Comparisons of Two- and Three-Dimensional Convection in Type I X-ray Bursts

M. Zingale[1], C. M. Malone[2,4], A. Nonaka[3], A. S. Almgren[3], J. B. Bell[3]

michael.zingale@stonybrook.edu

ABSTRACT

We perform the first detailed three-dimensional simulation of low Mach number convection preceding thermonuclear ignition in a mixed H/He X-ray burst. Our simulations include a moderate-sized, approximate network that captures hydrogen and helium burning up through rp-process breakout. We look in detail at the difference between two- and three-dimensional convective fields, including the details of the turbulent convection.

Subject headings: convection—hydrodynamics—methods: numerical—stars: neutron—X-rays: bursts

1. Introduction

X-ray bursts (XRBs) are the thermonuclear runaway in a H/He layer on the surface of a neutron star. These transient events can be used to probe the structure of neutron stars and the equation of state of dense material (Steiner et al. 2010; Özel et al. 2010). Furthermore, they are also the sites of rp-process nucleosynthesis (Schatz et al. 2001). For these reasons, understanding the dynamics of the explosion has seen substantial research interest in the past years.

One-dimensional studies (Taam 1980; Taam et al. 1996; Woosley et al. 2004) can reproduce the observed energies, durations, and recurrence timescales for XRBs, but use a parameterized model for convection, namely mixing length theory. An open question is whether

[1]Dept. of Physics & Astronomy, Stony Brook University, Stony Brook, NY 11794–3800

[2]CCS-2, Los Alamos National Laboratory, Los Alamos, NM 87545

[3]Center for Computational Sciences and Engineering, Lawrence Berkeley National Laboratory, Berkeley, CA 94720

[4]Nicholas C. Metropolis Fellow

a fully-turbulent convective velocity field can modify the nucleosynthesis. Additionally, the convection may dredge up heavy element ash to the photosphere (in 't Zand & Weinberg 2010; Bhattacharyya et al. 2010) thereby altering the opacity of the atmosphere, which affects the inference of neutron star mass and radius. These are inherently three-dimensional problems.

Previously, we performed two-dimensional simulations, focusing on pure He bursts first (Malone et al. 2011), and later, mixed H/He bursts (Malone et al. 2014). The latter study used an approximate network to capture the hot-CNO, triple-α, and initial rp-process breakout burning. There we found that we needed a spatial resolution of about 6 cm zone^{-1} in order to accurately model the burning. In this Letter, we extend our studies by performing the first three-dimensional model of convective burning in a H/He XRB, using the reaction network from Malone et al. (2014). This initial study compares to our two-dimensional results, and discusses the computational requirements for a more extensive study.

2. Numerical Method

We use the publicly-available[1] MAESTRO code (Nonaka et al. 2010), which solves the equations of low Mach number hydrodynamics by reformulating the reactive Euler equations to filter soundwaves while retaining compressibility effects due to stratification and local heat release. By filtering dynamically unimportant soundwaves, MAESTRO enables efficient simulation of slow convective flows, such as those in XRBs (Malone et al. 2011, 2014), various progenitors of Type Ia supernovae (Zingale et al. 2011; Nonaka et al. 2012; Zingale et al. 2013), and in the cores of massive stars (Gilet et al. 2013). Also important for simulations like these is that the low Mach number formulation analytically enforces hydrostatic equilibrium of the base state, allowing us to maintain a hydrostatic atmosphere in the simulation code without the development of large spurious velocities (see, e.g., Zingale et al. 2002).

All of the MAESTRO options and microphysics used in our two-dimensional study of XRBs in Malone et al. (2014) are retained for this study. In particular, we use the new energy formulation variant of MAESTRO, based on the ideas in Klein & Pauluis (2012) and Vasil et al. (2013), which improves energy conservation and our treatment of gravity waves. Our reaction network contains 10 species, approximating hot CNO, triple-α, and rp-breakout burning up through ^{56}Ni, using the ideas from Wallace & Woosley (1981), but with modern reaction rates from ReacLib (Cyburt et al. 2010) where available (see the discussion in Malone et al. 2014 for more details). We use the Helmholtz equation of state (EOS) from Timmes &

[1]MAESTRO can be obtained from `http://bender.astro.sunysb.edu/Maestro/`

Swesty (2000), which includes an ideal gas of nuclei, a photon gas, and an electron/positron gas with arbitrary degeneracy and relativistic parameters, and Coulomb corrections.

We use the same parametrized initial model as in our two-dimensional study. Briefly, the model consists of a 1.4 $M_\odot$, 10 km neutron star, of which we model the outer $\sim 1.4 \times 10^3$ cm as an isothermal ($T = 3 \times 10^8$ K), pure ^{56}Ni gas. On top of the neutron star is a warm accreted layer of mainly H/He fuel that is slightly metal-rich compared to solar, with CNO metals tied up in ^{14}O and ^{15}O in a ratio comparable to their respective β-decay lifetimes. A smooth transition is applied between the density ($\rho = 2 \times 10^6$ g cm^{-3}) and temperature ($T = 9.5 \times 10^8$ K) at the base of the accreted layer and the surface of the neutron star. The accreted layer is given an isentropic profile, making it convectively unstable, and the temperature decreases until a cutoff temperature is reached. The original extent of the convective region is $\lesssim 2 \times 10^3$ cm. The reader is referred to the Appendix of Malone et al. (2014) for more details of our model construction procedure.

In this Letter, we perform a single three-dimensional calculation to assess the dynamics of the convective flow. We model the XRB using a plane-parallel geometry on a uniform grid of $256 \times 256 \times 768$ zones, with 6 cm zone^{-1} spatial resolution—the same resolution used in our two-dimensional study. As the simulation evolves, the one-dimensional hydrostatic base state that MAESTRO carries is allowed to expand due to the heating, following the procedure described in Almgren et al. (2006).

3. Results

In order to understand how dimensionality affects our results, we compare to the two-dimensional calculations from Malone et al. (2014). In particular, we use the 6 cm, standard-width calculation from that paper. Figure 1 shows the standard deviation of temperature (compared to other zones at the same height) as a function of height for the two- and three-dimensional runs, both at $t = 0.04$ s. The overall trend is the same for the two calculations, with the magnitude of the temperature fluctuations in the convective region ($\sim$ 1400 cm to 3550 cm) $\delta T/\langle T \rangle \sim 10^{-3}$ to 10^{-4}.

Figure 2 shows the peak temperature and peak Mach number as a function of time for the two runs. We see that they closely track one another, but that in the three-dimensional simulation the peak Mach number is slightly higher on average than the two-dimensional case. At the start of the calculation, there is always a period of transient behavior as the heating needs to set up a consistent convective velocity field, but the flow quickly settles down. For both simulations, the average Mach number after the transient is less than 0.1; in

the longer-duration two-dimensional case, the Mach number asymptotes to 0.05. We did not run the three-dimensional calculation as long as the two-dimensional calculation, to conserve computational resources.

The convective velocity structure of the three-dimensional simulation is shown in Figure 3, highlighting the vertical velocity. These four images are representative of the flow throughout the simulation. We do not see the tight layering that was apparent in the two-dimensional simulations (especially for narrower domains; see Figures 6 and 7, and the discussion in Section 4.2.1 of Malone et al. 2014). To better understand the difference in the nature of the convective flow, we need to examine the turbulent structure.

Turbulence is known to behave differently between two and three dimensions (see, e.g. Ouellette 2012). To get a feel for the turbulent nature of the convection in these simulations, we look at the kinetic energy power spectrum. Following the discussion regarding turbulence in stratified flows in Nonaka et al. (2012) and references therein, we calculate a generalized kinetic energy density spectrum as

$$E_n(k) = \frac{1}{\Omega} \int_{\boldsymbol{S}(\boldsymbol{k})} \frac{1}{2} \widehat{\boldsymbol{V}}_n(\boldsymbol{k}) \cdot \widehat{\boldsymbol{V}}_n^{\star}(\boldsymbol{k}) d\boldsymbol{S}, \tag{1}$$

where $\widehat{\boldsymbol{V}}_n(\boldsymbol{k})$ is the Fourier transform of $\boldsymbol{V}_n(\boldsymbol{x}) = \rho^n(\boldsymbol{x})\widetilde{\boldsymbol{U}}(\boldsymbol{x})$ with n specifying the density weighting, $\boldsymbol{S}(\boldsymbol{k})$ is the surface defined by $|\boldsymbol{k}| = k$, and the $\star$ denotes complex conjugation. We note that here we use $\widetilde{\boldsymbol{U}}$, the local velocity on the grid, instead of explicitly calculating the turbulent velocity fluctuations from the full velocity field, including the base state expansion, $\boldsymbol{U} = \widetilde{\boldsymbol{U}} + w_0 \boldsymbol{e}_r$, because $\widetilde{\boldsymbol{U}}$ is essentially the velocity perturbations on top of an otherwise hydrostatic background state (Nonaka et al. 2010). The volume, Ω, and surface element, $d\boldsymbol{S}$, are based on the dimensionality of the problem. The goal is to find the proper scaling of the energy density spectrum with wavenumber for both two- and three-dimensional flow.

The units of $\widehat{\boldsymbol{V}}_n$ are [g^n cm^{1+D-3n} s^{-1}], where the extra power of D on the length scale comes from the integral over $d\boldsymbol{x}$ in the definition of the Fourier transform. In Equation (1), the integral is done in k-space, such that $d\boldsymbol{S} \sim d^{D-1}k$ with units [cm^{1-D}], whereas the normalization is in real-space, so that Ω has units of [cm^D]. Upon integration of Equation (1), the dimensionality, D, drops out of the equation, and the units of the generalized kinetic energy density spectrum become [g^{2n} cm^{3-6n} s^{-2}] for both two- and three-dimensional configurations. For turbulent flows that have density variation (i.e. compressible or stratified flows), the typical Kolmogorov energy dissipation rate, $\epsilon(l)$, at a given length scale l should be weighted by the mass density (see Fleck 1983, 1996, for example): $\epsilon(l) = \rho U^3(l)/l$, which has units of [g cm^{-1} s^{-1}]. The arguments of Nonaka et al. (2012) then apply to any dimension: the only combination of $\epsilon^\alpha k^\beta E_n(k)$ that yields a dimensionless quantity is when $\alpha = -2/3$, $n = 1/3$, and $\beta = 5/3$. If the physics of two dimensional and three dimensional

turbulence were the same (this is likely not the case), then the spectrum defined in Eq. (1) should scale as $k^{-5/3}$ for both two- and three-dimensional flows.

In evaluating Eq. (1), we create equally-spaced radial wavenumber bins, k_i, ranging from the smallest physical wavenumber, $1/L$, to the highest meaningful wavenumber, $1/(2\Delta x)$, where L is the domain width. The Fourier transform of the kinetic energy density gives us

$$\hat{K}(k_x, k_y, k_z) = \frac{1}{2}(\widehat{\boldsymbol{V}}_n \cdot \widehat{\boldsymbol{V}}_n^{\star}) \tag{2}$$

For each of points in the three-dimensional $\hat{K}$ array, we define $|k| = \sqrt{k_x^2 + k_y^2 + k_z^2}$ and determine which of the radial bins, k_i, this falled into and add the value of $\hat{K}$ to that bin's sum. Done this way, we are integrating up in spherical shells in k-space, using our discrete bins. The same procedure is done in two dimensions, but now we are working in the k_x-k_y plane, and are integrating up over annular regions in that plane, again defined by our discrete bins, k_i. We do not worry about the $1/\Omega$ normalization, since we will normalize each spectrum such that its peak value is 1.

Figure 4 shows the power spectrum of the two-dimensional and three-dimensional XRB simulations at $t = 0.04$ s. For this analysis, we restrict the domain to just the vicinity of the convective region, including only the vertical range 1300 cm $< z <$ 3550 cm. For the three-dimensional case, we see that we have about a decade in wavenumber where we achieve a $k^{-5/3}$ power-law scaling, indicative of Kolmogorov turbulence. We note that the region we are studying is not periodic in the vertical direction, but an FFT assumes periodicity, so the discontinuity through the vertical boundary may affect the behavior at high wavenumbers, perhaps accounting for the slow fall in the 3-d spectrum. For the two-dimensional case, there does not appear to be a $k^{-5/3}$ scaling; the spectrum starts off with a shallower slope but then becomes steeper at moderate wavenumbers. Such a break in the power law scaling for two-dimensional turbulence is predicted for very idealized turbulence where the steeper part of the curve has a k^{-3} scaling attributed to a cascade of enstrophy (e.g. Kraichnan 1967; Leith 1968; Batchelor 1969). Numerical simulation cannot achieve the idealized conditions (e.g. infinte domain and infinite Reynolds number) assumed in the k^{-3} derivation, and sometimes achieve a steeper power law (e.g. the review by Gkioulekas & Tung 2006, and references therein). In our 2-d simulation, we see a small range in the spectrum after the break where we look consistent with k^{-3}.

We have also seen such a difference in scaling between two- and three-dimensional turbulence on smaller scales in reactive Rayleigh-Taylor simulations (Zingale et al. 2005), where we saw a spectrum that appeared to follow the $k^{-11/5}$ scaling predicted by Bolgiano-Obukhov statistics for a 2-d cascade (Niemeyer & Kerstein 1997). In that study, we found that a wide domain, giving more statistics, was essential to see this scaling. The difference in the scaling

we observe in the present simulations suggests that there is a fundamental difference in how the cascade takes place between the two- and three-dimensional convection in XRBs.

Figure 4 also shows that there is relatively more power in small scale (higher wavenumber, k) features for the three-dimensional simulation compared to the two-dimensional calculation. This is made more explicit by looking at a colormap plot of the enstrophy density $\eta = \left|\nabla \times \widetilde{\boldsymbol{U}}\right|^2/2$, as is shown in Figure 5, where the left (right) panel shows the two-dimensional (three-dimensional) simulation. For the three-dimensional simulation, the plot shows a slice through the center of the domain. The two-dimensional simulation plot appears to be dominated by moderate-sized vortices throughout the domain, while in three dimensions, we see structure on a much wider range of scales. This is similar to the results seen in comparisons of two- and three-dimensional simulations of novae (Kercek et al. 1998, 1999), although our two-dimensional results do not show as severe of a dominance of vortices as reported there.

The panels of Figure 5 also show that in two-dimensional flow the convective motions penetrate deeper into the underlying neutron star than in the three-dimensional case. This can have implications for the amount of metal-rich material that can be dredged up into the atmosphere, potentially polluting the photosphere and adjusting the opacity. We attempt to quantify the difference between the two simulations by calculating the total amount of ^{56}Ni in the atmosphere above $r = 10^3$ cm. For the two dimensional case (if extended in the third dimension to the size of the three-dimensional simulation) we have a total mass of $m(^{56}\mathrm{Ni}) = 9.81 \times 10^{15}$ g at $t = 0.04$ s; the three-dimensional simulation has $m(^{56}\mathrm{Ni}) = 9.80 \times 10^{15}$ g at the same time. The difference of this quantity between our simulations is quite small, however there are a couple of caveats. First, the convection is just getting started and it is unclear if this difference will increase with time. Second, we have not done a detailed subtraction of the amount of ^{56}Ni *produced* from nuclear burning within the region considered. Given that the flow structure is so different between the two simulations, it is unlikely that—even in an average sense—parcels of fluid in the three-dimensional simulation would have the same thermodynamic trajectory as in the two-dimensional case.

4. Discussion and Conclusions

We described the first three-dimensional model of convective burning in an XRB. While the peak temperature and Mach number behave qualitatively the same as our two-dimensional calculations, the structure of the convective velocity field differs substantially, both in the global appearance and in the turbulent statistics. This is illustrated well by the difference in appearance of the enstrophy density, which in three dimensions shows the typical cas-

cade to small scales. Since convective mixing is expected to distribute the synthesized nuclei throughout the atmosphere, potentially bringing some to the photosphere, modeling the convection accurately is important. Based on the differences seen between the two-dimensional and three-dimensional flows, this suggests that three-dimensional models should be the focus of our future simulation efforts.

The calculation presented here paves the way for a more detailed study of convective burning in XRBs. This calculation used a modest 400,000 CPU-hours on the NERSC Edison system (running with 192 MPI tasks and 6 OpenMP threads per task). To better capture the randomness of the convective process, we need a larger domain—preferably twice the size in each lateral dimension. At this same resolution, this would bring the computation cost to 1.6 M CPU-hours. While MAESTRO can use AMR, in these simulations we would refine the entire convective region, so the cost savings would be small.

Our future calculations will push for increased realism of the reaction network. As detailed in Malone et al. (2014), the approximate network used here reasonably captures the overall energy release, but displays odd features (namely convective layering in the longer-duration two-dimensional simulations) in the temperature and density regime where the $^{18}\mathrm{Ne}(,\mathrm{e}^{+}\nu)^{18}\mathrm{F}$ rate is comparable to the $^{18}\mathrm{Ne}(\alpha,p)^{21}\mathrm{Na}$ rate for hot CNO burning-breakout. We plan to both improve the nuclear reaction network with a more clever selection of isotopes for an approximate network, and to investigate using larger networks whose integration can be accelerated using highly parallel hardware accelerators, such as GPUs or Intel Xeon Phi processors.

We also plan to push our calculations to larger scales. For the near future, however, these sort of calculations will be limited to convection-in-a-box studies. Capturing the range of length scales necessary to follow a laterally propagating burning front, while resolving the energy-generation region, is not currently possible. The complementary approach to ours are the calculations by Cavecchi et al. (2012), which used wide-aspect ratio zones and did not perform hydrodynamics in the vertical direction. Ultimately these two methods can inform one-another to build a picture of nucleosynthesis and dynamics of the burning front in XRBs.

Visualization was done with yt (Turk et al. 2011). The power spectrum calculation followed the "Making a Turbulent Kinetic Energy Power Spectrum" recipe in the yt Cookbook. The git-hashes of the codes used are MAESTRO: `d82b99255277...` and BoxLib: `98426666046f...`. The reaction network is distributed with MAESTRO in `MAESTRO/Microphysics/network`

We thank Frank Timmes for making his equation of state publicly available. The work at Stony Brook was supported by DOE/Office of Nuclear Physics grants No. DE-FG02-

06ER41448 and DE-FG02-87ER40317 to Stony Brook. Work at LANL was done under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396. The work at LBNL was supported by the Applied Mathematics Program of the DOE Office of Advance Scientific Computing Research under U.S. Department of Energy under contract No. DE-AC02-05CH11231. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

REFERENCES

Almgren, A. S., Bell, J. B., Rendleman, C. A., & Zingale, M. 2006, ApJ, 649, 927, paper II

Batchelor, G. K. 1969, Phys. Fluids, 12, II

Bhattacharyya, S., Miller, M. C., & Galloway, D. K. 2010, 401, 2, 0908.4245

Cavecchi, Y., Watts, A. L., Braithwaite, J., & Levin, Y. 2012, ArXiv e-prints, 1212.2872

Cyburt, R. H., et al. 2010, ApJS, 189, 240

Fleck, Jr., R. C. 1983, ApJ, 272, L45

—. 1996, ApJ, 458, 739

Gilet, C., Almgren, A. S., Bell, J. B., Nonaka, A., Woosley, S. E., & Zingale, M. 2013, ApJ, 773, 137

Gkioulekas, E., & Tung, K. 2006, Journal of Low Temperature Physics, 145, 25

in 't Zand, J. J. M., & Weinberg, N. N. 2010, Astronomy and Astrophysics, 520, A81+, 1001.0900

Kercek, A., Hillebrandt, W., & Truran, J. W. 1998, A&A, 337, 379, astro-ph/9801054

—. 1999, A&A, 345, 831, astro-ph/9811259

Klein, R., & Pauluis, O. 2012, Journal of Atmospheric Sciences, 69, 961

Kraichnan, R. H. 1967, Phys. Fluids, 10, 1417

Leith, C. E. 1968, Phys. Fluids, 11, 671

Malone, C. M., Nonaka, A., Almgren, A. S., Bell, J. B., & Zingale, M. 2011, ApJ, 728, 118, 1012.0609

Malone, C. M., Zingale, M., Nonaka, A., Almgren, A. S., & Bell, J. B. 2014, ApJ, 788, 115

Niemeyer, J. C., & Kerstein, A. R. 1997, New A, 2, 239

Nonaka, A., Almgren, A. S., Bell, J. B., Lijewski, M. J., Malone, C. M., & Zingale, M. 2010, ApJS, 188, 358, paper V

Nonaka, A., Aspden, A. J., Zingale, M., Almgren, A. S., Bell, J. B., & Woosley, S. E. 2012, ApJ, 745, 73, 1111.3086

Ouellette, N. T. 2012, Physics Today, 65, 68

Özel, F., Baym, G., & Güver, T. 2010, Phys. Rev. D, 82, 101301, 1002.3153

Schatz, H., et al. 2001, Physical Review Letters, 86, 3471, arXiv:astro-ph/0102418

Steiner, A. W., Lattimer, J. M., & Brown, E. F. 2010, Astrophysical Journal, 722, 33, 1005.0811

Taam, R. E. 1980, Astrophysical Journal, 241, 358

Taam, R. E., Woosley, S. E., & Lamb, D. Q. 1996, Astrophysical Journal, 459, 271

Timmes, F. X., & Swesty, F. D. 2000, ApJS, 126, 501, source code obtained from `http://cococubed.asu.edu/codes/eos.shtml/helmholtz.tbz`

Turk, M. J., Smith, B. D., Oishi, J. S., Skory, S., Skillman, S. W., Abel, T., & Norman, M. L. 2011, ApJS, 192, 9, 1011.3514

Vasil, G. M., Lecoanet, D., Brown, B. P., Wood, T. S., & Zweibel, E. G. 2013, ArXiv e-prints, 1303.0005

Wallace, R. K., & Woosley, S. E. 1981, ApJS, 45, 389

Woosley, S. E., et al. 2004, Astrophysical Journal Supplement, 151, 75

Zingale, M., Nonaka, A., Almgren, A. S., Bell, J. B., Malone, C. M., & Orvedahl, R. J. 2013, ApJ, 764, 97, 1212.4380

Zingale, M., Nonaka, A., Almgren, A. S., Bell, J. B., Malone, C. M., & Woosley, S. E. 2011, ApJ, 740, 8

Zingale, M., Woosley, S. E., Bell, J. B., Day, M. S., & Rendleman, C. A. 2005, Journal of Physics Conference Series, 16, 405, astro-ph/0507040

Zingale, M., et al. 2002, Astrophysical Journal Supplement, 143, 539

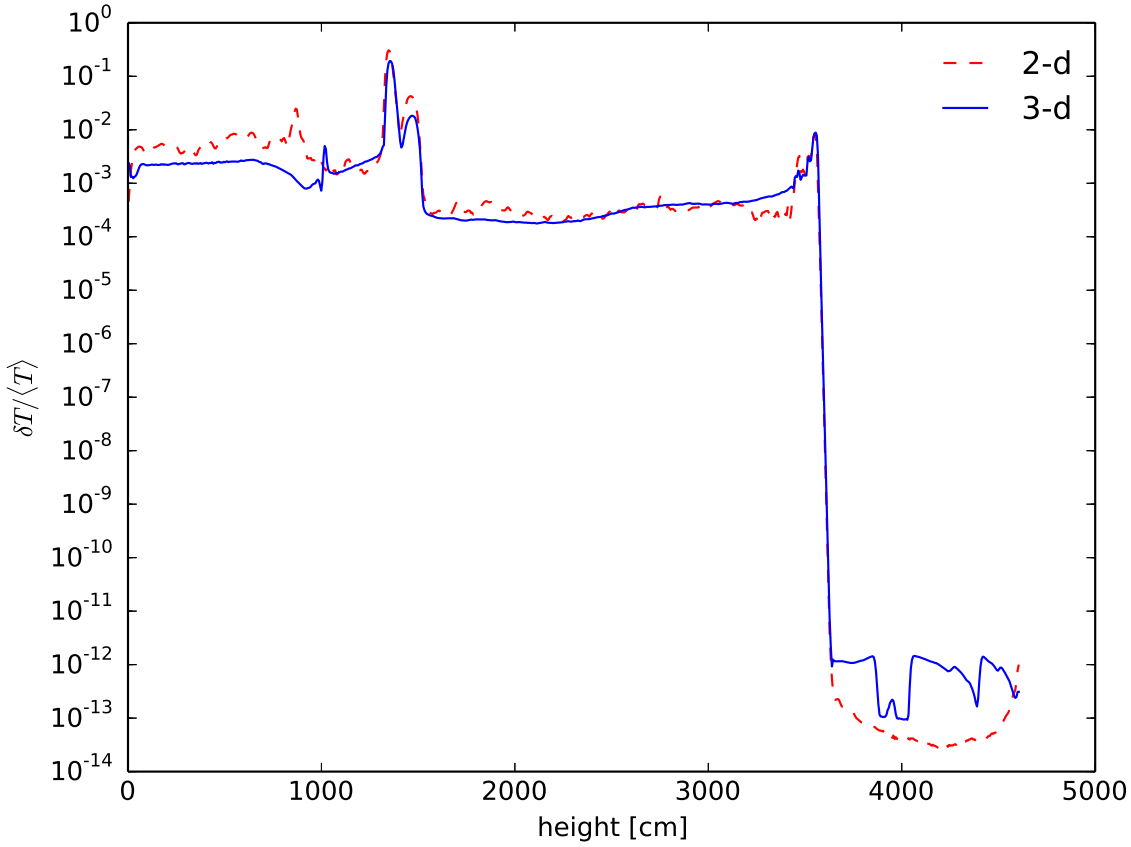

Fig. 1.— Variance of T normalized by its average at a given height for both our two-dimensional and three-dimensional simulations at $t = 0.04$ s. Within the convective region, $1600 \lesssim$ height $\lesssim 3200$, the temperature fluctuations between the two simulations are quite similar. We also note that at the edges of the convective region, due to overshoot/undershoot, there are local spikes in the average temperature fluctuations. Below the convective region, the two-dimensional simulation shows temperature fluctuations that are about four times larger than in the three-dimensional counterpart. This is likely due to the larger amount of convective undershoot present in the two-dimensional simulation compared to the three-dimensional simulation.

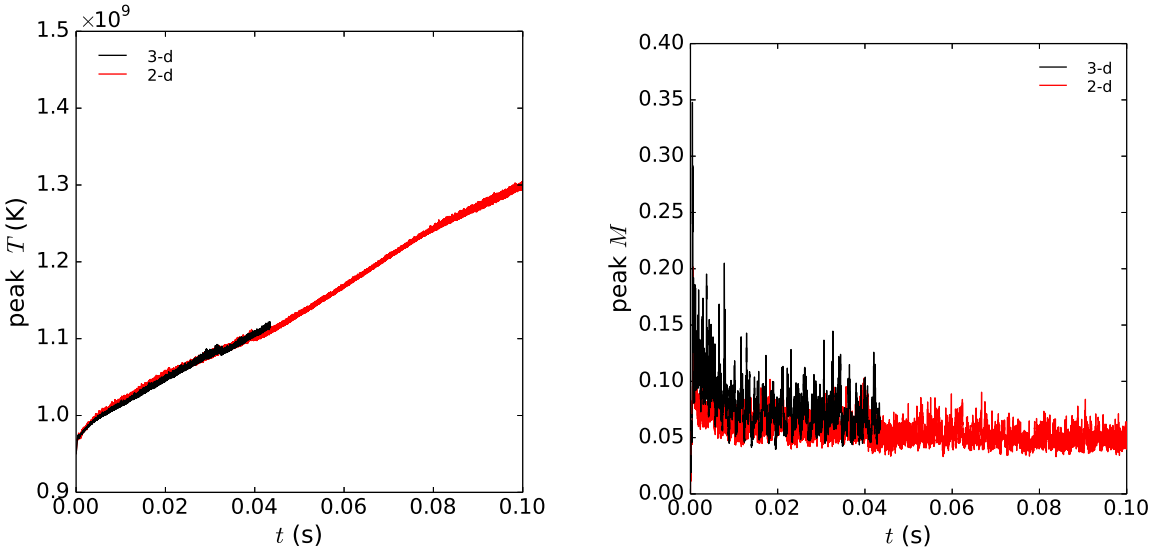

Fig. 2.— Comparison of the peak temperature vs. time between the two- and three-dimensional simulations (left) and the peak Mach number vs. time (right). Both simulations agree quite well in this context, however the three-dimensional simulation has slightly larger time-averaged peak Mach number. Both simulations experience an initial short-duration transient spike in Mach number as the system creates a convective flow field able to carry away the energy generated from nuclear reactions.

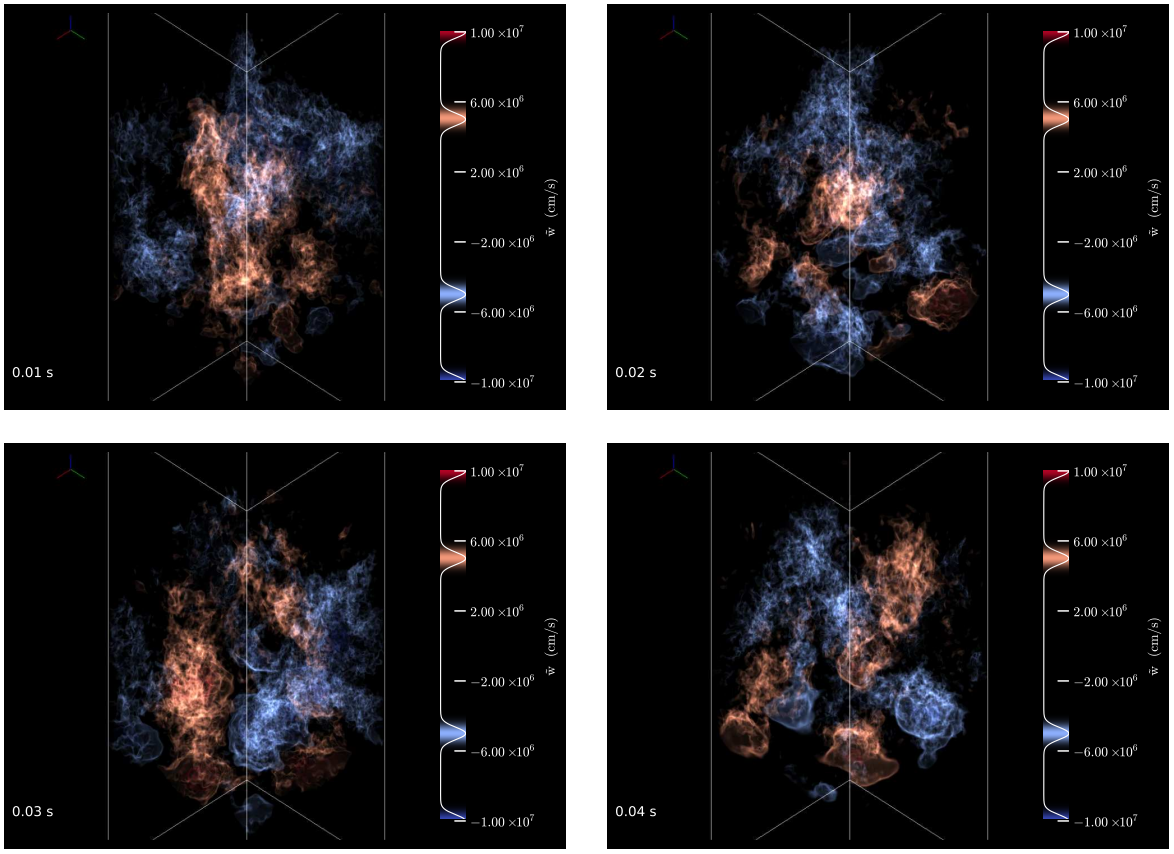

Fig. 3.— Volume renderings of the vertical velocity field at (from left to right, top to bottom) t =0.01, 0.02, 0.03, and 0.04 s. Upward moving fluid is in red and downward moving is blue.

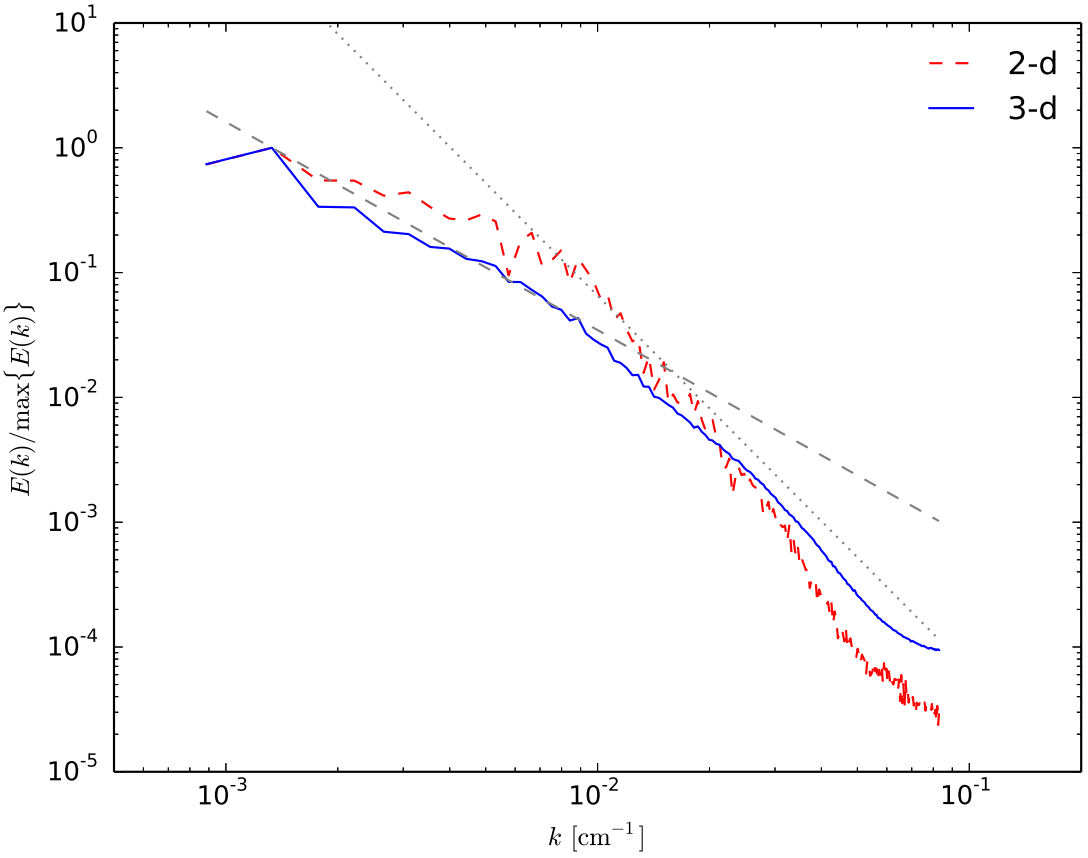

Fig. 4.— Kinetic energy power spectrum for the two- and three-dimensional simulations at $t = 0.04$ s. The dashed gray line is a $k^{-5/3}$ power-law and the dotted line is a k^{-3} power-law. A density weighting of $\rho^{1/3}$ was used for both two and three dimensions. The power is normalized so the two spectra have the same peak. There is about a decade in wavenumber where the three-dimensional simulation obeys the standard Kolmogorov turbulent cascade. The two-dimensional simulation displays a characteristic change in power-law scaling, having sections that are both shallower and steeper than what Kolmogorov predicts for three dimensions.

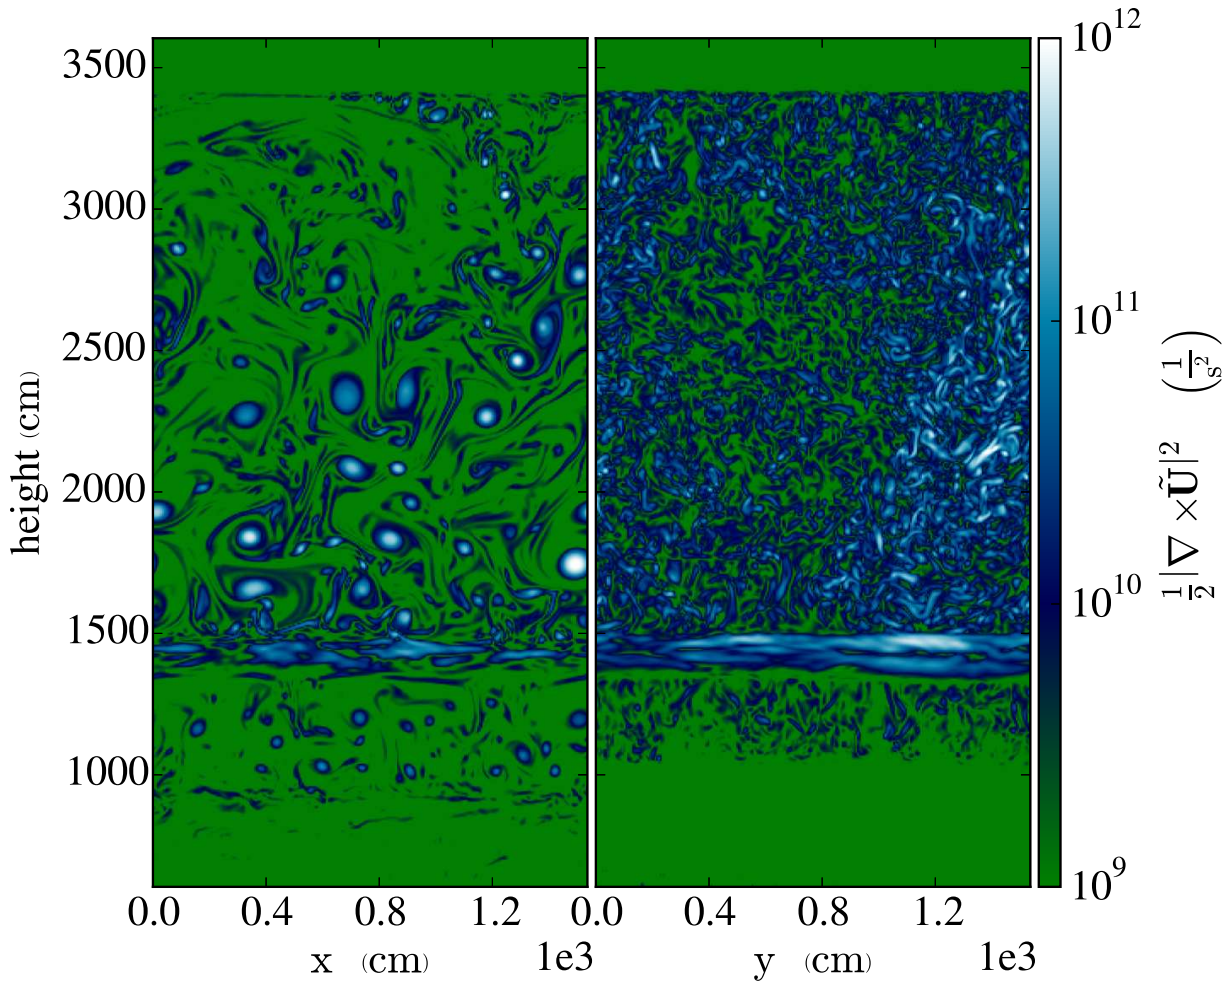

Fig. 5.— Enstrophy density of the turbulent flow in both the two-dimensional (left) and three-dimensional (right) simulations at $t = 0.04$ s. The plot for the three-dimensional simulation is a slice through the center of the domain. There are clear differences between two- and three-dimensional flows with the three-dimensional simulation showing much more small scale features. This is consistent with the presence of relatively more power at larger wavenumber for the three-dimensional case compared to the two-dimensional case, as seen in Figure 4.