Error analysis of the Thornado supernova simulation code

Kristopher Andrew[1] and Eirik Endeve[2]

[1]University of Kentucky
[2]Oak Ridge National Lab

August 8, 2018

Abstract

The iron core collapse of a supernova throws 1.5 solar masses at speeds up to 35% the speed of light into a space with radius 50km, which creates a high density environment capable of generating neutrinos, producing gravitational waves, and driving chemical evolution of matter past iron. By designing a simulation of the process, we can better understand the origins and properties of these rare phenomena which are not seen anywhere else in the universe. Typically, supernovas are simulated with a finite volume method, but we want to incorporate a Nodal Discontinuous Galerkin Method to better utilize high performance computing because the DG method allows us to increase precision without increasing the number of computational elements. This project focuses on the fluid dynamics during collapse and the resulting shockwaves. To understand this, we improved on an existing simulation by comparing benchmark tests against published and theoretical results to identify and correct sources of error. The DG Method creates discontinuous functions across computational elements by solving the Euler fluid equations computed over Runge-Kutta time steps. The DG method allows for discontinuous functions to be built from polynomials defined in individual

elements; however, it can lead to unphysical oscillation near discontinuities. We apply the characteristic slope limiter to reduce these affects, but when applied globally it leads to a 1% increase in absolute error. By introducing a Trouble Cell Indicator, which marks certain fluid elements based on the rate of change of observables, we can constrain the limiter to elements near shocks and manage the resulting absolute error. By constraining the limiter to discontinuous regions, we can remove what we know will be unrealistic without impacting naturally occurring oscillations.

Contents

I Introduction

Thornado aims to calculate the explosion of a core-collapse supernova using the Nodal Discontinuous Galerkin method (NDG). Following the description in [3], a supernova undergoes collapse because the fusion process increases the atomic mass of its fuel until it forms iron. Iron accumulates at the core which supports the weight of the star until the gravitational force of accumulating mass overcomes the force of electron degeneracy. The electrons and protons annihilate allowing the core to shrink into a much smaller space, and the matter falling into the remains of the core produces a shockwave which propagates upward into

the surrounding matter. Driven by neutrino production, circulation of the fluid allows new material to crash into the core and fuel the explosion process.

The NDG method subdivides a region of space into finite elements and measures properties in the elements based on the flux of properties through their faces. We are concerned with how shocks propagate through a medium, and so, our code needs to handle discontinuous solutions. The NDG method defines observable properties as a polynomial in an element but allows for discontinuities across element boundaries. We define an equation to find flux at these points and solve the Euler equations.

Discontinuities frequently result in oscillation, and so, we apply a limiting procedure to "smooth" the function into a more continuous line. The characteristic slope limiters try to reduce the magnitude of slope and flux between elements; however, while this process will eliminate unrealistic oscillations, it tends to be very expensive computationally and adds absolute error near regions containing a shock. According to the explanation in [4], the choice of limiter is highly problem dependent, and for the following error analysis, we see that the Thornado limiters will increase absolute error near shocks in attempts to dissipate unrealistic oscillation. To address this error, the Troubled Cell Indicator (TCI) [2]identifies regions of rapid change and flags them to be smoothed by the limiting procedures so the limiter will only be applied to discontinuities.

We have identified that the limiter is introducing absolute error at oscillations and discontinuities; however, this error can be managed by reducing the area impacted by the limiter with an appropriate TCI cut-off. In Section II, we introduce the Euler equations from fundamental laws and rewrite them in the form of conservation equation. In Section III, we explain how the Euler equations can be approximated by the NDG method. In Section IV, we talk about the simulations we ran to identify possible sources of error, and how they support our finding that the TCI is capable of reducing the absolute error introduced by the characteristic slope limiter.

II Mathematical model

We can express conservation of mass, Newton's second law, and the first law of thermodynamics in terms of the following volume integrals,

$$\partial_t \int_V \rho dV = -\oint_{\partial V} \rho \vec{v} \cdot d\vec{A}, \tag{1}$$

$$d_t \int_{V(t)} \vec{m} dV = -\int_{\partial V(t)} p d\vec{s} + \int_{\partial V(t)} \vec{f} dV, \tag{2}$$

$$dE = -pdV + Tds, \tag{3}$$

These integrals relate density ρ, momentum $\vec{m}$, energy E, and pressure p in three equations with momentum being defined by $\vec{m} = \rho\vec{v}$ with $\vec{v}$ being velocity. Volume V is held constant by setting the fluid elements to be the same size for all times. The change in entropy ds is assumed to be zero in the Euler equations, and so, to solve, we use an equation of state to relate pressure to internal energy e.

We can use Gauss's law to simplify the equations to show conservation of mass, conservation of momentum, and conservation of energy in the form of the following Euler equations.

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho\vec{v}) = 0 \tag{4}$$

$$\frac{\partial \vec{m}}{\partial t} + \nabla \cdot (\vec{m} \otimes \vec{v} + pI) = -\rho\nabla\phi \tag{5}$$

$$\frac{\partial E_F}{\partial t} + \nabla \cdot ((E_F + \rho)\,\vec{v}) = -\rho\vec{v}\nabla\phi \tag{6}$$

These conservation equations are more practical in the simulation as we are able to maintain values across element boundaries. These equations provide important constrains on the density, momentum, and free energy E_F across the boxes. This also demonstrates that the external gravitational potential field ϕ will add momentum and energy to the fluid as it moves in certain directions.

III Numerical method

Physical properties are calculated using the Nodal Discontinuous Galerkin Method with Runge-Kutta time steps. The NDG method provides a solution to discontinuous Partial Differential Equations by representing the solution in the form of a summation of k polynomials where k equals the number of elements. Following from the above equations 4, 5, and 6, we start with equations in the form

$$\partial_t u + \partial_x(f(u)) = s(u) \tag{7}$$

where u is the property we want to measure, $f(u)$ represents the flux of the given property, and s represents a source or sink of the given property. By defining the test space $\mathbb{V}^k$ of all polynomials up to degree k,

$$\mathbb{V}^k = \{v : v \in \mathbb{P}^k, \forall K \in D\} \tag{8}$$

we multiply Equation 7 by an element of $\mathbb{V}^k$ and integrate over K to reach the following solution,

$$\int_K \partial_t u_h v dx + [\hat{f}(u_h)v^- \mid_{x_H} - \hat{f}(u_h)v^+ \mid_{x_L}] - \int_K f(u_h)\partial_x v dx = \int_K s(u_h)v dx \tag{9}$$

where the flux over a boundary is given by,

$$\hat{f}_{H/L} = \frac{1}{2}(f(a) + f(b) - \alpha(b-a)), \alpha = max(|\lambda(a)|, |\lambda(b)|) \tag{10}$$

The definition of flux is an important because discontinuities in u would result in different values of $f(u)$ when coming from different directions. a and b represent the limit of u as it approaches either side of a boundary, and $\lambda(a)$ and $\lambda(b)$ represent Eigen values of the diagonal of matrix $\frac{df}{du}$. Discontinuous flux would result in a failure to maintain the Euler conservation laws, and so, this definition ensures that flux values are equal at the boundaries of elements.

Additionally, we can expand the above solution into vector components in the following fashion,

$$\sum_{j=1}^{N} \int_K v_i \phi_j dx \partial_t \hat{u}_j = -(\hat{f}_H v_{i,H}^- - \hat{f}_L v_{i,L}^+) + \int_K f \partial_x v_i dx + \int_K s v_i dx \tag{11}$$

where $\sum_{j=1}^{N}$ iterates over the dimensions of u. Equations 22 and 11 allow us to measure observable properties in elements and at their boundaries. In order to progress this solution through time, we use the Runge-Kutta method defined below,

$$\hat{u}^{(i)} = \sum_{j=0}^{i-1} \alpha_{ij}(u^{(j)} + \beta_{ij}\delta t M^{-1} G(\hat{u}^{(j)})) \tag{12}$$

where i designates a specific time step, M^{-1} designates the diagonalized matrix, and $G(\hat{u}^{(j)})$ represents the right hand side of equation 11. α_{ij} and β_{ij} are chosen constants, and for the 1-stage first order method they follow,

$$\alpha_{ij}, \beta_{ij} \geq 0, \tag{13}$$

$$\sum_{j=1}^{m} \alpha_{ij} = 1 \tag{14}$$

This equation estimates the solution at all time steps by summing the areas of rectangles of width β_{ij}. This method gives a rough estimate of the integral, and the change in the

solution over time.

IV Numerical experiments

This section describes the benchmark tests we used to compare the simulation against other published simulations.

A Sod's shock tube

Sod's Shock tube is a simulation that begins as a one dimensional tube divided in two by an imaginary boundary. The boundary results in a discontinuous function in density and pressure, and at t = 0, the boundary is removed. The fluid flows into the lower pressure region and forms three distinct shocks. The fluid flowing into the lower pressure region forms a shock wave that is followed by a contact discontinuity. The first is a discontinuity in the density, velocity, and internal energy, but the second is a discontinuity shown only in the density. The third shock is the rarefaction wave which propagates into the higher pressure region. This wave is a smooth slope of the elements accelerating into the lower pressure region.

The simulation runs until $t = 0.2$ in the region $x = 0$ to $x = 1$. For initial conditions of

$$(\rho, v, p) = (1.0, 0.0, 1.0) \text{ for } x \in [0, 0.5]; \qquad (\rho, v, p) = (0.125, 0, 0.1) \text{ for } x \in (0.5, 1].$$

From Figure 1, we can clearly see the formation of three shocks matching the expected types. These results match findings from [2], and suggest Thornado is capable of simulating important types of shock waves. From Figure 3, we can see that removing the limiters reduces absolute error, and the shock boundaries are the greatest sources of error. This is confirmed visually in Figure 2 where the shocks are better captured at the cost of oscillation.

Below we see the percent difference between 256 element simulations with different TCI cut-offs showing that the absolute error decreases with a higher cut-off.

TCI cut-off	% Error
0	0.170
0.03	0.168
0.10	0.160
None	0.048

B Shock-entropy wave interaction

The Shock-Entropy Wave Interaction describes a shock propagating into a region of sinusoidal density. This is meant to dissipate energy in a way that might be expected due to entropy and models how a shock might react in a turbulent region of the supernova. As the shock travels into the sinusoidal region, the density behind the shock begins to oscillate in two distinct patterns: a sawtooth wave and an erratic region close to the shock.

The simulation runs until $t = 1.8$ in the region $x = -5$ to $x = 5$ with initial conditions matching Fu's setup in [2],

$$(\rho, v, p) = (3.857143, 2.629369, 10.333333) \text{ for } x \in [-5, -4]; \tag{15}$$

$$(\rho, v, p) = (1 + 0.2 * \sin 5 * x, 0, 1) \, for x \in (-4, 5]; \tag{16}$$

Figure 4 shows simulations run with an increasing number of elements compared to a simulation run at 2048 elements. Based on the reference compared to the results of [2], it would appear we have a simulation that approaches the analytical solution. From Figure 5 we can see the error decreases for all conditions as the number of elements increases; however, error is again reduced with the removal of the slope and characteristic limiter. We can see that the difference between these two is reduced with the use of the TCI. This is obvious when one considers that the TCI removes boxes from the limiting process, and so, as less of the solution is limited, the difference between the limiting and non-limiting solutions should decrease. The boxes that the TCI selects for limiting can be seen in Figure 8. Below we see the percent difference between each 256 element simulation with different TCI cut-off values. This shows how absolute error is managed by the TCI.

TCI cut-off	% Error
0	1.75
0.03	0.82
0.10	0.75
None	0.43

C Rayleigh Taylor instability

The Rayleigh Taylor Instability is a problem where a region of high density fluid is suspended over a region of low density fluid in a uniform gravitational potential field. A perturbation at the boundary allows development into a lower energy state. Without perturbation, the

following initial conditions are stable, but with perturbation, the lower density fluid will bubble up as the higher density fluid pushes itself into the region of lower potential. Using the region from $[x, y, z] = [-0.25, -0.75, 0]$ to $[x, y, z] = [0.25, 0.75, 1]$ and simulating to $t = 8.5$ with the conditions,

$$(\rho, v, p) = (1, < 0.0\hat{x}, 0.0\hat{y}, 0.0\hat{z} >, 1) \text{ for } y \in [-0.75, 0]; \tag{17}$$

$$(\rho, v, p) = (2, < 0.0\hat{x}, 0.0\hat{y}, 0.0\hat{z} >, 1) \text{ for } x \in (0, 0.75]; \tag{18}$$

we get the density distribution shown in Figure 9 by changing the initial velocity field to

$$v = \frac{A_s}{4}[1 + \cos\frac{2\pi x}{L_x}] \times [1 + \cos 3\pi y]\hat{y} \tag{19}$$

This follows the set up of [1] and we see similar findings. In Figure 9 we see the resolution improving with increasing number of elements; however, following the behavior seen in the previous sections, removing the limiter in Figure 10 greatly increases the resolution. We can see that the swirls are much better defined, and another Rayleigh Taylor plume is developing on top. In Figure 11, we see increasing the TCI cut-off begins to improve resolution, but the high cut-offs diverge due to strong, early oscillations. Increasing the cut-off appears to cause chaotic divergence, and is likely due to the limiter ignoring some but not all of these oscillations.

To better understand this divergence, we looked at the maximum kinetic energy as a function of time. From Figure 12, we see full limiting and no limiting provide bounds for kinetic energy of tests performed with the TCI. Full and no limiting appear to provide bounds except for the range of $t = 5$ to $t = 6$. It is possible that this violation is what caused the chaotic deviation in Figure 11.

D Stellar core collapse

This simulation gives the development of a supernova undergoing core collapse. We set up the problem from a data set generated from a self similar solution proposed by Yahil in [5] and simulate almost to the transition to nuclear matter density and core bounce. The simulation fails to account for the change in equation of state at nuclear matter density, and so, the simulation is stopped before it reaches $t = 0$, when density becomes infinite. We ran simulations at 64, 128, and 256 elements with an outer boundary of 100,000km. For each simulation, an appropriate zoom factor was applied so that each element could increase in size until the innermost element was 1km or 2km depending on the simulation.

In Figure 13 we see that the simulation matches the reference solution as long as a time shift is applied. Because the solution is self similar, an arbitrary time shift is applied for each file so that the final density will match. By looking at the density everywhere within the computational domain at this time, we can see that the solution maintains the same shape for all times matching our expectation for the solution to be self similar.

κ is meant to be the constant ratio $\kappa = \frac{p}{\rho^\gamma}$, but in Figure 14, we can see how this ratio is dependent on position and time implying that entropy is not being held constant like we assumed for Equation 3. Below we see the percent difference in the average value that κ differs from the right most element of κ which appears to change the least over time. All values are from simulations with 256 elements and a 2km inner radius.

TCI cut-off	% Diff in Kappa
0	119.15
0.03	0.4710
0.10	0.1372
None	0.1372

Additionally, we were concerned with energy conservation. Because our simulation region does not include the entire star, it is possible for total mass and energy to increase, and so, we want to know how much mass and energy is introduced and whether Thornado correctly accounts for this flux. In Figure 15, we can visually see how flux and energy balance each other over time. Below we see the percent difference in mass and energy between Thornado and the Yahil equations.

% Diff in Mass	% Diff in Energy
1.077	28.14

To find this difference in energy, we write a conservation of energy following Equation 7 as the following,

$$\partial_t E_T + \frac{1}{r^2}\partial_r(r^2 F_T) = 0, \tag{20}$$

where E_T is the total sum of internal energy, kinetic energy, and gravitational potential energy, and F_T is the flux defined by,

$$F_T = (e + p + e_{kin} + p\Phi)v \tag{21}$$

where e is the internal energy, e_{kin} is the kinetic energy, and Φ is the gravitational potential. We can re-write Equation 20 as

$$\partial_t 4\pi \int_0^R E_T r^2 dr = -4\pi \int_0^R \partial_r (r^2 F_T) dr \tag{22}$$

by integrating over all space and multiplying by r^2. R is the outer radius. We see that the partial cancels with the integral, and the right hand side is simplified to

$$-4\pi R^2 F_T \mid_R \tag{23}$$

and we rewrite the left hand side to

$$\varepsilon_T = 4\pi \int_0^R E_T r^2 dr \tag{24}$$

allowing us to integrate over all time, and find

$$\varepsilon_T(t) = \varepsilon_T(0) - 4\pi \int_0^t (R^2 F_T) \mid_R dt \tag{25}$$

which would imply that the difference in energy should match the flux at the outer boundary seeming to match our observations in Figure 15

V Conclusion

In conclusion, we have identified that the limiter is introducing absolute error at oscillations and discontinuities; however, this error can be managed by reducing the area impacted by the limiter with an appropriate TCI cut-off. In future simulations, we will have more options to increase precision without increasing run time. We expect there to be some cost in variance as the TCI cut-off increases, and so, we still need to find an ideal cut-off value to minimize all forms of error. These findings also fail to compare error against the finite volume method that we want to replace. In future studies, we hope to see an optimal cut-off value that out computes other methods.

VI Acknowledgements

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.
This work was supported in part by the U.S. Department of Energy, Office of Science, Office

of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship program.

References

[1] C. Cardall, R. Budiardja, E. Endeve, and A. Mezzacappa. GENASIS: GENERAL ASTROPHYSICAL SIMULATION SYSTEM. I. REFINABLE MESH AND NONRELATIVISTIC HYDRODYNAMICS. *The Astrophysical Journal Supplement Series*, 210:1–29, 2014.

[2] C. Fu, G. Shu. A new troubled-cell indicator for discontinuous Galerkin methods for hyperbolic conservation laws. *Journal of Computational Physics*, 347:305–327, 2017.

[3] H. Janka. Explosion Mechanisms of Core-Collapse Supernovae. *Annual Review of Nuclear and Particle Science*, 62:407–451, 2012.

[4] K. Schaal, A. Bauer, P. Chandrashekar, R. Pakmor, C. Klingenberg, and V. Springel. Astrophyscical hydrodynamics with a high order discontinuous Galerkin scheme and adaptive mesh refinement. *MNRAS*, 0:1–23, 2015.

[5] A. Yahil and J. Lattimer. *Supernovae: A Survey of Current Research*. D. Reidel Publishing Company, 1982.

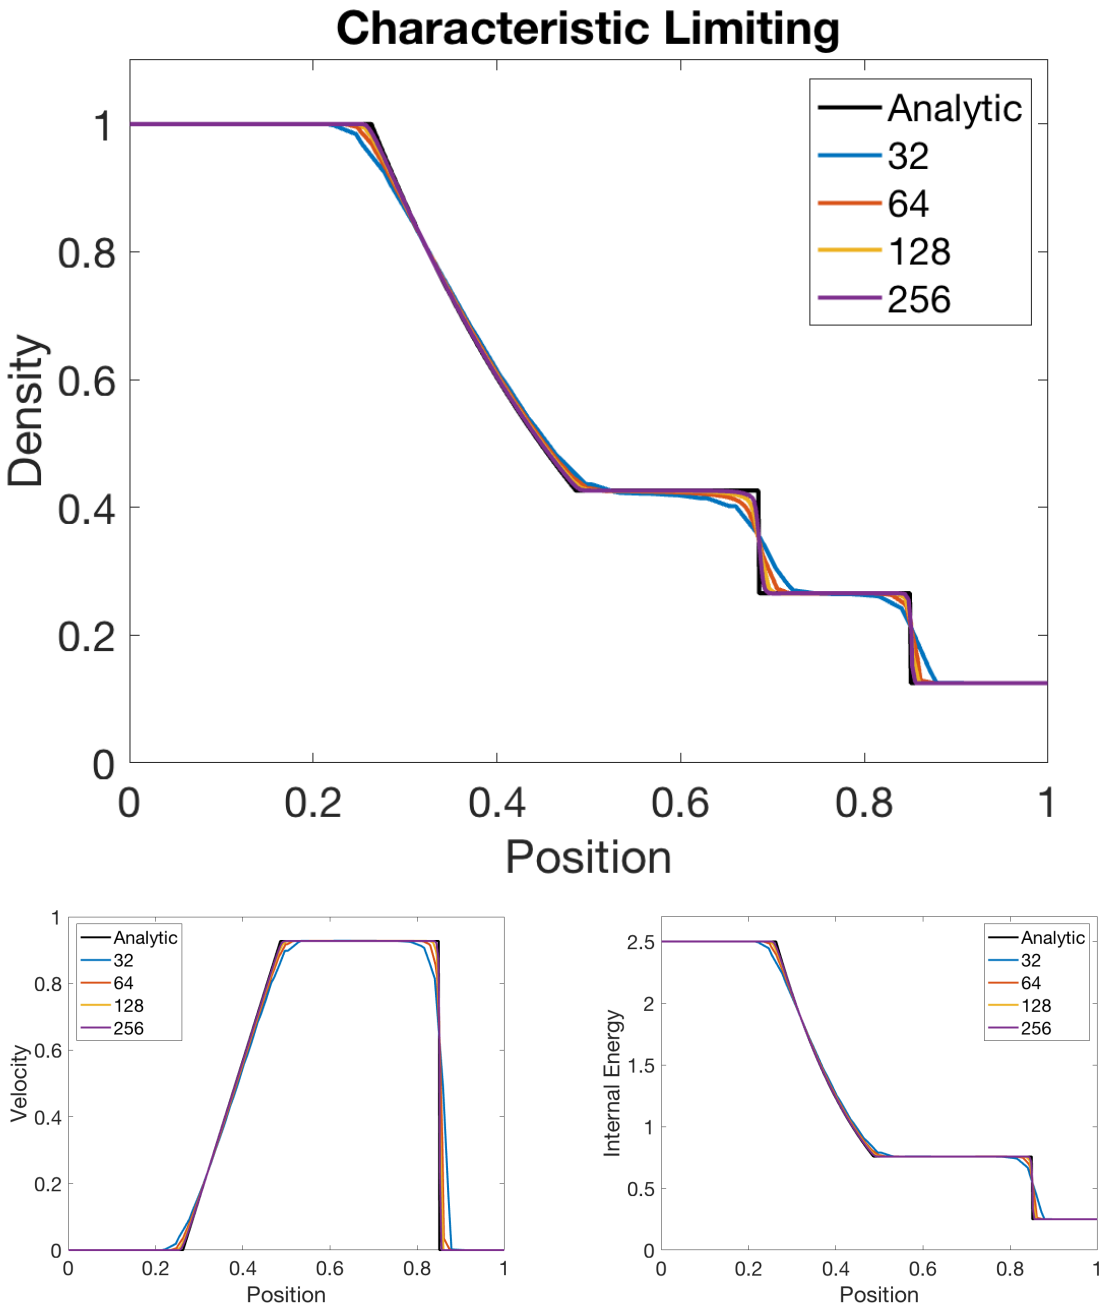

Figure 1: A plot of the Sod shock tube done for density, velocity, and internal energy. All subject to slope and characteristic limiting.

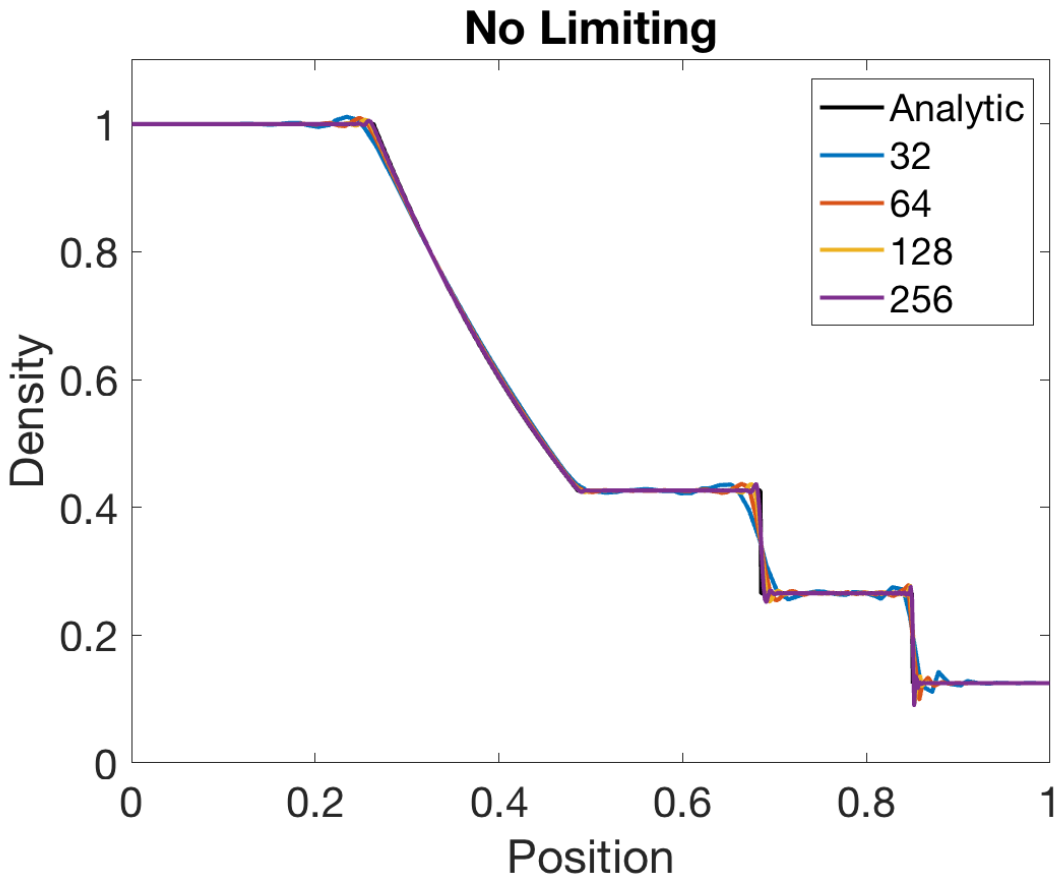

Figure 2: A plot of a one dimensional Riemann Problem computed numerically and analytically with increasing precision.

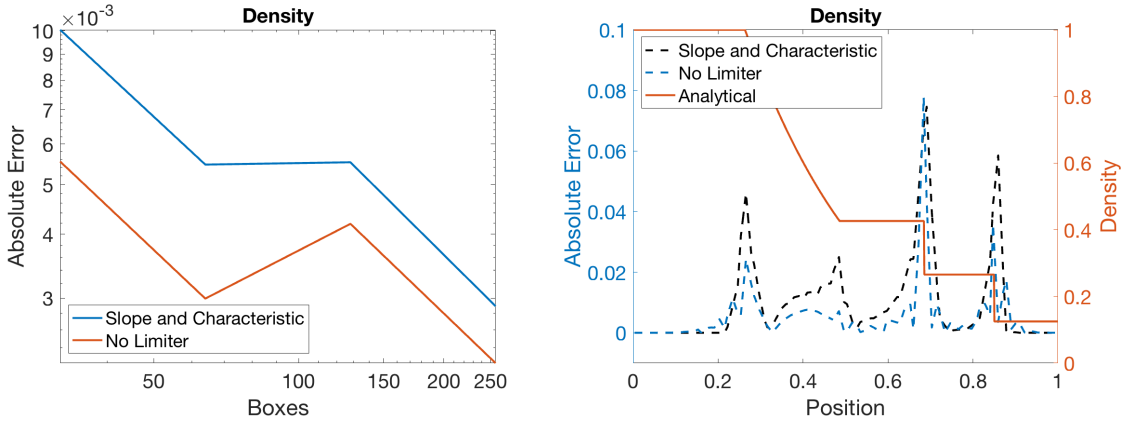

Figure 3: Here we see the error of various methods. We can see that no limiting has less error, and error occurs most often at the shockwaves. The right image is overlain with the analytical density plot.

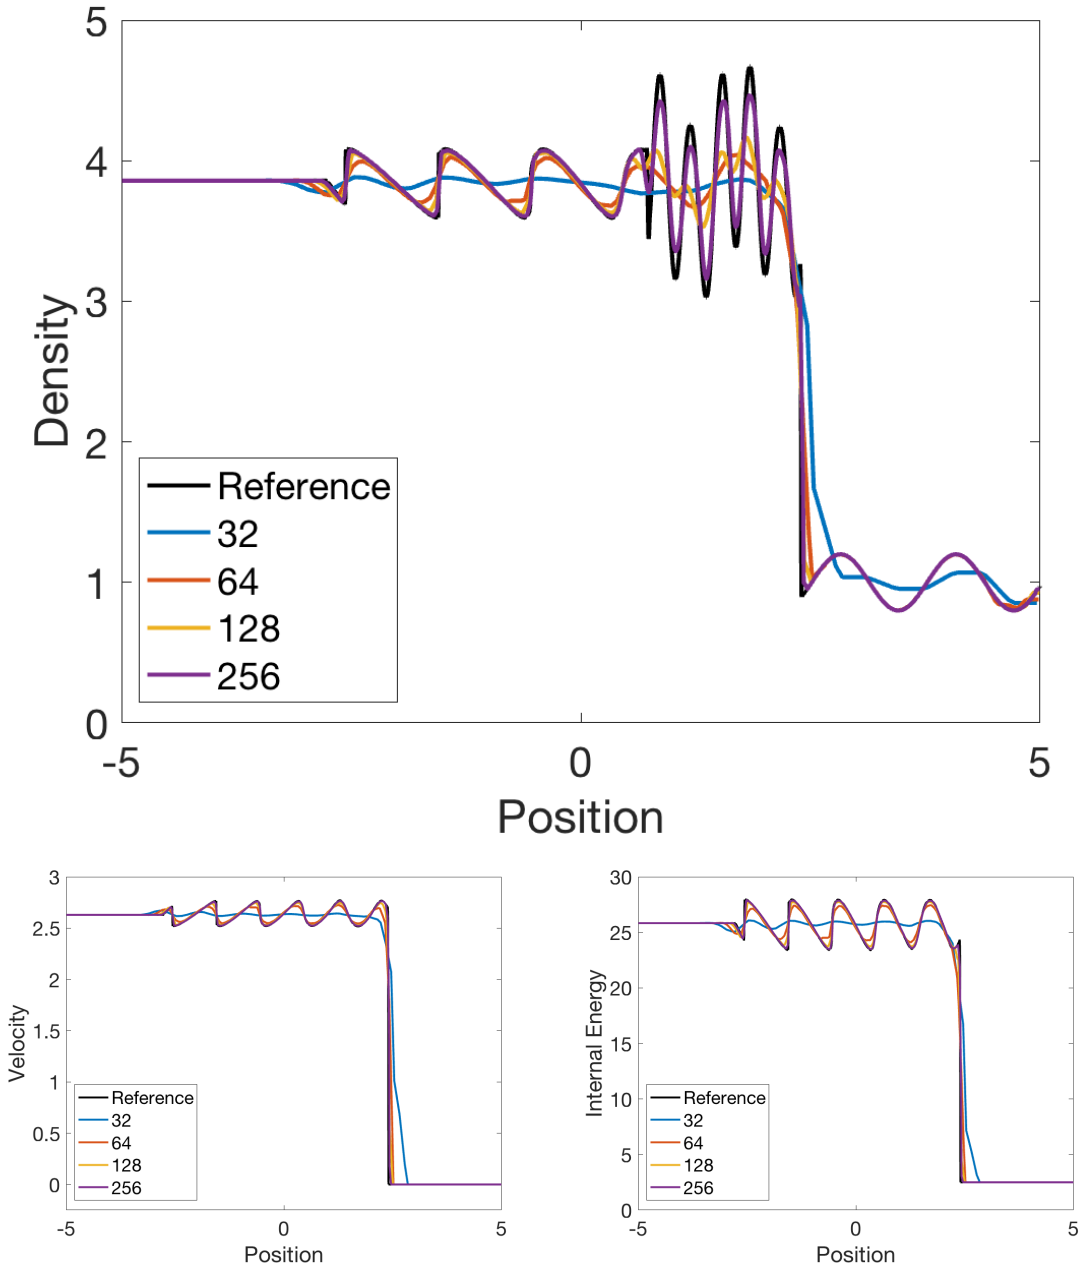

Figure 4: Density, Velocity, and Internal Energy plots. Analytic solution obtained from 2048 boxes.

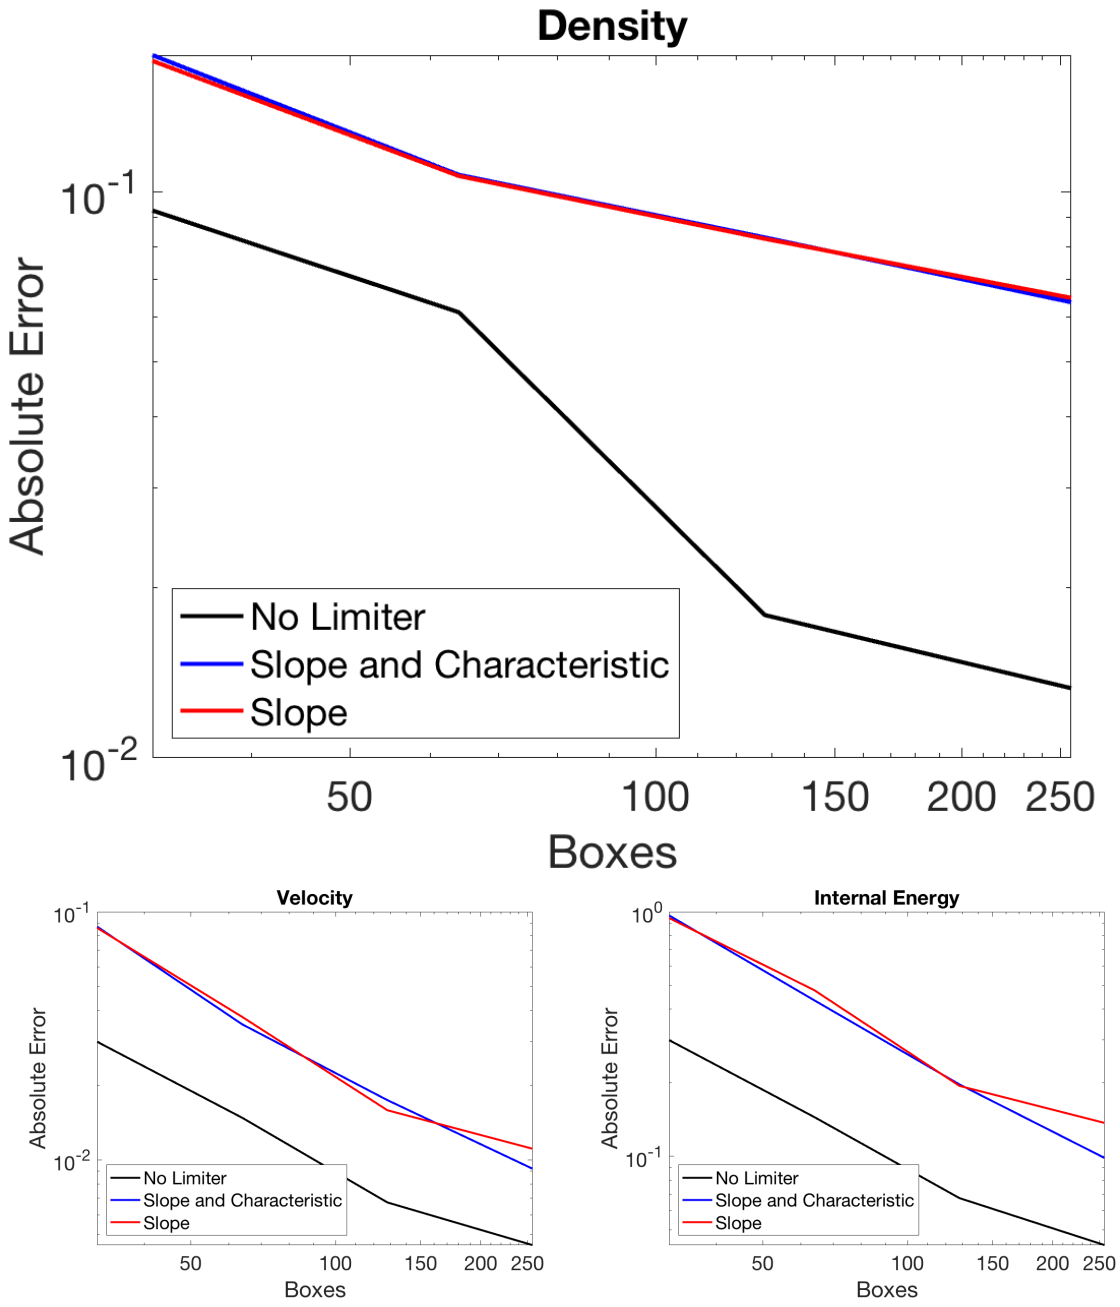

Figure 5: The Error as number of boxes is increased. We can see that due to the oscillatory nature of the problem, error is lower when no limiting is applied.

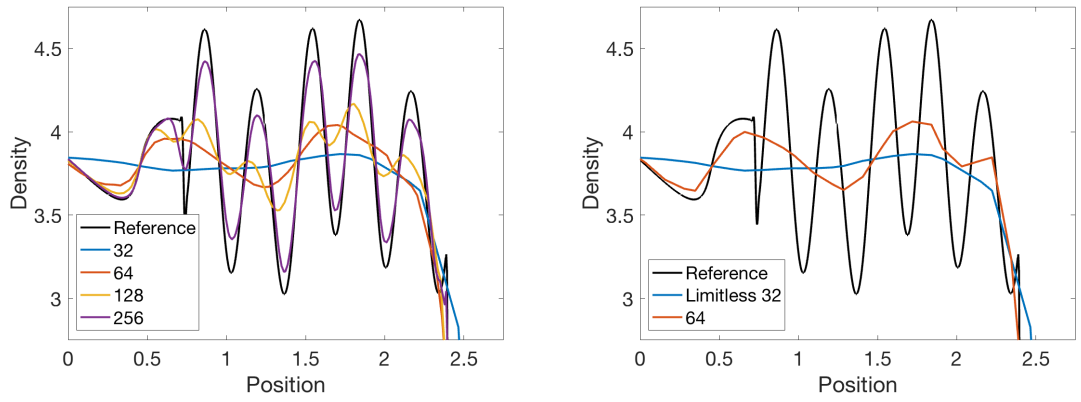

Figure 6: Here we see a larger image of the region of intense oscillation and the same region with a solution of 32 boxes with and without the slope and characteristic limiter. Here we can see visually see how the second solution more closely matches the limited 64 box solution. Matching the findings in our error.

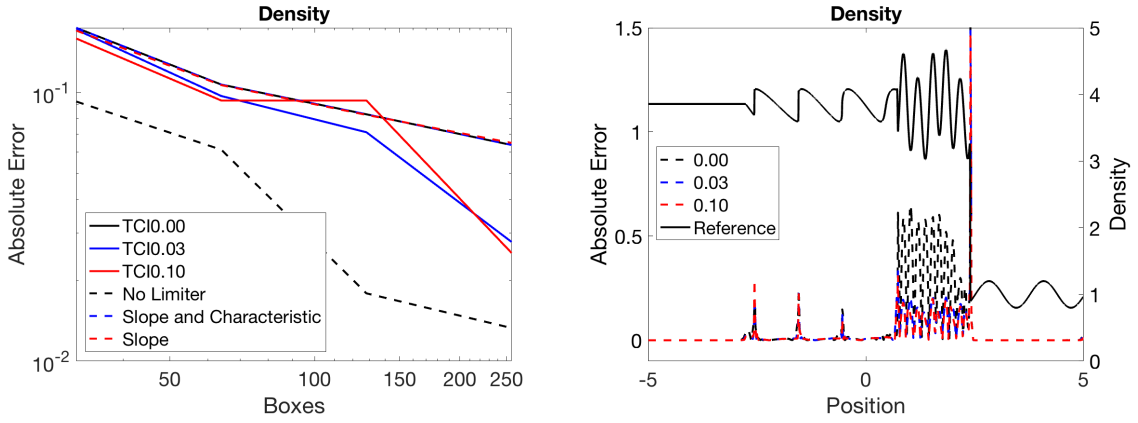

Figure 7: Here we see the error of various methods. We can see that having no limiter is the most accurate, and as the TCI removes boxes from the limiting process, the solution gets more accurate. Looking at error versus position, we can see that the TCI has the most problems with the region of high frequency.

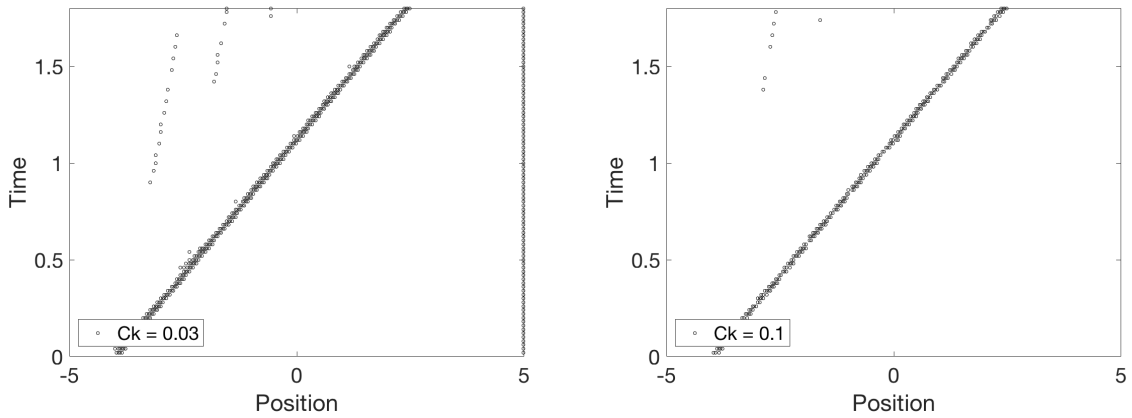

Figure 8: Here we can see which locations the Trouble Cell Indicator has flagged for limiting.

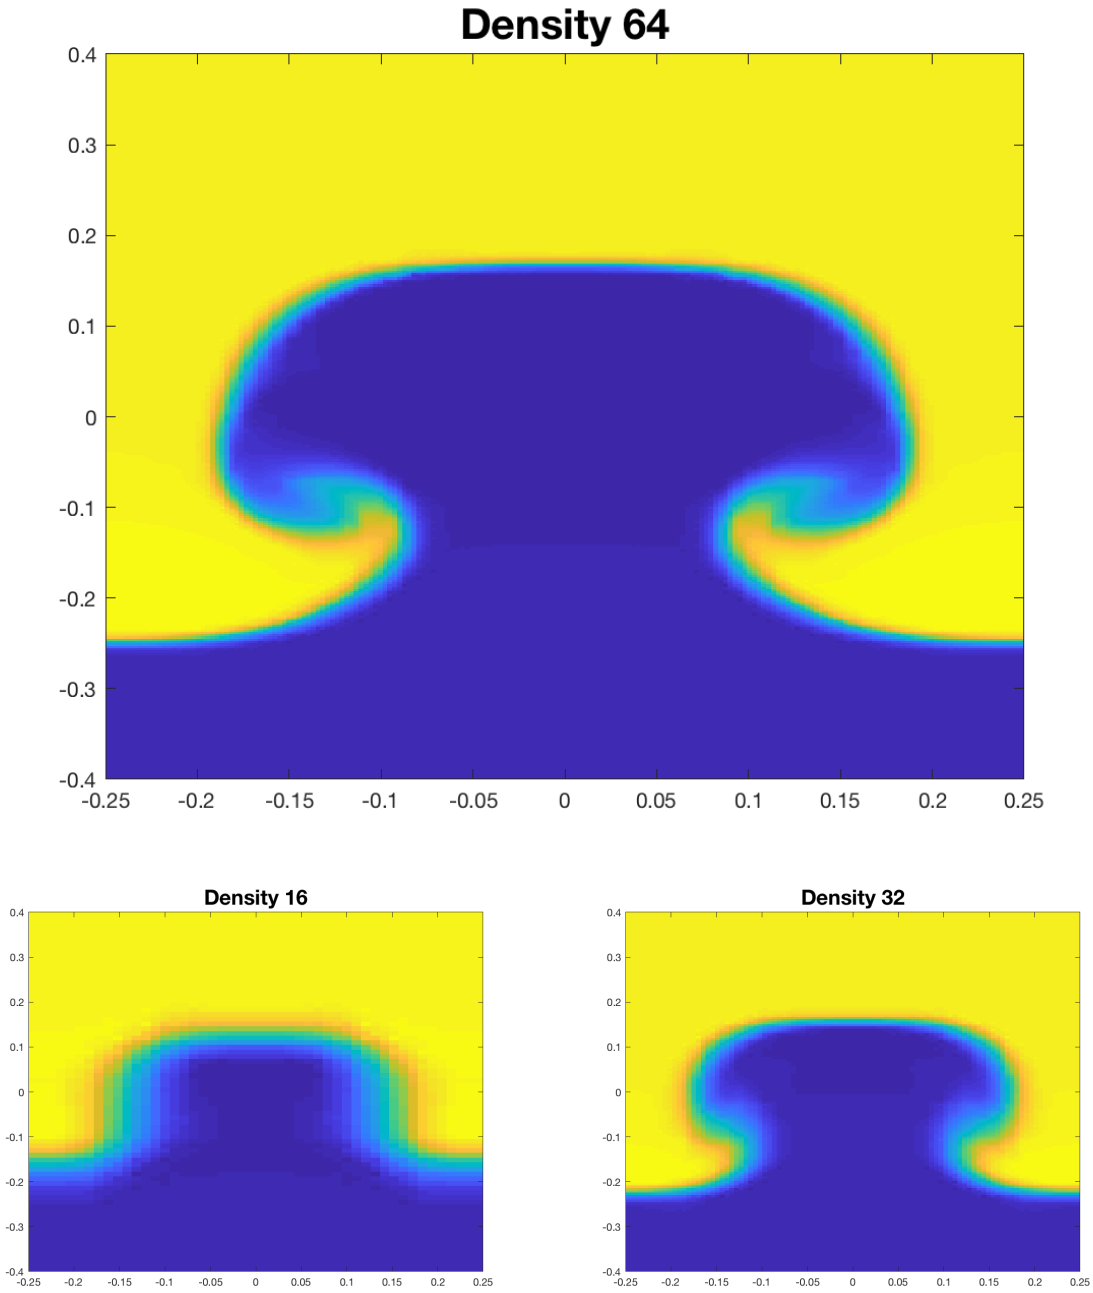

Figure 9: Here we see the Rayleigh Taylor Instability with decreasing element size. The top is 64 (wide) by 192 (tall) elements. The left is 16 by 48 elements. The right is 32 by 96 elements.

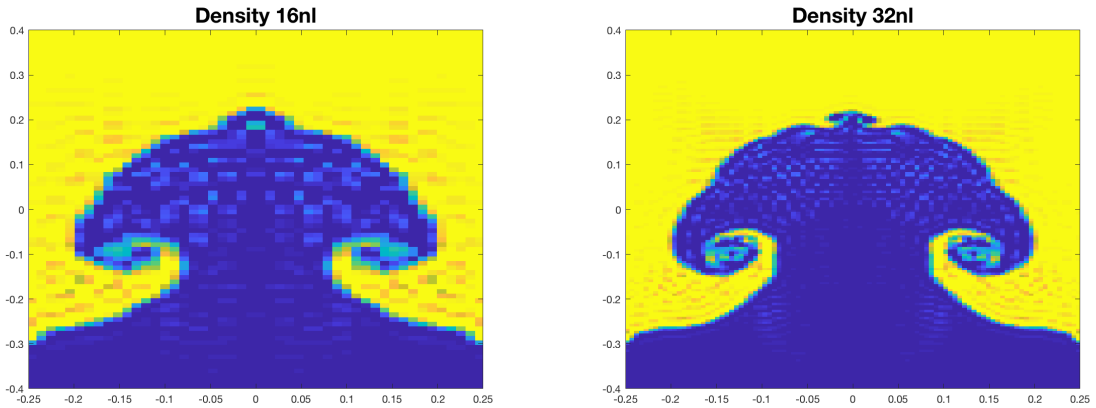

Figure 10: The Rayleigh Taylor Instabilty, but without limiting. The left is 16 by 48. The right is 32 by 96.

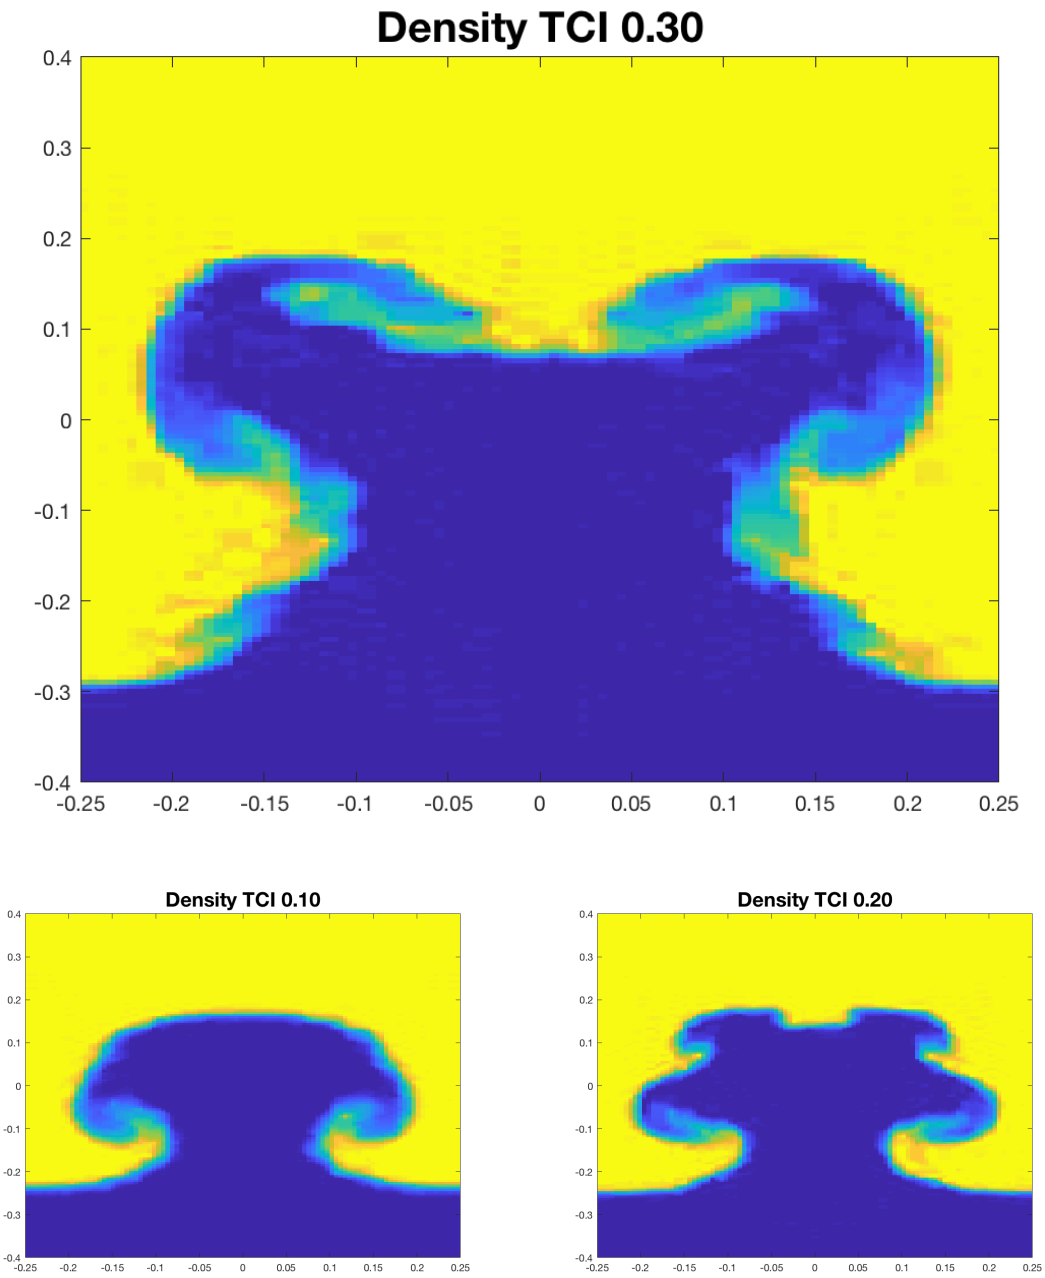

Figure 11: Here we increase the TCI cut-off. All tests are 32 by 96 elements. Top has a TCI cut-off of 0.3, left has TCI cut-off = 0.1, and right has a TCI cut-off of 0.2

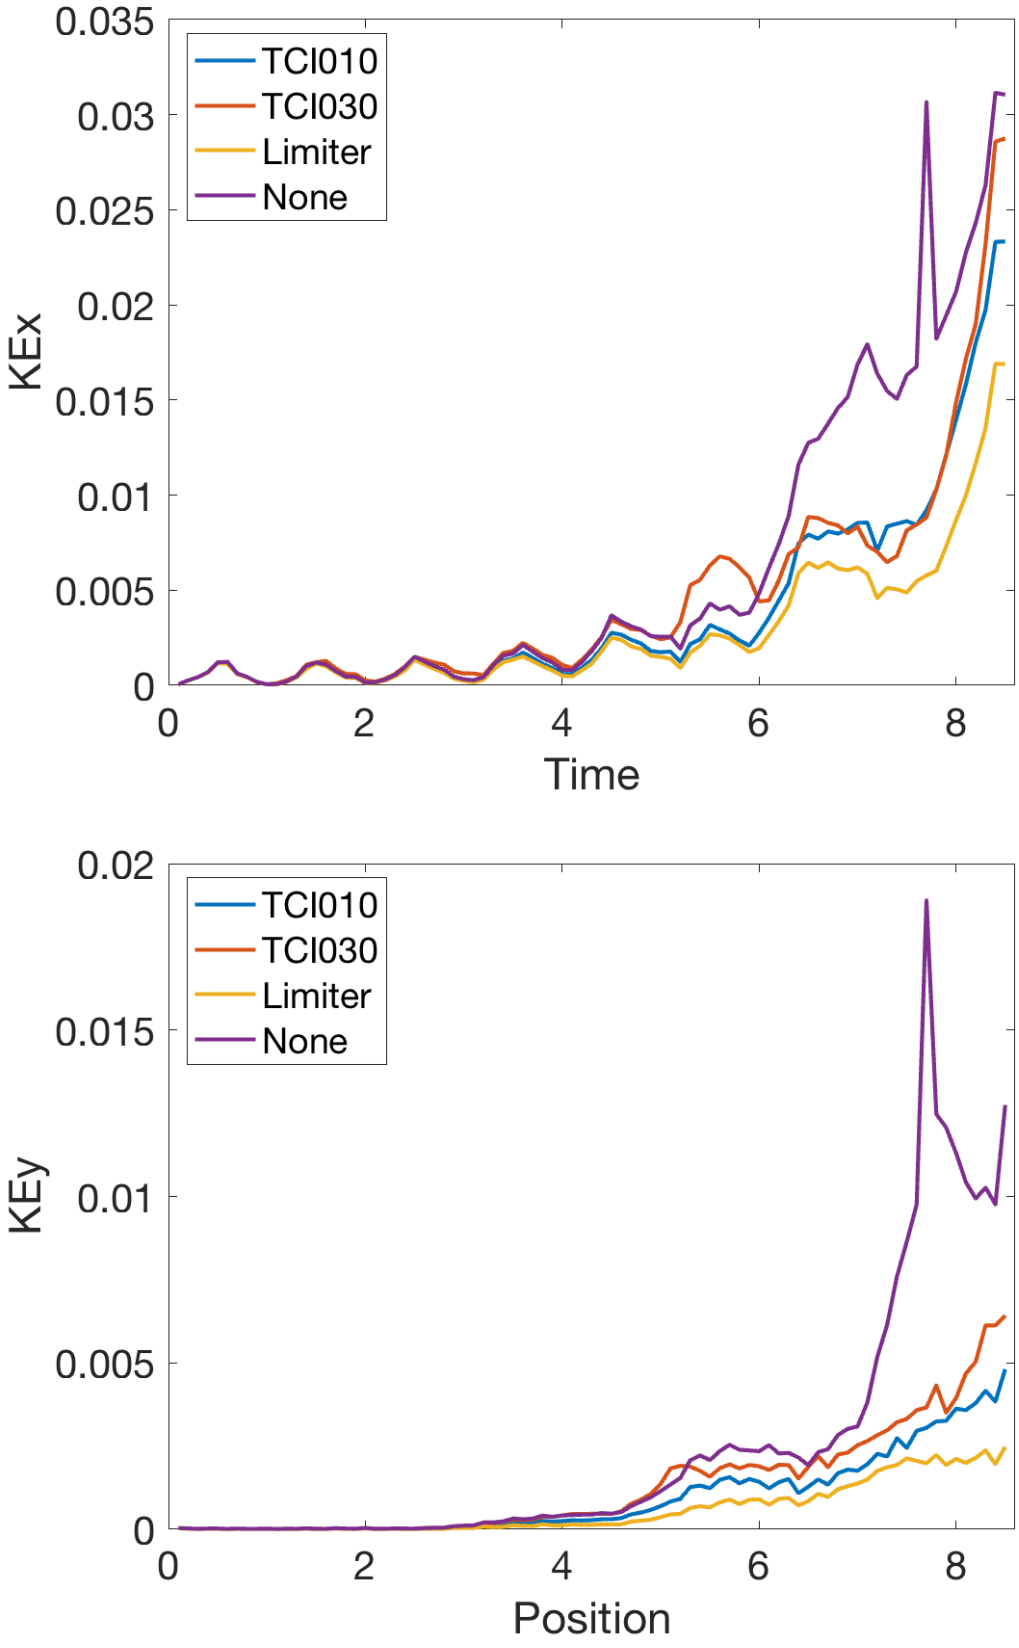

Figure 12: The kinetic energy of various tests.

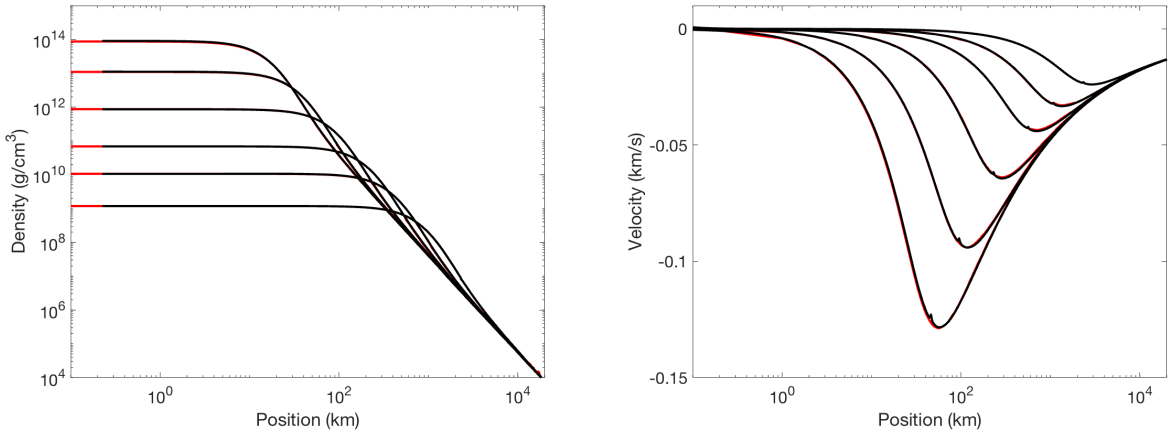

Figure 13: The simulation results in black with the data set reference in red. This simulation was run with 256 elements, outer boundary of 10^5 km and inner radius of 2 km. Notice how the bump in the velocity, an artifact of integration, is smoothed out in the simulation.

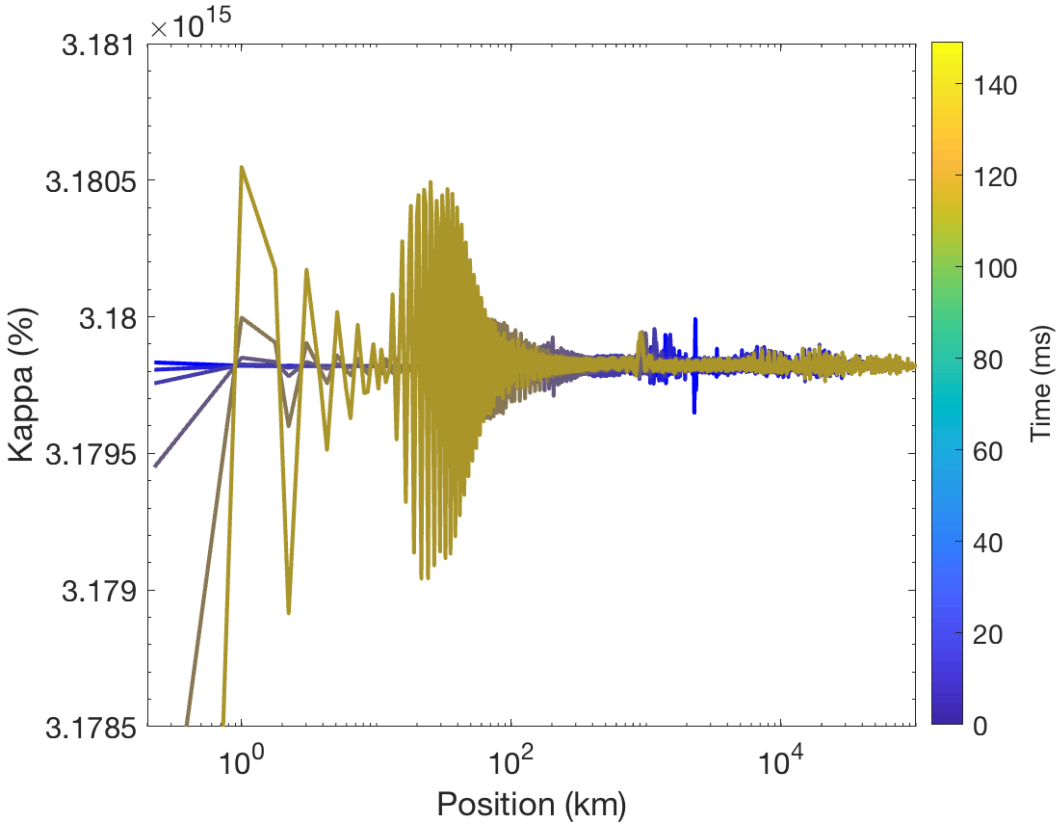

Figure 14: Here we see kappa at all positions at all times. Time is given by the color going from blue at early time to yellow at later times. $t = 150$ ms being the point of infinite density. This simulation was run with 256 elements, outer boundary of 10^5 km and inner radius of 2km.

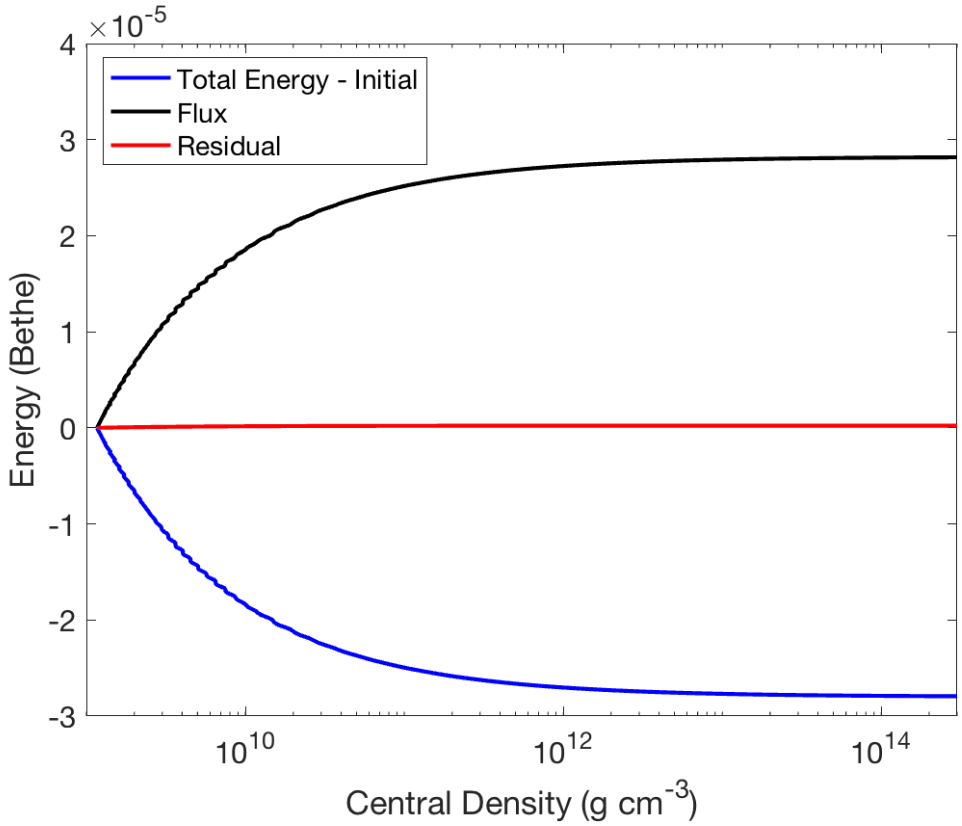

Figure 15: Here we see the summation of the flux compared to the total energy in the system minus the initial energy. From this, we can see how the flux generally accounts for the change in energy. This graph is generated from a simulation of 256 elements, outer boundary of 10^5 km, and inner radius of 2 km