

© Copyright Brian William O'Shea, 2005. All rights reserved.

THE FORMATION OF POPULATION III STARS AND THEIR EFFECT ON
COSMOLOGICAL STRUCTURE IN THE EARLY UNIVERSE

BY

BRIAN WILLIAM O'SHEA

B.S., University of Illinois at Urbana-Champaign, 2000

M.S., University of Illinois at Urbana-Champaign, 2001

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2005

Urbana, Illinois

THE FORMATION OF POPULATION III STARS AND THEIR EFFECT ON COSMOLOGICAL STRUCTURE IN THE EARLY UNIVERSE

Brian William O’Shea, Ph.D.
Department of Physics
University of Illinois at Urbana-Champaign, 2005
Michael L. Norman, Advisor

The first generation of stars to form in the universe have a profound impact on their environment. These stars are responsible for beginning the universe’s transition from a “cosmic dark age” where no sources of visible light existed, to the bright universe seen today. Additionally, these stars were believed to be the first sources of all elements heavier than lithium, which strongly affected the ability of gas to cool and permanently changed how star formation occurred.

In this dissertation I present results from numerical simulations of the formation of the first generation of stars to form in the universe (“Population III” stars) and their effects on later structure formation. I compare Enzo, the adaptive mesh refinement cosmology code used to perform all of the simulations in this work, to GADGET, a smoothed particle hydrodynamics cosmology code. Nearly identical results can be obtained when using two extremely different numerical methods, which helps to verify the correctness of both codes and strengthen the confidence of predictions made with these tools.

I perform high dynamical range calculations of the formation of an ensemble of Population III stars, varying multiple simulation parameters, in a standard cold dark matter cosmology as well as with a soft ultraviolet background and in a generic warm dark matter

cosmology. I find that the accretion rates of primordial protostars have been systematically overestimated by previously published work, which has profound implications for later structure formation and the reionization of the universe. Additionally, the presence of a soft ultraviolet background and warm dark matter serves to delay the onset of star formation. I propose limits on the possible mass of a warm dark matter particle.

I also present results of simulations which demonstrate the effects of the HII regions and metal enrichment from Population III stars. It appears that HII regions from these stars may hasten the formation of later generations of stars, and even the weakest supernova can spread material over large distances. Further calculations indicate that even with the most optimistic assumptions these stars cannot be responsible for the metals observed at low densities in the Lyman-Alpha forest.

To my family – past, present, and future.

Acknowledgments

I would like to extend my deepest thanks to my advisor, Michael Norman. His support and guidance has made the completion of this research possible. In addition, I would like to thank my thesis committee, Brian Fields, George Gollin and Telemachos Mouschovias, for their participation and useful comments.

I would like to thank my collaborators and scientific peers for useful discussion. In particular, I would like to thank Tom Abel, James Bordner, Stirling Colgate, David Collins, Nathan Currier, Chris Fryer, John Hayes, Alexander Heger, Falk Herwig, Tridivesh Jena, David Kirkman, Alexei Kritsuk, Mike Kuhlen, Hui Li, Wen-Ching Lin, Ken Nagamine, John O'Meara, Paolo Padoan, Pascal Paschos, Gabe Rockefeller, Frank Timmes, Matthew Turk, David Tytler, Dan Whalen, and John Wise. Many thanks go to Greg Bryan for useful discussions, for writing Enzo, and for all of the technical and scientific support that he has given me over the years. Without Enzo and his support I would not have been able to complete this work. Many thanks also to Robert Harkness for a huge amount of Enzo support and useful discussion about supercomputing, and to the computer consultants at the National Center for Supercomputing Applications and the

San Diego Supercomputing Center for their invaluable assistance, which often took place at extremely non-standard work hours. I would also like to thank Cheryl Matson, Dayna Gallegos, and Wendy Wimmer for their extremely efficient administrative support.

I would like to thank Richard Cyburt for providing me with his L^AT_EXthesis templates, which helped tremendously in the timely completion of this work. I would also like to thank my friends in San Diego and on #tvd for their support and encouragement. Special thanks go to Mark Wroblewski for his long-standing friendship, and for repeatedly helping me move my personal possessions from one end of the country to the other.

I would particularly like to thank my family for their support and encouragement (and for believing in me even though they have no idea what I actually do for a living). Last but most certainly not least, I would like to thank Melissa Wurtzel for being there for me every time I needed someone to talk to, for encouraging me when times were rough, and most importantly, for loving me and agreeing to marry me!

This work was completed with support from the University of Illinois Physics and Astronomy departments, the Center for Astrophysics and Space Sciences and the Department of Physics at the University of California in San Diego, NASA grant NAG5-12140, NSF grant AST-0307690 and NRAC allocation MCA98N020. In addition, this work has been supported under the auspices of the U.S. Department of Energy by its contract W-7405-ENG-36 to the Los Alamos National Laboratory. Some of the contents of this thesis have been published previously [1, 2, 3, 4].

Contents

Acknowledgments	vi
1 Introduction	1
1.1 Big-Bang cosmology and the FRW universe	2
1.2 Cosmological structure formation	5
1.3 Formation of the first generation of stars	7
1.3.1 A brief history of research regarding the first stars	7
1.3.2 The role of molecular hydrogen	8
1.3.3 The role of magnetic fields	9
1.3.4 Numerical simulations of Pop III star formation	11
1.3.5 The Initial Mass Function of primordial stars	13
1.4 Feedback from the first generation of stars	16
1.4.1 Physical properties of Population III stars	16
1.4.2 Radiative feedback	17
1.4.3 Chemical feedback	20
1.5 Flaws in the Λ CDM paradigm	21
2 Simulation Methodology	23
2.1 The generation of cosmological initial conditions	23
2.1.1 Creating Gaussian random fields	24
2.1.2 Creating initial conditions for high dynamical range simulations	27
2.1.3 Numerical limitations of cosmological initial conditions	30
2.2 The Enzo Simulation Code	32
2.2.1 The AMR machinery	33
2.2.2 Hydrodynamics with the piecewise parabolic method	38
2.2.3 Hydrodynamics with the ZEUS hydrodynamic method	39
2.2.4 Gravity Solver	40
2.2.5 Radiative processes and non-equilibrium primordial chemistry	42
2.2.6 Star Formation and Feedback Algorithms	45

3	A Comparison of AMR and SPH cosmology codes	50
3.1	Summary	50
3.2	Motivation	51
3.3	The GADGET smoothed particle hydrodynamics code	53
3.3.1	Hydrodynamical method	53
3.3.2	Gravitational method	55
3.4	The simulation set	57
3.5	Simulations with dark matter only	62
3.5.1	Dark matter power spectrum	62
3.5.2	Halo dark matter mass function and halo positions	63
3.5.3	Halo dark matter substructure	68
3.6	Adiabatic simulations	69
3.6.1	Unperturbed adiabatic expansion test	69
3.6.2	Differential distribution functions of gas properties	72
3.6.3	Cumulative distribution functions of gas properties	73
3.6.4	Phase diagrams	76
3.6.5	Mean gas temperature and entropy	80
3.6.6	Evolution of kinetic energy	82
3.6.7	The gas fraction in halos	82
3.7	The role of artificial viscosity	86
3.8	Timing & memory usage	91
3.8.1	Initial comparison on a distributed memory machine	92
3.8.2	More recent comparison on a shared memory machine	92
3.8.3	Mass resolution in Enzo and GADGET	94
3.9	Discussion and conclusions	99
4	The Formation of Population III stars in a ΛCDM universe	103
4.1	Summary	103
4.2	Motivation	104
4.3	Problem setup	105
4.4	Results	106
4.4.1	Collapse of a representative primordial star	106
4.4.2	Angular momentum transport	118
4.4.3	Consistency of result across multiple realizations	127
4.4.4	The formation of a Population III star in the presence of a soft UV background	141
4.5	Discussion	151

5	Formation of primordial stars in a WDM universe	161
5.1	Summary	161
5.2	Motivation and theory	162
5.3	Problem setup	165
5.4	Results	166
5.4.1	Comparison of realizations with many WDM particle masses . . .	166
5.4.2	Comparison of the evolution of two representative WDM realizations	168
5.5	Discussion	177
6	The formation of second-generation primordial objects	179
6.1	Summary	179
6.2	Motivation	180
6.3	Simulation Setup	180
6.4	Results	182
6.4.1	Comparison of the First and Second Stars	183
6.4.2	Black Hole Accretion	185
6.5	Discussion	185
7	Formation of metal-enriched 2nd generation objects	188
7.1	Summary	188
7.2	Motivation	188
7.3	Problem setup	189
7.4	Preliminary results	191
7.5	Discussion	195
8	Pre-galactic enrichment of the IGM	199
8.1	Summary	199
8.2	Motivation	200
8.3	Problem Setup	200
8.4	Results	202
8.5	Discussion	206
9	Summary and future Work	209
9.1	Summary	209
9.2	Future Work	212
A	The Enzo Primordial Chemistry Reaction Network	215
	Bibliography	219
	Author’s Biography	231

Chapter 1

Introduction

Cosmology is defined as “*The study of the physical universe considered as a totality of phenomena in time and space.*”[5] As one might expect from this lofty definition, the exploration of the nature of the universe has long been the province of poets, philosophers and religious thinkers – indeed, for the majority of the history of humanity, the field of cosmology has been dominated by the attempt to understand mankind’s role in the universe and his relationship with a god or gods. Most religions have some sort of creation myth that explain how the earth and the universe came to be and how all forms of life appeared. Typically these myths describe the universe as being created by a deity of some sort, who is also responsible for the creation of the earth and of mankind. These myths often foretell the end of the universe in great (and often gory) detail.

In the past century, physics has come to play a central role in shaping our understanding of the universe (though not necessarily our place in it). The development of the field of “physical cosmology” has been driven almost entirely by the improvement in technology used in astronomical observations and by Einstein’s theory of general relativity. The theory of “Big Bang” cosmology (described in Section 1.1) was proposed by Georges Lemaître, building upon Einstein’s ideas, and was later confirmed by Hubble’s observations of the recession of distant galaxies in 1924 [6] and by the discovery of the cosmic microwave background (CMB) by Penzias and Wilson in 1964 [7, 8].

Despite these advances, the field of cosmology was starved for data until quite recently. In the past decade, massive statistical surveys of galaxies and large-scale structure such as the Two-Degree Field (2dF) survey and the Sloan Digital Sky Survey (SDSS), combined with high-resolution observations of the cosmic microwave background by the Wilkinson Microwave Anisotropy Probe (WMAP) satellite and of distant supernovae, have allowed us to constrain essentially all of the cosmological parameters, such as the amount of matter in the universe, the rate of expansion, and the existence and rough properties of a cosmological constant, to within a few percent.

This new epoch of “precision cosmology” has also fostered a renaissance in compu-

tational structure formation. In the past, numerical simulations of large-scale structure (such as galaxies and clusters of galaxies) have been primarily used to constrain cosmological parameters and to rule out ideas such as the concept of “hot dark matter.” However, with the rapid advances in observations and the corresponding constraint of the basic cosmological parameters, numerical simulations of large-scale structure formation can now be used in a predictive sense, to study the formation of distant and highly nonlinear objects which are too complicated to approach purely analytically.

My thesis presents the results of high-resolution numerical simulations of structure formation in the early universe. In order to verify that Enzo, the cosmology simulation code used for the work presented here, is working correctly, I perform a comparison between it and the SPH code GADGET. I then predict properties of the first generations of stars in both the fiducial cold dark matter cosmology and also in a universe with a generic “warm dark matter” cosmology. I also show how the first generation of stars in the universe affects following generations of star formation, and present constraints on how much of an impact these “first stars” can have on the feedback of metals into the low-density intergalactic medium. In this chapter, I will review the basic principles of cosmology and cosmological structure formation, and also discuss the current state of literature on the formation of the first generation of stars and their effects on the universe via feedback processes.

1.1 Big-Bang cosmology and the FRW universe

The Big Bang theory of cosmology rests on two theoretical pillars: Einstein’s theory of General Relativity and what is known as the Cosmological Principle. Einstein’s theory provides a mathematical framework for describing gravity as a distortion of space and time and is a generalization of Newton’s theory of gravity. The Cosmological Principle assumes that, on very large scales, the universe is homogeneous and isotropic – there is no preferred reference frame, and the universe looks the same no matter where an observer is within it. Additionally, it is generally assumed that the laws of physics are the same at all places and at all times.

The content of the universe is described by the “standard model” of particle physics. This model, coupled with the understanding of how the universe expands from general relativity, provides predictions of the primordial composition of the universe (via Big Bang Nucleosynthesis, or BBN) and has been observationally confirmed to very high precision. Observations indicate that the universe is mostly composed of some sort of vacuum energy (“dark energy”), and a form of matter that only appears to interact with baryonic matter via gravitational coupling (“dark matter”). These observations indicate that at the present epoch baryons comprise approximately 4% of the total energy density of the universe, dark matter comprises approximately 23%, and vacuum energy comprises

roughly 73% of the total. The total energy density adds up to a “critical density” which suggests that the universe is geometrically flat (as discussed below).

This combination of theory and observations forms the entire theoretical basis of Big Bang cosmology, and produces very specific predictions for observable properties of the universe. These predictions have been heavily tested and appear to be accurate, suggesting that the Big Bang cosmological model is an accurate description of the universe. Given this model of the universe as a starting point, we can then make predictions about the nonlinear processes that follow, such as the formation of large-scale structure.

In principle one can use the theory of general relativity to predict the properties of any kind of universe. However, when one assumes that the universe is isotropic and homogeneous, the only sort of movement that is allowed is a uniform expansion of the universe as a whole. Under these assumptions, Einstein’s field equations reduce to the following pair of independent equations:

$$\left(\frac{\dot{a}}{a}\right)^2 - \frac{8\pi G}{3}\rho = -\frac{kc^2}{a^2} + \frac{\Lambda c^2}{3} \quad (1.1)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(\rho + \frac{3p}{c^2}\right) + \frac{\Lambda c^2}{3} \quad (1.2)$$

where a is the cosmological expansion parameter, G is the gravitational constant, ρ and p are the mass-energy density and pressure of the universe, c is the speed of light, k is the curvature constant, and Λ is the cosmological constant. Equation 1.1 is commonly referred to as *Friedmann’s equation*, and has the general form of an energy equation. Equation 1.2 has the form of a force equation, and is sometimes referred to as *Friedmann’s acceleration equation*.

These equations bear further examination. The first and second terms on the left hand side of Equation 1.1 look like kinetic and gravitational potential energies, respectively, and the right hand side is effectively a total energy. Before continuing further, it is useful to define the “Hubble parameter,” $H \equiv \dot{a}/a$, which has the value H_0 at the present epoch. Similarly, we define a “critical density,” which is the matter density at the present epoch in a universe with $k = 0$ (a geometrically flat universe), and is defined as $\rho_c \equiv 3H_0^2/8\pi G$ (with a cgs value at the present day of $1.8788 \times 10^{-29} \text{ h}^2 \text{ g cm}^{-3}$, where h is the Hubble parameter in units of 100 km/s/megaparsec), and relative densities $\Omega_i \equiv \rho_i/\rho_c$. At the present epoch ($a = 1$ and $\dot{a} = H_0$) with $\Lambda = 0$ and $\rho \equiv \Omega\rho_c/a^3$, we get $H_0^2 = \Omega H_0^2 - kc^2$, or $k = H_0^2/c^2(\Omega - 1)$. This provides a clear relationship between the curvature of the universe and the total mass-energy density of the universe: if $k = 0$, $\Omega = 1$. Likewise, if $k < 0$, $\Omega < 1$ and if $k > 0$, $\Omega > 1$. Also, it can be shown that for large values of the scale factor ($a \rightarrow \infty$) Equation 1.1 reduces to:

$$\dot{a}^2 = \frac{H_0^2}{c^2}(1 - \Omega) \quad (1.3)$$

In the absence of a cosmological constant, the curvature parameter k determines whether the universe has a net positive, negative or zero energy. If $k = +1$ ($\Omega > 1$) the universe is said to have “positive curvature,” meaning that the universe is effectively closed – the kinetic energy term is always dominated by the potential energy term. Practically speaking, this results in a universe that expands, attains some maximum size, and then contracts again to a point, or a “Big Crunch.” If $k = -1$ ($\Omega < 1$), the universe is said to have “negative curvature” and the kinetic energy term on the left hand side of Equation 1.1 always dominates. This means, effectively, that the universe will expand forever with some positive kinetic energy. This is referred to as an “open” universe. If $k = 0$ the universe is geometrically flat and (from a strict interpretation of Equations 1.1 and 1.3) its expansion will coast to a halt as $a \rightarrow \infty$.

The addition of a cosmological constant (Λ term) complicates matters somewhat. Einstein originally added this constant to maintain a steady-state universe. However, with Hubble’s announcement of the observation of an expanding universe, Einstein abandoned the idea of a cosmological constant, referring to it as “the greatest blunder of my life.” [9]. However, current observations suggest that roughly 70% of the energy density of the universe at the present epoch is due to a mysterious “dark energy” that behaves like a positive cosmological constant, making the cosmological constant an issue once again.

Examination of Equation 1.2 shows that, even in the absence of any other source of mass-energy (e.g. $\rho = p = 0$), the existence of a positive cosmological constant indicates that the expansion of the universe is accelerating – essentially, the vacuum energy is acting as a repulsive force. Also, if $k = 0$, Equations 1.1 and 1.2 can be reworked to show that at the present epoch $\Omega_{tot} = 1.0$, where $\Omega_{tot} = \Omega_m + \Omega_{rad} + \Omega_\Lambda$, the sum of all of the constituents of the mass-energy density of the universe. Here Ω_m is the total matter content of the universe, Ω_{rad} is the total relativistic particle content (including photons), and Ω_Λ is the energy density of the cosmological constant, all in units of the critical density. At the present epoch the approximate values of Ω_{rad} , Ω_m and Ω_Λ are $\sim 10^{-4}$, 0.27 and 0.73, respectively (with $\Omega_m = \Omega_b + \Omega_{dm}$, where $\Omega_b = 0.04$ and $\Omega_{dm} = 0.23$).

For the purposes of clarity, we can simplify Equation 1.1 to be in terms of the components of the mass-energy contents of the universe at the present epoch. The proper energy density of matter scales with the cosmological constant as a^{-3} due to pure geometrical dilution. The proper energy density of relativistic particles such as photons scales as a^{-4} – a factor of a^{-3} due to geometric dilution and an additional factor of $1/a$ due to the redshifting of a particle’s momentum. The proper energy density of the cosmological constant is unchanged at all times, by definition. Refactoring Equation 1.1 gives us:

$$\left(\frac{\dot{a}}{a}\right)^2 + \frac{kc^2}{a^2} = H_0^2 \left(\frac{\Omega_{rad}}{a^4} + \frac{\Omega_m}{a^3} + \Omega_\Lambda \right) \quad (1.4)$$

The different scaling factors for each of the terms on the right hand side suggest that each one dominates at different epochs, with radiation dominating first (at very small

a), then matter, then finally the cosmological constant at late times (the epoch that we are currently entering into). The universe expands as $a(t) \sim t^{1/2}$ during the radiation-dominated epoch, $a(t) \sim t^{2/3}$ during the matter-dominated epoch, and $a(t) \sim e^t$ during the cosmological constant-dominated epoch.

For more detailed discussion of general relativity, big bang nucleosynthesis and related topics, the following references may be of use: [10, 11, 12, 13, 14].

1.2 Cosmological structure formation

The current paradigm describing the formation of large-scale structure is referred to as “hierarchical structure formation.” During the epoch of inflation, quantum mechanical effects manifested themselves as very tiny density perturbations in an otherwise homogeneous universe. As the universe expanded these perturbations grew via gravitational instability and eventually became gravitationally bound halos, which grew by a sequence of mergers into the galaxies and other large scale structure observed today.

This scenario is demonstrated analytically by the Press-Schechter (PS) formalism [15, 16, 17]. The PS formalism very accurately describes many properties of the dark halo population in the Λ CDM cosmology, and has been verified both observationally and numerically to be accurate on large scales. An example of the use of the PS formalism to describe the growth of large-scale structure is shown in Figure 1.1. This is a plot of cumulative dark matter halo mass functions for several redshifts. At very early times the universe is sparsely populated with gravitationally bound objects. As time goes by (redshift decreases), dark matter halos grow more numerous and the maximum halo mass increases via merger of smaller halos. It is intriguing to note that at the current epoch ($z = 0$) there are actually *fewer* low mass halos than at higher redshifts (earlier times). This supports the idea of hierarchical mergers of dark matter halos.

Though the PS formalism provides a good description of the dark matter halo properties in a Λ CDM cosmology, it only tells the simplest part of the story, and only in a purely statistical sense. Also, by definition dark matter is not directly observable. Baryons in stellar and gaseous form comprise all of the visible matter in the universe, and are much more complicated to model. It is difficult, if not impossible, to analytically model the effects that halo mergers would have on the properties of the baryons that are gravitationally bound to the dark matter halos. The range of physics that is involved – radiative cooling, star formation, the feedback of radiation and metals – combine together with the unique merger histories of individual halos to produce the galaxies, groups and clusters that are observed in the universe today.

It is for this reason that three-dimensional numerical simulations of the formation and evolution of large scale structure are exceedingly useful. Essentially all of the baryon physics described above can be modeled, either from first principles or through con-




Figure 1.1: Cumulative halo dark matter mass function calculated using the Press-Schechter formalism for a cosmology with fiducial cosmological parameters ($\Omega_m = 0.3, \Omega_\Lambda = 0.7, h = 0.7, \sigma_8 = 0.9, n = 1$). Plot is of comoving number density of halos greater than a given mass M vs. mass for several redshifts. The number to the right of each line indicates the redshift at which the mass function is calculated.

strained phenomenological models, within a cosmological framework. The entire scope of cosmological structure formation simulations is vast and continually expanding with increases in computing power and the creation of new simulation techniques. For a somewhat dated review, see the 1998 Bertschinger Annual Reviews article [18].

1.3 Formation of the first generation of stars

The first generation of stars (also known as Population III stars or Pop III stars) formed in a very different, and far simpler, physical environment than present-day stars. According to BBN, the primordial gas consisted of primarily hydrogen and helium (76.2%/23.8% by mass, respectively) and trace amounts of deuterium and lithium - an extremely simple mix of elements with well-understood and easily modeled gas chemistry! [19]. Unlike the present-day universe, there were no metals or dust grains - two factors which complicate the solution of the modern-day star formation problem tremendously. The main source of cooling in the early universe was molecular hydrogen, which is inefficient below 200 K, as opposed to current star formation, where cooling via the heavier elements allows temperatures in the molecular clouds where star formation takes place to drop to $\simeq 10$ K [20, 21].

The universe during the epoch of formation of the first stars was a very dull place. By definition, there were no other stars – meaning that there were no sources of radiation, winds or supernovae that could affect star formation in any way. Significantly, this also means that there were no sources of intense ultraviolet radiation to disrupt the formation of molecular hydrogen and no cosmic rays to ionize hydrogen. Also, there were no sources to sustain turbulent motion, as long as the density perturbations remained linear. Only after the explosion of the first supernovae, and the associated input of mechanical and thermal energy, was this state of quiescence bound to change [22, 23].

1.3.1 A brief history of research regarding the first stars

The history of research regarding the first generation of stars is long and full of conflicting results. Peebles & Dicke [24] were among the pioneers in the field. They suggested in 1968 that globular clusters may have originated as gravitationally bound gas clouds before the galaxies form. Their idea follows from what was then called the primitive-fireball picture (and is now referred to as the Big Bang theory) and they showed that the first bound systems to have formed in the expanding universe were gas clouds with mass and shape similar to the globular star clusters observed around the Milky Way and nearby galaxies. They also argued that only a small fraction of the total cloud mass would fragment into stars and they also discussed the influence of molecular hydrogen on cooling and fragmentation. A year later, Hirasawa [25] performed similar calculations but claimed

that his results suggested that collapsing hydrogen clouds would result in supermassive black holes.

Palla, Salpeter & Stahler published an important work in 1983 discussing the role of molecular hydrogen (H_2) in star formation [20]. They discuss (as I will in Section 1.3.2) the cooling of a collapsing cloud of hydrogen gas via molecular hydrogen (H_2) formation and explore the importance of the three-body reaction for creating H_2 . They also suggest that the Jeans mass is higher for stars without metals to cool them but argue that cooling leads to a rapidly dropping Jeans mass, resulting in fragmentation which would lead to low-mass stars regardless. In the same year, Silk published a work on Population III stars showing that large density fluctuations of $\sim 0.1 M_\odot$ ($M_\odot =$ “solar mass”) arise in any collapsing cloud with extremely low metallicity ($Z_{cloud} \leq 10^{-3} Z_\odot$) [26]. Gravitational instability ensures that many of the clumps coagulate to form protostars of masses extending up to the Jeans mass at the time when the fluctuations start to develop, roughly $\sim 100 M_\odot$. He argues that the primordial IMF would have spanned the mass range from ~ 0.1 - $100 M_\odot$ but could have been dominated by the more massive stars.

This disagreement in theoretical studies seems quite surprising. However, the first bound objects in the universe formed via the gravitational collapse of a thermally unstable reactive medium, which naturally makes conclusive analytical calculations difficult.

1.3.2 The role of molecular hydrogen

Molecular hydrogen (H_2) in primordial gas clouds is produced at low densities primarily by these coupled gas-phase reactions:



This pair of reactions depends on free electrons to act as a catalyst, and even a small mass fraction of molecular hydrogen ($f_{H_2} \sim 10^{-3}$) can contribute significantly to the cooling of a cloud via the rotational and vibrational transitions of the hydrogen molecule, allowing primordial gas to cool efficiently below $\sim 10^4$ K, which is the lowest temperature gas can radiatively cool to due to atomic hydrogen line transitions. Significant amounts of molecular hydrogen can cool gas down to $\simeq 200$ K. Below this temperature it is relatively ineffective as a coolant, as can be seen from Figure 1.2, which shows the molecular hydrogen cooling function for gas at three different densities. At all densities, the cooling rate of a gas of primordial composition decreases sharply below $\simeq 200$ K. Metal-enriched gas can cool to much lower temperatures efficiently, due to the presence of many closely-spaced line transitions in the various molecules and dust grains that exist.

At high densities in primordial gas ($n \geq 10^8 \text{ cm}^{-3}$) the 3-body reaction for the formation of molecular hydrogen becomes dominant:



and is so efficient that virtually all of the atomic hydrogen at that density or above can be converted to molecular hydrogen before it is dissociated, allowing rapid cooling and contraction of the gas cloud [20, 21].

The properties of the hydrogen molecule are extremely important in studying the formation of the first stars. In the absence of metals, the properties of H_2 completely control the size and formation times of the first objects. Therefore, it is extremely important to include the effects of H_2 formation in simulations of the formation of the first stars.

For an excellent review of the chemistry of the early universe (including all deuterium and lithium chemistry, which we have ignored here) see the paper by Galli & Palla [27], and for analysis of the effects of H_2 cooling on structure formation, see Tegmark et al. [28].

1.3.3 The role of magnetic fields

It is believed that the magnetic fields that existed at the epoch of first star formation were dynamically unimportant at large scales (though they may be relevant to angular momentum transport in primordial protostars). This stands in sharp contrast to the local universe, where magnetic fields play a critical role in star formation. Observations of the cosmic microwave background (CMB) provide a strong upper limit of $B \leq 3 \times 10^{-8} \text{ G}$ (as measured at the present epoch) for large-scale (megaparsec) coherent magnetic fields at the time of recombination [29], with theory suggesting that limits could be obtained which are as stringent as $\sim 1 \text{ nG}$ [30]. This limit is poor enough that it does not disprove that magnetic fields are dynamically important at the epoch of Population III star formation – see Section 4.5 for more discussion of this issue. An examination of the possible (known) sources of magnetic fields in the pre-structure formation era provides two likely candidates. A discontinuous (i.e., first order) phase transition at the time of the QCD or electroweak phase transitions could create significant coherent magnetic fields before recombination. However, the mechanisms involved are highly speculative and predictions of the possible magnetic field strengths are unreliable [31]. Intriguingly, the standard picture of cosmology predicts that large-scale magnetic fields were created at recombination due to Thompson scattering differentially accelerating electrons and ions. However, the strengths of the resulting fields are on the order of 10^{-20} G [32]. More recent work was done by Matarrese et al. [33], who derive the *minimum* magnetic field that invariably arises prior to recombination. They show that a weak magnetic field is




Figure 1.2: Molecular hydrogen cooling rate as a function of temperature for gas at three different densities. Black line: $n = 1 \text{ cm}^{-3}$ (proper). Red line: $n = 10^4 \text{ cm}^{-3}$ (proper). Blue line: $n = 10^8 \text{ cm}^{-3}$ (proper).

generated in the radiation dominated era by Harrison’s mechanism [124], which occurs in regions of non-vanishing vorticity. They show that this vorticity is generated by the 2nd order nonlinear coupling of primordial density fluctuations, resulting in a differential rotational velocity between ions and electrons that produces a small magnetic field. The power spectrum of this field is determined entirely by the power spectrum of primordial density perturbations. The RMS amplitude of this field at recombination is predicted to be $B \simeq 10^{-23} (\lambda/\text{Mpc})^{-2}$ G on comoving scales $\lambda \geq 1$ Mpc. The magnetic fields are suppressed at smaller scales via Silk damping, which is simply diffusion of photons on sub-horizon scales from high density regions to low density regions prior to recombination. Electrons are “dragged” with the photons via Compton interactions and then proceed to carry protons along with them via the Coulomb interaction. This smoothes out the matter density at small scales, which has the effect of damping out CMB fluctuations (and thus B-field creation) on those scales. The estimated B-fields, while not strong enough to be dynamically important in the formation of the first stars, could be amplified via the dynamo effect during protogalaxy formation to strengths that are significant today [35].

There is no observational evidence against strong magnetic fields at the time of recombination (only strong upper limits), but there are also no theoretical arguments demanding strong fields, so it seems reasonable to assume that magnetic fields are negligible for the first treatment of the problem. As a side note, I would direct the reader who is interested in speculations of the origin of the magnetic fields that are of such great importance in present day star formation to a paper by Kulsrud et al. [36], who discuss one possible scenario for the creation and amplification of dynamically important magnetic fields during the epoch of protogalaxy formation. In addition, Gnedin et al. [37] discuss the generation of magnetic fields by the Biermann battery mechanism in cosmological ionization fronts propagating through a dense, irregular medium. Though their estimates suggest that the magnetic fields generated are small ($\sim 10^{-19} - 10^{-18}$ G), this is a lower bound and could be amplified significantly during the following epochs of structure formation. Similar work by Langer et al. [38] presents a model of magnetic field generation based on local charge separation provide by an anisotropic, inhomogeneous radiation pressure. This process would also take place during reionization ($z \geq 7$), and would produce fields on the order of $\sim 10^{-12} - 10^{-11}$ G. They also show that these fields are generated preferentially on large (> 1 kpc) scales, and strongly suppress coherent B-fields on smaller scales. Though this is interesting, and a reasonable seed mechanism for galactic magnetic fields, it occurs at far too late of an epoch to be useful for Population III stars.

1.3.4 Numerical simulations of Pop III star formation

The most detailed *ab initio* simulations to date have been done by Abel, Bryan & Norman, henceforth referred to as ABN [39]. They perform a cosmological simulation of

the collapse of a cosmological density perturbation into a protostellar core utilizing the adaptive mesh refinement (AMR) technique to obtain very high dynamical range. They discover that the collapsing halo is initially characterized by a period of rapid cooling and infall. This corresponds to an increase in H_2 mass fraction in the center of the halo to $\sim 0.1\%$, which is sufficient to rapidly cool the gas down to ~ 200 K. Following this, the central density of the halo increases to 10^4 cm^{-3} and, at which point the cooling time becomes density-independent, so the temperature again increases coupled with a rise in the cooling rate. This causes an increase in inflow velocities, and by $z \simeq 18.5$, the central $100 M_\odot$ of gas exceeds the Bonnor-Ebert critical mass at that radius, which is indicative of unstable collapse [40, 41]. Interestingly enough, it is found that the collapsing cloud does not fragment into multiple cores, which is the result expected by analytical treatment. [26, 20] Instead, they find that a single protostar of $\sim 1 M_\odot$, made completely of molecular hydrogen, forms at the center of the $100 M_\odot$ core. In addition, the core is not rotationally supported, meaning that it will collapse on a timescale determined by the cooling processes of the gas. The final mass of the star remains unclear, since the simulations lack the necessary physics to compute how much of the available cool material surrounding the protostar will accrete or at what point feedback from the protostar will limit further accretion. At the rate of infall when the simulation stopped, roughly $70 M_\odot$ of matter would be accreted in the following 10^4 years, with a maximum of $600 M_\odot$ in the following 5×10^6 years. Though the maximum value of $600 M_\odot$ is exceedingly unlikely (the main-sequence lifetime of a star weighing $\sim 100 M_\odot$ is much less than 5×10^6 years), this does point towards a top-heavy initial mass function for Population III stars.

Bromm, Coppi & Larson [42] pursue the same avenue of research with a complementary method (smoothed particle hydrodynamics) and find somewhat different results. They initialize their simulations with a top-hat overdensity using similar cosmological parameters to ABN, and set this initial top-hat configuration into rigid-body rotation with a given angular velocity in order to simulate tidal interactions with nearby clumps. They find that their gas clumps evolve similarly to those of ABN, which is unsurprising as the dynamics of the halo collapse is dominated by the physics of H_2 formation and cooling. However, they find that their halo develops a very lumpy, filamentary substructure with several sub-clumps, each of which which individually evolve in a manner comparable to the single halo in ABN and end up with a $\sim 100 M_\odot$ core of cool gas in a state of semi-free fall in the center of each sub-clump. More detailed information on the core is unavailable due to lack of simulation resolution. Bromm et al. also suggest that Population III star formation might have favored very massive stars.

More recently, Gao et al. [43] and Reed et al. [44] have performed a series of dark matter-only simulations where they use a sequence of nested N-body simulations to follow the growth of progenitors of the most massive object in a $\sim 500 \text{ Mpc}/h$ volume. They use a sequence of nested re-simulations to “zoom in” on this object to study the environment

and merger history of the halos in that area. The first object capable of forming stars is believed to collapse at $z \simeq 47$, when the mass of this halo is $\simeq 2.4 \times 10^5 M_{\odot} h^{-1}$. Halos forming in this environment are significantly overabundant and also undergo rapid mergers compared to a more “average” part of the universe. This leads to the rapid growth of halos in this region – the largest reaches a mass of $\sim 5 \times 10^7 M_{\odot}$ at $z = 29$. These authors suggest that by $z = 30$ a substantial population of primordial objects are capable of forming Population III stars, and that by this time small “galaxies” with $T_{vir} > 10^4$ K (that are able to cool effectively by atomic hydrogen) will also exist. These authors also note that halo populations, merger rates and correlation scales in their simulations are well-modeled by the extended Press-Schechter formalism at all times, which is promising.

It should be noted that, as stated above, the simulations discussed by Gao et al. and Reed et al. are N-body calculations and do not contain baryonic physics. The results that predict, e.g., the redshift of first star formation, should be taken to be approximations only. An important lesson, however, is that the simulations which include baryonic physics use box sizes that are most likely too small to adequately model the scatter in star formation times (which will be discussed in later sections of this work). This is a reasonable and correct issue, and is investigated in this thesis.

The reader desiring a more thorough review of Population III star formation is directed to reviews by Bromm & Larson [45] and Ciardi & Ferrara [46].

1.3.5 The Initial Mass Function of primordial stars

One of the most interesting open questions relating to the issue of primordial stars concerns their initial mass function (IMF). If these stars are very massive, they will be copious emitters of UV radiation and produce large amounts of metals, as discussed in Section 1.4. This problem, however, is a very difficult one to solve, owing to the range of physics involved.

Abel et al. [39] are unable to follow the evolution of the fully molecular protostar that forms in the center of their halo to the point where it moves onto the main sequence. Their simulations are terminated due to a lack of appropriate physics – namely, the optically thin radiative cooling approximation for primordial gas breaks down at $\sim 10^{14} \text{ cm}^{-3}$. This can be extended another few orders of magnitude using analytical approximations to the primordial cooling function [47], but eventually full radiation transport will be necessary. Regardless, Abel et al. estimate the mass range of the protostar by examining the spherically-averaged accretion rate onto the $1 M_{\odot}$, fully molecular protostar that had formed by the end of their simulation. Based on the observed accretion rates, they observe that at least $30 M_{\odot}$ of gas will have accreted onto the central core in a few thousand years, which is much shorter than expected protostellar evolution times. Approximately $200 M_{\odot}$ of gas will accrete in $\sim 10^5$ years, and a total of $600 M_{\odot}$ will accrete in 5×10^6

years. It is implausible that this mass will be reached, since the lifetime of primordial stars in the mass range of hundreds of solar masses is only 2 – 3 million years. They suggest that a more reasonable mass range is 30 – 300 M_{\odot} , and no attempt is made to state the possible distribution of masses in this range.

The other 3D simulations that have been performed are by Bromm et al. Their earlier simulations do not have the mass and spatial resolution to estimate the masses of these stars – however, they suggest that the first generation of stars may have been with massive, with $m_* \geq 100 M_{\odot}$ [42]. A later simulation by Bromm & Loeb [48] improves upon this calculation and follows the evolution of a primordial protostar down to a scale of ~ 100 AU. They find a conservative upper limit of $m_* \leq 500 M_{\odot}$ and suggest that the actual stellar mass is likely to be significantly lower than that due to feedback from the protostar.

Though useful in many ways, and the final arbiter of the Population III IMF, 3D calculations of Pop III star formation in a cosmological context are limited by computational costs and the physics packages currently implemented in them. Fully 3D calculations of accretion onto the primordial protostar, including all relevant physics such as multifrequency radiation transport, accurate models of the primordial protostar, magnetohydrodynamics, and a full nonequilibrium chemical reaction network are in principle technically feasible, but the computational cost for doing such a calculation is prohibitive at best. One could wait for computational resources to increase to the point where this sort of calculation is reasonable, but more impatient (and practical) researchers have resorted to analytical and one and two-dimensional numerical models.

Tan & McKee and Tan & Blackman [49, 50] have created theoretical models of the evolution of the Population III protostar as it moves onto the main sequence. They combine a range of assumptions about the strength of magnetic fields generated in the protostellar disk (as well as their efficiency in transporting angular momentum) with estimates of the disk structure, gas infall rates and protostellar evolution models to gain some understanding of the radiative feedback from the protostar and its role in shutting off accretion. Based on accretion rates from Abel et al. [39] and from their calculations showing feedback is dynamically unimportant for protostars with masses $< 30 M_{\odot}$, they conclude that the masses of these primordial stars should be at least $30 M_{\odot}$.

Omukai & Nishi [51] performed calculations modeling the hydrodynamical evolution of primordial, spherically symmetric clouds taking into account chemistry as well as continuum and molecular hydrogen line radiative transfer. They find (similarly to Abel et al.) that a $\sim 1 M_{\odot}$, fully molecular protostar forms in the inner region of their calculation. However, they see that as accretion continues and central densities climb, the molecular hydrogen in the core dissociates and a hydrostatic core with mass $M_{core} \sim 5 \times 10^{-3} M_{\odot}$ forms at the center of the cloud, with gas accreting onto it at $\sim 10^{-2} M_{\odot}/\text{year}$. The accretion rate declines with time. They make no estimate of the final range of stellar masses. Later work by Omukai and various collaborators [52, 53, 54] predict upper mass

limits for massive primordial stars that range from $300 - 1000 M_{\odot}$. All of these works assume spherical symmetry and high, time-dependent accretion rates ($\sim 10^{-2} M_{\odot}/\text{year}$ initially, decreasing as a function of time), with the upper limit depending strongly on assumptions regarding the evolution of the protostar, the strength and efficiency of radiation from the star in halting accretion, and the accretion rates onto the stars.

The results of Omukai et al. are in sharp contrast to the calculations of Nakamura & Umemura [55, 56, 57]. They perform one and two-dimensional hydrodynamic simulations coupled with nonequilibrium primordial chemistry and follow the evolution of the clouds from a central density of $\sim 10^2 \text{ cm}^{-3}$ up to $\sim 10^{13} \text{ cm}^{-3}$. They observe that the star-forming clouds tend to fragment out of filaments, and therefore choose to simulate these objects using cylindrical symmetry. In their earlier work [55] they perform one-dimensional cylindrically symmetric hydrodynamic calculations that neglect all deuterium-related chemistry and cooling and observe that the typical mass of their central object is $\sim 3 M_{\odot}$ (though they state that it could grow to be approximately five times that mass via accretion) over a wide range of input assumptions about cloud temperature and other properties. Later calculations [56] were performed in both 1D and 2D, again assuming axial symmetry, and improve upon the previous result. These calculations show that the initial density of the filaments in their problem setup strongly affects the scale at which fragmentation occurs, and they posit that the IMF of Population III stars is likely to be bimodal, with peaks at ~ 1 and $\sim 100 M_{\odot}$, with the relative numbers of stars in each peak possibly being a function of the collapse epoch. They also perform 1D calculations including deuterium chemistry [57] and show that due to the enhanced cooling from the HD molecule there is still a bimodal distribution with a low-mass peak of $\sim 1 - 2 M_{\odot}$, though the high mass peak can now be somewhere between $10 - 100 M_{\odot}$, depending on the initial filament density and H_2 abundance.

This section has shown that there is both agreement and disagreement between different groups' results. All of the research discussed here indicates that the IMF of Population III stars is wildly different than the IMF of stars in our galaxy at the present day, with the mean stellar mass being significantly higher in primordial stars than at the present epoch. The disagreement lies in both the shape of the Population III IMF and in the mean mass of the primordial stars.

What is the root of this discrepancy? The simulations performed by Abel et al. [39] (and myself, as discussed later in this thesis) show that the accretion onto the protostar is not inherently one-dimensional – we typically see the formation of a generally spherical core forming in the center of the halo, though there is evidence for angular momentum-transporting turbulence within this core. Also, the cosmological structures that these halos form out of are inherently aspherical. This suggests that 1D models are missing crucial physics. Additionally, it seems apparently that modeling the interplay between radiation from the growing protostar and the accreting gas is going to be very important, and this must be done carefully. It also appears that complete modeling of

the primordial gas (including deuterium, lithium and the various molecules they form between themselves and with hydrogen) may be important in 2D and 3D simulations. It may be that magnetic fields also play a significant role in angular momentum transport at scales corresponding to the size of the forming protostar, so they must be included in simulations and analytical models. Though expensive, these calculations are possible in 2D, and will be feasible in 3D in a few years, assuming that the power and availability of computing resources continues to grow at similar rates to today. At that point, accurately predicting the IMF of primordial stars may be a tractable problem.

1.4 Feedback from the first generation of stars

The formation of the first stars marks the transformation of the universe from its almost completely smooth initial state to its clumpy, complicated current state. As the first luminous objects to form in the universe, Population III stars play an extremely important role by fundamentally changing the environment that later cosmological structures form in, through radiative, mechanical, and chemical feedback. The literature discussing the feedback properties of Population III stars is vast and rapidly evolving. This section will provide only a brief overview of these properties, and the interested reader is encouraged to refer to the 2001 review article by Barkana & Loeb [58] or the much more recent review by Ciardi & Ferrara [46].

1.4.1 Physical properties of Population III stars

As discussed in Section 1.3, recent analytical work and numerical simulations using a range of initial conditions and assumptions about relevant physics suggest that Population III stars may have an exceedingly top-heavy IMF compared to stars in our galaxy. Massive primordial stars are believed to have several interesting properties that distinguish them from stars with a significant fraction of metals: these stars are extremely compact and, as a result, have very high effective temperatures (approximately 10^5 K). As a result, Population III stars have rather hard spectra and produce large numbers of both hydrogen and helium-ionizing photons. Additionally, due to the lack of metals in these objects, they are expected to have little mass loss near the end of their main-sequence lifetime due to line-driven winds. See Schaerer [59] and references therein for a more complete review.

The final fate of these stars is also quite remarkable. Recent one-dimensional, nonrotating simulations of the evolution of massive primordial stars suggest that at the end of their lives, the more massive of these stars ($M_* \geq 30M_\odot$) which typically collapse directly into a black hole ($M_* \sim 30 - 100$, $M_* > 260M_\odot$) or explode in a massive pair instability supernova (PISN; $M_* \sim 140 - 260 M_\odot$), which would completely destroy the star, leaving no compact remnant behind. These supernovae can be almost two orders of magnitude

more energetic than a standard Type II supernovae and also leave behind a very distinct nucleosynthetic signature [60]. A middle range of extremely massive stars is believed to have an energetic pulsational instability that causes the ejection of much of its envelope before collapsing into a black hole. The low-mass end of the Population III stellar IMF would see behavior more comparable to that seen by dying stars in the local universe – collapse to a white dwarf preceded by asymptotic giant branch-type activity or a Type II supernova resulting in a neutron star or black hole compact remnant. See Heger et al. [61] and Figure 1.3 for more information on the the fates of Population III stars over a large mass range. The black hole remnants of extremely massive Population III stars have been suggested as seeds for the super massive black holes (SMBHs) that have been observed in the centers of essentially all large galaxies [62, 63].

1.4.2 Radiative feedback

As mentioned in Section 1.4.1, recent numerical work indicates that Population III stars are copious emitters of ultraviolet radiation. In addition, it has been noted that the black hole remnants from these stars may produce large amounts of x-rays. What are the possible effects of the radiative feedback from these stars and their remnants?

Observations of the polarization of the cosmic microwave background by the Wilkinson Microwave Anisotropy Probe (WMAP) satellite have detected excess power on large angular scales compared to predictions based solely on the temperature power spectrum. This result is consistent with a period of partial reionization of the intergalactic medium taking place at redshifts of $11 < z < 30$ [64]. Observations of high-redshift quasars have shown that the universe was fully ionized at $z=5.8$ [65]. Recent calculations suggest that complete reionization occurred some time between $z = 7$ and $z = 12$ [66, 67, 68], though the exact epoch of reionization in these simulations is sensitive to a number of highly uncertain parameters, such as the formation efficiency of stars and quasars and the escape fraction of ionizing photons produced by these sources. Also, it is apparent that the topology of reionization is very complex - the epoch of reionization starts as “patches” of ionized material around the first stellar objects and spreads as structure evolves and more UV-emitting massive stars form [68]. In addition, the regions of highest gas density (which harbor the stars producing ionizing radiation) also contain significant amounts of neutral gas. Though the first stars to form are prodigious UV emitters, they are relatively few in number and quite short-lived. This makes the most likely scenario one where Population III stars are responsible for a partial reionization of the universe, and the structures that form from material polluted by these stars, namely, the first galaxies and pre-galactic objects (PGOs), is responsible for the final reionization of the universe, an idea that was first put forth by Cen [69] and supported by Hui & Haiman, who show that an early epoch of reionization would have to be followed by some cooling and recombination, or else the IGM would have a significantly different temperature than is




Figure 1.3: Stellar endpoints for zero metallicity stars, as a function of mass. This images were calculated using the 1D stellar evolution code KEPLER, and were done assuming a nonrotating model. Results may quantitatively change when rotation, convection and other physical effects are added. Figure courtesy of Alex Heger (Theoretical Astrophysics (T-6), Los Alamos National Laboratory).

observed today [70].

The primary cooling agent in the formation of massive primordial stars is molecular hydrogen. This molecule is quite fragile – its formation is dependent on the availability of H^- , the formation of which in turn depends on the availability of free electrons, and it can be easily destroyed by radiation in the Lyman-Werner band, which ranges from 11.18 to 13.6 eV (corresponding to the soft UV band). Since this is below the ionization threshold of atomic hydrogen, photons in this energy band can propagate great distances in the IGM. Also, since Population III stars appear to be prolific emitters of UV radiation, they will build up a background of this soft UV light, which may cause the overall dissociation of molecular hydrogen, halting the epoch of Population III star formation until more massive halos, whose virial temperatures are high enough that the gas can cool effectively by atomic hydrogen (e.g. $T_{vir} > 10^4$ K) have time to form [72, 73]. At this point, the clouds can continue to collapse and eventually produce primordial stars [74], which may have a different mass spectrum than Population III stars that form in minihalos with masses of $\sim 10^6 M_\odot$ [75]. Simulations have shown that this soft UV radiation is quite effective in suppressing the formation of Population III stars in halos with masses of $\sim 10^5 - 10^6 M_\odot$ [71].

The HII regions produced by massive primordial stars may also have a significant effect on star formation. Whalen et al. [76] show that the I-fronts from massive primordial stars can propagate several proper kpc in the high-redshift intergalactic medium, ionizing large volumes of space. These stars will also heat the gas in their parent halo, typically resulting in the majority of the baryons in a $10^6 M_\odot$ halo to be driven out of the halo at speeds of up to ten times the escape velocity of that halo. Oh & Haiman [77] suggest that these HII regions will suppress the formation of any further stars in that region. However, work presented in this thesis shows that the high electron fraction produced as a result of ionization actually promotes the formation of molecular hydrogen, and in halos with a density above some critical density this can actually result in a positive feedback process where stars would form in halos that otherwise would not experience star formation. It has also been demonstrated that HD (deuterium hydride) can be a significant source of cooling in star-forming sites in fossil HII regions, allowing the gas temperature to drop even lower than via cooling due to molecular hydrogen alone would allow (below 100 K) and possibly resulting in low mass Population III stars [78].

Accretion onto the black holes formed by the collapse of Population III stars may be a source of significant x-ray radiation in the early universe. It has been suggested that this radiation is at least partially responsible for the WMAP polarization result [79, 80, 81]. Additionally, the soft x-ray background will produce a significant free electron fraction, that may result in a positive feedback effect on the formation of Population III stars by spurring the creation of molecular hydrogen [82]. This is not a certainty – depending on assumptions about the associated soft UV background and the hardness of the x-ray spectrum from the Population III black holes, the feedback effects may only be a weak

positive effect, or even negative [83, 77].

1.4.3 Chemical feedback

Observations of quasar absorption spectra show that the universe at the present day is uniformly polluted with metals, even at the lowest observed column densities, which correspond to regions of very low overdensity commonly referred to as the Lyman- α forest [84, 85]. The primordial composition of the universe is well understood, and post-BBN nucleosynthesis is believed to take place only in stars and the cataclysmic events associated with them. Because of this, it is apparent that this period of enrichment must have taken place between the epoch of first star formation and the present day.

As with reionization, it is unclear which objects are responsible for the majority of metals in the low-overdensity universe: The most massive galaxies in the early universe, often referred to as Lyman Break Galaxies (LBGs), are sites of vigorous star formation and metal production and characterized by strong, metal-rich galactic outflows and high luminosities [86]. However, being massive, they have deep potential wells which might serve to trap the ejected materials. Also, since these objects are the most massive bound objects in the early universe, the theory of hierarchical structure formation tells us (and observations support the assertion) that these galaxies are few and far between, so metals produced by them would have to be transported cosmologically significant distances in order to be as homogeneously distributed (as indicated by observations). The other possible candidate for homogeneous metal enrichment in a Λ CDM scenario would be pre-galactic objects and the first dwarf galaxies. While they are much smaller than LBGs, with correspondingly smaller star formation rates, they have shallower potential wells which would allow outflowing material to escape much more easily [87]. In addition, these smaller objects begin to form much earlier and there are many more of them than LBGs, so metal outflowing as winds from these galaxies or released via ram pressure stripping during the frequent galaxy mergers demanded by the hierarchical structure formation scenario would have more time to be distributed and also be required to travel a much more reasonable distance from their point of origin to obtain the observed relatively homogeneous distribution of metals [88].

The metals produced by Population III supernovae would have a very important effect on the following generations of stars. They enhanced the cooling properties of the gas significantly – molecular hydrogen is a relatively poor coolant compared to dust grains, which are believed to be produced in significant quantities by both primordial Type II supernovae and pair-instability supernovae, with the fractional quantity of dust increasing as the stellar mass increases [89, 90, 91]. Very little metal is required for gas to cool efficiently – analytical work and simulations suggest that the presence of carbon and oxygen at levels $10^{-4} - 10^{-3} Z_{\odot}$ would be sufficient for enhanced fragmentation of collapsing gas clouds, signifying a change from the top-heavy Population III IMF to

a mass function resembling that observed in the galaxy today [93, 94]. As with their HII regions, the metals ejected from Population III supernovae, particularly if the stars fall into the mass range that produces highly energetic pair-instability supernovae, can propagate to great distances – simulations indicate that the ejecta from a massive PISN can eject metal into a sphere of ~ 1 kpc diameter at $z \sim 20$, producing a metallicity floor above that needed for enhanced cooling to take place [95].

It is doubtful that individual Population III stars can be observed directly during their main sequence lifetime, even by the James Webb Space Telescope (JWST), which is scheduled to be launched in 2011. However, it has been suggested that extremely massive Population III stars may be the progenitors of gamma ray bursts, and as such may be observable to very high redshifts [96]. The predicted rates of Population III supernovae suggest that their gamma ray bursts may be observable by the SWIFT satellite at the rate of approximately one per year, and that Population III supernovae may be observable by JWST at the rate of $4 \text{ deg}^{-2} \text{ year}^{-1}$ at $z \sim 15$, with a high level of uncertainty in these calculations [97, 98]. The nucleosynthetic yields of these stars may have already been detected in observations of the abundance ratios of two extremely metal poor stars, which have Fe/H ratios of $\sim 10^{-5.5}$ [99, 100]. Both of these stars show extreme overabundances of carbon and nitrogen with respect to iron, which suggests a similar origin of the abundance patterns. However, these abundance patterns do not agree with theoretical predictions for yields of Population III supernovae, so their origin is uncertain, though it has been suggested that these results can be naturally explained as the concurrent pollution of at least two supernovae of relatively low mass [101]. Finally, the coalescence of black hole remnants of Population III stars may be directly detected by gravitational interferometers such as Advanced LIGO [102], and indirectly by their contribution to the near-infrared background excess [103, 104].

1.5 Flaws in the Λ CDM paradigm

It is acknowledged that there appear to be flaws in the Λ CDM scenario. Observations of gravitational lensing by galaxy clusters indicate that the dark matter profile in the center of these clusters forms a smooth core, while theory and simulations using the Λ CDM model suggest that there should be a dark matter ‘cusp’ [105]. This cusp is also predicted to be seen in galaxies, but has not been observed [106, 107]. In addition, the CDM model predicts the formation of a large number of dwarf galaxies, and also suggests that these tiny galaxies will form in the cosmic voids - a prediction that has not been verified observationally [108, 109]. In addition, there have been observations of significant numbers of dwarf galaxies forming after the larger Lyman break galaxies, which is not what one would expect in a hierarchical clustering scenario [110].

This is not to say, however, that the CDM model is fatally flawed. The model has

done an excellent job of predicting the formation and evolution of large scale structure [111], as shown in recent years by extensive surveys of the local universe (such as the 2dF and SDSS surveys). The flaws in the dark matter models are on the sub-galactic scale [106, 112] – too small to affect the formation and evolution of large-scale structure, but certainly significant for the evolution of the first structures in the universe. Also, some work has been done recently which shows that the dark matter cusp issue in galaxies can be resolved [113]. A complete abandonment of the CDM paradigm seems premature.

There are several proposed solutions to the problems with the Λ CDM model, which include decaying dark matter [114], warm dark matter [115], collisional dark matter [116], annihilating dark matter [117] and fuzzy dark matter [118]. The essential feature of the majority of these models is that they suppress the formation of low-mass cosmological halos and can significantly alter predictions for early structure formation. The impact of the suppression of small-scale power by a generic warm dark matter model on cosmic structure formation in the early universe is explored in Chapter 5.

Chapter 2

Simulation Methodology

2.1 The generation of cosmological initial conditions

Creating a set of initial conditions for simulations of cosmological structure formation is, on the surface, a very straightforward task. One specifies a background cosmological model, typically described as a spatially flat or open Robertson-Walker spacetime. Following that, perturbations are imposed upon this background assuming a triply periodic, finite simulation volume. In reality, numerous approximations must be made which should be carefully considered, and are discussed in some detail below.

The specification of background cosmology requires several inputs: The amount and nature of dark matter, the Hubble parameter H_0 , and possibly the amount of baryonic matter and cosmological constant in the universe. Most of these quantities are typically specified in terms of the standard cosmological parameters: Ω_m , Ω_{dm} and Ω_b , which correspond to the total overall mass density, dark matter density and baryon density at the present epoch as a fraction of the critical density, $\rho_c \equiv 3H_0^2/8\pi G$. The cosmological constant is specified as Ω_Λ , which is the vacuum energy density at the current epoch as a fraction of the critical energy density, which is simply $\rho_c c^2$. The perturbations of dark matter and baryons are specified as a power spectrum which has one or more parameters. At the very least, a power spectrum has an index n which specifies the shape of the primordial power spectrum.

At the epoch of baryon-photon decoupling ($z \sim 1100$), small-amplitude (“linear”) fluctuations in density are already present in all of the components that contribute to the energy density of the universe (such as baryons, dark matter, photons, and neutrinos). The statistical nature of these fluctuations depends on their origin. There are two general classes of early universe models that are considered to provide reasonable mechanisms for perturbations: topological defects [119] and inflation [120]. Inflation predicts fluctuations that obey Gaussian statistics and defect models predict non-Gaussian fluctuations.

Gaussian fluctuations are simple since they are specified completely by a single func-

tion, the power spectrum $P(k)$. In Gaussian models the perturbations are set down almost immediately after the Big Bang (during the inflationary epoch in the canonical model of inflationary cosmology, roughly 10^{-35} seconds after the Big Bang) and evolve according to the linearized Fokker-Planck equation. In real space, the probability distribution of density fluctuations is a multidimensional Gaussian, and it is very easy to sample a Gaussian random field by sampling its Fourier components on a Cartesian lattice, which is the technique that will be discussed here. For more information on other methods, see Bertschinger’s review on simulations of cosmological structure formation [121].

Non-Gaussian models are much more complicated to model. Not only do they require more initial information than a simple power spectrum, they also are more costly in a computational sense. Typically, topological defects induce matter density fluctuations from the time of their creation in the early universe to the present day, and the dynamics of their formation and evolution are relativistic and nonlinear. For more information on creating initial conditions of topological defects, see Bertschinger [121] or Durrer et al. [119].

A simple test of the Gaussianity of the primordial density perturbations can be made by examining the power spectrum of the cosmic microwave background. The power spectrum of temperature fluctuations in the CMB was imposed at a very early time ($z \sim 1100$, approximately 300,000 years after the Big Bang) when all perturbations were linear (and are at most $\delta\rho/\bar{\rho} \sim 10^{-5}$, where $\bar{\rho}$ is the mean density and $\delta\rho \equiv |\rho - \bar{\rho}|$), significantly before the epoch of cosmological structure formation (which is when density perturbations become nonlinear and gravitationally self-bound). Recent observations of the CMB using the Wilkinson Microwave Anisotropy Probe (WMAP) satellite have examined the CMB for signs of non-Gaussianity and have presented upper limits on the amplitude of non-Gaussian primordial fluctuations using two separate tests [122]. These observations show to a high level of confidence that the WMAP data is consistent with Gaussian primordial fluctuations.

Due to the preponderance of evidence in favor of a Gaussian perturbation spectrum, all simulations described in this thesis will assume this spectrum. Non-Gaussian initial conditions (i.e. topological defect models) will most likely produce significantly different results for structure formation on the scales of interest.

2.1.1 Creating Gaussian random fields

The creation of cosmological initial conditions using a Gaussian random field is relatively straightforward. Given a power spectrum $P(k)$, the linear density fluctuation field is calculated at some initial time when density perturbations are highly linear ($\delta\rho/\bar{\rho} \ll 1$; a typical choice of starting redshift is $z \sim 100$ for high-resolution cosmological simulations). From this, dark matter particle positions and velocities are determined, along with baryon

velocity fields, as described below.

The Linear Density Fluctuation Field

The first step towards creating a Gaussian density field is to specify a power spectrum, which is a function that relates amplitudes of density perturbations with wavelength λ to their wavenumber k , where $k \equiv 2\pi/\lambda$. The power spectrum of the fractional density fluctuations at the redshift $z = z_{eq}$ when the energy density in matter is equal to that in radiation, can be related to the primordial power spectrum by $P(k, z_{eq}) \sim T^2(k)P_p(k)$, where $T(k)$ is the matter transfer function as a function of wave number and describes the processing of the initial density perturbations during the radiation dominated era (see Padmanabhan [123]) and $P_p(k)$ is the primordial matter power spectrum, which typically has a power law form, i.e., $P_p(k) \sim k^n$, where n is the index of the primordial power spectrum. This index is equal to unity for Harrison-Zel'dovich scale-invariant spectra, a typical model [124, 125]. The power spectrum at any redshift z in the matter dominated era may then be written in the form

$$\frac{k^3}{2\pi^2}P(k, z) = \left(\frac{ck}{H_0}\right)^{3+n} \delta_H^2 T^2(k) D_g^2(z)/D_g^2(0), \quad (2.1)$$

where the D_g 's are the linear growth factor for perturbations, which is defined in Peebles [126] and many other places. A closed-form fitting function (much more appropriate for computation) is given in Eisenstein & Hu [194]. δ_H is a constant describing the amplitude of density fluctuation, which can be provided from observations of the CMB or from large scale structure, or can be normalized by comparing to, e.g., σ_8 , which is the RMS amplitude of the mass fluctuations in the universe when smoothed using a top-hat function with characteristic radius of $8 h^{-1}$ Mpc.

Once $P(k)$ has been determined, we then proceed to calculate δ_k , namely, the density fluctuations in k -space. To simplify matters, we choose a cubic three-dimensional Cartesian grid with N grid points per dimension, though in principle this set of operations can be calculated for any rectangular solid lattice with some small additional complications. Each of the grid points has a unique identifier (n_x, n_y, n_z) associated with its location along the (x, y, z) axis. We sample the power spectrum $P(k)$ discretely at each grid location (n_x, n_y, n_z) , obtaining k in this manner:

$$k^2 = (n_x^2 + n_y^2 + n_z^2)dk^2 \quad (2.2)$$

with $dk = 2\pi/L_{box}$, and L_{box} is the size of the simulation box in h^{-1} Mpc. It is important to note that the method of sampling $P(k)$ at discrete points, while convenient when doing simulations on a Cartesian grid, assumes that $P(k)$ is a smoothly and slowly varying function of k on the physical scales relevant to the simulation (namely, on scales

encompassing the simulation volume and the size of the grid lattice). If this is not true, large errors will be introduced unless $P(k)$ is sampled in a different way. We also assume that the simulation volume is a finite size and assume periodicity - i.e., an object that travels out of one side of the box re-enters it on the other side. This assumption implicitly states that the simulation volume is an “average” patch of the universe.

The k-space fluctuation δ_k is a complex value with random phase and amplitude, where the distribution of amplitudes fit a Gaussian distribution with a mean of $P(k)$. One method to calculate it is to generate a phase angle θ , which is randomly selected in a uniform manner in the interval $[0, 2\pi]$, and an amplitude A such that

$$A = \sqrt{-\log(R) * P(k)} \quad (2.3)$$

where R is randomly selected in a uniform manner in the interval $(0,1)$. δ_k is then

$$\delta_k = Ae^{i\theta} \quad (2.4)$$

We then perform a Fourier transform on the grid of δ_k values, giving δ_x , the relative density fluctuation at each spatial grid point in the simulation volume. The actual physical (meaning, in position space rather than k-space) dark matter density is then

$$\rho_{DM}(\mathbf{x}) = (1 + \delta_x(\mathbf{x}))\bar{\rho}_{DM} \quad (2.5)$$

where $\bar{\rho}_{DM}$ is the mean dark matter density in the simulation. The perturbations in the dark matter and baryon densities are assumed to be coupled, which is a reasonable assumption in the linear regime, so the baryon density at any position is

$$\rho_b(\mathbf{x}) = \frac{\Omega_b}{\Omega_{DM}}\rho_{DM}(\mathbf{x}) \quad (2.6)$$

where Ω_b and Ω_{DM} are the ratios of the present-day mean baryon and dark matter densities to the critical density of the universe, i.e., $\Omega_b = \rho_b/\rho_c$. The assumption of complete baryon-dark matter coupling, though reasonable, is not strictly true at the redshifts at which cosmological simulations are typically initialized due to the dissipative nature of the baryons. Therefore, it is more accurate to use separate transfer functions for the dark matter and baryon components, though in practice this only makes a small difference for a certain class of simulations. For this reason we typically ignore it when generating our initial conditions.

Position and Velocity Fields

In cosmological simulations, dark matter is usually represented by discrete particles rather than a uniform grid of densities. In that case, these particles must be perturbed so as to reflect the density fluctuations described in Section 2.1.1. The standard approach for the dark matter is to displace equal-mass particles from a uniform Cartesian lattice using the Zel'dovich approximation [127]:

$$\mathbf{x} = \mathbf{q} + D(a)\psi(\mathbf{q}) \quad (2.7)$$

and

$$\mathbf{v} = a \frac{dD}{dt} \psi = aHfD\psi, \quad (2.8)$$

where \mathbf{q} labels the unperturbed lattice position, $D(a)$ is the growth factor of the linear growing mode, H is the Hubble parameter, a is the cosmological scale factor, and $f = d \ln D / d \ln a$ is its logarithmic growth rate [126]. The irrotational (curl-free) displacement field ψ is computed by solving the linearized continuity equation,

$$\nabla \cdot \psi = -\frac{\delta_x}{D(t)}, \quad (2.9)$$

Since the equation is linearized, ψ can be found by taking the Fourier transform of $-i\delta_k \hat{k} / kD(k)$. This approximation is then also used to calculate the baryon velocities (using the baryon perturbation field if a different transfer function is used, or a scaling factor for δ_x and δ_k otherwise).

2.1.2 Creating initial conditions for high dynamical range simulations

Compromises must be made when one is performing cosmological simulations of large-scale structure formation. In order to get a statistically accurate representation of large scale structure, one must use as large of a spatial volume as possible. The dark matter particle mass and ultimate spatial resolution of the simulation must be chosen based on the structures of interest, so that any given object (e.g. galaxy, galaxy cluster, etc.) is resolved by a large number of dark matter particles and by many grid cells. Unfortunately, given finite computational resources (in terms of both memory and processing power) these two criteria are often in direct competition. For example, in the simulations discussed later in this thesis, it is important to choose a box that is of large enough size to contain at least one $5 \times 10^5 M_\odot$ dark matter halo, which gives us a minimum box size that is approximately $0.3 h^{-1}$ Mpc on a side. We also need to resolve the region around

this halo with dark matter particles that are small compared to the halo mass and the final protostar, typically on the order of a few solar masses for a calculation of this type. Given that $M_{dm} = \Omega_{cdm}\rho_c(L_{box}/N_{grid})^3 M_\odot/h$ (where L_{box} is the box size in Mpc/h and N_{grid} is the number of grid cells along one edge of the box), this results in a volume with approximately 1024^3 grid cells and dark matter particles. This is computationally wasteful, since we only need to resolve a few percent of the box with high spatial and mass resolution. It is important to resolve tidal forces from other structure in the simulation volume, but for this particular application they do not have to be resolved with either the spatial or the mass resolution of our object of interest.

The solution to this is to use “nested” initial conditions, which are created by nesting subvolumes of higher spatial and mass resolution within a low-resolution parent grid. The low-resolution parent grid provides necessary tidal forces and a cosmological volume which provides reasonable statistics, and the nested subgrid, which is typically only a small fraction of the overall volume, provides higher mass and spatial resolution in the area of interest. This can be done recursively and can significantly save computational resources, particularly when the grids are further refined with adaptive mesh only in the most highly refined subvolume. See Figure 2.1 for an example.

An example of the overall cost savings can be shown very simply. Let us assume that the situation is as described for the simulation discussed above, where the overall ratio of box size to spatial resolution needed in the region of interest is $\sim 10^3$. Furthermore, we will assume that each subgrid has twice the spatial resolution as the one above it, suggesting that a grid at level L has to take twice as many timesteps as its parent grid in order to advance the same physical time. This assumption about time steps is reasonable for simulations using an explicit hydrodynamics algorithm, where $\Delta t = \kappa\Delta x/c_s$, where Δx is the grid size, c_s is the local sounds speed, and κ is a “Courant parameter” which has a value $0 < \kappa \leq 1$ and a typical value of $\kappa \simeq 0.3 - 0.5$. For the sake of comparison, we will compare the amount of time that it takes a simulation generated with a single, 1024^3 cell and particle initial condition to advance a fixed amount of time compared to a simulation with a 128^3 “root grid” and three 128^3 grids that are nested inside of each other, where the “root grid” resolves the same volume as the monolithic 1024^3 grid, and the most highly refined grid in the nested grid gives an equivalent resolution of 1024^3 (i.e. the same as the large monolithic grid). This nested calculation is similar to those discussed later in this thesis.

Since the simulation with 1024^3 grid cells has 8 times better spatial resolution than the root grid of the nested calculation, it will take 8 timesteps of the monolithic grid calculation to equal a single timestep on the root grid of the nested calculation (approximately). This means that the monolithic calculation will need to calculate $N_{mon} = 8 \times 1024^3 \simeq 8 \times 10^9$ individual cell-steps (which is a reasonable, though crude, proxy for computational cost). In order to evolve the nested grid simulation the same amount of physical time, it needs to take one timestep on its root grid, two on the first subgrid,




Figure 2.1: Example of a set of nested-grid initial conditions. The image is of a slice through the initial conditions that is one grid cell thick on the finest level ($1/1024$ of the thickness of the entire box), and shows log baryon density. The overall fluctuations are at the 20% level throughout the volume. The root grid (“level 0”) is 128^3 grid cells, and each successive nested grid is shown by a black square. The Level 3 grid (central square) is the highest level of resolution and has the same resolution as a calculation with identical spatial extent to the entire simulation volume but 1024^3 grid cells. In calculations such as these adaptive mesh refinement is only allowed to take place within the highest level of nested-grid resolution, which in this case corresponds to only $1/64$ of the overall volume.

four on the second subgrid, and 8 timesteps on the most highly refined subgrid. This comes out to a total of $N_{nest} = 128^3 + 2 \times 128^3 + 4 \times 128^3 + 8 \times 128^3 \simeq 3.15 \times 10^7$ cell-steps. The relative nested grid calculation takes $N_{nest}/N_{mon} = 0.0037$ as many cell-steps as the monolithic simulation, which is over a factor of 200 in computational costs! The reduction in memory needed is also substantial: The nested grid calculation has a total of $4 \times 128^3 \simeq 8 \times 10^6$ cells, whereas the monolithic grid calculation has $1024^3 \simeq 10^9$ cells. This is a savings of approximately 100 in memory.

The cost savings of these simulations is obvious. Fortunately, it is extremely easy to create nested initial conditions. One simply recursively generates initial conditions as described above covering the same physical volume, but at higher and higher spatial resolutions. The coarsest set has the desired spatial resolution of the “root grid,” and each additional set of initial conditions is twice as highly resolved in space (resulting in dark matter particles that are 1/8 the mass of the particles in the next coarser set of initial conditions), and the subvolume of interest is extracted and put into a separate file. This method of generating initial conditions results in each grid having the appropriate perturbation spectra, and a much smaller overall simulation size in terms of computation and memory use.

The technique of “nesting” boxes of successively higher spatial and mass refinement described here is only limited by the available memory in which the necessary FFTs are performed. This limitation can be avoided by using an out-of-core initial conditions generation technique [128], or a more sophisticated method which does not require large, global FFTs to be performed [129]. These methods are costly in other ways, primarily in CPU time, but are necessary to be able to do calculations whose initial conditions have even higher dynamical range.

2.1.3 Numerical limitations of cosmological initial conditions

The standard methods and assumptions pertaining to simulations of early structure formation and the generation of cosmological initial conditions result in limitations which the careful simulator (or consumer of simulation-based data) must be cognizant of. These limitations can generally be broken up into large-scale effects which are due to finite box size and periodicity, and small scale effects due to finite resolution. Small-scale effects are dominated by the methods used in computation of forces, and the evolution of baryon fields. Therefore, we leave discussion of small-scale errors to later in this chapter (where the relevant numerical methods are discussed) and concentrate on large-scale effects here.

Truncation of the power spectrum due to effects of finite sampling causes issues at both large and small scales. Essentially, wave modes larger than half of the simulation ($L_{box}/2$) or smaller than twice the grid resolution ($2\Delta x$) are not represented correctly, which can have very important statistical effects. Gelb & Bertschinger [130] showed that truncation of power at the scale of the simulation box can result in estimates of σ_8 , which

corresponds to the size of density fluctuations when smoothed by a sphere $8 \text{ h}^{-1} \text{ Mpc}$ in radius, which are 40% lower than the input value for a simulation with a $50 \text{ h}^{-1} \text{ Mpc}$ box when using the previously-described method for generating initial conditions. In order to reduce the error to 10% the simulation volume had to double in size, to $\sim 100 \text{ Mpc}$. Pen [131] showed a method to fix this statistical issue, which generates initial conditions by convolving perturbations in real space instead of k-space. His method also allows the simple creation of nested-grid ICs without using large amounts of memory. This method was later implemented into a publicly-available package by Bertschinger [129]. An alternate method to correctly model large-scale power is discussed by Sirko [132], though in order to obtain correct statistics this method requires a large number of simulations to be performed, which is somewhat undesirable when performing large simulations due to overall computational expense.

A more concrete understanding of the effect of finite box sizes and periodicity on simulation results is provided by Barkana & Loeb [133]. They show that the majority of cosmological simulations, which are limited to relatively small volumes (particularly for study of objects in the early universe), largely omit cosmic scatter. This results in effects such as the epoch of reionization being shorter in simulations than has been observed and predicted analytically, non-physical biases in the halo mass function and halo correlations (i.e. lack of halo clustering), and overall mis-estimation of the halo mass function in small simulation volumes. This points towards a failing in small boxes of modeling the overall star formation rate at high redshift, as well as estimates of the amount and effects of feedback due to the ultraviolet radiation background at high redshift.

Though all of the issues discussed above are valid, we still must perform simulations of the formation of the first generations of stars in small simulation volumes. This is due entirely to the finite (though ever-growing) computational resources available to us. Extension of our computation capabilities by using nested grids (as described in Section 2.1.2) and techniques such as adaptive mesh refinement (described later in this chapter) allow us to increase our simulation volumes significantly, but not to the point where we can state with confidence that we are adequately resolving structure formation in the high-redshift universe with a high degree of statistical accuracy. This is not such a handicap, though, since we are interested in the dynamics of individual star forming clouds, and not the global star formation rate.

This begs the question, “if we can’t trust our statistics, what can be learned from this sort of simulation?” Fortunately, quite a bit. The simulations still accurately model the hydrodynamics, primordial gas chemistry, and gravitational interactions between baryons and dark matter (as well as self-interaction of these components). This allows us to simulate the evolution of individual star-forming halos very accurately. Additionally, we can obtain qualitative estimates of the reliability of our calculations by simulating a range of simulation volumes and extrapolating to larger box size. One can study the effects of, for example, feedback from a UV background on an individual halo, using a

single random realization but varying other parameters, to get an idea of the effects of various feedback processes in a general sense. In general, quite a bit can be done with these small-box simulations as long as one is mindful of possible effects due to finite box size, and assesses their results in light of this.

2.2 The Enzo Simulation Code

‘Enzo’ is an adaptive mesh refinement cosmological simulation code developed by Greg Bryan and his collaborators [134, 135, 136, 137, 2]. This code couples an N-body particle-mesh (PM) solver [138, 139] used to follow the evolution of collisionless dark matter with an Eulerian AMR method for ideal gas dynamics by Berger & Colella [140], which allows extremely high dynamic range in gravitational physics and hydrodynamics in an expanding universe.

Enzo solves the equations of ideal gas dynamics in a coordinate system that is comoving with the expanding universe:

$$\frac{\partial \rho_b}{\partial t} + \frac{1}{a} \mathbf{v}_b \cdot \nabla \rho_b = -\frac{1}{a} \rho_b \nabla \cdot \mathbf{v}_b \quad (2.10)$$

$$\frac{\partial \mathbf{v}_b}{\partial t} + \frac{1}{a} (\mathbf{v}_b \cdot \nabla) \mathbf{v}_b = -\frac{\dot{a}}{a} \mathbf{v}_b - \frac{1}{a \rho_b} \nabla p - \frac{1}{a} \nabla \phi \quad (2.11)$$

$$\frac{\partial E}{\partial t} + \frac{1}{a} \mathbf{v}_b \cdot \nabla E = -\frac{\dot{a}}{a} \left(3 \frac{p}{\rho_b} + \mathbf{v}_b^2 \right) - \frac{1}{a \rho_b} \nabla \cdot (p \mathbf{v}_b) - \frac{1}{a} \mathbf{v}_b \cdot \nabla \phi + \Gamma - \Lambda \quad (2.12)$$

Where ρ_b is the comoving baryon density, \mathbf{v}_b is the baryon velocity, p is the pressure, ϕ is the modified gravitational potential (in comoving coordinates, which is related to the potential in proper coordinates Φ by $\phi \equiv \Phi + 0.5a\ddot{\mathbf{x}}^2$) and a is the “expansion parameter” which describes the expansion of a smooth, homogeneous universe as a function of time. This expansion parameter is related to the redshift: $a \equiv 1/(1+z)$. All derivatives are in comoving coordinates. E is the specific energy of the gas (total energy per unit mass), and Γ and Λ represent radiative heating and cooling processes as described below. Equations 2.10, 2.11 and 2.12 represent the conservation of mass, momentum and total (e.g., kinetic plus thermal) fluid energy.

The equations above are closed with three more equations:

$$E = p/[(\gamma - 1)\rho_b] + \mathbf{v}^2/2 \quad (2.13)$$

$$\nabla^2 \phi = \frac{4\pi G}{a} (\rho_b + \rho_{dm} - \rho_0) \quad (2.14)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3a^3} (\rho_0 + 3p_0/c^2) + \Lambda/3. \quad (2.15)$$

where ρ_{dm} is the comoving dark matter density, ρ_0 is the comoving background density ($\rho_0 \equiv \Omega_{matter}\rho_{crit}$) and p_0 is the background pressure, γ is the ratio of specific heats and Λ is the cosmological constant. Equations 2.13, 2.14 and 2.15 are the equation of state, Poisson’s equation in comoving form and an equation that describes the evolution of the comoving coordinates (i.e. the formula for the expansion of an isotropic, homogeneous universe). All particles in the simulation are governed by Newton’s equations in comoving coordinates:

$$\frac{d\mathbf{x}_{part}}{dt} = \frac{1}{a}\mathbf{v}_{part} \quad (2.16)$$

$$\frac{d\mathbf{v}_{part}}{dt} = -\frac{\dot{a}}{a}\mathbf{v}_{part} - \frac{1}{a}\nabla\phi \quad (2.17)$$

Where \mathbf{x}_{part} and \mathbf{v}_{part} refer to the position and peculiar velocity of any particles in the system. Note that the system of equations 2.10-2.17 is valid only in regimes where relativistic effects are not important ($v_b, v_{dm} \ll c$, where c is the speed of light) and where cosmological curvature effects are unimportant, meaning that the simulation volume is much smaller than the radius of curvature of the universe, as defined as $r_{hub} \equiv c/H_0$, where c is the speed of light and H_0 is the Hubble constant.

Two different hydrodynamic methods are implemented in Enzo: the piecewise parabolic method (PPM) developed by Woodward & Colella [144] and extended to cosmology by Bryan et al. [145], and the hydrodynamics method used in the magnetohydrodynamics (MHD) code ZEUS [146, 147]. Below, we describe both of these methods in turn (PPM in Section 2.2.2 and the ZEUS method in Section 2.2.3), noting that PPM is viewed as the preferred choice for cosmological simulations since it is higher-order-accurate and is based on a technique that does not require artificial viscosity to resolve shocks. Enzo solves for the gravitational potential using the adaptive particle-mesh (APM) method, which is described in Section 2.2.4. The primordial chemistry and radiative cooling/heating packages are described in Section 2.2.5, and the star formation and feedback algorithms are described in Section 2.2.6.

2.2.1 The AMR machinery

Unlike moving mesh methods [141, 142] or methods that subdivide individual cells [143], Berger & Collella’s AMR (also referred to as *structured* AMR) utilizes an adaptive hierarchy of grid patches at varying levels of resolution. Each rectangular grid patch (referred to as a “grid”) covers some region of space in its *parent grid* which requires higher resolution, and can itself become the parent grid to an even more highly resolved *child grid*. ENZO’s implementation of structured AMR places no restrictions on the number of grids at a given level of refinement, or on the number of levels of refinement. However, owing to

limited computational resources it is practical to institute a maximum level of refinement ℓ_{\max} .

The Enzo implementation of AMR allows arbitrary integer ratios of parent and child grid resolution. However, in practice refinement factors (defined as the ratio of parent grid resolution to child grid resolution) of two or four are typically used, with a refinement factor of two being most commonly used for cosmological simulations for reasons of efficiency. The ratio of box size to the maximum grid resolution of a given simulation is therefore $L/e = N_{\text{root}} \times 2^{\ell_{\max}}$, where N_{root} is the number of cells along one edge of the root grid, ℓ_{\max} is the maximum level of refinement allowed, and L and e are the box size and grid resolution of the most highly refined region, respectively.

The AMR grid patches are the primary data structure in Enzo. Each patch is treated as an individual object which can contain both field variables and particle data. Individual grids are organized into a dynamic, distributed hierarchy of mesh patches. Every processor keeps a description of the entire grid hierarchy at all times, so that each processor knows where all grids are. However, baryon and particle data for a given grid only exists on a single processor. See Figure 2.2 for a schematic example of this, and Figure 2.3 for an example of an Enzo simulation performed using this AMR scheme where both the baryon density and AMR grid hierarchy are shown together. The code handles load balancing on a level-by-level basis such that the workload on each level is distributed as uniformly as possible across all processors. Spatial locality of grids is not forced during message passing, for maximum flexibility (though not necessarily maximum efficiency). The MPI message passing library¹ is used to transfer data between processors.

Each grid patch in Enzo contains arrays of values for baryon and particle quantities. The baryon quantities are stored in arrays with the dimensionality of the simulation itself, which can be 1, 2 or 3 spatial dimensions. Grids are partitioned into a core of *real zones* and a surrounding layer of *ghost zones*. The real zones store field values and ghost zones are used to temporarily store values which have been obtained directly from neighboring grids or interpolated from a parent grid. These zones are necessary to accommodate the computational stencil of the hydrodynamics solvers (Sections 2.2.2 and 2.2.3) and the gravity solver (Section 2.2.4). The hydro solvers typically require ghost zones which are three cells deep and the gravity solver requires 6 ghost zones on every side of the real zones. This can lead to significant memory and computational overhead, particularly for smaller grid patches at high levels of refinement.

Since the addition of more highly refined grids is adaptive, the conditions for refinement must be specified. The criteria of refinement can be set by the threshold value of the overdensity of baryon gas or dark matter in a cell (which is really a refinement on the mass of gas or DM in a cell), the local Jeans length, the local density gradient, or local pressure and energy gradients. A cell reaching any or all of these criteria will then be flagged for refinement. Once all cells at a given level have been examined, rectangu-

¹<http://www-unix.mcs.anl.gov/mpi/>




Figure 2.2: *Left:* Example of a simple, distributed AMR hierarchy showing real and ghost grids. *Right:* Example 2D Enzo grid showing real and ghost zones, as needed for the PPM hydro stencil. Image courtesy of James Bordner (Cent. Astrophysics and Space Sciences, UCSD).




Figure 2.3: Example of an Enzo simulation showing the AMR grid hierarchy. This is a simulation of the collapse of a single massive halo with a 32^3 root grid and two 32^3 static nested grids. AMR is only allowed within the most highly refined static nested grid. Log baryon density is shown by colored isocontours with values corresponding to the legend at center left. The rectangular solid wire frames correspond to individual AMR grids, with different colors corresponding to level as indicated by the legend at top left.

lar boundaries are determined which minimally encompass the flagged cells. A refined grid patch is introduced within each such bounding rectangle. Thus the cells needing refinement, as well as adjacent cells within the patch which do not need refinement, are refined. While this approach is not as memory efficient as cell-splitting AMR schemes, it offers more freedom with finite difference stencils. For example, PPM requires a stencil of seven cells per dimension. This cannot easily be accommodated in cell-splitting AMR schemes. In the simulations discussed in this thesis we typically use baryon and dark matter overdensities as our refinement criteria, though for some higher-resolution simulations we also refine on other criteria as needed.

In Enzo, resolution of the equations being solved is adaptive in time as well as in space. The timestep in Enzo is satisfied on a level-by-level basis by finding the largest timestep such that multiple criteria are satisfied on each level. The timestep criteria used by Enzo are (showing the one-dimensional case for clarity):

$$\Delta t_{hydro} = \min \left(\kappa_{hydro} \frac{a \Delta x}{c_s + |v_x|} \right)_L, \quad (2.18)$$

$$\Delta t_{dm} = \min \left(\kappa_{dm} \frac{a \Delta x}{v_{dm,x}} \right)_L, \quad (2.19)$$

$$\Delta t_{exp} = f_{exp} \left(\frac{a}{\dot{a}} \right), \quad (2.20)$$

$$\Delta t_{accel} = \min \left(\sqrt{\frac{\Delta x}{\mathbf{g}}} \right)_L \quad (2.21)$$

In equations 2.18-2.21, the $\min(\dots)_L$ formalism means that this value is calculated for all cells on a given level L and the minimum overall value is taken as the timestep. Equation 2.18 ensures that all cells satisfy the Courant-Freidrichs-Levy (CFL) condition for accuracy and stability of an explicit finite difference discretization of the Euler equations. Effectively this condition forces the timestep to be small enough such that any changes in the fluid propagate less than a single grid spacing, Δx . In this equation, κ_{hydro} is a numerical constant with a value of $0 < \kappa_{hydro} \leq 1$ (with a typical value of $\kappa_{hydro} \sim 0.3 - 0.5$) that ensures that the CFL condition is always met, and c_s and v_x are the sound speed and peculiar baryon velocity in a given cell. Equation 2.19 is analogous to Equation 2.18 and ensures accuracy in the N-body solver by requiring that no dark matter particle travels more than one cell width. The parameter κ_{dm} is analogous to κ_{hydro} , with an identical range of values. Equation 2.20 limits the timestep such that the expansion parameter a can only change by a fractional amount of $f_{exp} = \Delta a/a$, where f_{exp} is a user-defined parameter and has typical values of $f_{exp} = 0.01 - 0.02$. This is required for the stability of the PPM algorithm in comoving coordinates, and typically limits the timestep only at

high redshifts when densities are relatively homogeneous. Equation 2.21 is supplementary to equation 2.18 in that it takes into account the possibility of large accelerations causing numerical instabilities by violating the Courant condition. In this equation, \mathbf{g} is the gravitational acceleration in each cell on level L .

For simplicity and stability, all cells on a given level are advanced using the same timestep, which is taken to be the minimum value of Equations 2.18-2.21. Once a level L has been advanced in time by Δt_L , all grids at level $L + 1$ are advanced, using the same criteria for timestep calculation described above, until they reach the same physical time as the grids at level L . At this point grids at level $L + 1$ exchange flux information with their parents grids, providing a more accurate solution on level L . This step, controlled by the parameter *FluxCorrection* in Enzo, is extremely important, and can significantly affect simulation results if not used in an AMR calculation. At the end of every timestep on every level each grid updates its ghost zones by exchanging information with its neighboring grid patches (if any exist) and/or by interpolating from a parent grid. In addition, cells are examined to see if they should be refined or de-refined, and the entire grid hierarchy is rebuilt at that level (including all more highly refined levels). The timestepping and hierarchy advancement/rebuilding process described here is repeated recursively on every level to the specified maximum level of refinement in the simulation.

2.2.2 Hydrodynamics with the piecewise parabolic method

The primary hydrodynamic method used in Enzo is based on the piecewise parabolic method (PPM) of Woodward & Colella [144], which has been significantly modified for the study of cosmological fluid flows. The method is described in Bryan et al. [145], but we provide a short description here for clarity.

PPM is an explicit, higher order-accurate version of Godunov’s method for ideal gas dynamics with third order-accurate piecewise parabolic monotonic interpolation and a nonlinear Riemann solver for shock capturing. It does an excellent job of capturing strong shocks in at most two cells. Multidimensional schemes are built up by directional splitting and produce a method that is formally second order-accurate in space and time which explicitly conserves mass, linear momentum, and energy. The conservation laws for fluid mass, momentum and energy density are written in comoving coordinates for a Friedman-Robertson-Walker space-time, as described previously in Equations 2.10 through 2.12. Both the conservation laws and the Riemann solver are modified to include gravity, which is solved using an adaptive particle-mesh (PM) technique (see Section 2.2.4). The terms due to cosmological expansion, as well as primordial chemistry and radiative heating and cooling, are solved in a separate step because they have different numerical behavior, and therefore must be treated differently to ensure stability. Note that unlike the ZEUS hydro scheme, PPM does not need to use artificial viscosity to resolve shocks.

The system of equations described above works well for systems with relatively low

Mach numbers, as long as these systems are well resolved. However, cosmology is replete with situations where there are bulk hypersonic flows. In these situations, the ratio of kinetic to thermal energy can be very high – in some situations up to $10^6 - 10^8$. This implies that the thermal energy is an extremely tiny fraction of the kinetic energy, which can cause numerical problems when one is interested in just the thermal energy of the gas, since Equation 2.12 solves for the total energy. In this system of equations, the thermal energy E_{therm} is calculated as $E - E_{kin}$, where E is the total specific energy as calculated in equation 2.12 and E_{kin} is the specific kinetic energy, $0.5\mathbf{v}_b^2$. In hypersonic flows E and E_{kin} are nearly the same, and any number calculated as the difference of these is going to be strongly affected by numerical error. To avoid this problem, Enzo also solves the internal energy equation in comoving coordinates:

$$\frac{\partial e}{\partial t} + \frac{1}{a}\mathbf{v}_b \cdot \nabla e = -\frac{3(\gamma - 1)\dot{a}}{a}e - \frac{p}{a\rho}\nabla \cdot \mathbf{v}_b \quad (2.22)$$

In this equation e is the internal energy and the other terms are as described previously. The code still conserves total energy (E) as well. In order to maintain consistency, both equations are solved at all times in all cells, with the equation for the total energy (eqtn. 2.12) being used for hydrodynamics routines and the internal energy e being used when temperature is required. When pressure is required for dynamic purposes, the total energy is used if the ratio of thermal energy to total energy is less than some threshold value η , and the internal energy is used for values of the ratio larger than η . A typical value of this parameter is 10^{-3} . This *dual energy formulation* ensures that the method produces the correct entropy jump at strong shocks and also yields accurate pressures and temperatures in cosmological hypersonic flows.

2.2.3 Hydrodynamics with the ZEUS hydrodynamic method

As a check on PPM, Enzo also includes an implementation of the finite-difference hydrodynamic algorithm employed in the compressible magnetohydrodynamics code ‘ZEUS’ [146, 147]. Fluid transport is solved on a Cartesian grid using the upwind, monotonic advection scheme of van Leer [148] within a multistep (operator split) solution procedure which is fully explicit in time. This method is formally second order-accurate in space but first order-accurate in time. It is important to note that this method conserves internal energy rather than total energy, so the “dual energy formulation” discussed in Section 2.2.2 is unnecessary.

Operator split methods break the solution of the hydrodynamic equations into parts, with each part representing a single term in the equations. Each part is evaluated successively using the results preceding it. The individual parts of the solution are grouped into two steps, called the *source* and *transport* steps.

The ZEUS method uses a von Neumann-Richtmeyer artificial viscosity to smooth shock discontinuities that may appear in fluid flows and can cause a break-down of finite-difference equations. The artificial viscosity term is added in the source terms as

$$\rho \frac{\partial \mathbf{v}}{\partial t} = -\nabla p - \rho \nabla \phi - \nabla \cdot \mathbf{Q} \quad (2.23)$$

$$\frac{\partial e}{\partial t} = -p \nabla \cdot \mathbf{v} - \mathbf{Q} : \nabla \mathbf{v}, \quad (2.24)$$

where \mathbf{v} is the baryon velocity, ρ is the mass density, p is the pressure, e is the internal energy density of the gas and \mathbf{Q} is the artificial viscosity stress tensor, such that:

$$Q_{ii} = \begin{cases} Q_{\text{AV}} \rho_b (a \Delta v_i + \dot{a} \Delta x_i)^2, & \text{for } a \Delta v_i + \dot{a} \Delta x_i < 0 \\ 0 & \text{otherwise} \end{cases} \quad (2.25)$$

and

$$Q_{ij} = 0 \quad \text{for } i \neq j. \quad (2.26)$$

Δx_i and Δv_i refer to the comoving width of the grid cell along the i -th axis and the corresponding difference in gas peculiar velocities across the grid cell, respectively, and a is the cosmological scale factor. Q_{AV} is a constant with a typical value of 2. We refer the interested reader to Anninos et al. [149] for more details.

The limitation of a technique that uses an artificial viscosity is that, while the correct Rankine-Hugoniot jump conditions are achieved, shocks are broadened over 6-8 mesh cells and are thus not treated as true discontinuities. This may cause unphysical pre-heating of gas upstream of the shock wave, as discussed in Anninos et al. [149].

2.2.4 Gravity Solver

There are multiple methods to compute the gravitational potential (which is an elliptic equation in the Newtonian limit) in a structured AMR framework. One way would be to model the dark matter (or other collisionless particle-like objects, such as stars) as a second fluid in addition to the baryonic fluid and solve the collisionless Boltzmann equation, which follows the evolution of the fluid density in six-dimensional phase space. However, this is computationally prohibitive owing to the large dimensionality of the problem, making this approach unappealing for the cosmological AMR code.

Instead, Enzo uses a particle-mesh N-body method to calculate the dynamics of collisionless systems [139]. This method follows trajectories of a representative sample of individual particles and is much more efficient than a direct solution of the Boltzmann equation in most astrophysical situations. The gravitational potential is computed by solving the elliptic Poisson's equation (Eqn. 2.14).

These equations are finite-differenced and for simplicity are solved with the same timestep as the hydrodynamic equations, as discussed in Section 2.2.1. The dark matter particles are distributed onto the grids using the cloud-in-cell (CIC) interpolation technique to form a spatially discretized density field (analogous to the baryon densities used to calculate the equations of hydrodynamics). After sampling dark matter density onto the grid and adding baryon density if it exists (to get the total matter density in each cell), the gravitational potential is calculated on the periodic root grid using a fast Fourier transform. In order to calculate more accurate potentials on subgrids, Enzo resamples the dark matter density onto individual subgrids using the same CIC method as on the root grid, but at higher spatial resolution (and again adding baryon densities if applicable). Boundary conditions are then interpolated from the potential values on the parent grid (with adjacent grid patches on a given level communicating to ensure that their boundary values are the same), and then a multigrid relaxation technique is used to calculate the gravitational potential at every point within a subgrid. Forces are computed on the mesh by finite-differencing potential values and are then interpolated to the particle positions, where they are used to update the particle’s position and velocity information. Potentials on child grids are computed recursively and particle positions are updated using the same timestep as in the hydrodynamic equations. Particles are stored in the most highly refined grid patch at the point in space where they exist, and particles which move out of a subgrid patch are sent to the grid patch covering the adjacent volume with the finest spatial resolution, which may be of the same spatial resolution, coarser, or finer than the grid patch that the particles are moved from. This takes place in a communication process at the end of each timestep on a level.

At this point it is useful to emphasize that the effective force resolution of an adaptive particle-mesh calculation is approximately twice as coarse as the grid spacing at a given level of resolution. The potential is solved in each grid cell; however, the quantity of interest, namely the acceleration, is the gradient of the potential, and hence two potential values are required to calculate this. In addition, it should be noted that the adaptive particle-mesh technique described here is very memory intensive: in order to ensure reasonably accurate force resolution at grid edges the multigrid relaxation method used in the code requires a layer of ghost zones which is very deep – typically 6 cells in every direction around the edge of a grid patch. This greatly adds to the memory requirements of the simulation, particularly because subgrids are typically small (on the order of $12^3 - 16^3$ real cells for a standard cosmological calculation) and ghost zones can dominate the memory and computational requirements of the code.

2.2.5 Radiative processes and non-equilibrium primordial chemistry

Though the equations of hydrodynamics described above are a closed system, they are still missing a crucial piece of physics: radiative heating and cooling. Radiative cooling is extremely important in many astrophysical situations, as is the heating of gas from some sort of radiation background. Enzo has a very simple Sutherland and Dopita equilibrium cooling function [150] implemented, which uses a cooling table assuming a fixed metallicity of $Z = 0.3Z_{\odot}$, and also a nonequilibrium heating/cooling model that assumes gas of primordial composition exposed to a uniform metagalactic ultraviolet (UV) background that varies with time [151].

The simulations discussed in this thesis almost exclusively use the nonequilibrium routines, described in great detail by Abel et al. and Anninos et al. [152, 153] and summarized in Appendix A. These routines follow the non-equilibrium chemistry of a gas of primordial composition with 9 total species: H , H^+ , He , He^+ , He^{++} , H^- , H_2^+ , H_2 , and e^- . The code also calculates radiative heating and cooling, following atomic line excitation, recombination, collisional excitation, free-free transitions, molecular line excitations, and Compton scattering of the cosmic microwave background, as well as any of approximately a dozen different models for a metagalactic UV background that heat the gas via photoionization and photodissociation. The multispecies rate equations are solved out of equilibrium to properly model situations where, e.g., the cooling rate of the gas is much shorter than the hydrogen recombination time. The effect of this nonequilibrium cooling is to leave behind a much larger fraction of residual free electrons than one would expect if the assumption of equilibrium were being made. The practical effect of this is that more H^- is formed, which subsequently produces hydrogen molecules. If large amounts of H_2 is formed it can greatly increase the cooling rate of primordial gas at relatively low temperatures ($T \leq 10^4$ K). This can efficiently cool the gas to approximately 200 K, which significantly reduces the Jeans mass of the gas. Correct modeling of the formation of molecular hydrogen is crucial to the study of star formation in a primordial gas.

A total of 9 kinetic equations are solved, including 29 kinetic and radiative processes, for the 9 species mentioned above. See Table 2.1 for the collisional processes and Table 2.2 for the radiative processes solved.

The chemical reaction equation network is technically challenging to solve due to the huge range of reaction timescales involved – the characteristic creation and destruction timescales of the various species and reactions can differ by many orders of magnitude. As a result, the set of rate equations is extremely stiff, and an explicit scheme for integration of the rate equations can be exceptionally costly if small enough timesteps are taken to keep the network stable. This makes an implicit scheme much more preferable for such a set of equations. However, an implicit scheme typically require an iterative procedure

Collisional Processes

(1)	H	+	e ⁻	→	H ⁺	+	2e ⁻
(2)	H ⁺	+	e ⁻	→	H	+	γ
(3)	He	+	e ⁻	→	He ⁺	+	2e ⁻
(4)	He ⁺	+	e ⁻	→	He	+	γ
(5)	He ⁺	+	e ⁻	→	He ⁺⁺	+	2e ⁻
(6)	He ⁺⁺	+	e ⁻	→	He ⁺	+	γ
(7)	H	+	e ⁻	→	H ⁻	+	γ
(8)	H ⁻	+	H	→	H ₂	+	e ⁻
(9)	H	+	H ⁺	→	H ₂ ⁺	+	γ
(10)	H ₂ ⁺	+	H	→	H ₂	+	H ⁺
(11)	H ₂	+	H ⁺	→	H ₂ ⁺	+	H
(12)	H ₂	+	e ⁻	→	2H	+	e ⁻
(13)	H ₂	+	H	→	3H		
(14)	H ⁻	+	e ⁻	→	H	+	2e ⁻
(15)	H ⁻	+	H	→	2H	+	e ⁻
(16)	H ⁻	+	H ⁺	→	2H		
(17)	H ⁻	+	H ⁺	→	H ₂ ⁺	+	e ⁻
(18)	H ₂ ⁺	+	e ⁻	→	2H		
(19)	H ₂ ⁺	+	H ⁻	→	H ₂	+	H
(20)	3H			→	H ₂	+	H

Table 2.1: Collisional processes solved in the Enzo nonequilibrium primordial chemistry routines.

Radiative Processes

(21)	H	+	γ	→	H ⁺	+	e ⁻
(22)	He	+	γ	→	He ⁺	+	e ⁻
(23)	He ⁺	+	γ	→	He ⁺⁺	+	e ⁻
(24)	H ⁻	+	γ	→	H	+	e ⁻
(25)	H ₂	+	γ	→	H ₂ ⁺	+	e ⁻
(26)	H ₂ ⁺	+	γ	→	H	+	H ⁺
(27)	H ₂ ⁺	+	γ	→	2H ⁺	+	e ⁻
(28)	H ₂	+	γ	→	H ₂ [*]	→	2H
(29)	H ₂	+	γ	→	2H		

Table 2.2: Radiative processes solved in the Enzo nonequilibrium primordial chemistry routines.

to converge, and for large networks (such as this one) an implicit method can be very time-consuming, making it undesirable for a large, three-dimensional simulation.

Enzo solves the rate equations using a method based on a backwards differencing formula (BDF) in order to provide a stable and accurate solution. This technique is optimized by taking the chemical intermediaries H^- and H_2^+ , which have large rate coefficients and low concentrations, and grouping them into a separate category of equations. Due to their fast reactions these species are very sensitive to small changes in the more abundant species. However, due to their low overall concentrations, they do not significantly influence the abundance of species with higher concentrations. Therefore, reactions involving these two species can be decoupled from the rest of the network and treated independently. In fact, H^- and H_2^+ are treated as being in equilibrium at all times, independent of the other species and the hydrodynamic variables. This allows a large speedup in solution as the BDF scheme is then applied only to the slower 7-species network on timescales closer to those required by the hydrodynamics of the simulation. Even so, the accuracy and stability of the scheme is maintained by subcycling the rate solver over a single hydrodynamic timestep. These subcycle timesteps are determined so that the maximum fractional change in the electron concentration is limited to no more than 10% per timestep.

It is important to note the regime in which this model is valid. According to Abel et al. and Anninos et al. [152, 153], the reaction network is valid for temperatures between $10^0 - 10^8$ K. The original model discussed in these two references is only valid up to $n_H \sim 10^4 \text{ cm}^{-3}$. However, addition of the 3-body H_2 formation process (equation 20 in Table 2.1) allows correct modeling of the chemistry of the gas up until the point where collisionally-induced emission from molecular hydrogen becomes an important cooling process, which occurs at $n_H \sim 10^{14} \text{ cm}^{-3}$. A further concern is that the optically thin approximation for radiative cooling breaks down, which occurs before $n_H \sim 10^{16} - 10^{17} \text{ cm}^{-3}$. Beyond this point, modifications to the cooling function that take into account the non-negligible opacity in the gas must be made, as discussed by Ripamonti & Abel [47]. Even with these modifications, a completely correct description of the cooling of this gas will require some form of radiation transport, which will greatly increase the cost of the simulations.

Several processes are neglected. The deuterium atom and its processes are completely ignored, which may have some effect. Recent work shows that HD is a more effective coolant than previously thought [154]. However, the fractional abundance of HD is so low that under circumstances relevant to the formation of a Population III star in an un-ionized region it should be sub-dominant. However, the enhanced electron fraction in fossil HII regions (as discussed later in this thesis) could result in the HD molecule becoming a dominant cooling mechanism at relatively low (\sim few hundred K) temperatures, and could potentially cool the gas down to below 100 K, which can enhance fragmentation and could have important consequences for the IMF of primordial stars

forming in a relic HII region [78].

Aside from deuterium, the chemical reactions involving lithium are also neglected. According to Galli & Palla [27], these are not important for the density and temperature regimes explored by the simulations discussed in this thesis. However, at higher densities it is possible that there are regimes where lithium can be an important coolant.

2.2.6 Star Formation and Feedback Algorithms

While the physics discussed previously is all crucial to the study of cosmological structure formation, most cosmological observations are of stars and related phenomena. Additionally, stars eject energy and metal-enriched gas throughout their lives, drastically modifying their own environment. The formation of galaxies and clusters cannot be completely modeled without including the feedback of energy and metals. In particular, it is thought that feedback is crucial for suppressing the large numbers of dwarf-like galaxies that CDM theories predict [16, 155]. An early burst of star formation could remove a large fraction of cold gas from such systems [156, 157]. Also, the unexpectedly low luminosities of small clusters and groups (relative to the richest clusters) is commonly explained through feedback [158]. Energy ejected during the formation of the cluster ellipticals lowers the central density and hence the X-ray luminosity of such clusters [159]. Therefore, the inclusion of star formation and the feedback of energy and metals into the intergalactic medium in a cosmological code is crucial for many reasons.

We have extended the Enzo code to include multiple models for star formation and feedback. It is difficult to directly simulate the formation of individual stars in the context of simulations of galaxy formation and evolution due to the immense computational cost. Therefore, we adopt parametric methods which attempt to model star forming regions. One model is based on Kravtsov's method [160], and the other is based on the method of Cen & Ostriker [170], which has been modified for use in an AMR code. The basic ideas behind the methods are straightforward and observationally motivated. Similar (though somewhat more advanced) methods have been employed for the smoothed particle hydrodynamics (SPH) method, most recently by Springel & Hernquist [162, 163, 164]. This method assumes a multiphase IGM and has been shown to accurately reproduce the cosmic star formation rate, and will be implemented into Enzo in the near future.

In the following sections we will detail both the Kravtsov and Cen & Ostriker star formation methods separately. Tests of the star formation and feedback algorithms will not be shown, since this work is currently in progress (and an extension of the AMR/SPH code comparison described in Chapter 3).

The Kravtsov star formation and feedback algorithm

Kravtsov’s method of star formation is designed to reproduce the global Schmidt law of star formation [160, 165]. This algorithm is deliberately minimal, and is explicitly geared towards modeling star formation in a phenomenological way on kiloparsec scales. Stars are assumed to form with a characteristic gas timescale τ_* such that $\dot{\rho}_* = \rho_{gas}/\tau_*$ where ρ_{gas} and ρ_* are the baryon gas and stellar densities, respectively. This “constant efficiency” model on the scale of star formation regions is well motivated observationally [166, 167]. Star formation is only allowed to take place in very dense regions with $\rho_{gas} \geq \rho_{SF}$, where ρ_{SF} is a constant proper density threshold above which star formation is allowed to occur. No other criteria are imposed. Kravtsov’s typical choices for τ_* and ρ_{SF} are $\tau_* = 4$ Gyr and $\rho_{SF} = 1.64 M_\odot \text{ pc}^{-3}$ ($n_H \sim 50 \text{ cm}^{-3}$). The adopted timescale is derived from the observationally-determined normalization of the Schmidt law, and the density threshold is determined by observations of star forming regions on ~ 100 pc scales. Note that the density threshold is in proper, not comoving, units.

Algorithmically, the star formation events in Kravtsov’s code are assumed to occur once every global time step $\Delta t_0 \leq 10^7$. In cells where star formation is determined to occur (i.e. $\rho_{gas} \geq \rho_{SF}$) a collisionless “star particle” is assumed to form, with a mass $m_* = \dot{\rho}_* V_{cell} \Delta t_0$, where $\dot{\rho}_*$ is described above and V_{cell} is the volume of the mesh cell. These star particles are formed instantaneously, with all m_* of gas being taken out of the cell and immediately deposited into the star particle. This particle is then given the velocity of the gas in the cell, and thereafter treated as a collisionless particle. The Enzo implementation of this algorithm is similar, except that instead of forming stars only at the root grid time step, we allow stars to form at the time step of the highest level of resolution at any particular point in space. As can be seen from the equation for m_* above, this can result in very small stellar masses. To avoid memory and processor time issues related to having very large numbers of star particles we impose a threshold mass $M_{*,min}$ such that a star particle only forms if $m_* \geq M_{*,min}$. An appropriate choice of $M_{*,min}$ does not significantly change the star overall star formation history of a simulation, though it may delay the onset of star formation in a given cell relative to a simulation without a particle mass threshold.

Each “star particle” is assumed to represent an ensemble of stars and is treated as a single-age stellar population. Kravtsov assumes that the IMF is described by a Miller & Scalo functional form with stellar masses between 0.1 and 100 M_\odot [168]. All stars in this IMF with $M_* > 8 M_\odot$ deposit 2×10^{51} ergs of thermal energy and a mass $f_z M_*$ of metals into the cell in which they form without delay, with $f_z \equiv \min(0.2, 0.01 M_* - 0.06)$ (i.e. instantaneous deposition of metals). The definition of f_z is a rough approximation of the results of Woosley & Weaver [169].

Kravtsov reports that simulations with this algorithm reliably reproduce the star formation rate-gas surface density relation of the Schmidt law, and are not particularly

sensitive to numerical parameters [160]. He also notes that this is surprisingly insensitive to the presence or absence of feedback and details of the cooling and heating properties of the gas, which suggests that the global star formation rate is determined by gravitationally-driven supersonic turbulence (on large scales) rather than stellar feedback or thermal instabilities on small scales.

The Cen & Ostriker star formation algorithm

The Cen & Ostriker method is a heuristic model of star formation on galactic scales. This method, first described in a 1992 paper [170], is similar in some ways to the Kravtsov algorithm but has more complex criteria for determining where stars should be formed. In this method, cells that form stars must have a baryon overdensity higher than some threshold $\rho_b/\bar{\rho}_b \geq \eta$ where ρ_b is the baryon density in that cell, $\bar{\rho}_b$ is the mean baryon density in the simulation, and η is the user-defined overdensity threshold. Additionally, the gas in the cells must be contracting, cooling rapidly, and gravitationally unstable, e.g.:

$$\nabla \cdot \mathbf{v}_b < 0, \quad (2.27)$$

$$t_{cool} < t_{dyn} \equiv \sqrt{3\pi/32G\rho_{tot}}, \quad (2.28)$$

$$m_b > m_J \equiv G^{-3/2}\rho_b^{-1/2}C^3 \left[1 + \frac{\delta\rho_d}{\delta\rho_b} \right]^{-3/2} \quad (2.29)$$

where \mathbf{v} is the velocity of the gas in the cell, ρ_b and ρ_d are the cell's baryon and dark matter density, respectively, $\rho_{total} = \rho_b + \rho_d$, m_b and m_j are the baryonic mass in the cell and jeans mass of the cell, and C is the isothermal sound speed in the cell. If all of these criteria are met, the mass of a star particle is calculated as $m_* = m_b \frac{\Delta t}{t_{dyn}} f_{*eff}$, where f_{*eff} is the star formation efficiency parameter.

If m_* is greater than a minimum star mass m_{*min} , a particle is created and given several attributes: Mass, a unique index number, the time of formation t_{form} , the local dynamical free-fall time t_{dyn} and the metallicity fraction of the baryon gas in the cell f_{Zb} . There is a user-defined minimum dynamical time $T_{dyn,min}$ which is observationally motivated and affects the feedback rates (see below). The particle is placed in the center of the cell and given the same peculiar velocity as the gas in the cell, and is then treated in the same manner as the dark matter particles. An amount of baryon gas corresponding to the new particle's mass is removed from the cell.

In addition, we have added a stochastic star formation algorithm that keeps track of all of the sub-mass threshold stars that should have been created and when the total mass of uncreated stars is greater than the minimum mass, a star particle is created.

The feedback of energy and metals into the baryon gas is similar to the Kravtsov feedback, with some important differences. The star formation algorithm creates each star particle instantaneously. However, feedback should take place over a significant timescale, as all of the stars contained within the “star particle” would form over a long period of time. Therefore, we assume that for the purposes of feedback that the mass of stars formed at a time t with finite timestep Δt is:

$$\begin{aligned} \Delta m_{sf} = & m_* \left[\left(1 + \frac{t - t_{form}}{t_{dyn}} \right) \exp \left(\frac{-(t - t_{form})}{t_{dyn}} \right) \right. \\ & \left. - \left(1 + \frac{t + \Delta t - t_{form}}{t_{dyn}} \right) \exp \left(\frac{-(t + \Delta t - t_{form})}{t_{dyn}} \right) \right] \end{aligned} \quad (2.30)$$

which can be represented more clearly in integral form:

$$\int_t^{t+Dt} \frac{dM}{dt} dt = \int_t^{t+Dt} m_* \frac{dt}{t_{dyn}} \left(\frac{t - t_{form}}{t_{dyn}} \right) \exp \left(- \frac{t - t_{form}}{t_{dyn}} \right) \quad (2.31)$$

During this timestep, we assume that the star particle feeds back metal-enriched gas and thermal energy from supernovae and from stellar winds. Since massive stars have very short lifetimes, we assume that there is an immediate feedback of some fraction f_{SN} of the rest energy from the stars created that timestep into the baryon gas, such that $E_{add} = f_{SN} \Delta m_{sf} c^2$, where c is the speed of light. In addition, a fraction f_{Z*} of the metal from the star particle is fed back into the baryon gas, which takes into account the effects of metal recycling. Finally, a fraction of the mass f_{m*} from all stars (rather than just supernovae) is fed back into the gas along with momentum in order to simulate the mass ejection from non-exploding stars via stellar winds.

There are six user-defined parameters in this algorithm: three deal with the star formation (η , m_{*min} and $T_{dyn,min}$), and three deal with feedback (f_{SN} , f_{Z*} and f_{m*}). Some of these parameters are completely free, while others can be guided by observation or theory. For example, the supernova feedback parameter, f_{SN} , can be constrained assuming that, for every $200 M_\odot$ of stars created, one supernova occurs, and this event feeds back approximately 10^{51} ergs of thermal energy, giving:

$$f_{SN} = \frac{10^{51} \text{ergs}}{200 M_\odot c^2} \simeq 3 \times 10^{-6} \quad (2.32)$$

The metal yield f_{Z*} , defined as the mass in metals produced per unit mass of stars created, can be constrained by a theoretical model of Woosley & Weaver [169]. This model suggests that $f_{Z*} = 0.02$ is an appropriate number. The minimum dynamical time is set to be $T_{dyn,min} = 10^7$ years to agree with timescales seen in nearby OB associations.

The other parameters, such as the overdensity threshold η , minimum star mass m_{*min} , and mass ejection fraction f_{m*} are not well constrained either theoretically or observationally. Indeed, m_{*min} is a purely numerical parameter designed to keep the code from producing too many star particles, and thus has no observational or theoretical counterpart. These parameters have to be set by performing parameter studies and comparing to observations. Unfortunately, the range of parameter space is large, and the results may be degenerate for some combinations of these parameters.

Chapter 3

A Comparison of AMR and SPH cosmology codes

3.1 Summary

We compare two cosmological hydrodynamic simulation codes in the context of hierarchical galaxy formation: the Lagrangian smoothed particle hydrodynamics (SPH) code ‘GADGET’, and the Eulerian adaptive mesh refinement (AMR) code ‘Enzo’. Both codes represent dark matter with the N-body method but use different gravity solvers and fundamentally different approaches for baryonic hydrodynamics. The SPH method in GADGET uses a recently developed ‘entropy conserving’ formulation of SPH, while for the mesh-based Enzo two different formulations of Eulerian hydrodynamics are employed: the piecewise parabolic method (PPM) extended with a dual energy formulation for cosmology, and the artificial viscosity-based scheme used in the magnetohydrodynamics code ZEUS. In this paper we focus on a comparison of cosmological simulations that follow either only dark matter, or also a non-radiative (‘adiabatic’) hydrodynamic gaseous component. We perform multiple simulations using both codes with varying spatial and mass resolution with identical initial conditions.

The dark matter-only runs agree generally quite well provided Enzo is run with a comparatively fine root grid and a low overdensity threshold for mesh refinement, otherwise the abundance of low-mass halos is suppressed. This can be readily understood as a consequence of the hierarchical particle-mesh algorithm used by Enzo to compute gravitational forces, which tends to deliver lower force resolution than the tree-algorithm of GADGET at early times before any adaptive mesh refinement takes place. At comparable force resolution we find that the latter offers substantially better performance and lower memory consumption than the present gravity solver in Enzo.

In simulations that include adiabatic gas dynamics we find general agreement in the distribution functions of temperature, entropy, and density for gas of moderate to high

overdensity, as found inside dark matter halos. However, there are also some significant differences in the same quantities for gas of lower overdensity. For example, at $z = 3$ the fraction of cosmic gas that has temperature $\log T > 0.5$ is $\sim 80\%$ for both Enzo/ZEUS and GADGET, while it is $40 - 60\%$ for Enzo/PPM. We argue that these discrepancies are due to differences in the shock-capturing abilities of the different methods. In particular, we find that the ZEUS implementation of artificial viscosity in Enzo leads to some unphysical heating at early times in preshock regions. While this is apparently a significantly weaker effect in GADGET, its use of an artificial viscosity technique may also make it prone to some excess generation of entropy which should be absent in ENZO/PPM. Overall, the hydrodynamical results for GADGET are bracketed by those for Enzo/ZEUS and Enzo/PPM, but are closer to Enzo/ZEUS. This chapter has been previously published as a paper in the *Astrophysical Journal* [4].

3.2 Motivation

Within the currently leading theoretical model for structure formation small fluctuations that were imprinted in the primordial density field are amplified by gravity, eventually leading to non-linear collapse and the formation of dark matter (DM) halos. Gas then falls into the potential wells provided by the DM halos where it is shock-heated and then cooled radiatively, allowing a fraction of the gas to collapse to such high densities that star formation can ensue. The formation of galaxies hence involves dissipative gas dynamics coupled to the nonlinear regime of gravitational growth of structure. The substantial difficulty of this problem is exacerbated by the inherent three-dimensional character of structure formation in a Λ CDM universe, where due to the shape of the primordial power spectrum a large range of wave modes becomes nonlinear in a very short time, resulting in the rapid formation of objects with a wide range of masses which merge in geometrically complex ways into ever more massive systems. Therefore, direct numerical simulations of structure formation which include hydrodynamics arguably provide the only method for studying this problem in its full generality.

Hydrodynamic methods used in cosmological simulations of galaxy formation can be broken down into two primary classes: techniques using an Eulerian grid, including ‘Adaptive Mesh Refinement’ (AMR) techniques, and those which follow the fluid elements in a Lagrangian manner using gas particles, such as ‘Smoothed Particle Hydrodynamics’ (SPH). Although significant amounts of work have been done on structure/galaxy formation using both types of simulations, very few detailed comparisons between the two simulation methods have been carried out [171, 172], despite the existence of widespread prejudices in the field with respect to alleged weaknesses and strengths of the different methods.

Perhaps the most extensive code comparison performed to date was the *Santa Barbara*

cluster comparison project [172], in which several different groups ran a simulation of the formation of one galaxy cluster, starting from identical initial conditions. They compared a few key quantities of the formed cluster, such as radially-averaged profiles of baryon and dark matter density, gas temperature and X-ray luminosity. Both Eulerian (fixed grid and AMR) and SPH methods were used in this study. Although most of the measured properties of the simulated cluster agreed reasonably well between different types of simulations (typically within $\sim 20\%$), there were also some noticeable differences which largely remained unexplained, for example in the central entropy profile, or in the baryon fraction within the virial radius. Later simulations by Ascasibar et al. [173] compare results from the Eulerian AMR code ART [174] with the entropy-conserving version of GADGET. They find that the entropy-conserving version of GADGET significantly improves agreement with grid codes when examining the central entropy profile of a galaxy cluster, though the results are not fully converged. The GADGET result using the new hydro formulation now shows an entropy floor – in the Santa Barbara paper the SPH codes typically did not display any trend towards a floor in entropy at the center of the cluster while the grid-based codes generally did. The ART code produces results that agree extremely well with the grid codes used in the comparison. The observed convergence in cluster properties is encouraging, but there is still a need to explore other systematic differences between simulation methods.

The purpose of the present study is to compare two different types of modern cosmological hydrodynamic methods, SPH and AMR, in greater depth, with the goal of obtaining a better understanding of the systematic differences between the different numerical techniques. This will also help to arrive at a more reliable assessment of the systematic uncertainties in present numerical simulations, and provide guidance for future improvements in numerical methods. The codes we use are ‘GADGET’¹, an SPH code developed by Springel et al. [175], and ‘Enzo’², an AMR code developed by Bryan et al. [134, 135]. In this paper, we focus our attention on the clustering properties of dark matter and on the global distribution of the thermodynamic quantities of cosmic gas, such as temperature, density, and entropy of the gas. Our work is thus complementary to the Santa Barbara cluster comparison project because we examine cosmological volumes that include many halos and a low-density intergalactic medium, rather than focusing on a single particularly well-resolved halo. We also include idealized tests designed to highlight the effects of artificial viscosity and cosmic expansion.

The present study is the first paper in a series that aims at providing a comprehensive comparison of AMR and SPH methods applied to the dissipative galaxy formation problem. In this paper, we describe the general code methodology, and present fundamental comparisons between dark matter-only runs and runs that include ordinary ‘adiabatic’ hydrodynamics. This paper is meant to provide statistical comparisons between simula-

¹<http://www.MPA-Garching.MPG.DE/gadget/>

²<http://www.cosmos.ucsd.edu/enzo/>

tion codes, and we leave detailed comparisons of baryon properties in individual halos to a later paper. Additionally, we plan to compare simulations using radiative cooling, star formation, and supernova feedback in forthcoming work.

The organization of this paper is as follows. We provide a short overview of the GADGET code in Section 3.3 (Enzo is described in detail in Section 2.2), and then describe the details of our simulations in Section 3.4. Our comparison is then conducted in two steps. We first compare the dark matter-only runs in Section 3.5 to test the gravity solver of each code. This is followed in Section 3.6 with a detailed comparison of hydrodynamic results obtained in adiabatic cosmological simulations. We then discuss effects of artificial viscosity in Section 3.7, and the timing and memory usage of the two codes in Section 3.8. Finally, we conclude in Section 3.9 with a discussion of our findings.

3.3 The GADGET smoothed particle hydrodynamics code

In this study, we compare Enzo with a new version of the parallel TreeSPH code GADGET [176], which combines smoothed particle hydrodynamics with a hierarchical tree algorithm for gravitational forces. Codes with a similar principal design [177, 178, 179, 180] have been employed in cosmology for a number of years. Compared with previous SPH implementations, the new version GADGET-2 used here differs significantly in its formulation of SPH (as discussed below), in its timestepping algorithm, and in its parallelization strategy. In addition, the new code optionally allows the computation of long-range forces with a particle-mesh (PM) algorithm, with the tree algorithm supplying short-range gravitational interactions only. This ‘TreePM’ method can substantially speed up the computation while maintaining the large dynamic range and flexibility of the tree algorithm.

3.3.1 Hydrodynamical method

SPH uses a set of discrete tracer particles to describe the state of a fluid, with continuous fluid quantities being defined by a kernel interpolation technique if needed [181, 182, 183]. The particles with coordinates \mathbf{r}_i , velocities \mathbf{v}_i , and masses m_i are best thought of as fluid elements that sample the gas in a Lagrangian sense. The thermodynamic state of each fluid element may either be defined in terms of its thermal energy per unit mass, u_i , or in terms of the entropy per unit mass, s_i . We in general prefer to use the latter as the independent thermodynamic variable evolved in SPH, for reasons discussed in full detail by Springel & Hernquist [184]. In essence, use of the entropy allows SPH to be formulated so that both energy and entropy are manifestly conserved, even when adaptive smoothing lengths are used. [185] In the following we summarize the “entropy formulation” of SPH,

which is implemented in GADGET-2 as suggested by Springel & Hernquist [184].

We begin by noting that it is more convenient to work with an entropic function defined by $A \equiv P/\rho^\gamma$, instead of directly using the thermodynamic entropy s per unit mass. Because $A = A(s)$ is only a function of s for an ideal gas, we will simply call A the ‘entropy’ in what follows. Of fundamental importance for any SPH formulation is the density estimate, which GADGET calculates in the form

$$\rho_i = \sum_{j=1}^N m_j W(|\mathbf{r}_{ij}|, h_i), \quad (3.1)$$

where $\mathbf{r}_{ij} \equiv \mathbf{r}_i - \mathbf{r}_j$, and $W(r, h)$ is the SPH smoothing kernel. In the entropy formulation of the code, the adaptive smoothing lengths h_i of each particle are defined such that their kernel volumes contain a constant mass for the estimated density; i.e. the smoothing lengths and the estimated densities obey the (implicit) equations

$$\frac{4\pi}{3} h_i^3 \rho_i = N_{\text{sph}} \bar{m}, \quad (3.2)$$

where N_{sph} is the typical number of smoothing neighbors, and \bar{m} is the average particle mass. Note that in traditional formulations of SPH, smoothing lengths are typically chosen such that the number of particles inside the smoothing radius h_i is equal to a constant value N_{sph} .

Starting from a discretized version of the fluid Lagrangian, one can show [184] that the equations of motion for the SPH particles are given by

$$\frac{d\mathbf{v}_i}{dt} = - \sum_{j=1}^N m_j \left[f_i \frac{P_i}{\rho_i^2} \nabla_i W_{ij}(h_i) + f_j \frac{P_j}{\rho_j^2} \nabla_i W_{ij}(h_j) \right], \quad (3.3)$$

where the coefficients f_i are defined by

$$f_i = \left[1 + \frac{h_i}{3\rho_i} \frac{\partial \rho_i}{\partial h_i} \right]^{-1}, \quad (3.4)$$

and the abbreviation $W_{ij}(h) = W(|\mathbf{r}_i - \mathbf{r}_j|, h)$ has been used. The particle pressures are given by $P_i = A_i \rho_i^\gamma$. Provided there are no shocks and no external sources of heat, the equations above already fully define reversible fluid dynamics in SPH. The entropy A_i of each particle simply remains constant in such a flow.

However, flows of ideal gases can easily develop discontinuities where entropy must be generated by microphysics. Such shocks need to be captured by an artificial viscosity technique in SPH. To this end GADGET uses a viscous force

$$\left. \frac{d\mathbf{v}_i}{dt} \right|_{\text{visc.}} = - \sum_{j=1}^N m_j \Pi_{ij} \nabla_i \bar{W}_{ij}. \quad (3.5)$$

For the simulations of this paper, we use a standard Monaghan-Balsara artificial viscosity Π_{ij} [186, 187], parameterized in the following form:

$$\Pi_{ij} = \begin{cases} [-\alpha c_{ij} \mu_{ij} + 2\alpha \mu_{ij}^2] / \rho_{ij} & \text{if } \mathbf{v}_{ij} \cdot \mathbf{r}_{ij} < 0 \\ 0 & \text{otherwise,} \end{cases} \quad (3.6)$$

with

$$\mu_{ij} = \frac{h_{ij} \mathbf{v}_{ij} \cdot \mathbf{r}_{ij}}{|\mathbf{r}_{ij}|^2}. \quad (3.7)$$

Here h_{ij} and ρ_{ij} denote arithmetic means of the corresponding quantities for the two particles i and j , with c_{ij} giving the mean sound speed. The symbol \overline{W}_{ij} in the viscous force is the arithmetic average of the two kernels $W_{ij}(h_i)$ and $W_{ij}(h_j)$. The strength of the viscosity is regulated by the parameter α , with typical values in the range 0.75 – 1.0. Following Steinmetz [188], GADGET also uses an additional viscosity-limiter in Eqn. (3.6) in the presence of strong shear flows to alleviate angular momentum transport.

Note that the artificial viscosity is only active when fluid elements approach one another in physical space, to prevent particle interpenetration. In this case, entropy is generated by the viscous force at a rate

$$\frac{dA_i}{dt} = \frac{1}{2} \frac{\gamma - 1}{\rho_i^{\gamma-1}} \sum_{j=1}^N m_j \Pi_{ij} \mathbf{v}_{ij} \cdot \nabla_i \overline{W}_{ij}, \quad (3.8)$$

transforming kinetic energy of gas motion irreversibly into heat.

We have also run a few simulations with a ‘conventional formulation’ of SPH in order to compare its results with the ‘entropy formulation’. This conventional formulation is characterized by the following differences. Equation (3.2) is replaced by a choice of smoothing length that keeps the number of neighbors constant. In the equation of motion, the coefficients f_i and f_j are always equal to unity, and finally, the entropy is replaced by the thermal energy per unit mass as an independent thermodynamic variable. The thermal energy is then evolved as

$$\frac{du_i}{dt} = \sum_{j=1}^N m_j \left(\frac{P_i}{\rho_i^2} + \frac{1}{2} \Pi_{ij} \right) \mathbf{v}_{ij} \cdot \nabla_i \overline{W}_{ij}, \quad (3.9)$$

with the particle pressures being defined as $P_i = (\gamma - 1) \rho_i u_i$.

3.3.2 Gravitational method

In the GADGET code, both the collisionless dark matter and the gaseous fluid are represented by particles, allowing the self-gravity of both components to be computed with

gravitational N-body methods. Assuming a periodic box of size L , the forces can be formally computed as the gradient of the periodic peculiar potential ϕ , which is the solution of

$$\nabla^2\phi(\mathbf{x}) = 4\pi G \sum_i m_i \left[-\frac{1}{L^3} + \sum_{\mathbf{n}} \tilde{\delta}(\mathbf{x} - \mathbf{x}_i - \mathbf{n}L) \right], \quad (3.10)$$

where the sum over $\mathbf{n} = (n_1, n_2, n_3)$ extends over all integer triples. The function $\tilde{\delta}(\mathbf{x})$ is a normalized softening kernel, which distributes the mass of a point-mass over a scale corresponding to the gravitational softening length ϵ . The GADGET code adopts the spline kernel used in SPH for $\tilde{\delta}(\mathbf{x})$, with a scale length chosen such that the force of a point mass becomes fully Newtonian at a separation of 2.8ϵ , with a gravitational potential at zero lag equal to $-Gm/\epsilon$, allowing the interpretation of ϵ as a Plummer equivalent softening length.

Evaluating the forces by direct summation over all particles becomes rapidly prohibitive for large N owing to the inherent N^2 scaling of this approach. Tree algorithms such as that used in GADGET overcome this problem by using a hierarchical multipole expansion in which distant particles are grouped into ever larger cells, allowing their gravity to be accounted for by means of a single multipole force. Instead of requiring $N - 1$ partial forces per particle, the gravitational force on a single particle can then be computed from just $\mathcal{O}(\log N)$ interactions.

It should be noted that the final result of the tree algorithm will in general only represent an approximation to the true force described by Eqn. (3.10). However, the error can be controlled conveniently by adjusting the opening criterion for tree nodes, and, provided sufficient computational resources are invested, the tree force can be made arbitrarily close to the well-specified correct force.

The summation over the infinite grid of particle images required for simulations with periodic boundary conditions can also be treated in the tree algorithm. GADGET uses the technique proposed by Hernquist et al. [189] for this purpose. Alternatively, the new version GADGET-2 used in this study allows the pure tree algorithm to be replaced by a hybrid method consisting of a synthesis of the particle-mesh method and the tree algorithm. GADGET's mathematical implementation of this so-called TreePM method [190, 191, 192] is similar to that of Bagla [193]. The potential of Eqn. (3.10) is explicitly split in Fourier space into a long-range and a short-range part according to $\phi_{\mathbf{k}} = \phi_{\mathbf{k}}^{\text{long}} + \phi_{\mathbf{k}}^{\text{short}}$, where

$$\phi_{\mathbf{k}}^{\text{long}} = \phi_{\mathbf{k}} \exp(-\mathbf{k}^2 r_s^2), \quad (3.11)$$

with r_s describing the spatial scale of the force-split. This long range potential can be computed very efficiently with mesh-based Fourier methods. Note that if r_s is chosen slightly larger than the mesh scale, force anisotropies that exist in plain PM methods can be suppressed to essentially arbitrarily small levels.

The short range part of the potential can be solved in real space by noting that for $r_s \ll L$ the short-range part of the potential is given by

$$\phi^{\text{short}}(\mathbf{x}) = -G \sum_i \frac{m_i}{r_i} \text{erfc}\left(\frac{r_i}{2r_s}\right). \quad (3.12)$$

Here $r_i = \min(|\mathbf{x} - \mathbf{r}_i - \mathbf{n}L|)$ is defined as the smallest distance of any of the images of particle i to the point \mathbf{x} . The short-range force can still be computed by the tree algorithm, except that the force law is modified according to Eqn. (3.12). However, the tree only needs to be walked in a small spatial region around each target particle (because the complementary error function rapidly falls for $r > r_s$), and no corrections for periodic boundary conditions are required, which together can result in a very substantial performance gain. One typically also gains accuracy in the long range force, which is now basically exact, and not an approximation as in the tree method. In addition, the TreePM approach maintains all of the most important advantages of the tree algorithm, namely its insensitivity to clustering, its essentially unlimited dynamic range, and its precise control about the softening scale of the gravitational force.

3.4 The simulation set

In all of our simulations, we adopt the standard concordance cold dark matter model of a flat universe with $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$, $\sigma_8 = 0.9$, $n = 1$, and $h = 0.7$. For simulations including hydrodynamics, we take the baryon mass density to be $\Omega_b = 0.04$. The simulations are initialized at redshift $z = 99$ using the Eisenstein & Hu [194] transfer function. For the dark matter-only runs, we chose a periodic box of comoving size $12 h^{-1}$ Mpc, while for the adiabatic runs we preferred $3 h^{-1}$ Mpc to achieve higher mass resolution, although the exact size of the simulation box is of little importance for the present comparison. Note however that this is different in simulations that also include cooling, which imprints additional physical scales. We place the unperturbed dark matter particles at the vertices of a Cartesian grid, with the gas particles offset by half the mean interparticle separation in the GADGET simulations. These particles are then perturbed by the Zel'dovich approximation for the initial conditions. In Enzo, fluid elements are represented by the values at the center of the cells and are also perturbed using the Zel'dovich approximation.

For both codes, we have run a large number of simulations, varying the resolution, the physics (dark matter only, or dark matter with adiabatic hydrodynamics), and some key numerical parameters. Most of these simulations have been evolved to redshift $z = 3$. We give a full list of all simulations we use in this study in Tables 3.1 and 3.2 for GADGET and Enzo, respectively; below we give some further explanations for this simulation set.

We performed a suite of dark matter-only simulations in order to compare the gravity solvers in Enzo and GADGET. For GADGET, the spatial resolution is determined by

the gravitational softening length ϵ , while for Enzo the equivalent quantity is given by the smallest allowed mesh size e (note that in Enzo the gravitational force resolution is approximately twice as coarse as this: see Section 2.2.4). Together with the box size L_{box} , we can then define a dynamic range L_{box}/e to characterize a simulation (for simplicity we use L_{box}/e for GADGET as well instead of L_{box}/ϵ). For our basic set of runs with 64^3 dark matter particles we varied L_{box}/e from 256 to 512, 1024, 2048 and 4096 in Enzo. We also computed corresponding GADGET simulations, except for the $L_{\text{box}}/e = 4096$ case, which presumably would already show sizable two-body scattering effects. Note that it is common practice to run collisionless tree N-body simulations with softening in the range $1/25 - 1/30$ of the mean interparticle separation, translating to $L_{\text{box}}/e = 1600 - 1920$ for a 64^3 simulation.

Unlike in GADGET, the force accuracy in Enzo at early times also depends on the root grid size. For most of our runs we used a root grid with 64^3 cells, but we also performed Enzo runs with a 128^3 root grid in order to test the effect of the root grid size on the dark matter halo mass function. Both 64^3 and 128^3 particles were used, with the number of particles never exceeding the size of the root grid.

Our main interest in this study lies, however, in our second set of runs, where we additionally follow the hydrodynamics of a baryonic component, modeled here as an ideal, non-radiative gas. As above, we use 64^3 DM particles and 64^3 gas particles (for GADGET), or a 64^3 root grid (for Enzo), in most of our runs, though as before we also perform runs with 128^3 particles and root grids. Again, we vary the dynamic range L_{box}/e from 256 to 4096 in Enzo, and parallel this with corresponding GADGET runs, except for the $L_{\text{box}}/e = 4096$ case.

An important parameter of the AMR method is the mesh-refinement criterion. Usually, Enzo runs are configured such that grid refinement occurs when the dark matter mass in a cell reaches a factor of 4.0 times the mean dark matter mass expected in a cell at root grid level, or if it has a factor of 8.0 times the mean baryonic mass of a root level cell, but several runs were performed with a threshold density set to 0.5 of the standard values for both dark matter and baryon density. All that the “refinement overdensity” criteria does is set the maximum gas or dark matter mass which may exist in a given cell before that cell must be refined based on a multiple of the mean cell mass on the root grid. For example, a baryon overdensity threshold of 4.0 means that a cell is forced to refine once a cell has accumulated more than 4 times the mean cell mass on the root grid.

When the refinement overdensity is set to the higher value discussed here, the simulation may fail to properly identify small density peaks at early times, so that they are not well resolved by placing refinements on them. As a result, the formation of low-mass DM halos or substructure in larger halos may be suppressed. Note that lowering the refinement threshold results in a significant increase in the number of refined grids, and hence a significant increase in the computational cost of a simulation; i.e., one must tune

GADGET simulations

Run	L_{box}/e	N_{part}	m_{DM}	m_{gas}	ϵ	notes
L12N64_dm	2048	64^3	5.5×10^8	—	5.86	DM only
L12N128_dm	3840	128^3	6.9×10^7	—	3.13	DM only
L12N256_dm	7680	256^3	8.6×10^6	—	1.56	DM only
L3N64_3.1e	256	2×64^3	7.4×10^6	1.1×10^6	11.7	Adiabatic
L3N64_3.2e	512	2×64^3	7.4×10^6	1.1×10^6	5.86	Adiabatic
L3N64_3.3e	1024	2×64^3	7.4×10^6	1.1×10^6	2.93	Adiabatic
L3N64_3.4e	2048	2×64^3	7.4×10^6	1.1×10^6	1.46	Adiabatic
L3N128	3200	2×128^3	9.3×10^5	1.4×10^5	0.78	Adiabatic
L3N256	6400	2×256^3	1.2×10^5	1.8×10^4	0.39	Adiabatic

Table 3.1: List of GADGET cosmological simulations that are used in this study. L_{box}/e is the dynamic range, and N_{part} is the particle number (in the adiabatic runs there are identical numbers of dark matter and gas particles). m_{DM} and m_{gas} are the masses of the dark matter and gas particles in units of [$h^{-1} M_{\odot}$]. ϵ is the Plummer-equivalent gravitational softening length in units of [$h^{-1} \text{kpc}$], but the GADGET code adopts the spline kernel. See Section 3.3.2 for more details.

the refinement criteria to compromise between performance and accuracy.

We also performed simulations with higher mass and spatial resolution, ranging up to 2×256^3 particles with GADGET, and 128^3 dark matter particles and a 128^3 root grid with Enzo. For DM-only runs, the gravitational softening lengths in these higher resolution GADGET runs were taken to be 1/30 of the mean dark matter interparticle separation, giving a dynamic range of $L_{\text{box}}/e = 3840$ and 7680 for 128^3 and 256^3 particle runs, respectively. For the adiabatic GADGET runs, they were taken to be 1/25 of the mean interparticle separation, giving $L_{\text{box}}/e = 3200$ and 6400 for the 128^3 and 256^3 particle runs, respectively. All Enzo runs used a maximum refinement ratio of $L_{\text{box}}/e = 4096$.

As an example, we show the spatial distribution of the projected dark matter and gas mass in Figure 3.1 from one of the representative adiabatic gas runs of GADGET and Enzo. The mass distribution in the two simulations are remarkably similar for both dark matter and gas, except that one can see slightly finer structures in GADGET gas mass distribution compared to that of Enzo. The good visual agreement in the two runs is very encouraging, and we will analyze the two simulations quantitatively in the following sections.

Enzo simulations

Run	L_{box}/e	N_{DM}	N_{root}	notes
64g64d.6l_dm_hod	4096	64^3	64^3	DM only, high od
128g64d.5l_dm_hod	4096	64^3	128^3	DM only, high od
128g128d.5l_dm_hod	4096	128^3	128^3	DM only, high od
64g64d.6l_dm_lod	4096	64^3	64^3	DM only, low od
128g64d.5l_dm_lod	4096	64^3	128^3	DM only, low od
128g128d.5l_dm_lod	4096	128^3	128^3	DM only, low od
64g64d.6l_z	4096	64^3	64^3	Adiabatic, ZEUS
64g64d.6l_z_lod	4096	64^3	64^3	Adiabatic, ZEUS, low OD
64g64d.6l_q0.5	4096	64^3	64^3	Adiabatic, ZEUS, $Q_{\text{AV}} = 0.5$
128g64d.5l_z	4096	64^3	128^3	Adiabatic, ZEUS
128g64d.5l_z_lod	4096	64^3	128^3	Adiabatic, ZEUS, low OD
128g128d.5l_z	4096	128^3	128^3	Adiabatic, ZEUS
64g64d.2l_ppm	256	64^3	64^3	Adiabatic, PPM
64g64d.3l_ppm	512	64^3	64^3	Adiabatic, PPM
64g64d.4l_ppm	1024	64^3	64^3	Adiabatic, PPM
64g64d.5l_ppm	2048	64^3	64^3	Adiabatic, PPM
64g64d.6l_ppm	4096	64^3	64^3	Adiabatic, PPM
64g64d.6l_ppm_lod	4096	64^3	64^3	Adiabatic, PPM, low OD
128g64d.5l_ppm	4096	64^3	128^3	Adiabatic, PPM
128g64d.5l_ppm_lod	4096	64^3	128^3	Adiabatic, PPM, low OD
128g128d.5l_ppm	4096	128^3	128^3	Adiabatic, PPM
128g128d.5l_ppm_lod	4096	128^3	128^3	Adiabatic, PPM, low OD

Table 3.2: List of Enzo simulations used in this study. L_{box}/e is the dynamic range (e is the size of the finest resolution element, i.e. the spatial size of the finest level of grids), N_{DM} is the number of dark matter particles, and N_{root} is the size of the root grid. ‘ZEUS’ and ‘PPM’ in the notes indicate the adopted hydrodynamic method. ‘low OD’ means that the low overdensity threshold for refinement were chosen (cells refine with a baryon overdensity of 4.0/dark matter density of 2.0). ‘ Q_{AV} ’ is the artificial viscosity parameter for the ZEUS hydro method when it is not the default value of 2.0.




Figure 3.1: Projected dark matter (top row) and gas mass (bottom row) distribution for GADGET and Enzo in a slab of size $3 \times 3 \times 0.75 (h^{-1} \text{ Mpc})^3$. For GADGET (left column), we used the run with 2×64^3 particles. For Enzo (right column), the run with 64^3 dark matter particles and 128^3 root grid was used.

3.5 Simulations with dark matter only

According to the currently favored theoretical model of the CDM theory, the material content of the universe is dominated by as of yet unidentified elementary particles which interact so weakly that they can be viewed as a fully collisionless component at spatial scales of interest for large-scale structure formation. The mean mass density in this cold dark matter far exceeds that of ordinary baryons, by a factor of $\sim 5 - 7$ in the currently favored Λ CDM cosmology. Since structure formation in the Universe is primarily driven by gravity it is of fundamental importance that the dynamics of the dark matter and the self-gravity of the hydrodynamic component are simulated accurately by any cosmological code. In this section we discuss simulations that only follow dark matter in order to compare Enzo and GADGET in this respect.

3.5.1 Dark matter power spectrum

One of the most fundamental quantities to characterize the clustering of matter is the power spectrum of dark matter density fluctuations. In Figure 3.2 we compare the power spectra of DM-only runs at redshifts $z = 10$ and 3. The short-dashed curve is the linearly evolved power spectrum based on the transfer function of Eisenstein & Hu [194], while the solid curve gives the expected nonlinear power spectrum calculated with the Peacock & Dodds [195] scheme. We calculate the dark matter power spectrum in each simulation by creating a uniform grid of dark matter densities. The grid resolution is twice as fine as the mean interparticle spacing of the simulation (i.e. a simulation with 128^3 particles will use a 256^3 grid to calculate the power spectrum) and densities are generated with the triangular-shaped cloud (TSC) method. A fast Fourier transform is then performed on the grid of density values and the power spectrum is calculated by averaging the power in logarithmic bins of wavenumber. We do not attempt to correct for shot-noise or the smoothing effects of the TSC kernel.

The results of all GADGET and Enzo runs with 128^3 root grid agree well with each other at both epochs up to the Nyquist wavenumber. However, the Enzo simulations with a 64^3 root grid deviate on small scales from the other results significantly, particularly at $z = 10$. This can be understood to be a consequence of the particle-mesh technique adopted as the gravity solver in the AMR code, which induces a softening of the gravitational force on the scale of one mesh cell (this is a property of all PM codes, not just Enzo). To obtain reasonably accurate forces down to the scale of the interparticle spacing, at least two cells per particle spacing are therefore required at the outset of the calculation. In particular, the force accuracy of Enzo is much less accurate at small scales at early times when compared to GADGET because before significant overdensities develop the code does not adaptively refine any regions of space (and therefore increased force resolution to include small-scale force corrections). GADGET is a tree-PM code –

at short range, forces on particles are calculated using the tree method, which offers a force accuracy that is essentially independent of the clustering state of the matter down to the adopted gravitational softening length (see Section 3.3.2 for details).

However, as the simulation progresses in time and dark matter begins to cluster into halos, the force calculation by Enzo becomes more accurate as additional levels of grids are adaptively added to the high density regions, reducing the discrepancy seen between Enzo and GADGET at redshift $z = 10$ to something much smaller at $z = 3$.

3.5.2 Halo dark matter mass function and halo positions

We have identified dark matter halos in the simulations using a standard Friends-of-Friends algorithm with a linking length of 0.2 in units of the mean interparticle separation. In this section, we consider only halos with more than 32 particles. We obtained nearly identical results to those described in this section using the HOP halo finder [196].

In Figure 3.3, we compare the cumulative DM halo mass function for several simulations with 64^3 , 128^3 and 256^3 dark matter particles as a function of L_{box}/e and particle mass. In the bottom panel, we show the residual in logarithmic space with respect to the Sheth-Tormen mass function, i.e., $\log(N>M) - \log(\text{S\&T})$. The agreement between Enzo and GADGET simulations at the high-mass end of the mass function is reasonable, but at lower masses there is a systematic difference between the two codes. The Enzo run with 64^3 root grid contains significantly fewer low mass halos compared to the GADGET simulations. Increasing the root grid size to 128^3 brings the low-mass end of the Enzo result closer to that of GADGET.

This highlights the importance of the size of the root grid in the adaptive particle-mesh method based AMR simulations. Eulerian simulations using the particle-mesh technique require a root grid twice as fine as the mean interparticle separation in order to achieve a force resolution at early times comparable to tree methods or so-called P³M methods [138], which supplement the softened PM force with a direct particle-particle (PP) summation on the scale of the mesh. Having a conservative refinement criterion together with a coarse root grid in AMR is not sufficient to improve the low mass end of the halo mass function because the lack of force resolution at early times effectively results in a loss of small-scale power, which then prevents many low mass halos from forming.

We have also directly compared the positions of individual dark matter halos identified in a simulation with the same initial conditions, run both with GADGET and Enzo. This run had 64^3 dark matter particles and a $L_{\text{box}} = 12h^{-1}$ Mpc box size. For GADGET, we used a gravitational softening equivalent to $L_{\text{box}}/e = 2048$. For Enzo, we used a 128^3 root grid, a low overdensity threshold for the refinement criteria, and we limited refinements to a dynamic range of $L_{\text{box}}/e = 4096$ (5 total levels of refinement).

In order to match up halos, we apply the following method to identify “pairs” of halos




Figure 3.2: Dark matter power spectra at $z = 10$ and $z = 3$ for both Enzo and GADGET simulations with 64^3 dark matter particles, $L_{\text{box}} = 12 h^{-1} \text{Mpc}$ (comoving) and varying spatial resolution. The short-dashed curve in each panel is the linear power spectrum predicted by theory using the transfer function of Eisenstein & Hu [194]. The solid curve in each panel is the non-linear power spectrum calculated with the Peacock & Dodds method. [195] Arrows indicate the largest wavelength that can be accurately represented in the simulation initial conditions ($k = 2\pi/L_{\text{box}}$) and those that correspond to the Nyquist frequencies of 64^3 , 128^3 , and 256^3 Enzo root grids.




Figure 3.3: Cumulative mass functions at $z = 3$ for dark matter-only Enzo & GADGET runs with 64^3 particles and a comoving box size of $L_{\text{box}} = 12h^{-1}\text{Mpc}$. All Enzo runs have $L_{\text{box}}/e = 4096$. The solid black line denotes the Sheth & Tormen [198] mass function. In the bottom panel, we show the residual in logarithmic space with respect to the Sheth & Tormen mass function, i.e., $\log(N>M) - \log(\text{S\&T})$.




Figure 3.4: *Left column:* Probability distribution function of the number of dark matter halos as a function of the separation of the matched halo pair in corresponding Enzo and GADGET simulations (see text for the details of the runs used in this comparison). The separation is in units of the initial mean interparticle separation, Δ . The shaded region in the distribution function shows the quartiles on both sides of the median value (which is shown by the arrows) of the distribution. *Right column:* Separation of each pair (in units of Δ) vs. mean halo mass of each pair. The top row is of pairs whose masses agree to within 10% (i.e. $f_M = 1.1$) and the bottom row is of pairs whose masses agree to within a factor of two (i.e. $f_M = 2.0$).

with approximately the same mass and center-of-mass position. First, we sort the halos in order of decreasing mass, and then select a halo from the massive end of one of the two simulations (i.e. the beginning of the list). Starting again from the massive end, we then search the other list of halos for a halo within a distance of $r_{\max} = f_R \Delta$, where Δ is the mean interparticle separation (1/64 of the box size in this case) and f_R is a dimensionless number (chosen here to be either 0.5 or 1.0). If the halo masses are also within a fraction f_M of one another, then the two halos in question are counted as a ‘matched pair’ and removed from the lists to avoid double-counting. This procedure is continued until there are no more halos left that satisfy these criteria.

In the left column of Figure 3.4, we show the distribution of pair separations obtained in this way. The arrow indicates the median value of the distribution, and the quartile on each side of the median value is indicated by the shaded region. The values of r_{\max} and f_M are also shown in each panel. A conservative matching-criterion that allows only a 10% deviation in halo mass and half a cell of variation in the position (i.e. $r_{\max} = 0.5\Delta$, $f_M = 1.1$) finds only 117 halo pairs (out of ~ 292 halos in each simulation) with a median separation of 0.096Δ between the center-of-mass positions of halos. Increasing r_{\max} to 1.0Δ does very little to increase the number of matched halos. Keeping $r_{\max} = 0.5\Delta$ and increasing f_M to 2.0 gives us 252 halo pairs with a median separation of 0.128Δ . Increasing f_M any further does little to increase the number of matched pairs, and looking further away than $r_{\max} = 1.0\Delta$ produces spurious results in some cases, particularly for low halo masses.

This result therefore suggests that the halos are typically in almost the same places in both simulations, but that their individual masses show somewhat larger fluctuations. Note however that a large fraction of this scatter simply stems from noise inherent in the group sizes obtained with the halo finding algorithms used. The friends-of-friends algorithm often links (or not yet links) infalling satellites across feeble particle bridges with the halo, so that the numbers of particles linked to a halo can show large variations between simulations even though the halo’s virial mass is nearly identical in the runs. We also tested the group finder HOP [196], but found that it also shows significant noise in the estimation of halo masses. It may be possible to reduce the latter by experimenting with the adjustable parameters of this group finder (one of which controls the “bridging problem” that the friends-of-friends method is susceptible to), but we have not tried this.

In the right panels of Figure 3.4, we plot the separation of halo pairs against the average mass of the two halos in question. Clearly, pairs of massive halos tend to have smaller separations than low mass halos. Note that some of the low mass halos with large separation ($L/\Delta > 0.4$) could be false identifications. It is very encouraging, however, that the massive halos in the two simulations generally lie within 1/10 of the initial mean interparticle separation. The slight differences in halo positions may be caused by timing differences between the two simulation codes.

3.5.3 Halo dark matter substructure

Another way to compare the solution accuracy of the N-body problem in the two codes is to examine the substructure of dark matter halos. The most massive halos in the 128^3 particle dark matter-only simulations discussed in this paper have approximately 11,000 particles, which is enough to marginally resolve substructure. We look for gravitationally-bound substructure using the SUBFIND method described in Springel et al. [197], which we briefly summarize here for clarity. The process is as follows: a Friends-of-Friends group finder is used (with the standard linking length of 0.2 times the mean interparticle spacing) to find all of the dark matter halos in the simulations. We then select the two most massive halos in the calculation (each of which has at least 11,000 particles in both simulations) and analyze them with the subhalo finding algorithm. This algorithm first computes a local estimate of the density at the positions of all particles in the input group, and then finds locally overdense regions using a topological method. Each of the substructure candidates identified in this way is then subjected to a gravitational unbinding procedure where only particles bound to the substructure are kept. If the remaining self-bound particle group has more than some minimum number of particles it is considered to be a subhalo. We use identical parameters for the Friends-of-Friends and subhalo detection calculations for both the Enzo and GADGET dark matter-only calculations.

Figure 3.5 shows the projected dark matter density distribution and substructure mass function for the two most massive halos in the 128^3 particle DM-only calculations for both Enzo and GADGET, which have dark matter masses close to $M_{\text{halo}} \sim 10^{12} M_{\odot}$. Bound subhalos are indicated by different colors, with identical colors being used in both simulations to denote the most massive subhalo, second most massive, etc. Qualitatively, the halos have similar overall morphologies in both calculations, though there are some differences in the substructures. The masses of these two parent halos in the Enzo calculation are $8.19 \times 10^{11} M_{\odot}$ and $7.14 \times 10^{11} M_{\odot}$, and we identify total 20 and 18 subhalos, respectively. The corresponding halos in the GADGET calculation have masses of $8.27 \times 10^{11} M_{\odot}$ and $7.29 \times 10^{11} M_{\odot}$, and they have 7 and 10 subhalos. Despite the difficulty of Enzo in fully resolving the low-mass end of the halo mass function, the code apparently has no problem in following dark matter substructure within large halos, and hosts larger number of small subhalos than the GADGET calculation. Some corresponding subhalos in the two calculations appear to be slightly off-set. Overall, the agreement of the substructure mass functions for the intermediate mass regime of subhalos is relatively good and within the expected noise.

It is not fully clear what causes the observed differences in halo substructure between the two codes. It may be due to lack of spatial and/or dark matter particle mass resolution in the calculations – typically simulations used for substructure studies have at least an order of magnitude more dark matter particles per halo than we have here. It is

also possible that systematics in the grouping algorithm are responsible for some of the differences.

3.6 Adiabatic simulations

In this section, we start our comparison of the fundamentally different hydrodynamical algorithms of Enzo and GADGET. It is important to keep in mind that a direct comparison between the AMR and SPH methods when applied to cosmic structure formation will always be convolved with a comparison of the gravity solvers of the codes. This is because the process of structure formation is primarily driven by gravity, to the extent that hydrodynamical forces are subdominant in most of the volume of the universe. Differences that originate in the gravitational dynamics will in general induce differences in the hydrodynamical sector as well, and it may not always be straightforward to cleanly separate those from genuine differences between the AMR and SPH methods themselves. Given that the dark matter comparisons indicate that one must be careful to appropriately resolve dark matter forces at early times unless relatively fine root grids are used for Enzo calculations, it is clear that any difference found between the codes needs to be regarded with caution until confirmed with AMR simulations of high gravitational force resolution.

Having made these cautionary remarks, we will begin our comparison with a seemingly trivial test of a freely expanding universe without perturbations, which is useful to check conservation of entropy (for example). After that, we will compare the gas properties found in cosmological simulations of the Λ CDM model in more detail.

3.6.1 Unperturbed adiabatic expansion test

Unperturbed ideal gas in an expanding universe should follow Poisson’s law of adiabatic expansion: $T \propto V^{\gamma-1} \propto \rho^{1-\gamma}$. Therefore, if we define entropy as $S \equiv T/\rho^{\gamma-1}$, it should be constant for an adiabatically expanding gas.

This simple relation suggests a straightforward test of how well the hydrodynamic codes described in Chapter 2 and Section 3.3 conserve entropy [185]. To this end, we set up unperturbed simulations for both Enzo and GADGET with 16^3 grid cells or particles, respectively. The runs are initialized at $z = 99$ with uniform density and temperature $T = 10^4$ K. This initial temperature was deliberately set to a higher value than expected for the real universe in order to avoid hitting the temperature floor set in the codes while following the adiabatic cooling of gas due to the expansion of the universe. The box was then allowed to expand until $z = 3$. Enzo runs were performed using both the PPM and ZEUS algorithms and GADGET runs were done with both ‘conventional’ and the ‘entropy conserving’ formulation of SPH.




Figure 3.5: Dark matter substructure in both Enzo and GADGET dark matter-only calculations with 128^3 particles. The Enzo simulations use the “low overdensity” refinement parameters. Left column: data from the most massive halo in the simulation volume. Right column: second most massive halo. Top row: Projected dark matter density for halos in the Enzo simulation with substructure color-coded. Middle row: projected dark matter density for GADGET simulations. Bottom row: Halo substructure mass function for each halo with both Enzo and GADGET results plotted together, with units of number of halos greater than a given mass on the y axis and number of particles on the x axis. In these simulations the dark matter particle mass is $9.82 \times 10^7 M_\odot$, resulting in total halo masses of $\sim 10^{12} M_\odot$. In the top and middle rows subhalos with the same color correspond to the most massive, second most massive, etc. subhalos. In the Enzo calculation all subhalos beyond the 10th most massive are shown using the same color. Both sets of halos have masses of $\sim 10^{12} M_\odot$. The x and y axes in the top two rows are in units of comoving kpc/h. In the bottom row, Enzo results are shown as a black solid line and GADGET results are shown as a red dashed line.




Figure 3.6: Fractional deviations from the expected adiabatic relation for temperature, comoving gas density and entropy as a function of redshift in simulations of unperturbed adiabatic expansion test. *Left column:* the ‘entropy conserving’ formulation of SPH (top panel) and the ‘conventional’ formulation (bottom panel). *Right column:* The PPM (top panel) and ZEUS (bottom panel) hydrodynamic methods in Enzo. Error bars in all panels show the variance of each quantity. The short-long-dashed line in the bottom right panel shows the case where the maximum timestep is limited to be 1/10 of the default maximum. Note the difference in scales of the y axes in the bottom row.

In Figure 3.6 we show the fractional deviation from the expected adiabatic relation for density, temperature, and entropy. The GADGET results (left column) show that the ‘entropy conserving’ formulation of SPH preserves the entropy very well, as expected. There is a small net decrease in temperature and density of only $\sim 0.1\%$, reflecting the error of SPH in estimating the mean density. In contrast, in the ‘conventional’ SPH formulation the temperature and entropy deviate from the adiabatic relation by 15%, while the comoving density of each gas particle remains constant. This systematic drift is here caused by a small error in estimating the local velocity dispersion owing to the expansion of the universe. In physical coordinates, one expects $\nabla \cdot \mathbf{v} = 3H(a)$, but in conventional SPH, the velocity divergence needs to be estimated with a small number of discrete particles, which in general will give a result that slightly deviates from the continuum expectation of $3H(a)$. In our test, this error is the same for all particles, without having a chance to average out for particles with different neighbor configurations, hence resulting in a substantial systematic drift. In the entropy formulation of SPH, this problem is absent by construction.

In the Enzo/PPM run (top right panel), there is a net decrease of only $\sim 0.1\%$ in temperature and entropy, whereas in Enzo/ZEUS (bottom right panel), the temperature and entropy drop by 12% between $z = 99$ and $z = 3$. The comoving gas density remains constant in all Enzo runs. In the bottom right panel, the short-long-dashed line shows an Enzo/ZEUS run where we lowered the maximum expansion of the simulation volume during a single timestep (i.e. $\Delta a/a$, where a is the scale factor) by a factor of 10. This results in a factor of ~ 10 reduction of the error, such that the fractional deviation from the adiabatic relation is only about 1%. This behavior is to be expected since the ZEUS hydrodynamic algorithm is formally first-order-accurate in time in an expanding universe.

In summary, these results show that both the Enzo/ZEUS hydrodynamic algorithm and the conventional SPH formulation in GADGET have problems in reliably conserving entropy. However, these problems are essentially absent in Enzo/PPM and the new SPH formulation of GADGET.

3.6.2 Differential distribution functions of gas properties

We now begin our analysis of gas properties found in full cosmological simulations of structure formation. In Figures 3.7 and 3.8 we show mass-weighted one-dimensional differential probability distribution functions of gas density, temperature and entropy, for redshifts $z = 10$ (Figure 3.7) and $z = 3$ (Figure 3.8). We compare results for GADGET and Enzo simulations at different numerical resolution, and run with both the ZEUS and PPM formulations of Enzo.

At $z = 10$, effects owing to an increase of resolution are clearly seen in the distribution of gas overdensity (left column), with runs of higher resolution reaching higher densities

earlier than those of lower resolution. However, this discrepancy becomes smaller at $z = 3$ because lower resolution runs tend to ‘catch up’ at late times, indicating that then more massive structures, which are also resolved in the lower resolution simulations, become ever more important. One can also see that the density distribution becomes wider at $z = 3$ compared to those at $z = 10$, reaching to higher gas densities at lower redshift.

At $z = 3$, both Enzo and GADGET simulations agree very well at $\log T > 3.5$ and $\log S > 21.5$, with a characteristic shoulder in the temperature (middle column) and a peak in the entropy (right column) distributions at these values. This can be understood with a simple analytic estimate of gas properties in dark matter halos. We estimate the virial temperature of a dark matter halo with mass $10^8 M_\odot$ ($10^{11} M_\odot$) at $z = 3$ to be $\log T = 3.7$ (5.7). Assuming a gas overdensity of 200, the corresponding entropy is $\log S = 21.9$ (23.9). The good agreement in the distribution functions at $\log T > 3.5$ and $\log S > 21.5$ therefore suggests that the properties of gas inside the dark matter halos agree reasonably well in both simulations. The gas in the upper end of the distribution is in the most massive halos in the simulation, with masses of $\sim 10^{11} M_\odot$ at $z = 3$. Enzo has a built-in temperature floor of 1 Kelvin, resulting in an artificial feature in the temperature and entropy profiles at $z = 3$. GADGET also has a temperature floor, but it is set to 0.1 Kelvin and is much less noticeable since that temperature is not attained in this simulation. Note that the entropy floor stays at the constant value of $\log S_{\text{init}} = 18.44$ for all simulations at both redshifts.

However, there are also some interesting differences in the distribution of temperature and entropy between Enzo/PPM and the other methods for gas of low overdensity. Enzo/PPM exhibits a ‘dip’ at intermediate temperature ($\log T \sim 2.0$) and entropy ($\log S \sim 20$), whereas Enzo/ZEUS and GADGET do not show the resulting bimodal character of the distribution. We will revisit this feature when we examine two dimensional phase-space distributions of the gas in Section 3.6.4, and again in Section 3.7 when we examine numerical effects due to artificial viscosity. In general, the GADGET results appear to lie in between those obtained with Enzo/ZEUS and Enzo/PPM, and are qualitatively more similar to the Enzo/ZEUS results.

3.6.3 Cumulative distribution functions of gas properties

In this section we study cumulative distribution functions of the quantities considered above, highlighting the quantitative differences in the distributions in a more easily accessible way. In Figures 3.9 and 3.10 we show the mass-weighted cumulative distribution functions of gas overdensity, temperature and entropy at $z = 10$ (Figure 3.9) and $z = 3$ (Figure 3.10). The measurements parallel those described in Section 3.6.2, and were done for the same simulations.

We observe similar trends as before. At $z = 10$ in the GADGET simulations, 70% of the total gas mass is in regions above the mean density of baryons, but in Enzo, only 50%




Figure 3.7: Probability distribution functions of gas mass as functions of gas overdensity (left column), temperature (middle column) and entropy (right column) at $z = 10$. For GADGET, runs with 2×64^3 (red solid line), 2×128^3 (red short-dashed line) and 2×256^3 (red long-dashed line) particles are shown. The dynamic range of the Enzo simulations were fixed to $L_{\text{box}}/e = 4096$, but the particle numbers and the root grid size were varied between 64^3 and 128^3 . Both the ZEUS and PPM hydro methods were used in the Enzo calculations. The Enzo line types are: 128g/128dm PPM lowod (black dash-dotted line), 128g/128dm ZEUS (black dotted line), and 64g/64dm PPM lowod (black long dash-short dashed line). In the bottom panels, the residuals in logarithmic scale with respect to the GADGET N256 run are shown.




Figure 3.8: Probability distribution functions of gas mass as functions of gas overdensity (left column), temperature (middle column) and entropy (right column) at $z = 3$. For GADGET, runs with 2×64^3 , 2×128^3 and 2×256^3 particles were used. The dynamic range of the Enzo simulations were fixed to $L_{\text{box}}/e = 4096$, but the particle numbers and the root grid size were varied between 64^3 and 128^3 (e.g. 64dm/128grid means 64^3 DM particles and 128^3 root grid). Both the ZEUS and PPM hydro methods were used in the Enzo calculations, as shown in the figure key. Lines are identical to those in Figure 3.7. In the bottom panels, we show the residuals in logarithmic scale with respect to the GADGET N256 run.

is in such regions. This mass fraction increases to 80% in GADGET runs, and to 70% in Enzo runs at $z = 3$, as more gas falls into the potential wells of dark matter halos.

More distinct differences can be observed in the distribution of temperature and entropy. At $z = 10$, only 10 – 20% of the total gas mass is heated to temperatures above $\log T = 0.5$ in Enzo/PPM, whereas this fraction is 70 – 75% in Enzo/ZEUS, and 35 – 55% in GADGET. At $z = 3$, the mass fraction that has temperature $\log T > 0.5$ is 40 – 60% for Enzo/PPM, and $\sim 80\%$ for both Enzo/ZEUS and GADGET. Similar mass fractions can be observed for gas with entropy $\log S > 18.5 - 19.0$.

In summary, these results show that both GADGET and particularly Enzo/ZEUS tend to heat up a significant amount of gas at earlier times than Enzo/PPM. This may be related to differences in the parameterization of numerical viscosity, a topic that we will discuss in more detail in Section 3.7.

3.6.4 Phase diagrams

In Figure 3.11 we show the redshift evolution of the mass-weighted two-dimensional distribution of entropy vs. gas overdensity for redshifts $z = 30, 10$ and 3 (top to bottom rows). Two representative GADGET simulations with 2×64^3 and 2×256^3 particles are shown in the left two columns. The Enzo simulations shown in the right two columns both have a maximum dynamic range of $L_{\text{box}}/e = 4096$ and use 128^3 dark matter particles with a 128^3 root grid. They differ in that the simulation in the rightmost column uses the PPM hydrodynamic method, while the other column uses the ZEUS method.

The gas is initialized at $z = 99$ at a temperature of 140 K and cools as it adiabatically expands. The gas should follow the adiabatic relation until it undergoes shock heating, so one expects that there should be very little entropy production until $z \sim 30$, because the first gravitationally-bound structures are just beginning to form at this epoch. Gas that reaches densities of a few times the cosmic mean is not expected to be significantly shocked; instead, it should increase its temperature only by adiabatic compression. This is true for GADGET and Enzo/PPM, where almost all of the gas maintains its initial entropy, or equivalently, it stays on its initial adiabat. At $z = 30$, only a very small amount of high-density gas departs from its initial entropy, indicating that it has undergone some shock heating. However, in the Enzo/ZEUS simulation, a much larger fraction of gas has been heated to higher temperatures. In fact, it looks as if essentially *all* overdense gas has increased its entropy by a non-negligible amount. We believe this is most likely caused by the artificial viscosity implemented in the ZEUS method, a point we will discuss further in Section 3.7.

As time progresses, virialized halos and dark matter filaments form, which are surrounded by strong accretion shocks in the gas and are filled with weaker flow shocks [199]. The distribution of gas then extends towards much higher entropies and densities. However, there is still a population of unshocked gas, which can be nicely seen as a flat




Figure 3.9: Cumulative distribution functions of gas mass as functions of comoving gas overdensity (left column), temperature (middle column) and entropy (right column) at $z = 10$. The simulations and the line types used here are the same as in Figures 3.7 and 3.8. In the bottom panels, we show the residuals in logarithmic scale with respect to the GADGET N256 run.




Figure 3.10: Cumulative distribution functions of gas mass as functions of comoving gas overdensity (left column), temperature (middle column) and entropy (right column) at $z = 3$. The simulations and the line types used here are the same as in Figures 3.7 and 3.8. In the bottom panels, we show the residuals in logarithmic scale with respect to the GADGET N256 run.




Figure 3.11: Redshift evolution of the two dimensional mass-weighted distribution of gas entropy vs. gas overdensity for four representative Enzo and GADGET simulations. Rows correspond to (from top to bottom) $z = 30$, 10 and 3 . In each panel six contours are evenly spaced from 0 to the maximum value in logarithmic scale, with the scale being identical in all simulations at a given redshift to allow for direct comparison. *Column 1*: GADGET, 2×64^3 particles, $L_{\text{box}}/e = 2048$. *Column 2*: GADGET, 2×256^3 particles, $L_{\text{box}}/e = 6400$. *Column 3*: Enzo/ZEUS hydro, 128^3 DM particles, 128^3 root grid, $L_{\text{box}}/e = 4096$. *Column 4*: Enzo/PPM hydro, 128^3 DM particles, 128^3 root grid, $L_{\text{box}}/e = 4096$. The increasing minimum entropy with decreasing overdensity in the Enzo results is an artifact of imposing a temperature floor—a numerical convenience.

constant entropy floor in all the runs until $z = 10$. However, the Enzo/ZEUS simulation largely loses this feature by $z = 3$, reflecting its poor ability to conserve entropy in unshocked regions. On the other hand, the GADGET ‘entropy conserving’ SPH-formulation preserves a very well defined entropy floor down to $z = 3$. The result of Enzo/PPM lies between that of GADGET and Enzo/ZEUS in this respect. The 1 Kelvin temperature floor in the Enzo code results in an artificial increase in the entropy “floor” in significantly underdense gas at $z = 3$.

Perhaps the most significant difference between the simulations lies however in the ‘bimodality’ that Enzo/PPM develops in the density-entropy phase space. This is already seen at redshift $z = 10$, but becomes clearer at $z = 3$. While Enzo/ZEUS and GADGET show a reservoir of gas around the initial entropy with an extended distribution towards higher density and entropy, Enzo/PPM develops a second peak at higher entropy, i.e. intermediate density and entropy values are comparatively rare. The resulting bimodal character of the distribution is also reflected in a ‘dip’ at $\log T \sim 2.0$ seen in the 1-D differential distribution function in Figures 3.7 and 3.8.

We note that the high-resolution GADGET run with 256^3 particles exhibits a broader distribution than the 64^3 run because of its much larger dynamic range and better sampling, but it does not show the bimodality seen in the Enzo/PPM run. We also find that increasing the dynamic range L_{box}/e with a fixed particle number does not change the overall shape of the distributions in a qualitative way, except that the gas extends to a slightly higher overdensity when L_{box}/e is increased.

3.6.5 Mean gas temperature and entropy

In Figure 3.12 we show the mass-weighted mean gas temperature and entropy of the entire simulation box as a function of redshift. We compare results for GADGET simulations with particle numbers of 64^3 , 128^3 and 256^3 , and Enzo runs with 64^3 or 128^3 particles for different choices of root grid size and hydrodynamic algorithm.

In the temperature evolution shown in the left panel of Figure 3.12, we see that the temperature drops until $z \sim 20$ owing to adiabatic expansion. This decline in the temperature is noticeably slower in the Enzo/ZEUS runs compared with the other simulations, reflecting the artificial heating seen in Enzo/ZEUS at early times. After $z = 20$ structure formation and its associated shock heating overcomes the adiabatic cooling and the mean temperature of the gas begins to rise quickly. While at intermediate redshifts ($z \sim 40 - 8$) some noticeable differences among the simulations exist, they tend to converge very well to a common mean temperature at late times when structure is well developed. In general, Enzo/PPM tends to have the lowest temperatures, with the GADGET SPH-results lying between those of Enzo/ZEUS and Enzo/PPM.

In the right panel of Figure 3.12, we show the evolution of the mean mass-weighted entropy, where similar trends as in the mean temperature can be observed. We see that




Figure 3.12: Mass-weighted mean gas temperature and entropy for Enzo and GADGET runs as a function of redshift. The runs used are the same as those shown in Figures 3.7 and 3.8.

a constant initial entropy ($\log S_{\text{init}} = 18.44$) is preserved until $z \sim 20$ in Enzo/PPM and GADGET. However, an unphysical early increase in mean entropy is observed in Enzo/ZEUS. The mean entropy quickly rises after $z = 20$ owing to entropy generation as a result of shocks occurring during structure formation.

Despite differences in the early evolution of the mean quantities calculated we find it encouraging that the global mean quantities of the simulations agree very well at low redshift, where temperature and entropy converge within a very narrow range. At high redshifts the majority of gas (in terms of total mass) is in regions which are collapsing but still have not been virialized, and are hence unshocked. As we show in Section 3.7, the formulations of artificial viscosity used in the GADGET code and in the Enzo implementation of the ZEUS hydro algorithm play a significant role in increasing the entropy of unshocked gas which is undergoing compression (though the magnitude of the effect is significantly less in GADGET), which explains why the simulations using these techniques have systematically higher mean temperatures/entropies at early times than those using the PPM technique. At late times these mean values are dominated by gas which has already been virialized in large halos, and the increase in temperature and entropy due to virialization overwhelms heating due to numerical effects. This suggests that most results at low redshift are probably insensitive to the differences seen here during the onset of structure formation at high redshift.

3.6.6 Evolution of kinetic energy

Different numerical codes may have different numerical errors per timestep, which can accumulate over time and results in differences in halo positions and other quantities of interest. It was seen in the Santa Barbara cluster comparison project that each code calculated the time versus redshift evolution in a slightly different way, and overall that resulted in substructures being in different positions because the codes were at different “times”. In our comparison of the halo positions in Section 3.5.2 we saw something similar – the accumulated error in the simulations results in our halos being in slightly different locations. Since we do not measure the overall integration error in our codes (which is actually quite hard to quantify in an accurate way, considering the complexity of both codes) we argue that the kinetic energy is a reasonable proxy because the kinetic energy is essentially a measure of the growth of structure - as the halos grow and the potential wells deepen the overall kinetic energy increases. If one code has errors that contribute to the timesteps being faster/slower than the other code this shows up as slight differences in the total kinetic energy.

In Figure 3.13 we show the kinetic energy (hereafter KE) of dark matter and gas in GADGET and Enzo runs as a function of redshift. As expected, KE increases with decreasing redshift. In the bottom panels, the residuals with respect to the GADGET 256³ particle run is shown in logarithmic units (i.e., $\log(\text{KE}_{\text{others}}) - \log(\text{KE}_{256})$). Initially at $z = 99$, GADGET and Enzo runs agree to within a fraction of a percent within their own runs with different particle numbers. The corresponding GADGET and Enzo runs with the same particle/mesh number agree within a few percent. These differences may have been caused by the numerical errors during the conversion of the initial conditions and the calculation of the KE itself. It is reasonable that the runs with a larger particle number result in a larger KE at both early and late times, because the larger particle number run can sample the power spectrum to a higher wavenumber, therefore having more small-scale power at early times and more small-scale structures at late times. The 64³ runs both agree with each other at $z = 99$, and overall have about 1% less kinetic energy than the 256³ run. At the same resolution, Enzo runs show up to a few percent less energy at late times than GADGET runs, but their temporal evolutions track each other closely.

3.6.7 The gas fraction in halos

The content of gas inside the virial radius of dark matter halos is of fundamental interest for galaxy formation. Given that the Santa Barbara cluster comparison project hinted that there may be a systematic difference between Eulerian codes (including AMR) and SPH codes (Enzo gave slightly higher gas mass fraction compared to SPH runs at the virial radius), we study this property in our set of simulations.

In order to define the gas content of halos in our simulations we first identify dark




Figure 3.13: Kinetic energy of dark matter and gas as a function of redshift, and the residuals in logarithmic units with respect to the 256^3 particle Gadget run (red long-dashed line) is shown in the bottom panels. Red short-dashed line is for GADGET 128^3 particle run, and red solid line is for GADGET 64^3 particle run. Black lines are for Enzo runs: 128g128dm PPM, lowod (dot-short dash), 128g128dm Zeus (dotted), 64g64dm PPM, lowod (short dash-long dash).

matter halos using a standard friends-of-friends algorithm. We then determine the halo center to be the center of mass of the dark matter halo and compute the “virial radius” for each halo using Equation (24) of Barkana & Loeb [200] with the halo mass given by the friends-of-friends algorithm. This definition is independent of the gas distribution, thereby freeing us from ambiguities that are otherwise introduced owing to the different representations of the gas in the different codes on a mesh or with particles. Next, we measure the gas mass within the virial radius of each halo. For GADGET, we can simply count the SPH particles within the radius. In Enzo, we include all cells whose centers are within the virial radius of each halo. Note that small inaccuracies can arise because some cells may only partially overlap with the virial radius. However, in significantly overdense regions the cell sizes are typically much smaller than the virial radius, so this effect should not be significant for large halos.

In Figure 3.14 we show the gas mass fractions obtained in this manner as a function of total mass of the halos, with the values normalized by the universal mass fraction $f_{\text{gas}} \equiv (M_{\text{gas}}/M_{\text{tot}})/(\Omega_b/\Omega_m)$. The top three panels show results obtained with GADGET for 2×64^4 , 2×128^3 , and 2×256^3 particles, respectively. The bottom 9 panels show Enzo results with 64^3 and 128^3 root grids. Simulations shown in the right column use the ZEUS hydro algorithm and the others use the PPM algorithm. All Enzo runs shown have 64^3 dark matter particles, except for the bottom row which uses 128^3 particles. The Enzo simulations in the top row use a 64^3 root grid and all others use a 128^3 root grid. Grid and particle sizes, overdensity threshold for refinement and hydro method are noted in each panel.

For well-resolved massive halos, the gas mass fraction reaches $\sim 90\%$ of the universal baryon fraction in the GADGET runs, and $\sim 100\%$ in all of the Enzo runs. There is a hint that the Enzo runs seem to give values a bit higher than the universal fraction, particularly for runs using the ZEUS hydro algorithm. This behavior is consistent with the findings of the Santa Barbara comparison project. Given the small size of our sample, it is unclear whether this difference is really significant. However, there is a clear systematic difference in baryon mass fraction between Enzo and GADGET simulations. Examining the mass fraction of simulations to successively larger radii show that the Enzo simulations are consistently close to a baryon mass fraction of unity out to several virial radii, and the gas mass fractions for GADGET runs approaches unity at radii larger than twice the virial radius of a given halo.

The systematic difference between Enzo and GADGET calculations, even for large masses, is also somewhat reflected in the results of Kravtsov et al. [201]. They perform simulations of galaxy clusters done using adiabatic gas and dark matter dynamics with their adaptive mesh code and GADGET. At $z = 0$ their results for the baryon fraction of gas within the virial radius converge to within a few percent between the two codes, with the overall gas fraction being slightly less than unity. It is interesting to note that they also observe that the AMR code has a higher overall baryon mass fraction than




Figure 3.14: Gas mass fraction normalized by the universal baryon mass fraction $f_{\text{gas}} = (M_{\text{gas}}/M_{\text{tot}})/(\Omega_b/\Omega_m)$ is shown. The top 3 panels are for GADGET runs with 2×64^3 , ($L_{\text{box}}/e = 2048$), 2×128^3 ($L_{\text{box}}/e = 3200$), and 2×256^3 ($L_{\text{box}}/e = 6400$) particles. The bottom panels are for the Enzo runs with 64^3 or 128^3 grid, and 64^3 or 128^3 dark matter particles. The ZEUS hydrodynamics method is used for one set of the Enzo simulations (right column) and the PPM method is used for the rest. All Enzo runs have $L_{\text{box}}/e = 4096$.

GADGET, though still slightly less than what we observe with our Enzo results.

Note that the scatter of the baryon fraction seen for halos at the low mass end is a resolution effect. This can be seen when comparing the three panels with the GADGET results. As the mass resolution is improved, the down-turn in the baryon fraction shifts towards lower mass halos, and the range of halo masses where values near the universal baryon fraction are reached becomes broader. The sharp cutoff in the distribution of the points corresponds to the mass of a halo with 32 DM particles.

It is also interesting to compare the cumulative mass function of gas mass in halos, which we show in Figure 3.15 for adiabatic runs. This can be viewed as a combination of a measurement of the DM halo mass function and the baryon mass fractions. In the lower panel, the residuals in logarithmic scale are shown for each run with respect to the Sheth & Tormen [198] mass function (i.e., $\log(N[>M]) - \log(\text{S\&T})$).

As with the dark matter halo mass function, the gas mass functions agree well at the high-mass end over more than a decade of mass, but there is a systematic discrepancy between AMR and SPH runs at the low-mass end of the distribution. While the three SPH runs with different gravitational softening agree well with the expectation based on the Sheth & Tormen mass function and an assumed universal baryon fraction at $M_{\text{gas}} < 10^8 h^{-1} M_{\odot}$, the Enzo run with 64^3 root grid and 64^3 DM particles has fewer halos. Similarly, the Enzo run with 128^3 grid and 128^3 DM particles has fewer low mass halos at $M_{\text{gas}} < 10^7 h^{-1} M_{\odot}$ compared to the GADGET 128^3 DM particle run. Convergence with the SPH results for Enzo requires the use of a root grid with spatial resolution twice that of the initial mean interparticle separation, as well as a low-overdensity refinement criterion. We also see that the PPM method results in a better gas mass function than the ZEUS hydro method at the low-mass end for the same number of particles and root grid size.

3.7 The role of artificial viscosity

In Section 3.6.4 we found that slightly overdense gas in Enzo/ZEUS simulations shows an early departure from the adiabatic relation towards higher temperature, suggesting an unphysical entropy injection. In this section we investigate to what extent this effect can be understood as a result of the numerical viscosity built into the ZEUS hydrodynamic algorithm. As the gas in the pre-shocked universe begins to fall into potential wells, this artificial viscosity causes the gas to be heated up in proportion to its compression, potentially causing a significant departure from the adiabat even when the shock has not occurred yet; i.e. when the compression is only adiabatic.

This effect is demonstrated in Figure 3.16, where we compare two-dimensional entropy-overdensity phase space diagrams for two Enzo/ZEUS where the strength of the artificial viscosity was reduced from its “standard” value of $Q_{\text{AV}} = 2.0$ to $Q_{\text{AV}} = 0.5$. These runs




Figure 3.15: Cumulative halo gas mass function at $z = 3$. For reference, the solid black line is the Sheth & Tormen [198] mass function multiplied by the universal baryon mass fraction Ω_b / Ω_m . In the bottom panel, the residuals in logarithmic scale with respect to the Sheth & Tormen mass function are shown for each run (i.e., $\log(N[>M]) - \log(\text{S\&T})$).




Figure 3.16: Two-dimensional distribution functions of gas entropy vs. gas overdensity for two Enzo runs performed with the ZEUS hydrodynamics algorithm, varying with redshift. Rows correspond to (top to bottom) $z = 30$, 10 and 3. In each panel, six contours are evenly spaced from 0 to the maximum value in equal logarithmic scale. Two different values of the ZEUS artificial viscosity parameter are used: $Q_{\text{AV}} = 0.5$ (left column) and $Q_{\text{AV}} = 2.0$ (right column). Both runs use 64^3 dark matter particles and a 64^3 root grid and have a maximum spatial resolution of $L_{\text{box}}/e = 4096$. The standard value of the artificial viscosity parameter is $Q_{\text{AV}} = 2.0$.

used 64^3 dark matter particles and 64^3 root grid, and the $Q_{AV} = 2.0$ corresponds to the case shown earlier in Figure 3.11.

Comparison of the Enzo/ZEUS runs with $Q_{AV} = 0.5$ and 2.0 shows that decreasing Q_{AV} results in a systematic decrease of the unphysical gas heating at high redshifts. Also, at $z = 4$ the $Q_{AV} = 0.5$ result shows a secondary peak at higher density, so that the distribution becomes somewhat more similar to the PPM result. Unfortunately, a strong reduction of the artificial viscosity in the ZEUS algorithm is numerically dangerous because the discontinuities that can appear owing to the finite-difference method are then no longer smoothed sufficiently by the artificial viscosity algorithm, which can produce unstable or incorrect results.

An artificial viscosity is needed to capture shocks when they occur in both the Enzo/ZEUS and GADGET SPH scheme. This in itself is not really problematic, provided the artificial viscosity is very small or equal to zero in regions without shocks. In this respect, GADGET's artificial viscosity behaves differently from that of Enzo/ZEUS. It takes the form of a pairwise repulsive force that is non-zero only when Lagrangian fluid elements approach each other in physical space. In addition, the strength of the force depends in a non-linear fashion on the *rate of compression* of the fluid. While even an adiabatic compression produces some small amount of (artificial) entropy, only a compression that proceeds *rapidly* with respect to the sound-speed, as in a shock, produces entropy in large amounts. This can be seen explicitly when we analyze equations (3.6) and (3.8) for the case of a homogeneous gas which is uniformly compressed. For definiteness, let us consider a situation where all separations shrink at a rate $q = \dot{r}_{ij}/r_{ij} < 0$, with $\nabla \cdot \mathbf{v} = 3q$. It is then easy to show that the artificial viscosity in GADGET produces entropy at a rate

$$\frac{d \log A_i}{d \log \rho_i} = \frac{\gamma - 1}{2} \alpha \left[\frac{-q h_i}{c_i} + 2 \left(\frac{q h_i}{c_i} \right)^2 \right]. \quad (3.13)$$

Note that since we assumed a uniform gas, we here have $h_i = h_{ij}$, $c_i = c_{ij}$, and $\rho_i = \rho_{ij}$. We see that only if the compression is fast compared to the sound-crossing time across the typical spacing of SPH particles, i.e. for $|q| > c_i/h_i$, a significant amount of entropy is produced, while slow (and hence adiabatic) compressions proceed essentially in an isentropic fashion. On the other hand, the artificial viscosity implemented in Enzo/ZEUS produces entropy irrespective of the sound-speed, depending only on the compression factor of the gas.

We have also investigated the pre-shock entropy generation in Enzo/ZEUS using another simple test, the collapse of a one-dimensional Zel'dovich pancake. The initial conditions of this test are simple and described in full detail by Bryan et al. [145] A one-dimensional simulation volume is set up in an expanding coordinate system in a flat cosmology with an initially sinusoidal density perturbation with a peak at $x = 0.5$ and a corresponding perturbation in the velocity field with nodes at $x = 0.0, 0.5, \text{ and } 1.0$.




Figure 3.17: Final results of the Zel'dovich pancake test. *Panel (a)*: Log baryon overdensity vs. position. *Panel (b)*: Log entropy vs. position. *Panel (c)*: Cumulative mass-weighted entropy distribution function. *Panel (d)*: Pre-shock entropy evolution in the 256-cell Zeus run. All Zel'dovich pancake simulations are performed in one dimension using the Enzo code. *Panels (a), (b), (c)*: The line types are for PPM/1024 (short-dashed), PPM/256 (solid), ZEUS/1024 (dotted), and ZEUS/256 (long dashed) where 256 and 1024 are the number of cells used in the calculation. All data is at the end of the run ($z = 0$). *Panel (d)*: Entropy evolution of the 256-cell ZEUS and PPM runs for redshifts $z = 20$ (solid line), 10 (dotted line), 5 (long-dashed line), 2.5 (dot-dashed line) and 2 (dot-long-dashed line). All PPM results overlay the ZEUS initial conditions ($z = 20$). Note that the x-axis range for panel (d) is different from that of panels (a) and (b).

A temperature perturbation is added such that gas entropy is constant throughout the volume.

In Figure 3.17 we show the density and entropy profiles as a function of position, at a time when the non-linear collapse of the pancake is well underway. We also show the pre-shock evolution of the entropy profile for both algorithms. We compare runs using 256 and 1024 grid cells with both the ZEUS and PPM formulations of Enzo.

As the matter falls in onto the density peak at $x = 0.5$, accretion shocks on either side form, clearly marked by the jumps in density, entropy, and temperature. Note that the dip in the temperature at $x = 0.5$ is physical – the gas sitting there is unshocked and only adiabatically compressed, and therefore has relatively low temperature. Reassuringly, both the ZEUS and PPM hydrodynamical methods reproduce the qualitative behavior of the Zel’dovich pancake quite well, but there are also some systematic differences at a given resolution. This can be seen most clearly in the mass-weighted cumulative entropy distribution in the bottom left panel of Figure 3.17. We see that the Enzo/ZEUS calculations show a broader distribution than Enzo/PPM for a given spatial resolution. This can be interpreted as another sign of pre-shock entropy generation by the artificial viscosity in ZEUS. In contrast, the Riemann solver used in PPM can capture shocks such that they are resolved as true discontinuities, which avoids this problem.

More concrete evidence of spurious entropy generation in the artificial viscosity-based scheme can be seen by examining the pre-shock evolution of entropy in these simulations (as seen in panel (d) of Figure 3.17). No entropy should be generated before the twin shocks form to the left and right of $x = 0.5$ (as can be seen in panel (b) of the same figure). The simulations using PPM (black solid line in panel d) produce no spurious entropy. The simulations using the ZEUS scheme, however, produce significant amounts of entropy in the infalling (but unshocked) gas. Note that the magnitude of the entropy generation is relatively small compared to the final entropy produced in the shocks (as seen in panel (b)), but the values are still significant.

While this test showed only comparatively small differences between the different methods, it is plausible that the effects of pre-shock entropy generation become much more important in three-dimensional cosmological simulations, where galaxies form hierarchically through complicated merger processes that involve extremely complex shock patterns. We thus speculate that this effect may be the key reason for the systematic differences between the Enzo/PPM runs and the ZEUS and GADGET simulations.

3.8 Timing & memory usage

An important practical consideration when assessing the relative performance of computational methods or simulation codes is the amount of computational resources they require to solve a given problem. Of primary importance are the total amount of memory

and the CPU time that is needed. However, it is not always easy to arrive at a meaningful comparison, particularly for very different methods such as AMR and SPH. For example, the variable number of grid cells owing to the adaptive nature of AMR is an important complication, making the number of resolution elements change over time, while the particle number stays constant in the SPH method. An additional layer of complexity is added when considering parallel codes. The parallelization strategies that are used for AMR applications can be significantly different than those used in SPH codes, and the performance of an individual simulation code can heavily depend on the specific computer architecture and implementation of MPI (or other software used for parallelization) chosen. Therefore we caution the readers to take all of the timing information discussed in this section as results for a particular problem setup and machine architecture, and not to extrapolate directly to different types of cosmological simulations (e.g., with cooling and star formation) and machines.

3.8.1 Initial comparison on a distributed memory machine

When we started this project, we initially performed our comparison runs on the IA-64 Linux cluster Titan at the National Center for Supercomputing Applications (NCSA). It had 134 dual processor nodes with 800 MHz Intel Itanium 1 chips, 2.5 GB memory per node, and Myrinet 2000 network interconnect. Our initial comparison on Titan showed that the GADGET code was faster than Enzo by a factor of 40 (15) for a 64^3 (128^3) particle DM-only run when Enzo was using a low overdensity criteria for grid refinement. The low overdensity refinement criterion was required for Enzo in order to obtain a DM halo mass function comparable to that of GADGET at low-mass end. GADGET used a factor of 18 (4) less amount of memory than Enzo for a 64^3 (128^3) particle DM-only run. For the adiabatic runs, GADGET was faster than Enzo by a factor of 2.5 for a 64^3 DM particles and 64^3 gas particles (a 64^3 root grid for Enzo). A GADGET run with 128^3 dark matter and gas particles completed 8 times faster than an Enzo simulation with a 128^3 root grid and 64^3 DM particles. These performance results were gathered using Linux-based Beowulf-style clusters with relatively slow inter-node communication networks. Since the AMR code performs load balancing by passing grids between processors, it was expected that the performance of Enzo would improve on a large shared-memory machine. The disparity is most significant for DM-only simulations, so improvement of the Enzo N-body solver could significantly increase the performance of the AMR code.

3.8.2 More recent comparison on a shared memory machine

During the course of this comparison study, both GADGET and Enzo evolved, and the performance of both codes have greatly improved. Therefore, we repeated the performance comparison with our updated codes using the IBM DataStar machine at the San

Diego Supercomputing Center.³ The portion of the machine used for these timing tests is composed of 176 IBM p655 compute nodes, each of which has eight 1.5 GHz IBM Power4 processors. These processors are super-scalar, pipelined 64 bit chips which can execute up to 8 instructions per clock cycle and up to four floating point operations per clock cycle, with a theoretical peak performance of 6.0 GFlop per chip. Processors in a single node share a total of 16 GB of memory. All nodes are connected by an IBM Federation switch, which provides processor-to-processor bandwidth of approximately 1.4 GB/s with 8 microsecond latency when using IBM’s MPI library. Each node is directly connected to a parallel filesystem through a Fibre Channel link.

We first compare the series of dark matter-only runs discussed in Section 3.5. A GADGET simulation with 64^3 dark matter particles takes total wall-clock time of 225 seconds on 8 cpus (total 1800 seconds CPU time) and requires 270 MB of memory. 24% of the total computational time was spent doing interprocessor message-passing. The corresponding Enzo simulation with 64^3 particles and a 64^3 root grid requires 1053 seconds on 8 cpus (total 8424 seconds CPU time) when refining on a dark matter overdensity of 2.0, and requires 1.21 GB of memory total. 34% of the total computational time was spent in interprocessor communication. This is a factor of 4.7 slower than the corresponding GADGET simulation, and requires roughly 4.5 times more memory. Raising the refinement criteria to a dark matter overdensity of 4.0 (at a cost of losing low-mass DM halos) reduces the wall clock time to 261 seconds on 8 processors (total 2088 seconds CPU time) and decreases the total amount of memory needed to 540 MB, which is comparable to the GADGET simulation. A 128^3 DM particle GADGET adiabatic run takes a total of 2871 seconds to run on 8 cpus (total 22,968 seconds CPU time) and requires 1.73 GB of memory. An Enzo simulation with 128^3 particles and a 128^3 root grid that refines on a dark matter overdensity of 2.0 needs approximately 34,028 seconds on 8 processors (total 272,224 CPU seconds) and 5.6 GB of memory. This is a factor of 12 slower and 3.2 times more memory than the equivalent GADGET run. The same calculation run with refinement overdensities of 4.0 or 8.0 completes in 13,960 and 3839 seconds, respectively, which are factors of 4.9 and 1.3 slower than the equivalent GADGET run. The reason for the huge change in computational speeds is due to the low overdensity threshold used in the first simulation, which results in a huge number of grids to be instantiated and a great deal of time to be spent regriding the simulation. Raising the overdensity criteria suppresses the formation of halos at the low mass end of the mass function, though higher-mass halos are unaffected. This timing comparison suggests that if one is interested in simulating the full spectrum of dark matter halos at a reasonable computational cost, GADGET would be a wiser choice than Enzo for this application. If one was interested in only the high-mass end of the mass function, the codes have comparable performance.

Comparison of the adiabatic gas + N-body cosmological simulations in Section 3.6

³http://www.sdsc.edu/user_services/datastar/

is also quite informative. The 64^3 dark matter particle/ 64^3 gas particle GADGET calculation takes 1839 seconds to run on 8 processors (total 14,712 seconds CPU time) and requires 511 MB of memory. The equivalent Enzo simulation with 64^3 particles and a 64^3 root grid using the low overdensity refinement criteria (refining on a baryon overdensity of 4.0 and a dark matter overdensity of 2.0) requires 6895 seconds on 8 processors (55,160 seconds total) and 2.5 GB of memory. This is 3.7 times slower and 4.9 times more memory than the corresponding GADGET run. Raising the overdensity thresholds by a factor of two decreases the computational time to 2168 seconds on 8 processors and the memory required to 1.28 GB. The GADGET calculation with 128^3 dark matter and baryon particles requires 35,879 seconds on 8 cpus (287032 seconds total CPU time) and 5.4 GB of memory, and an Enzo calculation with 128^3 particles on a 128^3 root grid which refines on a baryon overdensity of 8.0 and a dark matter overdensity of 4.0 requires 64,812 seconds and 8 GB of memory. Enzo simulations using the PPM and Zeus hydro methods require comparable amounts of simulation time.

3.8.3 Mass resolution in Enzo and GADGET

It is clear from Sections 3.8.1 and 3.8.2 that Enzo, at present, is significantly slower and requires more memory than GADGET when one demands convergence on all scales for a simulation of a given size. If one requires convergence only at the high-mass end of the mass function the relative performance of the two codes becomes much more comparable. However, it is unclear that raw computational time and memory is a fair assessment of the performance of the two codes. As discussed previously, the number of dark matter and gas particles in the GADGET simulations remain constant always. In the AMR simulations, the number of dark matter particles is fixed, but the adaptive nature of the code adds more cells in areas of high overdensity, so that the number of root grid cells (N_{grid} in Table 3.2) is a lower bound for the total number N_{rez} of cells used to solve the hydrodynamics in an Enzo simulation, which becomes typically larger than the number of root grid cells by a factor of at least a few once structure has developed.

Note in this context that the refinement criterion presently used in Enzo tries to roughly keep the baryonic mass per cell constant, which is in principle similar to the Lagrangian behavior of GADGET, where a constant mass resolution is imprinted by construction. This is seen more clearly in Figure 3.18, where the mean gas mass in cells in Enzo simulations is shown as a function of gas overdensity. The AMR simulations show a nearly flat behavior for a broad range of densities, i.e. by and large they distribute their resolution elements similarly as a function of density, except at very high and low density. At low densities, the baryonic mass resolution tends to become better compared with a purely Lagrangian code, the prime reason being that the root grid never becomes coarser. In contrast, the mass resolution tends to become worse for high densities, because the imposed limit on the maximum level of refinements prevents the placing of additional

refinements. Note however that one conceptual strength of AMR is that the refinement criteria are flexible, and that they do not necessarily have to mimic Lagrangian behavior as has been the case here. The SPH particle masses are shown as horizontal short-dashed lines in this figure. It can be clearly seen that the mean gas mass resolved in a single cell in the Enzo calculations is roughly a factor of 8 better than the equivalent GADGET calculation (i.e. when comparing simulations with the same root grid size as number of particles).

Another way of looking at this is to examine the distribution of baryon masses in cells in the Enzo calculations. Figure 3.19 shows a plot of the number of cells (per mass bin) as a function of cell baryon mass. The (fixed) baryon particle masses for the GADGET calculations are shown as vertical arrows. This plot shows that, for simulations with the same root grid size/number of baryon particles, the median cell mass in the Enzo calculations are approximately an order of magnitude less than the corresponding GADGET runs, and that the “low overdensity” Enzo simulations have a median mass resolution that is roughly a factor of 2 – 3 better than the corresponding “high overdensity” calculation. This is unsurprising, considering that a low overdensity threshold for refinement directly translates into a lower cell mass. Given that the mean cell mass stays relatively constant as a function of overdensity (as shown in Figure 3.18), this implies that there is a large amount of scatter in the baryon mass contained in cells at a given spatial resolution or overdensity.

Another interesting comparison that can be made between the Enzo and GADGET calculations concerns the total number of unique resolution elements and their distribution as a function of overdensity (which is directly related to spatial resolution). Figure 3.20 shows the distribution of unique resolution elements as a function of overdensity for both Enzo and GADGET simulations. The curves shown in this figure are not normalized, meaning that integrating the area under the curve gives the total number of resolution elements. For a given simulation size, the Enzo calculations have at least a factor of three more resolution elements overall than the corresponding GADGET simulation, with the “low overdensity” calculations having roughly an order of magnitude more unique resolution elements overall. The overall distributions of resolution elements are comparable between the two different codes, though there is a peak at low overdensities in the Enzo simulations, which is due to the root grid, which has a fixed (and large) number of grid cells. In terms of overall resolution of baryons as a function of overdensity, the Enzo 64^3 “low overdensity” simulation has a comparable number of resolution elements to the 128^3 GADGET run at overdensities greater than one, and the 128^3 Enzo low overdensity run is comparable to the 256^3 GADGET calculation, with the caveat in both cases that the GADGET calculations resolve to slightly higher overdensities overall. There are very few resolution elements at the highest densities, so the significance of this is unclear.

Though it seems that a given Enzo calculation has much better overall baryon mass resolution than the equivalent GADGET simulation, the significance of this is unknown.




Figure 3.18: Mean cell mass as a function of overdensity for four representative Enzo calculations. All simulations use the PPM hydro algorithm – calculations done with the ZEUS algorithm are essentially the same as the corresponding PPM simulation. We show results for the 64^3 and 128^3 root grid Enzo calculations (black and red lines, respectively), and use only simulations that have the same number of dark matter particles as root grid cells. Results for both high and low-overdensity calculations are shown. The baryon particle mass for the equivalent GADGET simulations are shown as horizontal black dashed lines corresponding to (from top to bottom) the 64^3 , 128^3 and 256^3 particle GADGET simulations.




Figure 3.19: Number of cells as a function of cell baryon mass for four representative Enzo simulations. All simulations use the PPM hydro algorithm – calculations done with the ZEUS algorithm are essentially the same as the corresponding PPM simulation. We show results for the 64^3 and 128^3 root grid Enzo calculations (black and red lines, respectively), and use only simulations that have the same number of dark matter particles as root grid cells. Results for both high and low-overdensity calculations are shown. The baryon particle mass for the equivalent GADGET simulations are shown as vertical arrows corresponding to (left to right) the 256^3 , 128^3 and 64^3 particle GADGET simulations.




Figure 3.20: Number of resolution elements as a function of baryon overdensity for representative Enzo and GADGET calculations. All Enzo simulations use the PPM hydro algorithm – calculations done with the ZEUS algorithm are essentially the same as the corresponding PPM simulation. We show results for the 64^3 and 128^3 root grid Enzo calculations (blue and red lines, respectively), and use only simulations that have the same number of dark matter particles as root grid cells. Results for both high and low-overdensity calculations are shown. Resolution element distributions for the GADGET simulations with 64^3 , 128^3 and 256^3 particles are shown as black lines (with the total number of particles increasing from bottom to top).

Given that the dark matter particle masses in the Enzo calculations are fixed (and the same as in the corresponding GADGET simulation), the difference in baryon resolution will have essentially no effect on the large scale structure in the simulation. However, within a given halo better baryon mass resolution implies that shocks and the internal dynamics of the halo are more highly resolved, which may become important in simulations with more complicated physics, such as radiative cooling and star formation and feedback.

3.9 Discussion and conclusions

This chapter presents initial results of a comparison of two state-of-the-art cosmological hydrodynamic codes: Enzo, an Eulerian adaptive mesh refinement code, and GADGET, a Lagrangian smoothed particle hydrodynamics code. These codes differ substantially in the way they compute gravitational forces and even more radically in the way they treat gas dynamics. In cosmological applications structure formation is driven primarily by gravity, so a comparison of the hydrodynamical methods necessarily involves an implicit comparison of the gravitational solvers as well. In order to at least partially disentangle these two aspects we have performed both a series of dark matter-only simulations and a set of simulations that followed both a dark matter and an adiabatic gaseous component.

Our comparison of the dark matter results showed good agreement in general provided we chose a root grid resolution in Enzo at least twice that of the mean interparticle separation of dark matter particles together with a relatively conservative AMR refinement criterion of dark matter overdensity of 2. If less stringent settings are adopted, the AMR code shows a significant deficit of low mass halos. This behavior can be readily understood as a consequence of the hierarchical particle-mesh algorithm used by Enzo for computing gravitational forces, which softens forces on the scale of the mesh size. Sufficiently small mesh cells are hence required to compete with the high force-resolution tree-algorithm of GADGET. In general, we find excellent agreement with the results of Heitmann et al.[202], particularly with regards to systematic differences in the power spectrum and low-mass end of the halo mass function between mesh and tree codes. Our results are complementary in several ways – Heitmann et al. use simulations run with the “standard” parameters for many codes (using the same initial conditions) and then compare results without any attempt to improve the quality of agreement, whereas we examine only two codes, but systematically vary parameters in order to understand how the codes can be made to agree to very high precision.

Examination of the dark matter substructure in the two most massive halos in our 128^3 particle dark matter-only calculations shows that while both codes appear to resolve substructure (and obtain substructure mass functions that are comparable) there are some differences in the number and the spatial distribution of subhalos between the

two codes. While the origin of these differences are not fully clear, it may be due to a lack of spatial (i.e. force) or dark matter mass resolution, or possible due in part to systematics in the grouping algorithm used to detect substructure. The observed differences in substructure are not surprising when one considers how dissimilar the algorithms that Enzo and GADGET use to calculate gravitational accelerations on small scales are, and a further study with much higher resolution is necessary.

We also found broad agreement in most gas quantities we examined in simulations which include adiabatic gas evolution, but there were also some interesting discrepancies between the different codes and different hydrodynamical methods. While the distributions of temperature, density, and entropy of the gas evolved qualitatively similarly over time, and reassuringly converged to the same mean temperature and entropy values at late times, there were clearly some noticeable differences in the early evolution of the gas and in the properties of intermediate density gas.

In particular, in the Enzo/ZEUS simulations we found an early heating of collapsing or compressed gas, caused by injection of entropy by the artificial viscosity in this code. This resulted in substantial pre-shock entropy generation in the Enzo/ZEUS runs. While GADGET also uses an artificial viscosity to capture shocks, effects of pre-shock entropy generation are substantially weaker in this code. This reflects its different parameterization of artificial viscosity, which better targets the entropy production to shocked regions.

Considering the entropy-density distribution in more detail, we found that Enzo/PPM calculations show a marked trend towards a segregation of gas into a low-entropy reservoir of unshocked low density gas and a pool of gas that has been shocked and accumulated high entropy when it reached higher density regions. Such a bimodality is not apparent in the Enzo/ZEUS and GADGET runs at $z = 3$. Instead, there is a smoother transition from low- to high-entropy material; i.e. more gas of intermediate entropy exists. It is possible that this intermediate-entropy gas is produced by the artificial viscosity in pre-shock regions, where entropy generation should not yet take place. Some supporting evidence for this interpretation is provided by the fact that the distributions of temperature and entropy of Enzo/ZEUS become somewhat more similar to those of Enzo/PPM when we reduce the strength of the artificial viscosity.

Perhaps the most interesting difference we found between the two methods lies in the baryon fraction inside the virial radius of the halos at $z = 3$. For well-resolved halos Enzo results asymptote to slightly higher than 100% of the cosmic baryon fraction, independent of the resolution and hydro method used (though note that the results using the ZEUS method appear to converge to a marginally higher value than the PPM results). This also shows up as an overestimate of gas mass function $M_{gas} > 10^8 h^{-1} M_{\odot}$ compared to the Sheth & Tormen function multiplied by (Ω_b/Ω_M) . In contrast, GADGET halos at all resolutions only reach $\sim 90\%$ of the cosmic baryon fraction. This result is not easily understood in terms of effects due to artificial viscosity since the ZEUS method used in

Enzo produces more artificial viscosity than either of the other methods, yet the results for the two hydro methods in Enzo agree quite well. The systematic difference between Enzo and GADGET results in this regime provides an interesting comparison to Kravtsov et al. [201], who examine the enclosed gas mass fraction at $z = 0$ as a function of radius of eight galaxy clusters in adiabatic gas simulations done with the ART and GADGET codes. They see that at small radii there are significant differences in enclosed gas mass fraction, but at distances comparable to the virial radius of the cluster the mass fractions converge to within a few percent and are overall approximately 95% of the universal mass fraction. It is interesting to note that the enclosed gas mass fraction at the virial radius produced by the ART code is higher than that of GADGET by a few percent, and the ART gas mass fraction result would be bracketed by the Enzo and GADGET results, overall. This suggests that it is not clear that a universal baryon fraction of $\sim 100\%$ is predicted by AMR codes, though there seems to be a clear trend of AMR codes having higher overall baryon mass fractions in halos than SPH codes to, which agrees with the results of Frenk et al. [172]

It is unclear why our results with the GADGET code differ from those seen in Kravtsov et al. (with the net gas fraction in our calculations being approximately 5% lower at the virial radius), though it may be due entirely to the difference in regime – we are examining galaxy-sized halos with masses of $\sim 10^9 - 10^{10} M_\odot$ at $z = 3$, whereas they model $\sim 10^{13} - 10^{14} M_\odot$ galaxy clusters at $z = 0$. Regardless, the observed differences between the codes are significant and will be examined in more detail in future work.

It should be noted that the hydrodynamic results obtained for the GADGET SPH code are typically found to be bracketed by the two different hydrodynamic formulations implemented in the AMR code. This suggests that there is no principle systematic difference between the techniques which would cause widely differing results. Instead, the systematic uncertainties within each technique, for example with respect to the choice of shock-capturing algorithm, appear to be larger than the intrinsic differences between SPH and AMR for the quantities of interest in this paper. We also note that some of the differences we find in bulk simulation properties are likely to be of little relevance for actual simulations of galaxy formation. For example, in simulations including more realistic physics, specifically a UV background, the low temperature gas that is affected most strongly by artificial early heating in Enzo/ZEUS will be photoionized and thus heated uniformly to approximately 10^4 K, so that many of the differences in temperature and entropy at low overdensity owing to the choice of hydrodynamical method will disappear. We will investigate such effects of additional physics in the future.

We have also examined the relative computational performance of the codes studied here, using metrics such as the total CPU time and memory consumption. If one simply compares simulations which have the same number of particles and grid cells at the start of the simulation, GADGET performs better; i.e. it finishes faster, uses less memory, and is more accurate at the low-mass end of the halo mass function. However, much of this

difference is caused by the slowly increasing number of cells used by the AMR code to represent the gas, while the Lagrangian code keeps the number of SPH particles constant. If the consumed resources are normalized to the number of resolution elements used to represent the gas (cells or particles), they are roughly comparable. Unfortunately, the lower gravitational force-resolution of the hierarchical particle-mesh algorithm of Enzo will usually require the use of twice as many root grid cells as particles per dimension for high-accuracy results at the low-mass end of the mass function, which then induces an additional boost of the number of needed cells by nearly an order of magnitude with a corresponding impact on the required computational resources. As a consequence of this, the gas will be represented more accurately, and this is hence not necessarily a wasted effort. However given that the dark matter mass resolution is not also improved at the same time (unless the DM particle number is also increased), it is probably of little help to make progress in the galaxy formation problem, where the self-gravity of dark matter is of fundamental importance. It is also true that the relative performance of the codes is dependent upon the memory architecture and interprocessor communication network of the computer used to perform the comparison as we discussed in Section 3.8.

It is encouraging that, with enough computational effort, it is possible to achieve the same results using both the Enzo and GADGET codes. In principle both codes are equally well-suited to performing dark matter-only calculations (in terms of their ability to obtain results that both match analytical estimates and also agree with output from the other code), but practically speaking the slower speed of the AMR code makes it undesirable as a tool for doing large, high-resolution N-body calculations at the present day. It should be noted that solving Poisson's equation on an adaptive mesh grid is a relatively new technique, particularly compared to doing N-body calculations using tree and PM codes, and much can be done to speed up the Enzo Poisson solver and decrease its memory consumption. The GADGET N-body solver is already very highly optimized. If the speed of the Enzo N-body solver can be increased by a factor of a few, an improvement which is quite reasonable to expect in the near future, the overall speed that the codes require to achieve solutions with similar dark matter force resolutions and mass functions will be comparable.

In future work it will be important to understand the origin of the small but finite differences between Enzo/ZEUS, Enzo/PPM, and SPH at a more fundamental level. These differences will most likely be seen (and the reasons for the differences identified) when making direct comparisons of the formation and evolution of individual dark matter halos and the gas within them. Additionally, isolated idealized cases such as the Bertschinger adiabatic infall solution [203] will provide useful tests to isolate numerical issues. Examination of individual halos may also point the way to improved parameterizations of artificial viscosity (and/or diffusivity) which would then also be beneficial for the SPH method. Simultaneously, we plan to investigate the differences of the current generation of codes when additional physical effects are modeled.

Chapter 4

The Formation of Population III stars in a Λ CDM universe

4.1 Summary

In this chapter I discuss aspects of primordial star formation in a Λ CDM universe. The collapse of gas in a representative halo which will contain a single Population III protostar is described in detail in Section 4.4.1, emphasizing the critical role that the chemistry and cooling properties of molecular hydrogen gas (H_2) play. Our results are both qualitatively and quantitatively similar to those discussed in Abel, Bryan & Norman [39] despite being in a different cosmology. Section 4.4.2 describes the role of angular momentum in the formation of primordial protostars for the same halo, showing that the gas which collects in the halo core preferentially has low angular momentum compared to gas that does not collapse into the center of the halo. I analyze angular momentum transport in the halo using standard disk methods as well as by using Lagrangian “tracer particles” and show that some angular momentum transport also appears to take place during the collapse of the halo through turbulent transport. The disk approximation is a poor one, however, since the molecular cloud-like objects that form in the center of the halos are spheroidal in shape and generally have circular velocities that are far lower than the Keplerian circular velocity.

In Section 4.4.3, I examine the consistency of radial properties of halos which form Population III stars when the simulation volume and the large scale structure are varied. A dozen simulations are examined using three different box sizes. The simulations are compared when the gas at the center of the halo has reached $\sim 10^{11} \text{ cm}^{-3}$. Simulations with the same box size but different large scale structure show comparable evolution of the first protostar to form in each calculation,, though with significant scatter in properties such as the core temperature and accretion rate onto the central protostar. Simulations with different box sizes show a systematic trend towards higher molecular

hydrogen fractions, lower baryon temperatures, and lower overall accretion rates onto the primordial protostar with increasing box size. This result is robust and suggests that previous estimates of the Population III mass range from three-dimensional cosmological calculations may overestimate the Population III mass function by a significant amount.

Section 4.4.4 explores the formation of Population III stars in the presence of a soft UV (SUV) background. I use a single cosmological realization and vary the SUV background over a wide range of values. I show that for a significant range of the SUV background fluxes the main effect of this background is to delay the collapse of the primordial cloud by inhibiting the formation of molecular hydrogen. The final properties of the Population III protostar are quite similar regardless of the UV background strength, though final accretion rates vary nonlinearly with the SUV background flux. Halos subject to very high soft UV background fluxes do not form enough H_2 to collapse during the simulation, which implies that a different mode of primordial star formation must take place when the UV background is very high.

4.2 Motivation

As discussed in Section 1.3, many unresolved issues remain concerning the formation of Population III stars. Exploration of the detailed properties of these objects via purely analytical work is essentially impossible due to the wide range of physics involved, which includes the dynamics of dark matter systems, hydrodynamics, and the nonequilibrium chemistry and radiative cooling of the primordial gas out of which these objects form. Similarly, simulations have shown that the formation process of Population III stars is not inherently symmetrical, and that the formation of these stars takes place in a cosmological context. This suggests that in order to correctly model the formation of Population III stars, we need three-dimensional simulations with extremely high spatial and temporal dynamical range.

Previously published three dimensional, high dynamical range cosmological simulations of the formation of Population III stars (Abel, Bryan and Norman [39], hereafter ABN) are an important step towards understanding these objects. This work was performed in a $\Omega_m = 1$ universe, and derives results about the possible mass function of Population III stars using only one cosmological realization in a relatively small ($128 \text{ h}^{-1} \text{ kpc}$) simulation volume. In this chapter I explore the formation of Population III protostars in a ΛCDM universe, using multiple box sizes and cosmological realizations, in order to determine the robustness of the predictions in ABN.

Another important scenario for the formation of Population III stars involves the presence of a soft ultraviolet (SUV) background. Massive primordial stars are copious emitters of ultraviolet radiation, particularly in the Lyman-Werner energy band (11.18 – 13.6 eV) which is responsible for the photodissociation of molecular hydrogen. Since

this radiation is below the ionization energy of atomic hydrogen it is probable that photons in the Lyman-Werner band form a background of soft ultraviolet light, which could significantly affect the formation of later generations of Population III stars via the dissociation of molecular hydrogen. Previous work has been done on this subject by Machacek, Bryan & Abel [71] – however, the work presented here uses higher resolution calculations and examines more fully the evolution of a single halo.

4.3 Problem setup

All of the simulations described in this chapter are performed using the adaptive mesh cosmology code Enzo, which is described in detail in Section 2.2. They are initialized at $z = 99$ assuming a “concordance” cosmological model: $\Omega_m = 0.3$, $\Omega_b = 0.04$, $\Omega_{CDM} = 0.26$, $\Omega_\Lambda = 0.7$, $h = 0.7$ (in units of 100 km/s/Mpc), $\sigma_8 = 0.9$, and using an Eisenstein & Hu power spectrum [194] with a spectral index of $n = 1$. The selection of CDM power spectrum is unimportant, as at the length scales being explored in these calculations the power spectrum effectively becomes a power law with $P(k) \sim k^{-3}$ for all plausible CDM power spectra. Twelve simulations are generated using a separate random seed for each, meaning that the large-scale structure that forms in each of the simulation volumes is statistically independent of the others. These simulations are divided into sets of four simulations in three different box sizes: 0.3, 0.45, and 0.6 h^{-1} Mpc (comoving). The first halo to form in each simulation with a mass of $\sim 10^6 M_\odot$ is found using a dark matter-only calculation with 128^3 particles on a 128^3 root grid with a maximum of 4 levels of adaptive mesh, refining on a dark matter overdensity criterion of 8.0. The initial conditions are then regenerated with both dark matter and baryons for each of the simulation volumes such that the Lagrangian volume in which the halo formed is now resolved at much higher spatial and mass resolution using the nested initial condition method described in Section 2.1.2. These simulations have a 128^3 root grid and three static nested grids, for an overall effective grid size of 1024^3 in the region where the most massive halo will form. The highest resolution grid in each simulation is 256^3 grid cells, and corresponds to a volume (75, 112.5, 150) h^{-1} comoving kpc on a side for the (0.3, 0.45, 0.6) h^{-1} Mpc box. The dark matter particles in the highest resolution grid are (1.81, 6.13, 14.5) $h^{-1} M_\odot$ and the spatial resolution of cells on these grids are (293, 439, 586) h^{-1} parsecs (comoving). Though the simulations have a range of initial spatial resolutions and dark matter masses, we find that the final simulation results are converged – the spatial and mass resolution of the 0.3 h^{-1} Mpc volume simulations can be degraded to that of the 0.6 h^{-1} Mpc without significantly changing the results.

The simulations are then started at $z = 99$ and allowed to evolve until the collapse of the gas within the center of the most massive halo, which occurs at a range of redshifts (as shown in Section 4.4.3). The equations of hydrodynamics are solved using the PPM

method with a dual energy formulation, as described in Section 2.2.2 (the results are the same when the ZEUS hydrodynamic method is used). The nonequilibrium chemical evolution and optically thin radiative cooling of the primordial gas is modeled as described in Section 2.2.5, following 9 separate species including molecular hydrogen (but excluding deuterium). Adaptive mesh refinement is turned on such that cells are refined by factors of two along each axis, with a maximum of 22 total levels of refinement. This corresponds to a maximum resolution of (115, 173, 230) h^{-1} astronomical units (comoving) at the finest level of resolution, with an overall spatial dynamical range of 5.37×10^8 . To avoid effects due to the finite size of the dark matter particles, the dark matter density is smoothed on a comoving scale of ~ 0.5 pc. This is reasonable because at that radius in all of our calculations the gravitational potential is dominated by the baryons.

Grid cells are adaptively refined based upon several criteria: baryon and dark matter overdensities in cells of 4.0 and 8.0, respectively, checks to ensure that the pressure jump and/or energy ratios between adjoining cells never exceeds 5.0, that the cooling time in a given cell is always longer than the sound crossing time of that cell, and that the Jeans length is always resolved by at least 16 cells. This guarantees that the Truelove criterion [217], which is an empirical result showing that in order to avoid artificial fragmentation in numerical simulations the Jeans length must be resolved by at least 4 grid cells, is always maintained by a comfortable margin. Simulations which force the Jeans length to be resolved by a minimum of 4 and 64 cells produce results which are essentially identical to when the Jeans length is resolved by a minimum of 16 cells.

The simulations described in Section 4.4.4 take one of the models described previously (with a box size of $0.3 h^{-1}$ Mpc) and resimulate it assuming a range of unevolving soft UV backgrounds with intensities in the Lyman-Werner band of $F_{LW} = 0.0, 10^{-24}, 10^{-23}, 3 \times 10^{-23}, 10^{-22}, 10^{-21}$ and 10^{-20} $\text{erg s}^{-1} \text{cm}^{-2} \text{Hz}^{-1}$. This covers a much wider range of parameter space than the results described in Machacek, Bryan & Abel [71]. As with the other calculations, these are initialized at $z = 99$ and evolved until the collapse of the core of the largest halo, which occurs at a range of redshifts. The simulations with the two highest SUV fluxes do not collapse before $z = 10$, the point at which the simulations are stopped.

4.4 Results

4.4.1 Collapse of a representative primordial star

In this section we describe in detail the collapse of a single primordial protostar out of the ensemble discussed in Section 4.4.3. This simulation was selected at random out of the four simulations performed in a $0.3 h^{-1}$ Mpc comoving volume. The results described here are qualitatively similar for all of the calculations described in Section 4.4.3, though there is some scatter in the exact evolution of each halo due to differences in large scale

structure and the detailed merger history of the halo. However, since the collapse is essentially controlled by the chemistry of molecular hydrogen formation, the result is general.

Figures 4.1, 4.2, and 4.3 zoom in on the central gas core in each halo at the redshift of collapse by factors of four, showing projections of log baryon density, log baryon temperature, and maximum refinement level, respectively. The largest-scale panel shows a projection of a volume of the universe 1320 proper parsecs across and deep, and zooms in to approximately 1.3 pc across. Each panel is centered on the collapsing protostar. At large scales it is apparent from Figure 4.1 that the halo in which the first star in the simulation volume forms is at the intersection of two cosmological filaments, a distinctly asymmetrical situation. Examination of Figure 4.2 shows that the filaments and majority of the volume of the halo are relatively hot (~ 1000 Kelvin), due primarily to accretion shocks formed by gas raining onto the filaments and into the halo. However, as we zoom in towards the center of the halo we can see that the high-density gas is at a much lower temperature (a few hundred Kelvin) due to cooling by the significant quantity of molecular hydrogen that is formed in the halo. The gas within the halo is not particularly spherical until scales of a few parsecs are reached, where a slightly warmer core of gas forms with an overall mass of a few thousand solar masses, harboring a fully-molecular protostar with a mass of $\sim 1 M_{\odot}$. The central core is generally spheroidal due to gas pressure and is not rotationally supported at any point. Figure 4.3 shows how the adaptive mesh refinement is used to resolve the cosmological structure by concentrating refinement only where it is needed. This method is extremely effective at conserving computational resources - the level 16 grids, which are the highest level of resolution shown in Figure 4.3, only encompass $\sim 2.5 \times 10^{-17}$ of the overall volume!

Figures 4.4 through 4.7 show the evolution of radial profiles of several spherically averaged, mass-weighted baryon quantities of a representative primordial protostar from approximately the onset of halo collapse until the formation of a fully molecular protostar. The halo begins its collapse at $z = 18.05$ (approximately 2.04×10^8 years after the Big Bang) and ends its collapse 6.294×10^6 years later, at $z = 17.67$. Figure 4.4 shows the spherically-averaged baryon number density, temperature, and enclosed mass as a function of radius. Figure 4.5 shows the molecular hydrogen fraction, electron fraction, and H^- fraction as a function of radius. Figure 4.6 shows the evolution of angular momentum as a function of enclosed mass, baryon radial velocity as a function of radius, and circular velocity of the cloud as a function of radius. Figure 4.7 shows the ratios of gas cooling time to sound crossing time, cooling time to system dynamical time, and sound crossing time to dynamical time as a function of radius. The lines in all of these plots are color coded such that the same line color and type corresponds to the same time in each panel.

We begin to follow the evolution of the halo at $z = 18.05$, when the central hydrogen number density has reached $n \sim 10^3$ particles per cubic centimeter (black solid line in




Figure 4.1: Zoom on projected mass-weighted baryon density by factors of four for a representative Population III protostar formation calculation at the last simulation output. At this redshift ($z = 19.28$), the maximum density was $\sim 10^{12} \text{ cm}^{-3}$ with a cosmic mean density of $\simeq 0.003 \text{ cm}^{-3}$, for an overall density increase of 15 orders of magnitude. Top left: view is 1320 pc across. Top center: 330 pc. Top right: 82.5 pc. Bottom left: 20.6 pc. Bottom center: 5.2 pc. Bottom right: 1.29 pc. Note that all sizes are in proper parsecs at $z = 19.28$. In all panels yellow represents high densities and blue represents low density, with the color table relative in each frame.




Figure 4.2: Zoom on projected mass-weighted baryon temperature by factors of four in a representative Population III protostar formation calculation at the last simulation output. The collapse redshift is $z = 19.28$ and the simulation and spatial sizes of each panel are the same as in Figure 4.1. In all panels white represents high temperatures and dark colors represent low temperatures. The color table is relative in each frame.




Figure 4.3: Zoom on projected maximum level in a representative Population III protostar formation calculation at the last simulation output. The spatial scale for each panel and simulation are the same as in Figure 4.1. The maximum projected level in each panel is as follows. Top left: Level 6. Top center: Level 8. Top right: Level 10. Bottom left: Level 12. Bottom center: Level 14. Bottom right: Level 16. In each panel the highest level grid is represented in red, second highest in yellow, and so on. The highest level of resolution at this time is $L = 22$.

all plots). This corresponds to a core with a radius of ~ 1 parsec and a mass of a few thousand solar masses, which is accreting gas at a significant rate. The molecular hydrogen fraction within this core is slightly less than 10^{-3} but is still enough to rapidly cool the center of the halo to ~ 200 Kelvin at a cooling rate proportional to the square of the gas density. The gas cannot cool below this temperature because of the sharp decrease in the cooling rate of molecular hydrogen below $\simeq 200$ Kelvin. This core is the high-redshift equivalent of a molecular cloud core. The halo “loiters” for approximately six million years as the amount of molecular hydrogen is slowly built up to a mass fraction of a few times 10^{-3} and the central density increases. As the gas density passes roughly $n \sim 10^4 \text{ cm}^{-3}$ the ro-vibrational levels of H_2 are populated at their equilibrium value and the cooling rate again becomes independent of density, which corresponds to an increase in gas temperature with increasing density (as can be seen by the blue and green solid lines in the temperature vs. radius plot in Figure 4.4). As the temperature increases the cooling rate again begins to rise, leading to an increase in the inflow velocities of gas. Examination of the plot of enclosed mass vs. radius in Figure 4.4 shows that at this point the enclosed gas mass has exceeded the Bonnor-Ebert critical mass, which is defined as $M_{BE} = 1.18M_{\odot}(c_s^4/G^{3/2})P_{ext}^{-1/2}$, where c_s is the local sound speed and G is the gravitational constant. This is the critical mass at which an isothermal sphere of gas with an external pressure P_{ext} becomes unstable and undergoes collapse. This occurs in this halo at a mass scale of $\sim 1000 M_{\odot}$.

When the central density of the cloud core becomes sufficiently large ($n \sim 10^8 \text{ cm}^{-3}$) the three-body H_2 formation process takes over, resulting in a rapid increase in the molecular hydrogen fraction from a few times 10^{-3} to essentially unity. This causes a huge increase in the cooling rate, which results in a rapid drop in temperature of the center of the halo, allowing it to contract and causing an increase in central density of $n \sim 10^{15} \text{ cm}^{-3}$ in only another $\sim 2 \times 10^4$ years, with a corresponding increase in the inflow rates. At a mass scale of $\sim 1 M_{\odot}$ a protostellar core forms which is completely molecular and has gas accreting onto it supersonically, producing a protostellar accretion shock at ~ 100 astronomical units from its center. At this point the optical depth of the halo core becomes close to unity to molecular hydrogen ro-vibrational line emission, so we terminate the simulation because the assumption of optically thin radiative cooling used in our code is no longer correct.

It is useful to examine the relevant chemical, thermal and dynamical timescales of the collapsing halo. The ratios of cooling time to sound crossing time (calculated in spherically averaged radial shells) as a function of radius, cooling time to dynamical time, and sound crossing time to dynamical time are plotted in Figure 4.7. Within the core of the halo ($r \sim 1$ parsec) the sound crossing time (t_{cross}) is slightly less than the dynamical time (t_{dyn}) for the majority of the evolution time of the halo, while the cooling time (t_{cool}) is somewhat longer than both of these timescales (but generally only by a factor of a few). If $t_{cross} \ll t_{dyn}$ the halo is stable against collapse because the halo

can easily equilibrate its pressure to compensate for collapsing gas. If $t_{cross} \gg t_{dyn}$, the system cannot come into equilibrium and is in free-fall. In this case, $t_{cross} \approx t_{dyn} < t_{cool}$, and the system is undergoing a quasistatic collapse. This can also be seen by examining the evolution of the radial infall velocity as a function of radius in Figure 4.6, where the radial infall velocity is subsonic until the very last portion of the core’s evolution, when it becomes locally supersonic. This corresponds to a dramatic increase in the molecular hydrogen fraction, and a corresponding rapid decrease in the cooling time. In the center of the halo at the last few data outputs, the cooling time becomes shorter than both the dynamical time and sound crossing time, creating a situation where gas is essentially free-falling onto the central protostar.

As in ABN, we carefully examine the forming protostellar core for signs of fragmentation. This might be expected due to chemothermal instabilities caused by the rapid formation of molecular hydrogen via the 3-body process and the resulting rapid increase in cooling rate. However, the sound crossing time within the core is less than the H_2 formation timescale until the last output time, allowing mixing to take place and preventing the formation of large density contrasts. By the time that the H_2 formation timescale becomes shorter than the sound crossing time, the core is fully molecular and therefore stable against this chemothermal instability.

As discussed previously, at the time that the simulation is stopped (due to a breakdown in the assumption of optically thin radiative cooling at the center of the protostellar cloud) a fully-molecular protostar with a mass of $\sim 1 M_\odot$ has formed and is accreting gas supersonically. The spherically-averaged accretion time at the last output timestep, plotted as a function of enclosed gas mass, is shown as the red solid line in Figure 4.8. The accretion time is defined as $t_{acc} \equiv M_{enc}/\dot{M}$, where M_{enc} is the enclosed baryon mass and $\dot{M} \equiv 4\pi r^2 \rho(r)v(r)$, with $\rho(r)$ and $v(r)$ being the baryon density and velocity as a function of radius, and $v(r)$ defined as being positive towards the center of the halo. The green solid line the accretion time as determined by taking the local accretion rate from the Shu isothermal collapse model, $\dot{M}_{Shu} = m_0 c_s^3/G$, where m_0 is a dimensionless constant of order unity, c_s is the sound speed, and G is the gravitational constant. This value of \dot{M} is calculated in each bin and the accretion time is plotted as M_{enc}/\dot{M}_{Shu} . The dot-long dashed line is the Kelvin-Helmholz time for a Population III star with a mass identical to the enclosed mass, as calculated from the results given by Schaerer [59]. The dot-short dashed line is the baryon accretion time for the result in Abel, Bryan, and Norman.

The agreement between the spherically-averaged accretion rate and that estimated by the Shu isothermal collapse model is striking. As shown by Shu [204], as long as the densities in a condensing molecular cloud core span several orders of magnitude before a stage of dynamic instability is reached, the subsequent collapse properties of the cloud should resemble those of an isothermal sphere. The lack of characteristic time and length scales results in a self-similar wave of infalling gas which propagates outward




Figure 4.4: Evolution of spherically-averaged values for baryon number density (top left), baryon temperature (top right), and enclosed baryon mass (bottom) as a function of radius of a representative primordial protostar. The solid black line in each panel corresponds to spherically averaged radial profile of each quantity the onset of halo collapse, at $z = 18.05$ (approximately 2.04×10^8 years after the Big Bang). Solid blue line: the state of the halo 8.73×10^5 years after that. Solid green line: 5.103×10^6 years later. Solid red line: 2.99×10^5 years later. Solid cyan line: 16,660 years later. Dashed black line: 2267 years later. Dashed blue line: 310 years later. Dashed green line: 91 years later. Dashed red line: 31 years later, at a final output redshift of $z = 17.673$. The total time spanned by the lines in these panels is 6.294×10^6 years. The black dot-dashed line in the bottom left panel is the Bonnor-Ebert critical mass calculated at the last timestep.




Figure 4.5: Evolution of spherically averaged radial profiles of molecular hydrogen fraction (top left), electron fraction (top right), and H^- fraction (bottom) as a function of radius of a representative primordial protostar. The lines correspond to the same times as in Figure 4.4 and are of the same simulation.




Figure 4.6: Evolution of radial profiles of spherically-averaged baryon angular momentum as a function of enclosed mass of a representative primordial protostar (top left panel) and radial velocity and circular velocity as a function of radius (top right and bottom panels, respectively). The lines correspond to the same times as in Figure 4.4 and are of the same simulation. The black dot-dashed line in the plot of radial velocity as a function of radius is the sound speed calculated using the local baryon temperature in each radial bin at the last simulation timestep. The black dot-dashed line in the plot of baryon circular velocity vs. time is the Newtonian circular velocity computed from the radius and enclosed baryon mass at the last timestep.




Figure 4.7: Evolution of the ratio of gas cooling time to sound crossing time (top left panel), gas cooling time to system dynamical time (top right panel), and sound crossing time to system dynamical time (bottom panel) as a function of radius of a representative primordial protostar. These quantities are mass-weighted and spherically-averaged, and the lines correspond to the same times as in Figure 4.4 and are of the same simulation.

at the speed of sound, resulting in the accretion rate described above. This accretion rate can be derived in a more intuitive way by considering the properties of a cloud of gas with radius R and mass M_{cl} which is marginally unstable. The accretion rate of this gas must be given (as an order of magnitude estimate) by $\dot{M} \sim M_{cl}/t_{dyn}$, where $t_{dyn} = R/a$, where a is the characteristic velocity associated with the cloud (the virial velocity). If this cloud was originally marginally supported against its self-gravity, then $a^2 \sim GM_{cl}/R$ (where G is the gravitational constant), which can be substituted into the expression for \dot{M} to give $\dot{M} \sim a^3/G$, independent of R . In the case of this quasi-statically collapsing cloud, the virial speed is comparable to the sound speed c_s , giving $\dot{M} \sim c_s^3/G$. While the Shu model assumes that the entire cloud is of a constant temperature, our calculations have a varying temperature as a function of radius, and a radially-varying accretion rate based on this temperature is an excellent fit. This is reasonable because the isothermal collapse model assumes that the infall wave propagates at the local sound speed, assuming that the cloud is not supported by any other means. In this calculation we completely neglect the effects of magnetic fields, and it can be seen from Figure 4.6 that rotation contributes an insignificant amount of support, resulting in gas pressure being the sole means of support of the cloud.

Unfortunately, the final mass of the star remains unclear. This simulation (as well as all of the other calculations discussed in this chapter) lacks any of the necessary feedback processes that might halt accretion onto the protostar, making it impossible to accurately determine the main-sequence mass of the star. However, rough bounds on the mass of this object can be determined from examining Figure 4.8 and applying similar arguments to those used in ABN. A one solar mass protostar evolves far too slowly to halt accretion, particularly considering the high rates at which mass is falling onto the star ($\sim 10^{-2} M_{\odot}/\text{year}$ at $M_{enc} \leq 100 M_{\odot}$). Approximately $20 M_{\odot}$ of gas will be accreted within 10^3 years, and $\sim 60 M_{\odot}$ of gas will be accreted within 10^4 years. A conservative minimum mass limit to this star should be approximately $20 M_{\odot}$, since a thousand years is much shorter than the observed evolution timescales of galactic protostars. This is unclear, however, since little is known about the evolution of Population III protostars, whose evolution timescales may bear little relation to those of their modern-day brethren. A reasonable maximum mass is obtained by taking the mass at the maximum accretion time, which corresponds to $\sim 10^3 M_{\odot}$ at 2×10^6 years. However, the estimated main sequence lifetime of a $\sim 300 M_{\odot}$ primordial star is approximately 2×10^6 years, so it is implausible that the star could grow to be much larger than that. From these rough estimates, we obtain bounds on the mass of this Population III star of $\sim 20 - 300 M_{\odot}$. As discussed previously, however, this is a rough estimate at best – radiative feedback from the forming protostar will strongly affect the dynamics of the gas accreting onto the protostar. It is unlikely that photon pressure will significantly contribute to halting the accretion onto the protostar. However, as the star becomes luminous, production of ultraviolet light will cause photodissociation of molecular hydrogen in the accreting gas,

drastically affecting its cooling properties and possibly dramatically reducing or stopping the accretion of gas onto the protostar.

4.4.2 Angular momentum transport

One of the most pressing issues in modern-day star formation is the transport of angular momentum. The average rotational velocity observed in stars forming in the disk of our galaxy is several orders of magnitude smaller than one would assume if angular momentum of the collapsing molecular cloud out of which the star was formed were conserved, implying that somehow a huge amount of angular momentum is transported away from the center of a collapsing molecular cloud core. The mechanisms responsible for this in the galaxy are generally believed to be the fragmentation of molecular cloud cores into multiple clumps (with the majority of the angular momentum going into their orbits around each other) and transport via magnetic fields.

The scenario with Population III stars is significantly different. Examination of the plot of radial velocity vs. radius in Figure 4.6 shows that the collapse of the cloud core is never stopped by rotational support. The reasons for this can be seen by the plot of angular momentum vs. enclosed gas mass in the same figure. At the onset of collapse, the core of the gas cloud has a very low specific angular momentum, with the overall distribution being generally well described by a power law. This is a natural consequence of dark matter halos which are produced by gravitational collapse, as explained by Quinn & Zurek [207], and is a result of efficient mixing and angular momentum transport during the relaxation of the halo. They show that for well-mixed halos the angular momentum distribution can be shown to have a power-law behavior, as is observed in this simulation. Examination of the halo in which the first protostar forms shows it to be well-mixed, with little substructure, and with the dark matter density profile of a halo in virial equilibrium. All of these clues suggest that one might expect the angular momentum distribution described above.

Even though the gas which ends up in the protostellar core starts out with very little angular momentum, it is clear from the plot of specific angular momentum vs. enclosed mass in Figure 4.6 that there is still some angular momentum transport occurring. This can be seen even more clearly in Figure 4.9, which shows the mean angular momentum evolution of the individual parcels of gas which end up at various radii in the dark matter halo which contains the first protostar. Due to the Eulerian nature of Enzo, we cannot in general directly follow the evolution of a specific parcel of gas. However, the code was modified to include “tracer particles,” which are Lagrangian particles with zero mass that are placed in the baryon gas and simply follow along with the flow of the fluid. These particles were distributed evenly in the gas at the initialization of the simulation and were given unique numerical identifiers, allowing us to track each of their courses individually. These particles followed the flow of the baryon gas until the end of the simulation. At




Figure 4.8: Baryon gas accretion time as a function of enclosed baryon mass for a representative primordial protostar. This is defined as M_{enc}/\dot{M} , where M_{enc} is the enclosed baryon mass and $\dot{M} \equiv 4\pi r^2 \rho(r) v(r)$, with $\rho(r)$ and $v(r)$ being the baryon density and velocity as a function of radius, and $v(r)$ defined as being positive towards the center of the halo. The red solid line is the baryon accretion time for this simulation. The green solid line is the accretion time as determined by taking the accretion rate from the Shu isothermal collapse model, $\dot{M}_{Shu} = m_0 c_s^3 / G$, where m_0 is a dimensionless constant of order unity, c_s is the sound speed, and G is the gravitational constant. This value of \dot{M} is calculated in each bin and the accretion time is plotted as M_{enc}/\dot{M}_{Shu} . The dot-long dashed line is the Kelvin-Helmholtz time for a Population III star with a mass identical to the enclosed mass, as calculated from the results given by Schaerer. The dot-short dashed line is the baryon accretion time for the result in Abel, Bryan & Norman. The plot here corresponds to the last output dataset, corresponding to the red dashed line in Figures 4.4 through 4.7.

this time, we then bin the particles as a function of their final radius (separating the particles into bins spaced in roughly logarithmic intervals), and work backwards in time, calculating the mean angular momentum and radius of each bin of particles as a function of time. The results are displayed in the top panels of Figure 4.9. This shows that gas which is always outside of the core of the halo experiences proportionally very little transport of angular momentum or net movement inward, whereas the particles that end up in the innermost bins typically lose a factor of more than 50 in specific angular momentum and decrease their mean radius by a factor of more than 10. Figure 4.10 shows the evolution of cumulative angular momentum as a function of enclosed mass for this halo for each component of the angular momentum, as well as for the overall angular momentum of the system. This plot shows that the baryons within the cosmological halo conserve angular momentum overall. However, it is clear that within the innermost few thousand solar masses of gas, which corresponds to the quasistatically collapsing core of the halo, angular momentum is being redistributed outward.

What causes this transport of angular momentum? In this situation there are three plausible mechanisms for angular momentum transport in the absence of magnetic fields: Tidal coupling between gas in the core and the dark matter halo (or other nearby cosmological structures), non-axisymmetric perturbations in the baryons themselves (which generally serve to transport angular momentum outwards), and turbulent transport of angular momentum.

The first mechanism can be easily ruled out by examining the relative distributions of asymmetries in the baryons and dark matter. While the baryons that experience significant angular momentum transport show deviations from spherical symmetry on a scales of a few parsecs (see Figure 4.1), the dark matter is roughly spherical on scales smaller than 10 parsecs. This implies little tidal coupling between the two components of the halo. Tidal coupling from nearby cosmological structures can be ruled out by noting that the different components of the angular momentum appear to be redistributed in a similar manner, and also on a very rapid timescale. Tidal coupling from objects that are very far away and much larger than the halo core would cause relatively small changes which would take place much more gradually.

The analytical examination of non-axisymmetric perturbations requires us to treat the halo core as a disk. Using the formalism discussed by Hachisu et al. [208], we define a parameter $\beta \equiv T/|W|$, where T is the rotational energy and W is the gravitational energy, calculated in cylindrical shells around the axis of rotation. They show that an incompressible gas in axisymmetric equilibrium first experiences a dynamical gravitational instability that forms a non-axisymmetric perturbation (e.g. spiral density wave) at $\beta = 0.27$. The gas in our calculation is manifestly compressible, and Hachisu et al. show in a later paper that gas with a polytropic equation of state experiences these instabilities for $\beta \simeq 0.44$ [209]. The parameter β is calculated as a function of radius and shown in the top right panel of Figure 4.11 for the same series of simulation outputs




Figure 4.9: Evolution of mean baryon properties (as measured using baryon tracer particles) as a function of time. Left: Mean particle specific angular momentum as a function of time. Right: Mean particle radius as a function of time. Each line corresponds to the mean quantity for a number of particles particles that are in a specific spherical shell at the end of the calculation, traced backwards through the simulation. Black solid line: all particles with $r_{part} \leq 0.25$ pc at the end of the simulation. Blue solid line: particles with $0.25 < r_{part} \leq 0.5$ pc. Red solid line: particles with $0.5 < r_{part} \leq 1$ pc. Green solid line: particles with $1 < r_{part} \leq 2$ pc. Cyan solid line: particles with $2 < r_{part} \leq 4$ pc. Black dashed line: particles with $4 < r_{part} \leq 8$ pc. Blue dashed line: particles with $8 < r_{part} \leq 16$ pc. Red dashed line: particles with $16 < r_{part} \leq 32$ pc. Green dashed line: particles with $32 < r_{part} \leq 64$ pc.




Figure 4.10: Cumulative angular momentum as a function of enclosed mass for several data outputs. Top left panel: absolute value of x component of angular momentum. Top right panel: absolute value of y component of angular momentum. Bottom left panel: absolute value of z component of angular momentum. Bottom right panel: total angular momentum. The lines correspond to the same times as in Figure 4.4 and are of the same simulation.

described in Section 4.4.1. The critical value of β for noncompressible and polytropic gases are shown as horizontal black lines. This plot shows that the only gas that could be susceptible to non-axisymmetric perturbations is at very large radii which, while it may experience some small amount of angular momentum transport over the the evolution of the halo, is at a distance where the overall rotational period of the gas is significantly longer than the the evolution time of the core of the halo. The beta values for the gas which is seen to be undergoing angular momentum transport are far too small for non-axisymmetric perturbations to be responsible for this transport. Additionally, visual inspection of the core of the halo does not show any sort of spiral arm-type structures which are the classical manifestation of non-axisymmetric perturbations.

One can analytically examine the turbulent transport of angular momentum in an accretion disk using the method described by Li et al. [210]. They show that the radial flux of the angular momentum can be calculated by integrating the off-axis component of the Reynolds stress tensor such that $F_L = \langle \Sigma \delta v_r \delta v_\phi \rangle$, where F_L is the angular momentum flux, Σ is the surface density of the gas, $\delta v_r \equiv v_r - \langle v_r \rangle$ is the deviation from the mean radial velocity, $\delta v_\phi \equiv v_\phi - \langle v_\phi \rangle$ is the deviation from the mean circular velocity, and $\langle \dots \rangle$ indicates averaging over the azimuthal component, ϕ . The results are plotted in the bottom panel of Figure 4.11 in units of specific angular momentum transport (for consistency) and with line colors corresponding to previous plots in this chapter. A positive value of F_l indicates angular momentum transport outwards. There is some evidence for turbulent transport of angular momentum in this analysis.

One problem with the analytical results discussed above is that the approximation of a self-gravitating “thin disk” is very poor. Analysis of thin disks generally assumes that the scale height of the disk, h , is much smaller than the radius at any given point, r . The ratio of the scale height to the radius is typically calculated as $h/r \equiv c_s/v_{circ}$, where c_s is the sound speed and v_{circ} is the circular velocity. This implies that $c_s \ll v_{circ}$, which is not the case in this situation. In fact, for the radii in question, $c_s \simeq v_{circ}$, implying that the cloud core is spheroidal. This can be shown more quantitatively by examining the moment of inertia of the cloud core. The diagonal components of the moment of inertia have values that are similar. If the core were disk-like, one of the three components of the moment of inertia would be much smaller than the other two. Additionally, one cannot apply standard analytics that use a “thick disk” approximation because in standard situations where a thick disk is relevant, such as the accretion disks around compact objects, there is a central source (such as a black hole) whose gravity overwhelmingly dominates the self-gravity of the disk, which is not the case in this situation. Furthermore, analytical models of thin disks typically assume that the disk is rotating at a Keplerian orbital velocity (i.e. is rotationally supported) and that there is a central object such as a star or black hole dominating the gravitational potential. In the situation examined here, the overall circular velocity is far less than the Keplerian velocity and there is no central object dominating the potential. This casts further doubts on the validity of

using analysis techniques for thin disks on this particular situation.

Even though the validity of the analytical estimates for angular momentum transport discussed previously are dubious, this does not mean that all of the scenarios discussed previously are completely ruled out. Figure 4.13 shows the spherically-averaged, mass-weighted evolution of the radial velocity, RMS (root-mean-square) velocity of the baryon gas, and RMS Mach number of gas in the collapsing halo. Examination of this figure shows that at all radii where angular momentum transport is significant, the RMS velocity of the baryon gas is much greater than the average radial velocity, and that the average radial velocity is much smaller than the sound speed at most radii, while the RMS Mach number vs. radius plot shows that the baryon velocity is significantly larger than the sound speed at radii where angular momentum is being transported. This suggests that the gas in the core of the halo is experiencing significant supersonic turbulent motion, which may be responsible for the transport of angular momentum. In a turbulent medium in a gravitational potential, a given volume of space contains packets of gas that have a large distribution of angular momentum with respect to the center of the well. This turbulent medium effectively transports angular momentum outwards by allowing packets of gas with low angular momentum to sink towards the center of the potential well, replacing gas with higher angular momentum. This mechanism is only effective when the cooling time of the gas is longer than the dynamical time (e.g. when gas pressure plays a significant role). This scenario is given credibility by Figure 4.12, which shows a scatter plot of the specific angular momentum of tracer particles (with respect to the maximum baryon density) at the onset of halo collapse to the distance of that particle from the baryon density maxima at the point at which the simulation is terminated. Gas within the innermost few parsecs (which is the region undergoing quasistatic collapse and angular momentum transport) shows a distinct (though noisy) relationship between initial angular momentum and final distance with respect to the halo center.

A further possible source of transport of angular momentum could be due to numerical shear viscosity. In order to determine whether this is the case, a parameter study was carried out where we varied the effective resolution of the simulation by a factor of 16 along each grid axis by enforcing that the Jeans length be resolved by a minimum of 4, 16 or 64 cells. The properties of the halo in all cases were very similar. Additionally, we used two different hydrodynamical methods (both PPM and the method used in the ZEUS code) with different orders of accuracy, and found no significant differences between the two calculations. The ZEUS hydro method has an explicit artificial viscosity for shock capturing purposes, and the PPM method uses a Riemann solver which has no explicit numerical viscosity, and they get the same result. While this is not a formal proof of the lack importance of numerical viscosity, it is highly suggestive that the observed angular momentum transport is not caused by numerical effects.

Quantifying the magnitude of numerical viscosity is difficult due to the range of spatial and temporal resolutions. The most straightforward way to do so would be to simulate




Figure 4.11: Evolution of quantities related to the angular momentum as a function of radius at several output times. Top left panel: Specific angular momentum as a function of radius. Top right panel: Evolution of the β parameter as a function of radius. Bottom panel: specific angular momentum flux as a function of radius. In the top right panel the critical values of β for compressible and noncompressible fluids are shown by the top and bottom horizontal black lines, respectively. The line colors correspond to the same output times as are described in previous figures.




Figure 4.12: Scatter plot of initial specific angular momentum of tracer particles with respect to the maximum density of the collapsing halo core versus the particle’s final distance from the halo core at the point that the simulation is terminated.

a problem with an analytical solution that explicitly includes viscosity and solve it on an adaptive grid using the Euler equations, which implicitly assume no viscosity. Then one would observe the evolution of the problem and post facto estimate the numerical viscosity. The numerical viscosity in an adaptive mesh code may be highly dependent on the test problem.

4.4.3 Consistency of result across multiple realizations

Another important issue relating to studies of the formation of Population III stars in a cosmological context is the consistency of the results over a range of simulation parameters. As discussed in Section 4.2, previously published high dynamical range calculations of Pop III star formation have concentrated upon a single cosmological realization. While this is an important first step, it neglects possible systematic effects relating to simulation box size and other parameters, and also allows for error due to small number statistics.

In this section I attempt to address some of these issues. Twelve simulations are set up as described in Section 4.3. Each simulation has a different cosmological realization (i.e. large scale structure). Four simulations in each of three box sizes (0.3, 0.45, and 0.6 h^{-1} comoving Mpc) are performed, with the results shown in Figures 4.14 through 4.23.

Figures 4.14 - 4.16 display several mean properties of the halos. In each of the panels in these graphs the information for each separate simulation is plotted as a filled-in square which is color coded according to box size as described in the figure captions. The colored crosses correspond to mean values for all simulations of a given box size (with colors again corresponding to the box size), and the green circle corresponds to the mean of all twelve of the simulations together.

The top left panel of Figure 4.14 shows the virial mass of each halo at the time of protostellar cloud collapse plotted against the redshift of collapse. Though there is a large amount of scatter in virial mass overall (with the smallest halo having a mass of $1.36 \times 10^5 M_{\odot}$ and the largest 6.92×10^5), the average virial mass in each box size is virtually identical. The mean virial mass of all twelve of the halos is $3.63 \times 10^5 M_{\odot}$, which is significantly lower than the halo mass of $7 \times 10^5 M_{\odot}$ in ABN. In contrast to the virial mass, there is a strong trend in earlier collapse time (large collapse redshift) as a function of box size, with the 0.45 and 0.6 h^{-1} Mpc boxes collapsing at a mean redshift of $z \simeq 27.5$ and the 0.3 h^{-1} Mpc boxes collapsing at a mean redshift of $z \simeq 22$. This can be understood as a result of the distribution of power as a function of simulation volume. Since the power as a function of wave number falls off as $P(k) \sim k^{-3}$ at large wave numbers (small box sizes), doubling the box size significantly increases the power on large scales. The net results of this is that the most massive halo in each box forms significantly earlier, out of a density peak that is statistically denser overall, and also undergoes a more rapid merger history. An effect of this can be seen in the top




Figure 4.13: Top left panel: spherically-averaged, mass-weighted radial velocity as a function of radius for several simulation output times. Top right panel: spherically-averaged, mass-weighted root-mean-square velocity as a function of radius for several simulation output times. Bottom panel: spherically-averaged, mass-weighted RMS Mach number as a function of radius for several simulation output times. The line colors correspond to the same output times as are described in previous figures.

right panel of Figure 4.14, which shows the mean baryon temperature in each halo as a function of collapse redshift. As with the plot of virial mass vs. collapse redshift, there is a significant amount of scatter in the results, but a clear trend of increased overall halo temperature with box size is apparent. This is explainable in terms of competing heating and cooling effects. The gas in a given halo is heated by shocking during merger events, and cools radiatively (primarily due to line cooling from molecular hydrogen at these temperatures for a gas of primordial composition). A higher overall merger rate results in a warmer overall halo temperature assuming that the halo merger rate is comparable to the overall gas cooling rate (which is true for halos in the early universe). However, the rate of molecular hydrogen production at low densities increases as a function of temperature (due to the increased availability of free electrons) at the temperatures relevant to Population III halos, so even though the halos that form in larger boxes have higher overall mean temperatures, they also produce more molecular hydrogen in their cores, leading to an overall cooler halo core when collapse takes place. This will be discussed in more detail later in this section.

The mean halo baryon temperature vs. halo virial mass is plotted in the bottom left panel of Figure 4.14, and the mean halo temperature versus the halo virial temperature is plotted in the bottom right hand panel. The dashed line in the plot of halo mean temperature vs. virial mass scales as $T \sim M^{2/3}$, arbitrarily scaled to intersect the mean temperature and mass value for all simulations. This is the scaling relation expected from gas in a halo forming in a situation where radiative cooling is unimportant, and is commonly used to understand the mass-temperature relationship of the intracluster medium in galaxy clusters. There is some relationship between mean temperature and virial mass that generally conforms to this power law, which is due to the relatively poor cooling properties of molecular hydrogen. Still, there is a great deal of scatter in the relationship. The amount of scatter is reduced when considering the mean halo temperature vs. halo virial radius (shown in the bottom right hand panel). The halo virial temperature is a function of both halo mass and of redshift, with $T_{vir} \sim M_{vir}^{2/3}(1+z)$. The reduction in scatter is primarily due to the general trend of halos collapsing at higher redshifts having higher overall gas temperatures for an unevolving mean virial mass, which results in overall higher virial temperatures at high redshift. Though there is a clear relationship between halo baryon temperature and virial temperature, the mean baryon temperature in all cases is significantly lower than the halo virial temperature, suggesting that radiative cooling plays a non-negligible role in the overall temperature of the halo despite its generally poor cooling properties at low temperatures. If radiative cooling were completely unimportant the mean halo baryon temperature would be approximately the virial temperature.

Figure 4.15 shows the relationship of the angular momentum in the halos with various quantities. The angular momentum of a cosmological halo can be described as a function of the dimensionless spin parameter, $\lambda \equiv J|E|^{1/2}/GM^{5/2}$, where J is angular momentum,




Figure 4.14: Plots of basic halo quantities for 12 different cosmological random realizations. Top left panel: halo virial mass vs. protostar collapse redshift. Top right panel: mean mass-weighted halo baryon temperature vs. collapse redshift. Bottom left panel: mean mass-weighted halo vs. halo virial mass. Bottom right panel: halo mean baryon temperature vs. halo virial temperature. In each plot, black, blue and red squares correspond to simulations with $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ comoving box sizes. Each colored “x” corresponds to the average value for simulations with that box size, with colors corresponding to the same simulation box sizes as the squares. The green circle corresponds to the average for all simulations together. The dashed line in the plot of halo mean temperature vs. virial mass (bottom left) scales as $T \sim M^{2/3}$, arbitrarily scaled to intersect the mean temperature and mass value for all simulations.

E is the total energy, G is the gravitational constant and M is the halo mass. This is roughly equivalent to the ratio of the angular momentum of the halo to the angular momentum needed for the halo to be rotationally supported. Typical values for the halo spin parameter are 0.02–0.1, with a mean value of 0.05 [221, 222]. The gas spin parameter is somewhat lower than the dark matter spin parameter, which this is a function of the way in which the total energy and masses are calculated, and not an indication that there is less angular momentum per unit mass in the baryons as opposed to the dark matter.

The top left panel of Figure 4.15 shows the gas and dark matter spin parameters plotted against each other for the halo in each simulation that forms the Population III protostar, at the time of collapse. The mean value of the dark matter spin parameter is approximately 0.05, and both this and the overall range and distribution of the spin parameters agree with previous analytical and numerical results [221, 222]. The baryon gas spin parameter is lower overall (this is an effect of the scaling of the parameter and should be taken as a renormalization), and the distribution agrees with previous work. There appears to be some overall positive correlation between the dark matter and baryon spin parameters (e.g. halos with higher overall dark matter angular momentum tend to have higher overall baryon angular momentum) but there is considerable scatter. In all cases the spin parameters are much less than one, which suggests that the halos have little overall angular momentum. This is a general property of cosmological halos, and is consistent with previous analytical and numerical work, as well as the properties of angular momentum in the representative halo discussed in Section 4.4.1.

The top right hand panel of Figure 4.15 plots dark matter spin parameter vs. collapse redshift of the halo. There is no evidence for evolution of the spin parameter with redshift. The bottom left and right panels of Figure 4.15 plot the baryon and dark matter spin parameters against the halo virial mass. As with the other quantities examined, there is considerable scatter in the distributions, but no evidence for a clear relationship between halo virial mass and either gas or dark matter spin parameter. In all of the panels of this figure there is no evidence for any systematic effect due to box size.

Figure 4.16 plots the angle between the overall dark matter and baryon angular momentum vectors (θ) versus several different quantities. The top left panel plots θ vs. halo virial mass at the time of formation of the Pop III protostar in each halo. Overall, the average value for θ is approximately 25 degrees, which is consistent with recent numerical simulations. There is a great deal of scatter in θ , which is also consistent. There is no evidence for correlation between θ and halo virial mass. The top right panel plots θ vs. collapse redshift for each simulation, and the bottom left and right panels plot the gas and dark matter spin parameters vs. θ , respectively. There appears to be no correlation between θ and collapse redshift or the gas or dark matter spin parameters, and no evidence of there being any systematic effect due to box size.

In addition to plots of mean halo properties, it is very useful to look at more detailed information about each halo. Figures 4.17 through 4.23 show spherically-averaged, mass-




Figure 4.15: Plots of basic halo quantities for 12 different cosmological random realizations. Top left panel: gas spin parameter vs. dark matter spin parameter. Top right panel: dark matter spin parameter vs. halo collapse redshift. Bottom left: gas spin parameter vs. halo virial mass. Bottom right: dark matter spin parameter vs. halo virial mass. In each plot, black, blue and red squares correspond to simulations with $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ comoving box sizes. Each colored “x” corresponds to the average value for simulations with that box size, with colors corresponding to the same simulation box sizes as the squares. The green circle corresponds to the average for all simulations together.




Figure 4.16: Plots of basic halo quantities for 12 different cosmological random realizations. Top left panel: Theta (angle between gas and dark matter angular momentum vectors) vs. halo virial mass. Top right panel: theta vs. halo collapse redshift. Bottom left panel: gas spin parameter vs. theta. Bottom right panel: dark matter spin parameter vs. redshift. In each plot, black, blue and red squares correspond to simulations with $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ comoving box sizes. Each colored “x” corresponds to the average value for simulations with that box size, with colors corresponding to the same simulation box sizes as the squares. The green circle corresponds to the average for all simulations together.

weighted radial profiles of several baryon quantities in eleven of the twelve simulations (one simulation crashed and could not be restarted before reaching a high enough density). Since the cores of the most massive halo in each simulation collapse at a range of redshifts, it is not worthwhile to compare each halo at a specific point in time. Instead, we choose to compare them at a fixed point in the halo's evolution, as measured by the peak central baryon density in the protostellar cloud, which is roughly analogous to a constant point in the evolution of the protostar. In each of the figures discussed below, the top left panel shows radial profiles for all of the simulations plotted together. The top right panel shows only the results for the $0.3 \text{ h}^{-1} \text{ Mpc}$ box, the bottom left panel shows only results for the $0.45 \text{ h}^{-1} \text{ Mpc}$ box, and the bottom right panel shows only results for the $0.6 \text{ h}^{-1} \text{ Mpc}$ box. Line of a given color and line type correspond to the same simulation in all figures.

Figure 4.17 shows the plots of number density as a function of radius for eleven simulations, shown at approximately the same point in their evolution. There is remarkably little scatter in the density profiles for all of the simulations, and the density profiles all tend towards $\rho(r) \sim r^{-2}$. It was shown by Bodenheimer & Sweigart [205] that for a cloud of gas that is nearly isothermal and slowly rotating and which has negligible support from a magnetic field, the subsonic evolution of the gas will tend to produce a $1/r^2$ density distribution as long as the thermal pressure remains approximately in balance with the gravitational field. In particular, Chandrasekhar [206] showed that a molecular cloud core which forms at subsonic speeds will tend towards the density distribution of a singular isothermal sphere,

$$\rho(r) = \frac{c_s^2}{2\pi G r^2} \quad (4.1)$$

where $c_s \equiv (kT/m)^{1/2}$ is the isothermal sound speed, T , k , and m are the temperature, Boltzmann's constant, and mean molecular weight of the gas, respectively, and r is the radius. Since the halos generally have low angular momentum (as seen in Figure 4.15) and magnetic fields are completely neglected in these simulations, it is reasonable that the density go as $\rho(r) \sim r^{-2}$ in all of the simulations. The overall normalization of the density profiles also agrees very well. This can be understood as a result of the cooling properties of hydrogen gas. Each of the halos examined in this figure has the same composition, and therefore is cooled by the same mechanism. Only a small amount of molecular hydrogen is needed to cool the gas relatively efficiently, suggesting that in a halo that is in a somewhat stable equilibrium the gas temperature at low densities should be approximately constant for different halos, independent of the molecular hydrogen fraction. At densities above approximately 10^4 cm^{-3} the cooling rate becomes independent of density and the overall evolution of the gas happens very rapidly, so small differences in the initial molecular hydrogen fraction become magnified (as discussed later in this section).




Figure 4.17: Mass weighted, spherically-averaged baryon number density as a function of radius for 11 different cosmological random realizations, chosen at an output time where peak baryon density values are approximately the same. There are three box sizes: $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ (comoving). One of the $0.45 \text{ h}^{-1} \text{ Mpc}$ simulations has been omitted since the simulation crashed before reaching a comparable density value. Top left panel: All 11 simulations plotted together. Top right panel: $0.3 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom left panel: $0.45 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom right panel: $0.6 \text{ h}^{-1} \text{ Mpc}$ box simulations.

Figure 4.18 shows the baryon temperature as a function of radius. At radii outside of ~ 1 parsec, the temperature profiles are similar between all of the simulations, though halos forming in larger simulation volumes tend to have a higher overall temperature. At smaller radii there is significant scatter in core temperature of the simulations (for a fixed density), with a systematic trend towards halos forming in larger boxes having a lower overall core temperature. Examination of Figure 4.19 (molecular hydrogen mass fraction as a function of radius) shows that halos which form in a larger simulation volume have systematically larger H_2 mass fractions, though this effect is much more pronounced in the core of the halo than in the envelope. This difference in molecular hydrogen fraction can be understood as a result of the overall halo temperature. The rate at which molecular hydrogen is produced at low densities is limited by the availability of free electrons, as described in Section 1.3.2. The mean fraction of free electrons available in the primordial gas is a function of baryon temperature, with larger temperatures corresponding to larger electron fractions. On the other hand, the rate at which molecular hydrogen forms via the H^- channel declines at high temperatures. Since the limiting reaction in the formation of molecular hydrogen via the H^- channel is the formation of free electrons, this reaction dominates, and it can be shown using a simple one-zone calculation following the nonequilibrium primordial chemistry that molecular hydrogen production is maximized at ~ 1000 K. Halos with higher overall baryon temperatures will have systematically higher molecular hydrogen fractions. Once the core of the halo begins to collapse to very high densities small differences in the molecular hydrogen fraction are amplified, resulting in a general trend towards halos with higher overall baryon temperatures having higher H_2 fractions in their cores, and thus lower central temperatures.

Figures 4.20 and 4.21 show the enclosed baryon mass and circular velocity in each halo as a function of radius. The plot of enclosed mass versus radius shows very little scatter between the different simulations. This is to be expected since this is essentially another way of showing that the overall density distributions of the halos has little scatter (as in Figure 4.17), and is a result of the cooling properties of a primordial gas. The plot of circular velocity as a function of radius shows much more scatter, though there is no clear trend with simulation box size. In all cases the overall circular velocity is significantly less than the Keplerian orbital velocity, which agrees with our previous observation that the halos have generally low angular momentum, and that during the collapse of the core angular momentum is transported outward by turbulence.

It is useful to examine Figures 4.22 and 4.23 together, since they are essentially two different ways of looking at the same data. Figure 4.22 is the mean baryon radial velocity as a function of radius, and Figure 4.23 is the baryon accretion time as a function of enclosed mass. Figure 4.22 shows that there is a clear systematic effect present, where halos forming in simulations with larger boxes having a significantly lower overall radial velocity at small radii. This translates directly to a lower overall accretion rate




Figure 4.18: Mass-weighted, spherically-averaged baryon temperature as a function of radius for 11 different cosmological random realizations, chosen at an output time where peak baryon density values are approximately the same. There are three box sizes: $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ (comoving). One of the $0.45 \text{ h}^{-1} \text{ Mpc}$ simulations has been omitted since the simulation crashed before reaching a comparable density value. Top left panel: All 11 simulations plotted together. Top right panel: $0.3 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom left panel: $0.45 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom right panel: $0.6 \text{ h}^{-1} \text{ Mpc}$ box simulations.




Figure 4.19: Mass-weighted, spherically-averaged molecular hydrogen fraction as a function of radius for 11 different cosmological random realizations, chosen at an output time where peak baryon density values are approximately the same. There are three box sizes: $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ (comoving). One of the $0.45 \text{ h}^{-1} \text{ Mpc}$ simulations has been omitted since the simulation crashed before reaching a comparable density value. Top left panel: All 11 simulations plotted together. Top right panel: $0.3 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom left panel: $0.45 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom right panel: $0.6 \text{ h}^{-1} \text{ Mpc}$ box simulations.




Figure 4.20: Mass-weighted, spherically-averaged enclosed baryon mass as a function of radius for 11 different cosmological random realizations, chosen at an output time where peak baryon density values are approximately the same. There are three box sizes: $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ (comoving). One of the $0.45 \text{ h}^{-1} \text{ Mpc}$ simulations has been omitted since the simulation crashed before reaching a comparable density value. Top left panel: All 11 simulations plotted together. Top right panel: $0.3 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom left panel: $0.45 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom right panel: $0.6 \text{ h}^{-1} \text{ Mpc}$ box simulations.




Figure 4.21: Mass-weighted, cylindrically-averaged baryon circular velocity as a function of radius for 11 different cosmological random realizations, chosen at an output time where peak baryon density values are approximately the same. There are three box sizes: $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ (comoving). One of the $0.45 \text{ h}^{-1} \text{ Mpc}$ simulations has been omitted since the simulation crashed before reaching a comparable density value. Top left panel: All 11 simulations plotted together. Top right panel: $0.3 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom left panel: $0.45 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom right panel: $0.6 \text{ h}^{-1} \text{ Mpc}$ box simulations.

onto protostellar forming in halos in larger simulation volumes, which can be seen in Figure 4.23. For a wide range of enclosed mass, the average accretion rate of halos in the $0.6 \text{ h}^{-1} \text{ Mpc}$ simulations is more than an order of magnitude less than that of halos in the $0.3 \text{ h}^{-1} \text{ Mpc}$ boxes. As discussed in Section 4.4.1, this can be understood using the Shu isothermal sphere model, where subsonic collapse of gas onto the core of the sphere occurs at a rate controlled by the sound speed. Since the core temperatures are lower overall in the large simulation volumes, this translates to a lower sound speed and overall lower accretion rate. The implications of this are extremely significant – a lower accretion rate implies a lower overall Population III IMF in larger boxes. Applying the same estimates for the bounds of the stellar masses used in Section 4.4.1, we obtain a mass range of roughly $10 - 500 M_{\odot}$ for all four of the halos that form in simulations with a box size of $0.3 \text{ h}^{-1} \text{ Mpc}$, and $\sim 10 - 100 M_{\odot}$ for the halos that form in simulations with a box size of $0.6 \text{ h}^{-1} \text{ Mpc}$, though the mean maximum mass (based on the Kelvin-Helmholtz time) in the smaller box is $\simeq 200 M_{\odot}$ and in the larger box is $\simeq 30 M_{\odot}$.

4.4.4 The formation of a Population III star in the presence of a soft UV background

Another important scenario for the formation of Population III stars involves the presence of a soft ultraviolet (SUV) background. As discussed in Section 4.2, massive primordial stars are copious emitters of ultraviolet radiation, particularly in the Lyman-Werner band ($11.18 - 13.6 \text{ eV}$) which is responsible for the photodissociation of molecular hydrogen. Since this radiation is below the ionization energy of atomic hydrogen it is predicted that photons in the Lyman-Werner band would form a nearly uniform background of soft ultraviolet light, which could significantly affect the formation of later generations of Population III stars due to the dissociation of molecular hydrogen. Previous work has been done on this subject by Machacek, Bryan & Abel [71] – however, the work presented here uses higher resolution calculations and examines more fully the evolution of a single halo.

The simulations are set up as described in Section 4.3 in an $0.3 \text{ h}^{-1} \text{ Mpc}$ box. A single cosmological realization is resimulated assuming a constant Lyman-Werner soft UV background with intensities of $F_{LW} = 0, 10^{-24}, 10^{-23}, 3 \times 10^{-23}, 10^{-22}, 10^{-21}$ and $10^{-20} \text{ ergs s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1}$, which covers a much wider range of parameter space than the results described by Machacek et al. The simulations are initialized at $z = 99$ and are evolved until the collapse of the core of the largest halo, which occurs at a range of redshifts. The simulations with the two highest SUV fluxes do not collapse before $z = 10$, when these two simulations are stopped.

Figure 4.24 shows mean halo quantities for several of these simulations at the redshift of collapse of the halo core. The top left panel shows the Lyman-Werner flux vs. halo collapse redshift for all of the simulations whose halos actually collapsed. The top right




Figure 4.22: Mass-weighted, spherically-averaged baryon radial velocity as a function of radius for 11 different cosmological random realizations, chosen at an output time where peak baryon density values are approximately the same. There are three box sizes: $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ (comoving). One of the $0.45 \text{ h}^{-1} \text{ Mpc}$ simulations has been omitted since the simulation crashed before reaching a comparable density value. Top left panel: All 11 simulations plotted together. Top right panel: $0.3 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom left panel: $0.45 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom right panel: $0.6 \text{ h}^{-1} \text{ Mpc}$ box simulations.




Figure 4.23: Mass-weighted, spherically averaged baryon accretion time as a function of radius for 11 different cosmological random realizations, chosen at an output time where peak baryon density values are approximately the same. There are three box sizes: $0.3 \text{ h}^{-1} \text{ Mpc}$, $0.45 \text{ h}^{-1} \text{ Mpc}$ and $0.6 \text{ h}^{-1} \text{ Mpc}$ (comoving). One of the $0.45 \text{ h}^{-1} \text{ Mpc}$ simulations has been omitted since the simulation crashed before reaching a comparable density value. Top left panel: All 11 simulations plotted together. Top right panel: $0.3 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom left panel: $0.45 \text{ h}^{-1} \text{ Mpc}$ box simulations. Bottom right panel: $0.6 \text{ h}^{-1} \text{ Mpc}$ box simulations. The baryon accretion time is defined as $T_{acc} \equiv M_{enc}/\dot{M}$, where M_{enc} is the enclosed baryon mass and $\dot{M} \equiv 4\pi r^2 \rho(r) v(r)$, with $\rho(r)$ and $v(r)$ being the baryon density and velocity as a function of radius, and $v(r)$ defined as being positive towards the center of the halo. The dot-long dashed line in each panel is the Kelvin-Helmholtz time for a Population III star with a mass identical to the enclosed mass, as calculated from the results given by Schaerer. The dot-short dashed line in each panel is the baryon accretion time for the result in Abel, Bryan & Norman. The upper and lower diagonal solid black lines correspond to constant accretion rates of 10^{-3} and $10^{-3} \text{ M}_{\odot}/\text{yr}$, respectively.

panel shows the Lyman-Werner flux vs. virial mass of the eventual halo that formed for each simulation. Finally, the bottom panel shows the virial mass vs. collapse redshift for each of these calculations. The collapse redshift of the “control” simulation ($F_{LW} = 0$) is shown as a vertical blue dashed line in the top two panels and as a blue square in the bottom panel.

This figure shows that there is a clear relationship between the Lyman-Werner flux intensity and the collapse redshift and virial mass of the halo. A larger Lyman-Werner flux results in a later collapse time because the halo must be larger in order to have core densities high enough that significant amounts of molecular hydrogen can form in them. The final mass of the halo in the simulation with $F_{LW} = 10^{-22}$ is approximately five times that of the control simulation, and collapses significantly later in time. This agrees qualitatively with the results seen by Machacek, Bryan and Abel [71], who suggest that there is a “minimum halo mass” which is a function of the strength of the UV background. The practical effect of this is that as the Lyman-Werner UV background builds up the minimum halo mass which is necessary to form significant amounts of molecular hydrogen climbs, causing an overall suppression of the formation of Population III stars in halos with masses that are $\sim 10^6 M_{\odot}$. When the Lyman-Werner flux becomes extremely large, the formation of molecular hydrogen is almost entirely suppressed, resulting in termination of Pop III star formation in halos in this mass range. In this situation, primordial star formation cannot occur in halos this small – rather, the star formation must occur in halos that are massive enough that the mean gas temperature in the halo is at least 10^4 Kelvin, at which point the gas can cool efficiently via atomic line cooling. When a high enough density is reached through this cooling mechanism, the formation of molecular hydrogen can take place essentially independent of the strength of the UV background, allowing primordial star formation to occur. This mode of star formation has not been explored in depth by any published numerical calculations (though Bromm & Loeb [211] use a relatively poorly resolved calculation of the evolution of these halos to examine the possible formation of the first supermassive black holes), and the resulting IMF is unknown.

It is worth noting that we do not see the same functional form for the “threshold mass” of Machacek et al. They predict that the lowest halo mass that can collapse in a simulation with a uniform Lyman-Werner flux is

$$M_{TH}(M_{\odot}) = 1.25 \times 10^5 + 8.7 \times 10^5 \left(\frac{F_{LW}}{10^{-21}} \right)^{0.47} \quad (4.2)$$

Where F_{LW} is the Lyman-Werner flux in the same units as above. Our calculations agree with this result qualitatively: The Machacek fitting form is of the lowest possible halo mass that can collapse, whereas our halo is more representative of a “typical” halo, in the sense that its mass is approximately the average halo mass (as determined from previous sections). Of the halos that do collapse in our calculations, the halo mass at the

time of collapse is well-fit by the power law $M(F_{LW}) = 4.35 \times 10^5 (F_{LW}/10^{-24})^{0.27} M_{\odot}$ (excluding the $F_{LW} = 0$ case). Both of these functional forms are plotted in the top right panel of Figure 4.24, with the Machacek et al. result as a green dashed line and our fitting form as a red dashed line. The apparent lack of agreement is due to our simulation of a “typical” halo, whereas Machacek plots a threshold mass. Additionally, they use a cosmological model with somewhat different values for σ_8 , Ω_m , Ω_{Λ} , Ω_b and h , which may contribute to the lack of agreement.

Figure 4.25 through 4.27 show spherically-averaged, mass weighted radial profiles of several baryon quantities of all simulations, including those whose largest halo did not collapse. Radial profiles for simulations which did undergo halo collapse are plotted when the central density reaches a set density ($n \simeq$ a few times 10^{10} cm^{-3}). Simulations which did not undergo halo collapse are shown at the last data output, $z = 10$.

Figure 4.25 shows the spherically-averaged radial profiles of baryon number density, baryon temperature, and enclosed mass as a function of radius. As in Section 4.4.3, the baryon density profiles and enclosed mass profiles are similar between the different calculations, which can be explained by the cooling properties of a primordial gas. The temperature profiles of the collapsed gas are similar at large radii, but in the center of the halo there is a trend towards simulations with higher Lyman-Werner fluxes having a higher core temperatures, with the results being separated essentially into two populations. The overall core temperature of the population with fluxes $F_{LW} \leq 10^{-24} \text{ ergs s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1}$ is roughly 500 Kelvin, while the population with fluxes of $10^{-23} \leq F_{LW} \leq 10^{-22}$ has a central core temperature of approximately 1000 K. Simulations with $F_{LW} > 10^{-22}$ do not collapse by the end of the simulation and have significantly higher overall temperatures. The reason for this can be seen by examination of Figure 4.26, which plots the molecular hydrogen fraction, H^- fraction, and electron fraction of the baryon gas as a function of radius. The overall H_2 fraction in simulations with $F_{LW} \leq 10^{-24} \text{ ergs s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1}$ is a factor of a few higher than the population with $10^{-23} \leq F_{LW} \leq 10^{-22}$, and simulations with a higher UV flux have very little H_2 overall. This trend is due to the photodissociation of molecular hydrogen by the soft UV background, which delays collapse of the halos by affecting cooling rates. Once the center of a halo reaches some critical density (which is a function of F_{LW}), H_2 formation can continue as before. However, at lower densities the H_2 fraction reaches some equilibrium with the UV background, where the rates of photodissociation and creation of H_2 are equal. This can be clearly seen outside of $r \sim 0.1$ parsec in the plot of H_2 fraction as a function of radius, where there is a monotonic decrease in the H_2 fraction outside the core with increasing UV flux.

As one might expect, the spread in halo core central temperatures produces somewhat different end results. Figure 4.27 shows the radial and circular baryon velocities as a function of radius, and the accretion time as a function of enclosed mass. Predictably, simulations with a lower core temperature tend to have a lower overall infall velocity of gas onto the central protostar, though there is some scatter in the result. This is most likely




Figure 4.24: Mean halo quantities for several simulations with the same cosmic realization but a range of Lyman-Werner molecular hydrogen photodissociating flux backgrounds. Top left: Lyman-Werner flux vs. halo collapse redshift. Top right: Lyman-Werner flux vs. halo virial mass at collapse. Bottom: halo virial mass vs. collapse redshift. Simulations with values for the soft UV background of $F_{LW} = 10^{-21}$ and 10^{-20} did not collapse and are not shown. In the top two panels the collapse redshift/virial mass of the $F_{LW} = 0$ “control” result are shown as vertical blue dashed lines. In the top right panel the green dashed line corresponds to the fitting function for threshold mass from Machacek et al. (eqtn. 8), and the red dashed line corresponds to a simple power law, $M(F_{LW}) = 4.35 \times 10^5 (F_{LW}/10^{-24})^{0.27} M_{\odot}$.




Figure 4.25: Mass-weighted, spherically-averaged baryon quantities as a function of radius for 7 simulations of the same cosmological realization and a range of soft UV backgrounds. Top left: number density as a function of radius. Top right: temperature as a function of radius. Bottom: enclosed mass as a function of radius. In all panels, the black long-dashed line corresponds to $F_{LW} = 0$, the black solid line to $F_{LW} = 10^{-24}$, the blue solid line to $F_{LW} = 10^{-23}$, the green solid line to $F_{LW} = 3 \times 10^{-23}$, the red solid line to $F_{LW} = 10^{-22}$, the black short-dashed line to $F_{LW} = 10^{-21}$, and the blue short-dashed line to $F_{LW} = 10^{-20}$. All simulations with $F_{LW} \leq 10^{-22}$ are shown at their redshift of collapse, while simulations with values for the soft UV background of $F_{LW} = 10^{-21}$ and 10^{-20} did not collapse and are shown at the last available redshift.




Figure 4.26: Mass-weighted, spherically-averaged baryon quantities as a function of radius for 7 simulations with the same cosmological realization and a range of soft UV backgrounds. Top left: molecular hydrogen fraction as a function of radius. Top right: H^- fraction as a function of radius. Bottom: electron fraction as a function of radius. In all panels, the black long-dashed line corresponds to $F_{LW} = 0$, the black solid line to $F_{LW} = 10^{-24}$, the blue solid line to $F_{LW} = 10^{-23}$, the green solid line to $F_{LW} = 3 \times 10^{-23}$, the red solid line to $F_{LW} = 10^{-22}$, the black short-dashed line to $F_{LW} = 10^{-21}$, and the blue short-dashed line to $F_{LW} = 10^{-20}$. All simulations with $F_{LW} \leq 10^{-22}$ are shown at their redshift of collapse, while simulations with values for the soft UV background of $F_{LW} = 10^{-21}$ and 10^{-20} did not collapse and are shown at the last available redshift.

due to the fact that the halos in these calculations do not exist in isolation – during the significant amount of time that the collapse of the halo core is delayed, structure formation is still taking place and the overall halo mass is being substantially increased by accretion, which may have some nonlinear effects on the halo temperature, molecular hydrogen fraction, and as a result the accretion rate onto the protostar. Regardless, the final estimate for the mass of the star, using the same criteria as in Section 4.4.1, suggests that the overall spread of final stellar masses is significantly enlarged by the addition of a soft UV background. Halos with larger SUV fluxes tend to have higher temperatures, and thus higher accretion rates and possibly higher stellar masses.

Figures 4.28 through 4.30 show two-dimensional, mass weighted distribution functions of several quantities which illustrate the overall effects of the photodissociating flux. All panels are shown at their redshift of collapse or (for those that do not collapse) the final output of the simulation at $z = 10$. Each panel has 10 contours spaced equally in logarithmic intervals between the lowest and highest values in the panel.

Figure 4.28 shows the two-dimensional distribution function of baryon temperature vs. baryon overdensity (δ) for all of the simulations discussed in this section, and Figure 4.29 shows the two-dimensional distribution of molecular hydrogen fraction as a function of overdensity. These two plots illustrate the important density thresholds related to the chemistry and cooling properties of molecular hydrogen. The “knee” seen in the temperature-overdensity plot at $\delta \sim 100$ corresponds to a critical density at which the molecular hydrogen formation time scale becomes less than the Hubble time. This is due to the molecular hydrogen formation rate increasing with density. There is a corresponding “knee” in the H_2 - baryon overdensity distribution function at this overdensity. At overdensities between 10^2 and 10^5 the temperature decreases with increasing density, and the molecular hydrogen fraction continues to grow. At an overdensity of approximately 10^5 the ro-vibrational levels of H_2 are populated at their equilibrium value and the cooling rate becomes independent of density, which corresponds to an increase in gas temperature with increasing density. Finally, at number densities of $n \sim 10^9 - 10^{10} \text{ cm}^{-3}$ (overdensities of $\sim 10^{11}$) the three-body molecular hydrogen formation becomes dominant and H_2 is formed very rapidly with increasing density. This can be seen as another “knee” in the $H_2 - \delta$ distribution function at overdensities of $\sim 10^{11}$, though the cooling properties of the gas still remains density independent. Simulations with a Lyman-Werner background flux that is high enough to completely suppress the formation of molecular hydrogen ($F_{LW} \geq 10^{21}$) cannot cool efficiently and therefore cannot collapse to overdensities higher than $\sim 10^3$.

From the standpoint of Population III star formation, the practical effects of the cooling properties of H_2 can be summed up in Figure 4.30, which is the two-dimensional mass-weighted distribution function of the Jeans mass (which scales as $\sim T^{3/2}/\rho^{1/2}$) versus overdensity. The Jeans mass (or, more precisely, the Bonnor-Ebert critical mass, which differs from the Jeans mass by a numerical constant) controls the mass scale at




Figure 4.27: Mass-weighted baryon quantities for 7 simulations with the same cosmological realization and a range of soft UV backgrounds. Top left: spherically-averaged radial velocity as a function of radius. Top right: cylindrically-averaged circular velocity as a function of radius. Bottom: Accretion time (defined as in Figure 4.23) as a function of enclosed mass. In all panels, the black long-dashed line corresponds to $F_{\text{LW}} = 0$, the black solid line to $F_{\text{LW}} = 10^{-24}$, the blue solid line to $F_{\text{LW}} = 10^{-23}$, the green solid line to $F_{\text{LW}} = 3 \times 10^{-23}$, the red solid line to $F_{\text{LW}} = 10^{-22}$, the black short-dashed line to $F_{\text{LW}} = 10^{-21}$, and the blue short-dashed line to $F_{\text{LW}} = 10^{-20}$. All simulations with $F_{\text{LW}} \leq 10^{-22}$ are shown at their redshift of collapse, while simulations with values for the soft UV background of $F_{\text{LW}} = 10^{-21}$ and 10^{-20} did not collapse and are shown at the last available redshift.

which fragmentation of the gas via gravitational instability can occur. At low densities, the overall temperature is high and therefore the Jeans mass is quite high – at least $10^5 M_\odot$ for $\delta \sim 10^2 - 10^3$. However, at higher densities the temperature drops rapidly up to $\delta \sim 10^5$, resulting in a corresponding rapid drop in the Jeans mass. At these density scales the Jeans mass is still too high for efficient fragmentation. Once the cooling rate becomes independent of density the temperature begins to climb as a function of overdensity, but at a fairly low rate, so the Jeans mass continues to drop, though not as rapidly. Finally, at an overdensity of $\sim 10^{11}$ (after the 3-body process begins to dominate) the Jeans mass drops below $100 M_\odot$, which is roughly equivalent to the total amount of gas that has reached that overdensity. In principle one might suppose that the gas in the halo core could begin to fragment at this point, but no evidence of fragmentation has been found in any of the simulations that we have examined.

4.5 Discussion

In this chapter we have explored several aspects of the formation of Population III stars in a Λ CDM universe. This section summarizes some of the processes neglected in our calculations and also attempts to put some of the results in context.

The results presented in Section 4.4.3 demonstrate that there is a great deal of scatter between the bulk halo properties such as overall virial mass, collapse redshift, and mean halo temperature among the twelve simulations shown. However, the final state of the density profile is extremely similar between all of the calculations. This is entirely due to the chemical and cooling properties of the primordial gas – the minimum temperature of the gas (which is determined by its chemical composition) creates a density profile that goes as r^{-2} for any gas cloud which is only supported by thermal pressure. This seems to be true for the gas contained in the halos out of which Population III stars form, so it is reasonable to expect consistent density profiles on halo scales.

Though there is consistency in the bulk halo properties, a detailed examination of the gas properties which may contribute significantly to the final Pop III star mass, such as the core baryon temperature and accretion rate onto the forming primordial protostar, show a tremendous amount of scatter. This scatter appears to be due to variations in the molecular hydrogen content of the halo on large scales, which is brought on by differences in halo temperature as a result of varied merger rates between simulations. There appears to be a systematic effect between the simulation box size and the mean temperature, with larger boxes (which have more large-scale power and overall a more rapid merger history) having higher overall halo temperatures and lower accretion rates. The higher temperatures result in somewhat larger molecular hydrogen mass fractions, which cause the halo core to cool more rapidly during its eventual collapse. Since the accretion onto the primordial protostar is primarily subsonic, the accretion rate depends on the sound




Figure 4.28: Mass-weighted two-dimensional distribution functions of baryon temperature vs. baryon overdensity for seven simulations with the same cosmological realization and a range of soft UV backgrounds. The strength of the UV background is marked in each panel. All simulations with $F_{LW} \leq 10^{-22}$ are shown at their redshift of collapse, while simulations with values for the soft UV background of $F_{LW} = 10^{-21}$ and 10^{-20} did not collapse and are shown at the last available redshift. Each panel has 10 contours equally spaced in logarithmic intervals between the maximum and minimum values in that panel.




Figure 4.29: Mass-weighted two-dimensional distribution functions of molecular hydrogen fraction vs. baryon overdensity for seven simulations with the same cosmological realization and a range of soft UV backgrounds. The strength of the UV background is marked in each panel. All simulations with $F_{LW} \leq 10^{-22}$ are shown at their redshift of collapse, while simulations with values for the soft UV background of $F_{LW} = 10^{-21}$ and 10^{-20} did not collapse and are shown at the last available redshift. Each panel has 10 contours equally spaced in logarithmic intervals between the maximum and minimum values in that panel.




Figure 4.30: Mass-weighted two-dimensional distribution functions of cell Jeans mass vs. baryon overdensity for seven simulations with the same cosmological realization and a range of soft UV backgrounds. The strength of the UV background is marked in each panel. All simulations with $F_{LW} \leq 10^{-22}$ are shown at their redshift of collapse, while simulations with values for the soft UV background of $F_{LW} = 10^{-21}$ and 10^{-20} did not collapse and are shown at the last available redshift. Each panel has 10 contours equally spaced in logarithmic intervals between the maximum and minimum values in that panel.

speed cubed, with lower core temperatures directly resulting in lower accretion rates.

After the onset of collapse, the evolution of the core of the halo (roughly the inner few thousand solar masses) becomes effectively decoupled from the halo envelope since the time scales become much shorter within the halo core. This tells us that while the formation of the initial primordial protostellar cloud is strongly coupled to the time scales associated with cosmological structure formation, once the cloud has collapsed we can treat the core of the halo separately from the rest of the calculation. This decoupling will become highly useful when more detailed calculations of the evolution of Population III protostars, including more complicated physics such as radiative transfer and protostellar accretion models, are performed, and will save us significant computational cost.

The observation that the rate of accretion onto the primordial protostar varies systematically as a function of box size, with larger box sizes having an overall lower accretion rate, has significant implications for both reionization and metal enrichment of the early universe. The accretion rate results cannot be proven to be converged yet, due to noise and small number statistics, though the 0.45 and $0.6 \text{ h}^{-1} \text{ Mpc}$ boxes seem to have overall similar accretion rates. We make an estimate of the minimum possible accretion rate by observing that molecular hydrogen is only effective at cooling the primordial gas down to approximately 200 Kelvin, which gives us an accretion rate of $\dot{M}_* \simeq 5 \times 10^{-4} \text{ M}_\odot/\text{year}$, which is reasonably close to the lower envelope of accretion rates observed in the $0.6 \text{ h}^{-1} \text{ Mpc}$ box calculations. Though this implies convergence, it would be prudent to perform another suite of calculations at an even larger box size to be sure.

If in fact a lower overall accretion rate results in a less massive population of stars, these objects will be much less effective at ionizing the intergalactic medium (since they produce overall fewer UV photons per baryon) and will produce a completely different nucleosynthetic signature. This is important because the measurement of the polarization of the cosmic microwave background by the WMAP satellite implies early reionization, which possibly implies a significant contribution from extremely massive Population III stars, whereas observations of ultra metal poor stars in the galactic halo see abundance ratios that do not agree with numerical predictions for the abundance ratios of extremely massive primordial stars. At this point it is difficult to say what accretion rate is most common during Population III star formation. In particular, once a soft ultraviolet background begins to build up and cause the photodissociation of molecular hydrogen the overall halo core temperatures may rise, causing an increase in the baryon accretion rate. Wise & Abel [97] use Press-Schechter models of Population III star formation to predict a slowly rising Lyman-Werner background which provides some support to this idea. This suggests that further calculations including larger simulation volumes as well as a soft UV background will be necessary to make a definitive statement about the most common accretion rates. Additionally, these calculations completely neglect the mode of primordial star formation that takes place in halos whose virial temperatures are above

10^4 K. Cooling in these systems is dominated by atomic hydrogen line emission and, particularly in the presence of a strong soft UV background, may result in a much larger amount of cold gas distributed in a different manner than in the systems simulated in this thesis, which have mean virial masses of a few times $10^5 M_\odot$ and virial temperatures of around 1000 K.

In Section 4.4.2 we examined angular momentum transport in the collapsing halo core. This appears to be a robust result, and the primary mode of angular momentum redistribution appears to be due to turbulence. Unfortunately it is somewhat difficult to analyze this effect using the standard analytical formalism for angular momentum transport in accretion disks since the collapsing halo core is approximately spherical. It seems that a more complicated method of analyzing the turbulence properties within the halo core, such as structure functions, may be in order. Additionally, the efficiency of turbulent angular momentum transport in a quasistatically collapsing spheroid has not been studied analytically or numerically, and a better physical understanding of the situation may result from developing new analytics and idealized numerical test problems.

It is not completely clear what drives the turbulence in the halo core. There is accretion of gas into the core of the halo, and this provides a possible mechanism for driving turbulence. Also, during the quasistatic collapse of the halo core the gas is subject to significant cooling, which may drive turbulence via thermal instabilities. This seems less likely, because the sound speed is comparable to the speed of collapse, which serves to smooth the overall density perturbation. More investigation is necessary, perhaps using idealized numerical simulations, to understand the precise mechanism for driving the turbulence seen in the collapsing halo core.

Unlike galactic star formation at the present epoch, the collapsing cosmological halo core has little angular momentum from the outset – the “angular momentum problem” that plagues present-day star formation simply doesn’t appear to be an issue in the Population III star formation scenario. During the quasi-static collapse of the halo core the gas is never rotationally supported – the gas is essentially completely held up by thermal pressure. It appears that the small amount of angular momentum that is actually transported is not a critical factor in the cloud core’s collapse.

As discussed in Section 4.4.2, the angular momentum transport result does not appear to be due to numerical effects. Unfortunately, it is particularly difficult to quantify numerical viscosity due to finite resolution in an adaptive mesh refinement code, since the overall grid resolution can vary tremendously. Essentially the only reliable way to quantify numerical viscosity in an AMR code is to pose a problem with an analytical solution that explicitly includes a physical viscosity and then simulate it with adaptive mesh code using the Eulerian equations for fluid flow. Since these equations assume zero physical viscosity, any viscous effects observed are completely numerical, and the analytical problem can be used to extract a value for the numerical viscosity. This is a challenging task, and the measured numerical viscosity is most likely dependent on the

details of the test problem and can vary strongly depending on the physical system being modeled.

The primordial chemistry model used in these calculations ignores the effects of deuterium, lithium, and the various molecules they form between these elements and ordinary hydrogen. Deuterium and lithium have been shown to be unimportant in the temperature and density regimes that we have examined in this chapter. However, it is possible that they may be relevant in other situations of importance to Population III star formation – in particular, regions which have been ionized to very high electron fractions may experience significant cooling from the HD molecule, which due to its permanent dipole moment makes it more than 100 times more effective as a cooling agent than molecular hydrogen (per molecule), and has the potential to cool gas down to approximately the temperature of the cosmic microwave background, which scales with redshift as $T_{cmb}(z) = 2.73 \times (1 + z)$ K. This gives a minimum baryon temperature of approximately 55 Kelvin at $z = 20$ and could further reduce the minimum accretion rate onto a primordial protostar by a factor of two, to $\dot{M}_{min} \simeq 2.5 \times 10^{-4} M_{\odot}/\text{year}$.

The effects of magnetic fields are completely ignored in the simulations discussed in this thesis. We can justify this by examining the magnetic field necessary to delay the collapse of the halo core. If one assumes that the halo core can be represented reasonably well by an isothermal sphere of constant density (which is reasonable at the onset of halo collapse), we can use the virial theorem to estimate the strength of the magnetic field which is necessary to support the collapse of the halo against gravity. Assuming flux freezing and a uniform magnetic field, a magnetically critical isothermal sphere has a mass-to-flux ratio of

$$\frac{M_{cl}}{\Phi_B} = \frac{1}{\sqrt{31G}} \quad (4.3)$$

Where M_{cl} is the mass of the halo, $\Phi_b = \pi R_{cl}^2 B_{cl}$ is the magnetic flux in the cloud (with R_{cl} and B_{cl} being the cloud radius and magnetic field strength, respectively), and G is the gravitational constant. Reasonable values for M_{cl} and R_{cl} are $\simeq 2 \times 10^3 M_{\odot}$ and 4 parsecs, respectively, which gives us a value of the magnetic field of $B_{cl} = 1.21 \times 10^{-5}$ G. The mean density of the cloud is $n_{cl} \simeq 300 \text{ cm}^{-3}$ and the mean density of the universe at $z = 18$ (the redshift that our cloud collapses) is $\simeq 0.003 \text{ cm}^{-3}$, so if we assume a spherical collapse from the mean cosmic density assuming flux freezing, we see that the ratio of the magnetic field in the cloud to the mean universal magnetic field is

$$\frac{B_{cl}}{B_{igm}} = \left(\frac{n_{cl}}{n_{igm}} \right)^{2/3} \quad (4.4)$$

This gives us a mean magnetic field of $B_{IGM} \simeq 3.5 \times 10^{-9}$ G at $z \simeq 18$. Since there are no known objects that may produce magnetic fields between recombination

($z \sim 1100$) and the epoch of Pop III star formation, and the magnetic field scales with the expansion of the universe as $(1+z)^2$, we estimate that in order for magnetic fields to be dynamically important in the formation of Population III stars the magnetic field strength at recombination must be $B_{rec} \sim 10^{-5}$ G. The current observational upper limit to magnetic field strength at recombination (albeit at large scales) is $B \leq 3 \times 10^{-8}$ G as measured at the present epoch [29], which corresponds to a magnetic field at recombination of approximately 4×10^{-2} G. This is three orders of magnitude higher than needed to be dynamically relevant for Population III star formation! However, there are no known mechanisms that can produce a magnetic field of that magnitude that have not been ruled out due to other observational limitations. Currently, the most plausible mechanisms for creating magnetic fields at recombination suggest that field strengths are on the order of 10^{-23} G at recombination [33]. Given the observational uncertainty, it seems reasonable to ignore this effect, though future simulations will certainly include magnetic fields with a variety of field strengths and physical scales..

Assuming that the magnetic field at that epoch was strong enough to be dynamically important, we can calculate the effect that this has on the collapse of the star. Due to the low electron fraction in the halo core (which has electron densities comparable to or lower than that observed in present-day molecular cloud cores), the assumption of flux freezing in the magnetic field is not valid. Magnetic fields couple to charged particles (electrons and ions), and these charged particles interact with the neutral medium. At high levels of ionization collisions between charged and neutral particles are frequent, implying that the magnetic field is strongly coupled to the gas. However, at low levels of ionization there are few charged particles, and the coupling with the neutral gas is weak. In an object that is subject to a gravitational acceleration this produces a relative drift of charged and neutral particles which allows the neutral gas to decouple from the magnetic field. This effect is known as “ambipolar diffusion,” and is believed to be an extremely important process in galactic star formation. The retardation effect that ambipolar diffusion may have on the collapse of the halo core can be estimated by examining the relative timescales of ambipolar diffusion and halo collapse. The ambipolar diffusion timescale can be estimated as

$$\tau_{AD} = \frac{L}{v_D} \simeq 2 \times 10^6 \frac{x_i}{10^{-7}} \text{ years} \quad (4.5)$$

Where L and v_D are a characteristic length scale and the neutral-ion relative drift velocity, respectively, and x_i is the overall ionization fraction. A proxy for the halo collapse time scale is the free fall time, which for a spherical system is

$$\tau_{ff} = \left(\frac{3\pi}{32G\rho} \right)^{1/2} \simeq \frac{5 \times 10^7}{n^{1/2}} \text{ years} \quad (4.6)$$

where n is the particle number density in particles per cubic centimeter and G is the gravitational constant. The relevance of ambipolar diffusion can be estimated by taking the ratio of these two quantities, which is known as the “collapse retardation time,” ν_{ff} . Substituting in equations 4.5 and 4.6, we see that

$$\nu_{ff} \equiv \frac{\tau_{AD}}{\tau_{ff}} \simeq 4 \times 10^5 x_i n^{1/2} \quad (4.7)$$

Examination of figures 4.4 and 4.5 show that at the final timestep in the calculation, the number density can be fitted by a power law and is roughly $n(r) \simeq 10^3 (r/pc)^{-2} \text{ cm}^{-3}$ while the ionization fraction scales roughly as $x_i(r) \simeq 10^{-6} (r/pc)$. Plugging these into equation 4.7 shows that $\nu_{ff} \simeq 13$ is constant with radius. This is only a crude approximation, since the free fall time really should depend on the mean number density instead of the number density at a given radius. However, considering the rapid falloff of density, $n(r)$ is a reasonable approximation of \bar{n} – strictly speaking, for a cloud with a density profile that scales as r^{-2} over many orders of magnitude in radius, the mean density is equal to $3 \bar{n}$, so our estimate of the free fall time is too high by a factor of $\sqrt{3}$. Plugging this in to the equation, we get that $\nu_{ff} \simeq 23$ everywhere, which indicates significant delay in collapse with respect to the free fall time. However, the relevant time scale in this case is more appropriately the quasistatic collapse time, which is approximated as $\tau_{qs} \simeq L/v_r$. Figure 4.6 shows that the mean radial velocity at the scales of interest ($\sim 2 - 3$ parsecs) is roughly 0.5 km/s. This corresponds to $\tau_{qs} \sim 4 \times 10^6$ and scales linearly with the radius. Comparison with the ambipolar diffusion time scale shows that τ_{AD} and τ_{qs} are within a factor of two of each other, which suggests that the presence of a magnetic field would not significantly impede the collapse of the halo core for the quasistatic collapse case.

Section 4.4.4 discusses the formation of Population III stars in simulations with a constant soft UV background. Our results agree well qualitatively with that of Machacek, Bryan & Abel [71] – we both find that a soft UV background can delay the cooling and collapse of small ($\sim 10^6 M_\odot$) halos, and that increasing the soft UV background increases the minimum halo mass required for a halo to collapse. Machacek et al. derived a mass threshold for collapse as a function of the Lyman-Werner background flux that agrees well with our simulations, though the halo masses in our calculation are significantly higher. This is due to the halo that we examine being an average halo rather than at the plausible halo minimum mass. If we perform these calculations for large number of stars in a range of cosmological realizations it seems reasonable that the results from our lower-mass halos would agree more completely with their work. Future large volume, high-resolution calculations will further test the results of Machacek et al. at large boxes and for more halos. Our work is an improvement upon that of Machacek et al. because our simulations are much more highly resolved and we examine the evolution of a single halo over a much wider range of soft UV background fluxes.

In the calculations that we performed using a constant soft UV background we completely ignore self-shielding by molecular hydrogen. Though this effect could in principle

be important, the actual column densities of molecular hydrogen are typically far too small to actually block the soft UV flux. According to Machacek et al. (and references therein), a column density of $5 \times 10^{14} \text{ cm}^{-2}$ is enough for shielding to become important. However, this was derived for a static distribution of H_2 , while the Lyman-Werner band consists of hundreds of individual lines whose width in this case is dominated by Doppler broadening. It is useful to note in this case that the average line width is $\sim 2 \text{ km/s}$ and the RMS baryon velocity in our calculations are $\sim 4 \text{ km/s}$. In order for self-shielding to be important in the case of a turbulent medium the column density must be much higher. Typical maximum H_2 column densities in our calculations are on the order of $10^{16-17} \text{ cm}^{-2}$, but these occur late in the collapse of the core, and in the highest density regions the cooling and H_2 production times are much shorter than the photodissociation time scale, at which point self-shielding becomes unimportant.

As discussed previously, the simulation volumes used by our soft UV background calculations are rather small. The results of this is that the largest halo in our calculation has a virial temperature which is significantly below 10^4 K , which is when cooling via atomic hydrogen lines becomes an effective means of cooling halos. It is possible to produce halos that are this large without them having undergone previous epochs of star formation by having a very strong soft UV background. When the center of this type of halo cools via atomic lines to a large enough density, rapid molecular hydrogen formation will take place essentially independent of the strength of the soft UV background and the gas will be able to cool down to $\simeq 200 \text{ K}$ very quickly. Since these halos are more than an order of magnitude larger than those that we have considered in this project, presumably the reservoir of cold, dense gas in the center of the halo will also be correspondingly large. In this case, the amount of cold, dense gas will almost certainly exceed the Jeans mass by large factors, which in principle would make it possible for multiple Population III stars to form in a single halo, as opposed to the single star that we see forming at the center of halos in the simulations analyzed in this work. Further work will be necessary to understand how these larger halos form stars, and the potential IMF of the stars forming in them.

Chapter 5

Formation of primordial stars in a WDM universe

5.1 Summary

In this chapter I discuss aspects of primordial star formation in a universe with a generic warm dark matter (WDM) cosmology. After describing the theory and general effects of a warm dark matter power spectrum, I compare the results of simulations done with a single cosmological realization but using a wide range of warm dark matter particle masses which have not yet been ruled out by observation. The main effect of the WDM is that the addition of a warm dark matter component to the initial power spectrum results in a delay in the collapse of the gas at the center of the halo which forms a primordial protostar and an increase in the virial mass of the halo at the onset of collapse. Both of these effects become more pronounced as the WDM particle mass becomes smaller. A cosmology using a warm dark matter spectrum assuming a particle mass of $m_{WDM} \simeq 40$ keV is effectively indistinguishable from the cold dark matter case, and a reasonable lower limit to a warm dark matter particle mass of $\simeq 15$ keV is suggested in order for Population III stars to contribute significantly to the polarization result observed by the WMAP satellite. There is remarkably little scatter in the final properties of the primordial protostar which forms at the center of the halo, possibly due to the overall low rate of halo mergers which is a result of the WDM power spectrum. I also describe the detailed evolution of the collapsing halo core in two representative WDM cosmologies. Once the gas in the center of the halo reaches relatively high densities ($n \sim 10^5$ cm⁻³), the overall evolution is essentially identical in the two calculations.

5.2 Motivation and theory

As discussed in Section 1.5, there are some apparent flaws in the Λ CDM paradigm at small scales. These include a lack of observation of dark matter “cusps,” which are predicted by theory and numerical simulations in the CDM universe, the absence of the very large number of dwarf galaxies predicted by the cold dark matter paradigm, and others. However, at large scales (greater than 1 megaparsec or so) the CDM model seems to describe the evolution of the universe and the structure within it incredibly well. The problem, therefore, is to come up with some physical explanation for the apparent lack of power on small scales while retaining the desirable qualities of the CDM model on large scales.

Many models have been proposed that do this, with a wide variety of mechanisms. A generic feature of these models is that they suppress the cosmological power spectrum on small scales, while leaving large scales alone. In this chapter I will discuss the ramifications of the most general possible model, referred to as “warm dark matter.” The effects of a general warm dark matter cosmology are discussed by Bode, Ostriker & Turok [115], who derive the relevant linear perturbation theory and perform several N-body calculations of warm dark matter cosmologies to understand the general effects of suppression of power on small scales. They find that replacing cold dark matter with warm dark matter results in the smoothing of massive halo cores, which lowers core densities and increases core radii, lowers the characteristic density of low-mass halos, reduces the overall total number of low-mass halos, suppresses the number of low-mass satellite halos in high-mass halos, and results in the formation of low-mass halos almost entirely within caustic sheets or filaments connecting larger halos – voids are almost completely empty, in contrast to CDM. They also find that low-mass halos tend to form at late times, in a top-down process (as opposed to the bottom-up process of halo formation one would expect from a CDM cosmology), and that halo formation is suppressed overall at early times (high redshift), with an increased evolution of halos at low redshifts relative to the CDM model. Furthermore, they suggest that a reasonable minimum warm dark matter particle mass would be 1 keV.

Other constraints on the minimum mass of a warm dark matter particle have been placed by various groups. Barkana, Haiman & Ostriker [212] use an extended Press-Schechter model to constrain warm dark matter based on constraints from cosmological reionization. They calculate that in order for super massive black holes to exist at $z \simeq 6$, and if massive galaxies are responsible for the nearly complete reionization of the universe by the same redshift, a reasonable minimum mass for a warm dark matter particle is $m_{WDM} \geq 1.2$ keV. Dalal & Kochanek [213] show that constraints on small-scale linear power using strong gravitational lensing find agreement with agreement with Λ CDM models, and place a minimum bound on a possible warm dark matter mass of $m_{WDM} > 5.2$ keV. Additionally, Yoshida et al. perform a SPH cosmological simulation of

structure formation in the early universe assuming a WDM particle mass of 10 keV. They find that this calculation is inconsistent with the high Thomson optical depth observed by the WMAP satellite, and suggest that any successful WDM model will have a particle mass greater than 10 keV [214].

Bode et al. derive a formula for the power spectrum cutoff due to the existence of a generic warm dark matter particle. They provide the following transfer function that models the smoothing of small-scale density perturbations:

$$T_k^X = [1 + (\alpha k)^2]^{-5} \quad (5.1)$$

Where $\alpha = 0.05(\Omega_x/0.4)^{0.15}(h/0.65)^{1.3}(keV/m_X)^{1.15}(1.5/g_x)^{0.29}$ and k is in units of $h \text{ Mpc}^{-1}$. In this equation Ω_X is the contribution of the warm dark matter species to the energy density of the universe, in units of the critical density, m_X is the WDM particle mass in keV, h is the Hubble constant in units of 100 km/s/Mpc, and g_x is a parameter meant to represent the effective number of relativistic species present at decoupling, and is taken to be 1.5 for light neutrinos. This is considered to be the fiducial value for g_x for the warm dark matter particle. This transfer function corresponds to a strong rollover in the power spectrum, corresponding to a smoothing scale of:

$$R_s \simeq 0.31 \left(\frac{\Omega_x}{0.3}\right)^{0.15} \left(\frac{h}{0.65}\right)^{1.3} \left(\frac{keV}{m_X}\right)^{1.15} h^{-1} \text{Mpc} \quad (5.2)$$

This smoothing scale corresponds to the comoving half-wavelength of the mode for which the linear perturbation amplitude is suppressed by a factor of 2. This results in a characteristic mass scale below which structure forms by the top-down fragmentation of halos, rather than by the bottom-up hierarchical structure formation associated with the cold dark matter paradigm. This mass scale can be quantified as:

$$M_s = 10^{10} \left(\frac{\Omega_x}{0.3}\right)^{1.45} \left(\frac{h}{0.65}\right)^{3.9} \left(\frac{keV}{m_X}\right)^{3.45} h^{-1} M_\odot \quad (5.3)$$

Figure 5.1 contains several panels demonstrating the effects of the warm dark matter cosmology discussed above. The top left panel shows the cosmological power spectrum $P(k)$ at $z = 0$ for a CDM cosmology as well as for when the WDM transfer function has been applied for several different warm dark matter masses ranging from 0.1 – 100 keV. The top right panel shows the dimensionless linear power $\Delta^2(k) \sim k^3 P(k)$ with the same particle masses. The bottom left panel shows the suppression mass as a function of radius, and the bottom right panel shows the comoving smoothing scale. In the bottom two panels the red dashed line indicates the mass and radius corresponding to a halo of mass $4 \times 10^5 M_\odot$, which is approximately the mean halo mass of all of the CDM simulations discussed in Chapter 4. This mass scale corresponds to a WDM particle mass of ~ 15 keV.




Figure 5.1: Plots showing various effects of a generic warm dark matter particle. Top left panel: The dark matter power spectrum $P(k)$ as a function of wavenumber $k \equiv 2\pi/\lambda$ for a CDM cosmology and a range of warm dark matter masses. Top right panel: The dimensionless power, $\Delta^2(k) \sim k^3 P(k)$, versus k . Bottom left panel: The “suppression mass” as a function of WDM particle mass. Bottom right panel: The “smoothing scale” as a function of WDM particle mass. In the bottom two plots the red dashed line indicates the mass and radius corresponding to a halo of mass $4 \times 10^5 M_{\odot}$, which is approximately the mean halo mass from the simulations discussed in Chapter 4.

5.3 Problem setup

The simulations discussed in this chapter use a similar setup to those in Chapter 4. A single cosmological realization in a box size of $0.3 \text{ h}^{-1} \text{ Mpc}$ is chosen at random from the four available CDM calculations. This calculation is initialized at $z = 99$ assuming a “concordance” cosmological model: $\Omega_m = 0.3$, $\Omega_b = 0.04$, $\Omega_{DM} = 0.26$, $\Omega_\Lambda = 0.7$, $h = 0.7$ (in units of 100 km/s/Mpc), $\sigma_8 = 0.9$, and using an Eisenstein & Hu power spectrum [194] with a spectral index of $n = 1$. At this point, we generate several sets of initial conditions with the same large-scale structure by smoothing the CDM initial conditions with the warm dark matter transfer function described in Equation 5.1, assuming $\Omega_x = \Omega_{DM} = 0.26$, $g_x = 1.5$, and warm dark matter masses of $m_x = 10, 12.5, 15, 17.5, 20, 25, 30, 35,$ and 40 keV . The initial conditions are generated with both dark matter and baryons such that the Lagrangian volume in which the halo in the CDM case formed is resolved at high spatial and mass resolution using the nested initial condition method described in Section 2.1.2, with a 128^3 root grid and three static nested grids, for an overall effective grid size of 1024^3 . The highest resolution grid is 256^3 grid cells, and corresponds to a volume 75 h^{-1} comoving kpc on a side. The dark matter particles in the highest resolution grid are $1.81 \text{ h}^{-1} M_\odot$ and the spatial resolution of the highest resolution grid is 293 h^{-1} parsecs (comoving). Previous work shows that this particle mass resolution is more than adequate to fully resolve the collapse of the halo.

All simulations are performed using the adaptive mesh cosmology code Enzo, described in detail in Section 2.2. The simulations are started at $z = 99$ and allowed to evolve until the collapse of the gas within the center of the most massive halo, which occurs at a range of redshifts. The equations of hydrodynamics are solved using the PPM method with a dual energy formulation, as described in Section 2.2.2. The nonequilibrium chemical evolution and optically thin radiative cooling of the primordial gas is modeled as described in Section 2.2.5, following 9 separate species including molecular hydrogen (but excluding deuterium). Adaptive mesh refinement is used such that cells are refined by factors of two along each axis, with a maximum of 22 total levels of refinement. This corresponds to a maximum spatial resolution of 115 h^{-1} astronomical units (comoving) at the finest level of resolution, with an overall spatial dynamical range of 5.37×10^8 . To avoid effects due to the finite size of the dark matter particles, the dark matter density is smoothed on a comoving scale of $\sim 0.5 \text{ pc}$. This is reasonable because at that radius in all of our calculations the gravitational potential is dominated by the baryons.

Grid cells are adaptively refined based upon several criteria: baryon and dark matter overdensities in cells of 4.0 and 8.0, respectively, as well as criteria to ensure that the pressure jump and/or energy ratios between adjoining cells never exceeds 5.0, that the cooling time in a given cell is always longer than the sound crossing time of that cell, and that the Jeans length is always resolved by at least 16 cells. This guarantees that the Truelove criterion [217] is always maintained by a comfortable margin.

5.4 Results

5.4.1 Comparison of realizations with many WDM particle masses

In this section we discuss the results of a comparison of all of the WDM simulations, along with the cold dark matter “control” simulation. Figure 5.2 shows bulk properties of the halo in which the Population III protostar forms as a function of the warm dark matter particle mass. The top left and right panels plot the WDM particle mass versus the redshift at which the halo core collapses. The CDM result in each case is shown by a vertical blue line. These panels demonstrate that decreasing the warm dark matter particle mass delays the formation of the protostar. The 40 keV calculation forms at essentially the same time as the CDM simulation, while collapse of the halo core in the calculation assuming a 12.5 keV WDM particle mass is delayed by approximately 130 million years. The simulation with a 10 keV particle mass does not collapse by $z = 10$ (the end of the simulation) and is not shown here. The delay of the halo collapse appears to be smoothly varying as a function of WDM particle mass.

The bottom panel of Figure 5.2 shows the virial mass of the halo (at the redshift of collapse) as a function of WDM particle mass. A reduction in the WDM particle mass leads to an increase in the halo virial mass, which is related to the delay in collapse of the halo core – by the time the halo core collapses in the lower particle mass simulations, the halo has had time to accrete more mass.

Figures 5.3 and 5.4 show several spherically-averaged, mass-weighted radial profiles of baryon quantities as a function of radius or enclosed mass of the simulations. All profiles are taken at a constant point in the evolution of the protostellar cloud (when the central number density is $n \sim 10^{10} \text{ cm}^{-3}$) rather than at a constant point in time, since the halos collapse over a wide range of redshifts. Figure 5.3 shows the baryon number density as a function of enclosed mass (top left panel), baryon temperature as a function of enclosed mass (top right panel), molecular hydrogen fraction as a function of enclosed mass (bottom left panel), and enclosed mass as a function of radius (bottom right panel). As expected, the number density profiles of all of the simulations are very similar over the entire range of WDM (and CDM) particle masses. This is a result of the cooling properties of a gas of primordial composition, as explained in Chapter 4. The simulation assuming a WDM particle mass of 10 keV does not collapse by the time the simulation is stopped at $z = 10$, and the density profile at the last output time is shown. The plot of enclosed mass as a function of radius shows a strong similarity between the different calculations as well, which is to be expected since it is essentially another way of viewing the number density plot. The plots of temperature and molecular hydrogen fraction vs. enclosed mass show a significant amount of scatter. Ignoring the 10 keV case, the overall spread in temperature in the core of the halo is a factor of ~ 3 and the spread in molecular hydrogen fraction is roughly 1.5 orders of magnitude. Interestingly




Figure 5.2: Dark matter halo properties as a function of WDM particle mass for several simulations with the same cosmological realization but different warm dark matter particle masses. Top left: WDM particle mass vs. collapse redshift of halo core. Top right: WDM particle mass vs. collapse time of halo core (measured in millions of years after the big bang). Bottom: halo virial mass at collapse vs. WDM particle mass. In the top two panels the collapse redshift/time of the cold dark matter (CDM) simulation is shown by a vertical, dashed blue line. In the bottom panel the virial mass of the halo in the CDM simulation is shown by a solid blue circle, while the WDM simulations are represented by open black squares.

enough, the CDM “control” simulation has one of the higher core temperatures.

Figure 5.4 shows the specific angular momentum as a function of enclosed mass (top left panel), circular velocity as a function of radius (top right panel), radial velocity as a function of enclosed mass (bottom left panel), and accretion time as a function of enclosed mass (bottom right panel). The angular momentum distributions are extremely similar for all of the calculations (disregarding the 10 keV case since it does not collapse), as is the circular velocity. The Keplerian orbital velocity is plotted in this panel (upper thin black line), and all of the simulations display circular velocities that are significantly below this velocity. The plot of radial velocity as a function of enclosed mass shows that the CDM simulation has the greatest infall velocity at the output time in question, which corresponds to the largest accretion rate overall (as shown in the plot of accretion time vs. enclosed mass). The rest of the calculations have similar infall velocities and accretion rates, except for the 15 keV model, which has a much lower overall infall velocity and accretion rate. The reason for this is not obvious at the present time. The overall accretion rates for the WDM calculations are slightly less than that of the CDM calculation, suggesting that the final stellar masses may be slightly lower.

5.4.2 Comparison of the evolution of two representative WDM realizations

In this section we compare the evolution of two representative warm dark matter simulations. We somewhat arbitrarily choose the calculations with WDM particle masses of 12.5 and 25 keV. Figure 5.5 shows mass-weighted projections of dark matter density, baryon density, and baryon temperature at $z = 20.38$ for the two representative WDM calculations and a CDM calculation of the same cosmological realization. All panels show a volume that is ~ 300 pc (proper) across and are centered on the point in space where the first Population III protostar will form. There is a huge difference between the calculations at a fixed point in time – the cold dark matter calculation (right column) shows a great deal of clumpy dark matter structure, including knots along the cosmological filaments and even dark matter halos in void regions, with corresponding variety in the baryon density and temperature plots. The 25 keV calculation shows the effects of smoothing - two halos are forming, but there are no halos in the voids, and no substructure around the halos that form. This is reflected in the baryon temperature and density plots, where the accretion shocks onto the filaments show little small-scale structure and the gas is quite smooth. The 12.5 keV calculation is an even more striking example of the effects of small-scale smoothing – though an overdensity in the dark matter is apparent, no halos are visible at this redshift and there is no smaller scale structure whatsoever. This particle mass corresponds to a smoothing scale of a few times $10^6 M_{\odot}$, below which top-down fragmentation takes place. This mass is roughly equivalent to the coalescing halo shown in this image.




Figure 5.3: Mass-weighted, spherically-averaged baryon quantities for several simulations with the same cosmological realization but different warm dark matter particle masses. Top left: baryon number density as a function of enclosed mass. Top right: baryon temperature as a function of enclosed mass. Bottom left: molecular hydrogen fraction as a function of enclosed mass. Bottom right: enclosed baryon mass as a function of radius. Output times are chosen such that the peak baryon density in each simulation is approximately the same. In each panel, the CDM simulation is represented by a solid black line. Solid red line: $M_{\text{WDM}} = 35$ keV. Solid blue line: $M_{\text{WDM}} = 30$ keV. Solid green line: $M_{\text{WDM}} = 25$ keV. Dashed black line: $M_{\text{WDM}} = 20$ keV. Dashed red line: $M_{\text{WDM}} = 17.5$ keV. Dashed blue line: $M_{\text{WDM}} = 15$ keV. Dashed green line: $M_{\text{WDM}} = 12.5$ keV. Dot-dashed black line: $M_{\text{WDM}} = 10$ keV. The halo in the $M_{\text{WDM}} = 10$ keV does not collapse by the end of the simulation and is shown at the last available output time.




Figure 5.4: Mass-weighted baryon quantities for several simulations with the same cosmological realization but different warm dark matter particle masses. Top left: spherically-averaged baryon angular momentum as a function of enclosed mass. Top right: cylindrically-averaged circular velocity as a function of radius. Bottom left: spherically-averaged radial velocity as a function of enclosed mass. Bottom right: spherically-averaged accretion rate as a function of enclosed mass. Output times are chosen such that the peak baryon density in each simulation is approximately the same. In each panel, the CDM simulation is represented by a solid black line. Solid red line: $M_{WDM} = 35$ keV. Solid blue line: $M_{WDM} = 30$ keV. Solid green line: $M_{WDM} = 25$ keV. Dashed black line: $M_{WDM} = 20$ keV. Dashed red line: $M_{WDM} = 17.5$ keV. Dashed blue line: $M_{WDM} = 15$ keV. Dashed green line: $M_{WDM} = 12.5$ keV. Dot-dashed black line: $M_{WDM} = 10$ keV. The halo in the $M_{WDM} = 10$ keV does not collapse by the end of the simulation and is shown at the last available output time. In the bottom right panel the dot-long dashed line is the Kelvin-Helmholtz time calculated from Population III stellar properties from Schaerer and the upper and lower diagonal solid black lines correspond to constant accretion rates of 10^{-3} and $10^{-3} M_\odot/\text{yr}$, respectively.




Figure 5.5: Mass-weighted projections of dark matter density, baryon density and baryon temperature for 3 simulations with the same cosmological realization and a range of warm dark matter (WDM) particle masses at $z = 20.38$. The field in each calculation is the same, though the color tables are relative for each panel in order to highlight density differences. Left column: $M_{WDM} = 12.5$ keV. Center column: $M_{WDM} = 25$ keV. Right column: Cold dark matter realization (corresponds to $M_{WDM} \rightarrow \infty$). Top row: projected dark matter density. Middle row: projected baryon density. Bottom row: projected baryon temperature. The spatial scale is ~ 300 pc (proper) in each volume.

Figure 5.6 shows the same quantities and spatial volume as Figure 5.5, though instead of the outputs all being at the same point in time, they are at the time when the halo core collapses in each simulation. This corresponds to $z = 18.001$ for the CDM calculation, $z = 16.54$ for the WDM calculation with $m_{WDM} = 25$ keV, and $z = 12.09$ for the WDM calculation with $m_{WDM} = 12.5$ keV. At the time of collapse the 12.5 keV calculation has formed a halo which is more massive than the CDM halo by a factor of ~ 5 (and collapses approximately 130 million years later). Very little substructure is evident in the projected dark matter distribution of the 12.5 keV calculation. Some is apparent in the 25 keV run, but not nearly as much as in the CDM calculation. As predicted by Bode et al., the warm dark matter calculations have suppressed substructure and satellite halos, and there is significant evidence that the halo in the 12.5 keV calculation forms by top-down fragmentation of a filament rather than hierarchical merging of smaller halos.

Figures 5.7 through 5.9 show the time evolution of several spherically averaged, mass-weighted radial quantities for the two representative warm dark matter calculations. The CDM run is the one discussed in Section 4.4.1, and its evolution is shown in Figures 4.4 through 4.8. The plots are chosen such that the central densities of the collapsing halo core are matched between the two calculations.

Figure 5.7 shows the evolution of number density as a function of enclosed mass for the 12.5 keV and 25 keV WDM calculations. The lowest-density line corresponds to $z = 13.16$ for the 12.5 keV run and $z = 18.05$ for the 25 keV calculation. Intriguingly, it takes the 12.5 keV calculation about 4×10^7 years to advance to a core baryon number density of $n \sim 10^6 \text{ cm}^{-3}$, while the 25 keV calculation only requires $\sim 2 \times 10^7$ years to get to that point. However, once the calculations reach $\sim 10^6 \text{ cm}^{-3}$ they take extremely similar amounts of time to evolve to the highest number density shown. As discussed in previous sections, this reflects the fact that the halo evolution on small scales is controlled by the chemistry and cooling properties of the primordial gas, which is the same in the two calculations.

Figure 5.8 shows the evolution of baryon temperature and molecular hydrogen fraction as a function of enclosed mass for the two WDM simulations. The overall temperature evolution is very similar between the two calculations, though the calculation with $m_{WDM} = 12.5$ keV ends up with a slightly lower molecular hydrogen fraction and slightly higher central temperature. The evolution of radial infall velocity and angular momentum as a function of enclosed mass (shown in Figure 5.9) is also quite similar between the two calculations. As shown in Figure 5.4, the final accretion rates are also essentially the same.

The purpose of this section was to demonstrate that the evolution of the halo collapse, and the resulting protostar, is quite similar for two simulations with significantly different warm dark matter particle masses. The large-scale structure evolves somewhat differently in these two cases – the halo that forms in the 12.5 keV calculation is approximately the same mass scale as the suppression mass, meaning that this dark matter halo is roughly




Figure 5.6: Mass-weighted projections of dark matter density, baryon density and baryon temperature for 3 simulations with the same cosmological realization and a range of warm dark matter (WDM) particle masses at the redshift at which the Population III protostar collapses in each simulation. The comoving size of the projected volume in each calculation is the same, though the color tables are relative for each panel in order to highlight density differences. Left column: $M_{WDM} = 12.5$ keV, $z_{coll} = 12.09$. Center column: $M_{WDM} = 25$ keV, $z_{coll} = 16.54$. Right column: Cold dark matter realization (corresponds to $M_{WDM} \rightarrow \infty$), $z_{coll} = 18.001$. Top row: projected dark matter density. Middle row: projected baryon density. Bottom row: projected baryon temperature. The spatial scale is ~ 300 pc (proper) for the CDM and $M_{WDM} = 25$ keV WDM simulation and ~ 450 pc (proper) for the $M_{WDM} = 12.5$ keV WDM simulation; the comoving scales are the same in each panel.




Figure 5.7: Evolution of spherically-averaged, mass-weighted baryon number density as a function of enclosed mass in halos with two representative warm dark matter simulations. The cosmological realization is the same for each calculation, and output times are chosen such that the baryon densities are approximately the same. Left column: simulation with $M_{WDM} = 12.5$ keV. Right column: simulation with $M_{WDM} = 25$ keV. The realization is the same as the simulation discussed in Section 4.4.1 and these panels are directly comparable to Figures 4.4 through 4.8. Lines for the $M_{WDM} = 12.5$ keV (25 keV) simulations as follows. Black solid line: $t = 319$ Myr/ $z = 13.163$ ($t = 204$ Myr/ $z = 18.05$). Red solid line: 3.12×10^7 years later (1.04×10^7 years later). Green solid line: 8.15×10^6 years later (1.04×10^6 years later). Blue solid line: 98,345 years later (5.73×10^6 years later). Black short-dashed line: 2.86×10^5 years later (2.63×10^5 years later). Red short-dashed line: 1.25×10^5 years later (82,433 years later). Green short-dashed line: 45,152 years later (38,738 years later). Blue short-dashed line: 22,697 years later (24,865 years later). Black long-dashed line: 2691 years later (3332 years later).




Figure 5.8: Evolution of spherically-averaged, mass-weighted baryon temperature (top row) and molecular hydrogen fraction (bottom row) as a function of enclosed mass in halos with two representative warm dark matter simulations. The cosmological realization is the same for each calculation, and output times are chosen such that the baryon densities are approximately the same. Left column: simulation with $M_{WDM} = 12.5$ keV. Right column: simulation with $M_{WDM} = 25$ keV. The realization is the same as the simulation discussed in Section 4.4.1 and these panels are directly comparable to Figures 4.4 through 4.8. The lines are at the same times as in Figure 5.7.




Figure 5.9: Evolution of spherically-averaged, mass-weighted baryon properties in halos with two representative warm dark matter simulations. The cosmological realization is the same for each calculation, and output times are chosen such that the baryon densities are approximately the same. Left column: simulation with $M_{WDM} = 12.5$ keV. Right column: simulation with $M_{WDM} = 25$ keV. Top row: baryon radial velocity as a function of enclosed baryon mass (velocity is positive away from the center of the halo). Bottom row: baryon angular momentum as a function of enclosed mass. The realization is the same as the simulation discussed in Section 4.4.1 and these panels are directly comparable to Figures 4.4 through 4.8. The lines are at the same times as in Figure 5.7.

the smallest object that can directly form at that mass scale. The halo that forms in the 25 keV WDM model is significantly larger than the suppression mass, implying that it formed out of the merging of smaller objects. Despite this, the final protostellar properties are similar, which is due to the collapse dynamics being controlled at small scales primarily by the properties of the primordial gas rather than by the large scale structure.

5.5 Discussion

In this chapter we show how the suppression of small-scale power, which is meant to mimic the effect of a warm dark matter cosmological model, affects the formation of a Population III protostar. We use an identical cosmological realization, but apply smoothing to the initial conditions according to the WDM transfer function given by Bode et al. We find that, for a wide range of warm dark matter particle masses, the main effect of the smoothing is to delay the collapse of the halo core, while the properties of the protostar that forms in the center of the halo remains largely unaffected, and appears to have approximately the same mass range as the reference CDM calculation.

Warm dark matter models are somewhat analogous to CDM calculations with a soft UV background, in that both cause an overall delay in collapse of the halo core and result in halos with a somewhat larger virial mass (corresponding to the later collapse time). This is due to different physical reasons, of course. One striking difference in the warm dark matter calculations is that for WDM particle masses below $\simeq 15$ keV, the suppression mass is actually at the mass of the halo in which the primordial protostar forms (at a few times $10^5 M_\odot$) so a different paradigm for structure formation occurs: the halos at this scale will form by town-down fragmentation of larger objects rather than bottom-up formation via hierarchical mergers.

Examination of the delay of halo collapse allows us to introduce a new constraint on the warm dark matter particle mass. A warm dark matter cosmology with a particle mass of $m_{WDM} \simeq 15$ keV delays the formation of the first star in the simulation volume by approximately 10^8 years (compared to the CDM case with the same cosmological realization). If this delay is representative of the overall delay in structure formation experienced due to small-scale smoothing from the warm dark matter then the 15 keV case is still marginally acceptable when compared against polarization measurements of the CMB from the WMAP satellite, which suggests that the universe was at least partially reionized by $z = 17 \pm 5$ [64], suggesting that a reasonable constraint on the minimum warm dark matter mass from Population III star formation is $m_{WDM} \geq 15$ keV. However, this estimate is somewhat crude, and a large number of cosmological realizations with varied warm dark matter masses should be run in order to test this hypothesis. Our simulations also show that a warm dark matter model with $m_{WDM} = 40$ keV is

indistinguishable from the CDM case, making this something of an upper limit of particle masses that have any conceivable effect on large-scale structure.

Chapter 6

The formation of second-generation primordial objects

6.1 Summary

There has been considerable theoretical debate over whether photoionization and supernova feedback from the first Population III stars facilitate or suppress the formation of the next generation of stars. We present results from an Eulerian adaptive mesh refinement simulation demonstrating the formation of a primordial star within a region ionized by an earlier nearby star. Despite the higher temperatures of the ionized gas and its flow out of the dark matter potential wells, this second star formed within 23 million years of its neighbor’s death. The enhanced electron fraction within the HII region catalyzes rapid molecular hydrogen formation that leads to faster cooling in the subsequent star forming halos than in the first halos. This “second generation” primordial protostar has a much lower accretion rate because, unlike the first protostar, it forms in a rotationally supported disk of $\sim 10 - 100 M_{\odot}$. This is primarily due to the much higher angular momentum of the halo in which the second star forms. In contrast to previously published scenarios, such configurations may allow binaries or multiple systems of lower mass stars to form. These first high resolution calculations offer insight into the impact of feedback upon subsequent populations of stars and clearly demonstrate how primordial chemistry promotes the formation of subsequent generations of stars even in the presence of the entropy injected by the first stars into the IGM. This chapter has been previously published as a paper in the *Astrophysical Journal* [3].

6.2 Motivation

Calculations performed by Abel, Bryan and Norman ([39], hereafter referred to as ABN02) show that rapid accretion rates driven by molecular hydrogen cooling cause the formation of solitary massive protostars in the range of 30 to 300 M_\odot in minihalos of $10^5 - 10^6 M_\odot$ at redshifts $\gtrsim 20$. Simulations indicate that the hard UV spectra of these 10^5 K zero-metallicity stars will envelop them in large HII regions several kiloparsecs in diameter [76, 215]. Over the main sequence lifetime of the central star (on the order of 2-6 Myr for the range of 30 – 300 M_\odot) half of the baryons within the minihalo are driven beyond its virial radius by ionized flows that quickly steepen into shocks. These shocks exhibit expansion rates of up to ten times the escape velocity of the halo. After the death of the central star, cooling and recombination are out of equilibrium in the ionized gas, which results in significant electron fractions even after its temperature has dropped to 1000 - 2000 K after 20 - 50 Myr. One dimensional, nonrotating calculations [61] predict two possible fates for the primordial stars themselves: complete destruction by the pair instability ($140 M_\odot < M_* < 260 M_\odot$) which is very energetic and leaves no remnant, or direct collapse to black holes above and below this mass range, with the added possibility of SN-like precollapse mass ejection by pulsational pair instabilities from 100-140 M_\odot stars [60].

An important question is whether later generations of stars can efficiently form in the relatively high temperatures and ionization fractions of the relic HII regions left by the first stars. One analytical study [73] found that the first stars injected sufficient entropy into the early IGM by photoheating and supernova explosions to prevent further local star formation in their vicinity. Lyman-Werner SUV background radiation is also thought to have contributed negative feedback by photodissociating primordial H_2 and quenching the molecular hydrogen cooling processes allowing the first stars to form [216, 71]. In this chapter we present fully resolved simulations which show that a second primordial star can form in the relic HII region of an earlier Pop III star. We determine its properties, considering the effect of Lyman-Werner radiation from the resultant black hole assuming accretion rates consistent with the density fields left by ionized outflows from the parent minihalo.

6.3 Simulation Setup

We carried out simulations using Enzo, an Eulerian adaptive mesh refinement (AMR) hydrodynamics + N-body code described in Chapter 2. We initialized a box of size $300 h^{-1}$ kpc at $z = 99$ for a cosmology with $(\Omega_M, \Omega_\Lambda, \Omega_B, h, \sigma_8, n) = (0.3, 0.7, 0.04, 0.7, 0.9, 1)$. We first ran a simulation with 128^3 dark matter particles in a 128^3 root grid with 6 total levels of adaptive mesh, refining on a dark matter overdensity of 4.0. This model was run with dark matter alone in order to identify the most massive halo that evolves in the

simulation volume, which at $z \sim 18$ had a mass $\sim 5 \times 10^5 M_\odot$.

We then re-initialized the calculation in the original simulation volume at $z = 99$ with both baryons and dark matter using a 128^3 root grid and three static nested subgrids, each of which was twice as refined as its parent grid and was centered on the Lagrangian volume of the peak that later evolved into the identified halo. The effective root grid resolution was 1024^3 in this volume, which corresponds to a comoving spatial resolution of $\sim 300 h^{-1}$ pc and a dark matter particle mass of $1.8 h^{-1} M_\odot$ in the most highly refined region. Every dark matter particle that later enters into dark matter halos of interest was within this most highly refined grid at the start of the simulation.

We started the simulation with this set of initial conditions at $z = 99$ and followed the collapse of the first star, which occurred at a redshift of 17.76. As a refinement criteria we used a baryon overdensity of 4.0 and a dark matter overdensity of 8.0. In addition, to ensure appropriate simulation resolution we mandated that the Jeans length must be resolved by at least 16 cells at all times, which exceeds the Truelove criterion by a factor of 4 along each axis [217]. At the collapse redshift the three dimensional structure was resolved with 8727 grids on nine levels containing a total of 49,641,744 unique resolution elements.

To compute the extent of the HII region of the $120 M_\odot$ Pop III star assumed to form in the collapse, we interpolated the density, energy, and velocity fields from the entire Enzo simulation volume at the formation redshift of this star onto a three dimensional grid of fixed resolution with 256^3 cells for import into a static radiative transfer code. The code utilizes the ionization front tracking technique of Abel [218] to calculate the boundary of the HII region along rays cast outward from the central star by the adaptive ray tracing technique of Abel & Wandelt [219]. Within the HII region we set the ionization fraction to unity and the H_2 and H^- fractions to zero. We assume that the mean energy of ionization for the gas is 2.4 eV, which results in a post-ionization temperature of $\sim 18,000$ K when calculated in our multispecies ZEUS simulations. This is somewhat cooler than one might expect due to the relatively hard spectrum of massive primordial stars, and is a result of our use of monochromatic radiative transfer in the ZEUS code, which underestimates the UV photoheating of the halo by not taking into account contributions from very high energy photons. Whalen et al. [76] show that an increase in post-front temperatures results in somewhat higher sound speeds. These yield higher shock speeds that promote the photoevaporative flow of gas from the halo in which the first star is formed and could in principle affect the dynamics of nearby halos. We show below that in this case the outflow of gas has a negligible effect on the formation of a second primordial star, which suggests that our result is at worst only weakly affected by post-front temperature. Higher post front temperatures will not significantly retard the cooling and recombination crucial to the formation of molecular hydrogen.

We approximated the dynamics of the HII region by imposing the one dimensional velocity, ionization, density and temperature profiles for a $120 M_\odot$ star at the end of its

main sequence lifetime from Whalen et al. [76] along every line of sight from the central star. We modified baryon densities and velocities out to ~ 120 pc (corresponding to the location of the shock wave in the 1D calculation) but changed only ionization fractions and temperatures beyond this radius out to the boundary of the HII region determined by the ray tracing code. We then mapped this HII region onto the full hierarchy of grids in the Enzo calculation, centering it on the location of the first protostar. This state corresponds to only 2.5 million years after the initial star formed ($z \simeq 17.4$), so we assume that instantaneous ionization is a reasonable approximation for all gas outside the first halo (which has had the hydro profiles from the 1D simulations imposed in it). An important question is whether the satellite halos are also ionized by the I-front propagating outward from the first star, an issue investigated in detail at later redshifts by Shapiro et al. [220]. Simulations we performed in 1D in ZEUS-MP indicate that the neighboring halos are photoionized by the parent star by the end of its main sequence lifetime.

We then continued the simulation until the collapse of the next protostar, which occurs at $z = 16.44$, 22.8 million years later. The final time that we analyzed contains 10,710 grid patches on 24 levels with 54,996,560 unique resolution elements. In this calculation we neglect the pulsational pair instability that may eject the hydrogen envelope for this star [60].

As a check on our simulation setup we also ran a simulation where we simply instantaneously ionized the entire simulation volume by raising the baryon temperature to $\sim 10,000$ K and setting ionization fractions to one and H_2 fractions to zero. This simulation tests whether the addition of the one dimensional radial profiles from the Whalen et al. [76] calculations changed the properties of the second protostar appreciably. We find that the collapse time and accretion rate of the protostar formed in this simulation are essentially identical to the results of our full setup, and only discuss results from the full setup in the rest of this chapter.

6.4 Results

The second primordial protostar forms in a neighboring minihalo approximately 265 proper parsecs from the location of the halo in which the first star formed (and where the HII region originated). The halo in which this second protostar forms was completely ionized by the first star to a temperature of $\sim 1.7 \times 10^4$ K. Due to its relatively high density, the center of this halo cools very rapidly and molecular hydrogen formation is catalyzed by the extremely high electron fraction. After only a few million years the core of the halo has a molecular hydrogen fraction of $\sim 5 \times 10^{-3}$, well above what one would expect for a halo which has not been ionized. This halo is significantly smaller than the first: $\sim 2 \times 10^5 M_\odot$ rather than $\sim 5 \times 10^5 M_\odot$.

6.4.1 Comparison of the First and Second Stars

Figure 6.1 compares the mass accretion times of the initial and second Population III stars formed in this simulation. In addition, this figure shows the mass accretion time of the halo in ABN02 and an estimate of the Kelvin-Helmholtz timescale as a function of mass, using values of luminosity and effective temperature taken from Schaerer [59]. The upper and lower dotted lines correspond to an object with constant accretion rates of 10^{-3} and $10^{-2} M_{\odot}/\text{year}$, respectively. Our calculation of accretion timescales for the initial protostar agrees well with that of ABN02. The fact that the two results are in good agreement even though the ABN02 calculations assumed a lower baryon fraction supports the analysis of Ripamonti and Abel [47] showing that all mass scales in these calculations are set by molecular physics. Comparison of the accretion rates to the Kelvin-Helmholtz timescale provides an estimate of $\sim 200 M_{\odot}$ for the upper bound of the mass of the star. The accretion timescales suggest a reasonable lower bound of $\sim 80 M_{\odot}$, since this much gas will accrete in 10^4 years, an insufficient time for fusion to begin. In contrast, the accretion rate of the second protostar is over an order of magnitude lower. This is because the second protostar has a much more pronounced thick disk structure than the first protostar. The disk is rotationally supported past a radius ~ 0.01 pc (corresponding to an enclosed mass of $\sim 10 M_{\odot}$), whereas the disk around the first star in the volume is not. Similar accretion timescale arguments as before suggest a mass of $\sim 5 - 20 M_{\odot}$ for the second star, although accretion physics will ultimately determine the true mass, particularly given the presence of this more pronounced disk.

Examination of the net angular momentum of the two halos is illuminating. The angular momentum of a cosmological halo can be described by the dimensionless spin parameter: $\lambda \equiv J|E|^{1/2}/GM^{5/2}$ where J is angular momentum, E is the total energy, G is the gravitational constant and M is the halo angular momentum. This is roughly equivalent to the ratio of the angular momentum in the halo to the angular momentum needed for the halo to be completely rotationally supported [123]. Typical values of the spin parameter for cosmological halos are $\sim 0.02 - 0.1$, with a mean of $\lambda \simeq 0.05$ [221, 222]. We find that the halo in which the first primordial protostar forms has a spin parameter for the gas and dark matter of $(\lambda_{gas}, \lambda_{dm}) = (0.0275, 0.0363)$, which is slightly lower than the mean. The spin parameter of the second halo is $(\lambda_{gas}, \lambda_{dm}) = (0.1079, 0.1607)$, which is atypically high. Examination of the evolution of angular momentum in the gas of the halos as the two protostars form shows that the angular momentum distributions are different in the two clouds, and if angular momentum is conserved one would expect to see a centrifugally supported disk that is approximately four times larger in the second halo.




Figure 6.1: Mass accretion time $t_a = M(r)/\dot{M} \equiv M(r)/(4\pi r^2 \rho v_r)$ as a function of enclosed gas mass. This is at the final output corresponding to $z = 16.437$. The dashed line is the corresponding data dump of the initial star which had formed at $z = 17.67$. The red dashed line corresponds to the first star to form in this simulation. The blue dot-dashed line corresponds to the first star calculated in ABN02. The solid black line corresponds to the second star forming in this simulation, and the green long-dashed line corresponds to the Kelvin-Helmholtz time of a representative star. The upper and lower black dotted lines correspond to objects with constant mass accretion rates of 10^{-2} and $10^{-3} M_{\odot}/\text{yr}$, respectively.

6.4.2 Black Hole Accretion

Here we consider whether accretion onto a relic black hole could generate enough photodissociative radiation to inhibit H_2 formation in the second star’s halo. We assume Bondi-Hoyle accretion [223] for the $120 M_\odot$ black hole that forms after the collapse of the first star to estimate the Lyman-Werner flux from its accretion. This rate depends on the mass of the accretor as well as the local gas temperatures, densities, and relative velocities it encounters. To sample the local environment the black hole would traverse over the duration of the simulation, we followed the 40 dark matter particles closest to the first protostar (within ~ 0.1 proper pc) from the end of its main sequence lifetime until the collapse of the second protostar. We tallied the cell quantities they crossed to compile the accretion rate history each particle would have if it were the black hole. The histories for the 40 black hole proxies appear in Figure 6.2. The mass accretion rates grow from $10^{-11} M_\odot/\text{yr}$ to $10^{-8.5} M_\odot/\text{yr}$ for most of the tracer particles.

To estimate the effect of Lyman-Werner radiation from the black hole on molecular hydrogen formation in nearby halos we assume a canonical 10% radiative efficiency for the accretion. The uppermost accretion curve yields $2.2 \times 10^{37} (M/100 M_\odot)$ erg/s ($\sim 4500 L_\odot$) for an upper limit to the total luminosity (which is much lower than the Eddington luminosity of this object, 1.5×10^{40} erg/s, or $\sim 4 \times 10^6 L_\odot$). Taking this to be a blackbody spectrum, the flux in the Lyman-Werner band (11.1-13.6 eV) reaching the second protostar is $\sim 1.6 \times 10^{-25} (M/100 M_\odot)$ erg s $^{-1}$ cm $^{-2}$ Hz $^{-1}$, resulting in photodissociation rates that are significantly lower than the formation rates of molecular hydrogen there. The expulsion of gas by ionized flows from the first halo prevents higher accretion rates and greater Lyman-Werner fluxes. A star in this mass range may shed its envelope just prior to collapse, resulting in a smaller black hole and making the results discussed here an upper limit.

6.5 Discussion

This first high resolution three dimensional simulation of the evolution of gas within a primordial HII region demonstrates the crucial role of H_2 chemistry driven by photoionization in the formation of the next generation of stars. While this has been addressed in previous work [224] our simulations are the first with sufficient resolution to directly examine the formation of individual protostars. Further investigation will be necessary to determine if the lower accretion rates leading to the smaller mass of the second star are a coincidental feature of this calculation or a general trend of early star formation in halos preprocessed by HII regions. The low accretion rate that we observe in this calculation is primarily due to the high initial angular momentum of the second halo.

One possible source of error lies in the method and assumptions determining whether the neighboring halos are photoionized. While our 1D results indicate that these halos




Figure 6.2: Bondi-Hoyle mass accretion rate around the black hole calculated from the local gas temperature, density and relative velocity. Integration of these curves lead to estimates of growth of the black hole (initially $120 M_{\odot}$) of that range from 0.009 to $0.032 M_{\odot}$ over 23 Myrs

will be ionized, this issue merits further investigation with fully 3D simulations. We further assume that this ionization occurred instantaneously and simply ionize the gas outside of the initial halo without changing the total density or velocity profiles of nearby halos. Instantaneous ionization appears to be a reasonable approximation since the sound crossing time of all of the ionized halos is longer than the main-sequence lifetime of the parent star. Again, full 3D radiation photo-evaporation simulations will be necessary to determine whether the hydrodynamic evolution of these halos during the main sequence lifetime of the parent star is unimportant.

We note that our HII region enveloped roughly a dozen minihalos similar to the one that formed the second star. More calculations will be required to see if these too form stars. The evolution of the massive disk also merits examination to ascertain whether it breaks up into a multiple system or fully accretes to form a single star. The situation realized in our cosmological simulation may lead to objects with initial conditions similar to the cases studied by Saigo et al. [225]. Lower mass second generation stars or the possibility of binaries or multiple systems of primordial stars would have strong implications for the observability of such objects and their impact on subsequent structure formation. Less massive stars might have different nucleosynthetic signatures than those of the pair-instability supernovae that may occur in the first generation of primordial stars. The immense size of early HII regions could also make the scenario of primordial stars forming in a relic HII region much more common than extremely massive stars forming in pristine halos. These two facts taken together may account for the lack of detection of the characteristic odd-even abundance pattern from pair-instability supernovae expected in observations of ultra metal poor halo stars ([226] and references therein). How HII regions from the first stars may regulate local star formation by suppressing the collapse of gas in local halos which have not reached relatively high densities also remains to be explored.

Chapter 7

Formation of metal-enriched 2nd generation objects

7.1 Summary

In this chapter we discuss very preliminary results concerning the dispersal of metals into the IGM and nearby halos from two representative Population III supernovae: a $30 M_{\odot}$ star with an explosion energy of $E_{SN} = 1.2 \times 10^{51}$ ergs, and a $250 M_{\odot}$ star with an explosion energy of $E_{SN} = 8 \times 10^{52}$ ergs. Despite the amount of gas in the Population III halo and the relatively high density of the IGM at $z \simeq 18$, the explosion of the $30 M_{\odot}$ star spread metals over a region several hundred proper parsecs in radius within fifty million years. The simulation of the supernova from the $250 M_{\odot}$ star is ongoing, but due to the preprocessing of the halo in which the star forms by the star's HII region and the much higher explosion energy suggests that the more massive star will spread metals over a much larger volume of space.

7.2 Motivation

Observations of quasar absorption spectra show that the universe at the present day is uniformly polluted with metals, even at the lowest observed column densities, which correspond to regions of very low overdensity which are known as the Lyman- α forest [84, 85]. The primordial composition of the universe is well understood, and post-BBN nucleosynthesis is believed to take place only in stars and cataclysmic events associated with them. Because of this, it is apparent that this period of enrichment must have taken place between the epoch of first star formation and the present day.

The metals produced by Population III supernovae would have a very important effect on the following generations of stars. The addition of metals greatly enhances the cooling

properties of the gas – molecular hydrogen is a relatively poor coolant compared to dust grains and metal-based molecules such as CO, both of which are produced in significant quantities by both primordial Type II supernovae and pair-instability supernovae, with the fractional quantity of dust increasing as the stellar mass increases [89, 90, 91]. Very little metal is required for gas to cool efficiently – analytical work and simulations suggest that the presence of carbon and oxygen at levels $10^{-4} - 10^{-3} Z_{\odot}$ would be sufficient for enhanced fragmentation of collapsing gas clouds, signifying a change from the top-heavy Population III IMF to a mass function resembling that observed in the galaxy today [93, 94]. As with their HII regions, the metals ejected from Population III supernovae, particularly if the stars fall into the mass range that produces highly energetic pair-instability supernovae, can propagate to great distances – simulations indicate that the ejecta from a massive PISN can eject metal into a sphere of ~ 1 kpc diameter at $z \sim 20$, producing a metallicity floor above that needed for enhanced cooling to take place [95]. However, these calculations were done using SPH calculations with poor spatial and mass resolution, and cannot accurately resolve such important properties such as the degree to which the supernova ejecta mixed with the primordial gas and also the dispersion of metallicities in the enriched gas. Furthermore, previous calculations did not follow the ejected gas to the point where the metals reached nearby halos, and due to their use of SPH, even if they did allow the calculation to progress for a long enough time they would not have been able to make any statements about the degree to which metals may have enriched nearby halos.

This chapter presents very preliminary results from two calculations done with the Enzo AMR code (described in detail in Chapter 2). These calculations are of the explosions of two Population III stars: one with a mass of $30 M_{\odot}$ and an explosion energy of 1.2×10^{51} ergs (corresponding to the Population III equivalent of a standard Type II supernova), and a second with a mass of $250 M_{\odot}$ and an explosion energy of 8×10^{52} ergs (corresponding to a massive pair-instability supernova). These two calculations bracket the plausible range of Population III supernova energies presented by Heger et al. [60, 61]. The purpose of the calculations is to model the ejection of metals from these representative Pop III supernovae into the IGM and to follow the mixing history of the supernova ejecta. The ultimate goal is to predict the overall dispersal of metals into the IGM by these stars, and also to attempt to understand the formation site of the first generation of metal-enriched stars.

7.3 Problem setup

In order to obtain the initial conditions for our supernova calculations, we perform a standard Population III star formation calculation, as described in detail in Chapter 4. This calculation was set up in a $0.3 \text{ h}^{-1} \text{ Mpc}$ box, as described in Section 4.3, with

identical cosmological parameters, initialization strategy, and spatial and mass resolution. In this calculation we use the ZEUS hydro method (Section 2.2.3) instead of the PPM method (Section 2.2.2) because it is more stable in situations where the overall pressure jump between adjacent cells is extremely high. This does not change the results of the star formation calculation appreciably. The simulation was followed until $z \simeq 18$, at which point the halo core collapsed and began to form a Population III protostar. At this point, we stop the calculation and put a Pop III supernova into the calculation by hand in the simulation volume and the simulation is then allowed to continue.

The $30 M_{\odot}$ supernova is set up as follows: We assume that the evolution of the star along the main sequence occurs much more rapidly than the evolution of cosmological structure, and place the supernova into the calculation at the redshift at which we stopped the calculation. One dimensional calculations by Dan Whalen (private communication) suggest that in the spherically symmetric case, the HII region produced by a $30 M_{\odot}$ primordial star will not escape the halo – it remains as a completely bound HII region with a maximum radius of ~ 0.1 pc (proper). For this reason, we do not concern ourselves with the HII region, since it will not affect the dynamics of the supernova remnant in any significant way. Using the mean density and temperature of the inner few parsecs of the halo at this time, we calculate the Sedov blast profile of the supernova at a radius of 2 proper parsecs assuming $E_{SN} = 1.2 \times 10^{51}$ ergs. The mean density of the inner two parsecs of the halo is $n_{mean} \simeq 1800 \text{ cm}^{-3}$ and the mean temperature is approximately 500 K. At this radius the velocity of the supernova shock is 171 km/s. We assume that the supernova is spherically symmetric, and overlay the Sedov blast profile over the inner two parsecs of the halo, centered on the cell with the highest baryon density at the time of collapse. The calculation is then allowed to evolve for approximately 50 million years, with snapshots of the state of the simulation taken at regular intervals over this time.

The $250 M_{\odot}$ supernova is set up in a somewhat different way. One dimensional calculations by Whalen, Abel and Norman [76] show that in the spherically symmetric case the HII region from a $250 M_{\odot}$ primordial star will extend out to several proper kiloparsecs at $z = 20$, and additionally the increase in gas pressure in the cosmological halo will cause most of the gas in the halo to be ejected by the end of the main sequence lifetime of the star. At the end of the star’s life, this gas will be in a thick, dense shell at roughly the virial radius with a velocity of approximately 30 km/s, which is roughly a factor of ten greater than the halo’s escape velocity. Due to this, we set up our supernova calculation assuming that the $250 M_{\odot}$ supernova goes off within a HII region which is set up using an identical method to that described in Section 6.3. We then calculate the Sedov blast profile of a supernova with $E_{SN} = 8 \times 10^{52}$ ergs exploding in an ambient medium with a density of $\simeq 0.4 \text{ cm}^{-3}$ and a temperature of 17,000 K, which is approximately the mean density and temperature of the gas left over in the halo by the HII region at the end of the star’s main sequence lifetime. At this density, the supernova will sweep up its own mass of gas in the ambient medium at $r \simeq 20$ pc (which

is approximately the mass scale at which it makes the transition from the free-expansion to Sedov-Taylor phase), and will not encounter significant density increases until the gas reaches ~ 60 pc. We somewhat arbitrarily decide to put the Sedov profile in with a radius of 40 proper parsecs. At this radius the supernova explosion is propagating outward at a speed of $\simeq 1800$ km/s.

Both supernova calculations are then restarted, assuming a maximum refinement level of 10 (corresponding to a spatial resolution of approximately 0.17 proper parsecs at $z = 18$) and allowed to continue until the velocity of the supernova ejecta drops below the sound speed of the ambient medium, which approximately corresponds to the point at which the supernova ejecta and ambient medium reach pressure equilibrium. This takes slightly more than fifty million years for the $30 M_{\odot}$ calculation. At the time that this chapter is being written the $250 M_{\odot}$ calculation has not yet been completed, since the overall computational cost is much greater, due to the higher shock velocities and hotter average temperature of the ambient medium due to the HII region. Therefore, we only present results from the $30 M_{\odot}$ supernova calculation below.

7.4 Preliminary results

Due to time constraints, only very preliminary results are presented in this chapter. Figures 7.1 and 7.2 show the evolution of the supernova remnant over approximately 60 million years. Figure 7.1 shows the evolution of the remnant over approximately the first half million years after the supernova occurs. In this figure, the left, center and right columns show the initial conditions for our supernova (the Sedov blast profile), 2.68×10^5 years after the initial explosion, and 5.51×10^5 years after the initial explosion, respectively. The rows correspond to (from top to bottom) log baryon density, log metallicity, and log baryon temperature. In all cases the color table is fixed such that all panels in a given row can be directly compared. In the plots of log metallicity the color table spans a range of metallicities from $10^{-4} Z_{\odot}$ and below (black, essentially unenriched gas) to 10 times solar metallicity (bright yellow). The majority of gas contained within the actual remnant in the second two timesteps is at a few times solar metallicity. The temperature color table ranges from 100 K (black) to 10^4 K (bright white). The mean temperature of the gas within the supernova remnant at late times is a few thousand K. The box is a projection of a field of view 93 proper parsecs on a side (at $z \simeq 18$) and 93 proper parsecs deep. The development of non-perturbations can be seen very clearly as the supernova remnant expands. Though the initial conditions for the supernova assume spherical symmetry at small scales, inhomogeneities in baryon density outside of the halo core introduce asymmetries which are magnified as the remnant expands, resulting in small “bullets” of gas and significant variations in density, temperature and metallicity in the remnant shell. Unfortunately, some numerical artifacts can be seen –

the supernova remnant shows some signs of a cross-like symmetry, aligned with the grid axes. Most of the “bullets” of high density gas also appear to be preferentially traveling along the axes. The effect is relatively subtle, fortunately, and may not have a significant effect on the results.

Figure 7.2 shows the evolution of the remnant at later times and on a much larger scale. The projected volume is now 1.5 proper kpc on a side and 1.5 proper kpc deep at $z \sim 18$. The columns correspond to (from left to right) 5.51×10^5 , 5.79×10^6 , and 6.3×10^7 years after the supernova is initialized, respectively. The final timestep corresponds to $z \simeq 15$. The rows and color tables are identical to Figure 7.1. As the supernova remnant expands it becomes less spherically symmetric, eventually becoming pressure-balanced with the IGM after metal has spread several hundred parsecs from the initial halo. Comparison of the leftmost and rightmost columns shows that metal has spread to at least the distance of the nearest neighboring halos, though it is unclear at present how much metal has managed to get to the center of the neighboring halos, where one presumes star formation would take place.

Figures 7.3 and 7.4 show several spherically-averaged, mass weighted baryon quantities as a function of radius. The lines correspond to the five time outputs shown in Figures 7.1 and 7.2. Figure 7.3 shows the number density, enclosed mass, radial velocity and baryon temperature as a function of radius. The shock starts out at $r = 2$ pc with a maximum speed of approximately 100 km/s and an overall temperature of 10^5 K (which is too high to see on this plot). The r^{-2} density profile of the halo can be clearly seen at larger radii. As time goes by, the shock rapidly cools and slows down, and a shell of gas with a thickness of several parsecs forms that is rushing out of the halo. The extreme thickness of this shell may be due to the poor cooling properties of the gas – in this calculation, we assume a primordial cooling function for simplicity. The relatively low molecular hydrogen fraction (discussed below) and temperature ($\sim 10^3$ K) suggest that the gas is not cooling efficiently. The use of a cooling function that assumes a metallicity consistent with that observed in the remnant (also discussed below) would result in much colder gas and presumably a thinner shell. It is also interesting to note the feature in the radial velocity plot at early times where the outgoing gas has a double-peaked profile, with an inner shock that has a slower velocity than the outer one. The inner shock is actually traveling backwards in the rest frame of the main shock, and corresponds to gas which has caught up to the decelerating outer shock and has reflected off of it. This feature is common to calculations of supernova remnant expansion.

Figure 7.4 shows the time evolution of spherically averaged radial profiles of the mass-weighted metallicity and metal dispersion (in units of solar metallicity), as well as the molecular hydrogen mass fraction. In this case what we call the “metal dispersion” is really the square root of the mass-averaged metal dispersion calculated in spherical shells, so that it also has units of solar metallicity (instead of the square of solar metallicity). As one might expect, the metallicity drops rapidly as the remnant expands and mixes with




Figure 7.1: Early time evolution of the supernova at $t = t_{SN}$ (left column), $t = t_{SN} + 2.68 \times 10^5$ years (center column) and $t = t_{SN} + 5.51 \times 10^5$ years (right column). Rows correspond to (from top to bottom) projected, mass-weighted baryon density, metallicity, and temperature. The projected region is 93 proper parsecs across at $z \simeq 18$. Images of all quantities have fixed color tables and the columns are directly comparable.




Figure 7.2: Late time evolution of the supernova at $t = t_{SN} + 5.51 \times 10^5$ years (left column), $t = t_{SN} + 5.79 \times 10^6$ years (center column) and $t = t_{SN} + 6.3 \times 10^7$ years later (right column). Rows correspond to (from top to bottom) projected, mass-weighted baryon density, metallicity, and temperature. The projected region is $\simeq 1.5$ proper kiloparsecs across at $z \simeq 18$. Images of all quantities have fixed color tables and the columns are directly comparable.

the pristine gas. However, after the initial time shown the variance of the metallicity remains on the same order of magnitude as the actual metallicity, suggesting that the metallicity of the gas in spherical shells varies strongly, indicating that the gas is poorly mixed. This is clearly true at late times, since the supernova remnant is no longer spherically symmetric.

7.5 Discussion

This chapter presents results from a representative Population III supernova. Two calculations were initialized (supernovae from $30 M_{\odot}$ and $250 M_{\odot}$ primordial stars) but only the $30 M_{\odot}$ case was completed by the time that this thesis was due. The $250 M_{\odot}$ is a significantly more costly calculation since the higher temperatures and blast velocities directly correspond to a lower Courant time, making the simulation more than an order of magnitude more expensive overall.

These calculations make several simplifying assumptions that may affect the overall results. In both of the calculations we make the assumption that the supernova remnant is spherically symmetric at the radius at which we put in the Sedov blast profile. While a reasonable assumption for a first calculation of Population III supernovae, this is clearly not true for supernovae from rotating stars, and may affect the overall results. The trend in a rotating star would be towards more metal mixing early in the explosion, resulting in a greater overall distribution of metallicities in the gas polluted by these supernovae, making the scatter in metal distributions in these calculations a lower limit. This being said, the supernovae encounter many thousands of times their own initial mass in gas outside of the radius at which the Sedov profile is put in the gas. This gas is not of uniform composition, and causes a great deal of metal mixing. This can be seen in the figures in Section 7.4, where the remnant shows distinctly non-spherical features and a high degree of metal mixing at both early and late times, so even though the initial setup assumes spherical symmetry the result may not be significantly different than if the initial setup was of an aspherical supernova.

An additional simplifying assumption concerns the cooling rate of the supernova remnant. In this calculation we assume a nonequilibrium primordial chemistry, which predicts a cooling rate that is far lower than that of a metal-enriched gas. This lower cooling rate results in a thicker remnant shell, which then becomes less likely to be affected by the Vishniac instability [227]. This instability is significant in that it can cause fragmentation in the shell of the remnant, which may lead to the development of a second generation of stars which are metal-enriched. This level of fragmentation is not seen in this calculation, due to the choice of cooling function. Future work using simulations whose cooling tables assume metal line cooling will be done, and will show whether this effect is significant.




Figure 7.3: Evolution of spherically-averaged, mass weighted baryon properties of the supernova remnant as a function of time. Black line: $t = t_{SN}$. Red line: $t = +2.68 \times 10^5$ years. Blue line: $t = t_{SN} + 5.51 \times 10^5$ years. Green line: $t = t_{SN} + 5.79 \times 10^6$ years. Cyan line: $t = t_{SN} + 6.3 \times 10^7$ years. Top left panel: number density as a function of radius. Top right: enclosed mass as a function of radius. Bottom left: radial velocity as a function of radius. Bottom right: temperature as a function of radius. The x-axis in all plots is identical.




Figure 7.4: Evolution of spherically-averaged, mass weighted baryon properties of the supernova remnant as a function of time. Black line: $t = t_{SN}$. Red line: $t = +2.68 \times 10^5$ years. Blue line: $t = t_{SN} + 5.51 \times 10^5$ years. Green line: $t = t_{SN} + 5.79 \times 10^6$ years. Cyan line: $t = t_{SN} + 6.3 \times 10^7$ years. Top left panel: metallicity (in units of solar metallicity) as a function of radius. Top right: metal dispersion as a function of radius. Bottom: molecular hydrogen mass fraction as a function of radius. The x-axis in all plots is identical.

It is interesting to note that the $30 M_{\odot}$ supernova completely disrupts the halo that it occurs in, and drives the majority of the baryons in the halo into the intergalactic medium. This can be understood simply by comparing the energy of the supernova with the binding energy of the halo. The $30 M_{\odot}$ supernova has a total energy of 1.2×10^{51} ergs, more than 95% of which is kinetic energy at the time at which the Sedov blast profile is assumed. At this point in time the total mass of the halo is $4.15 \times 10^5 M_{\odot}$, with $4.7 \times 10^4 M_{\odot}$ of this being in baryons. The virial radius of a halo with this mass is $\simeq 125$ proper parsecs at $z = 18$. Altogether, this corresponds to a binding energy of 7×10^{49} ergs for the entire halo. However, the dark matter does not directly feel the effects of the supernova explosion, so the relevant component of the halo is simply the baryon gas. This component of the halo has a binding energy of approximately 8×10^{48} ergs, which is more than two orders of magnitude less than the supernova explosion energy! This calculation of the binding energy makes the simplifying assumption that the density of the halo is uniform, and somewhat underestimates the overall binding energy of the halo, which has a density profile that scales as r^{-2} . Still, we obtain a reasonable order-of-magnitude estimate that agrees with the simulation results and provides a simple physical understanding of the situation.

The calculations described in this section are considered to be preliminary “scoping” calculations for a more advanced series of simulations involving output from calculations looking at the three-dimensional evolution of supernovae from rotating Population III stars done by Fryer & Rockefeller. Additionally, these calculations will include a cooling table which correctly incorporates cooling due to metal molecules and dust, and will be done at higher spatial resolution to more accurately determine metal mixing in the intergalactic medium and nearby halos, and to examine the formation of the first generation of metal-enriched stars.

Chapter 8

Pre-galactic enrichment of the IGM

8.1 Summary

We examine the dynamical evolution and statistical properties of the supernova ejecta of massive primordial stars in a cosmological framework to determine whether this first population of stars could have enriched the universe to the levels and dispersions seen by the most recent observations of the Lyman- α forest. We evolve a Λ CDM model in a 1 Mpc^3 volume to a redshift of $z = 15$ and add “bubbles” of metal corresponding to the supernova ejecta of the first generation of massive stars in all dark matter halos with masses greater than $5 \times 10^5 M_\odot$. These initial conditions are then evolved to $z = 3$ and the distribution and levels of metals are compared to observations. In the absence of further star formation the primordial metal is initially contained in halos and filaments. Photoevaporation of metal-enriched gas due to the metagalactic ultraviolet background radiation at the epoch of reionization ($z \sim 6$) causes a sharp increase of the metal volume filling factor. At $z = 3$, $\sim 2.5\%$ of the simulation volume ($\approx 20\%$ of the total gas mass) is filled with gas enriched above a metallicity of $10^{-4} Z_\odot$, and less than 0.6% of the volume is enriched above a metallicity of $10^{-3} Z_\odot$. This suggests that, even with the most optimistic prescription for placement of primordial supernova and the amount of metals produced by each supernova, this population of stars cannot entirely be responsible for the enrichment of the Lyman- α forest to the levels and dispersions seen by current observations unless we have severely underestimated the duration of the Pop III epoch. However, comparison to observations using carbon as a tracer of metals shows that Pop III supernovae can be significant contributors to the very low overdensity Lyman- α forest. This chapter has been previously published as a paper in the *Astrophysical Journal* [1].

8.2 Motivation

Recent observations by Schaye et al. and Aguirre et al. [84, 85] have shown that the Lyman- α forest is polluted with metals at very low densities. The distribution of metal is very strongly dependent on overdensity, with median metallicity values ranging from $[C/H] = -4.0$ at $\log \delta = -0.5$ (where δ is defined as $(\delta \equiv \rho/\bar{\rho})$) to $[C/H] = -2.5$ at $\log \delta = 2.0$ using their fiducial UV background model. Their observations show little evidence for metallicity evolution of the Lyman- α forest over the redshift range $z = 1.5 - 4.5$.

The lack of observed evolution in metallicity is suggestive of a very early epoch of stellar evolution. Recent observations by the Wilkinson Microwave Anisotropy Probe suggest an epoch of star formation in the redshift range of $z = 11 - 30$ [64], which is consistent with the simulation results of Abel et al. [39] and Bromm et al.[42], which suggest that the first generation of stars (known as Population III, or Pop III) formed in the redshift range $z = 20 - 30$. The Abel et al. results, which are the highest-resolution simulations of formation of the first generation of primordial stars to date, also suggest that Pop III stars are very massive - on the order of $\sim 200M_{\odot}$. Stars that are in this mass range will die in extremely energetic pair-instability supernovae and can eject up to $57 M_{\odot}$ of ^{56}Ni [60, 61]. The formation site of Pop III stars is in halos with total masses of $\sim 10^6 M_{\odot}$. [39, 228]. These halos have escape velocities which are on the order of a few km/s. Due to the shallowness of the potential wells that Pop III stars form in, Ferrara [229] suggests that ejecta from a massive Pop III supernova can propagate to very large distances (far greater than the halo virial radius), a result which is supported in simulations performed by Bromm et al.[95].

In this chapter we describe the results of cosmological hydrodynamic simulations which address whether or not a population of massive primordial stars can be responsible for metal enrichment of the Lyman- α forest to the level and dispersion seen today. We examine the most optimistic possible scenario for Pop III star formation and enrichment in order to establish an upper limit on metal enrichment of the Lyman- α forest due to Population III stars.

8.3 Problem Setup

The simulations were set up using the concordance cosmological model ($\Omega_m = 0.3, \Omega_{\Lambda} = 0.7, \Omega_b = 0.04$ and a Hubble parameter of $h = 0.7$ where h is in units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$). Initial perturbations in the dark matter and baryon density were created with an Eisenstein & Hu [194] power spectrum with $\sigma_8 = 0.9$ and $n = 1$ using the Zel'dovich approximation [125] in a box which is $0.7 h^{-1}$ comoving Mpc on a side. In our simulations we use a computational box with 256^3 grid cells with a comoving spatial resolution of $2.7 h^{-1} \text{ kpc}$ and a dark matter mass resolution of $1477 h^{-1} M_{\odot}$.

The simulation was initialized at $z=99$ and was allowed to evolve to $z=15$ using the Eulerian adaptive mesh refinement code Enzo (as described in Chapter 2). The simulation was stopped at $z = 15$ and all halos with dark matter mass $M_{DM} \geq 5 \times 10^5 M_{\odot}$ were found using the HOP halo-finding algorithm [196]. This yielded 184 halos, each of which we assume produces one pair instability supernova. Note that we are ignoring negative feedback, which raises the minimum halo mass which can form Pop III stars [?]. This is consistent with our desire to simulate a best-case scenario. We discuss this and other assumptions in Section 8.5.

At this point, two separate cases are considered. In the first case (referred to as Case ‘A’), spheres of uniform metal density 1 kpc (proper) in radius with $127 M_{\odot}$ of metal are placed in the simulation, centered on the cell of highest dark matter density of each halo. This corresponds to a mass averaged metallicity in the volume of $\langle Z \rangle \equiv M_Z/M_B = 4.02 \times 10^{-5} Z_{\odot}$ (where M_Z and M_B are total metal and baryon masses in the simulation volume), which remains constant throughout the simulation. No other modifications to the data were made – in particular, the baryon density, temperature and peculiar velocities and dark matter density were unmodified.

In the second case (Case ‘B’), the spheres of uniform metal density are placed down in the same manner. In addition, the baryon gas in the corresponding volume is smoothed to the cosmic mean ($\langle \rho \rangle = \Omega_b \rho_c$), and the temperature of the baryon gas is set to 10^4 K. This corresponds to the net smoothing and heating of baryons in primordial halos due to pair-instability supernovae. The mass averaged metallicity in this case is slightly greater, $\langle Z \rangle = 4.11 \times 10^{-4} Z_{\odot}$. This is due to a small net loss in baryon mass when the densities are smoothed. As in Case A, peculiar velocities and dark matter density are unmodified.

The simulations are then evolved to $z=3$, following the evolution of the gas using non-equilibrium chemistry and cooling from hydrogen and helium and their ions [152, 153]. Metal line cooling is ignored. At $z = 7$ we initialize a uniform metagalactic ultraviolet background field [230] with a slope of $q_{\alpha} = -1.8$ that is ramped to a pre-computed value by a redshift of $z = 6$. Zhang et al. [231] determined that such a prescription can reproduce the Gunn-Peterson effect in the hydrogen Ly- α lines by a redshift of 5.5. Self-shielding of the gas is not explicitly included in the assumed UV background.

It is important to note that the box size, at $0.7 h^{-1}$ Mpc, is somewhat small. Statistical results are reliable at $z = 15$. However, by $z = 3$ (when the simulations are terminated) the box is too small to be statistically reliable. We also performed a simulation in eight times the volume at the same mass and spatial resolution (i.e. 512^3 cells/particles), which gave results indistinguishable from what follows. Nonetheless, all results at late times should be considered qualitative.




Figure 8.1: Projected log metal density (Case A). The area viewed is a projection of the entire simulation volume. The four panels correspond to (from left to right) $z=15$, 7, 5 and 3. The Pop III supernova remnants are placed in the volume at $z=15$ and advect along the filaments. Photoevaporation of gas in the filaments, driven by the metagalactic UV background, causes the volume filling factor of the metal-enriched gas to increase substantially by $z=3$.

8.4 Results

The general evolution of our simulation after injection of metal bubbles at $z = 15$ is as follows: Beginning at $z = 15$, the bubbles of metal track the flow of gas onto and along dark matter filaments. The competing effects of advection along the filaments and the collapse of filaments during this period essentially cancel out, with little net effect on the fraction of the volume occupied by metal enriched gas (also referred to as the volume filling factor, or VFF). Regions of relatively high metallicity ($Z \geq 10^{-3} Z_{\odot}$), corresponding to the densest regions of filaments, decrease their volume filling factor significantly from $z=15$ to $z=6$. Case A shows more of a decrease in VFF at relatively high metallicities (from a VFF of $10^{-2.5}$ to $10^{-3.2}$) than Case B does (which has a minimum VFF of $10^{-2.9}$) due to the higher initial densities and lower initial temperatures of polluted regions in Case A. Figure 8.1 shows snapshots of the metal distribution in the simulation volume taken at 4 different redshifts, and Figure 8.2 shows the volume filling factor for three metallicity thresholds as a function of redshift.

A uniform metagalactic ultraviolet background is switched on at $z = 7$. Photoheating raises the mean temperature of the baryon gas. In the Lyman- α forest overdensity regime ($1 \leq \delta \leq 10$), which roughly corresponds to the filamentary structure observed in Figure 8.1, the temperature is raised to $\sim 10,000$ K. The local thermal speed of the baryon gas then exceeds the escape velocity of the filaments, resulting in significant expansion of the volume occupied by the gas in those filaments. This includes the gas polluted by metals. This effect can be clearly seen in the third and fourth panels of Figure 8.1. Figure 8.2 shows the sharp increase in metal VFF in a more quantitative way. Gas




Figure 8.2: Volume filling factor. The lines describe the fraction of the simulation volume filled to a metallicity of at least $10^{-6} Z_{\odot}$ (black), $10^{-4} Z_{\odot}$ (blue) and $10^{-3} Z_{\odot}$ (red). Panel (a) corresponds to the simulation where spheres of uniform metal density are added and no other changes are made. Panel (b) corresponds to the simulation where, in addition to uniform spheres of metal density, the baryon density in the corresponding volume is smoothed to the mean density of the simulation and the gas temperature of the sphere is raised to 10^4 K.

which has been polluted above a metallicity of $10^{-6}Z_{\odot}$ (corresponding to essentially the entire volume of gas polluted by metals) increases in VFF to 0.048 for both Case A and 0.046 for Case B. This corresponds to 28% (26%) of the baryon mass being enriched for Case A (Case B). The two cases are essentially indistinguishable with regards to the total mass and volume of gas polluted by metals. Examination of gas with higher metallicity ($Z \geq 10^{-3}Z_{\odot}$) shows some difference between the cases, with a maximum VFF of 5.2×10^{-3} for Case A and 5.8×10^{-3} for Case B (corresponding to 1.5% and 1.9% of the total gas mass, respectively). This difference is due to the initial smoothing and heating of baryons in Case B.

In Figure 8.3 we estimate the amount of carbon contained within the primordial metallicity field at $z = 3$. We assume that the carbon abundance in the metal density is equal to $X_C = 0.027$, which is taken from the supernovae metal yield of a massive primordial star of approximately $260 M_{\odot}$ as computed by Heger et al. [?]. We then compute the mean and median carbon metallicity, $[C/H] = \log(n_C/n_H) - [C/H]_{\odot}$ (where $[C/H]_{\odot} = -3.45$) in bins of constant logarithmic overdensity between $\log \delta = -0.5 - 2.0$ and plot this in Panel A of Figure 8.3. Altering the effects of X_C results in this figure being scaled along the y-axis by a factor of $\log(X_C/0.027)$. The solid green line is the fit to the observations of Schaye et al. [84] using their fiducial model. The dashed-green line is the fit of the lower bound lognormal scatter of their data. The results of our simulation yield that the Population III carbon content in the IGM at $z = 3$ is below the observed limits across the entire overdensity range.

In Panel B of Figure 8.3 we plot the probability distribution function (PDF) for both simulation cases in the overdensity range $\log \delta = -0.5$ to $+2.0$. The vertical lines correspond to the mean (solid) and median (dashed) values of the $[C/H]$ values. The distributions corresponding to the two cases are statistically indistinguishable within one standard deviation. The small variation between the two cases in the range $[C/H]=-4$ to -2 is due to the difference in their mean initial metallicity per bubble (caused by the difference in treatment of of the baryon density field in the region initially polluted by metals in the two cases).

In order to determine the observability of the primordial metallicity field, we post-process our data to compute the fraction of CIV to neutral hydrogen (HI) for each cell in our computational volume. Obtaining the quantity $\log F = \log(n_{CIV}/n_{HI})$ within constant overdensity bins, in the range $\log \delta = -0.5 - 2.0$, allows the determination of the lognormal average $\langle \log F \rangle$ at each redshift. At $z=3$ our analysis yields $\langle \log F \rangle = -3.78 \pm 0.91 (-3.78 \pm 0.92)$ for Case A (B). We can then approximate the CIV optical depth due to Pop III stars as $\tau_{CIV}^{popIII} \approx \left(\frac{f_{CIV}\lambda_{CIV}}{f_{HI}\lambda_{HI}}\right) \cdot 10^{\langle \log F \rangle} \tau_{HI}$, where $f_{CIV,HI}$ and $\lambda_{CIV,HI}$ are the oscillator strengths and rest-frame absorption wavelengths for the two species (CIV & HI). Using our mean values for $\langle \log F \rangle$ at $z=3$ we obtain an estimated optical depth, due to CIV, of $\tau_{CIV}^{popIII} = 0.97_{0.12}^{7.86} \times 10^{-4} \tau_{HI}$ ($0.97_{0.11}^{8.13} \times 10^{-4} \tau_{HI}$). Schaye et al. [84] computed $\tau_{CIV} = 10^{-2} \tau_{HI}$ in the Ly- α forest overdensity range. If




Figure 8.3: In panel (a) we plot the volume averaged $[C/H]$ in solar units within constant logarithmic overdensity bins in the range of $1 \leq \delta \leq 10^2$ at $z = 3$. The profiles for the two cases discussed in the text show no statistical difference. In panel (b) we plot the probability distribution function (PDF) of $[C/H]$ within the overdensity range $1 \leq \delta \leq 10^2$. The y-axis measures the fraction of the metal polluted volume with $[C/H]$ values between $[C/H]$ and $[C/H] + \Delta[C/H]$. The result is computed with $X_C = 0.027$. The two distributions yield $\langle [C/H] \rangle = -4.68 \pm 0.80$ for Case A and $\langle [C/H] \rangle = -4.68 \pm 0.81$ for Case B. Median values for the two PDFs are -4.53 and -4.56 respectively. The two cases are indistinguishable within one standard deviation.

the above value corresponds to the total metallicity at $z=3$ then the contribution of the primordial component to the total optical depth of CIV is about $\tau_{CIV}^{popIII} = 0.01 \tau_{CIV}$. This result is somewhat sensitive to the shape of the UV background – see Schaye et al. [84] for more details. Statistical correlations between Pop III CIV and HI absorbers and more detailed examination of spectra due to ejecta from Population III stars will be discussed in a forthcoming paper.

8.5 Discussion

In this chapter we use cosmological hydrodynamic simulations to examine the evolution of metals ejected by an early population of massive primordial stars. We show that, in the absence of further star formation, photoevaporation of baryons bound to dark matter filaments during reionization is the most important mechanism in determining the volume filling fraction by $z = 3$. Our two study cases, although different in their initial setup, give the same results for the global distribution of the primordial metal field by $z = 3$, suggesting that our result is insensitive to small-scale dynamics.

Comparison of our results to observations of carbon in the Lyman- α forest by Schaye et al. [84] show that at $z = 3$ the median value of the Pop III carbon metallicity for both cases considered fall within the low end of the scatter range of the observed data for $\log \delta \leq 0$. For $\log \delta \geq 0$ the Population III carbon metallicity is below the observed values, with the Schaye result showing a much stronger increase in metallicity with overdensity, resulting in the median value of $[C/H]^{PopIII}$ becoming an increasingly smaller fraction of the total observed $[C/H]$.

Our results depend strongly on two factors, namely, the total number of Population III stars formed in our volume and the metal yield per star. In these simulations we make the assumption that all halos with mass $M_{DM} \geq 5 \times 10^5 M_{\odot}$ form a massive primordial star by $z = 15$, which was guided by the simulations performed by Abel et al. [?] and Yoshida et al. [228], which show that this is the characteristic dark matter mass of a halo which forms a star in the early universe. Simulations by Machacek et al. [71] and semianalytical calculations by Wise & Abel [97] show that a soft UV background produced by the first Pop III stars effectively dissociates H_2 , which is the primary cooling mechanism in primordial star formation. This so-called negative feedback effect raises the minimum halo mass that can form a primordial star within it and therefore reduces the number of halos which will form Population III stars at a given epoch. Wise & Abel [97] find that negative feedback reduces the number of star forming halos by a factor of 5-10 relative to what we used. On the other hand, suppression of Pop III star formation by negative feedback would be compensated by an extended epoch of Pop III star formation. At present, we do not know when Pop III star formation ceases. We view our choice of $M_{min} = 5 \times 10^5 M_{\odot}$ at $z = 15$ as a hedge between competing effects.

The decision to place spheres of metal in the simulation volume at $z = 15$ was guided primarily by the WMAP polarization results [64]. This choice may have resulted in an underestimation of the number of Population III stars (and therefore metal pollution due to Pop III supernovae) because there are dark matter halos which form after $z = 15$ but may still be unpolluted by metals. However, results by Bromm et al. [93] suggest the existence of a “critical metallicity” of $\sim 5 \times 10^{-3} Z_{\odot}$ above which a solar IMF dominates, and it has been argued that this metallicity is reached by $z \sim 15 - 20$ [232, 233]. The choice of $z = 15$ for our epoch of instantaneous metal enrichment seems to be a reasonable compromise.

The physical properties of the metal “bubbles” can have a possible effect on our results. The choice of a 1 kpc (proper) radius for the metal bubbles is somewhat arbitrary. Several calculations have been performed that suggest that ejecta from the most massive pair-instability supernovae can propagate to large distances [95, 234], but the maximum propagation distance is unclear. Additionally, Bromm et al. [95] suggest that the ejecta from pair-instability supernovae still has substantial peculiar velocities (~ 50 km/s) at 500 pc. The metal spheres in this calculation have no initial outward peculiar velocity, which may result in a smaller volume filling factor than if this were taken into account.

The second factor that strongly affects our result is the choice of the amount of metals created per Population III supernova. Abel et al. [39] and Bromm et al. [42] both suggest that the first population of stars will be very massive. The mass function of the first generation of stars is unclear, due to lack of resolution and appropriate physics. The main-sequence mass of the star strongly affects its ultimate fate: Stars with the range of $\sim 140 - 260 M_{\odot}$ detonate in pair instability supernovae, which are much more energetic (up to $\sim 10^{53}$ ergs compared to 10^{51} ergs for a standard Type I or Type II supernova) and produce more metal (up to $57 M_{\odot}$ of ^{56}Ni and almost $130 M_{\odot}$ of total metals for a $260 M_{\odot}$ primordial star). However, stars between $\sim 50 - 140 M_{\odot}$ and above $\sim 260 M_{\odot}$ form black holes without first ejecting significant quantities of nucleosynthesized material [61]. The amount of metals placed into the simulation volume is scalable - if the mean amount of metals ejected by Population III stars were lower (due to some substantial fraction collapsing directly into black holes, for instance), all of the results shown in Figure 8.2 and Panel A of Figure 8.3 scale linearly with the mean amount of metal produced per star.

Our results for $[\text{C}/\text{H}]$ vs. overdensity (using carbon as a proxy for metallicity) agree with the results of Schaye et al. [84] to within one standard deviation for the lowest observed densities ($\log \delta < 0$). These are the densities that are the most likely to remain unpolluted by later generations of stars, which form in deeper potential wells. Further study of the lower density regions of the Lyman- α forest could yield more constraints on the mass and total number density of massive Population III stars.

An additional factor to consider is that the nucleosynthetic yields of very massive primordial stars are much different than that of metal-enriched stars [61]. Due to this,

it may be possible to disentangle the effects of massive primordial stars and their metal-polluted descendants, as discussed by Oh et al. [235].

Due to our choices of low minimum halo mass and high metal yield per supernova, our result is a strong upper limit on the pollution of the Lyman- α forest due to Population III stars, unless we have severely underestimated the duration of the Population III epoch.

The simulation volume is relatively small. Though a reasonable statistical representation of the universe at $z = 15$, the results obtained at later times ($z \sim 3$) should be considered qualitative due to the small box size. A much larger simulation volume is required for adequate statistics at $z \sim 3$. However, simulating a much larger volume which would still have reasonable spatial and dark matter mass resolution on a single grid is computationally prohibitive at the present time.

All of the simulation results described in this paper are performed without further star formation or feedback. Having a single episode of star formation at $z = 15$ means that metal evolution after that time is passive, whereas in reality there would be continuous star formation and feedback. A logical extension of this work is the inclusion of later epochs of star formation and their resulting feedback of metals and energy into the IGM. These results will be presented in a forthcoming paper.

Chapter 9

Summary and future Work

9.1 Summary

Population III stars play an important role in the formation of large-scale structure in the universe through their feedback of metals, kinetic energy, and radiation, and are possible seeds for the super massive black holes (SMBHs) that are observed at the centers of most massive galaxies. These stars are the first luminous objects to form after the Big Bang, and are believed to play a significant role in the ensuing generations of star formation by preprocessing large volumes of the universe with ionizing radiation and metals, which greatly enhance the ability of gas to cool, radically changing the properties of the gas out of which later generations of stars form. Their critical role in structure formation in the early universe makes an understanding of the Population III mass function and the feedback properties of these stars crucial to be able to model the formation of the first generation of galaxies from first principles. At present there are no direct observational constraints on the mass function of Population III stars. Some indirect evidence has been obtained from observations of extremely metal poor stars in the galactic halo, the extragalactic infrared background, and measurements of polarization in the cosmic microwave background. Though future observations by facilities such as the Square Kilometer Array (SKA), the Low Frequency Array (LOFAR) and the James Webb Space Telescope (JWST) may directly observe these stars, at present the only way to study their properties directly is by the use of numerical simulations.

In this thesis I have used the adaptive mesh refinement cosmology code Enzo (described in detail in Chapter 2) to study aspects of the formation and feedback processes of the first generation of stars in the universe. Chapter 3 describes the results of an extensive comparison of Enzo with the smoothed particle hydrodynamics (SPH) code GADGET. This comparison, in the context of cosmological simulations of galaxy formation, shows that it is possible to achieve very similar results with methodologies that are extremely different when one is performing simulations with only dark matter, and also

when one includes non-radiative (“adiabatic”) hydrodynamics. This comparison helps to verify the correctness of both codes and also lends credibility to simulation results performed in regimes where direct observational evidence is not currently available (e.g. structure formation in the early universe).

Chapter 4 presents results from high dynamical range simulations of the formation of Population III stars in a Λ CDM universe. I performed simulations varying both the simulation box sizes and random seeds (effectively changing the pattern of large scale structure in the simulation), as well as choosing a single realization and varying simulation parameters, in order to obtain further constraints on the formation and ultimate stellar mass of Population III stars. I show that the mean formation redshift of these stars, as well as the overall accretion rate onto the primordial protostar, varies systematically with the volume of the cosmological simulation, with larger box sizes tending towards lower overall accretion rates, and in principle lower overall stellar masses. The implications of this are significant – the mass estimates of Population III stars that are most commonly used in the literature have been based primarily upon the results of three-dimensional cosmological simulations using very small simulation volumes, and as a result much work has been based on assumptions of a Population III mass range that may be more massive than is implied by this work. The range of supernova energies, as well as nucleosynthetic yields, varies significantly between the inferred mass range for the largest simulation volumes described in this work ($\sim 10\text{--}100 M_{\odot}$) and from previous work ($\sim 30\text{--}300 M_{\odot}$).

Chapter 4 also presents results from simulations of the formation of Population III stars assuming a constant Lyman-Werner background. This soft ultraviolet background will photodissociate molecular hydrogen, the primary coolant in primordial gas. I demonstrate that increasing the strength of this UV background will delay the onset of Population III star formation, and may even completely prevent the formation of primordial stars in halos whose virial temperatures are less than $\sim 10^4$ K.

In Chapter 5, I discuss results from simulations of the formation of Population III stars assuming a simple warm dark matter (WDM) model. This model is quite generic in that it makes no assumptions about the mechanism which suppresses power on small physical scales. I apply suppression to the power spectrum over a range of assumed warm dark matter particle masses in order to present a more reasonable lower limit on a possible warm dark matter particle mass, and also run a cold dark matter version as a control sample. This suppression is applied on multiple simulations which all have the same large scale structure (all other simulation parameters are also identical) so that we can directly compare results. I show that simulations of the formation of Population III stars in a warm dark matter cosmology are effectively identical to the results from the fiducial Λ CDM case at $m_{WDM} \simeq 35$ keV, and that observations of polarization in the cosmic microwave background (which imply partial reionization of the universe at $z \simeq 17 \pm 5$) suggest a lower limit on the warm dark matter particle mass of ~ 15 keV, which is a factor of ~ 3 greater than previous published lower limits. This significantly

tightens constraints on a possible warm dark matter particle mass.

I present results of the formation of a second generation of Population III stars in a halo which has been ionized by the formation of a nearby extremely massive primordial star in Chapter 6. This scenario is a likely one, as $\sim 10^6 M_\odot$ halos are highly clustered in the early universe, and it is plausible that several nearby halos will be ionized by the first Pop III star to form in a given cluster of halos, assuming that this star is very massive (at least $\sim 100 M_\odot$). In this chapter I show that the ionization of gas due to the first star to form in a region actually enhances the formation of molecular hydrogen in nearby halos (after the death of the first star) and allows star formation to take place in a halo that would not otherwise form a Population III. This suggests a possible positive feedback effect due to Population III star formation. Additionally, the star that forms in the ionized halo has a much lower accretion rate than the first star to form in the simulation volume, which is most likely due to the higher overall angular momentum of the halo that this star forms in. This lower accretion rate implies a relatively low stellar mass ($\sim 5 - 20 M_\odot$) and presents an additional plausible scenario for the formation of Population III stars with masses lower than the “fiducial” range of $30 - 300 M_\odot$ that is typically assumed in the literature.

Chapter 7 presents very preliminary results of adaptive mesh refinement simulations of the supernovae from Population III stars. We examine the enrichment of metals in a cosmological context by applying a standard Sedov-Taylor supernova solution to the core of a halo in which a Population III protostar forms for two representative stellar masses, 30 and $250 M_\odot$, which effectively bracket the range of possible Population III supernova energies ($10^{51} - 10^{53}$ ergs). The less massive star is capable of spreading metals throughout a region almost a proper kiloparsec across (at $z \sim 18$), and it seems reasonable that the more energetic supernova will have a proportionally greater effect. Initial analysis of the $30 M_\odot$ supernova simulation indicates that though a very large mass of gas is enriched by the supernova, the level of enrichment varies widely. It is unclear at present if the cores of neighboring halos (which are presumably the sites of the next generation of star formation) are appreciably enriched with metals. Though these results are qualitative at present, this method is a very promising approach to determine the metallicities of the first generation of metal-enriched stars.

The results of simulations which model the feedback of metals from a large number of Population III stars are presented in Chapter 8. This work, which attempts to model the effects of Population III stars on the overall metallicity of the Lyman- α forest, assumes a very optimistic scenario for Population III metal enrichment: every halo with a dark matter mass of at least $5 \times 10^5 M_\odot$ or higher is assumed to form a Population III star with a mass of $260 M_\odot$ at $z = 15$. These stars are assumed to produce $\sim 130 M_\odot$ of metals which are spread uniformly over a sphere one proper kiloparsec in radius. The evolution is then followed to $z = 3$ and the metallicity of the IGM is compared to current observations of the Lyman- α forest. This work shows that, even with an

extremely optimistic scenario, it is impossible to explain the metallicity of the Lyman- α forest (as inferred by measurements of CIV absorption lines) at densities higher than the cosmological mean density by metal enrichment from Population III stars alone.

The work presented in this thesis, and summarized in the previous paragraphs, contributes directly to the greater understanding of the formation of Population III stars and the effects that they have on their surroundings via radiative, chemical, and mechanical feedback. This work is critical to the accurate modeling of the formation and evolution of galaxies from first principles, since Population III stars are responsible for preprocessing regions of space which will eventually become galaxies. Additionally, Population III stars may be the seeds of the super massive black holes that are observed at the centers of many galaxies. Thus, understanding the mass range of these stars is also significant for research in both the formation and evolution of super massive black holes, as well as work with active galactic nuclei (AGN), which are believed to be powered by these black holes. Finally, understanding the Pop III IMF and the resulting feedback from these stars can be useful in interpreting observations the extragalactic infrared background as well as the recent WMAP observations of polarization in the cosmic microwave background. Population III stars may also be the progenitors of extremely high redshift gamma ray bursts (GRBs), and predictions of the IMF may place constraints on the frequency of Pop III GRB events which are observable by the current and future generations of gamma ray satellites.

9.2 Future Work

The work presented in this thesis is extendible in many ways. The code comparison project in Chapter 3 discusses simulations that have dark matter only, or have dark matter with a non-radiative gas component. While an important first step, galaxy formation (which is the context of the comparison) is heavily dependent upon the cooling processes of gas in the collapsing dark matter halos as well as the formation of stars and their feedback of metal and thermal energy into the interstellar and intergalactic medium. An obvious extension of this project would be to incrementally include radiative cooling, star formation, and then star formation plus feedback, with the goal at each step of understanding and attempting to reconcile the differences between the two simulation methods. A portion of this work (simulations that include radiative cooling) is already in progress.

The study of Population III stars in a Λ CDM universe discussed in this thesis (Chapter 4) is limited by physics when the gas in the dark matter halo that the stars formed in reaches approximately 10^{12} particles per cubic centimeter. At this point one of the main assumptions concerning the cooling in the gas, namely, that the gas is optically thin to cooling radiation, becomes invalid. Recent work by Ripamonti & Abel [47] has

shown that for several more orders of magnitude in density the effects of opacity are local, meaning that minor corrections to the chemical rate equations are all that is necessary to accurately model the collapsing cloud. These fixes are already implemented into Enzo, and are a logical extension of the work presented here. In addition, all of the work in this thesis completely ignores the effects of magnetic fields. While it is widely assumed that B-fields are dynamically unimportant in the formation of Population III stars, this can be tested. The equations of magnetohydrodynamics are currently being implemented into the Enzo code, and the examination of the effect of magnetic fields on the formation of a Pop III protostar will be one of the first applications of this code. Finally, it is unclear if the observed relationship between increased box size and decreased protostellar accretion rate has converged at the largest box size studied in this work. Extension of the study to even larger box sizes would help to determine whether this result is converged.

The simulations of the collapse of a halo in a cosmology with warm dark matter (Chapter 5) are of a single cosmological realization. A straightforward (though computationally expensive) extension to make the predictions in Chapter 5 more robust would be to perform the same set of simulations for multiple random realizations and for several different box sizes, as was done in Chapter 4. This would serve to reduce issues involving small-number statistics and would also provide an attempt to disentangle the effects of the suppression of the power spectrum and issues involving the box size. Similarly, the results involving the formation of a second generation of primordial stars (Chapter 6) should be extended by performing multiple random realizations in order to increase the statistical robustness of the results. The HII region models that are used are also very simple, and rely on the one-dimensional output from a radiation transport code. This presents an inherent problem, since the universe is manifestly three-dimensional! Also, to accurately model the evolution of the HII region (particularly when the ionization front encounters nearby halos) it is important to model the time-dependent three dimensional evolution of the HII region rather than using a static density field, since hydrodynamic effects may be significant, particularly in the neighboring halos which are of most interest in this situation. Unfortunately, this is currently beyond the capabilities of our cosmology code (not to mention incredibly expensive computationally, even if the tools existed). However, one can perform computationally feasible models in two dimensions and possibly even in somewhat idealized three dimensional situations that will lend insight to the problem.

The simulations of Pop III supernovae shown in Chapter 7 are first attempts, and can stand to be improved in many ways. In particular, the models that we are using for the supernovae are somewhat simplistic, being Sedov-Taylor blast waves. Work is in progress to take the output from three dimensional models of supernova explosions and use this as the initial conditions for cosmological simulation. This work is already in progress. Additionally, the metal-enriched gas from these supernovae would have significantly enhanced cooling compared to gas of primordial composition. Implementing cooling models that

include the cooling effects from metals would allow one to study the enhanced fragmentation of enriched gas, and improve our estimates of the properties of second generation metal-enriched stars. Finally, the existence of several grid-related numerical artifacts suggests that these calculations are somewhat under-resolved. Improving the spatial resolution of the calculations, though computationally costly, is important if we wish to accurately model the mixing of metals into the intergalactic medium and in neighboring halos, which is crucial to the understanding the properties and metal distribution of the first generation of metal-enriched stars.

Finally, the simulations of the pre-galactic enrichment of the IGM by Population III stars discussed in Chapter 8 can be extended in several significant ways. The simulation volumes used in the study presented are somewhat small, and it is computationally feasible to increase the size of the volume to 1024^3 cells at the same physical resolution, allowing us to increase the overall volume by a factor of 64 and significantly improving our statistics. In addition, it would be useful to explore scenarios where the simulations are enriched by Pop III stars at multiple epochs rather than only once. Finally, since we know that star formation took place after the Population III epoch was over, it would be useful to do similar calculations with the effects of post-Pop III star formation included so that one can understand their relative effects on metal enrichment of the IGM, as well as to explore the possibility of differences in QSO spectra that might allow observers to attempt to separate Lyman- α clouds which have been enriched only by Population III stars from those which have been enriched by a combination of Pop III and later generations of star formation. Much of this work is already in progress, and the results will be presented in a forthcoming paper.

Appendix A

The Enzo Primordial Chemistry Reaction Network

The primordial chemistry network implemented in Enzo is discussed in Section 2.2.5 and also in much more detail by Abel et al. [152] and Anninos et al. [153]. These papers describe the chemistry and cooling behavior of low-density primordial gas ($n \simeq 10^4$ and below), as well as the steps that are necessary to obtain fast and accurate numerical solutions of the nonequilibrium chemical reaction rate network. This network is extended by Abel, Bryan & Norman [39] to include the 3-body molecular hydrogen creation process, which becomes important at higher densities, and extends the validity of the reaction network several more orders of magnitude in density, essentially until the gas becomes optically thick to cooling by H_2 line emission. For the sake of completeness, and because the properties the primordial gas are so crucial to the results that are discussed in this work, we describe the chemistry network in this appendix.

Tables 2.1 and 2.2 summarize the collisional and radiative processes solved in the Enzo nonequilibrium chemistry routines. Abel et al. [152] show that accurate results can be obtained if several unnecessary reactions are eliminated and the reaction network is reduced to the following:

$$\frac{dn_{\text{H}}}{dt} = k_2 n_{\text{H}^+} n_{\text{e}} - k_1 n_{\text{H}} n_{\text{e}} + 2k_{31} n_{\text{H}_2} \quad (\text{A.1})$$

$$\frac{dn_{\text{H}^+}}{dt} = k_1 n_{\text{H}} n_{\text{e}} - k_2 n_{\text{H}^+} n_{\text{e}} \quad (\text{A.2})$$

$$\frac{dn_{\text{He}}}{dt} = k_4 n_{\text{He}^+} n_{\text{e}} - k_3 n_{\text{He}} n_{\text{e}} \quad (\text{A.3})$$

$$\frac{dn_{\text{He}^+}}{dt} = k_3 n_{\text{He}} n_{\text{e}} + k_6 n_{\text{He}^{++}} n_{\text{e}} - k_4 n_{\text{He}^+} n_{\text{e}} \quad (\text{A.4})$$

$$\frac{dn_{\text{He}^{++}}}{dt} = k_5 n_{\text{He}^+} n_{\text{e}} - k_6 n_{\text{He}^{++}} n_{\text{e}} \quad (\text{A.5})$$

$$\frac{dn_{\text{H}_2}}{dt} = k_8 n_{\text{H}^-} n_{\text{H}} + k_{22} n_{\text{H}}^3 - n_{\text{H}_2} (k_{31} + k_{11} n_{\text{H}^+} + k_{12} n_{\text{e}}), \quad (\text{A.6})$$

where the number density of H^- is given by the equilibrium condition

$$n_{\text{H}^-} = \frac{k_7 n_{\text{H}} n_{\text{e}}}{k_8 n_{\text{H}} + k_{16} n_{\text{H}^+} + k_{14} n_{\text{e}}}. \quad (\text{A.7})$$

The H^- number density can be calculated in equilibrium because the timescale at which the reactions controlling its number density occur are much shorter than the rest of the reactions in this system. The rate coefficients used in the equations above are defined as follows:

$$\begin{aligned} k_1 = & \exp[-32.71396786 + 13.536556 \ln(T) - 5.73932875 \ln(T)^2 \\ & + 1.56315498 \ln(T)^3 - 0.2877056 \ln(T)^4 + 3.48255977 \times 10^{-2} \times \ln(T)^5 \\ & - 2.63197617 \times 10^{-3} \times \ln(T)^6 + 1.11954395 \times 10^{-4} \ln(T)^7 \\ & - 2.03914985 \times 10^{-6} \ln(T)^8] \text{ cm}^3 \text{ s}^{-1} \end{aligned} \quad (\text{A.8})$$

$$\begin{aligned} k_2 = & \exp[-28.6130338 - 0.72411256 \ln(T) - 2.02604473 \times 10^{-2} \ln(T)^2 \\ & - 2.38086188 \times 10^{-3} \ln(T)^3 - 3.21260521 \times 10^{-4} \ln(T)^4 \\ & - 1.42150291 \times 10^{-5} \ln(T)^5 + 4.98910892 \times 10^{-6} \ln(T)^6 \\ & + 5.75561414 \times 10^{-7} \ln(T)^7 - 1.85676704 \times 10^{-8} \ln(T)^8 \\ & - 3.07113524 \times 10^{-9} \ln(T)^9] \text{ cm}^3 \text{ s}^{-1} \end{aligned} \quad (\text{A.9})$$

$$\begin{aligned} k_3 = & \exp[(-44.09864886 + 23.91596563 \ln(T) - 10.7532302 \ln(T)^2 \\ & + 3.05803875 \ln(T)^3 - 0.56851189 \ln(T)^4 + 6.79539123 \times 10^{-2} \ln(T)^5 \\ & - 5.00905610 \times 10^{-3} \ln(T)^6 + 2.06723616 \times 10^{-4} \ln(T)^7 \\ & - 3.64916141 \times 10^{-6} \ln(T)^8) \text{ cm}^3 \text{ s}^{-1} \end{aligned} \quad (\text{A.10})$$

$$k_{4r} = 3.925 \times 10^{-13} T^{-0.6353} \text{ cm}^3 \text{ s}^{-1} \quad (\text{A.11})$$

$$\begin{aligned} k_{4d} = & 1.544 \times 10^{-9} T^{-\frac{3}{2}} \exp\left(-\frac{48.596 \text{ eV}}{T}\right) \times \\ & \left[0.3 + \exp\left(\frac{8.10 \text{ eV}}{T}\right)\right] \text{ cm}^3 \text{ s}^{-1} \end{aligned} \quad (\text{A.12})$$

$$\begin{aligned}
k_5 = & \exp[-68.71040990 + 43.93347633 \ln(T) - 18.4806699 \ln(T)^2 \\
& + 4.70162649 \ln(T)^3 - 0.76924663 \ln(T)^4 + 8.113042 \times 10^{-2} \ln(T)^5 \\
& - 5.32402063 \times 10^{-3} \ln(T)^6 + 1.97570531 \times 10^{-4} \ln(T)^7 \\
& - 3.16558106 \times 10^{-6} \ln(T)^8] \text{ cm}^3 \text{ s}^{-1}
\end{aligned} \tag{A.13}$$

$$k_6 = 3.36 \times 10^{-10} T^{-\frac{1}{2}} \left(\frac{T}{1000 \text{ K}} \right)^{-0.2} \left(1 + \left(\frac{T}{10^6 \text{ K}} \right)^{0.7} \right)^{-1} \text{ cm}^3 \text{ s}^{-1} \tag{A.14}$$

k_7 for $T \leq 6000 \text{ K}$:

$$k_7 = 1.429 \times 10^{-18} T^{0.7620} T^{0.1523 \log_{10}(T)} T^{-3.274 \times 10^{-2} \log_{10}^2(T)} \text{ cm}^3 \text{ s}^{-1} \tag{A.15}$$

k_7 for $T > 6000 \text{ K}$:

$$\begin{aligned}
k_7 = & 3.802 \times 10^{-17} T^{0.1998 \log_{10}(T)} \\
& \text{dex} \left(4.0415 \times 10^{-5} \log_{10}^6(T) - 5.447 \times 10^{-3} \log_{10}^4(T) \text{ cm}^3 \text{ s}^{-1} \right)
\end{aligned} \tag{A.16}$$

$$\begin{aligned}
T > 0.1 \text{ eV} : k_8 = & \exp[-20.06913897 + 0.22898 \ln(T) + 3.5998377 \\
& \times 10^{-2} \ln(T)^2 - 4.55512 \times 10^{-3} \ln(T)^3 - 3.10511544 \times 10^{-4} \ln(T)^4 \\
& + 1.0732940 \times 10^{-4} \ln(T)^5 - 8.36671960 \times 10^{-6} \ln(T)^6 + 2.23830623 \\
& \times 10^{-7} \ln(T)^7] \text{ cm}^3 \text{ s}^{-1}.
\end{aligned} \tag{A.17}$$

$$T < 0.1 \text{ eV} : k_8 = 1.428 \times 10^{-9} \text{ cm}^3 \text{ s}^{-1} \tag{A.18}$$

$$\begin{aligned}
\ln(k_{11}) = & -24.24914687 + 3.40082444 \ln(T) - 3.89800396 \ln(T)^2 \\
& + 2.04558782 \ln(T)^3 - 0.541618285 \ln(T)^4 + 8.41077503 \times 10^{-2} \ln(T)^5 \\
& - 7.87902615 \times 10^{-3} \ln(T)^6 + 4.13839842 \times 10^{-4} \ln(T)^7 \\
& - 9.36345888 \times 10^{-6} \ln(T)^8 \text{ cm}^3 \text{ s}^{-1}
\end{aligned} \tag{A.19}$$

$$k_{12} = 5.6 \times 10^{-11} T^{\frac{1}{2}} \exp\left(-\frac{102,124K}{T}\right) \text{cm}^3 \text{s}^{-1} \quad (\text{A.20})$$

$$\begin{aligned} k_{14} = & \exp[-18.01849334 + 2.3608522 \ln(T) - 0.28274430 \ln(T)^2 \\ & + 1.62331664 \times 10^{-2} \ln(T)^3 - 3.36501203 \times 10^{-2} \ln(T)^4 \\ & + 1.17832978 \times 10^{-2} \ln(T)^5 - 1.65619470 \times 10^{-3} \ln(T)^6 \\ & + 1.06827520 \times 10^{-4} \ln(T)^7 - 2.63128581 \times 10^{-6} \ln(T)^8] \text{cm}^3 \text{s}^{-1} \end{aligned} \quad (\text{A.21})$$

$$k_{16} = 7 \times 10^{-8} \left(\frac{T}{100\text{K}}\right)^{-\frac{1}{2}} \text{cm}^3 \text{s}^{-1} \quad (\text{A.22})$$

k_{22} For $T < 300$ K:

$$k_{22} = 1.3 \times 10^{-32} (T/300\text{K})^{-0.38} \text{cm}^6 \text{s}^{-1} \quad (\text{A.23})$$

k_{22} for $T \geq 300$ K:

$$k_{22} = 1.3 \times 10^{-32} (T/300\text{K})^{-1} \text{cm}^6 \text{s}^{-1} \quad (\text{A.24})$$

The photodissociation of molecular hydrogen by Lyman-Werner radiation (as discussed in Section 4.4.4) is controlled by the k_{31} parameter:

$$k_{31} = 1.13 \times 10^8 F_{\text{LW}} |t| \quad (\text{A.25})$$

where $|t| \equiv (4\pi G \Omega_m \rho_c (1+z_i)^3)^{-1/2}$ and has units of seconds, and F_{LW} is the Lyman-Werner background flux in units of $\text{erg s}^{-1} \text{cm}^{-2} \text{Hz}^{-1}$.

Bibliography

- [1] Norman, M. L., O'Shea, B. W., & Paschos, P. 2004, ApJ, 601, L115
- [2] O'Shea, B.W., Bryan, G., Bordner, J., Norman, Michael L., Abel, T., Harkness, R. and Kritsuk, A. 2004, In "Adaptive Mesh Refinement - Theory and Applications", Eds. T. Plewa, T. Linde & V. G. Weirs, Springer Lecture Notes in Computational Science and Engineering
- [3] O'Shea, B. W., Abel, T., Whalen, D., & Norman, M. L. 2005, ApJ, 628, L5
- [4] O'Shea, B. W., Nagamine, K., Springel, V., Hernquist, L. & Norman, M.L. 2005, ApJ, in press. astro-ph/0312651
- [5] The American Heritage Dictionary of the English Language, Fourth Edition. (2000) Houghton Mifflin Co.
- [6] Hubble, E. 1929, Proceedings of the National Academy of Science, 15, 168
- [7] Penzias, A.A. & Wilson, R.W. 1965, ApJ, 142, 419
- [8] Dicke, B.H. & Peebles, P.J.E. 1965, ApJ, 142, 414
- [9] Gamow, G. 1970. Viking Press, New York, New York, U.S.A. "My world line"
- [10] Misner, C.W., Thorne, K.S. & Wheeler, J.A. 1973, W.H. Freeman and Company, New York, New York, USA, "Gravitation"
- [11] Weinberg, S. 1972, John Wiley & Sons, New York, New York, "Gravitation and Cosmology"
- [12] Peacock, J.A. 1999, Cambridge University Press, Cambridge, "Cosmological Physics"
- [13] Peebles, P.J.E. 1993, Princeton University Press, Princeton, New Jersey, "Principles of Physical Cosmology"

- [14] Kolb, E.W. & Turner, M.S. 1990, Perseus Books, Reading, Massachusetts, “The Early Universe”
- [15] Press, W.H. & Schechter, P. 1974, ApJ, 187, 425
- [16] White, S.D.M. & Frenk, C.D. 1991, ApJ, 379, 52
- [17] Mo, H.J. & White, S.D.M. 2002, MNRAS, 336, 112
- [18] Bertschinger, E. 1998, ARAA, 36, 599
- [19] Copi, C.J., Schramm, D.N. & Turner, M.S. 1995, Phys.Rev.Lett., 75, 3981
- [20] Palla, F., Salpeter, E.E., & Stahler, S.W. 1983, ApJ, 271, 632
- [21] Abel, T. & Haiman, Z. 2000 Invited review in *H₂ in Space*, Paris, France 1999, ed. F. Combes and G. Pineau des Forets
- [22] Loeb, A. & Haiman, Z. 1997, ApJ, 490, 571
- [23] Miralda-Escude, J. & Rees, M.J. 1997, ApJ, 478, 57
- [24] Peebles, P.J.E. & Dicke, R.H. 1968, ApJ, 154, 891
- [25] Hirasawa, T. 1969, Prog. Theor. Phys., 42, 523
- [26] Silk, J. 1983, MNRAS, 205, 705
- [27] Galli, D. & Palla, F. 1998, Astronomy & Astrophysics, 335, 403
- [28] Tegmark, M., Silk, J., Rees, M. J., Blanchard, A., Abel, T., & Palla, F. 1997, ApJ, 474, 1
- [29] Jedamzik, K., Katalinik, V., Olinto, A.V. 2000, Phys.Rev.Lett., 85, 700
- [30] Kahniashvili, T. et al. 2000, In *Cosmology and Particle Physics*, eds. J. Garcia-Bellido, R. Durrer, and M. Shaposhnikov
- [31] Sigl, G., Olinto, A.V. & Jedamzik, K. 1997, Phys. Rev. D, 55, 4582
- [32] Hogan, C. 2000, astro-ph/0005380
- [33] Matarrese, S., Mollerach, S., Notori, A. & Riotto, A. 2005, Phys. Rev. D, 71, 043502
- [34] Harrison, E.E. 1970, MNRAS, 147, 279
- [35] Kulsrud, R.M. 1999, ARAA, 37, 37

- [36] Kulsrud, R.M., Cen, R., Ostriker, J.P., & Ryu, D. 1997, *ApJ*, 480, 481
- [37] Gnedin, N.Y., Ferrara, A. & Zweibel, E.G. 2000, *ApJ*, 539, 505
- [38] Langer, M., Puget, J.L., Aghanim, N. 2003, *Phys. Rev. D*, 67, 3505
- [39] Abel, T., Bryan, G.L. & Norman, M.L. 2002, *Science*, 295, 93
- [40] Ebert, R. 1955, *Zs. Ap.* 217
- [41] Bonnor, W.B. 1956, *MNRAS*, 116, 351
- [42] Bromm, V., Coppi, P.S. & Larson, R.B. 2002, *ApJ*, 564, 23
- [43] Gao, L., White, S.D.M., Jenkins, A., Frenk, C. & Springel, V. 2005, *MNRAS* submitted; astro-ph/0503003
- [44] Reed, D., Bower, R., Frenk, C.S., Gao, L., Jenkins, A., Theuns, T. & White, S.D.M. 2005, *MNRAS* submitted; astro-ph/0504038
- [45] Bromm, V. & Larson, R.B. 2004, *ARAA*, 42, 79
- [46] Ciardi, B., & Ferrara, A. 2005, *Space Science Reviews*, 116, 625
- [47] Ripamonti, E. & Abel, T. 2004, *MNRAS*, 348, 1019-1034
- [48] Bromm, V., & Loeb, A. 2004, *New Astronomy*, 9, 353
- [49] Tan, J. & McKee, C. 2004, *ApJ*, 603, 383
- [50] Tan, J. & Blackman, E.G. 2004, *ApJ*, 603, 401
- [51] Omukai, K. & Nishi, R. 1998, *ApJ*, 508, 141
- [52] Omukai, K. & Palla, F. 2001, *ApJ*, 561, L55
- [53] Omukai, K. & Inutsuka, S. 2002, *MNRAS*, 332, 59
- [54] Omukai, K. & Palla, F. 2003, *ApJ*, 589, 677
- [55] Nakamura, F. & Umemura, M. 1999, *ApJ*, 515, 239
- [56] Nakamura, F. & Umemura, M. 2001, *ApJ*, 548, 19
- [57] Nakamura, F. & Umemura, M. 2002, *ApJ*, 569, 549
- [58] Barkana, R., & Loeb, A. 2001, *Physics Reports*, 349, 125

- [59] Schaerer, D. 2002, *Astronomy & Astrophysics*, 382, 28
- [60] Heger, A., & Woosley, S. E. 2002, *ApJ*, 567, 532
- [61] Heger, A., Fryer, C. L., Woosley, S.E., Langer, N. & Hartmann, D. H. 2003, *ApJ*, 591, 288
- [62] Madau, P. & Rees, M.J. 2001, *ApJ*, 551, L27
- [63] Shapiro, S. 2005, *ApJ*, 620, 59
- [64] Kogut, A., et al. 2003, *ApJS*, 148, 161
- [65] Fan, X. et al. 2000, *AJ*, 120, 1167
- [66] Haiman, Z. & Loeb, A. 1998, *ApJ*, 503, 505
- [67] Gnedin, R. & Ostriker, J.P. 1997, *ApJ*, 486 581
- [68] Gnedin, R. 2000, *ApJ*, 535, 530
- [69] Cen, R. 2003, *ApJ*, 591, 12
- [70] Hui, L., & Haiman, Z. 2003, *ApJ*, 596, 9
- [71] Machacek, M. E., Bryan, G. L., & Abel, T. 2001, *ApJ*, 548, 509
- [72] Haiman, Z., Rees, M.J. & Loeb, A. 1997, *ApJ*, 476, 458
- [73] Oh, S.P. & Haiman, Z. 2002, *ApJ*, 569, 558
- [74] Omukai, K. 2001, *ApJ*, 546, 635
- [75] Omukai, K. & Hoshi, Y. 2003, *ApJ*, 599, 746
- [76] Whalen, D., Abel, T., & Norman, M. L. 2004, *ApJ*, 610, 14
- [77] Oh, S. P., & Haiman, Z. 2003, *MNRAS*, 346, 456
- [78] Nagakura, T. & Omukai, K. 2005, *MNRAS*, submitted; astro-ph/0505599
- [79] Ricotti, M., & Ostriker, J. P. 2004, *MNRAS*, 352, 547
- [80] Ricotti, M., Ostriker, J. P., & Gnedin, N. Y. 2005, *MNRAS*, 357, 207
- [81] Madau, P., Rees, M., Volonteri, M., Haardt, F. & Oh, S.P. 2004, *ApJ*, 604, 484
- [82] Kuhlen, M. & Madau, P. 2005, *MNRAS*, submitted

- [83] Machacek, M. E., Bryan, G. L., & Abel, T. 2003, MNRAS, 338, 273
- [84] Schaye, J., Aguirre, A., Kim, T., Theuns, T., Rauch, M., & Sargent, W. L. W. 2003, ApJ, 596, 768
- [85] Aguirre, A., Schaye, J., Kim, T., Theuns, T., Rauch, M., & Sargent, W. L. W. 2004, ApJ, 602, 38
- [86] Adelberger, K. L., Steidel, C. C., Shapley, A. E., & Pettini, M. 2003, ApJ, 584, 45
- [87] Mac Low, M., & Ferrara, A. 1999, ApJ, 513, 142
- [88] Scannapieco, E. 2005, ApJ, 624, L1
- [89] Todini, P. & Ferrara, A. 2001, MNRAS, 325, 726
- [90] Nozawa, T. & Kozasa, T. 2003, ApJ, 598, 785
- [91] Schneider, R., Ferrara, A., & Salvaterra, R. 2004, MNRAS, 351, 1379
- [92] Herwig, F. 2005, ARAA, in preparation
- [93] Bromm, V., Ferrara, A., Coppi, P. S., & Larson, R. B. 2001, MNRAS, 328, 969
- [94] Bromm, V., & Loeb, A. 2003, Nature, 425, 812
- [95] Bromm, V., Yoshida, N., & Hernquist, L. 2003, ApJL, 596, L135
- [96] Loeb, A. 2003, In Proc. of IAU Colloquium 192 on "Supernovae", April 2003, Valencia, Spain, eds. J. M. Marcaide and K. W. Weiler, astro-ph/0307231
- [97] Wise, J. & Abel, T. 2005, ApJ, submitted; astro-ph/0411558
- [98] Weinmann, S.M. & Lilly, S.J. 2005, ApJ, 624, 526
- [99] Christlieb, N., et al. 2002, Nature, 419, 904
- [100] Frebel, A., et al. 2005, Nature, 434, 871
- [101] Limongi, M., Chieffi, A., & Bonifacio, P. 2003, ApJ, 594, L123
- [102] Wyithe, J. S. B., & Loeb, A. 2004, ApJ, 612, 597
- [103] Cooray, A., & Yoshida, N. 2004, MNRAS, 351, L71
- [104] Madau, P., & Silk, J. 2005, MNRAS, 359, L37

- [105] Navarro, J.F., Frenk, C.S., & White, S.D.M. 1997, *ApJ*, 490, 493
- [106] Moore et al. 1999, *MNRAS*, 310, 1147
- [107] Ghigna et al. 2000, *ApJ*, 544, 616
- [108] Peebles, P.J.E. 2000, in *ASP Conf. Ser.* 200, 377
- [109] Peebles, P.J.E. 2001 preprint [astro-ph/0101127]
- [110] Metcalfe et al. 2001, *MNRAS*, 323, 795
- [111] Bahcall, N.A. et al. 1999, *Science*, 284, 1481
- [112] Klypin et al. 1999, *ApJ*, 522, 82
- [113] Weinberg, M.D. & Katz, N. 2002, *ApJ*, 580, 627
- [114] Cen, R. 2001, *ApJ*, 546, L77
- [115] Bode, P., Ostriker, J.P. & Turok, N. 2001, *ApJ*, 556, 93
- [116] Spergel, D.N. & Steinhardt, P.J. 2000, *PhRvL*, 84, 3760
- [117] Craig, M.W. & Davis, M. 2001, *NewA*, 6, 425
- [118] Hu, W., Barkana, R., & Gruzinov, A. 2000, astro-ph/0003365
- [119] Durrer et al. 2002, *Phys.Rept.*, 1-81
- [120] Guth, A. 1981, *Phys. Rev. D*, 23, 347-356
- [121] Bertschinger, E. 1998, *ARAA*, 36, 599-654
- [122] E. Komatsu et al. 2003, *ApJS*, 148, 119
- [123] Padmanabhan, T. 1993, *Structure Formation In The Universe*, Cambridge University Press, Cambridge, England
- [124] Harrison, E. 1970, *Phys Rev. D.*, 1, 2726-30
- [125] Zel'Dovich, Y. 1972, *MNRAS*, 160, 1-3
- [126] Peebles, P. 1980, Princeton University Press, Princeton, N.J., "The Large Scale Structure of the Universe"
- [127] Zel'Dovich, Y. 1970, *Astron. Astrophys.*, 5, 84-89

- [128] Trac, H. & Pen, U.L. 2004, ApJ submitted; astro-ph/0402444
- [129] Bertschinger, E. 2001, ApJS, 137, 1
- [130] Gelb, J.M. & Bertschinger, E. 1994, ApJ, 436, 491
- [131] Pen, U.L. 1997, ApJ, 490, L127
- [132] Sirko, E. 2005, ApJ submitted; astro-ph/0503106
- [133] Barkana, R. & Loeb, A. 2004, ApJ, 609, 474
- [134] Bryan, G. L. & Norman, M. L. 1997, in Computational Astrophysics; 12th Kingston Meeting on Theoretical Astrophysics, proceedings of meeting held in Halifax; Nova Scotia; Canada October 17-19; 1996, ed. D. A. Clarke & M. Fall (ASP Conference Series # 123)
- [135] Bryan, G. L. & Norman, M. L. 1999, in Workshop on Structured Adaptive Mesh Refinement Grid Methods, ed. N. Chrisochoides (IMA Volumes in Mathematics No. 117), 165
- [136] Norman, M. L. & Bryan, G. L. 1999, in Numerical Astrophysics : Proceedings of the International Conference on Numerical Astrophysics 1998 (NAP98), held at the National Olympic Memorial Youth Center, Tokyo, Japan, March 10-13, 1998., ed. S. M. Miyama, K. Tomisaka, & T. Hanawa (Kluwer Academic), 19
- [137] Bryan, G. L., Abel, T., & Norman, M. L. 2001, in Supercomputing 2001 (IEEE, <http://www.sc2001.org/papers/>)
- [138] Efstathiou, G., Davis, M., White, S. D. M., & Frenk, C. S. 1985, ApJS, 57, 241
- [139] Hockney, R. W. & Eastwood, J. W. 1988, Institute of Physics Publishing, "Computer Simulation Using Particles"
- [140] Berger, M. J. & Colella, P. 1989, J. Comp. Phys., 82, 64
- [141] Pen, U.-L. 1995, ApJS, 100, 269
- [142] Gnedin, N. Y. 1995, ApJS, 97, 231
- [143] Adjerid, S. & Flaherty, J. E. 1998, SIAM J. of Sci. and Stat. Comp., 5, 792
- [144] Woodward, P. R. & Colella, P. 1984, J. Comput. Physics, 54, 174
- [145] Bryan, G. L., Norman, M. L., Stone, J. M., Cen, R., & Ostriker, J. P. 1995, Comp. Phys. Comm, 89, 149

- [146] Stone, J. M. & Norman, M. L. 1992, ApJ, 80, 753
- [147] Stone, J. M. & Norman, M. L. 1992, ApJ, 80, 791
- [148] van Leer, B. 1977, J. Comp. Phys., 23, 276
- [149] Anninos, W. Y. & Norman, M. L. 1994, ApJ, 429, 434
- [150] Sutherland, R.S. & Dopita, M.A. 1993, ApJS, 88, 253-327
- [151] Katz, N., Weinberg, D.H. & Hernquist, L. 1996, ApJ, 105, 19
- [152] Abel, T., Anninos, P., Zhang, Y. & Norman, M.L. 1997, New Astronomy, 2, 181-207
- [153] Anninos, P., Zhang, Y., Abel, T. & Norman, M.L. 1997, New Astronomy, 2, 209-224
- [154] Lipovka, A., Núñez-López, R. & Avila-Reese, V. 2005, MNRAS, submitted; astro-ph/0505599
- [155] Lacey, C., Silk, J. 1991, ApJ, 381, 14
- [156] White, S.D.M. & Rees, M.J. 1978, MNRAS, 183, 341
- [157] Cole, S. 1991, ApJ, 367, 45
- [158] Kaiser, N. 1991, ApJ, 383, 104
- [159] Cavaliere, A., Menci, N., & Tozzi, P. 1997, ApJ, 484, L21
- [160] Kravtsov, A. 2003, ApJ, 590, L1
- [161] Cen, R.Y. and Ostriker, J.P. 1998, ApJ, 399, L113.
- [162] Springel, V. & Hernquist, L. 2003, MNRAS, 339, 312-334
- [163] Springel, V. & Hernquist, L. 2003, MNRAS, 339, 289-311
- [164] Hernquist, L. & Springel, V. 2003, MNRAS, 341, 1253-1267
- [165] Schmidt, M. 1959, ApJ, 129, 243
- [166] Young, J.S., Allen, L., Kenney, J.D.P., Lesser, A. & Rownd, B. 1996, AJ, 112, 1903
- [167] Wong, T. & Blitz, L. 2002, ApJ, 569, 157
- [168] Miller, G.E. & Scalo, J.M. 1979, ApJS, 101, 181
- [169] Woosley, S.E. & Weaver, T.A. 1995, ApJS, 101, 181

- [170] Cen, R. & Ostriker, J.P. 1992, ApJ, 399, L113
- [171] Kang, H., Ostriker, J. P., Cen, R., Ryu, D., Hernquist, L., Evrard, A. E., Bryan, G. L., & Norman, M. L. 1994, ApJ, 430, 83
- [172] Frenk, C. S., White, S. D. M., Bode, P., Bond, J. R., Bryan, G. L., Cen, R., Couchman, H. M. P., Evrard, A. E., Gnedin, N., Jenkins, A., Khokhlov, A. M., Klypin, A., Navarro, J. F., Norman, M. L., Ostriker, J. P., Owen, J. M., Pearce, F. R., Pen, U.-L., Steinmetz, M., Thomas, P. A., Villumsen, J. V., Wadsley, J. W., Warren, M. S., Xu, G., & Yepes, G. 1999, ApJ, 525, 554
- [173] Ascasibar, Y., Yepes, G., Müller, V. & Gottlöber, S. 2003, MNRAS, 346, 731
- [174] Kravtsov, A.V., Klypin, A. & Hoffman, Y. 2002, ApJ, 571, 563
- [175] Springel, V., Yoshida, N., & White, S. D. M. 2001, New Astronomy, 6, 79
- [176] Springel, V. 2005, MNRAS submitted (astro-ph/0505010)
- [177] Hernquist, L. & Katz, N. 1989, ApJ, 70, 419
- [178] Navarro, J. F. & White, S. D. M. 1993, MNRAS, 265, 271
- [179] Katz, N., Weinberg, D. H., & Hernquist, L. 1996, ApJS, 105, 19
- [180] Davé, R., Dubinski, J., & Hernquist, L. 1997, New Astronomy, 2, 277
- [181] Lucy, L. B. 1977, AJ, 82, 1013
- [182] Gingold, R. A. & Monaghan, J. J. 1977, MNRAS, 181, 375
- [183] Monaghan, J. J. 1992, ARA&A, 30, 543
- [184] Springel, V. & Hernquist, L. 2002, MNRAS, 333, 649
- [185] Hernquist, L. 1993, ApJ, 404, 717
- [186] Monaghan, J. J. & Gingold, R. A. 1983, J. Comp. Phys., 52, 374
- [187] Balsara, D. S. 1995, J. Comp. Phys., 121, 357
- [188] Steinmetz, M. 1996, MNRAS, 278, 1005
- [189] Hernquist, L., Bouchet, F. R., & Suto, Y. 1991, ApJS, 75, 231
- [190] Xu, G. 1995, ApJS, 98, 355

- [191] Bode, P., Ostriker, J. P., & Xu, G. 2000, *ApJS*, 128, 561
- [192] Bagla, J. S. 2002, *J. of Astrophysics and Astronomy*, 23, 185
- [193] Bagla, J. S. & Ray, S. 2003, *New Astronomy*, 8, 665
- [194] Eisenstein, D. & Hu, P. 1999, *ApJ*, 511, 5
- [195] Peacock, J. A. & Dodds, S. J. 1996, *MNRAS*, 280, 19
- [196] Eisenstein, D. J. & Hut, P. 1998, *ApJ*, 498, 137
- [197] Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001, *MNRAS*, 328, 726
- [198] Sheth, R. K. & Tormen, G. 1999, *MNRAS*, 308, 119
- [199] Ryu, D., Kang, H., Hallman, E., & Jones, T. W. 2003, *ApJ*, 593, 599
- [200] Barkana, R. & Loeb, A. 2000, *Physics Reports*, 349, 125
- [201] Kravtsov, A.V., Nagai, D. & Vikhlinin, A. A. 2005, *ApJ* submitted (astro-ph/0501227)
- [202] Heitmann, K., Ricker, P. M., Warren, M. S. & Habib, S. 2005, *ApJ*, accepted (astro-ph/0411795)
- [203] Bertschinger, E. 1985, *ApJS*, 58, 39
- [204] Shu, F. H. 1977, *ApJ*, 214, 488
- [205] Bodenheimer, P., & Sweigart, A. 1968, *ApJ*, 152, 515
- [206] Chandrasekhar, S. 1939. University of Chicago Press, Chicago, IL, USA, “An Introduction to Stellar Structure”
- [207] Quinn, P. J., & Zurek, W. H. 1988, *ApJ*, 331, 1
- [208] Hachisu, I., Tohline, J. E., & Eriguchi, Y. 1987, *ApJ*, 323, 592
- [209] Hachisu, I., Tohline, J. E., & Eriguchi, Y. 1988, *ApJS*, 66, 315
- [210] Li, H., Colgate, S. A., Wendroff, B., & Liska, R. 2001, *ApJ*, 551, 874
- [211] Bromm, V., & Loeb, A. 2003, *ApJ*, 596, 34
- [212] Barkana, R., Haiman, Z., & Ostriker, J. P. 2001, *ApJ*, 558, 482

- [213] Dalal, N. & Kochanek, C. 2002, PRL, submitted; astro-ph/0202290
- [214] Yoshida, N., Sokasian, A., Hernquist, L., & Springel, V. 2003, ApJL, 591, L1
- [215] Kitayama, T., Yoshida, N., Susa, H., & Umemura, M. 2004, ApJ, 613, 631
- [216] Haiman, Z., Abel, T. & Rees, M. 2000, ApJ, 534, 11
- [217] Truelove, J.K. et al. 1998, ApJ, 495, 821
- [218] Abel, T. 2000, Revista Mexicana de Astronomia y Astrofisica Conference Series, 9, 300
- [219] Abel, T. & Wandelt, B. D. 2002, MNRAS, 330, L53
- [220] Shapiro, P. R., Iliev, I. T., & Raga, A. C. 2004, MNRAS, 348, 753
- [221] Barnes & Efstathiou (1987), ApJ, 319, 575
- [222] Gardner, J.P. 2001, ApJ, 557, 616
- [223] Bondi, H. & Hoyle, F. 1944, MNRAS, 104, 273
- [224] Ricotti, M., Gnedin, N.Y. & Shull, J.M. 2002 ApJ, 575, 49
- [225] Saigo, K., Matsumoto, T., & Umemura, M. 2004, ApJ, 615, L65
- [226] Umeda, H. & Nomoto, K. 2005, ApJ, 619, 427
- [227] Mac Low, M.-M., & Norman, M. L. 1993, ApJ, 407, 207
- [228] Yoshida, N., Abel, T., Hernquist, L. & Sugiyama, N. 2003, ApJ, 592, 645
- [229] Ferrara, A. 1998, ApJ, 499, L17
- [230] Haardt, F. & Madau, P. 1996, ApJ, 461, 20
- [231] Zhang, Y., Anninos, P., Norman, M. L., & Meiksin, A. 1997, ApJ, 485, 496
- [232] Mackey, J., Bromm, V., & Hernquist, L. 2003, ApJ, 586, 1
- [233] Schneider, R., Ferrara, A., Natarajan, P., & Omukai, K. 2002, ApJ, 571, 30
- [234] Madau, P., Ferrara, A., & Rees, M. 2001, ApJ, 555, 92

[235] Oh, S. P., Nollett, K. M., Madau, P., & Wasserburg, G. J. 2001, ApJ, 562, L1

This thesis was prepared using the \LaTeX typesetting language [236, 237].

[236] L. Lamport, 1985 Addison-Wesley, Boston, “ \LaTeX : A Document preparation System”

[237] D. E. Knuth, 1985 Addison-Wesley, Boston, “The \TeX book”

Author's Biography

Brian William O'Shea was born in Oak Lawn, Illinois on August 16, 1978. He graduated from the University of Illinois in Urbana-Champaign in 2000 with a Bachelor of Science degree (cum laude) in Engineering Physics. He then stayed at the University of Illinois for graduate school, earning his Master of Science degree in Physics in December 2001. During the course of his graduate research, Brian spent most of his time at the Laboratory for Computational Astrophysics, which is part of the Center for Astrophysics and Space Sciences at the University of California in San Diego, and several months at the Theoretical Astrophysics Group (T-6) at Los Alamos National Laboratory. Following the completion of his PhD he will begin work as a postdoctoral researcher in the Theoretical Astrophysics Group at LANL.