The Astrophysical Journal, 798:49 (26pp), 2015 January 1 doi:10.1088/0004-637X/798/1/49
© 2015. The American Astronomical Society. All rights reserved.

HYDRODYNAMIC SIMULATIONS OF H ENTRAINMENT AT THE TOP OF He-SHELL FLASH CONVECTION

Paul R. Woodward[1], Falk Herwig[2], and Pei-Hung Lin[1]
[1] LCSE & Department of Astronomy, University of Minnesota, Minneapolis, MN 55455, USA; paul@lcse.umn.edu
[2] Department of Physics & Astronomy, University of Victoria, Victoria, BC V8P5C2, Canada; fherwig@uvic.ca
Received 2013 July 14; accepted 2014 August 12; published 2014 December 19

ABSTRACT

We present the first three-dimensional, fully compressible gas-dynamics simulations in 4π geometry of He-shell flash convection with proton-rich fuel entrainment at the upper boundary. This work is motivated by the insufficiently understood observed consequences of the H-ingestion flash in post-asymptotic giant branch (post-AGB) stars (Sakurai's object) and metal-poor AGB stars. Our investigation is focused on the entrainment process at the top convection boundary and on the subsequent advection of H-rich material into deeper layers, and we therefore ignore the burning of the proton-rich fuel in this study. We find that for our deep convection zone, coherent convective motions of near global scale appear to dominate the flow. At the top boundary convective shear flows are stable against Kelvin–Helmholtz instabilities. However, such shear instabilities are induced by the boundary-layer separation in large-scale, opposing flows. This links the global nature of thick shell convection with the entrainment process. We establish the quantitative dependence of the entrainment rate on grid resolution. With our numerical technique, simulations with 1024^3 cells or more are required to reach a numerical fidelity appropriate for this problem. However, only the result from the 1536^3 simulation provides a clear indication that we approach convergence with regard to the entrainment rate. Our results demonstrate that our method, which is described in detail, can provide quantitative results related to entrainment and convective boundary mixing in deep stellar interior environments with very stiff convective boundaries. For the representative case we study in detail, we find an entrainment rate of $4.38 \pm 1.48 \times 10^{-13}\ M_\odot\ \mathrm{s}^{-1}$.

Key words: convection – hydrodynamics – stars: AGB and post-AGB – stars: evolution – stars: interior – turbulence

1. INTRODUCTION

Convection in stars is the most important mixing process. It is essential for heat transport and mixing of elements. Therefore, convection is an essential ingredient in the evolution of stars and in the formation of the elements in stars and stellar explosions (Woosley et al. 2002; Herwig 2005). The treatment of convective boundaries is a key uncertainty in present stellar evolution calculations. In particular, mixing at the rather stiff (see below) boundaries of convection in the deep stellar interior has been related to numerous observational properties of stars. Determining reliable quantitative mixing properties of convection boundaries remains an unsolved problem. Overshooting at the boundary of H-core convection has been investigated for the past 40 yr (e.g., Maeder 1976; Schaller et al. 1992; Ventura et al. 1998; Deupree 2000; VandenBerg et al. 2006). Evidence for convective boundary mixing (CBM) at the bottom of the solar convection zone has been presented based on helioseismology (Christensen-Dalsgaard et al. 2011) and hydrodynamic simulations (e.g., Freytag et al. 1996; Rogers et al. 2006). For thermal-pulse asymptotic giant branch (AGB) stars CBM is important for the occurrence of the third dredge-up (for example Herwig et al. 1997; Herwig 2000; Mowlavi 1999; Weiss & Ferguson 2009), the reproduction of observed properties of AGB stars (Karakas et al. 2010), and in post-AGB stars (Herwig et al. 1999; Miller Bertolami et al. 2006; Werner & Herwig 2006). More recently it has been shown that CBM can explain the observed enrichment of nova ejecta (Denissenkov et al. 2012a). Furthermore, CBM at the bottom of C- and Ne-burning shells in super-AGB stars and transition mass objects is critical for the progenitors of supernova Ia and possible core-collapse supernova (Denissenkov et al. 2013; Jones et al. 2013).

In this paper we present results for simulations of He-shell flash convection in AGB stars, including CBM at the upper boundary. We would like to investigate the global properties of convective flows in this type of convection, and the CBM in general for a case in which the boundary is very stiff. Specifically we focus here upon the entrainment of hydrogen-rich material into the convection zone above a helium burning shell. Such situations have been addressed by Mocák et al. (2010). Mocák et al. (2011) and Stancliffe et al. (2011), and similar entrainment events by the convection zone above an oxygen burning shell and the H-burning core convection have been treated by Meakin & Arnett (2007). Viallet et al. (2013) also discuss such simulation results with a goal of producing improved mixing length models for use in one-dimensional (1D) stellar evolution simulations. Our own work has to this point focused on the entrainment problem (Herwig et al. 2006, 2007; Woodward et al. 2008a, 2008b), separated from the complicating factors of nuclear burning of the ingested fuel. We wish to establish that our computational methods are capable of accurately simulating the entrainment process before we add in the nuclear burning and back reaction on the flow of the energy liberated in that process. Nevertheless, we have produced a detailed analysis using 1D stellar evolution techniques with a parameterized treatment of CBM that indicates the dynamical behavior that a proper three-dimensional (3D) analysis of hydrogen ingestion and burning is likely to yield in the specific, well observed case of Sakurai's object (Herwig et al. 2011). A followup to this article will present results obtained using our PPMstar code for that particular case (see Herwig et al. 2014 for preliminary results).

The present paper is intended to describe the challenges that CBM poses to numerical simulation, set out in the context of

a representative hydrogen ingestion problem. The upper convective boundary of a He-shell flash convection zone about to make contact with the H-rich layer above is much stiffer compared, for example, to the cases investigated by Meakin & Arnett (2007), and consequently we find a much smaller entrainment rate in this situation. We describe how we meet the challenge of converged entrainment simulations for a very stiff convective boundary by a combination of advanced numerical techniques and high grid resolution. We present evidence that for the aspect of these problems that we judge to be the most critical—the hydrogen entrainment rate at the upper convective boundary—our simulations are able to produce results that converge upon grid refinement. Because this ingestion is expected to have dramatic effects upon the star by releasing more energy than the shell-burning source producing the convection that causes it, it is very important that it be computed accurately. We cannot comment upon the accuracy of simulations by others, but we note that the results that are obtained vary greatly. In work by others, the conditions of the CBM differ substantially from our own and from each other, so that direct comparisons are not possible. Nevertheless, we have observed in our own work that the all-important entrainment rate depends strongly upon the grid resolution until a sufficiently fine grid is used. All investigators have of course reported results obtained on different grids. Stancliffe et al. (2011) find luminosities that differ by three orders of magnitude for simulations of the same problem performed on grids differing in linear resolution by a factor of two. Mocák et al. (2011) present results for only a single 3D grid. It is arranged in a 45° wedge geometry that precludes the development of large-scale or global modes. They establish their grid requirement from simulations performed at different resolutions in two dimensions. Viallet et al. (2013) present results for a series of simulations on progressively finer grids, which cover a sector of the convection zone above an oxygen burning shell. They state that they do not observe convergence under grid refinement in the mean fields in the narrow region of steep gradients at the base of their oxygen burning shell. Because we will base our subsequent studies of hydrogen ingestion flash events upon our PPMstar code, discussed here, and because previously reported results establish the difficulty of accurate simulation of such events, our goals in this paper are twofold. In addition to presenting our quantitative results on entrainment for our particular problem and the general properties of He-shell flash convection in 4π geometry, we set out in this paper our computational techniques and the evidence that they are equal to the task we have set for ourselves in our subsequent work.

Regions in stars are convectively unstable if the radiative energy transport is less efficient than convective transport (Kippenhahn & Weigert 1990). The Schwarzschild condition determines instability against convection if the radiative temperature gradient is larger than the adiabatic temperature gradient. The radiative temperature gradient is given by

$$-\left.\frac{dT}{dr}\right|_{\rm rad} = \frac{3}{16\pi ac}\frac{\kappa L\rho}{r^2T^3}, \tag{1}$$

where κ is the Rosseland mean opacity or the electron conduction coefficient, L is the luminosity, T the temperature, r the radius, ρ the density, a is the radiation-density constant, and c the speed of light.

Equation (1) shows that convective instability can be caused by either large opacity or by large luminosity (see, for example, Figure 3 in Paxton et al. 2011). The former is the case for surface convection, for example in the shallow surface convection of A-type stars and white dwarfs (Freytag et al. 1996), in the sun (e.g., Stein & Nordlund 1998; Miesch et al. 2000; Robinson et al. 2003; Miesch et al. 2008), or in the deep convective envelopes of giant stars (Porter & Woodward 2000; Porter et al. 2000; Freytag & Höfner 2008). Convection in the deep interior is usually driven by high luminosity, as for example in He-shell flashes in thermal pulse AGB stars (Herwig et al. 2006, 2007), in oxygen and carbon shell burning in the advanced evolution phases of massive stars (Asida & Arnett 2000; Meakin & Arnett 2007) or in He-core flashes (Mocák et al. 2008, 2009).

Convection in the deep interior at high densities is usually very efficient and the temperature gradient is therefore nearly adiabatic. However, in surface convection, where in some cases a significant fraction of the energy is transported by radiation, the actual temperature stratification is often superadiabatic, reflecting the inefficiency of convection. Along with this difference between near-surface and envelope convection on the one side, and deep interior convection on the other side goes generally speaking a marked difference in stiffness of the convective boundary. This stiffness reflects the ratio of the degree of acceleration in the convectively unstable region to the degree of deceleration in the stable region.

The stratification can be approximated piecewise by polytropes with

$$P = K_{\rm s}\rho^{\gamma_{\rm s}}, \tag{2}$$

where $K_{\rm s}$ is a constant, P and ρ are the pressure and the density and $\gamma_{\rm s}$ is given by the polytropic index $n_{\rm s} = (1/(\gamma_{\rm s} - 1))$. The stiffness $\mathcal{S}$, or relative stability, can then be expressed in terms of the polytropic index for the adiabatic stratification ($n_{\rm ad} = 3/2$ for $\gamma_{\rm ad} = 5/3$), the convectively stable ($n_1 > n_{\rm ad}$) and the convectively unstable ($n_2 < n_{\rm ad}$) stratification by (Hurlburt et al. 1994):

$$\mathcal{S} = -\frac{n_{\rm ad} - n_1}{n_{\rm ad} - n_2}\,. \tag{3}$$

The stability of the stratification of a monatomic ideal gas in terms of the polytropic index $n_{\rm s}$ follows from the definition of the entropy

$$S = c_{\rm v}\log(p/\rho^{\gamma}) + \text{constant} \tag{4}$$

considering that the entropy gradient dS/dr is zero for an adiabatic stratification, >0 for a stable (subadiabatic) and <0 for an unstable (superadiabatic) stratification.

As may be expected, the degree of penetration and overshooting of flows across the convective boundary and the associated degree of mixing of thermodynamic quantities and species concentrations is larger for smaller stiffness (Brummell et al. 2002; Rogers & Glatzmaier 2005). At small stiffness, such as in shallow surface convection (Freytag et al. 1996), coherent convective systems will cross the convective boundary and only start to decelerate on the stable side due to buoyancy effects.

In the deep stellar interior convection zones, high density implies effective convective transport. The denominator in Equation (3) is small ($n_{\rm ad} - n_2 \ll 1$) and therefore $\mathcal{S} \gg 1$. In this situation the convective flows nearing the boundary are "feeling" via the building pressure the almost impenetrable boundary already from a distance (Rogers & Glatzmaier 2005), and in order to obey mass, continuity will start to turn around already inside the convectively unstable layer. A requirement for this behavior is also of course convective velocities well below the local speed of sound. In the He-shell flash convection situations of interest to us, this requirement is certainly fulfilled, at least in the absence of positive feedback from the burning of

ingested hydrogen. Even while no coherent convective systems cross the convective boundary, mixing at the boundary will still occur. In this case the shear flows induced by turning-around convection flows will induce entrainment (Meakin & Arnett 2007), mostly via the Kelvin–Helmholtz instability in the boundary layer.

In summary, a range of hydrodynamic processes and instabilities contribute to mixing and entrainment at convective boundaries (see Viallet et al. 2013 for a similar but more detailed discussion). For deep stellar interior mixing the CBM layer is very small. Another example of this type of CBM is found in the context of the thermonuclear-runaway-driven convection in novae (Casanova et al. 2011; Denissenkov et al. 2012b). While averaged properties pertaining to the inside of the convection zone seem to converge already at modest numerical resolutions (e.g., Herwig et al. 2006; Viallet et al. 2013) the narrow convective boundary layers constitute a challenge in terms of reaching numerical convergence (Viallet et al. 2013).

A narrow and very stiff convective boundary layer can be found at the top of the He-shell flash convection zone in AGB and post-AGB stars at the moment when they make contact with the H-rich material above. Such a situation has been encountered in stellar evolution simulations of AGB stars with very low metal content with [Fe/H] $\leqslant -2$ (e.g., Fujimoto et al. 2000), post-AGB stars (Iben & MacDonald 1995). It leads to violent convective-reactive events involving rapid nuclear energy release from $^{12}\mathrm{C}+p$ reactions in convective regions with large mass fractions of primary $^{12}\mathrm{C}$. Numerous investigations of this scenario have been subsequently carried out in the spherically symmetric stellar evolution framework (e.g., Herwig et al. 1999; Lawlor & MacDonald 2003; Herwig 2003; Iwamoto et al. 2004; Campbell & Lattanzio 2008; Lau et al. 2009; Suda & Fujimoto 2010; Cristallo et al. 2009; Campbell et al. 2010). However, protons ingested into the He-shell flash convection zone on the convective timescale will be advected downward to increasing temperatures until the nuclear reaction timescale of the $^{12}\mathrm{C}(p,\gamma)^{13}\mathrm{N}$ rate equals the hydrodynamic flow timescale. The resulting significant energy feedback on the convective flow timescale violates the assumptions of the mixing length theory, concerning time and spatial averages, and a treatment in spherical symmetry adopted in one-dimensional stellar evolution codes may become problematic. Evidence for the failure of mixing-length theory based stellar evolution models arises from the analysis of the light curve (Herwig 2001) and the highly non-solar, anomalous abundance distribution (Herwig et al. 2011) in the H-ingestion, very late thermal pulse post-AGB object Sakurai's object.

This post-AGB star belongs to the ∼20% of all single post-AGB stars that suffer a final He-shell flash after they have already left the AGB. Among those, the very late thermal pulse stars (Werner & Herwig 2006), but especially Sakurai's object, provide a unique opportunity to investigate the H–^{12}C-combustion regime in stellar evolution because it is a nearby object that has been observed in real-time (Duerbeck & Benetti 1996; Hajduk et al. 2005 and references therein) so that detailed information about the light curve is available. Even more importantly, abundance information of the post-flash evolution has been obtained at a time after the outburst (Asplund et al. 1999) and this has been shown to provide strong constraints on the hydrodynamic processes of combined convection and rapid nuclear burning, i.e., the physics of H–^{12}C-combustion in stellar evolution (Herwig et al. 2011). We are therefore fortunate to have a very powerful validation case for simulations of H–^{12}C combustion, and therefore our attention is directed to carefully check if our simulations pass this validation test.

As mentioned earlier, hydrodynamic simulation results of H-ingestion presented so far are either unlikely to be converged and therefore must be considered questionable (Stancliffe et al. 2011), or it is not clear if they indeed constitute a convective-reactive H-ingestion case. We suspect the latter to be the case for the simulations presented by Mocák et al. (2011) who adopt an initial setup with an artificially shifted H-profile. The emerging H-burning convection features a long nuclear timescale for the $^{12}\mathrm{C}(p,\gamma)^{13}\mathrm{N}$ reaction due to the relatively low temperature in that zone. The resulting Damköhler number is $D < 0.01$, much lower than the value around unity that characterizes a convective–reactive regime. We may therefore conclude that previous work on CBM as it pertains to the H-ingestion events, which are our interest in this paper, provides only limited guidance.

Our strategy is to attack the problem of H-ingestion into He-shell flash convection zones by constructing accurate 3D hydrodynamic simulations. To verify these simulations we show that they converge upon grid refinement in important respects, such as the entrainment rate studied here. In the future we will take advantage of the validation opportunity provided by Sakurai's object to compare simulation results with observations. Our goal at this point is not to develop models of mixing and burning in H-ingestion that can be applied in stellar evolution calculations to all possible realizations of the H-ingestion regime. As a first step to improve upon the rather qualitative simulations presented in Herwig et al. (2011), we would rather like to investigate specifically the entrainment process at the upper boundary of He-shell flash convection that is just about to connect to the H-rich envelope.

In H-ingestion cases, like in Sakurai's object, the onset of the He-shell flash expands the shell layers before the upper boundary reaches the H-rich layers. Typically the aspect ratio $\Delta r/r_{\mathrm{bot}} > 1$, where Δr is the geometric width of the convection zone and r_{bot} is the radius of the bottom of the convection zone. We found in our previous studies (Porter & Woodward 2000; Porter et al. 2000) that the convective flows are dominated by large scale and global modes for such large aspect ratios. We therefore consider it necessary to choose a full 4π geometry for our simulations. One may also expect that the interaction of H-enriched downflows with high-temperature conditions and the ensuing nuclear burning will depend on the entrainment process and vice versa.

For this study we set up our initial conditions to produce shell convection that is typical of the conditions encountered during initial mixing of H-rich material from above the He-shell flash convection zone into the material of the convection zone. We expect the entrainment rate that develops in such a simulation to depend upon this initial base state, on the heating rate corresponding to the He-burning driving the convection, as well as, of course, on the grid resolution of the simulation. We might expect the size of the velocity of the fluid of the convection zone near the location of H-rich material entrainment to play an important role in determining the entrainment rate. Therefore, we also investigate the dependence of these velocities in our simulations upon the grid resolution.

With this study we wish to introduce our updated star simulation methodology and establish that simulations of this type can provide good estimates of the entrainment rate at affordable cost before we proceed to augment these simulations with a treatment of the burning of the ingested H-rich material

and the back reactions upon the flow dynamics that this burning produces (see Herwig et al. 2014 for preliminary results).

This paper is organized as follows. In Section 2 we will briefly describe the numerical techniques of the PPMstar hydrodynamics code and the initial setup used for the simulations presented here. The properties of the hydrodynamic flow and results for entrainment as a function of grid size will be described in Section 3. We close with a discussion (Section 4). The results presented in this paper depend critically on the deployment of the PPB multifluid advection scheme, which is described in detail in the Appendix.

2. METHOD

2.1. The Challenge

To properly simulate the hydrogen ingestion flash, our numerical techniques must be capable of very accurately treating the dynamics within the small range in radius where the Kelvin–Helmholtz shear instabilities act to entrain the stably stratified, more buoyant gas above the helium shell flash convection zone into the convection flow. These instabilities result in breaking waves that cause small puffs of more buoyant gas to become incorporated into the convection flow and subsequently dragged downward into the convection zone. These breaking waves are observed as trains of small eddies near the top of the convection zone. These trains of eddies turn downward where opposing horizontal flows associated with adjacent large-scale convection cells meet (see also the detailed discussion of the entrainment process in Section 3.2). Here buoyant fluid from above the convection zone that has become incorporated into the eddies is pulled downward into the convection zone with the descending sheets of cooler gas (see Section 3.1). The final panel of Figure 1 shows these large convection cells separated by regions shown in blue of descending gas carrying entrained, buoyant fluid along with it. Trains of eddies peeling off from the top of the convection zone can also be seen in Figure 2, which shows the magnitude of the vorticity in a thin slice through the star.

Not only must our method be able to describe the growth and breaking of these Kelvin–Helmholtz waves at the top of the convection zone, but it must also be capable of carefully tracking the entrained buoyant fluid as it is dragged downward and progressively mixed into the surrounding gas of the convection zone due to turbulence. This is no small feat because the boundary between the convectively unstable and stable gas at the top of the convection zone is quite thin. This is clear from Figure 2, where this transition region appears as an almost perfectly circular line, at which the behavior of the flow changes suddenly and radically, for reasons that were discussed earlier. Although the wavelengths of Kelvin–Helmholtz shear modes will generally be much larger than the thickness of this transition layer, the incorporation of buoyant material at this transition is hard to describe accurately in a numerical treatment unless the thickness of the layer can be resolved on the computational grid. Even deciding what the thickness of this layer would be is not straightforward. The physical thickness of the transition region where the concentration of hydrogen falls to essentially zero is initially determined by the character of hydrogen burning in the star before the helium shell flash occurs. As an example that we have adopted here, based on 1D simulations (Herwig 2001) this thickness is found to be about 500 km (Figure 3) for the post-AGB model before the H-shell flash convection zone. In the lower panel of that figure, the transition region is shown in a zoomed-in view, and individual grid cell intervals are also indicated. On our finest grid of 1536^3 cells, this transition occurs over 11 grid cell widths (along the grid direction—we use a uniform Cartesian grid to describe this region, which has the topology of a thin spherical shell). However, the hydrodynamic properties are not only described by the mean-molecular weight gradient. The entropy profile matters as well. This cannot be known for the convective boundary region from one-dimensional stellar evolution calculations. Our initialization strategy is described in Section 2.3. This is an important point because of the concern that our initial setup may somehow be special and not representative of the conditions in a real star.

We face an additional computational challenge because our simulations lie in a regime that is right about at the break-even cost Mach number for implicit relative to explicit numerical methods. The peak Mach numbers in our flows range around 0.03. We have addressed this flow regime with an explicit gas dynamics scheme, relying on the high computational efficiency of explicit schemes and the 12%–23% of the machine's peak floating point computation rate (depending upon the machine) that our code achieves (Woodward et al. 2010a, 2012) to make this approach competitive with implicit methods.

2.2. The Code

Because the entropy plays such a key role in the convective stability or instability of the gas in our problem, we find that to obtain accurate results, it is necessary to solve a conservation equation for the entropy rather than for the total energy. In the absence of shocks and nuclear reactions, the entropy of our gas is conserved along streamlines. At our low flow Mach numbers we do not expect shocks, and we can represent nuclear reactions by source terms in an entropy conservation law. With our gamma-law equation of state, with $\gamma = 5/3$, we use a conservation law for the adiabatic constant $A = p/\rho^{\gamma}$ (Equation (5)). This is a technique employed in meteorological codes, where a quantity called the potential temperature takes the place of the entropy. In weather prediction, Mach numbers in our same general range are typically encountered, so it is natural that similar numerical techniques are useful.

The explicit gas-dynamics code is the same as the one we used in Herwig et al. (2011, Appendix A.2) and is described in full in Woodward (2006). Of particular importance for the entrainment properties at the stable–unstable interface at the top of the convection zone with a mean-molecular weight gradient is the PPB moment-conserving advection scheme (see Woodward et al. 2008a).

Details of the features of our version of the Piecewise-Parabolic Method (PPM; Woodward & Colella 1981, 1984; Colella & Woodward 1984; Woodward 1986, 2006) that make this explicit scheme highly accurate in this flow regime are given in the Appendix. We also describe the Piecewise-Parabolic Boltzmann (PPB) scheme (Woodward 1986, 2005; Woodward et al. 2010b) as well as the features that enable our PPM scheme to be coupled in a natural way to the PPB scheme to describe multifluid fractional volume advection there.

As in previous work (e.g., Porter & Woodward 2000; Herwig et al. 2006) we adopt a monatomic ideal gas equation of state, which represents the conditions in advanced He-shell flash convection at high densities well. In particular, this equation of state provides a good representation of the conditions in the post-AGB He-shell flash that has occurred in Sakurai's object. In some other cases, such as He-shell flashes in the first thermal

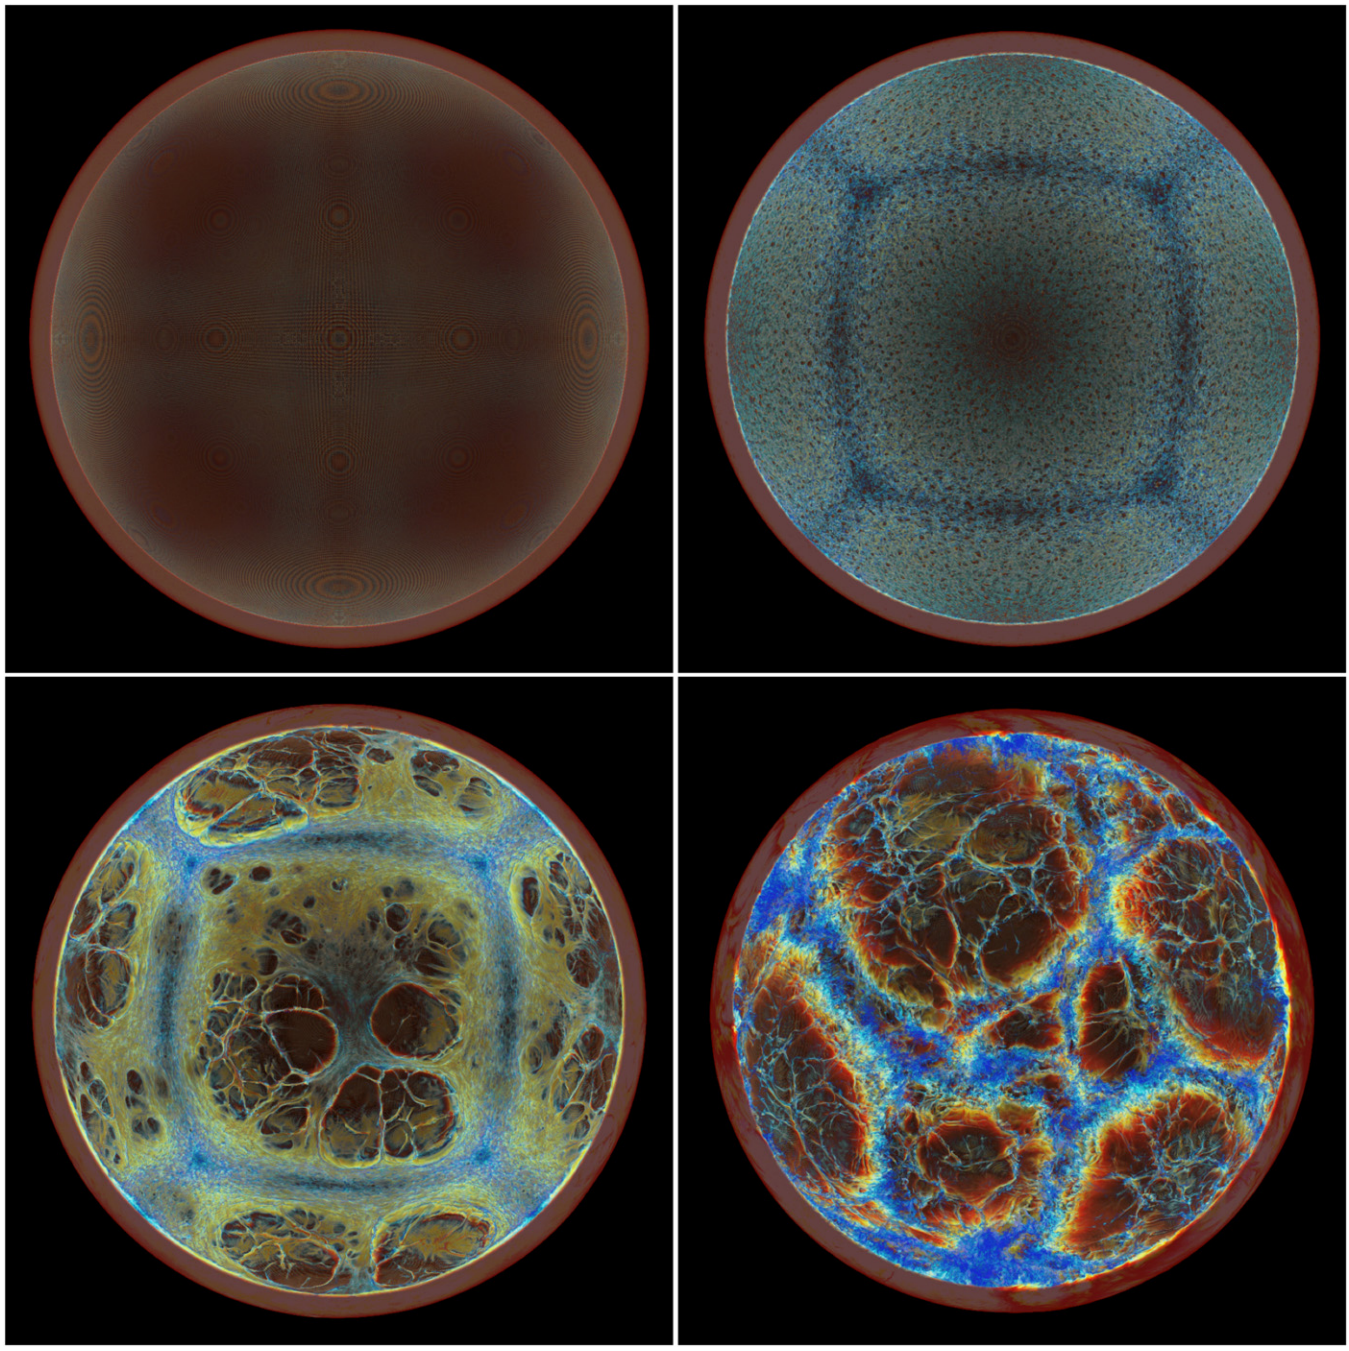


Figure 1. Evidence of initial grid imprints and how they are overwhelmed when the physical flow gets established in the 1536^3 grid simulation. Shown is the logarithm of the fractional volume of the "H+He" fluid originally located only in the stable layer above (Section 2.3), in the rear 75% of the simulation domain (the 25% of the front are cut away). The opacity for the volume rendering is chosen such that levels above $\sim 10^{-4}$ are transparent. The observer is therefore looking into the open shell and the top convection boundary and the mixing interface at that location is seen from the inside. Left top: $t = 1$ minute, right top: $t = 20$ minutes, left bottom: $t = 30$ minutes, right bottom: $t = 75$ minutes. Each of the visualizations appears to have a red ring around it is outer circumference. This is due to the fact that the portion of the 4π convection zone is a bit more than one half. The red ring represents the outside of the convection zone in the small extra portion of the sphere that is closest to the observer.

pulse of low-mass and low metallicity stars, radiation pressure may be important. The code uses appropriately scaled code units for all physical quantities.

2.3. Setup of Star Simulations

We consider a two-fluid setup of a convectively unstable shell and a stable layer below with initially one fluid, and a stable layer above the convection zone with a second fluid with lower mean molecular weight. The simulations are performed on a uniform Cartesian grid with a range of grid sizes up to 1536^3. The full 4π shell of the convectively unstable layer is included in the simulation, which ensures that any global or large-scale motions can be captured. The stratification described below covers 8.3 pressure scale heights in the entire simulation domain, and 4.9 pressure scale heights in the convection zone.

As in many other multi-dimensional investigations of the hydrodynamics of convection (e.g., Porter & Woodward 1994; Hurlburt et al. 1994; Herwig et al. 2006, and many of the works

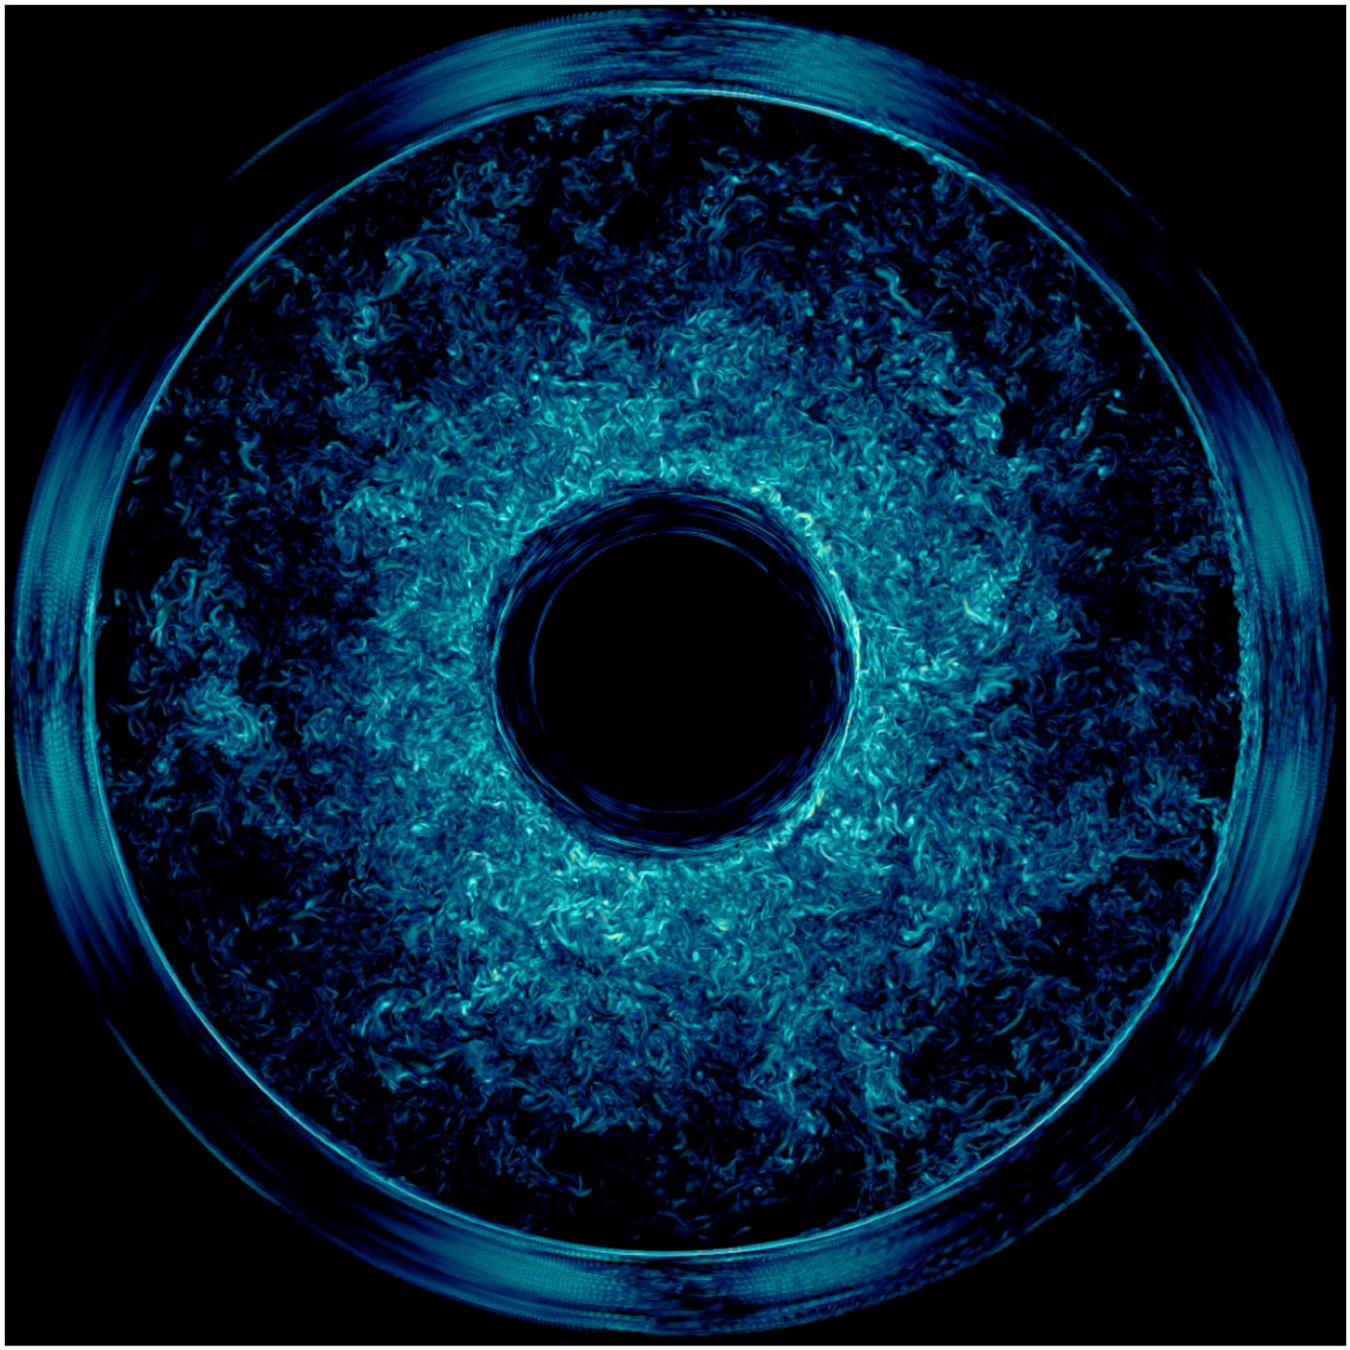


Figure 2. Vorticity of simulation with 1536^3 grid at $t = 480$ minutes. The volume rendering shows a thin slice of the central ~5% of the full 4π simulation domain. The inert core appears as a black disk in the center.

mentioned in Section 1), we construct the radial stratification with a set of connected, piecewise polytropic layers in which the polytropic constant and the polytropic index are chosen in such a way that the overall representation resembles a typical situation of a luminosity-driven shell convection layer bounded by a stable layer above and below, just like He-shell flash convection. This approach has the advantage that the initial state is in very good numerical hydrostatic equilibrium, and initial transients are minimized. Mapping a 1D profile from a stellar evolution code (as done, e.g., by Meakin & Arnett 2007) would provide little in additional accuracy in our case where the microphysics is exceptionally simple. In Figure 4 we show the resulting initial stratification in terms of ρ and in terms of the adiabatic constant A, which is related to the entropy (Equation (4))

$$\log A = \frac{1}{c_v}(S + \text{constant}) = \log(p/\rho^\gamma), \tag{5}$$

with $\gamma = 5/3$ for the monatomic ideal gas. The sound speed decreases from ~1500 km s^{-1} at the bottom of the convection zone to ~200 km s^{-1} at the top of the convection zone.

The dimension and stratification details of He-shell flash convection in AGB stars are variable, as a function of mass, metallicity, early versus late thermal pulse and AGB versus post-AGB thermal pulse. In addition, and quite obviously, the conditions also change dramatically for a given thermal pulse as a function of time due to the large energy deposit from the

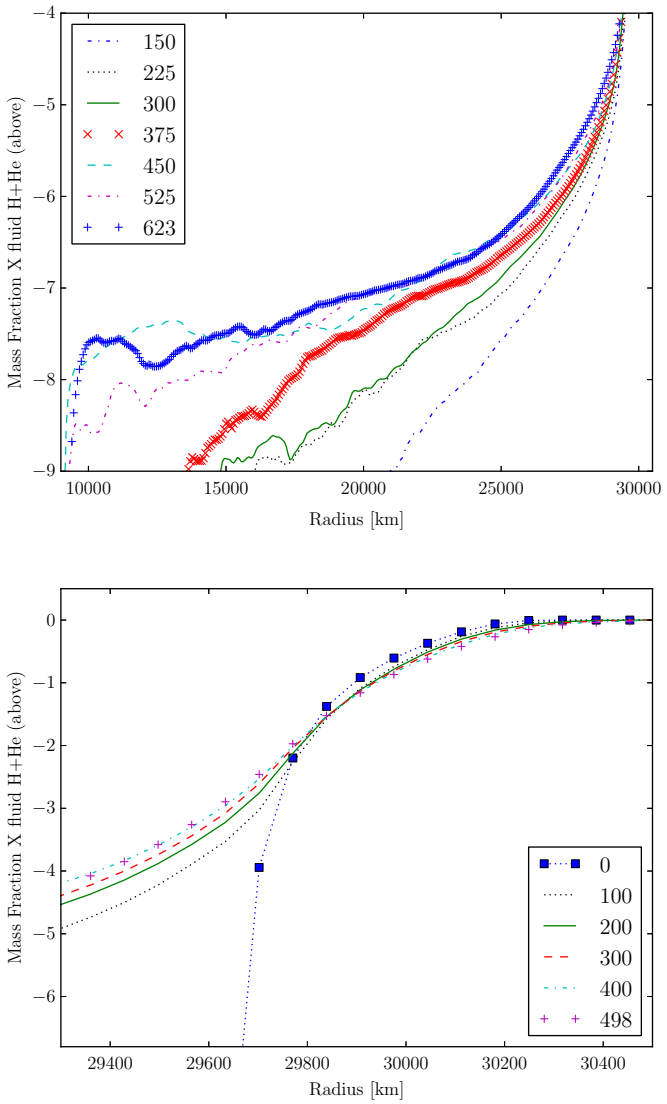


Figure 3. Spherically averaged abundance profiles in the convection zone for a series of dumps. Each dump corresponds to ≈1 minute star time. The markers correspond to grid zones. Top: 1024^3 grid simulation for the entire convection zone. Bottom: spherically averaged abundance profiles at top convection boundary for the 1536^3 grid simulation, at different times, line labels correspond to minutes of simulated time. Minute 0 represents the initial profile reflecting the smooth transition according to the former H-shell burning at that location (see the text).

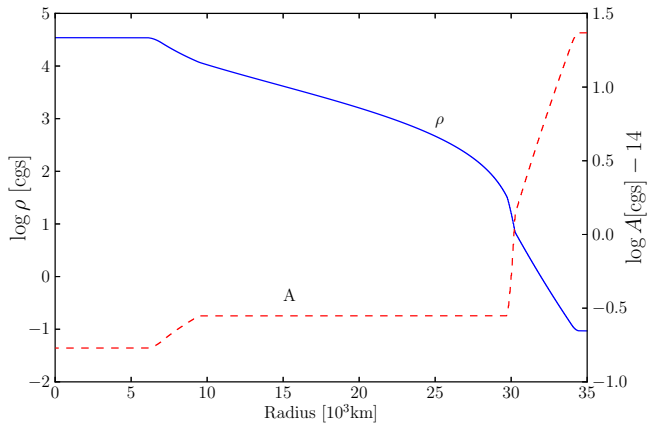


Figure 4. Initial stratification for all simulations presented in this paper. The entropy-like quantity A is defined in Equation (5).

He-shell flash (see, for example, Figures 1 and 2 in Herwig et al. 2006). The variety of conditions is also reflected in the degree in which radiation pressure contributes to the total pressure. Initially, when the flash starts, the density is so high that the entire pressure is provided by the gas. As the He-shell layer absorbs the peak-flash luminosity of several $10^7\,L_\odot$ it rapidly expands. The temperature decreases more slowly and the radiation pressure contribution increases. The lowest value for the gas pressure fraction $\beta = P_{\rm gas}/P$ is found at the bottom of the He-shell flash convection zone where it decreases to $\beta \sim 0.94$ at the time when the He-burning luminosity reaches its maximum for thermal pulses in a stellar evolution sequence of $M_{\rm ini} = 2\,M_\odot$ and $Z = 0.02$ with core masses ranging from 0.5 to 0.6 $M_\odot$.

Herwig et al. (2006) adopted a plane-parallel modeling approach, which is appropriate only for a small ratio of the geometric thickness of the convection shell and the radius of the underlying core. They have therefore chosen a time of the flash just before the peak luminosity is reached. We simulate the full 4π geometry of the shell. This is necessary because the entrainment starts only when the peak luminosity has been reached, or just passed. At that later time in the thermal pulse the convection zone has already greatly expanded and is now thicker than the radius of the underlying core. We have therefore chosen to reproduce in our setup the geometric dimensions encountered at $t = 0.25$ yr of the model sequence in Figures 1 and 2 shown in Herwig et al. (2006) where $t = 0$ yr corresponds to the time when He-burning has reached the flash-peak luminosity. At this time the density would have further decreased, but still $\beta_{\rm min} \gtrsim 0.8$.

Thus, we are not simulating here a specific metal-poor AGB star, or a post-AGB star. Rather than trying to reproduce the specific conditions of a particular case we investigate more generally the typical behavior of shell flash convection with very stiff convective boundaries, which do not depend on the exact numerical value of any of the setup parameters.

The unstable layer. The density, the pressure, the gravity and the radius at the bottom of the convection zone are $\rho_{\rm bot} = 1.174 \times 10^4$ g cm^{-3}, $P_{\rm bot} = 1.696 \times 10^{20}$ g (cm s^{-2}), $g_{\rm bot} = 4.9545 \times 10^7$ cm s^{-2} and $r_{\rm bot} = 9.5 \times 10^8$ cm, which imply a core mass below the convection zone of $M_{\rm bot} = 0.337\,M_\odot$. This mass is lower than the typical core mass of thermal pulse AGB stars, but it allows us to accommodate our choice of an ideal gas equation of state and the geometric dimension of the convection zone while ignoring for the moment a moderate fraction of radiation pressure.

The top of the convection zone is located at $r_{\rm top} = 3.00 \times 10^9$ cm. The equations for hydrostatic equilibrium, mass conservation and Equation (2) with $\gamma_{\rm s} = \gamma_{\rm ad}$ and $K_{\rm s}$ given by $\rho_{\rm bot}$ and $P_{\rm bot}$, are numerically integrated to provide the density and pressure stratification in the convection zone.

The upper boundary and the stable layers. The layer above the convection zone is stable with $\gamma_{\rm s-top} = 1.01$. This region is populated by a fluid representing the H- and He-dominated envelope of AGB and post-AGB stars with a mean molecular weight $\mu_{\rm H+He} = 0.7$ similar to a mostly unprocessed envelope abundance distribution that consists predominantly of H and He. This fluid is therefore labeled "H+He". This compares to $\mu_{\rm conv} = 1.58$ of the fluid in the convection zone (labeled "conv"), which contains a mix of ^{4}He, ^{12}C and ^{16}O that is typical for the He-shell flash convection zone in AGB or post-AGB stars.

The abundance interface between these two layers is not a discontinuity, but instead it reflects the smooth transition resulting from H-shell burning with varying efficiency across the temperature profile at that location. The resulting H-profile of a post-AGB stellar evolution model shown in Figure 5 in Herwig et al. (1999) is representative for this situation. Therefore, the H-profile at the bottom of the H-rich envelope in a radiative layer, before the He-shell flash convection approaches from below, will vary smoothly from a mass fraction of $X(\mathrm{H}) \sim 0.7$ to zero over a width of the former H-burning shell of $\Delta r_{\mathrm{H\text{-}shell}} \sim 5 \times 10^7$ cm (bottom panel, Figure 3). The width of this gradual transition from H-rich to H-free depends on the temperature gradient in the H-burning shell, and may vary depending on the particular type of thermal pulse as well as the mass of the underlying degenerate core and the metallicity.

When continuing the integration of the stratification into the stable layer we let both γ and the concentration of the two fluids vary smoothly over a width of $\Delta r_{\mathrm{H\text{-}shell}}$ centered at r_{top}. That means that the constant entropy region of the convective unstable zone ends at $r_{\mathrm{top}} - \Delta r_{\mathrm{H\text{-}shell}}/2 = 29,750$ km where the entropy gradient becomes positive. In the 1536^3 grid simulation the interface width $\Delta r_{\mathrm{H\text{-}shell}}$ corresponds to 11 cells along a grid axis. The shape of the resulting profile is shown in the bottom panel of Figure 3.

When integrating the initial stratification downward from the bottom of the convection zone into the lower stable layer we adopt $\gamma_{\mathrm{s-bot}} = 1.2$ and a transition layer of 2.5×10^7 cm centered at r_{bot}. The fluid type is the same as in the convection zone.

Luminosity driving the convection. The flash constitutes a thermonuclear runaway that is the result of an unstable balance between thermodynamic, hydrodynamic and nuclear physics components. When mapping a 1D profile into a 3D code we cannot expect the thermonuclear runaway to continue on the same trajectory as in 1D, with a multitude of small differences between the two simulation approaches. Instead, we use the piecewise-polytropic stratification as a background and add a constant volume heating at the bottom of the He-shell flash convection zone, that has a total luminosity equivalent to the He-burning luminosity in a 1D stellar evolution model. The constant heating is applied over a heating region of width $\Delta r_{\mathrm{heat}} = 10^8$ cm, the lower boundary of which is offset from the bottom of the convection zone by 5×10^7 cm to avoid any entrainment of low-entropy material from below the convection zone. The heating is applied in the 4π shell with a smooth bell-shaped radial intensity distribution. This arrangement is providing a robust heat source that is independent of the resolution and represents the He-burning luminosity from the triple-α reaction. The heating rate in the standard case corresponds to $L_{\mathrm{He}} = 4.20 \times 10^7\, L_\odot$.

On the absence of radiation in our simulations. Unlike in envelope convection and shallow surface convection simulations, the convection is not driven by a subtle degree of superadiabaticity that is only reached in a quantitatively correct sense over the rather long thermal timescale due to the effect of radiative transport at the bottom and top layer of the convection zone. Instead, the convection is very efficient and driven by the large luminosity from the underlying thermonuclear runaway.

In the deep interior case of He-shell flash convection, radiation transport is not important on the dynamic or convection turnover timescale (but radiation pressure may be). The mean free path of a photon is $l_{\mathrm{ph}} = 1/\rho\kappa$ where the $\rho \sim 10^4$ to 10 g cm^{-3} in the convection zone. The opacity is of order unity, which means that $l_{\mathrm{ph}} < 0.1$ cm. The characteristic length scale

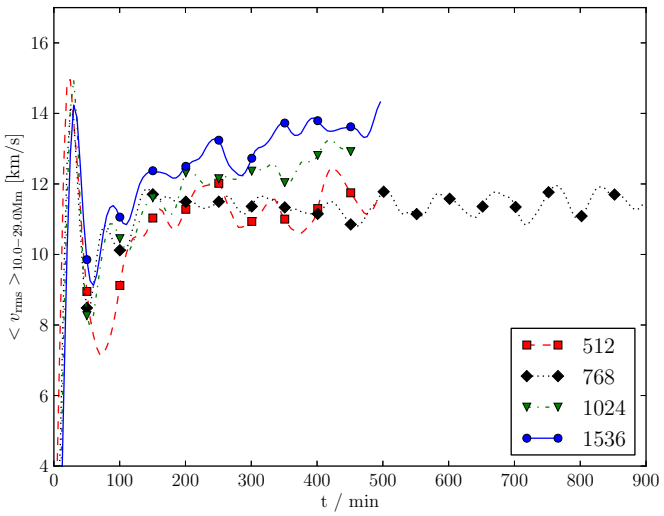


Figure 5. Spherically and radially averaged rms velocity (averaged over the sphere and over convection region between $r = 10 \times 10^6$ and 29×10^6 m) as a function of time for runs with different grid sizes n^3. The number of grid points n in one direction is shown in the legend.

for radiation transport can be expressed as $l_{\mathrm{rad}} = 2\sqrt{(Dt)}$ where $D = (1)/(3)cl_{\mathrm{ph}} \sim 10^9$ cm^2 s^{-1}. With a radial domain size of $R = 3.5 \times 10^9$ cm the highest resolution runs presented in this paper (with 1536/2 radial zones) have a grid size of $dx = 4.6 \times 10^6$ cm. If we take as a convective timescale the radial extent of the convection zone divided by the typical velocity in the convection zone (Figure 5) we obtain $t_{\mathrm{conv}} \sim 1500$ s and $l_{\mathrm{rad}} \sim 2 \times 10^6$ cm.

Therefore $l_{\mathrm{rad}} \approx dx$, or in other words, within the time that it takes convection to flow radially across the entire convection zone radiation transport crosses one simulation cell in the highest resolution case. However, in the outermost layers of the simulation domain the density decreases to 0.1 g cm^{-3} and here indeed radiation diffusion will eventually become relevant over the entire course of the simulated time (typically $\lesssim$1000 minutes). If these lowest density regions will become important in future simulations, or if the simulated time will further increase or if we further increase the resolution, then indeed diffusive radiation transport should be included. For the present set of simulations this physics aspect can, however, safely be neglected (see Meakin & Arnett 2007 for a similar conclusion for their O-shell burning simulations).

2.4. Tests and Simulations

We started this investigation with numerous test simulations at various resolutions, experimenting, for example, with different assumptions for the treatment of the convection boundaries, the properties of the heating zone, peculiarities of different Fortran compilers, and the effect of using single versus double precision arithmetic. For the simulations shown here we are adopting the Intel 12 Fortran compiler and double precision arithmetic. We have also run test simulations on different hardware platforms.

After we were satisfied with the results of these tests, we performed a grid of simulations for the following grid resolutions[3]: 512^3, 768^3, 1024^3, 1536^3. Animations of these

[3] We did tests with a 384^3 grid, which showed obvious numerical artifacts and were not further considered. A particularly troublesome feature were entropy shelves that formed above and below the proper convection zone. These were small constant-entropy spheres, detached from the big convection zone that formed and survived in the initially stable layers. These features went entirely away once we adopted higher grid resolutions and double-precision arithmetic.

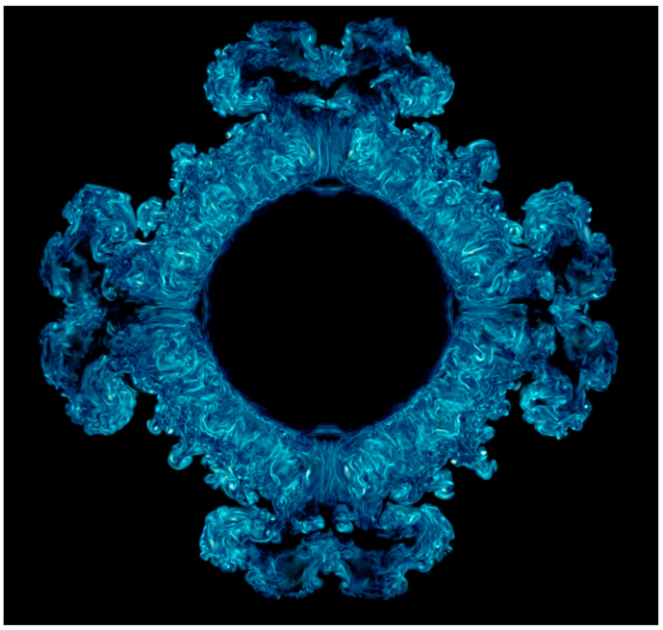


Figure 6. Zoomed in view of the vorticity during the initial transient period ($t = 20$ minutes) of the run with 1536^3 grid. Same view as in Figure 2 except that a smaller portion of the simulation domain is shown. The heads of the upwelling plumes along the grid axis have advanced approximately half the distance between the bottom and top of the convection zone.

simulations can be found at http://www.lcse.umn.edu/movies. The simulations on the coarsest grids were performed on the 25 node workstation cluster at the Laboratory for Computational Science & Engineering (LCSE) at the University of Minnesota. Those on 768^3 grids were performed at either the Minnesota Supercomputing Institute (MSI) or on the Canadian WestGrid Orcinus and Lattice high-performance computers, runs on 1024^3 grids were performed at MSI, while runs at the highest resolution were performed on the NSF's Kraken supercomputer at the National Institute for Computational Sciences (NICS). Our run on a grid of 1536^3 cells used 98,312 CPU cores on Kraken and consumed 5.2 million CPU-core hours. The delivered performance was about 120 Tflop/s in 64 bit precision.

We will demonstrate in the next section that the 512^3 grid gives an entrainment rate that is a factor of almost 10 larger than the asymptotic entrainment rate, while the 768^3 grid provides enough resolution to reproduce the asymptotic entrainment almost within a factor of two. We therefore consider the latter grid size the absolute minimum to perform, for example, differential investigations of the effects of various input assumptions.

3. RESULTS

3.1. General Properties of the Convection Simulations

Initial transients and grid effects. We do not seed the initial setup with perturbations. Instead, the grid provides asymmetries that allow initial perturbations, fueled by the continuous heating at the bottom of the convection zone, to form and grow along the axis directions of the grid (Figure 6). The initial start-up phase of any simulation like this is prone to artificial, initial transients that are also observed in the time evolution of global properties, such as the spherically and radially averaged rms velocity (Figure 5). The velocity maximum at early times is due to preferential rise of the first plumes along the grid axis. These initial transients do not matter since we run the simulations for long enough, until we reach a convective steady state, which is dominated by the physically relevant and realistic fluid motions. The amount of H-rich material entrained during the initial transients is also very small. Even if its burning were to be included in the simulation, this small amount of entrainment would not generate a dynamically relevant amount of energy.

During the initial transient the first convective plumes are launched along the grid directions. This is shown in Figure 6, which shows that at time 20 minutes these plumes have traveled through about half of the radial distance to the upper convective boundary. In all cases an initial peak of the transient velocity subsides within ≈50 minutes, shortly after the first plumes have first reached the upper convection boundary (bottom, left panel Figure 1). Velocities adopt a local minimum before the heating luminosity takes over the fluid flow dynamics and the rms velocities increase again. Abundance profiles at 100 minutes (Figure 3) show the action of the initial convective turnovers.

In some previous investigations such initial transients have been damped before or while the convection driving luminosity was turned on (e.g., Herwig et al. 2006). We do not do anything in particular to dampen the initial transients, but instead study how they become smaller in magnitude with increasing resolution. Indeed, close inspection of Figure 5 shows that for the highest resolution run (1536^3), the peak rms velocity of the initial transient is about as high as the steady-state value approached after 200 minutes. Another concern may be that the initial transient is shaking the interface so much that we get artificial smearing. The profile evolution at the top convection boundary (bottom panel, Figure 3) shows that this is not the case.

The initial transients can also be observed in the entrainment visualization shown in Figure 1. These images show the "H+He" fluid that initially occupies only the region above the convection zone (see bottom panel, Figure 3). Abundances $\gtrsim 10^{-4}$ are assigned to be transparent in the volume rendering. Levels down to $\approx 10^{-7}$ are assigned with opacities of different color, from red to yellow to blue in decreasing order of concentration. The upper left panel in Figure 1 shows the initial state in which by design the visible interface is represented by only about one grid zone (see Figure 3) because we only show concentrations below $\approx 10^{-4}$. Because the thickness of the visible interface is comparable with the size of a grid zone interference, rings and patterns appear at the location of the Cartesian grid directions.

During the initial transient phase (top right and bottom left panel, Figure 1) rectangular grid imprint patterns are clearly visible. They correspond to very small perturbations at this interface when the convectively driven motions at that boundary are not fully established yet. However, already at $t = 75$ minutes any evidence of grid-imposed patterns has vanished. This time corresponds to the end of the initial transient phase when the radially and spherically averaged rms velocities have a minimum (Figure 5). From this point onward the luminosity-driven convective motions are strong enough to dominate any grid-imposed numerical noise in our simulations.

Once the convective motions dominate the flow in the entire constant-entropy shell the patterns at the upper boundary imply large-scale features. Typically two to four coherent patches (or granules) can be identified per hemisphere. These represent areas in which fluids are collectively upwelling. The situation shown in the lower right panel in Figure 1 is representative for later stages (see also Figure 7), which rather

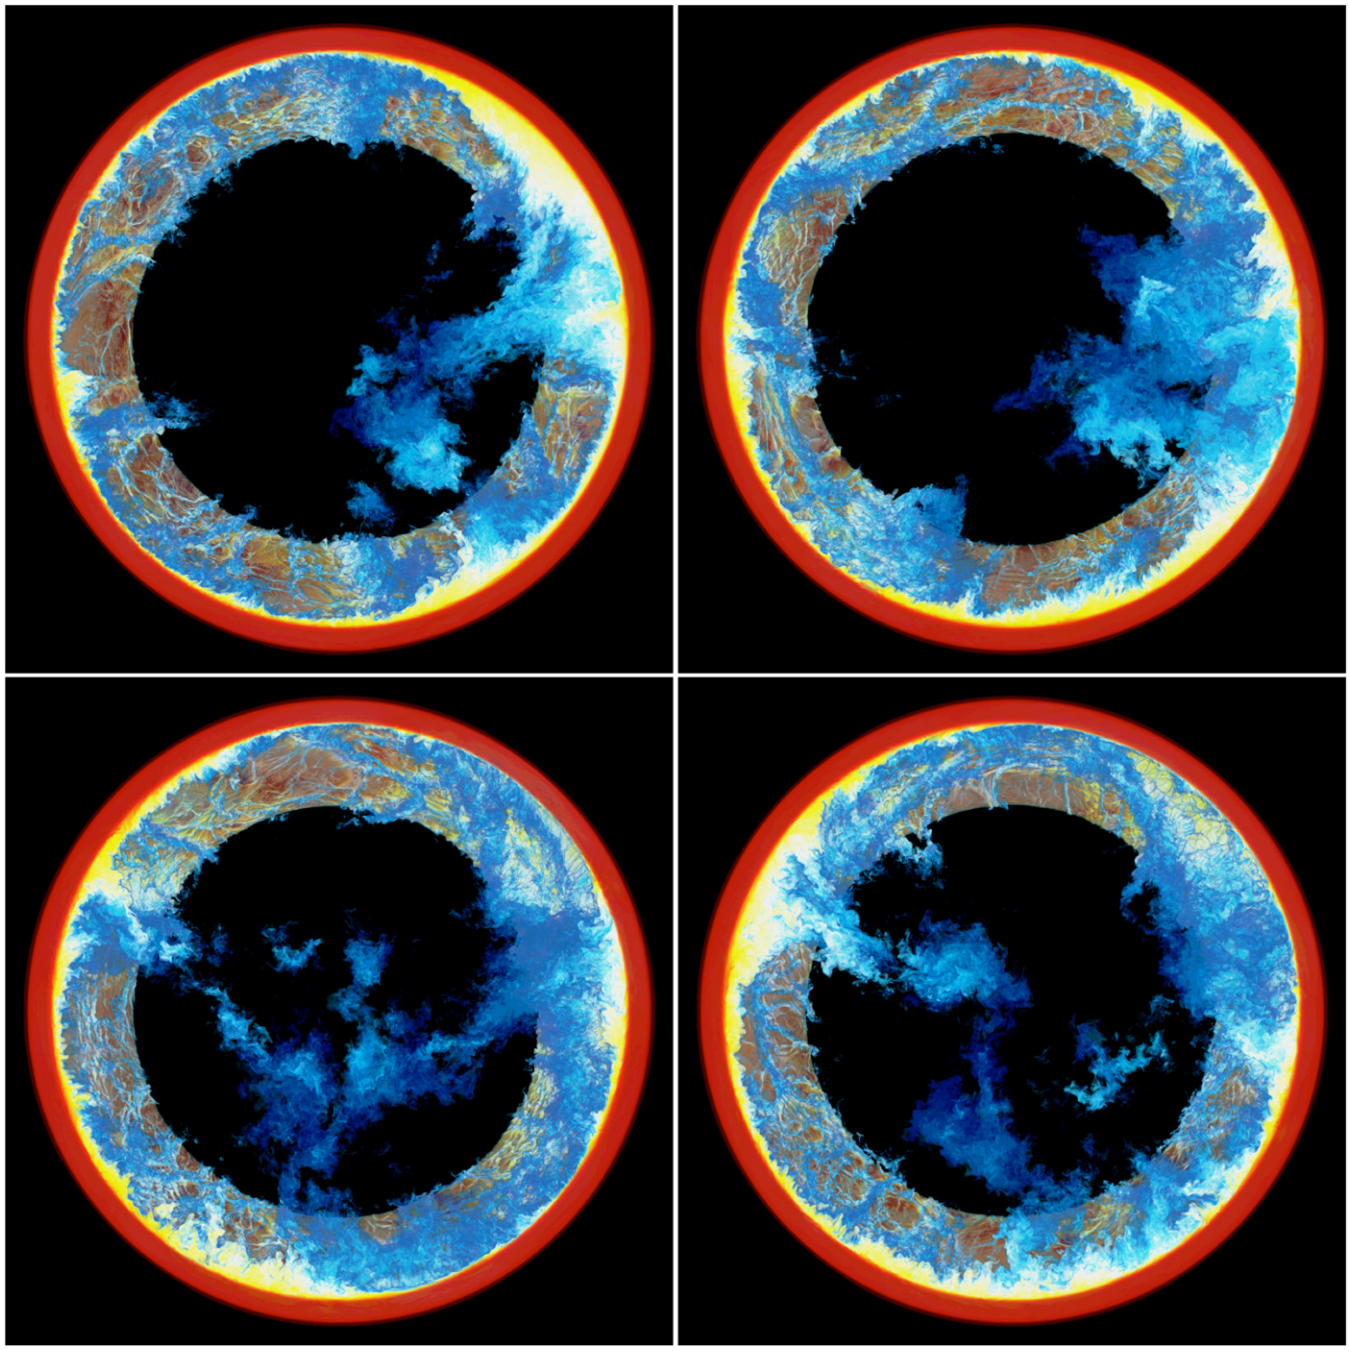


Figure 7. Same as in Figure 1: four snapshots of the logarithm of the fractional volume of the "H+He" fluid is shown as it is entrained into the convection zone. The lighter and darker blue color corresponds to concentration of 10^{-5}–10^{-6}. The front and rear 25% are cut away in order to show the interior distribution of entrained material. Left top: $t = 510$ minutes, right top: $t = 555$ minutes, left bottom: $t = 625$ minutes, right bottom: $t = 640$ minutes.

tend to show fewer coherent systems and therefore a more global nature of the convection flows. These global modes have as their limit the global dipole mode reported for non-rotating, fully convective configurations (Porter & Woodward 2000; Kuhlen et al. 2006) and are correlated with the aspect ratio $\Delta r_{\rm shell}/r$ of shell convection. Large aspect ratios lead to larger cells, while small aspect ratios have smaller cells. In any case, convection cells fill the entire radial extent of the convection zone, and approximately the radial extent determines the horizontal extent.

Aspect ratio and geometry of simulation domain. Shell-flash convection typically starts out with a small aspect ratio. This regime is suitable even for plane-parallel simulations (Herwig et al. 2006). As the thermo-nuclear runaway proceeds, much of the energy released goes into expansion work and lifts the outer layers of the convection zone and above leading to an increasing aspect ratio. Stellar evolution models suggest that the energy release from ingested hydrogen may induce a further dramatic expansion of the convection zone and increase of the aspect ratio. For large aspect ratios geometric properties of the simulation domain become important. We would expect that entrainment processes be driven by the velocity field at the upper convective boundary, which depends on the realistic simulation of the large-scale cells. In this situation a 3D 4π

simulation domain is therefore likely to give the most realistic representation of convection flows and therefore entrainment.

Scales and turbulence. Images of the abundance distribution clearly show the large-scale nature of convectively driven entrainment shear flows at the upper convective boundary (bottom right panel, Figure 1) and the subsequent downward advection in coherent cloud systems (Figure 7). The picture of convective cells, however, has limits. The downward directed H-rich plumes are highly turbulent and have significant substructure. This is supported by the vorticity visualization (Figure 2), which shows that the convection zone is highly turbulent throughout, and that the large-scale features that can be identified in the mixing images (Figure 7) carry along with them a full range of smaller-scale motions. Vorticity, arising from derivatives of the velocity field, accentuates the smallest scales, and therefore, looking at the vorticity alone would make the notion of dominant large-scale structures less assertive. Nevertheless, the largest features are those corresponding to the bulk convective motion from the bottom to the top of the convection zones, and can be identified to span the radial extent of the convection zone (≈19,000 km). The largest turbulent scale can be determined by finding the largest scale at which the direction of the convective large-scale motions cannot be identified. This scale is ≈2,800 km.

The velocity field and the lower convective boundary. We return to the discussion of velocities in our simulations. We follow the simulations for several hundred minutes past the initial transient phase (see above), where the convection timescale is ≈25 minutes (Figure 5). From the spherically and radially averaged rms velocity, we obtain an overall velocity scale for our particular convection setup of $v_{\rm rms} \approx 11$–13.5 km s^{-1} depending on resolution.

The spherically averaged profiles of the rms velocity, as well as the vertical and horizontal components (Figure 8), reveal the dominance of the horizontal over the vertical convective velocities near the convective boundary. This was described for the upper convection zone in the previous paragraph, but it applies equally to the bottom convection zone for the same reasons. Also here the radial velocities decrease well inside the convective boundaries while the horizontal velocities that are associated with the turnaround of fluid elements peak right above the convective boundaries.

3.2. The Entrainment Process

We noted above the importance of the large-scale motions for the entrainment process. Our simulations show for this setup consistently that large upwelling features occupy a full quadrant or a larger fraction of an entire hemisphere. When these upwelling fluid elements approach the upper boundary they are turning around. The vertical velocities are decreasing while the horizontal motions increase (Figure 8). The continuity equation demands that fluid elements start turning around some distance away from the convection boundary (which they cannot cross in highly efficient convection in the deep stellar interior, see Section 1). Thus, the large-scale flows cause convection-induced shear flows at the convective boundary. The Richardson number is

$$J = -\frac{g}{\rho}\frac{d\rho/dz}{(dU/dz)^2}, \tag{6}$$

where, for our case, U is the horizontal velocity, z is the radial coordinate, and g and ρ have their usual meaning. For shear

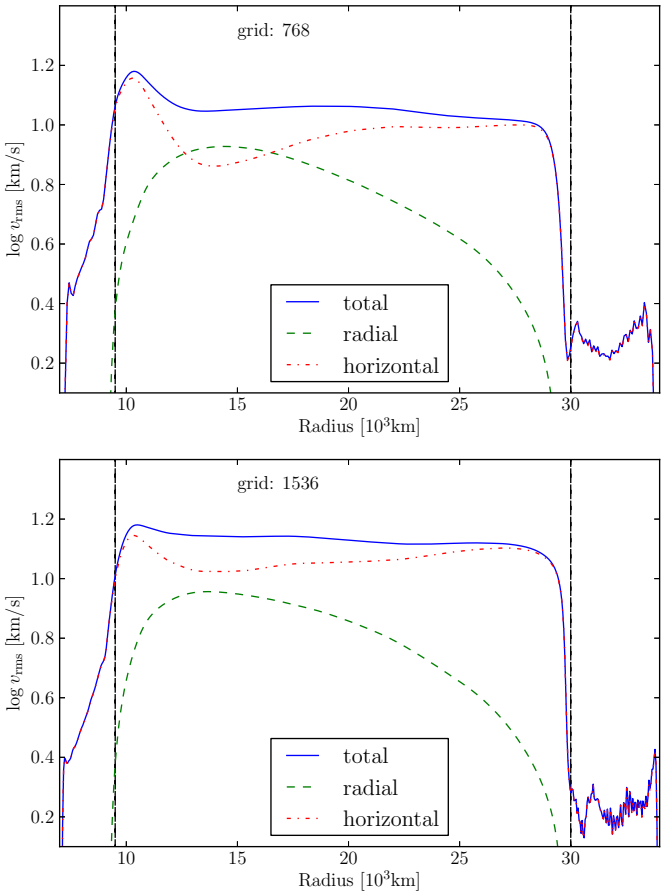


Figure 8. Logarithm of spherically averaged rms velocity, averaged from minute 400 to 430, for two resolutions. The dashed vertical lines are placed at the convective boundaries according to the initial setup.

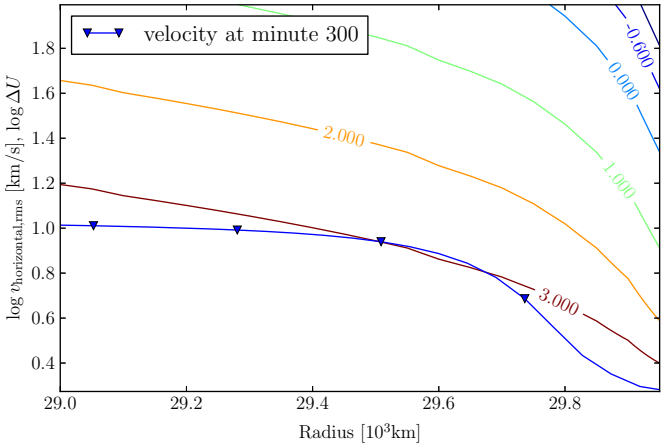


Figure 9. Lines of constant Richardson number log J as a function of distance from the boundary and assumed velocity difference U between that distance and the midpoint of the boundary at 30.0×10^6 m. The line with triangle markers shows the spherically averaged profile of the horizontal velocity component. For stratified shear flow that is unstable against Kelvin–Helmholtz instabilities log $J < -0.6$ (see the text).

flows to be unstable against Kelvin–Helmholtz instabilities it is necessary that

$$J < \frac{1}{4}$$

(Chandrasekhar 1961). From a 1D profile it is not obvious how large the transition region should be assumed to be. In the region

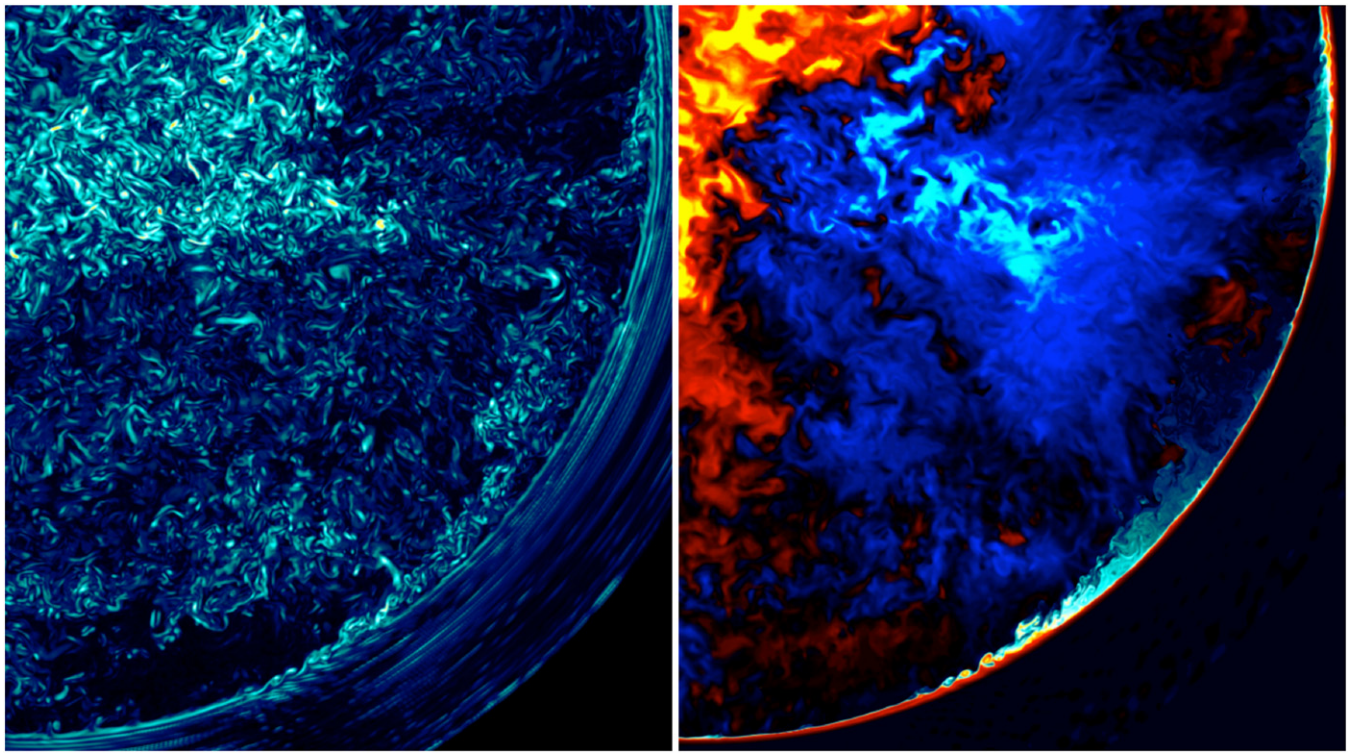


Figure 10. Vorticity magnitude (left), and (right) the radial (blue inward, yellow/red outward) velocity and the fractional volume of the H+He fluid (visible only near the outer boundary of the convection zone) in a slice through the domain with a thickness of 0.2% of the domain for time 408 minutes from the 1536^3 grid simulation.

of the upper convection boundary a pressure scale height is $H_p \approx 1000$ km. The actual velocities of the shear flows will have a range around the value given by the spherically averaged 1D profile. Therefore, we show in Figure 9 the Richardson number as a function of an assumed shear velocity difference ΔU for positions r within the region of a pressure scale height below the upper convective boundary at $r_{top} = 30 \times 10^6$ km. We approximate dz in Equation (6) with $\Delta r = r - r_{top}$ (where r is the radius), $d\rho$ with $\Delta\rho = \rho(r) - \rho(r_{top})$, dU with $\Delta U = U(r) - U(r_{top})$. We also provide the profile of the spherically averaged velocities at time 300 minutes. Even taking into account that the local horizontal velocities are fluctuating around the spherical average the Richardson number indicates that the convectively driven horizontal shear flows are stable against Kelvin–Helmholtz instabilities.

This is consistent with the flow images of the convective boundary where the horizontal flows take the form of coherent gusts that sweep along over the inside of the large upwelling patches visible at the top convection boundary. Larger patches of such apparently stable and very thin boundary layers can be seen in Figure 13 in locations (here, e.g., the lower left) where radially upwelling flows first encounter the convective boundary and then are deflected.

Figure 10 shows at the bottom left a sliced view of another Kelvin–Helmholtz stable part of the horizontal flow. In the region near the center of this image, two such horizontal flows directed toward each other meet and are deflected downward. This is the region where the entrainment of H-rich fluid occurs. The deflection of the flow downward, away from the top of the convection zone, is akin to the well-known phenomenon of boundary-layer separation. However, in this instance, we do not have a solid wall, as in engineering flows, but a strongly stable transition layer from one fluid to another. Boundary-layer separation flows are highly non-linear and cannot be described by simple relations (e.g., Simpson 1989). Quantitative characterizations of such flows do therefore rely entirely on sufficiently resolved simulations or suitable experiments.

In Figure 11 we see zoomed-in views of the vorticity magnitude in the flow near the top of the convection zone, taken from the larger images in Figure 10. In these zoomed-in views, we indicate with prominent white double arrows the 500 km thickness (11 grid cell widths) of the initial transition region from the fluid of the convection zone to the more buoyant, stably stratified "H+He" fluid above it. In visualizing the concentration of the "H+He" fluid in the images at the bottom of Figure 11, we have used the 10 moments of this variable (see the Appendix) that are updated by the PPB advection scheme to generate 8 octant averages in each grid cell. This permits the full resolution of the PPB scheme, with its subcell information, to be seen in this double resolution visualization process. For the image of the "H+He" fluid concentration, the white double arrow is thus 22 voxels wide. In both zoomed-out and zoomed-in views we see that the transition layer at the top of the convection zone is strongly stable to the Kelvin–Helmholtz instability. This is due to the great strength of gravity and the Rayleigh–Taylor stable stratification of this layer. Nevertheless, at the tip of the wedge of turbulent gas that separates the two large convection cells and that serves to deflect their flows downward, we clearly see that Kelvin–Helmholtz roll-up indeed occurs.

At these locations, we no longer have the parallel flow in a single direction assumed in the textbook Kelvin–Helmholtz stability analysis. Instead we have two colliding flows that must, at these locations, generate vertical motions perpendicular to the transition layer that marks the top of the convection zone, in order to satisfy the continuity equation for this essentially incompressible flow. A classic series of breaking Kelvin–Helmholtz waves is formed at each tip of this wedge of relatively

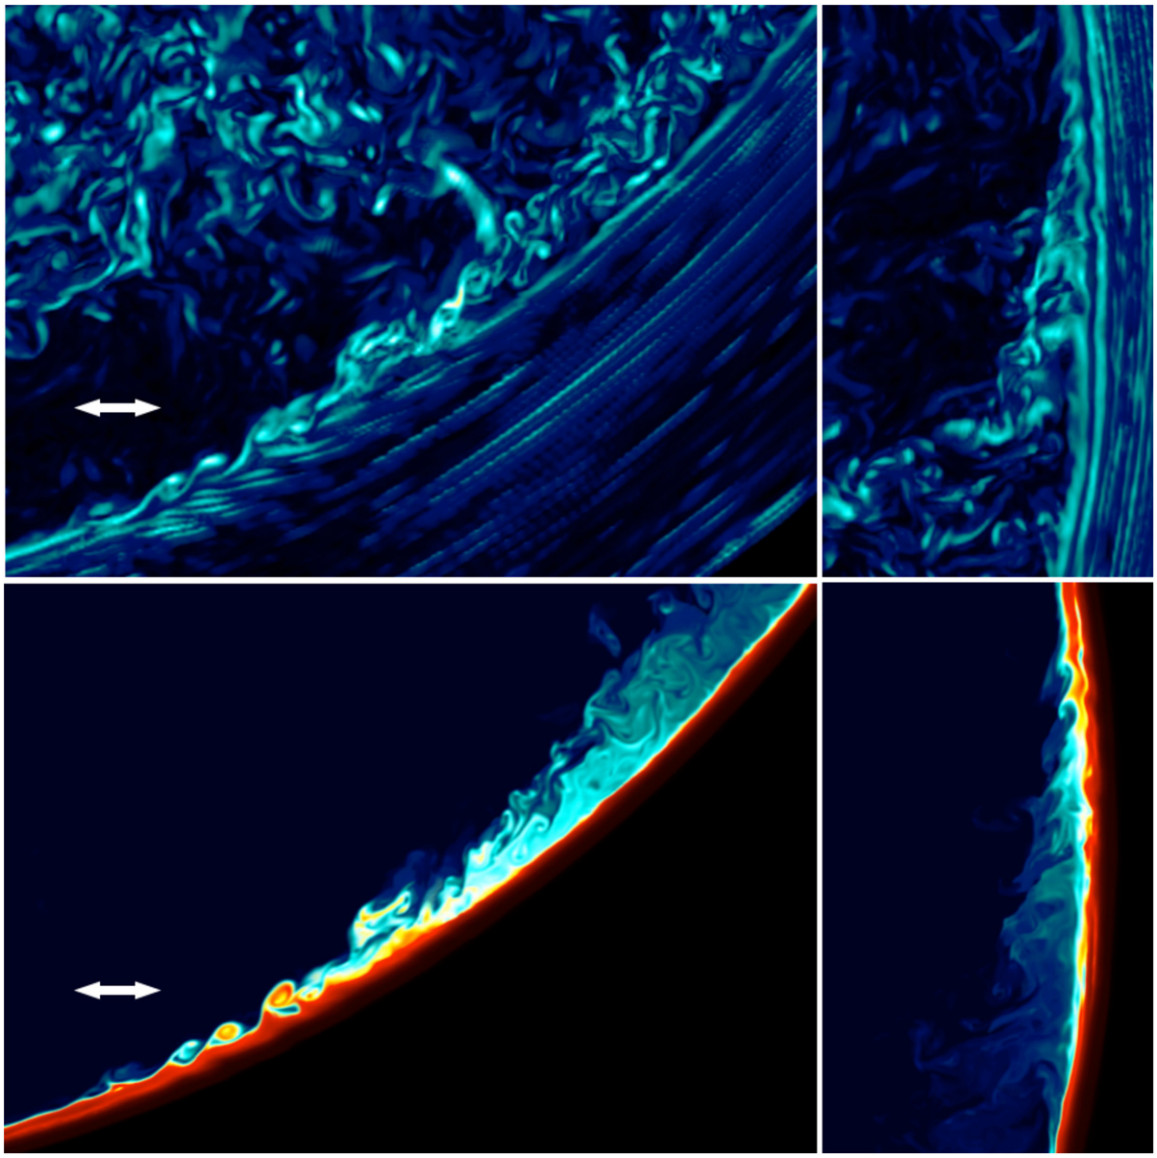


Figure 11. Zoomed-in view of Kelvin–Helmholtz unstable breaking waves in the boundary-separation layer shown in Figure 10. The white double arrows indicate the 500 km thickness of the initial transition layer between the gas of the convection zone and the more buoyant, stably stratified "H+He" gas above it.

stagnant gas between the large convection cells. Not only is low-concentration H-rich gas entrained into the gas of this wedge, but also along the shear layer between the wedge and convection cell gases entrainment into the rapidly moving convection cell gas occurs. This process is perhaps more clearly seen in the zoomed-in views of the concentration of H-rich gas in Figure 11. In these zoomed-in images of the H-rich gas concentration, the slice through the domain is three times thicker, namely 0.6% of the domain, and therefore the thickness of the red region in these images of more concentrated H-rich gas is exaggerated slightly by being seen in projection through the transparent region of pure H-rich gas. These images clearly show the strong decrease in H-rich gas concentration as we go away from the wedge tip and away from the top of the convection zone. Nevertheless, the turbulent eddies in this region bring about mixing, and there is also some small amount of entrainment from the transition layer even well away from the wedge tip. Further down, as the horizontal velocity further decreases and the inward radial velocity component further increases, this leads to elongated downward entrainment curtains that carry clouds of material somewhat enriched with the "H+He" fluid from above down into the convection zone (Figure 7).

The results of our measured amounts of entrainment in our series of runs with increasing grid resolution, given in Section 3.3, indicates that the extent to which we have resolved this process in the highest resolution simulation that is shown in Figure 11 is likely to be sufficient to compute the time and space averaged entrainment rate with reasonable precision.

As the fluid elements enriched with "H+He" fluid are dragged down inside the entrainment curtain, horizontal turbulence soon disperses and dilutes the "H+He" fluid, and its fractional volume falls below the rendering cut-off chosen for the visualization. At that point the "H+He" fluid becomes invisible in Figure 7. The cumulative entrainment evolution is shown by means of spherically averaged abundance profiles in Figure 3 and by zoomed-in views of spherically averaged entropy profiles in Figure 12. In Figure 3 we have transformed the fractional volume, which is the primary quantity to express abundances in the hydrodynamics code, into mass fractions, which are

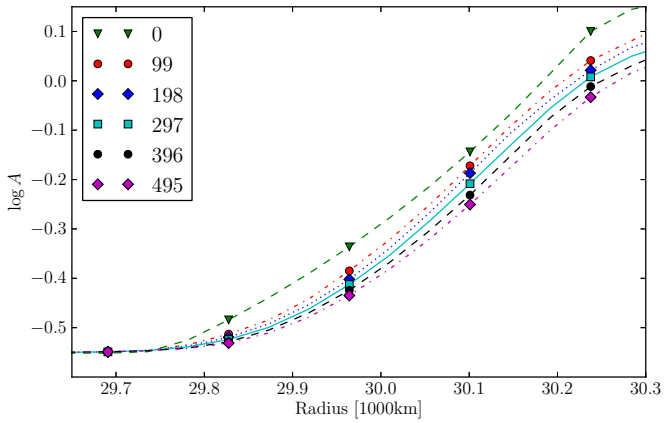


Figure 12. Entropy (Equation (5)) evolution at the top of the convection zone for the 1536^3 simulation. Logarithm of entropy in the transition layer for the initial setup (dump 0) and subsequent evolution (labels in minutes simulated time). A line mark is placed every third grid point.

customary in stellar evolution codes. Even at the end of the simulation, after 12 convective turnover times (24 convection timescales), the abundance profiles show significant variations in the lower part of the convection zone. These features are also highly variable in time and represent individual clouds of "H+He"-enriched material to arrive at the still mostly pristine lower part of the convection zone.

We can summarize that the entrainment process in our simulations consists of three closely interacting components. Large-scale convective upwelling motions create fast horizontal flows over big patches of the convective boundary surface. The horizontal convective flows are initially stable against Kelvin–Helmholtz instabilities. However, where such opposing flows meet, the boundary layer separates from the stiff convective boundary and Kelvin–Helmholtz instabilities are encountered in this highly unstable environment. Finally, the interacting and colliding horizontal gusts lead to coherent downward flows that contain "H+He" fluid. Therefore, the entrainment process is a combination of the local properties at the convective boundaries and the global properties of the convection that depend on all aspects of the particular situation throughout the convection zone.

3.3. Entrainment Simulations for a Range of Grid Sizes

An important concern in any multi-dimensional simulation is the dependence of results on the adopted grid resolution. Here we focus on a qualitative discussion of the velocity field and a more quantitative analysis of the entrainment process. But first we would like to comment on the effect of numerical viscosity.

Numerical viscosity. One might expect the numerical viscosity in the simulation to affect the details of the entrainment process. The image shown in Figure 2 has 1121^2 pixels, and the underlying simulation has a 1536^3 grid. Smallest features that can be identified have a scale of ≈100 km, which corresponds to only a few grid zones. The train of Kelvin–Helmoltz waves in the boundary-separation layer (Section 3.2) has a typical length of 1000 km but the width of this feature is only a fraction of that. In Figure 11 we show zoomed-in views of such boundary-separation layers. The double arrows indicate $\Delta r_{\text{H-shell}} = 500$ km, the width of the transition layer over which the mixing fraction of hydrogen-rich gas changes from nearly zero to nearly unity. From these close-up views we can see that the simulation's 1536^3 grid is resolving vorticity features that are considerably smaller than 11 grid cells. The double resolution views of the "H+He" concentration, made possible by the subcell resolution of the PPB advection scheme (see the Appendix) show that still smaller features in the advected concentration can be made out. This subcell resolution feature of PPB advection can be seen in the simple advection test problem shown in Figure 2 in Woodward et al. (2008a).

Detailed simulations of shear layer instability in isolation, as well as our earlier simulations of this sort of convection-driven shear in small sections of such a flow (Woodward et al. 2008a) indicate that the development of the unstable layer tends to converge when we are able to resolve the largest relevant

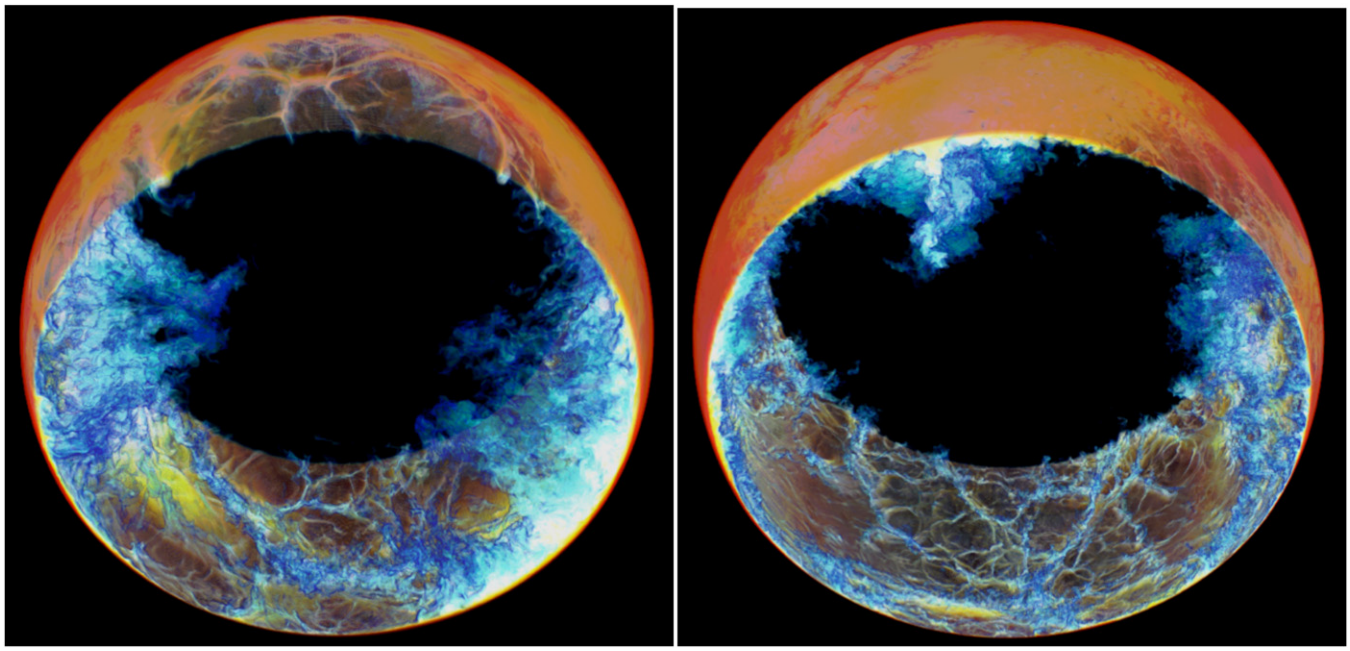


Figure 13. Logarithm of fractional volume of "H+He" fluid shown for a slice of 3D domain at $t = 491$ minutes. The smallest visible values are 10^{-6}–10^{-7}. The blue descending plumes trace out the convection cells. Left: 768^3 grid, right: 1536^3 grid.

scales as well as scales of wavelengths smaller by about a factor of three. This is just a rule of thumb, of course. Only quantitative convergence under grid refinement of a quantity of interest, which in our present investigation is the entrainment rate for the hydrogen-rich gas, can be our guide to the adequacy or inadequacy of our computational grid. The appearance of resolved swirls on multiple scales in the figure is an indication that our grid may be fine enough, but it is the actual observed convergence of the entrainment rate that is our more reliable guide. The principal causes of error in the entrainment in our numerical simulation are the limitations imposed by the grid cell size on the smallest resolvable overturning waves that can produce mixing.

The behavior of PPM gas dynamics simulations of such complex, nonlinear, turbulent flows, with and without PPB advection of multifluid volume fractions, has been studied in comparison to matching simulations with Navier–Stokes viscosity terms added in Sytine et al. (2000) and Woodward et al. (2010b).

Porter & Woodward (1994) carefully investigated the numerical viscosity of our PPM gas dynamics scheme and quantified it as consisting of two terms, the effects of each of which tend toward zero as a disturbance wavelength increases. The effects of physical, Navier–Stokes viscosity tend toward zero as λ^{-2}, where λ is the disturbance wavelength. The actual viscosity in the star is smaller than could be represented in our simulation. Instead, it is the numerical viscosity of PPM that is dominant in the simulation. This numerical viscosity tends to zero as λ^{-3} or λ^{-4}, depending upon which of our two numerical terms is larger under the particular circumstances. For the numerical viscosity, it is the disturbance wavelength measured in grid cell widths that matters. When doubling the grid resolution the impact of the numerical viscosity on the flow, as measured by effective Reynolds numbers for wavelengths of physical importance, is diminished by a factor of four or a factor of eight, depending upon which of our two terms dominates. Therefore, when our results converge under grid refinement, it must be true that our numerical viscosity no longer matters. The physical reason why this can happen is that the entrainment due to breaking Kelvin–Helmholtz waves (Section 3.2) is not a viscous phenomenon. Viscosity determines the wavelengths at which shear instabilities are damped away, but it is the largest wavelengths, largely unaffected by viscous damping, that are chiefly responsible for the entrainment. These we may hope to capture in a well-resolved simulation. Our results in Section 3.3 indicate that our finest grid of 1536^3 cells has just marginally captured these important waves accurately. One could attempt to double the grid resolution once more to lay any doubt on this point to rest, but we have judged that the cost of doing this on present computing equipment makes this more definite demonstration of convergence impractical.

Radial profiles of velocity. The overall velocity scale, shown for the total rms velocity in Figure 5 increases somewhat with resolution. Inspection of the separate components of spherically averaged velocities reveals that this difference is mostly in the horizontal velocities that span a range between 12 km s^{-1} for the 512^3 grid and 17 km s^{-1} for the highest resolution, while the radial velocities for all resolutions fall in a narrow band between $v_{\rm radial,rms} \approx 8$ and 10 km s^{-1} (Figure 14). The overall difference between horizontal and vertical velocity components is larger for higher resolution. However, as Figure 8 shows, the difference between the horizontal and radial velocity profile shape is larger for lower resolution. In particular, the 768^3 grid horizontal velocity profile shows a much more pronounced peak

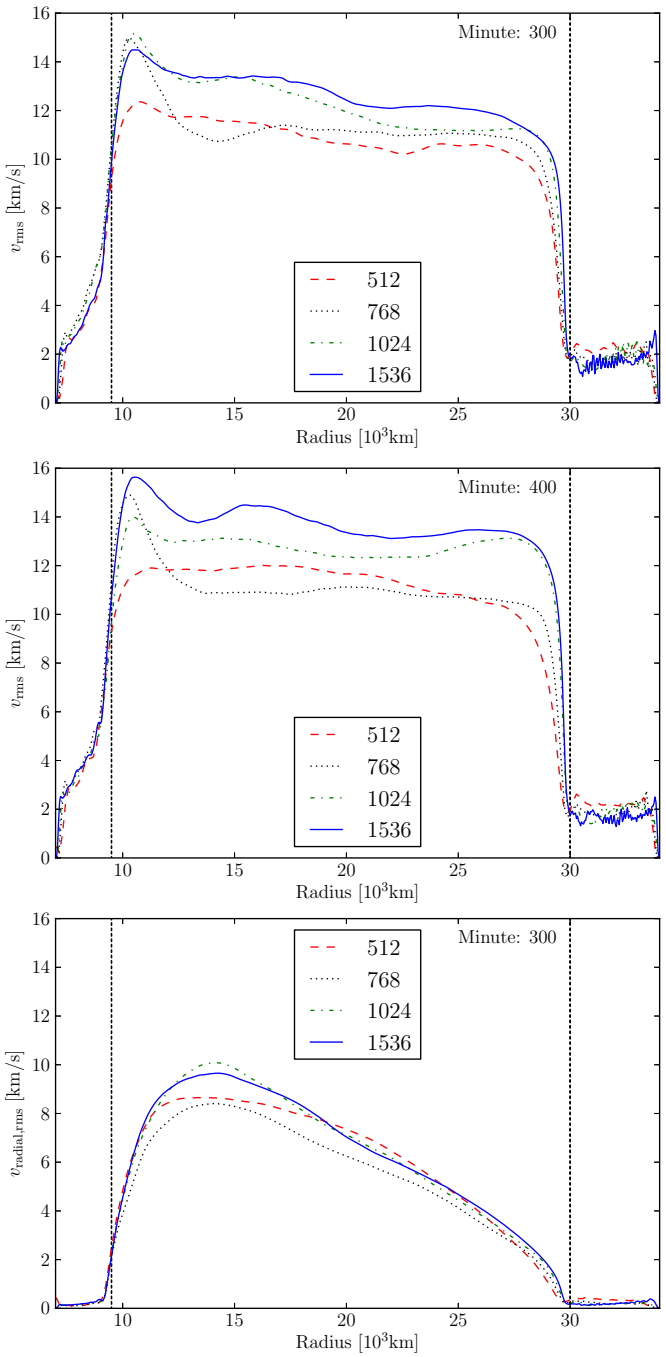


Figure 14. Spherically averaged rms velocity for 300 and 400 minutes (not averaged in time), and just the radial component for 300 minutes, for all four grid size simulations. The dashed vertical lines are placed at the convective boundaries according to the initial setup.

at the lower convective boundary and a significant depression just above the heating zone where the radial velocities have a maximum.

Figure 8 also shows that the gradient of the horizontal velocities becomes steeper for the higher resolution runs. In the upper part of the convection zone the total rms velocity is dominated by the horizontal component. This steeper horizontal velocity gradient can therefore also be seen in Figure 14 at different times. In the inner part of the convection zone the velocities at a given radius are not always strictly correlated with resolution. This is partly due to long-term convective breathing

modes that are out of phase for the different runs, as well as chaotic fluctuations. However, near the top boundary, where the horizontal velocity drops sharply, we always found velocity profiles to be strictly ordered by resolution, with the highest velocity and the steepest velocity gradients found for the highest resolution runs.

Entrainment process and rate. Eventually we are interested to study the burning of the entrained H-rich fuel via the $^{12}C(p,\gamma)^{13}N$ reaction and the feedback of energy from that process into the hydrodynamic flow. The amount of entrainment will determine the timescale and type of feedback from the energy release of the nuclear burning. This, in turn, will determine the nucleosynthesis in reactive-convective environments, including the n-capture process in *i*-process conditions (Herwig et al. 2011). The entrainment rates will be affected by the hydrodynamic feedback from the ingested H. The entrainment rate and the way in which the entrained H will be advectively distributed throughout the He-shell flash convection zone will determine the properties of the subsequent convective–reactive H-combustion, its feedback into the convective flow and the conditions for nucleosynthesis.

A first test is the visual inspection and comparison of 3D images of the entrainment process for simulations with different resolutions. We find no obvious artifacts or differences in scales or patterns when comparing a 768^3 and 1536^3 simulation (Figure 13). In both cases the large-scale upwelling flows that turn around at the upper boundary of the convection zone cover large fractions of a hemisphere, just as described in the previous section. Also the scraping horizontal motions associated with boundary-layer separation of converging horizontal flow at the convective boundary are evident in both cases in very similar patterns. If anything, the interfacial structures and the entrainment curtains have a somewhat smaller scale in the higher resolution case. This is consistent with the steeper gradient of the horizontal velocity found in the higher resolution runs (see above).

We have described above how the entrained "H+He" fluid is cumulatively filling the convection zone (Figure 3). In Figure 15 we show the spherically averaged fractional volumes[4] for all four grids at three times. The entrained mass can be determined from these spherically averaged abundance profiles by integrating over the convection zone. The profiles already reveal that the entrained mass is similar in the two highest resolution runs, and that the lowest grid case (512^3) has at all times a significantly larger entrained mass. This pattern prevails through all times, which is evident from the animation view of these profiles available at http://www.lcse.umn.edu/3Dstar-convection-entrainment. The profiles, especially in the animated presentation, are perhaps the strongest indication of convergence of the entrainment rate in our simulations. We can attempt to summarize this finding in the following way.

The entrainment rate is just the amount of mass of fluid "H+He" mixed into the convection zone per unit time. The entrained mass is shown for the four grid choices as a function of time in Figure 16. The entrainment rate is therefore the slope of these lines.

Figure 16 is showing in another form the same result as the profile evolution. With the two smaller grids the entrainment

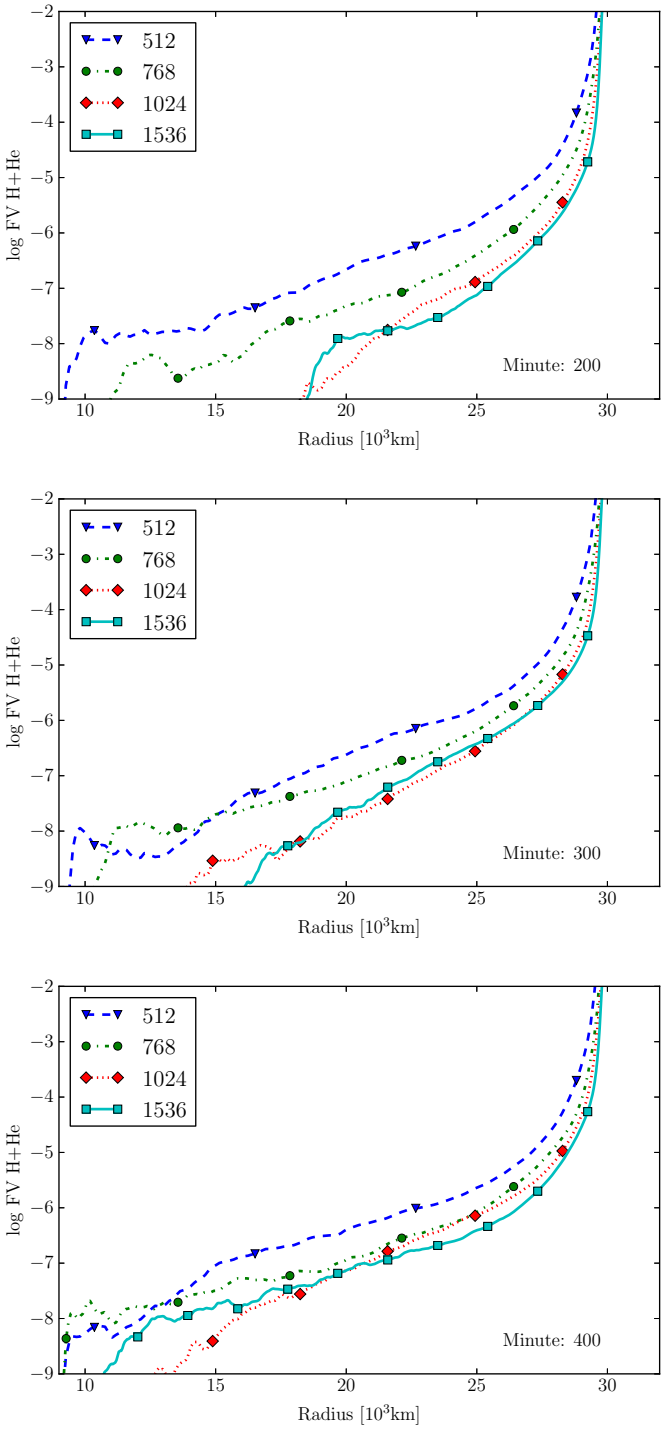


Figure 15. Profile of the entrained fractional volume of the H+He fluid for the different grid resolutions at (from top to bottom) 200, 300, and 400 minutes.

rate is much higher compared to the high-resolution runs. A larger amount of entrained H would lead to a large H-burning luminosity, and this seems to be the case in the simulations presented by Stancliffe et al. (2011, Figure 3) in which the low-resolution run has a 1000 times higher H-burning luminosity compared to the high-resolution case.

The entrainment proceeds in a gusty or intermittent manner, with episodes of larger than average entrainment for periods of ≈30–60 minutes followed by phases of smaller entrainment rates. The irregular periodicity of the entrainment fluctuation is of the order of 100 minutes. When subtracting

[4] The fractional volume FV_1 is the primary variable to represent concentration of fluid 1 in the PPMstar hydrodynamics code. It is related to the mass fraction X by $X_1 = \rho_1 * FV_1/\rho$, where ρ_1 is the density of fluid 1 and ρ is the density.

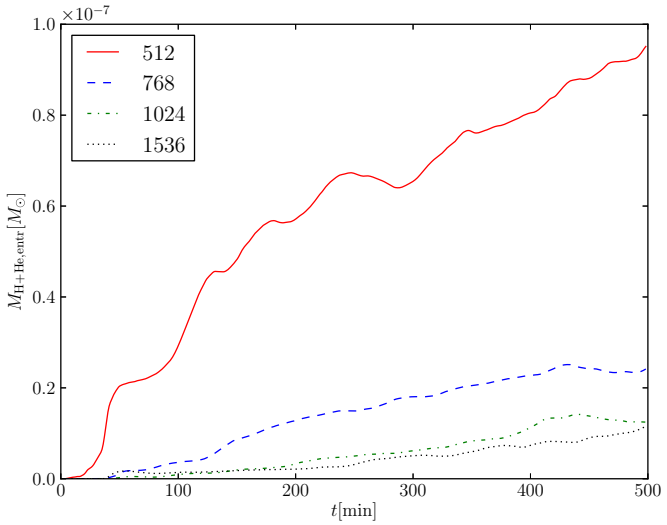


Figure 16. Entrained mass, integrated between the radial coordinates 10,100 km and 29,200 km, as a function of time. Simulations for different grid sized n^3 are shown. The number of grid points n in one direction is shown in the legend.

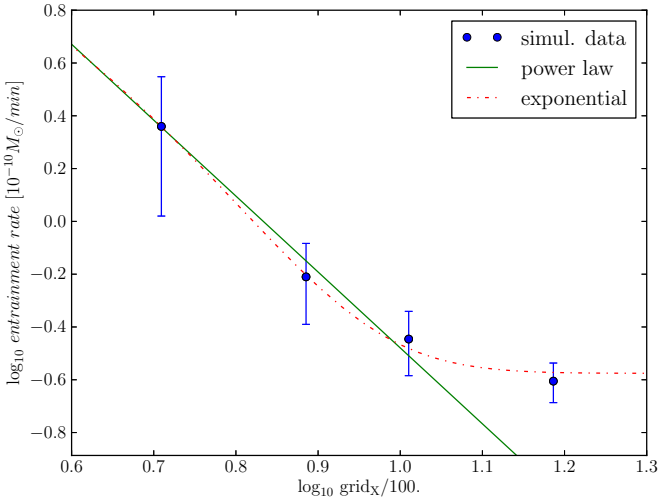


Figure 17. Logarithm of mean entrainment rates with error bars representing 99% confidence intervals (see the text) as a function of grid size. The mean entrainment rates have been fitted with a power law and an exponential according to Equation (7).

Table 1
Entrainment Rates ($10^{-10}M_\odot$ minute^{-1}) for each Run with Different Grid Resolution, and the Asymptotic Entrainment Rate under the Assumption of a Fit According to Equation (7)

Grid Size/Case	Mean	Std. Dev.	99% CI
512^3	2.290	1.668	±1.2420
768^3	0.616	0.281	±0.2089
1024^3	0.358	0.132	±0.0979
1536^3	0.248	0.057	±0.0424
$c_{\rm entr}$	0.263	0.089	±0.0659

Note. $c_{\rm entr}$ is the asymptotic limit variable given in Equation (7).

the first ≈150–200 minutes, which represent the initial transient and the settling into a steady-state convective flow, we are only left with ≈300 minutes that can be used for the entrainment analysis, since the highest-resolution run (1536^3) has been computed only to 498 minutes. Therefore, the time averaging of the entrainment rate can be performed only over a small multiple of the typical period over which the entrainment rate fluctuates, which introduces a statistical error in determining the entrainment rates for each grid size. One of the simulations, the 1024^3 run, has a particularly strong entrainment gust from $t = 410$ minutes to $t = 460$ minutes preceded and followed by a period of lower than average entrainment. Such individual events, and the way in which they are included or excluded in the integration will inevitably introduce some error.

This error is estimated by performing 12 fits of the entrainment rates with different assumptions on the selection of the range of cycles over which a linear least-square fit of the entrained mass as a function of time is performed to determine the entrainment rate, the choice of the weight of the zero-point of the entrainment evolution as well as the upper boundary of the integration (see discussion at the end of this section). Together these 12 cases represent the range of reasonable choices that could be made when integrating the entrained mass, and they represent the sample of measurements on which the measurement error is based. The resulting entrainment rates are provided in Table 1 and plotted in Figure 17.

We adopt two simple expressions that represent the non-convergence and the convergence case. If the entrainment rate does not converge for the range of grids adopted here it may follow a power law $f_{\rm pow}(x) = ax^b$ as a function of the number of grid points in one direction $x = {\rm grid}_X$. However, if the entrainment does converge, the entrainment rate would instead show signs of approaching an asymptotic limit. In this case, the dependence of the entrainment rate on the grid may be represented by an exponential plus a constant of the form

$$f_{\rm exp}(x) = ae^{bx} + c_{\rm entr}\ . \qquad (7)$$

In that case the constant $c_{\rm entr}$ would represent an estimate of the true entrainment rate based on the assumed functional form of Equation (7). Figure 17 shows both alternatives fitted to the entrainment rate data provided in Table 1. The data is better represented by the exponential function, and the asymptotic limit approximated by $c_{\rm entr}$ agrees within the adopted uncertainty with the entrainment rate of the 1536^3 simulation, and represents the converged entrainment rate (for our assumption on the integration boundaries).

We repeated the same procedure just with the three lower resolution runs, and find in that case $c_{\rm entr} = 0.311$. This is within the 99% CI of the asymptotic entrainment. But the fits do not convincingly show convergence by a clear preference for the exponential fit. However, if we use the knowledge that the simulations do converge we could obtain, for a similar case, a very good result with ≈1/5th of the computational cost, which is the fraction of the cost of the three lower resolution runs combined compared to the cost of all runs.

Limitations of the convergence analysis. The entrained mass shown in Figure 16 is determined by integrating up to 29,200 km. This is 550 km below the formal convection boundary defined as where in the initial setup the entropy gradient becomes positive. This choice is necessary to exclude the actual entrainment interface from the integration. This interface is not resolved sufficiently in the lower-resolution runs. The horizontal velocity profile, for example, depends on the resolution (Figure 14). At the same time the region nearest to the formal convection zone has the highest abundance of the "H+He" material from the stable layer, and would therefore dominate the

integration of the entrained material. Including the boundary layers would therefore rather measure the amount of material present in the 550 km or so within the range of the boundary. We cannot show convergence in the way described above for the material in the 550 km layer inside the convection boundary. However, that material will not react with ^{12}C and release energy. We are therefore not interested whether or not the amount of material in the boundary layer is converged. Instead, we only ask if the material that actually enters the region below this boundary layer, i.e., below 29,200 km can be quantitatively simulated. That, indeed, can be accomplished with simulations at the resolutions used here.

4. DISCUSSION AND CONCLUSIONS

The main goals of this paper are to investigate the properties of He-shell flash convection in 4π geometry and determine the entrainment properties of H-rich material at the top of the convection zone. This entrainment process is the basis of the H-ingestion flash in low-Z AGB stars and in post-AGB stars, and the results of this investigation are the foundation of our future work on this topic (see Herwig et al. 2014 for preliminary results). The top convection boundary in this He-shell flash convection situation is much stiffer compared to the O-shell burning convection or the core convection investigated by Meakin & Arnett (2007), as visual comparison of their Figures 3 and 10 and our Figure 7 immediately reveals. It is therefore very consistent that their entrainment rates for O shell burning ($1.1 \times 10^{-4}\,M_\odot$ s^{-1}) and core convection ($2.72 \times 10^{-7}\,M_\odot$ s^{-1}, with a driving luminosity 10× natural) are much larger compared to our converged entrainment rate of $4.38 \times 10^{-13}\,M_\odot$ s^{-1}. There exists a significant range of entrainment rates, even between convection boundaries of different deep stellar interior convection regimes. He-shell flash convection boundaries with H-ingestion are much stiffer compared to core convection and in turn O shell burning. Large entrainment rates have been reported as well by Mocák et al. (2010) based on a single 3D simulation.

In order to demonstrate that the fidelity of our approach matches the challenge of the problem, we describe key elements of our method that we will apply in future work. The key capability that we demand from a simulation is to reproduce the total amount and rate of entrainment of material from above the convection zone into the unstable layer. It will be this entrainment amount that, through the nuclear reactions between protons and ^{12}C, generates energy at a rate that is large enough to alter the global flow properties. The accretion of unburned material through the upper convection boundary of the He-shell flash convection zone will therefore determine all subsequent H-ingestion flash simulation steps. Our simulations show that despite the enormously small entrainment rate, the H-rich material that enters the He-shell flash convection zone is advected into the deeper layer in distinct and coherent clouds. The visual inspection of the fractional volume images demonstrate that these clouds reach the burning layer where protons will react with ^{12}C. This is an important result, as without our simulations one may have plausibly assumed that strong horizontal turbulence would shred any downward directed H-rich advection. In such a case, a spherically symmetric simulation approach of the H-ingestion flash may have seemed justifiable. Our simulations, however, show significant inhomogeneities in the depth of the convection zone where the ^{12}C + p reaction would take place. Our results therefore provide evidence that global 4π simulations are indeed needed to simulate the H-ingestion flash accurately.

How does the entrainment rate we obtain compare to the H-ingestion flash constraints we already have established for Sakurai's object? Herwig et al. (2011) demonstrated that the anomalous observed abundance distribution could be reproduced in the 1Dimensional multi-zone nucleosynthesis models with a static mixing-length based convective diffusion parameter if an entrainment rate of $5.3 \times 10^{-10}\,M_\odot$ s^{-1} is assumed. The hydrodynamic entrainment rate that we have determined is significantly smaller than the nucleosynthesis-constrained entrainment rate. We may therefore conclude that the hydrodynamic feedback from the nuclear energy release due to H ingestion will lead to enhanced entrainment rates. Herwig et al. (2011) did also find that the neutron flux must be terminated at some point, and speculated that this would be due to the hydrodynamic feedback mechanism induced by the nuclear energy release from the ingested protons.

The ingestion rate of protons that is necessary to cause significant flow variations to the convective velocity field established by the underlying He-burning luminosity can be estimated by noticing that for this to happen the luminosity from ^{12}C + p reactions must equal or exceed the He-burning luminosity $L_{\rm He} = 4.2 \times 10^7\,L_\odot$. This is the case if

$$\dot{M}_{\rm entr} \gtrsim \frac{L_{\rm He} m_{\rm p}}{Q} = 4.9 \times 10^{-11}\,M_\odot\ {\rm s}^{-1}, \qquad (8)$$

where $m_{\rm p}$ is the mass of the proton and Q is the energy released in each ^{12}C$(p,\gamma)^{13}$N reaction. Our entrainment rate is significantly lower than this limit, and we therefore do not expect an immediate and dramatic effect of the burning of the ingested protons. In fact, one may wonder how any relevant hydrodynamic feedback can arise with such a low entrainment rate. As we will demonstrate in detail in a forthcoming publication H-ingestion entrainment simulations for Sakurai's object with energy feedback from H ingestion taken into account show a long quiescent phase of entrainment and nuclear burning initially without noticeable effect on the hydrodynamic flow pattern, followed eventually by violent and non-spherical eruptions that significantly elevate the entrainment rate. The details of these events will be reported elsewhere.

A key property of the convective boundary layer is the entropy profile immediately above the Schwarzschild boundary. In Figure 12 we observe a self-similar evolving entropy profile. Thus, the numerical method is able after a short while to establish an entropy profile that does not appreciably change shape as it translates in response to the entrainment.

We give evidence that the entrainment rate is a quantitative simulation property that shows convergence. However, some other quantities cannot be considered to be converged. It seems that the spherically averaged radial velocity components are rather similar for all resolutions, while the horizontal velocity components still show some dependence on resolution. Depending on how the horizontal component is averaged, we roughly estimate that horizontal velocity components may differ by 30% between our runs with lowest versus the highest resolution. Although it would be desirable to get rid of these differences as well, one should keep in mind that the modeling uncertainties introduced by the traditional 1D spherically symmetric simulation approach based on the MLT and some CBM algorithm have been shown to be much larger and even qualitative in some instances, as demonstrated via a nucleosynthesis-based validation analysis by Herwig et al. (2011).

The fact that there are aspects of the simulations that are not fully converged is interesting but not important, as long as we can show that convergence is or can be obtained for those quantities that do matter. In our case this is the entrainment rate. Based on this result, we may now move ahead and include the nuclear burn of H with ^{12}C and the associated energy release.

P.R.W. acknowledges support from contracts from the Los Alamos (LANL) and Sandia National Laboratories and NSF PRAC grant OCI-0832618. The development of the PPB scheme for multifluid fractional volume advection was supported by PRW's LANL contract and DoE grant DEFG02-03ER25569. F.H. acknowledges NSERC Discovery Grant funding. The hydrodynamics simulations were performed on the LCSE workstation cluster, supported by NSF CRI grant CNS-0708822, on NSF's Kraken supercomputer at the National Institute for Computational Sciences (NICS), on the Itasca machine at the Minnesota Supercomputing Institute (MSI), and the Canadian WestGrid Lattice and Orcinus high-performance computers.

APPENDIX

THE PPB MULTIFLUID ADVECTION SCHEME

A.1. PPM Gas Dynamics

Our simulation code for the hydrogen ingestion flash is based on the Piecewise-Parabolic Method (PPM; Woodward & Colella 1981, 1984; Colella & Woodward 1984; Woodward 1986, 2006). We use a version of this numerical scheme that has been modified in a few minor, yet substantive, ways over the years. It is described in complete detail in Woodward (2006), and the effective numerical viscosity of PPM has been quantitatively characterized by Porter & Woodward (1994). The most relevant modification of this scheme from the present perspective is that in the interpolation process that determines the parabolae we use to describe the subcell structure, monotonicity constraints are applied only if the local behavior of the function to be interpolated is judged not to be sufficiently smooth. A measure of function smoothness is constructed that has a five-cell stencil in the direction of the present 1D pass. In our low Mach number flows, we do not expect to find shocks, so that we expect the flow to be smooth under most conditions. In this case, we do not expect the elaborate interpolation of the values of the function at each of the two interfaces of the cell with its nearest neighbors in the direction of the 1D pass to be modified in any way by monotonicity constraints. These cell interface values are interpolated by evaluating at these points the unique cubic polynomials that assume the prescribed cell averages in the two grid cells to the left and the two to the right of the interface in question. These interface values are therefore one full formal order more accurate than the parabola used to describe the distribution of the interpolated variable inside the grid cell. That parabola, when not modified by monotonicity constraints, passes through the two cell interface values and has the prescribed cell-averaged value. To find the amount of this variable that is advected across an interface of the grid cell, we use this parabola in the upstream cell. For low Mach number flows, only a thin sliver of the cell is advected across the interface. Therefore the average value of the interpolated variable within this sliver is nearly equal to the value of the variable at the interface. Because this value is one order more accurate than the parabola as a whole, the PPM scheme in low Mach number flow regimes has the very important property that it becomes more accurate as the Mach number, and hence the Courant number, is decreased. Even though more time steps are then required to arrive at a given time level, the higher formal accuracy completely offsets the loss in accuracy that we would have expected at low Courant numbers. This behavior is counter to the behavior of most numerical schemes. It is a feature of PPM that is a result of its history. PPM was developed from an earlier, more accurate numerical scheme in which the cell interface values were stored and updated as primary data for the method. Because this data structure was incompatible with large production codes at the Lawrence Livermore National Lab in the late 1970s, the PPM scheme was designed as a replacement. It generated the previously independent cell interface data by interpolation from the cell averages. To retain as much of the earlier method's accuracy as possible, this interpolation of cell interface values was made as accurate as seemed practical or necessary at the time.

In our stellar hydrodynamic flows, we are concerned that structure in the flow with a characteristic scale size of 20–40 grid cells can be advected in our circulating convection zone for, say, 100 scale lengths with minimal amplitude damping and acceptable phase error. We can assess this capability using 1D advection of a sine wave of 40 cells. Using a Courant number of 0.03125, advecting this sine wave for 100 wavelengths with our PPM scheme, we find that the wave amplitude is damped by a factor of 0.9985, while the final phase is off by 0.078 cell widths, which is just 1 part in 50,000. A more stringent test is to advect a Gaussian pulse with a FWHM of 10 grid cells a distance of 100 times this width, or 1000 cell widths, with a Courant number of 0.03125. Doing this results in a pulse with height damped by a factor of 0.9383 and with a phase error of just 0.723 grid cell widths, which is just 1 part in 1384. The behavior of the scheme with a Courant number of one-third has slightly greater amplitude error but less phase error. Running the same sine wave and Gaussian advection tests at this Courant number, we find amplitude damping by factors of 0.9923 and 0.9320 with phase errors of 0.02 and 0.234 grid cell widths. It should therefore be clear that with PPM advection we suffer very little accuracy loss in going from the typical range of Courant numbers encountered in explicit gas dynamics simulations to those we encounter in our present simulations of the hydrogen ingestion flash.

A.1.1. PPB Multifluid Volume Fraction Advection

Our multifluid PPM code uses the much more accurate Piecewise-Parabolic Boltzmann (PPB, described below) scheme to advect the critically important cell volume fraction occupied by the hydrogen-rich fluid, which is located initially just above the top boundary of the helium shell flash convection zone. For comparison with PPM, PPB running the same advection experiments just mentioned at a Courant number of 0.03125 for the same sine wave and Gaussian produces damping by factors of 0.999937 and 0.99961 with phase errors of 0.938 and 0.02 cell widths. The PPB phase error may seem large until one realizes that, as discussed in van Leer (1977) and Woodward (1986), PPB incurs one-time errors related to going over to the piecewise-parabolic representation, after which there is extremely little error accumulation. Consequently, advecting the same sine wave at the same small Courant number for an additional 100 wavelengths causes the phase error to increase from 0.938 cell widths after 100 wavelengths to only 1.066 cell widths after a second 100 wavelengths. As a general rule of thumb, PPB advection is about as accurate as PPM advection if PPM is given

three times more grid cells (and three times more time steps) to work with. This is because in 1D, PPB has three times as much independent information on any given grid. PPB advection has fifth-order formal accuracy (van Leer 1977; Woodward 1986), while PPM advection is only third-order accurate. Nevertheless, our rule of thumb holds for practical problems, since we are then essentially never in the range of grid resolutions where strict formal convergence rates hold. For the advection of sharp jumps, PPB cannot yield a thinner numerical representation of the jump than PPM, with its contact discontinuity detection and steepening because PPM's representation is the thinnest possible when parabolae are used for the subgrid structures. Nevertheless, PPB advection is still superior even in this case because PPB generates far fewer numerical glitches in rare pathological circumstances. Once again, this result is due to PPB having much more independent data to work with.

PPB achieves its very high accuracy by conserving to machine accuracy the 10 lower-order moments—$\langle f\rangle$, $\langle f\tilde{x}\rangle$, $\langle f\tilde{y}\rangle$, $\langle f\tilde{z}\rangle$, $\langle f\tilde{x}\tilde{x}\rangle$, $\langle f\tilde{x}\tilde{y}\rangle$, $\langle f\tilde{x}\tilde{z}\rangle$, $\langle f\tilde{y}\tilde{y}\rangle$, $\langle f\tilde{y}\tilde{z}\rangle$, $\langle f\tilde{z}\tilde{z}\rangle$—of the fractional volume, f, of a tracked fluid within a local volume element. We define these moments with respect to a set of Cartesian coordinates, $\tilde{x}$, $\tilde{y}$, and $\tilde{z}$, centered on the grid cell of interest, aligned with our grid directions, and for which the width of the cell in each dimension is unity. Using these cell-centered coordinates, we define the moments as

$$\langle f\tilde{x}^l\tilde{y}^m\tilde{z}^n\rangle = \int_{-1/2}^{1/2} f\ \tilde{x}^l\ \tilde{y}^m\ \tilde{z}^n\ d\tilde{x}\ d\tilde{y}\ d\tilde{z}.$$

Our use of the cell-centered and normalized coordinates restricts us to the use of cubical grid cells. However, given the proliferation of adaptive mesh refinement techniques, we view this restriction as entirely acceptable for all practical calculations. For our work with stars, cubical grid cells provide the least distortion of the underlying flow geometry that is possible with a Cartesian mesh. The centered and normalized coordinates give us a tremendous advantage; they allow us to perform our calculations at double speed in 32 bit precision. On the very fine grids permitted by the tremendous power of today's computing systems (we have run problems related to inertial confinement fusion on over one trillion grid cells; Woodward et al. 2012), it can easily consume four digits of precision just for a grid cell to know where it is located in a global sense. Injecting this global location through moments of x, y, and z in place of $\tilde{x}$, $\tilde{y}$, and $\tilde{z}$ is simply wasteful. It forces the machine to do twice as much work for no good reason, since the global location of a grid cell is irrelevant in specifying the distribution of a variable within it. One might ask if the fifth-order formal accuracy of the scheme, to whatever extent it really matters in a practical computation, is preserved in 32 bit precision. Our experience indicates that it is, with the single exception of the IBM Cell processor, which did not perform rounding in its 32 bit arithmetic.

Our PPM scheme is directionally split. We update the flow for each time step in three 1D passes, using a symmetrized sequence *xyzzyx* in each time step pair. We therefore require a directionally split version of the PPB advection scheme. This delivers an immense simplification. The PPB scheme is built upon van Leer's Scheme VI (van Leer 1977), which is a 1D scheme with no monotonicity or other constraints. In the early 1980s, this scheme was made into a two-dimensional (2D), directionally split, and constrained scheme, described in Woodward (1986). The version of PPB that we use in our work is simplified from this early work, with a minimal set of moments and a simplified method of updating them, described below. It is also enhanced by a more elaborate and useful method of constraining each 1D pass of the algorithm. This scheme has only been described in detail in internal reports (Woodward 2005) and in broad terms (Woodward et al. 2008a, 2010b, 2012) to this date. A code module containing this PPB scheme for 3D computation was delivered to the Los Alamos XRAGE code in 2004, and this module, combined with our version of PPM, was included in the official 2005 release of the Los Alamos XRAGE code. Here we give a complete, but brief, description of the PPB advection scheme, adapted for use in the advection of a fluid concentration, so that f is forced to remain within the range of values from 0 to 1. A more voluminous description can be found in the LCSE internal report (Woodward 2005). The antecedents of this PPB scheme go back to the 1970s and 1980s, and Woodward has taught it in his course on numerical methods at the University of Minnesota in the late 1980s and 1990s. A parallel development of numerical schemes, beginning around the same time as the work of van Leer but in the finite element community, has been given the name Discontinuous Galerkin (see recent books; Cockburn et al. 2000; Hesthaven & Warburton 2000). These schemes are similar, but they involve more computational labor as a result of their use of Runge–Kutta techniques. They also apply the methodology over the full set of hydrodynamical equations, although there is little evidence of which we are aware that this yields any significant benefit for flows in which shocks and contact discontinuities are involved, so that the very high formal order of these techniques might be justified. Here we apply a simpler form stemming from van Leer's original work (van Leer 1977), and we cut to the bone the computational labor involved. We also apply an elaborate set of constraints that address the special problem of multifluid flow. Despite the complexity of these constraints, our formulation in 1D passes allows us to apply these constraints only to the three moments $\langle f\rangle$, $\langle fx\rangle$, and $\langle fxx\rangle$, where x is the direction of our 1D pass. This delivers an enormous simplification of the scheme, to the extent that it now involves only about three times the computational labor of the much less accurate PPM advection scheme.

We can use the 10 moments of f to construct a quadratic polynomial describing the subgrid behavior of f within a grid cell:

$$f(\tilde{x},\tilde{y},\tilde{z}) = f_{000} + f_{100}\tilde{x} + f_{010}\tilde{y} + f_{001}\tilde{z} + f_{200}\tilde{x}^2 + f_{110}\tilde{x}\tilde{y} + f_{101}\tilde{x}\tilde{z} + f_{020}\tilde{y}^2 + f_{011}\tilde{y}\tilde{z} + f_{002}\tilde{z}^2,$$

where

$$f_{000} = 4.75\ \langle f\rangle - 15(\langle f\tilde{x}\tilde{x}\rangle + \langle f\tilde{y}\tilde{y}\rangle + \langle f\tilde{z}\tilde{z}\rangle)$$
$$f_{100} = 12\ \langle f\tilde{x}\rangle \quad f_{010} = 12\ \langle f\tilde{y}\rangle \quad f_{001} = 12\ \langle f\tilde{z}\rangle$$
$$f_{110} = 144\ \langle f\tilde{x}\tilde{y}\rangle \quad f_{101} = 144\ \langle f\tilde{x}\tilde{z}\rangle \quad f_{011} = 144\ \langle f\tilde{y}\tilde{z}\rangle$$
$$f_{200} = 15\ (12\ \langle f\tilde{x}\tilde{x}\rangle - \langle f\rangle) \quad f_{020} = 15\ (12\ \langle f\tilde{y}\tilde{y}\rangle - \langle f\rangle)$$
$$f_{002} = 15\ (12\ \langle f\tilde{z}\tilde{z}\rangle - \langle f\rangle).$$

From this single polynomial, we can construct the following separate functions of $\tilde{x}$:

$$f_x(\tilde{x}) = f_{000} + f_{100}\tilde{x} + f_{200}\tilde{x}^2, \quad f_y(\tilde{x}) = f_{010}\tilde{y} + f_{110}\tilde{x}\tilde{y}$$
$$f_z(\tilde{x}) = f_{001}\tilde{z} + f_{101}\tilde{x}\tilde{z}$$
$$f_{yz}(\tilde{x}) = f_{011}\tilde{y}\tilde{z} \quad f_{yy}(\tilde{x}) = f_{020}\tilde{y}^2 \quad f_{zz}(\tilde{x}) = f_{002}\tilde{z}^2.$$

In a 1D pass, we perform separate advection operations on the above functions of $\tilde{x}$. Only the first of these six functions,

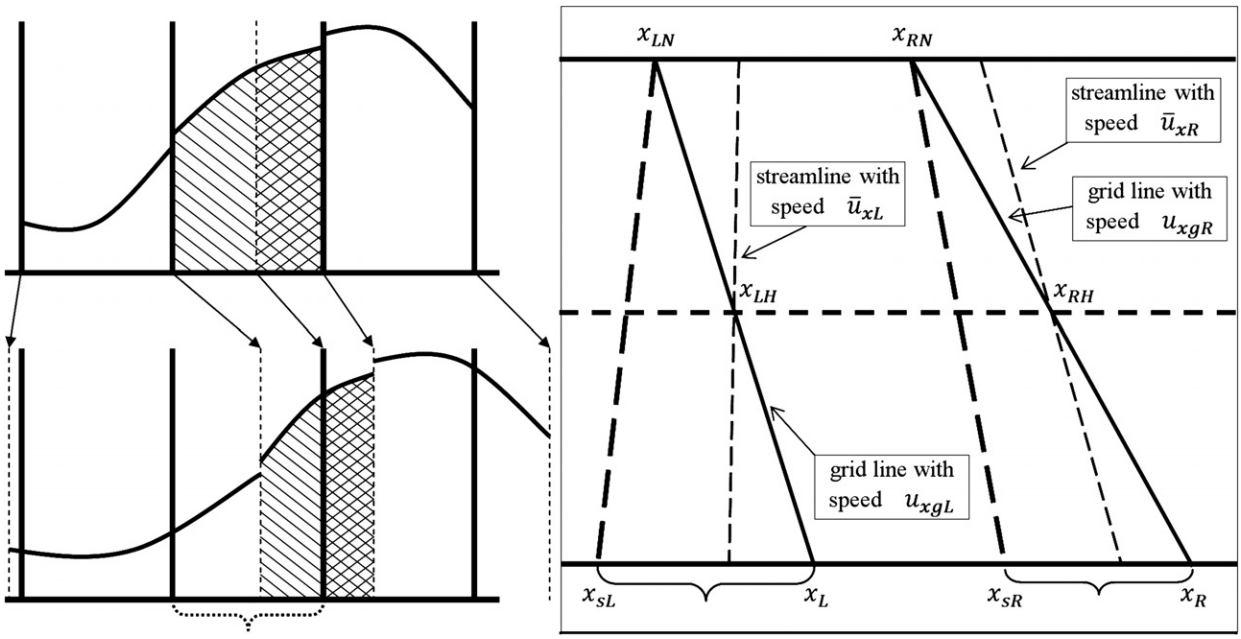


Figure 18. Left panel: the interpolation parabolae, $f_x(\tilde{x})$, for three grid cells are shown at the beginning of the x-pass in the upper part of the figure. The motion of the cell interfaces is indicated, and the new stretched or squashed parabolae are shown in the lower part of the figure. The two portions of the central cell that become parts of the new central cell and of its neighbor on the right are indicated by the diagonal and cross-hatched shading patterns in both parts of the figure. To obtain the new interpolation parabola for the central cell, we must evaluate the moment integrals over the cell domain in x, which is indicated by the bracket at the bottom of the figure. Right panel: space–time diagram showing time levels, moving grid trajectories and fluid streamlines, with coordinate and speed values indicated that are used in the evaluation of moment integrals over the advected portions of grid cells, which are marked with brackets.

$f_x(\tilde{x})$, requires the application of any constraints. The other five functions can be advected without concern for constraints because the results will be constrained in later 1D passes in the $\tilde{y}$ and $\tilde{z}$ directions. We will discuss the constraints last, and first give the formulae for updating the 10 moments of f in the cell during an x-pass. We begin with the most difficult formulae, which result from the advection of $f_x(\tilde{x})$. Once we see how to advect $f_x(\tilde{x})$, all the remaining functions will be trivial by comparison, although a slight complication results from the dependence of f_{020} and f_{002} upon $\langle f \rangle$. A further complication arises from changes in volume of the grid cells during the time step. Conceptually, we draw each of our six functions of $\tilde{x}$ in the grid cells according to the values of the coefficients in their polynomial representations given above. Then we demand that the value of each function at a particular point must remain unchanged as that point moves with its time-averaged velocity in x over the time step. This construction is illustrated for $f_x(\tilde{x})$ in the left panel of Figure 18. Finally, we integrate the resulting function segments over the new grid cell volume to obtain the new x-moments for each of our six functions, which are then the new 10 low-order moments of f in the new grid cell.

Many applications require a moving (radially collapsing or expanding) Eulerian grid, such as the rapid expansion of the outer convection zone in a H-ingestion flash during the first thermal pulse of a low-Z AGB star. Although we have not used this feature in the entrainment simulations presented here, it is an integral part of our PPB implementation and we prefer to describe here the general formalism that allows for a co-moving grid, and that will be used in future stellar convection simulations. The diagram in the right panel of Figure 18 is intended to clarify the coordinates that we use. Each grid cell interface has a constant velocity proportional to its distance from the origin. This is indicated in the diagram. Our PPM gas dynamics scheme computes the time-averaged velocity, $\overline{u}_{xL}$, of the fluid at each moving grid cell interface L. This is indicated in the diagram by the slopes of the fluid streamlines crossing the cell interfaces at the time level half-way through the time step (subscript "H"). We denote the new time level, at the end of the time step, by the subscript "N."

We perform these moment integrals with respect to the new grid cell's own cell-centered and scaled coordinates. To do this, we relate the new cell's coordinates to those in the old cell.

$$\begin{aligned}\tilde{x}_N &= (x_N - x_{MN}) \,/\, \Delta x_N \\ &= \left[(x_M - x_{MN}) + \tilde{x}\ \Delta x + (\overline{u}_{xM} + \tilde{x}_H\ \Delta\overline{u}_x)\Delta t\right]/\Delta x_N .\end{aligned}$$

Here we have used the subscript "M" to denote the middle of the cell, so that $\overline{u}_{xM} = (\overline{u}_{xL} + \overline{u}_{xR})/2$ and $u_{xgM} = (u_{xgL} + u_{xgR})/2$. Also we denote the difference in the time-averaged fluid velocity across the cell by $\Delta\overline{u}_x = \overline{u}_{xR} - \overline{u}_{xL}$. The difference across the cell for the grid velocity is $\Delta u_{xg} = u_{xgR} - u_{xgL}$. We approximate each streamline as having a constant velocity equal to its time average, interpolated linearly in space at the half-time level. Therefore, $\tilde{x}_H = (\tilde{x} + \tilde{x}_N)/2$. After some algebra, we write:

$$\tilde{x}_N = f_{\rm exp}\tilde{x} + \sigma_M , \tag{A1}$$

where

$$f_{\rm exp} = [1 + 0.5\ \Delta\overline{u}_x\ \Delta t/\Delta x]/[1 + (\Delta u_{xg} - \Delta\overline{u}_x/2)(\Delta t/\Delta x)]$$

and

$$\sigma_M = (\overline{u}_{xM} - u_{xgM})(\Delta t/\Delta x)/[1 + (\Delta u_{xg} - \Delta\overline{u}_x/2)(\Delta t/\Delta x)].$$

In the spirit of conservation laws, we generate fluxes at the interfaces of the grid cells. Extra terms arise because each cell has its own set of coordinates. We assume a uniform grid of cells with widths Δx. Thus for streamlines such as

those shown in the right panel of Figure 18 that cross grid cell interfaces, we must either add or subtract 1 from the expression for $\tilde{x}_N$ (Equation (A1)). As stated earlier, this extra complexity is necessary to permit us to use 32 bit arithmetic in the computation. We now focus our attention on the cell interface L, which is the left-hand interface of our cell of interest. The velocity $\overline{u}_x$ at the cell interface that is *upstream* from interface L is $\overline{u}_{xL\text{up}}$, and the sign of the velocity $\overline{u}_{xL}$ is s_L. Thus, $s_L = 1$ if $\overline{u}_{xL} \geqslant 0$, and $s_L = -1$ otherwise. In the cases $s_L = \pm 1$, the streamline that reaches the cell interface at the end of the time step begins at $\tilde{x}_{sL\text{up}} = (s_L/2 - \sigma_{ML\text{up}})/f_{\text{exp}L\text{up}}$. Here we have used the subscript Lup to denote more clearly that $\tilde{x}_{sL\text{up}}$ is evaluated using the scaled coordinate at the beginning of the time step in the grid cell upstream from interface L. To denote the cell downstream from interface L we will use the subscript Ldn. We run subscripts together in a Germanic style because they translate naturally into suffixes for variable names in a code. We will need to evaluate moment increments that result from advection of the segments indicated in Figure 18 into adjacent cells. We denote these moment increments by $d\mathcal{M}_{kL}$. These will be integrals over the indicated segments of the upstream cells at the beginning of the time step, but using weight factors corresponding to the new grid cell's scaled coordinates at the new time. Thus:

$$d\mathcal{M}_{kL\text{dn}} = s_L \int_{\tilde{x}_{sL\text{up}}}^{s_L/2} d\tilde{x}_{L\text{up}} \left(d\tilde{x}_{NL\text{dn}}/d\tilde{x}_{L\text{up}}\right) f_{L\text{up}}(\tilde{x}_{L\text{up}})\, \tilde{x}_{NL\text{dn}}^k.$$

We now recognize that $(d\tilde{x}_{NL\text{dn}}/d\tilde{x}_{L\text{up}})$ is just $f_{\text{exp}L\text{up}}$ and that

$$\tilde{x}_{NL\text{dn}} = \tilde{x}_{NL\text{up}} - s_L = f_{\text{exp}L\text{up}}\, \tilde{x}_{L\text{up}} + \sigma_{ML\text{up}} - s_L.$$

We define the Courant number for the cell interface L as $\sigma_L = 1/2 - s_L \tilde{x}_{sL\text{up}}$. We should be careful to advect this same fraction of the upstream cell across the interface when, in the PPM algorithm, computing the fluxes of conserved quantities. We now write for the moment contributions:

$$d\mathcal{M}_{kL\text{dn}} = s_L\, f_{\text{exp}L\text{up}} \int_{\tilde{x}_{sL\text{up}}}^{s_L/2} d\tilde{x}_{L\text{up}}\, f_{L\text{up}}(\tilde{x}_{L\text{up}}) \times (f_{\text{exp}L\text{up}}\, \tilde{x}_{L\text{up}} + \sigma_{ML\text{up}} - s_L)^k.$$

It is easy to express these integrals in terms of simpler ones that are more readily evaluated:

$$d\mathfrak{M}_{kL\text{up}} = s_L \int_{\tilde{x}_{sL\text{up}}}^{s_L/2} d\tilde{x}_{L\text{up}}\, f_{L\text{up}}(\tilde{x}_{L\text{up}})\, \tilde{x}_{L\text{up}}^k = \sigma_L\, (s_L^k f_{0L\text{up}} D_{Lk} + s_L^{k+1} f_{1L\text{up}} D_{Lk+1} + s_L^{k+2} f_{2L\text{up}} D_{Lk+2}).$$

The $f_{0L\text{up}}$, $f_{1L\text{up}}$, and $f_{2L\text{up}}$ are the coefficients of the parabola representing f as a function of $\tilde{x}_{L\text{up}}$, and the constants D_{Lk} can be evaluated recursively via $D_{L0} = 1$ and

$$D_{Lk} = [2^{-k} + s_L\, k\, \tilde{x}_{sL\text{up}}\, D_{Lk-1}]/[k+1].$$

We may now express the desired moment integrals as follows:

$$\begin{aligned} d\mathcal{M}_{0L\text{dn}} &= f_{\text{exp}L\text{up}} d\mathfrak{M}_{0L\text{up}}, \\ d\mathcal{M}_{1L\text{dn}} &= f_{\text{exp}L\text{up}}^2 d\mathfrak{M}_{1L\text{up}} + (\sigma_{ML\text{up}} - s_L)\, d\mathcal{M}_{0L\text{dn}} \\ d\mathcal{M}_{2L\text{dn}} &= f_{\text{exp}L\text{up}}^3 d\mathfrak{M}_{2L\text{up}} + (\sigma_{ML\text{up}} - s_L) \times (f_{\text{exp}L\text{up}}^2 d\mathfrak{M}_{1L\text{up}} + d\mathcal{M}_{0L\text{dn}}). \end{aligned}$$

In these formulae we take special care to expose a highly efficient strategy for evaluating these quantities in a computer code.

We are evaluating contributions to conservation laws, even though this is somewhat obscured by our use of different coordinates in different cells and at different times. Nevertheless, it is easiest to evaluate the moment integrals corresponding to the portions of the fluid that remain in their original cells by instead evaluating total moment integrals over these cells and then subtracting off the advected portions. We denote these total moment integrals by $d\mathcal{M}_{kT}$:

$$d\mathcal{M}_{kT} = f_{\text{exp}} \int_{-1/2}^{1/2} d\tilde{x}\; f(\tilde{x})(f_{\text{exp}}\tilde{x} + \sigma_M)^k.$$

These are quite easily evaluated. Expressing the results recursively, we find

$$\begin{aligned} d\mathcal{M}_{0T} &= f_{\text{exp}} \langle f \rangle, \quad d\mathcal{M}_{1T} = f_{\text{exp}}^2 \langle f\tilde{x} \rangle + \sigma_M\, d\mathcal{M}_{0T} \\ d\mathcal{M}_{2T} &= f_{\text{exp}}^3 \langle f\tilde{x}^2 \rangle + \sigma_M\, (f_{\text{exp}}^2 \langle f\tilde{x} \rangle + d\mathcal{M}_{1T}). \end{aligned}$$

To obtain the contributions to the new moments in this original cell from the fluid that remains there, we need to subtract from the above total moment integrals the portions corresponding to the fluid that is advected into neighboring cells. These integrals are very closely related to those, $d\mathcal{M}_{kL\text{dn}}$, evaluated above. They only lack the term $-s_L$ in the bracket in the integrand that is raised to the k^{th} power. Thus

$$d\mathcal{M}_{kL\text{up}} = s_L\, f_{\text{exp}L\text{up}} \int_{\tilde{x}_{sL\text{up}}}^{s_L/2} d\tilde{x}_{L\text{up}}\, f_{L\text{up}}(\tilde{x}_{L\text{up}}) \times (f_{\text{exp}L\text{up}}\, \tilde{x}_{L\text{up}} + \sigma_{ML\text{up}})^k.$$

Because we have done most of the work in evaluating these slightly different integrals already, we find

$$\begin{aligned} d\mathcal{M}_{0L\text{up}} &= d\mathcal{M}_{0L\text{dn}}, \quad d\mathcal{M}_{1L\text{up}} = d\mathcal{M}_{1L\text{dn}} + s_L\, d\mathcal{M}_{0L\text{up}} \\ d\mathcal{M}_{2L\text{up}} &= d\mathcal{M}_{2L\text{dn}} + s_L\, d\mathcal{M}_{1L\text{dn}} + s_L\, d\mathcal{M}_{1L\text{up}}. \end{aligned}$$

There are now four cases to consider in evaluating the moments in the cells at the new time level:

$$\langle f\tilde{x}^k \rangle_N = \begin{cases} d\mathcal{M}_{kT} - d\mathcal{M}_{kL\text{up}} + d\mathcal{M}_{kR\text{dn}} & \text{if } s_L = -1 \text{ and } s_R = -1 \\ d\mathcal{M}_{kT} - d\mathcal{M}_{kL\text{up}} - d\mathcal{M}_{kR\text{up}} & \text{if } s_L = -1 \text{ and } s_R = +1 \\ d\mathcal{M}_{kT} + d\mathcal{M}_{kL\text{dn}} + d\mathcal{M}_{kR\text{dn}} & \text{if } s_L = +1 \text{ and } s_R = -1 \\ d\mathcal{M}_{kT} + d\mathcal{M}_{kL\text{dn}} - d\mathcal{M}_{kR\text{up}} & \text{if } s_L = +1 \text{ and } s_R = +1. \end{cases}$$

All these computations are easily implemented in 100% vectorized loops making use of vector logic (the *cvmgm* or conditional move hardware instructions).

We have now described the most difficult of the six advection calculations. It is easy to see how the moments of the functions $f_y(\tilde{x})$ and $f_y(\tilde{x})$, defined earlier, are computed. We use the same method that was just described for the function $f_x(\tilde{x})$, but we simply drop the highest-order terms. Updating the moments of the functions $f_{yz}(\tilde{x})$, $f_{yy}(\tilde{x})$, and $f_{zz}(\tilde{x})$ is simpler still, but we must take care to handle the dependence of the last two of these functions on $\langle f \rangle$.

We now discuss the very important topic of how we constrain the initial functions that we advect as described above. Because

we will perform 1D passes in all three grid dimensions, there is no need to constrain any of the functions in the x-pass except for $f_x(\tilde{x})$. Since van Leer introduced the 1D unconstrained version of this advection scheme in the mid-1970s, we have learned a great deal about constraining interpolation parabolae. The constraints that were introduced for the PPB scheme (Woodward 1986) have been augmented and improved over many years of using PPB advection. The modifications have sought to eliminate the generation of tiny bits of a fluid that can otherwise precede or follow a moving region of a given fluid. The modifications also seek to keep a sharp multifluid interface sharp when it moves perpendicular to itself while letting the interface spread when that is appropriate. Behavior of the numerical representation of the interface between two fluids when that interface is physically unstable is particularly important. This is a delicate business, and that is why the algorithm described below is complex.

Constraints that we apply take two general forms. The first, and the simplest, is blending with $f_x(\tilde{x})$ a fraction of the constant function $\langle f \rangle$, so that an extremum of the composite function that was negative or exceeded unity before this blending operation just attains the value 0 or 1 after blending. We never blend in the constant function $\langle f \rangle$ for any other reason or in any other circumstances. This is very important. The PPB advection scheme is so accurate that any degradation arising from tampering with the functions it advects can be devastating, removing essentially all benefit of its elaborate advection computation. Therefore we only blend in constant functions in order to remove unphysical implied values inside a grid cell. Because we use PPB to advect multifluid concentrations, values outside the interval from 0 to 1 must be eliminated. However, blending in portions of the constant function is not the only way, and is often not the best way, to enforce the specialness of the values 0 and 1. Nevertheless, if $\langle f \rangle$ is negative or exceeds unity, we must set the function $f_x(\tilde{x})$ to a constant at either the value 0 or 1 as appropriate. This is the first constraint that we apply, and we can see no alternative to it.

Second, we reset edge values that lie outside the allowed range from 0 to 1 (0.000002 to 0.999998 when using 32 bit arithmetic), and we mark cells where we have done this. We also reset the edge values to 0 or 1 if in the cell right across this cell interface we have pure fluid, with $\langle f \rangle$ equal to 0 or 1, and we mark these cells along with the previous set. We need to keep track of these marked cells because in these cells and these cells only we will constrain the implied parabolae in the cells so that they do not assume minimum or maximum values within the cell. A difficulty in treating cells adjacent to pure fluid regions arises from the scheme's tendency to insert in such cells interpolation parabolae that have extrema in those cells. Even if we flatten these parabolae to the point that these extrema are constrained to lie in the allowed range of values, we will have directly adjacent to the pure fluid region a tiny pocket of ever so slightly mixed fluid. When this tiny pocket translates, the PPB scheme can begin to treat it with an interpolation parabola having a permitted extremum within the cell. In such an event, the little pocket of mixed fluid will be preserved, which is an undesired behavior. We therefore wish to detect such cells and force the scheme to use monotone interpolation parabolae in them for which, if possible, the edge value adjacent to the pure fluid region is unity or zero. Our detection must be very carefully done, or else we will mistakenly find such cells all over the grid, apply monotone interpolations in them, and destroy the resolving power of the PPB scheme. Ideally, we would like the scheme to treat physical diffusion of edges properly, while still keeping non-diffusing edges sharp. This implies that we would like the scheme to sense the difference between physical and numerical diffusion, allowing the former and disallowing the latter. This is a tall order. We will assume that edges, as opposed to truly smooth transitions, are sharper than the Gaussian curve in their approach to either special value 0 or 1. This is a property that we can successfully detect. We test this by demanding that our constraints do not introduce sharp jumps artificially into a Gaussian pulse of height unity and with full width at half maximum that is 25% of a periodic domain resolved by between 9 and 27 grid cells when this Gaussian is advected for 100 transits of this periodic domain at a Courant number of 0.15 or more. This is not an empty requirement. The PPM advection scheme, with its contact discontinuity and steepening feature enabled, will turn such a Gaussian pulse into a square wave under these circumstances. Higher grid resolution than this (six cells across the pulse at half maximum) is needed to distinguish the Gaussian from a square wave using PPM. Disabling the discontinuity detection and steepening in PPM leaves the Gaussian a Gaussian, but diminishes its amplitude substantially at low grid resolutions. PPB advection, with the constraints laid out below, easily distinguishes Gaussians from square waves, and vice versa, and performs the appropriate interpolations for each. PPB needs only as few as six cells across a square wave pulse to preserve its full amplitude while advecting it 400 times its width. Square waves with widths of only two or three cells are turned into Gaussian-type pulses by the PPB scheme. This seems a more benign behavior than performing the reverse metamorphosis, so long as the pulse amplitude is adequately maintained. For PPB this is the case, due to its subcell resolution provided through its moment data. In marginal cases, where grid resolution is insufficient, we are forced to make a choice between square and Gaussian-type pulse shapes. Our constraints will choose the square pulse shape in cases that cannot be clearly decided from the data provided, but the subcell information PPB maintains allows us to apply such arbitrary decisions exclusively on very highly under-resolved signals. Also, we arbitrarily favor sharp interpolated features only when we have transitions to one of the special function values of 0 or 1. The design of our constraints below therefore means that an interface that is originally sharp will remain so unless sufficiently strong physical diffusion is applied in a single time step. We assume that this is a desired behavior. Originally smooth multifluid interface transitions will become sharp only if they are stretched in the dimensions parallel to the transition surface, so that Liouville's theorem implies that the transition region must become thinner. This, we believe, is also a highly desirable behavior. We note that our experience with many multifluid applications of the PPB advection scheme indicates that very few numerical glitches arise from the scheme's switching between constrained and unconstrained behaviors (see Woodward et al. 2012, 2010a), since this occurs only when the multifluid interfaces involved have numerical representations of about three cells in thickness or less.

We will perform our testing and resetting of the left and right interface values, f_L and f_R, in a given cell according to the following logical tree. Note that because this resetting of both values causes a monotone interpolation parabola to result, we can avoid unnecessary later work in these cells by flagging them as already handled. Using subscripts ZL and ZR to denote the zone (or grid cell) to the left or right of the cell of interest, we

execute the following logical tree:

$$\begin{aligned}
&\text{if } \langle f\rangle_{ZL} < 0.001 \quad \text{and} \quad \langle f\rangle_{ZR} > 5\langle f\rangle \\
&\quad \text{then} \quad f_L = 0 \quad \text{and} \quad f_R = 3\langle f\rangle \\
&\text{if } \langle f\rangle_{ZR} < 0.001 \quad \text{and} \quad \langle f\rangle_{ZL} > 5\langle f\rangle \\
&\quad \text{then} \quad f_R = 0 \quad \text{and} \quad f_L = 3\langle f\rangle \\
&\text{if } \langle f\rangle_{ZL} > 0.999 \quad \text{and} \quad \langle f\rangle_{ZR} < 5\langle f\rangle - 4 \\
&\quad \text{then} \quad f_L = 1 \quad \text{and} \quad f_R = 3\langle f\rangle - 2 \\
&\text{if } \langle f\rangle_{ZR} > 0.999 \quad \text{and} \quad \langle f\rangle_{ZL} < 5\langle f\rangle - 4 \\
&\quad \text{then} \quad f_R = 1 \quad \text{and} \quad f_L = 3\langle f\rangle - 2.
\end{aligned}$$

We now locate cells containing extrema. At this point in our sequence of operations, we are working with the three quantities f_L, f_R, and $\langle f\rangle$, which together are sufficient to define our interpolation parabola. We now compute the coefficients, f_0, f_1 and f_2, of the interpolation parabola's terms in $\tilde{x}^0$, $\tilde{x}^1$, and $\tilde{x}^2$:

$$\begin{aligned}
&f_1 = f_R - f_L \ , \\
&f_2 = 3(f_L + f_R - 2\langle f\rangle) \quad \text{and} \quad f_0 = \langle f\rangle - \ f_2/12 \ .
\end{aligned}$$

It is easily shown that we must have a minimum inside the cell if $f_2 > |f_1|$, and that we must have a maximum inside the cell if $-f_2 > |f_1|$. Such an extremum must occur at $\tilde{x}_{\rm ext} = -f_1/(2f_2)$, where the parabola assumes the value

$$f_{\rm ext} = f_0 \ - f_1^2/(4f_2) = \langle f\rangle - \ (f_2/12)[1 + 3(f_1/f_2)^2].$$

The reduction factor that we apply to f_1 and f_2 in these cells in order to bring $f_{\rm ext}$ to the desired, allowed value $\langle f\rangle_{OK}$ is then

$$f_{\rm reduce} \ = 12\,(\langle f\rangle - \langle f\rangle_{OK})/\left(f_2\ \left[1 + 3(f_1/f_2)^2\right]\right).$$

We can make this process of testing computationally efficient, avoiding unnecessary divisions and evaluations by first computing only f_1 and f_2 from our values of f_L and f_R, and testing and resetting according to the following:

$$\begin{aligned}
&\text{if} \quad |f_2| > |f_1| \quad \text{and} \quad f_2^2 + 3f_1^2 > 12\ f_2\ (\langle f\rangle - \langle f\rangle_{OK}), \\
&\quad \text{then} \quad f_{\rm reduce} = \frac{12\ f_2\ (\langle f\rangle - \langle f\rangle_{OK})}{(f_2^2 + 3f_1^2)}.
\end{aligned}$$

The reduction factor $f_{\rm reduce}$ is applied to f_1 and f_2 only in these cells. We can see that we do not actually need to evaluate the coefficient f_0. We note that in the set of cells that we have already marked for the application later of constraints forcing the parabolae within these cells to be monotone, we must take care to assure that $f_{\rm reduce}$ is set to unity. For these cells, we will keep f inside the range from 0 to 1 in a different fashion.

The idea here is that we reduce the magnitudes of f_1 and f_2 together, flattening the interpolation parabola, in cases where our grid cell is likely to be located in between regions where f is either 0 or 1. Then we are likely to have in our cell a segment of a thin, unresolved strip in which the distribution f is either 1 or 0 (respectively). It is entirely appropriate to describe this unresolved strip segment using a parabola that has an extremum inside our grid cell. We must however be careful not to let the extreme value poke outside of our allowable range. If, however, we have reset either edge value to one of our limiting allowable values, that is, to either 0 or 1, then we expect that our cell is located next to a region in which the distribution f is either 0 or 1 (respectively). In this case, it would be inappropriate to simply flatten our interpolation parabola, which would cause the value at the cell edge to move away from the limiting value shared by the region adjacent to it. Instead, we would like to keep the edge value at this limit, if that is possible.

In order to prevent the generation of inappropriate subcell extrema in cells located next to regions of roughly constant f where this constant value is neither 0 nor 1, we will additionally demand that $f_{\rm reduce}$ be set to unity unless the cell average, $\langle f\rangle$, is an extremum relative to those on the left and right, $\langle f\rangle_{ZL}$ and $\langle f\rangle_{ZR}$, or the inferred extremum within the grid cell was within the allowable range. Here again the subscripts *ZL* and *ZR* signify "zone on the left" and "zone on the right." Thus we demand that

$$\begin{aligned}
&\text{if} \quad (\langle f\rangle - \langle f\rangle_{ZL})(\langle f\rangle_{ZR} - \langle f\rangle) \leqslant 0 \\
&\quad \text{and} \quad f_2 > |f_1| \quad \text{then} \quad f_{\rm reduce} = 1.
\end{aligned}$$

We add cells reset in this way and where $f_{\rm reduce}$ had previously been set to a value less than unity to the cells we have been marking as candidates for the application of PPM-style constraints that force the parabolae within these cells to be monotone. We may now safely apply the factor $f_{\rm reduce}$ to the present values of f_1 and f_2. We will not need to evaluate f_0.

Finally, we apply the constraint that the parabolae in cells that we have been marking must be monotone inside those cells. Since our constraint procedures are so complicated, it is worthwhile at this point to remember which cells are now marked. First, all cells for which an implied edge value, either f_L or f_R, was outside the allowable range and which, after these edge values were reset, had extrema inside the cell have been marked. Second, all cells with internal extrema outside the allowable range and for which the cell averages are not corresponding extrema relative to the cell averages in neighboring cells have been marked. Thus in all marked cells, the parabola describing the subcell structure is not monotone, but we will demand that it be monotone. In the marked cells, we therefore apply

$$\begin{aligned}
&\text{if} \ \ f_L = 0 \quad \text{or} \quad f_R = 1 \quad \text{and} \quad -f_2 > |f_1| \\
&\quad \text{then} \quad f_2 = -f_1 = 3\,(\langle f\rangle - f_R) \\
&\text{if} \ \ f_L = 1 \quad \text{or} \quad f_R = 0 \quad \text{and} \quad f_2 > |f_1| \\
&\quad \text{then} \quad f_2 = -f_1 = 3\,(\langle f\rangle - f_R) \\
&\text{if} \ \ f_L = 1 \quad \text{or} \quad f_R = 0 \quad \text{and} \quad -f_2 > |f_1| \\
&\quad \text{then} \quad f_2 = f_1 = 3\,(\langle f\rangle - f_L) \\
&\text{if} \ \ f_L = 0 \quad \text{or} \quad f_R = 1 \quad \text{and} \quad f_2 > |f_1| \\
&\quad \text{then} \quad f_2 = f_1 = 3\,(\langle f\rangle - f_L).
\end{aligned}$$

At the end of this resetting procedure, we recompute the moments of the distribution via

$$\langle f\tilde{x}\rangle = f_1/12 \quad \text{and} \quad \langle f\tilde{x}\tilde{x}\rangle \ = \ (\langle f\rangle + \ f_2/15)/12.$$

With this final operation, the application of our constraints is complete.

One may compare the above PPB advection scheme with multiple other variants of the scheme as well as with variations on PPM advection in both 1D and 2D by experimenting with the MS Windows applications provided on the LCSE Web site at http://www.lcse.umn.edu/two-stream-test, http://www.lcse.umn.edu/Gas1D, and http://www.lcse.umn.edu/WindTunnel. These Web URLs contain example applications, executable on any Windows PC, that were generated in support

of courses taught in the 1990s. The first solves a 1D gravitational two-stream instability problem in a 2D phase space. Multiple variations on PPB, several more accurate than the version we have described above and that we have used for the last several years, are enabled in this Windows program. The behavior of PPB can be compared with that of PPM, in multiple flavors, as well as a high-order Runge–Kutta advection scheme used in meteorological codes. The Gas1D and WindTunnel URLs host Windows applications that combine PPB multifluid fractional volume advection with PPM gas dynamics in much the way that they are combined in our present codes. The WindTunnel code performs a multifluid variation on the classic 2D wind tunnel test problem Woodward & Colella (1984). We do not guarantee that these downloadable applications will work on any PC, although they might, nor that we will maintain them forever, nor that we will answer inquiries concerning them. We also assert that one downloads these applications at his or her own risk, and they are provided "as is." Nevertheless, they do now work for us, and they are both instructive and fun to experiment with. Examples of the use of our PPB scheme for large-scale multifluid gas dynamics simulations with the PPM code can be found in Woodward et al. (2008a, 2008b, 2010a, 2012) and Herwig et al. (2011).

A.1.2. Combining PPB Advection with PPM Gas Dynamics

Once we have computed the time-and-space averaged velocities at the grid cell interfaces using our version of PPM, we perform the PPB advection computation to obtain the new values of the 10 moments of f without reference to the values of any other variables. This computation gives us advected volumes of the two fluids, but not advected masses. We obtain from PPB volumes of the two gases within the cells *at the beginning of the time step* that cross the cell interfaces into neighboring cells. To convert these advected volumes into advected masses, needed for strict mass conservation, we must introduce interpolations of the individual fluid densities as functions of cell volume coordinates at the beginning of the time step. We find these interpolation parabolae using our standard PPM procedure.

Here we invoke an important assumption. We assume strict pressure and temperature equilibrium inside each cell. This implies that the two fluids must have at each point in the cell a ratio of their densities that is given by that of their mean molecular weights—which we assume is constant over the entire duration of the problem. In our hydrogen ingestion flash simulations (Herwig et al. 2011), we take this ratio to be 2.26. The pressure and temperature equilibrium assumption permits us to derive the individual densities of the two fluids given only the averaged density of the mixture plus the mixing fraction, f. This is a huge simplification because we need not store densities and internal energies for each fluid. When we interpolate a parabola to represent the variation of the density of one of the fluids across a cell, this implies such a parabola for the other fluid. Together with our moments for the distribution of the mixing fraction, f, we can derive the implied distribution of density for the mixture. Because careful interpolation is very expensive, this represents a great saving in computational labor. There is still another advantage. Our equilibrium assumption means that even in a cell containing no gas of one type, a reasonable average density for that gas is implied. Consequently, we have no need to interpolate gas density across a discontinuity at a multifluid boundary: the density of each gas is well-behaved across such a boundary. It is instead the mixing fraction, f, that jumps suddenly across this boundary. However, we have for f our very much more accurate PPB description, with its 10 moments in each cell. This allows, as a practical matter, a smooth description of f across a multifluid boundary that is only about two grid cells thick. This representation is prevented from becoming too sharp, which would introduce numerical oscillations or glitches because it is forced to consist of a parabola extending all the way across each grid cell. Our interpolated distribution of f is therefore very sharp—only a couple of grid cells in thickness—and at the same time very smooth because it is defined by parabolae in these cells that are determined by subcell information that is operationally equivalent, as a rule of thumb, to a two- or three-fold grid refinement for PPM (see Woodward 2005) for just this single, all-important variable.

In slow-flow, Rayleigh–Taylor instability problems (see Woodward et al. 2010a, 2012), we find that the elaborate approach of the PPB scheme is sufficient to essentially eliminate the appearance of certain bad behaviors familiar to us from the PPM advection scheme when applied to multifluid problems (see for example Bassett & Woodward 1995a, 1995b). PPB advection, with its formal fifth-order accuracy, is capable of moving multifluid interfaces with very detailed structure great distances through the mesh with no noticeable diffusion. This behavior is possible because PPB consistently treats the internal structure of the multifluid interface transition from 0 to 1 in the mixing fraction, f, as a smooth transition. Unlike PPM advection, it does not switch between fundamentally different interpolation strategies dynamically. Switching strategies in this way can cause PPM to introduce small glitches, which can later become amplified by a physical instability to form large glitches. We have been using multifluid PPM+PPB for Rayleigh–Taylor problems in the weakly compressible regime since 2004, and find it very much superior to PPM alone for these problems, as the results reported in Woodward et al. (2010a, 2012) attest. Almgren et al. (2010) reported similar experience with high-order advection and Rayleigh–Taylor problems using their CASTRO code. This experience, we find, does not carry over to the much more violent inertial confinement fusion problems considered in Woodward et al. (2012). This can be seen in two dimensions for both our code and for CASTRO in code comparison work (Joggerst et al. 2014). We find that the interaction of strong shocks, as they are handled in PPM, with our very carefully treated multifluid interfaces produce familiar sorts of glitches, whose causes were explained decades ago (Woodward & Colella 1984). These issues are addressed and largely resolved by Woodward et al. (2012), which also describes our multifluid PPM code's highly scalable parallel implementation in some detail.

REFERENCES

Almgren, A. S., Beckner, V. E., Bell, J. B., et al. 2010, ApJ, 715, 1221
Asida, S. M., & Arnett, D. 2000, ApJ, 545, 435
Asplund, M., Lambert, D. L., Kipper, T., Pollacco, D., & Shetrone, M. D. 1999, A&A, 343, 507
Bassett, G. M., & Woodward, P. R. 1995a, JFM, 284, 323
Bassett, G. M., & Woodward, P. R. 1995b, ApJ, 441, 582
Brummell, N. H., Clune, T. L., & Toomre, J. 2002, ApJ, 570, 825
Campbell, S. W., & Lattanzio, J. C. 2008, A&A, 490, 769
Campbell, S. W., Lugaro, M., & Karakas, A. I. 2010, A&A, 522, L6
Casanova, J., José, J., García-Berro, E., Shore, S. N., & Calder, A. C. 2011, Natur, 478, 490
Chandrasekhar, S. 1961, Hydrodynamic and Hydromagnetic Stability (Oxford: Clarendon)
Christensen-Dalsgaard, J., Monteiro, M. J. P. F. G., Rempel, M., & Thompson, M. J. 2011, MNRAS, 414, 1158

Cockburn, B., Karniadakis, G. E., & Shu, C.-W. (ed.) 2000, Lecture Notes in Comput. Sci. Eng., Vol. 10, Discontinuous Galerkin Methods. Theory, Computation and Applications (New York: Springer)
Colella, P., & Woodward, P. R. 1984, JCoPh, 54, 174
Cristallo, S., Piersanti, L., Straniero, O., et al. 2009, PASA, 26, 139
Denissenkov, P. A., Herwig, F., Bildsten, L., & Paxton, B. 2012a, ApJ, 762, 8
Denissenkov, P. A., Herwig, F., Bildsten, L., & Paxton, B. 2012b, ApJ, 762, 8
Denissenkov, P. A., Herwig, F., Truran, J. W., & Paxton, B. 2013, ApJ, 772, 37
Deupree, R. G. 2000, ApJ, 543, 395
Duerbeck, H. W., & Benetti, S. 1996, ApJL, 468, L111
Freytag, B., & Höfner, S. 2008, A&A, 483, 571
Freytag, B., Ludwig, H. G., & Steffen, M. 1996, A&A, 313, 497
Fujimoto, M. Y., Ikeda, Y., Iben, I., & J. 2000, ApJL, 529, L25
Hajduk, M., Zijlstra, A. A., Herwig, F., et al. 2005, Sci, 308, 231
Herwig, F. 2000, A&A, 360, 952
Herwig, F. 2001, ApJL, 554, L71
Herwig, F. 2003, CNO in the Universe, 304, 318
Herwig, F. 2005, ARA&A, 43, 435
Herwig, F., Blocker, T., Langer, N., & Driebe, T. 1999, A&A, 349, L5
Herwig, F., Blocker, T. B., Schonberner, D. S., & El Eid, M. F. 1997, A&A, 324, L81
Herwig, F., Freytag, B., Fuchs, T., et al. 2007, in ASP Conf. Ser. 378, Why Galaxies Care About AGB Stars: Their Importance as Actors and Probes, ed. F. Kerschbaum, C. Charbonnel, & R. F. Wing (San Francisco, CA: ASP), 43
Herwig, F., Freytag, B., Hueckstaedt, R. M., & Timmes, F. X. 2006, ApJ, 642, 1057
Herwig, F., Pignatari, M., Woodward, P. R., et al. 2011, ApJ, 727, 89
Herwig, F., Woodward, P. R., Lin, P.-H., Knox, M., & Fryer, C. 2014, ApJL, 792, L3
Hesthaven, J. S., & Warburton, T. (ed.) 2000, Texts in Applied Mathematics, Vol. 54, Nodal Discontinuous Galerkin Methods, Algorithms, Analysis, and Applications (New York: Springer)
Hurlburt, N. E., Toomre, J., Massaguer, J. M., & Zahn, J.-P. 1994, ApJ, 421, 245
Iben, I., Jr., & MacDonald, J. 1995, in White Dwarfs, ed. D. Koester & K. Werner (Heidelberg: Springer), 48
Iwamoto, N., Kajino, T., Mathews, G. J., Fujimoto, M. Y., & Aoki, W. 2004, ApJ, 602, 378
Joggerst, C. C., Nelson, A., Woodward, P. R., et al. 2014, JCoPh, 275, 154
Jones, S., Hirschi, R., Nomoto, K., et al. 2013, ApJ, 772, 150
Karakas, A. I., Campbell, S. W., & Stancliffe, R. J. 2010, ApJ, 713, 374
Kippenhahn, R., & Weigert, A. 1990, Stellar Structure and Evolution (Berlin: Springer)
Kuhlen, M., Woosley, S. E., & Glatzmaier, G. A. 2006, ApJ, 640, 407
Lau, H. H. B., Stancliffe, R. J., & Tout, C. A. 2009, MNRAS, 396, 1046
Lawlor, T. M., & MacDonald, J. 2003, ApJ, 583, 913
Maeder, A. 1976, A&A, 47, 389
Meakin, C. A., & Arnett, D. 2007, ApJ, 667, 448
Meakin, C. A., & Arnett, D. 2007, ApJ, 667, 448
Miesch, M. S., Brun, A. S., De Rosa, M. L., & Toomre, J. 2008, ApJ, 673, 557
Miesch, M. S., Elliott, J. R., Toomre, J., et al. 2000, ApJ, 532, 593
Miller Bertolami, M. M., Althaus, L. G., Serenelli, A. M., & Panei, J. A. 2006, A&A, 449, 313
Mocák, M., Campbell, S. W., Müller, E., & Kifonidis, K. 2010, A&A, 520, A114
Mocák, M., Müller, E., Weiss, A., & Kifonidis, K. 2008, A&A, 490, 265
Mocák, M., Müller, E., Weiss, A., & Kifonidis, K. 2009, A&A, 501, 659
Mocák, M., Siess, L., & Müller, E. 2011, A&A, 533, A53
Mowlavi, N. 1999, A&A, 344, 617
Paxton, B., Bildsten, L., Dotter, A., et al. 2011, ApJS, 192, 3
Porter, D. H., & Woodward, P. R. 1994, ApJS, 93, 309
Porter, D. H., & Woodward, P. R. 2000, ApJS, 127, 159
Porter, D. H., Woodward, P. R., & Jacobs, M. L. 2000, in New York Academy Sciences Annals, Vol. 898, Astrophysical Turbulence and Convection, ed. J. R. Buchler & H. Kandrup, 1, online at http://www.lcse.umn.edu/convsph
Robinson, F. J., Demarque, P., Li, L. H., et al. 2003, MNRAS, 340, 923
Rogers, T. M., & Glatzmaier, G. A. 2005, ApJ, 620, 432
Rogers, T. M., Glatzmaier, G. A., & Jones, C. A. 2006, ApJ, 653, 765
Schaller, G., Schaerer, D., Meynet, G., & Maeder, A. 1992, A&AS, 96, 269
Simpson, R. L. 1989, AnRFM, 21, 205
Stancliffe, R. J., Dearborn, D. S. P., Lattanzio, J. C., Heap, S. A., & Campbell, S. W. 2011, ApJ, 742, 121
Stein, R. F., & Nordlund, A. 1998, ApJ, 499, 914
Suda, T., & Fujimoto, M. Y. 2010, MNRAS, 405, 177
Sytine, I. V., Porter, D. H., Woodward, P. R., Hodson, S. W., & Winkler, K.-H. 2000, JCoPh, 158, 225
van Leer, B. 1977, JCoPh, 23, 276
VandenBerg, D. A., Bergbusch, P. A., & Dowler, P. D. 2006, ApJS, 162, 375
Ventura, P., Zeppieri, A., Mazzitelli, I., & D'Antona, F. 1998, A&A, 334, 953
Viallet, M., Meakin, C., Arnett, D., & Mocák, M. 2013, ApJ, 769, 1
Weiss, A., & Ferguson, J. W. 2009, A&A, 508, 1343
Werner, K., & Herwig, F. 2006, PASP, 118, 183
Woodward, P., & Colella, P. 1981, in Lecture Notes in Physics, Vol. 141, Numerical Methods in Fluid Dynamics, ed. W. C. Reynolds & R. W. MacCormack (Berlin: Springer), 434
Woodward, P., & Colella, P. 1984, JCoPh, 54, 115
Woodward, P., Herwig, F., Porter, D., et al. 2008a, AIP Conf. Ser. 990, First Stars III, ed. B. W. O'Shea, A. Heger, & T. Abel (Melville, NY: AIP), 300
Woodward, P. R. 1986, in Astrophysical Radiation Hydrodynamics, ed. K.-H. Winkler & M. L. Norman (Dordrecht: Reidel), 245
Woodward, P. R. 2005, PPB: the Piecewise-Parabolic Boltzmann Scheme for Moment-Conserving Advection in 2 and 3 Dimensions, LCSE internal report, University of Minnesota, available at http://www.lcse.umn.edu/PPBdocs
Woodward, P. R. 2006, in Implicit Large Eddy Simulation, Computing Turbulent Fluid Dynamics, ed. L. M. F. Grinstein & W. Rider (Cambridge: Cambridge Univ. Press), 130, available at http://www.lcse.umn.edu/ILES/PPM-for-ILES-turb-appls-2-26-05-cup6a-10.pdf
Woodward, P. R., Jayaraj, J., Lin, P.-H., et al. 2010a, in Proc. NECDC 2010 Conf., Los Alamos, 2010 Oct., Los Alamos Natl. Lab. report LA-UR 11-00061, available at www.lcse.umn.edu/NECDC2010
Woodward, P. R., Jayaraj, J., Lin, P.-H., et al. 2012, in Proc. NECDC 2012 Conf., Livermore, CA, Los Alamos Natl. Lab. report LA-UR-13-20949, available at http://www.lcse.umn.edu/NECDC2012
Woodward, P. R., Porter, D., Dai, W., et al. 2010b, in The Piecewise Parabolic Boltzmann Advection Scheme (PPB) Applied to Multifluid Hydrodynamics, Los Alamos Natl. Lab. Report LA-UR 10-01823, available at http://www.lcse.umn.edu/PPMplusPPB
Woodward, P. R., Porter, D. H., Herwig, F., et al. 2008b, in Nuclei in the Cosmos (NIC X), (arXiv:0901.1414)
Woosley, S. E., Heger, A., & Weaver, T. A. 2002, RvMP, 74, 1015