Stony Brook University

Multidimensional Simulations of Convection Preceding a Type I X-ray Burst

A Dissertation Presented

by

Christopher Michael Malone

to

The Graduate School

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

in

Physics

Stony Brook University

August 2011

Stony Brook University

The Graduate School

Christopher Michael Malone

We, the dissertation committee for the above candidate for the Doctor of Philosophy degree, hereby recommend acceptance of this dissertation.

Dr. Michael Zingale – Dissertation Advisor
Assistant Professor, Department of Physics and Astronomy

Dr. Alan Calder – Chairperson of Defense
Assistant Professor, Department of Physics and Astronomy

Dr. James Lattimer
Professor, Department of Physics and Astronomy

Dr. Thomas Hemmick
Distinguished Teaching Professor, Department of Physics and Astronomy

Dr. Marat Khairoutdinov
Associate Professor, School of Marine and Atmospheric Sciences, Stony Brook University

This dissertation is accepted by the Graduate School.

Lawrence Martin
Dean of the Graduate School

Abstract of the Dissertation

Multidimensional Simulations of Convection Preceding a Type I X-ray Burst

by

Christopher Michael Malone

Doctor of Philosophy

in

Physics

Stony Brook University

2011

Type I X-ray bursts are thermonuclear explosions on the surfaces of neutron stars that can be used to determine the mass and radius of the underlying neutron star and hence help constrain the equation of state for dense matter. Particularly important is our physical understanding of how a localized, subsonic burning front ignites and spreads, the state of the material in which the burning front propagates, and the extent to which heat released from reactions expands the photosphere of the neutron star. Multidimensional simulation of low Mach number astrophysical flows, such as the propagation of a flame or the slow convective turnover, in such systems have been rather restricted in the past; fully compressible hydrodynamics algorithms have a time step size that is constrained by the propagation of acoustic waves, which can be neglected in low

Mach number flows of this type.

This thesis presents results of multidimensional, plane-parallel simulations of the convection preceding ignition in a Type I X-ray burst. I use a low Mach number hydrodynamics code, MAESTRO, based on a low Mach number approximation that filters acoustic waves from the system allowing for a larger time step size while retaining the important compressible features, such as expansion from local heating and composition change. This allows for performing long-term evolution of the system and characterizing the effects of convection on the atmosphere. In particular, the simulations presented here suggest that the convection dredges up some of the underlying ^{56}Fe neutron star material into the atmosphere, which may affect any subsequent subsonic burning front as well as the color correction factor used to infer the underlying neutron star's mass and radius.

*To my wife, for always having faith in me and providing *support*.*

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

The main body of work during this dissertation revolved around comfort — getting comfortable with running MAESTRO, getting MAESTRO comfortable with running XRB problems, being comfortable with "breaking" the code (and subsequently fixing it), discovering comfort in the open-endedness of researching where no answers are known *a priori*, and taking comfort in knowing that the MAESTRO algorithm is complicated enough that I'm not the only one amongst my collaborators who does not understand its nuances and limitations.

The process of going from uncomfortable to comfortable was facilitated with the help of many people. I want to thank my collaborators at LBL, John Bell, Ann Almgren, and especially Andy Nonaka for taking the time and having the patience to explain math-*ey* things to this sloppy physicist — this dissertation would not have been possible without your interest in astrophysics. I would also like to thank Frank Timmes not only for making publicly available his microphysics routines, but for also excellent discussions on thermonuclear reaction networks and thermodynamics in general — thank you for taking the time to answer my plethora of questions I sent via email. Additionally, I would like to thank some of my officemates (you know who you are :-) for allowing me to bounce ideas off of them, and for allowing me to be their wall to bounce their own ideas off of — I have learned a lot from these interactions.

I would like to acknowledge Andrew Cumming, Stan Woosley, and Alex Heger for providing initial models for the calculations presented in this dissertation. I also would like to thank my thesis committee members Alan Calder, Jim Lattimer, Tom Hemmick, and Marat Khairoutdinov

for taking the time to read my thesis and provide suggestions/corrections. I especially want to thank Alan Calder for his acting as a second mentor to me, both in academic and general life-related areas.

I should also acknowledge computing support from the following entities that allowed me to run simulations on their machines: nyblue at BNL, NERSC, ANL, and TeraGrid. Furthermore, some text and figures are reproduced with permission of the AAS from the paper "Multidimensional Modeling of Type I X-ray Bursts. I. Two-dimensional Convection Prior to the Outburst of a Pure ^{4}He Accretor" (Malone et al., 2011).

But above all I would like to thank my dissertation advisor, Mike Zingale, who took me from not being familiar with a Unix terminal to running massively parallel computer simulations on some of the fastest computers in the world. Mike has played a key role in my development as a scientist, both educating me in the ways of computational astrophysics and advising me in cromulent career decisions regarding conferences, publishing, postdoc positions, etc. I hope to someday have his Jedi-like ability to sense the underlying phyiscal/computational issues of a problem, and I thank him for taking me as an apprentice.

CHAPTER 1

X-RAY BURSTS

For excellent overviews and a complete historical background, please see the following x-ray burst review articles: Lewin & Joss (1981); Lewin et al. (1993); Bildsten (2000); Strohmayer & Bildsten (2006). Here, I introduce some of the more salient features pertaining to this dissertation. I also discuss some of the reasons why one should care about x-ray bursts in the first place. At the end of this chapter, I outline the remainder of this dissertation.

1.1 A Brief Historical Background

Cosmic x-ray astronomy began with the launch of the *Uhuru* satellite in 1970. There had previously been balloon-borne instruments or detectors attached to sounding rockets, but *Uhuru* was the first satellite dedicated to x-ray astronomy. *Uhuru* surveyed the entire sky and catalogued over 300 x-ray sources, with its two proportional counters sensitive to photons in the 2-20 keV energy range (Forman et al., 1978). One such source, `4U 1820-30` near the center of globular cluster NGC 6624, was later observed by Grindlay and Heise (1975) to exhibit "two intense x-ray bursts" with sharp ($\lesssim 1$ s) rise times, peak luminosities $\gtrsim 20$ times the quiescent luminosity, and exponential-like decays over ~ 10 s. At about the same time, Belian, Conner and Evans (1976) reported on observations of 10 very brief x-ray "count-rate enhancements" with luminosities $\gtrsim$

15σ above the background signal, originating from somewhere in the constellation Norma.

There were several ideas regarding what could be causing these short bursts of x-rays. For 4U 1820−30, Grindlay and Gursky (1976) attempted to explain the decay time of the outbursts via Compton scattering of the primary x-rays within a hot gas in equilibrium surrounding the x-ray source. The parameters derived for the hot gas implied a compact object with a mass $M \sim 10^3 M_\odot$ — a black hole. Canizares (1976) showed that the surrounding gas need not be in equilibrium, and that relaxing this assumption does not necessitate the need for a supermassive compact x-ray source. The following year, several bursts were observed that were well-fit by a cooling blackbody source of size comparable to a neutron star (e.g., Swank et al., 1977; Hoffman et al., 1977).

At about the same time, there was a proposal that favored a neutron star environment for an x-ray burst of thermonuclear origin. Building upon the work of Hansen and van Horn (1975), Maraschi and Cavaliere (1977) — and independently Woosley and Taam (1976) — started to relate the observed x-ray burst properties to the results of simple calculations of thermonuclear burning in an accreted atmosphere on the surface of a neutron star. The ratio of energy released from accretion to rest mass energy is $\sim GM/Rc^2 \approx 20\%$ for a neutron star of mass $M = 1.4 M_\odot$ and radius $R = 10$ km; this same ratio for nuclear burning is $\lesssim 0.7\%$. The disparity in efficiency of energy release implies that the neutron star must act as a "storage battery" for enough accreted nuclear fuel ($\sim 10^{22}$ g) until conditions are met to explosively burn the material and produce an observable luminosity above that of the accretion luminosity (Woosley & Taam, 1976).

1.1.1 X-ray Burst Taxonomy

The nuclear origin concept for x-ray bursts worked well except for one source, MXB1730−335, which was discovered (Lewin et al., 1976) just prior to the Woosley and Taam paper of 1976. MXB1730−335, later dubbed "The Rapid Burster," exhibited thousands of bursts per day. For

such a high cadence of outbursts under the nuclear origin idea, The Rapid Burster would need a very high accretion rate and, hence, a very large persistent accretion luminosity, which was not observed.

In late 1977, The Rapid Burster was observed to emit x-ray outbursts similar to all other observed x-ray bursts in addition to its eponymous outbursts. Two types of x-ray bursts were then proposed to exist: Type I bursts are those associated with thermonuclear flashes, while the rapid Type II bursts are thought to be due to instabilities in the accretion disk, which dump material onto the surface of the neutron star on short timescales (Hoffman et al., 1978). For the remainder of this dissertation — when I use the terms x-ray burst (hereafter XRB), burst, outburst, etc. — I will only be concerned with thermonuclear, Type I x-ray bursts.

1.2 A More Modern Look

The basic XRB paradigm takes place in a mass-transferring, low-mass X-ray binary (LMXB) system in which the neutron star's companion has filled its Roche lobe and is dumping H- and/or He-rich material onto the surface of the neutron star. Depending on the accretion rate and composition, there are several burning regimes that will trigger an XRB (see Bildsten (2000) for an overview). The general idea is that a column of accreted material — or heavier-element ash from prior stable burning of accreted material — builds up until the temperature sensitivity of the energy generation rate at the base of the layer exceeds that of the local cooling rate and a thin-shell thermal instability forms. The instability eventually causes a runaway of unstable burning, resulting in an outburst. The energy released during the outburst is not enough to unbind the star, and it is unlikely — given the high surface gravity — that a significant fraction of the accreted material is ejected. While returning back to its quiescent state, the neutron star may resume accreting material from its companion, and the process can begin anew.

One-dimensional hydrodynamic studies reproduce many of the observable features of XRBs such as burst energies ($\sim 10^{39}$ erg), rise times (seconds), durations (10's – 100's of seconds), and recurrence times (hours to days) (Woosley & Weaver, 1984; Taam et al., 1993; Heger et al., 2007a). By construction, however, one-dimensional models assume that the fuel is burned uniformly over the surface of the star, which is highly unlikely given the large disparity between the thermalization and burning timescales of the accreted material (Shara, 1982). Furthermore, the *Rossi X-ray Timing Explorer* satellite observed coherent oscillations in the lightcurves of ~ 20 outbursts from LMXB systems (first by Strohmayer et al., 1996; more recently by Altamirano et al., 2010 and references therein). The asymptotic evolution of the frequency of such oscillations suggests they are modulated by the neutron star spin frequency (Muno et al., 2002). Oscillations observed during the rising portion of an outburst lightcurve are therefore indicative of a spreading burning front being brought in and out of view by stellar rotation. Additionally, oscillations observed during the decay phase of the burst are thought to be caused by unstable surface modes that may depend critically on the local heating and cooling rates during the burst (Narayan & Cooper, 2007, and references therein). The manner in which the burning front spreads and propagates throughout the accreted atmosphere is not well known, and a proper multidimensional modeling of the conditions in the atmosphere prior to outburst is needed (e.g. Fryxell & Woosley, 1982b).

Prior to the actual outburst, the burning at the base of the ignition column drives convection throughout the overlying layers and determines the state of the material in which the burning front will propagate. One-dimensional simulations of XRBs usually attempt to parameterize the convective overturn and mixing using astrophysical mixing-length theory (Böhm-Vitense, 1958) or through various diffusive processes (see Heger et al. 2000 for a thorough discussion). Recent multidimensional simulations of stellar convection (see Arnett et al., 2009, and references therein), however, show a large discrepancy in, for example, the velocity of a typical convective eddy when compared to one-dimensional models in the case of stellar evolution codes that use mixing-length

theory. Indeed, there has recently been an effort put forth in the astrophysical community, the so-called Convection Algorithms Based on Simulations, or CABS, to derive from multidimensional simulations a more physically motivated prescription for handling convection in one dimension (Arnett et al., 2008). To date, such methods have not propagated into the XRB-simulation community and a proper treatment of convection, without assumptions, requires simulation in multiple dimensions.

Multidimensional simulations of any aspect of XRBs, however, have hitherto been rather restrictive. A burning front can propagate either supersonically as a detonation or subsonically as a deflagration. Full hydrodynamic XRB detonation models in the spirit of Fryxell & Woosley (1982a) or Zingale et al. (2001) require a thick ($\sim$ 100 m) accreted helium layer. Such deep layers are only produced by very low accretion rates, which are inconsistent with the majority of rates inferred from observations of XRBs, and therefore the burning front in most XRBs likely propagates as a deflagration. Deflagration models are difficult to compute with standard compressible hydrodynamics codes due to the long integration times required. One possible solution is to eliminate the effect of acoustic waves in the system, allowing the time step to be controlled by the fluid velocity, rather than the sound speed. Such a method can be derived using low Mach number asymptotics; classic examples of low Mach number approaches include the incompressible, anelastic (Ogura & Phillips, 1962) and Boussinesq (Boussinesq, 1903) approximations. To this end, Spitkovsky et al. (2002) used a simple, shallow-water, two-layer, incompressible fluid to model the vertical structure of a deflagration front and showed how rotation coupled with convection may play an important role in regulating the spread of the front over the surface of the neutron star.

1.3 Why Care About X-ray Bursts?

X-ray bursts provide an excellent laboratory for studying combustion and flame propagation under extreme, semi-degenerate conditions. From a nuclear physics standpoint, they are one of the few possible locations for the rapid proton, or rp-process (Wallace & Woosley, 1981), burning — a series of proton captures and β^+ decays along the proton-rich side of nuclear stability — that can produce heavy elements up to ^{110}Te (e.g. Schatz, 2011, and references therein). Furthermore, type I X-ray bursts are possibly the most frequent thermonuclear explosions in the universe and provide a large amount of observational data that can be used to determine the properties of matter near the surface of a neutron star. To make meaningful inferences about these properties from observational data, however, we must have a proper theoretical understanding of the bursting phenomena (Bhattacharyya, 2010).

During the outburst of some XRBs, the luminosity can reach (and sometimes exceed) a critical limit, called the Eddington luminosity, where the radiation pressure exerts a force on the atmosphere greater than the gravitational pull. The atmosphere lifts off the surface of the neutron star; such bursts are called photospheric-radius expansion (PRE) bursts. Observationally, this expansion is inferred from the spectral softening — evolution towards lower energies — of the lightcurve during the burst. As the photosphere expands, the blackbody temperature as observed at the Earth, decreases, reaches a minimum, and then increases as the photosphere contracts. At "touch down," the photosphere is thought to have returned to the stellar radius and the blackbody temperature reaches a maximum before decreasing slowly due to cooling. At touch down, the flux is thought to be close to the Eddington flux. Furthermore, because the system is no longer expanding or contracting, the normalized emitting area $A = F_{\mathrm{bb},\infty}/\sigma_{\mathrm{SB}} T_{\mathrm{bb},\infty}^4$ — where $F_{\mathrm{bb},\infty}$ and $T_{\mathrm{bb},\infty}$ are the blackbody flux and temperature as observed at the Earth, and σ_{SB} is the Stefan-Boltzmann constant — after touch down remains relatively constant. By comparing the observed flux *at* touch

down with the normalized area *after* touch down, along with some estimate of the distance to and atmospheric composition of the source, one can obtain independent measures for both the mass and radius of the underlying neutron star (Steiner et al., 2010, e.g.). This information of mass and radius can then be used along with the Tolman-Oppenheimer-Volkov equations, which determine relativistic stellar structure, to rule out some of the highly uncertain, myriad equations of state that exist for cold dense matter (see Lattimer & Prakash, 2004; Lattimer, 2007, for example).

1.4 Roadmap of this Dissertation

As this dissertation is concerned with numerical hydrodynamics simulations of XRBs, in Chapter 2 I give a brief introduction to simple hydrodynamics, finite volume methods, and low Mach number approximations. In Chapter 3, I describe the low Mach number approximation algorithm, `MAESTRO`, used for this work. XRBs would not exist without thermonuclear burning, so I give a very brief introduction to reactions and reaction networks in Chapter 4. In Chapters 5 and 6, I present the difficulties and results of performing two-dimensional simulations of various x-ray burst models. Finally, in Chapter 7 I summarize what has been accomplished during this dissertation as well as provide direction for future research in the field.

CHAPTER 2

HYDRODYNAMICS

In this Chapter, I will give a brief flavor of some hydrodynamics to help motivate the need for the low Mach number approximation method described in Chapter 3. Portions of this discussion have been adopted from LeVeque's excellent book *Finite Volume Methods for Hyperbolic Problems* (LeVeque, 2002), as well as from parts of Doug Swesty's superb graduate course on numerical radiation hydrodynamics taught in the Spring of 2011 at Stony Brook University.

2.1 Some Basic Equations

The equations of hydrodynamics are essentially conservation laws for mass, momentum, and energy. Mathematically, (the differential form of) conservation laws can take on the form of constant coefficient, homogeneous hyperbolic systems of equations

$$\frac{\partial \mathbf{q}}{\partial t} + \mathbf{A}\nabla \cdot \mathbf{q} = 0 \tag{2.1}$$

where $\mathbf{q}$ is an n-vector and $\mathbf{A}$ is an $n \times n$ coefficient matrix. For this system to be considered hyperbolic, the eigenvalues of $\mathbf{A}$ must be real and the eigenvectors of $\mathbf{A}$ must be non-trivial and linearly independent. A consequence of this is that $\mathbf{q}$ can be decomposed into n independently

travelling waves propagating with speeds equal to the eigenvalues.

2.1.1 The Advection Equation

As a simple example, we will consider the one-dimensional wave, or advection equation with constant speed:

$$\frac{\partial q}{\partial t} + c\frac{\partial q}{\partial x} = 0, \tag{2.2}$$

where c is the propagation speed. On an infinite spatial domain, this equation is an initial value problem with an analytic solution $q(x,t) = q_0(x-ct)$ for given initial condition $q(x,0) = q_0(x)$. The solution simply *advects* to the right (for $c > 0$) at speed c. Along the lines $x - ct = const.$, called the *characteristic* curves, the solution remains constant; stated differently, information about the system propagates along these lines at speed c.

For a finite spatial domain, $a \leq x \leq b$, this problem becomes an initial boundary value problem, and values for q at one of the boundaries, $q(a,t) = q_a(t)$ or $q(b,t) = q_b(t)$, needs to be specified along with the initial condition $q_0(x)$. If $c > 0$, then the characteristic curves are incoming at a and outgoing at b. From a physical standpoint, this means that q within the domain can only be affected by the q flowing into the domain at point a; mathematically, this means that for a stable numerical method of solution — *roughly*, a method whose global error does not grow with time — one must only use the incoming boundary condition (Gustafsson et al., 1972).

2.1.2 The Euler Equations for a Polytropic Gas

More generally, conservation laws are written in integral form relating the time rate of change of the total q in some volume V to the net flux through the surface, S, out of the volume:

$$\frac{\partial}{\partial t}\int_V q(\mathbf{x},t)\, dV = -\int_S \mathbf{F}(\mathbf{x},q,t)\cdot \mathbf{dS}, \tag{2.3}$$

where $\mathbf{F}$ is the flux vector and $\mathbf{dS} = \mathbf{n}dS$, with $\mathbf{n}$ being the outward normal vector to surface S. Applying the divergence theorem to 2.3, which assumes $\mathbf{F}$ is smooth, results in

$$\int_V \left[\frac{\partial q}{\partial t} + \nabla \cdot \mathbf{F} \right] \, dV = 0. \tag{2.4}$$

The choice of volume element over which to evaluate the integral is arbitrary, and therefore this equation must hold for all V and we have a differential form of the conservation law reminiscent of 2.1,

$$\frac{\partial q}{\partial t} + \nabla \cdot \mathbf{F} = 0. \tag{2.5}$$

Note that the one-dimensional advection equation, 2.2, is a special case where $\mathbf{F}(\mathbf{x}, q, t) = cq(x,t)\mathbf{e}_x$ where $\mathbf{e}_x$ is the unit vector in the x-direction.

For a gas, the conservation law for mass, the *continuity equation*, can be written as

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{U}), \tag{2.6}$$

where ρ is the mass density and $\mathbf{U}$ is the fluid velocity. Neglecting viscous and body forces, the conservation laws for linear momentum, the *momentum equations*, are

$$\frac{\partial (\rho \mathbf{U})}{\partial t} + \nabla \cdot [(\rho \mathbf{U}) \mathbf{U} + p] = 0, \tag{2.7}$$

where p is the pressure of the gas. If we assume, for simplicity, that the gas behaves according to a polytropic equation of state, where $p = p(\rho) = K\rho^{\gamma}$ with K and γ constants, then the nonlinear

system of equations — the Euler equations for a polytropic gas — is closed

$$\frac{\partial}{\partial t}\begin{bmatrix} \rho \\ \rho\mathbf{U} \end{bmatrix} + \nabla\cdot\begin{bmatrix} \rho\mathbf{U} \\ \rho\mathbf{U}\mathbf{U} + p(\rho) \end{bmatrix} = 0. \tag{2.8}$$

This system is said to be in *conservative form*, and looks similar to 2.5. It is instructive to take this equation out of conservative form to relate to the hyperbolicity discussion in the previous sections.

One way to put 2.8 into a constant coefficient form similar to 2.1 is to use the primitive variables: ρ, $\mathbf{U}$, p. These variables are more physically intuitive, and putting the system in this form will allow us to easily learn something about the characteristics. To do this, we will expand the derivatives such as

$$\nabla\cdot(\rho\mathbf{U}) = \rho\nabla\cdot\mathbf{U} + \mathbf{U}\cdot\nabla\rho, \tag{2.9}$$

and form a pressure equation by using the continuity equation, 2.6, and the equation of state relationship $p' = K\gamma\rho^{\gamma-1}\rho'$ where the prime notation means a derivative (see Section 14.7 of LeVeque, 2002, for example). The resulting system of equations in one dimension is

$$\frac{\partial}{\partial t}\begin{bmatrix} \rho \\ U \\ p \end{bmatrix} + \begin{bmatrix} U & \rho & 0 \\ 0 & U & 1/\rho \\ 0 & \gamma p & U \end{bmatrix}\cdot\frac{\partial}{\partial x}\begin{bmatrix} \rho \\ U \\ p \end{bmatrix} = 0. \tag{2.10}$$

The adiabatic speed of sound in a gas is defined by $c_s = (\partial p/\partial\rho)^{1/2}$ at constant entropy; for the polytropic equation of state, this gives $c_s = (p\gamma/\rho)^{1/2}$. The eigenvalues of the coefficient matrix in 2.10 are then: U, $U \pm c_s$ — information can advect with the fluid or propagate as acoustic waves relative to the fluid motion. Fluids that permit finite-speed acoustic waves are called compressible fluids. Note, again, these equations are not in conservative form — velocity, U, and pressure, p, are not conserved quantities.

2.2 Short Description of Numerical Solution

In order to solve the coupled, nonlinear equations of hydrodynamics numerically, we must discretize the problem, which involves breaking the domain of the problem into a grid of zones or cells, and then solving the conservation laws for each zone. Furthermore, we must discretize the problem in time — we break the evolution of the system into a series of small time steps. In the remainder of this dissertation, unless otherwise stated, we will use the notation q_i^n to associate some value of q in the ith grid zone at time $t = t_n$.

We could use *finite difference* methods, where q_i^n denotes the *value* of q at point x_i: $q_i^n \approx q(x_i, t^n)$. Such methods are fairly straight-forward to generate, using Taylor series expansions to approximate the derivatives in 2.5. For example, using Taylor expansions for q about x_i,

$$q(x,t) = q_i + \left.\frac{\partial q}{\partial x}\right|_{x=x_i} (x - x_i) + \frac{1}{2} \left.\frac{\partial^2 q}{\partial x^2}\right|_{x=x_i} (x - x_i)^2 + \ldots, \tag{2.11}$$

and likewise expanding in time, one can create the very simple — and unstable for $c > 0$ — explicit forward-time, forward-space method

$$q_i^{n+1} = q_i^n - \left(\frac{c\Delta t}{\Delta x}\right)\left(q_{i+1}^n - q_i^n\right) + \mathscr{O}\left(\Delta t^1, \Delta x^1\right). \tag{2.12}$$

Here, Δt is the time step size, and the notation $\mathscr{O}\left(\Delta t^1, \Delta x^1\right)$ means the method is first order accurate in space and time — while determining the derivatives, the Taylor series terms of order Δx and Δt, or higher, were neglected.

2.2.1 Finite Volume Methods

Alternatively, for a two-dimensional *finite volume* method we say that $q^n_{i,j}$ is associated with the spatial *average* of q over the (i,j)th grid zone at time t^n

$$q^n_{i,j} \approx \frac{1}{\Delta x \Delta y} \int_{x_{i-1/2}}^{x_{i+1/2}} \int_{y_{j-1/2}}^{y_{j+1/2}} q(x,y,t^n)\, dxdy, \tag{2.13}$$

where $\Delta x \Delta y$ is the area of the (i,j)th cell with $\Delta x = x_{i+1/2} - x_{i-1/2}$ and $\Delta y = x_{j+1/2} - x_{j-1/2}$ the grid spacing in the x and y directions, respectively. Figure 2.1 shows an example two-dimensional, uniform grid where the cell centers are indicated with integer indices while cell edges are denoted by half-integer indices.

The advantage of using the finite volume approach over the finite difference approach is that the variables we are storing — the $q_{i,j}$'s — are essentially the terms on the left-hand side of the integral form of the conservation law 2.3. For example, Equation 2.3 applied to the (i,j)th cell of Figure 2.1 gives

$$\frac{\partial q_{i,j}}{\partial t} = \frac{1}{\Delta x \Delta y} \left[\int_{y_{j-1/2}}^{y_{j+1/2}} \left(F_{x,i-1/2} - F_{x,i+1/2} \right)\, dy + \int_{x_{i-1/2}}^{x_{i+1/2}} \left(F_{y,j-1/2} - F_{y,j+1/2} \right)\, dx \right], \tag{2.14}$$

where $F_{x,i} = \mathbf{F}(x_i, y, q, t) \cdot \mathbf{e}_x$ and $F_{y,j} = \mathbf{F}(x, y_j, q, t) \cdot \mathbf{e}_y$ are the fluxes through the cell edges. Integrating this in time from t^n to t^{n+1} gives the basic form of a finite volume method for a two-dimensional conservation law on a uniform cartesian grid

$$q^{n+1}_{i,j} = q^n_{i,j} + \frac{\Delta t}{\Delta x} \left(\mathscr{F}^{n+1/2}_{x,i-1/2,j} - \mathscr{F}^{n+1/2}_{x,j+1/2,j} \right) + \frac{\Delta t}{\Delta y} \left(\mathscr{F}^{n+1/2}_{y,i,j-1/2} - \mathscr{F}^{n+1/2}_{y,i,j+1/2} \right), \tag{2.15}$$

where the flux functions, $\mathscr{F}$, are approximations to an average flux along an edge of a cell during

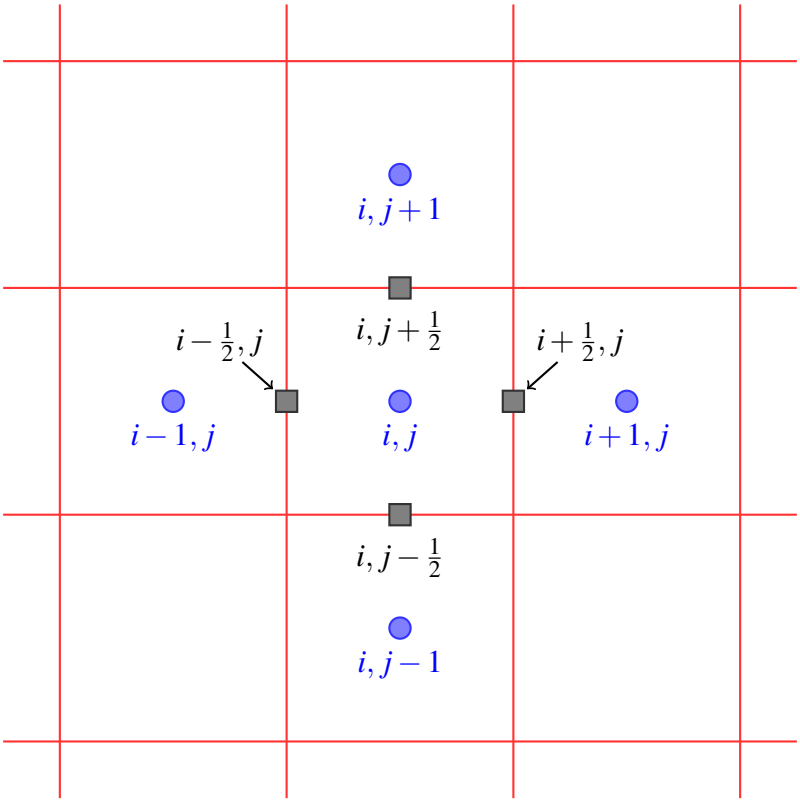

Figure 2.1: Example of a two-dimensional cartesian grid structure showing the locations of integer and half-integer indices used in finite volume methods.

a time step (c.f. Chapter 19 of LeVeque, 2002)

$$\mathscr{F}^{n+1/2}_{x,i-1/2,j} = \frac{1}{\Delta t \Delta y} \int_{t^n}^{t^{n+1}} \int_{y_{j-1/2}}^{y_{j+1/2}} F_{x,i-1/2}\, dydt \tag{2.16}$$

$$\mathscr{F}^{n+1/2}_{y,i,j-1/2} = \frac{1}{\Delta t \Delta x} \int_{t^n}^{t^{n+1}} \int_{x_{i-1/2}}^{x_{i+1/2}} F_{y,j-1/2}\, dxdt. \tag{2.17}$$

The flux functions 2.16 and 2.17 are typically unknown, but are constructed from the values of q^n surrounding the edge of interest. The manner in which these fluxes are reconstructed determines the order of accuracy for the method.

Godunov (1959) proposed a finite volume method where the flux functions are based on solu-

tions to the Riemann problem at cell edges. To set up the Riemann problem (in one dimension) at the edge $x_{i-1/2}$, the profile of $q(x,t)$ needs to be reconstructed within the $(i-1)$th and ith zones. Figure 2.2 shows this reconstruction process. In red is the (unknown) analytic function that needs reconstructed. The circles at cell centers are the q_i^n — the average of the analytic function over the ith zone — given by Equation 2.13. The black lines show Godunov's original method of using a piecewise constant reconstruction function, which results in a first-order accurate method. The Riemann problem at $x_{i-1/2}$ for this reconstruction is then the hyperbolic equation with initial left and right states given by the (black-labeled) $q_{i-1/2,L}$ and $q_{i-1/2,R}$, respectively.

To achieve a higher-order Godunov method, such as what is used in the `MAESTRO` code, a higher-order reconstruction of $q(x,t)$ is needed. The blue lines in Figure 2.2 show an example of a piecewise linear reconstruction and the accompanying left and right states for the Riemann problem at $x_{i+1/2}$ using this reconstruction. Higher-order reconstructions use values of q in neighboring cells to determine the profile — for example, in Figure 2.2, the piecewise linear reconstruction in the ith cell uses the values of q_i and q_{i+1} to determine the slope of the line. Other higher-order reconstructions use different combinations of neighboring points to generate the piecewise reconstruction function. Ideally, one would like the reconstructed value of $q_{i+1/2}$ to lie between the cell-centered values q_i and q_{i+1} to not introduce any local extrema that may grow with evolution of the system. A general reconstruction that satisfies Equation 2.13, does not automatically satisfy $q_i \leq q_{i+1/2} \leq q_{i+1}$, especially near sharp discontinuities in the initial data.. To enforce this constraint, one invokes a *slope limiter*, which restricts the slopes of the reconstruction to ensure that the algorithm is monotonicity-preserving and does not introduce oscillations (see Section 6.7 of LeVeque (2002), for example).

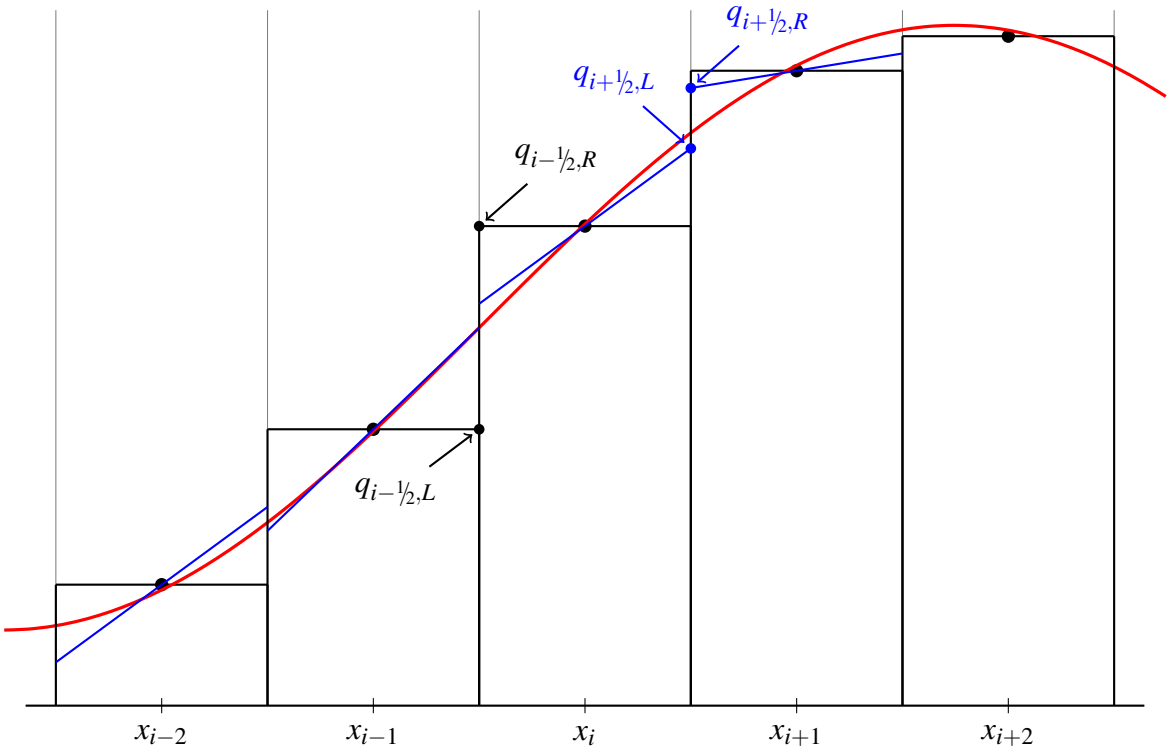

Figure 2.2: Example reconstruction of a function (red) using piecewise constant values (black) or a piecewise linear function (blue). The labeled points, $q_{*,L}$ and $q_{*,R}$, show the various edge states used in the Riemann solve for the fluxes in a Godunov method; the color of the label relates to the reconstruction method.

2.2.2 The CFL Condition

Up until now, we have not said anything about any restrictions placed upon the time step size. One might ask: "if I am interested in the evolution of the system from $t = t_0$ through some later time $t = t_{\text{end}}$, why not take a single time step of size $\Delta t = t_{\text{end}} - t_0$ and be done with it?"

In 1928, Courant, Friedrichs, and Lewy discovered a necessary condition — the so called CFL condition — for convergence of an explicit numerical method, in the limit of Δt and Δx go to zero: the domain of dependence of the original partial differential equation must be contained within the numerical domain of dependence of the method (Courant et al., 1928). The domain of dependence for a hyperbolic partial differential equation is essentially the set of points along the characteristic

curves that could affect the solution given the fact that information propagates at a finite speed. As an example, consider the one-dimensional advection equation of 2.2. The characteristic curves for this equation are the lines $x - ct = 0$; the domain of dependence at some point x_p and time t_p is then the point $x_p - ct_p$.

The numerical domain of dependence obviously depends on the choice of numerical method. If we apply the finite volume method from 2.15 to the advection equation of 2.2, realizing that the flux is simply $\mathbf{F} = cq\mathbf{e}_x$ then we get

$$q_i^{n+1} = q_i^n + \frac{c\Delta t}{\Delta x}\left(q_{i-1/2}^n - q_{i+1/2}^n\right). \tag{2.18}$$

For simplicity, let us define $q_{i+1/2}^n = 1/2\left(q_{i+1}^n + q_i^n\right)$. Therefore, numerically, the only points that can affect the solution of q_i in a single time step — $\{q_{i-1}, q_i, q_{i+1}\}$ — are within $x_i - \Delta x \leq x \leq x_i + \Delta x$. If we take $n = t_p/\Delta t$ time steps to reach time t_p, then the only points which can affect the solution at x_p — $\left\{q_{p-n}, q_{p-n+1}, \ldots, q_{p+n-1}, q_{p+n}\right\}$ — are within $x_p - n\Delta x \leq x \leq x_p + n\Delta x$. For the CFL condition to be satisfied and the method to be stable, we must have $x_p - n\Delta x \leq x_p - ct_p \leq x_p + n\Delta x$ or

$$\left|\frac{c\Delta t}{\Delta x}\right| \leq 1. \tag{2.19}$$

The ratio $|c\Delta t/\Delta x| \equiv \alpha_{\text{CFL}}$ is called the Courant or CFL number. Again, the CFL condition is a *necessary* condition for stability, and, in practice for such a method, a CFL number of no greater than about $\alpha_{\text{CFL}} \sim 0.9$ is used. Therefore, in practice the time step size is usually restricted by

$$\Delta t \lesssim \alpha_{\text{CFL}} \frac{\Delta x}{c}, \tag{2.20}$$

even for a right-hand side of 2.19 that is not equal to one.

2.3 Low Mach Number Approximation

For a system of equations such as the Euler equations, the CFL condition restricts the time step size by the maximum characteristic speed, $U + c_s$, where $U = |\mathbf{U}|$ is the magnitude of the velocity vector. Note that if the Mach number, $M \equiv U/c_s$, is small ($\ll 1$) — as is the case for the convection and burning occurring in XRBs — then the CFL restriction on the time step is entirely dominated by the propagation of sound waves

$$\Delta t \lesssim \alpha_{\mathrm{CFL}} \frac{\Delta x}{c_s (1+M)} \approx \alpha_{\mathrm{CFL}} \frac{\Delta x}{c_s}, \qquad \text{for } M \ll 1. \tag{2.21}$$

In low Mach number fluid flow, these sound waves may not be important, and the CFL condition may be overly restrictive for efficiently evolving the system.

To overcome this restriction, different forms of the so-called *low Mach number approximation* methods have been developed. These methods tend to filter the acoustic waves from the system, thus allowing the CFL time step to be restricted by the *dynamics* of the fluid rather than the propagation of sound waves. A classic example is the incompressible fluid, which has the constraint

$$\nabla \cdot \mathbf{U} = 0 \tag{2.22}$$

that doesn't allow for any form of compressibility effects along a streamline. Another example is the anelastic approximation (Ogura & Phillips, 1962), which allows for large-scale background stratification of ρ and p, but assumes that the thermodynamic perturbations about this background are small. The anelastic approximation appears as the constraint

$$\nabla \cdot (\rho_0 \mathbf{U}) = 0, \tag{2.23}$$

where ρ_0 is the background density profile. More complex methods that filter the acoustics but allow more compressible effects can be constructed.

More recently, Lin et al. (2006) developed and applied a low Mach number approximation method to the problem of convective burning at the base of an accreted layer in an XRB system. Their method, however, was first-order accurate in space and time and did not allow for the evolution of the hydrostatic base state, a feature that is needed to capture the expansion of the atmosphere in response to heating. Furthermore, Lin et al. did not model the surface of the accreted layer, which is vital to understanding bursts that exhibit photospheric radius expansion; such bursts are crucial in determining the stellar properties of neutron stars (see Section 1.3 and Steiner et al., 2010, and references therein).

CHAPTER 3

MAESTRO

In this Chapter, I give an overview of the low Mach number approximation method of the MAESTRO code used for the simulations comprising this dissertation. A series of papers (see Almgren et al. (2006a) — henceforth Paper I, Almgren et al. (2006b) — henceforth Paper II, Almgren et al. (2008) — henceforth Paper III, and Zingale et al. (2009) — henceforth Paper IV) describe the derivation of the low Mach number equation set, its algorithmic implementation, and the initial application to convection in a white dwarf preceding a Type Ia supernova. Here, I will describe the most recent version of the algorithm as outlined in our paper Nonaka et al. (2010) — henceforth Paper V. For the XRB problem, I am only concerned with the upper $\lesssim 20$ m of atmosphere on a ~ 10 km radius neutron star. Therefore, I use a plane-parallel atmosphere with x the horizontal coordinate and r the radial coordinate, which is how I will describe the algorithm below.

At the start of this dissertation work, the core of the MAESTRO algorithm had already been developed. My role was to take the lead on applying MAESTRO to the XRB problem, improving problem-specific aspects of the code to make such simulations meaningful, and validating the applicability of MAESTRO to this problem. Applying the code to flows under the extreme conditions found in a neutron star atmosphere really tested the limits of the MAESTRO algorithm. This

stressing of the code tended to discover unforeseen issues, the resolution of which lead towards a more robust algorithmic development and a deeper understanding of the low Mach number approximation method used in the code. Another of my roles involved implementing and validating the microphysics (e.g. the reaction networks described in section 4.4) that is used in MAESTRO for the XRB problem.

3.1 The Equations

The fully compressible equations for a non-viscous, reacting astrophysical fluid are

$$\frac{\partial(\rho X_k)}{\partial t} + \nabla\cdot(\rho X_k \mathbf{U}) = \rho\dot{\omega}_k, \tag{3.1}$$

$$\frac{\partial(\rho \mathbf{U})}{\partial t} + \nabla\cdot(\rho \mathbf{U}\mathbf{U}) + \nabla p = -\rho g \mathbf{e}_r, \tag{3.2}$$

$$\frac{\partial(\rho E)}{\partial t} + \nabla\cdot(\rho \mathbf{U} E + p\mathbf{U}) = \nabla\cdot(k_{\rm th}\nabla T) - \rho g(\mathbf{U}\cdot\mathbf{e}_r) + \rho H_{\rm nuc}, \tag{3.3}$$

where as before $\rho, \mathbf{U}$, and p are the mass density, velocity vector, and pressure respectively. X_k and $\dot{\omega}_k$ are the mass fraction and production rate of the kth species, $g(r)$ is the radial-dependent gravitational acceleration, $E = e + \frac{1}{2}\mathbf{U}\cdot\mathbf{U}$ is the total specific energy with e the specific internal energy, $k_{\rm th}$ is the thermal conductivity, T is the temperature, and $H_{\rm nuc}$ is the specific energy generation rate from nuclear reactions. The set of equations is closed by an equation of state

$$p = p(\rho, T, X_k). \tag{3.4}$$

Note that the equation of continuity, 2.6, is contained within Equation 3.1 because the density can be defined as $\rho = \sum_i \rho X_i$, and likewise $\sum_i \dot{\omega}_i = 0$.

In the low Mach number method presented here, we prefer to work in terms of specific enthalpy,

$h = e + p/\rho$, instead of total specific energy. Using low Mach number asymptotics (see Paper I) the total pressure, $p(x,r,t)$, is decomposed into a one-dimensional base state pressure, $p_0(r,t)$, and a perturbational, or dynamic, pressure, $\pi(x,r,t)$, such that $|\pi|/p_0 = \mathscr{O}(M^2)$, with M the Mach number. The one-dimensional base state density, $\rho_0(r,t)$, is in hydrostatic equilibrium (HSE) with the base state pressure such that $\nabla p_0 = -\rho_0 g \mathbf{e}_r$. With these modifications, Equations 3.2 and 3.3 become

$$\frac{\partial \mathbf{U}}{\partial t} = -\mathbf{U}\cdot\nabla\mathbf{U} - \frac{1}{\rho}\nabla\pi - \frac{(\rho-\rho_0)}{\rho} g\mathbf{e}_r, \tag{3.5}$$

$$\frac{\partial(\rho h)}{\partial t} = -\nabla\cdot(\rho h\mathbf{U}) + \frac{Dp_0}{Dt} + \rho H_{\text{nuc}} + \nabla\cdot(k_{\text{th}}\nabla T), \tag{3.6}$$

where $D/Dt = (\partial/\partial t + \mathbf{U}\cdot\nabla)$ is the Lagrangian derivative. The constraint imposed by the equation of state, 3.4, is re-expressed as a divergence constraint on the velocity field (see Paper III), analogous to the divergence constraints (2.22 and 2.23) of the low Mach number methods described in Section 2.3:

$$\nabla\cdot(\beta_0\mathbf{U}) = \beta_0\left(S - \frac{1}{\overline{\Gamma}_1 p_0}\frac{\partial p_0}{\partial t}\right), \tag{3.7}$$

where β_0 is a density-like variable,

$$\beta_0(r,t) = \rho(0,t)\exp\left(\int_0^r \alpha\frac{\partial p_0}{\partial r'}\,dr'\right). \tag{3.8}$$

Here, we define for convenience $\alpha \equiv \left(\overline{\Gamma}_1 p_0\right)^{-1}$ with $\Gamma_1(r) = (d\ln p/d\ln\rho)_s$, where the subscript s means the derivative is taken at constant entropy. The overline notation denotes a lateral average

$$\overline{\phi}(r,t) = \frac{\int \phi(x,r,t)\,dx}{\int dx}, \tag{3.9}$$

where the integration is taken over the entire domain width. The expansion term, S, accounts for

local compressibility effects resulting from nuclear burning, compositional changes, and thermal conduction:

$$S = \sigma H_{\rm nuc} + -\sigma \sum_k \xi_k \dot{\omega}_k + \frac{1}{\rho p_\rho} \sum_k p_{X_k} \dot{\omega}_k + \frac{\sigma}{\rho} \nabla \cdot (k_{\rm th} \nabla T), \tag{3.10}$$

where $\xi_k \equiv (\partial h / \partial X_k)_{p,T,(X_j, j \neq k)}$, $p_\rho \equiv (\partial p / \partial \rho)_{T,X_k}$, $p_{X_k} \equiv (\partial p / \partial X_k)_{T,\rho,(X_j, j \neq k)}$ and $\sigma \equiv p_T / (\rho c_p p_\rho)$ with $p_T \equiv (\partial p / \partial T)_{\rho,X_k}$, and $c_p \equiv (\partial h / \partial T)_{p,X_k}$.

It should be noted here that the original derivation of the form of α in Paper I assumed the system was in chemical equilibrium. I have since derived the form of α that relaxes this assumption and incorporates the effects of chemical potential. This derivation, along with an explanation of when this new non-equilibrium definition of α is important, can be found in Appendix A.

To capture the base state evolution, the velocity field is decomposed into a one-dimensional base state velocity, $w_0(r,t)$, and a local velocity, $\widetilde{\mathbf{U}}(x,r,t)$, that governs the local dynamics

$$\mathbf{U}(x,r,t) = w_0(r,t)\mathbf{e}_r + \widetilde{\mathbf{U}}(x,r,t), \tag{3.11}$$

with

$$w_0(r,t) = \overline{\mathbf{U} \cdot \mathbf{e}_r}, \tag{3.12}$$

and therefore $\overline{\widetilde{\mathbf{U}} \cdot \mathbf{e}_r} = 0$. Each term in the decomposition gets its own version of the velocity equation, 3.5, and constraint equation, 3.7. As in Paper II, a one-dimensional base state enthalpy, $(\rho h)_0$, is defined as the average of the full enthalpy, (ρh). The evolution of the base state density and enthalpy are then obtained by averaging Equations 2.6 and 3.6, respectively. The evolution of the base state pressure is given by the base state density evolution and the HSE constraint.

Thermal conduction was not present in Paper V, so a semi-implicit discretization for this term was developed for use in the pure ^{4}He accreting XRBs discussed in Chapter 5. I implemented into

MAESTRO a diffusion test problem, which is described in Appendix B.

Another addition to the MAESTRO algorithm from Paper V is the use of a "volume discrepancy" correction. Because Equation 3.7 is a linearization of the nonlinear constraint imposed by the equation of state, the thermodynamic pressure, $p_{\rm EOS} = p(\rho, h, X_k)$, may drift from the base state pressure, p_0, (Pember et al., 1998). To correct for this drift, Equation 3.7 is augmented with a term that drives the thermodynamic pressure back to that of the base state:

$$\nabla \cdot (\beta_0 \mathbf{U}) = \beta_0 \left(S - \frac{1}{\overline{\Gamma}_1 p_0} \frac{\partial p_0}{\partial t} - \frac{f}{\overline{\Gamma}_1 p_0} \frac{p_0 - p_{\rm EOS}}{\Delta t} \right), \tag{3.13}$$

where f is the volume discrepancy correction factor and $0 \leq f \leq 1$. In Section 5.2.4, we explore the effectiveness of this term at keeping the overall solution in thermodynamic equilibrium.

To summarize, the low Mach number equation set used in MAESTRO is given by Equations 3.1, 3.5, 3.6, and 3.13. As mentioned in Section 2.3, one key advantage of using a low Mach number approach is the increase of allowable time step size, which enables long-time integration. Standard compressible hydrodynamics codes for astrophysical applications, such as CASTRO (Almgren et al., 2010) or FLASH (Fryxell et al., 2000), evolve a fully compressible equation set, i.e., the Euler equations, which allows for the formation and propagation of shocks. Our low Mach number equation set does not contain acoustic waves, and therefore MAESTRO is able to take time steps constrained by the maximum fluid velocity, rather than the maximum sound speed. As an example, if the maximum Mach number of the flow is $M \sim 0.01$, — which is common (see, e.g. Figure 5.16) — we will obtain a factor of $1/M \sim 100$ increase in time step size compared to a standard compressible approach. In fact, I first started working on the XRB simulations using the FLASH code. I quickly realized that such calculations using a compressible hydrodynamics code were infeasible, due to the CFL-restricted time step size, which was on the order of nanoseconds! Another advantage of a low Mach number method is that the overall HSE of the state can be guaranteed by

the inclusion of a base state in HSE in the low Mach number equation set, thereby removing the difficulties of maintaining HSE commonly found in compressible hydrodynamics codes.

The equation of state used in the simulations presented in this dissertation was the general stellar equation of state of Timmes & Swesty (2000). This EOS contains contributions from an ideal gas of nuclei, radiation, and an electron-positron gas with an arbitrary degree of degeneracy and relativity. Conductivities are calculated using Frank Timmes's publicly available conductivity routine, which includes contributions from radiation and electron conduction processes as explained in Timmes (2000).

3.2 The Algorithm

Mathematically, MAESTRO solves a system of advection-reaction-diffusion equations with the equation of state formulated as an elliptic constraint on the velocity. MAESTRO uses a higher-order Godunov method to discretize the advective terms, Strang-splitting to couple the reaction terms to the advective terms, and a semi-implicit treatment of the diffusion terms. The diffusion term and the divergence constraint are formulated as linear systems which are solved iteratively using a multigrid technique. The evolution of the one-dimensional base state density is also computed. The base state density represents the average state of the atmosphere, and is coupled to the base state pressure via HSE. The base state density has its own evolution equation that computes the expansion of the atmosphere due to heating and is discretized using a higher-order Godunov method. MAESTRO is second-order accurate in space and time.

To enforce the divergence constraint on the velocity field, Equation 3.7 (or 3.13), MAESTRO uses projection methods. To illustrate how a projection operator works, consider an incompressible

fluid. The divergence constraint for an incompressible fluid is, from Equation 2.22

$$\nabla \cdot \mathbf{U} = 0.$$

The time evolution of the velocity field is governed by the momentum equation, the solution of which need not satisfy the incompressible divergence constraint. In its simplest incarnation, a projection method relies on the fact that any vector field can be decomposed into a divergence-free vector field and the gradient of a scalar field:

$$\mathbf{U} = \mathbf{U}^d + \nabla \phi, \tag{3.14}$$

where $\nabla \cdot \mathbf{U}^d = 0$ and therefore

$$\nabla^2 \phi = \nabla \cdot \mathbf{U}. \tag{3.15}$$

Given the incompresible fluid divergence constraint, what we want to do is *project* the velocity field solution of the momentum equation, $\mathbf{U}$, onto a divergence-free vector field, $\mathbf{U}^d$. This can be accomplished by solving Equation 3.15 for ϕ and then rearranging Equation 3.14

$$\mathbf{U}^{n+1} = \mathbf{U}^d = \mathbf{U} - \nabla \phi. \tag{3.16}$$

For the divergence constraint used in MAESTRO, Equation 3.7, the situation is complicated by a non-zero right-hand side. Nevertheless, the same approach applies: solve an elliptic equation for a scalar field, the gradient of which can be subtracted from the provisional velocity field to construct a velocity field that satisfies the divergence constraint. MAESTRO uses a multigrid solver for the solution of the elliptic equation — see Appendix B of Paper III for details on the exact form of the projection operator.

To begin a two-dimensional, plane-parallel simulation, a one-dimensional model is mapped laterally across the simulation domain. The evolution of this model through a single time step proceeds in a predictor-corrector fashion, solving for the full state quantities $\left(\widetilde{\mathbf{U}}, \rho, X, T, h\right)$ and base state quantities $(w_0, \rho_0, (\rho h)_0, p_0)$ at the new time. Experience has shown that the slope-limiters are more effective when perturbational quantities, denoted as $\rho' \equiv \rho - \rho_0$ and $\rho h' \equiv \rho h - (\rho h)_0$, are advanced instead of the base state. Evolution equations for perturbational quantities can be obtained by subtracting the base state evolution equations from the full state equations. The perturbational quantities are predicted to cartesian edges to calculate fluxes for the conservative update of the base state quantities.

In the predictor phase of the algorithm, a provisional, time-centered expansion term, S, is calculated and used to construct provisional advective velocities, which are forced to satisfy the divergence constraint via a projection operator. Using these velocities, cell edge-predicted perturbational quantities are constructed for use in a conservative update of the full and base states, yielding a preliminary estimate of the solution at the new time step. The corrector phase of the algorithm uses this preliminary solution to improve the time-centered expansion and advective velocity terms, which in turn are used to construct the final solution vector.

3.2.1 Handling Low Density Regions

In the atmosphere of a neutron star, the density spans several orders of magnitude over a short distance, and special care is needed in these regions. Due to conservation of momentum, large velocities in the upper atmosphere, which do not affect the solution in the higher-density region, may be generated. Unfortunately, these large velocities reduce the efficiency of `MAESTRO` method by reducing the allowable time step size. The first technique used to address this problem is the utilization of a cutoff density, ρ_{cutoff}, which is the value the density is fixed at outside the star.

The second technique is the use of an anelastic cutoff density, $\rho_{\rm anelastic}$, below which we determine β_0 by keeping the ratio β_0/ρ_0 constant in the divergence constraint in order to minimize spurious wave generation. Full implementation details for the cutoff densities are described in Appendix A.5 of Paper V. For the simulations in this dissertation, I use $\rho_{\rm cutoff} = \rho_{\rm anelastic} = 10^4$ g cm^{-1}.

The third technique adopted for the low density region is sponging (or damping), which is used to reduce gravity waves at the surface of the star. This technique is commonly used in the atmospheric modeling community as lateral boundary conditions of limited area simulations (see, for example, Kesel & Winninghoff 1972; Perkey & Kreitzberg 1976) as well as upper boundary conditions to reduce wave reflection off of sharp gradients in the atmospheric structure (see, for example, Durran & Klemp 1983; Durran 1990; Chen et al. 2005). In addition, the sponging technique in MAESTRO has successfully been used in the study of convection in the cores of white dwarfs (Paper IV). Full details for the sponge implementation in MAESTRO can be found in Papers III and IV, but in summary, the velocity is augmented with a forcing term, which effectively damps the velocity so that $\mathbf{U}^{\rm new} \rightarrow \mathbf{U}^{\rm new} * f_{\rm damp}$. In the XRB simulations of Chapters 5 and 6, I use the following formulation for the sponge:

$$f_{\rm damp} = \begin{cases} 1, & r \leq r_{\rm sp}, \\ \frac{1}{2}\left(1 - f_{\rm damp,min}\right)\cos\left[\pi\left(\frac{r-r_{\rm sp}}{r_{\rm tp}-r_{\rm sp}}\right)\right] + \frac{1}{2}\left(1 + f_{\rm damp,min}\right), & r_{\rm sp} < r \leq r_{\rm tp}, \\ f_{\rm damp,min}, & r_{\rm tp} < r, \end{cases} \tag{3.17}$$

where $r_{\rm sp}$ is the radius at which $\rho_0 = 25\rho_{\rm cutoff}$, $r_{\rm tp}$ is the radius at which $\rho_0 = \rho_{\rm cutoff}$, and $f_{\rm damp,min} = 0.01$.[1] Note that as the system evolves it is free to expand, thus changing the location of the density cutoffs and, consequently, the location and extent of the sponge. The inclusion of a sponge layer

[1] Note that the form of this sponge is similar to that presented in Section 4.3.1 of Paper III but with $\kappa\Delta t = 1$ at each time step.

does not strictly conserve kinetic energy in the sponged region; however, the material above the surface of the star is at a relatively low density compared to the material in the convective region, and therefore the total amount of energy non-conservation is small. Furthermore, as shown in Figure 4 of Paper III, the inclusion of a sponge layer in the low density region of a simulation does not affect the dynamics of the flow in the convective region of interest.

CHAPTER 4

THERMONUCLEAR REACTION NETWORKS

In this Chapter, I outline — from a simulation point of view — some of the important aspects of thermonuclear reactions and reaction networks. I discuss some of the important reaction rates for XRBs, as well as some of the reaction networks I have implemented into `MAESTRO`. For a more indepth discussion of thermonuclear reactions and the nuclear physics involved, see Clayton's classic book on stellar evolution (Clayton, 1983). For a more modern view of reaction networks, see either Dave Arnett's book on supernovae (Arnett, 1996) or Frank Timmes's paper on network integration (Timmes, 1999) and excellent set of notes (and references therein) online at `http://cococubed.asu.edu/talk_pages/nnpss.shtml`.

Throughout the remainder of this dissertation, we will use the words "species", "isotope", and "nuclei" interchangeably.

4.1 Basics of Thermonuclear Reactions

Consider a gas of various types of particles in thermodynamic equilibrium. Consider a reaction where the ith particle collides with the jth producing the two particles k and l; such a reaction can be written as

$$i + j \longrightarrow k + l, \tag{4.1}$$

or $i(j,k)l$ for short. We define the *cross section* of the reaction, σ, to be

$$\sigma = \frac{\text{\# of reactions / particle } j \text{ / s}}{\text{incident flux of } i\text{'s}}. \tag{4.2}$$

Note that the incident flux depends on the *relative* velocity, v, between the ith and jth particles, and therefore the cross section is a function of the relative velocity, $\sigma = \sigma(v)$. The relative velocity between the two particles will have some distribution, $\phi(v)$. We can then, using 4.2, write the total number of reactions per unit volume per second between particle i and j as

$$r_{i,j} = \int_0^\infty n_j \sigma(v) \underbrace{n_i v \phi(v)}_{\text{flux term}} \, dv = n_i n_j \langle v\sigma(v) \rangle_{i,j}, \tag{4.3}$$

where n_i is the number density of particle i, and the angled brackets denote the average over the distribution (c.f. Section 4.2 of Clayton, 1983). Equation 4.3 gives the total reaction rate per unit volume for reaction 4.1. Note that the product $n_i n_j$ is equal to the total number of unique pairs of particles to react; therefore, if the ith and jth particles are identical particles, the total number of pairs is $\left(n_i^2/2!\right)$, and 4.3 needs to be modified. The same type of modification for reactions involving three identical particles, such as the 3–α rate, need to be used, but there the coefficient is $(1/3!)$ instead of $(1/2!)$ (e.g. Section 4.1 Arnett, 1996). This reaction then gives rise to changes in the number density of each of the reactants and products of the form

$$\begin{aligned} \dot{n}_i &= -r_{i,j} \\ \dot{n}_j &= -r_{i,j} \\ \dot{n}_k &= r_{i,j} \\ \dot{n}_l &= r_{i,j}, \end{aligned} \tag{4.4}$$

where the dot notation means the full time derivative. In practice, we calculate the changes in the molar fractions, Y_i, which are related to the mass fractions

$$Y_i = \frac{n_i}{\rho N_{\rm A}} = \frac{X_i}{A_i}, \tag{4.5}$$

where A_i is the mass number — the sum of protons and neutrons — for isotope i, and $N_{\rm A}$ is Avogadro's number. Using this notation, the system of ordinary differential equations (ODEs) 4.4 can be written as

$$\begin{aligned}
\dot{Y}_i &= -Y_i Y_j \rho \langle N_{\rm A} v\sigma \rangle_{i,j} \\
\dot{Y}_j &= -Y_i Y_j \rho \langle N_{\rm A} v\sigma \rangle_{i,j} \\
\dot{Y}_k &= Y_i Y_j \rho \langle N_{\rm A} v\sigma \rangle_{i,j} \\
\dot{Y}_l &= Y_i Y_j \rho \langle N_{\rm A} v\sigma \rangle_{i,j}.
\end{aligned} \tag{4.6}$$

If instead of 4.1 we were dealing with a general reaction of the form

$$\alpha i + \beta j \longleftrightarrow \gamma k + \delta l \tag{4.7}$$

where α, β, γ, and δ were the stoichiometric coefficients of the reaction, with all species being unique, we would have the following system of ODEs:

$$\begin{aligned}
\dot{Y}_i &= -\alpha Y_i Y_j \rho \langle N_{\rm A} v\sigma \rangle_{i,j} + \alpha Y_k Y_l \rho \langle N_{\rm A} v\sigma \rangle_{k,l} \\
\dot{Y}_j &= -\beta Y_i Y_j \rho \langle N_{\rm A} v\sigma \rangle_{i,j} + \beta Y_k Y_l \rho \langle N_{\rm A} v\sigma \rangle_{k,l} \\
\dot{Y}_k &= \gamma Y_i Y_j \rho \langle N_{\rm A} v\sigma \rangle_{i,j} - \gamma Y_k Y_l \rho \langle N_{\rm A} v\sigma \rangle_{k,l} \\
\dot{Y}_l &= \delta Y_i Y_j \rho \langle N_{\rm A} v\sigma \rangle_{i,j} - \delta Y_k Y_l \rho \langle N_{\rm A} v\sigma \rangle_{k,l}.
\end{aligned} \tag{4.8}$$

A nuclear reaction *network* consists of a set of reactions of the form 4.7 describing the general evolution of the species via a system of ODEs similar to 4.8, but with contributions from every reaction included on the right-hand side of the ODEs.

All of the detailed information about quantum mechanical tunneling probabilities, nuclear resonances, degeneracy, etc. are all contained in the averaging of $\langle v\sigma \rangle_{i,j}$. Some nuclear physics experiments are designed around obtaining values for this quantity for a particular reaction. Most of the rates used in the networks in `MAESTRO` are from the Caughlan and Fowler paper, which is a large compilation of $\langle N_{\mathrm{A}} v\sigma \rangle$'s as functions of temperature for various nuclear reactions (Caughlan & Fowler, 1988). There are also online repositories, such as the JINA Reaclib database (Cyburt et al., 2010), which catalogues both experimental and theoretical rates; most of the "suggested rates" from the Reaclib database for the reactions we consider here are actually the rates of Caughlan and Fowler.

In an astrophysical setting, the gas of nuclei is embedded inside a dense gas of electrons, as the overall charge of the plasma is neutral. Each nucleus, then, is surrounded by a cloud of electrons, the net effect of which is to *decrease* the Coulomb barrier between two nuclei. The reduction of this barrier causes an increase in the probability that the two nuclei will interact via tunneling and boosts the overall reaction rate above that of the laboratory reaction rate (Clayton, 1983, Section 4.8). This *screening enhancement* manifests itself by applying a multiplicative factor to the laboratory reaction rate. In `MAESTRO` we currently use an admittedly dated screening routine for both the weak and strong screening regimes (based on Graboske et al., 1973; Alastuey & Jancovici, 1978; Itoh et al., 1979).

4.2 Solving Reaction Networks In MAESTRO

The solution to a reaction network over some time step is equivalent to solving the first-order, nonlinear system

$$\dot{\mathbf{Y}} = \mathbf{f}(\mathbf{Y}) \tag{4.9}$$

where the vector $\mathbf{Y}$ contains all the species. Furthermore, to avoid instabilities with the solution, an energy, entropy, or temperature equation should be solved simultaneously (see Mueller, 1986, for example); we choose to evolve a temperature equation because it will give us more consistent evaluation of the rates, which are explicit functions of temperature.[1] To further complicate things, the system of ODEs is *stiff* — one measure of stiffness is that the ratio of the real part of the maximum eigenvalue to the real part of the minimum eigenvalue of the Jacobian matrix, $\mathbf{J} \equiv \partial \mathbf{f}(\mathbf{Y})/\partial \mathbf{Y}$, is very large (Timmes, 1999). Interpreted physically, the stiffness of the Jacobian is an indication that some species is/are changing significantly faster than some other species.

To efficiently and stably solve this stiff system of ODEs, we use the double precision form of the VODE package (Brown et al., 1989). The VODE solver uses an implicit, multistep, Backward Differentiation Formula method, which typically solves the system to fifth order. The fact that the method is implicit implies some form of iterative scheme for solving the nonlinear system. The iterative scheme used is a modified Newton iteration, which involves evaluation of the Jacobian matrix $\mathbf{J}$. VODE is intelligent enough to adaptively decrease its step size or drop order in regions of computational difficulty to help ensure convergence.

At the end of the burn step, VODE gives the updated composition vector $\mathbf{Y}$. If there are no weak

[1] Note that in MAESTRO, at the end of the burn step we throw out any changes in temperature. The temperature is updated via an EOS call using either the pressure or enthalpy after they have been updated.

reactions, then the specific energy generation rate is simply

$$H_{\rm nuc} = \sum_i \frac{\Delta X_i q_i}{\Delta t}, \tag{4.10}$$

where ΔX_i is the change in *mass* fractions over the burn step, and q_i is the binding energy per gram of isotope i. If there are weak reactions, then neutrinos can carry away some fraction of the energy. In this case, an energy equation is also integrated to keep track of the neutrino loss rate at each step in `VODE`'s multistep procedure. Furthermore, the production rates used in the species equation, 3.1, are given by

$$\dot{\omega}_i = \frac{\Delta X_i}{\Delta t}. \tag{4.11}$$

4.3 Important Reactions for XRBs

Depending on the composition of the accreted material and the accretion rate, various reactions could dominate the burning and energy release (Bildsten, 2000, for example). For the case of a pure ^{4}He accreting system — such as `4U 1820-30` — considered in Chapter 5, the dominant burning mode is via the 3–α reaction $2\alpha\,(\alpha,\gamma)^{12}$C. Further α–chain reactions can occur, but the majority of the energy is liberated in the 3–α reaction. Furthermore, this burning tends to happen rapidly as there are no slow, weak reactions, and as a result the local luminosity can quickly reach — and sometimes exceed — the Eddington limit, giving rise to PRE bursts (see the discussion in Section 1.3).

For systems that accrete a solar-like composition, the nuclear reaction path is a bit more involved. At the high temperatures reached during an XRB ($T \gtrsim 10^8$ K), most hydrogen burns to helium via the temperature-independent, hot, β-decay-limited CNO cycle of Fowler & Hoyle (1965) shown in Figure 4.1. Helium burns mostly from the 3–α reaction, but at high enough

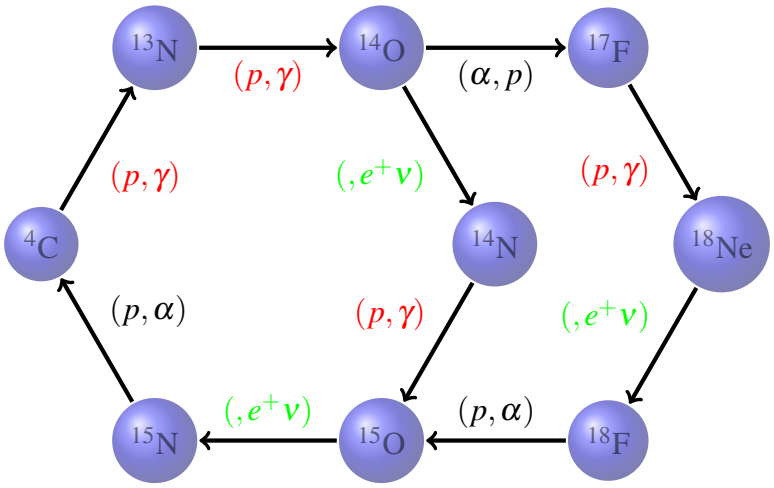

Figure 4.1: Diagram of the hydrogen-burning hot CNO cycle.

temperatures, further α–chain reactions occur. For temperatures $T \gtrsim 5 \times 10^8$ K, the α-capture reactions ^{15}O (α,γ) ^{19}Ne and ^{18}Ne (α,p) ^{21}Na occur more quickly than the β-decay reactions of ^{15}O and ^{18}Ne, respectively, causing the hydrogen burning to break out of the hot CNO cycle and begin burning along the rp-process of Wallace & Woosley (1981). From the long sequence of proton captures and β-decays, the rp-process produces elements far beyond the iron group.

4.4 Networks Added to MAESTRO

This section will cover some of the reaction networks I have added to the MAESTRO code, in part, for studying XRBs.

4.4.1 dvode_test network

This network is not actually based on a physical reaction network, but merely sets up a stiff system to test the solution of VODE method and its various tolerances. The system to be solved is

the simple reaction a $\longleftrightarrow$ b where the reaction rates are $r_{\mathrm{a,b}} \equiv \alpha$ and $r_{\mathrm{ba}} \equiv \beta$:

$$\frac{d}{dt}\begin{pmatrix} Y_{\mathrm{a}} \\ Y_{\mathrm{b}} \end{pmatrix} = \begin{pmatrix} -\alpha & \beta \\ \alpha & -\beta \end{pmatrix} \cdot \begin{pmatrix} Y_{\mathrm{a}} \\ Y_{\mathrm{b}} \end{pmatrix}. \tag{4.12}$$

Given the initial condition $Y_{\mathrm{b}}(t=0) = 0$, this system has an analytic solution of the form

$$\frac{Y_{\mathrm{b}}}{Y_{\mathrm{a}}} = \frac{e^{(\alpha+\beta)t} - 1}{\frac{\alpha}{\beta}e^{(\alpha+\beta)t} + 1}, \tag{4.13}$$

and tends toward a steady state solution of $(Y_{\mathrm{b}}/Y_{\mathrm{a}}) = (\alpha/\beta)$. Note that the matrix in Equation 4.12 is the Jacobian matrix, **J**, described in Section 4.2 for this reaction network. In that section, I described stiffness as a ratio of maximum to minimum eigenvalues. That method does not work here because one eigenvalue is zero and the other is positive. Using the physical interpretation of stiffness described in Section 4.2, I can set the stiffness of the system by fixing the ratio of the reaction rates to be a large number; for this test I set $(\alpha/\beta) = 2 \times 10^{15}$.

Figure 4.2 shows the time evolution of $Y_{\mathrm{b}}/Y_{\mathrm{a}}$ using a first-order explicit forward Euler method, a first-order implicit backward Euler method, the `VODE` solver, and the analytic solution. All of the methods tend towards the steady state solution, but the `VODE` solution follows the analytic solution the closest. Furthermore, the forward Euler method is very unstable for this stiff system, and wild oscillations and even negative numbers appear in the solution. The `dvode_test` network acts as an important unit test for the `VODE` solver, and provides assurance for all the other networks that utilize this solver.

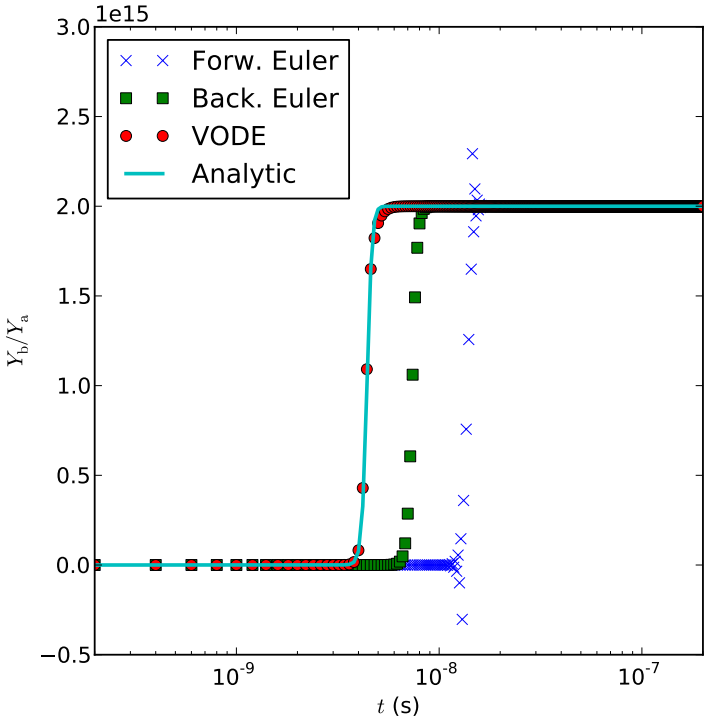

Figure 4.2: Evolution of the simple reaction network given by Equation 4.12. The Xs denote a first-order forward Euler method, the boxes a first-order backward Euler method, the circles the `VODE` solver method, and the analytic solution is the solid line.

4.4.2 `triple_alpha` and `fushiki_lamb` networks

For a pure ^{4}He accreting system discussed in Chapter 5, I implemented two variations of a 3–α burning network. The `triple_alpha` network uses the rates as compiled in Caughlan & Fowler (1988) with screening handled as described in Section 4.1. The `fushiki_lamb` reaction network, on the other hand, is based on a three-body S-matrix calculation of the 3–α reaction of Fushiki & Lamb (1987b). The S-matrix calculation uses experimental values for the widths of the nuclear states, and it includes screening effects in the interaction potential. Figure 4.3 shows a comparison between these two networks for the instantaneous energy generation rate in a pure ^{4}He gas — results from `triple_alpha` are shown as solid lines and the `fushiki_lamb` results are

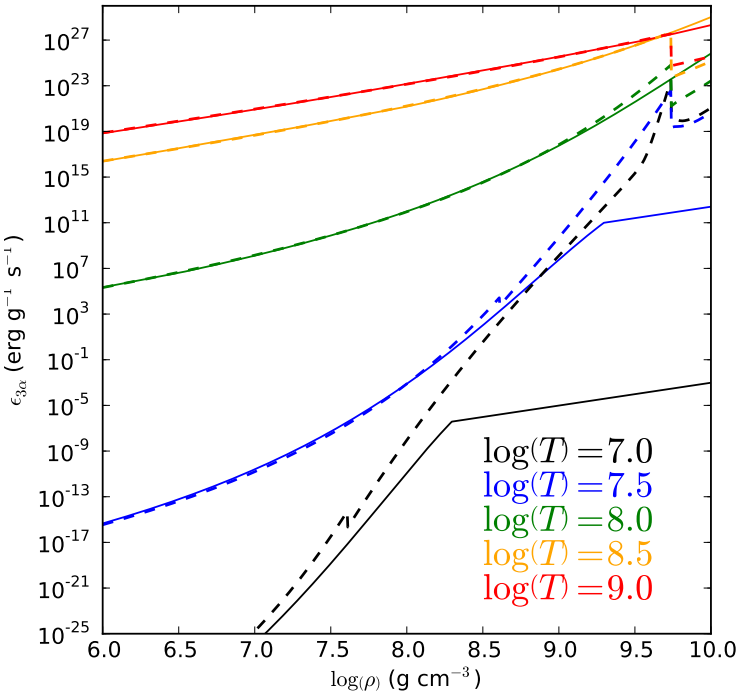

Figure 4.3: A comparison between the instantaneous energy generation rate for the `triple_alpha` and `fushiki_lamb` reaction networks used in `MAESTRO`. The solid lines are calculations with the `triple_alpha` network while the dashed lines are from the `fushiki_lamb` network; the color of the line indicates the temperature.

the dashed lines. The color of the line is related to the temperature at which the burning took place. Note that the `triple_alpha` network includes the reverse reaction $^{12}\mathrm{C} + \gamma \longrightarrow 3^{4}\mathrm{He}$, while the `fushiki_lamb` network does not because this reaction was not calculated in the Fushiki and Lamb paper. For the XRB problem, the density range is towards the far left $\left(\log[\rho] \lesssim 6.8 \text{ g cm}^{-3}\right)$ of the plot, and the temperature range is $\log[T] \gtrsim 8.0$ K. In this regime, the screening is weak and therefore the differences in screening methods between the two networks should not be significant, as is shown in the Figure. Nevertheless, there are several orders of magnitude difference in energy generation rate between the two networks at higher densities and/or lower temperatures, which is consistent with what other researchers have found (F. X. Timmes, 2007; private commu-

nication). With no obvious explanation for the differences, and because the rate of Fushiki and Lamb is not widely accepted in catalogues, such as the JINA Reaclib database, we choose to use the `triple_alpha` network over that of the `fushiki_lamb` for the calculations presented in Chapter 5.

4.4.3 `hotcno` and `approx8` networks

For the case of mixed H/He bursts, we need a more involved network, as described in Section 4.3. To this end I have modified Frank Timmes's public hot CNO and rp breakout network[2] to work with `MAESTRO` and the `VODE` solver. This `hotcno` reaction network consists of 21 species and 52 reactions, including all of those from Figure 4.1, plus the β-limited reactions of the normal CNO cycle, a couple of α-chain reactions, and some approximations to burn up through ^{56}Ni. Even in two dimensions, carrying 21 species in a network makes the solution of 4.9 very time consuming. Therefore, for the calculations of Chapter 6 we use Frank Timmes's publicly available aprox8 network,[3] which approximates hydrogen, helium, and carbon burning with only 8 species. This network was again modified and renamed the `approx8` network in `MAESTRO`. Unfortunately, this network does not hot CNO cycle or rp process burning, but it is much more tractable in a multidimensional simulation.

[2]Source code available here: `http://cococubed.asu.edu/codes/burn7/hotcno.tbz`
[3]Source code available here: `http://cococubed.asu.edu/codes/burn7/aprox8.tbz`

CHAPTER 5

TWO-DIMENSIONAL SIMULATION OF A PURE ^{4}HE BURST SOURCE

This chapter is based on the Malone et al. (2011) paper describing convection driven by helium burning in a pure ^{4}He layer accreted on top of a ^{56}Fe neutron star with a trace abundance (10^{-10}) of ^{12}C. A pure ^{4}He accretor was chosen both because the corresponding nuclear reaction network, `triple_alpha`, is simple compared to the slow, β-decay-limited burning processes in bursts involving H (see Chapter 4), and because ultra-compact XRB sources are possible pure ^{4}He accretors (`4U 1820-30`, for example; Cumming 2003).

5.1 Initial Models

To begin a simulation, a one-dimensional initial model of the accreted layer on the surface of a neutron star is needed. This model should be in both hydrostatic and thermal equilibrium. As mentioned at the beginning of Chapter 3, a plane-parallel geometry is assumed — that is, the gravitational acceleration, g, is assumed constant throughout the domain.

There are several approaches to one-dimensional model generation in the literature. In our approach, we begin with a semi-analytic initial model and then augment the model to account

for convective cooling. We also discuss proper mapping of the one-dimensional model to our multidimensional framework.

5.1.1 Semi-analytic Models

The semi-analytic approach to model generation involves integration of the heat equation and an entropy equation,

$$\frac{dT}{dy} = \frac{3\kappa F}{4acT^3} \tag{5.1}$$

$$\frac{dF}{dy} = 0, \tag{5.2}$$

where c is the speed of light, a the radiation constant, κ the opacity (including radiative and conductive contributions), T the temperature, F the outward heat flux and $dy = -\rho dr$ with $y(r)$ the column-depth (see Cumming & Bildsten 2000 for details of this method). Note that (5.1) can give a thermal profile that is superadiabatic, which would be convectively unstable — in practice, the thermal gradient is restricted to be $dT/dy \leq (dT/dy)_s$ where the subscript s means along an adiabat. Also note that for simplicity, equation (5.2) neglects any compressional heating contributions from the accretion itself and assumes the accreted material is not burning during the accretion phase — this is a steady-state configuration. There is, however, an outward heat flux from pycnonuclear reactions deep within the neutron star crust; we approximate this flux as a constant value throughout the accreted layer, $F = 200$ keV per nucleon. The integration starts at the top of the ^{4}He atmosphere (arbitrarily at $y_{\rm top} = 10^3$ g cm^{-2}) where a radiative zero solution is assumed, and continues until the thin shell instability condition (Fushiki & Lamb, 1987a),

$$\frac{d\varepsilon_{3\alpha}}{dT} > \frac{d\varepsilon_{\rm cool}}{dT}, \tag{5.3}$$

is reached at $y = y_{\rm base}$. The local cooling rate is typically approximated from (5.1) and (5.2) as

$$\varepsilon_{\rm cool} \approx \frac{acT^4}{3\kappa y^2}. \tag{5.4}$$

When (5.3) is attained, the composition for $y > y_{\rm base}$ is switched to ^{56}Fe and integration of (5.1) and (5.2) resumes until a thick enough substrate is formed such that $y_{\rm base}$ is sufficiently far from the bottom of the computational domain, $y(r = 0) = 10^{12}$ g cm^{-2} in our studies.

The approximation, (5.4), works well in one dimension because the only efficient way the system can cool (neglecting weak reactions) is via conduction and radiation, which enter through the opacity. When more spatial dimensions are added to the system and there is heating from below from nuclear reactions, the fluid is free to overturn and cool via convection. Now we have a situation where the local multidimensional cooling rate, $\varepsilon_{\rm cool,\ multi\text{-}d} = \varepsilon_{\rm cool} + \varepsilon_{\rm conv}$, exceeds the initial approximation and (5.3) may no longer be satisfied. Therefore, such a semi-analytic model is no longer close to runaway and to evolve the system in multiple dimensions until (5.3) is reached is intractable even with the advantages of a low Mach number approximation code.

5.1.2 `Kepler`-supplemented Models

One way to overcome the difficulties with evolving the model described in the previous section in multiple dimensions is to explicitly include an *effective* convective cooling term in the approximation to the local cooling given by equation (5.4). This *effective* convective cooling can be included via mixing-length theory typically found in stellar evolution codes. Using the semi-analytic model described above as initial conditions, the one-dimensional stellar evolution code, `Kepler` (Weaver et al., 1978), was used to construct the remainder of the underlying neutron star with $R_{\rm ns} = 10$ km and $M_{\rm ns} = 1.87 M_\odot$ (Woosley, 2010; private communication). The system is then allowed to evolve in one dimension whereupon nuclear burning heats the base of the layer, and the

convection prescription develops a well-mixed and nearly adiabatic region of ^{12}C ash overlying the ^{4}He base. This results in a model that is much closer to satisfying the thermal instability criterion, (5.3), when mapped into multiple dimensions.

It should be noted here that the opacities used in `MAESTRO` (from Timmes, 2000) are not the same as those used in `Kepler`. However, as the system evolves towards ignition, the dominant contribution to the local cooling comes from convection, and the role of thermal conduction is decreased. Even during the early-time simulations considered in this paper, the role of conduction is unclear — see Section 5.2.2. Furthermore, it is likely that the different methods give opacities that agree to within a factor of ~ 2 (see discussion in Heger et al., 2007b).

5.1.3 Mapping to Multiple Dimensions

The data from `Kepler` are given in a Lagrangian (mass) coordinate system and we need to convert them to an Eulerian (physical) coordinate system for use in `MAESTRO`. We use a procedure similar to that found in Zingale et al. (2002) to ensure our initial model is in HSE. Given the density, temperature, and composition from the `Kepler` evolution, we call the equation of state to get the pressure. We then discretize the HSE equation and solve for the non-uniform Eulerian grid spacing corresponding to the Lagrangian grid points,

$$r_i = r_{i-1} - \frac{1}{g}\frac{p_i - p_{i-1}}{\frac{1}{2}(\rho_i + \rho_{i-1})}, \tag{5.5}$$

where r is the radial coordinate, p the pressure, and ρ the density. We set $r_1 = 0$ to complete the description of the grid. The transition from the pure ^{56}Fe neutron star (at r_{trans}) to the ^{4}He atmosphere (at $r_{\text{trans}+1}$) is a step function as a result of the initial Lagrangian data. Such sharp transitions can be a source of numerical noise and oscillations as the solution evolves on an Eulerian grid. To minimize the numerical noise, we smooth the interface by adding n uniformly distributed coordi-

nate points between $r_{\rm trans}$ and $r_{\rm trans+1}$. The temperature at these new points is linearly interpolated between $T_{\rm trans}$ and $T_{\rm trans+1}$. Then $X(^4\mathrm{He})$ and $X(^{12}\mathrm{C})$ at the new points are filled with a tanh profile:

$$\phi_i = \alpha \tanh\left(\frac{r_i - r_{\rm c}}{\varphi}\right) + \phi_{\rm c} \tag{5.6}$$

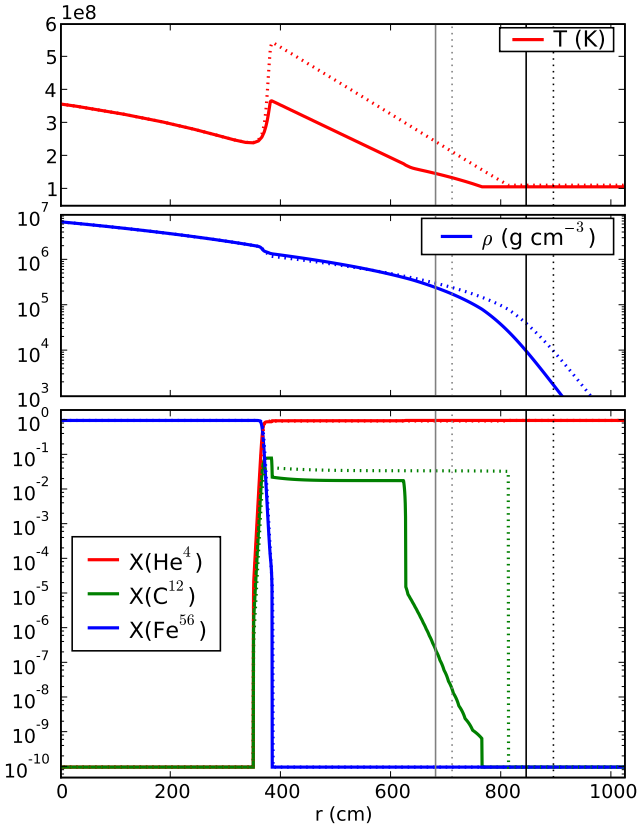

Figure 5.1: `Kepler`-supplemented `cold` (solid lines) and `hot` (dashed lines) models as described in the text. Energy release from nuclear burning at the base of the ^{4}He layer has caused the temperature to rise. The `cold` model is evolved to a peak $T_{\rm base} = 3.67 \times 10^8$ K and the `hot` model is evolved to a peak $T_{\rm base} = 5.39 \times 10^8$ K. The black vertical lines indicate the location of the anelastic cutoff while the grey vertical lines indicate the location of the beginning of our sponge forcing term for each of the models (see Section 3.2.1).

where $\alpha = (\phi_{\rm trans+1} - \phi_{\rm trans})/2$, $r_{\rm c} = (r_{\rm trans} + r_{\rm trans+1})/2$, $\phi_{\rm c} = (\phi_{\rm trans} + \phi_{\rm trans+1})/2$ and φ is a parameter to set the smoothness. $X(^{56}\mathrm{Fe})$ is then found from the constraint $\sum_k X_k = 1$, and p and ρ are found by using an iterative Newton-Raphson technique with the equation of state and (5.5) at these new points. This smoothed model is then linearly interpolated onto a completely uniform grid, with $r_i = r_{i-1} + \Delta r$, and is again put into HSE using (5.5) and the equation of state. Values of $n = 50$ and $\varphi = 3$ were used to smooth the models presented in this work.

Figure 5.1 shows the result of this procedure for two models that were evolved in `Kepler` until the base of the ^{4}He atmosphere had reached a temperature of 3.67×10^8 K (solid line, hereafter referred to as the `cold` model) and 5.39×10^8 K (dotted line, hereafter referred to as the `hot` model). The density at the base of the ^{4}He layer for the `cold` model is 1.4×10^6 g cm^{-3} and

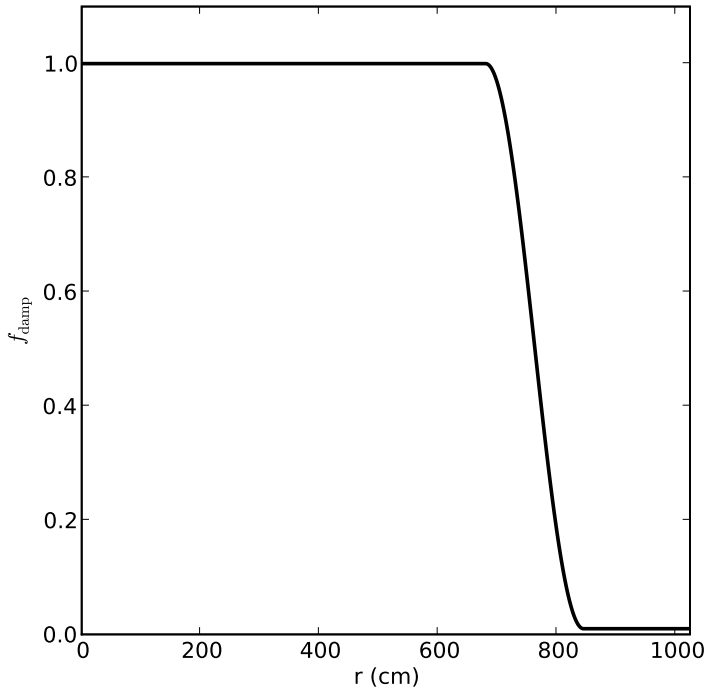

Figure 5.2: Initial sponge profile for the `cold` model where r_{sp} = 680 cm and r_{tp} = 844 cm.

is 1.2×10^6 g cm^{-3} for the `hot` model. For comparison, the initial model of Lin et al. (2006) had a base temperature and density of 2×10^8 K and 4×10^6 g cm^{-3}, respectively. The black and grey vertical lines correspond to the edge of the sponging term described in Section 3.2.1; the sponge profile for the `cold` model is shown in Figure 5.2. The `cold` model has a peak in ^{12}C production around $r = 382$ cm (i.e., the base of the ^{4}He layer in both models) that appears smoothed in the more evolved `hot` model. Both models, however, have an extended region of well-mixed ^{12}C that extends to $r = 624$ cm ($r = 812$ cm) for the `cold` (`hot`) model. These initial models contain no multidimensional velocity information from the `Kepler` simulations. We therefore make no assumptions about the nature of the convection when the models are mapped into multiple dimensions in `MAESTRO`.

5.2 Results of Pure ^{4}He–burning Simulations

We describe below the results of mapping the `Kepler`-supplemented models into `MAESTRO` in two dimensions, (x, r), and the system's subsequent evolution. Section 5.2.1 describes the resolution requirements needed to properly resolve the burning layer. In Section 5.2.2 we show how the inclusion of thermal diffusion affects the nuclear burning layer and its location. We show in Section 5.2.3 how utilizing a time-dependent base state allows us to capture the expansion of the atmosphere due to heating. Section 5.2.4 shows how including a volume discrepancy correction keeps the base state thermally consistent with the equation of state. Finally we discuss the extent and evolution of the convective region in Section 5.2.5.

To map the one-dimensional model into `MAESTRO`, we copy laterally across the domain such that $\phi(x,r,t=0) = \phi_{\text{one-d}}(r)$ for each variable ϕ in the model. In the following analysis, the subscript "max" refers to the maximum value of a quantity in the computational domain at a given time step. In two dimensions, we define the average as a function of radius,$\langle\phi\rangle = \langle\phi\rangle(r,t)$, of a

quantity ϕ by

$$\langle \phi \rangle_j^m = \frac{1}{N} \sum_{i=1}^{N} \phi_{i,j}^m \tag{5.7}$$

where $\phi_{i,j}^m = \phi(x_i, r_j, t^m)$ and N is the total number of grid zones in the lateral, x, direction at height r_j at time t^m.

The boundary conditions for all simulations are periodic in the x-direction to mimic a laterally extended convection region. The upper r boundary is outflow to allow for free expansion of the atmosphere. The lower r boundary of dense neutron star material is set to a wall with no normal flow. To solve the thermal diffusion contribution at the upper and lower boundaries, we impose the Neumann condition $dh/d\mathbf{n} = 0$, where $\mathbf{n}$ is the outward facing normal vector; the enthalpy boundary conditions are periodic in the lateral directions. We note that the upper and lower domain boundaries are sufficiently far from the burning layer so that they do not affect the dynamics of the convection. An advective CFL number of 0.7 was used in all of our simulations.

As previously mentioned, we do not obtain any multidimensional velocity information from the `Kepler` models; our system is initially static. For convection to begin, the symmetry of the system must, therefore, be broken. This can be accomplished either by placing a small perturbation at the base of the ^{4}He layer or by allowing numerical noise from the multigrid solver to seed the convective cells. For the simulations presented here, neither approach is advantageous over the other, both giving quantitatively similar steady-state convective flow fields; we utilize both approaches in our studies and when perturbing, we place a small ($\Delta T/T = 10^{-5}$) Gaussian temperature perturbation laterally centered at height $r = 384$ cm to break the initial symmetry of the problem.

5.2.1 Resolution Requirements

To date, the only other paper in the literature regarding multidimensional simulations of XRBs as deflagrations (Lin et al., 2006) used a finest resolution of 5 cm zone^{-1}. They presented multidimensional results at 5, 7.5, and 10 cm zone^{-1} resolutions and remarked that there is a "tendency toward convergence with increasing resolution" based on the time to reach the peak energy generation rate. It is important to note that our initial models are different from those of Lin et al. (2006). In particular, their models only considered two species — the accreted layer was pure ^{4}He and the underlying neutron star was composed entirely of ^{12}C. This caused their models to have a smaller jump in mean molecular weight across the neutron star/accreted layer boundary compared to our models. Furthermore, the initial conditions for their multidimensional studies were from the results of a one-dimensional diffusional-thermal code that evolved the system through several bursts. These differences from our method of initial model generation give the Lin et al. (2006) models an extended ($\sim$ 100 cm) thermal peak compared to our narrow ($\sim$ 10 cm) peak (compare our Figure 5.1 to their Figure 2).

The burning layer at the base of the accreted material is very thin; high resolution is required to properly model this region. The peak of the thermal profile for the `hot` model is broader than the corresponding peak in the `cold` model. Consequently, the burning layer in the `hot` model is thicker than that of the `cold` model — we therefore focus our study of resolution requirements on the more restrictive of the two, the `cold` model. The top panel of Figure 5.3 shows the $\langle H_{\rm nuc} \rangle$ profile at $t = 1$ ms for simulations of the `cold` model using the same resolutions as in the Lin et al. (2006) study. Even at this early time there is a 25% spread in the peak value of $\langle H_{\rm nuc} \rangle$ for these resolutions. The bottom panel shows the same profile but at several higher resolutions. The peak value of $\langle H_{\rm nuc} \rangle$ for the 4 cm zone^{-1} simulation is comparable to the peak values in the top panel, but we only see numerical convergence of the peak value as we go to higher spatial resolution. In

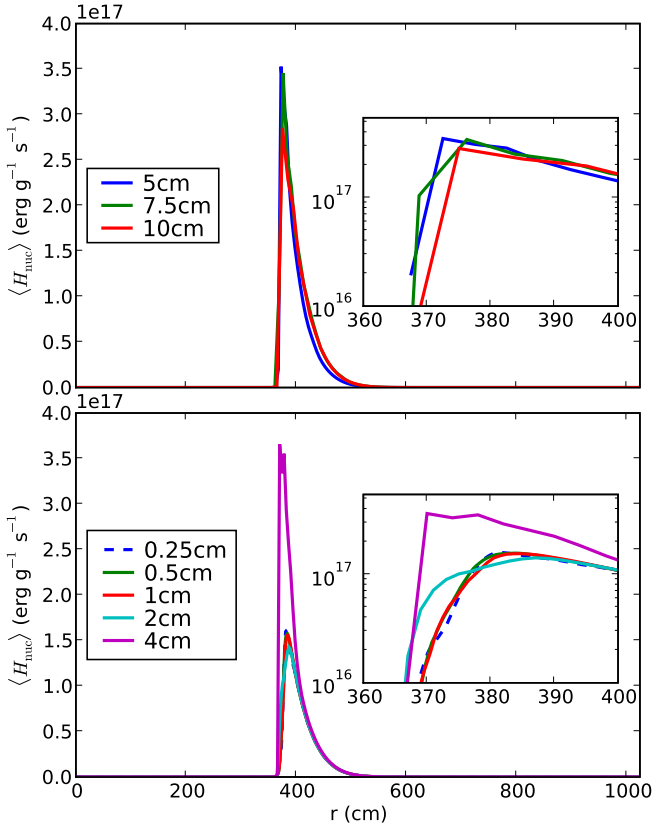

Figure 5.3: Average of $H_{\rm nuc}$ as a function of height for the `cold` model at various resolutions at $t = 1$ ms. Note that the vertical axes of the inset plots are in a logarithmic scale. For clarity, the top panel shows simulations which use the same resolutions as in the Lin et al. (2006) study and the bottom panel shows more resolved simulations. The peak of the profile at 0.5 cm zone^{-1} resolution is qualitatively similar to the peak of the profile at 0.25 cm zone^{-1} resolution.

addition, the shape of the profile near peak converges with increasing resolution; the 0.25 and 0.5 cm zone^{-1} resolution simulations look qualitatively similar. We therefore claim that the burning layer is not properly resolved in our models unless a resolution of 0.5 cm zone^{-1}, or finer, is used. It is important to note that even though our initial models differ, this resolution requirement is an order of magnitude higher than what has been previously presented in the literature and therefore

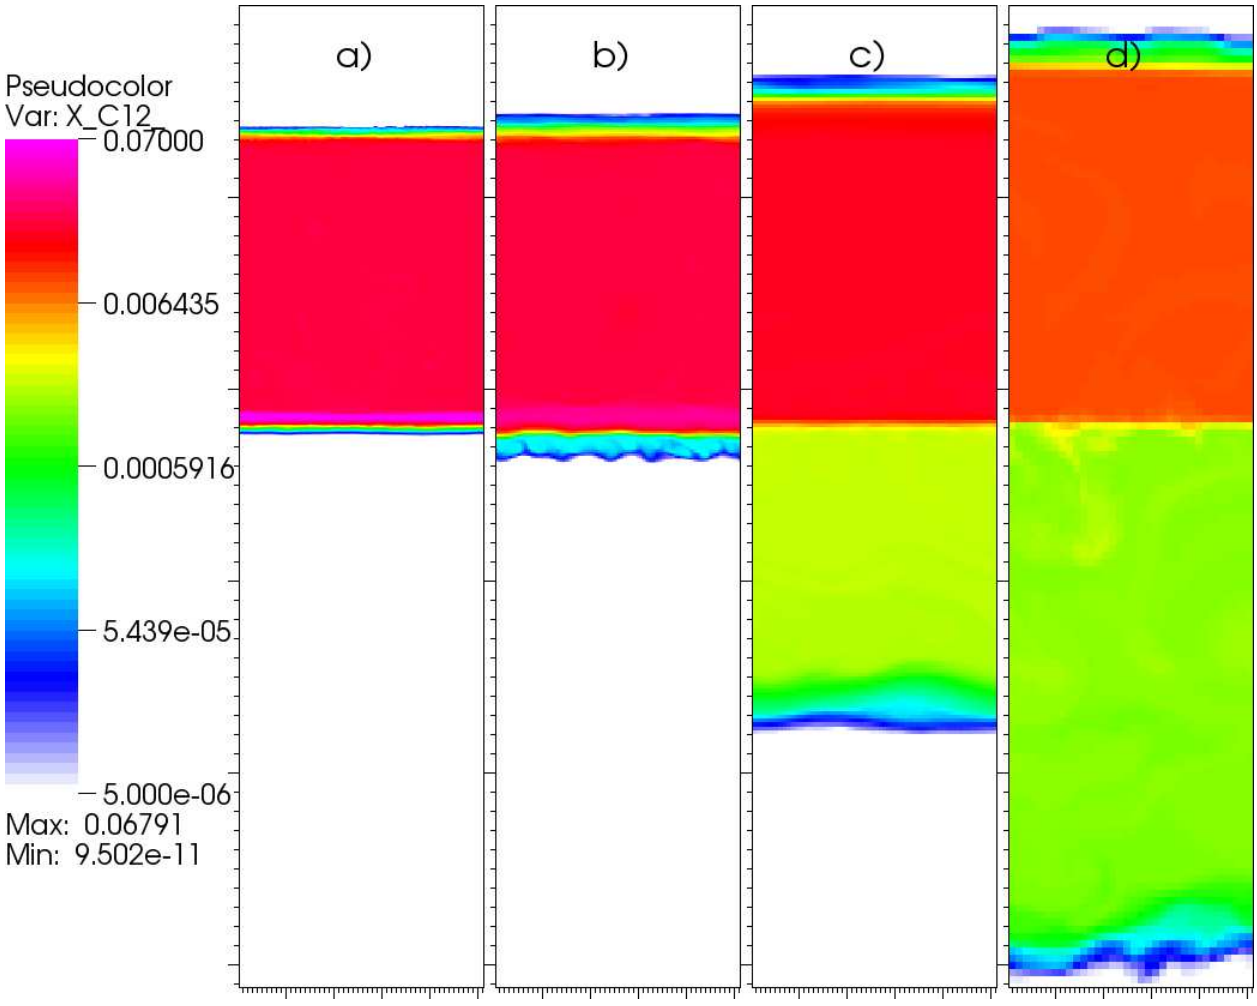

Figure 5.4: Effects of under-resolving convection for the `cold` model. Plotted is the ^{12}C mass fraction after 10 ms of evolution for various resolutions: `a`) 0.5, `b`) 2, `c`) 4, and `d`) 7.5 cm zone^{-1}. Each figure shows the same region of physical space and has dimensions 256 cm $\times$ 1024 cm. The coarse resolution simulations show an extended convective zone and a significant amount of convective undershoot.

significantly increases the computational cost of our XRB simulations.

Under-resolving the burning layer artificially boosts the energy generation rate, which in turn over-drives convection. Figure 5.4 shows a close-up of the ^{12}C mass fraction after 10 ms of evolution of the `cold` model at 0.5 (`a`), 2 (`b`), 4 (`c`), and 7.5 cm zone^{-1} (`d`) resolutions. The base of the burning layer is located in the bottom-most green region (just below the magenta) in panel

a. All four simulations give a well-mixed carbon region above the burning layer; the extent of the convective zone increases with decreasing resolution with the 7.5 cm zone^{-1} simulation's convective zone extending 30% further than the 0.5 cm zone^{-1} simulation's convective zone. The amount of convective undershoot — the tendency of material to penetrate below the burning layer — is much more sensitive to resolution. The 0.5 cm zone^{-1} simulation shows very little evidence of undershooting while the 7.5 cm zone^{-1} simulation has an undershoot region that is larger in physical extent than its corresponding convective region above the burning layer. For all of the studies described below, we use a resolution of 0.5 cm zone^{-1} in the burning layer.

5.2.2 Effects of Thermal Diffusion on the Burning Layer

As explained in Section 1.2, the burning front during an XRB likely propagates as a subsonic flame, the speed of which is regulated by the rate of thermal diffusion across the front. At the resolution required to resolve the thin burning layer (see previous section) it is currently intractable to evolve the system from our initial conditions to flame ignition. We can, however, investigate the effects of thermal diffusion on the stable burning that occurs in the burning layer. Here we focus on the `hot` model instead of the `cold` model because it has the larger thermal gradient — and hence diffusive heat flux — at the base of the accreted layer. For this simulation, we use the new adaptive mesh refinement (AMR) capability in `MAESTRO` (Nonaka et al., 2010), using two levels of refinement and ensuring that the entire convective region is at the finest level of refinement with resolution 0.5 cm zone^{-1}. Figure 5.5 shows these effects in $(H_{\rm nuc})_{\rm max}$ (solid lines) and the location of this maximum (dashed lines) as a function of time at early times both with (green) and without (blue) thermal diffusion. We note that the location of $(H_{\rm nuc})_{\rm max}$ is always at the finest level of refinement. The $(H_{\rm nuc})_{\rm max}$ evolution is similar for both cases with the magnitude in general being slightly larger for the case of no diffusion. The initial spike in $(H_{\rm nuc})_{\rm max}$ at $t \approx 0.25$ ms is due to the

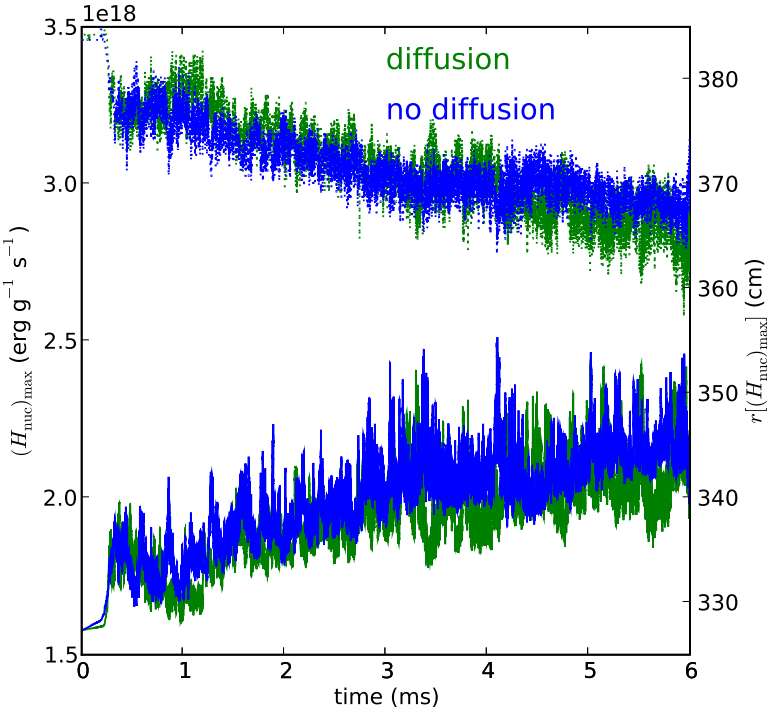

Figure 5.5: Evolution of $(H_{\rm nuc})_{\rm max}$ (solid lines) and its vertical location (dashed lines) as a function of time for the `hot` model both with (green) and without (blue) thermal diffusion.

fact that, initially, there is no established fluid flow that can advect away the energy released from nuclear reactions (see the discussion in Section 5.2.5). Over the next 3 ms, the location of $(H_{\rm nuc})_{\rm max}$ for both simulations moves radially inward at a rate of $\sim 2.9 \times 10^3$ cm s^{-1}. Around $t = 3.25$ ms, the inward radial progression of the location of $(H_{\rm nuc})_{\rm max}$ for the simulation with no diffusion significantly slows to ~ 900 cm s^{-1}. For the remainder of the simulation, the case with thermal diffusion shows no such slowdown — heat transported radially inward via diffusion expands the lower boundary of the convective zone, which mixes fresh fuel to slightly deeper layers. By the end of the simulations, the case that included diffusion had an $(H_{\rm nuc})_{\rm max}$ that occurred ~ 4 cm deeper within the atmosphere than in the case without diffusion. It should be noted that the typical standard deviation in the location of $(H_{\rm nuc})_{\rm max}$ for the case without diffusion is of order 2 cm; this

suggests that perhaps thermal diffusion plays a role in regulating the location of maximum nuclear burning, but further evolution is needed to make statistically significant claims.

5.2.3 Expansion of Base State due to Heating

Having a dynamical base state allows us to capture the large-scale expansion of the atmosphere due to heating from nuclear reactions. This differs from the work by Lin et al. (2006), which had a time-independent base state and did not model the top of the accreted atmosphere due to numerical complications with their algorithm. Figure 5.6 shows the ratio of the base state density to that of the initial ($t = 0$) base state density profile near the surface of the atmosphere for the `cold` model. We define the surface to be where $\rho_0 = \rho_{\text{cutoff}}$. The vertical dashed lines represent the location of the surface for each time-value. After 26.6 ms of evolution, the base state has responded to heating from nuclear reactions approximately 4.5 m below the surface by expanding 3.5 cm. The lower Mach number flow in the `cold` model compared to the `hot` model allows for longer-term evolution of the system and therefore larger expansion of the atmosphere.

The extent of the expansion is rather small at these early times. However, as the system progresses towards outburst the energy generation and, therefore, the rate of expansion increases. As the system expands, the p_0 profile changes, which can affect the dynamics in the convective region. Additionally, as the atmosphere expands, the burning layer becomes less degenerate, which may be important for the nucleosynthesis during the outburst. Furthermore, a proper modeling of this expansion during the peak of a PRE burst model may help pinpoint the location of the photosphere with respect to the stellar radius at touchdown, a quantity that plays an important role in using XRBs to measure the mass and radius of the underlying neutron star (Steiner et al., 2010, for example).

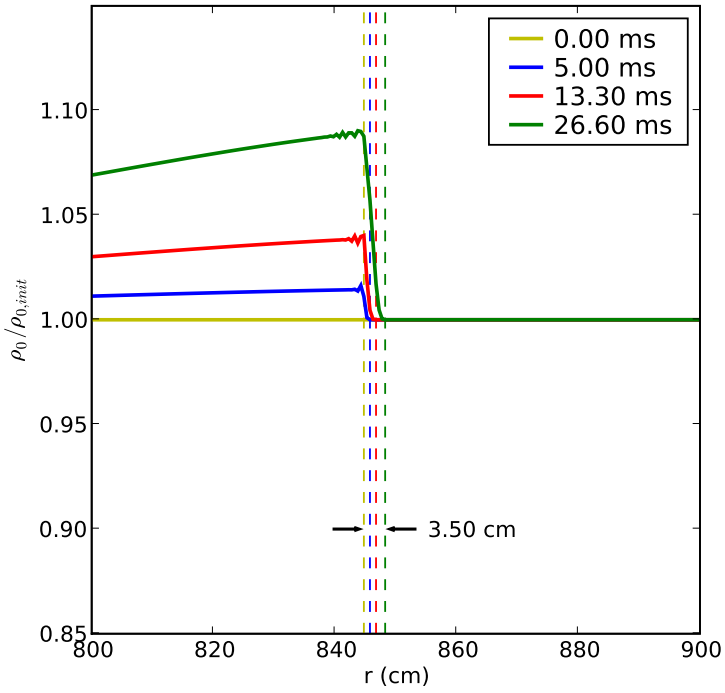

Figure 5.6: Expansion of the base state due to heating. Plotted is the ratio of base state density, ρ_0, to the initial ($t = 0$) base state density, $\rho_{0,\text{init}}$, near the surface of the atmosphere for the `cold` model. We define the surface to be where $\rho = \rho_{\text{cutoff}}$ and it is represented by the vertical dashed lines. The base state has expanded 3.5 cm in 26.6 ms of evolution.

5.2.4 Effects of the Volume Discrepancy Term

In Section 3.1 we explained that the thermodynamic pressure may drift from the base state pressure. To correct for this drift, we introduced the volume discrepancy term in equation (3.13), which drives the thermodynamic pressure towards the base state pressure. We focus our attention here on the `hot` model because it shows a more dramatic drift of the thermodynamic pressure from the base state pressure. Figure 5.7 shows the volume discrepancy term in action by examining the percent difference between the base state and thermodynamic pressures as a function of time for various values of f for the `hot` model. The top panel shows the maximum value whereas the

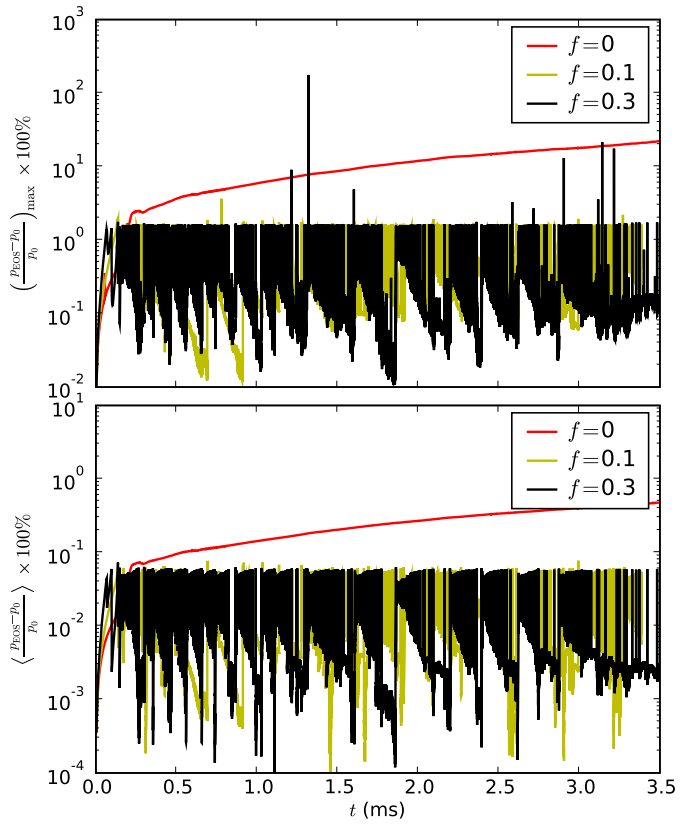

Figure 5.7: Effects of the volume discrepancy factor as characterized by the percent difference between the thermodynamic pressure as given by the equation of state, $p_{\rm EOS}$, and the base state pressure, p_0, for the `hot` model. The top panel shows the maximum value whereas the bottom panel shows the average value of the percent difference in the computational domain. Note the different vertical scales between the two plots

bottom panel shows the average value of this percent difference; both the peak and average values show the same trend for a given value of f. After the initial adjustment of the system, the average drift for the case of no volume discrepancy correction ($f = 0$) increases approximately linearly at $\sim 0.1\%$ per ms of evolution. Including the correction term restricts the temporal- and spatial-averaged value of the drift to $\lesssim 0.02\%$.

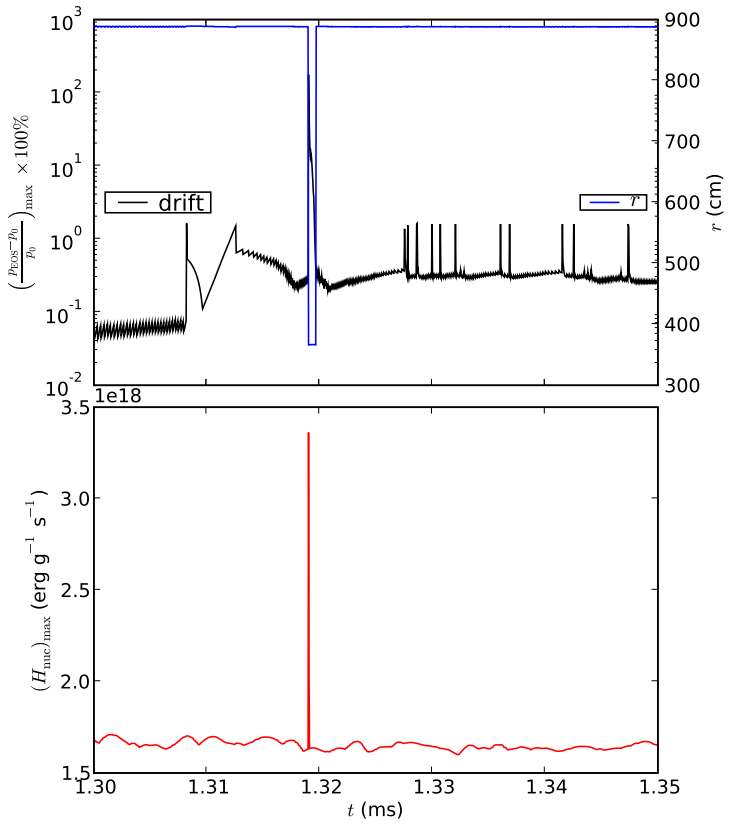

Figure 5.8: Closeup of the $\mathscr{O}(1)$ spike in the maximum value of the $f = 0.3$ drift as seen in the top panel of Figure 5.7. The top panel shows the drift value and its location in the domain; the bottom panel shows the maximum energy generation rate. The large amount of energy released from the burning spike causes the thermodynamic pressure to differ from the hydrostatic base state pressure and therefore a spike in the drift.

For nonzero f, the oscillatory behavior in the drift is due to the fact that the system may slightly over-correct the thermodynamic pressure in a given time step and then recover in the next step. A larger value of f causes a stronger driving of the drift, which tends to be more oscillatory. In addition, a larger value of f appears to be correlated with larger spikes in the drift. The top panel of Figure 5.8 shows a closeup of the $\mathscr{O}(1)$ error seen in the $f = 0.3$ curve in Figure 5.7. The

location of the maximum drift is also plotted in the top panel; the large spike in the drift occurs just below the burning layer at $r = 366.25$ cm. The bottom panel of Figure 5.8 shows the corresponding maximum energy generation rate, which also contains a spike that is coincident with the spike in the drift — a large deposit of energy on a short timescale causes the thermodynamic pressure to get out of sync with the hydrostatic base state pressure. The increase in the energy generation rate is due to a fluid parcel rich in ^{4}He fuel being brought into a region of high temperature via the turbulent convection. The duration of this transient behaviour is very short: nine time steps or $\sim 6.4 \times 10^{-7}$ s. The selection of an appropriate non-zero value for f is a problem-specific endeavor, but the chosen value has little effect on the dynamics of the convective flow field. For the simulations presented below we use a volume discrepancy correction value of $f = 0.3$, which is based on the results of several test runs and past experience with comparing the results to the $f = 0$ case. We will continue to study if and how the chosen value of f affects the long term development of the convective field for this specific problem.

5.2.5 Convective Dynamics

The adiabatic excess, $\Delta\nabla$, — with

$$\Delta\nabla = \nabla - \nabla_s, \tag{5.8}$$

where the actual thermal gradient is

$$\nabla = \frac{\partial \ln T / \partial r}{\partial \ln p / \partial r}$$

and the adiabatic thermal gradient is $\nabla_s = \left(d \ln T / d \ln p\right)_s$ with the subscript s meaning along an adiabat — is used to gauge the evolution of the convective zone for the Schwarzschild instability criterion. Under this criterion, a fluid element is unstable to thermally driven convection when

$\Delta\nabla > 0$ and is stable for $\Delta\nabla < 0$. The first term in (5.8) is calculated using finite differences of the temperature and pressure profiles along the radial direction. The second term in (5.8) depends solely on the thermodynamics of the equation of state. It is related to the second adiabatic exponent, Γ_2 (see Cox & Giuli (1968) Chapter 9):

$$\frac{\Gamma_2 - 1}{\Gamma_2} = \left(\frac{d\ln T}{d\ln p}\right)_s. \tag{5.9}$$

All three of the adiabatic exponents are related:

$$\frac{\Gamma_1}{\Gamma_3 - 1} = \frac{\Gamma_2}{\Gamma_2 - 1}, \tag{5.10}$$

where $\Gamma_3 - 1 = (d\ln T/d\ln\rho)_s$ and Γ_1 was defined in Section 3.1. Writing the equation of state as $p = p(\rho, T)$ and expanding the differential dp, we find the relation

$$\Gamma_3 - 1 = \frac{\Gamma_1 - \rho p_\rho / p}{T p_T / p} \tag{5.11}$$

along an adiabat. Our equation of state only returns Γ_1, but combining this with (5.10) and (5.11) allows us to solve for the adiabatic thermal gradient and hence the adiabatic excess.

Figure 5.9 shows the early evolution of $\Delta\nabla$ for the `cold` model. Each plot covers the spatial range (0 cm $\leq x \leq$ 256 cm, 350 cm $\leq r \leq$ 700 cm) to focus on the convective region. The stripes in the initial conditions, Figure 5.9(a), are due to small interpolation errors from mapping the initial data onto the two-dimensional grid. The initial adjustment of the system seen in Figure 5.9(b) causes a mixing of stable (blue) and unstable (red) fluid elements. This transient adjustment phase occurs for two reasons: 1) the initial conditions were based on a parameterization of convection in one dimension and the system now needs to adjust to a two-dimensional convective zone, and 2) the initial perturbation does not have an established convection zone and the system needs a

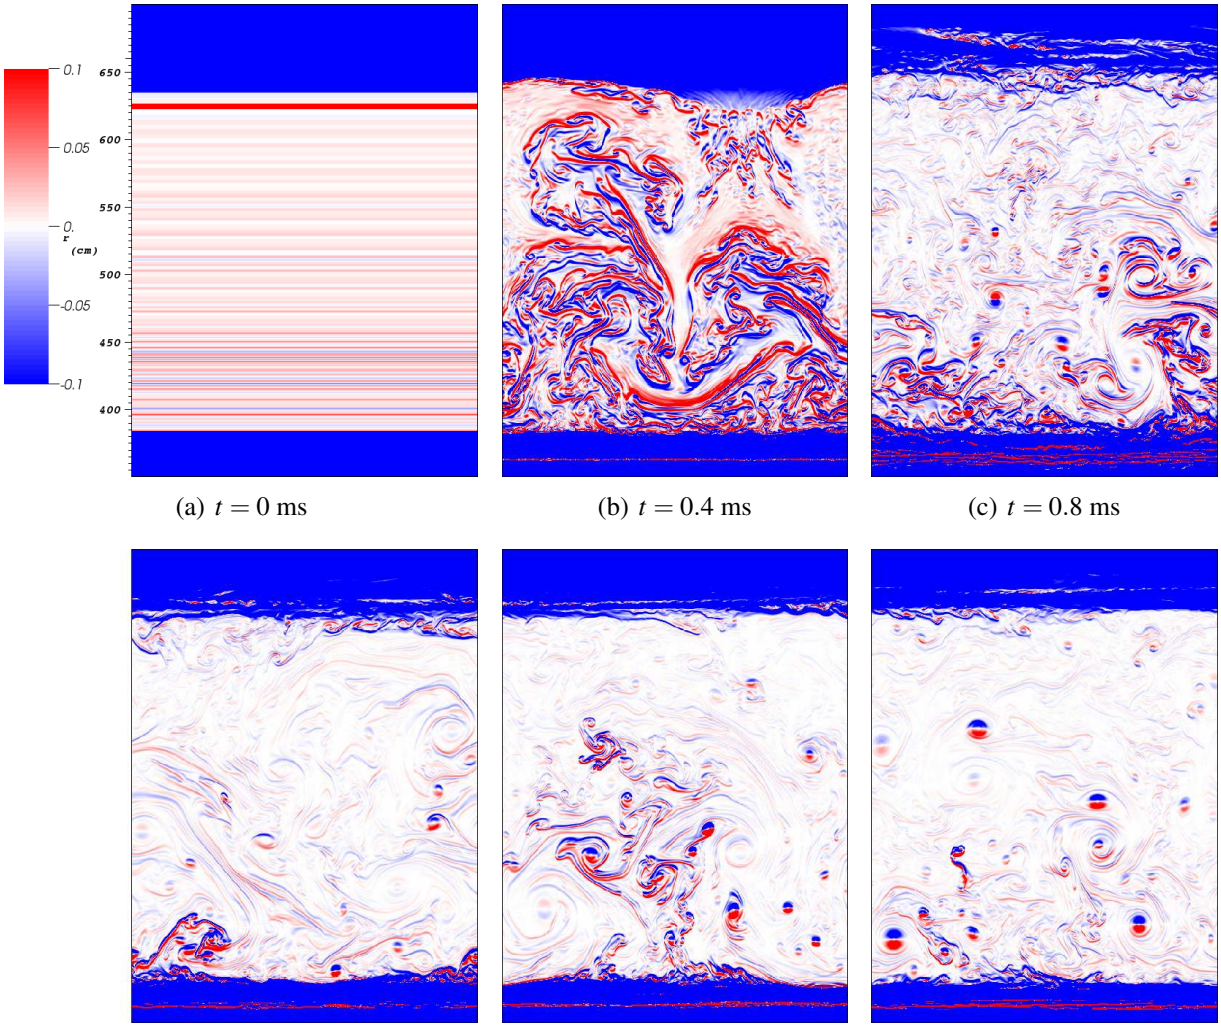

(a) $t = 0$ ms (b) $t = 0.4$ ms (c) $t = 0.8$ ms

(d) $t = 5$ ms (e) $t = 7.5$ ms (f) $t = 10$ ms

Figure 5.9: Colormap plot of the evolution of the adiabatic excess, $\Delta\nabla$, in the convective region for the `cold` model.

short amount of time to build up a flow pattern associated with the perturbation. This results in mixing that produces a region that is marginally convective ($\Delta\nabla \sim 0$; white) with localized pockets of stable and unstable fluid elements as seen in Figure 5.9(c). At later times, these pockets further localize into vortices whose circulation gives rise to roughly circular regions of nonzero adiabatic excess — with one hemisphere that is stable and the other which is unstable — that are advected

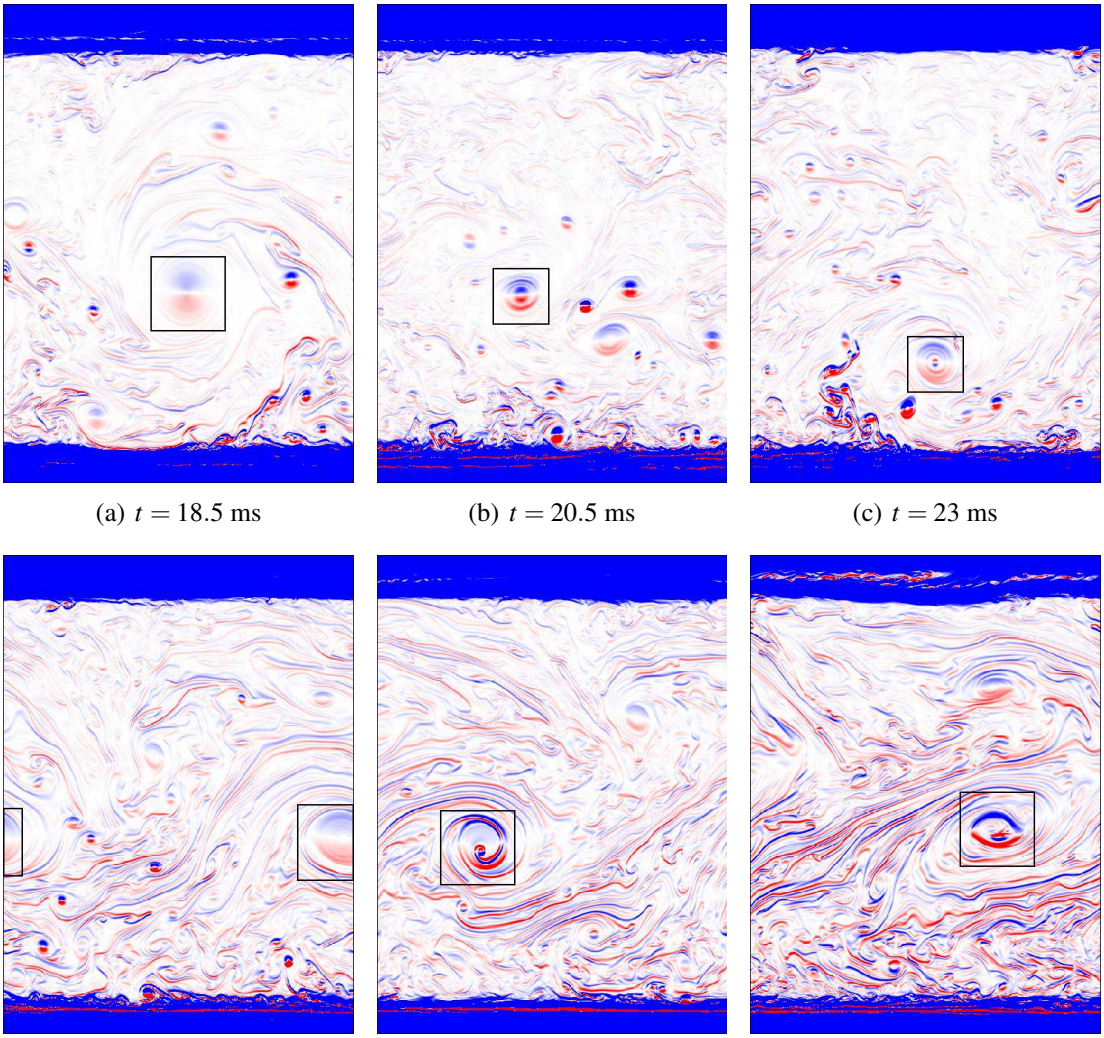

(a) $t = 18.5$ ms (b) $t = 20.5$ ms (c) $t = 23$ ms

(d) $t = 25$ ms (e) $t = 26$ ms (f) $t = 28$ ms

Figure 5.10: Same as Figure 5.9 but at later times. The boxes show the location of a single feature that, once formed, lasts for the remainder of the simulation.

with the flow before dispersing into the ambient medium on subconvective timescales, $\sim 10^{-4}$ s. The vortices are always associated with an adiabatic excess pattern that has an unstable (red) bottom and a stable (blue) top unless two vortices are merging and interacting, in which case the stability distribution becomes skewed.

Figure 5.10 shows $\Delta\nabla$ for the same simulation as in Figure 5.9 but at later times. The boxes

in these plots outline a single long-lived vortex that forms around $t = 18.5$ ms, Figure 5.10(a), and lasts throughout the remainder of the simulation. Formation of this vortex is correlated with the formation of stronger filamentary structures, which are clearest in Figures 5.10(d), 5.10(e) and 5.10(f). These filaments appear to wrap around the solitary vortex and restrict the main formation of smaller vortices to the lower boundary of the convective region.

Another way to quantify the convective region is to look at the ratio ∇/∇_s. From (5.8) we see that the system is unstable to convection under the Schwarzschild criterion when $\nabla/\nabla_s > 1$. The Schwarzschild criterion, however, does not consider the effects of composition gradients that may help stabilize the material against convection; for this we need to consider the Ledoux criterion for instability

$$\nabla - \nabla_L > 0, \tag{5.12}$$

where the Ledoux thermal gradient is (see, for example, Kippenhahn & Weigert (1994))

$$\nabla_L = \nabla_s - \sum_i \frac{\partial \ln X_i/\partial r}{\partial \ln p/\partial r}$$

and the second term above is evaluated via finite differences of the composition and pressure profiles. As with the Schwarzschild criterion, one can look at the ratio ∇/∇_L, which is greater than unity if the material is unstable to Ledoux convection. Figure 5.11(a) shows the above ratios for the average thermal gradients for the initial configuration (left) and after the system has evolved for $t =$ 23.5 ms (right); the black line is for the case of Schwarzschild criterion convection, while the red line is for Ledoux convection. The dashed horizontal line marks the boundary for stability against convection. Where the curves lie above this line, the configuration is unstable; when convection is efficient, the curves should lie very near the horizontal line. For both the initial condition and at late times, the Schwarzschild curve and the Ledoux curve are well matched except near the edges of the convective region where composition gradients cause the two curves to deviate slightly. This

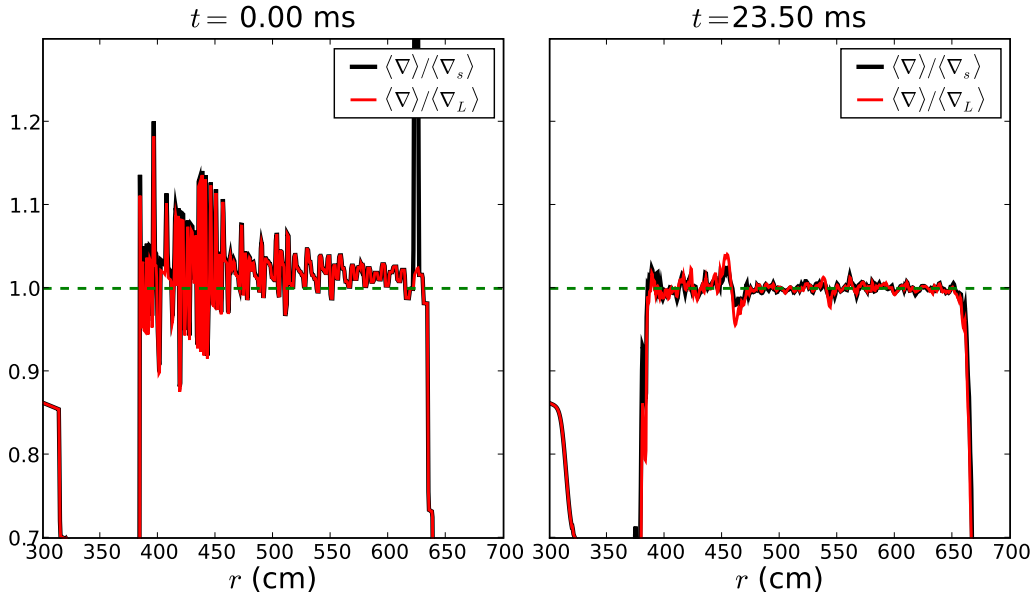

(a) Example Convective Profiles

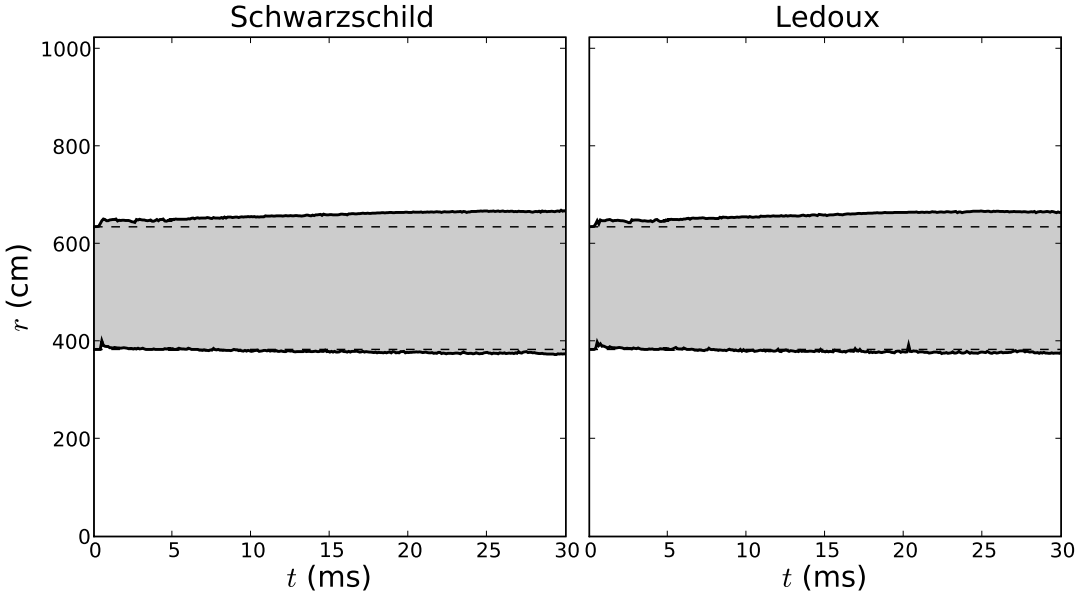

(b) Convective Region Extent

Figure 5.11: Analysis of the extent of the convective region. Panel (a) shows the convective profiles for both the Schwarzschild and Ledoux instability criteria at two different times. Panel (b) shows the extent of the convective region as a function of time as determined by both instability criteria.

deviation is most noticeable in the initial configuration at the upper boundary where there is a sharp jump in composition, which was not smoothed (see Figure 5.1). Of interest in the plot at $t = 23.5$ ms is the feature at $r = 450$ cm, which has an unstable bottom and a stable top; this is consistent with the vortices in Figures 5.9 and 5.10, which had red bottoms and blue tops. We define the edge of the convective region to be where $\nabla/\nabla_s, \nabla/\nabla_L = 0.75$. This particular value of 0.75 was chosen to be sufficiently small enough to rule out false positives from strong pockets of stability from, for example, vortices within the convective region, but also large enough to rule out any fluctuations at the boundaries due to overshoot. Figure 5.11(b) shows in grey the extent of the convective region as a function of time with respect to the full domain for both the Schwarzschild (left) and Ledoux (right) instability criteria. The horizontal dashed lines mark the initial location of the lower and upper boundaries. The overall expansion of the upper boundary for the Schwarzschild (Ledoux) criterion is 32.0 (29.5) cm in 30 ms of evolution; the lower boundary expands downward by 9.5 (6.0) cm in the same time. At late times, the upper boundary of the convective region has a much smoother composition transition than the lower boundary, therefore, the Schwarzschild and Ledoux criteria are much better matched at the upper boundary than the lower. Nevertheless, both the Ledoux and Schwarzschild criteria yield similar results when used to determine the extent of the convective region. In terms of column-depth, the convective zone after 30 ms of evolution spans the region 2.2×10^7 g cm$^{-2} \lesssim y \lesssim 2.6 \times 10^8$ g cm^{-2} for both instability criteria.

For comparison, Figures 5.12 and 5.13 show the ^{12}C mass fraction with velocity vectors for the same simulation and at the same times as in Figures 5.9 and 5.10, respectively. These figures clearly show the association of vortices with the circular regions of nonzero adiabatic excess seen in Figures 5.9 and 5.10. The initial adjustment of the system causes mixing that smooths the slight over-abundance of ^{12}C at the base of the accreted layer present in the initial model (see Figure 5.1). At late times, the convective region is very well-mixed, and the ^{12}C mass fraction is nearly laterally homogeneous. Furthermore, the circulation pattern associated with the long-lived vortex

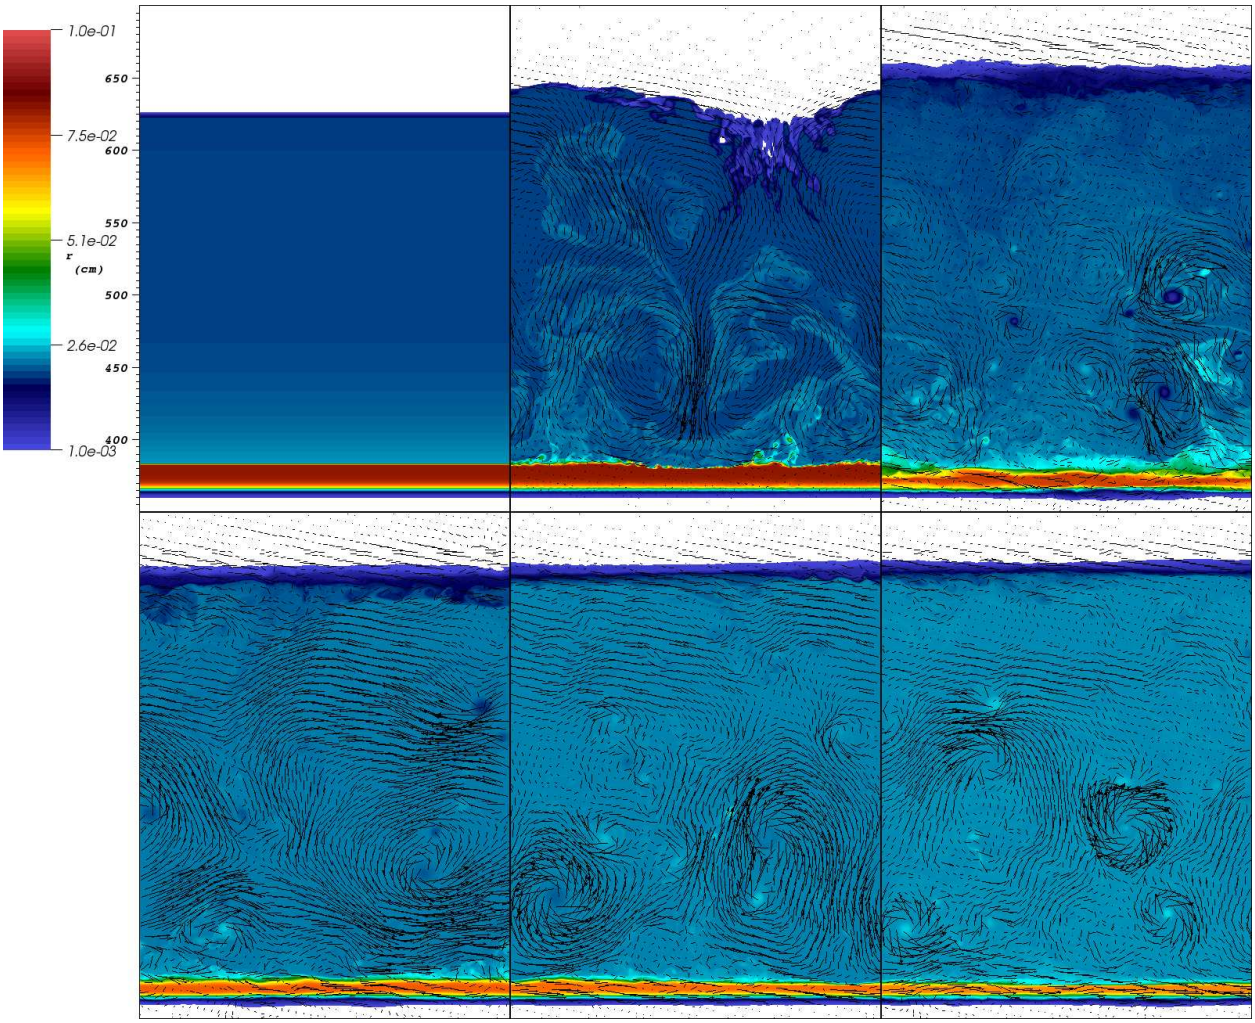

Figure 5.12: Colormap plot of ^{12}C mass fraction with velocity vectors for the same region and times as shown in Figure 5.9.

outlined in Figure 5.10 has grown to a large fraction of the convective zone and is self-interacting because of the periodic boundary conditions. The tendency of the system to form a single dominant vortex from smaller vortices is a feature of two-dimensional simulations. In three dimensions, the turbulent energy cascade moves from large to small scales; large vortical structures break down into smaller structures that are eventually dissipated by viscous effects. In two dimensions, as is the case here, the turbulent energy cascade is reversed — small vortical structures merge together

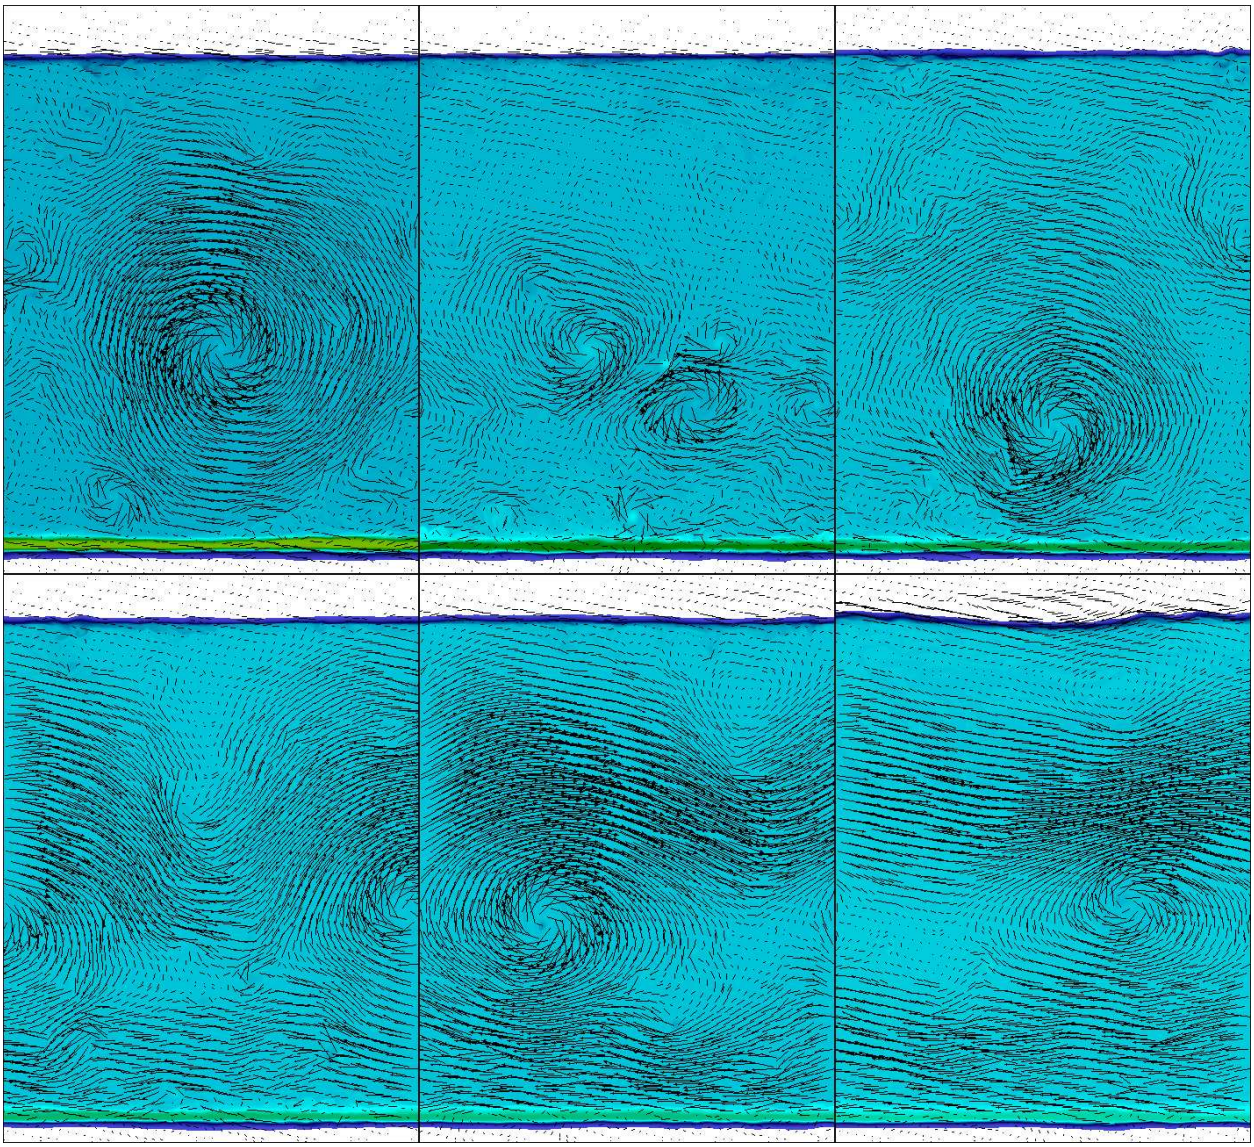

Figure 5.13: Colormap plot of ^{12}C mass fraction with velocity vectors for the same region and times as shown in Figure 5.10.

to form a single dominant vortex.

The circulation is counter-clockwise for the large, long-lived vortex; this causes a region with positive x-velocity below and a region of negative x-velocity above the vortex center. The positive x-velocity region extends all the way to the lower convective boundary where it causes shearing

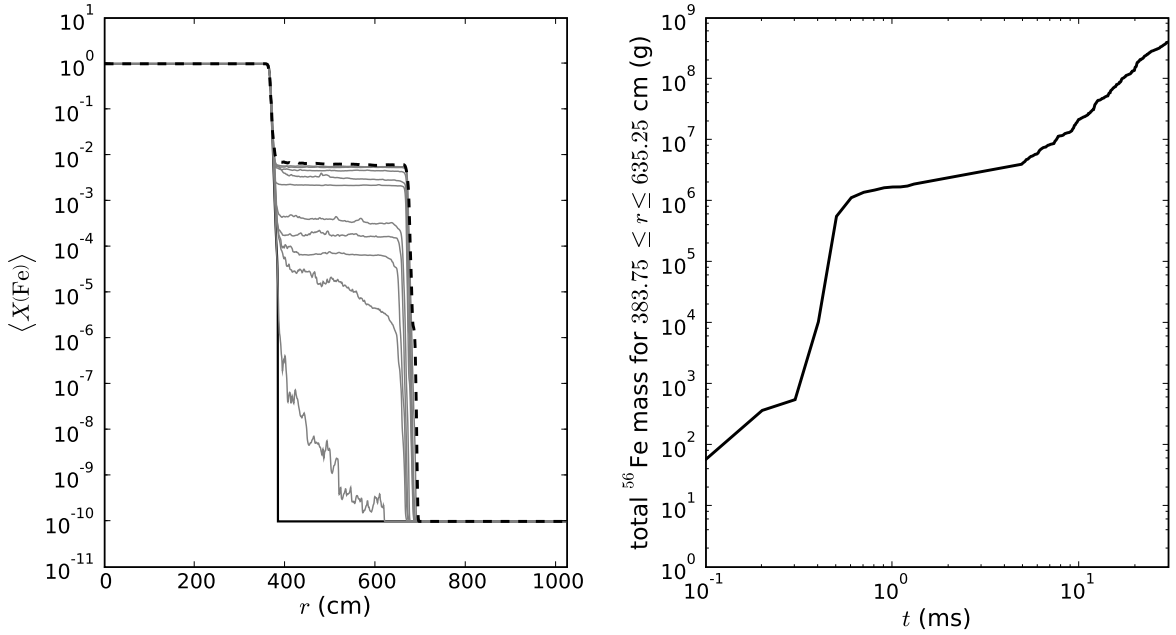

Figure 5.14: Plots showing the ^{56}Fe enrichment of the convective region. The left panel shows the evolution of the average ^{56}Fe mass fraction starting from the initial model distribution (solid thick line) and ending after 30 ms of evolution (dashed line); the thin grey lines show the evolution at the intermediate times shown in Figures 5.9 and 5.10. The right panel shows the total mass of ^{56}Fe in the convective region as a function of time. Note the log scale of the horizontal axis in the right plot.

of the ^{4}He /^{12}C-rich region with the underlying ^{56}Fe region. Consequently, Figure 5.14 shows that some of the underlying ^{56}Fe neutron star material is churned up into the convective region where it is mixed with the rest of the convective material. The left panel shows average ^{56}Fe mass fraction profiles starting with the initial model abundance (thick solid line) through $t = 30$ ms (thick dashed line); the intermediate thin solid lines show profiles at the times used in Figures 5.9 and 5.10. By $t = 5$ ms, the ^{56}Fe is fairly well-mixed in the convective region. The right panel shows the total mass of ^{56}Fe in the region defined by the initial convective zone. The greatest growth in the total mass occurs, as expected, in the initial adjustment ($t \lesssim 0.6$ ms) and then flattens until large enough structures form such that there is sufficient shearing occurring at the base of the

convective boundary. There is only a slight increase in the growth rate for the ^{56}Fe mass around $t = 18.5$ ms where the long-lived vortex first appears. This is due to the fact, mentioned above, that as the system evolves it goes from many small vortices to a few large, dominant vortices. It is only when the circulation pattern of a particular vortex is large enough to strongly interact with the lower convection boundary that we get the shearing and enrichment of the convective region; this occurs around $t \sim 5$ ms. The addition of ^{56}Fe to the convective region has a small but noticeable effect on the conductivity; for example, a displacement of $\sim 1\%$ ^{4}He for ^{56}Fe near the base of the accreted layer, with all other things being constant, gives a $\sim 4\%$ decrease in conductivity. This change in conductivity could play an important role in adjusting the flame speed once ignited.

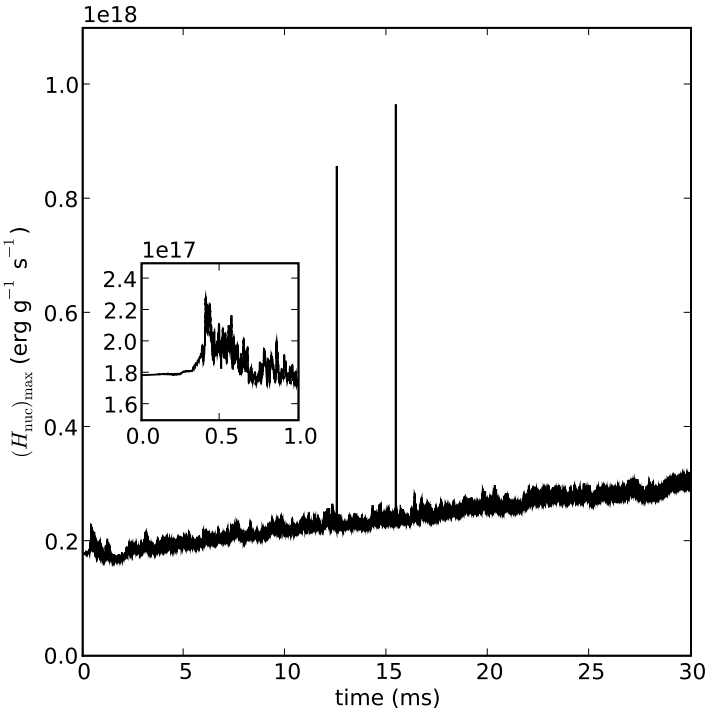

Figure 5.15: Plot of the maximum $H_{\rm nuc}$ in the `cold` model simulation as a function of time. The inset plot shows the early adjustment phase associated with Figures 5.9(b) and 5.9(c). The spikes are similar to that seen in the bottom panel of Figure 5.8, and are caused by the rapid burning of fresh fuel as it is brought into the burning layer by the turbulent convection.

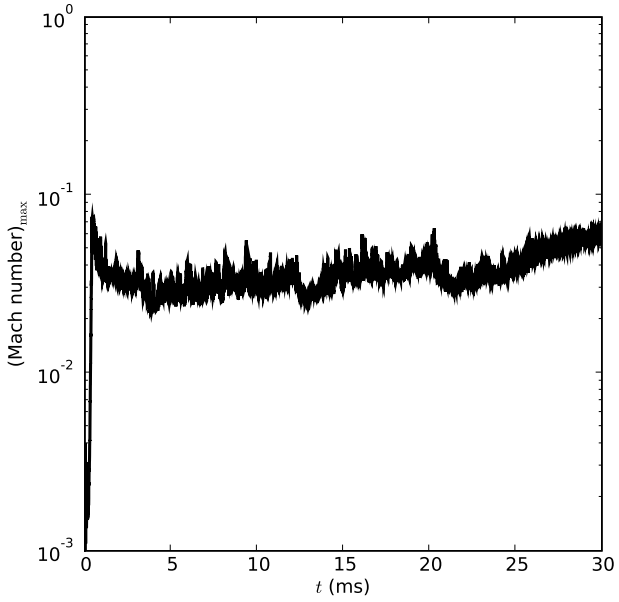

Figure 5.16: Plot of the maximum Mach number in the `cold` model simulation as a function of time. The slow convective flow justifies the use of a low Mach number approximation method.

Figure 5.15 shows the evolution of the maximum value of $H_{\rm nuc}$ throughout the duration of the `cold` model simulation. The inset plot shows the early adjustment phase mentioned above. The initial jump in $H_{\rm nuc}$ is due to the fact that there is no well established flow field that can efficiently advect away the energy released from reactions. Once the convective zone is well established, the energy generation rate relaxes before making its steady climb. We note that we have not yet achieved runaway — the rise in energy generation rate is still linear. This climb is temporarily interrupted by a couple of spikes similar to those seen in Figure 5.8 when fresh fuel is advected to a relatively hot region and burned quickly. Although well organized at later times, the convective fluid flow is slow with respect to the sound speed. Figure 5.16 shows the maximum Mach number in the computational domain as a function of time; this value never exceeds 0.08 in our simulation. The average value of the Mach number in the convective region rarely exceeds ~ 0.02 during our 30 ms simulation.

5.3 Conclusions

We have described some of the challenges and important concepts to keep in mind when performing multidimensional simulations of XRBs. The major results for this study of pure ^{4}He accretors, can be summarized as follows:

- To get a system that is much closer to thermal instability in multiple dimensions, the semi-analytic one-dimensional models should augment the local cooling rate estimate, (5.4), to include cooling due to convection.

- Properly resolving the burning layer using the initial models considered here requires a spatial resolution of 0.5 cm zone^{-1}, which is an order of magnitude higher than what has been presented in the literature for multidimensional models (Lin et al., 2006). It should be noted that our initial models differ in the underlying neutron star's composition — their ^{12}C opposed to our ^{56}Fe — and their models were evolved in one dimension through several bursts before being mapped into multiple dimensions.

- Under-resolving the burning layer leads to dramatic convective undershoot and the burning tends to die out.

- At the early times simulated here, the inclusion of thermal diffusion has little effect on the maximum energy generation rate, but does perhaps affect the depth at which this maximum occurs.

- The `MAESTRO` algorithm we use allows us to capture the expansion of the atmosphere due to heating, which will be important in the modeling of PRE burst sources.

- The average thermal gradient in the convective region is nearly adiabatic but there are localized pockets and filamentary structures that are either super- or sub-adiabatic.

- The strong convection interacts with and churns up the underlying neutron star material, which slightly alters the conductivity of the convective region.

The initial selection of a value to use for the volume discrepancy term in our simulations was based on experience with other applications. As we showed in Section 5.2.4, the value used for the long duration simulation in this paper, $f = 0.3$, may not the optimal choice for the XRB problem. Further investigation is required to determine which factors affect the appropriate value of f, and to determine if the spikes in the drift of the thermodynamic pressure from the base state pressure are simply numerical artifacts of a poor choice of f.

The width of the computational domain used in our simulations is adequate for the early evolution of the system; the size of any individual convective cell is initially small with respect to the width of the domain. As the system evolves and the convection becomes more established, the cells grow in size. The nature of vorticity in two dimensions is such that the smaller vortices merge to form a single vortex. In our simulations the cells grow to become a significant fraction of the domain width and the flow becomes dominated by a single vortex that interacts with itself through the periodic boundary conditions. By selecting a wider computational grid, we could delay the formation of a single, dominant vortex. Ideally the computational domain should be several pressure scale-heights wide so that we should form multiples of these convective cells that dominate the flow for an extended period of time before merging into a single vortex. Given our strict resolution requirements, such a setup was computationally infeasible.

CHAPTER 6

TWO-DIMENSIONAL SIMULATION OF A MIXED H/HE BURST SOURCE

In this chapter, I describe my ongoing study of thermonuclear burning and convection in a mixed H/He accreting XRB system. As this study is still very much a work-in-progress, I will present the material in a heuristic fashion.

The start of my XRB dissertation research began by studying a system that had accreted pure ^{4}He (Chapter 5). As stated previously, the idea behind starting with a pure ^{4}He system was that these should be "easy" to simulate, mainly because of the simple reaction network required (Section 4.3). Incidentally, these simulations turned out to be computationally more expensive than we previously imagined due to the high resolution required to resolve the burning layer (Section 5.2.1).

The bulk of the energy release in an XRB system with unstable ^{1}H-burning comes from reactions that are much less temperature-sensitive than the 3–α reaction. Therefore, the peak of the burning layer in a ^{1}H-burning XRB should be more extended than for a ^{4}He-burning, requiring less resolution. We would have a larger network to solve for each zone, but we should have fewer zones compared to the pure ^{4}He system studied in Chapter 5. In addition, by applying higher resolution

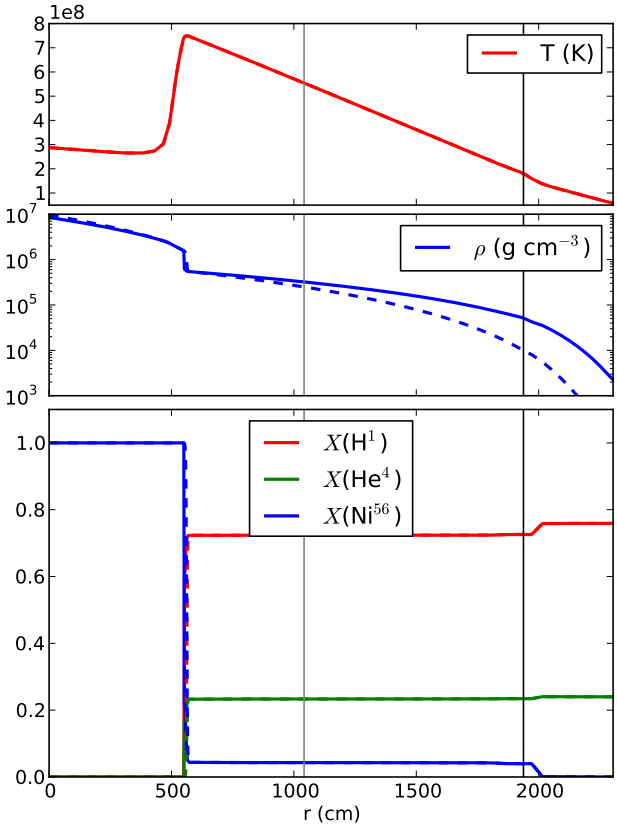

Figure 6.1: The mixed H/He burst models. The solid lines show the original Lagrangian data from `Kepler` provided by Alex Heger. The dashed lines show this model smoothed (see Section 5.1.3) and put into HSE with our EOS.

only where it was needed using the multilevel aspect of `MAESTRO` outlined in Paper V, which was not fully implemented at the beginning of the studies of Chapter 5, we thought we might end up with a simulation that had a computational cost less than or equal to that of the pure ^{4}He system.

6.1 Mixed H/He Initial Models

From our studies of a pure ^{4}He burst source, we knew that we needed an initial model that somehow contained an approximation to the cooling aspects of convection (see Section 5.1.1). During the excellent meeting "X-ray Bursts and Burst Oscillations," held at the Lorentz Center on the campus of the *Universiteit Leiden*, we approached Alex Heger about our need for new initial models. Within the hour, Alex graciously supplied me with approximately 3,000(!) output files from the evolution of a solar-composition-accreting XRB source using the `Kepler` code. These files followed the evolution of the system through three burst cycles.

From the models provided by Alex, I selected an initial model file that was ~ 0.3 s before

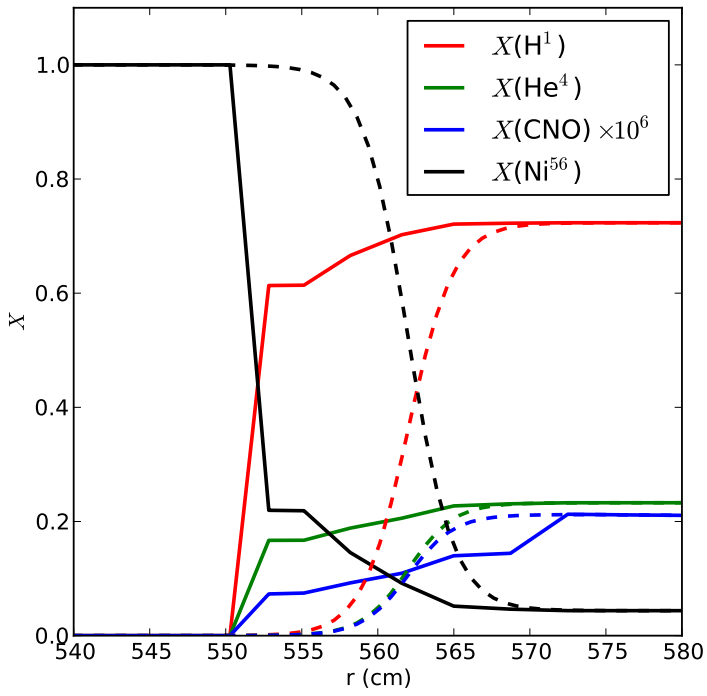

Figure 6.2: Closeup of the species in the smoothed region of Figure 6.1. Solid lines show the original `Kepler` data, while the dashed lines show our smoothed reconstruction.

the time when the temperature at the base of the accreted layer begins to significantly increase. Applying the same smoothing technique described in Section 5.1.3, I mapped the Lagrangian `Kepler` data into `MAESTRO`. Figure 6.1 shows both the original `Kepler` model (solid lines) and the smoothed model, which is in HSE with our EOS (dashed lines). The black and gray vertical lines are the boundaries of the sponge term of Equation 3.17. To see the smoothing in action, Figure 6.2 shows a closeup of Figure 6.1 on the smoothed region. Here, we also plot the total CNO mass fraction, $X(\mathrm{CNO}) = X(^{12}\mathrm{C}) + X(^{14}\mathrm{N}) + X(^{16}\mathrm{O})$, scaled to be of the same order as the ^{1}H , ^{4}He , ^{56}Ni mass fractions.

6.2 Issues with the Initial Model

When I first started running the model shown in Figure 6.1, I noticed odd oscillations in the temperature near the smoothed transition region. The oscillations tended to grow as the system was evolved, regardless of choice of reaction network `hotcno` or `approx8`. Figure 6.3 shows a closeup of such an oscillation after 0.1 ms of evolution in a simulation that used the same parameters as the simulations presented in Section 5.2.5 for a pure ^{4}He burst model.

What was so different about this model compared to the pure ^{4}He models that would cause bad numerical behavior? The temperature at the base of this mixed model is $\sim 40\%$ higher than that of the `hot` model of Chapter 5, but the hot CNO cycle — the main source of energy generation for this model — is not temperature-sensitive (Fowler & Hoyle, 1965) and therefore physically, the nuclear burning is probably not the source of the oscillations.

Numerically, however, there could be something that is causing the solution to behave in an unstable manner. One of the first tests we used to determine if the oscillations were caused by numerics was to run the same simulation but vary the advective CFL number, while keeping Δx fixed. Figure 6.4 shows the results of this test for the same simulation and at the same time as in

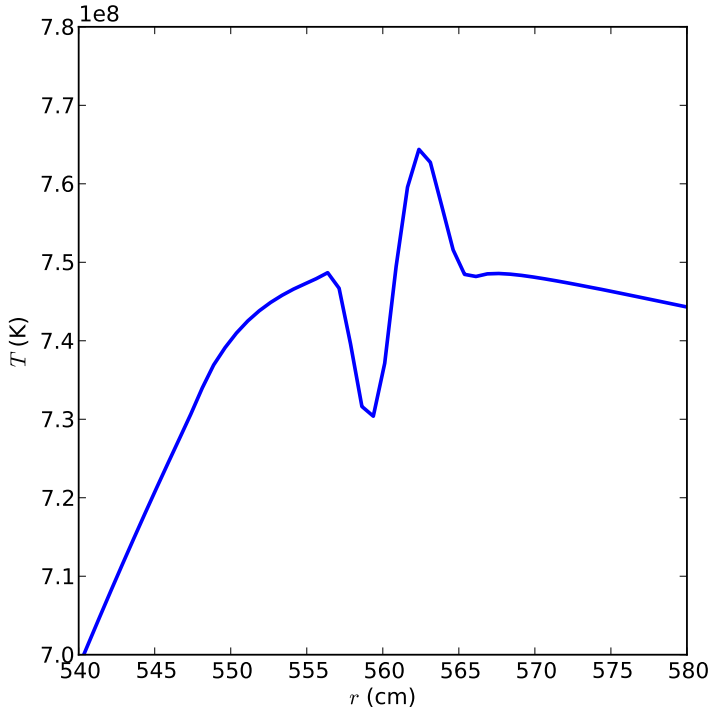

Figure 6.3: Closeup of a temperature oscillation near the base of the accreted layer after 0.1 ms of evolution. Such oscillations tend to grow out of control with evolution.

Figure 6.3, but with the value α_{CFL} changed from the default value of 0.7. Interestingly, we find that the amplitude of the oscillations decreases with decreasing CFL number. This result seemed to suggest that there was somehow a decoupling between the hydrodynamics and the thermonuclear burning.

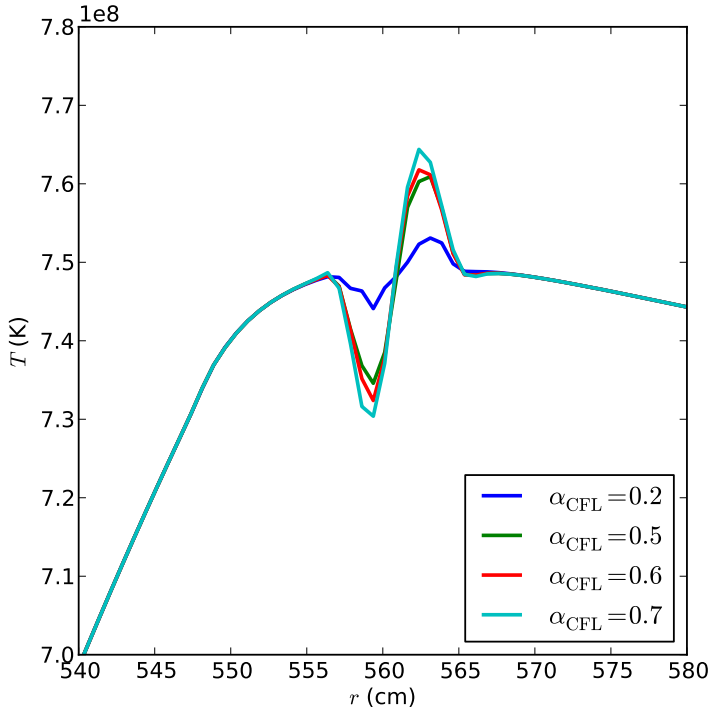

Figure 6.4: The effects of varying the CFL number for the mixed model. The $\alpha_{\mathrm{CFL}} = 0.7$ curve is the same as that in Figure 6.3.

6.3 Attempting Spectral Differed Corrections

Operator-split methods, such as that of Strang, are designed to break a complex problem into several more tractable problems. Take, for example, the enthalpy equation, 3.6,

$$\frac{\partial(\rho h)}{\partial t} = -\nabla\cdot(\rho h\mathbf{U}) + \frac{Dp_0}{Dt} + \rho H_{\mathrm{nuc}} + \nabla\cdot(k_{\mathrm{th}}\nabla T).$$

This equation could be rewritten as

$$\frac{\partial(\rho h)}{\partial t} = (\mathrm{AD})[\rho h] + (\mathrm{R})[\rho h], \tag{6.1}$$

where (AD) and (R) can be thought of as the Advection-Diffusion and Reaction *operators* acting on ρh, respectively. Note that these operators, in general, depend on the thermodynamic state of the material, as well as the velocity field. As mentioned in Section 3.2, MAESTRO uses Strang-splitting to couple the hydrodynamics and nuclear burning. This form of operator splitting is second-order accurate, and is equivalent to evolving ρh from initial data via the Reaction operator

$$\frac{\partial (\rho h)}{\partial t} = (\mathrm{R})\,[\rho h] \tag{6.2}$$

over a step size of $\frac{1}{2}\Delta t$, using the result from that step as initial conditions for evolving with the Advection-Diffusion operator

$$\frac{\partial (\rho h)}{\partial t} = (\mathrm{AD})\,[\rho h] \tag{6.3}$$

over a step size of Δt, and then using the result from the advection step as inputs into another reaction step, via Equation 6.2, over a step of size $\frac{1}{2}\Delta t$. If the order of operations of the Advection-Diffusion and Reaction operators does not matter — i.e. the operators commute — then there is no additional error introduced in the splitting of the right-hand side of the enthalpy equation (see Chapter 17 of LeVeque, 2002, for example).

Consider, however, a case where the operators would not commute. Imagine a fluid element in a high temperature region with strong local velocities — strong enough that the timescale for advection is comparable to the timescale for nuclear burning — pointed in a direction *away* from the high temperature region. Physically, then, the fluid should move away from the high temperature region before a significant amount of burning has occurred. The Strang-split prescription above, however, would have the fluid element at rest, vigorously burning because of the high temperature, *then* advect away, and burn some more a little less vigorously. The net result may be that too much burning has occurred, and now the fluid element has followed a slightly different thermodynamic trajectory. Repeatedly applying this process could then introduce significant errors, even if the

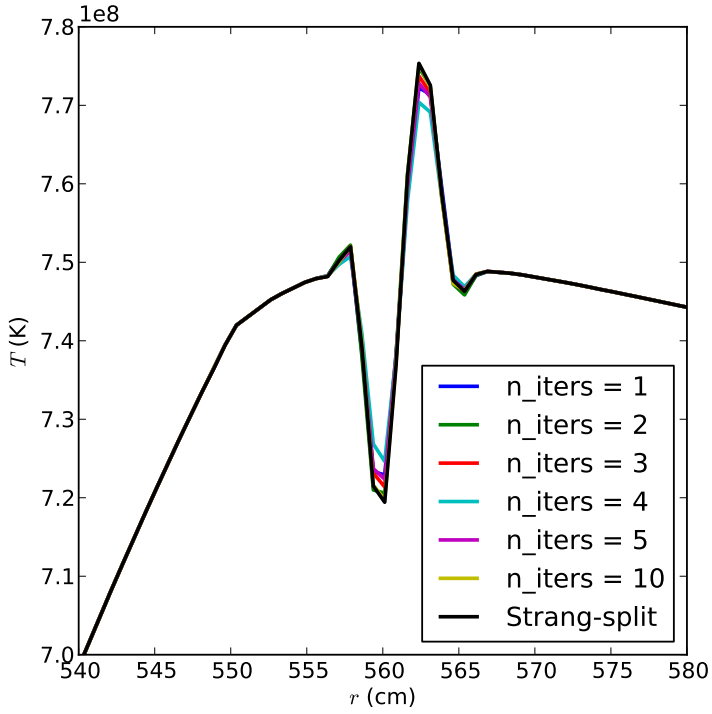

Figure 6.5: The effects of using the SDC method on the temperature oscillations observed in Figure 6.3. The curve labelled "Strang-split" is the same as the $\alpha_{\mathrm{CFL}} = 0.7$ curve in Figure 6.4. All of the SDC methods, labelled by the number of iteration steps, also used $\alpha_{\mathrm{CFL}} = 0.7$.

method is formally second-order accurate.

To remedy this in MAESTRO, my collaborators implemented a Spectral Differed Correction (SDC) method, with initial application to flame-propagation simulations (Nonaka et al.; in preparation). SDC methods were introduced by Dutt et al. (2000) to reduce splitting errors by applying a series of correction equations that attempt to couple the split operators. The basic idea is that an estimate for the solution at the new time step is determined using an operator-split method, and then iterative corrections are applied by applying the various operators to updated approximations of the solution. Practically, this means that the Advection-Diffusion and Reaction operators must know something about what the other operator is doing. This involved my creating a new reaction net-

work, cleverly named `approx8_SDC`, which incorporates enthalpy and species advection source terms into the right-hand side of ODEs in the network. It should be noted that the SDC method in `MAESTRO` does not currently support thermal diffusion, so the remainder of the simulations in this chapter are run without diffusion.

Recall that the whole reason for going to SDC was to verify our conclusions from the CFL tests of Section 6.2 — namely that there was a decoupling of the hydrodynamics from the nuclear reactions. Figure 6.5 shows the result of using the SDC method for the same problem as the $\alpha_{\mathrm{CFL}} = 0.7$ case as in Figure 6.4, but with thermal diffusion turned off. The label "n_iters" denotes the number of iterative correction steps used in the SDC method. The first, and most obvious, thing to note is that using the SDC method did not create much of an improvement — these oscillations still grow with evolution. It does appear, however, that the SDC algorithm is doing *something*, as increasing the number of corrective iterations generally reduces the amplitude of wiggles — at least up until n_iters= 4. Furthermore, comparing Figures 6.4 and 6.5 shows that diffusion seems to make things better as well. In any case, the result of this venture into SDC methods, although not correcting the problem for the mixed XRBs, has added a new method to the `MAESTRO` code for use in future problems.

6.4 Using Various Edge Prediction Methods

After all the work involved in setting up the SDC method, we were no closer to solving the temperature oscillation problem. So, again I blamed the initial model and again I asked: what is so different about this model compared to the pure ^{4}He models? I described in Section 6.2 that the temperature was similarly behaved in the two models. The density, on the other hand, *does* behave quite differently between the two models at the base of the accreted layer. The solid lines of Figure 6.6 show the density profile, normalized to the density at the base of the layer, for the `hot` model

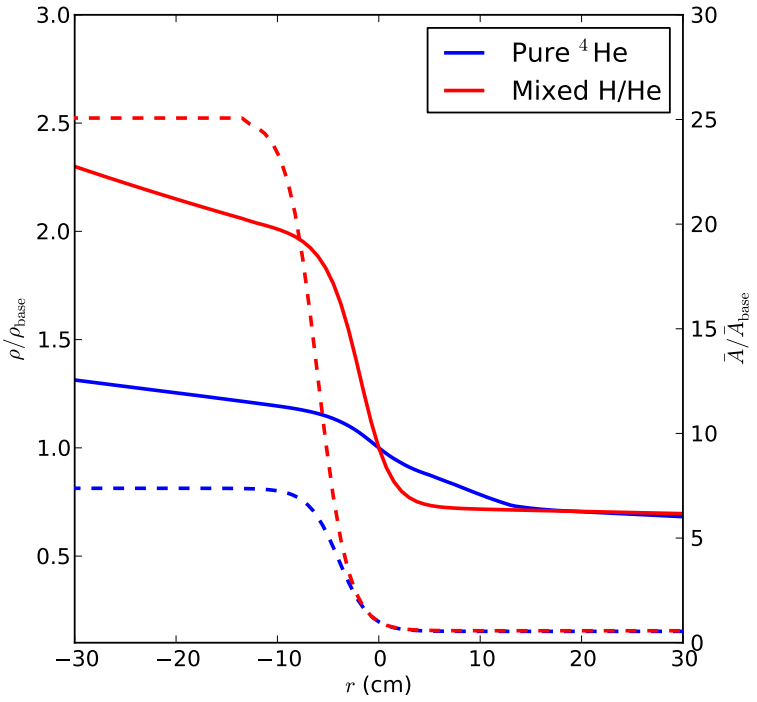

Figure 6.6: Density profiles normalized to the density at the base of the accreted layer (solid lines) along with the mean molecular weight profiles normalized to the value at the base (dashed lines). The curve labelled "Pure ^{4}He," is the `hot` model of Section 5.1.3. For easy comparison, the profiles have been shifted so that the base of the accreted layer lies at $r = 0$.

of Section 5.1.3 (blue) and the model of Figure 6.1 (red). Here, I have defined the base to be the radius at which

$$X_{\rm fuel} = \frac{X_{\rm fuel,max}}{2}, \tag{6.4}$$

where $X_{\rm fuel} = X\left(^4{\rm He}\right)$ for the pure ^{4}He model and $X_{\rm fuel} = X\left(^1{\rm H}\right)$ for the mixed H/He model. To facilitate comparison, the models have been shifted so that, in this plot, the base corresponds to $r = 0$. Figure 6.6 shows that going from the accreted atmosphere ($r > 0$) into the neutron star material ($r < 0$) results in a significantly larger jump in density for the mixed model compared to

the `hot` model. This is mainly due to the disparity in the transition in mean molecular weight,

$$\overline{A} = \frac{\sum_k X_k}{\sum_k (X_k/A_k)}, \tag{6.5}$$

which is shown as dashed lines in Figure 6.6.

Recall from the discussion of finite volume methods in Section 2.2.1 that slope limiters are used in higher-order Godunov methods to preserve monotonicity of a solution as it evolves. In regions with sharp profiles, as is the case for density in Figure 6.6, the limiting is more severe, and the reconstructed fluxes at cell edges can be significantly altered. I tried a few things, such as group-limiting the species — applying the same, maximum limiter to all the species – in the species advection equation, but this did not fix the problem.

As mentioned in Section 3.2, we use the slope-limited reconstruction to extrapolate perturbational quantities to cell edges for the conservative updates of the system. There are, however, numerous ways to combine perturbational and non-perturbational terms at the edges for a conservative update. For example, the default method of updating (ρX_i), and hence ρ, is by predicting ρ' and X_i to the edges and then constructing the forcing terms from these values. For the enthalpy update, the default is to predict $(\rho h)'$ to the edges, and use those values to construct the forcing terms. As another example, we note that in `MAESTRO`, the temperature can be determined either from the thermodynamic pressure or from the enthalpy, and therefore, the enthalpy can be updated by first predicting temperature to the edges, and then reconstructing enthalpy-related forcing terms from EOS calls.

In total, I investigated three approaches of edge-prediction for the species/density update and four methods for the enthalpy update. The 12 possible combinations along with qualitative descriptions of the temperature obtained from the pressure (red) or enthalpy (blue) are shown in Table 6.1. An amplitude deviation for a "good" method is $\lesssim 0.5\%$; for a "so-so" method it is between 0.5%

Table 6.1: Various edge prediction methods and a qualitative description of the temperature profiles determined from the pressure (red) or enthalpy (blue). The quantities listed in column/row headings are the variables that are predicted to the edge of cells. Figures 6.7 through 6.13 show the temperature profiles for each row and column of the table.

		enthalpy prediction methods			
		$(\rho h)'$	h	$T \rightarrow (\rho h)'$	$T \rightarrow h$
species prediction methods	ρ', X	bad	bad	bad	bad
		bad	so-so	good	good
	$\rho', \rho X$	good	so-so	good	good
		good	terrible	good	bad
	ρ, X	good	good	good	good
		good	good	good	good

and $\lesssim 2\%$; for a "bad" method it is between 2% and $\lesssim 3\%$; anything with a deviation greater than $\sim 3\%$ is considered "terrible. " The variables listed in the column and row headings are the variables that are predicted to the edges. Note that for the enthalpy updates, in the case of temperature predicted to the edge, an EOS call is made to get something that relates to the enthalpy. For completeness, I also show plots of the temperature profiles as determined from the pressure (red) or enthalpy (blue) for each row (Figures 6.7, 6.8, and 6.9) and column (Figures 6.10, 6.11, 6.12, and 6.13) of Table 6.1. There is quite a large range of behaviours between these different combinations of edge-prediction methods. In general, it appears that predicting a perturbational density, ρ', exhibits the worst behaviour, regardless of the enthalpy prediction method . Furthermore, predicting the full state density and species separately gives the most robust method, at least for the mixed H/He XRB problem. This finding is contrary to what was found using `MAESTRO` for full-star simulations of convection in a white-dwarf, where predicting perturbational quantities

resulted in better behaved profiles (Paper IV).

6.5 Current Status

All of the different methods for edge prediction should, in theory, converge to the same solution as the grid spacing, Δx, goes to zero. As a sharp transition gets spread over more and more zones, the slope limiters need to do less limiting in order to preserve monotonicity and the reconstruction function more closely approximates the true solution. In this limit, the differences in results between the various edge prediction methods should decrease, with all methods converging to the same solution. I am currently in the process of running simulations to numerically confirm the convergence of all the different methods.

For now, however, I have concluded that the proper choice of edge-prediction is problem specific, and likely depends on the initial model. It should be noted that an alternative way to get rid of the temperature oscillations was to extremely smooth the initial model profile over a physically unrealistic length scale. Currently, I have settled on using the combined species prediction (ρ, X) and enthalpy prediction $\left((\rho h)'\right)$ methods. I have ongoing simulations to confirm that this choice does not produce oscillations with further evolution, but preliminary studies suggest that this is the case. Eventually, these simulations will be carried out to long enough timescales to make comments about the convective dynamics, similar to Section 5.2.5 for the pure ^{4}He burst source.

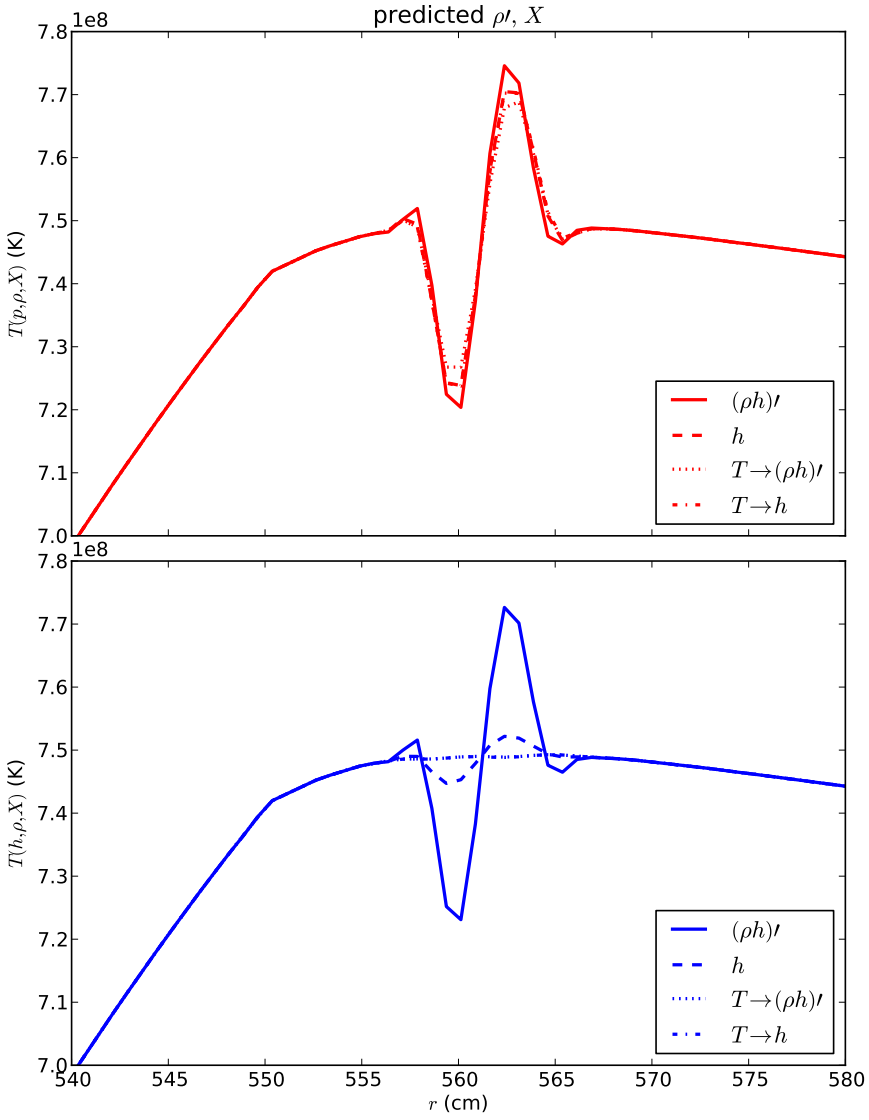

Figure 6.7: Results for the temperature as determined from the pressure (red) and enthalpy (blue), using the species edge-prediction row labelled "ρ', X" in Table 6.1 — the labels indicate which species prediction method was used. Aside from the choice of edge-prediction, the simulation parameters are the same as those from Figure 6.3.

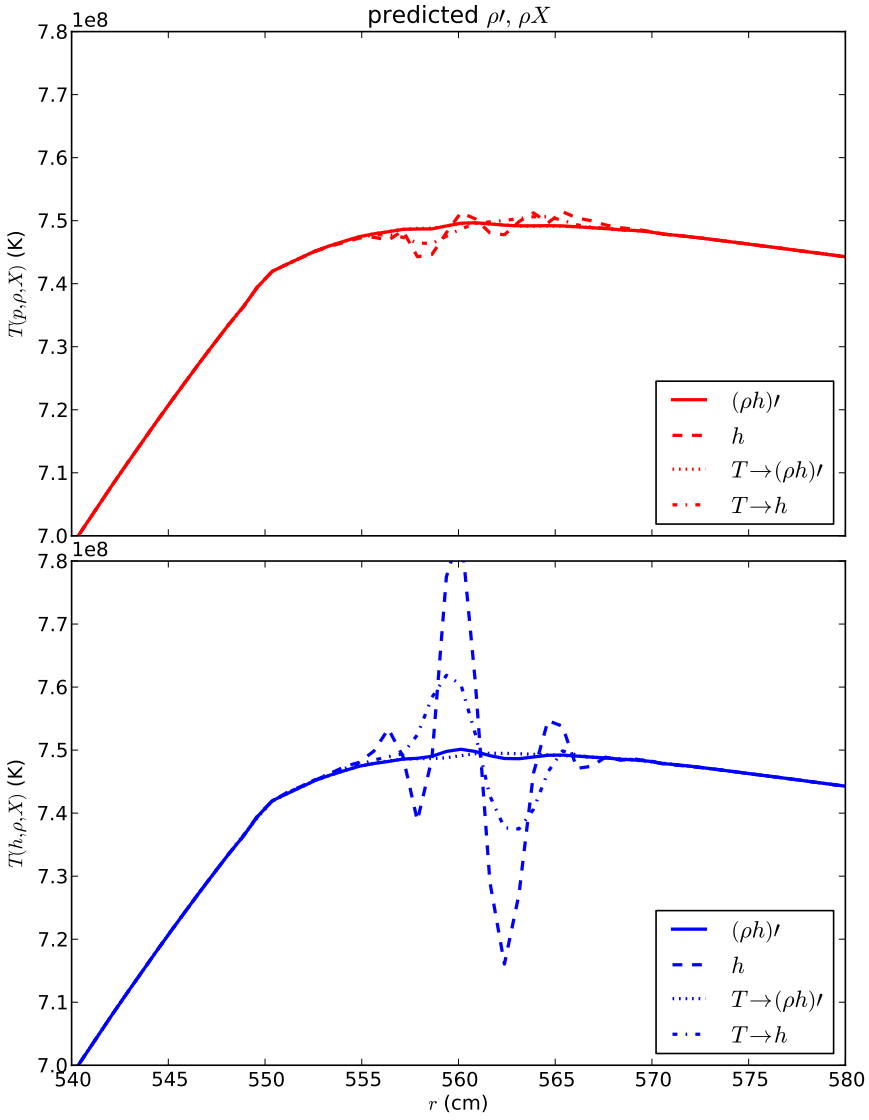

Figure 6.8: Results for the temperature as determined from the pressure (red) and enthalpy (blue), using the species edge-prediction row labelled "ρ', ρX" in Table 6.1 — the labels indicate which species prediction method was used. Aside from the choice of edge-prediction, the simulation parameters are the same as those from Figure 6.3.

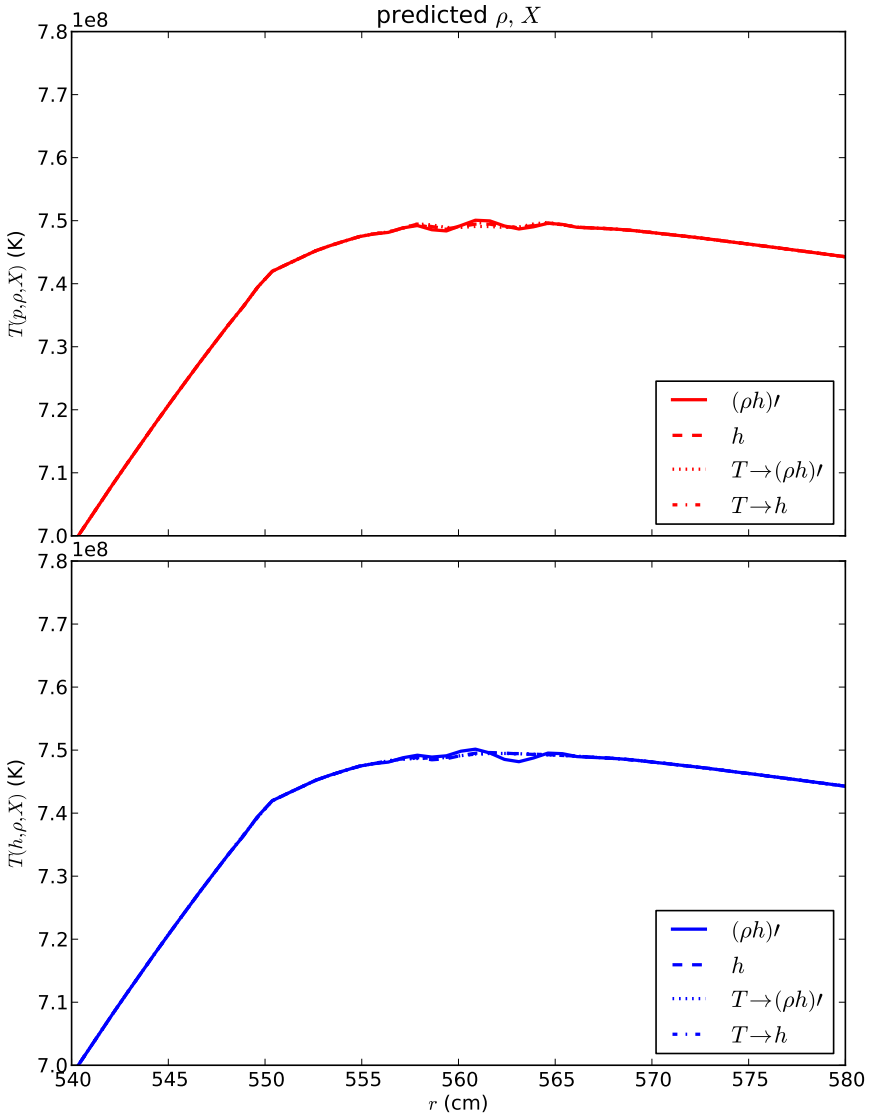

Figure 6.9: Results for the temperature as determined from the pressure (red) and enthalpy (blue), using the species edge-prediction column labelled "ρ, X" in Table 6.1 — the labels indicate which species prediction method was used. Aside from the choice of edge-prediction, the simulation parameters are the same as those from Figure 6.3.

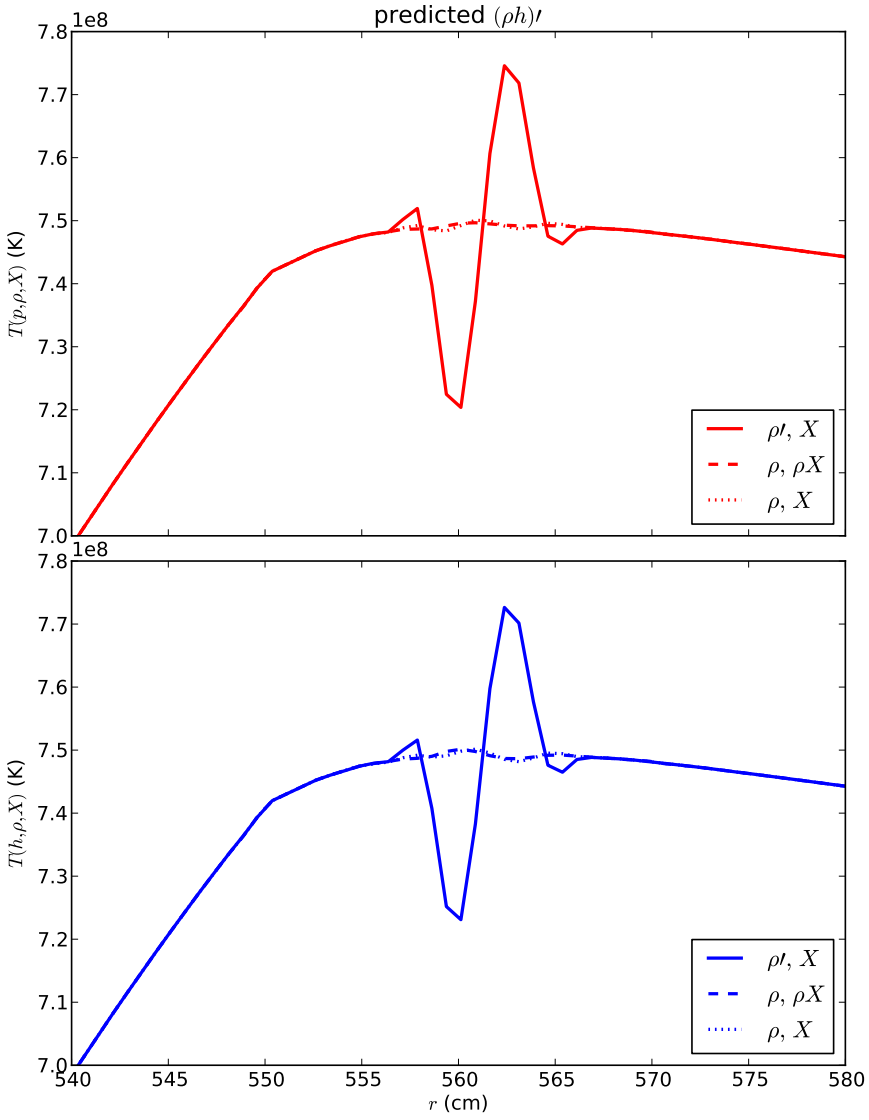

Figure 6.10: Results for the temperature as determined from the pressure (red) and enthalpy (blue), using the enthalpy edge-prediction column labelled "$(\rho h)'$" in Table 6.1 — the labels indicate which species prediction method was used. Aside from the choice of edge-prediction, the simulation parameters are the same as those from Figure 6.3.

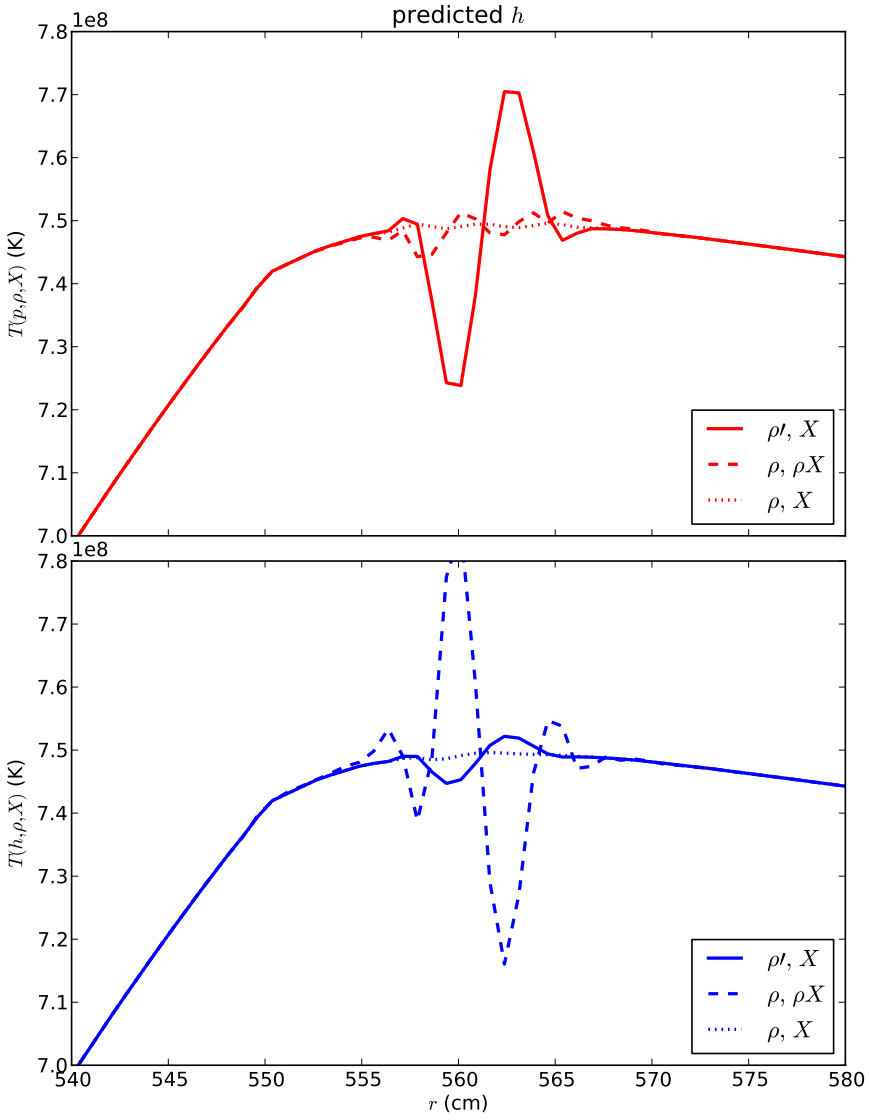

Figure 6.11: Results for the temperature as determined from the pressure (red) and enthalpy (blue), using the enthalpy edge-prediction column labelled "h" in Table 6.1 — the labels indicate which species prediction method was used. Aside from the choice of edge-prediction, the simulation parameters are the same as those from Figure 6.3.

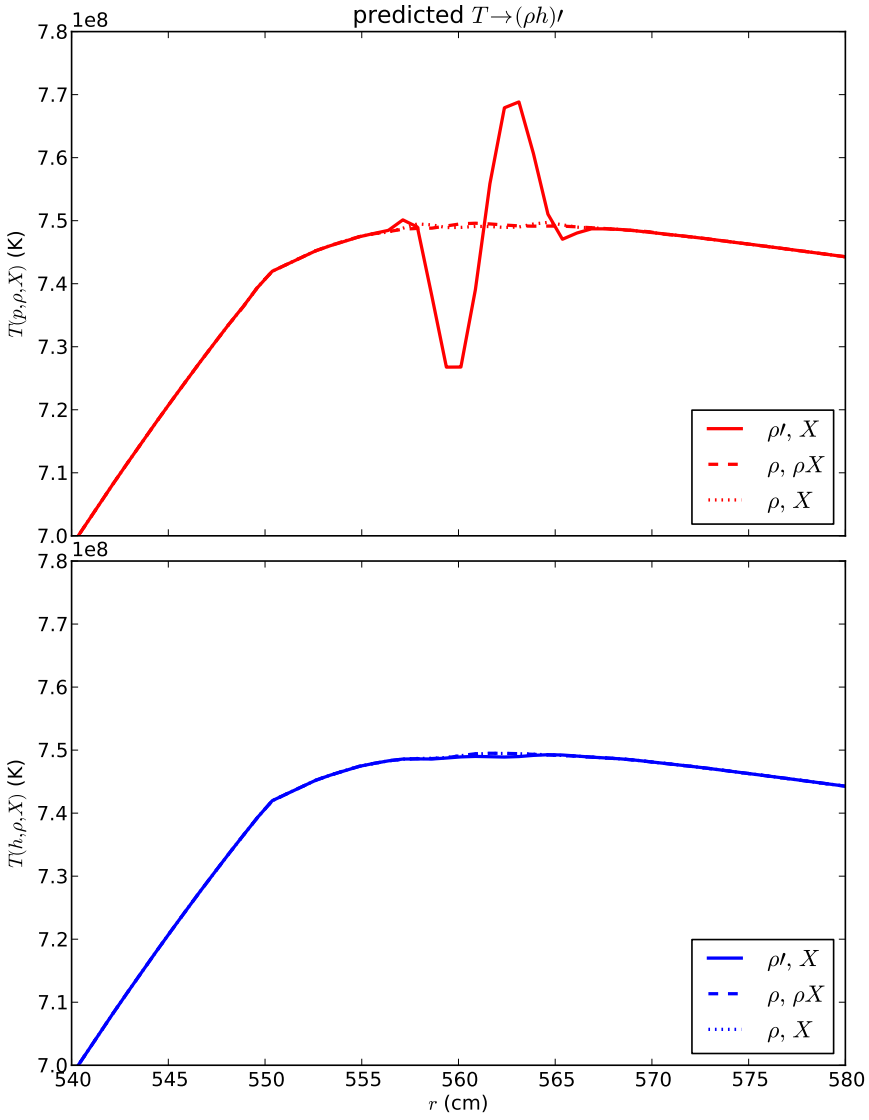

Figure 6.12: Results for the temperature as determined from the pressure (red) and enthalpy (blue), using the enthalpy edge-prediction column labelled "$T \rightarrow (\rho h)'$" in Table 6.1 — the labels indicate which species prediction method was used. Aside from the choice of edge-prediction, the simulation parameters are the same as those from Figure 6.3.

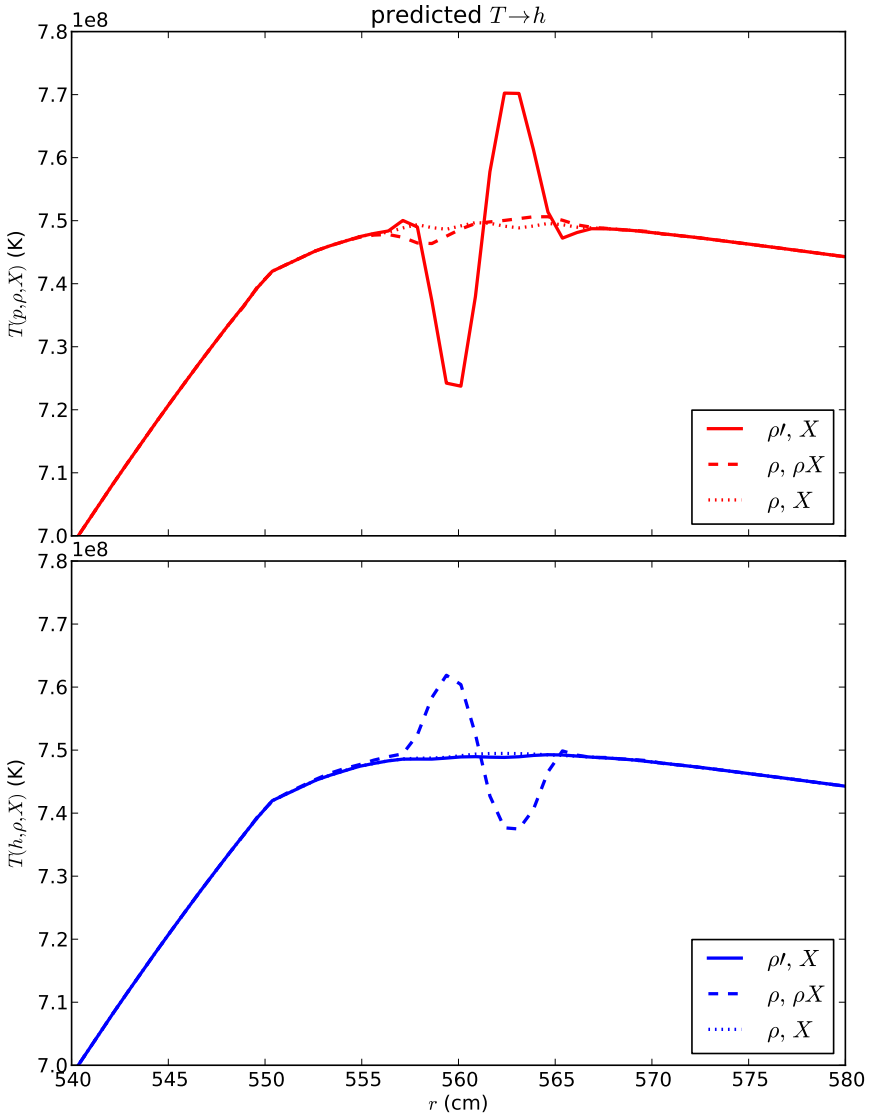

Figure 6.13: Results for the temperature as determined from the pressure (red) and enthalpy (blue), using the enthalpy edge-prediction column labelled "$T \rightarrow h$" in Table 6.1 — the labels indicate which species prediction method was used. Aside from the choice of edge-prediction, the simulation parameters are the same as those from Figure 6.3.

CHAPTER 7

SUMMARY AND FUTURE OUTLOOK

The research comprising this dissertation involved using a novel piece of technology, the MAESTRO code, to perform realistic, resolved multidimensional simulations of thermonuclear burning and the subsequent turbulent convection in a low Mach number framework. Lin et al. (2006) was the only other group in the literature to attempt such simulations, but their results, according to our simulations of a pure ^{4}He source in Chapter 5, were perhaps under-resolved. Furthermore, the Lin et al. (2006) numerical method was first-order accurate in space and time compared to MAESTRO's second-order accuracy, and it could not capture the expansion of the base state — as we did in Section 5.2.3 — or model the surface of the atmosphere without numerical problems, which the MAESTRO algorithm avoids by using a sponging technique as described in Section 3.2.1.

Along the way, I added several reaction networks to the MAESTRO code. I tested the VODE solver's ability to handle very stiff systems of equations in Section 4.4.1, and compared two different reaction rates for the 3–α reaction in Section 4.4.2. Furthermore, I implemented two of Frank Timmes's publicly available reaction networks into the MAESTRO code, and modified them to evolve the equations in a self-consistent, low Mach number framework. These networks are quite general, and can be applied to many other problems, thus adding to the capability of MAESTRO.

In Chapter 5, I described the successful application of the MAESTRO code to an XRB that had pure ^{4}He accretion. One of the main results of that study was that we find some of heavy-element neutron star material being dredged up into the atmosphere via the shearing occurring at the lower boundary of the convective region. This is already an important result, as it suggests that the opacities in the atmosphere could be significantly different from what is typically assumed when making inferences about bursts from observations. This will certainly be true if such a polluted convective zone can reach the photosphere. Our simulations were only run out to $t = 30$ ms — a factor of 200 further in time than the compressible detonation studies of Zingale et al. (2001) — but the edge of the convective region was still expanding, albeit slightly, at the end of our simulation (see Figure 5.11). If this result holds in three-dimensional simulations, in which the turbulent energy cascade behaves differently, then some of the assumptions used to infer the mass and radius of the underlying neutron star may need to be altered. Furthermore, we were not able to run this model to ignition of a deflagration, the propagation of which would be altered by the presence of heavy elements that change the conductivity. Given the resolution requirements we found in Section 5.2.1 for the initial model used here, significantly more computational resources will be needed to run these simulations out until ignition, or to investigate the three-dimensional effects.

I discussed simulations of a mixed H/He burst in Chapter 6 , and I presented some of the trials and tribulations of applying a new algorithm to a new problem, even though the problem at first glance did not seem all that different. I hope the reader has somewhat of an appreciation for the months-long "debugging" process involved — I know I certainly do. I say "debugging" in quotes because, in fact, the issues were not bugs in the code, but rather features of the algorithm. One advantage of working with a code that is completely novel is that these fun features seem to pop up in the most odd — and inopportune — places. However, applying several different approaches towards a solution, such as implementing new methods (SDC in Section 6.3) or using alternate

numerical methods which are analytically equivalent (edge-prediction methods in Section 6.4), results in a better understanding of where these features occur and when they are important. Based on my study of the evolution of the temperature oscillations in the mixed H/He burst model, I think we have found the correct combination of edge-prediction methods that results in a well-behaved solution. I am currently running these simulations out further for confirmation.

There is still plenty of research to be performed on the subject of multidimensional simulations of XRBs as deflagrations. Ideally, I would like to be able to carry out a simulation until ignition in two dimensions. This may be feasible with MAESTRO's adaptive mesh refinement (AMR) capability, which allows for placing more resolution only where it is needed, such as in a thin burning layer. By selectively refining the simulation in this fashion, the algorithm can spend less computational time in regions that are not of dynamical importance. This capability was not completely implemented at the start of the studies presented in Chapter 5, but they are currently being used with the mixed burst models of Chapter 6. AMR will save even more computer time needed to solve three-dimensional simulations, which are the next big step in multidimensional simulations of XRBs.

There are also more interesting physics that can be added to the problem. An obvious choice is a larger reaction network, but as I alluded to in Chapter 4, the trend in multiple dimensions is to forgo the detailed physics for a network that is more easily solved. One way around this is to implement tracer particles that simply advect with the fluid, but carry their thermodynamic history. One can then post-process these trajectories with a large reaction network to get more detailed nucleosynthesis. Such tracer particles have been preliminarily implemented into MAESTRO, and I have helped with some of the early analysis routines for this approach, but more work needs to be done. In addition to nuclear physics, there are also the effects of rotation and magnetic fields. Neutron stars are typically fast rotators, and the rotation may affect the propagation of a burning front Spitkovsky et al. (2002). I implemented rotation for planar-geometry problems into MAESTRO,

but more thought is needed on how to handle boundary conditions in a rotating frame of reference. Magnetic fields should, in principle, be tedious, but straightforward to implement into MAESTRO. One of the plagues of magnetohydrodynamics simulations is maintaining the divergence constraint on the magnetic field, $\nabla \cdot \mathbf{B}$, from Maxwell's equations. As discussed in Section 3.2, MAESTRO uses projection operators to satisfy its velocity divergence constraint, which is significantly more complicated than the magnetic field constraint of Maxwell. In essence, the "machinery" is already in the code for solving and enforcing such equations — the difficulty will lie in figuring out exactly how to fold such a constraint into MAESTRO's predictor-corrector formalism, or even the new SDC method. I will likely be working on this task during my postdoctoral appointment.

BIBLIOGRAPHY

Alastuey, A., & Jancovici, B. 1978, ApJ, 226, 1034

Almgren, A. S. et al. 2010, ApJ, 715, 1221

Almgren, A. S., Bell, J. B., Nonaka, A., & Zingale, M. 2008, ApJ, 684, 449, paper III

Almgren, A. S., Bell, J. B., Rendleman, C. A., & Zingale, M. 2006a, ApJ, 637, 922, paper I

—. 2006b, ApJ, 649, 927, paper II

Altamirano, D., Watts, A., Linares, M., Markwardt, C. B., Strohmayer, T., & Patruno, A. 2010, ArXiv e-prints

Arnett, D. 1996, Supernovae and Nucleosynthesis, ed. Arnett, D.

Arnett, D., Meakin, C., Starrfield, S., Timmes, F., & Young, P. 2008, in American Institute of Physics Conference Series, Vol. 1001, Evolution and Nucleosynthesis in AGB Stars, ed. R. Guandalini, S. Palmerini, & M. Busso, 287–294

Arnett, W. D., Meakin, C., & Young, P. A. 2009, ArXiv e-prints 0910.0821

Belian, R. D., Conner, J. P., & Evans, W. D. 1976, ApJ, 206, L135

Bhattacharyya, S. 2010, ArXiv e-prints 1001.1642

Bildsten, L. 2000, in American Institute of Physics Conference Series, Vol. 522, American Institute of Physics Conference Series, ed. S. S. Holt & W. W. Zhang, 359–369

Böhm-Vitense, E. 1958, ZAp, 46, 108

Boussinesq, J. 1903, Theorie analytique de la chaleur, Vol. 2 (Gauthier-Villars, Paris)

Brown, P. N., Byrne, G. D., & Hindmarsh, A. C. 1989, SIAM Journal on Scientific and Statistical Computing, 10, 1038

Canizares, C. R. 1976, ApJ, 207, L101

Caughlan, G. R., & Fowler, W. A. 1988, Atomic Data and Nuclear Data Tables, 40, 283, see also http://www.phy.ornl.gov/astrophysics/data/cf88/index.html

Chen, C.-C., Durran, D. R., & Hakim, G. J. 2005, Journal of the Atmospheric Sciences, 62, 3213

Clayton, D. D. 1983, Principles of stellar evolution and nucleosynthesis, ed. Clayton, D. D.

Courant, R., Friedrichs, K., & Lewy, H. 1928, Mathematische Annalen, 100, 32

Cox, J. P., & Giuli, R. T. 1968, Principles of stellar structure, ed. Cox, J. P. & Giuli, R. T.

Cumming, A. 2003, ApJ, 595, 1077

Cumming, A., & Bildsten, L. 2000, ApJ, 544, 453

Cyburt, R. H. et al. 2010, ApJS, 189, 240, http://groups.nscl.msu.edu/jina/reaclib/db/

Durran, D. 1990, Meteorol. Monographs, 23, 59

Durran, D. R., & Klemp, J. B. 1983, Monthly Weather Review, 111, 2341

Dutt, A., Greengard, L., & Rokhlin, V. 2000, Bit Numerical Mathematics

Forman, W., Jones, C., Cominsky, L., Julien, P., Murray, S., Peters, G., Tananbaum, H., & Giacconi, R. 1978, ApJS, 38, 357

Fowler, W. A., & Hoyle, F. 1965, Nucleosynthesis in massive stars and supernovae, ed. Fowler, W. A. & Hoyle, F.

Fryxell, B. et al. 2000, ApJS, 131, 273

Fryxell, B. A., & Woosley, S. E. 1982a, ApJ, 258, 733

—. 1982b, ApJ, 261, 332

Fushiki, I., & Lamb, D. Q. 1987a, ApJ, 323, L55

—. 1987b, ApJ, 317, 368

Godunov, S. K. 1959, Mat. Sb., 47, 271

Graboske, H. C., Dewitt, H. E., Grossman, A. S., & Cooper, M. S. 1973, ApJ, 181, 457

Grindlay, J., & Gursky, H. 1976, ApJ, 205, L131

Grindlay, J., & Heise, J. 1975, IAU Circ., 2879, 1

Gustafsson, B., Kreiss, H.-O., & Sundström, A. 1972, Mathematics of Computation, 26, pp. 649

Hansen, C. J., & van Horn, H. M. 1975, ApJ, 195, 735

Heger, A., Cumming, A., Galloway, D. K., & Woosley, S. E. 2007a, ApJ, 671, L141

Heger, A., Cumming, A., & Woosley, S. E. 2007b, ApJ, 665, 1311

Heger, A., Langer, N., & Woosley, S. E. 2000, ApJ, 528, 368

Hoffman, J. A., Lewin, W. H. G., & Doty, J. 1977, ApJ, 217, L23

Hoffman, J. A., Marshall, H. L., & Lewin, W. H. G. 1978, Nature, 271, 630

Itoh, N., Totsuji, H., Ichimaru, S., & Dewitt, H. E. 1979, ApJ, 234, 1079

Kesel, P. G., & Winninghoff, F. J. 1972, Monthly Weather Review, 100, 360

Kippenhahn, R., & Weigert, A. 1994, Stellar Structure and Evolution, ed. Kippenhahn, R. & Weigert, A.

Lattimer, J. M. 2007, Astrophysics and Space Science, 308, 371

Lattimer, J. M., & Prakash, M. 2004, Science, 304, 536

LeVeque, R. J. 2002, Finite Volume Methods for Hyperbolic Problems (Cambridge University Press)

Lewin, W. H. G. et al. 1976, ApJ, 207, L95

Lewin, W. H. G., & Joss, P. C. 1981, Space Sci. Rev. , 28, 3

Lewin, W. H. G., van Paradijs, J., & Taam, R. E. 1993, Space Sci. Rev. , 62, 223

Lin, D. J., Bayliss, A., & Taam, R. E. 2006, ApJ, 653, 545

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

Maraschi, L., & Cavaliere, A. 1977, in X-ray Binaries and Compact Objects, ed. K. A. van der Hucht, 127–128

Mueller, E. 1986, A&A, 162, 103

Muno, M. P., Chakrabarty, D., Galloway, D. K., & Psaltis, D. 2002, ApJ, 580, 1048

Narayan, R., & Cooper, R. L. 2007, ApJ, 665, 628

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

Ogura, Y., & Phillips, N. A. 1962, Journal of Atmospheric Sciences, 19, 173

Pember, R. B., Howell, L. H., Bell, J. B., Colella, P., Crutchfield, C. Y., Fiveland, W. A., & Jessee, J. P. 1998, Combustion Science and Technology, 140, 123

Perkey, D. J., & Kreitzberg, C. W. 1976, Monthly Weather Review, 104, 744

Schatz, H. 2011, Progress in Particle and Nuclear Physics, 66, 277

Shara, M. M. 1982, ApJ, 261, 649

Spitkovsky, A., Levin, Y., & Ushomirsky, G. 2002, ApJ, 566, 1018

Steiner, A. W., Lattimer, J. M., & Brown, E. F. 2010, ArXiv e-prints

Strohmayer, T., & Bildsten, L. 2006, Compact Stellar X-Ray Sources, ed. W. H. G. Lewin & M. van der Klis (Cambridge: Cambridge Univ. Press), 113–+

Strohmayer, T. E., Zhang, W., Swank, J. H., Smale, A., Titarchuk, L., Day, C., & Lee, U. 1996, ApJ, 469, L9+

Swank, J. H., Becker, R. H., Boldt, E. A., Holt, S. S., Pravdo, S. H., & Serlemitsos, P. J. 1977, ApJ, 212, L73

Swesty, F. D., & Myra, E. S. 2009, ApJS, 181, 1

Taam, R. E., Woosley, S. E., Weaver, T. A., & Lamb, D. Q. 1993, ApJ, 413, 324

Timmes, F. X. 1999, ApJS, 124, 241

—. 2000, ApJ, 528, 913, source code obtained from `http://cococubed.asu.edu/code_pages/kap.shtml`

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

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

Weaver, T. A., Zimmerman, G. B., & Woosley, S. E. 1978, ApJ, 225, 1021

Woosley, S. E., & Taam, R. E. 1976, Nature, 263, 101

Woosley, S. E., & Weaver, T. A. 1984, in American Institute of Physics Conference Series, Vol. 115, American Institute of Physics Conference Series, ed. S. E. Woosley, 273–+

Zingale, M., Almgren, A. S., Bell, J. B., Nonaka, A., & Woosley, S. E. 2009, ApJ, 704, 196

Zingale, M. et al. 2002, ApJS, 143, 539

—. 2001, ApJS, 133, 195

APPENDIX A

THE EFFECTS OF CHEMICAL POTENTIAL ON THE CONVENIENCE PARAMETER α

In MAESTRO, things are done in an operator split fashion — the hydro is de-coupled from the burning. This means that during the hydro parts of the algorithm (where β_0 is used), the system is fixed in chemical equilibrium. For completeness, however, here I describe the effects of the species' chemical potentials, which were neglected in the original derivation of β_0. Note that similar terms appear in the calculation of things such as specific heats, which *may* be important in the burning step — there appears to be very little about this in the literature, but everyone seems to assume it makes little difference.

A.1 Derivation of The Convenience Parameter α

In Paper I, α is defined as

$$\alpha \equiv -\left(\frac{(1-\rho h_p)p_T - \rho c_p}{\rho^2 c_p p_\rho}\right) \tag{A.1}$$

where

$$h_p \equiv \left(\frac{\partial h}{\partial p}\right)_{T,X}, \quad c_p \equiv \left(\frac{\partial h}{\partial T}\right)_{p,X}, \quad p_T \equiv \left(\frac{\partial p}{\partial T}\right)_{\rho,X}, \quad p_\rho \equiv \left(\frac{\partial p}{\partial \rho}\right)_{T,X} \tag{A.2}$$

where the subscript X means holding all X_i constant. In the absence of reactions, the X subscript can be dropped from all derivatives, and with the use of the equation of state $p = p(\rho, T)$, α can be written as $\alpha = \alpha(\rho, T)$. Such a system without reactions and in thermal equilibrium could be either a pure system of one species, or a system of many species in chemical (and therefore *thermodynamic*) equilibrium. Cox & Giuli (1968) (hereafter CG) call the former type of system a "simple system" and therefore the latter a "non-simple system" in chemical equilibrium. The

analysis in Paper I that reduced Equation A.1 to

$$\alpha = \frac{1}{\Gamma_1 p_0} \tag{A.3}$$

used CG's discussion of the various adiabatic Γ's. However, their discussion only pertains to "simple systems" or "non-simple systems" in chemical equilibrium. In general, nuclear reactions will be important and therefore this analysis needs to be reformed.

Even in the presence of reactions, A.1 can be rewritten as was done in Paper I:

$$\alpha = -\frac{1}{p\chi_\rho c_p}\left[\left(\frac{1}{\rho\chi_\rho} - \frac{\rho e_\rho}{p\chi_\rho}\right)\frac{p\chi_T}{T} - c_p\right], \tag{A.4}$$

where

$$\chi_\rho \equiv \left(\frac{\partial \ln p}{\partial \ln \rho}\right)_{T,X}$$
$$\chi_T \equiv \left(\frac{\partial \ln p}{\partial \ln T}\right)_{\rho,X}.$$

Following the results of Paper I, we want to find a relation between $p\chi_\rho$ and Γ_1.

For an equation of state $p = p(\rho, T, X)$ we have

$$d\ln p = \left(\frac{\partial \ln p}{\partial \ln \rho}\right)_{T,X} d\ln\rho + \left(\frac{\partial \ln p}{\partial \ln T}\right)_{\rho,X} d\ln T + \sum_i \left(\frac{\partial \ln p}{\partial \ln X_i}\right)_{\rho,T,(X_j, j\neq i)} d\ln X_i. \tag{A.5}$$

For brevity, we will shorten the subscript "$(X_j, j \neq i)$" to just "(X_j)." We define another logarithmic derivative

$$\chi_{X_i} \equiv \left(\frac{\partial \ln p}{\partial \ln X_i}\right)_{\rho,T,(X_j)}$$

and therefore

$$d\ln p = \chi_\rho\ d\ln\rho + \chi_T\ d\ln T + \sum_i \chi_{X_i}\ d\ln X_i. \tag{A.6}$$

From here we get the general statement

$$\frac{\partial \ln p}{\partial \ln \rho} = \chi_\rho + \chi_T \frac{\partial \ln T}{\partial \ln \rho} + \sum_i \chi_{X_i} \frac{\partial \ln X_i}{\partial \ln \rho} \tag{A.7}$$

which must hold for an adiabatic process as well, and therefore we have

$$\Gamma_1 = \chi_\rho + \chi_T\left(\Gamma_3 - 1\right) + \sum_i \chi_{X_i}\Gamma_{4,i} \tag{A.8}$$

where we use CG's definition of Γ_1 and Γ_3 and introduce a fourth gamma function:

$$\Gamma_1 \equiv \left(\frac{\partial \ln p}{\partial \ln \rho}\right)_{\rm AD}, \quad \Gamma_3 - 1 \equiv \left(\frac{\partial \ln T}{\partial \ln \rho}\right)_{\rm AD}, \quad \Gamma_{4,i} \equiv \left(\frac{\partial \ln X_i}{\partial \ln \rho}\right)_{\rm AD}, \tag{A.9}$$

where the subscript AD means along an adiabat. We now derive an expression for Γ_3.

The first law of thermodynamics can be written as

$$dQ = dE + pdV - \sum_i \mu_i dN_i \tag{A.10}$$

where $\mu_i = \left(\frac{\partial E}{\partial N_i}\right)_{\rm AD,\rho,(N_j, j\neq i)}$ is the chemical potential; or per unit mass we have

$$\begin{aligned} dq &= de - \frac{p}{\rho^2}d\rho - \sum_i \mu_i d\left(\frac{n_i}{\rho}\right) \\ &= de - \frac{p}{\rho^2}d\rho - \sum_i \left(\frac{\partial e}{\partial X_i}\right)_{\rho,{\rm AD},(X_j)} dX_i \end{aligned}$$

where we have used $X_i \equiv \rho_i/\rho = A_i n_i/\rho N_{\rm A}$ and the chemical potential has been replaced with $\mu_i = (A_i/N_{\rm A})\left(\partial e/\partial X_i\right)_{\rho,{\rm AD},(X_j)}$. Using this and expressing the specific internal energy as $e = e(\rho, T, X)$ we then have

$$dq = c_v dT + \left[\left(\frac{\partial e}{\partial \rho}\right)_{T,X} - \frac{p}{\rho^2}\right] d\rho + \sum_i \left[\left(\frac{\partial e}{\partial X_i}\right)_{\rho,T,(X_j)} - \left(\frac{\partial e}{\partial X_i}\right)_{\rho,{\rm AD},(X_j)}\right] dX_i \tag{A.11}$$

and

$$\begin{aligned} \left(\frac{d\ln T}{d\ln \rho}\right)_{\rm AD} \equiv \Gamma_3 - 1 = \frac{1}{c_v T}\Bigg[\frac{p}{\rho} - \left(\frac{\partial e}{\partial \ln \rho}\right)_{T,X} + \\ \sum_i \left[\left(\frac{\partial e}{\partial X_i}\right)_{\rho,{\rm AD},(X_j)} - \left(\frac{\partial e}{\partial X_i}\right)_{\rho,T,(X_j)}\right] X_i \Gamma_{4,i}\Bigg] \end{aligned} \tag{A.12}$$

Now we need to evaluate $\left(\partial e/\partial \ln \rho\right)_{T,X}$. Again using the first law and the fact that $ds = dq/T$

is an exact differential (i.e. mixed derivatives are equal) we have

$$\left(\frac{\partial}{\partial\rho}\left[\frac{c_v}{T}\right]\right)_{T,X} = \left(\frac{\partial}{\partial T}\left[\frac{1}{T}\left(\frac{\partial e}{\partial\rho}\right)_{T,X} - \frac{p}{T\rho^2}\right]\right)_{\rho,X}$$

$$\frac{1}{T}\left(\frac{\partial}{\partial\rho}\left(\frac{\partial e}{\partial T}\right)_{\rho,X}\right)_{T,X} = -\frac{1}{T^2}\left(\frac{\partial e}{\partial\rho}\right)_{T,X} + \frac{1}{T}\left(\frac{\partial}{\partial T}\left(\frac{\partial e}{\partial\rho}\right)_{T,X}\right)_{\rho,X} + \frac{p}{T^2\rho^2} - \frac{1}{T\rho^2}\left(\frac{\partial p}{\partial T}\right)_{\rho,X}$$

$$\therefore \quad \left(\frac{\partial e}{\partial \ln\rho}\right)_{T,X} = \frac{p}{\rho}\left(1-\chi_T\right), \tag{A.13}$$

exactly the same result if we were to exclude species information. Simlarly, we can find an expression for the derivative of energy with respect to composition

$$\left(\frac{\partial}{\partial X_i}\left[\frac{c_v}{T}\right]\right)_{\rho,T,(X_j)} = \left(\frac{\partial}{\partial T}\left[\frac{1}{T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,T,(X_j)} - \frac{1}{T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right]\right)_{\rho,X}$$

$$\frac{1}{T}\left(\frac{\partial}{\partial X_i}\left(\frac{\partial e}{\partial T}\right)_{\rho,X}\right)_{\rho,T,(X_j)} = \frac{1}{T^2}\left[\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)} - \left(\frac{\partial e}{\partial X_i}\right)_{\rho,T,(X_j)}\right] + \frac{1}{T}\left[\left(\frac{\partial}{\partial T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,T,(X_j)}\right)_{\rho,X} - \left(\frac{\partial}{\partial T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X}\right]$$

$$\therefore \quad \left(\frac{\partial e}{\partial X_i}\right)_{\rho,T,(X_j)} = \left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)} - \left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X}.$$

Plugging these back into A.12 we have

$$\Gamma_3 - 1 = \frac{1}{c_v T}\left[\frac{p}{\rho}\chi_T + \sum_i \left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i\Gamma_{4,i}\right], \tag{A.14}$$

or

$$c_v = \frac{1}{T(\Gamma_3-1)}\left[\frac{p}{\rho}\chi_T + \sum_i \left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i\Gamma_{4,i}\right]. \tag{A.15}$$

We can obtain an expression for the specific heat at constant pressure from the enthalpy

$$\begin{aligned} c_p \equiv \left(\frac{\partial h}{\partial T}\right)_{p,X} &= \left(\frac{\partial e}{\partial T}\right)_{p,X} - \frac{p}{\rho^2}\left(\frac{\partial \rho}{\partial T}\right)_{p,X} \\ &= \left(\frac{\partial e}{\partial T}\right)_{p,X} + \frac{p}{\rho^2}\left(\frac{\partial p}{\partial T}\right)_{\rho,X}\left(\frac{\partial \rho}{\partial p}\right)_{T,X} \\ &= \left(\frac{\partial e}{\partial T}\right)_{p,X} + \frac{p}{\rho T}\frac{\chi_t}{\chi_\rho}. \end{aligned}$$

The first term on the right-hand side can be obtained from writing $e = e(p,T,X)$ and $p = p(\rho,T,X)$:

$$\begin{aligned} de &= \left(\frac{\partial e}{\partial p}\right)_{T,X} dp + \left(\frac{\partial e}{\partial T}\right)_{p,X} dT + \sum_i \left(\frac{\partial e}{\partial X_i}\right)_{p,T,(X_j)} dX_i \\ dp &= \left(\frac{\partial p}{\partial \rho}\right)_{T,X} d\rho + \left(\frac{\partial p}{\partial T}\right)_{\rho,X} dT + \sum_i \left(\frac{\partial p}{\partial X_i}\right)_{\rho,T,(X_j)} dX_i \\ \therefore \left(\frac{\partial e}{\partial T}\right)_{\rho,X} &= \left(\frac{\partial e}{\partial p}\right)_{T,X}\left(\frac{\partial p}{\partial T}\right)_{\rho,X} + \left(\frac{\partial e}{\partial T}\right)_{p,X} \\ \Rightarrow \left(\frac{\partial e}{\partial T}\right)_{p,X} &= c_v - \left(\frac{\partial e}{\partial \rho}\right)_{T,X}\left(\frac{\partial \rho}{\partial p}\right)_{T,X}\left(\frac{\partial p}{\partial T}\right)_{\rho,X} \\ &= c_v - \frac{p\chi_T}{\rho T \chi_\rho}(1-\chi_T) \end{aligned}$$

and

$$c_p = \frac{p}{\rho T}\frac{\chi_T^2}{\chi_\rho} + c_v \tag{A.16}$$

Dividing this by A.15 and using the relation between the Γ's, A.8, we then have

$$\begin{aligned}
\gamma \equiv \frac{c_p}{c_v} &= 1 + \frac{p(\Gamma_3 - 1)}{\rho}\frac{\chi_T^2}{\chi_\rho}\left[\frac{p}{\rho}\chi_T + \sum_i \left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i\Gamma_{4,i}\right]^{-1} \\
&= 1 + \frac{p\chi_T\left(\Gamma_1 - \chi_\rho - \sum_i \chi_{X_i}\Gamma_{4,i}\right)}{p\chi_\rho\chi_T + \rho\chi_\rho\sum_i \left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i\Gamma_{4,i}} \\
&= \frac{p\chi_T\Gamma_1 + \sum_i \left[\rho\chi_\rho \left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i - p\chi_T\chi_{X_i}\right]\Gamma_{4,i}}{p\chi_\rho\chi_T + \rho\chi_\rho\sum_i \left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i\Gamma_{4,i}} \\
\Rightarrow p\chi_\rho &= \frac{1}{\chi_T\gamma}\left[p\chi_T\Gamma_1 + \sum_i \left[\rho\chi_\rho(1-\gamma)\left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i - p\chi_T\chi_{X_i}\right]\Gamma_{4,i}\right].
\end{aligned} \tag{A.17}$$

Plugging A.17 into A.4 and rewriting the partial derivative of e with the help of A.13 we have

$$\begin{aligned}
\alpha &= -\frac{1}{p\chi_\rho c_p}\left[\left(\frac{1}{\rho\chi_\rho} - \frac{\rho e_\rho}{p\chi_\rho}\right)\frac{p\chi_T}{T} - c_p\right] \\
&= \frac{\gamma}{c_p}\frac{c_p\chi_T + \left(\rho\left(\frac{\partial e}{\partial \ln\rho}\right)_{T,X} - p\right)\frac{\chi_T^2}{T\rho\chi_\rho}}{p\chi_T\Gamma_1 + \sum_i\left[\rho\chi_\rho(1-\gamma)\left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i - p\chi_T\chi_{X_i}\right]\Gamma_{4,i}} \\
&= \frac{\gamma}{\Gamma_1 p c_p}\left[\frac{c_p - \frac{p\chi_T^2}{T\rho\chi_\rho}}{1 + \sum_i\left[\frac{\rho\chi_\rho}{p\chi_T}(1-\gamma)\left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i - \chi_{X_i}\right]\frac{\Gamma_{4,i}}{\Gamma_1}}\right] \\
&= \left(\frac{1}{\Gamma_1 p}\right)\left[1 + \sum_i\left[\frac{\rho\chi_\rho}{p\chi_T}(1-\gamma)\left(\frac{\partial}{\partial \ln T}\left(\frac{\partial e}{\partial X_i}\right)_{\rho,\mathrm{AD},(X_j)}\right)_{\rho,X} X_i - \chi_{X_i}\right]\frac{\Gamma_{4,i}}{\Gamma_1}\right]^{-1}
\end{aligned}$$

$$\boxed{\alpha = \frac{1}{\Gamma_1 p}\left[1+\sum_i\left[\frac{\rho^2 p_\rho}{p p_T}(1-\gamma)\frac{N_\mathrm{A}}{A_i}\left(\frac{\partial \mu_i}{\partial T}\right)_{\rho,X} X_i - \chi_{X_i}\right]\frac{\Gamma_{4,i}}{\Gamma_1}\right]^{-1}} \tag{A.18}$$

A.2 Recalling The Derivation of Density-like Quantity β_0

Recall from Paper I that β_0 was derived from the equation

$$\nabla \cdot \mathbf{U} + \alpha \mathbf{U} \cdot \nabla p_0 = \tilde{S} \tag{A.19}$$

in such a fashion that we ended up with an equation of the form

$$\nabla \cdot (\beta_0(r)\mathbf{U}) = \beta_0 \tilde{S}. \tag{A.20}$$

The derivation in Appendix B of Paper I for a β_0 that satisfies A.20 automatically assumed $\alpha = (\Gamma_{1_0} p_0)^{-1}$. This would have to be modified with the above derivation of α to be correct in a non-operator split fashion.

APPENDIX B

THERMAL DIFFUSION TEST PROBLEM

This problem is designed to test the accuracy of our implementation of an implicit solver for the diffusion of a two-dimensional Gaussian enthalpy pulse. That is, we are only concerned with the diffusive term in (3.6):

$$\frac{\partial(\rho h)}{\partial t} = \nabla \cdot (k_{\text{th}} \nabla T)\,. \tag{B.1}$$

To easily compare with an analytic solution (see, for example Swesty & Myra (2009) for an analogous example for a radiation-hydrodynamics code) we assume the thermal conductivity to be constant: $k_{\text{th}} = 10^7 \text{ergKcm}^{-1}\text{s}^{-1}$. We also assume an ideal gas with $X(\text{He}^4) = 0.5, X(\text{C}^{12}) = X(\text{Fe}^{56}) = 0.25$ and ratio of specific heats $\gamma = 5/3$. Furthermore, we are not concerned with any hydrodynamic motions so we keep the density fixed. We can then express (B.1) in a simpler form:

$$\frac{\partial h}{\partial t} = D\nabla^2 h, \tag{B.2}$$

where $D = k_{\text{th}}/(\rho c_p)$ is the diffusion coefficient.

Given the initial conditions for the two-dimensional pulse,

$$h(\mathbf{r}, t = t_0) = (h_{\text{p}} - h_0) \times \exp\left(\frac{-|\mathbf{r} - \mathbf{r_0}|^2}{4Dt_0}\right) + h_0, \tag{B.3}$$

where h_p, h_0, $\mathbf{r_0} = (x_0, y_0)$, and t_0 are the peak enthalpy, ambient enthalpy, location of the center of the peak, and time from which the system has evolved respectively, the analytic solution takes on the form

$$h(\mathbf{r}, t) = \left(h_{\text{p}} - h_0\right)\left(\frac{t_0}{t + t_0}\right)\exp\left(\frac{-|\mathbf{r} - \mathbf{r_0}|^2}{4D\,(t + t_0)}\right) + h_0, \tag{B.4}$$

where t is the evolved time.

We solve this problem on a Cartesian grid of size 4 cm $\times$ 4 cm with the following parameters: $h_{\text{p}} = 10.0\text{ergg}^{-1}$, $h_0 = 1.0\text{ergg}^{-1}$, $\mathbf{r_0} = (2.0\text{cm}, 2.0\text{cm})$, $t_0 = 0.1\text{s}$, and $\rho = 1.0\text{gcm}^{-3}$. For the

density and composition used in this test, we obtain a diffusion coefficient of $D = 0.32\text{cm}^2\text{s}^{-1}$. Our implicit solve uses a Crank-Nicholson scheme that is second order accurate in space and time.

Figure B.1 shows an example of the initial enthalpy pulse and its evolution through $t = 0.4$ s on a 1024×1024 grid with fixed time step $\Delta t = 10^{-3}$ s. Note that as the pulse expands it begins to interact with the edges of the computational domain and the symmetry of the Gaussian peak is broken. Figure B.2 shows the computed average enthalpy as a function of radius (Xs) compared to the analytic solution (lines) for the same test problem shown in Figure B.1. Again, excepting boundary effects, the numerical and analytic solutions are well matched.

To check the convergence of the algorithm we ran simulations with various resolutions and compared the errors. To measure the error in the simulation, we use the L_1 norm of the difference between the analytic and numeric solutions normalized to the L_1 norm of the analytic solution, which we define as ε:

$$\varepsilon^m \equiv \frac{||h(r,t^m) - h^m||_{L_1}}{||h(r,t^m)||_{L_1}} = \frac{\sum_{i,j} \left| h(r_{i,j},t^m) - h^m_{i,j} \right|}{\sum_{i,j} |h(r_{i,j},t^m)|}, \tag{B.5}$$

where $h(r_{i,j},t^m)$ is the analytic solution at $r_{i,j} = \left((x_i - x_0)^2 + (y_j - y_0)\right)^{1/2}$ and time $t^m = m\Delta t$ and $h^m_{i,j}$ is the numeric solution at (x_i, y_j) and time t^m. We further define the convergence rate, α, by comparing the value of ε at the current resolution to the value of ε at a finer resolution simulation:

$$\alpha \equiv \log_2 \left(\frac{\varepsilon}{[\varepsilon]_{\text{finer}}} \right). \tag{B.6}$$

For our comparisons, we take "finer" to mean a simulation with twice the resolution; to compare the simulations at the same physical time, the finer simulation must have evolved through twice the number of time steps as the coarser simulation. If our algorithm truly is second-order accurate in space and time then α should equal 2. Table B.1 shows the values of ε and the convergence rate for various resolutions at $t = 0.08$ s; for α, the norm in the current column is compared to the norm of the finer resolution simulation in the column to its right. Our values of α agree very well with the expected value.

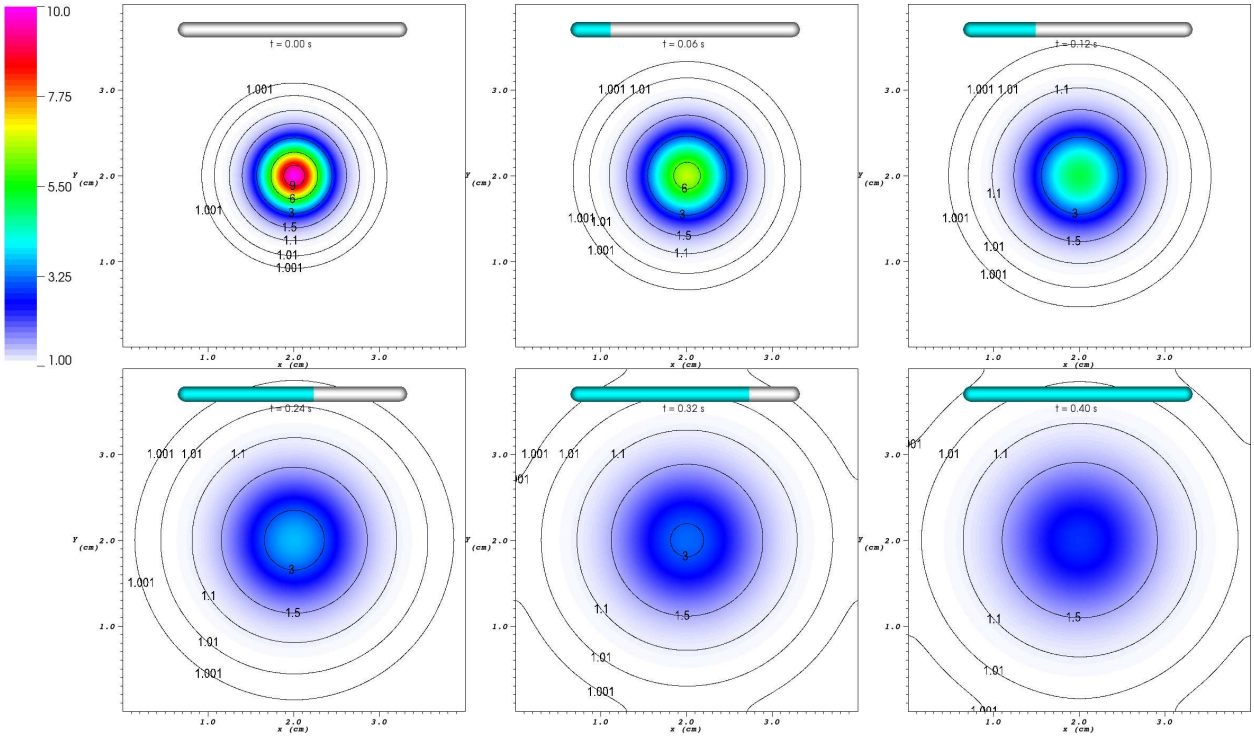

Figure B.1: Time evolution of the diffusion of a two-dimensional Gaussian pulse of enthalpy as described in the text. The value of time displayed is the evolution time, t. This simulation was run with a 1024×1024 grid with time step size $\Delta t = 0.001$ s. Excepting edge effects near the domain boundary, the numerical solution maintains its axisymmetric form about the center of the pulse at $(x, y) = (2.0, 2.0)$.

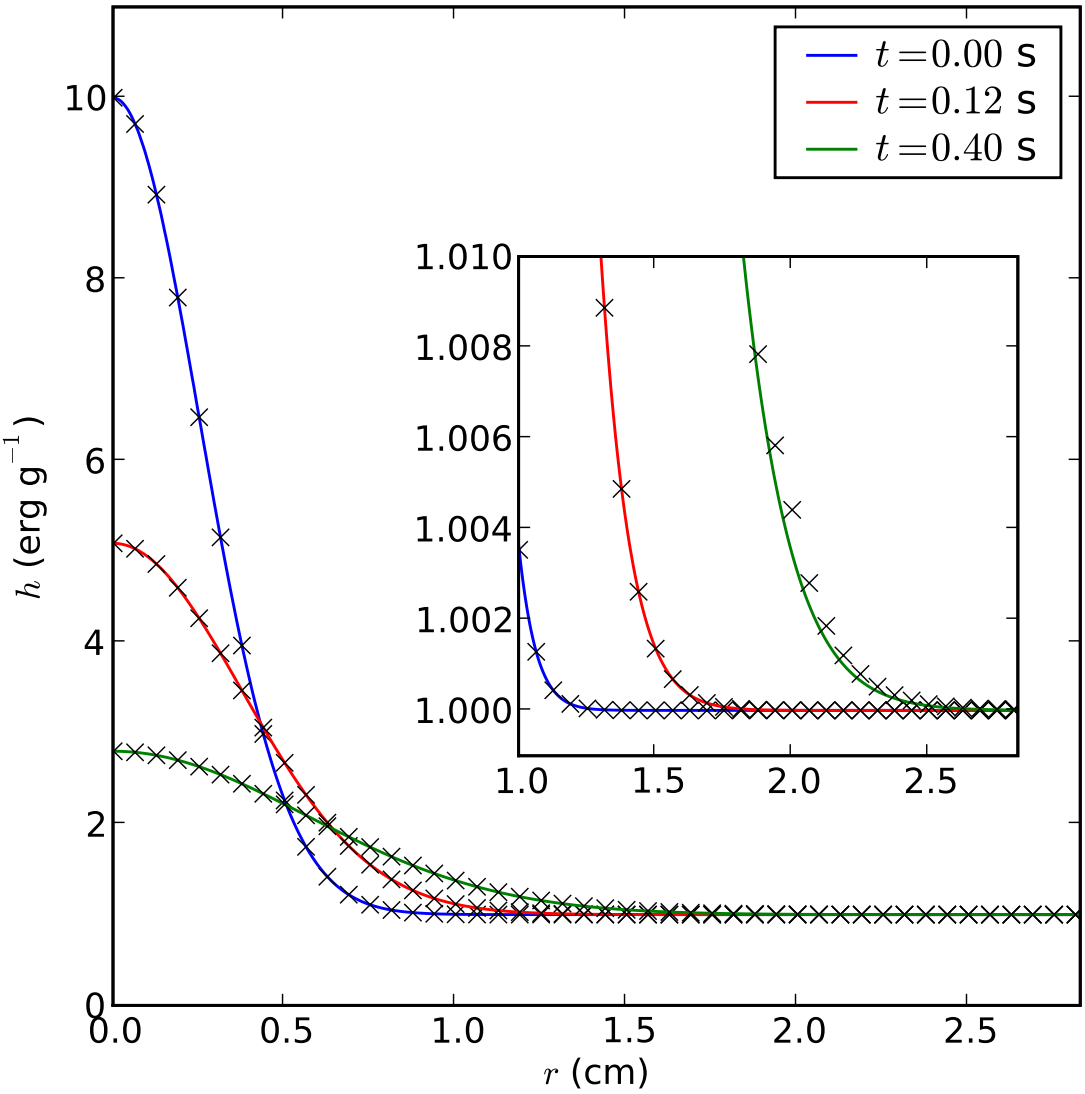

Figure B.2: The average of enthalpy as a function of radius from the center, $(x,y) = (2.0,2.0)$, of a two-dimensional Gaussian pulse. The X's are data from the numerical solution at the shown times. The lines represent the analytic solutions as given by (B.4). The numerical solution tracks the analytic solution very well except when the pulse has diffused enough that it begins to interact with the boundaries of the computational domain as seen in the inset plot.

Table B.1: Reduced L_1 norms and convergence rate for the diffusion test problem at $t = 0.08$ s.

	128×128 Error $\Delta t = 0.008$ s	256×256 Error $\Delta t = 0.004$ s	512×512 Error $\Delta t = 0.002$ s	1024×1024 Error $\Delta t = 0.001$ s
ε	8.64×10^{-5}	2.16×10^{-5}	5.39×10^{-6}	1.35×10^{-6}
α	2.0012	1.9999	1.9988	—