Out-of-core Hydrodynamic Simulations of Cosmological Structure Formation

by

Hy Trac

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Astronomy and Astrophysics
University of Toronto

Abstract

Out-of-core Hydrodynamic Simulations of Cosmological Structure Formation

Hy Trac
Doctor of Philosophy
Graduate Department of Astronomy and Astrophysics
University of Toronto
2004

Astrophysical and cosmological structure formation are challenging problems because they involve dynamical and hydrodynamical processes that can span a large range in scale, mass, and energy. Hydrodynamic and N-body simulations are powerful tools with which to solve the nonlinear physics, and their continuing development and application is the focus of this thesis.

I present a new approach to Eulerian computational fluid dynamics that is designed to work at high Mach numbers encountered in astrophysical simulations. The Eulerian conservation equations are solved in an adaptive frame moving with the fluid where Mach numbers are minimized. The Moving Frame code separately tracks local and bulk flow components, allowing thermodynamic variables to be accurately calculated in both subsonic and supersonic fluid.

An out-of-core hydrodynamic code has been developed for high resolution cosmological simulations. Out-of-core computation refers to the technique of using disk space as virtual memory and transferring data in and out of main memory at high I/O bandwidth. The code is based on a two-level mesh scheme where short-range physics is solved on a high-resolution, localized mesh while long-range physics is captured on a lower resolution, global mesh.

This thesis includes the first astrophysical application of Eulerian hydrodynamic simulations to model the formation of blue stragglers through stellar mergers. The off-axis collision of equal mass stars produces a single merger remnant. The merger of $n = 3$ polytropes results in substantial chemical mixing throughout the remnant, while the merger of realistic $M = 0.8 M_\odot$ main sequence stars produces significant mixing only outside of the core.

The Out-of-core Hydro code is applied to running the largest Eulerian hydrodynamic simulation to date for studying the thermal history of the high redshift $3 \leq z \leq 7$ intergalactic medium. The

temperature-density and gas-dark matter density relations, as well as the scatter in these relations, are robustly quantified. Reionization and shock heating are observed to influence the temperature of the photoionized gas.

Furthermore, a parallel particle-mesh N-body code is applied to simulating the clustering of dark matter halos. The PMFAST simulations show that that several bias parameters are consistent with being scale-invariant, a useful property for doing cosmology with galaxy clustering.

For my parents, Xuan Thi Hua and Sung Van Trac, who started with just $70 CAN *and worked so hard and sacrificed so much in order that their two children would have all the opportunities in life.*

Acknowledgements

This thesis signals the completion of the Ph.D. program and of Grade 21, here at the University of Toronto and the Canadian Institute for Theoretical Astrophysics. I owe my deepest gratitude to the many people who have contributed to making this scientific endeavour meaningful and memorable.

First and foremost, I would like to thank my supervisor, Ue-Li Pen, for his enthusiastic support and guidance, and for generously providing his time and resources. Working with you has been inspirational, rewarding, and an absolute pleasure.

I would like to thank Ray Carlberg and Bob Abraham for their invaluable support throughout the Ph.D. program. I am grateful to Hugh Couchman, Marten van Kerkwijk, and Peter Martin for the constructive critique of this thesis. I am indebted to Chris Loken for all his painstaking help with computing, to Howard Yee for his help in the M.Sc. program, and to Dick Bond for his support of my postdoctoral applications. I have learned and benefited much from the scientific collaborations with Pengjie Zhang, Hugh Merz, Alison Sills, and Uros Seljak. In addition, the many discussions with faculty and postdocs here at UofT and CITA have been both valuable and enjoyable.

The past five years of graduate studies and life in Toronto have been made exceptionally wonderful and memorable by my fellow graduate students. I will always remember sharing laughs, enjoying delicious meals and drinks, watching movies, playing volleyball and poker, and all things hockey. I simply can't name everyone, but I must thank Mark Brodwin for so much, Kris Blindert for being an amazing office mate, and Kevin Blagrave for being a good buddy.

I send my love and thanks to my parents, Xuan Thi Hua and Sung Van Trac, my brother, My Trac, and my grandma, Lang Thi Ma, who first taught me to print and to tell time. I would like to thank everyone at Augustana for being such a blessing in my life. To my friends back in Vancouver, thank you for all the laughs and the good times. I can't name you all and I can't say enough about each one of you. And finally, I would like to dearly thank Elaine Wang for the wonderful five years we have shared and all the rest to come.

This work has been funded by the Natural Sciences and Engineering Research Council of Canada, the Ontario Graduate Scholarship, the Walter C. Sumner Foundation, and the University of Toronto.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 The Cosmological Paradigm

Cosmology is commonly stated in brief as the study of the origin, evolution, and fate of the universe. This scientific field is primarily concerned with the dynamics of the expanding universe and the formation of the large-scale structure. One of the big tasks in modern cosmology is to determine how cosmic structure forms from a primordial mix of radiation, baryons, and dark matter. As the universe expanded and cooled, the protons, neutrons, and electrons eventually combined to make hydrogen and helium, the two most common elements in the universe. With the electrons locked up, the radiation became decoupled from the matter and since then have travelled mostly unimpeded through the universe, to be detected as the cosmic microwave background (CMB) today. The currently held cosmological paradigm is based mostly on observations of the CMB and of the numerous galaxies that have traditionally shaped our picture of the cosmos.

Presently, the favoured cosmological model of structure formation is a flat ΛCDM universe where less than one-third of the energy content is from matter ($\Omega_m = 0.3$) and more than two-thirds from dark energy ($\Omega_\Lambda = 0.7$). The mass density of the universe comes mainly from two components, where approximately five-sixths is cold dark matter ($\Omega_{\rm cdm} = 0.25$) and one-sixth is baryons ($\Omega_b = 0.05$). While there are on the order of 10^9 photons for every baryon in the universe, the mass density of the cold ($T_0 = 2.7$ K) radiation today is negligible. The expansion of the universe is traditionally described by the Hubble constant $H_0 = 100h$ km/s/Mpc and the modern value (h=0.7) sets the age of the universe to be approximately 14 billion years. This concordance or best-fit model is motivated primarily by results from the CMB, observations of galaxy clustering, and measurements of the Hubble expansion using supernovae (see Spergel et al., 2003).

At a given time or redshift, the expansion of the universe is driven by the relative strengths of dark energy and dark matter, while the dynamics is governed by the gravitational evolution of the mass. In the generic picture of gravitational collapse, matter flows out of underdense regions and falls into overdense regions to form structures like halos, filaments, and sheets that make up the

cosmic web. On large scales, the baryons trace the underlying dark matter and the hydrodynamics is dictated by the gravity of the mass. On smaller scales, the collisional gas and the collisionless dark matter set off on different evolutionary paths. The collisionless dark matter continues to collapse, forming virialized dark matter halos with deep potential wells. The baryons are pulled into the potential wells, but the collapse is initially halted by gas pressure. However, this collisional component can dissipate energy through radiative cooling and eventually collapse to very high densities to form stars, which assemble themselves into galaxies.

Star formation, galaxy assembly, stellar feedback through supernova, and the formation of planetary systems are some of the major links in the evolution of the universe that we still need to fully understand. However, the stars and interstellar medium (ISM) in galaxies as well as observable intracluster gas in clusters of galaxies only account for on the order of 10% of the total baryon budget. The majority of the missing baryons are believed to reside in the intergalactic medium (IGM).

The fundamental question concerning the distributions of baryons and dark matter still needs to be accurately addressed in order to understand the larger cosmological picture and to study in detail the evolution of cosmological systems. Hydrodynamic and N-body simulations are powerful tools with which to solve the nonlinear physics, and their continuing development is needed for studying structure formation at a level of accuracy on par with upcoming precision observations.

This thesis is primarily based on the development of new numerical techniques for cosmological and astrophysical simulations, with some initial scientific applications which are in progress. It is concentrated on three research areas: the development of new hydrodynamic algorithms, the improvement of computational techniques, and the application to cosmology and astrophysics.

1.2 Computational Fluid Dynamics

Astrophysical and cosmological structure formation involve nonlinear gasdynamical processes that, in general, cannot be modelled analytically but require numerical methods. Modelling structure formation is a challenging problem because the nonlinear gasdynamical processes can span a large range in scale, mass, and energy. Furthermore, strong shocks, gravitational collapse, and radiative feedback can occur and play an important role in the evolution of the gas. The evolution of complex systems is best modelled using numerical simulations. Computational fluid dynamics (CFD) is a powerful approach to simulating the complex fluid flow in astrophysical and cosmological hydrodynamics. Continuing advancements in the field is making it tractable to robustly probe the nonlinear physics, with emphasis on high-resolution capturing of shocks, prevention of numerical instabilities, and having high dynamic range.

A large class of astrophysical problems involve collisional systems where the mean free path is much smaller than all length scales of interest. Hence, one can appropriately adopt an ideal

fluid description of matter where the thermodynamical properties of the fluid obey well known equations of state. Conservation of mass, momentum, and energy allows one to write down the Euler equations that govern fluid mechanics (see Landau & Lifshitz, 1987). This formalism is a standard basis for simulating astrophysical fluids. Both Eulerian and Lagrangian methods have been developed to solve the temporal and spatial evolution of the fluid equations. While both of these methods share the same approach where time is discretized into discrete time steps, they differ in many other respects.

Lagrangian methods based on smoothed particle hydrodynamics (SPH; Gingold & Monahan, 1977; Lucy, 1977) consider a Monte-Carlo approximation to solving the fluid equations, somewhat analogous to N-body methods for the Vlasov equation. SPH schemes follow the trajectories of particles of fixed mass that represent fluid elements. The Lagrangian forms of the Euler equations are solved to determine smoothed fluid variables like density, velocity, and temperature. The particle formulation does not naturally capture shocks and artificial viscosity is added to prevent unphysical oscillations. However, the addition of artificial viscosity broadens shocks over several smoothing lengths and degrades the resolution. The Lagrangian approach has a large dynamic range in length but not in mass. It achieves good spatial resolution in high density regions but does poorly in low density regions. SPH schemes must smooth over a large number of neighbouring particles, making it computationally expensive and challenging to implement in parallel.

The standard approach to Eulerian methods is to discretize space into finite volume or cells. In the simplest case, the Euler equations are solved on a Cartesian grid by computing the flux of mass, momentum, and energy across grid cell boundaries. In conservative schemes, the flux taken out of one cell is added to the neighbouring cell and this ensures the correct shock propagation since mass, momentum, and energy are strictly conserved. Linear, first-order flux assignment schemes are, in general, diffusive and dispersive and much effort has been devoted to developing nonlinear, higher-order schemes that can more accurately capture gasdynamics. In particular, flux assignment schemes based on the total variation diminishing condition (TVD) condition (Harten, 1983) have been demonstrated to robustly provide high-resolution capturing of shocks. The TVD condition is a nonlinear stability condition that prevents the formation of unphysical oscillations and TVD schemes are second-order accurate and flux-conservative. The Eulerian approach has a large dynamic range in mass but not in length, opposite to that of Lagrangian schemes. In general, Eulerian algorithms are computationally faster by several orders of magnitude. They are also easy to implement and to parallelize.

This thesis is primarily concentrated on Eulerian CFD. Traditionally, it is recognized for its superior shock-capturing ability. In addition, Eulerian algorithms are computationally very fast and memory-friendly, allowing one to optimally use computing resources. These features make it an attractive method for simulating structure formation. However, there are current issues with standard Eulerian algorithms that need to be addressed in order to realize the full potential

of Eulerian simulations. In this thesis, the development of novel hydrodynamic algorithms and computational techniques is geared towards solving some of the known problems.

1.2.1 High Mach Number Hydrodynamics

One of the main challenges in simulating complex fluid flow is the capturing of shocks, which frequently occur and play an important role in gasdynamics. In particular, the capturing of shocks in the presence of high Mach numbers has proven to be exceptionally difficult. For example, in cosmological simulations of the intergalactic medium (IGM), there are large-scale velocity fields on the order of 1000 km/s and the typical sound speed in these bulk flows is ~ 10 km/s. At Mach numbers ~ 100, the ratio of the thermal energy to the kinetic energy is $\sim 10^{-4}$. In Eulerian CFD, the standard practice is to calculate the thermal energy by subtracting the kinetic energy from the total energy, but this calculation can be inaccurate, especially near shocks. The physical solutions for shock waves are discontinuous, making approximations near the shock fronts only first-order accurate. Thus, the cancellation error of subtracting the kinetic energy from the total energy will typically result in significant errors. The poor tracking of the thermal energy in supersonic regions is known as the high Mach number problem in Eulerian CFD. Lagrangian approaches have the advantage of being able to directly track the thermal energy and temperature, but SPH algorithms cannot accurately capture these relatively weak shocks.

In this thesis, I present a new approach that offers a solution to the high Mach number problem of Eulerian CFD. I have developed a new hybrid algorithm where the fluid conservation equations are solved on a grid, like in Eulerian schemes, but the reference frame is chosen to adaptively follow the fluid, like in Lagrangian schemes. The moving frame algorithm possesses the high-resolution shock-capturing ability of grid-based schemes, but it does not suffer from the high Mach number problem since local fluid variables can be directly tracked. I also present a new cosmological Hydro&N-body code where the moving frame algorithm is coupled to a particle-mesh (PM) N-body algorithm to simultaneously probe the interaction between the baryonic gas and the collisionless dark matter.

1.2.2 Out-of-core Computation

Numerical simulations must converge over a large range in scale, mass, and temperature. Large box sizes are required to capture large-scale correlations and power while high resolution is needed to resolve small-scale, nonlinear structures. Most simulations to date have sacrificed one for the other because of the limitations in computational resources. While parallelized numerical codes and optimized mathematical libraries have significantly reduced computation time on high performance computing (HPC) systems, memory limitations remain the bottleneck in expanding the size of simulations.

Conventional memory is an expensive commodity and at present, HPC systems typically have tens of gigabytes and some may have a few hundred gigabytes at most. However, disk space is relatively cheap and disk arrays can provide many more orders of magnitude in storage. In principle, very large simulations can be run by doing out-of-core computation, where disk space is used as virtual memory and data is transferred between disk and main memory at high I/O bandwidth. In order to be effective, out-of-core computation must avoid being I/O limited. SCSI disks can be arranged in striped chains and connected through high-bandwidth SCSI controllers to maximize the I/O performance. Striped SCSI disk arrays are capable of delivering on the order of $100 - 1000$ MB/s throughput and this is sufficient bandwidth to do out-of-core computation. However, the slow latency presents a major problem for codes requiring non-sequential disk access. Pioneering work has been done for computational astrophysics. Salmon & Warren (1997) developed a parallel, out-of-core, tree N-body code to run an 80 million particle simulation on a distributed-memory Pentium cluster with a total of 16 processors, 2 GB of memory, and 16 GB of disk space. The out-of-core paradigm remains a niche that has been relatively unexplored for numerical simulations, but advances in algorithmic development now make it feasible to do out-of-core computation effectively.

In this thesis, I present a new out-of-core hydro (OCH) code for high resolution cosmological simulations. The OCH code is developed based on the moving frame hydro and PM N-body algorithms for high-resolution capturing of shocks and high mass resolution at all scales.

1.3 Cosmological and Astrophysical Applications

This thesis contains three initial scientific applications which are under development: the thermal evolution of the high redshift IGM, the mass-dependent clustering and biasing of dark matter halos in a ΛCDM universe, and the astrophysical formation of blue straggler stars through stellar mergers. The novel numerical tools developed in this thesis provide the technique, accuracy, and resolution necessary for conducting these important studies of cosmological and astrophysical structure formation.

1.3.1 The Thermal Evolution of the High Redshift IGM

Probing the evolution of the intergalactic medium (IGM) is now a big endeavour in cosmology. At high redshifts, the IGM bridges the gap between reionization and galaxy formation, two very important phases in the history of the universe. Structure formation is primarily driven by the gravitational collapse of the dark matter and the hydrodynamic evolution of the baryons. However, reionization is expected to reset the initial conditions for the baryons in the IGM since this injection of thermal energy blows gas out of the mini-halos and generates an entropy floor that may limit further baryonic collapse and star formation. Thus, in the thermal history of the high redshift IGM, we can in principle discover the imprints of the first stage of star formation and of reionization and

we can also determine the initial conditions for galaxy formation.

Presently, the two main constraints on reionization from observations of a Gunn-Peterson trough in a SDSS quasar (Becker et al., 2001) and the excess correlation between the temperature and polarization in WMAP observations of the CMB (Kogut et al., 2003) demonstrate that reionization may be considerably complex. While the SDSS observations place reionization between redshifts $6 < z_{\rm reion} < 7$ and the WMAP measurements of the high optical depth $\tau \approx 0.17$ argue for a reionization redshift $z_{\rm reion} \sim 17$, they are not necessarily in conflict. CMB polarization comes from electron scattering and this measurement is more sensitive to the start of reionization because the electron number density is greater at higher redshifts. On the other hand, the photoionization or absorption cross-section is approximately 5 orders of magnitude larger than the electron scattering cross-section and an ionized IGM with a neutral fraction of roughly $10^{-3} - 10^{-4}$ can still be opaque. Therefore, the Gunn-Peterson trough signals the end of the reionization epoch.

While reionization may be a complex process, it does make the physics of the IGM more linear again and tractable to solve. It leaves an abundance of free electrons and residual neutral hydrogen which can be used to probe baryonic structure in the IGM through the Sunyaev-Zeldovich effect (Sunyaev & Zeldovich, 1980) and the Lyman alpha forest (see Rauch, 1998, for a review), respectively. These tracers can be used to probe the thermal history of the IGM and reionization itself. The Lyman alpha (Lyα) forest, in particular, is a promising tool for studying structure formation in the IGM.

The Lyα forest probes the distribution of residual neutral hydrogen in the IGM, seen through absorption lines in high redshift quasar spectra. High resolution Keck spectra and the large sample of quasar spectra from the SDSS have provided an abundance of high signal-to-noise data. These observations are complimented by theoretical modelling using hydrodynamic simulations (e.g. Cen et al., 1994; Zhang et al., 1995; Hernquist et al., 1996; Miralda-Escude et al., 1996; Theuns et al., 1998). The Lyα absorption is found to arise from a fluctuating IGM naturally experiencing gravitational instability and this picture has been coined the "fluctuating Gunn-Peterson effect" (Rauch et al., 1997). The forest is largely contributed to by photoionized gas with densities slightly higher than the mean baryonic density and temperatures of $\sim 10^4$ K. The physics of the gas is only mildly nonlinear and correlations between gas properties like density and temperature, as well as correlations with the underlying dark matter can in principle be quantified and used to study evolution and to do cosmology. However, this promising tool is currently limited by our understanding of how the high redshift IGM evolves.

In this thesis, the out-of-core hydro code is applied to running the largest Eulerian hydrodynamic simulation to date for studying the thermal evolution of the high redshift $3 \leq z \leq 7$ IGM, in both an adiabatic universe and a nonadiabatic one. In the adiabatic scenario, adiabatic cooling from the expansion of the universe and shock heating from gravitational collapse together determine the temperature of the IGM. In the nonadiabatic case, reionization is prescribed and UV heating and

atomic cooling are calculated to provide a more realistic model of the universe.

1.3.2 The Clustering and Biasing of Dark Matter Halos

Traditionally, galaxies have dominated our picture of the universe and have been used to map out the large-scale structure and to study the cosmological model. Modern galaxy redshift surveys such as the Sloan Digital Sky Survey (SDSS) and 2 Degree Field Galaxy Redshift Survey (2dFGRS), as well as upcoming ones like the CFHT Legacy Survey (CFHTLS) have tremendous statistical power, allowing them to be used in conjunction with CMB observations to break degeneracies in cosmological parameters of structure formation. Specifically, the measurement of galaxy clustering can be used to precisely determine the matter power spectrum. However, one inherent difficulty in this approach is that galaxies are indirect tracers of the underlying dark matter, and the distribution of light relative to mass is known to be biased.

Galaxy biasing is evident for a number of reasons. The biasing of density peaks in a Gaussian random field suggests that dark matter halos and galaxies are formed more strongly clustered than the underlying mass distribution (e.g. Kaiser, 1984; Bardeen et al., 1986). Observations show that galaxies of different types cluster differently (e.g. Strauss & Willick, 1995; Tegmark & Bromley, 1999) and therefore, not all galaxies can fairly trace the mass. Furthermore, cosmological simulations of galaxy formation show biasing with significant segregation among galaxies of different mass (e.g. Cen & Ostriker, 1992; Blanton et al., 2000).

In modern galaxy surveys, the galaxy power spectrum has been measured to unprecedented precision and the bias correction can be determined up to a multiplicative factor. Therefore, from just galaxy clustering alone, the shape of the matter power spectrum can be well constrained, though the amplitude is unknown. However, these measurements rest on several assumptions about the biasing relation which need to be tested rigourously in order for galaxy clustering to be used as a precision tool for doing cosmology.

Presently, large-scale hydrodynamic simulations of galaxy formation are not possible because very high dynamic range is required to capture nonlinear, small-scale physics in large volumes of the universe. A more feasible approach is to use N-body simulations to model the hierarchical formation of dark matter halos and to identify these halos as sites where galaxies may form. In principle, mass-dependent halo biasing can be translated into luminosity-dependent galaxy biasing, allowing one to use the measurement of galaxy clustering to extract both the shape and amplitude of the matter power spectrum. In this thesis, a new parallel particle-mesh N-body code named PMFAST (Merz, Pen, & Trac, 2004) is applied to simulating the mass-dependent and biasing of dark matter halos in a ΛCDM universe.

1.3.3 The Formation of Blue Stragglers through Stellar Mergers

The stellar density in the cores of globular and open clusters is high enough for stellar collisions to take place with significant frequency. These collisions can modify the stellar population by altering or destroying the interacting stars or by creating new merger remnants. In addition, these collisions can affect the cluster dynamics since mass loss changes the gravitational potential of the system. Stellar evolution and stellar dynamics in clusters are intertwined and it is necessary to understand their influence on each other in order to properly study globular and open clusters (Bailyn, 1995). In particular, it is important to study stellar collisions for the resulting collision products can be used as direct tracers of stellar interactions.

Current observations and simulations suggest that the merger of two main sequence stars near the turnoff can produce a blue straggler (e.g. Sills et al., 1997; Sandquist et al., 1997). The blue stragglers are out-lying main sequence stars that lie beyond the main sequence turnoff in the colour-magnitude diagram (CMD) of a stellar cluster. The blue stragglers are more massive, brighter, and bluer than the turnoff stars. Star formation in stellar clusters is believed to have taken place approximately all at once and since more massive stars evolve faster than lower mass stars, we do not expect to find any beyond the turnoff. Therefore, the blue stragglers must have formed more recently through some process different from the initial collective star formation.

Blue stragglers have been found in all globular clusters that have been studied in detail and the number found in each cluster is typically $10 - 20$, though it can range as high as 50. Given these facts, they must have typical lifetimes on the order of a billion years. In principle, the merger of two main sequence stars can produce a younger remnant star with a reasonably long lifetime provided that significant chemical mixing occurs in the process. The mixing must produce a higher hydrogen fraction in the core of the remnant than that of the parent stars, which being near the turnoff, have burnt most of the hydrogen to helium in their cores.

In this thesis, I present the first astrophysical application of three-dimensional Eulerian hydrodynamic simulations to model the formation of blue stragglers from mergers of main sequence stars. The simulations are used to specifically address the question of hydrogen mixing in the off-axis collision of equal mass main sequence stars found near the turnoff.

1.4 This Thesis

This thesis was carried out under the supervision of Prof. Ue-Li Pen (CITA) and some of the work also involved collaborations with others. I try to distinguish between my own work and that of collaborators where necessary. In this thesis, there are five chapters devoted to the development of numerical techniques and three chapters on the scientific application of numerical simulations to cosmology and astrophysics. Each chapter includes a short introduction to the layout and a summary of the key highlights. In addition, there is a final chapter that provides conclusions to the

thesis work and a discussion of future work. An outline of the layout of this thesis is given below.

Chapter 2 is a primer on Eulerian CFD, geared towards astrophysics. This pedagogical review was originally motivated by a 8-hour mini-course on Eulerian CFD given by Prof. Ue-Li Pen (CITA) and has been published in Trac & Pen (2003).

Chapter 3 describes the development and testing of a three-dimensional, self-gravitating Eulerian TVD hydro code for astrophysical applications. The TVD hydro algorithm from Chapter 2 is equipped with a gravity solver that is based on fast Fourier transforms (FFTs).

Chapter 4 is based on ongoing scientific work, where the three-dimensional Eulerian TVD hydro code is applied to simulating the astrophysical formation of blue stragglers through stellar mergers. This project has led to a collaborative project with Prof. Alison Sills (McMaster University), where we are comparing Eulerian and Lagrangian techniques for modelling collisions between main sequence stars. The work presented in this chapter is based on my own work with Eulerian simulations and the collaborative work will be presented in Trac, Sills, & Pen (2004, in prep.).

Chapter 5 describes the development and testing of the moving frame hydro algorithm for high Mach hydrodynamics. The algorithm was co-developed along with Prof. Ue-Li Pen (CITA), though the implementation was my own original work.

Chapter 6 describes the construction of a new cosmological code based on the moving frame algorithm and a PM N-body algorithm. A cosmological initial conditions generator, hydrodynamic tests, and sample cosmological simulations are also presented.

Chapter 7 describes the development and testing of a new out-of-core cosmological code and initial conditions generator. The code was also co-developed with Prof. Ue-Li Pen (CITA), but the implementation is my own original work.

Chapter 8 presents results from an ongoing cosmological study of the thermal evolution of the high redshift IGM, in both adiabatic and nonadiabatic scenarios. The out-of-core hydro (OCH) code is applied to running the largest Eulerian hydrodynamic simulation to date for this particular scientific study. The radiative cooling and heating algorithm that I have implemented for the OCH code is based on a version provided by Prof. Uros Seljak (Princeton).

Chapter 9 presents results from an ongoing cosmological study of the mass-dependent clustering and biasing of dark matter halos in high resolution simulations run with PMFAST. This MPI code was written by Hugh Merz (CITA) based on a prototype code written by me.

Chapter 10 provides conclusions to the thesis work and a discussion on future work.

Chapter 2

A Primer on Eulerian Computational Fluid Dynamics

2.1 Introduction

This chapter is a pedogogical review of some of the methods employed in Eulerian computational fluid dynamics (CFD). I briefly review the Euler equations and discuss the standard approach to discretizing conservation laws. Flux assignment methods based on traditional central differencing schemes such as the Lax-Wendroff scheme and more modern flux assignment methods such as the total variation diminishing (TVD) scheme are discussed in detail. I review the relaxing TVD scheme for solving systems of conservation laws such as the Euler equations and this algorithm is applied to hydrodynamic tests like the one-dimensional Sod shock tube and the three-dimensional Sedov-Taylor blast-wave test.

2.2 Eulerian Hydrodynamics

The Euler equations, which govern hydrodynamics, are a system of conservation laws for mass, momentum, and energy. In differential conservation form, the continuity equation, momentum equation, and energy equation are given as:

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j}(\rho v_j) = 0, \tag{2.1}$$

$$\frac{\partial (\rho v_i)}{\partial t} + \frac{\partial}{\partial x_j}(\rho v_i v_j + P\delta_{ij}) = 0, \tag{2.2}$$

$$\frac{\partial e}{\partial t} + \frac{\partial}{\partial x_j}[(e + P)v_j] = 0. \tag{2.3}$$

Gravitational and other source terms such as cooling and heating have been omitted here. The physical state of the fluid is specified by its density ρ, velocity field $\boldsymbol{v}$, and total energy density

$$e = \frac{1}{2}\rho v^2 + \epsilon. \tag{2.4}$$

In practice, the thermal energy density ϵ is evaluated by subtracting the kinetic energy density from the total energy density. For an ideal gas, the pressure $P(\epsilon)$ is related to the thermal energy density by the equation of state

$$P = (\gamma - 1)\epsilon, \tag{2.5}$$

where γ is the ratio of specific heats. The temperature T and the sound speed c_s are given by

$$T = \frac{P}{nk}, \tag{2.6}$$

$$c_s^2 \equiv \frac{\partial P}{\partial \rho} = \frac{\gamma P}{\rho}, \tag{2.7}$$

where n is the gas number density. The thermodynamical properties of an ideal gas obey well known equations of state.

The differential Euler equations require differentiable solutions and therefore, are ill-defined at jump discontinuities where derivatives are infinite. In the literature, nondifferentiable solutions are called weak solutions. The differential form gives a complete description of the flow in smooth regions, but the integral form is needed to properly describe shock discontinuities. In integral conservation form, the rate of change in mass, momentum, and energy is equal to the net flux of those conserved quantities through the surface enclosing a control volume. For simplicity of notation, I will continue to express the conservation laws in differential form, as a shorthand for the integral form.

2.3 Computational Fluid Dynamics

The standard approach to Eulerian computational fluid dynamics is to discretize time into discrete steps and space into finite volumes or cells, where the conserved quantities are stored. In the simplest case, the integral Euler equations are solved on a Cartesian cubical lattice by computing the flux of mass, momentum, and energy across cell boundaries in discrete time steps. Consider the Euler equations in vector differential conservation form,

$$\frac{\partial \boldsymbol{u}}{\partial t} + \frac{\partial \boldsymbol{F}_i(\boldsymbol{u})}{\partial x_i} = 0, \tag{2.8}$$

where $\boldsymbol{u} = (\rho, \rho v_x, \rho v_y, \rho v_z, e)$ contains the conserved physical quantities and $\boldsymbol{F}(\boldsymbol{u})$ represents the flux terms. In practice, the conserved cell-averaged quantities $\boldsymbol{u}_n \equiv \boldsymbol{u}(\boldsymbol{x}_n)$ and fluxes $\boldsymbol{F}_n$ are

defined at integer grid cell centres $\boldsymbol{x}_n$. The challenge is to use the cell-averaged values to determine the fluxes $\boldsymbol{F}_{n+1/2}$ at cell boundaries.

In the following sections, I describe flux assignment methods designed to solve conservation laws like the Euler equations. For ease of illustration, I begin by considering a one-dimensional scalar conservation law,

$$\frac{\partial u}{\partial t} + \frac{\partial F(u)}{\partial x} = 0, \tag{2.9}$$

where $F(u) = vu$ and v is a constant advection velocity. Equation (2.9) is referred to as a linear advection equation and has the analytical solution,

$$u(x,t) = u(x - vt, 0). \tag{2.10}$$

The linear advection equation describes the transport of the quantity u at a constant velocity v.

In integral flux conservation form, the one-dimensional scalar conservation law can be written as

$$\frac{\partial}{\partial t}\int_{x_1}^{x_2} u(x,t)dx + \int_{x_1}^{x_2} \frac{\partial F(u)}{\partial x}dx = 0, \tag{2.11}$$

where $x_1 \equiv x_{n-1/2}$ and $x_2 \equiv x_{n+1/2}$ for our control cells. Let $F^t_{n+1/2}$ denote the flux of u through cell boundary $x_{n+1/2}$ at time t. Note then that the second integral is simply equal to $F^t_{n+1/2} - F^t_{n-1/2}$. The rate of change in the cell-integrated quantity $\int u dx$ is equal to the net flux of u through the control cell. For a discrete time step, the discretized solution for the cell-averaged quantity u_n is given by

$$u_n^{t+\Delta t} = u_n^t - \left(\frac{F^t_{n+1/2} - F^t_{n-1/2}}{\Delta x}\right)\Delta t. \tag{2.12}$$

The physical quantity u is conserved since the flux taken out of one cell is added to the neighbouring cell that shares the same boundary. Note that equation (2.12) has the appearance of being a finite difference scheme for solving the differential form of the one-dimensional scalar conservation law. This is why the differential form can be used as a shorthand for the integral form.

2.4 Centered Finite-Difference Methods

Central-space finite-difference methods have ease of implementation but at the cost of lower accuracy and stability. For illustrative purposes, I start with a simple first-order centered scheme to solve the linear advection equation. The discretized solution is given by equation (2.12) where the fluxes at cell boundaries,

$$F^t_{n+1/2} = \frac{F^t_{n+1} + F^t_n}{2}, \tag{2.13}$$

are obtained by taking an average of cell-centered fluxes $F_n^t = vu_n^t$. The discretized first-order centered scheme can be equivalently written as

$$u_n^{t+\Delta t} = u_n^t - \left(\frac{F_{n+1}^t - F_{n-1}^t}{2\Delta x}\right)\Delta t. \tag{2.14}$$

In this form, the discretization has the appearance of using a central difference scheme to approximate spatial derivatives. Hence, centered schemes are often referred to as central difference schemes. In practice when using centered schemes, the discretization is done on the differential conservation equation rather than the integral equation.

This simple scheme is numerically unstable and one can show this using the *von Neumann* linear stability analysis. Consider writing $u(x,t)$ as a discrete Fourier series:

$$u_n^t = \frac{1}{N}\sum_{k=-N/2}^{N/2} c_k^t \exp\left(\frac{2\pi i k n}{N}\right), \tag{2.15}$$

where N is the number of cells in our periodic box. In plane-wave solution form, one can write this as

$$u_n^t = \frac{1}{N}\sum_{k=-N/2}^{N/2} c_k^\circ \exp\left[\frac{2\pi i(kn - \omega t)}{N}\right], \tag{2.16}$$

where c_k° are the Fourier series coefficients for the initial conditions $u(x,0)$. Equivalently, the time evolution of the Fourier series coefficients in equation (2.15) can be cast into a plane-wave solution of the form,

$$c_k^t = \exp\left(\frac{-2\pi i\omega t}{N}\right) c_k^\circ, \tag{2.17}$$

where the numerical dispersion relation $\omega(k)$ is complex in general. The imaginary part of ω represents the growth or decay of the Fourier modes while the real part describes the oscillations. A numerical scheme is linearly stable if $\mathrm{Im}(\omega) \leq 0$. Otherwise, the Fourier modes will grow exponentially in time and the solution will blow up.

The exact solution to the linear advection equation can be expressed in the form of equation (2.10) or by a plane-wave solution where the dispersion relation is given by $\omega_\circ = vk$. The waves all travel at the same phase velocity $\omega_\circ/k = v$ in the exact case.

The centrally discretized linear advection equation (Eq. 2.14) is exactly solvable. After m times steps, the time evolution of the independent Fourier modes is given by

$$c_k^{m\Delta t} = (1 - i\lambda\sin\phi)^m\, c_k^\circ, \tag{2.18}$$

where $\lambda \equiv v\Delta t/\Delta x$ and $\phi = 2\pi k\Delta x/N$. The dispersion relation is given by

$$\omega = \frac{N}{2\pi\Delta t}\left[\tan^{-1}(\lambda\sin\phi) + \frac{i}{2}\ln\left(1 + \lambda^2\sin^2\phi\right)\right], \tag{2.19}$$

For any time step $\Delta t > 0$, the imaginary part of ω will be > 0. The Fourier modes will grow exponentially in time and the solution will blow up. Hence, the first-order centered scheme is numerically unstable.

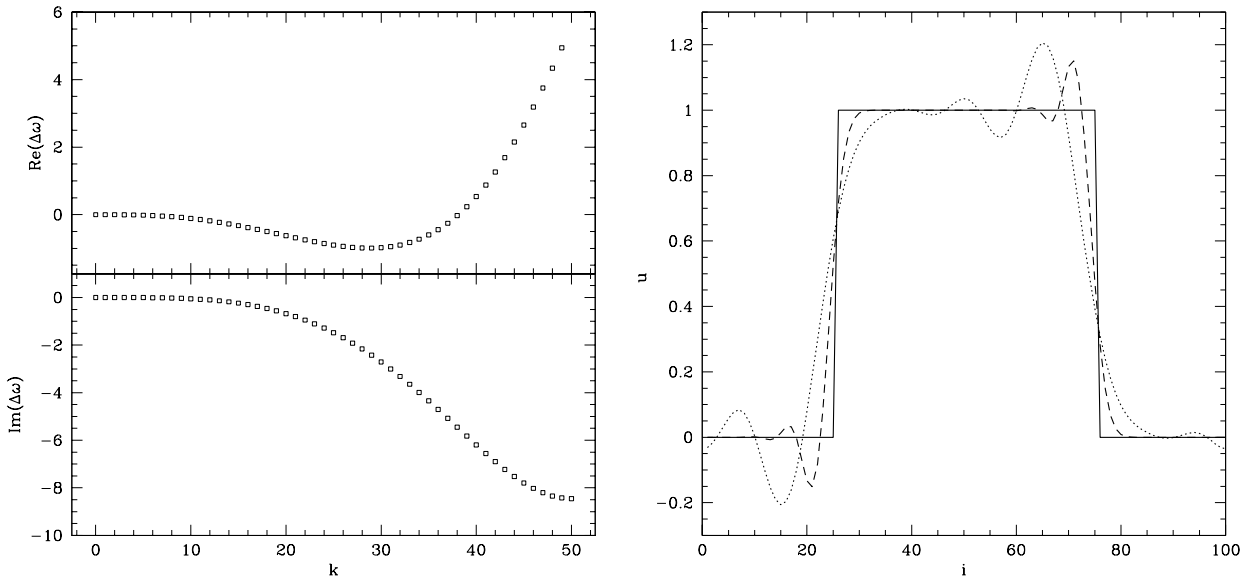

(a) **(b)**

Figure 2.1 Lax-Wendroff scheme with parameters $N = 100$, $v = 1$, and $\lambda = 0.9$. In (a), the phase error $\mathrm{Re}(\Delta\omega)$ and the amplification factor $\mathrm{Im}(\Delta\omega)$ are plotted. In (b), a square wave (solid line) is linearly advected once (dashed line) and 10 times (dotted line) through the box.

2.4.1 Lax-Wendroff Scheme

The Lax-Wendroff scheme (Lax & Wendroff, 1960) is second-order accurate in time and space, and the idea behind it is to stabilize the unstable first-order scheme from the previous section. Consider a Taylor series expansion for $u(x, t + \Delta t)$:

$$u(x, t + \Delta t) = u(x, t) + \frac{\partial u}{\partial t}\Delta t + \frac{\partial^2 u}{\partial t^2}\frac{\Delta t^2}{2} + \mathcal{O}(\Delta t^3), \tag{2.20}$$

and replace the time derivatives with spatial derivatives using the conservation law to obtain

$$u(x, t + \Delta t) = u(x, t) - \frac{\partial F}{\partial x}\Delta t + \frac{\partial}{\partial x}\left(\frac{\partial F}{\partial u}\frac{\partial F}{\partial u}\frac{\partial u}{\partial x}\right)\frac{\Delta t^2}{2} + \mathcal{O}(\Delta t^3). \tag{2.21}$$

For the linear advection equation, the eigenvalue of the flux Jacobian is $\partial F/\partial u = v$. Discretization using central differences gives

$$u_n^{t+\Delta t} = u_n^t - \frac{F_{n+1}^t - F_{n-1}^t}{2\Delta x}\Delta t + \left(\frac{F_{n+1}^t - F_n^t}{\Delta x} - \frac{F_n^t - F_{n-1}^t}{\Delta x}\right)\frac{v\Delta t^2}{2\Delta x}. \tag{2.22}$$

In conservation form, the solution is given by equation (2.12), where the fluxes at cell boundaries are defined as

$$F_{n+1/2}^t = \frac{1}{2}\left(F_{n+1}^t + F_n^t\right) - \left(F_{n+1}^t - F_n^t\right)\frac{v\Delta t}{2\Delta x}. \tag{2.23}$$

Compare this with the boundary fluxes for the first-order scheme (Eq. 2.13). The Lax-Wendroff scheme obtains second-order fluxes,

$$F^{(2)} = F^{(1)} - \frac{\partial F}{\partial u}\frac{\partial F}{\partial x}\frac{\Delta t}{2}, \tag{2.24}$$

by modifying the first-order fluxes $F^{(1)}$ with a second-order correction.

The stability of the Lax-Wendroff scheme to solve the linear advection equation can also be determined using the von Neumann analysis. The discretized Lax-Wendroff equation (Eq. 2.22) is exactly solvable and after m time steps, the Fourier modes evolve according to

$$c_k^{m\Delta t} = \left[1 - \lambda^2(1-\cos\phi) - i\lambda\sin\phi\right]^m c_k^\circ, \tag{2.25}$$

where $\lambda \equiv v\Delta t/\Delta x$ is called the *Courant* number and $\phi = 2\pi k\Delta x/N$. The dispersion relation is given by

$$\omega = \frac{N}{2\pi\Delta t}\tan^{-1}\left[\frac{\lambda\sin\phi}{1-\lambda^2(1-\cos\phi)}\right] + \frac{iN}{4\pi\Delta t}\ln\left[1 - 4\lambda^2(1-\lambda^2)\sin^4\left(\frac{\phi}{2}\right)\right]. \tag{2.26}$$

It is important to note three things. First, the Lax-Wendroff scheme is conditionally stable provided that $\mathrm{Im}(\omega) \leq 0$, which is satisfied if

$$\frac{v\Delta t}{\Delta x} \leq 1. \tag{2.27}$$

This constraint is a particular example of a general stability constraint known as the *Courant-Friedrichs-Lewyor* (CFL) condition. The Courant number λ is also referred to as the CFL number. Second, for $\lambda = 1$ the dispersion relation is exactly identical to that of the exact solution and the numerical advection is exact. This is a special case, however, and it does not test the ability of the Lax-Wendroff scheme to solve general scalar conservation laws. Normally, one chooses $\lambda < 1$ to satisfy the CFL condition. Lastly, for $\lambda < 1$ the dispersion relation $\omega(k)$ for the Lax-Wendroff solution is different from the exact solution where $\omega_\circ = vk$. The dispersion relation relative to the exact solution can be parametrized by

$$\Delta\omega \equiv \omega - \omega_\circ. \tag{2.28}$$

The second-order truncation of the Taylor series (Eq. 2.21) results in a phase error $\mathrm{Re}(\Delta\omega)$ which is a function of frequency. In the Lax-Wendroff solution, the waves are damped and travel at different speeds. Hence the scheme is both diffusive and dispersive.

In Figure 2.1(a), the phase error $\mathrm{Re}(\Delta\omega)$ and the amplification term $\mathrm{Im}(\Delta\omega)$ are plotted for the Lax-Wendroff scheme with parameters $N = 100$, $v = 1$, and $\lambda = 0.9$. A negative value of $\mathrm{Re}(\Delta\omega)$ represents a lagging phase error while a positive value indicates a leading phase error. For the chosen CFL number, the high frequency modes have the largest phase errors but they are highly damped. Some of the modes having lagging phase errors are not highly damped. We will subsequently see how this becomes important.

A rigourous test of the one-dimensional Lax-Wendroff scheme and other flux assignment schemes is the linear advection of a square wave. The challenge is to accurately advect this discontinuous function where the edges mimic Riemann shock fronts. Figure 2.1(b) shows how the Lax-Wendroff scheme does at advecting the square wave once (dashed line) and ten times (dotted line) through a periodic box of 100 grid cells at speed $v = 1$ and $\lambda = 0.9$. Note that this scheme produces numerical oscillations and is highly dispersive. Recall that a square wave can be represented by a sum of Fourier or sine waves. These waves will be damped and disperse when advected using the Lax-Wendroff scheme. The modes having lagging phase errors are not damped entirely away and the dispersion results in the oscillations seen in Figure 2.1(b). Note that when the Lax-Wendroff scheme is used to advect a sine wave, there will be no spurious oscillations due to dispersion since there is only one frequency mode, but a phase error will be present. For a comprehensive discussion on the family of Lax-Wendroff schemes and other centered schemes, see Hirsch (1990) and Laney (1998).

2.5 Upwind Methods

Upwind methods take into account the physical nature of the flow when assigning fluxes for the discrete solution. This class of flux assignment schemes, whose origin dates back to the work of Courant, Isaacson, & Reeves (1952), has been shown to be excellent at capturing shocks and also being highly stable.

To start, a simple first-order upwind scheme will be used to solve the linear advection equation. Consider the case where the advection velocity is positive and flow is to the right. The flux of the physical quantity u through the cell boundary $x_{n+1/2}$ will originate from cell n. The upwind scheme proposes that, to first-order, the fluxes $F^t_{n+1/2}$ at cell boundaries be taken from the cell-centered fluxes $F^t_n = vu^t_n$, which is in the upwind direction. If the advection velocity is negative and flow is to the left, the boundary fluxes $F^t_{n+1/2}$ are taken from the cell-centered fluxes $F^t_{n+1} = vu^t_{n+1}$. The first-order upwind flux assignment scheme can be summarized as follows:

$$F^t_{n+1/2} = \begin{cases} F^t_n & \text{if } v > 0, \\ F^t_{n+1} & \text{if } v < 0. \end{cases} \tag{2.29}$$

Unlike central difference schemes, upwind schemes are explicitly asymmetric.

The CFL condition for the first-order upwind scheme can be determined from the von Neumann analysis. Consider the case of a positive advection velocity. After m time steps, the Fourier modes evolve according to

$$c_k^{m\Delta t} = [1 - \lambda(1 - \cos\phi) - i\lambda \sin\phi]^m \, c_k^\circ, \tag{2.30}$$

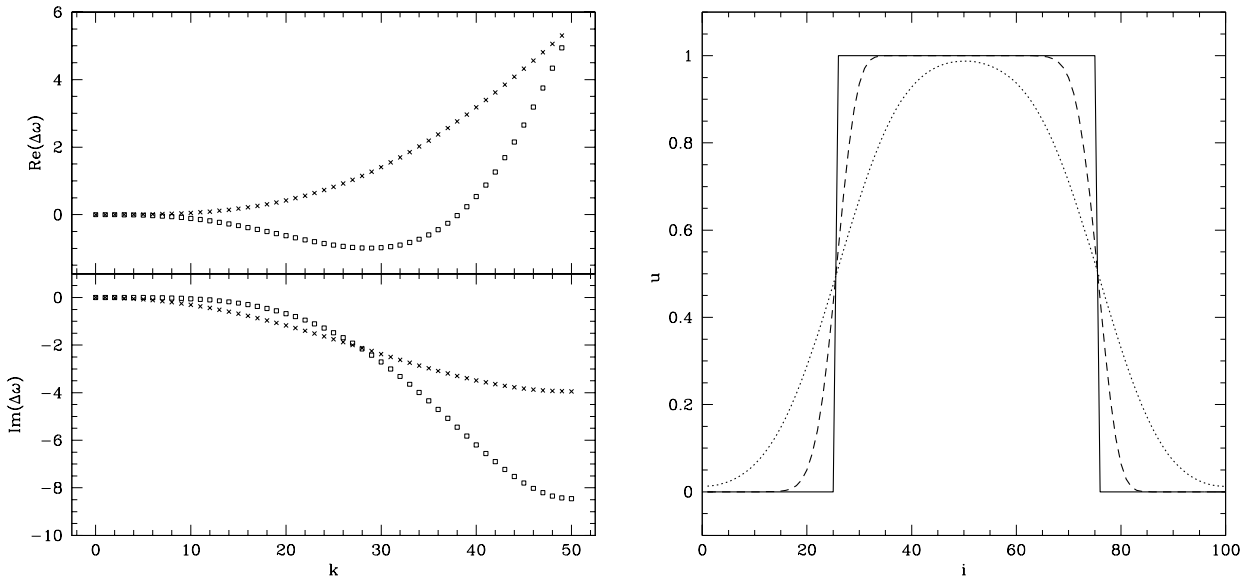

(a) **(b)**

Figure 2.2 First-order upwind scheme with parameters $N = 100$, $v = 1$, and $\lambda = 0.9$. In (a), the phase error $\mathrm{Re}(\Delta\omega)$ and the amplification factor $\mathrm{Im}(\Delta\omega)$ are plotted for the upwind scheme (crosses) and the Lax-Wendroff scheme (boxes) for comparison. In (b), a square wave (solid line) is linearly advected once (dashed line) and 10 times (dotted line) through the box

where $\lambda \equiv v\Delta t/\Delta x$ and $\phi = 2\pi k\Delta x/N$. The dispersion relation is given by

$$\omega = \frac{N}{2\pi\Delta t}\tan^{-1}\left[\frac{\lambda\sin\phi}{1-\lambda(1-\cos\phi)}\right] + \frac{iN}{4\pi\Delta t}\ln\left[1-4\lambda(1-\lambda)\sin^2\left(\frac{\phi}{2}\right)\right]. \tag{2.31}$$

The CFL condition for solving the linear advection equation with this scheme is to have $\lambda \leq 1$, identical to that for the Lax-Wendroff scheme. For $\lambda < 1$ the dispersion relation $\omega(k)$ for the first-order upwind scheme is different from the exact solution where $\omega_\circ = vk$. This scheme is both diffusive and dispersive. Since it is only first-order accurate, the amount of diffusion is large. In Figure 2.2(a) the dispersion relation of the upwind scheme is compared to that of the Lax-Wendroff scheme. The Fourier modes in the upwind scheme also have phase errors but they will be damped away. The low frequency modes that contribute to the oscillations in the Lax-Wendroff solution are more damped in the upwind solution. Hence, one does not expect to see oscillations resulting from phase errors.

Figure 2.2(b) shows how the first-order upwind scheme does at advecting the Riemann shock wave. This scheme is well-behaved and produces no spurious oscillations, but since it is only first-order, it is highly diffusive. The first-order upwind scheme has the property of having monotonicity preservation. When applied to the linear advection equation, it does not allow the creation of new extrema in the form of spurious oscillations. The Lax-Wendroff scheme does not have the property of having monotonicity preservation.

The flux assignment schemes that I have discussed so far are all linear schemes. Godunov (1959) showed that all linear schemes are either diffusive or dispersive or a combination of both. This is one part of Godunov's theorem. The Lax-Wendroff scheme is highly dispersive while the first-order upwind scheme is highly diffusive. Godunov's theorem also states that linear monotonicity preserving schemes are only first-order accurate. In order to obtain higher order accuracy and prevent spurious oscillations, nonlinear schemes are needed to solve conservation laws.

2.5.1 Total Variation Diminishing Schemes

Harten (1983) proposed the *total variation diminishing* (TVD) condition, which guarantees that a scheme have monotonicity preservation. According to Godunov's theorem, all linear TVD schemes are only first-order accurate. In fact, the only linear TVD schemes are the class of first-order upwind schemes. Therefore, higher order accurate TVD schemes must be nonlinear.

The TVD condition is a nonlinear stability condition. The total variation of a discrete solution, defined as

$$TV(u^t) = \sum_{i=1}^{N} |u_{i+1}^t - u_i^t|, \tag{2.32}$$

is a measure of the overall amount oscillations in u. The direct connection between the total variation and the overall amount of oscillations can be seen in the equivalent definition

$$TV(u^t) = 2\left(\sum u_{\max} - \sum u_{\min}\right), \tag{2.33}$$

where each maxima is counted positively twice and each minima counted negatively twice (see Laney, 1998). The formation of spurious oscillations will contribute new maxima and minima and the total variation will increase. A flux assignment scheme is said to be TVD if

$$TV(u^{t+\Delta t}) \leq TV(u^t), \tag{2.34}$$

which signifies that the overall amount of oscillations is bounded. In linear flux-assignment schemes, the von Neumann linear stability condition requires that the Fourier modes remain bounded. In nonlinear schemes, the TVD stability condition requires that the total variation diminishes.

I now describe a nonlinear second-order accurate TVD scheme that builds upon the first-order monotone upwind scheme described in the previous section. The second-order accurate fluxes $F_{n+1/2}^t$ at cell boundaries are obtained by taking first-order fluxes $F_{n+1/2}^{(1),t}$ from the upwind scheme and modifying it with a second order correction. First consider the case where the advection velocity is positive. The first-order upwind flux $F_{n+1/2}^{(1),t}$ comes from the averaged flux F_n^t in cell n. One can define two second-order flux corrections,

$$\Delta F_{n+1/2}^{L,t} = \frac{F_n^t - F_{n-1}^t}{2}, \tag{2.35}$$

$$\Delta F_{n+1/2}^{R,t} = \frac{F_{n+1}^t - F_n^t}{2}, \tag{2.36}$$

using three local cell-centered fluxes. Cell n and the cells immediately left and right of it are used. If the advection velocity is negative, the first-order upwind flux comes from the averaged flux F^t_{n+1} in cell $n+1$. In this case, the second-order flux corrections,

$$\Delta F^{L,t}_{n+1/2} = -\frac{F^t_{n+1} - F^t_n}{2}, \tag{2.37}$$

$$\Delta F^{R,t}_{n+1/2} = -\frac{F^t_{n+2} - F^t_{n+1}}{2}, \tag{2.38}$$

are based on cell $n+1$ and the cells directly adjacent to it. Near extrema where the corrections have opposite signs, no second-order correction is imposed and the flux assignment scheme reduces to first-order. A flux limiter ϕ is then used to determine the appropriate second-order correction,

$$\Delta F^t_{n+1/2} = \phi(\Delta F^{L,t}_{n+1/2}, \Delta F^{R,t}_{n+1/2}) \ , \tag{2.39}$$

which still maintains the TVD condition. The second-order correction is added to the first-order fluxes to get second-order fluxes. The first-order upwind scheme and second-order TVD scheme will be referred to as *monotone upwind schemes for conservation laws* (MUSCL).

Time integration is performed using a second-order Runge-Kutta scheme. A half time step is first performed,

$$u^{t+\Delta t/2}_n = u^t_n - \left(\frac{F^t_{n+1/2} - F^t_{n-1/2}}{\Delta x} \right) \frac{\Delta t}{2} \ , \tag{2.40}$$

using the first-order upwind scheme to obtain the half-step values $u^{t+\Delta t/2}$. A full time step is then computed,

$$u^{t+\Delta t}_n = u^t_n - \left(\frac{F^{t+\Delta t/2}_{n+1/2} - F^{t+\Delta t/2}_{n-1/2}}{\Delta x} \right) \Delta t \ , \tag{2.41}$$

using the TVD scheme on the half-step fluxes $F^{t+\Delta t/2}$.

I briefly discuss three TVD limiters. The minmod flux limiter chooses the smallest absolute value between the left and right corrections:

$$minmod(a,b) = \tfrac{1}{2}[\text{sign}(a) + \text{sign}(b)] \min(|a|, |b|) \ . \tag{2.42}$$

The superbee limiter (Roe, P. L., 1985) chooses between the larger correction and two times the smaller correction, whichever is smaller in magnitude:

$$superbee(a,b) = \begin{cases} minmod(a, 2b) & \text{if } |a| \geq |b|, \\ minmod(2a, b) & \text{otherwise.} \end{cases} \tag{2.43}$$

The Van Leer limiter (Van Leer, 1974) takes the harmonic mean of the left and right corrections:

$$vanleer(a,b) = \frac{2ab}{a+b}. \tag{2.44}$$

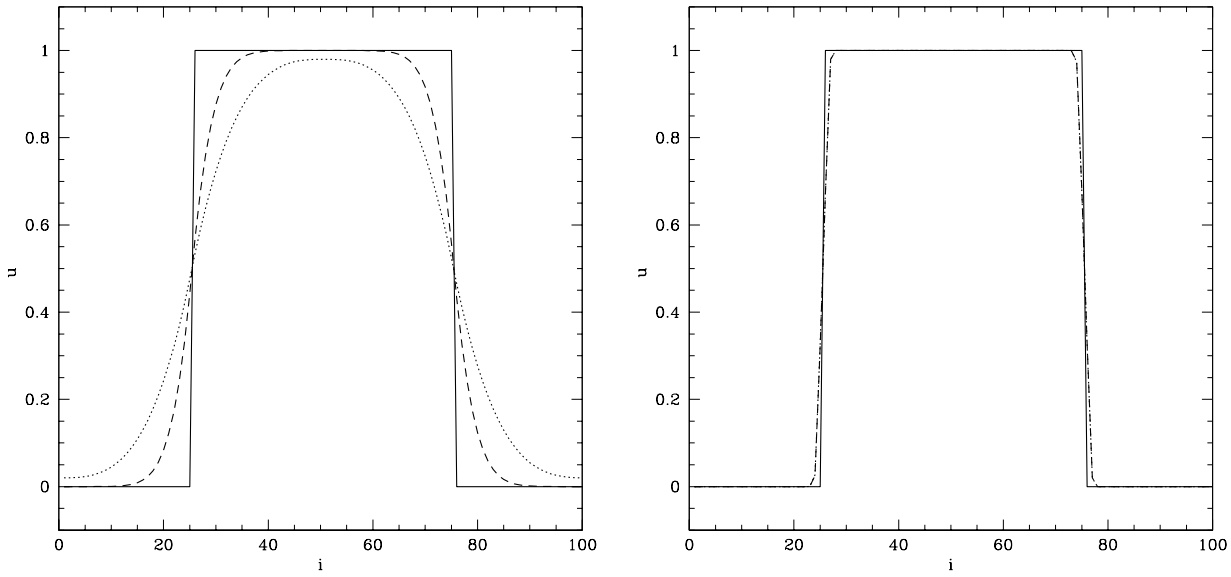

(a) **(b)**

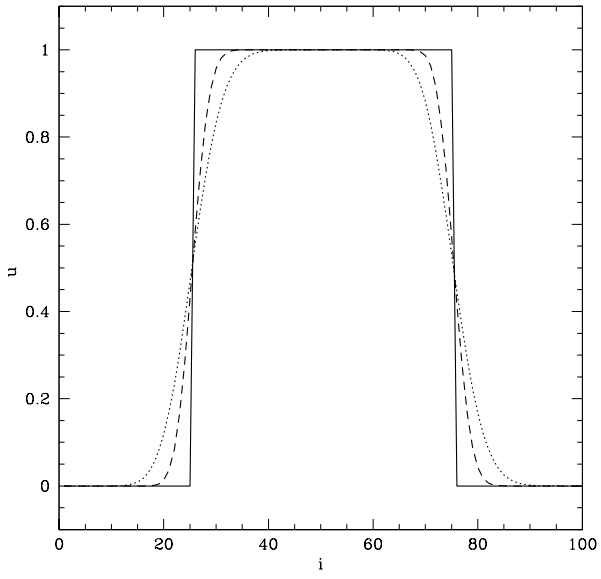

(c)

Figure 2.3 TVD scheme using the (a) minmod, (b) superbee, and (c) Van Leer flux limiters to advect a square wave (solid line) once (dashed line) and 10 times (dotted line) through a periodic box of $N = 100$ grid cells at speed $v = 1$.

The minmod limiter is the most moderate of all second-order TVD limiters. Figure 2.3(a) shows that the minmod limiter does not do much better than first-order upwind for the square wave advection test. Superbee chooses the maximum correction allowed under the TVD constraint. It is especially suited for piece-wise linear conditions and is the least diffusive for this particular test, as can be seen in Figure 2.3(b). Note that no additional diffusion can be seen by advecting the square wave more than once through the box. It can be shown that the minmod and superbee limiters are extreme cases that bound all other second-order TVD limiters. The Van Leer limiter differs from the previous two in that it is analytic. This symmetrical approach falls somewhere inbetween the other two limiters in terms of moderation and diffusion, as can be seen in Figure 2.3(c). It can be shown that the CFL condition for the second-order TVD scheme is to have $\lambda < 1$. For a comprehensive discussion on TVD limiters, see Hirsch (1990) and Laney (1998).

2.6 Relaxing TVD

I now describe a simple and robust method to solve the Euler equations using the MUSCL from the previous section. The relaxing TVD method (Jin & Xin, 1985) provides high resolution capturing of shocks using computationally inexpensive algorithms that are straightforward to implement and to parallelize. It has been successfully implemented for simulating cosmological astrophysical fluids by Pen (1998a) and Trac & Pen (2004a)

The MUSCL scheme is straightforward to apply to conservation laws like the advection equation since the velocity alone can be used as a marker of the direction of flow. However, applying the MUSCL scheme to solve the Euler equations is made difficult by the fact that the momentum and energy fluxes depend on the pressure. In order to determine the direction upwind of the flow, it becomes necessary to calculate the flux Jacobian eigenvectors using Riemann solvers. This step requires computationally expensive algorithms. The relaxing TVD method offers an attractive alternative.

2.6.1 One-dimensional Scalar Conservation Law

I first present a motivation for the relaxing scheme by again considering the one-dimensional scalar conservation law. The MUSCL scheme for solving the linear advection equation is explicitly asymmetric in that it depends on the sign of the advection velocity. The relaxing scheme is a symmetrical approach that applies to a general advection velocity.

The flow can be considered as a sum of a right-moving wave u^R and a left-moving wave u^L. For a positive advection velocity, the amplitude of the left-moving wave is zero and for a negative advection velocity, the amplitude of the right-moving wave is zero. In compact notation, the waves

can be defined as:

$$u^R = \left(\frac{1+v/c}{2}\right) u, \tag{2.45}$$

$$u^L = \left(\frac{1-v/c}{2}\right) u, \tag{2.46}$$

where $c = |v|$. The two waves are traveling in opposite directions with advection speed c and can be described by the advection equations:

$$\frac{\partial u^R}{\partial t} + \frac{\partial}{\partial x}(cu^R) = 0, \tag{2.47}$$

$$\frac{\partial u^L}{\partial t} - \frac{\partial}{\partial x}(cu^L) = 0. \tag{2.48}$$

The MUSCL scheme is straightforward to apply to solve equations (2.47) and (2.48) since the upwind direction is left for the right-moving wave and right for the left-moving wave. The one-dimensional relaxing advection equation then becomes

$$\frac{\partial u}{\partial t} + \frac{\partial F^R}{\partial x} - \frac{\partial F^L}{\partial x} = 0. \tag{2.49}$$

where $F^R = cu^R$ and $F^L = cu^L$. For the discretized solution given by equation (2.12), the boundary fluxes $F^t_{n+1/2}$ are now a sum of the fluxes $F^{R,t}_{n+1/2}$ and $F^{L,t}_{n+1/2}$ from the right-moving and left-moving waves, respectively. Note that the relaxing advection equation will correctly reduce to the linear advection equation for any general advection velocity.

Using this symmetrical approach, a general algorithm can be written to solve the linear advection equation for an arbitrary advection velocity. This scheme is indeed inefficient for solving the linear advection equation since one wave will have zero amplitude. However, the Euler equations can have both right-moving and left-moving waves with non-zero amplitudes.

2.6.2 One-dimensional Systems of Conservation Laws

I now discuss the one-dimensional relaxing TVD scheme and later generalize it to higher spatial dimensions. Consider a one-dimensional system of conservation laws,

$$\frac{\partial \boldsymbol{u}}{\partial t} + \frac{\partial \boldsymbol{F}(u)}{\partial x} = 0, \tag{2.50}$$

where for the Euler equations, we have $\boldsymbol{u} = (\rho, \rho v, e)$ and $\boldsymbol{F}(\boldsymbol{u})$ the corresponding flux terms. The vector conservation law is replaced with the relaxation system

$$\frac{\partial \boldsymbol{u}}{\partial t} + \frac{\partial}{\partial x}(c\boldsymbol{w}) = 0, \tag{2.51}$$

$$\frac{\partial \boldsymbol{w}}{\partial t} + \frac{\partial}{\partial x}(c\boldsymbol{u}) = 0, \tag{2.52}$$

where $c(x,t)$ is a free positive function called the freezing speed. I first consider the freezing speed to be spatially constant and then later generalize it to be spatially varying. The relaxation system contains two coupled vector linear advection equations. In practice, we set $\boldsymbol{w} = \boldsymbol{F}(\boldsymbol{u})/c$ and use it as an auxiliary vector to calculate fluxes. Equation (2.51) reduces to the one-dimensional vector conservation law and equation (2.52) is a vector conservation law for $\boldsymbol{w}$.

In order to solve the relaxed system, the equations are decoupled through a change of variables:

$$\boldsymbol{u}^R = \frac{\boldsymbol{u} + \boldsymbol{w}}{2}, \tag{2.53}$$

$$\boldsymbol{u}^L = \frac{\boldsymbol{u} - \boldsymbol{w}}{2}, \tag{2.54}$$

which then gives us

$$\frac{\partial \boldsymbol{u}^R}{\partial t} + \frac{\partial}{\partial x}(c\boldsymbol{u}^R) = 0 \ , \tag{2.55}$$

$$\frac{\partial \boldsymbol{u}^L}{\partial t} - \frac{\partial}{\partial x}(c\boldsymbol{u}^L) = 0 \ . \tag{2.56}$$

Equations (2.55) and (2.56) are vector linear advection equations, which can be interpreted as right-moving and left-moving flows with advection speed c. Note the similarity with their scalar counterparts, equations (2.47) and (2.48). The one-dimensional vector relaxing conservation law for $\boldsymbol{u}$ becomes

$$\frac{\partial \boldsymbol{u}}{\partial t} + \frac{\partial \boldsymbol{F}^R}{\partial x} - \frac{\partial \boldsymbol{F}^L}{\partial x} = 0 \ , \tag{2.57}$$

where $\boldsymbol{F}^R = c\boldsymbol{u}^R$ and $\boldsymbol{F}^L = c\boldsymbol{u}^L$. The vector relaxing equation can now be solved by applying the MUSCL scheme to Equations (2.55) and (2.56). Again, note the similarity between the vector relaxing equation and its scalar counterpart, Equation (2.49).

The relaxed scheme is TVD under the constraint that the freezing speed c be greater than the characteristic speed given by the largest eigenvalue of the flux Jacobian $\partial \boldsymbol{F}(\boldsymbol{u})/\partial \boldsymbol{u}$. For the Euler equations, this is satisfied for

$$c = |v| + c_s. \tag{2.58}$$

Jin & Xin (1985) considered the freezing speed to be a positive constant in their relaxing scheme. Time integration is again performed using a second-order Runge-Kutta scheme and the time step is determined by satisfying the CFL condition,

$$\frac{c_{\max}\Delta t}{\Delta x} \leq 1. \tag{2.59}$$

Note that a new freezing speed is computed for each partial step in the Runge-Kutta scheme. The CFL number $\lambda = c_{max}\Delta t/\Delta x$ should be chosen such that $c_{\max}$ will be larger than $\max(c_n^t)$ and $\max(c_n^{t+\Delta t/2})$.

The steps needed to numerically solve the one-dimensional Euler equations are now summarized. At the beginning of each partial step in the Runge-Kutta time integration scheme, one needs to calculate the cell-averaged variables defined at grid cell centres. First for the half time step, the fluxes $\boldsymbol{F}(u_n^t)$ and the freezing speed c_n^t are computed. One then sets the auxiliary vector $\boldsymbol{w}_n^t = \boldsymbol{F}(\boldsymbol{u}_n^t)/c_n^t$ and construct the right-moving waves $\boldsymbol{u}_n^{R,t}$ and left-moving waves $\boldsymbol{u}_n^{L,t}$. The half time step is given by

$$\boldsymbol{u}_n^{t+\Delta t/2} = \boldsymbol{u}_n^t - \left(\frac{\boldsymbol{F}_{n+1/2}^t - \boldsymbol{F}_{n-1/2}^t}{\Delta x} \right) \frac{\Delta t}{2}, \tag{2.60}$$

where

$$\boldsymbol{F}_{n+1/2}^t = \boldsymbol{F}_{n+1/2}^{R,t} - \boldsymbol{F}_{n+1/2}^{L,t}. \tag{2.61}$$

The first-order upwind scheme is used to compute fluxes at cell boundaries for the right-moving and left-moving waves. For the full time step, the right-moving waves $\boldsymbol{u}_n^{R,t+\Delta t/2}$ and left-moving waves $\boldsymbol{u}_n^{L,t+\Delta t/2}$ are constructed using the half-step values of the appropriate variables. The full time step,

$$\boldsymbol{u}_n^{t+\Delta t} = \boldsymbol{u}_n^t - \left(\frac{\boldsymbol{F}_{n+1/2}^{t+\Delta t/2} - \boldsymbol{F}_{n-1/2}^{t+\Delta t/2}}{\Delta x} \right) \Delta t, \tag{2.62}$$

is computed using the second-order TVD scheme. This completes the updating of $\boldsymbol{u}^t$ to $\boldsymbol{u}^{t+\Delta t}$.

In the general case of a spatially varying freezing speed, minor modifications to the implementation described above are needed. Consider writing the flux of the right-moving and left-moving waves as:

$$\boldsymbol{F}^R = c\boldsymbol{G}^R, \tag{2.63}$$

$$\boldsymbol{F}^L = c\boldsymbol{G}^L, \tag{2.64}$$

where $\boldsymbol{G}^R$ is the flux of $\boldsymbol{\mu}^R = \boldsymbol{u}^R/c$ and $\boldsymbol{G}^L$ is the flux of $\boldsymbol{\mu}^L = \boldsymbol{u}^L/c$. The linear advection equations for $\boldsymbol{\mu}^R$ and $\boldsymbol{\mu}^L$ are similar to equations (2.55) and (2.56), but where $\boldsymbol{u}^R$ is replaced with $\boldsymbol{\mu}^R$ and $\boldsymbol{u}^L$ with $\boldsymbol{\mu}^L$. For each partial step in the Runge-Kutta scheme, the net fluxes at cell boundaries are then taken to be

$$\boldsymbol{F}_{n+1/2} = c_{n+1/2}(\boldsymbol{G}_{n+1/2}^R - \boldsymbol{G}_{n+1/2}^L) \ , \tag{2.65}$$

where $c_{n+1/2} = (c_{n+1} + c_n)/2$. Note that this general implementation reduces to the original when the freezing speed is a constant.

2.6.3 Multi-Dimensional Systems of Conservation Laws

The one-dimensional relaxing TVD scheme can be generalized to higher dimensions using the dimensional splitting technique by Strang (1968). In three dimensions, the Euler equations can be dimensionally split into three separate one-dimensional equations that are solved sequentially. Let the operator $\boldsymbol{L}_i$ represent the updating of $\boldsymbol{u}^t$ to $\boldsymbol{u}^{t+\Delta t}$ by including the flux in the i direction. One first completes a forward sweep,

$$\boldsymbol{u}^{t+\Delta t} = \boldsymbol{L}_z \boldsymbol{L}_y \boldsymbol{L}_x \boldsymbol{u}^t, \tag{2.66}$$

and then performs a reverse sweep

$$\boldsymbol{u}^{t+2\Delta t} = \boldsymbol{L}_x \boldsymbol{L}_y \boldsymbol{L}_z \boldsymbol{u}^{t+\Delta t}, \tag{2.67}$$

using the same time step Δt to obtain second-order accuracy. The combination of the forward and reverse sweeps will be referred to as a double sweep.

A more symmetrical sweeping pattern can be used by permutating the sweeping order when completing the next double time step. The dimensional splitting or operator splitting technique can be summarized as follows:

$$\boldsymbol{u}^{t_2} = \boldsymbol{u}^{t_1+2\Delta t_1} = \boldsymbol{L}_x \boldsymbol{L}_y \boldsymbol{L}_z \boldsymbol{L}_z \boldsymbol{L}_y \boldsymbol{L}_x \boldsymbol{u}^{t_1}, \tag{2.68}$$

$$\boldsymbol{u}^{t_3} = \boldsymbol{u}^{t_2+2\Delta t_2} = \boldsymbol{L}_z \boldsymbol{L}_x \boldsymbol{L}_y \boldsymbol{L}_y \boldsymbol{L}_x \boldsymbol{L}_z \boldsymbol{u}^{t_2}, \tag{2.69}$$

$$\boldsymbol{u}^{t_4} = \boldsymbol{u}^{t_3+2\Delta t_3} = \boldsymbol{L}_y \boldsymbol{L}_z \boldsymbol{L}_x \boldsymbol{L}_x \boldsymbol{L}_z \boldsymbol{L}_y \boldsymbol{u}^{t_3}, \tag{2.70}$$

where Δt_1, Δt_2, and Δt_3 are newly determined time steps after completing each double sweep.

The CFL condition for the three-dimensional relaxing TVD scheme is similarly given by equation (2.59), but with

$$c_{\max} = \max[(c_x)_{\max}, (c_y)_{\max}, (c_z)_{\max}] \ . \tag{2.71}$$

where $c_i = |v_i| + c_s$. Note that since $\max(|v_i|)$ is on average a factor of $\sqrt{3}$ smaller than $\max(|\boldsymbol{v}|)$, a dimensionally split scheme can use a longer time step compared to an un-split scheme.

The dimensional splitting technique also has other advantages. The decomposition into a one-dimensional problem allows one to write short one-dimensional algorithms, which are easy to optimize to be cache efficient. A three-dimensional hydro code is straightforward to implement in parallel. When sweeping in the x direction, for example, one can break up the data into one-dimensional columns and operate on the independent columns in parallel.

2.7 Hydrodynamic Tests

In this section, the one-dimensional and three-dimensional versions of the relaxing TVD algorithm are applied to hydrodynamic tests like the Sod shock tube test and the Sedov-Taylor blast-wave test.

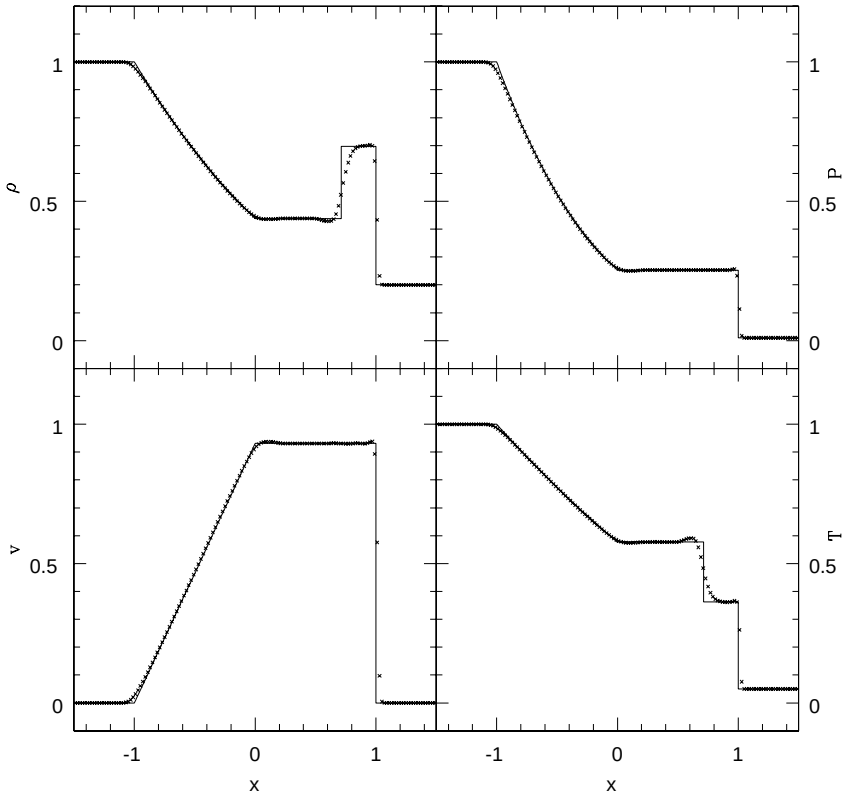

Figure 2.4 One-dimensional Sod shock tube test results obtained with the relaxing TVD algorithm. The shock ($x = 1$), contact ($x \sim 0.75$), and expansion ($x = -1$) sharply mark the different regions of the shock tube.

These standard tests are often presented in the literature to showcase the accuracy of hydrodynamic schemes in resolving strong shocks.

2.7.1 One-dimensional Sod Shock Tube Test

The Sod shock tube is a special case of the Riemann problem, where a jump discontinuity in the initial conditions leads to the development of shocks. Consider a one-dimensional tube containing two regions of fluid, initially at rest and separated by a membrane. The initial state to the right of the membrane is labelled with the subscript 1 and the initial state to the left with the subscript 5. Consider the case where $\rho_5 > \rho_1$, $P_5 > P_1$, and $v_5 = v_1 = 0$. At an initial time $t = 0$, the membrane is instantaneously removed and the pressure imbalance results in a contact discontinuity and shock wave propagating to the right towards the low-pressure region and an expansion or rarefraction fan moving to the left towards the high-pressure region. The shock, contact, and expansion each separate regions of steady flow that is labelled with the subscripts 2, 3, and 4 respectively.

The exact solution to the Riemann problem for the Euler equations uses the shock jump condi-

tions and self-similar arguments. A thorough derivation is presented in Laney (1998) and collected here is a summary of some useful relations. A contact discontinuity occurs when the velocity and pressure are continuous but other fluid properties are not. Therefore, the velocity and pressure on both sides of the contact are equal and constant:

$$v_2 = v_3, \tag{2.72}$$

$$P_2 = P_3. \tag{2.73}$$

The postshock pressure P_2 can be determined from the implicit equation

$$\frac{P_1}{P_5} = \frac{P_1}{P_2}\left[1 - \frac{c_1}{c_5}\left(\frac{\gamma-1}{2\gamma}\right)\left(\frac{P_2}{P_1}-1\right)\sqrt{\frac{2\gamma}{(\gamma+1)(P_2/P_1)+(\gamma-1)}}\right]^{2\gamma/(\gamma-1)}, \tag{2.74}$$

where $c_n = \sqrt{\gamma P_n/\rho_n}$ are the sound speeds. The postshock velocity v_2 is then given by

$$v_2 = \frac{c_1}{\gamma}\left(\frac{P_2}{P_1}-1\right)\sqrt{\frac{2\gamma}{(\gamma+1)(P_2/P_1)+(\gamma-1)}}. \tag{2.75}$$

From the Rankine-Hugoniot shock jump relations, we obtain the shock velocity

$$v_s = \frac{c_1^2}{\gamma v_2}\left(\frac{P_2}{P_1}-1\right). \tag{2.76}$$

The density jumps across the contact discontinuity and we have

$$\rho_2 = \rho_1 \frac{v_s}{v_s - v_2}, \tag{2.77}$$

$$\rho_3 = \rho_5 \left(\frac{P_3}{P_5}\right)^{1/\gamma}. \tag{2.78}$$

After time t, the contact discontinuity has propagated a distance of $x_c = v_3 t$ and the shock front is located at distance $x_s = v_s t$, relative to the initial membrane. Region 4 is a rarefraction fan that decreases the density and pressure as it expands leftward towards the high pressure region.

The one dimensional relaxing TVD algorithm is applied to the shock tube test with initial conditions: $\rho_5 = 1$, $P_5 = 1$, $\rho_1 = 0.2$, $P_1 = 0.01$. The initial conditions are identical to Shapiro et al. (1996) and Pen (1998a). The pressure ratio $P_5/P_1 = 100$ is 10 times larger than that in Sod (1978), allowing a rigourous test of the code's ability to capture strong shocks at high Mach numbers. The simulation is run until the shock front has propagated a distance of $x_s = 50$ grid cells. In Figure 2.4, the plots have been rescaled such that the initial discontinuity is placed at $x = 0$ and the shock front at $x = 1$. The grid spacing corresponds to $\Delta x = 0.02$. The relaxing TVD algorithm successfully captures the gas dynamics in this test. The shock, contact, and expansion sharply mark the different regions discussed previously. The shock front has been propagated accurately and resolved in roughly two grid cells, with no spurious oscillations. The contact discontinuity

has been degraded by diffusion, but this is inevitable when trying to advect such a front over approximately 35 grid cells. Different flux limiters for the TVD scheme can give different results. The superbee flux limiter has been found to be the least diffusive, but it is also the least stable. From experience, the van Leer limiter provides a preferable combination of accuracy and stability.

2.7.2 Three-dimensional Sedov-Taylor Blast-wave Test

A rigourous and challenging test for any three-dimensional Eulerian or Lagrangian hydrodynamic code is the Sedov-Taylor blast wave test. The simulation box is initially set up with a homogeneous medium of density ρ_1 and negligible pressure. A point-like supply of thermal energy $E_\circ$ is then injected at the centre of the box at time $t = 0$. The challenge is to accurately capture the strong spherical shock wave that sweeps along material as it propagates out into the ambient medium. The Sedov-Taylor test is used to model nuclear-type explosions. In astrophysics, it is often used as a basic setup to model supernova explosions and the evolution of supernova remnants (see Shu, 1992).

The analytical Sedov solution uses the self-similar nature of the blast wave expansion and the full analytical solution can be found in Landau & Lifshitz (1987). Collected here is a summary of some useful relations. Consider a frame fixed relative to the centre of the explosion. The spherical shock front propagates outward and the distance from the origin is given by

$$r_{sh}(t) = \xi_\circ \left(\frac{E_\circ t^2}{\rho_1} \right)^{1/5} , \tag{2.79}$$

where $\xi_\circ = 1.15$ for an ideal gas with $\gamma = 5/3$. The velocity of the shock $v_{sh} = \partial r_{sh}/\partial t$ is given by

$$v_{sh}(t) = \frac{2}{5} \frac{r_{sh}(t)}{t} . \tag{2.80}$$

Since the ambient medium has negligible pressure, the shocks will be very strong. The density ρ_2, velocity v_2, and pressure P_2 directly behind the shock front are:

$$\rho_2 = \left(\frac{\gamma + 1}{\gamma - 1} \right) \rho_1 , \tag{2.81}$$

$$v_2 = \left(\frac{2}{\gamma + 1} \right) v_{sh} , \tag{2.82}$$

$$P_2 = \left(\frac{2}{\gamma + 1} \right) \rho_1 v_{sh}^2 . \tag{2.83}$$

The immediate postshock gas density is constant in time, while the shocked gas velocity v_2 and pressure P_2 decrease as $t^{-3/5}$ and $t^{-6/5}$, respectively.

The three-dimensional relaxing TVD code based on the van Leer flux limiter is applied to capturing the Sedov-Taylor blast wave. A simulation box with 256^3 cells is initially set up with

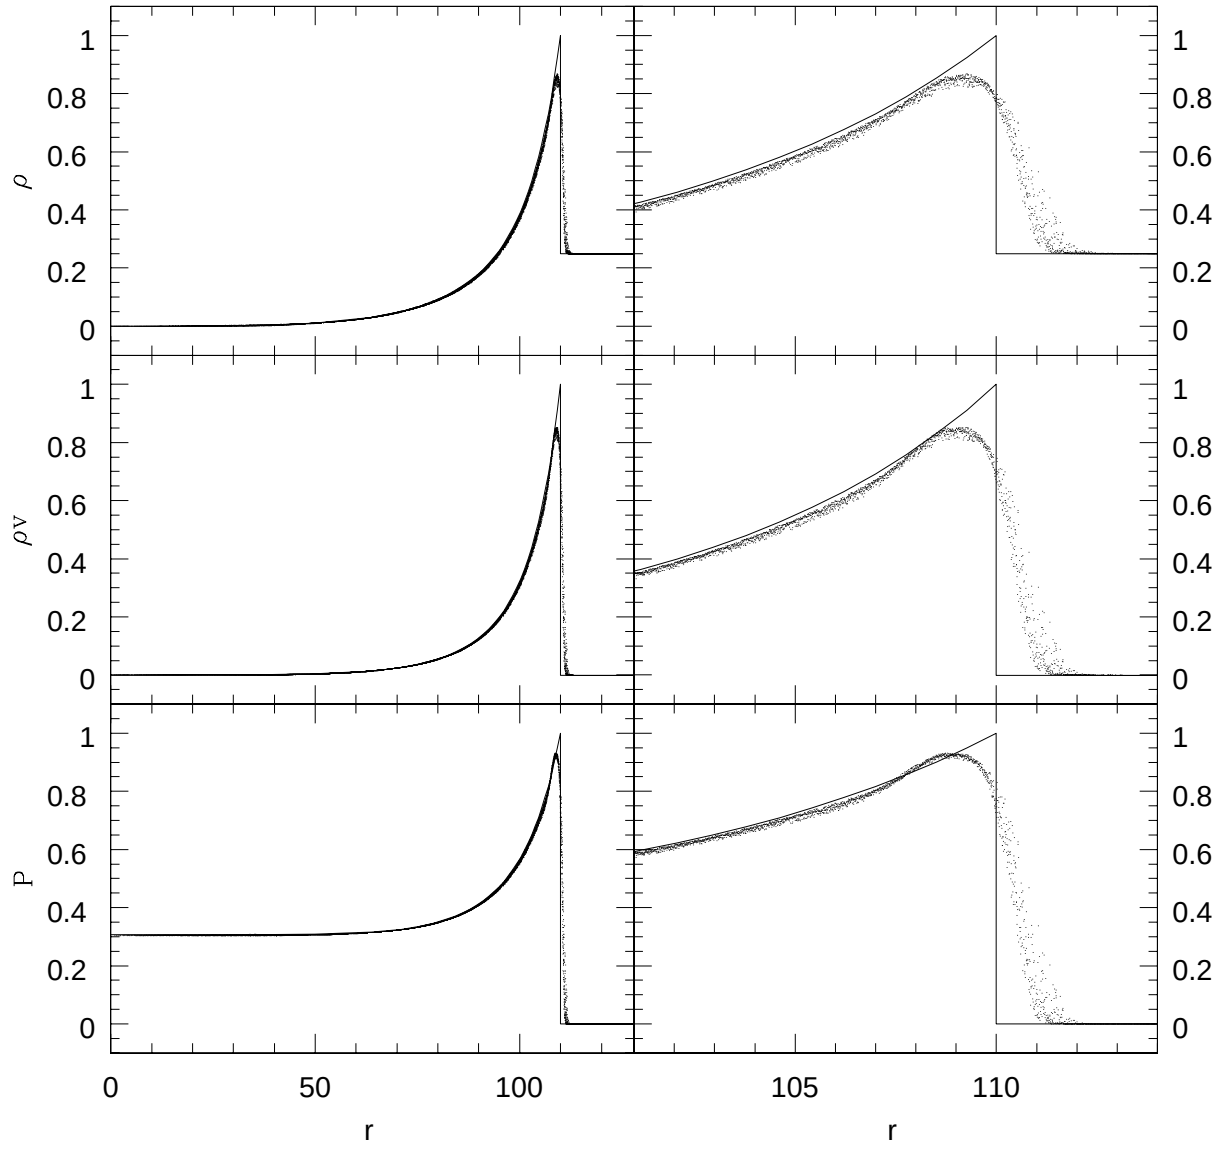

Figure 2.5 Three-dimensional Sedov-Taylor blast-wave test conducted with the relaxing TVD algorithm on a 256^3 grid. The data points are taken from a random subset of cells and the solid lines are the analytical self-similar solutions. The resolution of the shock front is roughly two grid cells and the anisotropic scatter is less than one grid cell, which is close to optimal for a grid-based scheme.

constant density $\rho_1 = 1$. At time $t = 0$, a supply of thermal energy $E_\circ = 10^5$ is injected into one cell at the centre of the box. The simulation is stopped at time $t = 283$, in which the shock front has propagated out to a distance of $r_{sh} = 110$ cells from the centre. Figure (2.5) contains plots of the radial distributions of density, momentum density, and pressure, normalized to ρ_2, $\rho_2 v_2$, and P_2 respectively. The data points are taken from a random subset of cells and the solid lines are the analytical Sedov-Taylor solutions. Despite solving a spherically symmetric problem on an explicitly non-rotationally invariant Cartesian grid, the anisotropic scatter is less than one grid cell, which is optimal considering the spatial discretization of the grid. The distance of the shock front from the centre of the explosion as a function of time is indeed given by equation (2.79), demonstrating that the three-dimensional relaxing TVD code ensures the correct shock propagation. The resolution of the shock front is roughly two grid cells. Again, this is close to optimal considering the spatial discretization. The anisotropic scatter and shock broadening are approximately fixed in grid units, regardless of the duration of the simulation. However, the numerical shock jump values of ρ_2, v_2, and P_2 are resolution dependent and come close to the theoretical values for our test with 256^3 cells. Note that since this is a self-similar problem, running the simulation longer is equivalent to increasing the resolution and therefore, the numerical shock jump values will converge on the theoretical values with time.

2.8 Summary

I have presented several numerical schemes for solving the linear advection equation and have discussed the CFL stability conditions for each scheme. All linear flux assignment schemes are either dispersive or diffusive, but nonlinear schemes like the second-order accurate TVD scheme can provide high resolution capturing of shocks while preventing spurious oscillations. The relaxing TVD scheme for solving the Euler systems of conservation laws is described and discussed in detail. The one-dimensional relaxing TVD algorithm can be generalized to higher dimensions using the dimensional splitting technique. A dimensionally split scheme can use longer time steps and is straightforward to implement in parallel. The second-order accurate scheme provides high resolution capturing of shocks, as has been demonstrated with the Sod shocktube test and the Sedov-Taylor blast-wave test. The research presented in this chapter is published in Trac & Pen (2003).

Chapter 3

Self-gravitating Hydrodynamics for Astrophysical Applications

3.1 Introduction

The focus of this chapter is on self-gravitating hydrodynamics and here I present the addition of a gravity solver to the three-dimensional Eulerian TVD algorithm described in Chapter 2. Poisson's equation is efficiently solved on a grid using fast Fourier transforms (FFTs) at a cost that scales as $\mathcal{O}(N \log N)$, where N is the number of grid cells. In one method, the gravitational forces can be solved for directly and in the other, the potential is first computed and then finite-differenced to obtain the forces. I discuss the advantages and disadvantages of both methods. A three-dimensional, self-gravitating Eulerian hydro code is developed and a non-trivial test involving the advection of a polytrope in hydrostatic equilibrium is performed. This code can be used to simulate astrophysical hydrodynamics and a scientific application is presented in Chapter 4.

3.2 Gravitating Hydrodynamics

For astrophysical applications, both hydrodynamical and gravitational forces are included. The gravitational forces arise from the self-gravity of the fluid and can also come from an external field. The Euler equations with the gravitational source terms included are given as

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j}(\rho v_j) = 0, \tag{3.1}$$

$$\frac{\partial (\rho v_i)}{\partial t} + \frac{\partial}{\partial x_j}(\rho v_i v_j + P\delta_{ij}) = -\rho \frac{\partial \Phi}{\partial x_i}, \tag{3.2}$$

$$\frac{\partial e}{\partial t} + \frac{\partial}{\partial x_j}[(e + P)v_j] = -\rho v_i \frac{\partial \Phi}{\partial x_i}. \tag{3.3}$$

where Φ is the gravitational potential. Poisson's equation

$$\nabla^2\Phi = 4\pi G\rho, \tag{3.4}$$

relates the gravitational potential to the density field. The general solution can be written as the convolution,

$$\Phi(\boldsymbol{x}) = \int \rho(\boldsymbol{x}')w(\boldsymbol{x}-\boldsymbol{x}')d^3x', \tag{3.5}$$

where the isotropic kernel is given by

$$w(r) = -\frac{G}{r}. \tag{3.6}$$

In the discrete case, the integral in equation (3.5) becomes a sum and Poisson's equation can be solved using fast Fourier transforms (FFTs) to do the convolution. The force terms $f_i \equiv -\partial\Phi/\partial x_i$ on the right hand side of the Euler equations are then calculated by finite differencing the potential. The real space kernel $w(r)$ is constructed on the grid with the zero point satisfying

$$f(1) = -\frac{w(2)-w(0)}{2\Delta x} = -G. \tag{3.7}$$

By choosing $w(0) = -2.5G$, the force at a separation of one grid cell is made to be exact. Alternatively, the forces can be calculated directly using the convolution,

$$f_i(\boldsymbol{x}) = \int \rho(\boldsymbol{x}')w_i(\boldsymbol{x}-\boldsymbol{x}')d^3x', \tag{3.8}$$

where the anisotropic force kernels are given by

$$w_i(\boldsymbol{x}) = -G\frac{x_i}{r^3}. \tag{3.9}$$

Both the potential and force methods require one FFT to forward transform the density field, but the potential method only requires one inverse transform to recover the potential field, while the force method requires three inverse transforms for the force components.

The accuracy of the potential and force methods can be quantified with a pair-wise force test. One cell in an empty grid is assigned unit density and the gravitational force field around this point source is measured. The force method comes at the cost of two extra FFTs, but it exactly reproduces the pair-wise inverse-square law on the grid. The potential method is computationally less expensive, but the finite differencing degrades the pair-wise force resolution by a few grid cells. The relative force error is defined as

$$\delta f \equiv \frac{|\boldsymbol{f}| - f(r)}{f(r)}, \tag{3.10}$$

where $\boldsymbol{f}$ is the measured force on the grid and $f(r) = G/r^2$ is the inverse-square law. Figures 3.1(a) and 3.1(b) are scatter plots of $|\boldsymbol{f}|$ and $|\delta f|$ in cells around the point source, calculated using

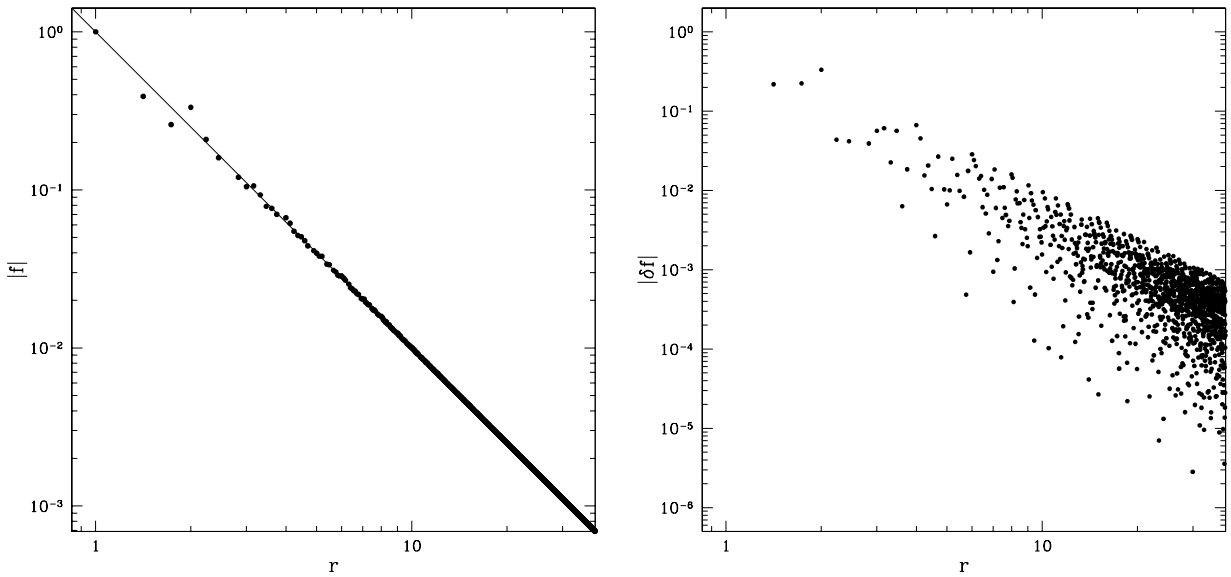

Figure 3.1 A pair-wise force test of the potential method. Shown in (a) is a scatter plot of the force in cells around a point source, while plotted in (b) is the relative error. The finite differencing degrades the pair-wise force resolution by a few grid cells.

the potential method. While the pair-wise force is not exact at small separations, the net force on any given cell is in general still highly accurate. This is demonstrated in the next section of this chapter. In practice, the relatively small accuracy trade-off of the potential method is preferred over the factor of 2 increase in computational work and time of the force method.

The addition of gravitational source terms in the Euler equations is easily handled using the operator splitting technique described in Chapter 2. Consider the double sweep,

$$\boldsymbol{u}^{t+2\Delta t} = \boldsymbol{L}_x\boldsymbol{L}_y\boldsymbol{L}_z\boldsymbol{L}_g\boldsymbol{L}_g\boldsymbol{L}_z\boldsymbol{L}_y\boldsymbol{L}_x\boldsymbol{u}^t, \tag{3.11}$$

where the operator $\boldsymbol{L}_i$ represents the updating of $\boldsymbol{u}$ by including the flux in the i direction and the operator $\boldsymbol{L}_g$ represents the gravitational acceleration of the fluid. In the gravity step, the density distribution $\rho^{t+\Delta t}$ does not change and only the fluid momenta and total energy density are updated. The forward and reverse gravity operations are consecutive and can be grouped into a single operation since the density field and the corresponding force fields do not change during the gravity step.

3.3 Self-gravitating Polytropes

Polytropes are self-gravitating gas spheres in hydrostatic equilibrium with a pressure and density relation given by

$$P = K\rho^{1+1/n}, \tag{3.12}$$

where n is the polytropic index and K is a constant. The density profile $\rho(r)$ is determined by solving the equations of hydrostatic equilibrium. In spherical coordinates, the equation of force balance, the continuity equation, and Poisson's equation are given as

$$\frac{dP}{dr} = -\frac{GM(r)}{r^2}\rho, \tag{3.13}$$

$$\frac{dM}{dr} = 4\pi r^2 \rho, \tag{3.14}$$

$$\frac{1}{r^2}\frac{d}{dr}\left(r^2\frac{d\Phi}{dr}\right) = 4\pi G\rho. \tag{3.15}$$

and they govern hydrostatic equilibrium. If one introduces the definitions

$$\rho = \rho_0\theta^n, \tag{3.16}$$

$$r = a\xi, \tag{3.17}$$

where

$$\rho_0 \equiv \rho(0), \tag{3.18}$$

$$a \equiv \left[\frac{(n+1)K}{4\pi G}\rho_0^{1/n-1}\right]^{1/2}, \tag{3.19}$$

then the equations of hydrostatic equilibrium can be reduced to the Lane-Emden equation

$$\frac{1}{\xi}\frac{d}{d\xi}\left(\xi^2\frac{d\theta}{d\xi}\right) = \frac{d^2\theta}{d\xi^2} + \frac{2}{\xi}\frac{d\theta}{d\xi} = -\theta^n. \tag{3.20}$$

Both θ and ξ are dimensionless. The Lane-Emden equation is subject to two boundary conditions. By definition $\theta(0) = 1$ since the central density is parametrized by ρ_0, and by symmetry $\theta'(0) = 0$ since the net gravitational force at the center of the polytrope is zero. Note that the radius of the polytrope is taken to be $R = a\xi_1$, where ξ_1 is the first root of $\theta(\xi)$, and the constant K is then given by

$$K \equiv \frac{4\pi G}{n+1}\rho_0^{1-1/n}\left(\frac{R}{\xi_1}\right)^2. \tag{3.21}$$

Analytical solutions can be found for polytropic index values of $n = 1$ and $n = 5$ while numerical techniques must be used for all other cases.

3.3.1 Net Force Test

Calculating the force directly reproduces the inverse-square law exactly on the grid, while computing the potential first and then finite differencing to get the force degrades the pair-wise force resolution by a few grid cells. While the pair-wise force is not exact for small separations, the net force is in general still highly accurate with the potential method. This is demonstrated with the following test where polytropes with $n = 3$ and radii $R = 16$, 32, 64, and 128 grid cells are constructed and the net force on each cell is measured. Figure 3.2 is a scatter plot of the relative force errors $|\delta f|$. At small radii, the relative force errors δf have the largest amplitude and are predominantly negative, as expected since the potential method is known to soften the force. For each successive doubling of the radius of a polytrope, the grid spacing decreases by a factor of 2 and the errors decrease by approximately a factor of 4. At sufficiently high resolution, the force errors become tolerably small and the potential method becomes preferable over the computationally intensive force method.

3.3.2 Advection Test

A non-trivial test of a self-gravitating Eulerian hydro code is the advection of an object in hydrostatic equilibrium. The challenge is to maintain the equilibrium profile as the object traverses the grid over a large number of time steps. Polytropes with $n = 3$ and radii $R = 48$ and $R = 96$ grid cells are placed on a grid with the some initial momentum in both the x and y direction, but zero momentum in the z direction. The lower resolution polytrope is advected a total distance of $128\sqrt{2}$ grid cells in the $x - y$ plane in over 800 time steps while the higher resolution one is advected a total distance of $256\sqrt{2}$ grid cells in over 1600 time steps. The polytropes traverse a distance of $(8\sqrt{2}/3)R$ in approximately 5 physical hours. In the next chapter, I will present simulations of two polytropes that are orbiting in the $x - y$ plane and on an off-axis collision course. These advection tests will be used to quantified the numerical limitations of the self-gravitating hydro code.

In Figure 3.3(a) the mass profiles for the lower resolution polytrope are shown. Hydrostatic equilibrium is broken and the polytrope diffuses outwards. The central density and pressure have decreased and the central entropy has increased. The problem lies not in the advection over the grid since even a static test produces similar results. It also does not arise from errors in the force calculations since the force is exact to a few tenths of a percent. The problem originates in constructing a spherically symmetric object on a Cartesian grid that is explicitly nonrotationally invariant. Hydrostatic equilibrium is broken because the balance of radial forces is offset on the Cartesian grid. However, this problem can be alleviated by increasing the resolution. In Figure 3.3(b) the mass profiles for the higher resolution polytrope have not changed by much after the advection step. Some small amount of diffusion is observed in the core of the polytrope. For mass $M < 5\%$, the decrease in density and pressure and the increase in entropy is $< 5\%$.

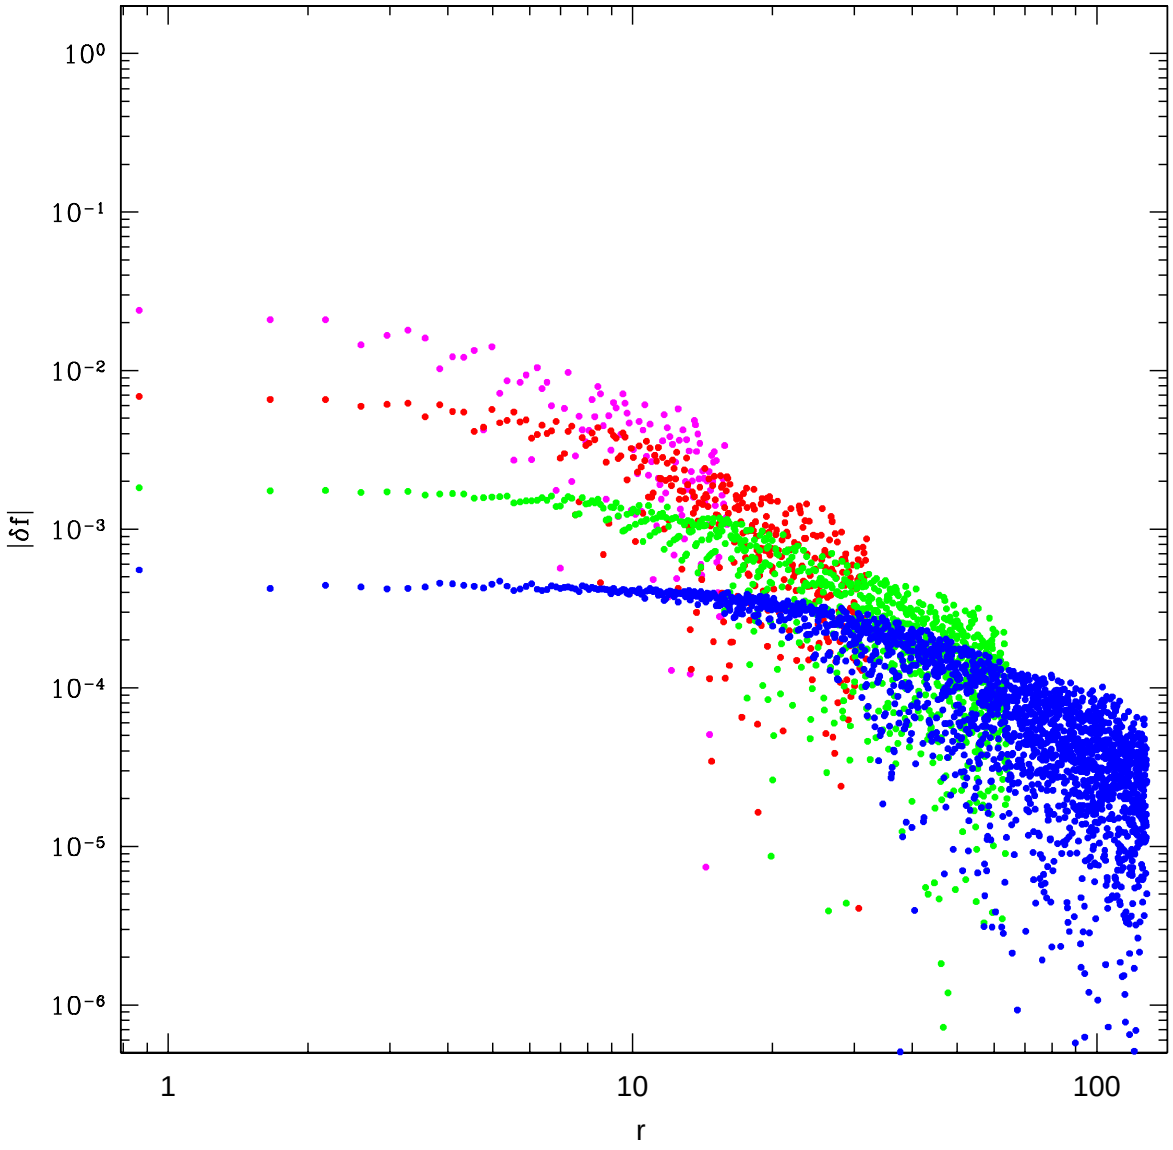

Figure 3.2 A net force test of the potential method. Polytropes with $n = 3$ and radii $R = 16$ (magenta), 32 (red), 64 (green), and 128 (blue) grid cells are constructed and the net force on each cell is measured. For each successive doubling of the radius of a polytrope, the grid spacing decreases by a factor of 2 and the relative force errors $|\delta f|$ decrease by approximately a factor of 4.

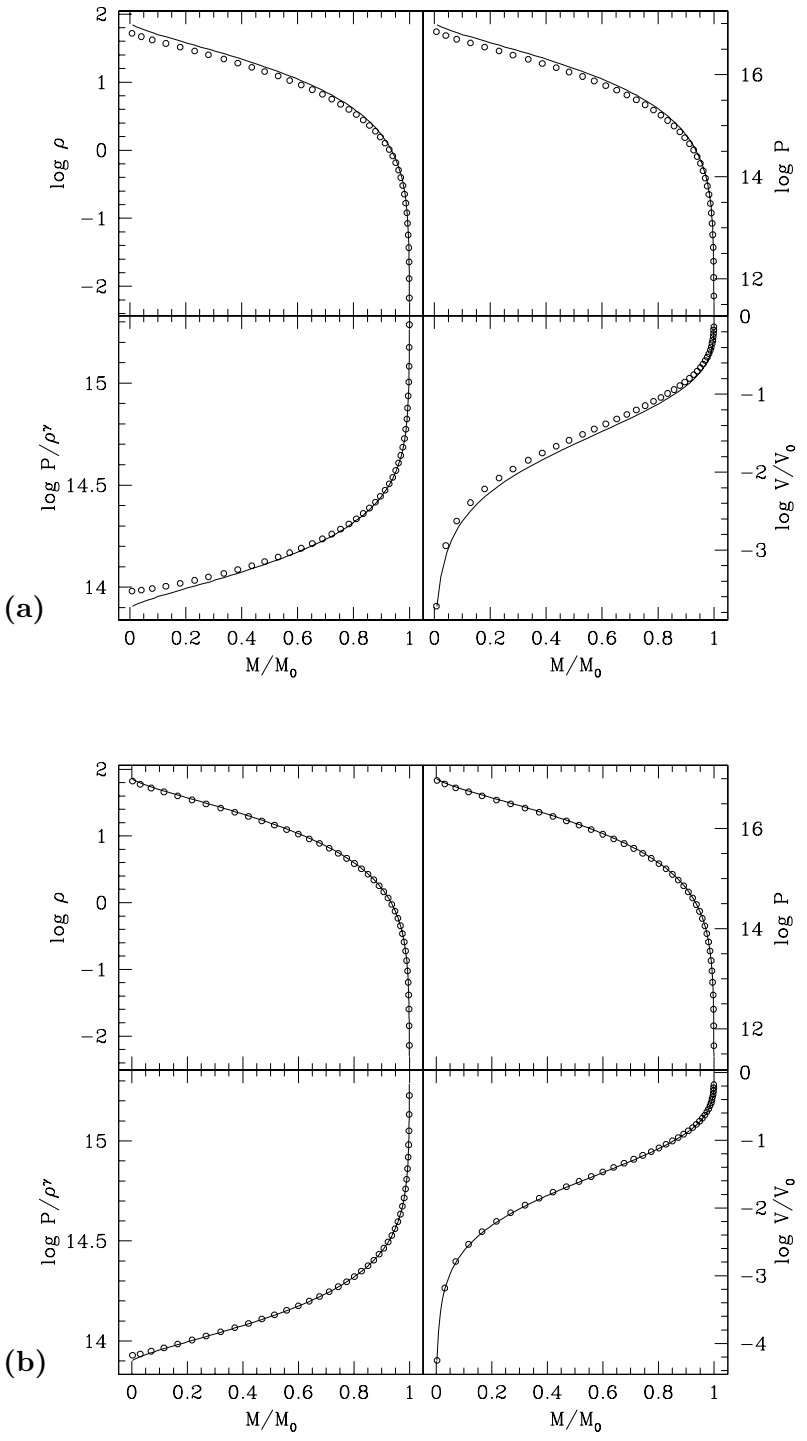

Figure 3.3 Advection of self-gravitating polytropes with $n = 3$ and radii $R = 48$ (a) and $R = 96$ (b) grid cells. The solid lines are the initial mass profiles while the open circles represent the mass profiles after the polytropes have been advected a total distance of $(8\sqrt{2}/3)R$ grid cells in the $x - y$ plane. Hydrostatic equilibrium is maintained in the higher resolution polytrope, but the lower resolution one experiences diffusion, particularly in the core.

3.4 Summary

I have described two methods for constructing a gravity solver for the three-dimensional Eulerian TVD hydro code. The potential method requires half the number of FFTs as does the force method, but the savings in computation also come with a reduction in the accuracy of the pair-wise force. While the pair-wise force is not exact at small separations, the net force on any given cell is in general still highly accurate. At sufficiently high resolution, the force errors become tolerably small and the potential method becomes preferable. The self-gravitating hydro code is robust and accurate. A self-gravitating polytrope is demonstrated to retain its hydrostatic equilibrium profile as it is advected over a large number of cells and time steps. The next chapter contains a sample astrophysical application of this code. Some of the work presented in this chapter is published in Trac & Pen (2003).

Chapter 4

Formation of Blue Stragglers through Stellar Mergers

4.1 Introduction

In this chapter, I apply the Eulerian TVD hydro code to simulating the formation of blue straggler stars through mergers of main sequence stars. To date, simulations of this nature have been performed using Lagrangian SPH codes. However, the differences between Eulerian and Lagrangian approaches may lead to very different results on mixing.

The first SPH simulations of stellar mergers was published in Benz & Hills (1987). They used low resolution SPH simulations with $\sim 10^3$ particles to simulate the merging of equal mass $n = 3/2$ polytropes and found that they fully mixed. However, medium resolution SPH simulations with $\sim 3 \times 10^4$ particles of $n = 3/2$ or $n = 3$ polytropes showed only weak mixing (Sills et al., 1997; Sandquist et al., 1997; Sills et al., 2001). It is worth noting that $n = 3/2$ polytropes are more representative of low mass main sequence stars with large convective envelopes, while $n = 3$ polytropes resemble main sequence stars near the turnoff that have little mass in their convective envelopes. High resolution SPH simulations involving $10^5 - 10^6$ particles have now been applied to simulating stellar collisions (Sills et al., 2002). While this particular work focussed mainly on numerical convergence and the rotation of merger remnants, it can be seen that simulations with $< 10^5$ particles do show weaker mixing than higher resolution simulations. Convergence is found for SPH simulations with $> 10^5$ particles.

SPH simulations are popular for studying stellar collisions mainly because the Lagrangian approach follows the fluid flow and all of the computational effort is spent on regions of interest. In addition, the particles of constant mass and known chemical composition can be directly tracked and this is useful when comparing the merger remnant to the parent stars. However, the Lagrangian particle scheme also has its disadvantages. In general, the merging stars process is mostly subsonic and strong shocks are not expected. The weak shocks that do occur are not well resolved by SPH

Table 4.1. Hydrodynamic Simulations of Off-axis Stellar Collisions

Stars	Grid Resolution	$\Delta x_{\rm grid}$ $(R_\odot)$	$N_{\rm particle}$	$m_{\rm particle}$ $(M_\odot)$
YREC main sequence	$256 \times 256 \times 192$	1.99×10^{-2}	2×64^3	3.05×10^{-6}
YREC main sequence	$512 \times 512 \times 384$	9.95×10^{-3}	2×128^3	3.81×10^{-7}
YREC main sequence	$768 \times 768 \times 576$	6.63×10^{-3}	2×192^3	1.13×10^{-7}
YREC main sequence	$1024 \times 1024 \times 768$	4.97×10^{-3}	2×256^3	4.77×10^{-8}
$n = 3$ polytropes	$256 \times 256 \times 192$	1.99×10^{-2}	2×64^3	3.05×10^{-6}
$n = 3$ polytropes	$512 \times 512 \times 384$	9.95×10^{-3}	2×128^3	3.81×10^{-7}
$n = 3$ polytropes	$768 \times 768 \times 576$	6.63×10^{-3}	2×192^3	1.13×10^{-7}

codes. In the absence of shocks, SPH particles will follow flow lines of constant entropy due to the Lagrangian nature of the method. As a result, the particles may experience sedimentation. In addition, the mixing can also depend on the adopted smoothing length and the form of artificial viscosity. For a SPH fluid, the Reynolds number is of order $(N_p/N_s)^{1/3}$, where N_p is the total number of particles and N_s is the number of particles over which the smoothing is done. For $N_p \sim 10^5$ and $N_s \sim 10^2$, the Reynolds number is ~ 10. However, a fluid with a low Reynolds number will tend to experience laminar flow and therefore, SPH simulations may under mix.

In this chapter, the Eulerian simulations used are unique in that the hydro algorithm is coupled to a particle-mesh algorithm that was designed to follow test particles of fixed mass and known chemical composition. This allows one to determine the chemical composition and the mixing of mass in the merger remnant. I describe the numerical methods and discuss numerical convergence. Both the gas and test particles are examined in detail to address the question of mixing during the merging process. Some preliminary work has already been published in Trac & Pen (2003) and detailed work will be presented in Trac, Sills, & Pen (2004). The latter is a comparison of Eulerian and Lagrangian hydrodynamic techniques for modelling collisions between main sequence stars.

4.2 Hydrodynamics Simulations

A three-dimensional Eulerian TVD hydro code is used to simulate the off-axis collision of two equal mass main sequence stars with $M_0 = 0.8 M_\odot$ and $R_0 = 0.955 R_\odot$. The main sequence stars are constructed as $n = 3$ polytropes and also with a realistic model calculated by Alison Sills, using the Yale stellar evolution code (YREC; Guenther et al., 1992). The main sequence stars are found

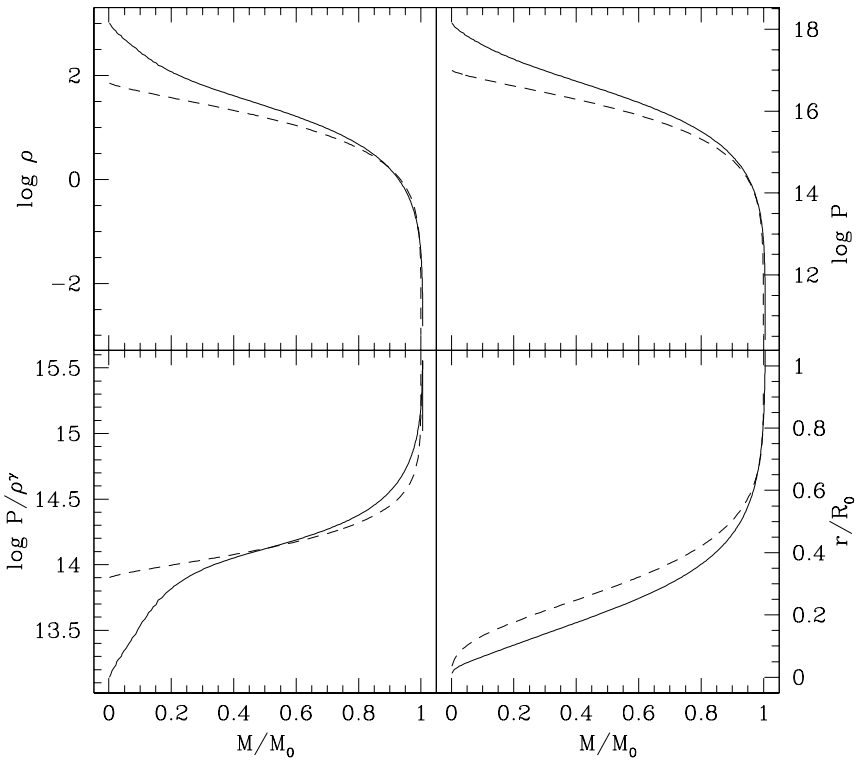

Figure 4.1 Mass profiles of a $M = 0.8M_{\odot}$ and $R = 0.955R_{\odot}$ main sequence star constructed as an $n = 3$ polytrope (dashed lines) and a realistic stellar model (solid lines) calculated with YREC. The realistic model is more centrally condensed with larger gradients than the polytropic one. In both models, approximately 90% of the total mass is contained within $r < 0.5R_0$. Units are cgs.

near the turnoff with an age of approximately 13.5 Gyr. In the central $\sim 25\%$ of the mass, formerly the hydrogen core, the hydrogen has been mostly depleted, leaving mainly helium. In Figure 4.1, the initial mass profiles for the models are compared. The realistic stellar model is more centrally condensed with steeper gradients than the polytropic one. Higher resolution will be required to resolve the core and maintain hydrostatic equilibrium in that case. In both models approximately 90% of the total mass is contained within $r < 0.5R_0$. Main sequence stars near the turnoff have relatively little mass in their convective envelopes. The dynamical time for the stars is defined to be

$$\tau_{dyn} \equiv \frac{1}{\sqrt{G\bar{\rho}}}, \tag{4.1}$$

where $\bar{\rho}$ is the average density and for the chosen parent stars the dynamical times is approximately 1 physical hour. Thus, the modelling of the collision is a hydrodynamic problem.

The collision is simulated on a non-periodic hydro grid with dimensions $nx \times ny \times nz$ cells. The orbital plane is taken to coincide with the $x - y$ plane and the dimensions are chosen such

that $nx = ny = (4/3)nz$. The collision is simulated at low resolution with $nx = 256$, at medium resolution with $nx = 512$, at medium-high resolution with $nx = 768$, and at high resolution with $nx = 1024$. Initially, each parent star has a radius $R = 48$, 96, 144, and 192 grid cells, respectively, for the 4 different resolutions. The simulations are summarized in Table 4.1. The stars are set up on zero-energy parabolic orbits with a pericenter separation equal to $0.25R$ and the initial trajectories are calculated assuming point masses. The chosen orbit is representative of a typical encounter. Note that the FFT Poisson solver can be made non-periodic by using a gravity grid with dimensions that are twice as large as the hydro grid and treating the padded cells as vacuum with zero density. However, this is both memory and work costly. Since the parent stars and remnant star only occupy a fraction of the hydro grid, a more optimal approach is to make the Poisson solver only pseudo non-periodic and use a gravity grid with dimensions that are 1.5 times as large as the hydro grid.

4.2.1 Particle-Mesh Scheme

To address the question of chemical mixing in the remnant star, the hydro code is implemented along with a particle-mesh scheme where test particles of fixed mass can be used to track passively advected quantities like chemical composition. Initially, each parent star is assigned a large number of test particles with known chemical composition and the particles are advected along the fluid velocity field lines. Each parent star is assigned $N_p = 64^3$, 128^3, 192^3, and 256^3 test particles, respectively, for the 4 different resolutions.

The particle-mesh scheme is coupled to the Eulerian TVD scheme in the following manner. Recall from Chapter 2 that the Eulerian TVD scheme is dimensionally split and in order to synchronize the particles with the fluid, the particle-mesh scheme must also be dimensionally split. When solving for fluid flow in the x direction, the second-order accurate solution for the mass continuity equation (Eq. [2.1]) is obtained via a Runge-Kutta time integration with a half step,

$$\rho_n^{t+\Delta t/2} = \rho_n^t + \left(\frac{F_{n+1/2}^t - F_{n-1/2}^t}{\Delta x} \right) \frac{\Delta t}{2}, \tag{4.2}$$

followed by a full step,

$$\rho_n^{t+\Delta t} = \rho_n^t + \left(\frac{F_{n+1/2}^{t+\Delta t/2} - F_{n-1/2}^{t+\Delta t/2}}{\Delta x} \right) \Delta t. \tag{4.3}$$

where the fluxes $F = F(\rho, v_x)$ are defined at cell boundaries rather than cell centers n. While the fluid velocity at cell centers can be defined by dividing the cell-averaged momentum by the cell mass, a more robust measure is that of the velocity at which mass is flowing across cell boundaries. The fluid velocity at cell boundaries can be calculated as

$$v_{x,n+1/2}^t = \frac{F_{n+1/2}^t}{(\rho_{n+1}^t + \rho_n^t)/2}. \tag{4.4}$$

and both the half-step velocity field $v(\boldsymbol{x})^{t}_{x,n+1/2}$ and full-step velocity field $v(\boldsymbol{x})^{t+\Delta t/2}_{x,n+1/2}$ are needed for the particle-mesh scheme. The particles are then advected in the x direction using a second-order Runge-Kutta time integration with a half step,

$$x^{t+\Delta t/2} = x^{t} + v(\boldsymbol{x}^{t})^{t}_{x}(\Delta t/2), \tag{4.5}$$

followed by a full step,

$$x^{t+\Delta t} = x^{t} + v(\boldsymbol{x}^{t+\Delta t/2})^{t+\Delta t/2}_{x}\Delta t. \tag{4.6}$$

The velocity of the particle at position $\boldsymbol{x}$ is interpolated from the fluid velocity field using a 'cloud-in-cell' (CIC) technique (Hockney & Eastwood, 1988). The CIC interpolation is also used to determine local quantities such as density, pressure, and entropy associated with each particle. The coupling of the particle-mesh algorithm to the Eulerian TVD code has the benefit of being able to track fluid variables like in an SPH code, while still being able to solve the fluid equations using the Eulerian approach.

4.3 Formation of Blue Stragglers

At time $t = 0$, the parent stars are placed on zero-energy parabolic orbits with a pericenter separation of $0.25R$. Initially, they are separated by a distance of $\Delta d = 3.75R = 3.58R_{\odot}$ and each are moving at a speed of ~ 150 km/s relative to the center of mass of the collision. Snapshots of the merging process are shown in Figure A.1 in Appendix A. The off-axis collision does indeed produce a single remnant star that establishes hydrostatic equilibrium in less than $10\tau_{dyn}$. During the merging process, the outer envelopes of the parent stars are shock-heated and approximately 10% of the total mass gets ejected into the interstellar medium (ISM). Figure 4.2 is a snapshot from the highest resolution simulations taken at $t = 10\tau_{dyn}$. The density contour plots are from thin slices 4 cells thick through the $x-y$, $x-z$, and $y-z$ mid-planes. For both the polytrope and main sequence star mergers, the resulting remnant is a rotating oblate that shows no indication of a disc.

4.3.1 Mass Profiles of the Gas

The mass profiles of the merger remnant from the off-axis collision of equal mass $M_0 = 0.8M_{\odot}$ $n = 3$ polytropes are shown in Figure 4.3. The medium (512) and medium-high (768) resolution simulations have converged everywhere except for the inner few percent of the mass. Thus the results from the medium-high resolution simulation are robust enough for studying the remnant mass profile of the merger of $n = 3$ polytropes. However, we will subsequently see that to address the question of mixing, we may require higher resolution for convergence in the core.

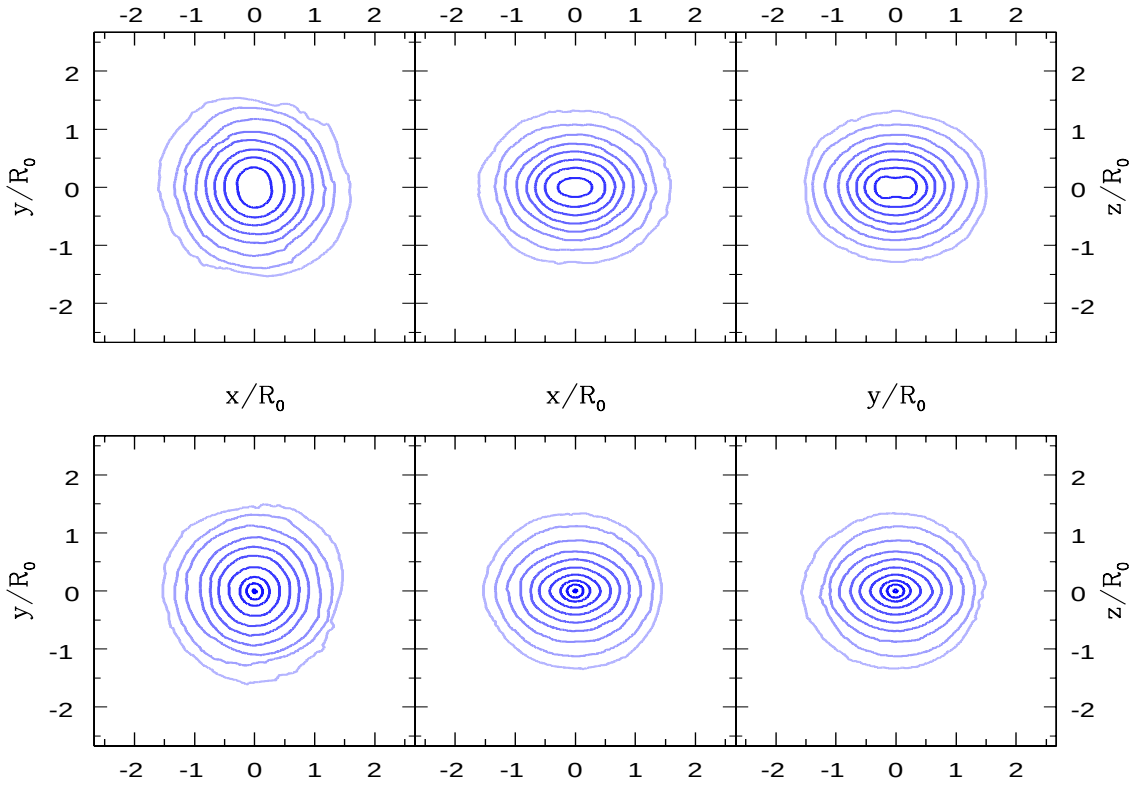

Figure 4.2 Projections of the merger remnant from the medium-high (768) resolution collision of polytropes (upper) and the high (1024) resolution collision of main sequence stars (lower). The projections are from thin slices 4 cells thick through the $x-y$, $x-z$, and $y-z$ mid-planes. Density contours are spaced logarithmically with two per decade, starting from a minimum density of 0.01 g/cm^3 at the edge. There is no indication of a disc in either simulations.

The merger of two $n = 3$ polytropes does not produce a polytrope. The central density and pressure in the core of the remnant have decreased by a factor of 1.9 and 1.8, respectively, relative to the corresponding values in the parent polytropes. The entropy floor has risen by factor of 1.6. This change in entropy can come from a combination of relaxation, shock heating, and mixing. For an $n = 3$ polytrope, only a small amount of relaxation of $< 5\%$ is found at medium resolution and negligible at higher resolutions. The mass profiles of the polytrope are quite featureless and hydrostatic equilibrium is robustly maintained for the 512 resolution and higher. Shock heating is expected to be minimal in the core, so the change in entropy suggests that some mixing has taken place. This question is examined in further detail in a later section.

The merger remnant has a mass of $M_* = 1.48 M_\odot$ and 7% of the mass from the parent polytropes have been ejected into the ISM. The remnant has significantly expanded and is less centrally concentrated. In the parent polytropes, 90% of the mass is found in the inner 13% of the total

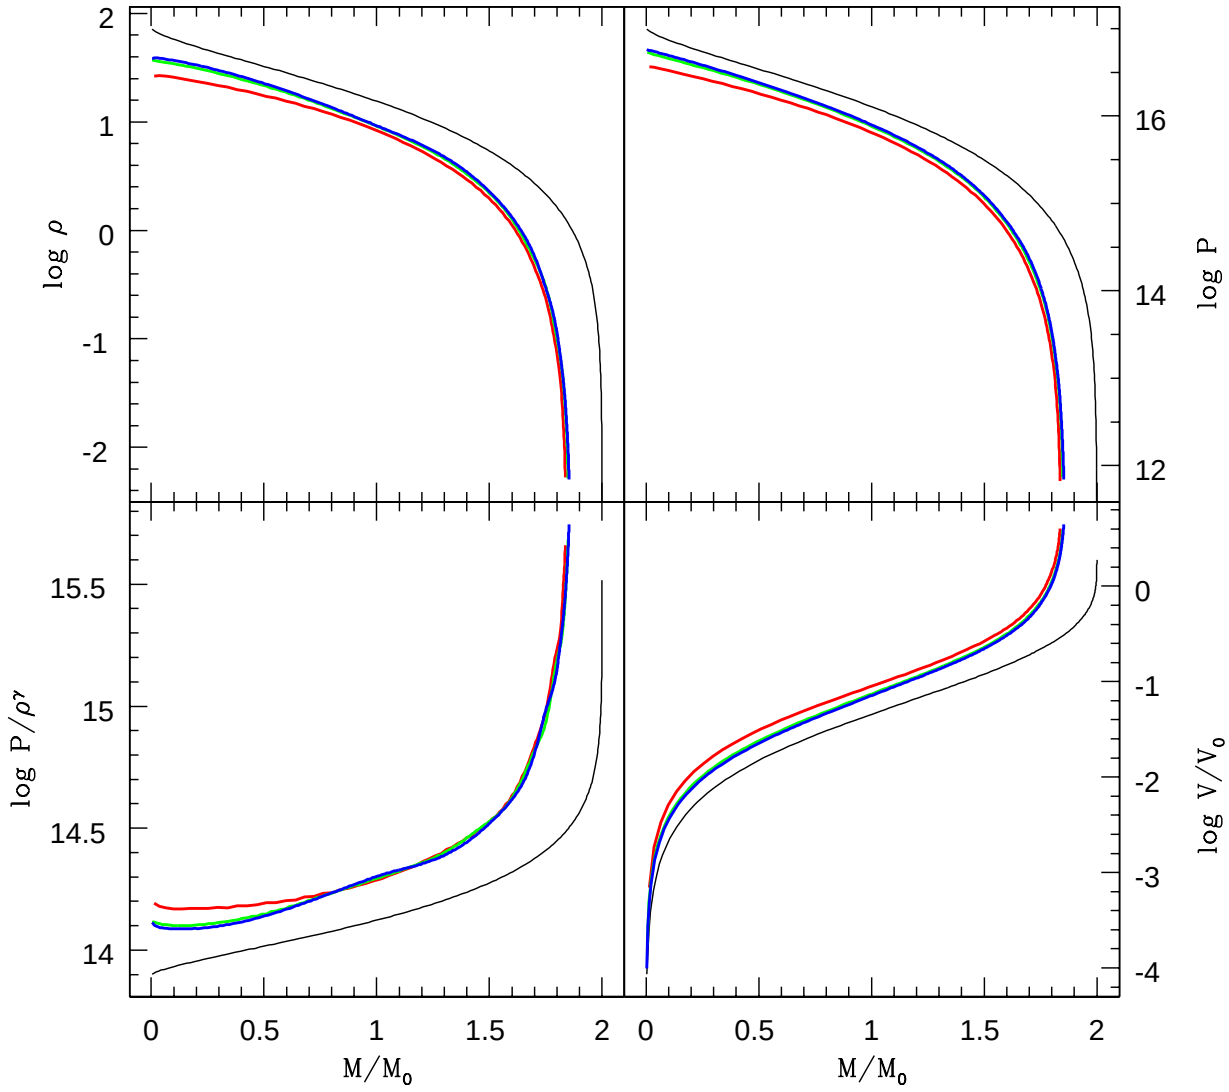

Figure 4.3 Mass profiles of the merger remnant from the off-axis collision of equal mass $M_0 = 0.8M_\odot$ $n = 3$ polytropes. The 256 (red), 512 (green), and 768 (blue) results are plotted in comparison with the initial mass profile of the parent polytropes (solid). Units are cgs. The mass profiles from the 512 and 768 simulations are very similar and the mass profiles from the 768 simulation have converged for $M/M_* \gtrsim 1\%$, where $M_* = 1.85M_0$ is the mass of the remnant.

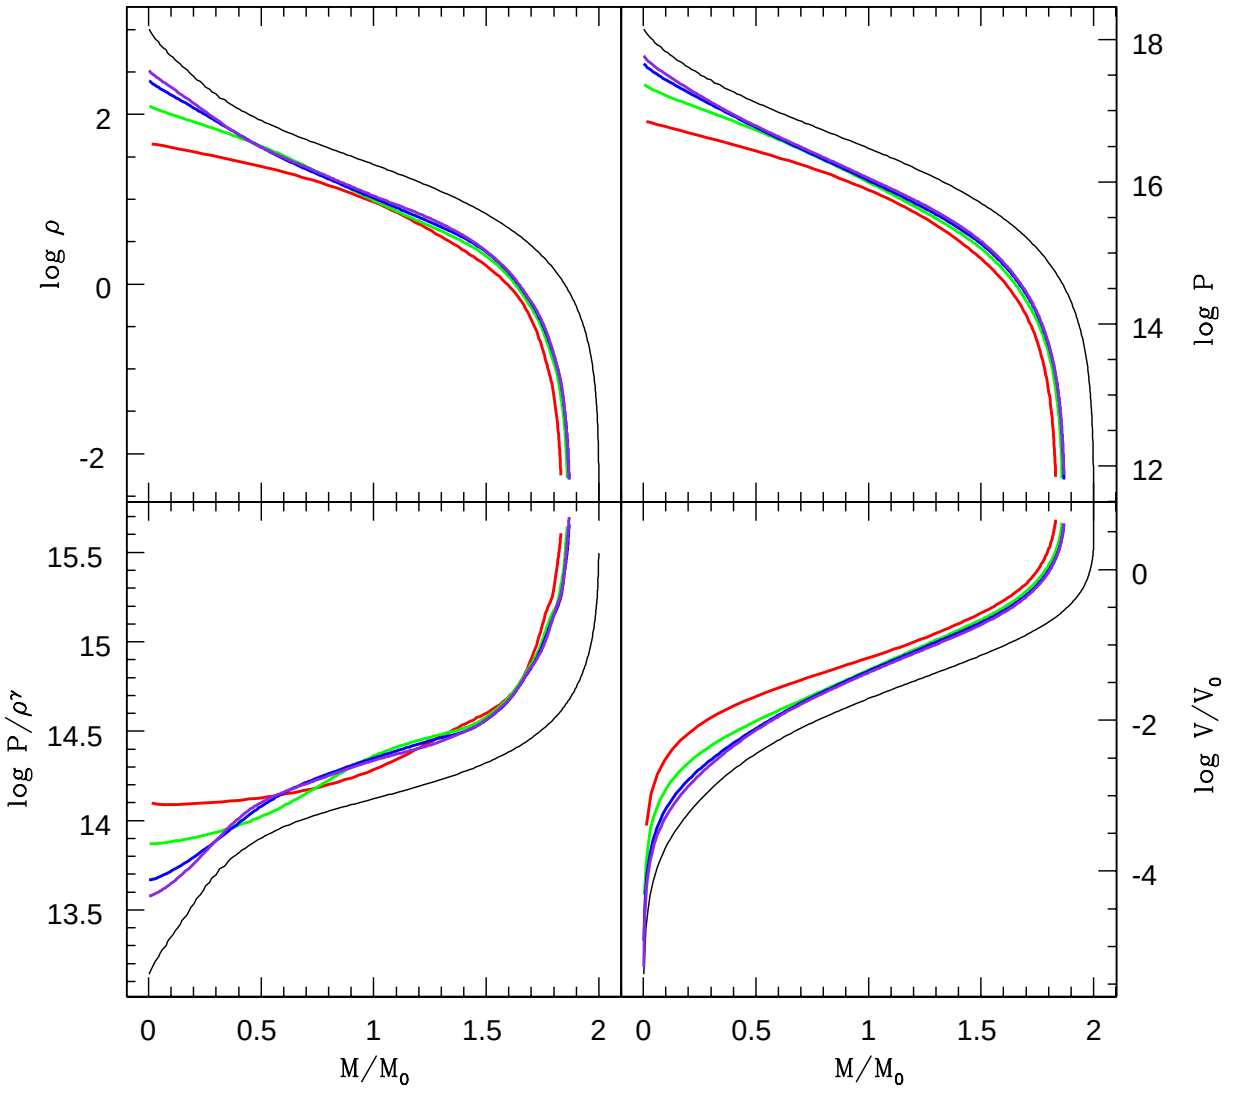

Figure 4.4 Mass profiles of the merger remnant from the off-axis collision of equal mass $M_0 = 0.8M_\odot$ main sequence stars. The 256 (red), 512 (green), 768 (blue), and 1024 (violet) results are plotted in comparison with the initial mass profile of the parent main sequence stars (solid). Units are cgs. The mass profiles from the highest resolution simulation are expected to have converged for $M/M_* \gtrsim 1\%$, where $M_* = 1.85M_0$ is the mass of the remnant.

volume, while in the remnant, 90% of the mass is found in the inner 40% of the total volume. Shock heating or mixing has significantly increased the entropy in the outer regions of the remnant.

Plotted in Figure 4.4 are the mass profiles of the merger remnant from the off-axis collision of equal mass $M_0 = 0.8M_\odot$ main sequence stars. Unlike the polytropes mergers, no convergence is found for the medium (512) and medium-high (768) resolution simulations. Inspection of the density and pressure profiles for signs of convergence is not as robust as using the entropy profile. At medium resolution, the density and pressure profiles appear to have converged for the outer $\sim 75\%$ of the mass, but the slight differences are magnified in the entropy profile, where the real convergence is found for the outer $\sim 55\%$. The medium-high entropy profile has converged for the outer $\sim 85\%$ of the mass when compared with the high (1024) resolution simulation, which is expected to be robust on the order of a percent for looking at mass profiles.

In the high resolution simulation, the entropy floor in the remnant has risen by a factor of 2.7, but some of the entropy structure from the core of the parent main sequence stars is still retained. This change in entropy is likely to come from a combination of relaxation, mixing, and weak shock heating. Relaxation is more significant in this case compared to the polytropes because of the sharp gradients in the mass profiles of the main sequence stars. The central density and pressure in the core of the remnant have decreased by a factor of 3.0 and 2.3, respectively, though some structure is also retained from that of the parent stars.

The merger remnant has a mass of $M_* = 1.49M_\odot$ and 7% of the mass from the parent stars have been ejected into the ISM, very similar to the polytropes merger. In the parent stars, 90% of the mass is found in the inner 10% of the total volume, which is slightly more centrally concentrated than the polytropes. In the remnant, 90% of the mass is found in the inner 40% of the total volume, similar to the remnant formed from the polytropes merger.

4.3.2 Mass Profiles of the Particles

In order to check if the particle-mesh scheme is working accurately and the test particles are indeed following the gas, the CIC interpolation technique is used to assign local quantities such as density, pressure, and entropy associated with each particle of fixed mass. If the particles are accurately tracing the gas, then the mass profiles should be the same.

In Figure 4.5, the gas and particles from both the highest resolution polytrope and main sequence stars simulations are compared. The particle-mesh algorithm does an excellent job of advecting the particles along the fluid velocity field lines. In fact, very similar accuracy is found for all lower resolution simulations, a demonstration of the robustness of the method. Thus, the particles can be used with high confidence to determine the chemical composition of the merger remnant.

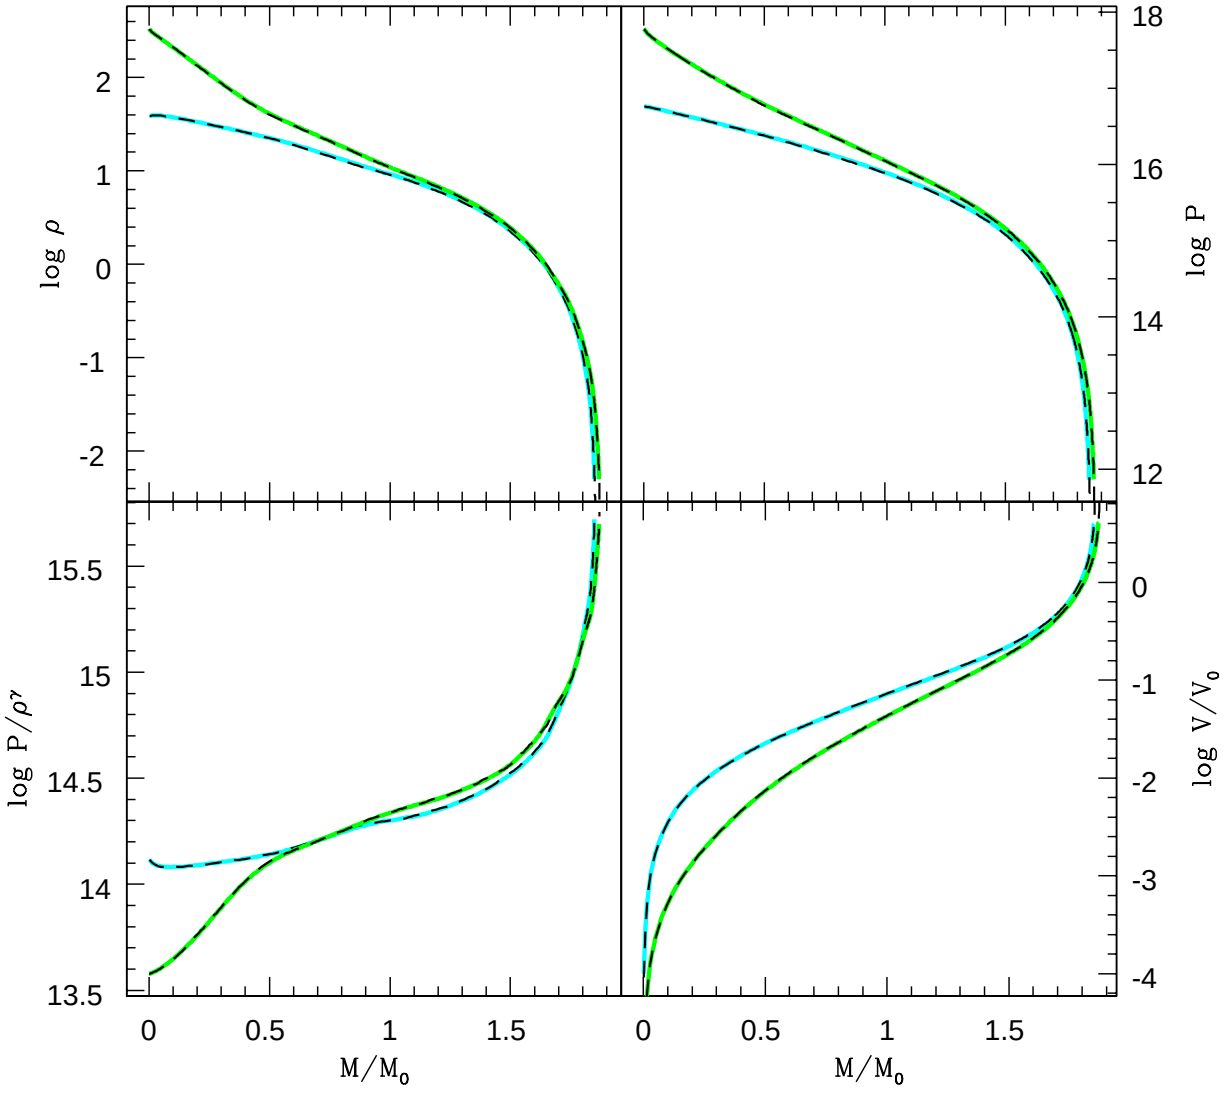

Figure 4.5 Mass profiles from the highest resolution polytrope (cyan) and main sequence stars (green) simulations. The mass profiles constructed for the particles (dashed) are in excellent agreement with the fluid for this resolution and for all lower resolution simulations as well. Units are cgs.

4.3.3 Mass Mixing

One of the major advantages of the Lagrangian SPH codes is that the particles of constant mass can be directly tracked allowing one to compare the initial and final entropies, the mixing of mass, and the final chemical composition of the merger remnant. This now is also possible for Eulerian hydro codes coupled with a particle-mesh algorithm. The following discussing on mixing is based on results from the medium-high (768) polytropes simulation and the high (1024) resolution main sequence stars simulation. Previously, I showed that the remnant mass profiles from these simulations have converged down to the inner $\sim 1\%$ of the mass.

In Figure 4.6 the initial and final entropies $A \equiv P/\rho^{\gamma}$ of the particles are compared. Only a subset of the particles are shown with the points in green. For the polytropes merger, the entropy in the center of the initial polytropes has a minimum value $A_{\rm min} = 10^{13.90}$, while the entropy in the center of the remnant has increased by a factor of 1.6. The minimum entropy in the core of the initial main sequence stars is $A_{\rm min} = 10^{13.15}$, but it can be seen in the plot that only a small fraction of the particles and mass are found in the core. In the remnant, the entropy floor has increased by a factor of 2.7, though this increase may be contributed to by relaxation. In general, the entropy is seen to increase everywhere. However, it is difficult to tell from this plot whether the entropy increase is due to mixing or shock heating. In Trac, Sills, & Pen (2004), these entropy plots from Eulerian simulations will be compared directly with corresponding results from Lagrangian SPH simulations to study how the treatment of shock heating and mixing differ.

The mixing of mass is closely examined with the following procedure. The particles can be ranked order using the effective potential or entropy and each particle is assigned a mass coordinate based on its ranking. This mass coordinate is simply the value of the cumulative mass.

For the highest resolution polytropes merger, the initial mass coordinate M_i is plotted against the final coordinate M_f in Figure 4.7. Only a small subset of the particles are shown by the points in green. The reversal of the axes is used because I wish to see where the mass in a given region of the remnant is coming from. All of the particles are also binned in terms of the final mass M_f and for each bin, the 50% value of M_i is found. The median curve is illustrated by the solid blue line. Similarly, the long dashed lines enclose 50% of the scatter, while the short dashed lines enclose 90% of the scatter.

Numerical scatter can arise from relaxation and advection and this amount of scatter needs to be quantified. This is done with the advection test described in Chapter 3, where a single polytrope is advected across the grid in multiple directions over a large number of time steps. The black points represent the numerical scatter expected even in the absence of mass mixing during the merger. In the absence of mixing the median curve should be a straight line lying in the center of the numerical scatter. And in the presence of complete mixing, the phase space should be uniformly filled and the median curve should be flat with a constant value of $M_i/M_* = 0.5$, where $M_* = 1.85M_0$ is the

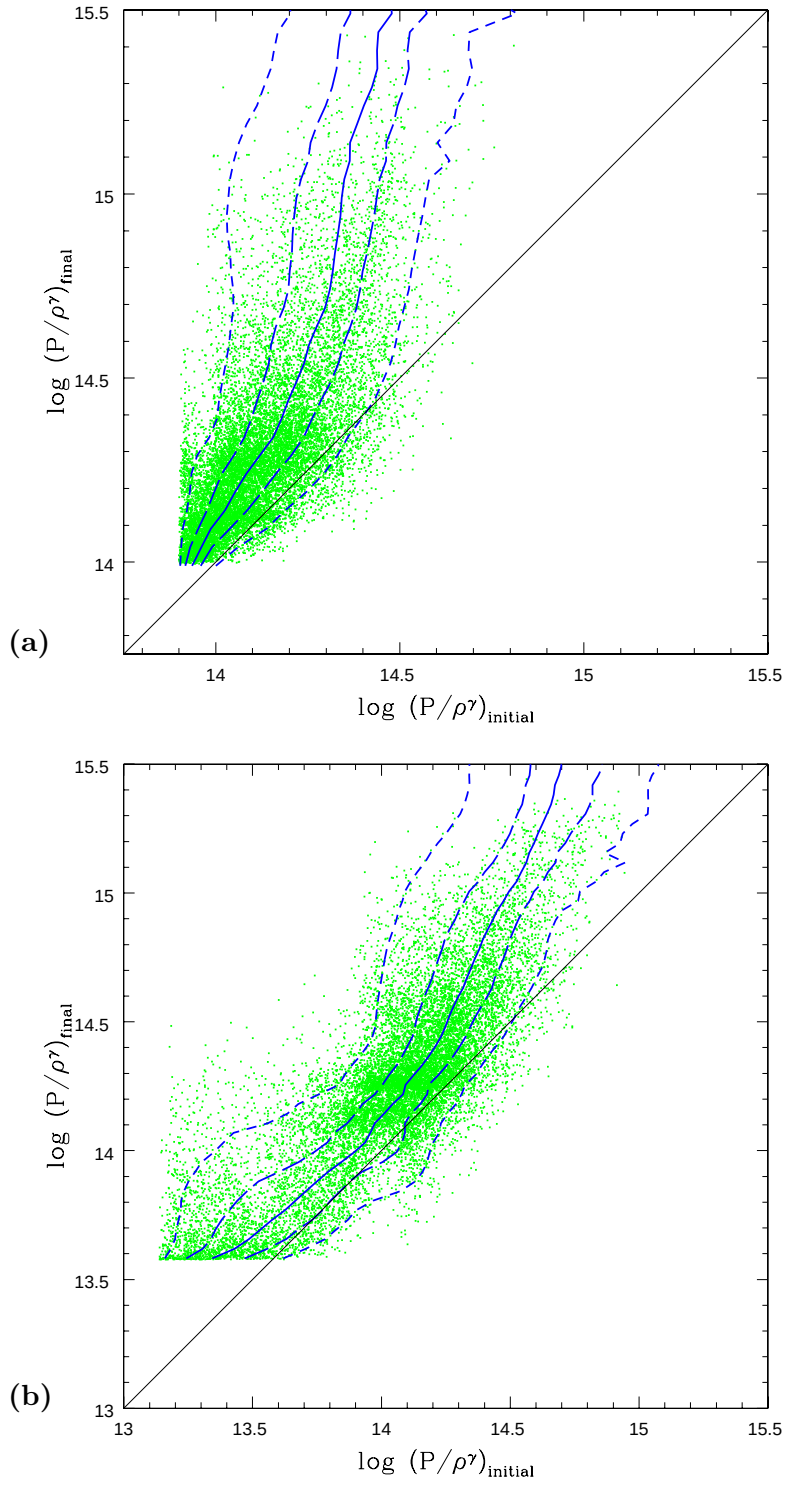

Figure 4.6 Entropy increase found in the merger of polytropes (a) and the merger of main sequence stars (b). The initial and final entropies are compared for a subset of particles shown in green. The median (solid), 50% scatter (long dashed), and 90% scatter (short dashed) in the entropy are given by the blue lines. In (b), only a small fraction of the particles and mass are found in the core where low entropy is found.

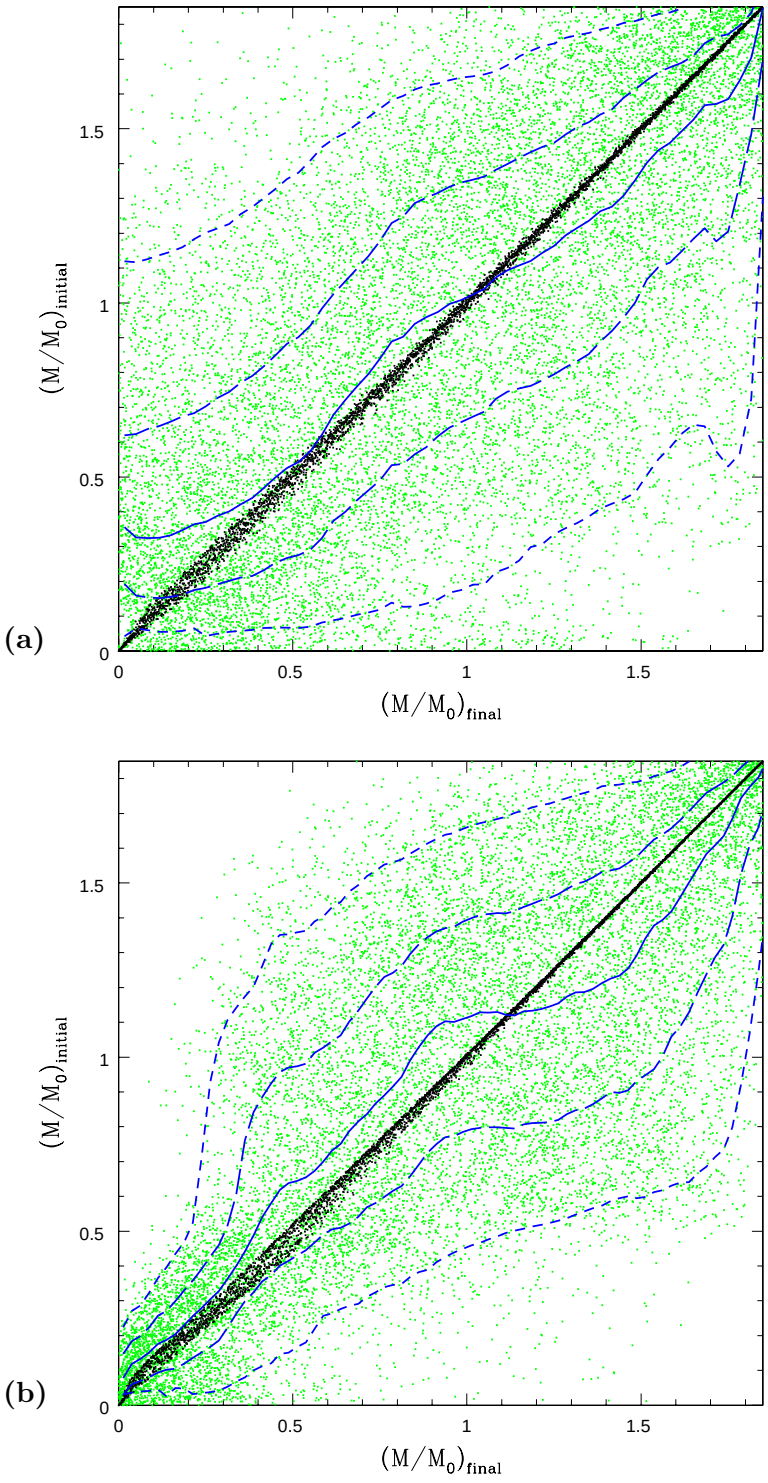

Figure 4.7 The mixing of mass in the merger of polytropes (a) and the merger of main sequence stars (b). The initial mass coordinate M_i is plotted against the final mass coordinate M_f for a subset of particles (green points) in the remnant. The median (solid), 50% scatter (long dashed), and 90% scatter (short dashed) in M_i are given by the blue lines. Numerical scatter estimated with the advection test are shown by the black points. Substantial mixing is seen everywhere for the polytropes merger while for the stellar merger, significant mixing is only seen outside of the core.

mass of the remnant.

The off-axis collision of equal mass $n = 3$ polytropes is observed to produce substantial mixing in Figure 4.7(a). For the inner region of the remnant given by $M_f/M_* < 0.25$, the median curve for M_i rises above the one-to-one line as mass is mixed into the center, mainly from above as the mixing from below has a hard limit. Considering this region as a whole, 44% of the contained mass is mixed in from outside, specifically from the mass range $0.25 < M_i/M_* < 0.60$ in the initial polytropes. For the mass range $0.25 < M_f/M_* < 0.75$ in the remnant, roughly the same amount of mass gets mixed in from above as from below. And for the outer $M_f/M_* > 0.75$, the median curve falls below the one-to-one line as more mass is mixed into it from below since mixing from above is physically limited by the finite star. The average amount of mixing for the entire remnant can be estimated using the average width ΔM_i of the zone enclosing half of the scatter, as bounded by the long dashed lines. The average width is measured to be $\Delta M_i = 30\%$ and for complete mixing the width is given by $\Delta M_i = 50\%$. Therefore, the remnant is estimated to be $\sim 60\%$ fully mixed by linearly interpolating.

In Figure 4.7(b), the results from the off-axis collision of equal mass main sequence stars are remarkably different from the polytropes merger. Rather weak mixing is seen in the core of the remnant. Recall that in the parent main sequence stars, the depleted hydrogen core contains $\sim 25\%$ of the mass and significant mixing of hydrogen into the central $\sim 25\%$ of the mass of the remnant is require. Considering the central region $M_f/M_* < 0.25$ as a whole, only 22% is mixed in from outside, compared to 44% for the polytropes merger. The average width of the zone enclosing half of the scatter is measured to be $\Delta M_i = 25\%$ and the remnant is estimated to be $\sim 50\%$ fully mixed, though most of the mixing is seen for $M_f/M_* > 0.25$.

Figure 4.8 is a plot showing the convergence of the mass mixing results. The average mixing percentage in the remnant is given by the short-dashed line and the percentage of mixing into the hydrogen core is given by the long-dashed line. For the main sequence stars merger, the medium-high (768) and high (1024) resolution simulations give very similar results within a few percent and therefore, the numbers previously presented are highly robust. For the polytropes merger, it is not apparent that the medium-high resolution (768) results have converged based on that plot alone since the high (1024) resolution simulation has not been run. However, it is clear from Figures 4.3 and 4.4 that the polytropes merger requires lower resolution than the main sequence stars merger. Since the mixing values for the latter at medium-high resolution is similar to those at high resolution, the mixing values for the polytropes merger at medium-high resolution should be quite robust. In addition, the difference between the medium (512) and medium-high (768) resolution results only differ by a few percent for the polytropes merger.

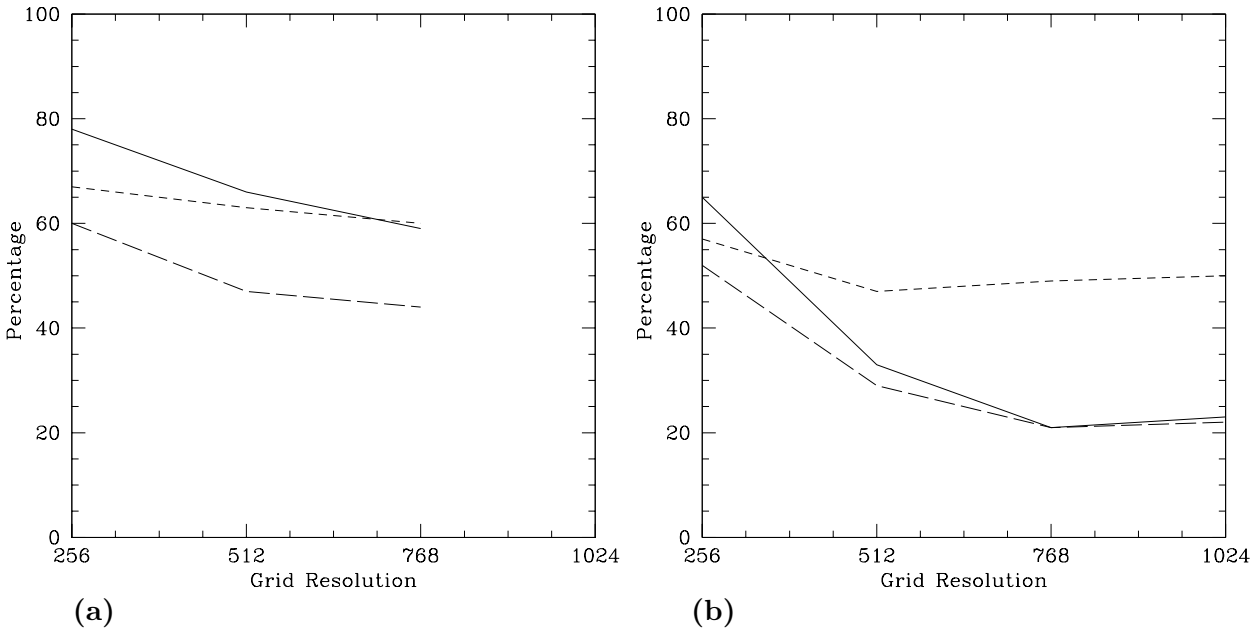

Figure 4.8 Numerical convergence for the polytropes merger (a) and the main sequence stars merger (b). The average mixing percentage (short-dashed lines), the percentage of mixing in the hydrogen core (long-dashed lines), and the relative increase in the hydrogen mass fraction in the hydrogen core of the remnant (solid lines) is plotted as a function of the grid resolution. For the main sequence stars merger, the medium-high (768) and high (1024) resolution simulations have converged for these particular mixing statistics.

4.3.4 Hydrogen mixing

Substantial mixing in the hydrogen core or central 25% of the mass of the remnant from the polytropes merger is expected to significantly change the chemical composition. This is shown to be true in Figure 4.9 where the hydrogen mass fraction X is plotted as a function of mass in the merger remnant and compared to the initial hydrogen mass fraction of the $M = 0.8M_{\odot}$ main sequence stars. For $M < 0.25M_{*}$, the hydrogen mass fraction is found to increase to a minimum value of $X = 0.4$. Considering this region as a whole, the hydrogen mass fraction has on average increased by 59%. This is consistent with the previous finding where for $M < 0.25M_{*}$, 44% of the contained mass is mixed in from outside. Note that the hydrogen profile from the 512 resolution simulation has converged for $M/M_{*} > 10\%$ and the profile from the 768 resolution is expected to be robust down to an inner mass value of order a percent, by extrapolation.

The weak mixing in the core of the remnant from the merger of main sequence stars has only changed the hydrogen mass fraction on average by 23% for the hydrogen core $M < 0.25M_{*}$, which is also consistent with the previous finding where 22% of the contained mass is mixed in from outside. The hydrogen profile from the 768 resolution simulation has only converged for $M > 0.25M_{*}$ and

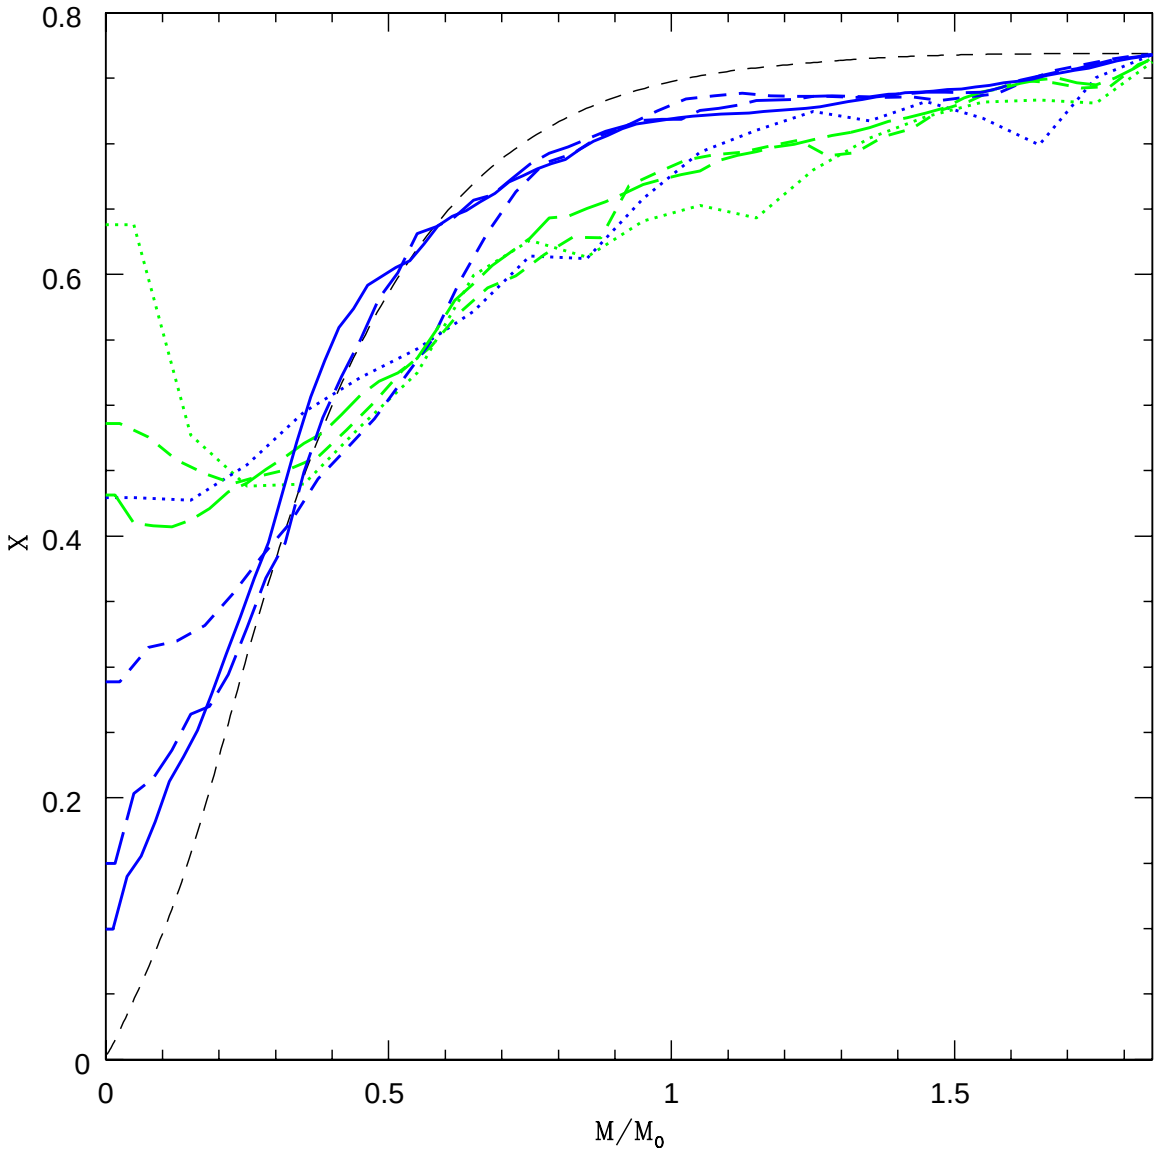

Figure 4.9 The hydrogen mass fraction in merger remnants from the collision of polytropes (green) and the collision of main sequence stars (blue). The 256 (dotted), 512 (short dashed), 768 (long dashed), and 1024 (solid) are all plotted for comparison of convergence. For $M < 0.25M_0$, the hydrogen fraction has increased by 59% and 23% for the polytropes merger and main sequence stars merger, respectively, relative to the initial hydrogen mass fraction (black).

it is difficult to extrapolate the convergence limit to the 1024 resolution simulation. It is apparent from Figure 4.9 that with increasing resolution, the hydrogen fraction in the inner few percent of mass decreases, but the hydrogen fraction in the outer half of the hydrogen core increases. However, considering the hydrogen core as a whole, the value of the average increase has converged, as can be seen in the solid line in Figure 4.8.

There must be enough hydrogen mixed into the hydrogen core, $M < 0.25M_*$, in order for the merger remnant to be long-lived enough to be seen as a blue straggler. Short-lived remnants do not remain in the blue straggler region of the colour-magnitude diagram (CMD) for long and therefore, they must have formed recently, no earlier than their lifetimes. Long-lived remnants are more likely to be seen, but they must be strongly mixed. Given the number of blue stragglers that have been found in all globular clusters that are studied in detail, the blue stragglers typically must have ages of at least a billion years. The lifetime of a $1.5M_\odot$ star is approximately 2 Gyr and in order for a $1.5M_\odot$ remnant to live a billion years, the hydrogen mass fraction in the hydrogen core must be approximately half of that found in a zero-age main sequence star. The merger remnant in the Eulerian simulations is found to have $\sim 34\%$ of the hydrogen found in the core of a zero-age main sequence star and therefore, it has an estimated lifetime on the order of 700 million years, lower than expected for a typical blue straggler. Note that the actual amount of hydrogen needed is difficult to estimate since the burning rate depends nonlinearly on the mass. This calculation will be performed in future work.

The difference in mixing seen in the polytropes merger and main sequence stars merger appears to be due to the fact that the main sequence stars are more centrally condensed and their cores are retained during the merger. Given the differences in results, it is important to use realistic stellar models rather than approximations like polytropes when simulating stellar collisions. Sills & Lombardi (1997) have also made the same argument.

4.4 Summary

A three-dimensional Eulerian TVD hydro code has been used to simulate the off-axis collisions of equal mass $M = 0.8M_\odot$ main sequence stars, modelled using $n = 3$ polytropes and with realistic calculations from YREC. The off-axis collision produces a single rotating oblate merger remnant that shows no indication of having a disc.

Substantial mixing is found for the polytropes merger. For the inner quarter of the mass $M < 0.25M_*$, 44% is mixed in from outside and the hydrogen mass fraction has increased by 59%. Strong mixing is seen throughout the remnant and on average, it is estimated to be $\sim 60\%$ fully mixed.

The remnant from the main sequence stars merger shows weaker mixing in the inner quarter of mass $M < 0.25M_*$, where only 22% is mixed in from outside and the hydrogen mass fraction has

increased by 23%. The rest of the star is in general much more strongly mixed and on average, it is estimated to be $\sim 50\%$ fully mixed.

Chapter 5

A Moving Frame Algorithm for High Mach Number Hydro

5.1 Introduction

In this chapter, I present a new Eulerian algorithm designed for capturing high Mach number hydrodynamics in cosmological and astrophysical simulations. It is known that Eulerian hydrodynamic codes that strictly conserve total energy suffer from the high Mach number problem. At high Mach numbers, extracting the thermal energy component from the much larger total energy is highly inaccurate. In the ZEUS code, Stone & Norman (1992) avoid the total energy equation and only solve the internal energy equation. This approach improves the accuracy of calculating the thermal energy in general, but it does not conserve the total energy. Other authors have proposed to solve an additional equation to complement the total energy conservation equation and near supersonic regions, the thermal energy is calculated using the new variable rather than the total energy. Ryu et al. (1993) solve an entropy conservation equation that is valid outside of shocks and they adopt this variable near supersonic regions and away from shocks. The TVD-entropy code was developed to correctly follow the adiabatic evolution and prevent spurious heating of cold, supersonic gas. Bryan et al. (1995) solve the internal energy equation and switch to this formulation only at high Mach numbers. We refer to their code as the PPM-Bryan code to distinguish it from the general PPM algorithm (Colella & Woodward, 1984).

All methods mentioned have been demonstrated to work when subjected to standard hydro tests, but these tests are static ones in which the medium is initially at rest. The static tests are idealized in that they involve strong shocks where the postshock fluid is subsonic. But in practice, regions experiencing shocks are often dynamical and can be embedded in fast-moving flows. For cosmological applications such as the study of the Lyman alpha forest, we often need to resolve modest shocks with pairwise velocities on the order of ~ 10 km/s in a large-scale flow moving $\sim 10 - 100$ times faster. In this case, both the preshock and postshock fluid are supersonic and

simulating the thermodynamic evolution is very difficult. The TVD-entropy code cannot capture the modest shocks in supersonic flows. The entropy equation is invalid since entropy is not conserved in shocks while the total energy formulation will encounter the high Mach number problem. The ZEUS and PPM-Bryan codes can in general capture such shocks, but the accuracy is limited by the fact that the thermal energy formulation does not conserve the total energy and this violates the shock jump conditions. A total energy conservation equation must be solved to achieve the correct shock jump conditions, but it is then difficult to calculate the thermal energy if the bulk kinetic component dominates the total energy.

This chapter describes the development and testing of the moving frame algorithm. I present a new formalism for Eulerian hydrodynamics and describe the implementation of the moving frame algorithm. This hybrid approach combines the high-resolution shock capturing ability of the Eulerian scheme with a low-diffusion Lagrangian advection scheme that does not suffer from the high Mach number problem. The algorithm is subjected to hydrodynamic tests like the one-dimensional Sod shock tube test and the three-dimensional Sedov-Taylor blast-waves. In addition, I have also modified these standard, static hydro tests to be dynamic in order to study the ability of the algorithm to capture shocks in fast-moving flows. The moving frame algorithm is designed for cosmological simulations and it is used to construct a cosmological code in Chapter 6.

5.2 Moving Frame Hydrodynamics

In the Eulerian approach, the fluid variables are defined with respect to the static grid frame of the simulation box. The high Mach number problem arises when the bulk flow components in the velocity field and total energy outweigh the corresponding local components. In the Lagrangian approach, the fluid equations are solved in the frame of the fluid and local variables can be naturally computed. In this new approach, a hybrid scheme is used. The Eulerian conservation equations are solved in a local frame moving with the flow and local quantities can be directly tracked. The moving frame approach will explicitly decompose the equations into a smooth flow, whose solution can be obtained by any simple finite-difference scheme, and a nonlinear component at low Mach numbers that can be solved by a traditional flux-conservative TVD solver. The TVD flux-conservative schemes are critical to capture shocks in the presence of discontinuities, where derivatives are ill-defined. This approach has the advantage of being able to effectively capture shocks like in standard grid-based schemes, but does not suffer from the high Mach number problem.

The Euler equations are the starting basis in this new formalism. For self-gravitating fluids,

the continuity equation, momentum equation, and energy equation are given as

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j}(\rho v_j) = 0, \tag{5.1}$$

$$\frac{\partial (\rho v_i)}{\partial t} + \frac{\partial}{\partial x_j}(\rho v_i v_j + P\delta_{ij}) = -\rho \frac{\partial \Phi}{\partial x_i}, \tag{5.2}$$

$$\frac{\partial e}{\partial t} + \frac{\partial}{\partial x_j}[(e + P)v_j] = -\rho v_i \frac{\partial \Phi}{\partial x_i}. \tag{5.3}$$

where the physical state of the fluid is specified by its density ρ, velocity field $\boldsymbol{v}$, and total energy density

$$e = \frac{1}{2}\rho v^2 + \epsilon. \tag{5.4}$$

In standard practice, the thermal energy density ϵ is evaluated by subtracting the kinetic energy density from the total energy density and this is the root of the high Mach number problem. For an ideal gas, the pressure P is related to the thermal energy density by the equation of state

$$P = (\gamma - 1)\epsilon, \tag{5.5}$$

where γ is the ratio of specific heats. Poisson's equation

$$\nabla^2 \Phi = 4\pi G\rho, \tag{5.6}$$

relates the gravitational potential Φ to the density field.

One can choose a frame that is locally moving with the fluid using a velocity decomposition technique. The velocity field of the fluid,

$$\boldsymbol{v} = \boldsymbol{\Delta v} + \tilde{\boldsymbol{v}}, \tag{5.7}$$

is decomposed into a local term $\boldsymbol{\Delta v}(\boldsymbol{x})$ and a smoothed background term $\tilde{\boldsymbol{v}}(\boldsymbol{x})$. The local velocity is a peculiar velocity with respect to the Eulerian grid and the smoothed background velocity is associated with the velocity of the grid cells. This construction results in a moving background reference frame that can approximate the bulk flow. The total energy density,

$$e = \Delta e + \tilde{e}, \tag{5.8}$$

can also be separated into a local term $\Delta e(\boldsymbol{x})$ and a grid term $\tilde{e}(\boldsymbol{x})$. The local total energy density,

$$\Delta e = \frac{1}{2}\rho \Delta v^2 + \epsilon, \tag{5.9}$$

is defined as the sum of the local kinetic energy density and the thermal energy density, analogous to equation (5.4). The grid energy density,

$$\tilde{e} = \frac{1}{2}\rho \tilde{v}^2 + \rho \boldsymbol{\Delta v} \cdot \tilde{\boldsymbol{v}}, \tag{5.10}$$

is uniquely defined in terms of the other fluid variables and is not a free variable. This decomposition allows the individual tracking of local and bulk terms. The thermal energy can be determined more accurately since there are no bulk flow components in equation (5.9).

In the standard Euler equations, the free fluid variables are the density ρ, velocity $\boldsymbol{v}$, and total energy density e. With the introduction of equations (5.7) to (5.10), the state of the fluid is now described by the density ρ, local velocity $\boldsymbol{\Delta v}$, grid velocity $\tilde{\boldsymbol{v}}$, and local total energy density Δe. The addition of the grid velocity requires an additional equation,

$$\frac{\partial \tilde{v}_i}{\partial t} + \Delta v_j \frac{\partial \tilde{v}_i}{\partial x_j} + \tilde{v}_j \frac{\partial \tilde{v}_i}{\partial x_j} = -\frac{\partial \Phi}{\partial x_i}, \tag{5.11}$$

to govern its time evolution. The grid velocity equation can be constructed arbitrarily and I choose a form that describes a simple physical picture and which simplifies the Euler equations. In equation (5.11), the grid velocity is advected at the global velocity. The advection equation is similar to Burger's equation, which describes the self-transport of a velocity field. In this frame decomposition, the Euler equations are then given by

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j}(\rho \Delta v_j) + \frac{\partial}{\partial x_j}(\rho \tilde{v}_j) = 0, \tag{5.12}$$

$$\frac{\partial (\rho \Delta v_i)}{\partial t} + \frac{\partial}{\partial x_j}(\rho \Delta v_i \Delta v_j + P\delta_{ij}) + \frac{\partial}{\partial x_j}(\rho \Delta v_i \tilde{v}_j) = 0, \tag{5.13}$$

$$\frac{\partial \Delta e}{\partial t} + \frac{\partial}{\partial x_j}[(\Delta e + P)\Delta v_j] + \frac{\partial}{\partial x_j}(\Delta e \tilde{v}_j) = -P\frac{\partial \tilde{v}_j}{\partial x_j}, \tag{5.14}$$

and completed by equation (5.11). The velocity decomposition is reflected in the structure of the expanded Euler equations (Eq. [5.11-5.14]). The first gradient term on the left hand side of each of the equations describes the flux of the free variables in the local frame where the fluid moves at the local velocity $\boldsymbol{\Delta v}$. These terms will be referred to as the Euler terms. The second gradient term describes the advection of the same quantities as the underlying background frame moves at the grid velocity $\tilde{\boldsymbol{v}}$. These terms will be referred to as the advection terms. The expanded Euler equations reduce to the standard ones when there is no background velocity field. In the case where the background velocity is a constant, the expanded Euler equations are just the standard ones with a simple Galilean transformation. For a nonuniform background velocity field, the energy equation (Eq. [5.14]) has an additional source term because the compression of the grid introduces a Coriolis term.

The gravitational acceleration of the fluid is exactly accounted for by the lone source term in the grid velocity equation. Gravity is a global process and it is more logical and advantageous to add the gravitationally induced changes to the grid rather than to the local quantities. An important point made by Ryu et al. (1993) is that since the gravitational energy can be comparable to the kinetic energy and much larger than the thermal energy, numerical errors in calculating the gravitational effects can significantly cause spurious heating of cold gas. Adding the gravitationally induced

changes to the grid can minimize the spurious heating problem. This is another example of how the velocity decomposition technique allows one to separate local and global components.

5.3 Moving Frame Algorithm

I now describe a hybrid algorithm for solving the expanded Euler equations. To simplify the notation, equations (5.12) to (5.14) are written in vector form:

$$\frac{\partial \boldsymbol{u}}{\partial t} + \frac{\partial \boldsymbol{F}_j^e}{\partial x_j} + \frac{\partial \boldsymbol{F}_j^a}{\partial x_j} = \boldsymbol{S}_j^g + \boldsymbol{S}_j^v, \tag{5.15}$$

where $\boldsymbol{u} = (\rho, \rho\boldsymbol{\Delta v}, \Delta e)$, $\boldsymbol{F}$ represents flux terms, and $\boldsymbol{S}$ represents source terms. The flux terms in the Euler and advection operations are represented by $\boldsymbol{F}^e$ and $\boldsymbol{F}^a$, respectively. Gravitational source terms are included in $\boldsymbol{S}^g$ while source terms arising from the velocity decomposition are found in $\boldsymbol{S}^v$. The expanded Euler equations are conveniently solved using an operator-splitting technique Strang (1968). Let the operator $\boldsymbol{L}$ represent the updating of $\boldsymbol{u}(\boldsymbol{x}, t)$ to $\boldsymbol{u}(\boldsymbol{x}, t + \Delta t)$ for a given operation. One first completes a forward sweep:

$$\boldsymbol{u}^{t+\Delta t} = \boldsymbol{L}^g \boldsymbol{L}^a \boldsymbol{L}^v \boldsymbol{L}^e \boldsymbol{u}^t, \tag{5.16}$$

and then performs a reverse sweep:

$$\boldsymbol{u}^{t+2\Delta t} = \boldsymbol{L}^e \boldsymbol{L}^v \boldsymbol{L}^a \boldsymbol{L}^g \boldsymbol{u}^{t+\Delta t}, \tag{5.17}$$

using the same time step Δt to obtain second-order accuracy. The operator-splitting technique is also useful for solving multi-dimensional problems. In three dimensions, the Euler equations can be dimensionally split into three separate one-dimensional equations that can be solved sequentially. The forward and reverse sweeps are

$$\boldsymbol{u}^{t+\Delta t} = \boldsymbol{L}^g (\boldsymbol{L}^a \boldsymbol{L}^v \boldsymbol{L}^e)_z (\boldsymbol{L}^a \boldsymbol{L}^v \boldsymbol{L}^e)_y (\boldsymbol{L}^a \boldsymbol{L}^v \boldsymbol{L}^e)_x \boldsymbol{u}^t, \tag{5.18}$$

$$\boldsymbol{u}^{t+2\Delta t} = (\boldsymbol{L}^e \boldsymbol{L}^v \boldsymbol{L}^a)_x (\boldsymbol{L}^e \boldsymbol{L}^v \boldsymbol{L}^a)_y (\boldsymbol{L}^e \boldsymbol{L}^v \boldsymbol{L}^a)_z \boldsymbol{L}^g \boldsymbol{u}^{t+\Delta t}, \tag{5.19}$$

and this splitting scheme is second-order accurate. The grid velocity $\tilde{\boldsymbol{v}}(\boldsymbol{x}, t)$ is updated similarly. Note that the gravity operator is not dimensionally split and that in the gravity step, the grid velocity is updated rather than the hydro vector $\boldsymbol{u}$.

5.3.1 The Euler operation

The Euler operation involves computing the flux of the free variables in the local frame where the fluid moves at the local velocity Δv. Here one solves the system of equations:

$$\frac{\partial \boldsymbol{u}}{\partial t} + \frac{\partial \boldsymbol{F}^e}{\partial x} = 0, \tag{5.20}$$

$$\frac{\partial \tilde{v}}{\partial t} + \Delta v \frac{\partial \tilde{v}}{\partial x} = 0. \tag{5.21}$$

Equation (5.20) is a vector conservation law that can be robustly solved using a second-order accurate TVD flux-assignment scheme with a second-order Runge-Kutta time integration technique (see Sec. 2.5.1). The general solution has the flux-conservative form

$$\boldsymbol{u}_n^{t+\Delta t} = \boldsymbol{u}_n^t - \left(\frac{\boldsymbol{F}_{n+1/2}^{t+\Delta t/2} - \boldsymbol{F}_{n-1/2}^{t+\Delta t/2}}{\Delta x} \right) \Delta t, \tag{5.22}$$

where the half-step fluxes $\boldsymbol{F}^{t+\Delta t/2}$ at cell boundaries are determined using the TVD flux assignment scheme. The procedure for assigning fluxes at cell boundaries based on cell-centered fluxes is difficult for the Euler operation. Due to the inclusion of gas pressure in the fluxes, the direction of flow depends on both the local velocity Δv and the local sound speed c_s and is not straightforward to determine. I have implemented the relaxing TVD scheme (Jin & Xin, 1985) with a van Leer flux limiter (Van Leer, 1974) to provide high-resolution capturing of shocks. The relaxing scheme offers an accurate and robust procedure for decomposing the flow into left and right moving waves, thereby making flux assignment straightforward. The relaxing TVD scheme has been successfully implemented for simulating astrophysical fluids by Pen (1998a) and Trac & Pen (2003).

Equation (5.21) is coupled to equation (5.20) via the local velocity Δv. The advection of the grid velocity can be solved using a second-order Runge-Kutta scheme:

$$\tilde{v}_n^{t+\Delta t} = \tilde{v}_n^t - \Delta v_n^{t+\Delta t/2} \left(\frac{\tilde{v}_{n+1}^{t+\Delta t/2} - \tilde{v}_{n-1}^{t+\Delta t/2}}{2\Delta x} \right) \Delta t, \tag{5.23}$$

where the half-step local velocity $\Delta v^{t+\Delta t/2}$ comes from the half-step value in the relaxing TVD solution. By construction, the grid velocity is smooth and a simple finite difference can be used to approximate the gradient. When fully expanded, equation (5.23) is similar to the second-order Lax-Wendroff equation (Lax & Wendroff, 1960). The Runge-Kutta scheme is chosen rather than the Lax-Wendroff scheme because the former provides an effective way to simultaneously solve equations (5.20) and (5.21).

In the energy equation (Eq. [5.14]), the source term arising from the velocity decomposition is associated with the Euler operation and not the advection operation. Therefore, the solution for the source term is calculated using a second-order Runge-Kutta scheme where the half-step pressure $P^{t+\Delta t/2}$ and half-step grid velocity $\tilde{v}^{t+\Delta t/2}$ comes from the half-step values in the solutions for equations (5.20) and (5.21). This additional Coriolis flux is from the compression of the grid. Since the grid velocity is smooth, the contributions from the source term will be small.

The total variation diminishing condition Harten (1983) is a nonlinear stability condition. The relaxing scheme for the Euler operation is TVD and stable provided that Courant or CFL number satisfies

$$\lambda \equiv \frac{\max(c_n)\Delta t}{\Delta x} < 1, \tag{5.24}$$

where $c \equiv |\Delta v| + c_s$ is called the freezing speed and c_s is the sound speed. The freezing speed is constructed to be greater than the largest eigenvalue of the flux Jacobian $\partial \boldsymbol{F}^e / \partial \boldsymbol{u}$. Depending on the application, the CFL number is normally chosen in the range $0.5 < \lambda < 0.9$.

5.3.2 The advection operation

The advection operation involves transporting the free variables as the underlying background frame moves at the grid velocity $\tilde{v}$. The advection system of equations:

$$\frac{\partial \boldsymbol{u}}{\partial t} + \frac{\partial \boldsymbol{F}^a}{\partial x} = 0, \tag{5.25}$$

$$\frac{\partial \tilde{v}}{\partial t} + \tilde{v}\frac{\partial \tilde{v}}{\partial x} = 0, \tag{5.26}$$

can be solved using the same method implemented for the Euler operation, but advection is more accurately solved using a Lagrangian approach to minimize numerical diffusion. Note that equation (5.26) has the form of Burger's equation, which describes the self-transport of a velocity field. In the Lagrangian step, the grid cells or zones are advected along lines of constant mass, where zone centres satisfy the equation

$$x_n^{t+\Delta t} = x_n^t + \tilde{v}_n^t \Delta t. \tag{5.27}$$

By construction, the starting distribution of Eulerian grid cells is regular, with equal volume cells. In general, the resulting distribution of Lagrangian zones will be irregular, with unequal volume zones. The Lagrangian zones are then remapped back onto the Eulerian grid for the next set of hydrodynamic operations.

Note that for a specific advection velocity in equation (5.27), zone centres can be transported to locations that coincide with grid cell centres, making the remap process straightforward. This observation is the central point of how the remap technique works. One can choose the advection velocity for the Lagrangian step to be

$$b_n^t \equiv \tilde{v}_n^t - a_n^t = \frac{\text{Round}(\tilde{v}_n^t \Delta t)}{\Delta t}, \tag{5.28}$$

where the residual term

$$a_n^t = \tilde{v}_n^t - \frac{\text{Round}(\tilde{v}_n^t \Delta t)}{\Delta t} \tag{5.29}$$

is chosen such that the distance a zone is transported is an integer multiple of the regular grid cell separation $\Delta x = 1$. Multiple zones can be mapped to the same grid cell and as a result, some grid cells may be left empty. To avoid having 2 zones be mapped to 2 cells with a gap inbetween, map the 2 zones to all 3 cells near divergent regions. Instead of advecting the left zone $n = l$ to the center of the left cell, increase the advection velocity b_l^t by $\Delta b_l = +1/3$ so that 2/3 of the zone falls

in the left cell and 1/3 in the center cell. Similarly, increase the advection velocity of the right zone $n = r$ by $\Delta b_r = -1/3$ so that 2/3 of the zone falls in the right cell and 1/3 in the center cell. The residual velocity a is re-adjusted by noting that $\Delta a = -\Delta b$.

First, the advection system of equations (eq. [5.25-5.26]) is solved, but for an advection velocity equal to a rather than $\tilde{v}$. This is done on the Eulerian grid using the Eulerian approach. Second, the Lagrangian step is carried out by transporting the quantities to their new location using the advection velocity b. The first step of the remap process can be solved using the Eulerian scheme devised for the Euler operation. However, one modification is made to speed up the algorithm. Since equation (5.25) is purely an advection equation, one can explicitly check the advection velocity to determine the direction of flow and assign fluxes accordingly. The relaxing scheme is replaced with a standard upwind TVD scheme that is computationally less expensive. Note that advection at the residual velocity a automatically satisfies the CFL stability condition because for any given time step, the residual velocity is chosen such that the fluid is moved no more than half a grid cell. In the second step, each Lagrangian zone is mapped to the corresponding Eulerian grid cell. It is necessary to impose the restriction that near divergent regions, the mapping process cannot result in the circumstance where an empty cell is adjacent to another empty cell. This criterion is satisfied if

$$\Delta t < \frac{1}{\max(|b_{n+1} - b_n|)}, \tag{5.30}$$

Since the grid velocity $\tilde{v}$ is smooth and its gradient small, this restriction rarely enforces the limiting time step. In almost all practical cases, the time step in the Euler operation will satisfy equation (5.30).

The two-step remap process is equivalent to advection at the velocity $\tilde{v}$, but this particular scheme offers some advantages over the alternative of solving the advection system using the Eulerian approach described for the Euler operation. Since the grid velocity is constructed to approximate the bulk flow, it can be relatively large compared to the local velocity Δv or the local sound speed c_s and would be a limiting factor in the length of the time step. However, we have seen that the two-step remap process effectively has no limiting time step when the grid velocity is smooth. Using longer time steps minimizes the numerical diffusion and shortens simulation run time.

5.3.3 Frame change

The procedure of constructing a new smoothed background velocity field will be referred to as a frame change. Again, the dimensional-splitting scheme effectively renders this process into a one-dimensional one. Since the flux sweep in the x direction only involves the gradients $\partial/\partial x$, the smoothing only has to be done in one direction at a time. I first illustrate the solution for a one-dimensional problem and then generalize it to higher dimensions.

At the beginning of both the forward and reverse sweeps, one constructs a smoothed background velocity field $\tilde{v}$ by smoothing the global velocity field v:

$$\tilde{v}_n = \frac{\sum_{m=1}^{N} v_m w_m s(x_n - x_m)}{\sum_{m=1}^{N} w_m s(x_n - x_m)}, \tag{5.31}$$

where N is the total number of grid cells, $w(x)$ is a weight function, and $s(x)$ is a smoothing kernel. The discrete convolution is computed using one dimensional fast Fourier transforms (FFTs). A Gaussian kernel with smoothing radius R,

$$s(x) = \frac{1}{\sqrt{2\pi}R} \exp\left(-\frac{x^2}{2R^2}\right), \tag{5.32}$$

will wash out the local peculiar velocity and leave the bulk velocity. A requirement of the algorithm is that the grid velocity be smooth and therefore, the smoothing radius must be set large enough to remove large gradients in the global velocity field. Typically, the largest gradients are found near shocks, which are resolved in approximately two grid cells, as demonstrated with the Sod shock tube test and the Sedov-Taylor blast-wave test in Chapter 2. In practice, the smoothing radius is chosen to be several times larger than the shock resolution. Based on hydrodynamic tests described later in this chapter and also in the next chapter, a value of $R = 8$ grid cells is found to be the best solution. Setting the smoothing radius too small explicitly affects the solutions to the grid velocity equation (Eq. [5.11]) and the local total energy equation (Eq. [5.14]), and this generally results in additional scatter and oscillations near shocks. On the contrary, choosing too large a value reduces the effectiveness of the algorithm for then, the magnitude of the grid velocity will decrease and the magnitude of the local velocity will be larger in general. Note that since the smoothing radius and shock broadening length are fixed in grid units, the effectiveness of the moving frame algorithm to reduce Mach numbers and resolve shocks increases with higher resolution.

The weight function can be set to unity in the simplest case, but it may be desirable to weight the contributions of the global velocity field not just on locality but also on the properties of the physical environment. In general, cold regions are more likely to suffer from the high Mach number problem simply due to the fact that the sound speed is small. Consider the case of a moving shock front. If we choose the reference frame to be that of the shock, then the cold gas in front of the shock will be moving supersonically in this frame and have high Mach numbers. However, if we choose the reference frame to be that of the cold ambient medium, then the ambient gas will be subsonic. While the postshock medium will be moving fast with respect to this reference frame, it will have low Mach numbers since the postshock gas is hot. Hence, a weight function

$$w_n = \frac{1}{\sqrt{\max(T_n, T_{min})}}, \tag{5.33}$$

where T is the temperature distribution of the gas is adopted. In practice, a minimum temperature T_{min} in equation (5.33) is imposed to prevent inaccurate weighting due to the possibility of unphysical temperatures or spurious oscillations in the temperature distribution.

During a frame change, mass is automatically conserved since the definition of mass is absolute. Momentum and energy conservation requires adjusting the local velocity Δv and local total energy density Δe. If the change in the grid velocity is $\Delta\tilde{v}_n$ for cell n, then the adjustments are

$$\Delta v_n^f = \Delta v_n^i - \Delta\tilde{v}_n, \tag{5.34}$$

$$\Delta e_n^f = \Delta e_n^i + \frac{1}{2}\rho_n(\Delta\tilde{v}_n)^2 - \rho\Delta v_n^i\Delta\tilde{v}_n. \tag{5.35}$$

where i and f label the initial and final states with respect to the frame change.

For multi-dimensional problems, all components of the global velocity $\boldsymbol{v}$ need to be decomposed, but the smoothing only has to be done in the direction of the flux sweep. For a three-dimensional problem, the splitting scheme given by equations (5.18) and (5.19) involves 6 flux sweeps, each of which is preceded by a frame change.

5.3.4 Gravity

Gravitationally induced accelerations are added to the grid velocity $\tilde{\boldsymbol{v}}$ rather than to the local quantities $\boldsymbol{u}$. The change in the global velocity and the global kinetic energy are entirely accounted for by updating the grid velocity solely. Poisson's equation (eq. [5.6]) relates the gravitational potential to the density field and the general solution can be written as

$$\Phi(\boldsymbol{x}) = \int \rho(\boldsymbol{x}')w(\boldsymbol{x}-\boldsymbol{x}')d^3x', \tag{5.36}$$

where the isotropic kernel is given by

$$w(r) = -\frac{G}{r}. \tag{5.37}$$

In the discrete case, the integral in equation (5.36) becomes a sum and Poisson's equation can be solved using FFTs to do the convolution. The force terms $f_i \equiv -\partial\Phi/\partial x_i$ on the right hand side of the Euler equations are calculated by finite differencing the potential. The real space kernel $w(r)$ is constructed on the grid with the zero point satisfying

$$f(1) = -\frac{w(2)-w(0)}{2\Delta x} = -G. \tag{5.38}$$

By choosing $w(0) = -2.5G$, the force at a separation of one grid cell can be made exact. Alternatively, the accelerations can be calculated directly using the convolution,

$$f_i(\boldsymbol{x}) = \int \rho(\boldsymbol{x}')w_i(\boldsymbol{x}-\boldsymbol{x}')d^3x', \tag{5.39}$$

where the anisotropic force kernels are given by

$$w_i(\boldsymbol{x}) = -G\frac{x_i}{r^3}. \tag{5.40}$$

The force method exactly reproduces the pair-wise inverse-square law on the grid, but comes at the cost of two extra FFTs. The potential method is computationally less expensive, but the finite differencing degrades the pair-wise force resolution by a few grid cells. While the pair-wise force is not exact at small separations, the net force on any given cell is in general still highly accurate. In practice, the relatively small accuracy trade-off of the potential method is preferred over the factor of 2 increase in computational work and time of the force method.

In the operator-splitting scheme, the forward and reverse gravity operations are consecutive and can be considered as one operation. For the gravity operation, we use the density field $\rho^{t+\Delta t}$ to determine the gravitational potential $\Phi^{t+\Delta t}$, which is then finite differenced to give the acceleration field $\boldsymbol{a}^{t+\Delta t}$. We impose the restriction that the time step must satisfy

$$\Delta t < \frac{1}{\sqrt{\max(a_n)}}. \tag{5.41}$$

In practice, the gravitational time step constraint rarely enforces the limiting time step. It tends to come into effect only when the fluid is initially at rest or moving relatively slowly so that the Euler time step is larger than the gravitational time step.

5.4 Hydrodynamics Tests

In this section, the one-dimensional and three-dimensional versions of the moving frame algorithm are applied to hydrodynamic tests like the Sod shock tube test and the Sedov-Taylor blast-wave test. The purpose is to demonstrate the strength of the moving frame algorithm to capture shocks in the presence of high Mach numbers. Two variants of each test are conducted. The standard version is a static one, where the initial ambient medium is at rest with respect to the simulation box. This version is often presented in the literature to showcase the accuracy of hydrodynamic schemes, but this idealized case minimizes the Mach number. In practice, regions experiencing shocks are often dynamical and not at rest with respect to the simulation box. In the modified version, the medium has an initial velocity boost and is rapidly moving through the periodic box. The boosted test mimics the situation where the shock is embedded in a bulk flow and both preshock and postshock Mach numbers are high.

5.4.1 One-dimensional Sod Shock Tube Test

The one-dimensional moving frame algorithm is applied to the shock tube test with initial conditions: $\rho_5 = 1$, $P_5 = 1$, $\rho_1 = 0.2$, $P_1 = 0.01$. The initial conditions are identical to Shapiro et al. (1996) and Pen (1998a) and that presented in Section 2.7.1. The pressure ratio $P_5/P_1 = 100$ is 10 times larger than that in Sod (1978), allowing a rigourous test of the algorithm's ability to capture strong shocks at high Mach numbers. In the static test, the fluid is initially at rest and $v_5 = v_1 = 0$. The simulation is run until the shock front has propagated a distance of $x_s = 50$ grid cells. In the

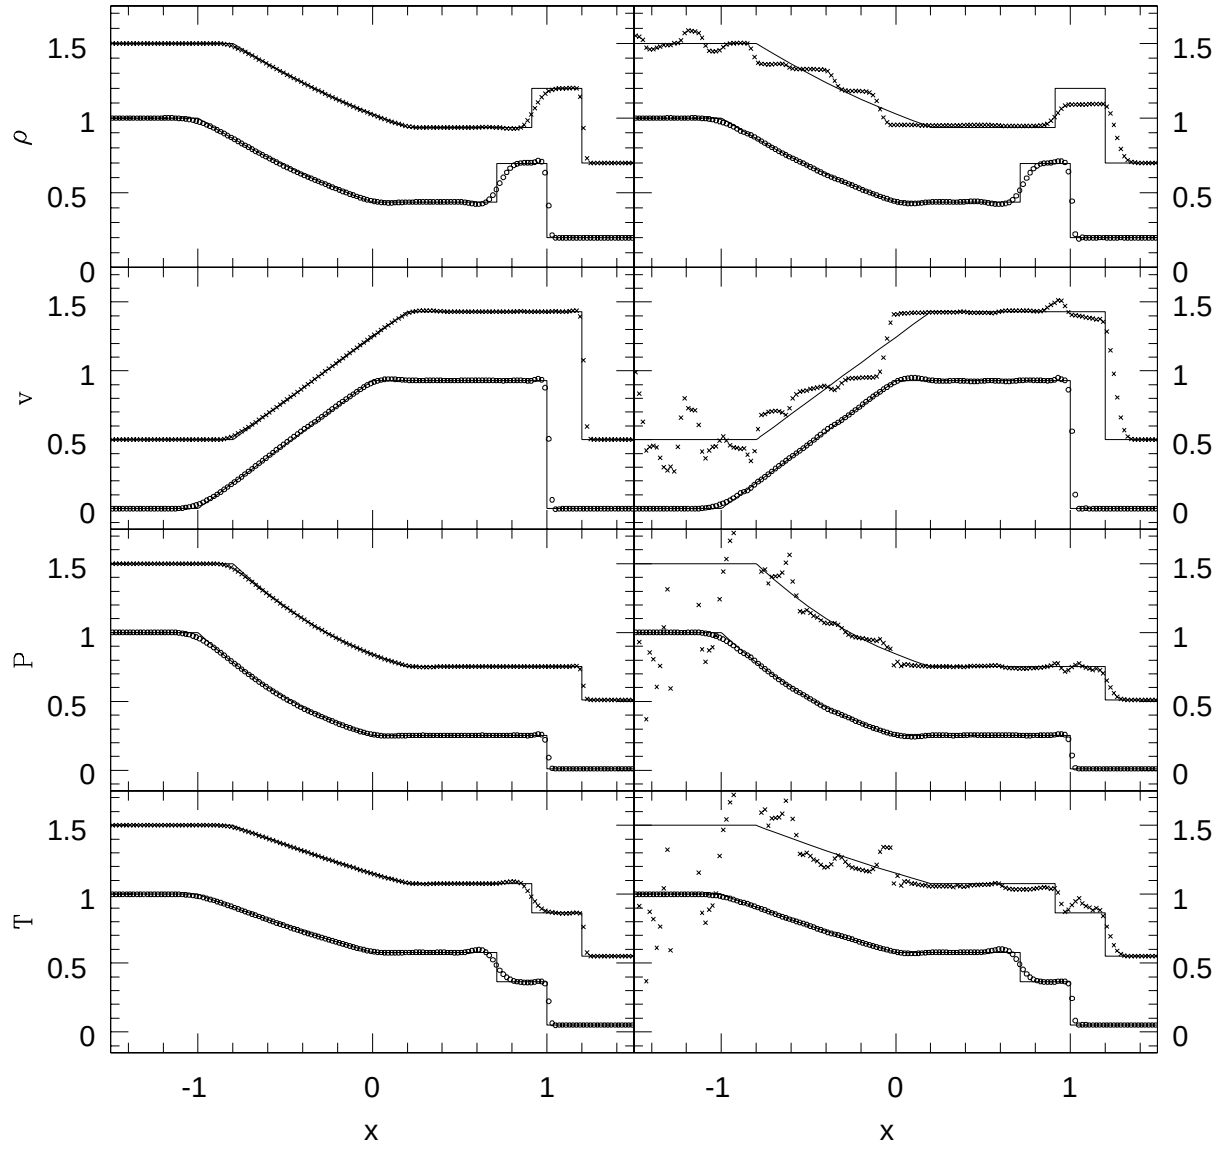

Figure 5.1 One-dimensional Sod Shock tube test results with the moving frame implementation turned on (open circles) and off (crosses). The standard tests are shown in (a) and the boosted test results are found in (b). In the boosted test, both pre-shock and post-shock Mach numbers are high.

boosted test, the fluid has uniform initial motion with $v_5 = v_1 = 100c_1$. In the elapsed time that it takes the shock to propagate a distance of $x_s = 50$ grid cells in its own frame, the entire region of interest will have moved through the box a total distance of more than 20 times x_s. The tests are conducted with the moving frame technique turned on and off. In the latter case, the algorithm reduces to the relaxing TVD algorithm.

In Figure 5.1(a), the static test results for moving frame algorithm (open circles) and relaxing TVD algorithm (crosses) are plotted against the exact solution (solid lines). The plots have been rescaled such that the initial discontinuity is placed at $x = 0$ and the shock front at $x = 1$. The grid spacing corresponds to $\Delta x = 0.02$. The relaxing TVD results have been shifted both horizontally and vertically for clarity. Both algorithms produce very similar results and are successful in capturing the shock, contact, and expansion of the shock tube.

The moving frame and relaxing TVD algorithms give very different results for the boosted test, as shown in Figure 5.1(b). The relaxing TVD algorithm suffers from the high Mach number problem in this test and its inability to accurately track the thermal energy is reflected in the relatively large jumps in the pressure and temperature curves. In addition, the large bulk velocity imposes a small time step, which in turn means that a larger number of steps is required for the simulation to evolve to the chosen final time. As a result, more diffusion is expected and both the shock and contact fronts are highly degraded. The moving frame algorithm does not suffer from the high Mach number problem here. Though the fluid is rapidly moving through the simulation box, the frame change in the moving frame algorithm is able to remove the bulk component, effectively reducing the problem so that the hydrodynamics can be solved like in the idealized static case.

5.4.2 Three-dimensional Sedov-Taylor Blast-wave Test

The three-dimensional moving frame algorithm is applied to two variants of the Sedov-Taylor blast-wave test. A simulation box with 128^3 cells is initially set up with constant density $\rho_1 = 1$. At time $t = 0$, a spike of thermal energy $E_0 = 10^5$ is injected into one cell at the centre of the box and let the simulation is run until final time $t_f = 35.59$ when the shock front has expanded to a radius of $r_s = 48$ cells. In the static test, the initial velocity of the medium is zero. In the boosted test, all three components of the initial velocity are chosen to be equal to 100 times the sound speed $c_{48} = (\gamma P_2/\rho_2)^{1/2}$ of the immediate postshock gas at $r_s = 48$ cells. In the elapsed time, the centre of the explosion will have moved through the periodic box a total distance of more than 20 times the radius of the shock.

In Figures 5.2(a) and 5.2(b), the results of the standard and boosted tests are compared. In the standard test, the centre of the explosion is at rest with respect to the simulation box and the grid velocity is approximately zero. In effect, the advection step contributes little to the fluid flow and the results are essentially the same compared to that produced by a standard relaxing TVD algorithm. The shock is propagated accurately and resolved in roughly two grid cells. There are no

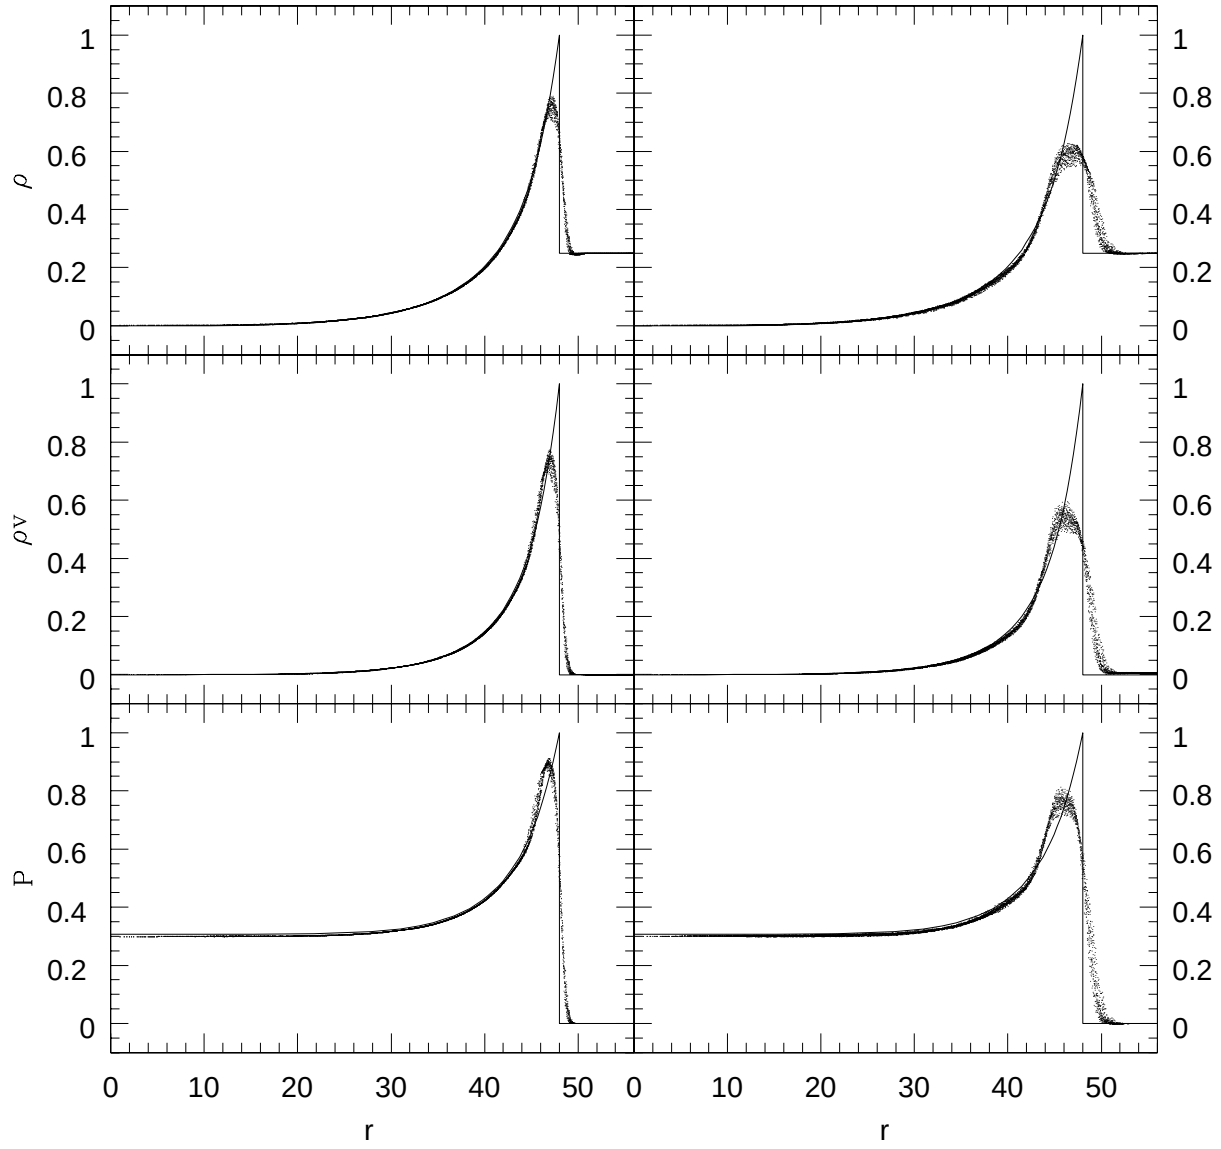

Figure 5.2 Three-dimensional moving frame algorithm applied to the standard (a) and boosted (b) Sedov-Taylor blast-wave test. In the boosted test, both pre-shock and post-shock Mach numbers are high.

spurious oscillations and little anisotropic scatter is seen. In the boosted test, the medium and the centre of the explosion are rapidly moving through the periodic box, but the frame change removes that bulk component, with the grid velocity being approximately equal to the initial velocity boost. As a result, in the Euler step, the hydrodynamics is solved in a frame at rest with respect to the centre of the explosion, like that in the standard test. The shock is correctly propagated despite the high Mach numbers. There is more diffusion and scatter in this case compared to the standard one because of the additional advection step. This is expected and inevitable when advecting over a large number of cells on an Eulerian grid. In fact, the diffusion is minimized by the hybrid Lagrangian-Eulerian advection scheme implemented in the moving frame algorithm. All previous Eulerian algorithms to date will fail this non-trivial boosted Sedov-Taylor test.

5.5 Summary

A new hybrid hydrodynamic algorithm has been designed to solve the Eulerian fluid conservation equations in an adaptive frame moving with the fluid where Mach numbers are minimized. Using a velocity decomposition technique, a local velocity and local total energy density can be defined and the bulk kinetic components can be stored in a smoothed background velocity field that is associated with the grid velocity. In addition, gravitationally induced kinetic changes are added to the grid rather than to the local quantities, thereby minimizing the spurious heating problem plaguing cold gas flows. Separately tracking local and bulk flow components allows thermodynamic variables to be accurately calculated in both subsonic and supersonic regions. A main feature of the algorithm is the ability to resolve shocks and prevent spurious heating where both the preshock and postshock Mach numbers are high.

The moving frame algorithm has been subjected to a one-dimensional Sod shock tube test and a three-dimensional Sedov-Taylor blast-wave test. These standard tests have also been modified to include an initial velocity boost. The modified tests mimic the situation where the shock is embedded in a supersonic flow where both preshock and postshock Mach numbers are high. The moving frame algorithm is able to accurately resolve the shocks without forming spurious oscillations in all tests conducted. Previous Eulerian implementations will fail the non-trivial boosted tests where the postshock Mach numbers remain high. The research presented in this chapter has been published in Trac & Pen (2004a).

Chapter 6

Cosmological Hydro&N-body Code

6.1 Introduction

In this chapter, the moving frame algorithm is adapted for cosmological applications in an expanding Friedman-Robertson-Walker (FRW) universe. In this new cosmological code, the hydrodynamic evolution of the baryons is captured using the moving frame algorithm while the gravitational evolution of the collisionless dark matter is tracked using a particle-mesh (PM) N-body algorithm. Cosmological initial conditions for the gas and dark matter can be generated from a Gaussian random field whose mass power spectrum is given by the initial matter power spectrum. In general, cosmological codes based on Eulerian hydro and PM N-body algorithms do not have high dynamic range in density as codes based on Lagrangian SPH and Tree N-body algorithms, but they do have higher dynamic range in mass and this inherent feature is often wanted for cosmological applications. This cosmological code is highly suited for simulating the evolution of the intergalactic medium (IGM) to study the Lyman alpha forest, the Sunyaev-Zeldovich effect, and the X-ray background. Hydrodynamic simulations run with the moving frame code have been used in Zhang, Pen, & Trac (2004a) and Zhang, Pen, & Trac (2004b).

I describe the cosmological code and initial conditions generator, along with the unit conversions between grid values and physical values. The PM algorithm is described and a pair-wise force test of the FFT-based Poisson solver is performed. Cosmological tests like the Zeldovich pancake test and the self-similar scaling test are conducted to demonstrate that the code accurately captures hydrodynamics, self-gravity, and expansion in cosmological fluids. In addition, a sample cosmological simulation is run and the numerical Mach number distribution is examined to study the effectiveness of the moving frame approach in solving the high Mach number problem.

6.2 Cosmological Hydrodynamics

In an expanding FRW background, the standard prescription is to convert proper coordinates x_p into comoving coordinates $x = x_p/a$, where the scale factor a is governed by the Friedman equation

$$\left(\frac{da}{dt}\right)^2 = a^2 H_0^2 (\Omega_m a^{-3} + \Omega_k a^{-2} + \Omega_\Lambda). \tag{6.1}$$

The cosmological constants H_0, Ω_m, Ω_k, and Ω_Λ are present epoch values and by choosing $a_0 = 1$, one obtains the relation

$$\Omega_m + \Omega_k + \Omega_\Lambda = 1. \tag{6.2}$$

In a cosmologically flat background, the curvature term Ω_k vanishes, leaving the definition $\Omega_{total} \equiv \Omega_m + \Omega_\Lambda = 1$. In comoving coordinates, the cosmological hydrodynamic equations are then written as

$$\frac{\partial \rho}{\partial t} + \frac{1}{a}\frac{\partial}{\partial x_j}(\rho v_j) = 0, \tag{6.3}$$

$$\frac{\partial (\rho v_i)}{\partial t} + \frac{1}{a}\frac{\partial}{\partial x_j}(\rho v_i v_j + P\delta_{ij}) = -\frac{\dot{a}}{a}\rho v_i - \frac{1}{a}\rho\frac{\partial \Phi}{\partial x_i}, \tag{6.4}$$

$$\frac{\partial e}{\partial t} + \frac{\partial}{\partial x_j}[(e+P)v_j] = -\frac{2\dot{a}}{a}e - \frac{1}{a}\rho v_i \frac{\partial \Phi}{\partial x_i}, \tag{6.5}$$

where x is the comoving coordinate, t is the proper cosmic time, ρ is the comoving density, v is the proper peculiar velocity, P is the comoving pressure, e is the comoving total peculiar energy density, and Φ is the proper peculiar gravitational potential.

The standard form of cosmological hydrodynamic equations do not preserve the time-invariant conservation form of the Euler fluid equations. However, with the appropriate time coordinate transformation, the conservation form can be maintained. Following Gnedin (1995), a new time variable is introduced,

$$d\tau = \frac{dt}{a^2}, \tag{6.6}$$

which preserves the conservation structure of the Euler equations. The cosmological moving frame hydrodynamic equations can then be written as

$$\frac{\partial \rho}{\partial \tau} + \frac{\partial}{\partial x_j}(\rho \Delta v_j) + \frac{\partial}{\partial x_j}(\rho \tilde{v}_j) = 0, \tag{6.7}$$

$$\frac{\partial (\rho \Delta v_i)}{\partial \tau} + \frac{\partial}{\partial x_j}(\rho \Delta v_i \Delta v_j + P\delta_{ij}) + \frac{\partial}{\partial x_j}(\rho \Delta v_i \tilde{v}_j) = 0, \tag{6.8}$$

$$\frac{\partial \Delta e}{\partial \tau} + \frac{\partial}{\partial x_j}[(\Delta e + P)\Delta v_j] + \frac{\partial}{\partial x_j}(\Delta e \tilde{v}_j) = -P\frac{\partial \tilde{v}_j}{\partial x_j}, \tag{6.9}$$

$$\frac{\partial \tilde{v}_i}{\partial \tau} + \Delta v_j \frac{\partial \tilde{v}_i}{\partial x_j} + \tilde{v}_j \frac{\partial \tilde{v}_i}{\partial x_j} = -\frac{\partial \Phi}{\partial x_i}, \tag{6.10}$$

where x and ρ remain the comoving coordinate and density, respectively, while the other hydro terms have new definitions. The time-dependence of the cosmological expansion is included in the gravitational source term, where the Newtonian gravitational constant G becomes aG. This change is reflected in Poisson's equation.

Under the spatial and time coordinate transformations, all density terms are for comoving volumes $d^3x = a^{-3}d^3x_p$. The mass density $\rho(\tau)$ is per comoving volume and related to the proper mass density by

$$\rho = a^3\rho_p. \tag{6.11}$$

The peculiar velocity $\boldsymbol{v}(\tau) \equiv d\boldsymbol{x}/d\tau$ used in the code is related to the proper peculiar velocity $\boldsymbol{v}_p$ by the expression,

$$\boldsymbol{v} = a\boldsymbol{v}_p \tag{6.12}$$

Correspondingly, the total energy density $e(\tau)$, pressure $P(\tau)$, and gravitational potential $\Phi(\tau)$ are

$$P = a^5P_p, \tag{6.13}$$

$$e = a^5e_p, \tag{6.14}$$

$$\Phi = a^5\Phi_p, \tag{6.15}$$

where P_p, e_p, and Φ_p are the proper pressure, proper peculiar total energy density, and proper peculiar potential. One factor of a^2 comes from the velocity transformation while the other factor a^3 comes from converting between comoving and proper volumes.

The relationship between the scalefactor a, proper time t, and new time τ can be determined by considering the alternative Friedman equation

$$\left(\frac{da}{d\tau}\right)^2 = a^3(\Omega_m H_0^2)\left[1 + a\frac{1-\Omega_m-\Omega_\Lambda}{\Omega_m} + a^3\frac{\Omega_\Lambda}{\Omega_m}\right]. \tag{6.16}$$

In an Einstein-de Sitter universe where $\Omega_m = 1$ and $\Omega_\Lambda = 0$, we have the analytical solution

$$a(t) = \left(\frac{3H_0}{2}\right)^{2/3} t^{2/3}, \tag{6.17}$$

$$a(\tau) = \left(\frac{4}{H_0^2}\right)(-\tau)^{-2}, \tag{6.18}$$

$$t(\tau) = \left(\frac{16}{3H_0^4}\right)\tau^{-3} \tag{6.19}$$

where $-\infty < \tau < 0$. For a general cosmology with non-zero cosmological constant Ω_Λ, there is no analytical solution. In the MACH code, equation (6.16) is integrated using a third-order Taylor expansion to obtain an accurate solution for $a(\tau)$. The cosmological expansion is restricted to

$$\frac{\Delta a}{a} < 0.02, \tag{6.20}$$

for every time step. At high redshifts where the gas velocity and acceleration are relatively small, the expansion time step is generally smaller than both the Euler and gravitational time steps.

Listed here are the unit conversions between grid values and physical values. The unit conversions are chosen so that the grid quantities are close to unity. The grid spacing is chosen to be $\Delta x = 1$ and this sets the length unit,

$$\mathcal{L} = a\frac{L}{N}, \tag{6.21}$$

where L is the physical length of the simulation box and N is the number of grid cells per side length. The mean comoving grid matter density is normalized to unity and this fixes the density unit,

$$\mathcal{D} = a^{-3}\frac{3\Omega_m H_0^2}{8\pi G}. \tag{6.22}$$

The length and density unit together restrict the mass unit to be $\mathcal{M} = \mathcal{D}\mathcal{L}^3$. The time unit is uniquely determined once the grid value of the gravitational constant G is set. We choose $G = 1/(6\pi)$ and obtain the time unit

$$\mathcal{T} = \frac{2a^2}{3}\frac{1}{\sqrt{\Omega_m H_0^2}}. \tag{6.23}$$

Note that this converts the grid time τ to physical time t. The three fundamental length, mass, and time units can be used to determine all other unit conversions. The velocity unit is given by $\mathcal{V} = \mathcal{L}/\mathcal{T}$ and the energy unit is $\mathcal{E} = \mathcal{M}(\mathcal{L}/\mathcal{T})^2$.

Pen (1998a) calls the new time variable τ a 'Newtonian' time because with this substitution, Newton's laws apply directly. Objects move in straight trajectories unless acted upon by an external force. The scaling relations for adiabatic expansion are automatically preserved. Consider a completely homogeneous universe with some initial density, velocity, and temperature. The velocity $\boldsymbol{v}(\tau)$ will remain constant and the proper peculiar velocity will automatically scale as $\boldsymbol{v}_p \propto a^{-1}$. Similarly, for constant $\rho(\tau)$ and $P(\tau)$, the temperature will scale as $T \propto a^{-3(\gamma-1)}$.

6.3 Particle-Mesh N-body Algorithm

The evolution of the collisionless dark matter is tracked using the standard approach of the particle-mesh (PM) N-body scheme (Hockney & Eastwood, 1988; Efstathiou et al., 1985). The collisionless particles are advected by solving Newton's equations of motion in an FRW universe. For each particle, its comoving position $\boldsymbol{x}(\tau)$ and velocity $\boldsymbol{v}(\tau)$ are stored and these adherent quantities can be updated if the acceleration $\boldsymbol{a}(\tau)$ is known. The gravitational force field is computed using the total matter density field, with contributions from both the baryonic gas and dark matter particles. Mass assignment of particles onto the grid is accomplished using the 'cloud-in-cell' (CIC) interpolation scheme (Hockney & Eastwood, 1988). In PM codes, the standard Poisson solver uses

FFTs to forward Fourier transform the density field and to inverse transform the potential field. The FFT-based gravity solver has been previously discussed in Chapter 3.

An explicit, second-order time integration technique based on the Runge-Kutta scheme is used. The gas and dark matter must be synchronized when the total density field is constructed. Recall that in the operator-splitting scheme, the gravity operation is carried out inbetween the forward and reverse sections of the hydro sweep and the fluid has been been advanced from u^{τ} to $u^{\tau+\Delta\tau}$. First, the gas and dark matter are synchronized by advancing the particles,

$$\boldsymbol{x}^{\tau+\Delta\tau} = \boldsymbol{x}^{t} + \boldsymbol{v}^{t}\Delta\tau, \tag{6.24}$$

where this half-step is similar to the half-step in a standard Runge-Kutta scheme. The particles are mapped to the grid using the CIC mass assignment scheme and the acceleration field is computed following the procedure outlined earlier in the section on gravitation (§). The acceleration $\boldsymbol{a}^{\tau+\Delta\tau}$ on each particle is interpolated from the grid using the CIC scheme. Using the same CIC scheme to do mass assignment and force interpolation guarantees that no fictitious self-force is imposed on the particles. The full-step in the Runge-Kutta scheme updates the position and velocity for each particle as follows:

$$\boldsymbol{x}^{\tau+2\Delta\tau} = \boldsymbol{x}^{\tau} + \boldsymbol{v}^{\tau}(2\Delta\tau) + \boldsymbol{a}^{\tau+\Delta\tau}\frac{(2\Delta\tau)^2}{2}, \tag{6.25}$$

$$\boldsymbol{v}^{\tau+2\Delta\tau} = \boldsymbol{v}^{\tau} + \boldsymbol{a}^{\tau+\Delta\tau}(2\Delta\tau). \tag{6.26}$$

Note that the acceleration $\boldsymbol{a}^{\tau+\Delta\tau}$ can be expanded as $\boldsymbol{a}^{\tau} + (d\boldsymbol{a}/d\tau)\Delta\tau$ and with this substitution, equations (6.25) and (6.26) will be identical to a Taylor expansion up to second-order.

This Runge-Kutta integration scheme for advancing the particles has several advantages. It provides a way to couple and synchronize the gas and dark matter during the gravity step. The time step $\Delta\tau$ can be adjusted every double sweep, just like in the operator-splitting scheme for hydrodynamics. In addition, the position and velocity are synchronized at the start and end of every double time step and not staggered temporally like in the leap-frog scheme, for instance.

6.3.1 Pair-wise Force Test

As discussed in Chapter 3, the FFT-based Poisson solver can be constructed using two different methods. In one method, the gravitational forces can be solved for directly and in the other, the potential is first computed and then finite-differenced to obtain the forces. The force method exactly reproduces the pair-wise inverse-square law on the grid, but not for particles that are mapped to the grid since the CIC mass assignment scheme smoothes the mass and pair-wise force near the grid scale. A pair-wise particle force test can be used to show that the force and potential methods give similarly accurate results. A particles is placed randomly in a box and the pair-wise forces for

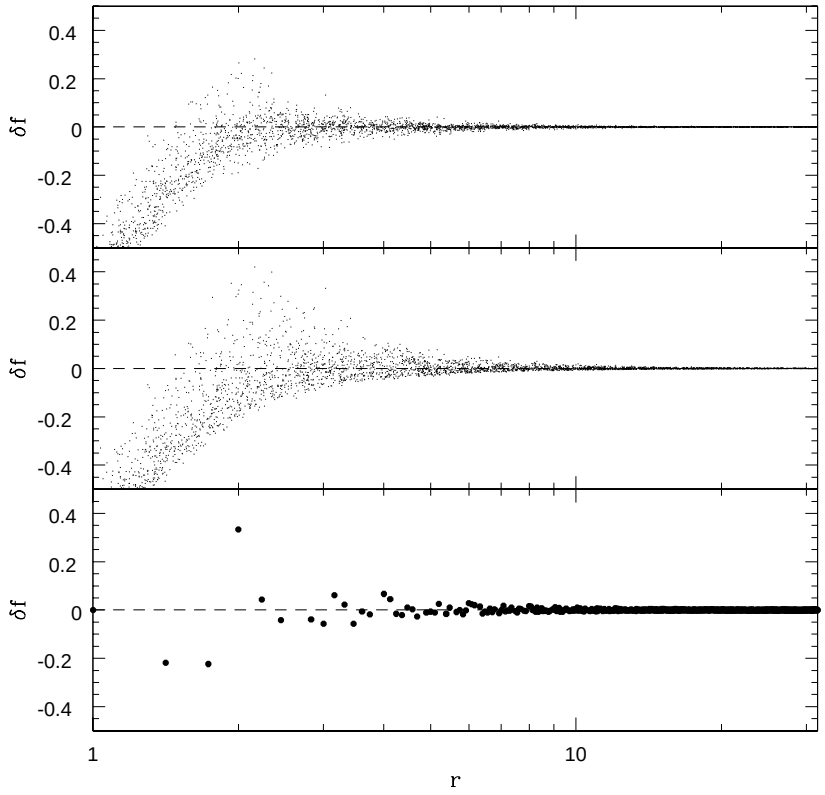

Figure 6.1 Fractional errors for the pair-wise force test of the PM gravity solver. The top and middle panels show the force errors for the particles from the force and potential methods, respectively, while the bottom panel shows the errors on the grid from the potential method.

a number of randomly placed test particles are measured. Plotted in Figure 6.1 is the fractional error

$$\delta f \equiv \frac{|\boldsymbol{f}| - f(r)}{f(r)}, \tag{6.27}$$

where $f(r) = G/r^2$ is the magnitude of the exact force at separation r. The top and middle panels show the fractional errors for the force and potential methods, respectively. The same random set of pairs, with equal numbers per logarithmic bin, are plotted for both methods for proper comparison. For comparison, the bottom panel shows the errors on the grid using the potential method, as taken from Figure 3.1.

For both methods, the fractional errors can approach unity near the grid scale because the cubical cloud shape of the CIC mass assignment scheme softens the mass resolution and introduces anisotropic scatter. Some additional softening and scatter is found for the potential method as a result of the finite differencing, but this difference between the two methods is relatively small. Hence, the computationally less expensive potential method is the preferred choice for the FFT-

based Poisson solver in PM codes.

6.3.2 Optimizations

The cosmological MACH code is computationally fast, memory friendly, and efficiently parallelized. The dimensional splitting technique for the hydrodynamics effectively decomposes the problem into a one-dimensional one and this has several advantages. First, this allows one to write highly optimized, one-dimensional algorithms that are cache efficient. Second, parallelization using OpenMP directives for shared-memory machines becomes straightforward. During a hydro sweep, the grid can be divided into one-dimensional columns and these independent columns can be efficiently distributed amongst the multiple processors. Third, the code is memory friendly since additional variables for the TVD and Runge-Kutta schemes only need to be stored temporarily in small arrays.

The PM N-body algorithm is also fully parallelized. For the first-half step in the Runge-Kutta scheme, the half-positions can be calculated in parallel with each processor responsible for a fraction of particles. No extra memory overhead is needed for storing the intermediate positions. The half-positions $\boldsymbol{x}^{\tau+\Delta\tau}$ can be stored in the same array as the initial positions $\boldsymbol{x}^{\tau}$. The fully updated positions can be calculated as

$$\boldsymbol{x}^{\tau+2\Delta\tau} = \boldsymbol{x}^{\tau+\Delta\tau} + \boldsymbol{v}^{\tau+2\Delta\tau}\Delta\tau, \tag{6.28}$$

where this equivalent alternative to equation (6.25) does not require the initial positions and velocities. The mass assignment and force interpolation steps can be done in parallel by using the standard procedure of constructing linked lists (Hockney & Eastwood, 1988) for the particles.

For a 1024^3 grid, the MACH code requires 20 GB to hold the hydro array $\boldsymbol{u}$. Since the grid velocity is smooth by construction, it only needs to be stored at half the resolution and 1.5 GB is needed for the array $\tilde{\boldsymbol{v}}$. It is expanded to full resolution for computation, but when working on one-dimensional columns of data, this requires relatively little memory overhead. Normally, simulations are run where the ratio of particles to grid cells is 1:8. For 512^3 particles, 3 GB are required to store both the positions and velocities and another 0.5 GB for the linked list. For the gravity step, the total mass density field and potential can be stored in the same array and this requires another 4 GB. In total for a 1024^3 grid and 512^3 particles, the MACH code uses 29 GB of memory to hold the large arrays and a few GB to hold smaller temporary ones.

6.4 Cosmological Initial Conditions

Cosmological initial conditions for Hydro&N-body simulations are generated from a Gaussian random field whose mass power spectrum is given by the initial matter power spectrum $P_i(k)$. The generic method involves a random sampling of the initial power spectrum to construct a Gaussian

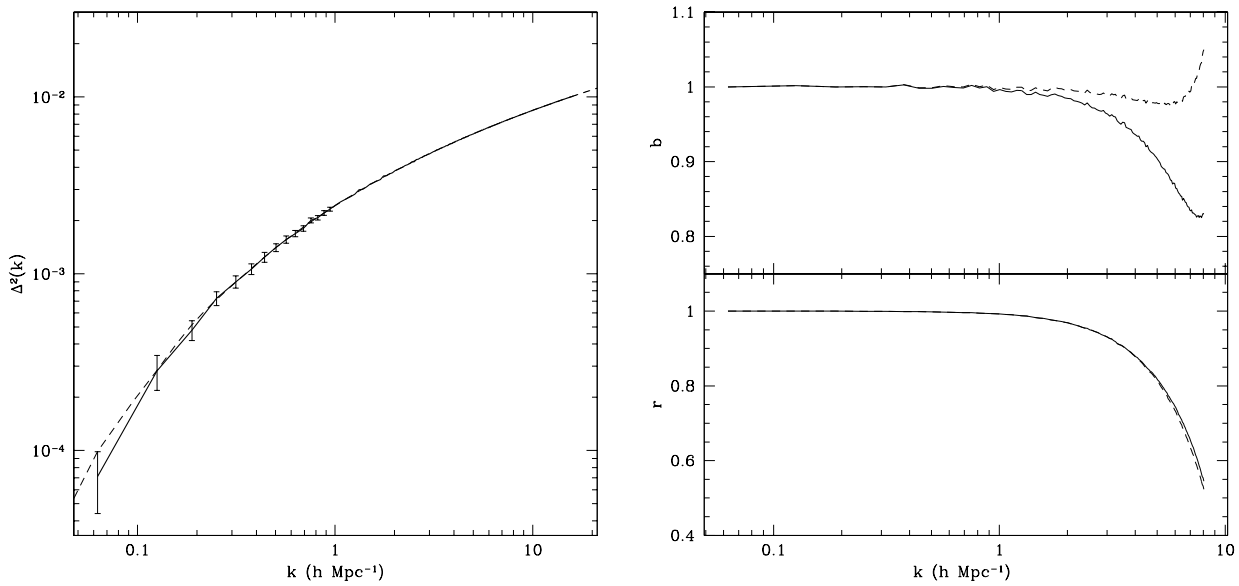

(a) **(b)**

Figure 6.2 In (a), the realization power spectrum (solid line) of the grid-constructed matter density field is plotted against the template CMBFAST power spectrum (dashed line), out to the grid Nyquist frequency. In (b), the bias and cross-correlation between the dark matter particles and the matter density field are plotted for both the potential method (solid lines) and the force method (dashed lines), out to the mean interparticle spacing.

random field $\delta(\boldsymbol{k})$. In Fourier space, the random sampling can be expressed as

$$\delta(\boldsymbol{k}) = n(\boldsymbol{k})w(\boldsymbol{k}), \tag{6.29}$$

where the kernel is given by

$$w(\boldsymbol{k}) \equiv \sqrt{P_i(k)}, \tag{6.30}$$

and $n(\boldsymbol{k})$ is the Fourier transform of a random white noise field $n(\boldsymbol{x})$. The white noise field has random phases and equal power at all scales, and is normalized so that $\langle|n(k)|^2\rangle = 1$. This realization Gaussian random field will also have random phases and a power spectrum,

$$\langle|\delta(\boldsymbol{k})|^2\rangle = P_i(k)\langle|n(\boldsymbol{k})|^2\rangle, \tag{6.31}$$

that is indeed given by the initial power spectrum. Note that the inverse transform of equation (6.29) can be written as the convolution

$$\delta(\boldsymbol{x}) = \int n(\boldsymbol{x'})w(\boldsymbol{x} - \boldsymbol{x'})d^3x'. \tag{6.32}$$

The Gaussian random field $\delta(\boldsymbol{k})$ is then inversed Fourier transformed to give the matter density field $\delta(\boldsymbol{x})$ on a grid.

Once the matter density field has been determined, the gravitational force field can be calculated and used to construct the comoving displacement field

$$\boldsymbol{dx}(\boldsymbol{x}) = -\frac{1}{4\pi G\bar{\rho}}\boldsymbol{\nabla}\Phi, \tag{6.33}$$

and the proper peculiar velocity field

$$\boldsymbol{v}_p(\boldsymbol{x}) \equiv a\frac{d\boldsymbol{x}}{dt} = a\frac{\dot{D}}{D}\boldsymbol{dx}, \tag{6.34}$$

where $\bar{\rho}$ is the comoving mean density, $D(t)$ is the linear growth factor, and $\dot{D}$ is its time derivative. Note that the comoving peculiar velocity $\boldsymbol{v}(\tau) \equiv d\boldsymbol{x}/d\tau$ used in the code is then given by $\boldsymbol{v}(\boldsymbol{x}) = a\boldsymbol{v}_p(\boldsymbol{x})$. The construction of the displacement and velocity fields require calculating the gravitational force field. This can be accomplished in two ways. In the potential method, the potential is first calculated and then finite-differenced to get the force field. In the force method, the force field is calculated directly. The dark matter particles are displaced from a uniform distribution and their velocities are determined by interpolating from the grid. The gas distribution is taken to trace the matter distribution and its initial conditions are readily obtained from the grid-constructed fields.

In Figure 6.2(a), the power spectrum (solid line) from a sample realization is plotted against the template initial power spectrum (dashed line). The dimensionless power spectrum $\Delta^2(k)$ is defined as

$$\Delta^2(k) = \frac{L^3}{2\pi^2}k^3P(k), \tag{6.35}$$

where $k = 2\pi/\lambda$ is the comoving wavenumber. In the linear regime, the power spectrum is characterized by $\Delta^2 \ll 1$. For $\Delta^2 \sim 1$, the overdensity reaches the mildly nonlinear value $\delta \sim 1$. The realization power $P(k)$ is calculated using FFTs and averaged in bins that span the range $k \pm \Delta k/2$ and have widths $\Delta k = 2\pi/L$. For a Gaussian random field, the fractional error in the power spectrum is given by

$$\frac{\Delta P}{P(k)} \sim \sqrt{\frac{2}{N(k)}}, \tag{6.36}$$

where $N(k)$ is the number of sampling modes for wavenumber bin k. In a finite-volume periodic box, the number of modes goes as $N(k) \propto k^2$. The grid-constructed matter density field has a power spectrum that statistically matches the template CMBFAST spectrum. At large wavenumbers or small scales, the agreement is very good due to the high number of sampling modes per bin. The deviations at small wavenumbers or large scales are expected due to sample variance. Averaging over multiple realizations, one reduces sample variance.

Figure 6.2(b) compares the realization matter power spectrum $P_{mm}(k)$ and the realization dark matter power spectrum $P_{pp}(k)$. To measure the dark matter power spectrum, the 256^3 particles

are first mapped onto the 512^3 grid using the CIC mass assignment scheme and then Fourier transformed. The statistical correlation between the particles and grid can be quantified with the bias

$$b(k) \equiv \sqrt{\frac{P_{pp}(k)}{P_{mm}(k)}}, \tag{6.37}$$

and cross-correlation

$$r(k) \equiv \frac{P_{mp}(k)}{\sqrt{P_{mm}(k)P_{pp}(k)}}, \tag{6.38}$$

where $P_{mp}(k) \equiv \langle \delta_m(k)\delta_p(k) \rangle$ is the cross spectrum. The cross-correlation coefficient or stochasticity parameter can have values $-1 \leq r \leq 1$. On large scales, the particles are perfectly correlated with the grid. No bias is found for either the potential method (solid lines) or force (dashed lines) method. On small scales, finite resolution reduces the correlation. The force method comes closest to achieving unity bias. The very small reduction in power is a result of the smoothing effect of the CIC mass assignment scheme. The uprise in power near the mean interparticle spacing is due to Poisson noise in the mapping process. In the potential method, the finite-differencing degrades the resolution, resulting in less power on small scales. Note that any gain from computing the particle displacement using the force method will be lost since the cosmological hydro code uses the potential method to calculate the gravitational force field. Therefore, it suffices to use the potential method for generating initial conditions, which are to be used in PM codes.

6.5 Cosmological Tests

The strengths and limitations of the cosmological moving frame code are quantified with the Zeldovich pancake test and the self-similar scaling test. The one-dimensional plane-parallel Zeldovich pancake problem is a rigourous test that involves both strong shock heating and adiabatic cooling in the presence of high Mach numbers. In the self-similar scaling test, scale-free initial conditions are expected to evolve in a self-similar fashion, but numerical limitations like finite grid resolution and finite simulation volume break the symmetry.

6.5.1 Zeldovich Pancake Test

A standard cosmological test is the Zeldovich pancake problem (Zeldovich, 1970). In the one-dimensional plane-parallel case, a sinusoidal perturbation evolves linearly initially and eventually collapses to form a caustic or pancake. This vigourous test involves both hydrodynamics and self-gravity and strong shocks are formed in the presence of high Mach numbers.

For a flat cosmology, the initial conditions and the evolution of the sinusoidal perturbation

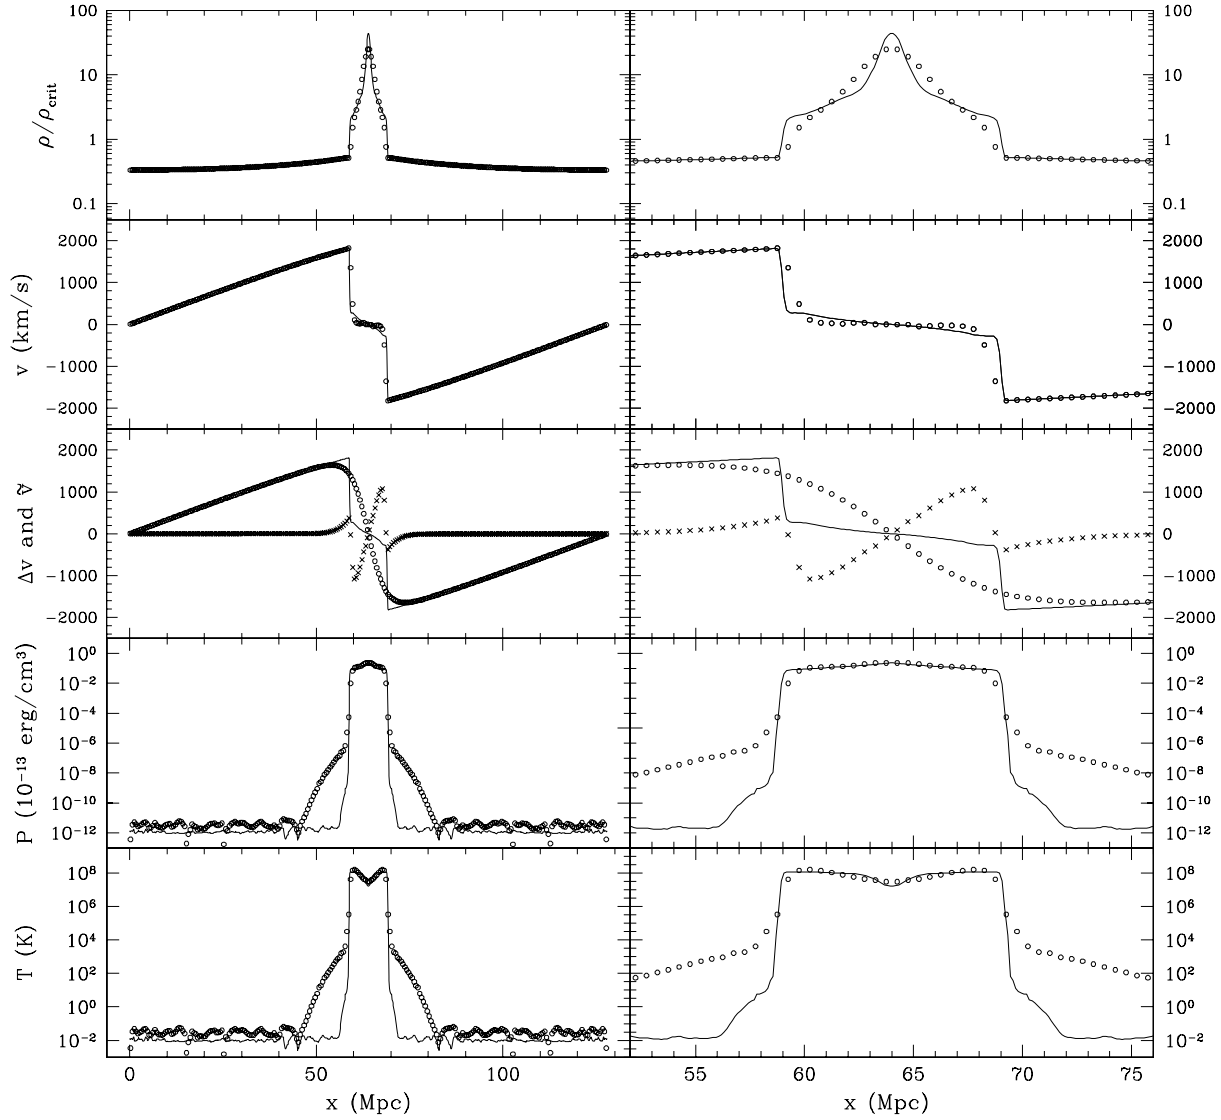

Figure 6.3 Cosmological pancake test results at redshift $z = 0$. A lower resolution pancake from 256 cells (open circles) is plotted against a high resolution pancake with 1024 cells (solid line). In the middle panels, the grid velocity (open circles) and the local velocity (crosses) are plotted to show the velocity decomposition.

(Zeldovich, 1970; Anninos & Norman, 1994) can be determined from the following equations:

$$x(q,z) = q - \frac{1+z_c}{1+z}\frac{\sin(kq)}{k}, \tag{6.39}$$

$$\rho(x,z) = \rho_0 \left[1 - \frac{1+z_c}{1+z}\cos(kq)\right]^{-1}, \tag{6.40}$$

$$v_p(x,z) = -H_0 \frac{1+z_c}{\sqrt{1+z}}\frac{\sin(kq)}{k}, \tag{6.41}$$

The solutions hold in both the linearly and moderately nonlinear regime prior to the redshift z_c where the gravitational collapse forms a pancake. For a given Eulerian comoving coordinate x, the corresponding Lagrangian coordinate q can be calculated by inverting equation (6.39). The comoving wavenumber $k = 2\pi/\lambda$ sets the wavelength of the perturbation. The density ρ is a comoving density and ρ_0 is the comoving average density. The proper velocity v_p has units of km/s. For an ideal adiabatic gas, the thermodynamic solution is given by:

$$T(x,z) = T_i \left[\left(\frac{1+z}{1+z_i}\right)^3 \frac{\rho(x,z)}{\rho_0}\right]^{\gamma-1} \tag{6.42}$$

$$P(x,z) = \frac{k_B}{\mu m_H}[(1+z)^3\rho(x,z)]T(x,z). \tag{6.43}$$

The temperature T is a monotonic function of the density and T_i is the initial average temperature at the initial redshift z_i. The proper pressure P is given by the equation of state for an ideal gas with mean molecular mass μ.

Following Bryan et al. (1995), the pancake initial conditions are generated with the parameters: $z_i = 100$, $z_c = 1$, $\Omega_m = 1$, $h = 0.5$, $\lambda = 64h^{-1}$ Mpc, $\rho_0 = \rho_{crit}$, and $T_i = 100$ K in a purely baryonic universe. Initially, the maximum Mach number is ~ 200 and this makes the test challenging for an Eulerian code. In Figure 6.3, the nonlinear results at $z = 0$ are plotted. A lower resolution pancake from 256 cells (open circles) is plotted against a high resolution pancake with 1024 cells (solid line). For both runs, a Gaussian smoothing radius of 8 grid cells is used to compute the smoothed grid velocity. The formation of the pancake is due to two cold bulk flows colliding and collapsing to form a caustic. The kinematics of the collapse is very well captured. In the lower resolution run, the pancake is under-resolved but the amount of diffusion is typical of that found by other second-order accurate TVD codes with fixed grids (Ryu et al., 1993; Pen, 1998a). In general for tests involving the formation of caustics, Eulerian schemes are more diffusive than Lagrangian schemes, which benefit from having high dynamic range in scale and density.

The more difficult part of the pancake test is the thermodynamic evolution. The initially cold gas is shock heated to very high temperatures near the caustic, while away from the shock the cold gas is expected to cool adiabatically as the universe expands. Numerically, it is difficult to capture the large range in temperature, over 10 orders of magnitude for this particular test. In

addition, conventional Eulerian codes suffer from a large amount of spurious heating due to the high Mach numbers involved. The moving frame approach allows us to remove the bulk flow and reduce the Mach numbers. Away from the shock, the grid velocity approximates the bulk flow very well with $\tilde{v} \approx v$. As a result, the Mach numbers are now small because the local velocity Δv is comparable to the low sound speed of the cold gas. No significant heating of the cold gas is found away from the shock. The spurious heating has been minimized by adding the gravitationally induced accelerations to the grid velocity. In the PPM-Bryan code (Bryan et al., 1995) there is an artificial minimum temperature floor of 1 K, but we impose no such restriction and the temperature floor in our simulations is consistent with the expected adiabatic cooling. The TVD-entropy code (Ryu et al., 1993) should be able to accurately capture the adiabatic evolution since no shocks are formed and entropy is conserved.

In the collapsing region, the grid velocity turns over and becomes zero at the center of the pancake. In the postshocked region, the local velocity is large but the gas is very hot and therefore the Mach numbers are low. The postshocked gas pressure and temperature are highly resolved. Immediately outside of the expanding shock front, some heating is seen in the lower resolution run but not in the high resolution pancake. In this region, the grid velocity turns over and the Mach numbers increase. In the TVD-entropy code, the entropy equation becomes invalid in the vicinity of shocks. The TVD-entropy code is expected to do worse than the moving frame code since the Mach numbers are even larger without the velocity decomposition. With the PPM-bryan code, the shock is highly resolved and the temperature profile is accurate, except at the base of the expanding shock front where it hits the imposed 1K temperature floor.

6.5.2 Self-similar Scaling Test

The self-similar evolution of scale-free cosmological initial conditions can be used to quantify the numerical limitations of hydrodynamic simulations. In an Einstein-de Sitter universe where $\Omega_m = 1$, scale-free initial conditions with power-law correlation function will evolve as

$$\xi(r) = \left(\frac{r_0}{r}\right)^{\gamma} \propto a^2. \tag{6.44}$$

Hence, the correlation function obeys the self-similar transformation $\xi(r, t_1) \rightarrow \xi(s, t_2)$ where

$$\frac{s}{r} = \left[\frac{a(t_2)}{a(t_1)}\right]^{-2/\gamma}. \tag{6.45}$$

In Fourier space, the power spectrum of the density modes will evolve as

$$P(k) \propto k^n \propto a^2, \tag{6.46}$$

where $n = \gamma - 3$. The power spectrum follows the the transformation $P(k, t_1) \rightarrow P(l, t_2)$ where

$$\frac{l}{k} = \left[\frac{a(t_2)}{a(t_1)}\right]^{2/n}. \tag{6.47}$$

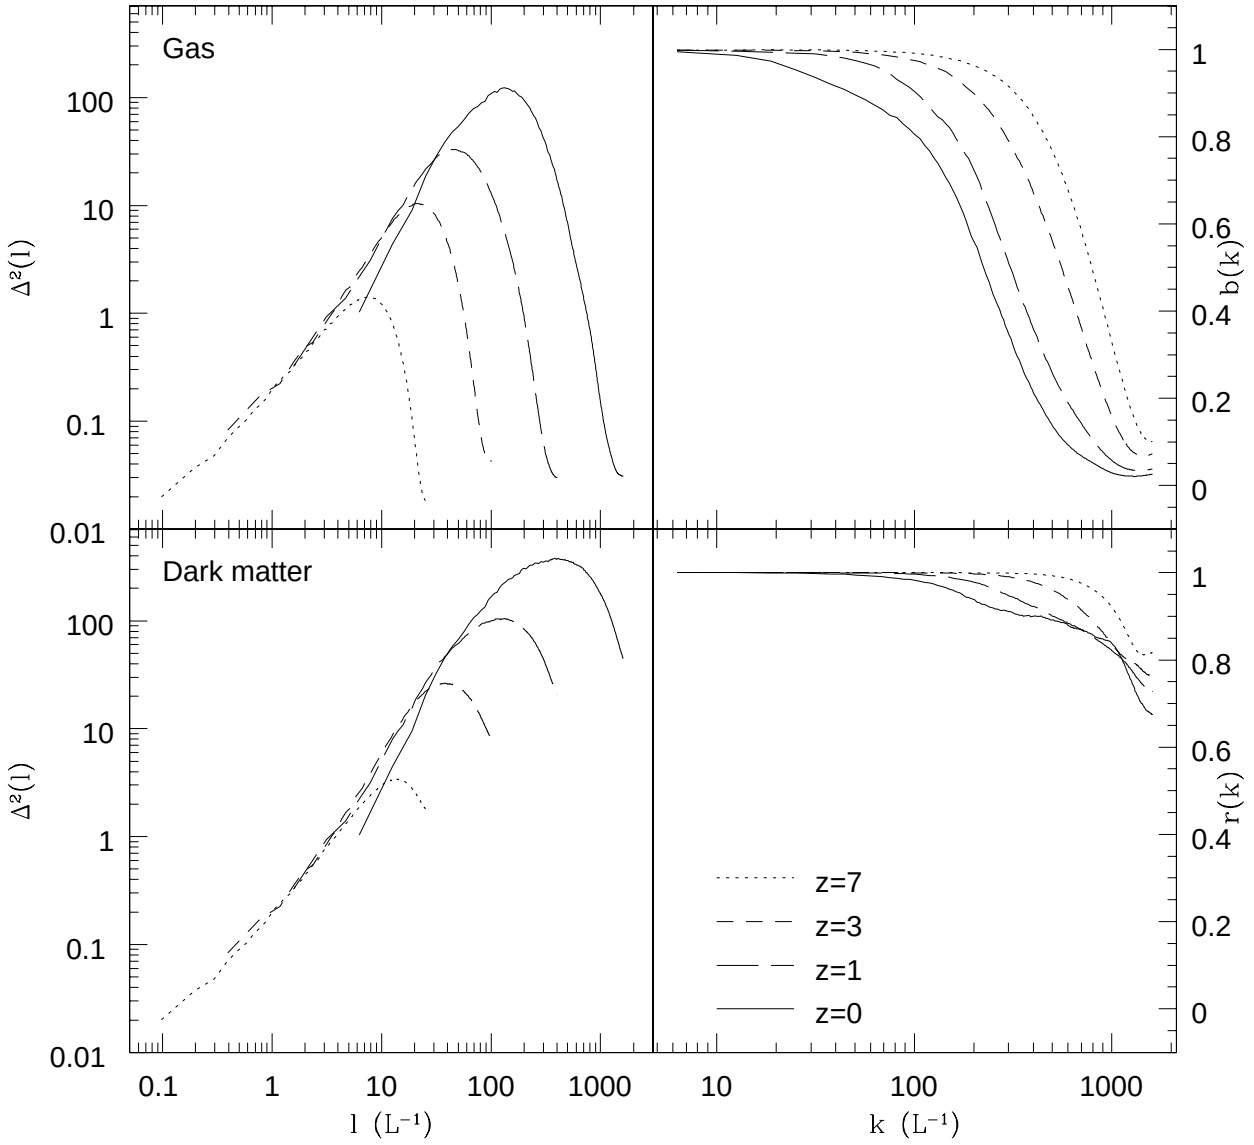

(a) **(b)**

Figure 6.4 Cosmological self-similar test for $n = -2$ scale-free initial conditions. The scaled, dimensionless gas and dark matter power spectra are plotted in (a) while the bias and cross-correlation coefficient are plotted in (b). The MACH code is run with 512^3 grid cells and 1:1 gas to dark matter ratio.

Numerical artifacts arising from the finite grid resolution and the finite simulation volume will cause the evolution to deviate from the self-similar scaling.

A three-dimensional self-similar test with initial spectral index $n = -2$ or equivalently $\gamma = 1$ is run to test the cosmological MACH code. The initial conditions are normalized such that the linearly extrapolated correlation function at redshift $z = 0$ has a correlation length r_0 equal to $1/4$ of the box length L. The gas to matter ratio Ω_b/Ω_m is chosen to be equal to the canonical ratio $1/6 = 0.05/0.30$. The simulation is run with 512^3 grid cells and 256^3 dark matter particles, starting from an initial redshift of $z = 120$, where the real space density field has $|\delta_{max}| < 1$. Plotted in Figure 6.4(a) are the scaled gas and dark matter power spectra. The self-similar scaling holds for a wide range of length and time scales. The finite volume has no significant effect on the clustering at redshift $z = 1$ where the correlation length is $r_0 = L/16$. At redshift $z = 0$ where the correlation length is expected to be $r_0 = L/4$, the clustering of large-scale structure is severely affected by finite volume effects and the self-similar scaling breaks down.

On large scales, the gas is expected to be highly correlated with the dark matter and the evolution is driven primarily by gravity. On small scales, gas pressure can be quite large due to gravitational collapse and shock heating and therefore, the gas is expected to deviate from the dark matter. One can define the linear bias parameter as

$$b(k) = \sqrt{\frac{P_{gas}(k)}{P_{cdm}(k)}}, \tag{6.48}$$

and parametrize the correlation between the gas and the dark matter using the cross-correlation coefficient

$$r(k) = \frac{P_{gm}(k)}{\sqrt{P_{gas}(k)P_{cdm}(k)}}, \tag{6.49}$$

where $P_{gm}(k) \equiv \langle \delta_{gas}(k)\delta_{cdm}(k) \rangle$ is the cross-spectrum. The cross-correlation coefficient or stochasticity parameter can have values $-1 \leq r \leq 1$. Figure 6.4(b) shows that the correlation is perfect at the largest scales and weakens slightly at smaller scales. The bias diminishes with decreasing scale since gas pressure counteracts gravitational infall and thereby reduces gas power. The bias and stochasticity curves should be identical at all times if self-similarity holds. However, self-similar scaling is broken at small scales because of finite grid resolution. In addition, gravity introduces a length scale since the CIC mass assignment scheme and the finite-differencing force scheme together soften the force resolution.

An interesting exercise is to plot a vector field diagram comparing the velocity field $\boldsymbol{v}$ and the grid velocity field $\tilde{\boldsymbol{v}}$. An arbitrary $x - y$ planar slice is chosen and the x and y components of the velocities are used to construct the vector field diagram shown in Figure 6.5. For the grid velocity, the x component is obtained by smoothing in the x direction and likewise for the y component. In Figure 6.5(a), the vector fields are plotted for the entire 512×512, but with subsampling so that

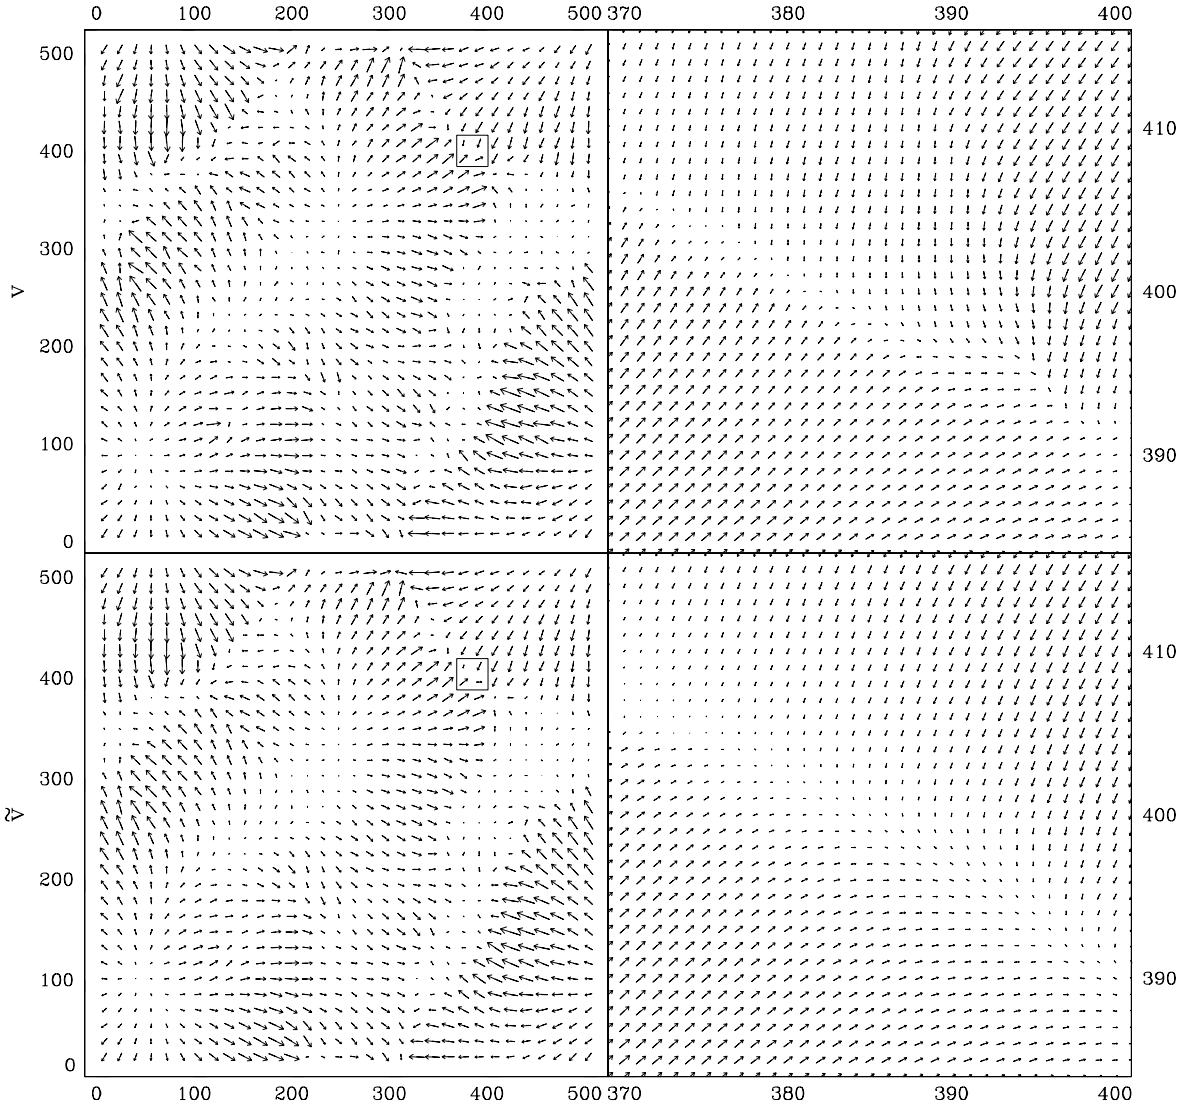

(a) **(b)**

Figure 6.5 Vector field diagram for a two-dimensional slice at z=3. In (a), the global velocity field and the smoothed grid velocity field on the entire 512×512 grid are compared. The subsampled 32×32 vector field shows the large-scale bulk flow. In (b) is a zoom in on a convergent region, shown at full resolution.

only 32×32 vectors are shown. The subsampling has the effect of displaying only the large-scale bulk flow. Overall the grid velocity approximates the bulk flow very well. Figure 6.5(b) zooms in a convergent region. At full resolution on smaller scales, the grid velocity appears very smooth, a requirement for the algorithm to work accurately.

6.6 Cosmological Simulations

Numerical simulations must converge over a large range in scale, mass, and temperature. High resolution is required to resolve small-scale, nonlinear structures. Large box sizes are needed to capture large-scale correlations and power. Cosmological simulations require large box sizes to capture the dynamical evolution of large-scale structure since the cosmological velocity field is dominated by large-scale flows spanning $\sim 10^2$ Mpc. This puts more demand on the Mach number dynamic range of Eulerian hydrodynamic codes. With the moving frame approach, the numerical Mach numbers are minimized and this makes capturing the thermodynamic evolution more tractable. The cosmological MACH code is highly suited for simulating the evolution of the IGM to study the Lyman alpha forest, the Sunyaev-Zeldovich effect, and the X-ray background.

A cosmological simulation with 512^3 grid cells and 256^3 dark matter particles in a 100 h^{-1} Mpc box is run to examine the nonlinear evolution of the large-scale structure and the Mach number distributions in the high redshift IGM. WMAP cosmology (Spergel et al., 2003) is chosen with parameters: $\Omega_m = 0.27$, $\Omega_\Lambda = 0.73$, $\Omega_b = 0.044$, $\sigma_8 = 0.84$, and $h_0 = 0.7$. The initial conditions are identical to the sample realization from Section 6.4 with gas and dark matter initial power spectra plotted in Figure 6.2.

6.6.1 Nonlinear Power Spectra

The gas and dark matter mass power spectra at redshifts $z = 0$, 1, 3, and 7 are computed using FFTs. The gas density field is readily defined on a periodic 512^3 grid. The clustered distribution of 256^3 dark matter particles is mapped onto a higher resolution 1024^3 grid to avoid reduction in power due to smoothing. Poisson noise have been subtracted from the dark matter power spectra, except at $z = 7$ where the power at small scales does not exceed the shot noise power.

In Figure 6.6, the simulated power spectra are plotted against the nonlinear fitting functions of Smith et al. (2003), which have been calibrated using P3M N-body simulations. For this redshift range and on scales $k \lesssim 10h$ Mpc^{-1}, the nonlinear fitting functions are similar to the nonlinear scaling formula of Peacock & Dodds (1996). On smaller scales, the latter under predicts the clustering power of the dark matter distribution. The simulated dark matter and predicted matter spectra are in good agreement, except at scales near the grid spacing. At redshift $z = 0$, the comparison is good to 10% on frequencies $k \lesssim 3h$ Mpc^{-1} or comoving wavelengths $\lambda \gtrsim 2h^{-1}$ Mpc. More than half the power is lost on scales less than 5 grid cells because the force resolution in the

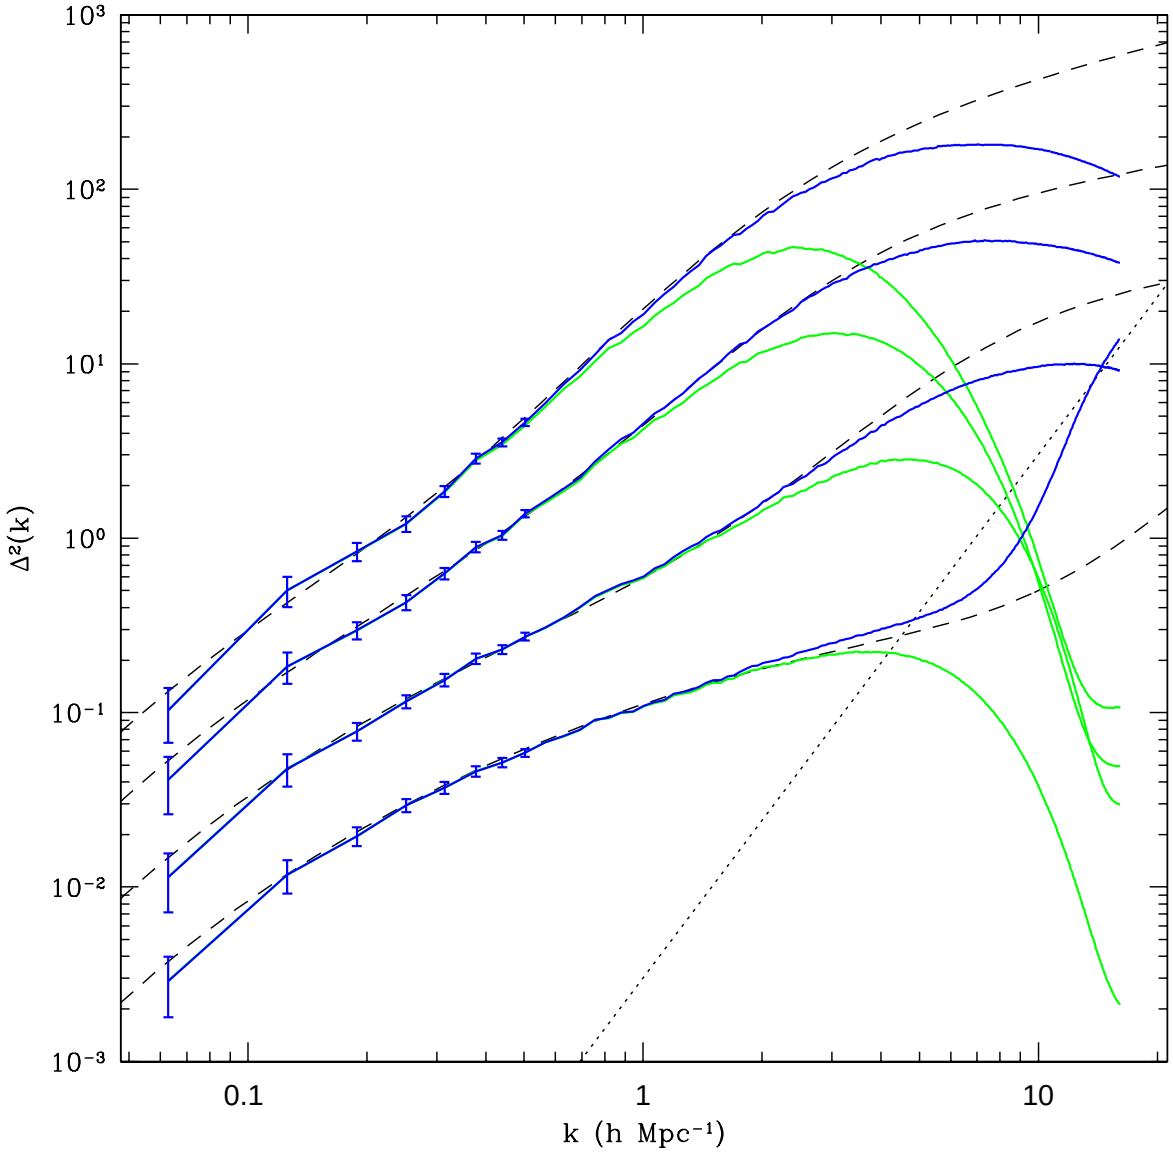

Figure 6.6 Nonlinear evolution of the dark matter (blue solid lines) and gas (green solid lines) power spectra at redshifts $z = 0$, 1, 3, and 7 from a cosmological moving frame hydro simulation with 512^3 grid cells and 256^3 dark matter particles in a 100 h^{-1}Mpc box. The simulated dark matter power spectra agree with the matter power spectra (dashed lines) predicted using the fitting functions of Smith et al. (2003). Poisson noise have been subtracted from the dark matter power spectra, except at $z = 7$ where the power does not exceed the shot noise power (dotted line).

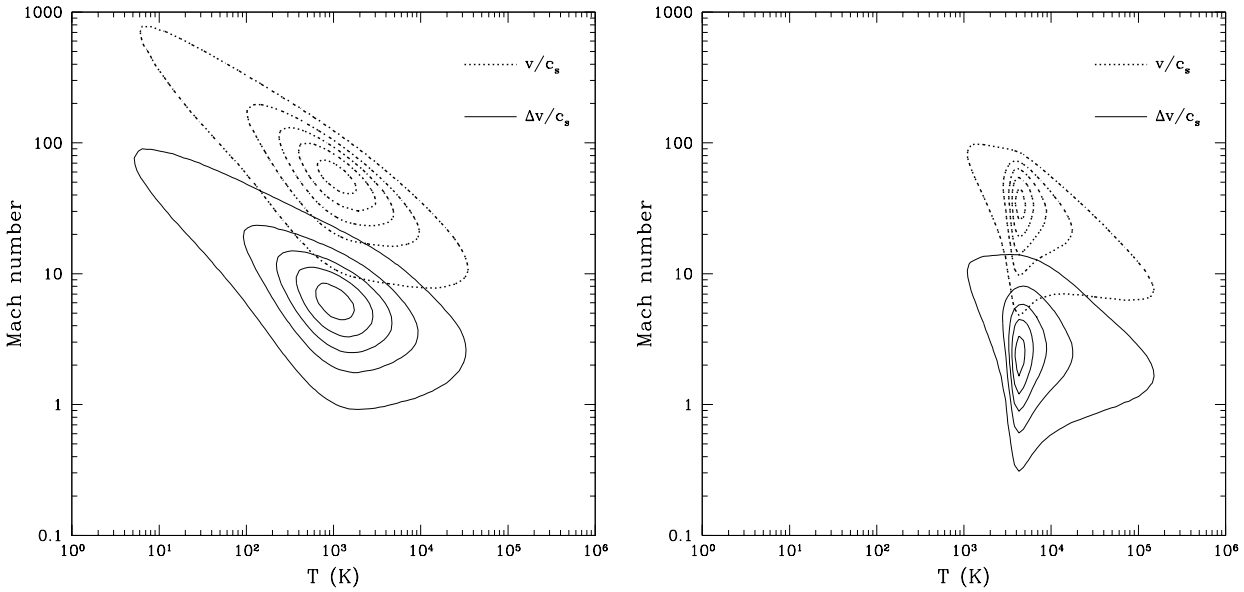

Figure 6.7 Volume-weighted Mach number distributions at $z = 7$ with no reionization (a) and at $z = 3$ with reionization (b). Contour levels are for 10%, 30%, 50%, 70%, and 90%. We plot both the local Mach number $M_l = \Delta v/c_s$ (solid lines) and the global Mach number $M_g = v/c_s$ (dotted lines) to show how the moving frame approach minimizes Mach numbers.

PM scheme is softened near the grid scale.

At large scales approaching the box size, the evolution is consistent with linear growth. The observed deviations are expected since the realization initial dark matter power spectrum has a deficit in power at large scales to begin with (see Figure 6.2). On linear and moderately nonlinear scales, the simulated gas and dark matter are highly correlated with no bias. On nonlinear scales, the gas loses power primarily because the pressure prevents gravitational collapse.

6.6.2 Mach Number Distribution

Plotted in Figure 6.7(a) is the Mach number distribution at redshift $z = 7$ with no reionization. In the moving frame code, the local Mach number $M_l = \Delta v/c_s$ is the quantity of interest whereas in a standard Eulerian code, it is the global Mach number $M_g = v/c_s$. The contours are drawn based on the probability density function $p(\log T, \log M)$. The peak of the local Mach number distribution is more than an order of magnitude smaller than the peak of the global Mach number distribution. At redshift $z = 7$, $\sim 20\%$ of the gas have $M_l < \sqrt{10}$ while $\sim 75\%$ of the gas have $M_l < 10$. In comparison, only ~ 3 of the gas have $M_g < 10$. At high redshifts prior to reionization, the Mach numbers are intrinsically high due to the low temperatures in the IGM.

A uniform and instantaneous reionization at redshift $z = 7$ is imposed, boosting the temperature

of the IGM by $T = 2 \times 10^4$ K. Figure 6.7(b) is a plot of the Mach number distribution at redshift $z = 3$, where the high redshift Lyman alpha forest is an important cosmological probe. Here, the peak of the local Mach number distribution is of order unity and close to 2 orders of magnitude smaller than the peak of the global Mach number distribution. At redshift $z = 3$, $\sim 65\%$ of the gas have $M_l < \sqrt{10}$ while $\sim 95\%$ of the gas have $M_l < 10$. In comparison, $\sim 10\%$ of the gas have $M_g < 10$. Hence, it is much more tractable to capture the thermodynamic evolution of the IGM with the moving frame hydrodynamic code.

6.7 Summary

For cosmological applications, the moving frame algorithm has been combined with a PM N-body scheme to simultaneously capture the hydrodynamic evolution of the baryonic gas and the gravitational evolution of the collisionless dark matter particles. In the cosmological pancake test, the code successfully captures both the nonlinear kinematic and thermodynamic evolution. Although the temperature ranges over 10 orders of magnitude for this particular test, the thermodynamic profiles are highly resolved. Away from the shock, the adiabatic cooling of the cold gas is accurately simulated with no spurious heating seen. The MACH code does not suffer from the high Mach number problem encountered in previous Eulerian hydrodynamic codes. Tests conducted with scale-free cosmological initial conditions show that the self-similar scaling laws hold over a wide range of length and time scales. The numerical Mach numbers in cosmological simulations are reduced using the moving frame algorithm.

Chapter 7

Out-of-core Cosmological Hydrodynamic Simulations

7.1 Introduction

This chapter describes a new out-of-core hydrodynamic (OCH) code for high resolution cosmological simulations. The code is based on the Hydro&N-body implementation described in Chapter 6, where the moving frame algorithm is used to solve the hydrodynamics of the baryonic gas and the particle-mesh (PM) N-body algorithm tracks the gravitational evolution of the collisionless dark matter. It is designed for cosmological applications where high mass resolution is required at all scales.

The challenges in doing out-of-core computation stem from the requirement that the data domain be decomposed into blocks which can fit in memory. The hydrodynamics of the ideal gas is a local process that is straightforward to solve on the decomposed domain with the addition of buffer regions. However, gravity is a global force and Poisson solvers operate on the global density field. Fast Fourier transforms (FFTs) are the optimal solvers in a standard PM scheme, but they involve global transposes and computing the transforms out-of-core will be intolerably slow. This problem is addressed using a two-level mesh scheme (Couchman, 1991) where the short-range force is solved on a high-resolution, localized mesh while the long-range force is captured on a lower resolution, global mesh. The two-level mesh gravity solver is memory-efficient and allows FFTs to be performed on data stored entirely in memory.

In this chapter, I describe the out-of-core algorithm and highlight the steps required for doing out-of-core computation, including optimizations to reduce I/O overhead by performing disk operations and numerical calculations concurrently. The two-level mesh gravity solver is described and its accuracy is compared to that of a standard one-level mesh solver. In addition, I present an out-of-core initial conditions generator that can be used to construct large data sets for high resolution cosmological simulations.

7.2 Out-of-core Algorithm

The three-dimensional out-of-core hydro (OCH) code is based on the Hydro&N-body implementation described in Chapter 6 and in Trac & Pen (2004a). The out-of-core implementation requires both the division of data into blocks that can fit in memory and the decomposition of physical laws with global dependence into short and long range components. The code uses a two-level mesh algorithm where short-range physics is solved on a high-resolution, localized mesh while long-range physics is captured on a lower resolution global mesh. The global mesh is 4 times coarser in each dimension, reducing memory usage by a factor of 64 when calculating the long-range interactions. The standard technique of having buffer regions around the local blocks ensures that the short-range interactions are calculated with no boundary effects imposed by the domain decomposition.

7.2.1 Hydro&N-body

The hydrodynamics of the baryonic gas is solved using an Eulerian algorithm based on the moving frame approach and the second-order accurate total variation diminishing (TVD) scheme (Harten, 1983). The Eulerian conservation equations are solved in an adaptive frame moving with the fluid where local fluid variables can be directly tracked. The moving frame approach minimizes the numerical Mach numbers, allowing high resolution capturing of shocks and preventing spurious oscillations in both subsonic and supersonic regions. Time integration is performed using a second-order accurate Runge-Kutta technique. The moving frame hydro algorithm uses an operator-splitting technique (Strang, 1968) to efficiently solve the various flux terms in the Euler equations. To get second-order accuracy, each hydro iteration is composed of a double time step and the order of operations is reversed in the second time step. A single gravity step is done between the forward and reverse hydro steps.

The gravitational evolution of the dark matter particles is tracked using the standard PM N-body scheme (Hockney & Eastwood, 1988). The collisionless particles are advected by solving Newton's equations of motion in an FRW universe. Time integration is performed using an explicit, second-order accurate Runge-Kutta technique that allows synchronization between the dark matter particles and the gas. Mass assignment onto the grid and force interpolation from the grid are conducted with the 'cloud-in-cell' (CIC) scheme.

The hydrodynamics of the ideal gas is a local process that is straightforward to solve on the decomposed domain with the addition of buffer regions. For the particles, the mass assignment and force interpolation are also local procedures. However, gravity is a global force and Poisson solvers operate on the global density field. Fast Fourier transforms (FFTs) are the optimal solvers in a standard PM scheme, but they involve global transposes and computing the transforms out-of-core will be intolerably slow. This problem is addressed using a two-level mesh scheme where the short-range force is solved on a high-resolution, localized mesh while the long-range force is captured on

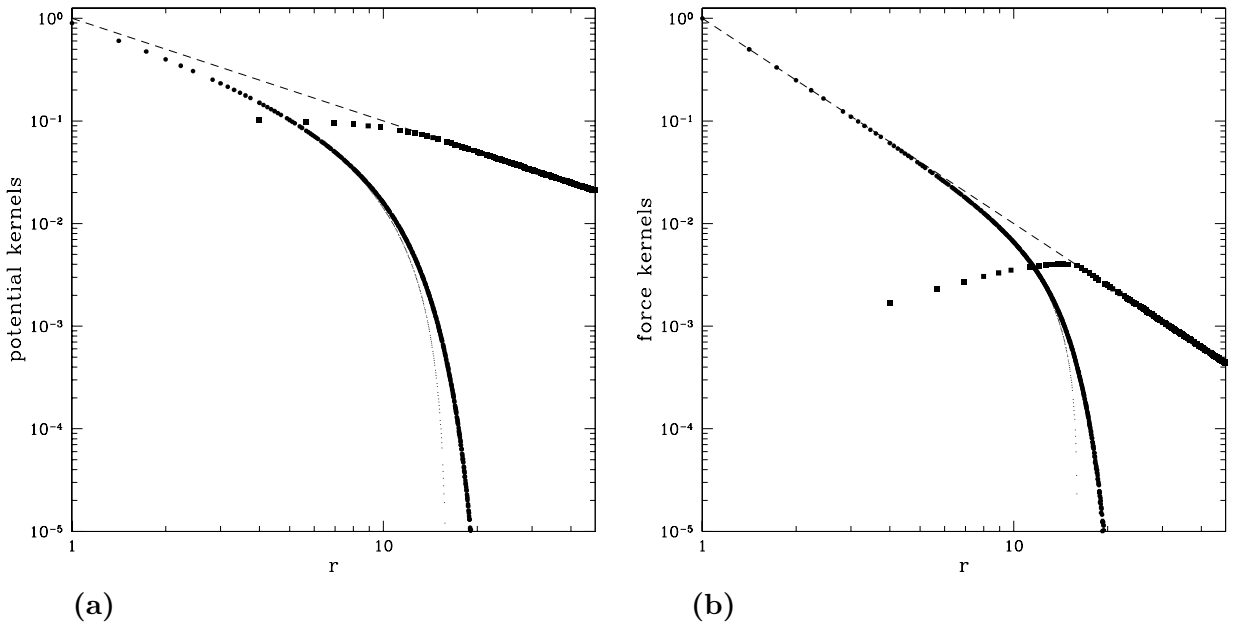

Figure 7.1 The decomposition of the potential kernel (a) and radial force kernel (b) for the two-level mesh gravity solver. The long-range (filled squares), uncorrected short-range (small dots), and corrected short-range (filled circles) components are plotted. For comparison, the $1/r$ isotropic potential kernel (a) and the $1/r^2$ isotropic force kernel are illustrated by the dashed lines. Here, the uncorrected short-range cutoff is $r_g = 16$ fine grid cells and the corrected cutoff is approximately 20 grid cells.

a lower resolution, global mesh. The two-level mesh gravity solver is memory-efficient and allows FFTs to be performed on data stored entirely in memory.

7.2.2 Two-level Mesh Gravity Solver

Gravity is a global force that is decomposed into short and long range components and solve with the two-level mesh scheme (Couchman, 1991). Poisson's equation

$$\nabla^2\Phi = 4\pi G\rho, \tag{7.1}$$

relates the gravitational potential Φ to the density field ρ and the general solution can be written as the convolution,

$$\Phi(\boldsymbol{x}) = \int \rho(\boldsymbol{x}')w(\boldsymbol{x}-\boldsymbol{x}')d^3x', \tag{7.2}$$

where the isotropic kernel is given by

$$w(r) = -\frac{G}{r}. \tag{7.3}$$

In the standard mesh approach, the convolution is done in Fourier space and fast Fourier transforms (FFTs) are used to convert the discrete density field and recover the discrete potential field. Since FFTs are highly nonlocal and involve global transposes, computing the transforms out-of-core using non-sequential disk accesses will be intolerably slow. However, the two-level mesh scheme allows us to apply the FFTs to data that are entirely stored in memory.

Th two-level mesh scheme decomposes the potential kernel $w(r)$ into a short-range component

$$w_s(r) = \begin{cases} w(r) - \alpha(r) & \text{if } r \leq r_g, \\ 0 & \text{otherwise,} \end{cases} \tag{7.4}$$

and a long-range term

$$w_l(r) = \begin{cases} \alpha(r) & \text{if } r \leq r_g, \\ w(r) & \text{otherwise,} \end{cases} \tag{7.5}$$

where the short-range gravity cutoff r_g is a free parameter. The function $\alpha(r)$ is chosen to be a polynomial,

$$\alpha(r) = G(a + br^2 + cr^4), \tag{7.6}$$

whose coefficients,

$$\begin{aligned} a &= -\frac{27}{16r_g}, \\ b &= \frac{7}{8r_g^3}, \\ c &= -\frac{3}{16r_g^5}, \end{aligned} \tag{7.7}$$

are determined from the conditions

$$\begin{aligned} \alpha(r_g) &= w(r_g), \\ \alpha'(r_g) &= w'(r_g), \\ \alpha''(r_g) &= w''(r_g). \end{aligned} \tag{7.8}$$

These restrictions ensure that the long-range kernel smoothly turns over near the cutoff and that the short-range term smoothly goes to zero at the cutoff.

The long-range potential $\Phi_l^c(\boldsymbol{x})$ is computed by performing the convolution over the coarse-grained global density field $\rho^c(\boldsymbol{x})$. The superscript c denotes that the discrete fields are constructed on a coarse grid. The total density field has contributions from both the dark matter particles and baryonic gas cells. Mass assignment onto the coarse grid is accomplished using the CIC interpolation scheme (Hockney & Eastwood, 1988) with cloud shape being the same as a coarse

cell. The long-range force field $\boldsymbol{f}_l^c(\boldsymbol{x})$ is obtained by finite differencing the long-range potential and force interpolation is carried out using the same CIC scheme to ensure no fictitious self-force.

Since the two-level mesh scheme uses grids at different resolutions, the decomposition given by equations (7.4) and (7.5) needs to be modified. In Fourier space, the long-range potential can be written as

$$\tilde{\Phi}_l(\boldsymbol{k}) = \tilde{\rho}^c(\boldsymbol{k})\tilde{w}_l^c(\boldsymbol{k}) = [\tilde{\rho}(\boldsymbol{k})\tilde{s}_\rho(\boldsymbol{k})][\tilde{w}_l(\boldsymbol{k})\tilde{s}_w(\boldsymbol{k})], \tag{7.9}$$

where $\tilde{s}_\rho(\boldsymbol{k})$ and $\tilde{s}_w(\boldsymbol{k})$ are the Fourier transforms of the mass smoothing window $s_\rho(\boldsymbol{x})$ and kernel sampling window $s_w(\boldsymbol{x})$, respectively. The mass smoothing window takes into account the CIC mass assignment scheme for constructing the coarse density field. The kernel sampling window corrects for the fact that the long-range kernel given by equation (7.5) is sampled on a coarse grid. In Fourier space, the corrected short-range potential kernel is now given by

$$\tilde{w}_s(\boldsymbol{k}) = \tilde{w}(\boldsymbol{k}) - \tilde{w}_l(\boldsymbol{k})\tilde{s}_\rho(\boldsymbol{k})\tilde{s}_w(\boldsymbol{k}), \tag{7.10}$$

and can be slightly anisotropic, particularly near the short-range cutoff.

For each local block in the domain decomposition, the short-range potential $\Phi_s^f(\boldsymbol{x})$ is computed by performing the convolution over the high-resolution density field $\rho^f(\boldsymbol{x})$. The superscript f denotes that the local fields are constructed on a fine grid. The short-range force field $\boldsymbol{f}_s^f(\boldsymbol{x})$ is then obtained by finite differencing the short-range potential.

Plotted in Figure 7.1(a) are the various components of the potential kernel for an initial short-range cutoff $r_g = 16$ fine grid cells. For comparison, the isotropic kernel $w(r)$ is illustrated by the dashed line. The long range kernel $w_l^c(\boldsymbol{x})$ is constructed on a coarse grid and the values in each cell are plotted with filled squares. The uncorrected and corrected short-range kernels $w_s^f(\boldsymbol{x})$ are constructed on a fine grid and plotted with small dots and filled circles, respectively. The corrected short-range kernel smoothly goes to zero at a radius now slightly larger than r_g. The new gravity cutoff $b_g = 20$ fine grid cells sets the size of the buffer for the local blocks.

Note that as an alternative to the potential method, the force can be calculated directed using the convolution,

$$\boldsymbol{f}_i(\boldsymbol{x}) = \int \rho(\boldsymbol{x}')w_i(\boldsymbol{x} - \boldsymbol{x}')d^3x', \tag{7.11}$$

where the anisotropic force kernels are given by

$$w_i(\boldsymbol{x}) = -G\frac{x_i}{r^3}. \tag{7.12}$$

The force kernels can be decomposed into short and long range components in a similar fashion to that described previously for the potential kernel. In Figure 7.1(b), the magnitude of the force kernels are plotted for a short-range cutoff $r_g = 16$ fine grid cells.

In the potential method, the finite differencing degrades the pair-wise force resolution by a few grid cells, but the net force on any given cell or particle is still highly accurate in general. The force method exactly reproduces the pair-wise inverse-square law on the grid, but not for particles since the CIC mass assignment scheme smoothes the mass and pair-wise force near the grid scale. Both the potential and force methods require one FFT to forward transform the density field, but the potential method only requires one inverse transform for the potential field while the force method requires three inverse transforms for the force components. Hence, the force method comes at the cost of two extra FFTs. For the high-resolution short-range force field, the relatively small accuracy trade-off of the potential method is preferred over the factor of 2 increase in computational work of the force method. For the lower-resolution long-range force field, the force method allows us to avoid finite differencing the coarse grid and further degrading the resolution of an already smoothed field. The extra factor of 2 increase in cost is affordable since the long-range force calculation is already a factor of 64 cheaper than the collective short-range force calculations.

In summary, the potential and force methods are combined to generate a two-level mesh gravity solver that is accurate, cost-effective, and memory-efficient. The short-range force is computed using the potential method while the long-range force is solved directly using the force method. The initial short-range cutoff $r_g = 16$ fine grid cells and a gravity cutoff $b_g = 20$ fine grid cells. A version of this two-level mesh gravity solver is used by Trac & Pen (2004b) and Merz et al. (2004).

7.2.3 Pair-wise Force Test

The accuracy of the two-level mesh gravity solver can be checked by performing a pair-wise force test. A particle is placed randomly in a box and the pair-wise forces for a number of randomly placed test particles are measured using both the short and long-range force fields. For this test, 128 random centers are chosen and for each center, 8192 test particles are randomly distributed with equal number per radial bin. Plotted in Figure 7.2 is the fractional error

$$\delta f \equiv \frac{|\boldsymbol{f}| - f(r)}{f(r)}, \tag{7.13}$$

where $f(r) = G/r^2$ is the magnitude of the exact force at separation r. The top and middle panels show the fractional errors for the uncorrected and corrected two-level mesh gravity solver while the bottom panel shows the errors for a standard one-level mesh gravity solver based on the potential method. The same random subset of pairs, with equal numbers per logarithmic bin, are plotted for all three cases for proper comparison.

In all three cases, the fractional errors can approach unity near the grid scale and the source of the problem is the CIC mass assignment scheme. The cubical cloud shape softens the mass resolution and introduces force anisotropies. In principle, higher-order mass assignment schemes (Hockney & Eastwood, 1988) like triangular-shaped clouds (TSC) can reduce the force fluctuations near the grid scale. However, Efstathiou et al. (1985) have conducted a detailed study of the

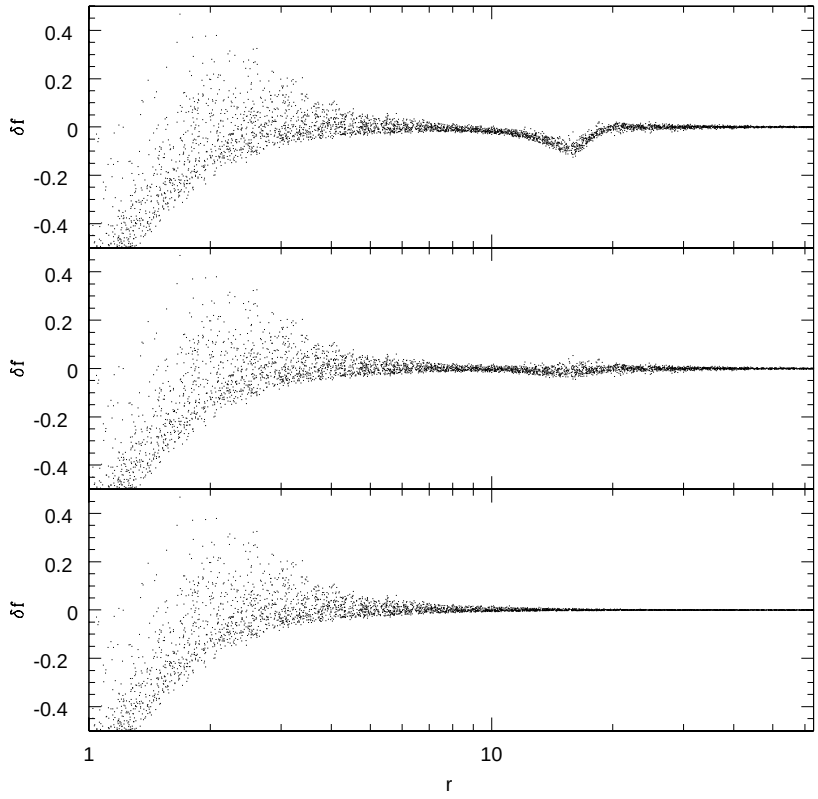

Figure 7.2 Fractional errors for the pair-wise force test. The top and middle panels show the errors for the uncorrected and corrected two-level mesh gravity solver while the bottom panel shows the errors for a standard single-mesh Poisson solver.

PM N-body method and have concluded that the CIC scheme offers the best balance between computational cost and force accuracy.

Near the initial cutoff $r_g = 16$ grid cells, the magnitude of the force is under predicted by the uncorrected gravity solver. The mass smoothing and kernel sampling on the coarse grid are responsible for the decrease in power. After correcting for these effects, the mean error for separations in the range $r = r_g \pm 4$ is approximately zero, the maximum error is $\sim 5\%$, and the rms error is less than 2%.

7.3 Out-of-core Computation

At the Canadian Institute for Theoretical Astrophysics (CITA) we have a HP/Compaq GS320 Alpha server with 32 processors and 64 GB of shared memory. The total floating-point speed has a peak value of 48 Gflops and the total memory bandwidth is 14 GB/s. The server has a total of 8 independent PCI busses, each running at 266 MB/s. In the current configuration, an array of 84

SCSI disks providing 3 TB of storage is connected to the server through 6 high bandwidth dual Ultra3 SCSI controllers, each with a peak bandwidth of 320 MB/s. To improve I/O performance, the disk array is striped using the HP/Compaq Logical Storage Manager software package. Read and write speeds of up to 500 MB/s are achieved when direct I/O is used.

Out-of-core computation requires dividing the global data domain into smaller local blocks that can fit entirely in memory. Each block is worked on sequentially: a data block is read from disks into memory, processed in parallel for a number of time steps, and then written back to the disk array. The following sections describe the domain decomposition, time stepping, and optimizations for doing out-of-core computation.

7.3.1 Domain Decomposition

The 3-dimensional global physical domain is decomposed into cubical regions and each local region is extended by buffer regions from its 26 neighbours. Associated with each extended local region is a high resolution hydro array and a dark matter particle list. The collisionless dark matter particles only interact gravitationally and for each gravity step, a buffer length of $b_g = 20$ grid cells is required to calculate the short-range force without introducing artificial boundary effects. The moving frame hydro algorithm requires 6 buffer cells per single time step and a hydro buffer $b_h = 12$ cells for a double time step. A total buffer size of $b_t = 32$ grid cells is required for every double time step. Only one gravity step is done per double hydro step.

The periodic global domain is represented solely by a total density field, constructed on a lower-resolution mesh. The global density field is used to calculate the long-range force field and it can also be used to provide another level of buffering for solving the short-range force field. To do multiple time steps while the data is in memory, one can increase the size of the buffer, but this is inefficient. A data block contains both the local physical region plus buffers and the maximum allowable size is limited by the amount of memory in the system. If the buffer size is increased, then the size of the local region must decrease, resulting in extra work overhead when trying to solve a fixed-size problem. An additional short-range gravity step can be performed by using the global density field to construct an additional buffer for the high-resolution local density field.

7.3.2 Time Stepping

A schematic Fortran code is presented in Figure 7.3 to illustrate the time stepping in the code. At the beginning of each stepping iteration, the global density field is read from the disks and the long-range force field is computed. The global data is stored on a lower-resolution mesh and fits entirely in memory. The global long-range force field is divided into local blocks and written back to the disks. This set of tasks is performed by the subroutine long_range_force. The global physical domain is decomposed into $N_a \times N_b \times N_c$ number of local regions and each region is worked on

```
call long_range_force

do c=1,Nc
  do b=1,Nb
    do a=1,Na
      call read_data
      call long_range_acceleration
      call hydro_forward
      call pm_gravity
      call hydro_reverse
      call hydro_forward
      call pm_gravity
      call hydro_reverse
      call write_data
    enddo
  enddo
enddo
```

Figure 7.3 A schematic Fortran code illustrating the time stepping in the OCH cosmological code. We loop over the decomposed domain and work on each local block sequentially: a data block is read from disks into memory, processed in parallel for a number of time steps, and then written back to the disk array.

sequentially. The subroutine read_data creates an extended local block by reading local and buffer data from files on disks. The velocities of the gas and dark matter particles are updated using the long-range force field. This is done with the subroutine long_range_acceleration.

A forward hydro sweep, gravity step, and reverse hydro sweep is then performed with the subroutines hydro_forward, pm_gravity, and hydro_reverse, respectively. In the gravity step, we construct a high resolution total density field from the synchronized gas and dark matter. The extended local block already contains a buffer of length b_t. In addition, the density field has another level of buffering drawn from the global density field. The gas velocities and the dark matter particle velocities and positions are advanced by two timesteps in the gravity step. In our multi time step scheme, the short-range tasks are repeated again.

The data for the next stepping iteration is prepared in the subroutine write_data. A new local block and new buffer blocks are written to individual files. In addition, the gas and dark matter are used to construct a portion of the lower resolution global density field to be used by the subroutine long_range_force in the next stepping iteration.

In total for each stepping iteration, we advance the gas and dark matter by 4 time steps while the data is in memory. To be second-order accurate, the value of the time step is fixed during a

double step and also between the last double step and the first double step of the next stepping iteration. The first criterion is imposed by the hydro algorithm and the short-range acceleration while the the latter criterion is required by the long-range acceleration. The value of the time steps between consecutive double steps in each stepping iteration can be different.

7.3.3 Optimizations

The OCH code is optimized in several ways. In order to be feasible, out-of-core computation must avoid being I/O limited. I/O overhead can be significantly reduced by performing disk operations concurrently with numerical calculations. In the OCH code, the first local block of data is read in and while computation is being done on it, the second local block is simultaneously read in. While doing computation on the second local block, the first block is written out and the third block is read in. One thread in the multi-processor machine is assigned to doing I/O. With this scheme, disk operations are effectively hidden and I/O overhead is reduced to less than 10%.

The amount of disk operations is also significantly reduced with the two-level mesh gravity solver and the multi-stepping scheme. The two-level mesh gravity solver allows FFTs to operate on data stored entirely in memory, which is much faster than the alternative of computing the transforms out-of-core through numerous non-sequential disk accesses. The multi-stepping scheme allows us to advance the gas and dark matter by 4 time steps before writing the data back to disks.

7.4 Cosmological Initial Conditions

Cosmological initial conditions for large OCH simulations are also generated out-of-core using a two-level mesh scheme (Pen, 1997). In the standard method, the overdensity field $\delta(\boldsymbol{x})$ is constructed from the convolution,

$$\delta(\boldsymbol{x}) = \int n(\boldsymbol{x}')w(\boldsymbol{x}-\boldsymbol{x}')d^3x', \tag{7.14}$$

of a random white noise field $n(\boldsymbol{x})$ with a density kernel $w(\boldsymbol{x})$, whose Fourier transform is the square root of the initial matter power spectrum $P_i(k)$. In the two-level mesh scheme, the density kernel is decomposed into short-range (Eq. [7.4]) and long-range (Eq. [7.5]) components. The density cutoff is denoted r_d to differentiate it from the gravity cutoff r_g. In this case, the function $\alpha(r)$ is chosen to be

$$\alpha(r) = (1 + ar^2 + br^4 + cr^6)w(r), \tag{7.15}$$

with coefficients,

$$a = -\frac{3}{r_d^2},$$
$$b = \frac{3}{r_d^4}, \tag{7.16}$$
$$c = -\frac{1}{r_d^6},$$

that satisfy the conditions given by equation (7.8). This parametrization is more generic in that the coefficients are independent of the global density kernel $w(r)$ and its derivatives.

After the matter density field has been determined, the gravitational force field is calculated and used to construct the comoving displacement field

$$\boldsymbol{dx}(\boldsymbol{x}) = -\frac{1}{4\pi G\bar{\rho}}\boldsymbol{\nabla}\Phi, \tag{7.17}$$

and the proper peculiar velocity field

$$\boldsymbol{v}_p(\boldsymbol{x}) \equiv a\frac{d\boldsymbol{x}}{dt} = a\frac{\dot{D}}{D}\boldsymbol{dx}, \tag{7.18}$$

where $\bar{\rho}$ is the comoving mean density, $D(t)$ is the linear growth factor, and $\dot{D}$ is its time derivative. The dark matter particles are displaced from a uniform distribution and their velocities are determined by interpolating from the grid. The gas distribution is taken to trace the matter distribution and its initial conditions are readily obtained from the grid-constructed fields.

The two-level cosmological initial conditions generator is tested with the following exercise. Consider a periodic simulation box of comoving length $L = 100h^{-1}$ Mpc that is discretized by a 512^3 grid. The global domain is divided into $2 \times 2 \times 2$ number of cubical local regions. Each local block is extended by buffers and has a length of $256 + 2b_t$ grid cells. The density cutoff and gravity cutoffs are chosen to be $r_d = 30$ and $r_g = 28$ grid cells, respectively, and a total buffer size $b_t = 32$ grid cells is needed after correcting the short-range kernels using equation(7.10). The gas is discretized on the 512^3 grid and the dark matter is represented by 256^3 particles.

In Fourier space, the isotropic density function $w(k)$ is taken to be the square root of the initial matter power spectrum $P_i(k)$. The matter transfer function is computed using CMBFAST (Seljak & Zaldarriaga, 1996) and with WMAP parameters: $\Omega_m = 0.27$, $\Omega_\Lambda = 0.73$, $\Omega_b = 0.044$, $n = 1$, $\sigma_8 = 0.84$, and $h_0 = 0.7$. The initial conditions are generated for an initial redshift $z = 50$ where the matter power spectrum is still linear and the real space density field has $|\delta_{max}| < 1$. The density function $w(k)$ is integrated to obtain the real space density function $w(r)$, which is then decomposed into short and long range terms for the two-level generator. The decomposition is shown in Figure 7.4. For $r_d = 30$ grid cells, the corrected and uncorrected short-range terms are very similar. The correction is only noticeable on scales very near the cutoff.

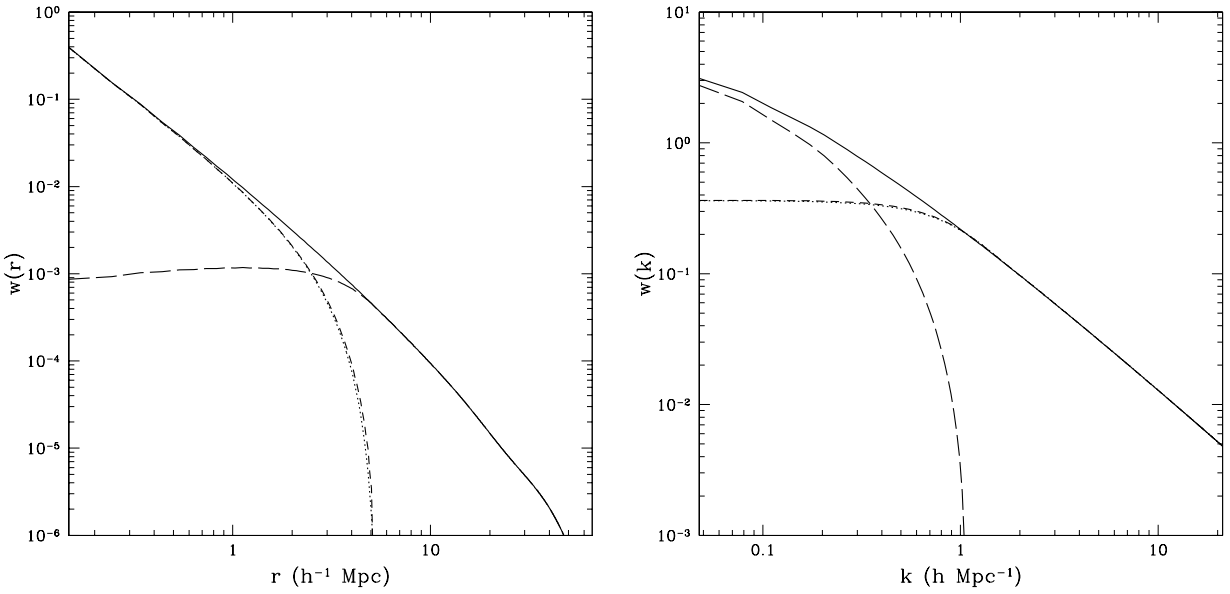

Figure 7.4 The decomposition of the density function (solid lines) for the two-level mesh initial conditions generator. The uncorrected short-range (dotted), corrected short-range (short dash), and long-range (long dash) components are plotted. The real space functions with a short-range cutoff $r_g = 30$ fine grid cells are shown in (a) and the Fourier space versions are in (b).

Plotted in Figure 7.5(a) is the power spectrum from a sample realization of the initial matter power spectrum. The grid-constructed matter density field has a power spectrum that statistically matches the template CMBFAST spectrum. At large wavenumbers or small scales, the agreement is very good due to the high number of sampling modes per bin. The deviations at small wavenumbers or large scales are expected due to sample variance. Figure 7.5(b) compares the realization matter power spectrum and the realization dark matter particles power spectrum. The bias (Eq. [6.37]) and cross-correlation (Eq. [6.38]) between the particles and the grid are plotted out to scales corresponding to the mean interparticle spacing. The dark matter particles are highly correlated with the matter density field. The turnover at small scales arises mainly because the CIC mass assignment scheme reduces power in the density field.

The two-level initial conditions generator can be directly compared to a standard one-level version by running the latter on the same noise field used in the former. The statistical correlation can be quantified with the bias

$$b(k) \equiv \sqrt{\frac{P_{22}(k)}{P_{11}(k)}}, \tag{7.19}$$

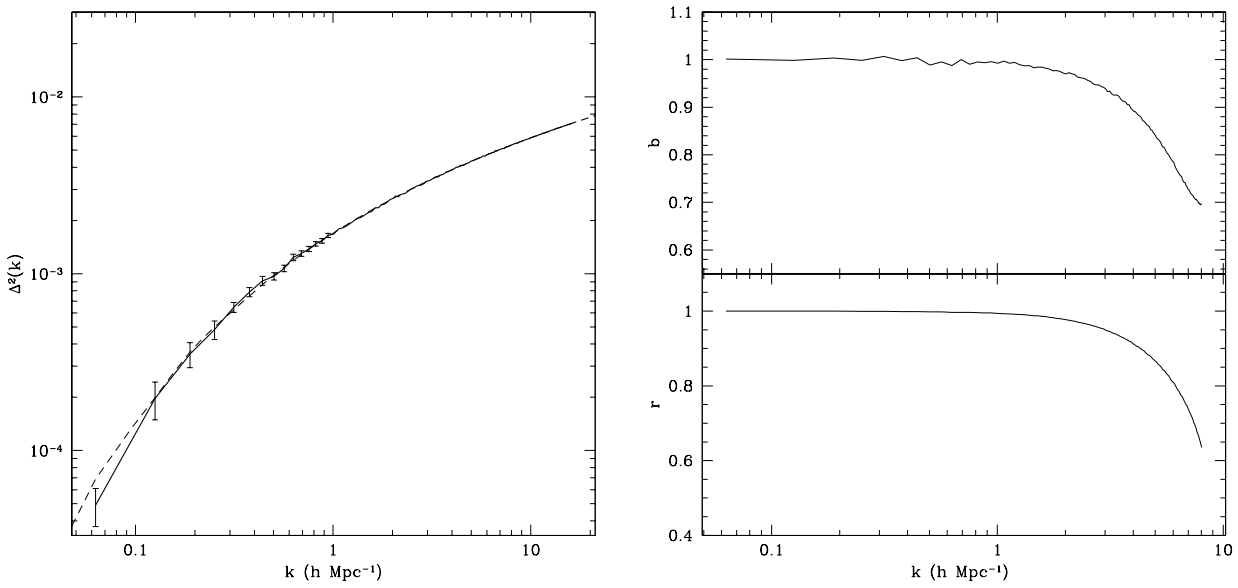

(a) **(b)**

Figure 7.5 Sample realization from the two-level mesh cosmological initial conditions generator. In (a), the realization power spectrum (solid line) of the grid-constructed matter density field is plotted against the template CMBFAST power spectrum (dashed line), out to the grid Nyquist frequency. In (b), the bias and cross-correlation between the dark matter particles and the matter density field are plotted out to the mean interparticle spacing.

and cross-correlation

$$r(k) \equiv \frac{P_{12}(k)}{\sqrt{P_{11}(k)P_{22}(k)}}, \tag{7.20}$$

where $P_{mn}(k) \equiv \langle \delta_m(k)\delta_n(k) \rangle$ and the subscripts denote the generators used. There is little statistical difference when the density kernel decomposition uses a short-range cutoff of $r_d = 30$ grid cells. For the grid-constructed matter density fields, both functions agree with unity at all wavenumbers. Some small scatter is found at scales approaching the box size or grid spacing, but it is less than 1%. For the dark matter particles the scatter is slightly higher but remains less than 3% for long and intermediate wavelengths. Near the interparticle spacing, the Poisson noise prevents a proper comparison.

7.5 Cosmological Simulations

The OCH cosmological code is applied to evolving the sample initial conditions generated in the previous section to demonstrate that it accurately simulates nonlinear structure formation in the universe. The configuration where the global domain is decomposed into 2 number of cubical

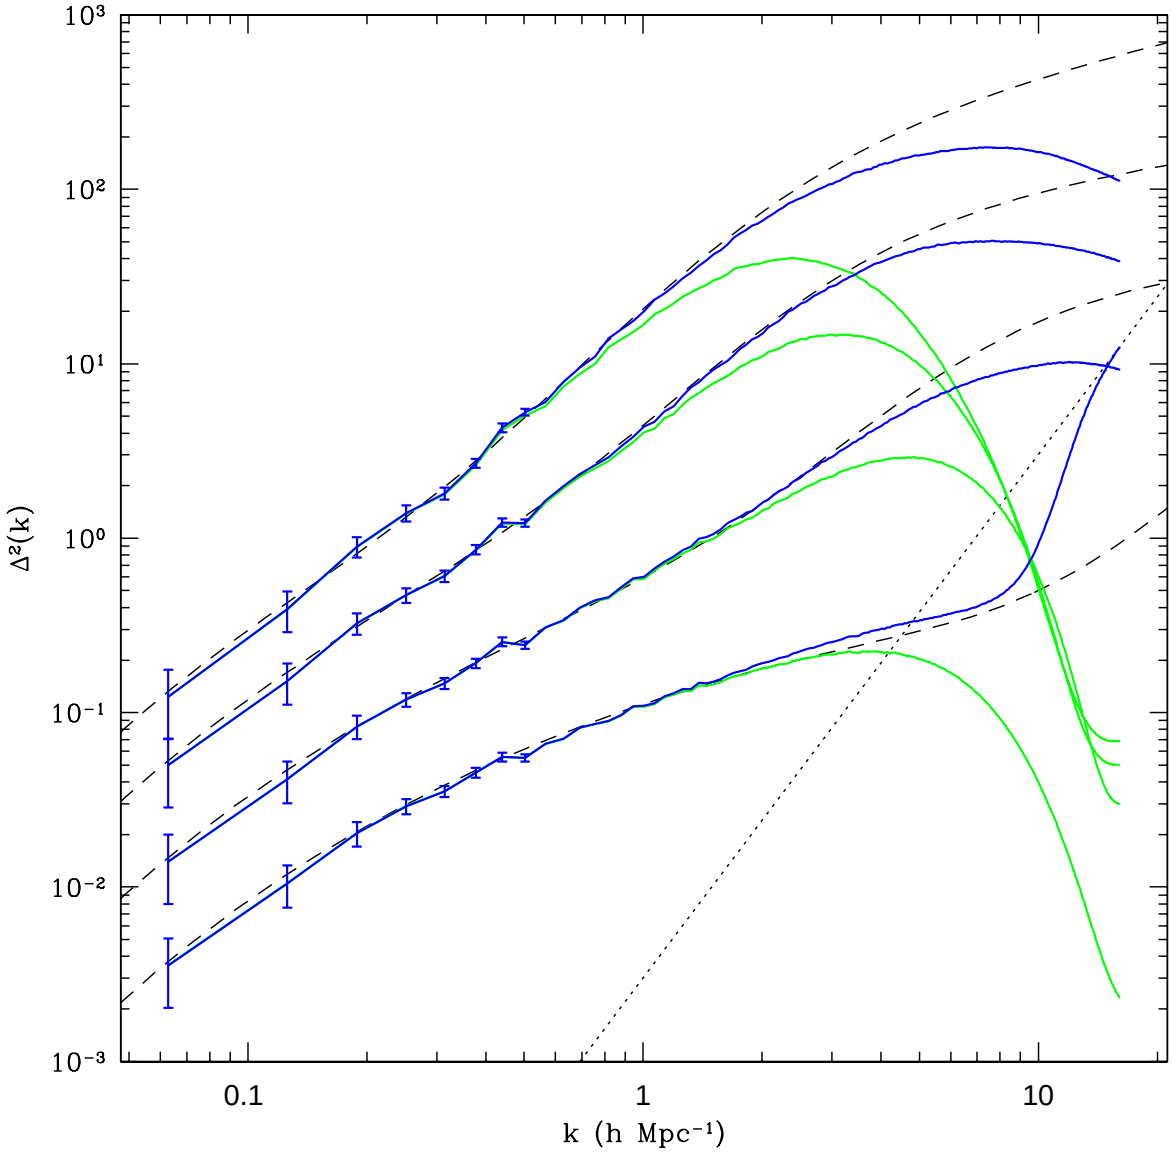

Figure 7.6 Nonlinear evolution of the dark matter (blue solid lines) and gas (green solid lines) power spectra at redshifts $z = 0$, 1, 3, and 7 from an out-of-core simulation with 512^3 grid cells and 256^3 dark matter particles in a 100 h^{-1}Mpc box. The simulated dark matter power spectra agree with the matter power spectra (dashed lines) predicted using the fitting functions of Smith et al. (2003). Poisson noise have been subtracted from the dark matter power spectra, except at $z = 7$ where the power does not exceed the shot noise power (dotted line).

blocks used in the generation of the initial conditions is retained for the out-of-core simulation. The dynamical evolution is checked by measuring the gas and dark matter mass power spectra at redshifts $z = 0$, 1, 3, and 7. Power spectra for the periodic density fields are computed using FFTs. For each redshift, the 8 local blocks are combined to construct the complete gas density field on a 512^3 grid. The clustered distribution of 256^3 dark matter particles is mapped onto a higher resolution 1024^3 grid to improve the accuracy of the measurements.

In Figure 7.6 the simulated power spectra are plotted against the nonlinear fitting functions of Smith et al. (2003). The simulated dark matter and predicted matter spectra are in good agreement, except at scales near the grid spacing. At redshift $z = 0$, the comparison is good on frequencies $k \lesssim 3h$ Mpc^{-1} or comoving wavelengths $\lambda \gtrsim 2h^{-1}$ Mpc. More than half the power is lost on scales less than 5 grid cells because the force resolution in the PM scheme is softened near the grid scale. The out-of-core results are very similar to that shown in Figure 6.6, obtained with the cosmological hydro code described in Chapter 6. The differences on large scales are due to sample variance in the initial conditions.

At large scales approaching the box size, the evolution is consistent with linear growth. The observed deviations are expected since the realization initial dark matter power spectrum has a deficit in power at large scales to begin with (See Figure 7.5). On linear and moderately nonlinear scales, the simulated gas and dark matter are highly correlated with no bias. On nonlinear scales, the gas loses power primarily because the pressure prevents gravitational collapse.

7.6 Summary

An out-of-core hydrodynamic code and an out-of-core initial conditions generator have been developed for high resolution cosmological simulations that require terabytes of memory. The OCH code utilizes a two-level mesh gravity solver and a multi-stepping scheme that significantly reduce the amount of disk operations. In addition, I/O overhead has been down to less than 10% by performing disk operations concurrently with numerical calculations. The code is cost-effective and memory-efficient. It has been demonstrated to accurately simulate the nonlinear structure formation in the universe and will provide high mass resolution for cosmological applications. The research presented in this chapter is collected in Trac & Pen (2004b).

Chapter 8

The Thermal Evolution of the High Redshift IGM

8.1 Introduction

The Lyman alpha (Lyα) forest traces the distribution of neutral hydrogen over a wide redshift range and can be used as a powerful probe of the structure and evolution of the IGM (see Rauch, 1998, for a review). It has opened a window to the high redshift universe, allowing one to study large-scale structure and the cosmological model of structure formation. For instance, the value of the mean transmitted flux in observed Lyα forest absorption spectra has been used to estimate the baryon density $\Omega_b h^2$ (e.g. Rauch et al., 1997). In addition, the measurement of the flux power spectrum can be used to determine the matter power spectrum, from which cosmological parameters can be extracted (e.g. Croft et al., 2002; Gnedin & Hamilton, 2002; McDonald et al., 2004a,b).

Measurements of the Lyα forest flux in quasar spectra rely critically on the temperature-density relation of the high redshift IGM. In the literature, often is made the claim that the temperature-density relation is well approximated by a power-law for the gas density range $0 < \Delta_b < 10$ relevant to studies of the Lyα forest. References are normally made to numerical work by Katz, Weinberg, & Hernquist (1996) and to analytical work by Hui & Gnedin (1997). It should be noted that the former is a paper describing the addition of a radiative transfer algorithm and star formation prescriptions to the TREESPH code Hernquist & Katz (1989). The temperature-density relation presented was meant to be an illustrative result of the code. In fact, the plotted relation shows considerable scatter and a single power-law cannot be accurately fitted to the data. Furthermore, in the analytical calculations, the power-law form is only valid for $\Delta_b < 5$ where shock heating is assumed to be weak and completely ignored. In general, the temperature-density relation in reality and in simulations can deviate significantly from a power-law (e.g. McDonald et al., 2001).

In this chapter, I present results from the largest Eulerian hydrodynamic simulation run to date for studying the thermal evolution of the high redshift $3 \leq z \leq 7$ IGM. I consider both adiabatic

and nonadiabatic evolution. The adiabatic evolution is simulated to quantify the effects of shock heating on raising the temperature of the IGM, even at low densities where it is numerically difficult to probe and therefore, normally ignored. In the nonadiabatic scenario, reionization is prescribed and radiative cooling and heating are implemented. As a working model, I consider a uniform and instantaneous reionization and an IGM in photoionization equilibrium with a uniform UV background computed by Haardt & Madau (1996). These standard conditions allow comparison with previous numerical simulations of the Lyα forest. The focus of this high resolution study is to robustly quantify the gas temperature-density relation and the gas-dark matter density relation. I discuss how the form of these relations, the scatter in the relations, and the shape of the scatter distributions reflect the underlying physics of the high redshift IGM.

8.2 Out-of-core Hydrodynamic Simulations

The out-of-core hydrodynamic (OCH) code described in Chapter 7 is applied to simulating the high redshift IGM in a WMAP cosmology: $\Omega_m = 0.27$, $\Omega_\Lambda = 0.73$, $\Omega_b = 0.044$, $n = 1$, $\sigma_8 = 0.84$, $h = 0.7$, and $Y = 0.24$. A high resolution simulation with 1536^3 grid cells and 768^3 dark matter particles in a 50 h^{-1}Mpc box is run from an initial redshift $z = 80$ down to $z = 3$. The comoving grid spacing is $\Delta x = 33$ h^{-1}kpc and the dark matter particle mass resolution is $\Delta m = 2.1 \times 10^7$ $h^{-1} M_\odot$. The initial conditions are prepared using the out-of-core initial conditions generator with an initial transfer function computed by CMBFAST.

The high resolution OCH simulation is effectively run on a 1536^3 grid, but with buffering the actual cost amounts to running on a 1664^3 grid. This corresponds to a work overhead of 30%. The OCH code uses approximately 45 GB of RAM and 210 GB of disk space. In comparison, the OpenMP MACH code described in Chapter 6 would require approximately 160 GB of shared memory. From redshift $z = 80$ down to $z = 3$, the simulation takes approximately 90 out-of-core sweeps, where each sweep is composed of 4 time steps and takes approximately 2 physical hours to complete on the HP/Compaq 32-processor GS320 Alpha server. The simulation finishes in 7.5 days and uses close to 6000 cpu hours.

8.2.1 Radiative Cooling and Heating

The non-adiabatic simulation of the thermal history of the high redshift IGM includes reionization and radiative heating and cooling. At redshift $z = 7$, the IGM is instantaneously and uniformly injected with thermal energy equivalent to $T_{\rm reion} = 25000$ K. A uniform photoionization background taken from Haardt & Madau (1996) is turned on thereafter. The non-adiabatic simulation is run from the $z = 7$ checkpoint in the adiabatic simulation. It takes approximately 70 out-of-core sweeps, 6 days to complete, and uses approximately 4500 cpu hours to run down to redshift $z = 3$.

The radiative cooling and heating algorithm implemented for the OCH code is based on a

version provided by Uros Seljak. Atomic processes within a photoionized medium of hydrogen and helium are simulated by tracking six species (H I, H II, He I, He II, He III, e^-) and computing the relevant cooling, heating, and ionization rates (see Cen, 1992; Theuns et al., 1998). The cooling rate Γ is calculated from cooling processes, which include collisional excitation, collisional ionization, recombination, Bremsstrahlung cooling, and inverse Compton cooling. The heating rate Λ is based on the Haardt-Madau photoheating rate. Assuming photoionization equilibrium, the equations describing the ionization evolution of the six species are solved using an iterative Jacobian method until the fractional change in all species abundance is below 0.001% (see Theuns et al., 1998).

8.3 The Gas-Dark Matter Density Relation

The bias (Eq. [6.48]) and stochasticity (Eq. [6.49]) of the gas relative to the dark matter is quite well understood. On large linear scales the evolution of the gas is driven primarily by the gravity of the dark matter and they are perfectly correlated with no bias. The gas continues to be an unbiased and highly correlated tracer even on moderately nonlinear scales where gravitational collapse starts to become important. On small nonlinear scales the gas pressure can be quite large due to gravitational collapse and shock heating and it in turn counteracts the gravitational infall. With decreasing scale, both the bias and stochasticity parameters deviate more below unity.

The gas-dark matter relation can be simply parametrized by the bias parameters, but the scatter in the relation has not been robustly quantified. Cosmological simulations based on Eulerian hydro algorithms coupled with Lagrangian N-body algorithms can have numerical scatter that are systematic due to the opposite nature of the schemes. The Eulerian approach has a large dynamic range in mass, but not in length, and has approximately equal resolution at all scales. On the other hand, the Lagrangian approach has a large dynamic range in length, but not in mass, and achieves good spatial resolution in high density regions, though poorly in low-density regions.

The gas-dark matter relation and its scatter is quantified using the following procedure. The dark matter density field,

$$\Delta_m \equiv (1 + \delta_m) = \frac{\rho_m}{\bar{\rho}_m}, \tag{8.1}$$

is constructed by mapping the particles onto the grid using the CIC mass assignment scheme. The baryonic or gas density field,

$$\Delta_b \equiv (1 + \delta_b) = \frac{\rho_b}{\bar{\rho}_b}, \tag{8.2}$$

is divided into logarithmic bins of width 0.1 dex and the mean and scatter in the dark matter density is found for each bin. The gas density field is binned rather than the dark matter density field because the former is well defined over a larger range than the latter, particularly at low densities. The mean value can be calculated by taking some weighted average, but when a relatively large

scatter is present, the calculated average can be quite biased. A more robust measure to use in the gas-density relation is the median. For each gas density bin, the cells that fall in the range are ranked order in terms of their value of the dark matter density. Starting from the lowest value, the cumulative fraction is calculated and the 16%, 50%, and 84% levels are found. The median or 50% value, aside from the standard definition, can also be thought of as some weighted mean where the weight decreases with increasing deviation from the 50% mark. The 16% and 84% levels bound the region in which 68% of the scatter is found. For a normal distribution, these levels correspond to the 1-σ spread.

In Figure 8.1 the gas-dark relation is plotted for the high redshift $3 \leq z \leq 6$ IGM. The slope of the relation is consistent with unity for $\Delta_m \lesssim 100$, demonstrating that the gas is unbiased relative to the dark matter everywhere except in high density regions like halos. For adiabatic evolution, the gravitational collapse of the gas is hindered by gas pressure and the slope decreases below unity at high densities $\Delta_m \gtrsim 100$. With radiative cooling the gas can undergo further collapse and for a fixed matter density, the gas density is higher with cooling than without.

Note that the gas-dark matter relation drops off rapidly at low densities just below the mean matter density. This numerical artifact arises because the Lagrangian approach of the N-body scheme poorly resolves the dark matter density field at low densities. For a particle to grid cell ratio η=1/8, the particle fraction in a cell with mean density is 1/8. From Figure 8.1 the dark matter density cutoff is given by $\Delta_m \approx 0.5$, which corresponds to a particle fraction of 1/16. In PM N-body simulations that use CIC for mass assignment, the minimum value of the dark matter density that is reliable can be estimated as

$$\Delta_{m,\min} \sim \frac{1}{16\eta}. \tag{8.3}$$

The "triangular shaped cloud" (TSC) mass assignment scheme is even smoother than CIC and the minimum reliable density is expected to be lower. For "nearest grid point" (NGP) schemes, the discontinuous mapping makes the minimum reliable limit greater than $1/\eta$.

8.3.1 Scatter in the Gas-Dark Matter Relation

In this discussion on scatter in the gas-dark matter density relation, let $\sigma_\pm$ represent the negative and positive errors in $\Delta_m \equiv (1 + \delta_m)$, defined as

$$\sigma_- \equiv \Delta_{m,16\%} - \Delta_{m,50\%}, \tag{8.4}$$

$$\sigma_+ \equiv \Delta_{m,84\%} - \Delta_{m,50\%}. \tag{8.5}$$

The negative and positive errors in $\log \Delta_m$ for each gas density bin are similarly defined as

$$\Sigma_- \equiv \log \Delta_{m,16\%} - \log \Delta_{m,50\%}, \tag{8.6}$$

$$\Sigma_+ \equiv \log \Delta_{m,84\%} - \log \Delta_{m,50\%}. \tag{8.7}$$

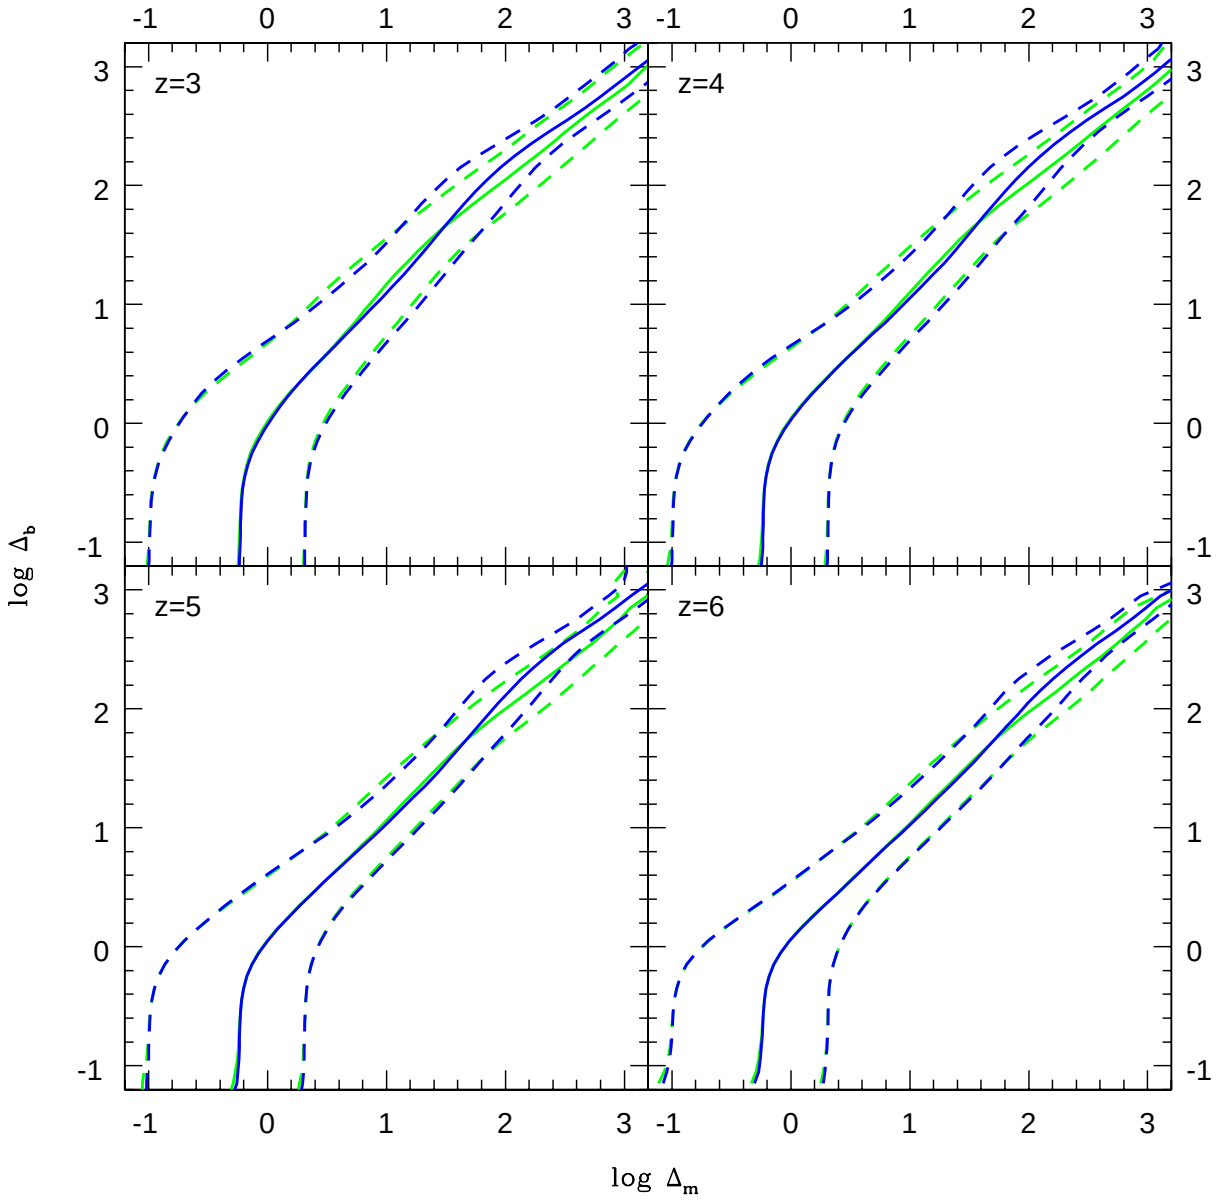

Figure 8.1 The gas-dark matter density relation in the high redshift IGM. The mean and 1-σ values are given by the solid and dashed lines, respectively. The slope of the relation is consistent with unity for $\delta \lesssim 100$. For adiabatic evolution (green) the gravitational collapse of the gas is hindered by gas pressure and the slope of the relation decreases below unity at high densities. With radiative cooling (blue) the gas can undergo further collapse.

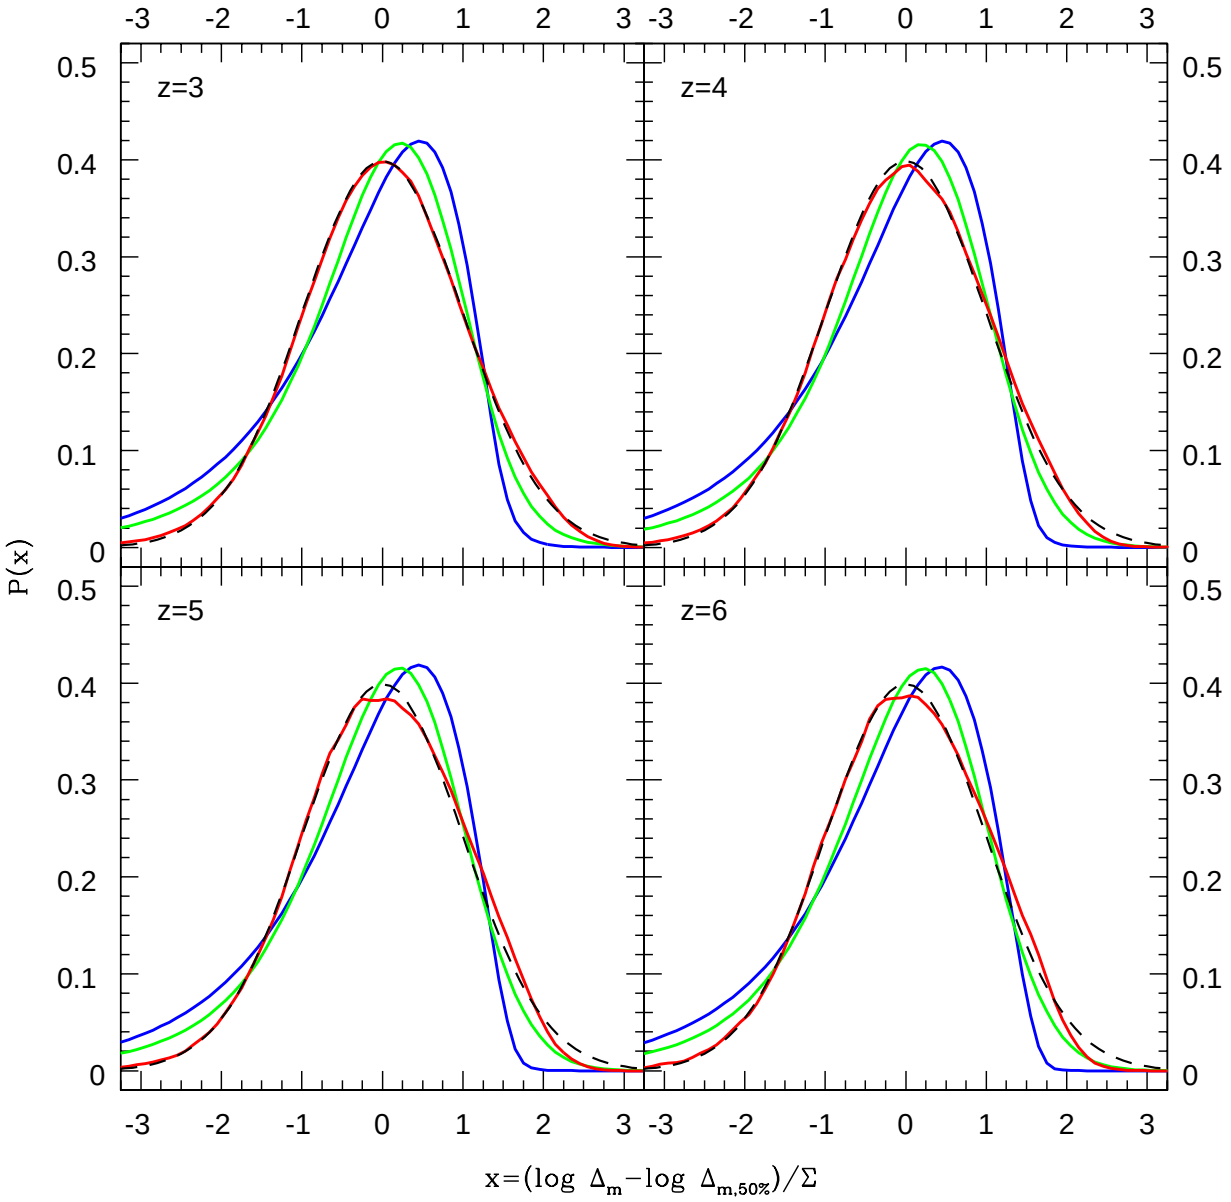

Figure 8.2 The distribution of scatter in $\log \Delta_m$ at fixed gas densities $\Delta_b = 1$ (blue), 10 (green), and 100 (red) for redshifts $z = 3$, 4, 5, and 6. At high densities, the distribution is Gaussian (dashed). With decreasing densities, the distribution becomes more skewed towards the left and this is a signature of the Lagrangian approach of the PM N-body scheme, which has poor mass resolution in low density regions.

In Figure 8.1 the negative and positive errors $\Sigma_{\pm}$ in $\log \Delta_m$ are approximately equal in magnitude. The distribution of scatter in Δ_m should then be skewed towards positive values. Consider the fractional errors in $(1 + \delta_m)$, given as

$$\varepsilon_{-} \equiv \frac{\Delta_{m,16\%}}{\Delta_{m,50\%}} - 1, \tag{8.8}$$

$$\varepsilon_{+} \equiv \frac{\Delta_{m,84\%}}{\Delta_{m,50\%}} - 1. \tag{8.9}$$

The negative fractional error ε_{-} has a value of -0.8 at low densities $\Delta_b \sim 1$ and it monotonically increases to -0.4 at high densities $\Delta_b \sim 1000$. The positive fractional error ε_{-} is always larger than the absolute value of the negative fractional error, with a value of 2.5 at low densities and monotonically decreasing to 0.8 at high densities. The distribution of scatter in Δ_m is highly skewed towards positive values.

Figure 8.2 shows the the distribution of scatter in $\log \Delta_m$ at fixed gas densities $\Delta_b = 1$, 10, and 100 for redshifts $z = 3$, 4, 5, and 6. The distributions have been rescaled and shifted so that the curves for different densities can be directly compared. The x-axis has been rescaled by Σ, defined as the symmetrical deviation for which the area under the probability distribution $p(x)$ in the range $-1 \leq x \leq 1$ is equal to 0.68. For the $\Delta_b = 100$ bin, the distribution of scatter in $\log \Delta_m$ is Gaussian and therefore, lognormal for Δ_m. For decreasing gas densities, the distribution becomes more skewed towards lower values. This skewness is also an artifact related to the poor particle mass resolution at low densities.

The intrinsic scatter in the gas-dark matter relation is particularly important to hydro-PM (HPM) simulations used for studying the Lyα forest (Gnedin & Hui, 1998; McDonald et al., 2001). In HPM simulations, a particle-mesh N-body code is used to simulate the dark matter density and velocity fields. The hydrodynamics of the gas is not actually simulated. The gas density and velocity fields are constructed by taking the corresponding dark matter fields and smoothing it to mimic the effects of pressure. To improve the realism of HPM simulations, the intrinsic scatter in the gas-dark matter density relation should be added. Measurements of the matter power spectrum from the Lyα forest often rely on HPM simulations for calibrations (e.g. McDonald et al., 2004b). Neglecting the intrinsic scatter may then result in underestimating the uncertainties in the measurement of cosmological parameters.

8.4 The Temperature-Density Relation

The temperature-density relation in the high redshift IGM is specifically used in calculating the mean transmitted flux and the flux power spectrum of the Lyα forest. Figure 8.3 shows the temperature-density relation in the high redshift IGM for both adiabatic gas evolution and with radiative cooling and heating. In the following sections, I discuss how the form of these relations,

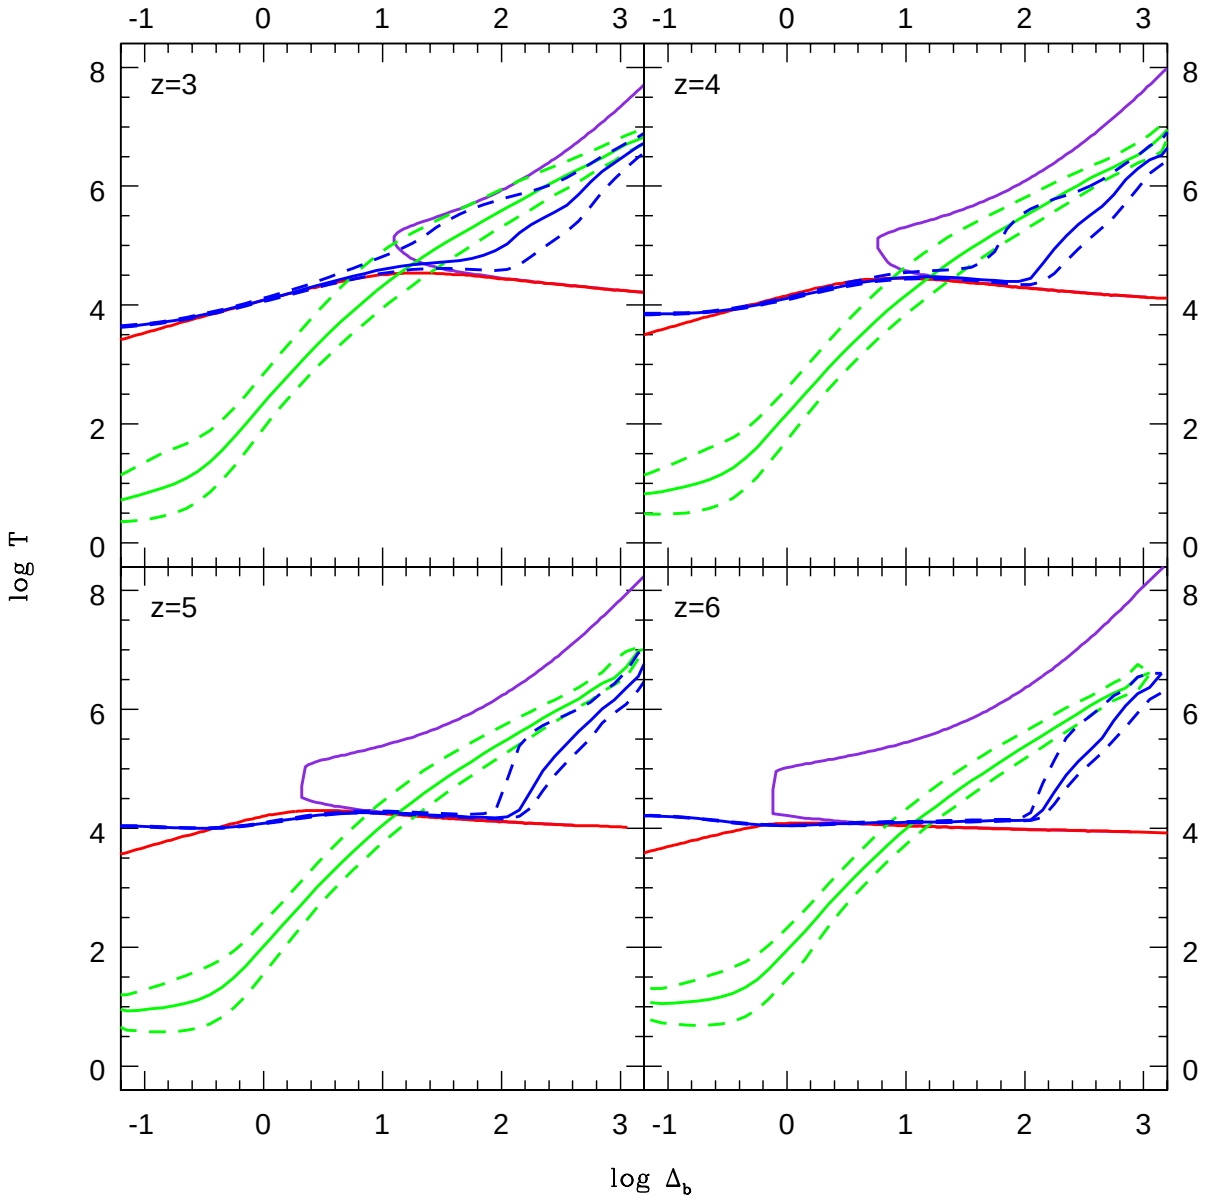

Figure 8.3 The temperature-density relation in the high redshift IGM for adiabatic gas evolution (green) and with reionization and radiative cooling and heating (blue). The mean and 1-σ values are given by the solid and dashed lines, respectively. Plotted for comparison are the heating track (red) where the heating time equals the Hubble time and the equilibrium track (violet) where the cooling rate balances the heating rate.

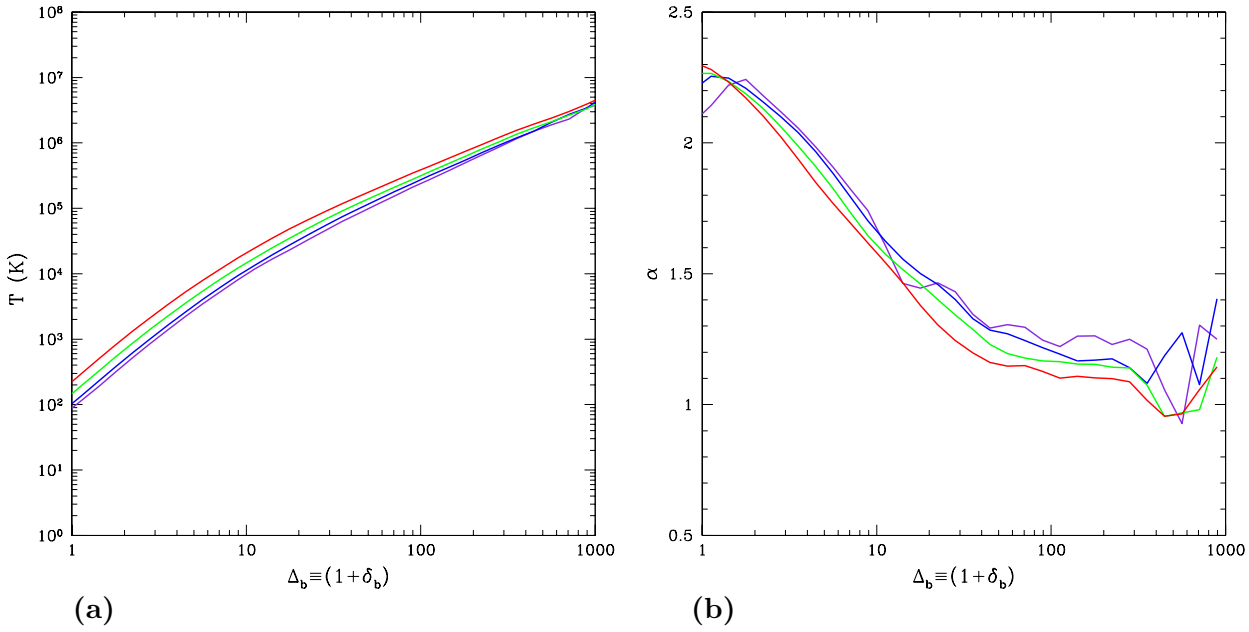

(a) **(b)**

Figure 8.4 The evolution of the adiabatic temperature-density relation (a) and the effective slope of the relation (b) for redshifts $z = 3$ (red), 4 (green), 5 (blue), and 6 (violet). Over this redshift range, the mean temperature is consistent with $T(z) \propto (1+z)^{\beta}$, where the slope $-2 < \beta < 0$ is a function of density. The effective slope α of the temperature-density relation approaches unity at high densities and low redshifts.

the scatter in the relations, and the shape of the scatter distributions reflect the underlying physics of the high redshift IGM.

8.4.1 Adiabatic Gas Evolution

For adiabatic evolution the thermal history of the IGM is primarily dictated by shock heating from gravitational collapse. The shock heating induces a temperature that monotonically increases with gas density. The adiabatic temperature-density relation in Figure 8.3 cannot be described by a single power-law however for it displays much curvature. At low densities $\Delta_b < 0.5$, the temperature floor $T_{\rm min} \sim 100$ K is a numerical limitation caused by spurious heating of cold gas. As discussed in Chapter 5, for cold gas flows the gravitational energy can be comparable to the kinetic energy and much larger than the thermal energy. At high mach numbers, numerical errors in calculating the gravitational effects can cause significantly heating on the order of $10^3 - 10^4$ K. The moving frame hydro algorithm is able to alleviate this problem by adding the gravitationally induced acceleration to the grid rather than to the fluid.

The evolution of the adiabatic temperature-density relation over the redshift range $3 \leq z \leq 6$

is displayed in Figure 8.4(a) and the effective slope of the relation,

$$\alpha(\Delta_b) \equiv \frac{\mathrm{d}\log T}{\mathrm{d}\log \Delta_\mathrm{b}}, \tag{8.10}$$

is plotted in Figure 8.4(b). In the literature, the effective slope is sometimes denoted by $\gamma - 1$ instead of α. The curvature in the relation is quite apparent in these two plots. Power-law fits are only consistent at high densities $\Delta_b > 100$ where the effective slope is approximately constant. Evolution is found primarily at lower densities and temperatures. At high densities there is relatively little change. Over this redshift range, the evolution of the temperature is consistent with the parametrization,

$$T(z) \propto (1+z)^\beta, \tag{8.11}$$

where both the temperature and the slope are functions of the gas density. For the gas density bins $\Delta_b = (1,\ 10,\ 100,\ 1000)$, the best-fit slopes are $\beta = (-1.70 \pm 0.14,\ -1.38 \pm 0.10, -0.86 \pm 0.04,\ -0.25 \pm 0.17)$ for $3 \leq z \leq 6$. It would highly interesting to determine how well the power-law parametrization holds at lower redshifts.

The effective slope α of the temperature-density relation rapidly changes at gas densities $\Delta_b < 10$. At higher densities $\Delta_b > 100$, it flattens out. For a fixed gas density, the effective slope is also seen to decrease with redshift. These observations are consistent with and reinforces the finding that $\alpha \approx 1$ and $T \propto \rho_b$ at low redshifts $z \sim 0$ (e.g. Zhang, Pen, & Trac, 2004b).

8.4.2 Radiative Cooling and Heating

A more realistic simulation of the thermal history of the high redshift IGM includes reionization and radiative cooling and heating. At redshift $z = 7$, the IGM is instantaneously and uniformly injected with thermal energy equivalent to $T_\mathrm{reion} = 25000$ K. A uniform photoionization background taken from Haardt & Madau (1996) is turned on thereafter. Atomic hydrogen and helium cooling is calculated assuming photoionization equilibrium.

The temperature-density relation obtained with radiative cooling and heating turned on will be referred as the Haardt-Madau (HM) temperature-density relation for it is specific to the choice of photoionization background. It is plotted against the adiabatic relation in Figures 8.3 and 8.5. In addition, two other relations, the cooling and heating tracks, are plotted for comparison. They are defined following Theuns et al. (1998). The cooling time is a function of both density and temperature and is defined as

$$t_\mathrm{cool} \equiv \frac{\epsilon}{\dot{\epsilon}} = \frac{(3k\rho_b T)/(2\mu m_H)}{\Gamma - \Lambda}, \tag{8.12}$$

where μ is the mean molecular weight, $X \equiv (1-Y)$ is the hydrogen mass fraction, and Γ and Λ are the cooling and heating rates, respectively. The mean molecular weight depends on the ionization fractions of hydrogen and helium, which in turn is both density and temperature dependent. For

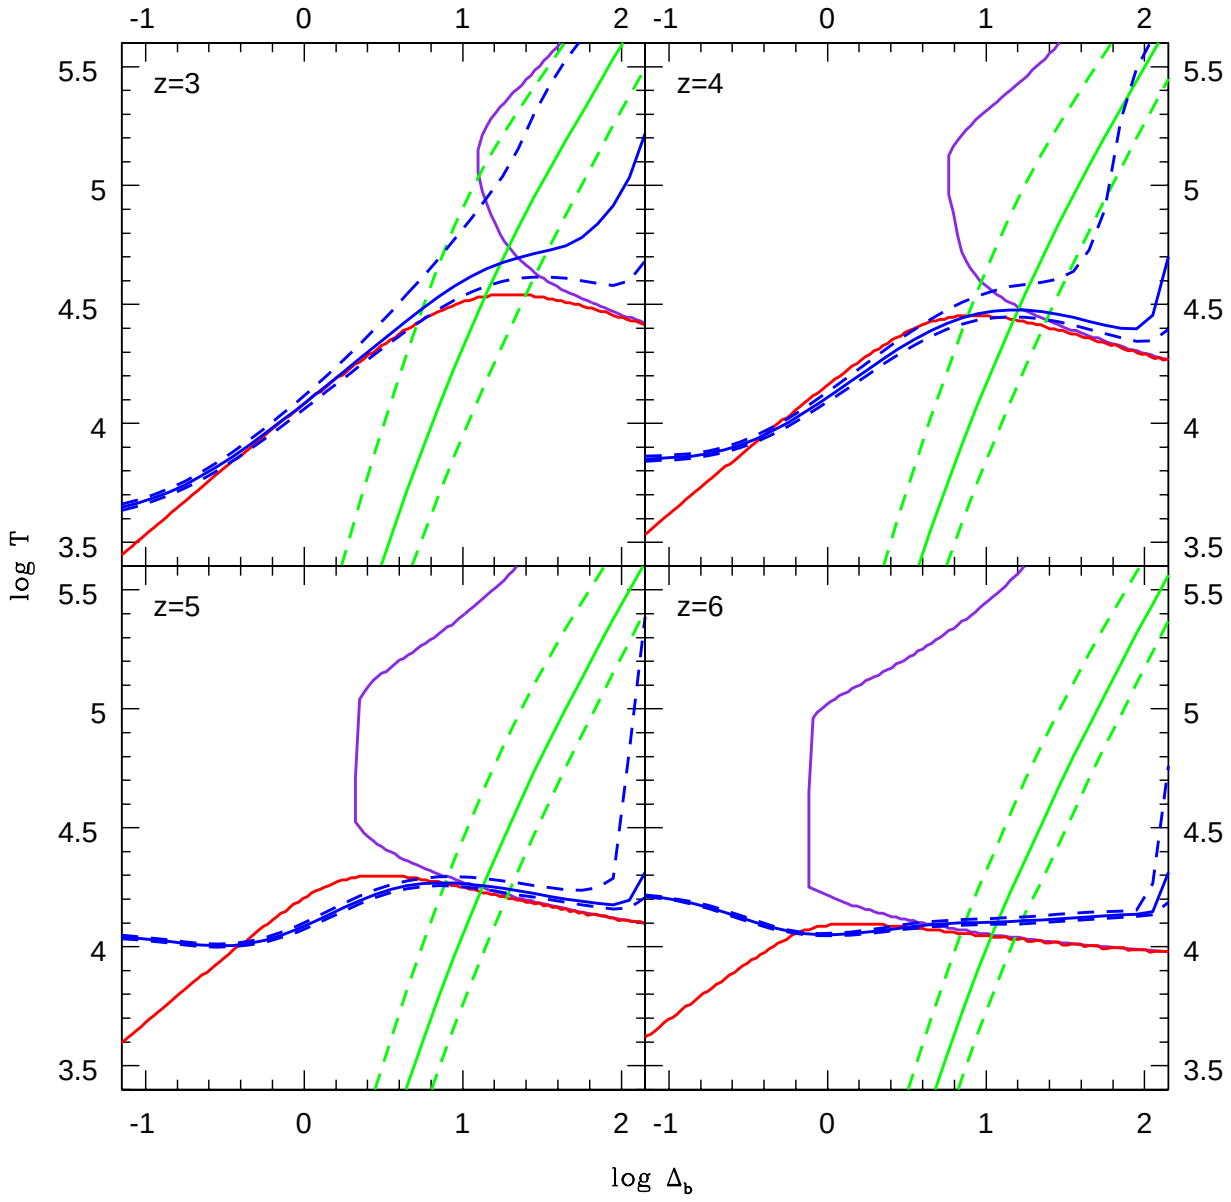

Figure 8.5 A zoomed view of Figure 8.3. The temperature-density relation in the high redshift IGM for adiabatic gas evolution (green) and with reionization and radiative cooling and heating (blue). The mean and 1-σ values are given by the solid and dashed lines, respectively. Plotted for comparison are the heating track (red) where the heating time equals the Hubble time and the equilibrium track (violet) where the cooling rate balances the heating rate.

$\Gamma > \Lambda$, cooling dominates and the cooling time is positive. For $\Lambda > \Gamma$, heating dominates and the heating time is taken to be $t_{\rm heat} = -t_{\rm cool}$. The cooling or equilibrium track is defined as the temperature-density relation for which the cooling rate equals the heating rate. The heating track is defined as the temperature-density relation for which the heating time is equal to the Hubble time,

$$t_{\rm Hubble} \frac{1}{\sqrt{6\pi G\bar{\rho}(z)}}. \tag{8.13}$$

At high redshifts, the universe is approximately Einstein-de Sitter, where equation (8.13) is valid.

The cooling and heating tracks divide the temperature-density phase space into three distinct regions. The region below the heating track is characterized by $t_{\rm heat} < t_{\rm Hubble}$. Gas that fall just below the heating track are efficiently photoheated back up to. Thus, the heating track approximately defines a minimum temperature curve. At higher densities, the cooling and heating tracks bound a region in the phase space characterized by $t_{\rm cool} < t_{\rm Hubble}$. Bremsstrahlung and line cooling dominates in this region, making the cooling time shorter than the Hubble time. Note that where the cooling and heating tracks intersect at high densities, the two tracks become one for all higher densities. Gas lying just above this track is efficiently cooled back on to it. In fact, gas lying in this region of phase space are quite rapidly cooled in general. Lastly, the remaining region of phase space, lying above the heating track and to the left of the cooling track, is characterized by $|t_{\rm cool}| > t_{\rm Hubble}$. Low density gas with high temperatures occupy this region where neither cooling nor heating is able to change the temperature by much.

Figure 8.5 is a zoomed view of Figure 8.3 for closer inspection of the HM temperature-density relation relative to the adiabatic temperature-density relation and to the cooling and heating tracks. At very low densities $\Delta_b \sim 0.1$, the HM temperature-density relation lies above the heating track and the gas temperatures at these densities are determined by a combination of adiabatic cooling and radiative cooling. For a reionization temperature of $T_{\rm reion} = 25000$ K at $z = 7$, adiabatic cooling alone should result in a temperature of $T = 12500$ K at $z = 3$. Radiative cooling has additionally cooled the gas down by ~ 8000 K.

Near mean densities $\Delta_b \sim 1$, the HM temperature-density relation approximately overlaps with the heating track over a narrow density range, particularly at lower redshifts. Between redshifts $z = 7$ and $z = 6$, the cooling time for the reionized gas is longer than the elapsed time and the gas is not able to cool onto the heating track. Between redshifts $z = 6$ and $z = 5$, the Haardt-Madau photoionization background changes relatively quickly and the gas cannot keep up. It is expected that by redshift $z \approx 2$, the temperature-density relation should overlap with the heating track for low densities $\Delta_b \lesssim 1$. In principle, the measurement of the temperature-density relation over that gas density range at redshift $z = 2$ could constrain the photoionization background. In practice, this is rather difficult since very low detectable limits are required to probe the very small neutral hydrogen column densities and optical depths.

In general, once gas has cooled onto the heating track at high densities, radiative cooling and heating should be able to maintain it at that level, almost regardless of shock heating. This is seen to be true between redshifts $z = 5$ and $z = 4$, at least to densities $\Delta_b \sim 100$, beyond which cooling is manually restricted to prevent overcooling in halos. However, at lower redshifts, the cooling track moves toward higher densities and shock heating can now play a role in heating the gas above the heating track. This is quite apparent at redshift $z = 3$ where shock heating is seen to raise the temperature of the gas at densities as low as $\Delta_b \sim 3$. The shock heating actually makes the HM temperature-density relation more power-law than the heating track. The fact that shock heating is important to the temperature-density relation for the gas density range $0 < \Delta_b < 10$ is often missed in the literature. Shock heating not only affects the mean temperature at a fixed density, but it also affects the shape of the scatter distribution. This subject will be discussed in more detail in the next section of this chapter.

In the literature, the temperature-density relation is often assumed to be well approximated by a power-law $T \propto \Delta_b^{\alpha}$ for the gas density range $0 < \Delta_b < 10$ relevant to studies of the Lyα forest. In Figure 8.6 the effective slope,

$$\alpha(\Delta_b) \equiv \frac{\mathrm{d}\log T}{\mathrm{d}\log \Delta_{\mathrm{b}}}, \tag{8.14}$$

of the HM temperature-density relation is plotted and compared with the effective slope of the heating track. At lower densities $\Delta_b \sim 1$, the differences between the two are primarily due to the inclusion of reionization in the HM temperature-density relation. At lower redshifts and lower densities, the slope of the heating track approaches an asymptotic value of $\alpha_{\mathrm{max}} = 0.58$ that is consistent with the predictions by Hui & Gnedin (1997).

The heating tracks can be approximately described by a double-power law. At higher redshifts $z \gtrsim 5$, it is the slope with the smaller value that gives the appearance of a power-law for the density range $0 < \Delta_b < 10$. At lower redshifts $z \lesssim 5$, the break in the double power-law occurs in the density range of interest. As this break moves towards higher densities at lower redshifts, the slope with the larger value then gives the appearance of a single power-law. However, the power-law approximation is more true at very low densities $\Delta_b \sim 0.1$ for curvature is already present at mean densities $\Delta_b \sim 1$. Over the density range $0 < \Delta_b < 10$, the effective slope has changed by $\Delta\alpha = -0.31$ at $z = 3$ and by $\Delta\alpha = -0.17$ at $z = 2$.

The effective slope of the HM temperature-density relation is more complicated, but in general still consistent with that from the heating track. At higher redshifts $z \gtrsim 6$, the slope is approximately constant in the density range of interest. This is indeed a result driven by the radiative cooling and heating. The HM temperature-density relation closely follows the heating track, at least on the high density side of the power-law break. At lower redshifts $z \lesssim 6$, the slope is approximately constant at mean densities $\Delta_b \sim 1$, but the range is quite small. It is somewhat surprising that at redshift $z = 3$ and for the gas density range $0 < \Delta_b < 10$, the HM temperature-density

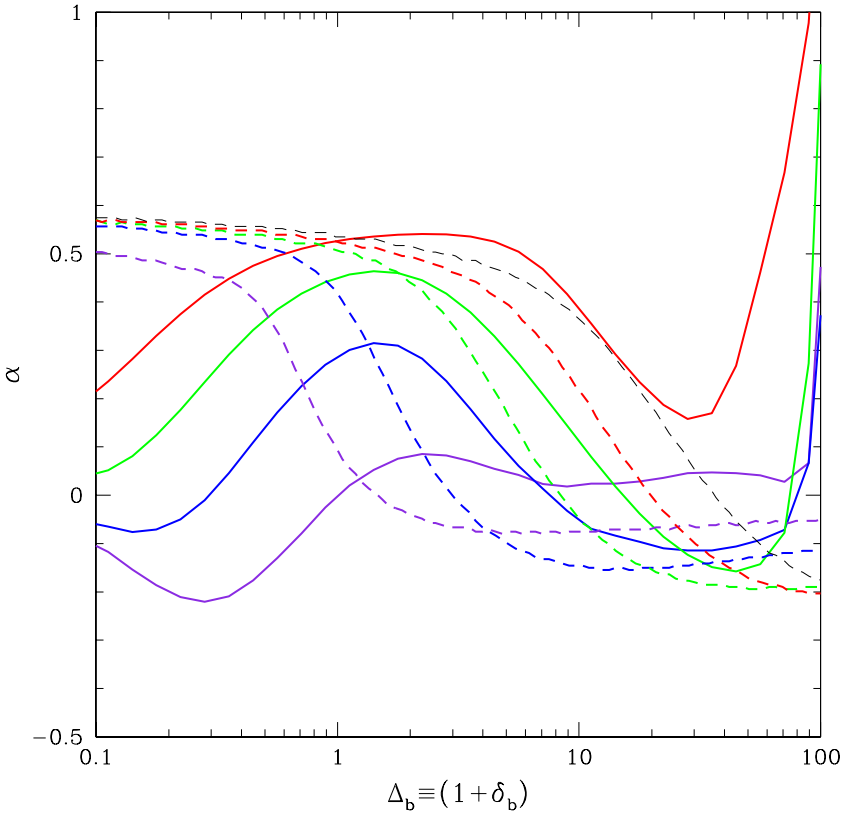

Figure 8.6 The effective slope of the HM temperature-density relation (solid) for redshifts $z = 3$ (red), 4 (green), 5 (blue), and 6 (violet). For comparison the effective slope of the heating track (dashed) is also plotted down to redshift $z = 2$ (black). The upturn in the curves is due to the manually reduced cooling at high densities $\Delta_b \sim 100$.

relation is better described by a power-law than the heating track and that the constancy of the effective slope is contributed to by shock heating at densities $\Delta_b \gtrsim 3$. It is expected that at even lower redshifts, the effective slope of the HM temperature-density relation will increase over the density range of interest as a result of shock heating.

Presently, cooling is effectively turned off at high densities $\Delta_b > 100$ to prevent overcooling in halos. This results in the upturn seen in the HM temperature-density relation and its effective slope. In future work, the reduced cooling will be imposed at higher densities to better see the relative effects of radiative cooling and shock heating for $\Delta_b < 100$.

In principle, the temperature-density relation can be modelled by first calculating the heating and cooling tracks for a given photoionization background and then modifying them with the adiabatic temperature-density relation to include the effects of shock heating. At high redshifts $z \sim 3$, the universe is approximately Einstein-de Sitter, and the measured shock heating should be approximately independent of cosmology. Thus, only a single high resolution simulation of the

adiabatic evolution, like this out-of-core simulation, can be used to quantify the shock heating. The heating and cooling tracks require negligible time to compute and many different photoionization background can be considered. Different reionization scenarios can be added using the method of Hui & Gnedin (1997). The predicted temperature-density relation can then be compared with the observed one from measurements of the Lyα forest to constrain reionization and the photoionization background.

8.4.3 Scatter in the Temperature-Density Relation

In this discussion on scatter in the adiabatic and HM temperature-density relations, let $\sigma_{\pm}$ represent the negative and positive errors in the temperature, defined as

$$\sigma_{-} \equiv T_{16\%} - T_{50\%}, \tag{8.15}$$

$$\sigma_{+} \equiv T_{84\%} - T_{50\%}. \tag{8.16}$$

The negative and positive errors in $\log T$ for each gas density bin are similarly defined as

$$\Sigma_{-} \equiv \log T_{16\%} - \log T_{50\%}, \tag{8.17}$$

$$\Sigma_{+} \equiv \log T_{84\%} - \log T_{50\%}. \tag{8.18}$$

The fractional errors in the temperature, given as

$$\varepsilon_{-} \equiv \frac{T_{16\%}}{T_{50\%}} - 1, \tag{8.19}$$

$$\varepsilon_{+} \equiv \frac{T_{84\%}}{T_{50\%}} - 1, \tag{8.20}$$

will also be quantified. Similar to the analysis on the gas-dark matter density relation, the temperature values in each gas density bin are ranked order and the 16%, 50%, and 84% marks in the cumulative fraction are found.

In Figure 8.3 the negative and positive errors $\Sigma_{\pm}$ in $\log T$ for the adiabatic temperature density relation are approximately equal in magnitude and therefore, the scatter in $\log T$ should approximately follow a normal distribution and the scatter in T given by a lognormal distribution. The fractional errors in the temperature, defined by $\varepsilon_{\pm}$ are plotted in Figure 8.7. The negative fractional error ε_{-} is similar across the redshift range $3 \leq z \leq 6$, with a value of -0.7 at mean densities $\Delta_g \sim 1$ and a value of -0.3 at high densities $\Delta_g \sim 1000$. The positive fractional error ε_{+} is seen to increase with decreasing redshifts, particularly with decreasing gas densities $\Delta_g < 10$. At mean densities $\Delta_g \sim 1$ it has values $\varepsilon = (2.0,\ 1.8,\ 1.6,\ 1.4)$ for redshifts $z = (3,\ 4,\ 5,\ 6)$, respectively. At high densities $\Delta_g \sim 1000$, the positive fractional error has a value of 0.4 that is independent of redshift. The change with redshift is consistent with those in the evolution of the adiabatic temperature-density relation in Figures 8.3 and 8.4. Redshift evolution appears to be more significant with decreasing gas densities and temperatures.

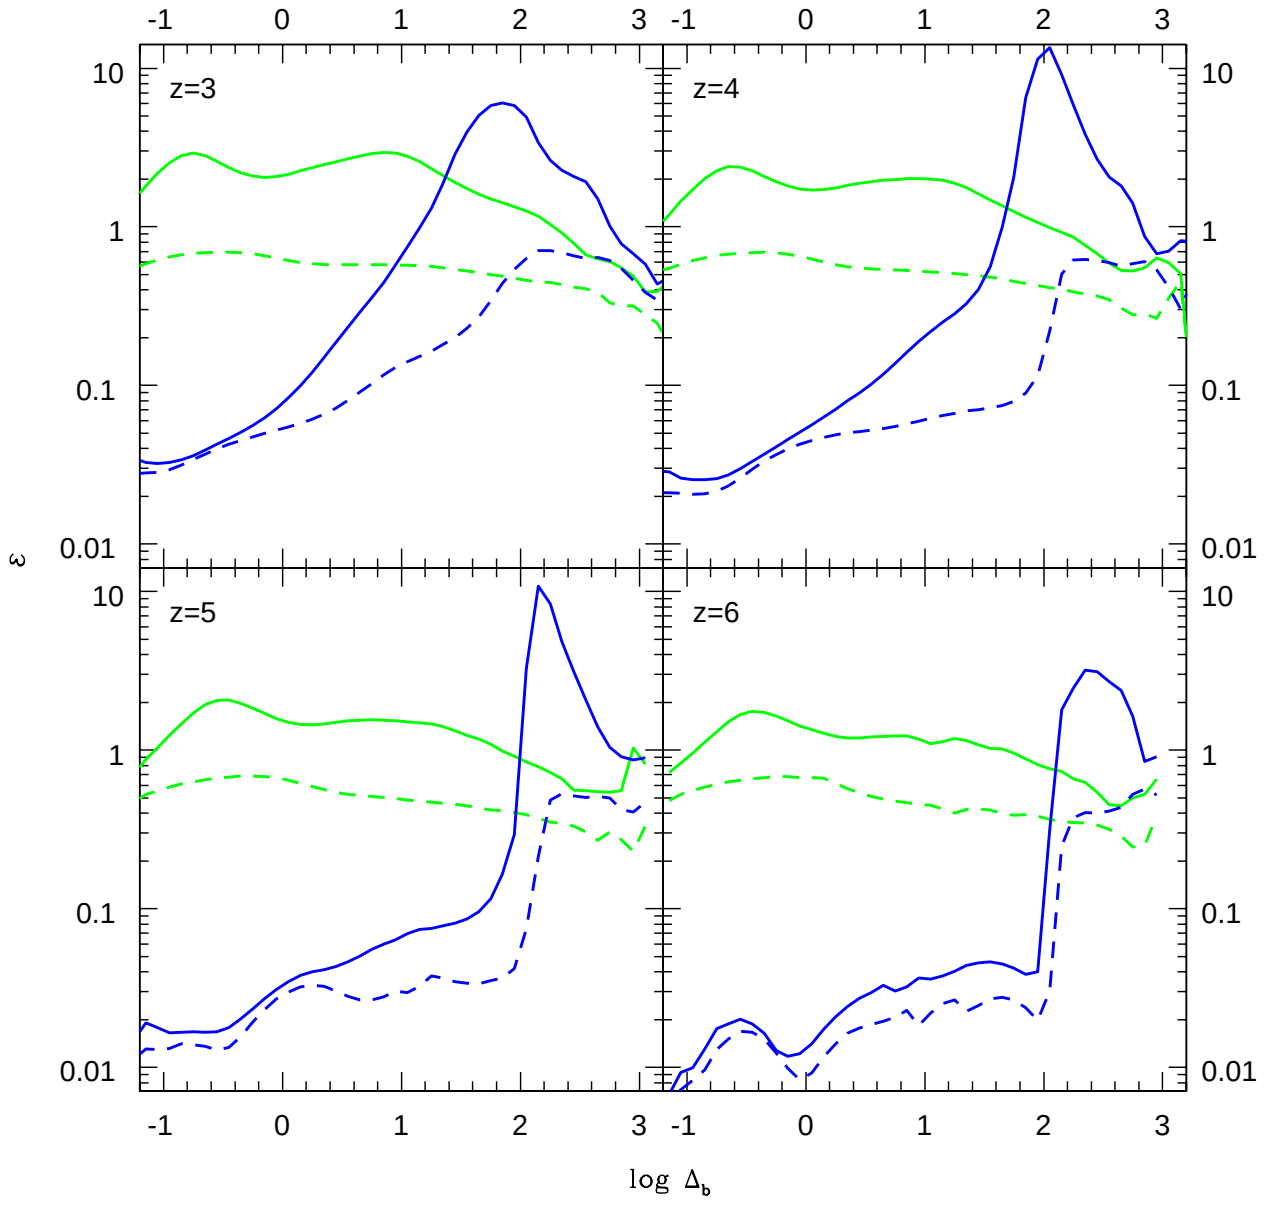

Figure 8.7 The fractional errors in the adiabatic (green) and HM (blue) temperature-density relations. The positive fractional error ε_+ (solid) is in general larger than the absolute value of the negative fractional error ε_- (dashed), showing that the distribution of scatter in T is skewed towards higher temperatures.

The shape of the scatter distribution for the adiabatic temperature-density relation is shown in Figure 8.8. The distribution of scatter in $\log T$ at fixed gas densities $\Delta_b = 1$, 3, 10, and 30 for redshifts $z = 3$, 4, 5, and 6 are plotted. The distributions have been rescaled and shifted so that the curves for different densities can be directly compared. The x-axis has been rescaled by Σ, defined as the symmetrical deviation for which the area under the probability distribution $p(x)$ in the range $-1 \leq x \leq 1$ is equal to 0.68. The redshift evolution of the scatter in the adiabatic temperature density relation is also apparent in Figure 8.8. More redshift evolution is seen with decreasing gas densities. In general the scatter in $\log T$ is not normally distributed. At higher densities $\Delta_b > 10$, the distribution is skewed towards higher temperatures and the shape of the distribution shows relatively little redshift evolution. At lower densities, the distribution changes from being skewed towards lower temperatures at $z = 6$ and to being skewed towards higher temperatures at $z = 3$. The skewness of T relative to normal is a signature of shock heating, which converts kinetic energy into thermal energy, thereby raising the gas temperature. The skewness of $\log T$ relative to normal appears to be a signature of stronger shock heating from gravitational collapse.

In Figure 8.7, the fractional errors in the HM temperature-density relation are dependent on the relative efficiencies of radiative cooling and heating and shock heating. For the chosen Haardt-Madau photoionization background, radiative cooling and heating is more strongly redshift dependent than shock heating. At higher redshifts $z \gtrsim 5$, both fractional errors $\varepsilon_{\pm}$ are very small in magnitude because the region bounded by the heating track and the equilibrium track covers a large fraction of the relevant portion of the temperature-density phase space (see Figure 8.3). Gas falling in this cooling region of phase space can efficiently cool down to the heating track while gas with temperatures just below the heating track are efficiently photoheated back up to it. Thus, at high redshifts, the scatter in the temperature-density relation is expected to be very small for uniform reionization and photoionization background, and photoionization equilibrium. Nonuniformity and nonequilibrium conditions will result in additional fluctuations in the temperature-density relation. In principle, these fluctuations can be extracted from higher order statistics of the Lyα forest (Zaldarriaga, Seljak, & Hui, 2001; Fang & White, 2004).

At lower redshifts, the cooling zone in the temperature-density phase space shifts to higher gas densities and the scatter becomes more appreciable, particular where shock heating is effective. The fluctuations from shock heating appear as an increase in the amplitude of the positive fractional error ε_+. By redshift $z = 3$, the temperature fluctuations for the gas density range $0 < \Delta_b < 10$, already exceed the systematic error of $\Delta T < 2000K$ that McDonald et al. (2001) have estimated for their measurement of the temperature-density relation from high resolution Keck spectra of Lyα absorption lines (Rauch et al., 1997).

The effects of reduced efficiency in cooling and an increase in shock heating on the fluctuations in the temperature-density relation are more apparent in the scatter distributions shown in Figure 8.9. The distribution of scatter in $\log T$ at fixed gas densities $\Delta_b = 1$, 3, 10, and 30 for redshifts

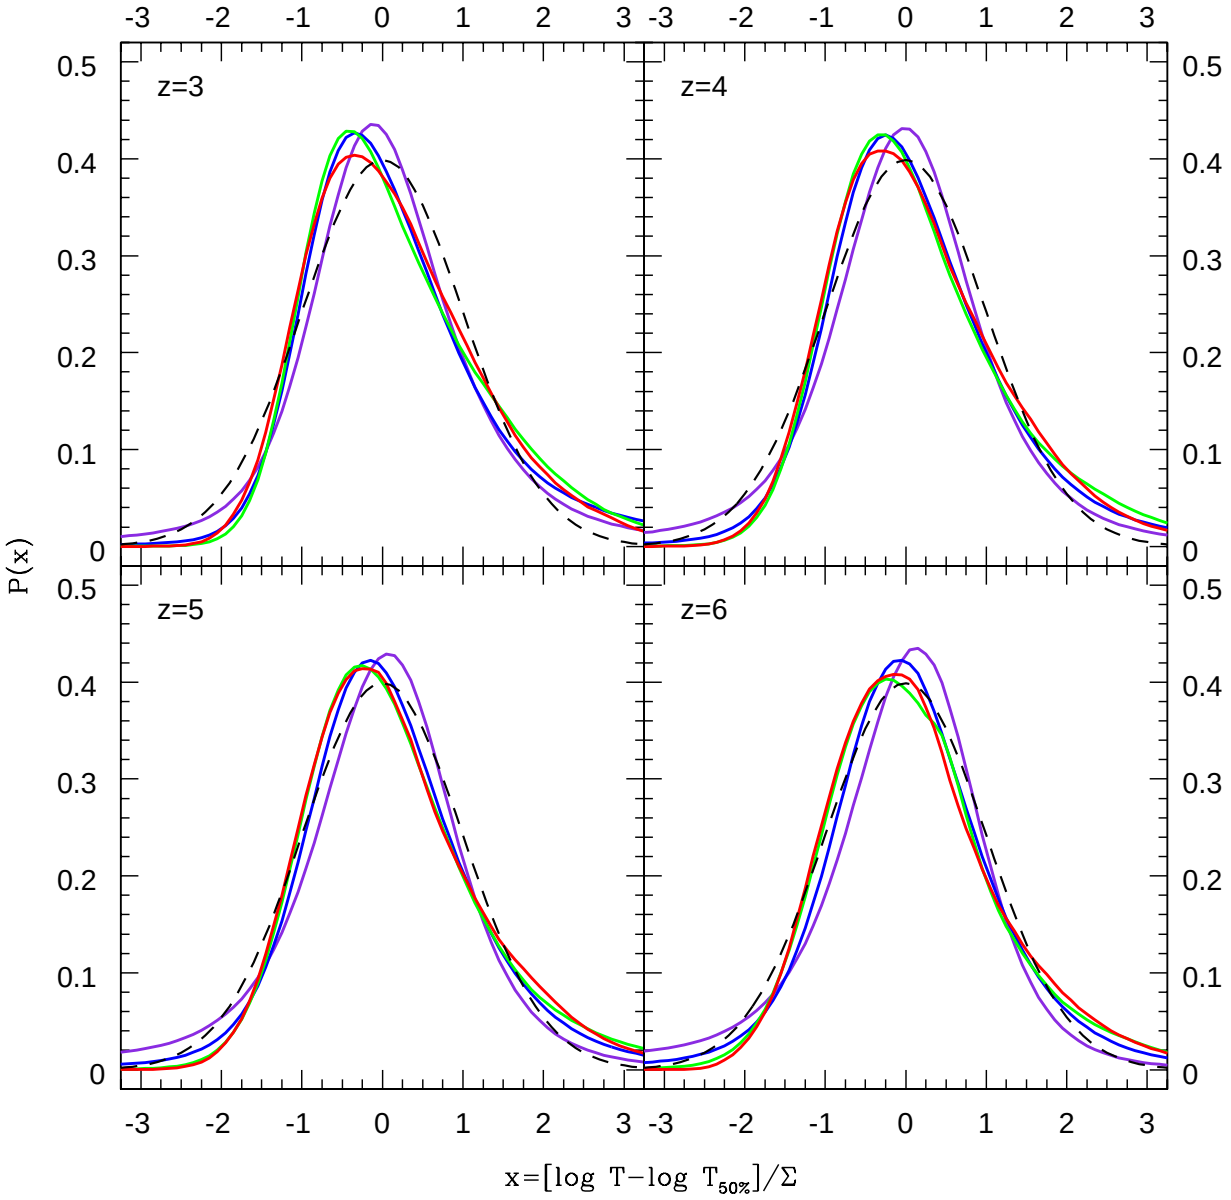

Figure 8.8 The shape of the scatter distribution for the adiabatic temperature-density relation. The distribution of scatter in $\log T$ at fixed gas densities $\Delta_b = 1$ (violet), 3 (blue), 10 (green), and 30 (red) for redshifts $z = 3$, 4, 5, and 6. More redshift evolution is seen with decreasing gas densities. In general the scatter is not normally distributed (dashed), though the two lower density bins change from being skewed towards the left at $z = 6$ to being skewed towards the right at $z = 3$ and the transition in between is close to being Gaussian.

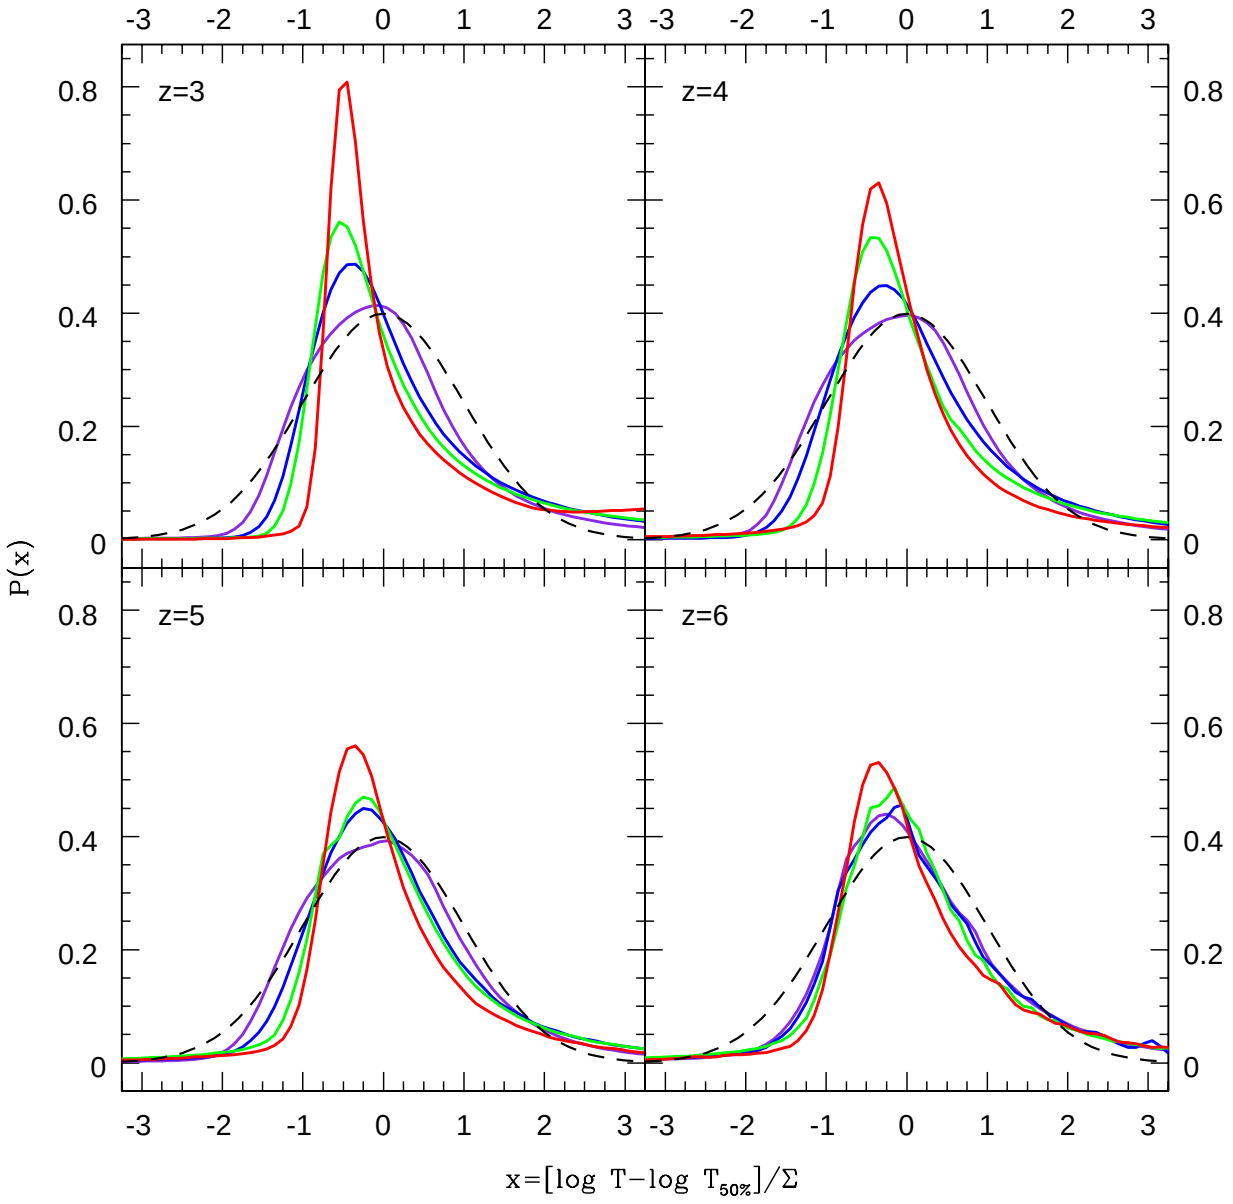

Figure 8.9 The shape of the scatter distribution for the HM temperature-density relation. The distribution of scatter in $\log T$ at fixed gas densities $\Delta_b = 1$ (violet), 3 (blue), 10 (green), and 30 (red) for redshifts $z = 3$, 4, 5, and 6. The distributions are all skewed relative to a Gaussian (dashed). At lower redshifts and higher gas densities, the relative efficiency of cooling decreases and shock heating increases, resulting in scatter distributions that are more skewed towards higher temperatures.

$z = 3$, 4, 5, and 6 are plotted. The distributions are quite smooth and none of them can be said to be Gaussian. At lower redshifts and higher gas densities, the relative efficiency of cooling decreases and shock heating increases, resulting in scatter distributions that are more skewed towards higher temperatures. Fang & White (2004) have assumed that the scatter in $\log T$ is normally distributed in their paper on using higher order moments of Lyα forest spectra to constrain scatter in the temperature-density relation. A more realistic analysis should use a distribution that becomes more skewed at lower redshifts and higher densities. In principle, the thermal history of the IGM is imprinted in the temperature-density relation and its scatter, and higher order statistics can provide information on reionization, star formation feedback, and radiative transfer, though the information may be degenerate. However, such studies, still in their infancy, will rely on detailed quantification of the shape of the scatter distribution. In future work, the skewness of the scatter distributions will be quantified by fitting the distributions with an Edgeworth expansion (e.g. Juszkiewicz et al., 1995).

8.5 Summary

The out-of-core hydro code has been used to run a high resolution cosmological simulation with 1536^3 grid cells and 768^3 dark matter particles in a 50 h^{-1}Mpc box. The thermal history of the high redshift $3 \leq z \leq 6$ IGM has been simulated in both an adiabatic universe and nonadiabatic one with reionization, UV heating, and atomic cooling.

In the coupled Eulerian hydro and PM N-body simulation, the gas is found to be unbiased relative to the dark matter everywhere except in high density regions like halos. The slope of the gas-dark matter density relation is consistent with unity for $\Delta_m \lesssim 100$. Both the scatter and the distribution of scatter has been precisely measured. The fractional error in the relation is found to increase monotonically with decreasing density. At high gas densities $\Delta_b \gtrsim 100$, the scatter in $\log \Delta_m$ is normally distributed. For decreasing gas density, the scatter distribution becomes more skewed towards lower values. The skewness is an artifact related to the poor particle mass resolution at low densities, a property of the Lagrangian N-body scheme. An empirical expression for the minimum reliable dark matter density in a PM N-body simulation has also been presented.

The adiabatic temperature-density relation is primarily regulated by shock heating from gravitational collapse. The shock heating induces a temperature-density relation where the effective slope decreases with increasing gas density and approaches a value close to unity at high gas densities $\Delta_b \gtrsim 100$. The distribution of scatter in $\log T$ is not given by a Gaussian, but rather it is slightly skewed towards higher temperatures. Over the redshift range $3 \leq z \leq 6$, evolution is found primarily at lower densities and temperatures. The importance of shock heating extends to lower densities as the universe evolves.

The Haardt-Madau temperature-density relation can be qualitatively explained using the cool-

ing and heating tracks, which divide the temperature-density phase space into three distinct regions. Gas with temperatures just below the heating track are quickly photoheated back up to it, while gas falling in the cooling zone are efficiently cooled back down to the heating track. However, at lower redshifts, the cooling zone recedes towards higher densities and shock heating can change the temperature of the gas and also the scatter distribution. Radiative cooling and heating imprints a scatter distribution in $\log T$ that is skewed towards higher temperatures. At lower redshifts and higher gas densities, the relative efficiency of cooling decreases and shock heating increases, resulting in scatter distributions that are more skewed towards higher temperatures.

The temperature-density relation defined by the heating track can be described by a double power-law, but the HM temperature-density relation is more complicated. Reionization changes the effective slope of underdense gas while shock heating can increase the effective slope. At lower redshifts $z \sim 3$, shock heating is found to increase the effective slope even in the gas density range $0 < \Delta_b < 10$.

Chapter 9

The Clustering and Biasing of Dark Matter Halos

9.1 Introduction

Kaiser (1987) originally introduced a local, linear biasing relation where the galaxy density field δ_g is related to the matter density field δ_m by the relation $\delta_g = b\delta_m$, where b is the linear bias parameter. In general, the bias parameter can be a function of scale in both real and Fourier space. Furthermore, to account for the possibly nonlinear (Fry & Gaztanaga, 1993) and nonlocal (Heyl et al., 1995) nature of galaxy formation, stochastic biasing has been proposed (Pen, 1998b; Tegmark & Peebles, 1998; Dekel & Lahav, 1999). In the linear stochastic biasing model, the linear bias parameter b describes the relative amplitude of the galaxy and matter density fields while the galaxy-matter cross-correlation coefficient r parametrizes the stochasticity in the biasing relation.

Linear perturbation theory furthermore predicts that the observed redshift space galaxy power spectrum in the linear regime will be amplified as a result of redshift distortions caused by large-scale peculiar velocity fields. On large scales, outflow out of underdense regions and infall towards overdense regions cause the clustering pattern to appear compressed along the line of sight, opposite to the fingers-of-god effect where the elongations along the line of sight are caused by small-scale peculiar velocities (see Hamilton, 1998, for a review). The large-scale distortions cause the redshift space galaxy power spectrum to be anisotropic in the linear regime and the amplitude of the distortion depends on a parameter $\beta = \Omega^0.6/b$ (Kaiser, 1987). The redshift distortions arise from gravitational induced peculiar velocity fields which directly probe the underlying mass distribution. In principle, by measuring the distortions present in redshift surveys, one can recover the real space galaxy power spectrum and also extract the bias parameters.

In very recent years, the focus is on measuring luminosity-dependent clustering and luminosity-dependent bias in modern surveys such as 2dFGRS (Norberg et al., 2001) and SDSS (Tegmark et al., 2004a). These observations suggest that the relative bias between galaxies of different luminosities

is approximately scale-invariant on large-scales $k < 0.1\ h/\mathrm{Mpc}$, allowing one to determine the bias $b(k)$ up to a multiplicative factor. In principle, the absolute bias can be determined through the measurement of the parameter β from redshift distortions, but in practice there is still considerable statistical uncertainty in its value.

The amplitude and shape of the bias can also be determined by comparing the luminosity-dependent galaxy bias measured from redshift surveys with the mass-dependent halo bias measured from N-body simulations. Analytical biasing models (Sheth & Tormen, 1999) and numerical simulations (Seljak & Warren, 2004) suggest that for low mass halos and faint galaxies, the large-scale bias approaches a minimum value that is independent of mass and luminosity. By measuring the minimum bias in galaxy redshift surveys and comparing it with that found in numerical simulations of dark matter halos, the amplitude of the matter power spectrum can be constrained (Seljak et al., 2004).

In this chapter, I present results on the mass-dependent clustering and biasing of dark matter halos in high resolution simulations run with PMFAST, a new particle-mesh N-body code developed by Merz, Pen, & Trac (2004). These simulations have the highest mass resolution for studies of this nature to date. Parameters of the linear stochastic bias model are measured and examined to determine where the assumption of scale-invariance is valid and how well correlated are the halos with each other and with the underlying dark matter. At present, the shape of the matter power spectrum determined from galaxy clustering is highly dependent on the assumption that the large-scale bias is scale-invariant.

9.2 PMFAST

PMFAST is a parallel particle-mesh N-body code developed by Merz, Pen, & Trac (2004) to run on distributed memory machines. It is based on a two-level mesh gravity solver where the gravitational forces are separated into long and short-range components (Couchman, 1991; Trac & Pen, 2004b), like the out-of-core gravity solver described in Chapter 7. The short-range forces are computed on a localized, high-resolution mesh using highly optimized vendor FFT libraries that operate on data stored locally on each node. The long-range forces are computed on a global mesh that is 4 times coarser in each dimension, reducing memory storage, communications, and computation by a factor of 64. This decomposition scheme minimizes communications costs and allows tolerance for slow networks between nodes on distributed memory clusters. The current version uses a slab decomposition geometry and has periodic boundary conditions for cosmological applications. The particle-to-mesh assignment and time integration of the equations of motions follow similarly to the out-of-core implementation described in Chapter 7. PMFAST is cost-effective, memory-efficient, and is made freely available to the public.

At the Canadian Institute for Theoretical Astrophysics (CITA) we have an IA-64 cluster consist-

ing of 8 nodes, each of which contains quad 733 MHz Itanium1 processors and 64 GB of RAM. The cluster has point-to-point gigabit ethernet connections between each node. PMFAST features support for both distributed and shared memory parallelization through the use of MPI and OpenMP, respectively. With 512 GB of distributed memory, large N-body simulations have been run using a 3712^3 effective mesh and 1856^3 particles. Each of the 8 slabs in the decomposed domain has a width of 464 grid cells with buffers of 24 cells on each end to make up a total width of 512 grid cells. The short-range forces are computed for 512^3 sections of the slab using the highly optimized Intel IPP FFT library. The long-range forces are computed on a global 928^3 mesh using the MPI FFTW library. PMFAST is highly memory-efficient, with the particle positions and velocities being the main cost, while the force arrays contribute relatively small overhead.

9.2.1 Cosmological N-body Simulations

Two cosmological N-body simulations have been run in periodic boxes with comoving side lengths of 200 and 400 h^{-1}Mpc. The PMFAST simulations have been run using a 3712^3 effective mesh and 1856^3 dark matter particles. The higher resolution run has a comoving grid spacing of $\Delta x = 54\ h^{-1}$kpc and a particle mass of $\Delta m = 9.4 \times 10^7\ h^{-1} M_\odot$. The N-body initial conditions are generated using the out-of-core cosmological initial conditions generator described in Chapter 7. The initial matter transfer function is computed using CMBFAST (Seljak & Zaldarriaga, 1996) and with WMAP parameters: $\Omega_m = 0.27$, $\Omega_\Lambda = 0.73$, $\Omega_b = 0.044$, $n = 1$, $\sigma_8 = 0.84$, and $h = 0.7$.

Dark matter halos are identified using a spherical overdensity algorithm similar to that described by Lacey & Cole (1994). Halo centers are located as the minima in the gravitational potential, which is computed at the same resolution as the PM grid. The halo mass is then calculated by taking the total mass within a sphere for which the mean overdensity is given by $\kappa = 178$. Note that subhalos found within larger halos are not counted. For the 200 and 400 h^{-1}Mpc simulations, the number of halos found with a minimum mass corresponding to a particle count of 100 are 1.4×10^6 and 1.9×10^6, respectively. The halo numbers in various mass bins are collected in Table 9.1.

The spherical overdensity mass functions are compared with the analytical prediction by Press & Schecter (PS; 1974) and the numerical results from Jenkins et al. (2001) in Figure 9.1. The PS prediction is known to overestimate the abundance of lower mass halos and underestimate the number of very massive halos. The higher resolution 200 h^{-1}Mpc simulation is in agreement to better than 20% with Jenkins et al. (2001) for masses $M > 10^{11}\ h^{-1} M_\odot$, while the 400 h^{-1}Mpc simulation is good for masses $M > 10^{12}\ h^{-1} M_\odot$. At high masses $M > 10^{14}\ h^{-1} M_\odot$, the small number of halos results in a relatively large Poisson error.

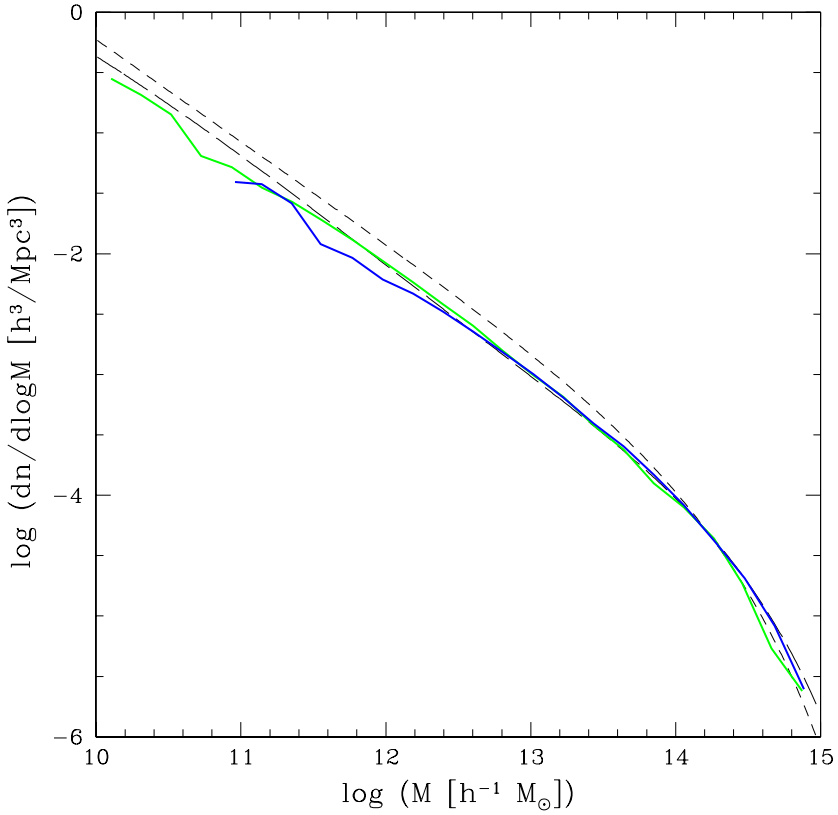

Figure 9.1 Spherical overdensity mass functions of dark matter halos from PMFAST N-body simulations of a WMAP cosmology. Results from the 200 $h^{-1}M_\odot$ (green) and 400 $h^{-1}M_\odot$ (blue) simulations are plotted down to a minimum mass corresponding to a particle count of 100. They are in better agreement with Jenkins et al. (2001) than with Press & Schechter (1974), where the difference relative to the former is $< 20\%$ for a particle count $\gtrsim 1000$.

9.3 Bias and Stochasticity of Dark Matter Halos

Linear stochastic bias of halos relative to the underlying dark matter can be completely described in terms of the the halo-halo power spectrum $P_{hh}(k)$, the dark matter power spectrum $P_{mm}(k)$, and the halo-matter cross spectrum $P_{hm}(k)$. In this model, the biasing relation is parametrized by the linear bias parameter,

$$b_{hm}(k) \equiv \sqrt{\frac{P_{hh}(k)}{P_{mm}(k)}}. \tag{9.1}$$

In the literature, the distribution or field is said to be biased for $b > 1$ and antibiased for $b < 1$, though the word antibiased is somewhat a misnomer. The statistical correlation in the linear biasing relation is specified by the cross-correlation or stochasticity parameter,

$$r_{hm}(k) \equiv \frac{P_{hm}(k)}{\sqrt{P_{hh}(k)P_{mm}(k)}}, \tag{9.2}$$

Table 9.1. Halo Numbers as a Function of Mass

Mass $x \equiv M/(h^{-1}M_\odot)$	Number (L =200 h^{-1}Mpc)	Density (h^3Mpc^{-3})	Number (L =400 h^{-1}Mpc)	Density (h^3Mpc^{-3})
$11.00 \leq \log x \leq 11.25$	73694	9.21×10^{-3}	622463	9.73×10^{-3}
$11.25 \leq \log x \leq 11.50$	52025	6.50×10^{-3}	402935	6.30×10^{-3}
$11.50 \leq \log x \leq 11.75$	33739	4.22×10^{-3}	158726	2.48×10^{-3}
$11.75 \leq \log x \leq 12.00$	21608	2.70×10^{-3}	122558	1.91×10^{-3}
$12.00 \leq \log x \leq 12.25$	13192	1.65×10^{-3}	80967	1.27×10^{-3}
$12.25 \leq \log x \leq 12.50$	8044	1.01×10^{-3}	54969	8.59×10^{-4}
$12.50 \leq \log x \leq 12.75$	4856	6.07×10^{-4}	35029	5.47×10^{-4}
$12.75 \leq \log x \leq 13.00$	2701	3.38×10^{-4}	21501	3.36×10^{-4}
$13.00 \leq \log x \leq 13.25$	1633	2.04×10^{-4}	13082	2.04×10^{-4}
$13.25 \leq \log x \leq 13.50$	898	1.12×10^{-4}	7319	1.14×10^{-4}
$13.50 \leq \log x \leq 13.75$	497	6.21×10^{-5}	4215	6.59×10^{-5}
$13.75 \leq \log x \leq 14.00$	248	3.10×10^{-5}	2242	3.50×10^{-5}

which can have values $-1 \leq r \leq 1$. A positive correlation is described by $r > 0$, while a negative or anticorrelation is specified by $r < 0$. If there is no stochasticity then $r = 1$ and biasing becomes deterministic. The linear bias parameter b and the stochasticity parameter r together completely describe the linear stochastic biasing model, but it is also useful to consider the cross-bias parameter,

$$c(k) \equiv \frac{b(k)}{r(k)} = \frac{P_{hm}(k)}{P_{hh}(k)}. \tag{9.3}$$

In principle, the cross-bias parameter for galaxies can be determined by measuring the monopole, decapole, and hexadecapole harmonics of the redshift space galaxy power spectrum (Hamilton, 1998). In practice, this is rather difficult because the hexadecapole distortion is much smaller than the other two harmonics. Only recently has this parameter been investigated in some detail with galaxy-galaxy lensing measurements in the SDSS (Sheldon et al., 2004) and N-body simulations of dark matter halos (Tasitsiomi et al., 2004). x

9.3.1 Halo Power Spectra as a Function of Mass

Power spectra for the simulated halo and dark matter distributions have been calculated in order to quantify the parameters of the linear stochastic biasing model. In Figure 9.2(a-b), the halo power

spectrum $P_{hh}(k)$ is plotted for halos in six mass bins that span the range $10^{11} < M/(h^{-1}M_{\odot}) < 10^{14}$ and have widths of 0.5 dex. The shading illustrates the 1-σ uncertainty, in the spirit of (Tegmark et al., 2004a), from sample variance alone, estimated using the fluctuations in each Fourier mode bin. The Poisson noise power has been subtracted from the halo spectra by removing the white noise power $P_N = 1/N$ due to all self-pairs. This subtraction is correct on large scales, but the Poisson noise is expected to fall below the white noise level on small scales where the finite extent of halos becomes important (Seljak, 2000). Halos can be no closer to each other than two times the viral radius and this introduces an anticorrelation that needs to be factored into the Poisson noise calculation. This problem becomes important and evident only at high masses where the halo abundance is small and the shot noise large.

The halo spectra for different mass bins are all consistent with being power-laws for wave numbers $k > 0.03$ h/Mpc where the dark matter power spectrum $P_{mm}(k)$ turns over. The halo spectra have approximately the same power-law shape and slope as each other. The amplitude of the clustering is clearly mass dependent and the bias parameter $b_{hm}(k)$ will be also. For comparison the simulated dark matter power spectrum is shown by the solid line and it highly agrees with the nonlinear fitting function (Smith et al., 2003) given by the dashed line, except at long wavelengths near the box size where sample variance is relatively large. The shot noise in the matter power spectrum is very small due to the very large number of dark matter particles. In Figure 9.2(c-d), the halo-matter cross spectrum $P_{hm}(k)$ are plotted for the same mass bins. There is no shot noise in the cross spectra since there are no self-pairs. The cross spectra are roughly power-law but there appears to be slightly more curvature in the shape, particularly at lower masses and smaller scales, compared to the halo spectra.

The halo spectra are well fit by power laws $P(k) \propto k^n$ in the wavenumber range $0.03 < k$ (h/Mpc) < 0.7. At low masses $M \sim 10^{11}$ $h^{-1}M_{\odot}$, the halo spectral slope has a value $n = -1.60 \pm 0.04$ and it increases monotonically to a value of $n = -1.39 \pm 0.06$ at high masses $M \sim 10^{14}$ $h^{-1}M_{\odot}$. For all halos in the mass range $10^{11} < M/(h^{-1}M_{\odot}) < 10^{14}$, the halo spectral slope is given by $n = -1.64 \pm 0.02$. Tegmark et al. (2004a) have measured power spectra for SDSS galaxies in six bins of absolute magnitude in the range $-23 < M_r < -17$ and found them to share similar power-law shape. Unfortunately, they did not publish the slopes of the spectra with which to compare to.

The cross spectral slope displays more spread, ranging from $n = -1.69 \pm 0.06$ at the low mass end to $n = -1.27 \pm 0.03$ at the high mass end. For all halos in the mass range $10^{11} < M/(h^{-1}M_{\odot}) < 10^{14}$, the cross spectral slope is given by $n = -1.74 \pm 0.03$. Sheldon et al. (2004) have used galaxy-galaxy lensing observations in the SDSS to measure the galaxy-matter correlation function $\xi_{gm}(r)$. They also found it to be consistent with being a power-law and the amplitude and logarithmic slope are seen to increase with galaxy luminosity. These trends are consistent with the mass dependence of the halo-matter cross spectra. In future work, the halo-matter correlation functions from the

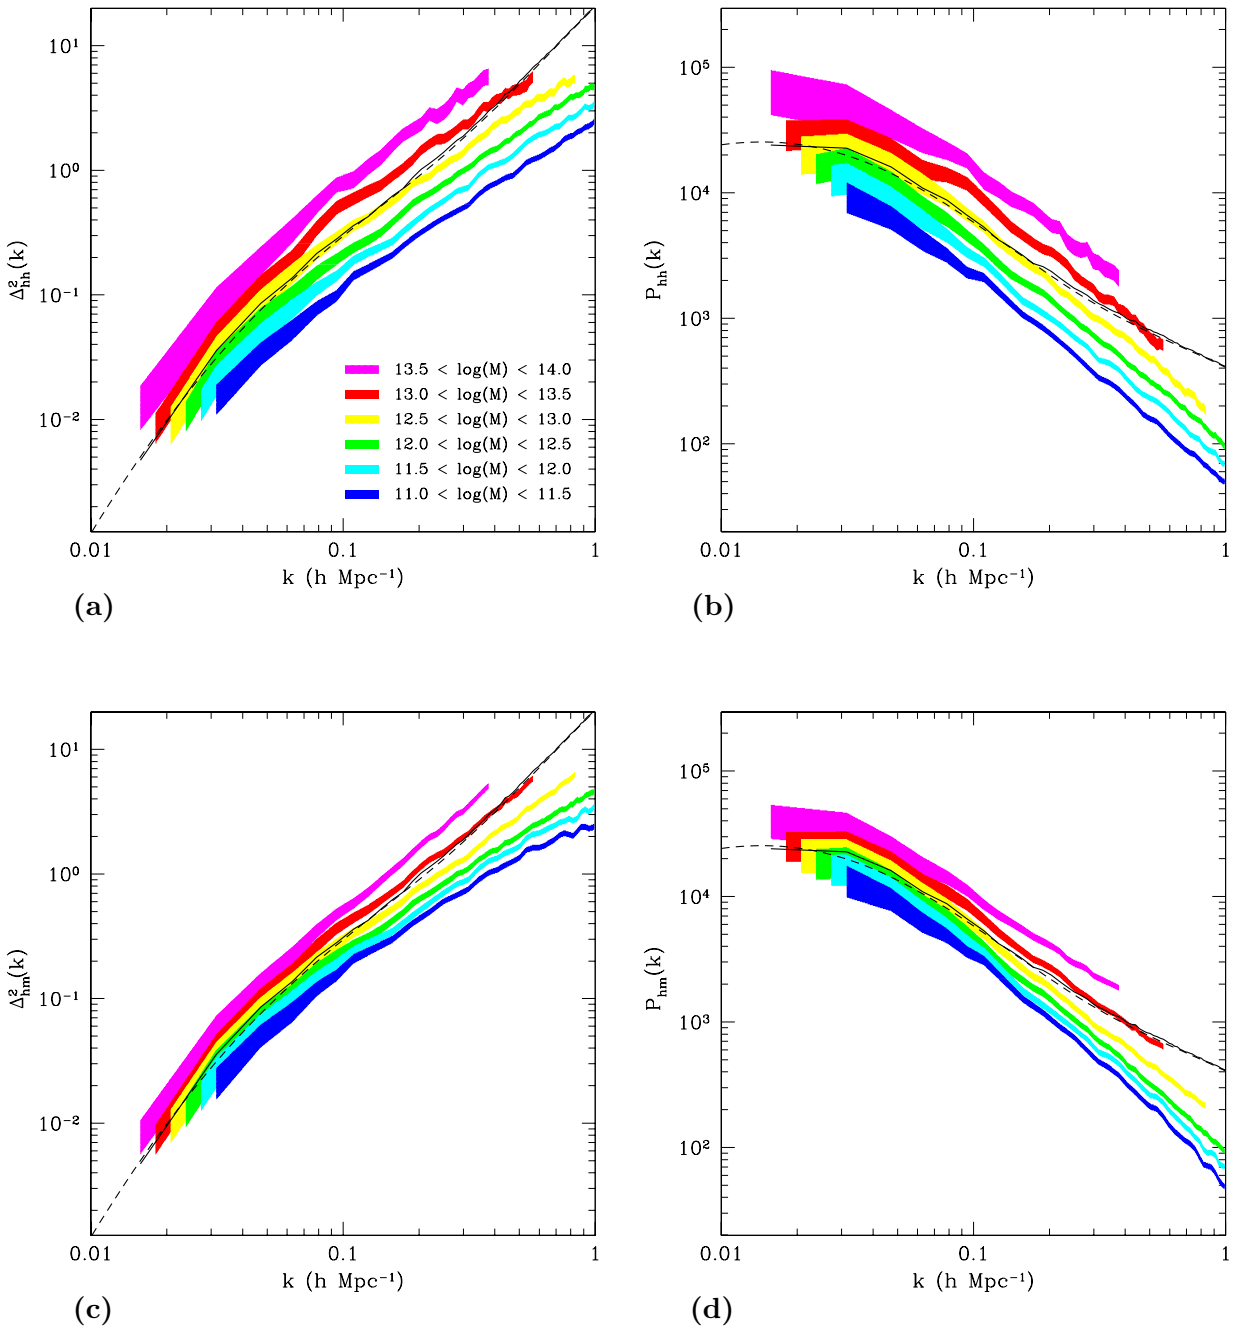

Figure 9.2 PMFAST simulated power spectra at redshift $z = 0$. In (a) and (b), the halo power spectrum $P_{hh}(k)$ is plotted for halos in six mass bins spanning the range $10^{11} < M/(h^{-1}M_{\odot}) < 10^{14}$. In (c) and (d), the halo-matter cross spectrum P_{hm}(k) for the same bins are shown. The shading illustrates the 1-σ uncertainty from sample variance. The solid line is the simulated dark matter power spectrum $P_{mm}(k)$ while the dashed line is from the nonlinear fitting functions of Smith et al. (2003).

PMFAST simulations will be calculated and compared with Sheldon et al. (2004).

9.3.2 Halo Bias as a Function of Mass

The parameters of the linear stochastic bias model are plotted in Figure 9.3. On large scales $k < 0.1$ h/Mpc, the bias parameter $b_{hm}(k)$ is approximately scale-invariant, but in general the power-law form for the halo spectra combined with a non power-law dark matter power spectrum leads to scale-dependent linear bias. This scale-dependence is real and not an artifact of mass-dependent clustering. Massive halos cluster are more strongly and dominate the large-scale power while the more abundant lower mass halos contribute more small-scale power. Thus, the power spectrum measured for a sample spanning a reasonably large range in mass will have a slope more negative than one where the mass is kept fixed to a narrow range. In Figure 9.3, the mass bins have reasonably small widths of 0.5 dex and further decreasing the size of the bins changes the results by relatively little. The scale-dependent linear bias is a real property that reflects the curvature in the nonlinear dark matter power spectrum.

On large scales $k < 0.1$ h/Mpc, there is no stochasticity with $r_{hm}(k) \approx 1$. Halos with masses $M > 10^{13}$ $h^{-1}M_{\odot}$ are highly correlated with the dark matter even at smaller scales. Note that the stochasticity parameter for the largest mass bin systematically exceeds unity because the shot noise subtraction overcorrects the actual Poisson noise. For the three lowest mass bins spanning the range $10^{11} < M/(h^{-1}M_{\odot}) < 10^{12.5}$, the stochasticity curves lie approximately above one another and this is suggestive of a minimum curve. In order to confirm this result, the analysis must be extended to lower masses. For future work, a PMFAST simulation with a 100 h^{-1}Mpc box will be run to probe down to mass scales $M \sim 10^{10}$ $h^{-1}M_{\odot}$.

The cross-bias parameter $c_{hm}(k)$ is approximately scale invariant, a reflection of the halo and cross spectra having similar power-law shapes and slopes. The developing picture is that as the correlation between the halos and dark matter decreases (parametrized by the stochasticity parameter), the relative amplitude of the clustering (parametrized by the bias parameter) also decreases, roughly proportionally.

In Figure 9.4 the parameters of the halo bias model are recalculated relative to halos of nonlinear mass, defined as the mass $M^* = (4/3)\pi\bar{\rho}R^3$ enclosed within a sphere of radius R for which the variance,

$$\sigma^2 = \int \frac{d^3k}{(2\pi)^3}\langle|\delta(\boldsymbol{k})|^2\rangle|W^R(\boldsymbol{k})|^2 = \delta_c^2, \tag{9.4}$$

of the linear density field δ, smoothed by a tophat filter W^R, equals the square of the critical overdensity threshold $\delta_c = 1.68$ for spherical collapse. For the chosen cosmology at redshift $z = 0$, the nonlinear mass is $M^* = 4.62 \times 10^{12}$ $h^{-1}M_{\odot}$.

A visual inspection of Figure 9.4(a) shows the approximately scale-invariant form of the relative bias b_{hh^*} (also equal to b_{hm}/b_{h^*m}). However, the relative bias is not exactly scale-invariant, but

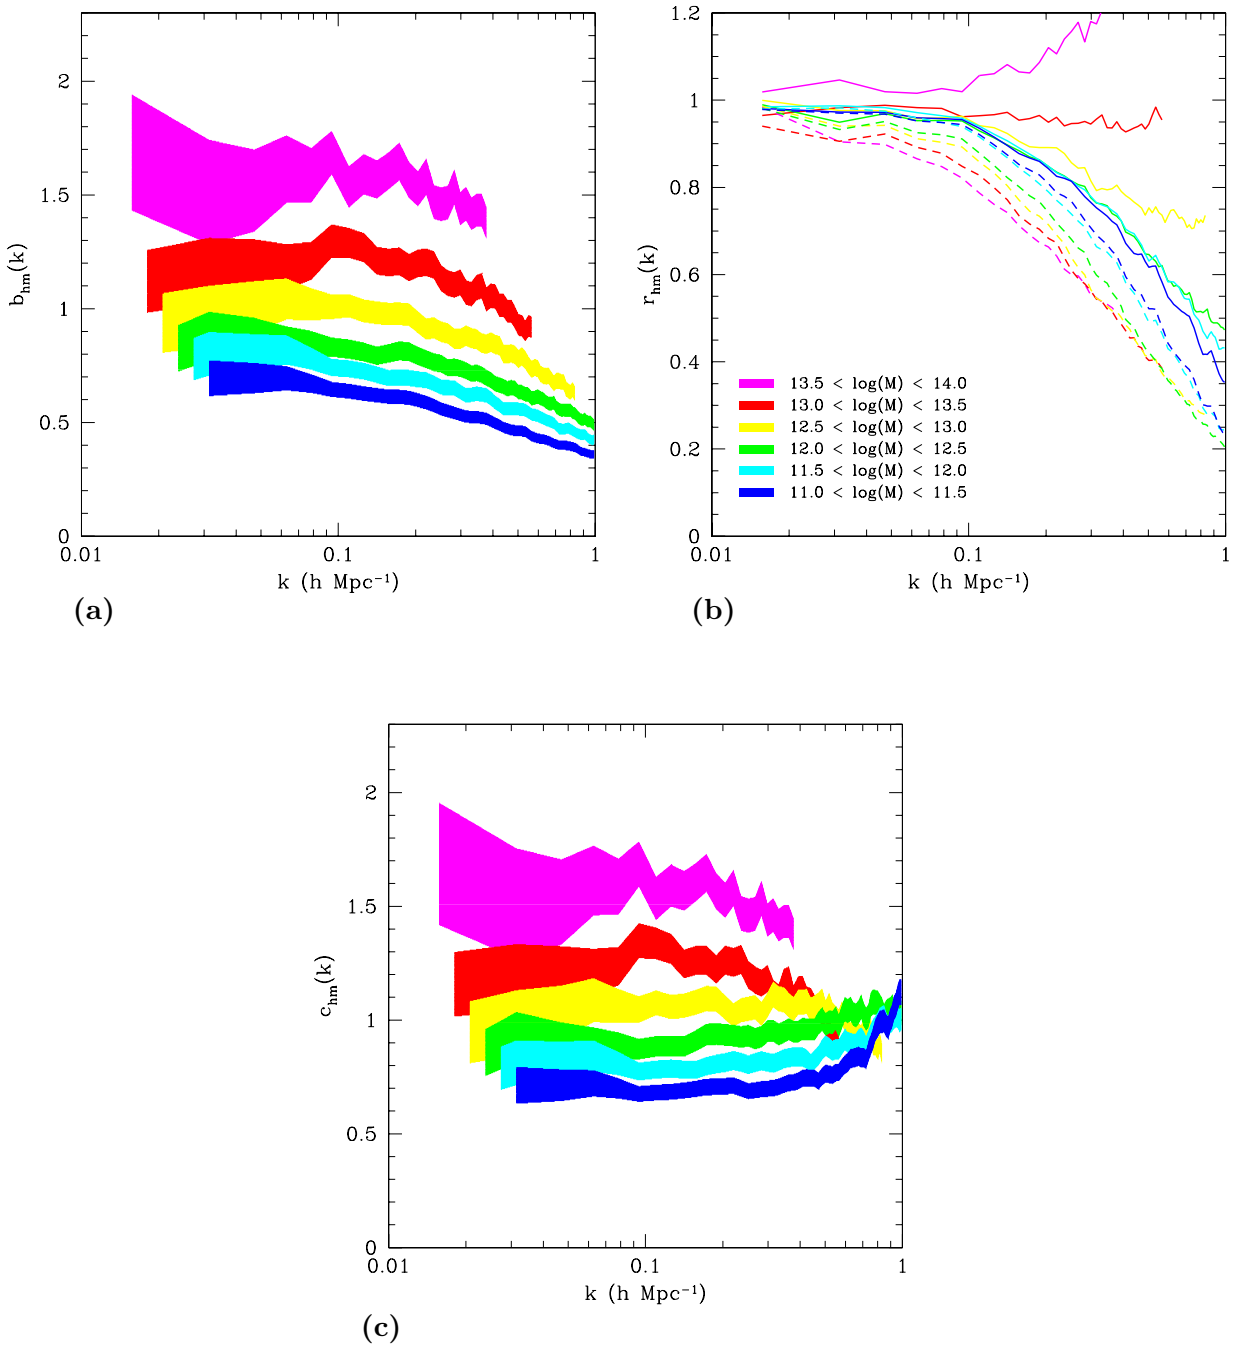

Figure 9.3 Parameters of the linear stochastic bias model. In (a) the power-law form for the halo spectra combined with a non power-law dark matter power spectrum leads to scale-dependent bias $b_{hm}(k)$. In (b) the stochasticity parameter $r_{hm}(k)$ is plotted prior to (dashed lines) and after (solid lines) shot noise subtraction. The 1-σ uncertainty shading is not shown for clarity. In (c) the cross-bias parameter $c_{hm}(k)$ is seen to be approximately scale-invariant.

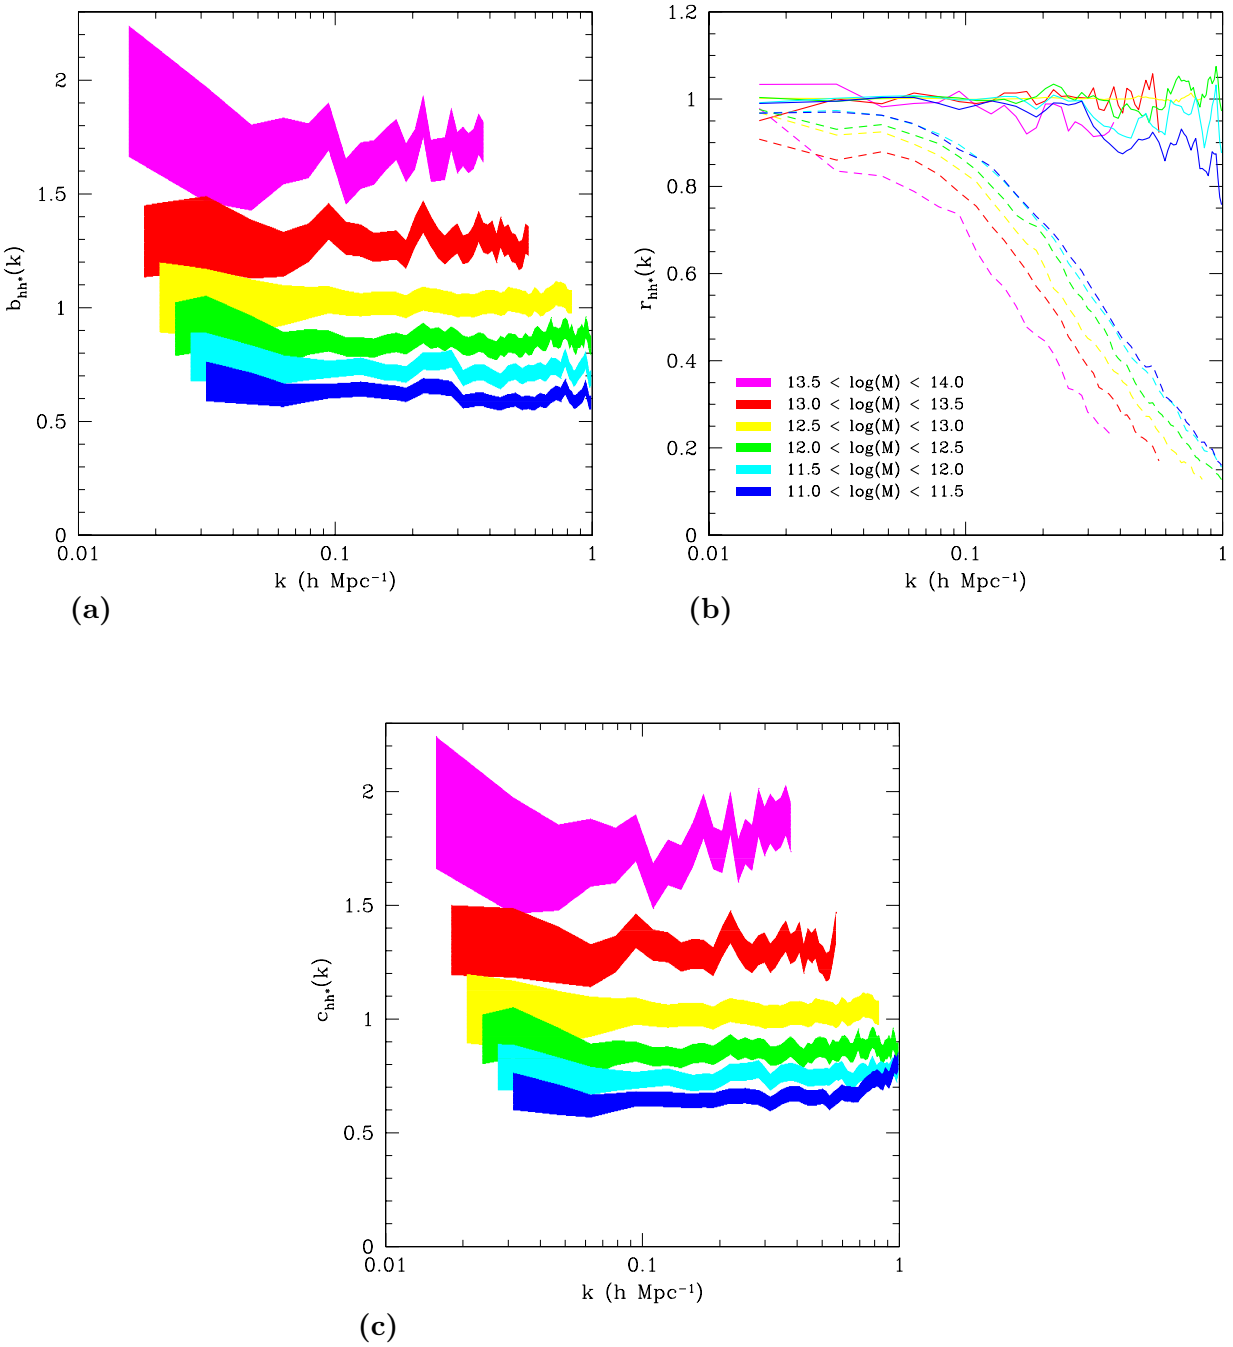

Figure 9.4 Similar to Figure 9.3, but here the halo bias is relative to halos of nonlinear mass $M^* = 4.62 \times 10^{12}\ h^{-1}M_\odot$. The halos have roughly the same spectral shape and therefore the relative bias is approximately scale-invariant. The halos are highly correlated with each other, but the stochasticity does increase as the mass difference increases.

rather a slowly varying function of scale k since halos of different masses have only approximately similar spectral shape. For the mass range $10^{11} < M/(h^{-1}M_{\odot}) < 10^{14}$, the spread in the halo spectral slope $n_{\rm eff} \equiv \mathrm{d}\log P_{hh}(k)/\mathrm{d}\log k$ is approximately 0.2 and the slope $\mathrm{d}b_{hm}/dk$ is expected to deviate more from unity as the absolute mass difference from the nonlinear mass increases. Since the nonlinear mass M^* falls approximately in the middle of the considered mass range, the deviation is small at both ends of the mass range. Figure 9.4(b) shows that the halos are more correlated with each other than with the underlying dark matter distribution, but the stochasticity increases slightly as the mass difference increases.

In Tegmark et al. (2004a), the relative galaxy bias on large scales $k < 0.1$ h/Mpc is also found to be approximately scale-invariant and consistent with the halo results. The derived shape of the matter power spectrum from the relative galaxy bias is therefore reliable, but with considerable uncertainties at large scales due to sample variance.

9.3.3 Scale-Invariant Bias Parameters

Some of the bias parameters are seen to be scale-invariant or only weakly scale-dependent over the relevant wavelength range of interest. Quantification of these parameters as a function of mass and luminosity is useful for doing cosmology with galaxy redshift surveys (Tegmark et al., 2004a; Seljak & Warren, 2004; Seljak et al., 2004).

In Figure 9.5(a) the large-scale bias parameter $\bar{b}_{hm}(M) \equiv \langle b_{hm}(k, M)\rangle$ for $k < 0.1$ h/Mpc is plotted as a function of mass and compared with an analytical prediction (Sheth & Tormen, 1999) and an empirical relation from numerical simulations (Seljak & Warren, 2004). According to Sheth & Tormen (1999), the large-scale bias is dependent on mass M but independent of scale k and is given by the expression,

$$b(\nu) = 1 + \frac{a\nu - 1}{\delta_c} + \frac{2p/\delta_c}{1 + (a\nu)^p} \tag{9.5}$$

where $\nu(M) \equiv [\delta_c/\sigma(M)]^2$ is a dimensionless mass variable, $\delta_c = 1.68$ is the critical overdensity for spherical collapse, and the coefficients $a = 0.707$ and $p = 0.3$ are fitted parameters of their modified PS mass function. Seljak & Warren (2004) used a parallel, Hashed Oct-tree (HOT) N-body code (Warren & Salmon, 1993) to run several simulations with $384^3 - 768^3$ particles and box sizes of $96 - 1152$ h^{-1}Mpc to probe the the large-scale clustering of dark matter halos. They measured halo bias relative to the linear matter density field, as opposed to the nonlinear matter density field, and derived a best-fit relation,

$$\bar{b}_0(x) = 0.53 + 0.39x^{0.42} + \frac{0.08}{40x + 1} + 10^{-4}x^{1.7}, \tag{9.6}$$

where $x \equiv M/M^*$ is mass normalized by the nonlinear mass, for a ΛCDM cosmology with $\Omega_m = 0.3$, $\Omega_b = 0.04$, $\sigma_8 = 0.9$, and $h = 0.7$. They also generalized this empirical relation by linearizing in

terms of cosmological parameters to obtain

$$\bar{b}(x) = \bar{b}_0(x) + \log(x)[0.4(\Omega_m - 0.3 + n_s - 1) + 0.3(\sigma_8 - 0.9 + h - 0.7) + 0.8\alpha_s], \tag{9.7}$$

where this correction is reasonable for $x > 0.1$, while at lower masses, a constant correction should be used by replacing the term $\log(x)$ with the constant -1.

The large-scale halo bias from the PMFAST simulations are in good agreement with the empirical relation from the HOT simulations, well within the uncertainties from sample variance, which is particularly large at large scales. Note that the halo bias in the HOT results are relative to the linear density field while the halo bias presented in this chapter are relative to the nonlinear density field. On large scales $k < 0.1\ h^{-1}$Mpc, the linear and nonlinear density fields are approximately the same and this comparison is still fair. At high masses $M > 5M^*$, both numerical results are similar to the analytical prediction while at lower masses, the analytical prediction is seen to flatten out at a higher bias value. In general, analytical predictions based on PS and extended PS formalisms are in good agreement with simulations at high masses, but the differences are seen to increase for decreasing mass. For example, the PS formalism overpredicts the abundance of low mass halos relative to that found in numerical simulations, as can be seen in Figure 9.1.

In Figure 9.5(a), the largest differences between the PMFAST and HOT derived bias are found near the nonlinear mass, though the numerical resolution near this mass scale is quite robust. In PMFAST, the nonlinear mass has a value $M^* = 4.62 \times 10^{12}\ h^{-1}M_\odot$, which is almost 2 times smaller than the HOT value of $M^* = 8.75 \times 10^{12}\ h^{-1}M_\odot$. As noted in Seljak & Warren (2004), the nonlinear mass calculated from the CMBFAST initial power spectrum can differ by up to 10% compared to the actual value in the realization and this would cause a horizontal shift in the $\bar{b}(M)$ curve by the same amount. It should also be pointed out that small differences in setting up the initial conditions can then lead to even bigger differences.

A more robust comparison, one that removes any potential offsets, is to consider the halo bias relative to halos of nonlinear mass. This relative bias is also independent of whether the linear or nonlinear density fields are used. In the PMFAST simulations, the relative bias is scale-invariant over a larger range in scale and therefore the average value can be calculated over more modes to decrease the sample variance. In Figure 9.5(b), the PMFAST relative bias are calculated using a conservative cut $k < 0.3\ h$/Mpc. The relative bias can be fitted with the expression,

$$\bar{b}_{hh^*}(x) = 0.695 + 0.305x^{0.460} + 0.085 \log x, \tag{9.8}$$

where $x \equiv M/M^*$ and the form of the fit is motivated by that used by Tegmark et al. (2004a) to fit the relative bias of SDSS galaxies. By definition, the relative bias is exactly unity for halos of nonlinear mass. The PMFAST and HOT simulations are now in very good agreement for $M > 0.5M^*$. The PMFAST bias values appear to be systematically smaller at lower masses. This difference could be real, though the actual magnitude of the difference could be exaggerated because

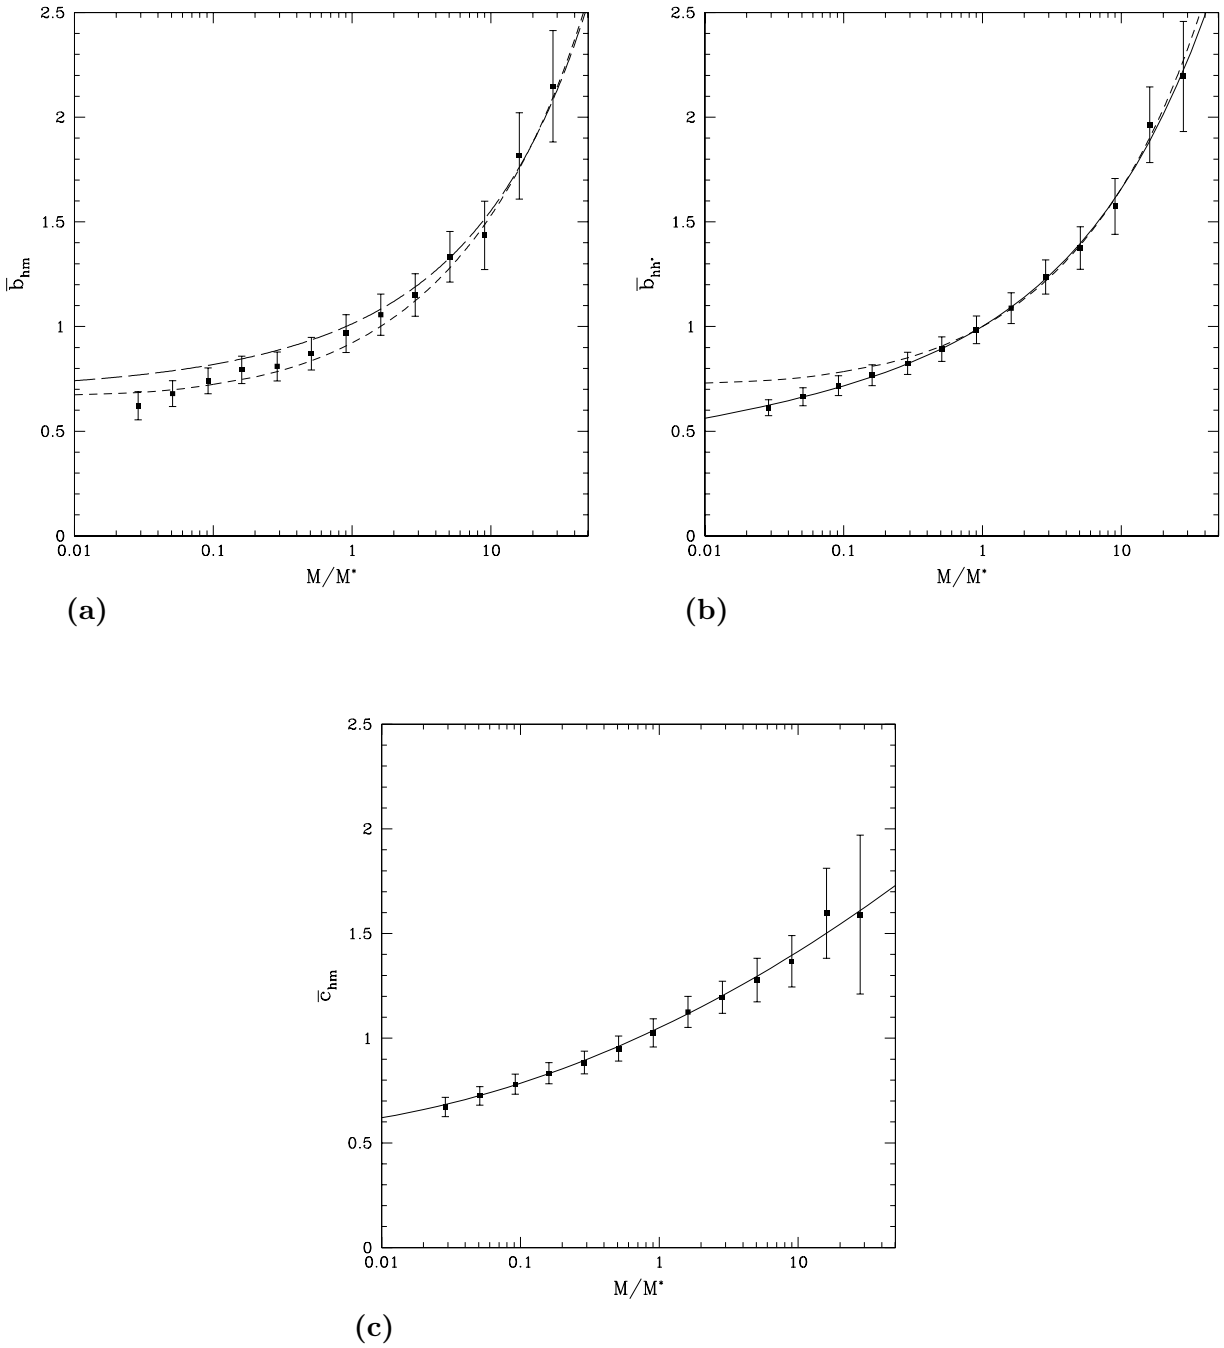

Figure 9.5 Scale-invariant bias parameters as a function of mass. In (a), the large-scale bias $\bar{b}_{hm}(M) \equiv \langle b_{hm}(k, M)\rangle$ for $k < 0.1\ h$/Mpc is compared with an analytical prediction (short dash; Sheth & Tormen, 1999) and an empirical relation from numerical simulations (long dashed; Seljak & Warren, 2004). In (b), the relative bias $\bar{b}_{hh^*}(M) \equiv \langle b_{hh}(k, M)\rangle$ for $k < 0.3\ h$/Mpc is compared with the numerical relation. In (c), the cross-bias $\bar{c}_{hm}(M) \equiv \langle c_{hm}(k, M)\rangle$ for $k < 0.3\ h$/Mpc is plotted with the best fit curve. The 1-σ uncertainties are from sample variance alone.

of the relatively poor numerical resolution in the PMFAST simulations at these mass scales. For a proper comparison at these low masses, a higher resolution simulation using a 100 h^{-1}Mpc box needs to be run.

It is especially important to determine if the bias plateau at low mass is real and if so, what is the minimum value of the bias curve. A fully self-consistent halo model, whether it be based on the PS formalism or not, should be able to correctly quantify this feature. Seljak et al. (2004) have proposed a new approach where the observed bias plateau can be compared with theoretical predictions to determine an overall normalization for converting the galaxy power spectrum into a matter power spectrum, from which one can extract cosmological parameters.

The similarity in shape and slope between the halo spectra and cross spectra combine to make the cross-bias very close to being scale invariant. In Figure 9.5, the cross-bias $\bar{c}_{hm}(M) = \langle c_{hm}(k, M)\rangle$ for $k < 0.3$ h/Mpc is remarkably well fit by the simple expression,

$$\bar{c}_{hm}(x) = 1.05 + 0.315x + 0.050x^2, \tag{9.9}$$

where $x \equiv \log(M/M^*)$. The ratio of the cross spectrum to the halo spectrum is equal for halos of mass $M \approx 0.7M^*$ and just slightly larger than unity for halos of nonlinear mass M^*. From galaxy-galaxy lensing measurements in the SDSS, Sheldon et al. (2004) found that $c_{gm}(r)$ is approximately scale-invariant over scales $0.2 - 6.7$ h^{-1}Mpc. It would be highly interesting to compare the mass-dependent halo cross-bias $c_{hm}(k, M)$ with the observed luminosity-dependent galaxy cross-bias $c_{gm}(k, L)$.

9.3.4 Constraints on the Mass-to-Light Ratio

The observed luminosity-dependent galaxy bias $\bar{b}_{gm}(L)$ from redshift surveys can be combined with the mass-dependent halo bias $\bar{b}_{hm}(M)$ measured from numerical simulations to place constraints on the mass-to-light ratio (MLR). The relative galaxy bias $b_{gg^*} = b_{gm}/b_{g^*m}$ measured from the 2dFGRS (Norberg et al., 2001) is fitted by the relation,

$$\bar{b}_{gg^*}(L) = 0.85 + 0.15\frac{L}{L^*}, \tag{9.10}$$

while the measurement from the SDSS (Tegmark et al., 2004a) is given by,

$$\bar{b}_{gg^*}(L) = 0.85 + 0.15\frac{L}{L^*} + 0.10\log\left(\frac{L}{L^*}\right). \tag{9.11}$$

These relations are very similar for galaxies brighter than L^*. For fainter galaxies, the 2dFGRS results are not as robust and the fitted relation also gives a rather high relative bias. The bias of L^* galaxies relative to the mass is found to be $b^* \equiv b_{g*m} = 1.04 \pm 0.11$ from the 2dFGRS (Verde et al., 2002) and $b^* = 0.99 \pm 0.10$ from the SDSS (Tegmark et al., 2004b), though these absolute values are in general not as robust as the relative measurements.

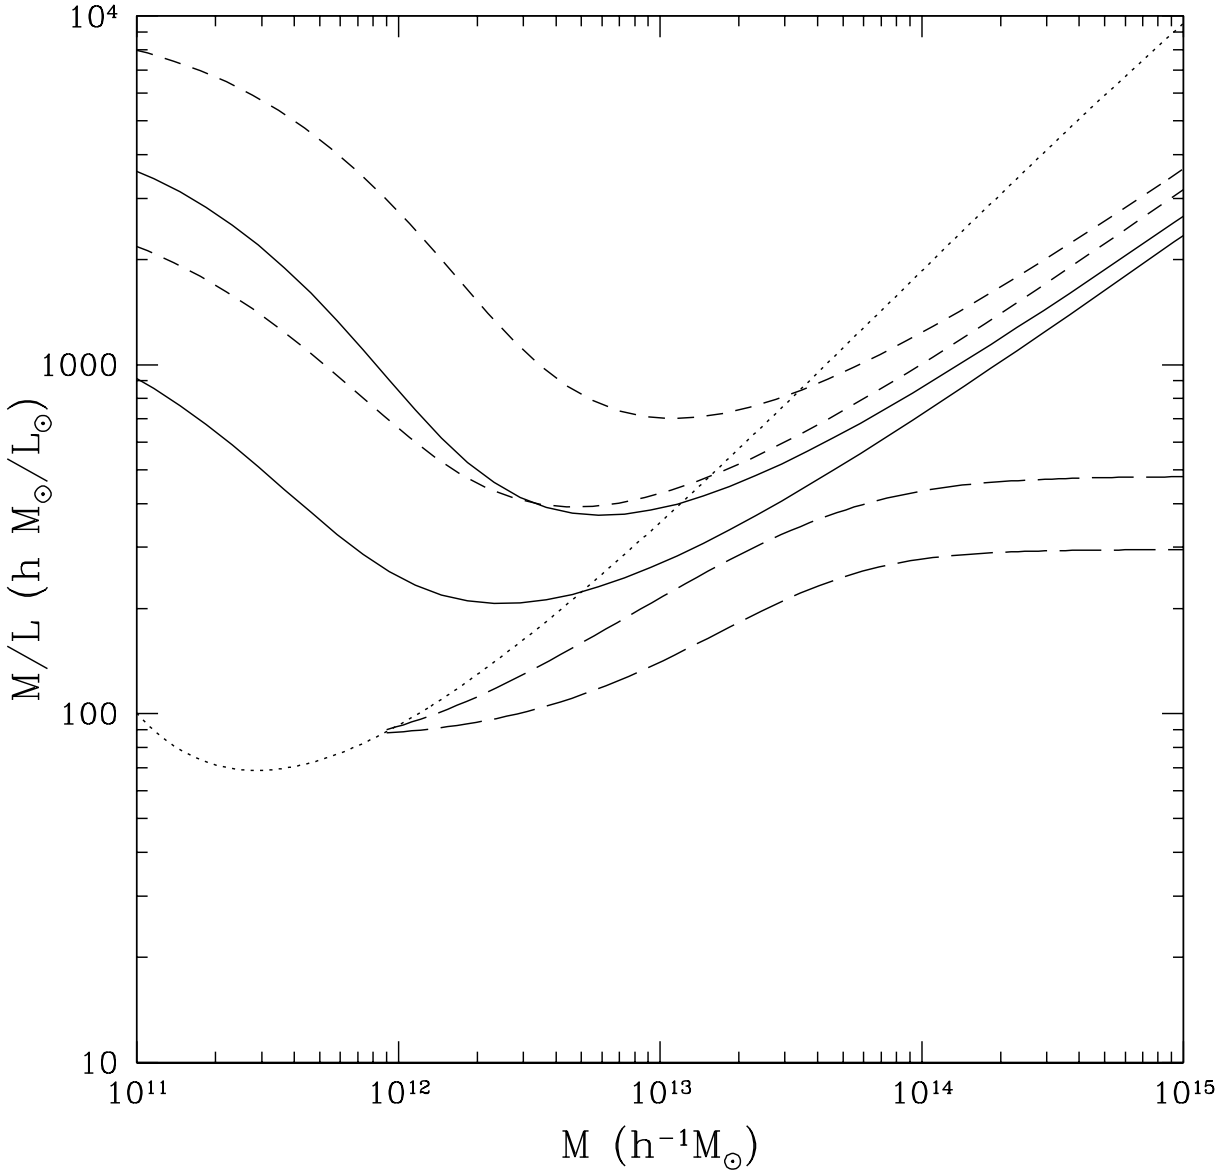

Figure 9.6 Constraints on the mass-to-light ratio from PMFAST (solid) and HOT (short dash) numerical measurements of halo bias, where the double lines represent the spread in the SDSS bias value $0.95 < b^* < 1.05$ (Tegmark et al., 2004b). For comparison, the mass-to-light ratio of galaxy groups from 2dFGRS (Eke et al., 2004) for the bJ and rF luminosity bands are given by the upper and lower long dashed lines, respectively. In addition, the dotted curve is the mass-to-light ratio obtained when the halo number density is matched to the SDSS Schechter luminosity function (Blanton et al., 2003).

Plotted in Figure 9.6 are the constraints on the mass-to-light ratio from the equality $\bar{b}_{hm}(M) = \bar{b}_{gm}(L)$. The solid lines are based on the PMFAST halo bias combined with the SDSS galaxy bias. The two lines are meant to account for some spread in the SDSS bias value $0.95 < b^* < 1.05$. Similarly, the short dashed lines are based on HOT halo bias results. Note that the SDSS bias is a function of L_r luminosities. For comparison, the mass-to-light ratio of galaxy groups from 2dFGRS (Eke et al., 2004) are given by the long dashed lines. The upper and lower lines are for the L_{bJ} and L_{rF} luminosities, respectively. The 2dFGRS mass-to-light ratios are valid for $10^{10} < L_{bJ}/(h^{-2}L_\odot) < 10^{12.5}$. In addition, the dotted curve is the mass-to-light ratio obtained when the halo number density is matched to the SDSS Schechter luminosity function (Blanton et al., 2003).

The mass-to-light ratios from the equality $\bar{b}_{hm}(M) = \bar{b}_{gm}(L)$ for both PMFAST and HOT are larger than the 2dFGRS values for galaxy groups. This is expected, particularly at higher masses, since the halos are not subdivided and only one galaxy is assigned per halo, whereas galaxy groups, by definition, have multiple galaxies. At higher masses, a more fair comparison would be to define an observed MLR where the mass is measured for groups or clusters and the luminosity from a single galaxy like the brightest cluster galaxy (BCG).

The HOT bias value is generally smaller than the PMFAST value for any given mass. For a fixed bias value then, the HOT mass is larger than the PMFAST value, resulting in a higher MLR. The HOT simulations predict a rather high MLR everywhere.

9.4 Summary

Two cosmological N-body simulations with box sizes of 200 and 400 h^{-1}Mpc, respectively, have been run with PMFAST using a 3712^3 effective mesh and 1856^3 dark matter particles. The mass-dependent clustering and biasing of dark matter halos have been measured for the mass range $10^{11} < M/(h^{-1}M_\odot) < 10^{14}$.

The halo and cross power spectra in all mass bins are well fit by power-laws in the wavenumber range $0.03 < k\ (h/\text{Mpc}) < 0.7$. The halo spectral slope increases monotonically with mass and a difference of 0.2 is found across the mass range $10^{11} < M/(h^{-1}M_\odot) < 10^{14}$. The same trend is found for the cross spectra slope, though the spread is twice as large.

On large scales $k < 0.1\ h/\text{Mpc}$, the bias parameter $b_{hm}(k)$ is consistent with scale-invariance and there is no stochasticity with $r_{hm}(k) \approx 1$. On smaller scales, the cross-correlation decreases and the power-law form for the halo spectra combined with a non power-law dark matter power spectrum leads to scale-dependent linear bias. However, the cross bias parameter $c_{hm}(k)$ is approximately scale-invariant over a larger range in scale, a reflection of the halo and cross spectra having similar power-law shapes and slopes. The halos are more correlated with each other than with the underlying dark matter. The relative bias $b_{hh^*}(k)$ is not exactly scale-invariant, but rather a slowly

varying function of scale k since halos of different masses have only approximately similar spectral shape.

The parameters of the linear stochastic biasing model that are only weakly scale-dependent have been quantified as a function of mass. These include the large-scale bias $\bar{b}_{hm}$, the relative bias b_{hh^*}, and the cross bias $\bar{c}_{hm}$. The large-scale bias curve $\bar{b}_{hm}(M)$ is consistent with the empirical relation from Seljak & Warren (2004), though the existence and value of the minimum bias plateau cannot be confirmed with the present set of simulations. For future work, two more simulations with box sizes of 100 and 800 h^{-1}Mpc, respectively, will be run to extend the mass range probed.

Chapter 10

Conclusions and Future Work

10.1 Conclusions

In this thesis, I have presented several new numerical techniques for cosmological and astrophysical simulations of structure formation. These notably include a moving frame algorithm for capturing hydrodynamics at high Mach numbers and an out-of-core hydrodynamic code for high resolution cosmological simulations. Three ongoing scientific applications of Eulerian hydrodynamic and N-body simulations have also been presented. In addition, the primer on Eulerian computation fluid dynamics (CFD) presented in Chapter 2 has also been published in Trac & Pen (2003) and this pedagogical review, aimed at astrophysics, is the first of its kind to be published in the astrophysical literature.

The moving frame hydrodynamic algorithm was designed to solve the high Mach number problem of Eulerian CFD. Standard Eulerian schemes that strictly conserve total energy suffer from this problem and proposed solutions to additionally solve the entropy or thermal energy equation still have their limitations. In the moving frame approach, I explicitly decompose the Euler equations into a nonlinear, local flow, where local fluid variables can be directly calculated, and a smooth background component, that by construction can approximate the bulk flow. This decomposition allows one to solve hydrodynamics in the appropriate reference frame where the ratio of the thermal energy to the kinetic energy is maximized, allowing thermodynamic variables to be accurately calculated. The moving frame algorithm is an improvement over all previous Eulerian implementations since it has the unique ability to resolve shocks and prevent spurious heating where both the preshock and postshock fluid are supersonic. The moving frame algorithm and the cosmological Hydro&N-body code based on it has been published in Trac & Pen (2004a).

The out-of-core hydro (OCH) code is based on the novel approach of using disk space as virtual memory in order to run large simulations that require more memory than is available or affordable. This computational technique has been relatively unexplored because previous numerical algorithms cannot overcome slow disk bandwidth and latency. Nowadays, disk arrays can be

striped to significantly improve bandwidth but latency still poses a major problem, particularly for algorithms that often access data non-sequentially. The OCH code has minimal I/O overhead for two main reasons. First, the two-level mesh scheme avoids having to do non-sequential disk access for all data transposes are done in memory. Second, I/O is mostly hidden for disk operations are done concurrently with numerical calculations. The OCH cosmological code, along with a new out-of-core initial conditions generator, is collected in Trac & Pen (2004b) and will be a unique and interesting addition to the literature.

Cosmological simulations with up to 4000^3 grid cells and 2000^3 particles can be run with the OCH code on a shared memory machine with 32 processors, 64 GB of memory, and 3 TB of striped disk space found here at the Canadian Institute for Astrophysics (CITA). However, such a simulation would require months of time because of the limited processing power of this computer. Cosmological simulations with 2000^3 grid cells and 1000^3 particles are much more feasible in practice. In time, memory may become cheaper, the typical shared memory machine may have several hundred gigabytes worth, and large simulations can be run in the standard approach. However, the out-of-core approach can always be used to run even larger simulations. In fact, as the memory in a system increases, the OCH code becomes more efficient. The domain decomposition can be accomplished with larger blocks and since the buffering length is fixed, this reduces the relative buffering overhead.

The OCH code has been applied to running the largest Eulerian hydrodynamic simulation to date, with 1536^3 grid cells and 768^3 dark matter particles, for the purpose of modelling the thermal evolution of the high redshift $3 \leq z \leq 7$ IGM. The primary focus of this work is to measure both the shape and the scatter of the temperature-density relation. Prior to this study, a power-law relation for gas densities near the mean baryonic density is often assumed in the literature, in particular for studies of the Lyman alpha forest. However, this assumption is only an approximation since shock heating is ignored. I have modelled the thermal evolution of the IGM in both an adiabatic scenario, where shock heating dominates, and a nonadiabatic one, where reionization, UV heating, and atomic cooling are also included.

UV heating and atomic cooling, in the absence of reionization and shock heating, is demonstrated to result in a temperature-density that is appropriately described by a double power-law, where the turnover unfortunately occurs in the density range $1 < \Delta_b \equiv (1+\delta_b) < 10$ that is relevant to studies of the Lyman alpha forest. This double power-law relation is a generic prediction based on efficient net heating at lower densities and efficient net cooling at higher densities relative to the turnover. When reionization and shock heating are included, the former decreases the slope of the temperature-density relation for densities lower than the turnoff while the latter increases the slope for densities higher than the turnover. This picture of the temperature-density relation and in addition, the discovery that shock heating is important at densities $\Delta_b \sim 3$ by redshift $z = 3$ has not been previously reported in the literature.

In principle, the thermal history of the IGM is imprinted in both the shape and scatter of the temperature-density relation. Most work to date have focussed on studying and measuring the shape, but the scatter has been predominantly ignored. In this thesis, I have measured the scatter and also examined the shape of the scatter distribution and showed how they reflect the underlying physics of the IGM. In addition to showing that shock heating changes the shape of the temperature-density relation even at densities as low as several times the mean baryonic density, I have also definitively demonstrated that shock heating leaves a unique and quantifiable signature on the shape of the scatter distribution.

A new parallel N-body code named PMFAST has been applied to running two very high mass resolution simulations for measuring the mass-dependent clustering and biasing of dark matter halos. Galaxy and halo bias have been studied for close to two decades now, but they have been done predominantly under the assumption that galaxies, halos, and the underlying dark matter are perfectly correlated. In addition, the bias on large-scales has been shown analytically to be scale-invariant and while this prediction is often used to do cosmology with galaxy clustering, it has not been conclusively tested with numerical simulations. In this thesis, I have measured the parameters of the linear stochastic biasing model as a function of both scale and mass.

It is known that a power-law form for the halo spectra combined with a non power-law form for the dark matter power spectrum produces scale-dependent bias. However, this work has shown that the effective slopes of the halo spectra are only weakly scale and mass dependent and thus, modelling the relative bias as scale-invariant is a good approximation considering the current uncertainties in the measurement of galaxy clustering. On large scales, the average bias and cross bias as functions of mass have been quantified and these statistics can be combined with results from galaxy clustering and weak lensing to constrain the mass-to-light ratio and to recover the matter power spectrum.

The astrophysical application of Eulerian hydrodynamic simulations to model the formation of blue stragglers through stellar mergers is the first to be reported in the literature (Trac & Pen, 2003). A three-dimensional Eulerian TVD hydro code, which uniquely includes a particle-mesh algorithm for tracking test particles, has been used to simulate the off-axis collision of equal mass $M_0 = 0.8M_\odot$ main sequence stars. The parent stars, being near the turnoff and having depleted hydrogen cores, have been modelled using $n = 3$ polytropes and with realistic calculations. Prior to this study, all mergers of this nature were simulated using Lagrangian SPH codes.

In the Eulerian simulations, the particle-mesh algorithm is used to track test particles of fixed mass and known chemical composition. The initial and final mass coordinates for the particles have been plotted to directly show how mass is mixed during the merging process. One can naturally do this with SPH particles as well, but no such mixing plots have been published in the literature to date. The Eulerian simulations show significant mixing and the remnant is estimated to be $\sim 50\%$ fully mixed on average, in contrast to the weak mixing found in SPH simulations. It is possible that since weak shocks are not accurately captured in SPH simulations, the particles may experience

sedimentation and undermix. The Eulerian results also definitively show that it is harder to mix hydrogen from surface layers into a dense core with sharp gradients. The parent main sequence stars do have sharp gradients and in the hydrogen core of the remnant, only $\sim 22\%$ is mixed in from outside and the hydrogen content has only increased by 23%. In contrast, the mass profiles of the $n = 3$ polytropes are smoother and for the resulting remnant, 44% is mixed in from outside and the hydrogen abundance has increased by 59%. These Eulerian mixing values will be compared with SPH results in future work.

10.2 Future Work

In Trac, Sills, & Pen (2004), the Eulerian TVD simulations will be compared with Lagrangian SPH simulations to determine how these different approaches handle shocks and mixing. The plots comparing initial and final mass coordinates are particularly useful for quantifying mixing, but no such plots have been made for SPH particles. The merger remnant will be run through the Yale stellar evolution code to determine its evolutionary track in the Hertzsprung-Russell (HR) diagram and to calculate its lifetime. In addition to the mixing problem, another long standing problem in the formation of blue stragglers from stellar mergers is that the remnants appear to be rotating too fast. In principle, a remnant with a surface convection zone can experience angular momentum loss through a magnetic wind (see Sills et al., 2002). In order to detect the presence or lack of a surface convection zone, high mass resolution is required to simulate the outer few percent of the mass of the remnant and this is ideally done with Eulerian codes. In addition to the off-axis collision of equal mass $M = 0.8M_{\odot}$ main sequence stars, we will also look at collisions between $M = 0.4M_{\odot}$ stars to determine how different masses and stellar profiles affect the structure and composition of the merger remnant.

PMFAST has been used to run two cosmological N-body simulations with box sizes of 200 and 400 h^{-1}Mpc and two more simulations with box sizes of 100 and 800 h^{-1}Mpc will be run to extend the scale and mass ranges in order to definitively address some of the findings in this thesis. The larger volume simulation is needed to measure bias at larger scales and to find halos in the mass range $M > 10^{14}h^{-1}M_{\odot}$, which have relatively low number densities. The smaller box simulation is needed in order to measure the clustering of halos in the mass range $10^{10} < M/(h^{-1}M_{\odot}) < 10^{12}$, where there is some indication that the large-scale bias is independent of both scale and mass. In addition, the stochasticity curves $r(k, M)$ appear to approach a minimum curve at low masses and it would be interesting to determine if such a curve exist, what does it look like, and why is there one. Halo bias will also be measured at other redshifts in order to study the evolution of the large-scale structure. This will compliment the wealth of observation data from the wide synoptic portion of the CFHT Legacy Survey, which has four 6×6 degree fields with galaxies out to redshift $z \sim 1$.

The thermal evolution of the high redshift IGM in a nonadiabatic universe where reionization, UV heating, and atomic cooling are included is model dependent, unlike the adiabatic scenario. An IGM modelled using uniform reionization, uniform UV background, and photoionization equilibrium is qualitatively consistent with the observed Lyman alpha forest (see Rauch, 1998, for a review), but in order to do precision cosmology with the latter, a more detailed treatment is required. Reionization may not be singular and it would interesting to impose separate events using the same prescription described in this thesis. Reionization can be singular but gradual in time over the redshift range $7 < z < 20$ and most likely then to be nonuniform. In this case, semi-analytical star formation and supernova feedback (e.g. Springel & Hernquist, 2003) can be implemented and calibrated against Lyman alpha forest flux statistics. Note that this method can also be used to set the UV background, which will be nonuniform in general. These various scenarios will be simulated with the cosmological Hydro&N-body code and the OCH code.

The Beowulf clusters are becoming the dominant high performance computing (HPC) systems in computational astrophysics because of their unrivalled speeds that can easily reach into the teraflops. These systems remain memory limited, but simulations can be run out-of-core using disk space as virtual memory to gain a few more orders of magnitude in capacity. IDE hard drives are relatively cheap and tens to hundreds of terabytes of disk space can be added at little cost relative to the cheap price of the clusters themselves. Hardware or software raiding can be used to prevent the loss of data in the expected event of disk failures. The current version of the out-of-core code is implemented using a two-level mesh decomposition and parallelized using OpenMP directives for shared memory machines. For the distributed memory clusters, the implementation will be based on a multi-level mesh or multigrid scheme with parallelization across nodes using Message Passing Interface (MPI) libraries. The multi-level mesh is necessary for domain decomposition and will minimize communication costs. We are well within reach of being able to run hydrodynamic simulations with a trillion grid cells and particles.

References

Anninos, W. Y., & Norman, M. L. 1994, ApJ, 429, 434

Bailyn, C. D. 1995, ARAA, 33, 133

Bardeen, J. M., Bond, J. R., Kaiser, N., & Szalay, A. S. 1986, ApJ, 304, 15

Becker, R. H., Fan, X., White, R. L., Strauss, M. A., Narayanan, V. K., Lupton, R. H., Gunn, J. E., & Annis, J. 2001, AJ, 122, 2850

Benz, W., & Hills, J. G. 1987, ApJ, 323, 614

Blanton, M., Cen, R., Ostriker, J. P., Strauss, M. A., & Tegmark, M. 2000, ApJ, 531, 1

Blanton, M. R., Hogg, D. W., Bahcall, N. A., Brinkmann, J., Britton, M., Connolly, A. J., Csabai, I., & Fukugita, M. 2003, ApJ, 592, 819

Bryan, G. L., Norman, M. L., Stone, J. M., Cen, R., & Ostriker, J. P. 1995, Comput. Phys. Commun., 89, 149

Cen, R. 1992, ApJS, 78, 341

Cen, R., Miralda-Escude, J., Ostriker, J. P., & Rauch, M. 1994, ApJL, 437, L9

Cen, R., & Ostriker, J. P. 1992, ApJL, 399, L113

Colella, P., & Woodward, P. R. 1984, J. Comput. Phys., 54, 174

Couchman, H. M. P. 1991, ApJL, 368, L23

Courant, R., Isaacson, E., & Reeves, M. 1952, Commun. Pure Appl. Math., 5, 243

Croft, R. A. C., Weinberg, D. H., Bolte, M., Burles, S., Hernquist, L., Katz, N., Kirkman, D., & Tytler, D. 2002, ApJ, 581, 20

Dekel, A., & Lahav, O. 1999, ApJ, 520, 24

Efstathiou, G., Davis, M., White, S. D. M., & Frenk, C. S. 1985, ApJS, 57, 241

Eke, V. R., Frenk, C. S., Baugh, C. M., Cole, S., Norberg, P., Peacock, J. A., Baldry, I. K., & Bland-Hawthorn, J. 2004, ArXiv Astrophysics e-prints

Fang, T., & White, M. 2004, ApJL, 606, L9

Fry, J. N., & Gaztanaga, E. 1993, ApJ, 413, 447

Gingold, R. A., & Monahan, J. J. 1977, MNRAS, 181, 375

Gnedin, N. Y. 1995, ApJ, 97, 231

Gnedin, N. Y., & Hamilton, A. J. S. 2002, MNRAS, 334, 107

Gnedin, N. Y., & Hui, L. 1998, MNRAS, 296, 44

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

Guenther, D. B., Demarque, P., Pinsonneault, M. H., & Kim, Y.-C. 1992, ApJ, 392, 328

Haardt, F., & Madau, P. 1996, ApJ, 461, 20

Hamilton, A. J. S. 1998, in ASSL Vol. 231: The Evolving Universe, 185

Harten, A. 1983, J. Comp. Phys., 49, 357

Hernquist, L., & Katz, N. 1989, ApJS, 70, 419

Hernquist, L., Katz, N., Weinberg, D. H., & Jordi, M. 1996, ApJL, 457, L51

Heyl, J. S., Cole, S., Frenk, C. S., & Navarro, J. F. 1995, MNRAS, 274, 755

Hirsch, C. 1990, Numerical computation of internal and external flows (New York: Wiley, 1990)

Hockney, R. W., & Eastwood, J. W. 1988, Computer simulation using particles (Bristol: Hilger, 1988)

Hui, L., & Gnedin, N. Y. 1997, MNRAS, 292, 27

Jenkins, A., Frenk, C. S., White, S. D. M., Colberg, J. M., Cole, S., Evrard, A. E., Couchman, H. M. P., & Yoshida, N. 2001, MNRAS, 321, 372

Jin, S., & Xin, Z. 1985, Comm. Pure and Applied Math, 48, 235

Juszkiewicz, R., Weinberg, D. H., Amsterdamski, P., Chodorowski, M., & Bouchet, F. 1995, ApJ, 442, 39

Kaiser, N. 1984, ApJL, 284, L9

—. 1987, MNRAS, 227, 1

Katz, N., Weinberg, D. H., & Hernquist, L. 1996, ApJS, 105, 19

Kogut, A., Spergel, D. N., Barnes, C., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., & Limon, M. 2003, ApJS, 148, 161

Lacey, C., & Cole, S. 1994, MNRAS, 271, 676

Landau, L. D., & Lifshitz, E. M. 1987, Fluid Mechanics, 2nd ed. (Oxford: Pergamon, 1987)

Laney, C. 1998, Computational gas dynamics (Cambridge: Cambridge Univ. Press, 1998)

Lax, P. D., & Wendroff, B. 1960, Commun. Pure Appl. Math., 10, 537

Lucy, L. B. 1977, AJ, 82, 1013

McDonald, P., Miralda-Escudé, J., Rauch, M., Sargent, W. L. W., Barlow, T. A., & Cen, R. 2001, ApJ, 562, 52

McDonald, P., Seljak, U., Burles, S., Schlegel, D. J., Weinberg, D. H., Shih, D., Schaye, J., & Schneider, D. P. 2004a, ArXiv Astrophysics e-prints

McDonald, P., Seljak, U., Cen, R., Weinberg, D. H., Burles, S., Schneider, D. P., Schlegel, D. J., & Bahcall, N. A. 2004b, ArXiv Astrophysics e-prints

Merz, H., Pen, U., & Trac, H. 2004, submitted to New Astronomy, astro-ph/0402443

Miralda-Escude, J., Cen, R., Ostriker, J. P., & Rauch, M. 1996, ApJ, 471, 582

Norberg, P., Baugh, C. M., Hawkins, E., Maddox, S., Peacock, J. A., Cole, S., Frenk, C. S., & Bland-Hawthorn, J. 2001, MNRAS, 328, 64

Peacock, J. A., & Dodds, S. J. 1996, MNRAS, 280, L19

Pen, U. 1997, ApJL, 490, L127

—. 1998a, ApJS, 115, 19

—. 1998b, ApJ, 504, 601

Press, W. H., & Schechter, P. 1974, ApJ, 187, 425

Rauch, M. 1998, ARAA, 36, 267

Rauch, M., Miralda-Escude, J., Sargent, W. L. W., Barlow, T. A., Weinberg, D. H., Hernquist, L., Katz, N., Cen, R., & Ostriker, J. P. 1997, ApJ, 489, 7

Roe, P. L. 1985, in Large-Scale Computations in Fluid Mechanics, Lectures in Applied Mathematics

Ryu, D., Ostriker, J. P., Kang, H., & Cen, R. 1993, ApJ, 414, 1

Salmon, J., & Warren, M. 1997, in Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing

Sandquist, E. L., Bolte, M., & Hernquist, L. 1997, ApJ, 477, 335

Seljak, U. 2000, MNRAS, 318, 203

Seljak, U., Makarov, A., McDonald, P., Anderson, S., Bahcall, N., Brinkmann, J., Burles, S., & Cen, R. 2004, ArXiv Astrophysics e-prints

Seljak, U., & Warren, M. S. 2004, ArXiv Astrophysics e-prints

Seljak, U., & Zaldarriaga, M. 1996, ApJ, 469, 437

Shapiro, P. R., Martel, H., Villumsen, J. V., & Owen, J. M. 1996, ApJS, 103, 269

Sheldon, E. S., Johnston, D. E., Frieman, J. A., Scranton, R., McKay, T. A., Connolly, A. J., Budavári, T., & Zehavi, I. 2004, AJ, 127, 2544

Sheth, R. K., & Tormen, G. 1999, MNRAS, 308, 119

Shu, F. H. 1992, Gas Dynamics (Mill Valley: University Science Books)

Sills, A., Adams, T., Davies, M. B., & Bate, M. R. 2002, MNRAS, 332, 49

Sills, A., Faber, J. A., Lombardi, J. C., Rasio, F. A., & Warren, A. R. 2001, ApJ, 548, 323

Sills, A., & Lombardi, J. C. 1997, ApJL, 484, L51

Sills, A., Lombardi, J. C., Bailyn, C. D., Demarque, P., Rasio, F. A., & Shapiro, S. L. 1997, ApJ, 487, 290

Smith, R. E., Peacock, J. A., Jenkins, A., White, S. D. M., Frenk, C. S., Pearce, F. R., Thomas, P. A., Efstathiou, G., & Couchman, H. M. P. 2003, MNRAS, 341, 1311

Sod, G. A. 1978, J. Comput. Phys., 27, 1

Spergel, D. N., Verde, L., Peiris, H. V., Komatsu, E., Nolta, M. R., Bennett, C. L., Halpern, M., & Hinshaw, G. 2003, ApJS, 148, 175

Springel, V., & Hernquist, L. 2003, MNRAS, 339, 289

Stone, J. M., & Norman, M. L. 1992, ApJS, 80, 753

Strang, G. 1968, SIAM J. Num. Anal., 5, 506

Strauss, M. A., & Willick, J. A. 1995, Phys. Rep., 261, 271

Sunyaev, R. A., & Zeldovich, Y. B. 1980, MNRAS, 190, 413

Tasitsiomi, A., Kravtsov, A. V., Wechsler, R. H., & Primack, J. R. 2004, ArXiv Astrophysics e-prints

Tegmark, M., Blanton, M. R., Strauss, M. A., Hoyle, F., Schlegel, D., Scoccimarro, R., Vogeley, M. S., & Weinberg, D. H. 2004a, ApJ, 606, 702

Tegmark, M., & Bromley, B. C. 1999, ApJL, 518, L69

Tegmark, M., & Peebles, P. J. E. 1998, ApJL, 500, L79

Tegmark, M., Strauss, M. A., Blanton, M. R., Abazajian, K., Dodelson, S., Sandvik, H., Wang, X., & Weinberg, D. H. 2004b, Phys. Rev. D, 69, 103501

Theuns, T., Leonard, A., Efstathiou, G., Pearce, F. R., & Thomas, P. A. 1998, MNRAS, 301, 478

Trac, H., & Pen, U. 2003, PASP, 115, 303

—. 2004a, New Astronomy, 9, 443

—. 2004b, submitted to New Astronomy, astro-ph/0402444

Trac, H., Sills, A., & Pen, U. 2004, to be submitted to MNRAS

Van Leer, B. 1974, J. Comput. Phys., 14, 361

Verde, L., Heavens, A. F., Percival, W. J., Matarrese, S., Baugh, C. M., Bland-Hawthorn, J., Bridges, T., & Cannon, R. 2002, MNRAS, 335, 432

Warren, M. S., & Salmon, J. K. 1993, in Supercomputing, 12–21

Zaldarriaga, M., Seljak, U., & Hui, L. 2001, ApJ, 551, 48

Zeldovich, Y. B. 1970, A&A, 5, 84

Zhang, P., Pen, U., & Trac, H. 2004a, MNRAS, 347, 1224

—. 2004b, ArXiv Astrophysics e-prints

Zhang, Y., Anninos, P., & Norman, M. L. 1995, Bulletin of the American Astronomical Society, 27, 1412

Appendix A

Stellar Mergers

A.1 Off-axis Collision of Main Sequence Stars

Two equal mass $M = 0.8M_{\odot}$ main sequence stars are placed on zero-energy parabolic orbits with a pericenter separation of $0.25R$. Initially, they are separated by a distance of $\Delta d = 3.75R = 3.58R_{\odot}$ and each are moving at a speed of $\sim$ 150 km/s relative to the center of mass of the collision. Snapshots of the merging process from the high resolution simulation are shown in Figure A.1. The density contour plots are from thin slices 4 cells thick through the $x-y$ orbital mid-plane. The contours are spaced logarithmically with two per decade, starting from a minimum density of 0.01 g/cm^3 at the edges. The off-axis collision does indeed produce a single remnant star that establishes hydrostatic equilibrium in less than $10\tau_{dyn}$. During the merging process, the outer envelopes of the parent stars are shock heated and approximately 10% of the total mass gets ejected into the interstellar medium (ISM).

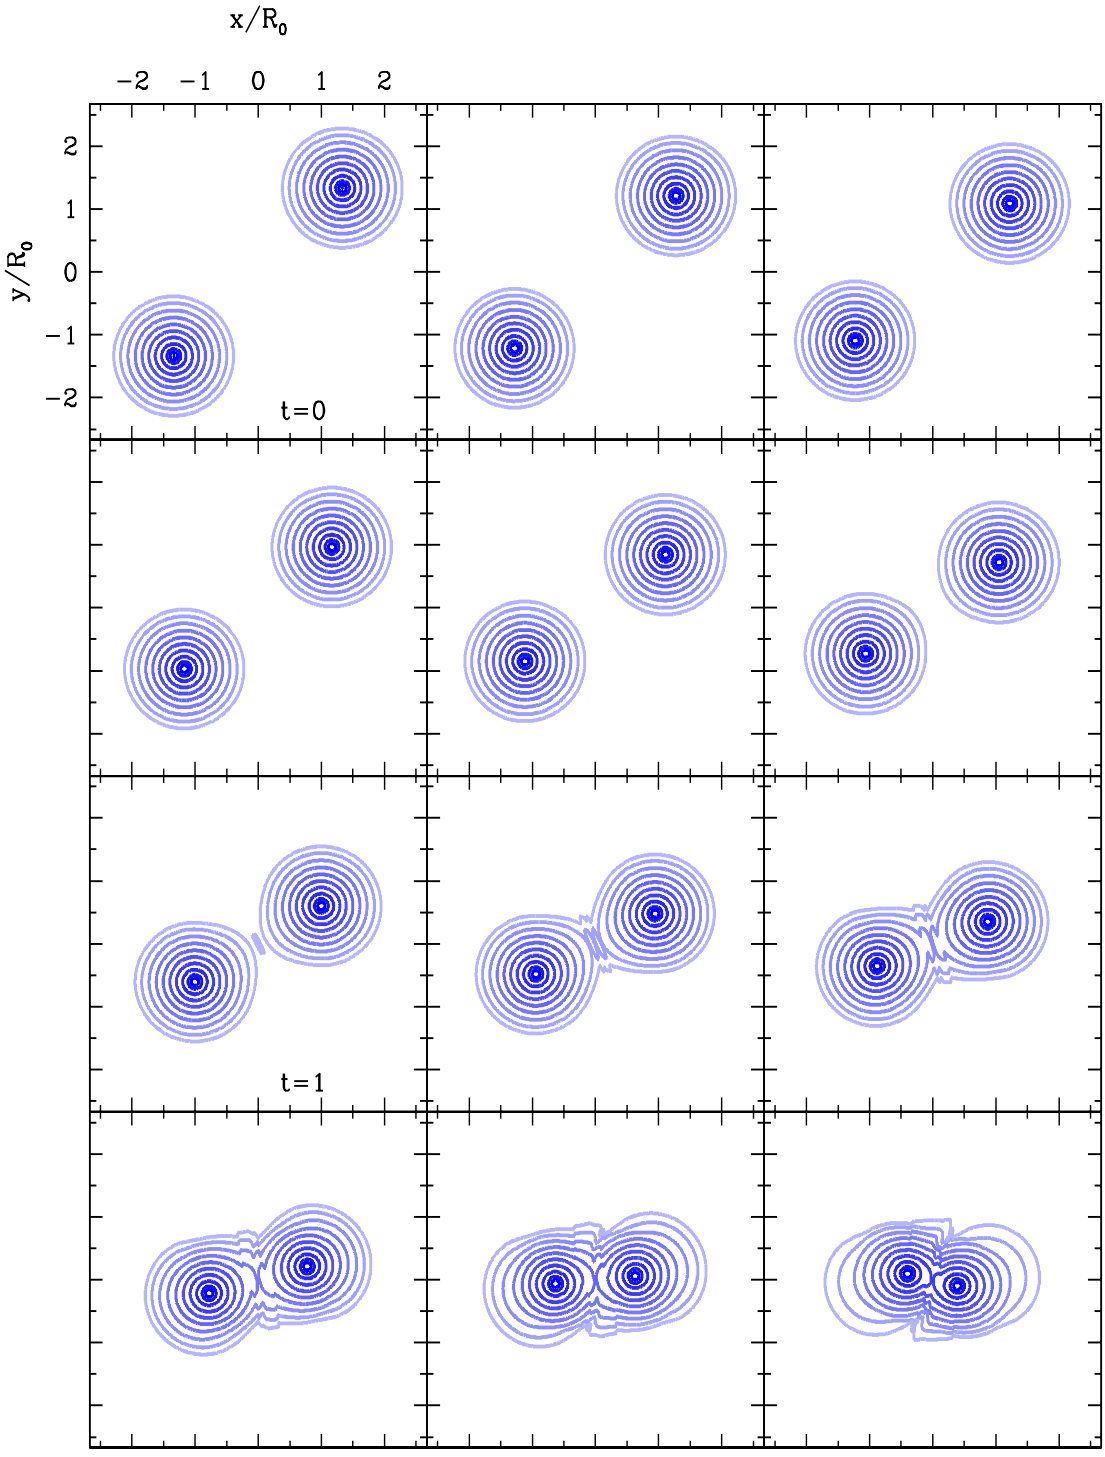
x/R0
−2 −1 0 1 2
y/R0
2
1
0
−1
−2
t=0
t=1

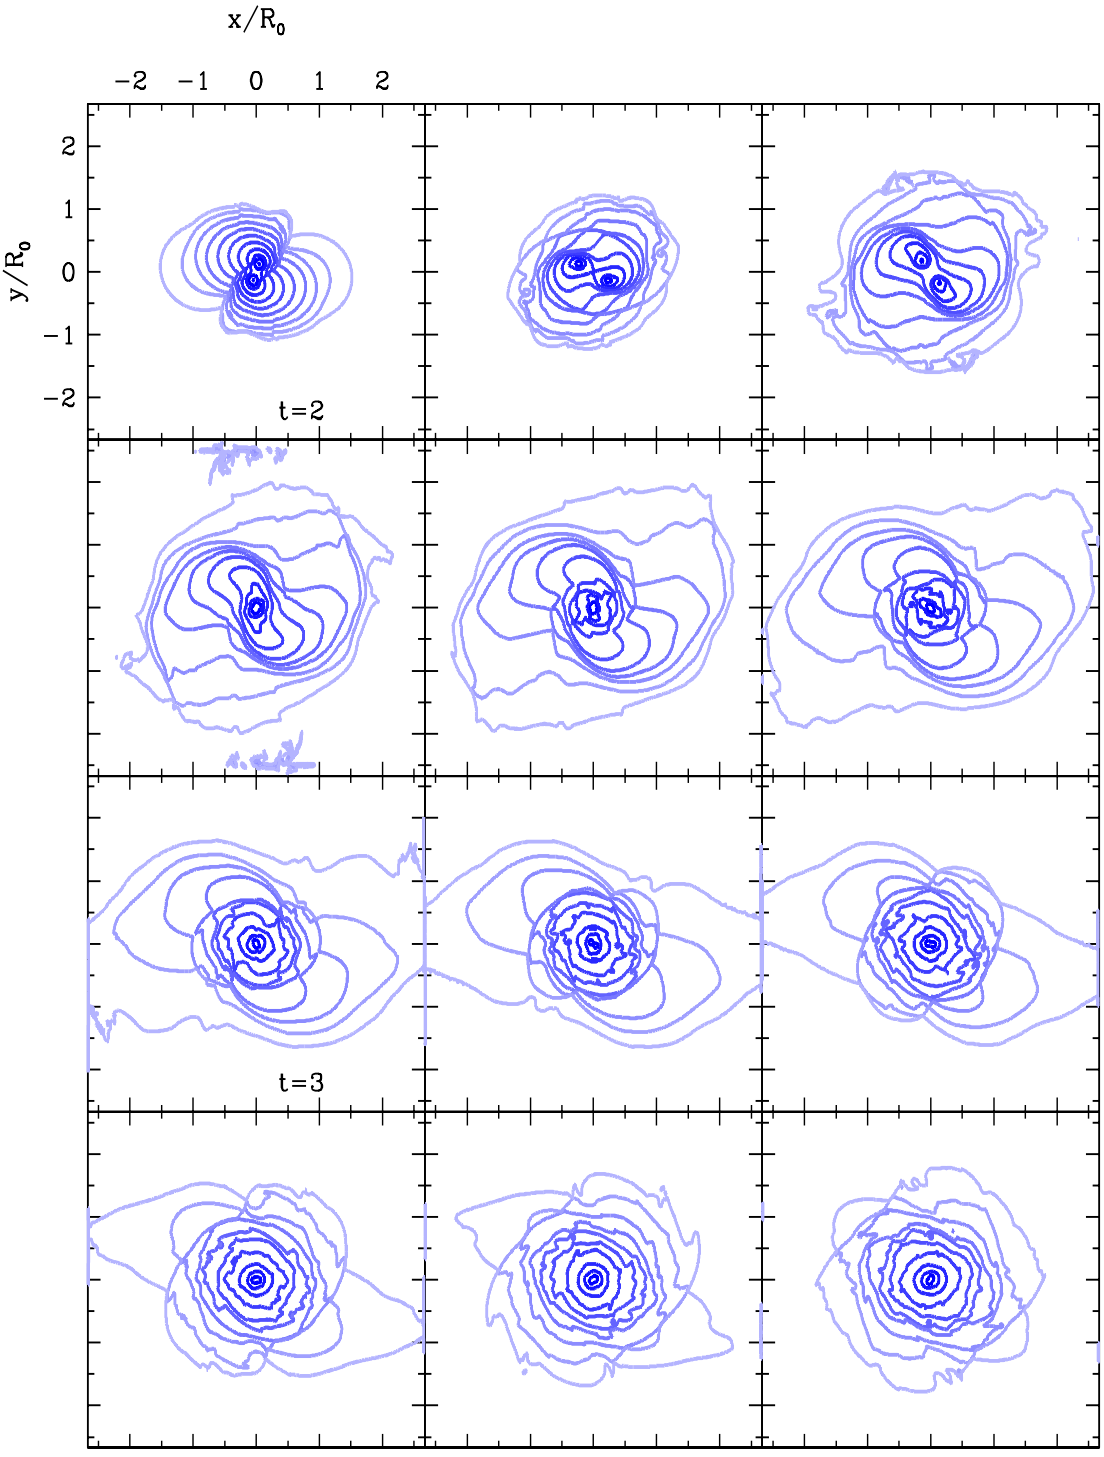
x/R0
−2 −1 0 1 2
y/R0
2
1
0
−1
−2
t=2
t=3

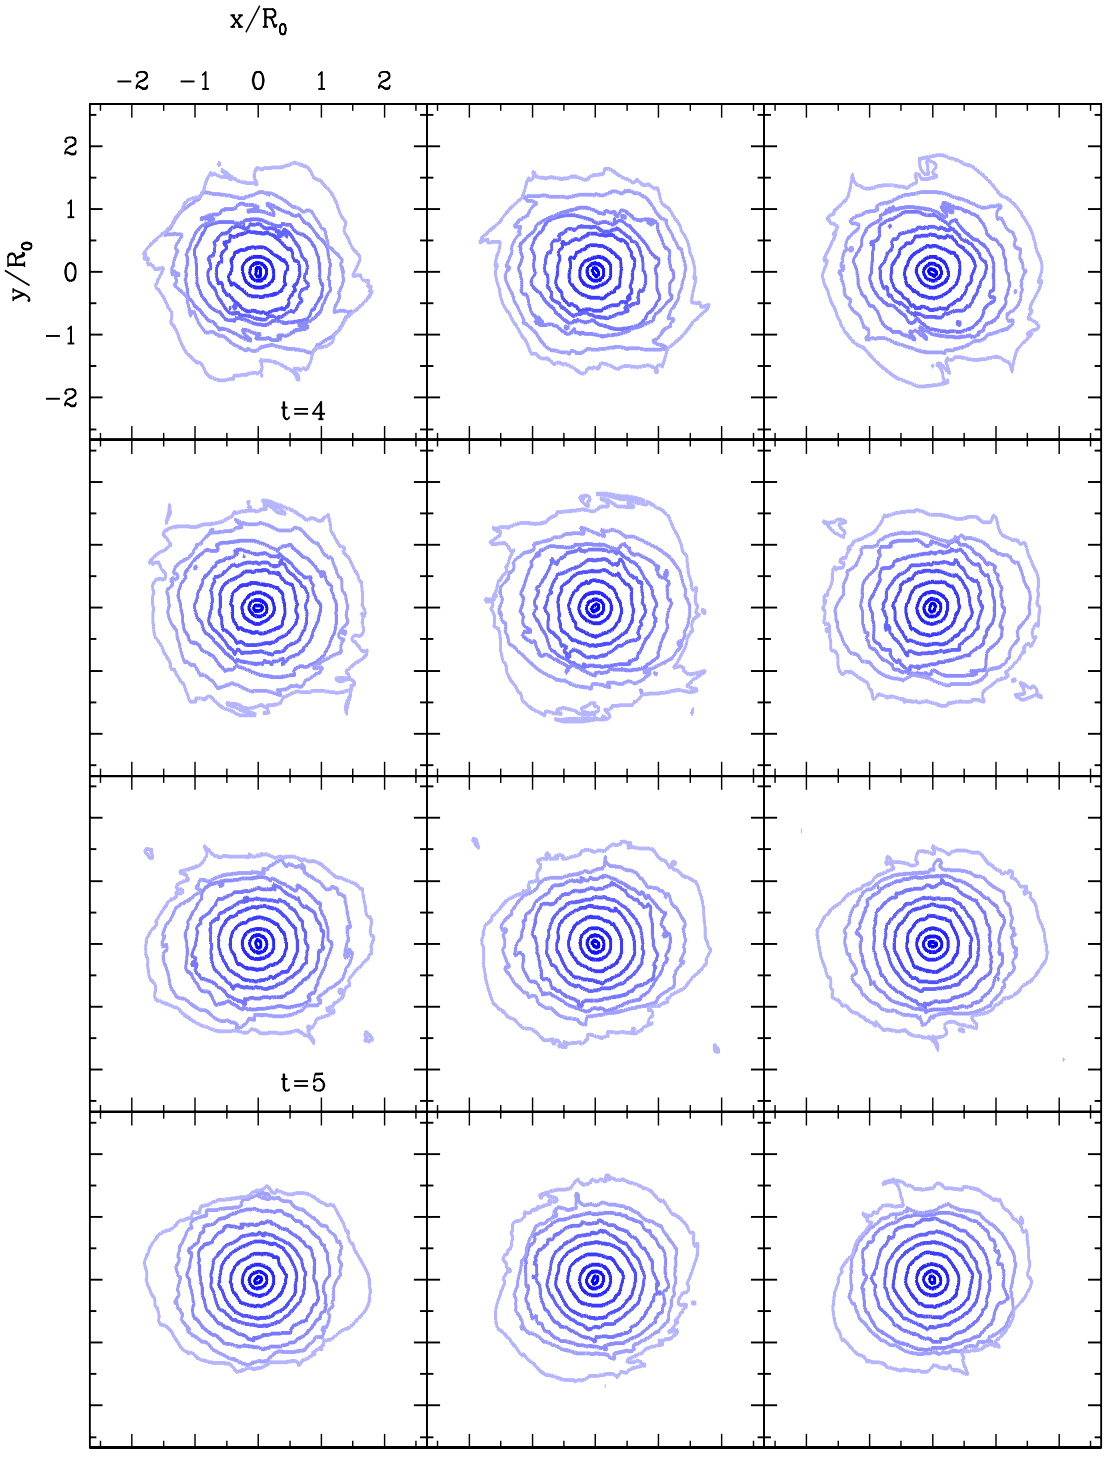
x/R0
−2 −1 0 1 2
y/R0
2
1
0
−1
−2
t=4
t=5

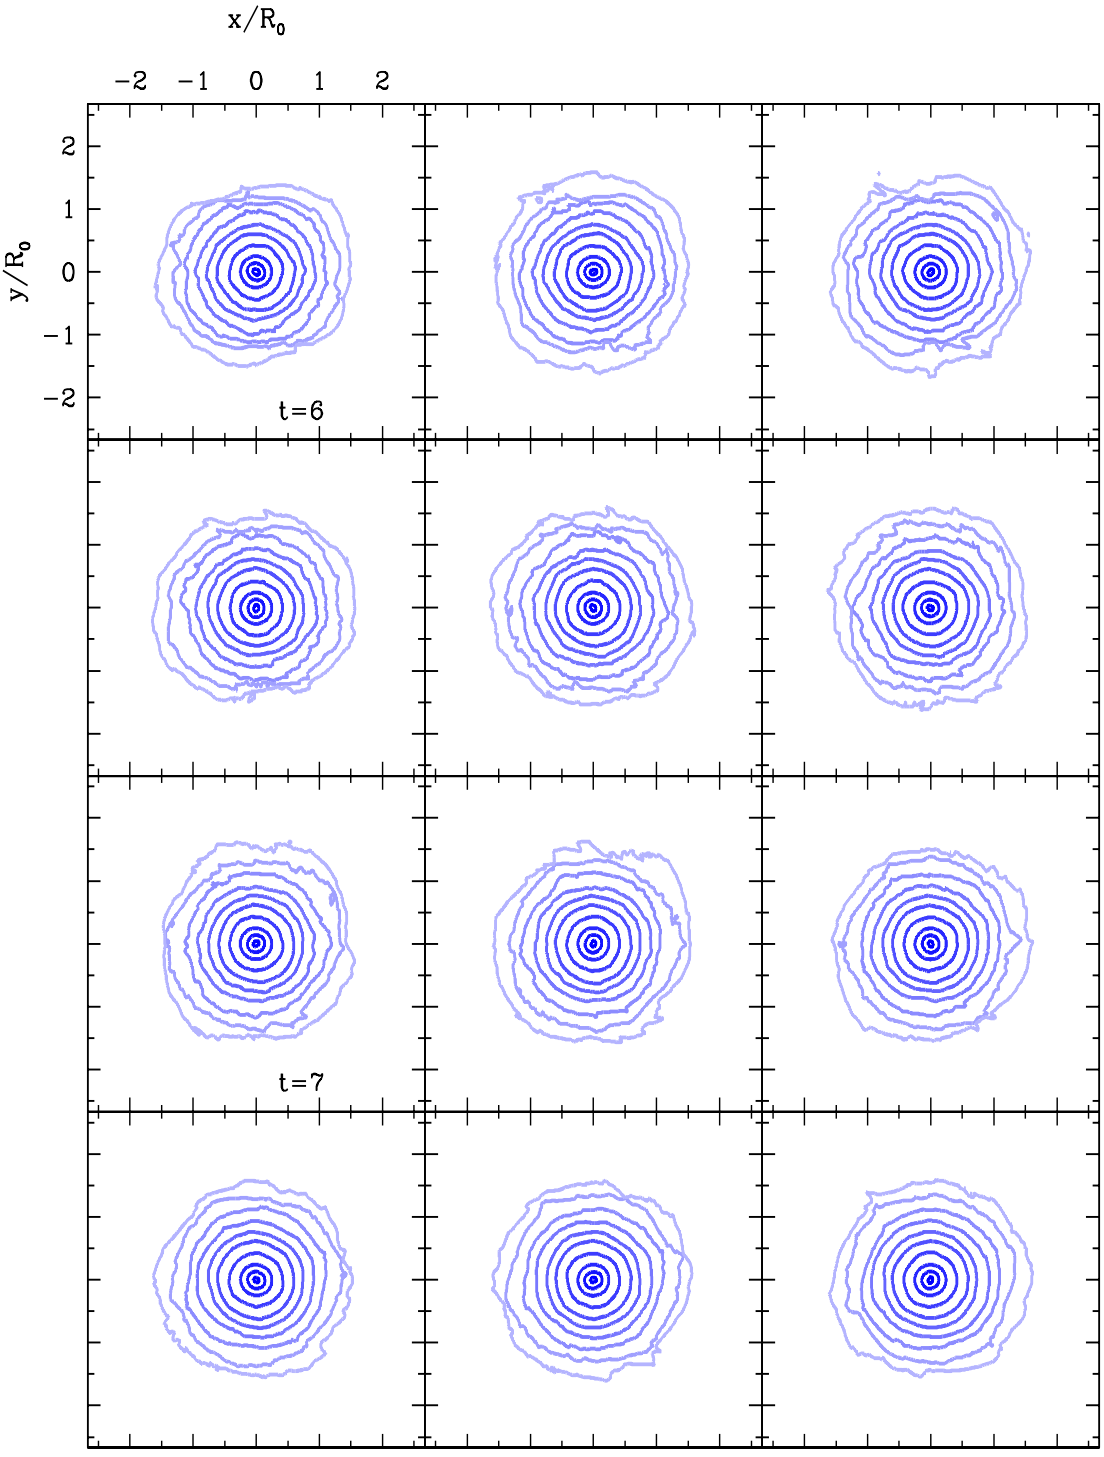
x/R0
−2 −1 0 1 2
y/R0
2
1
0
−1
−2
t=6
t=7

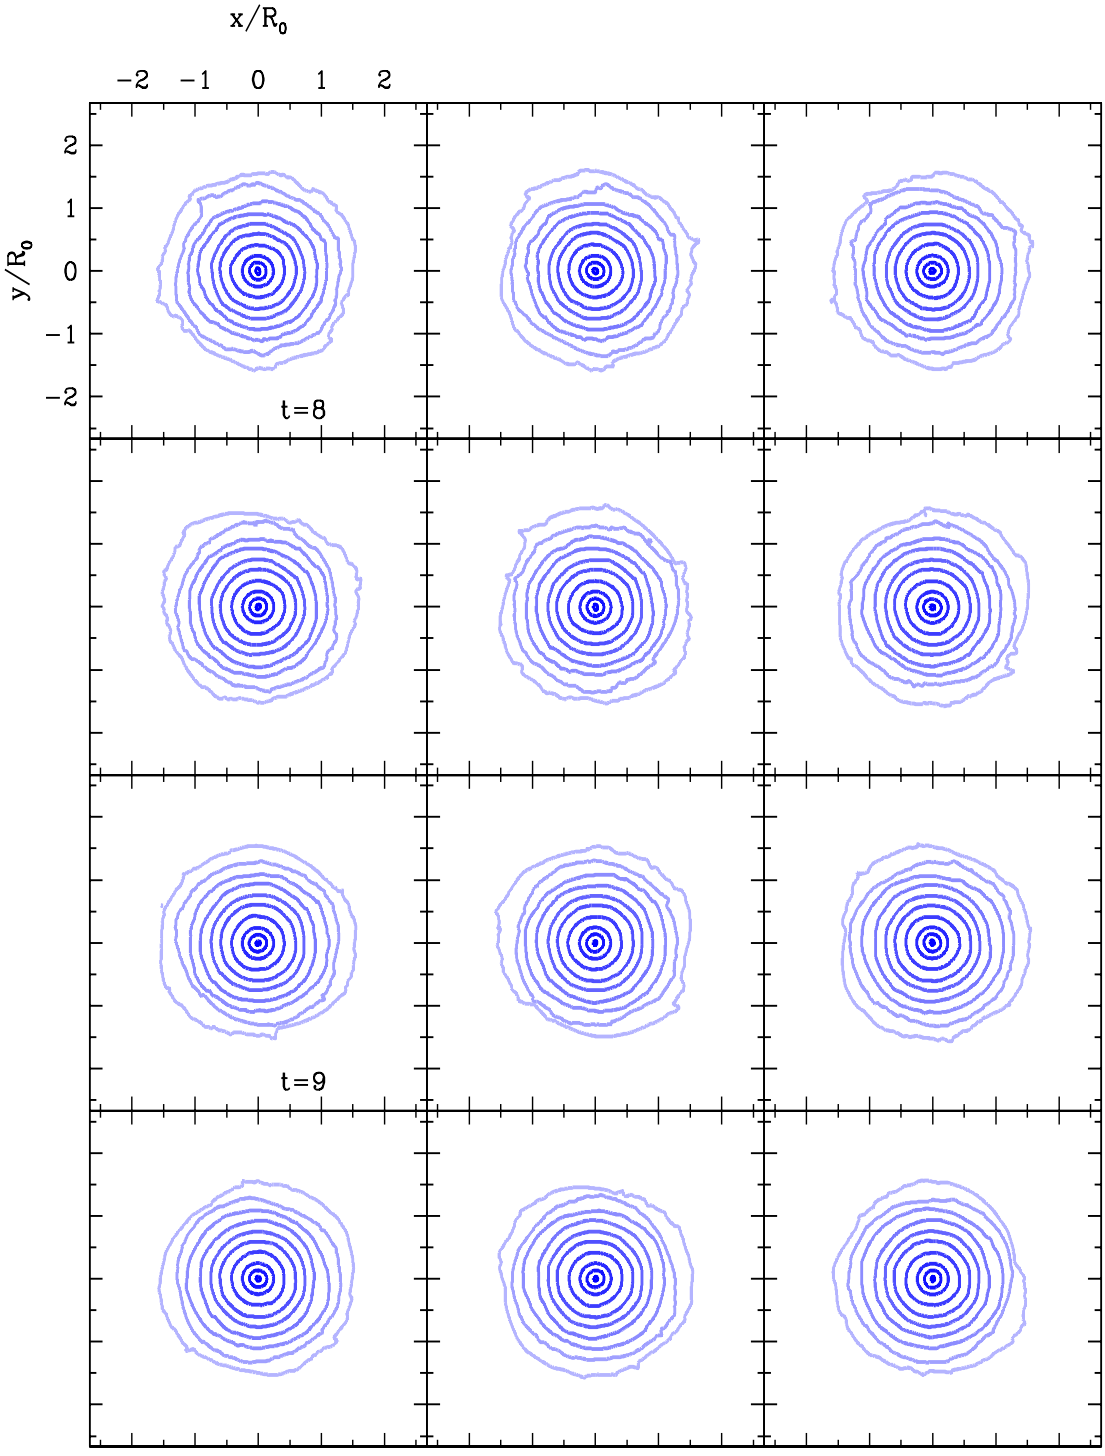

Figure A.1 Snapshots of the off-axis collision of two equal mass $M = 0.8M_{\odot}$ main sequence stars for time $0 \leq t/(\tau_{\rm dyn}) \leq 10$.

Appendix B

Adaptive Particle-to-Mesh Assignment

B.1 Adaptive CIC Scheme

I discussed in Chapter 6 how the "cloud-in-cell" (CIC) mass assignment scheme in PM N-body codes softens the mass and force resolution near the grid scale. In a pair-wise force test, the PM force on particles separated by one grid cell is reduced by approximately 50%. As a result, the net force in high density regions of a clustered distribution of particles tends to be underestimated in general. The "nearest grid point" (NGP) mass assignment scheme does produce a harder force, but this discontinuous mapping results in large anisotropic scatter in the pair-wise force. In high density regions, the scatter is averaged away and the NGP net force is more accurate than the CIC force. Unfortunately, the scatter poses a severe problem elsewhere.

One simple method of alleviating the softening is to use an adaptive CIC (ACIC) scheme where the volume of the cloud is inversely proportional to the density of the cell. In the CIC scheme, the size l of the cubical cloud is exactly equal to one grid spacing Δx, but in this adaptive scheme the size of the cloud is given by

$$l = \lambda \frac{\Delta x}{N}, \tag{B.1}$$

where N is the number of particles found in the cell and $\lambda \geq 1$ is a constant. For $\lambda > 1$, the cloud size could be restricted to be $l = \Delta x$ for $N < \lambda$. As N increases, this scheme smoothly changes over from CIC to NGP. Increasing the multiplicative factor λ decreases the anisotropic scatter. Note that the mass assignment and force interpolation must rely on the same scheme in order to prevent fictitious self-force. In principle, this simple fix should produce a harder force in high density regions.

The accuracy of both the CIC and ACIC scheme can be quantified by calculating the force errors that arise when the schemes are applied to measuring net forces on particles. First, a clustered

distribution of particles is generated by running a PM N-body simulation with 128^3 particles on a high resolution 1024^3 grid in a 10 h^{-1}Mpc box. The simulation is started from an initial redshift of $z = 60$ and outputs are saved at $z = 0$, 1, 3, and 7. For the simulation, the standard CIC scheme is used for mass assignment and force interpolation. The high grid to particle ratio improves the force measurements and reduces the softening in clustered regions.

For each redshift output, the 128^3 dark matter particles are then mapped to a 256^3 and the net force on each particle is measured using both assignment schemes. The true net force on each particle can be calculated exactly using the EWALD method, but this is computationally intensive and beyond the accuracy needed for this test. Instead, the true net force is taken to be that which has been calculated on the 1024^3 grid and the fractional force errors are defined as

$$\delta f \equiv \frac{|\boldsymbol{f}_{256}|}{|\boldsymbol{f}_{1024}|} - 1. \tag{B.2}$$

Plotted in Figure B.1 are histograms of the force errors at four redshifts. The green and blue lines show the fraction $F(\delta f)$ of particles with force errors δf for the CIC and ACIC schemes, respectively. The red and violet lines show the asymmetric cumulative fraction

$$C(\delta f) = \begin{cases} F(< \delta f) & \text{if } \delta f < 0, \\ F(> \delta f) & \text{if } \delta f > 0, \end{cases} \tag{B.3}$$

from summing the green and blue lines, respectively. At high redshifts, the clustering is weak and there is little difference between the two assignment schemes. At lower redshifts, the clustering increases and the fraction of particles with negative values of δf increases with the CIC scheme. The force errors are predominantly negative because the net force is underestimated in overdense regions. However, the ACIC scheme manages to reduce the force errors in high density regions and decrease the fraction of particles with $\delta f \approx -1$. At redshift $z = 0$, 10% of the particles have $\delta f \lesssim 0.7$ in the CIC scheme and $\delta f \lesssim 0.5$ in the ACIC scheme. The ACIC results shown are for a multiplicative factor $\lambda = 1$. Increasing λ reduces the differences between the ACIC and CIC schemes, as expected.

Comparing the force error distributions from the CIC and ACIC schemes suggests that the adaptive method is an improvement over the standard one. I then implemented the PM N-body code with the ACIC scheme for mass assignment and force interpolation and ran a simulation with the same 128^3 particles but on a 256^3 grid. If the ACIC scheme is indeed an improvement, then there should be less force softening in overdense regions and the simulated dark matter power spectrum should show less deficit in power at small scales relative to the CIC scheme. For redshifts $z = 0$, 1, 3, and 7 the particles are mapped onto a high resolution 1024^3 and the power spectra are computed using FFTs and plotted in Figure B.2. The ACIC scheme appears to lose more power at small scales, somewhat unexpected in light of the results from the force error test. The ACIC power spectra should actually be closer to the high resolution spectra from the 1024^3 simulation

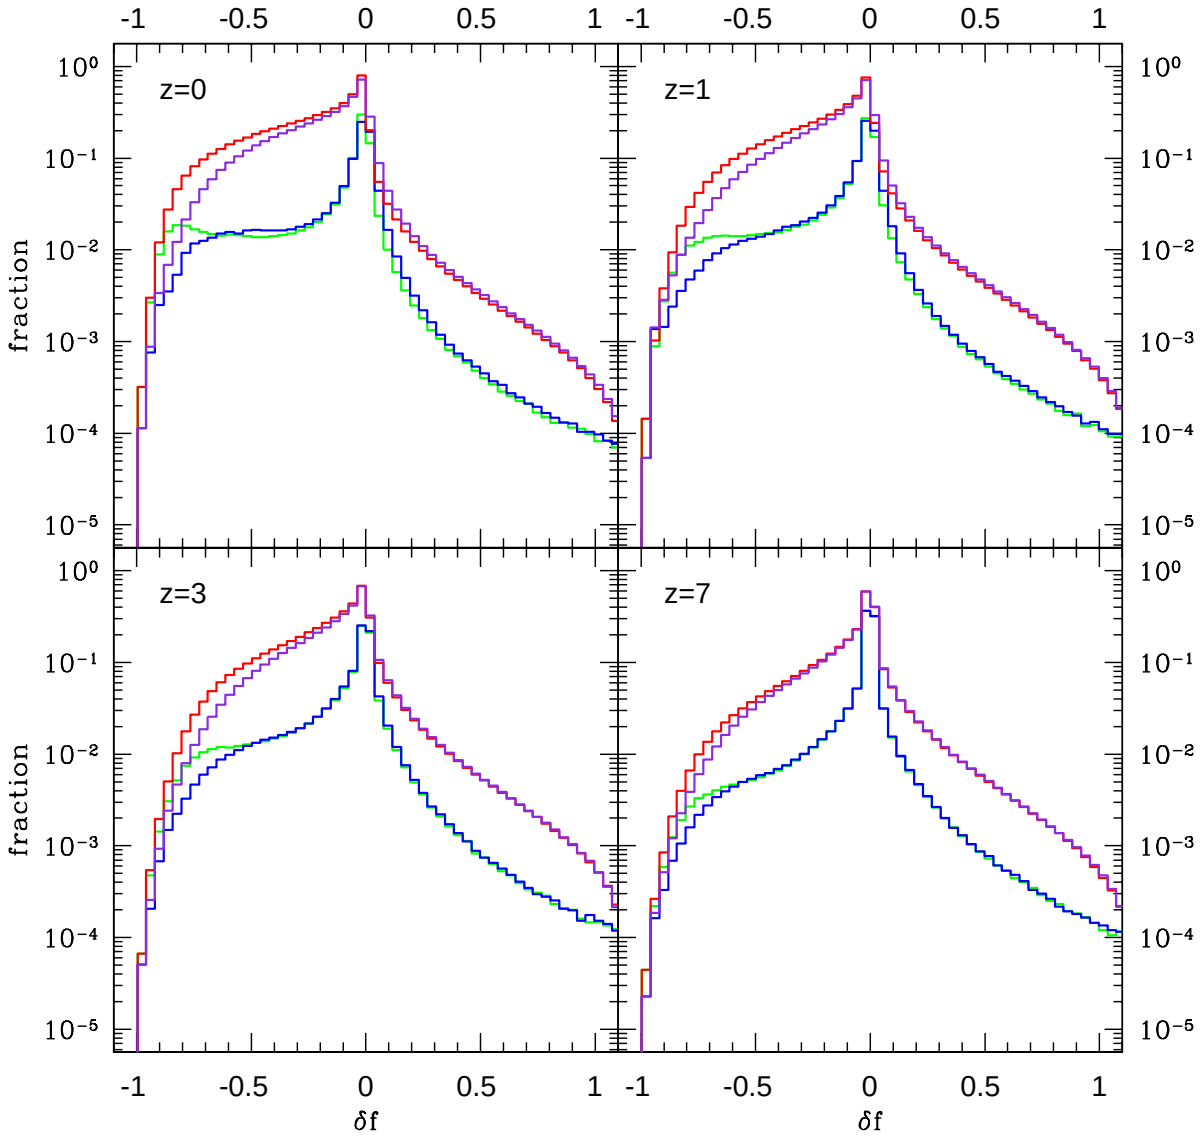

Figure B.1 Force error distributions comparing the CIC and ACIC assignment scheme. The green and blue lines show the fraction of particles with force errors δf for the CIC and ACIC schemes, respectively. The red and violet lines show the asymmetric cumulative fraction. The ACIC scheme with $\lambda = 1$ reduces the force errors in high density regions and decrease the fraction of particles with $\delta f \approx -1$.

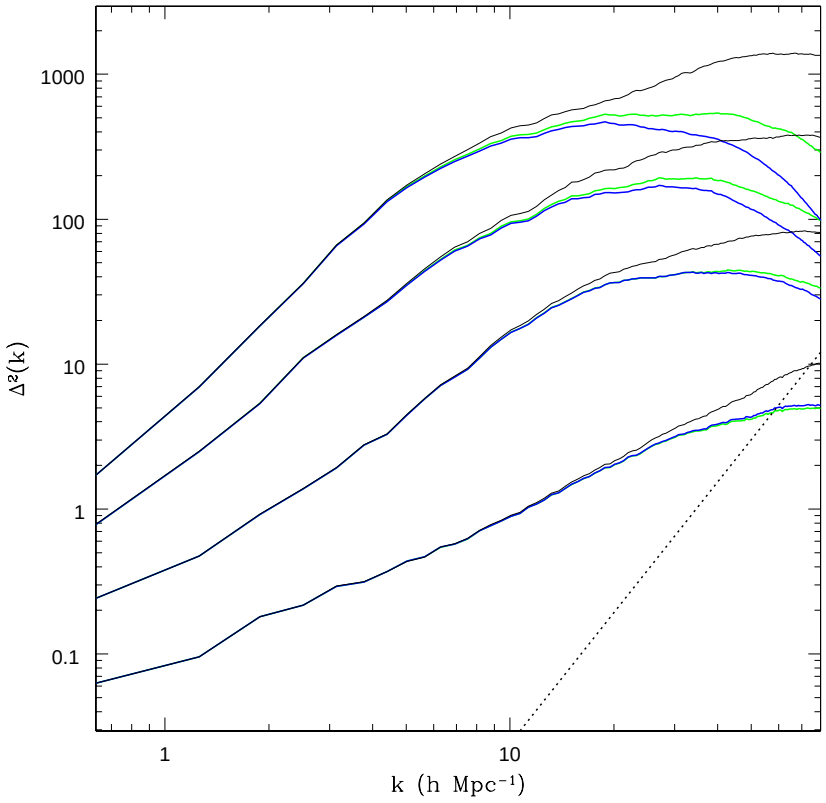

Figure B.2 Dark matter power spectra at redshifts $z = 0$, 1, 3, and 7 from a PM N-body simulation with 128^3 particles. The green and blue lines are from the 256^3 grid using the CIC and ACIC schemes, respectively. The solid line is from a high resolution 1024^3 grid while the dotted line is the Poisson noise. The ACIC scheme with $\lambda = 1$ unexpectedly loses more power on small scales compared to the CIC scheme.

than the CIC spectra. The additional anisotropic scatter could be the source of the power loss. By increasing the value of λ, the ACIC spectra lose relatively less power and approach the CIC spectra on small scales. However, the ACIC scheme does not improve upon the CIC scheme for any value of λ. More work is needed to understand the source of the discrepancy between the results of the force error test and the power spectrum test.