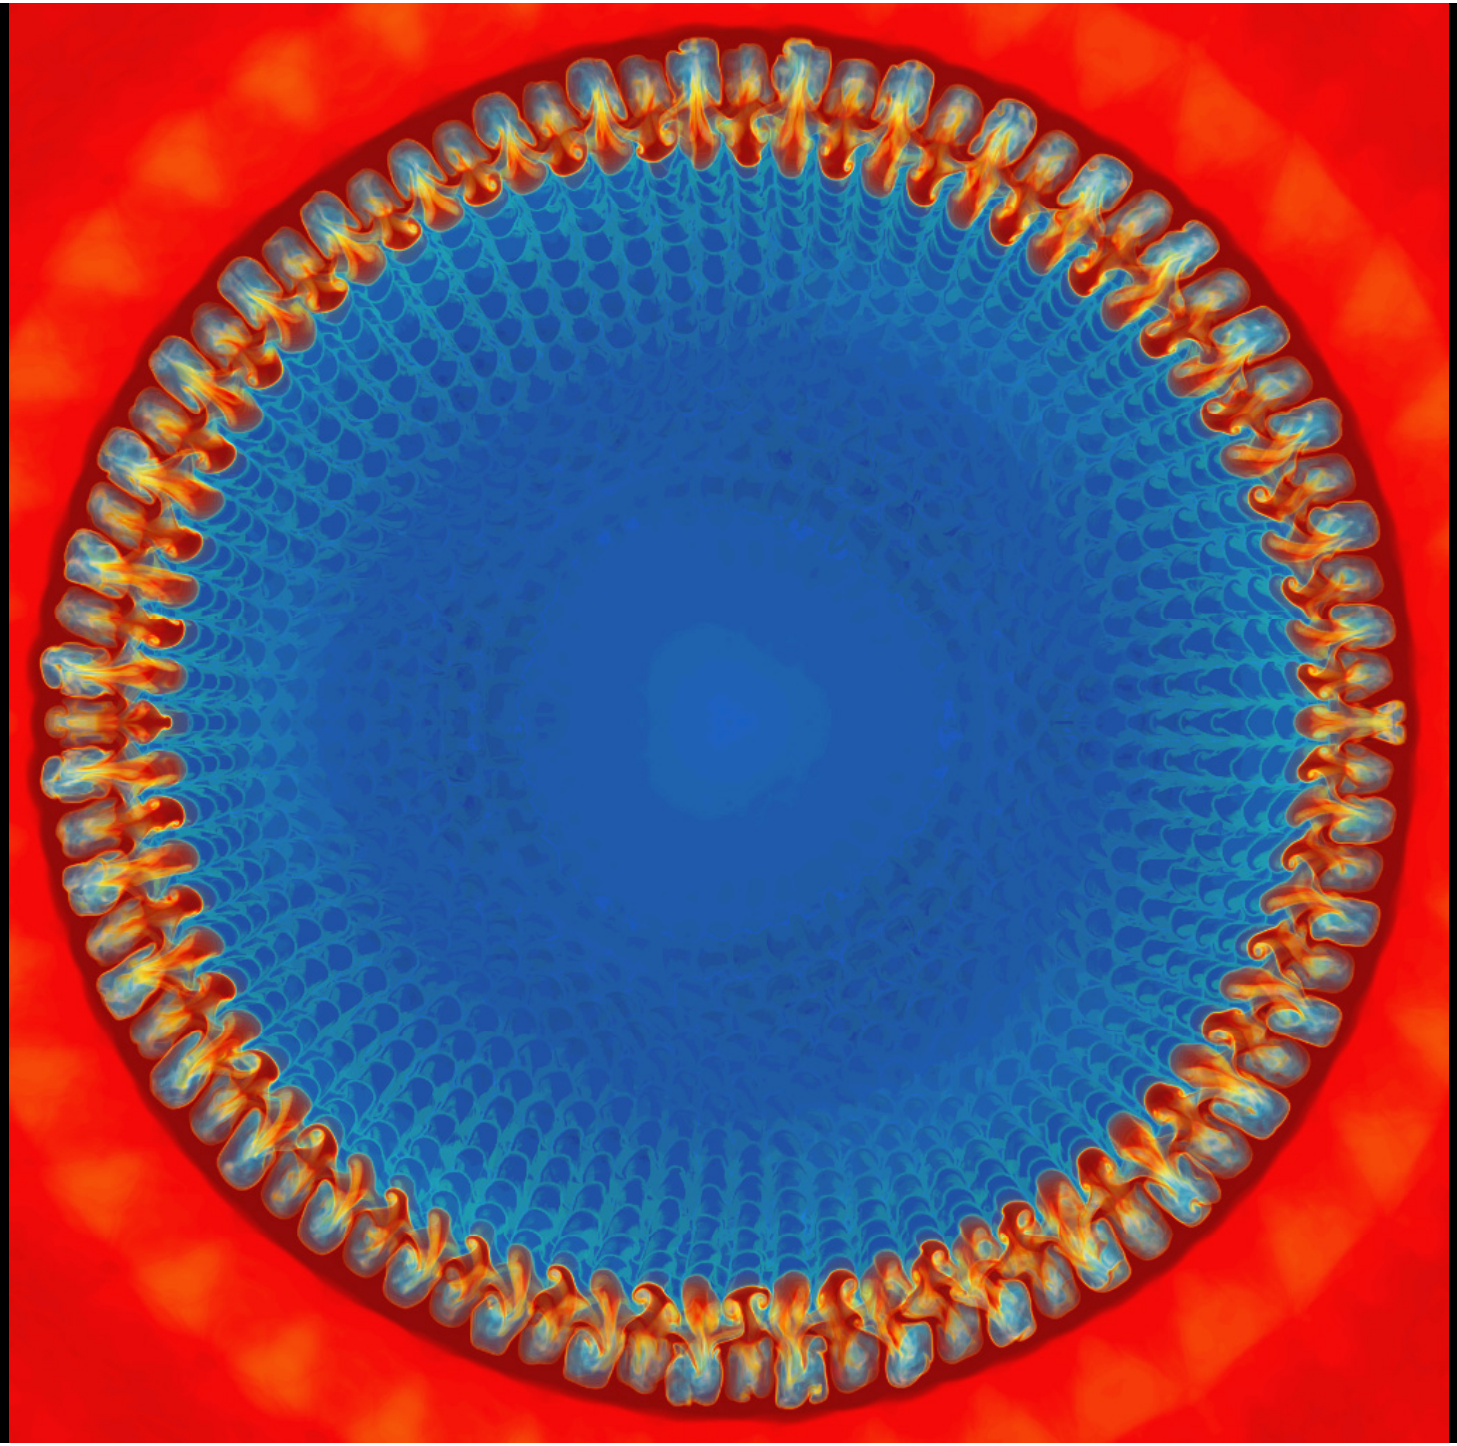

ρ

2048^3 grid

$3^{rd} - 4^{th}$ 8^{ths} in Y

opacity 7

distance from midplane 0.6

t = 0.0450

By now, many small wave-lengh modes have been intro-duced and the mixing region is be-coming chaotic. Note promi-nent mode 3.

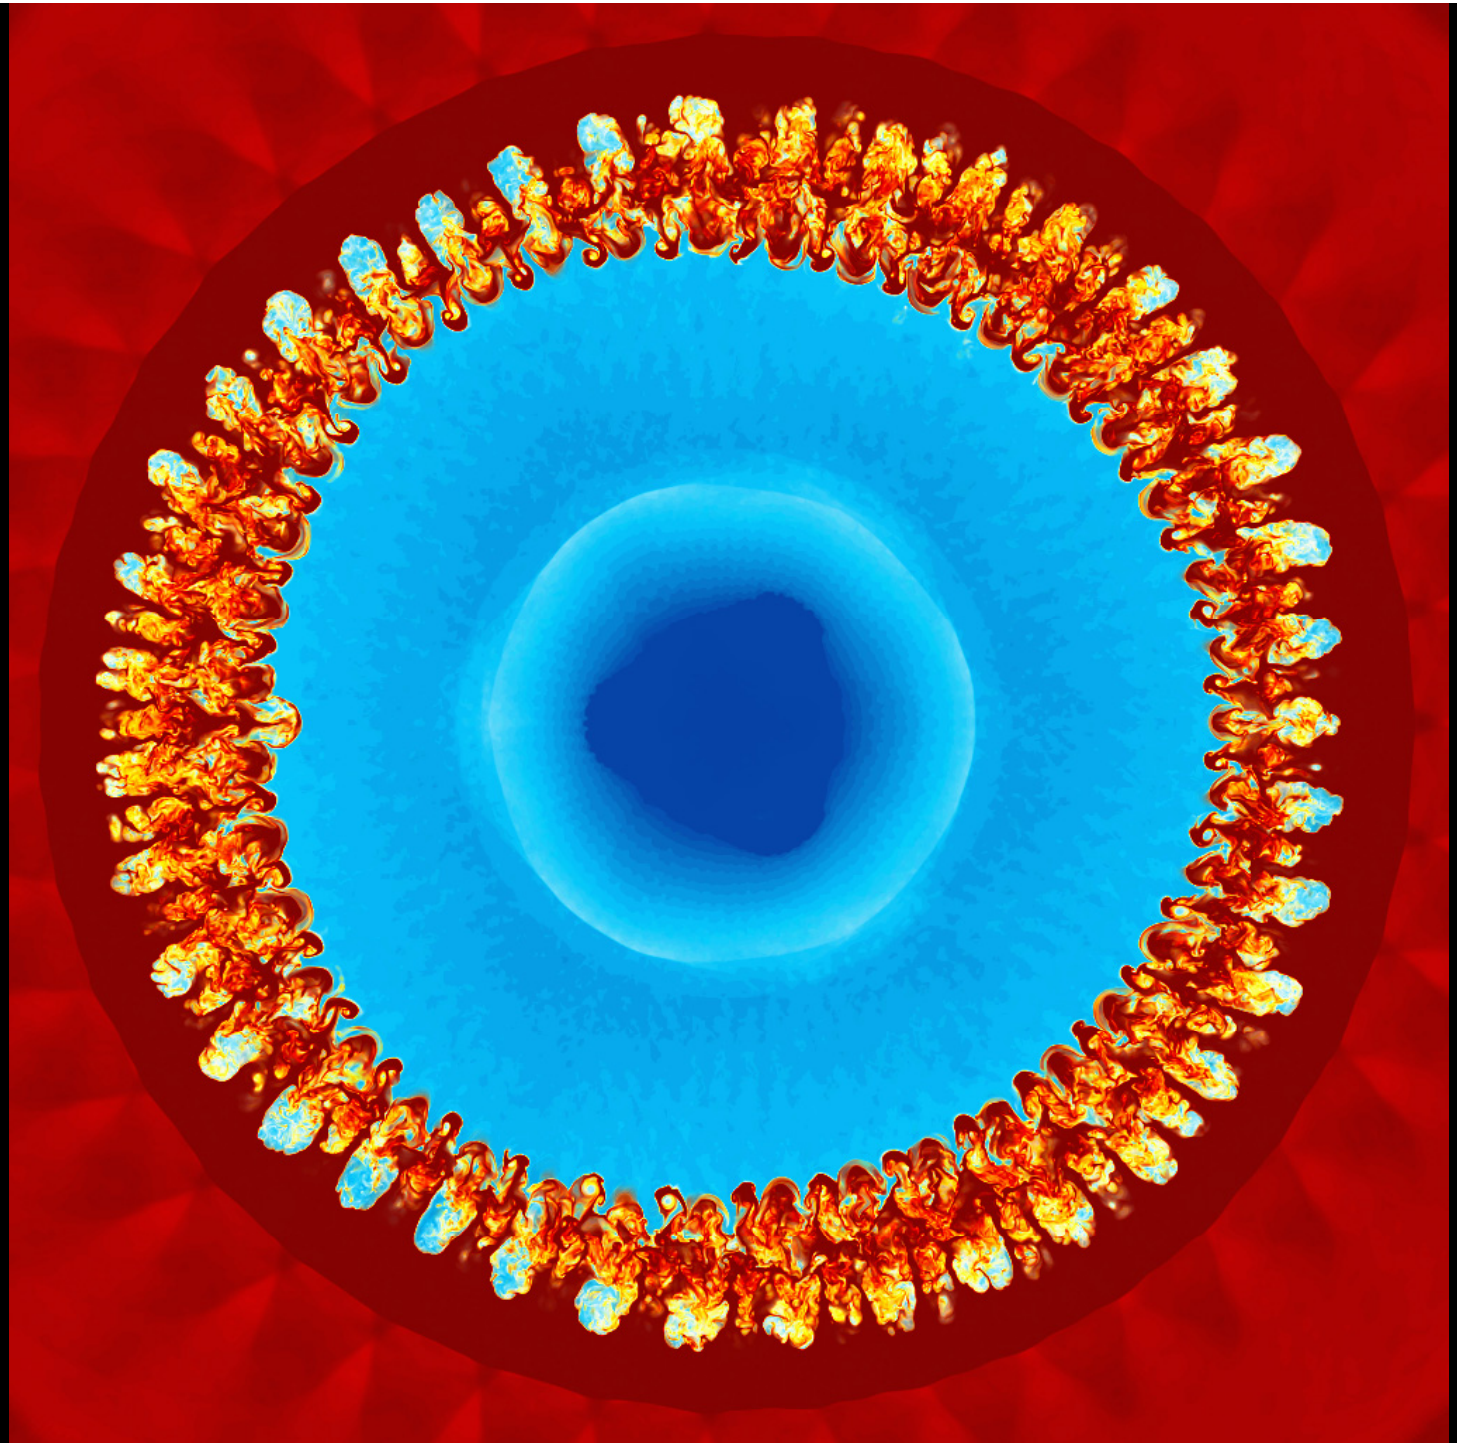

ρ

4160^3 grid

$7^{th} - 13^{th}$ 26^{ths} in Y

opacity 56

distance from midplane 0.6

t = 0.0450

This is the same view at the same time in the run at half the grid resolu-tion. The agree-ment is very good.

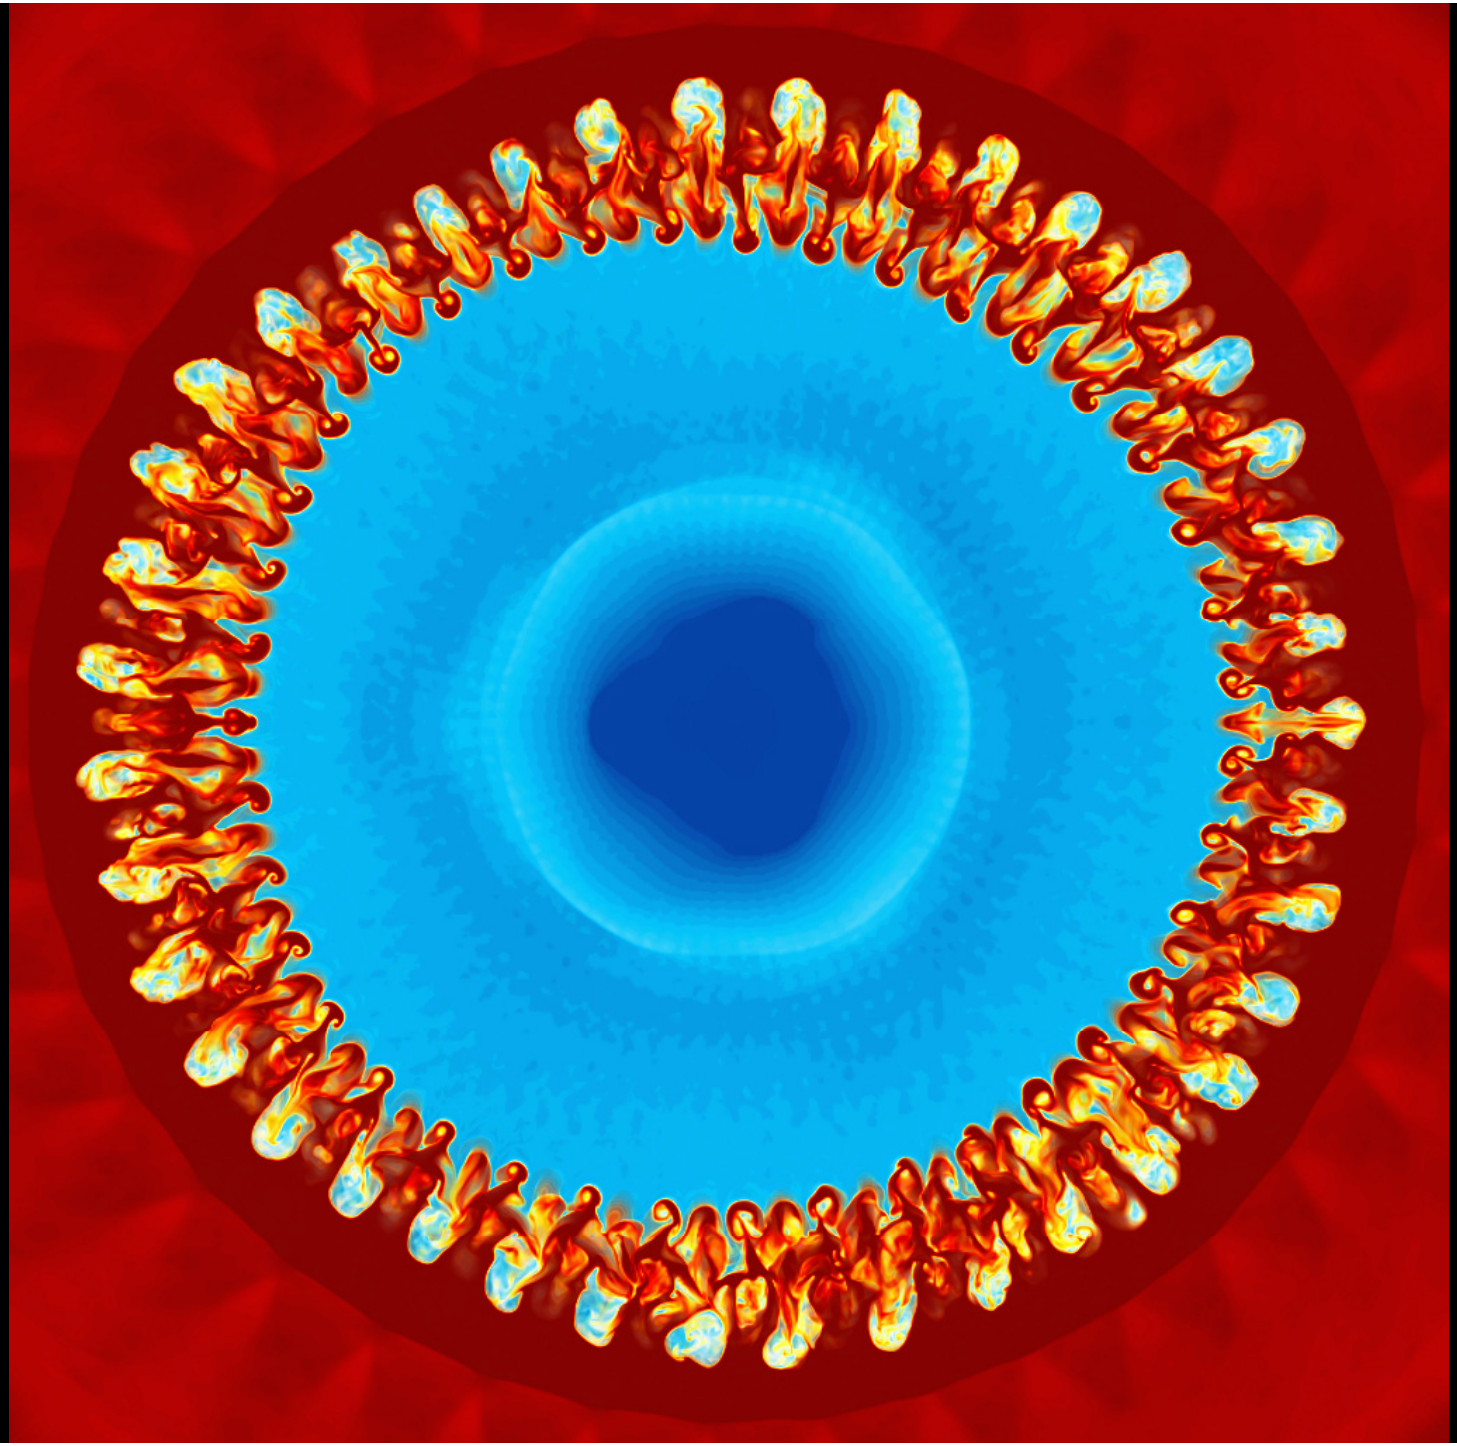

ρ

2048^3 grid

$3^{rd} - 4^{th}$ 8^{ths} in Y

opacity 28

distance from midplane 0.6

t =
0.0475

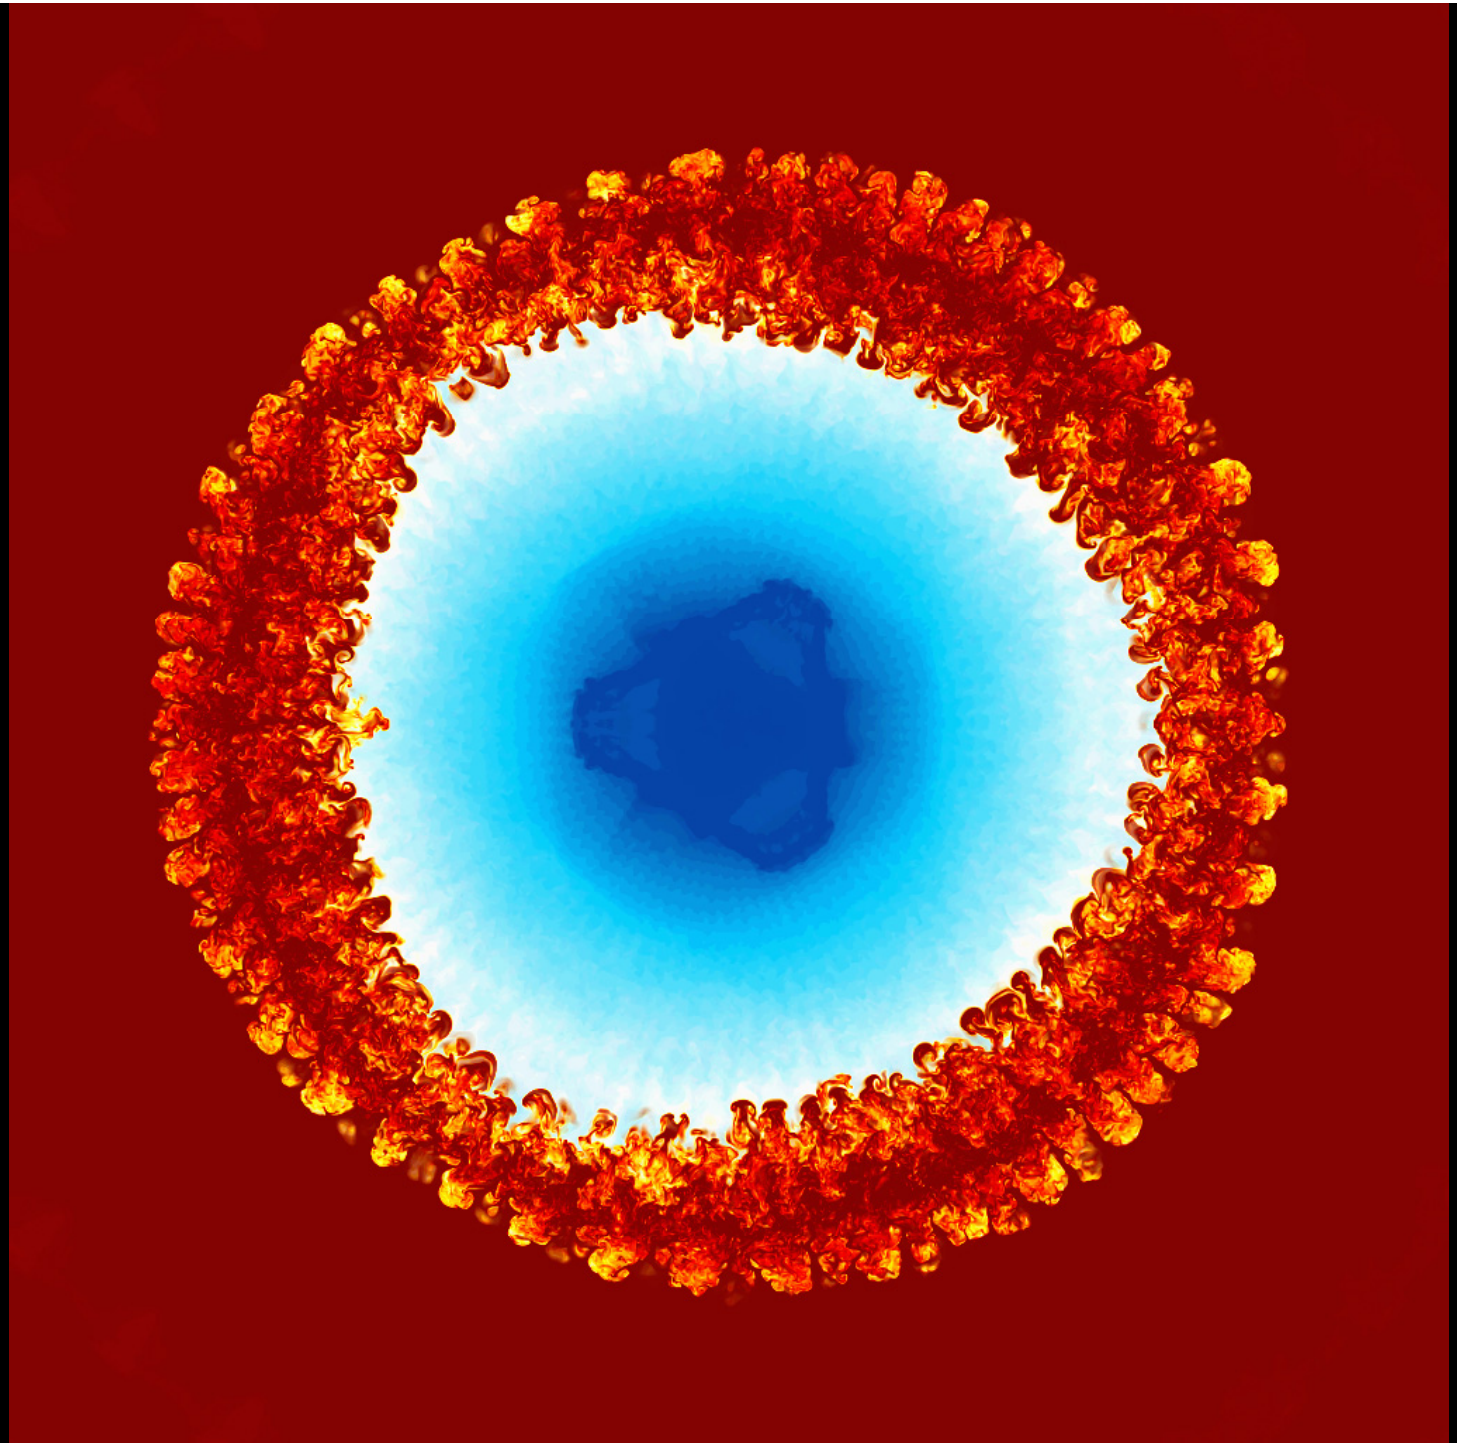

ρ

4160^3
grid

$7^{th} - 13^{th}$
26^{ths}
in Y

opacity
56

distance
from
midplane
0.6

t = 0.0500

At this very late time, we see the expected mode 3, but essentially no grid-induced modes 4 or 8. Only an insignificant jet pointing rightward is suspect.

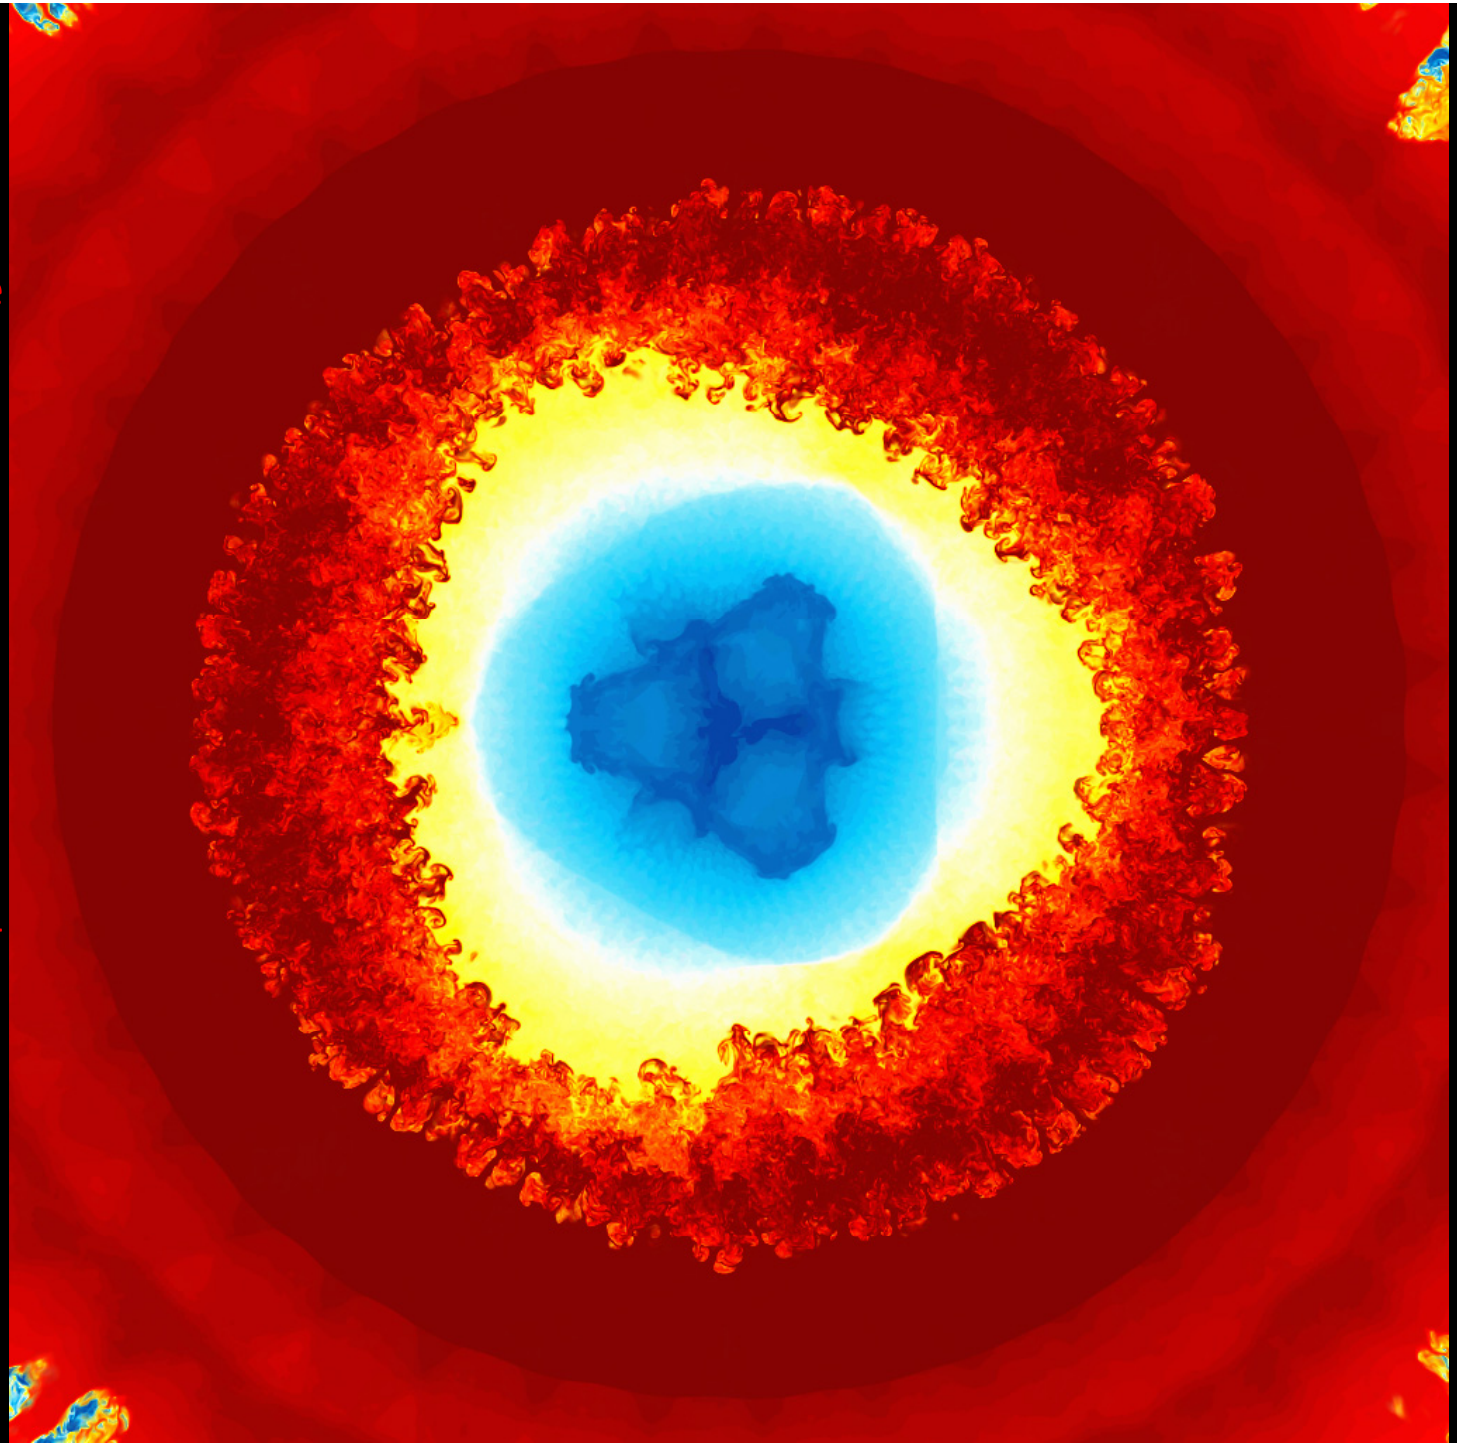

ρ

4160^3 grid

$7^{th} - 13^{th}$ 26^{ths} in Y

opacity 56

distance from midplane 0.6

t = 0.0500

Here the larger view includes the unstable outer surface. Grid effects are minor at this high resolution, made possible at hundreds of Tflop/s.

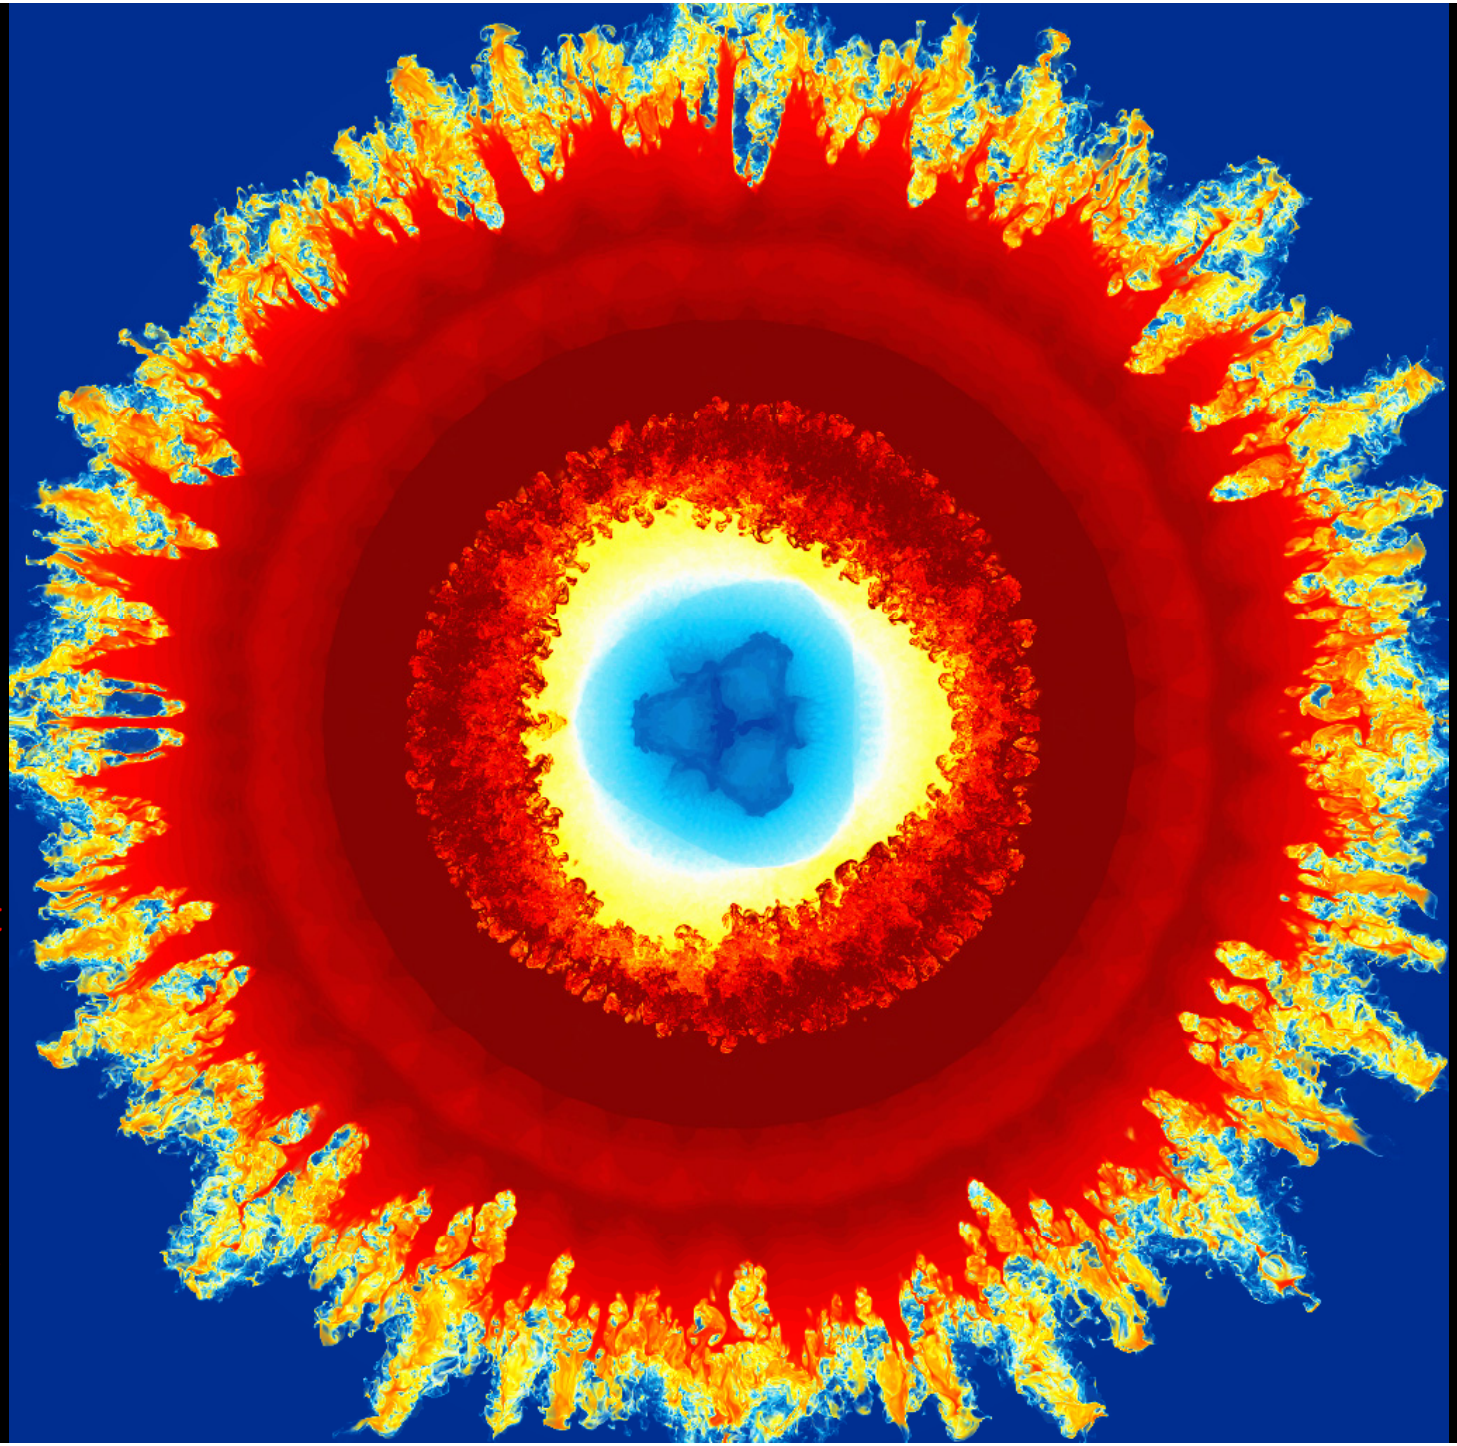

ρ

4160^3 grid

$7^{th} - 13^{th}$ 26^{ths} in Y

opacity 56

distance from midplane 1.0

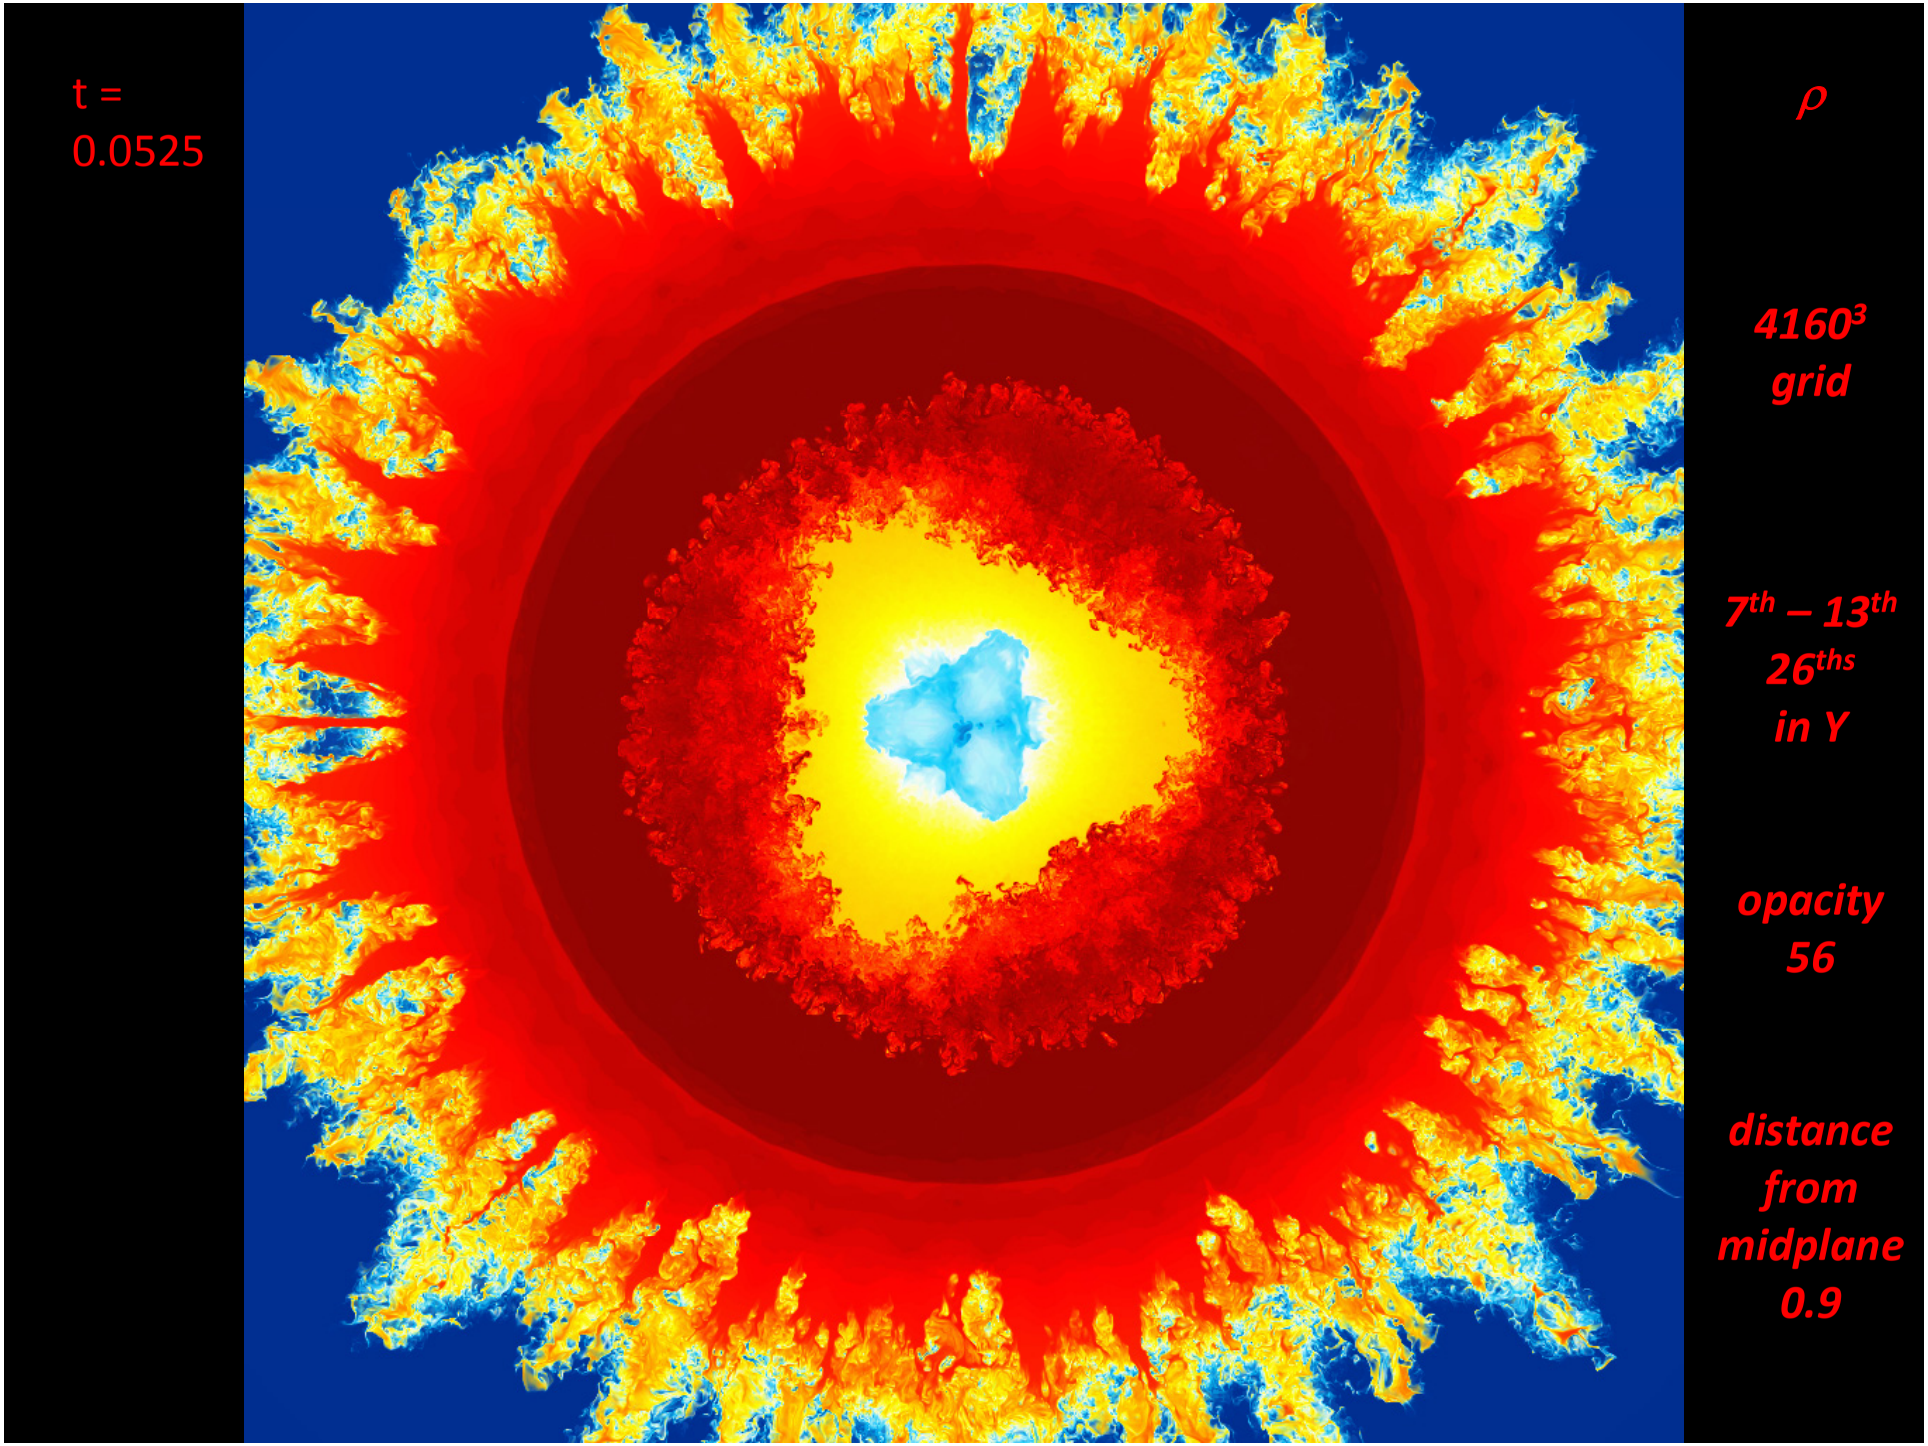
t =
0.0525
ρ
4160³
grid
7th – 13th
26ths
in Y
opacity
56
distance
from
midplane
0.9

t = 0.0525

In this half-resolution run, the 3-fold symmetry of the central hot spot is slightly distorted. Otherwise there is little to dislike.

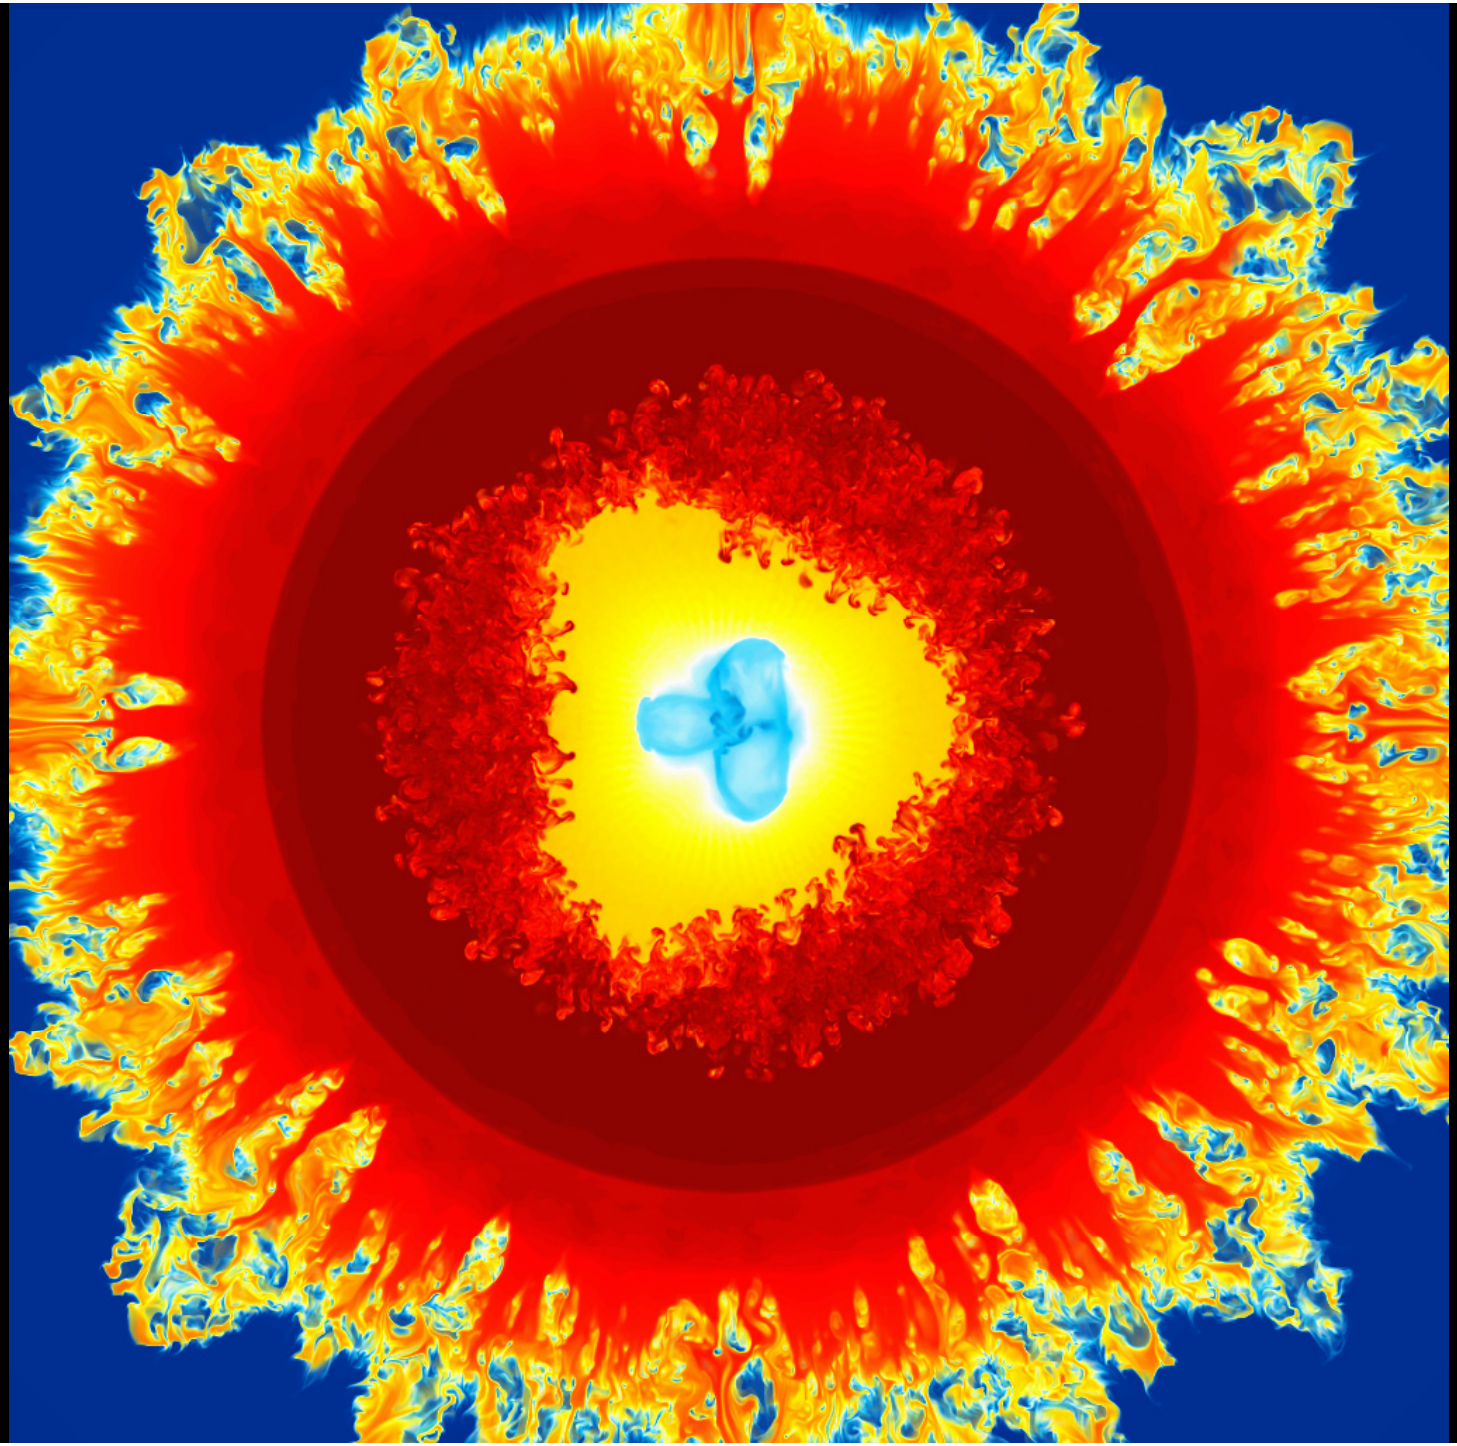

ρ

2048^3 grid

$3^{rd} - 4^{th}$ 8^{ths} in Y

opacity 28

distance from midplane 0.9

t = 0.0525

We have about a billion cells in the mixing region, enough for a good statist-ical descrip-tion to be extract-ed.

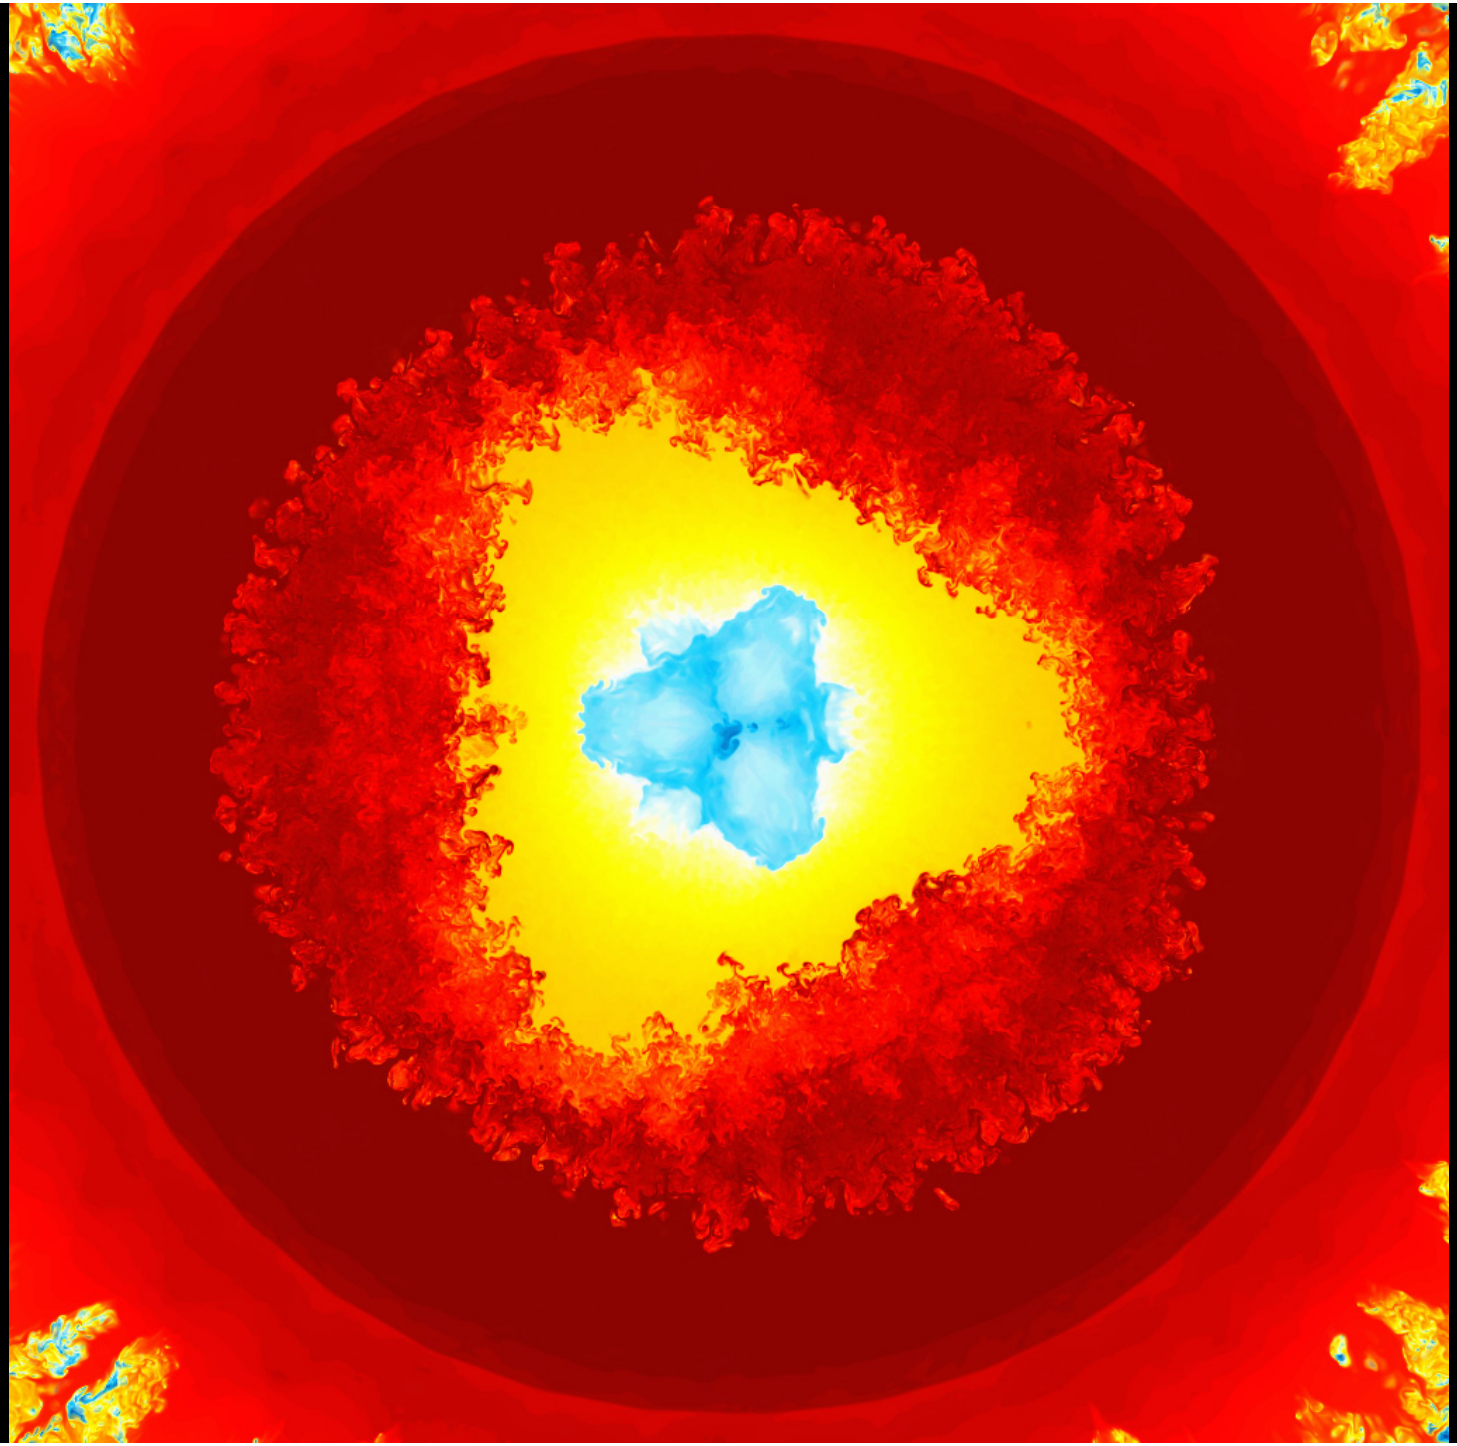

ρ

4160^3 grid

7^{th} – 13^{th} 26^{ths} in Y

opacity 56

distance from midplane 0.6

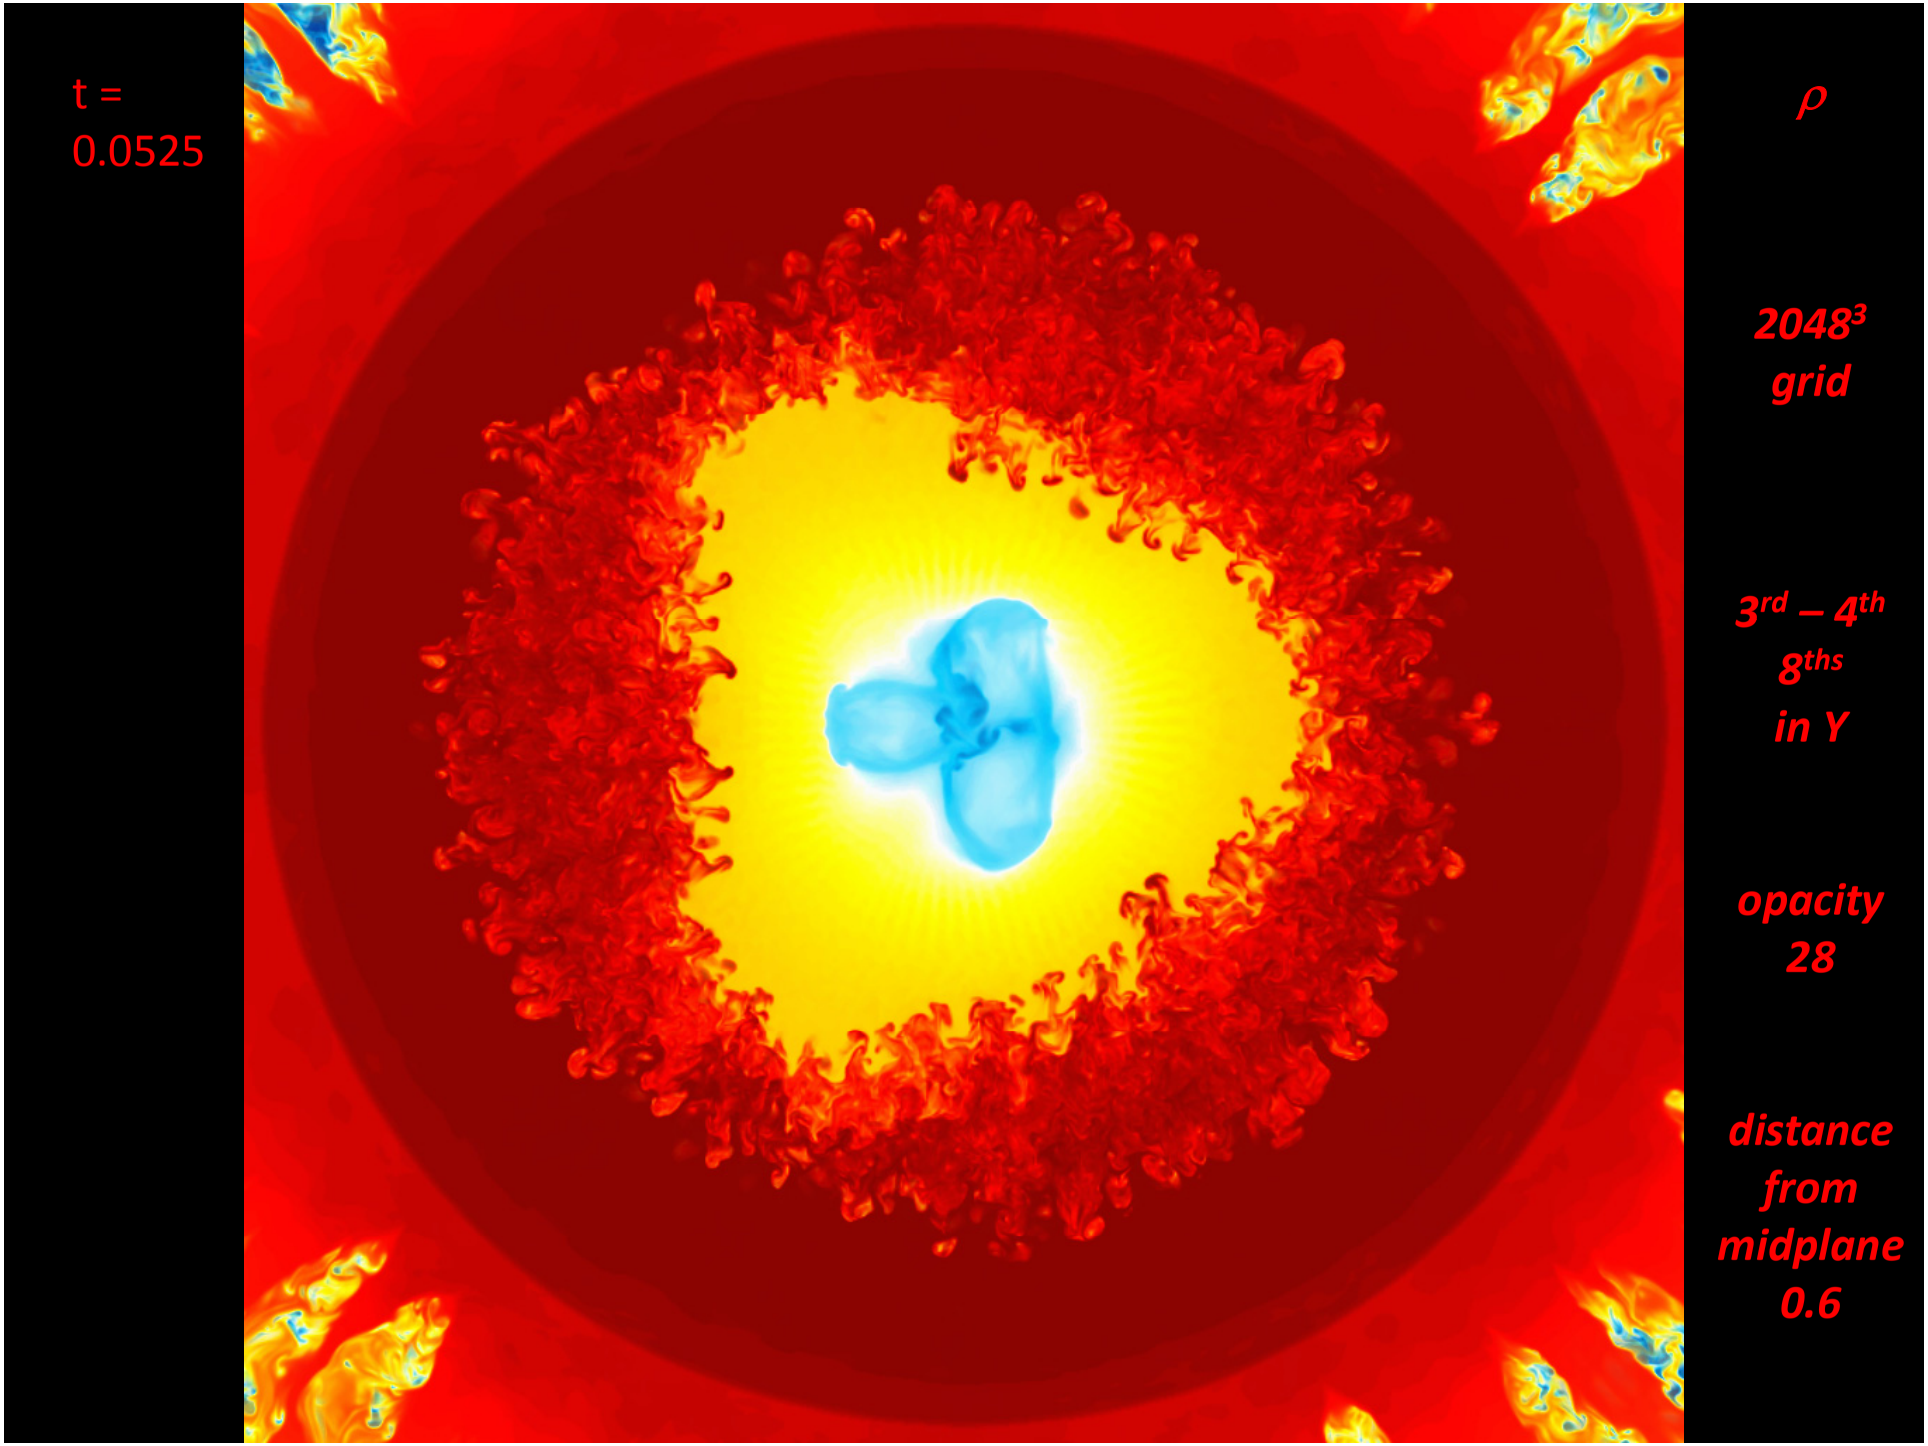
t =
0.0525
ρ
2048³
grid
3rd – 4th
8ths
in Y
opacity
28
distance
from
midplane
0.6

72 billion cell problem computed in only 8 hours on the Early Science Access Blue Waters machine at NCSA.

- It is the speed of the code that makes this problem practical.
- It is the accuracy of the numerical scheme that makes the solution worthwhile.
- Mixing region at time 0.0525 between radii of 0.0717 & 0.2326, that is, between radii of 408 and 661 grid cell widths.
- This is 2.5% of the problem domain volume, 1 billion cells.
- On a grid of 8320^3 cells, we would get 8 billion cells in the mixing region, and this would be sufficient for statistics of subgrid-scale models of compressible, turbulent mixing.
- This would take an 8-times larger machine (BW) 16 hours.

Accommodations for performance.

- Briquette data structure, with briquettes of 4^3 cells.
- Use of magic numbers of threads and briquettes.
- Reduced interpolation difference stencil from 7 to 5 cells.
- Strong shock diffusion no longer a separate step.
- Computation of strong shock diffusion only once per Δt.
- Special I/O processes, 8 bricks per MPI process, etc.

The solution converges, but how do we know it is right?

- We are limited of course by the assumed gamma-law equation of state, the assumed vanishing Navier-Stokes viscosity and heat conduction, the artificial boundary conditions, etc. This should go without saying.
- Otherwise, if we have no bugs in our code with any impact on the results, then convergence means the answer is right.
- We address the question of bugs by testing the code on standard problems, but, very significantly, we perform comparisons with other codes, which surely have *other* bugs.
- In 2011, we carried out a detailed comparison for 2-D test problems of this type with the following other codes: RAGE, Castro, and FLASH.
- The results (see next slides) are very encouraging.
- We gave the codes a chance to agree by using 4096^2 grids, polar coordinates for FLASH, and relatively large initial perturbation amplitudes.

Code similarities:

- PPM or MUSCL, in one form or another, was in every code.
- All but FLASH used Cartesian grids.
- All codes tracked the multifluid volume fraction without any detailed interface reconstruction (no VOF).
- All results were displayed in the same way.

Code differences:

- Different people wrote each code, so the chance that they have the same bugs is essentially zero.
- Different people ran each code on the problems, so the results could not be fudged to look alike.
- FLASH used a polar coordinate option, which is better adapted to this problem in 2D than a Cartesian grid.
- RAGE and Castro had AMR, implemented in different ways.
- PPM ran with 32-bit precision, while all other codes used 64-bit precision.
- PPM used PPB advection, with subcell resolution.
- Castro used piecewise-linear advection, and RAGE used this plus contact discontinuity detection and steepening.

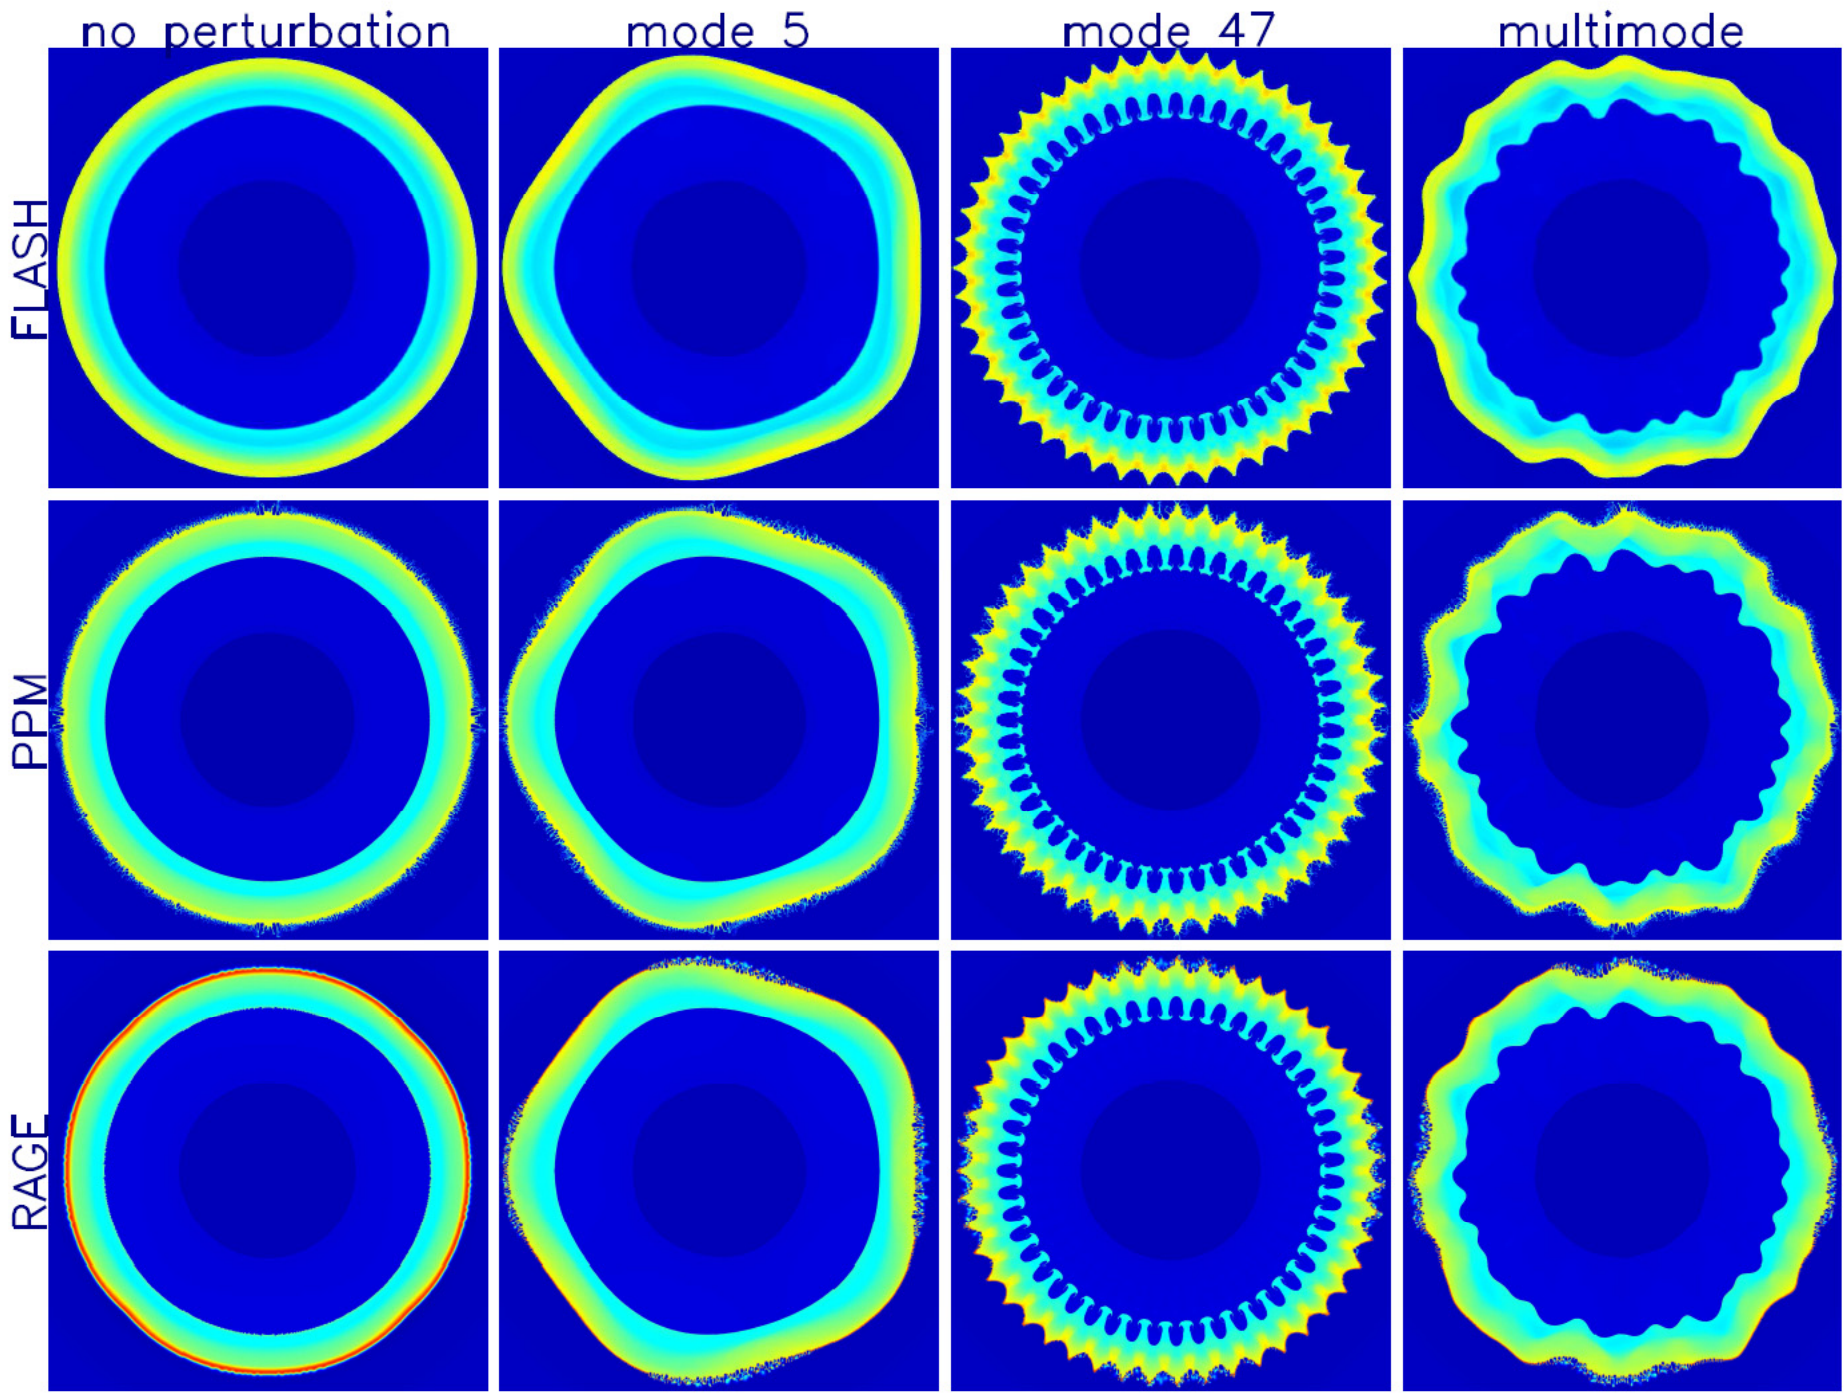
no perturbation
mode 5
mode 47
multimode
FLASH
PPM
RAGE

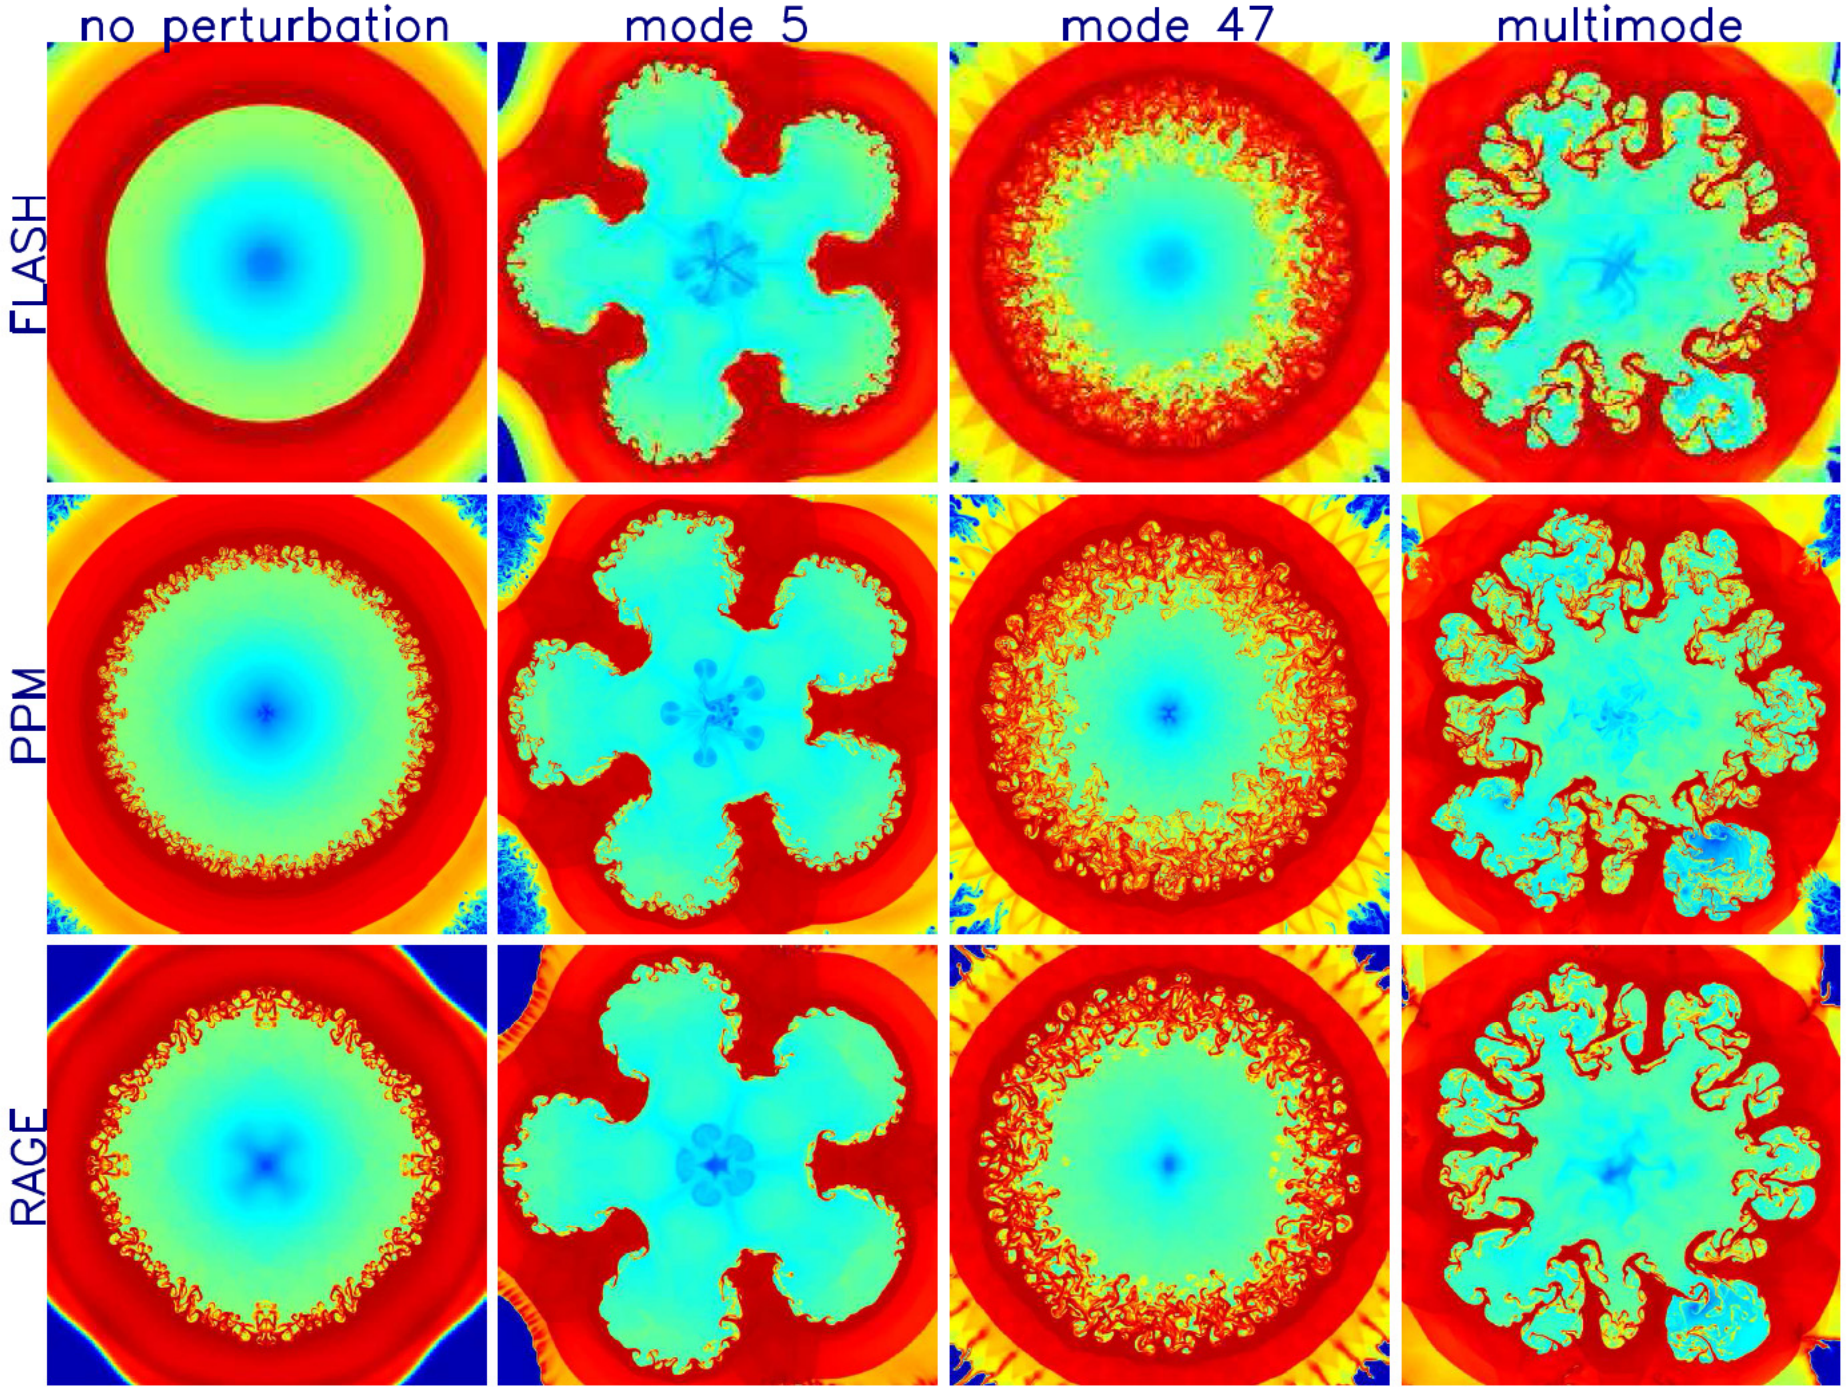
no perturbation
mode 5
mode 47
multimode
FLASH
PPM
RAGE

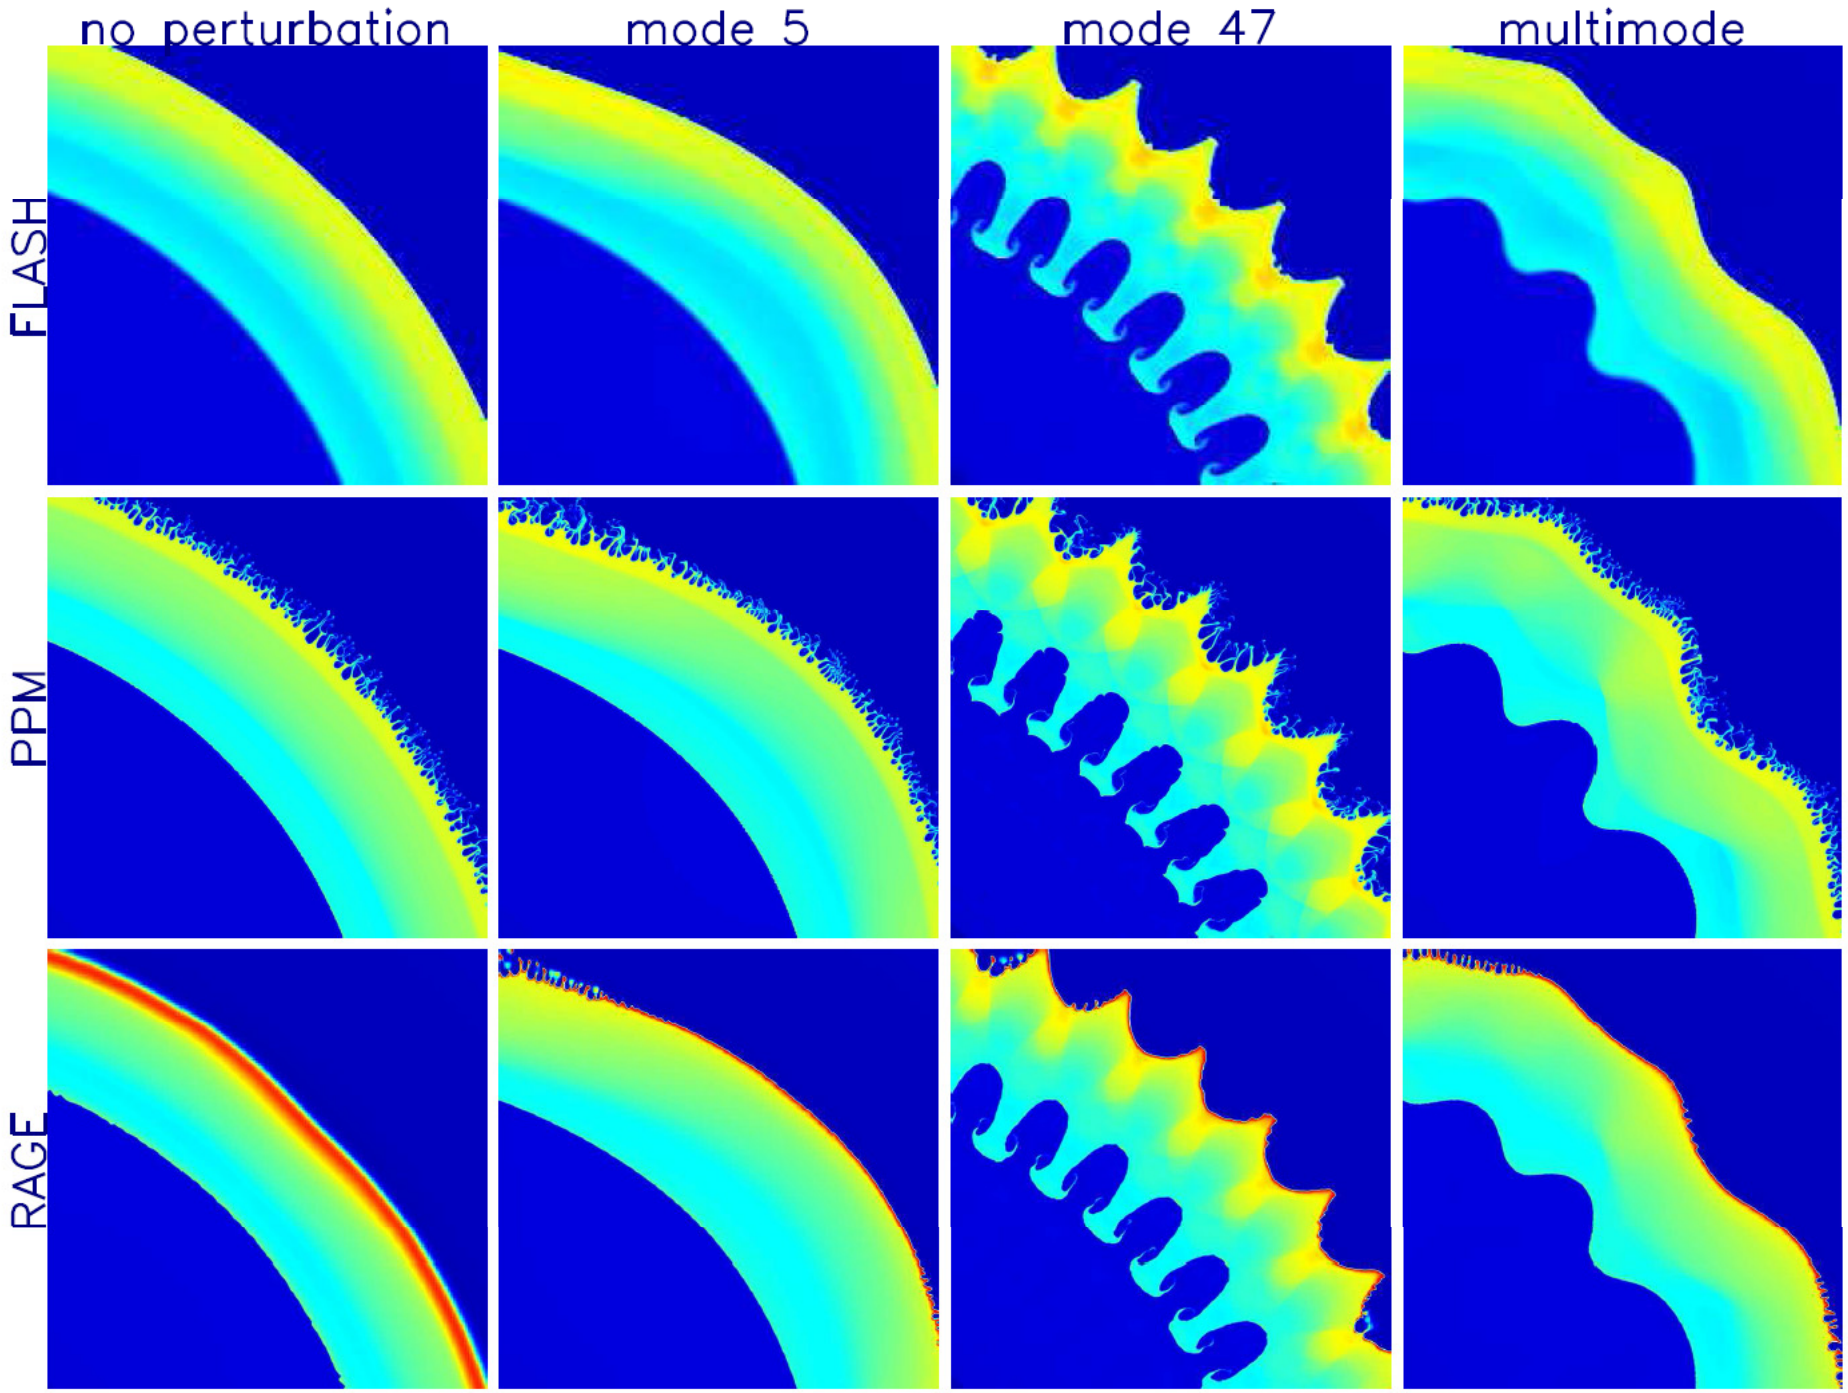
no perturbation
mode 5
mode 47
multimode
FLASH
PPM
RAGE

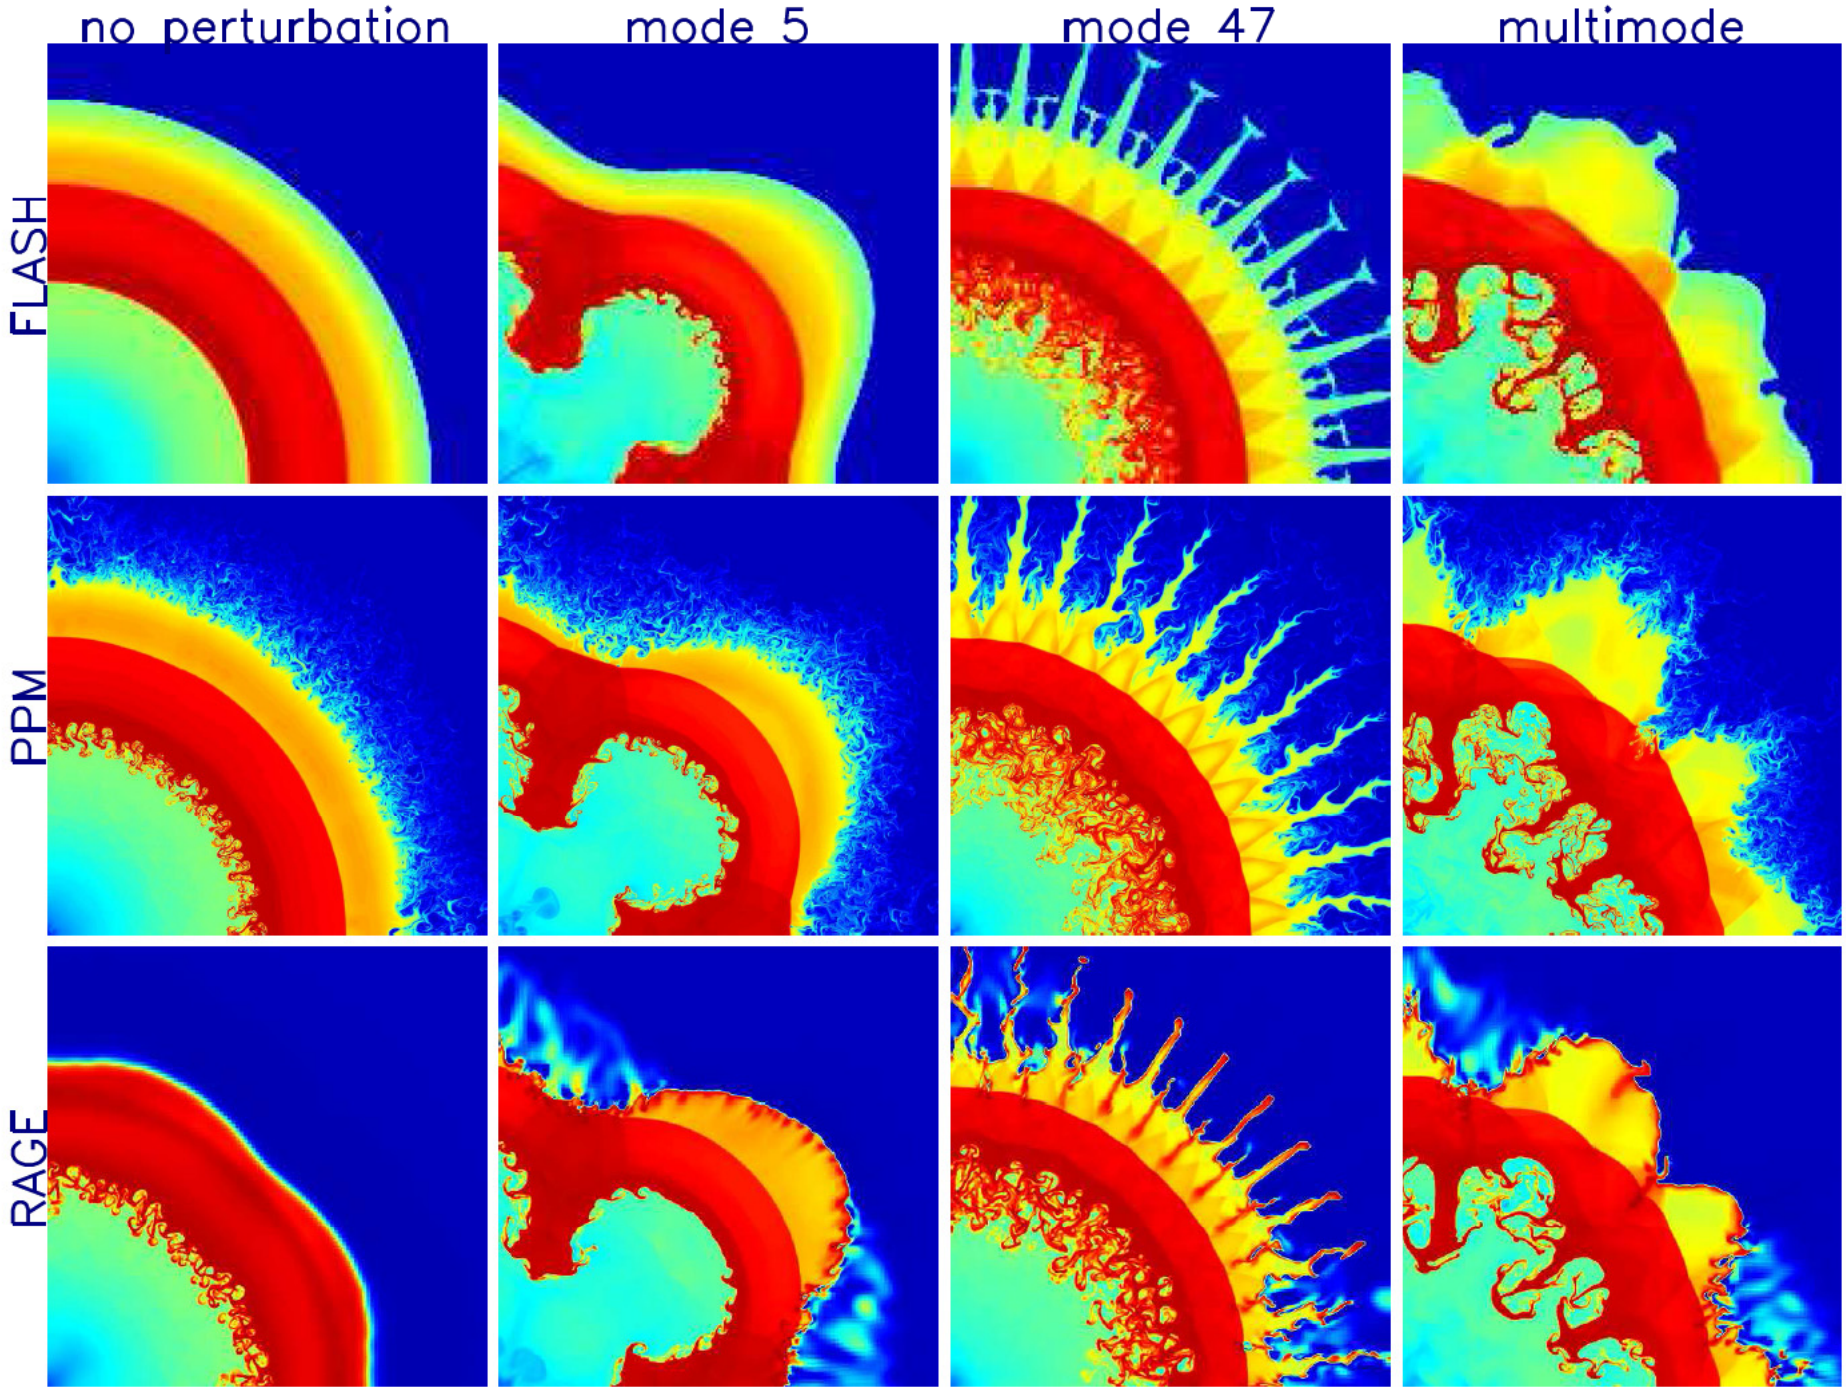
no perturbation
mode 5
mode 47
multimode
FLASH
PPM
RAGE

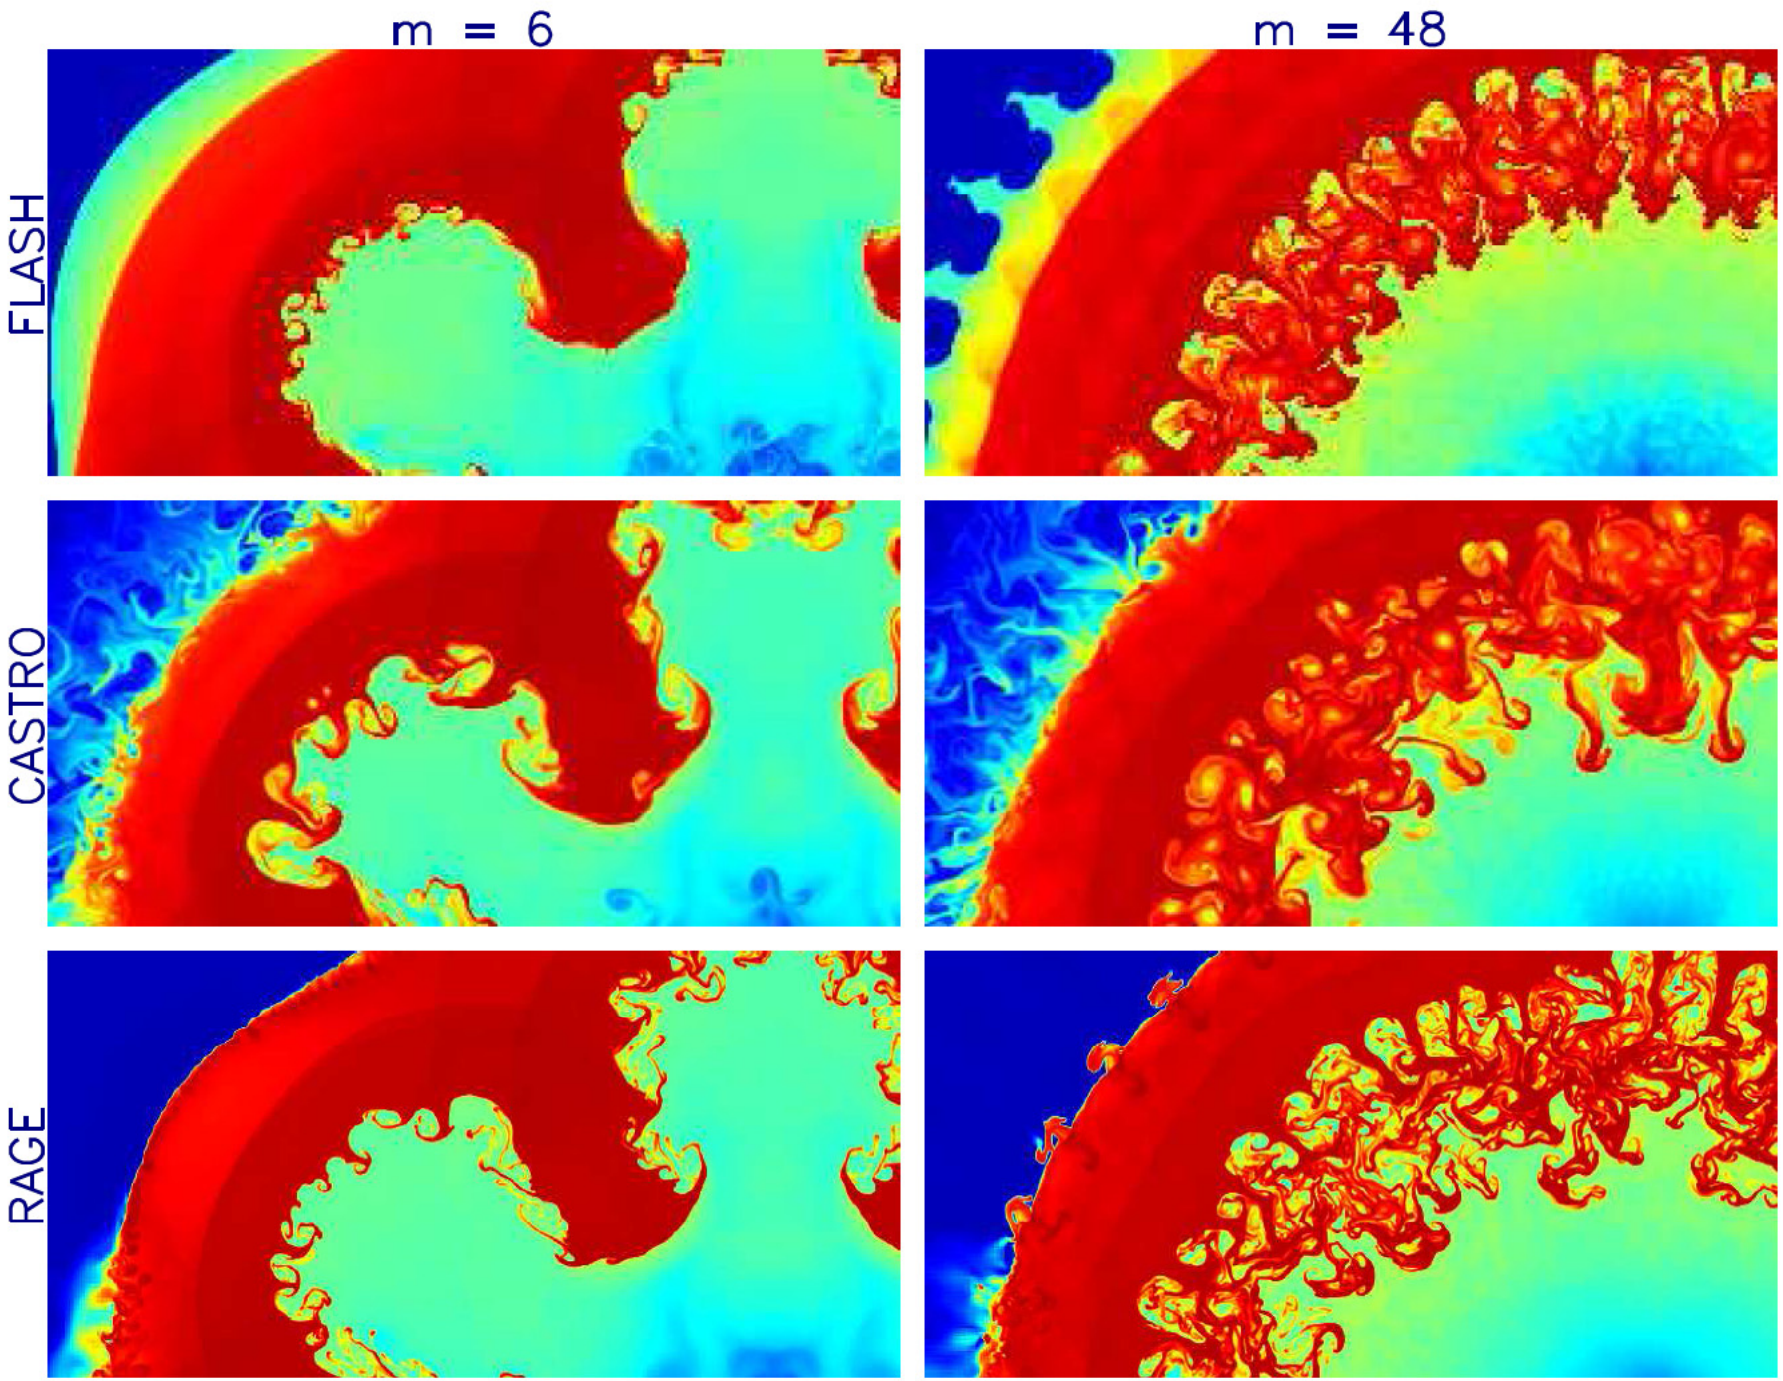
m = 6
m = 48
FLASH
CASTRO
RAGE

Remarks:

- Unstable interfaces display unstable behavior in all codes.
- This behavior depends upon the perturbation spectrum.
- The perturbations introduced through the grid seem all to come from the same cause.
- As a thin structure runs through a regular grid, there are beat frequencies set up with fundamental wavelengths equal to the grid plane crossing distance.
- Short wavelength disturbances are strongly damped, but wavelengths of 6 to 12 cells cannot be damped without significantly altering the solution in other parts of the flow. These are consequently amplified by the interface instabilities.
- Difficult to “solve” this problem without destroying the overall solution.
- Conclusion: embrace instability. It is real. No conceivable experiment in the real world could possibly contain only two modes. A computation doing that could not correspond to reality as we all know it.

Remarks Continued:

- One member of our team was convinced that the polar coordinate calculations with FLASH would blow away those of the other codes.
- This did not happen (you judge for yourself).
- The result of a physical problem cannot depend upon the coordinate system used to describe it. All codes in whatever coordinates, if they converge at all, must converge to the same answer in an appropriate statistical sense.
  Only the cost of the answer can differ, but not the answer.
- There are beat frequencies in every coordinate system, and hence grid-initiated unstable modes will always enter.
- What all such modes have in common is that they have high frequencies.
- Their appearance is suppressed if there is a real mode for the calculation to follow with comparable λ and a larger amplitude.
- The calculations seem to show that in the presence of low modes, high modes tend to be dynamically unimportant.

More Remarks:

- The 3-D ICF problem seems to be "easier" than the 2-D one.
- The inner surface of the shell, apart from a Richtmyer-Meshkov instability that is hard to get wrong, is R-T stable for the great bulk of the computation.
- When the inner surface does become R-T unstable, it is already strongly perturbed by the R-M modes.
- This means that the code has a strong signal to work with, and hence grid-induced modes are of very minor consequence at the inner dense shell surface.
- The R-M and R-T instabilities induce K-H instabilities naturally, and these bring about a broad spectrum of high-frequency modes leading to chaotic flow.
- Once these instabilities get going, it is hard to imagine any grid that would be superior to any other, except in its computational cost and programming convenience.
- This seems to be why the Cartesian grid works for this problem.

***If you have a Cartesian mesh code, how can you modify it to make it do ICF problems at the least cost*?**

- First, you could revise your data structures to build on grid briquettes, and you could process them in pipelined grid pencil updates, as we do.
- We are developing code translation tools to assist with this code restructuring process.
- The way our code works, at the level of a single thread (and it has run very efficiently on Blue Waters with over 730,000 threads on over 24,000 nodes), is illustrated by the diagram on the following slide, which is in turn followed on the next slide by a long-winded figure caption.

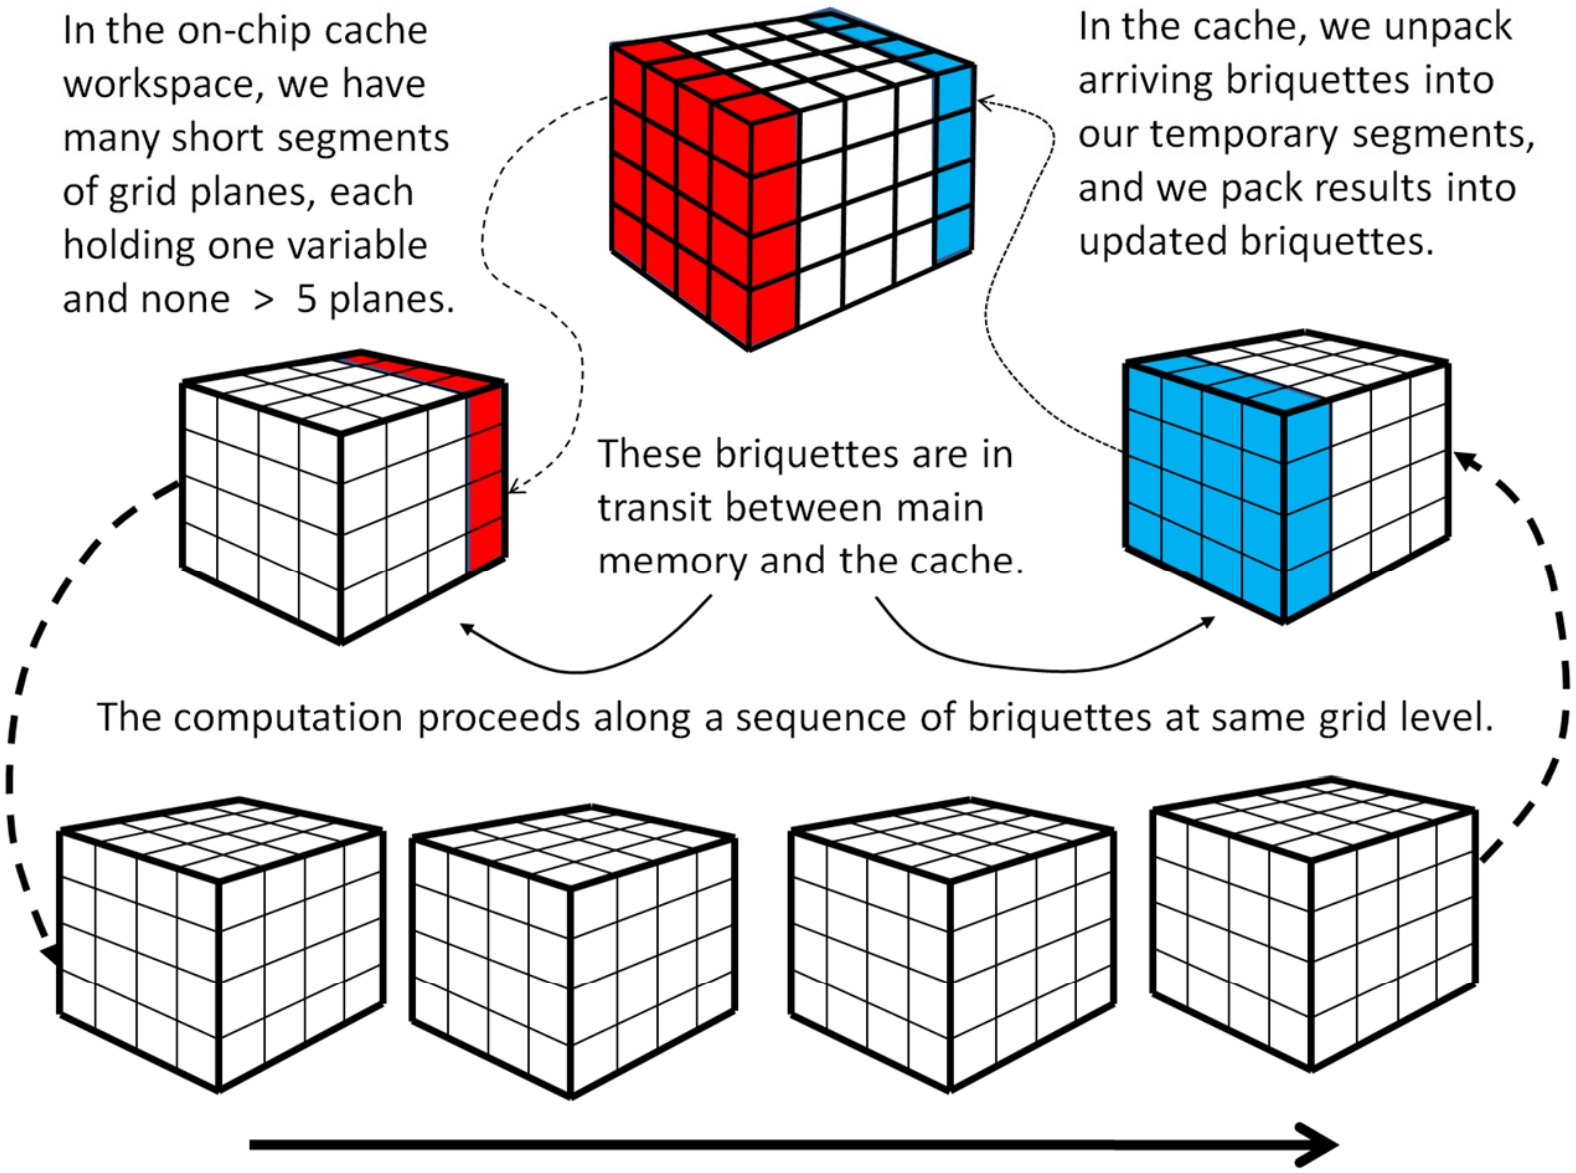

A proper caption for this figure is on the next slide

Caption for previous slide: An illustration of the manner in which our computation proceeds in a single 1-D pass of the numerical algorithm. The physical variables giving the fluid state – the density, pressure, 3 velocity components, and 10 moments of the tracked fluid fractional volume – are packed into contiguous records in main memory. Each record gives the values of these 15 variables in a grid briquette of 4^3 cells. Four such briquettes are represented at the bottom of this diagram, with an arrow indicating the direction of the 1-D pass. The briquette records are processed in a sequence indicated by this arrow. The dashed arrows indicate data movement between the CPU chip's cache and the main memory (at the bottom of the diagram). Briquette records are fetched one briquette ahead of the processing, so that this data transfer can occur simultaneously with computation. A briquette that has just arrived in the on-chip cache is shown at the middle right in the diagram. A shaded grid plane of this briquette is shown which is extracted from this record and placed into the appropriate element of a revolving buffer, just one of a great many such buffer structures in the on-chip cache. No such buffer need be larger than 5 grid planes, as shown here at the top center of the diagram. Most revolving buffers are shorter than this, so that despite needing over a hundred of them to hold all the intermediate results for the complete multifluid hydrodynamic update process, only about 60 KB of space is required for these buffers in the on-chip cache. All the routines of our hydrodynamic update computation are inlined to give a single, humongous subroutine, which has the revolving buffers as its stack memory, so that this data will be cache resident. This in-lining and the indexing of the revolving buffers is an error-prone process that we are automating with a code translator "front-end." This takes a Fortran-W input expression and automatically pipelines it into a single, enormous outer loop. In each iteration of this outer loop one grid plane (shown here in blue) is unpacked from a prefetched briquette record, and at the end of the loop one grid plane (shown in red) that is fully updated is packed into a new briquette record, shown at the center left in this diagram. When full, that new briquette record is written back to main memory. This pipe-lined code runs on an Intel Nehalem CPU core (2.93 GHZ) at 5.0 Gflop/s (21% of 32-bit peak).

***If you have a Cartesian mesh code, how can you modify it to make it do ICF problems at the least cost*?**

- What about the numerical methods?
- Increasing the initial perturbation amplitude to 4% of a wavelength, and reducing the grid resolution to 608^3 or 1216^3 cells, we show on the following slides the effects of turning off various aspects of our PPM code's numerical methods.
- This is all well below the convergence regime for the mixing layer details, but the overall problem convergence is not so bad (you be the judge).
- We investigate the importance of:
  - ✓ Grid jiggling.
  - ✓ Smart shock dissipation (a minor variation on that originally described for PPM).
  - ✓ PPB multifluid volume fraction advection.
  - ✓ Grid resolution, of course, and hence *speed*.

Dump 10
1216^3 grid

ρ

rho12-LUT
opacity 2

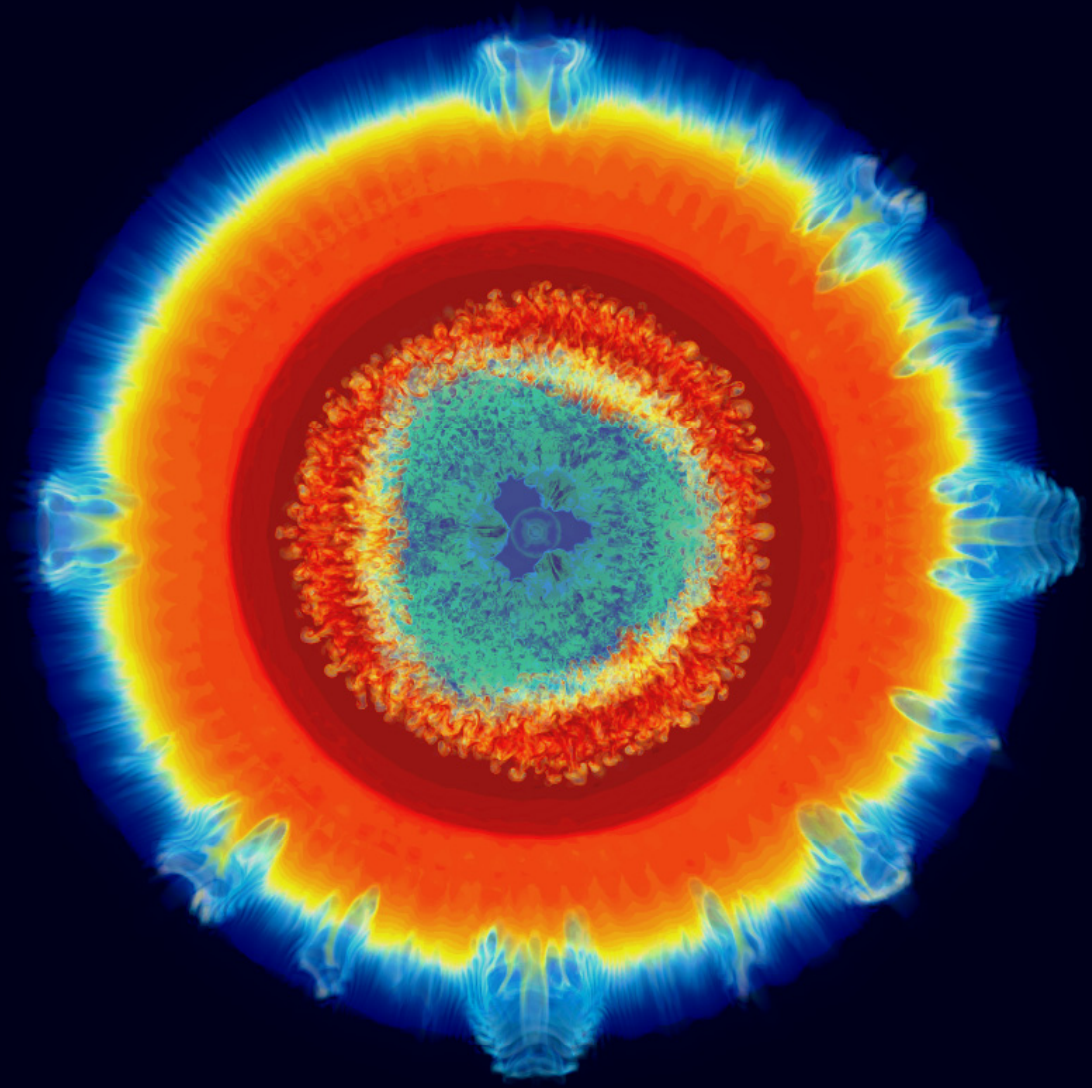

The "fountains" produced by grid-induced modes on the outer surface are more prominent at this grid resolution, but do not affect the inner surface dynamics.

Affect of Inadequate Resolution

Dump 10
608^3 grid

ρ

rho12-LUT
opacity 1

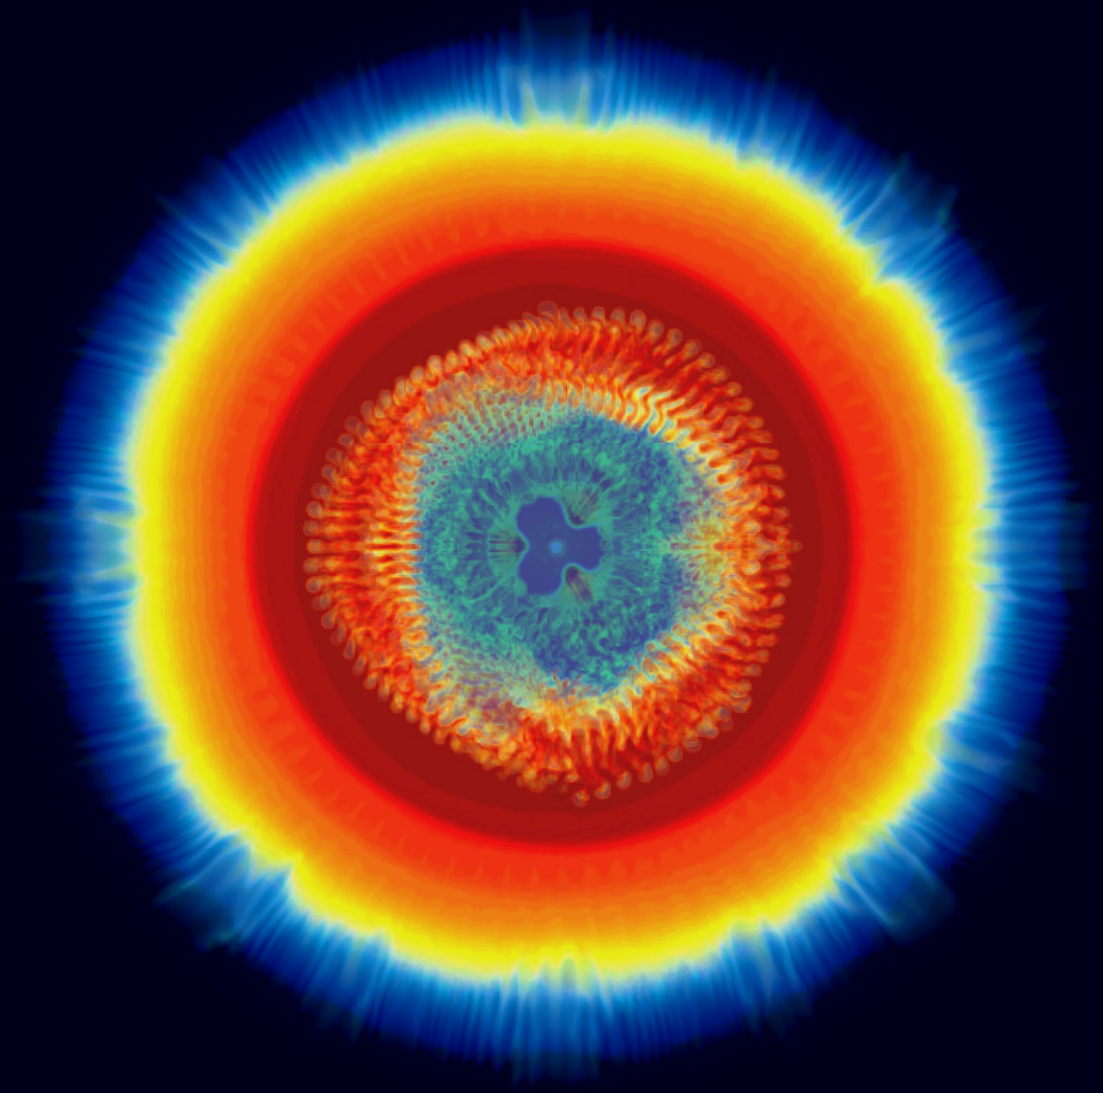

The "fountains" on the outer surface are less annoying here, but inconse-quential regardless. However, the inner shell surface is visibly distorted.

The 3-fold symmetry has been disrupted by a jet resulting from grid symmetry at the midplane.

Dump 10

1216^3 grid

ρ

rho12-LUT opacity 2

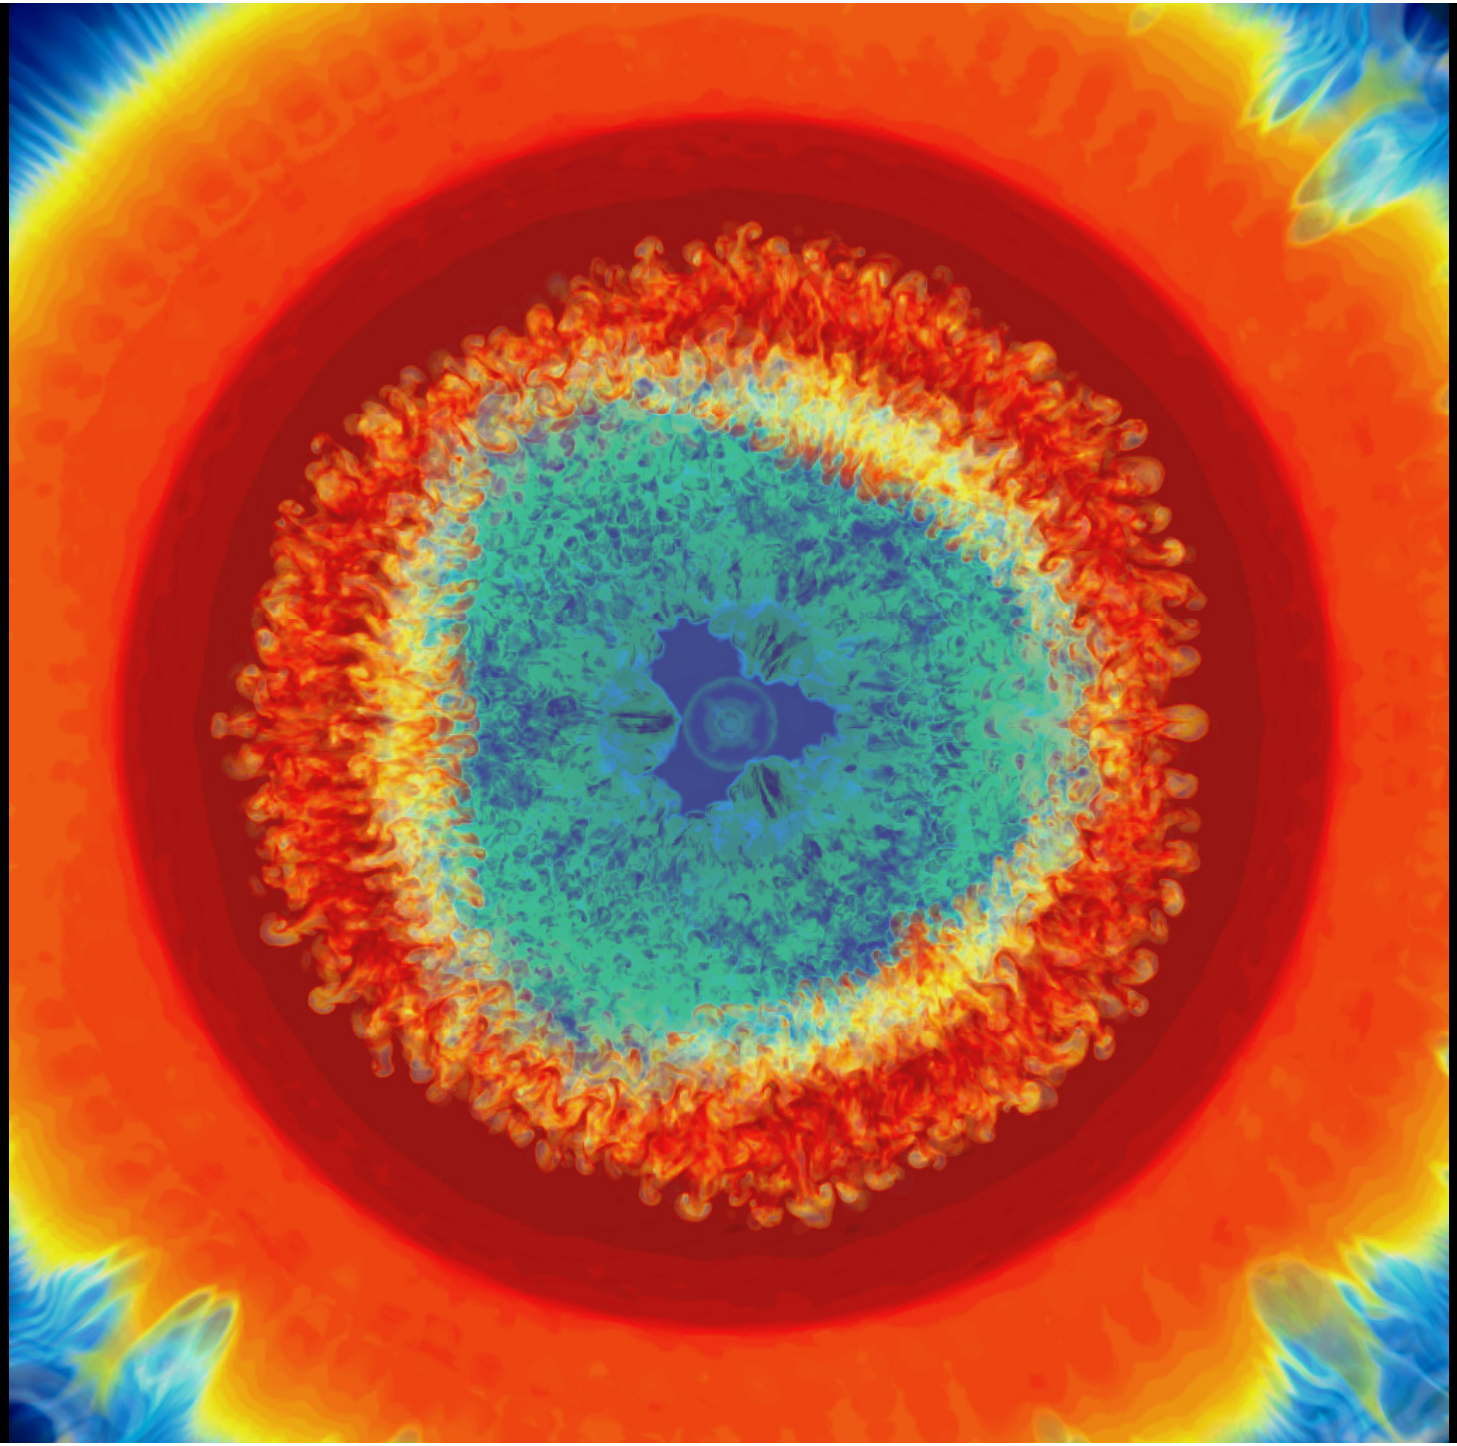

Dump 10
608^3 grid

ρ

rho12-LUT opacity 1

Affect of Inadequate Resolution

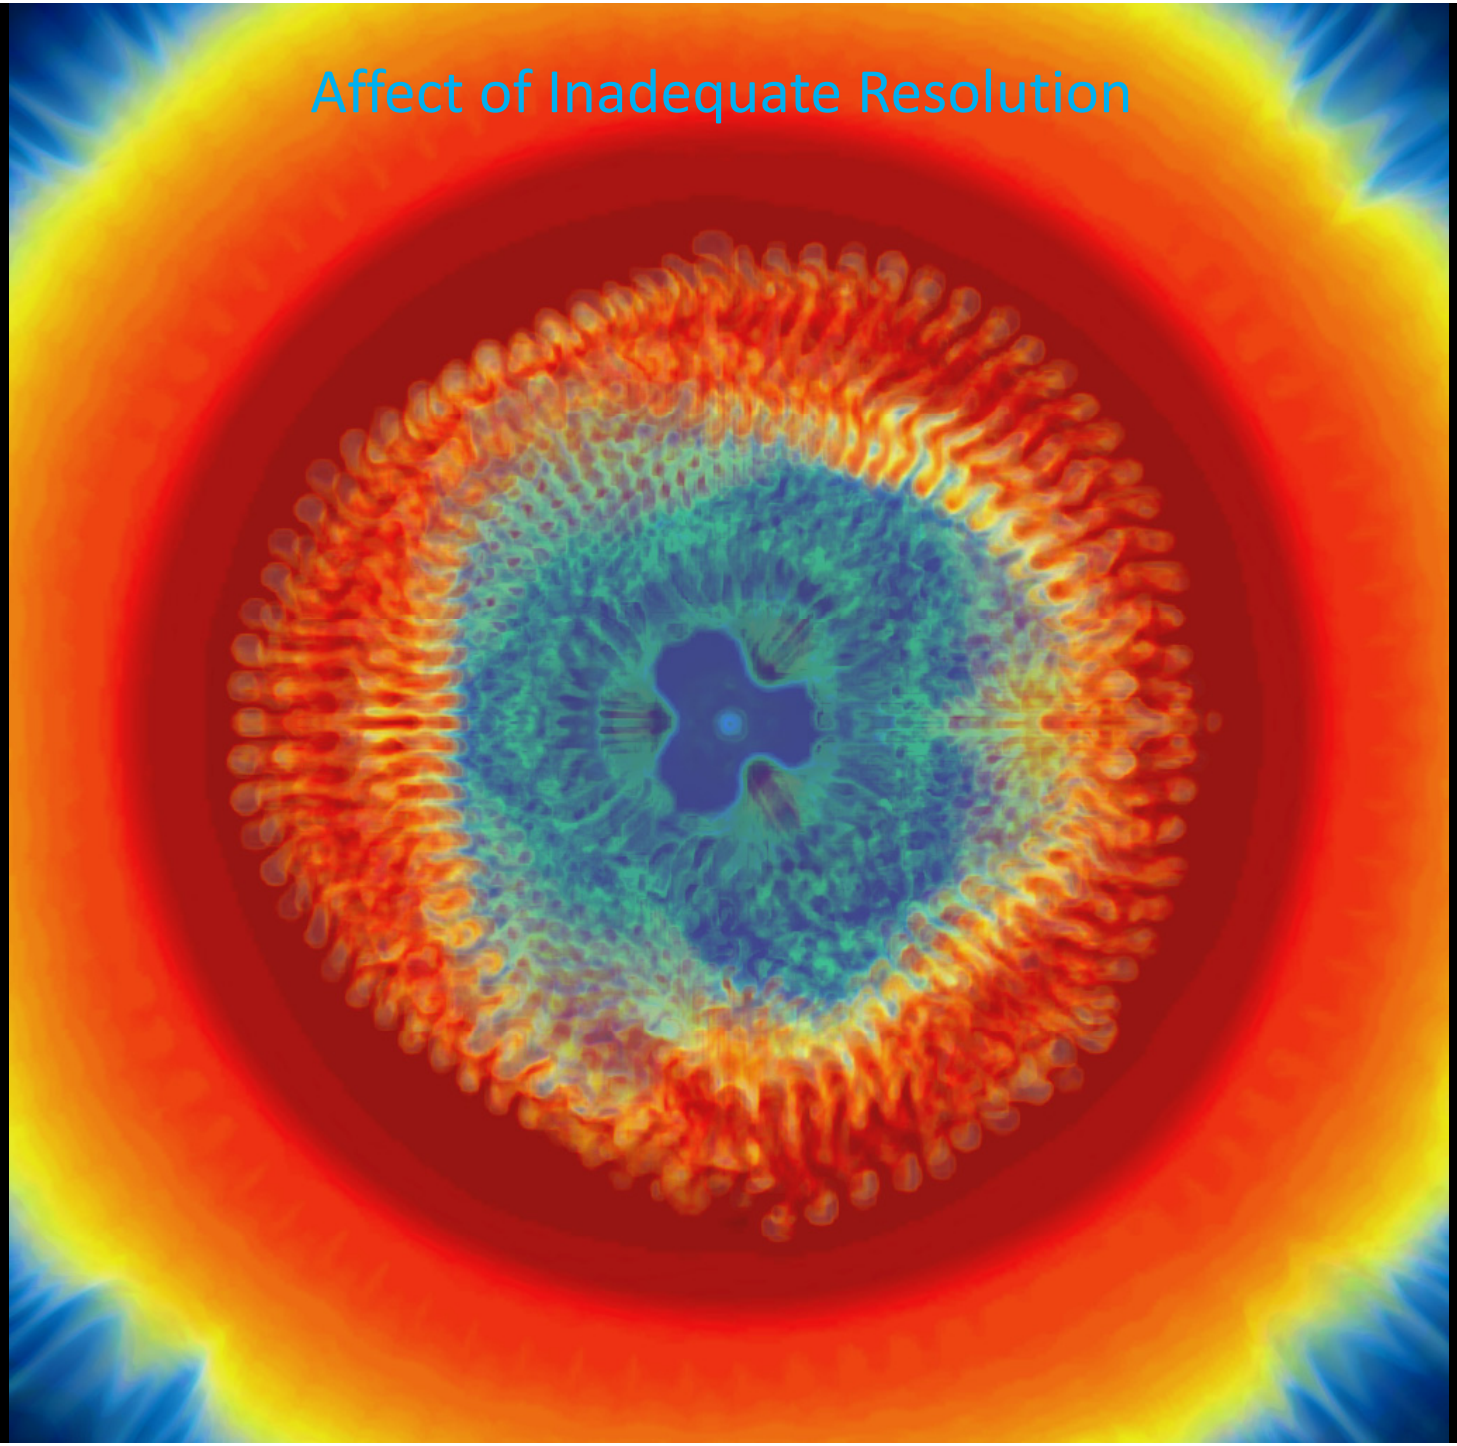

Inade-quate grid resol-ution accent-uates the jet visible at the mid-plane and dis-torts the shape of the cavity

Dump 10
608^3 grid

ρ

rho12-LUT opacity 1

NO GRID JIGGLING

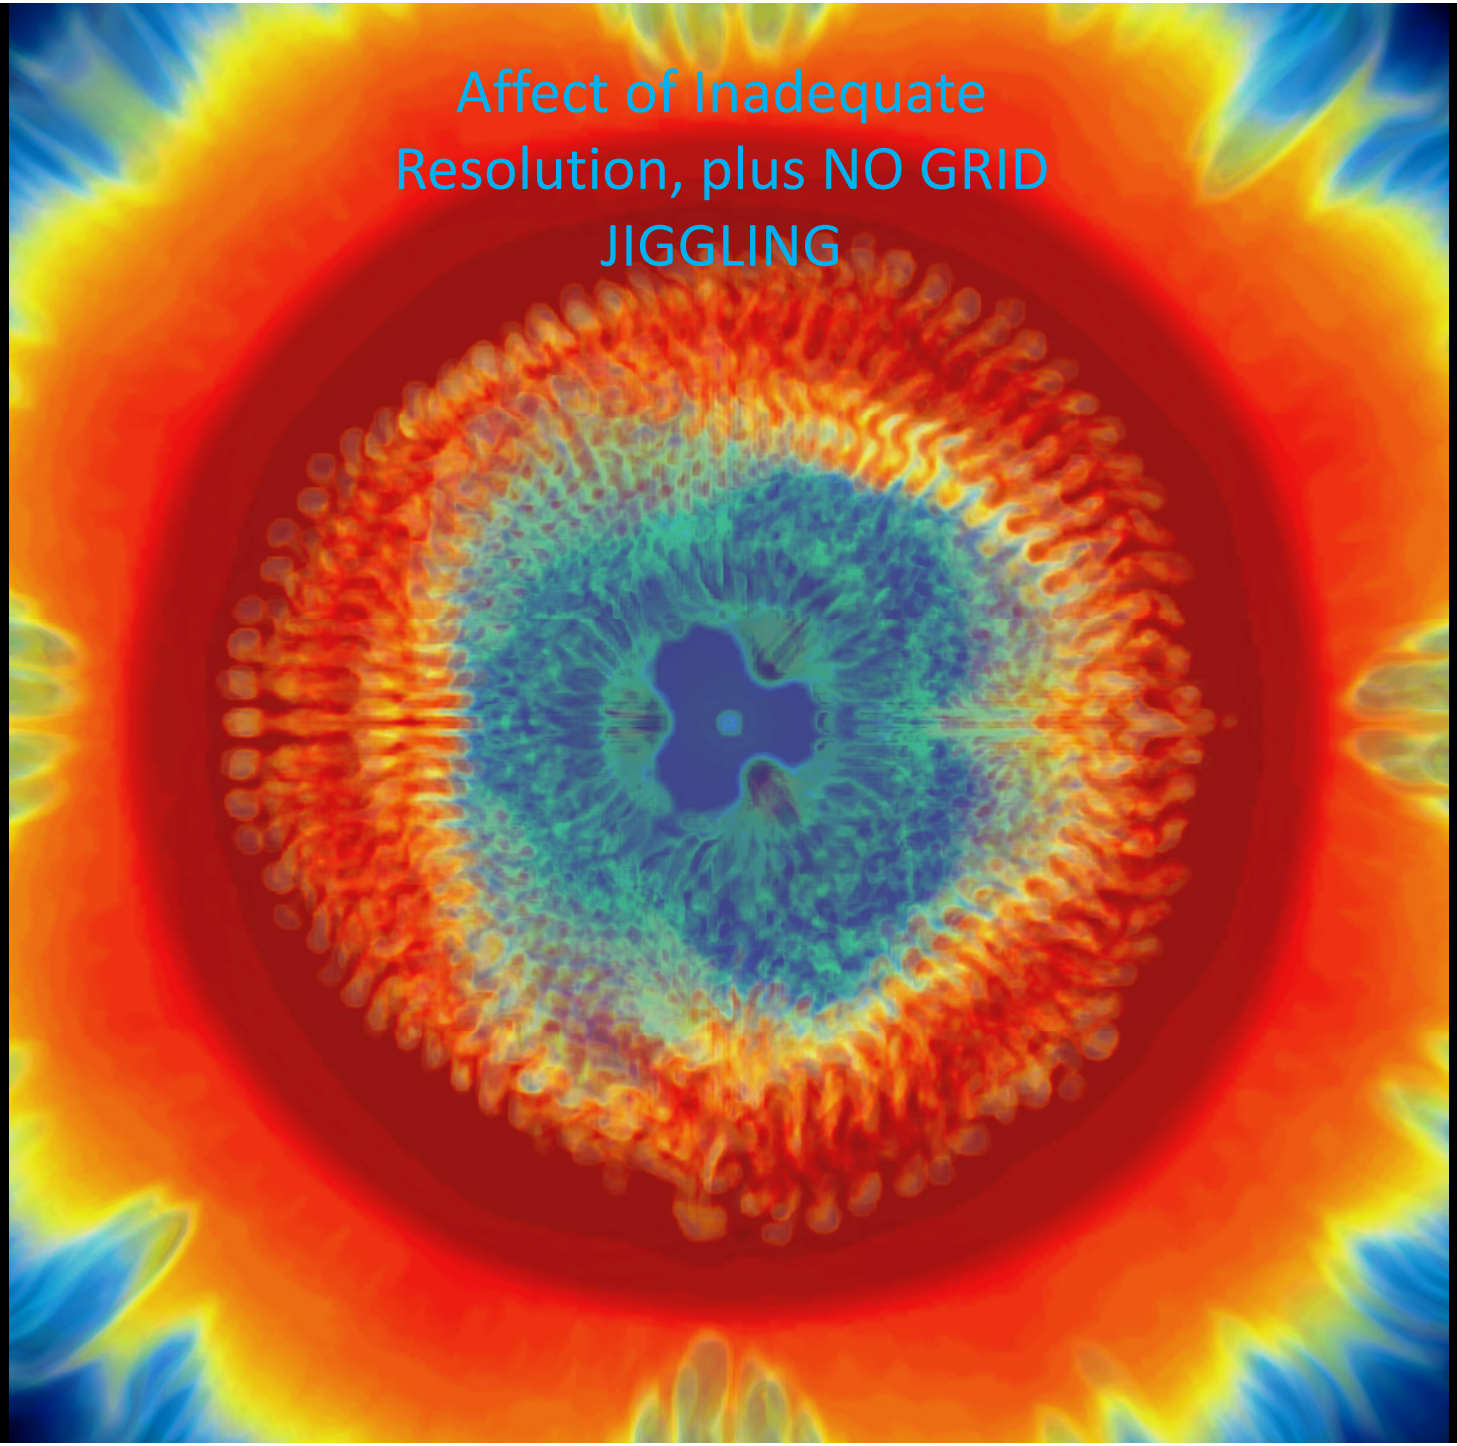

Jiggling the grid reduces the prominence of the "fountains" from grid-induced disturbances.

Dump 10
608^3 grid

ρ

rho12-LUT opacity 1

NO SMART SHOCK DISSIP-ATION

Affect of Inadequate Resolution, plus NO SMART SHOCK DISSIPATION

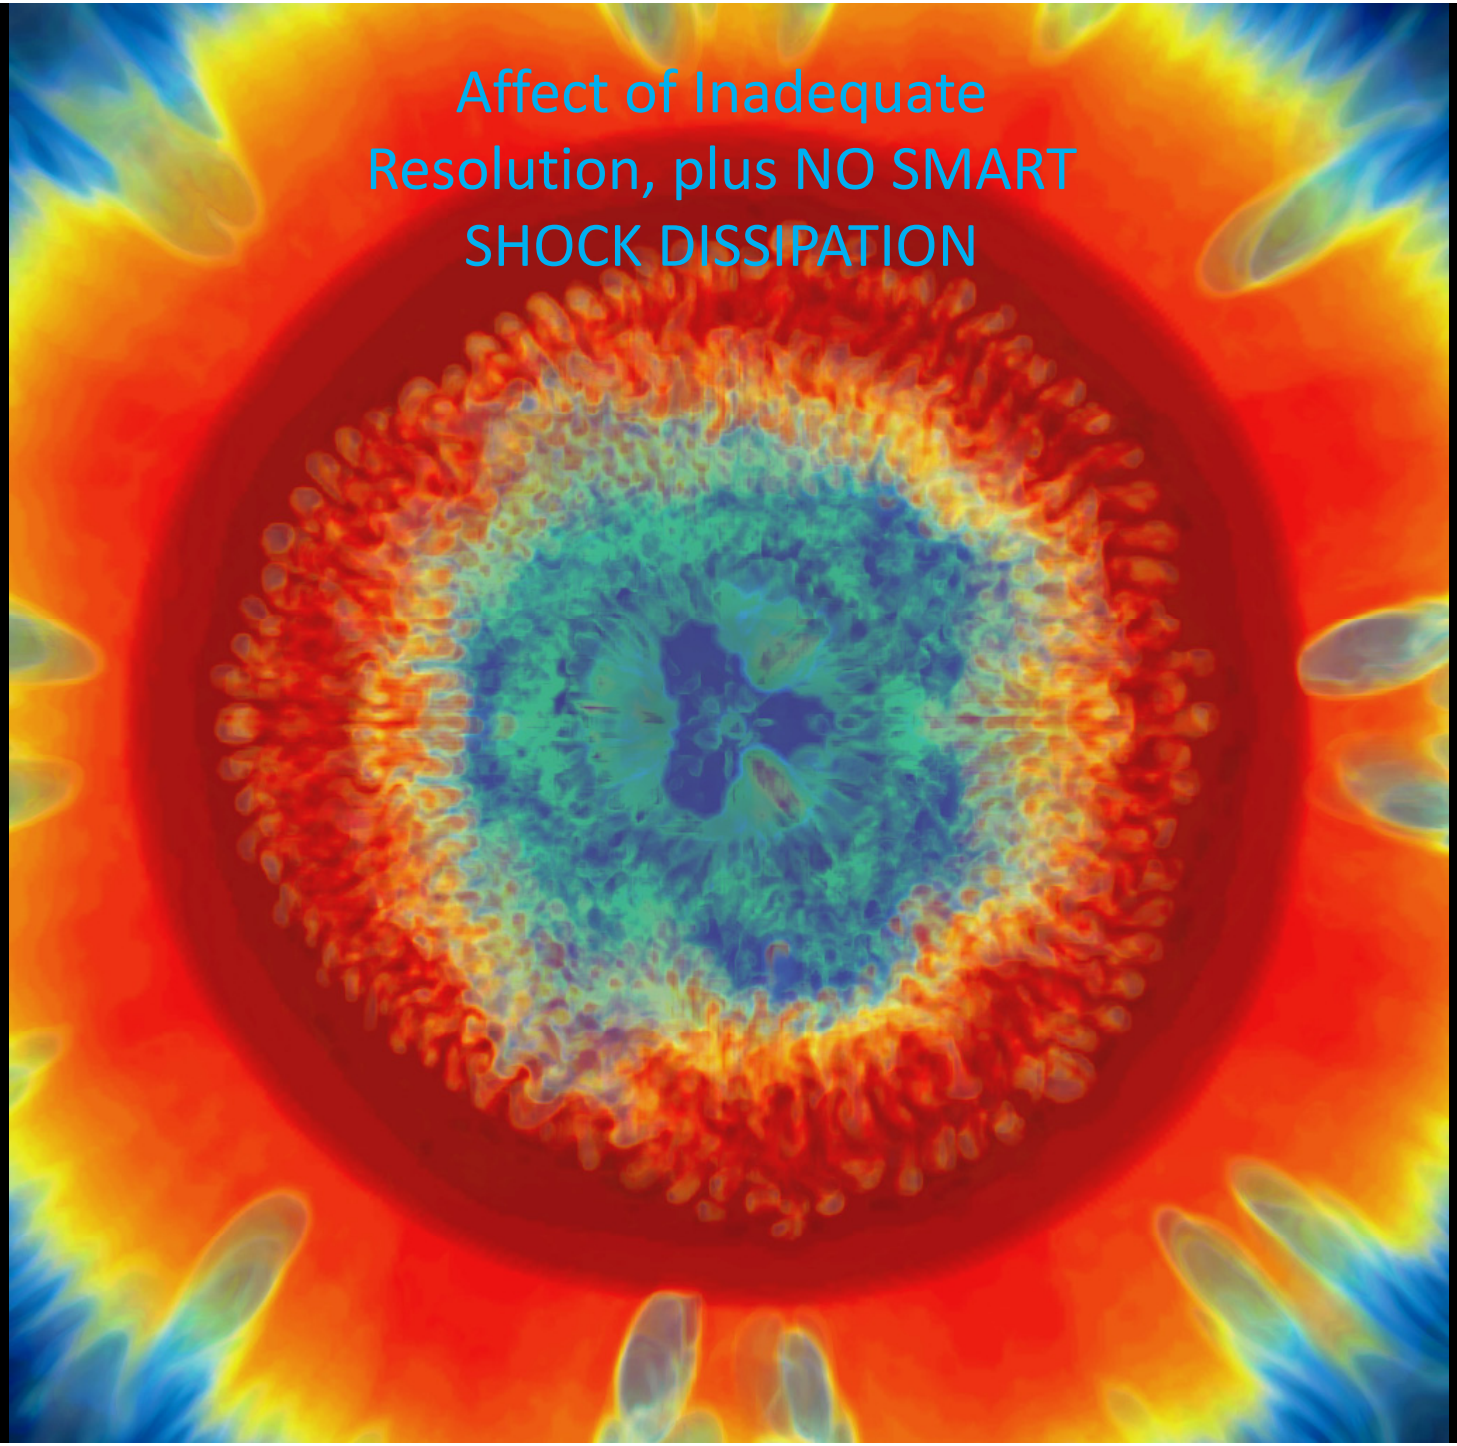

Adding smart dissip-ation to the treat-ment of strong shocks results in less con-tamin-ation by modes 4 and 8

***What about that PPB Multifluid Treatment*?**

- PPB advection conserves 10 moments of the sub-cell distribution exactly.
- It more than doubles the resolution of an advected quantity over PPM advection, which is already very good.
- For a sharp multifluid interface, however, the numerical representation cannot be thinner than one grid cell.
- For this reason, PPM advection with contact discontinuity detection and steepening is competitive.
- However, PPM advection is less well-behaved, because it has 10 times less information to work with.
- For advection of smooth distributions, PPM advection is far less accurate, but of course the multifluid volume fraction is not a smooth variable in this problem.
- We have observed that PPB has far fewer pathological behaviors; such behaviors occur, of course, only under rare sets of conditions (which nevertheless happen).

PPM with PPB multifluid advection compared with RTI-3D:

- As a part of a study of single-mode Rayleigh-Taylor instability, results for PPM with PPB multifluid volume fraction advection were compared to results from the RTI-3D code.
- In the *Physics of Fluids* paper on this work, "The late-time dynamics of the single-mode Rayleigh-Taylor instability," by P. Ramaprabhu, Guy Dimonte, P. Woodward, C. Fryer, G. Rockefeller, K. Muthuraman, P.-H. Lin, and J. Jayaraj, *Physics of Fluids* **24**, 074107 (2012), the RTI-3D code is described as follows:

The numerical codes used in this study have been described in detail in Refs. 30 and 31, and are briefly summarized below.

RTI-3D (Ref. 30) is an incompressible, multiphase 3D flow solver that employs a finite volume discretization of the governing Euler equations. For problems with sharp interfaces, flux-limiting is implemented through the third-order accurate Van Leer algorithm, preventing artificial overshoots and undershoots. Such numerical methods are classified as Monotonically integrated large eddy simulations (MILES) and have gained popularity in recent years due to their attractive properties in handling discontinuities. RTI-3D has been validated for single-mode RT and Kelvin-Helmholtz simulations,[30] and multimode RT in Ref. 10.

This figure has been taken from the article referenced on the previous slide.

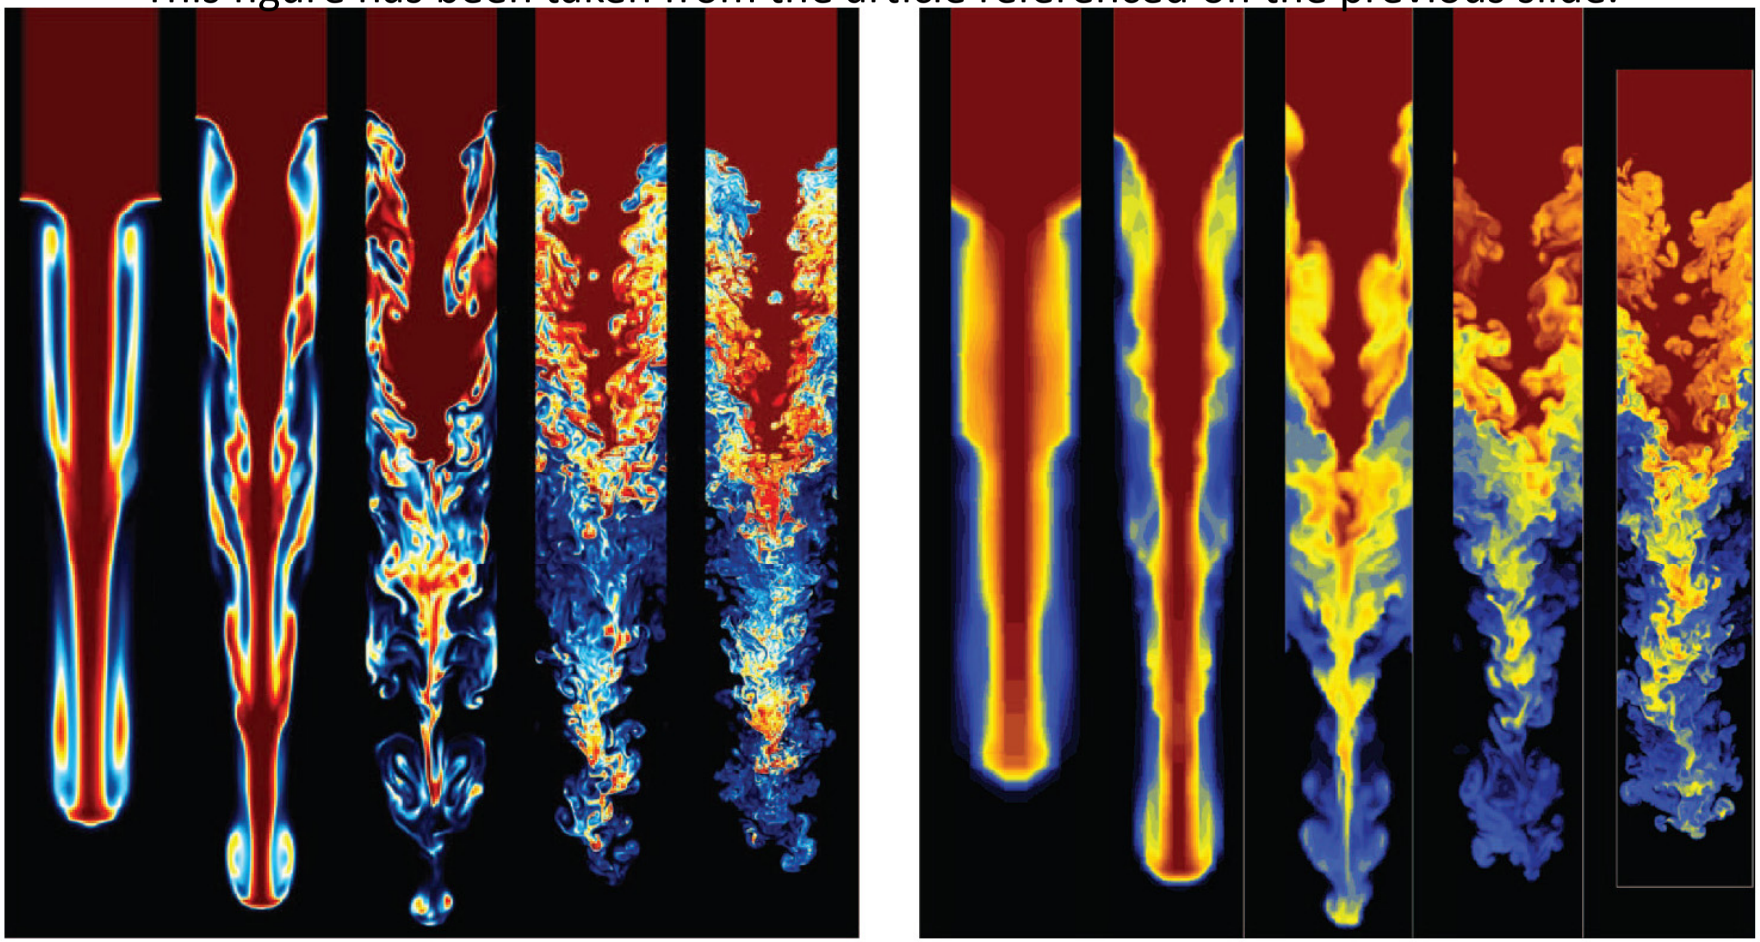

(a) (b)

FIG. 18. Late-time volume fraction contours from (a) PPM and (b) RTI-3D at 8, 16, 32, 64, and 128 zones, respectively, at $A = 0.15$. Images were chosen at times τ when bubble amplitudes are approximately the same.

PPM with PPB multifluid volume fraction advection is at the left.
RTI-3D, with a more standard treatment is at the right.

***What about Higher Modes and Lower Amplitudes*?**

- This particular case involves only azimuthal modes 36 and 39.
- It had seemed to have converged at 2048^3 resolution, and the run at 4160^3 resolution was meant to establish that.
- This case definitely does require a 2048^3 grid.
- It is not clear how much farther in mode number and in initial amplitude we can go on the 4160^3 grid, but it *is* clear that on modern machines, we can easily go to 8320^3 or even higher. After all, the 4160^3 run took only 8 hours on 144,000 AMD cores (and Intel cores run twice this fast on this code).
- Decades of experience indicate that convergence on these problems, whose solutions are not analytic, is linear at best.
- This implies that if we judge, by eye, that the description at 2048^3 cells is adequate for mode 39, then 4096^3 should be equally adequate for mode 78, and 8192^3 for mode 156.
- While the latest machines are still generally unavailable, we can get a sense of this by doing runs in 2D.

***What about AMR*?**

- This code is now being revised to incorporate a briquette-by-briquette AMR approach after the example of the RAGE code.
- We can get an idea of the potential effectiveness of AMR for this problem by observing the extent to which the grid imprint effects of the Cartesian mesh influence the results for coarsely zoned problems.
- Coarse problems, as shown on the previous slides, will tell us whether derefining portions of the problem domain would or would not destroy the accuracy of the finely zoned region.
- To see the behavior of the code on a coarse mesh of just 1216^3 cells, we double the initial perturbation amplitude.
- This gives us a fairly well converged general base flow, but of course the details of the mixing region cannot be well represented on this mesh. These are the natural regions for us to refine in any AMR treatment.
- Doubling the initial amplitude removes the distraction of inaccurate treatment of the unstable interface, which we would certainly refine in any AMR run.

Conclusions:

- We are encouraged by the results shown in the previous slides.
- The PPM+PPB multifluid approach was developed under contract from LANL to serve as a module of the RAGE code, and such a module was part of the 2005 XRAGE code release.
- In recent years, this numerical scheme was built as a stand-alone code targeting the Roadrunner machine at LANL, as well as other large systems based upon multicore CPUs.
- This extensive code restructuring work has paid off handsomely in delivered code performance.
- Without that performance, it would not make sense to attempt difficult ICF problems using uniform Cartesian grids.
- This PPM code is being extended to include RAGE-style AMR, but with briquettes of 4^3 cells playing the role of RAGE's cells.
- The code has already been adapted for the new Intel MIC co-processors in work in collaboration with Jim Ang, Richard Barrett, and Simon Hammond at Sandia.
- The techniques used in this code can be adapted to other codes, and should have equally beneficial results in those codes