Comparing Simulations and Observations of Galaxy Evolution: Methods for Constraining the Nature of Stellar Feedback

Cameron Hummels

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2012

ABSTRACT

Comparing Simulations and Observations of Galaxy Evolution: Methods for Constraining the Nature of Stellar Feedback

Cameron Hummels

Computational hydrodynamical simulations are a very useful tool for understanding how galaxies form and evolve over cosmological timescales not easily revealed through observations. However, they are only useful if they reproduce the sorts of galaxies that we see in the real universe. One of the ways in which simulations of this sort tend to fail is in the prescription of stellar feedback, the process by which nascent stars return material and energy to their immediate environments. Careful treatment of this interaction in subgrid models, so-called because they operate on scales below the resolution of the simulation, is crucial for the development of realistic galaxy models. Equally important is developing effective methods for comparing simulation data against observations to ensure galaxy models which mimic reality and inform us about natural phenomena.

This thesis examines the formation and evolution of galaxies and the observable characteristics of the resulting systems. We employ extensive use of cosmological hy-

drodynamical simulations in order to simulate and interpret the evolution of massive spiral galaxies like our own Milky Way. First, we create a method for producing synthetic photometric images of grid-based hydrodynamical models for use in a direct comparison against observations in a variety of filter bands. We apply this method to a simulation of a cluster of galaxies to investigate the nature of the red-sequence/blue-cloud dichotomy in the galaxy color-magnitude diagram. Second, we implement several subgrid models governing the complex behavior of gas and stars on small scales in our galaxy models. Several numerical simulations are conducted with similar initial conditions, where we systematically vary the subgrid models, afterward assessing their efficacy through comparisons of their internal kinematics with observed systems. Third, we generate an additional method to compare observations with simulations, focusing on the tenuous circumgalactic medium. Informed by our previous studies, we investigate the sensitivity of this new mode of comparison to hydrodynamical subgrid prescription. Finally, we synthesize the results of these studies and identify future avenues of research.

Contents

List of Figures

List of Tables

ACKNOWLEDGMENTS

The completion of this dissertation has required a large portion of my life, and this major accomplishment has only been possible through the generous mentorship and support of family, friends and colleagues who went beyond what was necessary to help me in my journey.

First, I would like to thank the many teachers I have had throughout my life who shaped me towards a scientific path: David Sherwood, my fourth-grade teacher who first taught me facts about the Universe which I now know were wrong; Larry Clow, my off-the-wall high-school physics instructor who first introduced me to physics; Joann Derie and Julie Nygaard, my most patient mathematics instructors for listening to my unending barrage of inane math questions; Ran Libeskind-Hadas for his provocative examples and demonstrations in class, providing me with the platonic ideal of excellent scientific pedagogy; and Bryan Penprase for his patience in allowing me to engage in astronomical projects night after night at the observatory.

Prior to my enrollment at Columbia, I received much encouragement toward my scientific aspirations from Steve Majewski and Ricky Patterson, helping me to lay the groundwork for the journey challenges ahead. Kathryn Johnston and Ed Moran exposed me to the rigors of graduate school, and they ignited my passion for performing scientific research instead of just coursework.

I found Columbia's astronomy department to be one of the most socially supportive

and friendly astrophysics departments in the country, with students and faculty alike willing to go out of their way to help me meet my goals. I appreciate all of the professors, but a few notable individuals are mentioned here explicitly. Despite their very busy schedules, David Helfand and Jacqueline van Gorkom always found a moment to listen to my problems, provide advice, and support me in my endeavors even when these sometimes conflicted with my studies. Ed Spiegel provided much intellectual and competitive support outside of the classroom with his demonstration of newtonian dynamics on the billiard table. Marcel Agueros and Mary Putman provided me with thoughtful feedback on research, outreach, and travel. Lastly, I was always abreast of news on the lunar front thanks to Arlin Crotts.

As for my immediate faculty advisors, I could not have asked for better. Greg Bryan is one of the most intelligent and patient men I know, respected by the community as much for his wisdom as for his honesty in assessing the truth. As a colleague recently related to me: "Everyone trusts Greg, because you know he doesn't have a side he is trying to promote." David Schiminovich is always available for discussions on any topic, sharing his vast knowledge with utter modesty. He's the first faculty member who ever deeply considered the ideas that I proposed in our meetings, truly thinking about them, and discussing them with me as though I were a peer. Together, these two men are the largest role models in my life. I hope to make them both proud of me and my accomplishments in the future.

All of the graduate students and post-docs in the department have come to my aid

one time or another throughout these last several years. I wish them thanks for helping me to figure out what I wanted to do, and how to do it. Gratitude is especially meant for graduate students: Sasha, Oishi, Ian, Roban, Stephanie, Maria, Emily, Taka, Neil, Kyle, Destry, Duane, Dan, Christine, Lia, and Josh; and post-docs: Matt, Summer, Josh, Ryan and Ally.

Obviously I wouldn't have made it this far without the support of my family. Thank you, Mom, Dad, and Allison for being behind me in my efforts 110%, and for always encouraging me to pursue my interests, even if they did not align with your own. Dad, I deeply appreciate the fact that you mentored me in my hobbies, and for introducing me to sidewalk astronomy at such a young age. I'll miss you.

Lastly, I owe much of my success the last few years to my wonderful girlfriend, Erika. She sets the bar high for accomplishments and the amount of work she can perform in a day, making it all look easy. Her can-do attitude and superhuman willingness to aid me in my various projects, academic and extracurricular, has enabled me to succeed in many difficult circumstances. I look forward to my opportunity to assist her in arriving at this same achievement in a couple of years.

Thank you everyone. I am endebted to you all.

Cameron Hummels

June 2012

New York City

Chapter 1

Introduction

Nearby galaxies such as Andromeda and the Magellanic Clouds have been observed since antiquity, but the nature of these objects as being distinct from our own Milky Way galaxy was not established until the early twentieth century. Studies at the time demonstrated that these "island universes" were moving too quickly (Slipher 1913) and were too distant (Hubble 1929) to be considered a part of our own Milky Way. Since that time, the study of galaxies has remained at the forefront of astrophysical research in an effort to understand how these massive cosmological building blocks form and evolve over the lifespan of the universe. This dissertation focuses on methods for comparing observational and numerical studies of galaxy evolution with a focus on massive, star-forming galaxies like our own Milky Way. While the methods presented herein can be applied to other systems including ellipticals and dwarf galaxies, Milky-Way-like objects are chosen as our focus for a number of reasons. One, we reside inside one such system, giving us

a unique, high-resolution perspective on its current conditions. Two, massive spiral galaxies are among the brightest objects in the visible universe, providing us with a large observational sample of these galaxies. Three, Milky-Way-like systems sit at the break in the galactic mass function, providing a medium-mass galaxy which can potentially act as a representative sample for a large number of galaxies. For these reasons, star-forming spiral galaxies make an excellent target for this comparison study.

I begin this thesis outlining different methods for modeling galaxies using computational simulations in Section 1.1. in Section 1.2, I cover the various implementations of subgrid models and their application to cosmological simulations of galaxy evolution, using star formation and supernovae feedback as an example. In Section 1.3, I describe the methods by which numerical simulations of galaxies are compared against observations, and I introduce the three avenues of comparison elaborated in the rest of this work. Finally, I summarize the remaining chapters of the thesis in Section 1.4.

1.1 Numerical Simulations of Galaxy Evolution

Because of their large size relative to their bulk and internal velocities, galaxies take a long time to move, change, and evolve. The dynamical time of the Milky Way is ~250 Myrs, far longer than the period over which humans have been observing galactic systems. With such a limited baseline of observations, coupled with our lack of understanding of many empirical laws governing galaxies, astronomers have been forced to turn to other means

for comprehending the formation and evolution of these systems. One such method is to employ computers to simulate them in order to trace their evolutionary path through the universe, both as counterparts to actual observable systems as well as statistically-accurate approximations of the general galactic population.

Early work focused on simply modeling the effects of gravity on galactic evolution, since gravity dominates at such large scales. Holmberg (1941) represented individual masses as light bulbs and photometers utilizing the r^{-2} dependence of both gravity and luminosity, calculating and moving masses by hand to approximate the time-dependent evolution of a stellar system of only 37 particles. The development of digital computers enabled larger studies incorporating hundreds of particles (von Hoerner 1960). One watershed moment for modeling was the demonstration by Toomre & Toomre (e.g. 1972) that mergers and gravitational interactions were responsible for many of the irregular galaxies which had previously baffled observers. Later work incorporated hydrodynamics (Larson 1969), the second most important effect behind gravity for galactic systems. Press & Schechter (1974) placed their simulations in a fully cosmological context, a crucial step in the development of numerical galaxy evolution. This allowed scientists for the first time to realistically construct galaxies from smaller structures in the early universe, as opposed to previous studies which assumed an artificial distribution of matter as the initial conditions before any subsequent evolution. While isolated galaxy studies of this type continue to be performed for their high spatial and mass resolution, they have limited utility for helping us comprehend the formation and early evolution of galaxies, since they

lack interactions with the galactic environment.

The simulation codes of today have progressed by leaps and bounds beyond these early attempts. The boom of the computer industry, the development of fast numerical algorithms, and the increasing use of computer clusters and supercomputers over the last 40 years has led to simulations with billions of resolution elements (e.g. Springel et al. 2005; Diemand et al. 2006; Springel et al. 2008) and refinements for chemistry and radiative transfer, in addition to the dominant effects of gravitation and gas dynamics.

Modern codes begin a simulation with initial conditions representing the high-redshift (i.e. $z \sim 100$) universe, typically employing the Zeldovich approximation (Zel'dovich 1970) to deposit mass in the form of dark matter and gas as small-scale stochastic fluctuations on an expanding volume following some preferred cosmological parameters. The dominant laws of physics (see below) are plugged into the code to account for as much physics as is computationally feasible, and the simulations are allowed to run forward in time, numerically solving the equations of motion and other related differential equations for the various components of the simulation (Bertschinger 1998). In the end, the simulation yields galaxies and other objects which should be statistically consistent with observations of the real universe, although they will not be copies of specific objects in our own universe.

Dark matter, a quantity which only seems to interact with other matter gravitationally, permeates galactic structures supplying some 90% of their mass. Unlike gas, it doesn't appear to follow the laws of hydrodynamics, and instead is properly treated as a colli-

sionless fluid represented by the six-dimensional Boltzmann equation. In nearly every modern code, the dark matter is sampled at discrete locations and represented as massive collisionless particles. Gravitational calculations are performed on the dark matter particles as well as on whatever form is chosen to represent the baryonic quantities (i.e. stars and gas). To avoid the computational expense of direct summation of gravitational effects on every pair of masses in a simulation, gravity calculations are typically performed using a tree method (Barnes & Hut 1986), particle-mesh method (i.e. PM method, Miller & Prendergast 1968), particle-particle / particle-mesh method (i.e. P^3M method, Efstathiou & Eastwood 1981), or multiresolution mesh method. The tree method is efficient, but memory intensive and challenging to use in periodic boundary conditions. On the other hand, the PM technique automatically incorporates periodic boundary conditions, useful for cosmological simulations, but performs poorly on small scales. The P^3M and multimesh techniques are a type of hybrid, which employ the PM on large scales but utilize corrections on smaller scales such as particle-particle calculations or higher-resolution grids, respectively.

Modern cosmological simulations are typically grouped based on their implementation of hydrodynamics, but ultimately all of these codes follow the same equations of hydrodynamics. In smoothed-particle hydrodynamics (SPH) (Lucy 1977), gas parcels are represented as particles in a Lagrangian approximation to the laws of hydrodynamics. Each parcel of gas is homogenous and indivisible, possessing density, temperature, and other fluid variables. To sample a hydrodynamical quantity at any location, one must

apply a smoothing kernel, smoothing over many nearby gas parcels. One can achieve high levels of resolution in regions of high density, where there are a lot of SPH particles but SPH methods suffer from low levels of resolution in voids. In addition, shocks which are omnipresent in cosmological astrophysical systems, tend to be reduced and lost during the smoothing process. SPH codes are common in cosmological modeling due to the long history of particle-based simulations for dealing with N-body gravitational problems. Popular astrophysical codes used today which employ the SPH technique are Gadget (Springel et al. 2001), Gasoline (Wadsley et al. 2003), and VINE (Wetzstein et al. 2009).

An alternative to SPH are grid-based codes, which represent the fluid equations in their Eulerian form with a fixed uniform grid. Each element of the grid, like the SPH particle, is an indivisible homogenous parcel of gas. At each timestep, gas quantities can pass through the grid interfaces to neighboring cells depending on the local conditions present. All grid-based codes which employ the total-variation diminishing (TVD) or piecewise parabolic method (PPM) for shock capture are much better than SPH codes at resolving shocks. The potential problem with grid codes is their limited spatial resolution. However, adaptive mesh refinement (AMR), a technique by which the grid elements are progressively refined in regions meeting some predefined criteria (e.g. of high density), has alleviated much of this problem, producing increased resolution at little extra computational expense. Enzo (Bryan & Norman 1997), ART (Kravtsov et al. 1997), RAMSES (Teyssier 2002), and FLASH (Fryxell et al. 2000) are all AMR-enabled grid-based codes

used in the astrophysical literature today.

There exists a hybrid of SPH and AMR methods called the moving mesh technique. It functions by producing an irregular grid to represent the fluid, which is subsequently allowed to move and deform with the fluid's evolution. The benefits of such a method provide high local resolution in regions of high-density (like SPH) but better shock capture (like AMR). Unfortunately, there can be challenges in implementing such a code to reconstruct an irregular grid at each timestep while conserving fluid quantities. Early attempts met limited application (e.g. Gnedin 1995), but recently a successful general implementation has been achieved, known as the AREPO code (Springel 2010).

In the limit of high resolution, SPH, AMR and moving mesh techniques should all result in identical behavior, since they are all approximations to the same fluid equations; however, in reality there are oftentimes differences demonstrated by individual algorithms (see, for example code comparison studies by O'Shea et al. 2005; Agertz et al. 2007; Scannapieco et al. 2011) due to shortcomings in individual code algorithms. This underlines the need to have many researchers working with different codes for correctly identifying the underlying physics present in all of them.

1.2 Subgrid Models

In each of these methods, resolution for both mass and spatial quantities remains one of the primary limiters to accuracy. Current computational constraints prevent simulations

incorporating cosmological structure (i.e. tens of megaparsecs on a side) from achieving better mass and spatial resolution than $\sim 10^3$ $M_\odot$ and ~ 10 comoving pc, respectively. Future increases in computing power will only slowly erode this problem (barring any major algorithmic advances). For example, in order to increase the spatial resolution by a factor of two in length, volume elements must be split into $2^3 = 8$ parts. If we fix the courant condition so quantities don't move too quickly through these smaller volumes, the timestep must also be shrunk by a factor of two. Thus, to double a simulation's spatial resolution requires $2^4 = 16$ times the time and resources, which may only be met by Moore's Law every six years from hardware advances. The use of AMR and SPH lightens this expense somewhat, depending on the filling factor of highly-refined regions, but the expense is significant nonetheless. It suggests that we will not be able to rely on advances in hardware technology to alleviate the need for alternative means for treatment of the physics at small spatial and mass scales in cosmological simulations. Unfortunately, there is a lot of important physics which occurs on such scales including star formation, stellar feedback, active galactic nuclei (AGN) feedback, turbulence, radiative transfer, etc. Furthermore, these processes appear to couple non-linearly to their larger environment and to the global evolution of a galaxy, requiring that they be accounted for in a relatively accurate manner.

Subgrid models are prescriptions for the behavior of gas below the resolution of a simulation, and they are created to mimic some of the aforementioned physics which is difficult to resolve in modern simulation codes. For example, stars form when molecular

clouds become cold and massive, undergoing gravitational collapse in their cores and fragmenting into a stellar population. While this behavior occurs well below the resolution of our simulation on sub-parsec scales, we can account for it by employing a star-formation subgrid model (e.g. Cen & Ostriker 1992; Katz 1992). In Cen & Ostriker (1992), four conditions must be met for a gas grid element to form stars: (i) its mass must have exceeded the Jeans Mass, (ii) it must be locally converging, (iii) it must have reached some user-defined density threshold, and (iv) its cooling time is less than its dynamical time. When these criteria are reached, some amount of the gas in that grid element is converted to a star "particle", representing a coeval population of stars. While this simple prescription makes sense given our current understanding of star formation, it raises more questions. What do we use as our density threshold? What initial mass function do we choose? How do we later treat a stellar "particle" as it continues to age and its component stars diverge in behavior?

Additional subgrid models are required to mimic how these stellar populations interact with their environment. As they grow older, the most massive stars in these particles produce copious amounts of ionizing radiation and then transform into core-collapse supernovae, ejecting material and energy into their local environment, whereas the lower-mass stars slowly become planetary nebulae, type Ia supernovae, and stellar remnants over longer timescales. Attempts at encapsulating this behavior, called stellar feedback models, deposit mass, metals and energy into the medium surrounding nascent stars for a period after a star particle's creation. However, scientists disagree on the

appropriate implementation of these complex processes, and numerous prescriptions exist for stellar feedback subgrid models. Slight modifications to these models can yield enormous changes in the observational properties of a galaxy, so accurate treatment is crucial to the future of galaxy evolution simulations. Constraining the nature of this interaction is one of the large, open questions in galaxy evolution today, and it is the at the heart of this dissertation (see Chapters 3 and 4).

The simplest method for returning the energy to the medium surrounding a stellar population particle is to deposit it as thermal energy, thereby increasing the temperature of the gas. The temperature increase causes the gas to expand, raising the Jeans mass in the local environment, theoretically leading to diminished star formation. Due to its ease of implementation, this method is used extensively in the literature (e.g. Governato et al. 2004; Stinson et al. 2006; Ceverino & Klypin 2009; Hummels & Bryan 2012), but when used in a simulation with low spatial resolution, it causes an artificial effect known as the angular momentum problem (Navarro & Benz 1991). This problem, the focus of Chapter 3 of this dissertation, is caused by the spurious mixing of adjacent hot and cold regions of the multiphase galactic medium, creating an aggregate warm component very efficient at radiating away excess energies (Katz 1992; Katz et al. 1996). When star formation and subsequent stellar feedback occurs in a cold medium, any thermal energy ejected into the gas to quench star formation is thus effectively radiated away, removing pressure support, contracting the gas, and allowing star formation to continue unabated. As a result, dense knots of cold gas and stars are formed in galactic structures, efficiently

transferring their orbital angular momenta to the dark matter particles inhabiting the halo through dynamical friction interactions (D'onghia et al. 2006). Thus, these knots of material lose their rotational support and tumble to the center of the system to produce an overdense cusp of material in the core of the simulated galaxy, in conflict with observations.

In response to this shortcoming of the thermal energy injection model, numericists have implemented additional subgrid models for stellar feedback. Gerritsen (1997) first presented an *ad hoc* method for effectively preventing cooling losses after star formation events by manually turning off radiative cooling in gas elements for a short period of 20 to 50 Myr. When coupled with thermal feedback injection, this method yields adequate results (e.g. Agertz et al. 2011; Guedes et al. 2011; Hummels & Bryan 2012). Later this model was justified as an application of the Sedov-Taylor blastwave solution for a supernova explosion (Taylor 1950; Sedov 1959), since there is a short period over which the blast wave of a Sedov solution does not radiate away its energy; however, the duration of this theoretical calculation is in the kiloyear range, far shorter and less effective, than the periods employed in current research. Cooling suppression, as this technique is known, remains an unphysical model, inappropriate as a long-term solution to the problem of stellar feedback.

Other implementations of stellar feedback include a method known as kinetic or momentum-driven feedback (Springel & Hernquist 2003a; Oppenheimer & Davé 2008), wherein energy is deposted as kinetic energy in the surround medium, capable of generating massive outflows of material in a variety of temperatures and densities. By

injecting energy in a kinetic form, it cannot be radiated away and easily lost as a result of low spatial resolution. Winds made up of gas in many phases have the potential of populating the cirgumgalactic medium with a more realistic multiphase structure (e.g. McKee & Ostriker 1977). Most star formation occurs in dense, cold environments where the effects of momentum impulse from a supernova would have little effect. Therefore, most current implementations of kinetic feedback resort to hydrodynamically decoupling "kicked" gas elements from their surrounding medium, transporting them to low-density environments, so as to intensify the nature of the feedback. As in the case of the cooling suppression model, researchers find some signs of success in generating realistic galaxies, but it is at the expense of using a physically-inconsistent prescription of feedback. We will discuss more realistic implementations of this method in Chapter 5.

Finally, computational algorithms modeling the deposition of energy as radiation from supernovae and young stars are now becoming feasible. This radiation pushes and heats the local medium, acting as a sort of hybrid between kinetic and thermal injection techniques. Radiation takes the form of ray-traced photons emitted isotropically from sources (e.g. Wise & Abel 2011), or as clusters of emitters applying kicks to only nearby gas clumps (e.g. Hopkins et al. 2011). Like the kinetic models, these techniques show a great deal of promise for producing realistic behaviors in evolving galaxies but are still in their infancy. More exotic techniques employ cosmic rays (e.g. Jubelgas et al. 2008) or magnetic fields as mechanisms for energy transfer from stars to the interstellar medium, but so far these have not been investigated in great detail.

1.3 Methods of Comparison

Computer simulations are only useful if they reproduce behavior that reflects what actually occurs in nature. Following the prescriptions for cosmological modeling of galaxies described in the last few sections, scientists are thus able to simulate cosmological environments and produce galaxies, but how accurately do they depict the galaxies we find in the universe?

To answer this question, we search for ways in which to compare our simulated systems against statistical samples of real galaxies. Numericists possess much more information about their galactic systems than observers do of their galactic samples, so many of the comparisons are made by downgrading, sampling, and/or converting simulated properties into those which can be easily observed. Some common methods for accomplishing this compare various bulk properties of galaxies like size, mass, and velocity dispersion, quantities relatively easy to determine through observations, but simplistic comparisons like these yield limited information. The main purpose of this dissertation is to identify additional methods of comparison between computational simulations of galaxies and galactic observations for not only making more effective subgrid models, but also for aiding in the understanding of the complex processes of aggregate star formation and stellar feedback and their effects on the overall dynamics of the host galaxy.

The most common way in which we observe galaxies in nature is through photometric images taken in the visible part of the electromagnetic spectrum. Much of our understanding of galaxies stems from this type of measurement, since a great deal of information

can be conveyed by photometric images of ultraviolet, visible and near-infrared wavelengths including star-formation histories, luminosities, dust content, merger history, etc. Utilizing this fact, we synthesize observations of computationally-generated galaxies for analysis . By picking a projection vector and numerically integrating the radiative transfer equations to account for emission from stars and absorption by intervening dust, one can produce a view which a virtual observer would see when viewing these modeled galaxies. The focus of the work in Chapter 2 is to produce synthetic observations of simulated galaxies in a variety of photometric filter bands, enabling comparisons of optically-determined galactic quantities like color, morphology, luminosity and concentration.

Research pioneered by Rubin et al. (1985) revealed that the radial distribution of transverse velocities at which material orbits galaxies, the so-called galactic rotation curve, is flat, implying the presence of a large quantity of unseen matter encircling all galaxies. Observers probe this rotation curve using the doppler-shifted spectral lines emitted by stars and gas as they orbit around a galaxy. One can easily calculate the equivalent galactic rotation curve for a fully numerical system, based solely on the radial distribution of mass according to the equation, $v_{\rm circ} = \sqrt{\frac{Gm(r)}{r}}$, where $v_{\rm circ}$ is the circular velocity of an orbiting object at radius r, with $m(r)$ the mass interior to radius r, and G as the gravitational constant. Recall that some subgrid models produced an effect known as the angular momentum problem, evidenced by a cusp of dense material at the centers of numerical galaxies. This massive cusp has a profound effect on the rotation curve of the corresponding galaxy, since it raises $m(r)$ at small r causing a huge boost in $v_{\rm circ}$ at small radii. Consequently, inspection

of rotation curves is a primary method of comparing numerical models to observations of galaxies, and it is used extensively in our analysis in Chapter 3.

In Chapter 4, we introduce a new avenue of comparison, probing the low-density gas surrounding galaxies known as the circumgalactic medium. This medium, a mix of cold, infalling, primordial material and hot gas outflows, is too tenuous for current techniques to observe it in emission, so observers have resorted to indirectly determining its characteristics as seen in absorption against bright background sources like quasars. Observations thus correlate the strength of an absorber, typically measured as a spectral equivalent width, with the proximity a quasar sightline passes to an intervening galaxy. There are numerous atomic transitions potentially detectable in absorption, depending on the detailed chemical makeup of the cirgumgalactic medium. Each atomic transition occurs at a fixed photon energy value, probing the metallicity, temperature and gas density necessary to produce that transition along any individual sightline. Compiling many observations of a galaxy, or certain type of galaxy, can reveal information about the phase of the CGM at different galactocentric radii. Synthetic observations can be produced in a similar manner, by passing random sightlines through a galaxy and its immediate environment. We achieve this comparison, and demonstrate its efficacy in placing tight constraints on numerical models in Chapter 4.

The purpose of this dissertation is three-fold: to produce numerical simulations of galaxies using various subgrid models for star formation and stellar feedback, to make detailed comparisons between our simulations and observations for gauging the effec-

tiveness of these models in producing realistic systems, and to gain insight into the nature of star formation and feedback. Each of the three chapters herein attempts to compare a different aspect of galaxies with the observational sample (as discussed above), in some cases creating entirely new infrastructure for enabling such a comparison. While the subgrid models presented do not fully match observations in their behavior, they place tight constraints on the next-generation of numerical models and lay the groundwork for future studies.

1.4 Thesis Summary

In Chapter 2, I implement a method for producing synthetic observations of numerical models, and apply them to a cluster simulation to address how well simulations reproduce the red and blue sequences of galaxies. Chaper 3 follows the implementation of several different subgrid models and subsequent testing of them using star formation rates, flat rotation curves and visual images as methods of comparison. In Chapter 4, I create an additional series of simulations with more detailed coverage of their radiation, chemistry and feedback; these are then processed using a novel technique to produce column density distributions in a variety of atomic species in the circumgalactic medium, enabling comparisons with several observational datasets. Finally, Chapter 5 summarizes the conclusions of this work and details the next steps for arriving at an accurate implementation of stellar feedback.

1.4.1 Chapter 2: A Photometric Comparison of Simulated and Observed Galaxies

We develop a method for producing photometric images from the stellar and gas components in the outputs of the hydrodynamics code `Enzo`. Population synthesis models are applied to the stellar particles to produce synthetic spectra and then these spectra are sampled over a set of filters to acquire broadband fluxes for each star particle. Projecting these stars particles' positions and reddening according to the column density of interlying gas and dust, we produce realistic two-dimensional images in different filter bands. We apply this method to a cosmological cluster simulation and extract photometric information pertaining to individual galaxies present in the cluster. The resulting dataset is compared with the photometric properties of the subsample of cluster galaxies found in the Abell Cluster Catalog that have been observed with GALEX, SDSS and 2MASS. Our analysis suggests some oversimplifications present in the `Enzo` code, as well as some perplexing differences between the observed and simulated populations of galaxies.

1.4.2 Chapter 3: A Kinematic Comparison of Simulated and Observed Galaxies for Constraining Subgrid Models

We carry out adaptive mesh refinement (AMR) cosmological simulations of Milky-Way mass halos in order to investigate the formation of disk-like galaxies in a Λ-dominated Cold Dark Matter model. We evolve a suite of five halos to $z = 0$ and find gaseous-disk

formation in all; however, in agreement with previous SPH simulations (that did not include a subgrid feedback model), the rotation curves of all halos are centrally peaked due to a massive spheroidal component. Our standard model includes radiative cooling and star formation, but no feedback. We further investigate this angular momentum problem by systematically modifying various simulation parameters including: (i) spatial resolution, ranging from 1700 to 212 pc; (ii) an additional pressure component to ensure that the Jeans length is always resolved; (iii) low star formation efficiency, going down to 0.1%; (iv) fixed physical resolution as opposed to comoving resolution; (v) a supernova feedback model which injects thermal energy to the local cell; and (vi) a subgrid feedback model which suppresses cooling in the immediate vicinity of a star formation event. Of all of these, we find that only the last (cooling suppression) has any impact on the massive spheroidal component. In particular, a simulation with cooling suppression and feedback results in a rotation curve that, while still peaked, is considerably reduced from our standard runs.

1.4.3 Chapter 4: A Comparison of Simulated and Observed Circumgalactic Media for Constraining Subgrid Models

Cosmological hydrodynamical simulations of galaxy evolution are increasingly able to produce realistic galaxies, but the largest hurdle remaining is in constructing subgrid models that accurately describe the behavior of stellar feedback. As an alternate way to test and calibrate such models, we propose to focus on the circumgalactic medium. To

do so, we generate a suite of adaptive-mesh refinement (AMR) simulations for a Milky-Way-massed galaxy run to $z = 0$, systematically varying the feedback implementation. We then post-process the simulation data to compute the absorbing column density for a wide range of common atomic absorbers throughout the galactic halo, including H I, Mg II, Si II, Si III, Si IV, C IV, N V, O VI, and O VII. The radial profiles of these atomic column densities are compared against several quasar absorption line studies, to determine if a particular feedback prescription is favored. We find that although our models match some of the observations (particularly those ions with lower ionization strengths), it is particularly difficult to match OVI observations. There is some indication that the models with increased feedback intensity are better matches. We demonstrate that sufficient metals exist in these halos to reproduce the observed column density distribution in principle, but the simulated circumgalactic medium lacks multiphase substructure and is generally too hot. Furthermore, we demonstrate the failings of inflow-only models (without feedback) at populating the CGM with adequate metals to match observations even in the presence of multiphase structure. Quasar absorption line observations of the gas around galaxies place a tight constraint on feedback models, rejecting galaxies able to meet traditional metrics of assessment (e.g. flat rotation curves, Kennicutt-Schmidt law) and suggesting a new avenue of comparison for cosmological galaxy simulations.

Chapter 2

A Photometric Comparison of Simulated and Observed Cluster Galaxies

2.1 Introduction

Astronomers have long sought to understand the process of galaxy formation and evolution. This search has uncovered numerous ways in which different physical properties of galaxies (e.g. mass, photometric color, luminosity, environment, metallicity, morphology, etc.) correlate with each other. Some examples include the Tully-Fisher relation for spiral galaxies (Tully & Fisher 1977), the fundamental plane for elliptical galaxies (Faber & Jackson 1976; Djorgovski & Davis 1987), the color-luminosity relation for ellipticals (Baum 1959), and the morphology-density relation (Oemler 1974; Dressler 1980). Only by ascertaining the reasons why these relationships exist, can astronomers hope to produce

a full picture of how galaxies form and evolve.

To aid in understanding the mechanisms responsible for galaxy evolution, astronomers employ numerical simulations of astrophysical systems. These models are powerful tools for applying what we understand about the physics of galaxy evolution and attempting to reproduce the complex behaviors observed. The closer that models are able to reproduce what is observed in the sky, the better the physics of the system is understood. Conversely, understanding the discrepancies between output from numerical simulations and observational data can lead to insight about the nature of the fundamental processes responsible for galaxy evolution.

Unfortunately, the output from numerical models is not always in a state comparable to data reduced by observers. Therefore, it is warranted to seek a transformation method for converting the simulated output into the form of the observational data for direct comparison. In our case, this means converting three-dimensional hydrodynamics outputs into broadband photometric images.

This chapter is organized as follows. Section 2.2 describes the method used for converting our simulated data into synthetic images. It additionally explains the procedure for extracting relevant quantities from the images, as well as the the properties of the observational sample use for comparison. Section 2.3 describes the form of our data analysis. It juxtaposes the observational and simulated datasets and discusses the similiarities and differences found between the two. Additionally, we mention possible explanations for the discrepancies observed and list future work to better constrain these hypotheses.

Section 2.4 cites the conclusions of this study and systematically describes further research for this topic.

2.2 Method

2.2.1 Simulated Data

Our simulated data consists of a cluster of galaxies evolved from $z = 40$ to present redshift ($z = .0004$) using the `Enzo` code. `Enzo` is an adaptive mesh refinement (AMR) hydrodynamics code for solving the fluid equations and gravity over a system consisting of dark matter, stars and gas (Bryan & Norman 1997; Norman & Bryan 1999; O'Shea et al. 2004). It treats dark matter and stellar components as collisionless point-like particles, whereas the gas data is evolved as cells in a fixed three-dimensional grid. When the gas becomes sufficiently cold and dense in a cell, `Enzo` splits the cell into a higher-resolution subgrid so that astrophysically-important features are well-resolved.

The cluster forms within a periodic simulation box which is 64 h^{-1} Mpc on a side, in a flat, cosmological-constant-dominated universe with the following parameters: $(\Omega_0, \Omega_\Lambda, \Omega_b, h, \sigma_8) = (0.3, 0.7, 0.045, 0.7, 0.9)$. We employ a multi-mass initialization technique in order to provide high resolution in the region surrounding the cluster, while evolving the rest of the box at low resolution. The dark-matter particle mass is $6.4 \times 10^8 \mathrm{M}_\odot$, with a gas mass resolution about five times better than this. The whole cluster has more than one million particles within the virial radius, and a typical L_* galaxy is resolved

by several thousand particles. The adaptive mesh refinement provides higher resolution in high density regions, giving a best cell size (resolution) of 3 kpc. This is sufficient to resolve the large galaxies in which we are interested. We also believe that it approximately reproduces effects such as ram-pressure stripping and tidal interaction, although the internal dynamics of galaxies are not well-resolved. While increased resolution is desired, we must find a balance in choosing our resolution and box size in order to produce a cluster of this mass.

The simulation includes radiative cooling using the Sarazin & White (1987) cooling curve and an approximate form of star formation and supernovae feedback following the Cen & Ostriker (1992) model. Briefly, the star formation method relies on identifying cold, collapsing, high-density clouds and forms stars at a rate proportional to the density of gas divided by the dynamical time, multiplied by an efficiency factor. This efficiency is taken to be 2%, which produces a Kennicutt-Schmidt relationship (see Figure 3.2.1.4) . See O'Shea et al. (2004) for a more complete discussion of the star formation algorithm. These stars then are represented as collisionless particles possessing mass, metallicity, and age. Because the amount of gas converted to stars in each formation event is large ($M_* = 10^5 M_\odot - 10^9 M_\odot$), it is better to think of each particle as an entire stellar population born in an instananeous starburst—what is known as a *simple stellar population* (SSP). The energy from Type II SNe in a stellar particle is returned to the surrounding gas in the form of thermal energy. Although this energetic output can be important, it is known that much of this energy is deposited in high-density gas where the cooling time

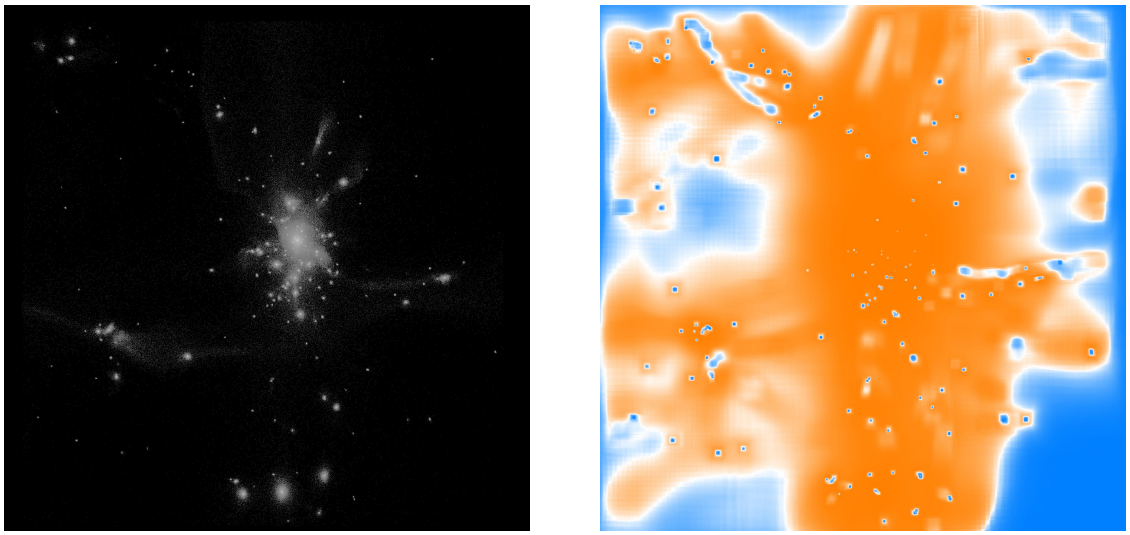

Figure 2.1 Stellar (left) and gas (right) components of simulated data. Brightness of a pixel in the stellar map is directly related to the mass of stars within it. There is a large halo of stellar matter circling around the central cD in the cluster. In the gas map, the red regions represent areas dominated by gas with temperatures $> 15,000$ whereas blue regions contain lots of gas mass with temperatures $\leq 15,000$. This map illustrates the hot ICM pervading the center of the cluster. It also shows several cool gas clouds collapsing onto star-forming galaxies.

is short, and so it is radiated away. This results in an "overcooling" problem (e.g. Balogh et al. 2001), and manifests itself in our simulations as somewhat overly massive galaxies, as well as a higher-than-observed ratio of stars and cool gas to hot gas. In addition, we observe hot intracluster gas cooling onto the centers of a few of the most massive galaxies (and in particular the central galaxy). This cooling is not observed in real clusters, probably because of feedback from supermassive black holes, which are not included in the simulation.

The relevant quantities we possess for each stellar particle are its position, mass, age, and metallicity, whereas for each gas grid element we have position, volume, mass,

temperature, and metallicity. Figure 2.1 is the set of projected images of the stellar particles and the gas cells. Each image has a resolution of 10 kpc pixels and as a whole it covers 20 Mpc on each side. The brightness of each pixel is directly related to the mass of the stars or gas present within it respectively. The gas map is color-coded such that gas below a temperature of 15,000 K is blue, and gas above this threshold is red. White areas just reflect an even amount of hot and cold gas. One can see there are numerous galaxies that have formed, and a large mass of stars surrounding the central region of the cluster. One can also see that the cold gas surrounds many of the galaxies in the periphery of the cluster, whereas hot gas pervades the interior of the cluster near the center. All of these findings are consistent with our understanding of galaxy evolution in the cluster environment.

2.2.2 Application of Population Synthesis Models

In order to convert simulated data into photometric images, we must first synthesize the emitted radiation from each stellar population particle. To accomplish this, we use the Bruzual-Charlot population synthesis code (`BC03`) (Charlot & Bruzual 1991; Bruzual & Charlot 1993, 2003), a code which assumes a base population of stars following an initial-mass function (IMF), evolves them according to the star formation history prescribed, and co-adds the spectra of each member of the final population. The user may specify the IMF, metallicity and star formation history of the population.

The code is calibrated with user-chosen stellar evolutionary tracks for ages 10^5 to 2×10^{10} years over a wide metallicity range. Output from the code is in the form of

an integrated spectrum for each stellar particle sampled to a resolution of 3 Å over the wavelength range $\lambda = 3200$ Å $- 9500$ Å. In addition, it produces lower resolution spectra over the wavelength range $\lambda = 91\text{Å} - 160\mu m$.

`BC03` is based on the *isochrone synthesis* technique, which utilizes the property that a stellar population with an arbitrary star formation history can be expanded into a collection of infinitescimal-duration starbursts, or SSPs. Recall that the stellar particles from the `Enzo` simulation are massive enough to approximate an entire population of stars, and they are created in instantaneous starbursts, thus making them good candidates for application of the `BC03` SSP models.

We use a Chabrier IMF (Chabrier 2003) and the Padova 1994 evolutionary tracks (Alongi et al. 1993; Bressan et al. 1993; Fagotto et al. 1994a,b; Girardi et al. 1996). These tracks are sampled 221 times over the age range 10^5 to 2×10^{10} years and 6 times in metallicity $Z = (0.0001, 0.0004, 0.004, 0.008, 0.02, 0.05)$. The `BC03` libraries correlate each sampling in age and metallicity with a high resolution spectra of a corresponding SSP. Figure 2.2 shows integrated spectra according to the Chabrier IMF and Padova 1994 models for a stellar population of $Z = Z_\odot = 0.02$ at different ages ($t_{\text{age}} = 0.001, 0.01, 0.1, 0.4, 1.0, 4.0, 13.0$ Gyr). As a population evolves, the UV side of the spectrum drops dramatically, whereas there is relatively little evolution on the IR side. For reference, the GALEX, SDSS, and 2MASS filter curves are plotted at the bottom.

Because the `BC03` interface is interactive and not made for running on a large set of individual stellar populations, we produce new code to sample the `BC03` data files

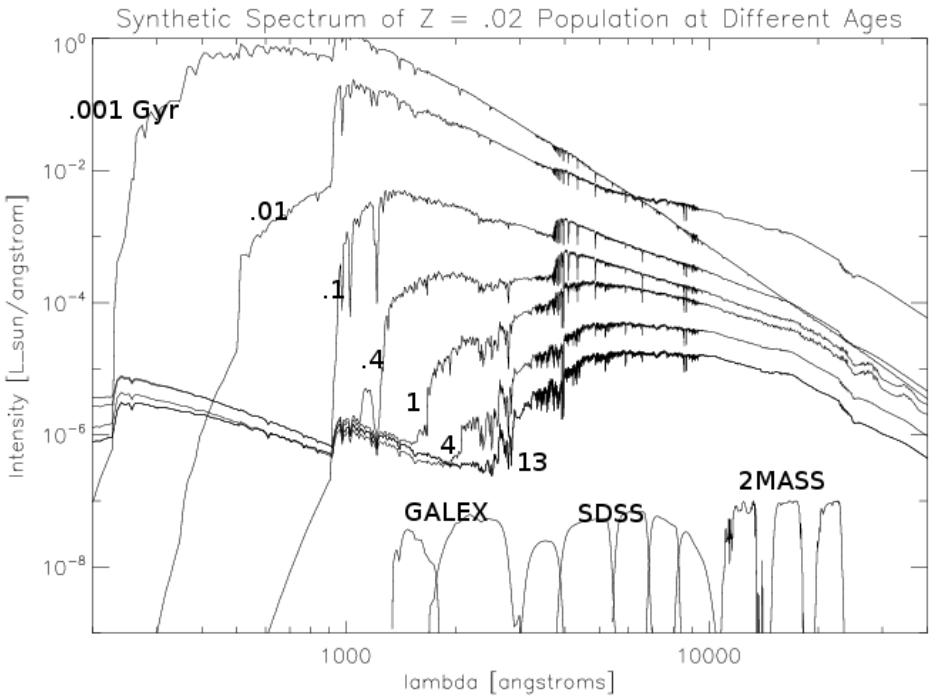

Figure 2.2 Synthetic spectra output by `BC03` using the Padova 1994 evolutionary tracks and the Chabrier IMF. This stellar population has a solar metallicity and each spectrum pertains to the population at a different time: age = 0.001, 0.01, 0.1, 0.4, 1.0, 4.0, 13.0 Gyrs. For reference, the GALEX, SDSS, and 2MASS filter functions are plotted at the bottom. The near IR portion of the spectrum is relatively insensitive to age of the population, whereas the near UV is very sensitive to it.

for the Padova evolutionary tracks. However, since our dataset contains particles over a continuum of stellar population ages and metallicities, our code needs to interpolate over the discretized 221×6 age-metallicity samplings. We apply a bilinear interpolation algorithm to do this efficiently, resulting in synthesized spectra interpolated over the full range of metallicity and age described above. Lastly, we assume a linear relationship between the mass of a stellar population and its luminosity. Thus, we multiply the intensity of each spectrum by the simulated mass of its star particle.

Since ultimately we want to compare the photometric properties of these simulated galaxies with their observational counterparts, we need to convert these synthesized spectra into photometric fluxes. We can effectively determine broadband fluxes for each spectra by convolving each with a selection of broadband filter of our choice. We want our filter list to include the filters from the Near UV, Optical and Near IR regimes where there exists a great deal of observed cluster galaxy data. We use the filters from the Galaxy Evolution Explorer (GALEX), the Sloan Digital Sky Survey (SDSS), and the Two-Micron All Sky Survey (2MASS) to cover this range in wavelength. Figure 2.2 shows from left to right the filter functions for the GALEX FUV & NUV filters, the SDSS U, G, R, I & Z filters, and the 2MASS J, H, & K filters.

2.2.3 Dust

The presence of dust in space significantly modifies the colors of the objects in its viscinity due to its ability to absorb and scatter radiation. In order to produce an accurate comparison of simulated data with an observational sample, we must provide some way of reddening and extinguishing the starlight from our simulation accordingly. So far, we only use the stellar particles from the `Enzo` simulation outputs, but there also exists the distribution of gas throughout the cluster space. By incorporating this gas into our calculations and using a few assumptions about the nature of it, we can make an approximation of its extinguishing effects on our simulated cluster sample.

Above a certain gas temperature, dust grains are not found in nature. We assume

that dust will only be found in gas below a temperature of 15,000 K. Our threshold value is somewhat arbitrary, but it is generally believed that at such temperatures grain-grain collisions, sputtering and photodestruction will remove dust particles from the gas (Osterbrock & Ferland 2006). Additionally, the majority of the mass of gas in the simulation is found in cells where the temperature is below this threshold, so making such an arbitrary cutoff has limited repercussions.

Calculation of the extinction from *all* of the gas in the simulation at the precision we desire is computationally unfeasible. The lion's share of this processing time would be spent calculating the extinguishing effect of low-density gas spread over huge volumes of the simulation. However, we believe this low-density gas has little impact on the light attenuation simply because it has so little mass to scatter and absorb radiation. Therefore, we ignore gas having a column density lower than a given threshold in our calculation of dust effects. Our cutoff column density is $N_{\mathrm{cutoff}} = 2 \times 10^9\ \mathrm{M}_\odot \cdot \mathrm{Mpc}^{-2}$. Although it is chosen somewhat arbitrarily, it corresponds to a transmission of .9999 in the V band, a very small effect. Even in the unlikely occurrence that 1000 such parcels of gas 10 kpc on a side lined up in a narrow column in front of a star particle, over 90 % of the original light would be transmitted to the observer in V band.

The chemical evolution of the `Enzo` code is relatively simplistic in that it only tracks metallicity in the form of a single index (Z = metals by mass content). This metallicity should be related to the amount of dust present in the gas, but in the current study all of the gas is considered to have the same metallicity. Future work will include incorporating

metallicity into the dust extinction calculation. Following the work of Bohlin et al. (1978), we assume the color excess (and hence the extinction factor) is approximately proportional to the column density of hydrogen (Bohlin et al. 1978; Kent et al. 1991). This takes the form:

$$E(B-V) = \frac{N_H}{5.8 \times 10^{21}\ \mathrm{cm}^{-2}} \tag{2.1}$$

The interstellar extinction curve has a well-defined slope near the V band, which provides us with a method of determining A_V from our color excess:

$$A_V = E(B-V) \times R_V \tag{2.2}$$

R_V quantifies the steepness of the extinction curve as it rises into the blue and UV. Astronomers still use the classical value for R_V derived by Schultz & Wiemer (1975); Sneden et al. (1978).

$$R_V \approx 3.1 \tag{2.3}$$

To convert the extinction coefficient in V to the broadband filters of GALEX, SDSS and 2MASS, we assume that the extinction curve of dust over our wavelength region of interest follows a $\lambda^{-0.7}$ profile as suggested in the work on intergalactic dust by Charlot & Fall (2000). Furthermore, we define each broadband filter as though it were a single wavelength λ_{eff} for the purposes of this relative-extinction approximation. For filter X, we calculate its extinction coefficient as:

Filter Set	Filter	$\lambda_{\rm eff}$ (Å)	A_X
GALEX[a]	FUV	1516	2.468
GALEX	NUV	2267	1.862
SDSS[b]	U	3520	1.368
SDSS	G	4800	1.101
Johnson[b]	V	5510	1.000
SDSS	R	6250	0.916
SDSS	I	7690	0.792
SDSS	Z	9110	0.703
2MASS[b]	J	12200	0.573
2MASS	H	16300	0.468
2MASS	K	21900	0.381

Table 2.1 Characteristics of the filters used in this study.

[a]Source for $\lambda_{\rm eff}$ for GALEX filters: http://galex.stsci.edu/GR2/?page=faq
[b]Source for $\lambda_{\rm eff}$ for SDSS, 2MASS & Johnson Filters: Binney & Merrifield (1998)

$$A_X = A_V \left(\frac{\lambda_{X,\rm eff}}{\lambda_{V,\rm eff}} \right)^{-0.7} \tag{2.4}$$

Table 2.1 displays the extinction coefficients for each filter used, as well as other pertinent filter data. As is expected, the extinction coefficient from the UV filters is significantly higher than for the optical and IR portion of the spectrum.

2.2.4 Production of Images

Now, we have determined the light output in a variety of broadband filters for stellar populations, and we have a prescription for the attenuation of the light due to the presence of dust. We must combine this information with the spatial coordinates of the star particles

and gas grid in order to create a projected image of the cluster. We produce a code that arbitrarily assumes the observer is looking into the simulation space from the +Z direction (in an X,Y,Z coordinate space), although it is trivial to execute the code from different observer orientations through simple trigonometric transformations.

The code divides the entire simulation space into slices in Z. For each slice in Z starting with Z=0, we project the light from each stellar particle present in the slice onto an X,Y array representing the final image that the observer sees. Next, we calculate the column density of gas present in the slice, and attenuate the light that has accumulated in the final image array according to our dust prescription, since this light comes from behind the gas relative to the observer. We conduct these two steps—adding starlight of a slice to the final image array, then attenuating existing starlight in the final array by the gas present in the slice—on each slice in Z of the simulation space from back to front. This produces an image of the simulation as an observer might see it.

The efficacy of this method is directly proportional to the thinness of our slices in Z. As one increases the slice thickness, there is less differentiation between the strata of gas and stars, resulting in overapplication of dust. Therefore, we choose our slices to have a thickness of 1 kpc, smaller than the the highest resolution gas cells to assure that a star's light is only reddened by the gas in front of it.

Figure 2.3 is a color image constructed from the SDSS I, G and GALEX NUV images as red, green and blue, respectively. Recall that the Near UV is much more sensitive to recent star formation than the Near IR, thus galaxies showing blue characteristics in

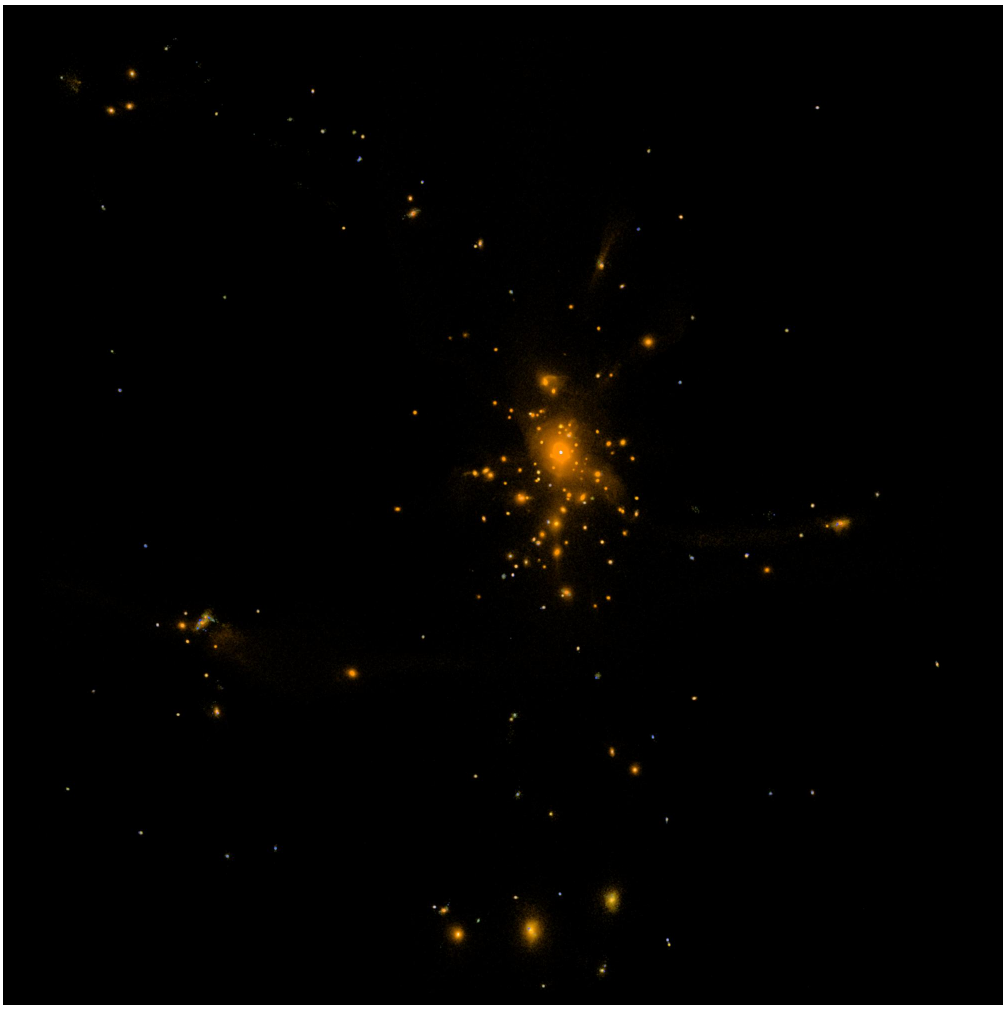

Figure 2.3 This color image of our cluster was taken from synthetic SDSS I, G and GALEX NUV bands images to represent red, green, and blue respectively. The central region of the cluster is dominated by older, red galaxies (lots of large ellipticals), whereas the outskirts have smaller, blue, irregularly-shaped galaxies with active star formation. There is an active merger in the middle of the left side.

the image suggest recent formation activity, whereas red objects are indicative of older stellar populations or extreme dust extinction. Notice that the blue objects tend to be found in the outer regions of the cluster, but the central portion possesses a redder population. One can even see an interacting system in the middle left portion of the

image, where two merging galaxies seem to be causing a burst of star formation. The cores are reddened by large column densities of cool gas, but there is blue star formation occurring around these cores. Also, the large ellipticals tend to be the reddest systems, as opposed to the smaller, irregular-looking, blue galaxies dotting the landscape consistent with observational trends.

2.2.5 Photometric Reduction of Simulated Data

After having produced the synthetic image of the cluster, it must be analyzed in a way similar to its observational counterpart. Therefore, we run Source Extractor (`SExtractor`) (Bertin & Arnouts 1996)—an aperture photometry program for measuring flux of objects in fits images. SExtractor is used in a variety of contexts such as in analyzing large galaxy surveys like SDSS or GALEX, where one seeks a standard, automated program for extracting information from a large collection of images. SExtractor works by first estimating the background noise in an image, then identifying objects above a certain detection threshold (e.g. 3σ), and lastly measuring the brightness of the sources using aperture photometry with an isophotal correction. This results in a catalog of objects and their corresponding luminosities.

By its nature, our simulation data is noise free, so it is likely that `SExtractor` would identify even very small clumps of stars as galaxy detections. However, these faint objects could never be measured in observational data due to the presence of background noise (e.g. sky background; read noise; discrete, unresolved sources; etc.) Therefore we add

a low level of gaussian noise to the background of each of our projected images of the simulated cluster prior to identifying photometric clumps. To determine the appropriate intensity of the noise, we go through each of the synthetic images and identify faint galaxies ($\sim 10^{9.5}$ $M_\odot$) that we think are the appropriate threshold of detection for `SExtractor`. Then we add gaussian noise with a full-width half-maximum (FWHM) equal to one-tenth the peak brightness of these faint galaxies. Once the gaussian noise is added to the images, `SExtractor` detects sources in the R band. Using these locations and apertures, it photometers each of the images accordingly (to ensure universal centroiding and identification of the same galaxies across all filterbands). The output is catalogued for comparison with the observational sample.

2.2.6 Observational Sample

Our observational sample of cluster galaxies comes from the Abell Cluster Catalog. We use the subsample of galaxies that have SDSS photometric and spectroscopic observations as well as photometric observations from the GALEX Mid-Imaging Survey (MIS). Additionally, the cluster must be somewhat centrally located in the GALEX images (in order to resolve all of its nearby galaxies). There are 87 clusters that fit these criteria. SDSS takes spectroscopic observations of galaxies brighter than r′ = 18.15, and the GALEX MIS resolves galaxies as faint as FUV,NUV = 23 mag AB.

Most of our galaxies are found at low redshift, making them appropriate for comparison with our simulated dataset at z = .0004. Figure 2.4 shows the redshift distribution of

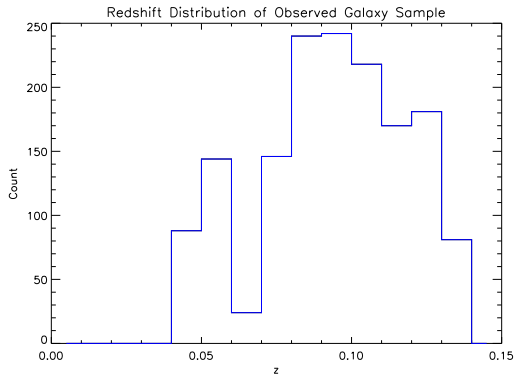

Figure 2.4 Redshift Distribution of Observed Galaxy Sample

the observational galaxy sample with its peak about $z = 0.1$.

The galaxies are selected out of this sample based on their proximity to the cluster center in both redshift and projected distance. Because many Abell clusters do not have measured redshifts, we use galaxies near each projected cluster center to look for a peak in the SDSS redshift distribution. We call this peak the "cluster velocity". This is done for all clusters. We select our sample galaxies from those galaxies within $z = .005$ ($\pm$1500 km/s) of their cluster velocity and those within 1.5^{o} projected distance ($\approx$ 10 Mpc at $z = .01$) of their cluster center to assure we get a good sample of galaxies. These constraints reduce the observational sample to just over 1600 galaxies, still nearly a factor of ten larger than our simulated cluster galaxy sample. Despite these selection criteria, field galaxies (foreground and background) do contaminate our cluster sample, because the field volume sampled is probably $\approx 10 - 20$ times the cluster volume. However, the redshift and astrometry uncertainties of both the galaxies and cluster centers are large enough to warrant this larger sample volume.

2.3 Results

2.3.1 Color-Magnitude Relation

We begin our analysis by comparing the distribution of the galaxies from the simulated and observational sample in color-magnitude space. Figure 2.5 displays four different color-magnitude diagrams (CMDs) using different filter pairs FUV-U, NUV-R, U-I, and G-Z. Here, the blue symbols represent the observed sample, whereas the simulated galaxies are plotted as red diamonds. In addition to this, Figure 2.6 displays several histograms of NUV-R color of the two populations. The plot at the top is a histogram of the colors of all the members of each dataset. Below this plot are six others histograms, each one the subset of galaxies within a single R-magnitude slice. These plots aid in analysis of the discrepancies between the two datasets, in particular the production of overluminious galaxies in the simulation. Here again, red represents the simulated galaxy population and blue is used for the observational dataset.

Initially, we see that there are some strong similarities between the observational and numerical samples. They tend to occupy the same region of color-magnitude space, and both exhibit the bimodality in color previously described in the literature (Strateva et al. 2001; Blanton et al. 2002). Clearly, the bulk behavior of the simulated and observed samples agree. That said, there are some discrepancies present which require further discussion.

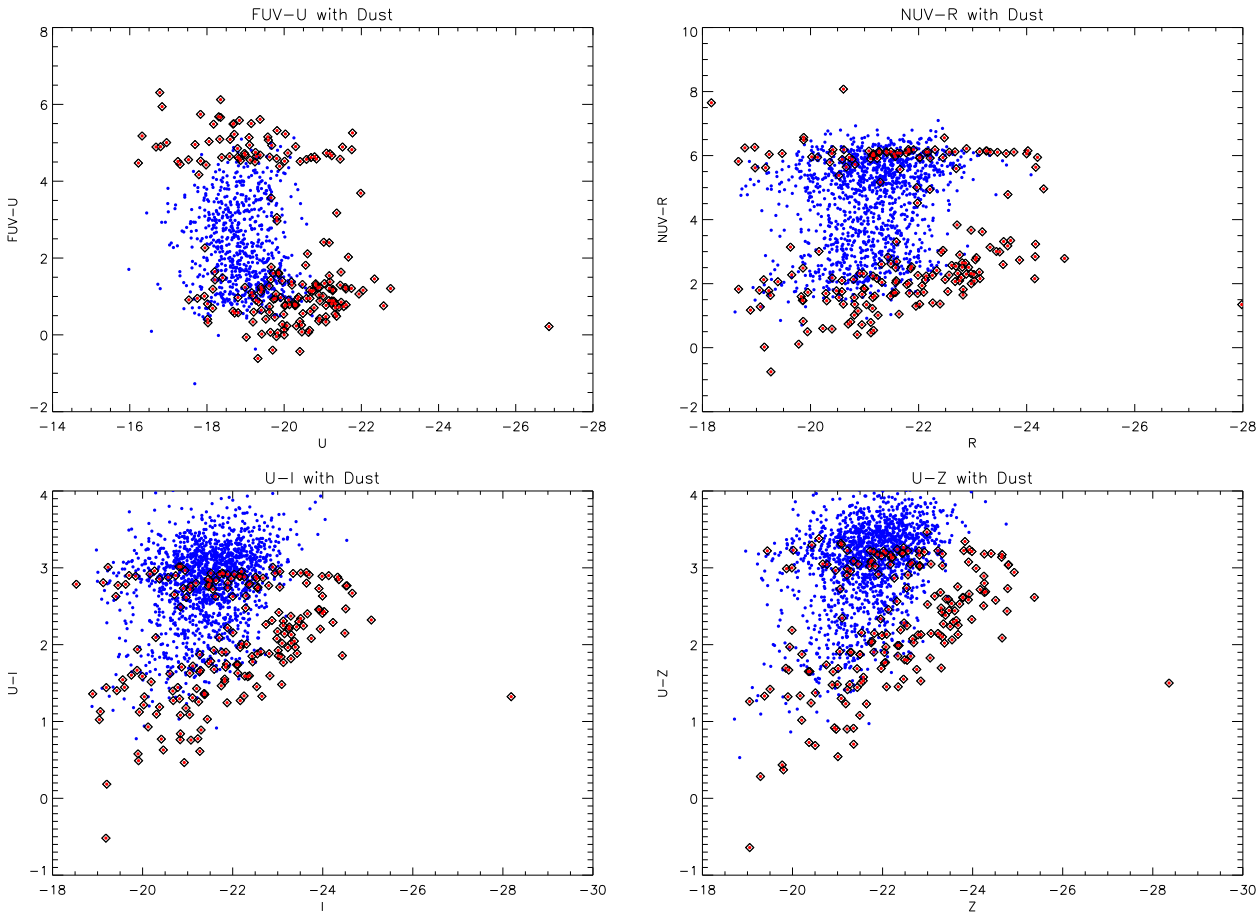

Figure 2.5 Color magnitude diagrams for the simulated and observed galaxy datasets in different color pairs. Going clockwise from upper left: FUV-U, NUV-R, U-Z, and U-I. Red diamonds represent the simulated data, and blue circles symbolize the observed galaxies.

2.3.1.1 Extreme Bimodality in Color

The most apparent difference between the two samples in the figures is the extreme bimodality found in the colors of the simulated galaxies. In each plot in Figure 2.5, the simulation data shows a population of red galaxies, a population of blue galaxies, and very few galaxies lying between the two color groups. In particular the simulated red-

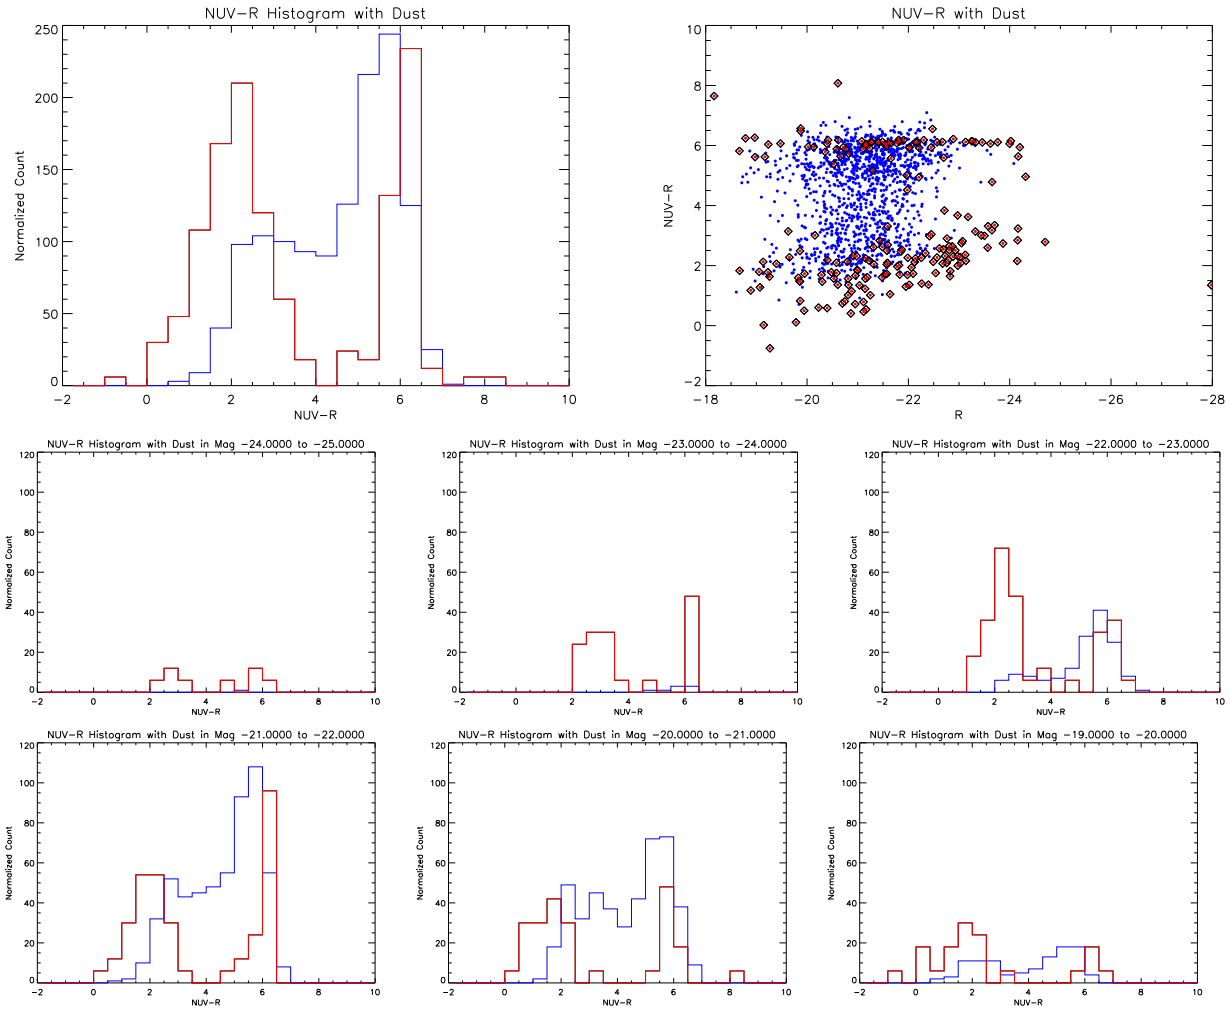

Figure 2.6 Histograms of NUV-R color in slices of magnitude for simulated and observed galaxy samples. Red symbolizes simulated data, and blue symbolizes observational data. The top two panels represent the entire samples of observed and simulated galaxies in color space. The bottom six panels plot histograms of the color of the galaxies in slices of R magnitude. Because the simulated data sample is approximately ten times smaller than the observed dataset, it has been normalized for an even comparison. The simulated dataset dominates the histograms in the bright galaxies ($-23 \geq R \geq -25$). Furthermore, it overdiscretizes the galaxy population into the red and blue sequences inconsistent with the observational sample.

galaxy group is defined by a very narrow range in color most easily seen in Figure 2.6, and the blue-galaxy population appears to be more blue than its observational counterpart. This bimodality is found, albeit to a lesser degree, in the observational data in all but the FUV-U plot of Figure 2.5. Because of the depth of the survey, all galaxies with $NUV - r < 5$ should be detected in this sample. However, our sample is less complete for some of the reddest galaxies at $NUV - r \approx 6$.

This bimodality of color is well-documented in the literature (Strateva et al. 2001; Blanton et al. 2002), and the respective loci are referred to as the red and blue sequences. While it is not well understood, it seems clear the red galaxies are older populations, usually ellipticals near the central part of a cluster, whereas the blue sequence is made up of galaxies with recent star formation found more commonly in low-galaxy-density environments. It remains to be seen how galaxies travel along these sequences, or what is the dominant mechanism for shifting them from one sequence to the other.

As was noted before, so far we have been unable to incorporate metallicity of the stellar populations into the population synthesis runs. The addition of this information will almost certainly produce an increased spread in the colors of the red-sequence galaxies, since it is well known that metallicity plays a distinct role in the color evolution of a stellar population (Bruzual & Charlot 2003). This may help to reduce the unrealistically tight color range of the red sequence in the simulated data and possibly move some red galaxies into the intersequence range.

Another reason for the severe bimodality exhibited by the simulated data is that the

blue galaxies appear to be too blue. The slope of the blue sequence tends to match that of the observed dataset, but it is blueward of the observational blue-sequence locus in all four CMDs. This may indicate that our dust models are not extinguishing enough light from young stellar system. As is discussed in Charlot & Fall (2000), it may be necessary to further attenuate light coming from young stellar particles (Age $\leq 10^7$ years) in order to mimic the obscuring effects of the dense clouds in which they are born. In principle, this should just push the locus of the simulated blue sequence upward in color space, and thereby reduce the severity of the bimodality in the simulation data.

Furthermore, the well-documented overcooling problem (e.g. Balogh et al. 2001) is likely responsible for most of the extreme bimodality in galaxy color. The lack of spatial resolution and treatment of stellar feedback in the simulation leads to an overproduction of stars in regions containing cold gas, yielding unnaturally high star formation rates in star-forming galaxies. Not only does this produce overly blue star-forming systems, but it means early star-forming galaxies effectively consume all of their fuel quickly, leaving no residual star formation in red systems at later epochs. Thus, we find blue galaxies are too blue, and red galaxies are too red, exactly the observed condition of the simulated sample.

It would take very little continued star formation to shift one of the red galaxies from their tight sequence and into the intersequence void, but without any gas as fuel in these systems, no such star formation is possible. Currently, there is no prescription in `Enzo` for mass loss from old population stars, such as asymptotic giant branch stars (AGB) or

planetary nebulae (PNe). If such a prescription were added, it would produce a small amount of gas in older galaxies, providing fuel for enough star formation to push some ellipticals into the intersequence region of the CMD.

The total lack of a bimodality found in the observational FUV-U CMD (upper left corner of Figure 2.5) is likely due to the *UV upturn*—the appearance of far UV light coming from old, elliptical galaxies. O'Connell (1999) determined the blue edge of the horizontal branch population (HB) is responsible for producing this effect. Thus, the old stellar populations present in what are normally red-sequence galaxies in most filter pairs, show up as blue populations because of their HB stars producing Far UV radiation. It is mysterious that our simulated data produces distinct red and blue sequences in this CMD, but resolving some of the previously stated issues may rectify this problem as well.

2.3.1.2 Overluminous Galaxies

Another major discrepancy between the observational sample and the simulated sample is the existence of simulated galaxies up to two magnitudes more luminous than those observed in nature. These overluminous galaxies are found in both the red and blue sequence of the simulated data as is illustrated in the brightest histograms in Figure 2.6. We believe that these are a result of a lack of star formation feedback mechanisms in the simulation, which causes overproduction of stars and hence overluminous galaxies in all colors. Insufficient feedback takes the form of supermassive blackholes (SMBHs) and supernovae in the simulation.

Even though there is a prescription for supernova feedback when stars are formed in the simulation, most of the energy that the SNe re-inject into the surrounding medium goes straight into warm gas where the cooling time is short and so it is efficiently radiated away. Because of the coarse resolution possible in our simulation, expanses of hot, ionized gas and patches of cold, dense gas in individual galaxies are not individually resolved. Instead, the insufficient spatial resolution merges them into cells of warm gas where cooling processes are very efficient (Osterbrock & Ferland 2006). If we had much higher spatial resolution (which is not currently computationally feasible) we could mimic the true behavior of the multiphase ISM at galactic scales, where much of the SNe energy is deposited in high-temperature gas. This plasma has a much slower cooling time than the warm gas, therefore it retains the thermal energy better, and this can lead to quenched star formation.

Furthermore, the lack of a prescription for supermassive blackholes (SMBHs) in the centers of galaxies only exacerbates the overcooling problem. The accretion disks around SMBHs are believed to be responsible for turning off star formation at certain epochs by heating up or blowing out galactic gas. In future generations of `Enzo` with a prescription for SMBH feedback and with higher spatial resolution, it is likely that these overluminous galaxies will not be produced.

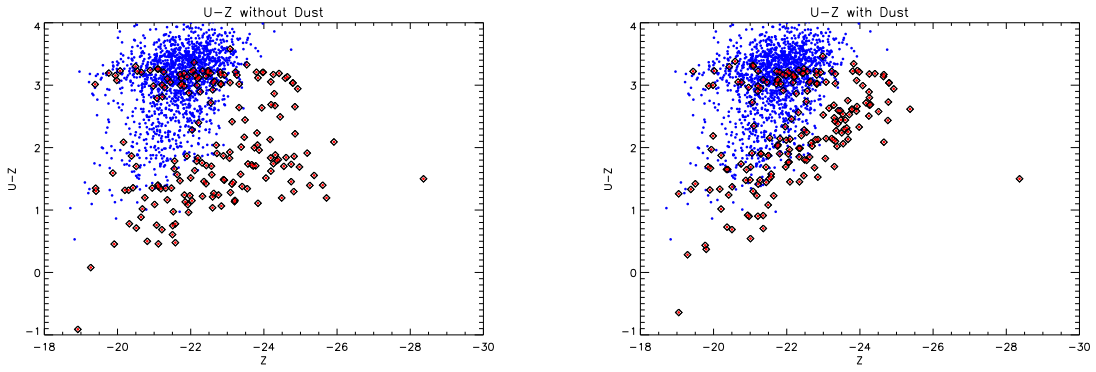

Figure 2.7 Color magnitude diagrams of U-Z color. Left side was produced without any dust processing on the simulated data; right side takes into account dust reddening effects. The slope of the blue sequence in the simulated data changes drastically with the dust-correction from near horizontal on the left to almost 45 degrees on the right. This indicates that there are no luminous blue galaxies without dusty gas enveloping them in the simulation and perhaps in nature.

2.3.1.3 Slope of Blue Sequence

The slope of the blue sequence varies with the filter pair chosen, such that filter pairs closer in wavelength to each other have more *horizontal* blue sequences. Is this simply due to the reddening effects of dust, or could it be some inherent behavior in the galactic population? In Figure 2.7, we plot two CMDs for the color pair U-Z, one before dust extinction processing and one after. Because these filters are significantly separated in the wavelength space, there is a steep slope to the blue sequence in the dust-processed CMD on the right. However, the CMD on the left without dust correction has an almost horizontal blue sequence. Therefore, we can state that the slope of the blue sequence is almost entirely due to the biasing effect of the dust in the simulations, which preferentially reddens the very luminous blue galaxies. This makes sense, since the brightest blue galaxies should

be ones that are actively forming stars with adequate supplies of cool gas available to them. This cool gas will redden their starlight. Furthermore, this observation implies there are very few luminous blue galaxies without any local cool gas. The observational population exhibits this same blue sequence slope as in the dust-corrected simulation data. This implies that, in nature, very few luminous blue galaxies lack localized cool gas.

2.3.2 Color-$r_{cluster}$ Relation

In Figure 2.8, we have plotted NUV-R color against the projected distance of each galaxy from its cluster center, $r_{cluster}$. The left plot shows the radial distribution in linear space, whereas the right plot displays the same information in log space. In both datasets, there is a trend that the red galaxies are found at smaller cluster distance than the blue galaxies, implying a paucity of galaxies with active star formation in the core of the cluster. This observation agrees with the evidence in Figure 2.1, where we find there is little cold gas found in the core region of the cluster.

Astronomers have observed this behavior before, but suggestions for its origin remain inconclusive. Some believe this behavior is a function of the harsh environment found in the cores of clusters, where the hot intracluster medium (ICM) and tidal harrassment strip out cool gas from galaxies and suppress star formation. Others suggest that these red galaxies are simply remnants of the galaxies that formed the cluster long ago, and hence they have already exhausted their cold gas fuel resources. It is likely that each of these plausible scenarios is responsible. Future work includes mapping the star formation

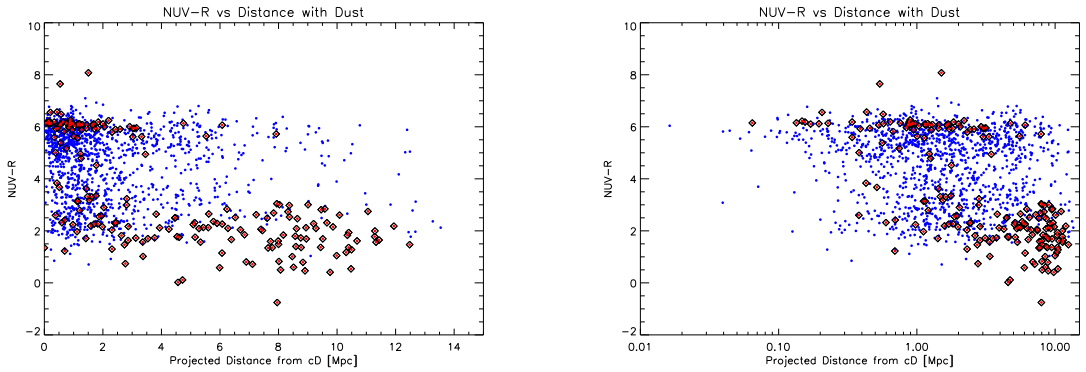

Figure 2.8 Color-$r_{cluster}$ diagrams in linear and logarithmic space. The blue-sequence galaxies are found more readily in the outer regions of the cluster, whereas the red-sequence galaxies are sequestered in the cluster core.

history for each of the simulated galaxies correllated with position relative to the cluster center in an attempt to pin down which possible mechanism dominates in producing the red-galaxy-sequestering behavior.

While there is general agreement of the two datasets in the plots, we find far too few simulated red galaxies in the outer regions ($r_{cluster} = 5 - 15$ Mpc) of the cluster as compared to the observational sample. Partly this may be due to field galaxies polluting the observational dataset, but it is likely to be a small effect. Evidently, the simulation underproduces red galaxies in the outskirts of the cluster compared to what is observed in nature. This effect could be another consequence of the overcooling problem discussed above, causing all galaxies outside the cluster center to accrete cool gas and spawn star formation and thus turn blue.

2.4 Conclusions

In this chapter, we have explained a method for comparing output from N-body and hydrodynamics simulations to photometric observational data. Application of this method to a cluster simulation proved fruitful for ascertaining the shortcomings of the `Enzo` simulation physics, as well as providing us with some insights into the nature of galaxy evolution in a cluster environment. Our major results are as follows:

- We wrote a code that converts simulation data into two-dimensional photometric images. The code takes into account the stellar populations present in the simulation by using the Bruzual-Charlot population synthesis code. Assuming a fixed gas-to-dust ratio, our code approximates radiative transfer through dust in the simulation and attenuates the starlight accordingly on its way to the simulated observer. This code is largely organized into a pipeline, such that it can be applied to any `Enzo` data output with little user interaction. This makes it useful in a variety of future applications in which one wants to see what simulated data might look like if it were observed photometrically. Additionally, we intend to generalize the simulation input procedure, so that it can accept inputs from other hydrodynamics simulations and possibly release this code into the public domain for other scientists to use freely.

- Using this code, we successfully produced projected images of an `Enzo` simulation of a cluster of galaxies and extracted galaxy information using `SExtractor`. We compared the extracted information to data from a statistically significant sample

of clusters, a subset of the Abell Cluster Catalog in the GALEX, SDSS and 2MASS filters. For analysis purposes, we produced color-magnitude diagrams, and plots of the distribution of galaxy color with distance from the cluster core.

- Our simulated data somewhat reproduce the red and blue sequences present in the distribution of observed galaxies in color-magnitude space. While there are numerous discrepancies between the distribution of the two datasets, the bulk behavior is similar, which makes us believe that many of the approximations we made in our method are valid. Similarly, the color-$r_{cluster}$ plot exhibits generally consistent behavior between the two samples, where blue galaxies tend to be found more often in the outskirts of clusters, while red galaxies are found more commonly in the cores.

- We discovered that the main source of the non-horizontal slope in the blue sequence is due to dust. By comparing the simulated sample before and after dust processing, we saw a marked change in the blue sequence slope. This implies that there are few luminous blue galaxies lacking massive cool gas envelopes and the dust found in them. Since a similar slope is detected in the observational sample of cluster galaxies, it seems plausible that the same behavior occurs in nature.

However, because there are some discrepancies between the two data samples in our diagrams, future work is warranted to: (i) better account for the overcooling problem through increased spatial resolution and better treatment of stellar and AGN feedback; (ii) integrate metallicity of stellar particles into our population synthesis routine; (iii)

increase the attenuation of light coming out of very young stellar populations to mimic the presence of dense dusty clouds; and (iv) create a prescription for mass-loss of highly evolved stars in `Enzo` in order to regenerate gas in older stellar populations.

Additional work also includes the application of these methods to hydrodynamical simulations in different contexts and scales including any simulation which produces stellar "particles" well characterized as an coeval population of stars. By creating synthetic observations of stellar or galactic systems undergoing specific evolutionary processes like tidal interaction or ram-pressure stripping, it may be possible to identify photometric signatures of certain galaxy evolution mechanisms.

Chapter 3

A Kinematic Comparison of Simulated and Observed Galaxies for Constraining Subgrid Models

3.1 Introduction

In cosmologies dominated by Cold Dark Matter (CDM), galaxy rotation is produced by gravitational tidal torques arising from the hierarchical collapse of structure. Analytic models and N-body simulations have shown that this can produce enough angular momentum to explain the observed sizes of disk galaxies (Fall & Efstathiou 1980; Mo et al. 1998). However, computational models including gas dynamics have struggled to reproduce realistic disk galaxies. Such models initially produced undersized disks with

low angular momentum (Navarro & Benz 1991; Navarro & White 1994). Later work did generate disks with approximately the correct extent (Steinmetz & Navarro 2002; Abadi et al. 2003), but these had oversized stellar spheroidal components and therefore unnaturally large core circular velocities (Steinmetz & Navarro 1999) far in excess of those found observationally. This *angular momentum problem* remains one of the major shortcomings of the CDM paradigm.

The origin of the angular momentum problem is not entirely understood, but it probably stems from the fact that such simulations do not achieve sufficient spatial and mass resolution to correctly model the appropriate physical processes. For example, insufficient spatial resolution leads to the spurious mixture of hot and cold components of the ISM, producing an artificial warm component which is very efficient at radiating away energy (Katz 1992; Katz et al. 1996; Steinmetz & Muller 1995). Thus, underresolved gas in such simulations can cool very quickly, and in cooling it loses its pressure support and collapses into dense knots of material. These knots interact with the galactic dark matter component through dynamical friction processes, and much of the angular momentum of the gas knots is transferred to the dark matter halo of a galaxy (D'onghia et al. 2006). Consequently, the knots tumble into the center of the galaxy to produce a dense cusp of material in the core of the simulated galaxy. Typically the buildup of these massive cores with cold gas and stars occurs rapidly, and even nascent simulated galaxies exhibit evidence of these cusps as early as $z \sim$ 4-5.

Energetic feedback from star formation events has the ability, in theory, to alleviate

this problem either by heating and "puffing" up the collapsing knots so they are more easily disrupted before losing their angular momentum (Weil et al. 1998), and/or by preferentially ejecting low angular momentum gas (e.g., Brook et al. 2011). However, existing simulations are unable to resolve the detailed structures of star-forming regions. Individual star-forming and stellar feedback events occur on parsec scales. Stellar feedback processes can be resolved in simulations confined to local regions of the ISM (e.g., Joung & Mac Low 2006) but not in cosmological simulations, which need to accurately co-evolve a galaxy's environment on scales of > 10 Mpc.

Therefore, cosmological simulations simplify and parameterize star formation and stellar feedback on scales more easily met by current computational resources (i.e. scales of $10^2 - 10^3$ pc). A variety of techniques have been suggested in order to achieve this puffiness within the confines of low-resolution models. In the most primitive prescription, stellar feedback is simply the return of energy from newly created star particles to their surrounding gas, usually in the form of thermal energy. Star "particles" in these simulations typically represent clusters of stars of mass $10^2 - 10^5$ $M_\odot$, so the thermal feedback is justified as the energy output from the most massive stars in the cluster becoming type II supernovae shortly after the creation of the star particle. In nature, this supernovae heating produces a small component of very hot gas surrounding the stellar population, so hot that its cooling time is very long. Unfortunately at the resolution currently achievable in cosmological simulations, this primitive thermal prescription for feedback deposits the same SN energy into a much larger reservoir of gas, which does not reach the same high

temperature as it should. This now-warm component of gas can easily radiate the excess energy away, cool further and proceed with star-formation runaway thus defeating the purpose of the feedback (e.g., Steinmetz & Navarro 1999).

Building from these failures, a number of research groups artificially turn off cooling in a gas parcel for a period of time ($t \sim 10^7$ yr) after a cluster of stars has formed out of it (e.g., Gerritsen 1997; Thacker & Couchman 2000; Sommer-Larsen et al. 2003; Stinson et al. 2006; Governato et al. 2007; Agertz et al. 2011; Colín et al. 2010; Piontek & Steinmetz 2011; Guedes et al. 2011). This method is justified as an application of the Sedov-Taylor blast wave solution for a Type II SN (Taylor 1950; Sedov 1959), which blows out any cold media from the immediate environment of a star formation event. Using this prescription, any gas in a galaxy which starts to collapse into knots will reach the star formation criteria, form a star, and then heat up without any allowed cooling, thus preventing further collapse. Not surprisingly, these research groups have found some success with this method, yielding simulated galaxies with reduced inner rotation curves due to less massive bulge components; however, gas parcel masses and sizes in cosmological simulations of this sort are typically too large for the Sedov-Taylor solution to apply (see Section 3.2.1.7). Thus despite the successes of the cooling suppression feedback model, the community continues to search for other more physically-motivated solutions.

Another subgrid model for feedback (i.e. on scales smaller than the true resolution of the simulation) is to inject kinetic energy directly into the gas; this can alleviate the problem of thermal energy being radiated away. For example, some studies (e.g. Springel

& Hernquist 2003b; Scannapieco et al. 2006; Oppenheimer & Davé 2008) using Smoothed Particle Hydrodynamics (SPH) give some of the SN energy to individual gas particles in the form of momentum. This method can result in significant mass outflows (by design), but at the cost of decoupling wind particles from hydrodynamic interaction for a period of time. An alternate approach, to keep wind particles coupled to the disk gas was explored by Schaye & Dalla Vecchia (2008). Both approaches help but, by themselves, do not appear to generate realistic rotation curves.

In addition, Truelove et al. (1997) showed that insufficient resolution in a simulation can lead to artificial fragmentation of the gas, perhaps resulting in a further overproduction of stars. One way to prevent artificial fragmentation is to add additional (numerical) pressure in high-density, low-temperature regions to ensure that the Jeans length is always resolved (Machacek et al. 2001; Robertson & Kravtsov 2008). This can be achieved by modifying the equation of state (EOS) itself, making it stiffer in order to provide an additional source of pressure to gas in denser regions (Schaye & Dalla Vecchia 2008; Ceverino & Klypin 2009; Agertz et al. 2011). A polytropic EOS ($P \propto \rho^{\Gamma}$) with $\Gamma = 4/3$ will keep the ratio of Jeans length to resolution length constant (assuming Lagrangian resolution such that the resolution length decreases as $\rho^{-1/3}$ – for fixed resolution $\Gamma = 2$ is required), but even stiffer relations have been used. For example, Agertz et al. (2011) ran simulations with such an equation of state, where in low-density regions it behaved as an ideal gas, but in high-density (star forming) regions it followed a polytropic equation of state with $\Gamma = 2$.

In this paper, we undertake an investigation of galaxy formation using an Adaptive Mesh Refinement (AMR) hydrodynamics code. The majority of work in this field has used SPH codes, and so this allows us to investigate the problem from a new angle. Although there has been some work with AMR codes (Joung et al. 2009b; Ceverino & Klypin 2009; Colín et al. 2010; Agertz et al. 2011), there has not been a clear demonstration that an equivalent AMR calculation (i.e. one without a subgrid feedback model) actually does reproduce the classic SPH result.

We begin by simulating a set of five halos without any feedback or subgrid model (except a minimum pressure support to prevent artificial fragmentation). We find, in agreement with SPH codes that a large, concentrated bulge is produced, resulting in a rotation curve that rises to ~ 500 km/s at $r \sim 1$ kpc. We then vary a number of numerical and physical parameters in order to understand how sensitive the result is to our a choice of parameters.

The paper is organized as follows. Section 3.2 describes the details of our hydrodynamics code, our initial conditions and the relevant parameters for this study. In Section 3.3, we present the results of our simulations including the five canonical runs, our resolution study, and our modified runs. Section 3.4 is a discussion of our results and their implications. Finally, Section 3.5 summarizes our conclusions and makes predictions for future solutions to the angular momentum problem.

3.2 Methodology

3.2.1 Code

Our simulations were performed using *Enzo*[1], an Eulerian, three-dimensional, grid-based hydrodynamics code that employs adaptive mesh refinement in order to achieve targeted regions of high resolution in a cosmological volume (Bryan & Norman 1997; O'Shea et al. 2004). Gas is discretized on the grid, but dark matter and stars are treated as particles. The *ZEUS* hydrodynamics code (Stone & Norman 1992) is used to evolve the gas on the grid. *Enzo* includes gas, self-gravity, a non-equilibrium model for H and He ionization and cooling, a metagalactic ultraviolet background (Haardt & Madau 1996), and equilibrium cooling due to metals (although for the runs described in this paper, we do not include metal cooling).

3.2.1.1 Star Formation

Star formation is modeled using a simple criteria based on Cen & Ostriker (1992). A grid cell will produce a star if: (i) the overdensity in that cell exceeds a given value ($\delta_{\rm SF}$), (ii) the mass of gas in the cell exceeds the local Jeans mass, (iii) there is locally convergent flow (i.e. the velocity divergence is negative) and (iv) the cooling time for the gas to collapse is less than the dynamical time in that cell (or the temperature is near the minimum allowed, around 10^4 K). If a grid cell meets all the previous criteria then gas is converted into a "star

[1]http://enzo-project.org

particle", as calculated using

$$m_* = \epsilon_{\rm SF} \frac{\Delta t}{t_{\rm dyn}} \rho_{\rm gas} \Delta x^3 \tag{3.1}$$

where $\epsilon_{\rm SF}$ is the star formation efficiency (more properly the efficiency per dynamical time), Δt is the size of the time step, $t_{\rm dyn} = (3\pi/32G\rho)^{1/2}$ is the dynamical time and $\rho_{\rm gas}$ is the gas density. If the resulting star particle would have a mass above a minimum mass, $M_{*,\rm min}$, it is formed immediately. This mass is chosen so that a large number of small star particles will not slow down the simulation. However, if the star particle would have a smaller mass, the probability that it will form is equal to the ratio of the mass of the projected star particle to $M_{*,\rm min}$. If the star particle is then formed, its mass is the minimum of $M_{*,\rm min}$ and 90% of the mass in the gas cell (see also Tasker & Bryan 2006). For all of the simulations in this study, we used a value of $M_{*,\rm min} = 10^5$ $M_\odot$. Our default value for $\epsilon_{\rm SF} = 10^{-2}$, although this was varied in some runs (see Table 3.1). We adopt $\delta_{\rm SF} = 1000$, corresponding approximately to a density threshold of 0.1 cm^{-3} at $z = 3$; this value simply ensures that star formation is limited to strongly non-linear regions.

3.2.1.2 Resolution

Because *Enzo* is an AMR code, it can dynamically refine the spatial resolution of the grid when certain criteria are met in a particular region of the simulated volume. We set the grid-refinement criteria to increase refinement whenever the dark matter mass in a cell in larger than four times the dark matter particle mass, with an equivalent criterion for the baryonic mass. When refined, the cell resolution is increased by a factor of 2. The

placement of these refinement regions is recalculated regularly throughout the simulation, to assure that moving or emerging regions of interest are always well-resolved. For the canonical runs, we cap this refinement when our cell sizes reach 9 levels of refinement or 425 comoving parsecs. In addition, we conduct a resolution study in Section 3.3.5.1 where we vary the resolution from 7 to 10 levels of refinement (212 - 1700 comoving parsecs).

3.2.1.3 Minimum Pressure Support

To prevent artificial Jeans fragmentation, for most runs we implement the minimum pressure support described in Machacek et al. (2001) such that the ratio of the the Jeans length to cell size, $J = L_J/\Delta_x$, is at least 8. We do this by adding an additional artificial pressure to the most highly refined grid cells such that this ratio is always maintained. The addition of this pressure is intended to prevent gas clouds from collapsing below the resolution scale of the simulation, which could cause spurious numerical effects, such as artificial fragmentation.

3.2.1.4 Star Formation Efficiency

For our canonical runs, we set the star formation efficiency per dynamical time $\epsilon_{\rm SF}$ to 1%, a parameter value that previous work found to approximately reproduce the Kennicutt-Schmidt relation for our chosen star formation law (Tasker & Bryan 2006). In Figure 3.1, we demonstrate that our canonical simulations roughly agree with the Kennicutt-Schmidt relation by plotting the gas surface density versus the recent star formation surface density

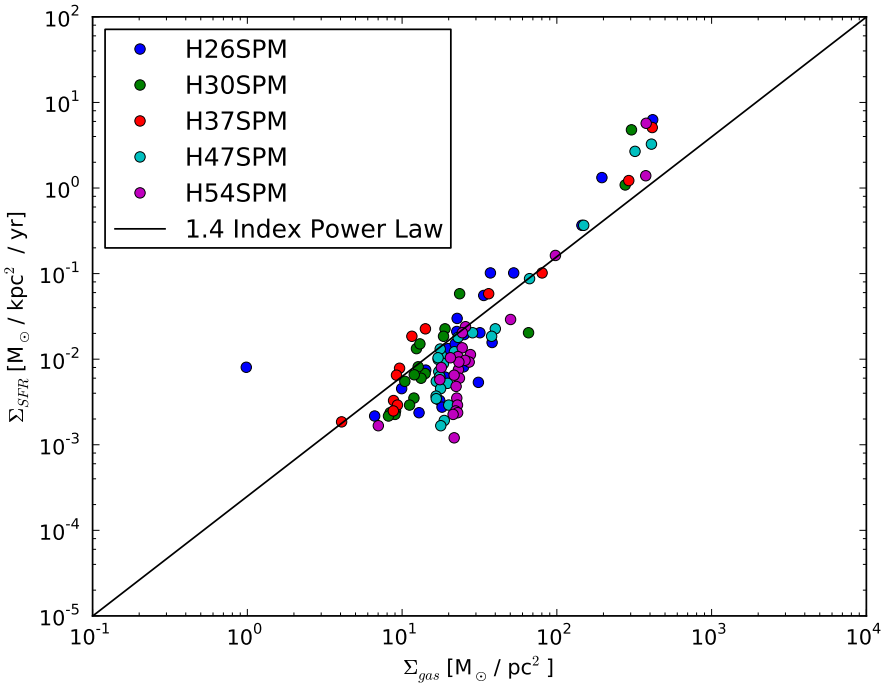

Figure 3.1 We plot gas surface density versus star formation rate surface density for our five canonical runs. We overplot the power law $\Sigma_{SFR} = 2.5 \times 10^{-4} \cdot \Sigma_{gas}^{1.4}$ (in these units) representing the Kennicutt-Schmidt local relation between star formation rate and gas surface density (Kennicutt 1989).

($\text{age}_{\text{star}} \leq 5$ Myr) in concentric annular cylinders of 0.25 kpc width and scale height 5 kpc each aligned with their respective gas disks. Overplot is a power-law with index 1.4, the accepted form of the Kennicutt-Schmidt law (Kennicutt 1989). Our canonical value for the parameter (the star formation efficiency per free fall time) is $\epsilon_{\text{SF}} = 10^{-2}$, which is in line with (but slightly lower than) typically suggested values (e.g. Krumholz & Tan 2007). Other galaxy formation simulations adopt different values, for example in the SPH regime Abadi et al. (2003) use a value of 3.3%, Governato et al. (2007) use 10%, and Stinson et al.

(2006) use 10% (after correcting for the fact that GASOLINE-based studies use a different definition of t_{dyn} than the other listed publications). Turning to AMR simulations, Agertz et al. (2011) use 1% for their most successful runs, Teyssier et al. (2010) also adopt 1%, while Gnedin & Kravtsov (2011) adopt 0.5%, arguing that it is a better fit to the observations of Bigiel et al. (2008). It is likely that the choice of star formation efficiency required depends both on the resolution of the simulation, as well as feedback prescription. In this paper, we explore variations in the efficiency, decreasing this parameter to explore the impact of slowing the conversion of gas into stars.

3.2.1.5 Fixed Physical Resolution

Most of the previous work on this topic was done using hydrodynamics codes in *physical* or *proper* units as opposed to *comoving* units. Our code uses comoving coordinates, which means that a fixed maximum resolution level therefore corresponds to a fixed comoving resolution at the finest refinement level (cold, dense gas is almost always refined at the maximum level in these simulations). Thus our models resolve more details in the early universe (by a factor of $1 + z$) relative to simulations employing the equivalent resolution in fixed physical units. In an effort to explore the impact of using fixed physical rather than comoving resolution, we modified the code to be able to run with a fixed physical resolution, by adjusting the maximum refinement level with redshift. We note that because of the discrete, factor of two changes in grid resolution, this only keeps the maximum refinement to a fixed physical resolution within a factor of two. We do not change the

refinement criteria, therefore our mass resolution remains the same (see Section 3.2.2). Section 3.3.5.4 describes the results of these runs.

3.2.1.6 Supernovae Feedback

Most simulations we discuss in this paper do not include explicit feedback; however, we do carry out a few runs which included prompt, energetic feedback from Type II SNe. We apply a simple prescription for stellar feedback to mimic the effects of type II supernovae. Because individual star-formation events produce star particles of $M_{\rm star} \sim 10^{4-5}$ $M_\odot$, we spread feedback over an extended period parameterized by

$$\frac{dM_{\rm form}}{dt} = M_0 \frac{t - t_0}{\tau} \exp \frac{-(t - t_0)}{\tau}, \tag{3.2}$$

where M_0 and t_0 are the initial star particle mass and star particle creation time (see also Cen & Ostriker 1992), and τ is the maximum of either the dynamical time of the gas out of which the star particle formed, or 10 Myr (to prevent unrealistically short dynamical times). We take a fraction of the rest mass energy $\epsilon_{\rm FB}$ of the forming stars as the available feedback energy. This parameter can be computed assuming an initial mass function and a minimum initial stellar mass for producing a Type II SN.

As described in Tasker & Bryan (2006), we simply add the thermal energy to the local grid cell. As the surrounding gas heats up, it increases the Jeans length and theoretically quenches star formation for a time. As was noted in Section 3.1, if the resolution of

the simulation is insufficient to clearly differentiate between a cold, neutral interstellar medium and a hot, ionized interstellar medium, then the pumping of this energy into the cold star-forming gas results in an unrealistic warm component (i.e. the mixing of hot and cold phases). This warm component sits near the peak of the cooling curve of the gas, and thus effectively radiates away its energy very quickly.

For the run which does include feedback, we take a value of $\epsilon_{\rm FB} = 3 \times 10^{-6}$, which corresponds to one 10^{51} erg SN for every 180 $M_\odot$ of stars formed. The value of this parameter is uncertain, as it depends on both the initial mass function, as well as assumptions about how the energy is radiated away immediately. Other values are used in the literature, for example Abadi et al. (2003) employed a feedback prescription which corresponds to, in our definition, a value of $\epsilon_{\rm FB} = 5.6 \times 10^{-6}$. They, like others, found that this energy was quickly radiated away, and so the value had little impact. As we discuss in more detail below, Stinson et al. (2006) added a cooling suppression model, and argued for a value of $\epsilon_{\rm FB} = 4.3 \times 10^{-7}$, although Governato et al. (2007) adopted $\epsilon_{\rm FB} = 2.6 \times 10^{-6}$, using a very similar model. Turning from SPH to AMR simulations, we note that Agertz et al. (2011) used $\epsilon_{\rm FB} = 5.6 \times 10^{-6}$ for their standard runs (also with a cooling suppression model). Neither Teyssier et al. (2010) nor Gnedin & Kravtsov (2011) appear to have used thermal feedback in their simulations. Finally, Cen (2011) used $\epsilon_{\rm FB} = 10 \times 10^{-6}$ (larger than the usual efficiency because of a postulated contribution from prompt Type Ia SN); this work did manage to drive winds without a cooling suppression model, although they added the energy to the nearest 27 cells weighted inversely by density (and also had somewhat

worse mass and spatial resolution than used here). In summary, we see that our chosen value is within the range used by other researchers.

3.2.1.7 Cooling Suppression Models

One way to prevent this energy from being quickly lost is to turn off radiative cooling in the region immediately surrounding newly formed stars. This was first attempted by Gerritsen (1997), but has since been explored by a range of simulations (e.g., Thacker & Couchman 2000; Sommer-Larsen et al. 2003; Stinson et al. 2006; Governato et al. 2007; Agertz et al. 2011; Colín et al. 2010; Piontek & Steinmetz 2011; Guedes et al. 2011). The idea is to use the Sedov-Taylor solution for a blast wave to model the subgrid shock physics that the code cannot resolve. It might seem most straightforward to use the length and time-scales of the energy-conserving phase (i.e. Sedov phase) to control where and when to turn cooling off, but these turn out to be too small and too short (a few 10^4 years) to make a significant difference, as acknowledged by previous studies (Stinson et al. 2006). Instead, the method employs the radius and time at the end of the momentum-conserving phase under the assumption that during this phase, much of the energy is conserved in kinetic motion (and hot, diffuse gas), which the code cannot model, and would otherwise dissipate. The larger length and timescales of the snowplow phase used are further justified as the combined forces of multiple supernovae in the star particle; however, although these supernovae may interact in a complex way, it is unlikely their effects will simply add constructively.

This prescription for cooling suppression feedback has generally been used in SPH codes, but recently some AMR codes have also adopted this technique (Colín et al. 2010; Agertz et al. 2011). The common prescription is to suppress cooling for a period of time ($30 - 50$ Myr) in the gas immediately around a star-formation event, regardless of where the star goes afterwards. In our implementation, we suppress cooling of the gas in the single cell in which the star particle resides. This is done for 50 Myr after the star particle is first created. Since both of these length and time scales correspond closely to those over which energy from the star particle is injected in the simulation (the feedback follows Eq. 4.2, above), this acts to suppress cooling in newly heated gas. Given our chosen cell size (425 comoving pc in most runs), these length and timescales are similar to the region and duration of influence adopted by other researchers (Stinson et al. 2006; Colín et al. 2010).

3.2.2 Initial Conditions

For this work, we use the WMAP 5-year results as our cosmological parameters (Komatsu et al. 2009), in particular we use: $\Omega_0 = 0.258$, $\Omega_\Lambda = 0.742$, $\Omega_{\rm baryon} = 0.044$, $\sigma_8 = 0.796$, $H_{\rm o} = 71.9$ km s^{-1}. We generate our initial conditions using *inits*, a program included in the *Enzo* suite. *Inits* sets up a 128^3 particle-mesh grid, and modifies the velocity and position of the dark matter particles in each grid cell as specified by linear perturbations with the required power spectrum at $z = 99$. The initial conditions are generated for a cubic volume of $L = 20h^{-1}$ comoving Mpc on a side with periodic boundary conditions. First, the

simulation is populated with only dark matter particles at low-resolution ($M_{\rm DM} = 3.2\times10^8$ $M_\odot$) and run to $z = 0$, where candidate Milky-Way-like halos are identified. Halos are chosen based on their final mass and accretion history, preferentially selecting halos with final masses $M_{200} \sim 10^{12}$ $M_\odot$, and those which have not undergone major mergers after $z \sim 2$ (M_{200} is the mass enclosed within a radius corresponding to a mean density of 200 times the critical density). Five such halos are identified, ranging in mass from $M_{200} = 4.8 \times 10^{11}$ to 1.2×10^{12} $M_\odot$.

For each halo, the component dark matter particles are traced back to their positions in the initial conditions at $z = 99$. This Lagrangian volume is further refined with two additional levels of refinement. It is here that the initial conditions are regenerated, and in these nested boxes, we additionally refine the dark matter particle masses by a factor of 8 for each region. The resulting high-resolution dark matter particle mass within the vicinity of each halo is $M_{\rm DM} = 4.9 \times 10^6$ $M_\odot$.

A series of new simulations are performed using these new initial conditions; each one focuses on a different halo. Baryons are included in these runs. The high-resolution regions are further refined dynamically with adaptively-placed grids, using the refinement scheme described earlier

3.2.3 Description of Simulations

We conducted five canonical simulations using identical simulation parameters and the initial conditions described above. These simulations are referred to as: H26SPM,

Run name	Halo	M_{200} (10^{12} $M_\odot$)	M_{DM} (10^{11} $M_\odot$)	M_{star} (10^{11} $M_\odot$)	M_{gas} [hot, cold] (10^{10} $M_\odot$)	r_{200} (kpc)	$v_{circ,max}$ (km/s)	Δx (pc)	ϵ_{SF}	$\frac{\Delta x}{L_J}$	ϵ_{FB}	t_{supp} (Myr)
H26SPM*	26	1.1	9.4	1.6	3.7 [3.4, 0.3]	210	560	425	10^{-2}	8	0	0
H30SPM	30	1.1	9.1	1.5	5.1 [4.8, 0.3]	210	460	425	10^{-2}	8	0	0
H37SPM	37	0.8	6.6	1.1	3.4 [3.2, 0.2]	190	480	425	10^{-2}	8	0	0
H47SPM	47	0.6	5.1	0.9	2.1 [1.3, 0.8]	170	470	425	10^{-2}	8	0	0
H54SPM	54	0.5	4.0	0.7	2.0 [1.5, 0.5]	160	410	425	10^{-2}	8	0	0
D7H26SPM	26	1.2	9.3	1.8	4.5 [2.3, 2.2]	210	440	1700	10^{-2}	8	0	0
D8H26SPM	26	1.2	9.4	1.8	4.2 [2.5, 1.6]	210	470	850	10^{-2}	8	0	0
D9H26SPM*	26	1.1	9.4	1.6	3.7 [3.4, 0.3]	210	560	425	10^{-2}	8	0	0
D10H26SPM	26	1.3	11.	1.4	4.8 [4.3, 0.5]	220	610	212	10^{-2}	8	0	0
H26S	26	1.1	9.5	1.6	3.8 [3.4, 0.4]	210	540	425	10^{-2}	0	0	0
H26SPML	26	1.2	9.5	1.7	4.4 [2.1, 2.3]	210	640	425	10^{-3}	8	0	0
H26SPMR	26	1.2	9.4	1.9	3.1 [2.5, 0.5]	210	550	425**	10^{-2}	8	0	0
H26SPMF	26	1.3	11.	1.4	6.3 [5.6, 0.8]	220	580	425	10^{-2}	8	3E-6	0
H26SPMC	26	1.2	9.4	1.7	8.6 [7.5, 1.0]	210	400	425	10^{-2}	8	0	50
H26SPMFC	26	1.2	9.5	1.2	10. [6.5, 3.5]	210	410	425	10^{-2}	8	3E-6	50
H26SPMFCR	26	1.2	9.5	1.1	11. [8.2, 2.4]	210	370	425**	10^{-2}	8	3E-6	50

Table 3.1 We present the bulk properties and simulation parameters for each halo and simulation. The first five halos are the canonical simulations, each following a different halo in the same overall volume. Additional simulations were created using the initial conditions for halo H26SPM with modifications to the encoded baryonic and star formation physics. *Halos H26SPM and D9H26SPM are the same simulation. **Halos H26SPMR and H26SPMFCR have a fixed *physical* resolution throughout the simulation, while all other halos have a fixed *comoving* resolution.

H30SPM, H37SPM, H47SPM and H54SPM. Additionally, several different runs were performed on the initial conditions for halo 26 (H26SPM), which systematically varied simulation and physical parameters to investigate the effects of each on galactic evolution. The parameters toggled (and their respective simulations) include: (i) excluding minimum pressure support [H26S]; (ii) changing the maximum spatial resolution [D7H26SPM, D8H26SPM, D10H26SPM]; (iii) using a constant physical resolution instead of a constant comoving resolution [H26SPMR, H26SPMFCR]; (iv) including thermal feedback [H26SPMF, H26SPMFCR]; (v) lowering the star-formation efficiency [H26SPML]; and (vi) suppressing cooling in star forming regions [H26SPMC, H26SPMFC, H26SPMFCR]. The details of the various simulations and the resulting galaxies are shown in Table 3.1.

3.3 Results

In this section, we present the results of our galaxy formation simulations, first describing the five canonical runs, which all contain identical physical prescriptions but track different galactic halos. Then, we explore variations in resolution as well as the numerical parameters we use to describe the gas and star formation.

3.3.1 Mass History

The mass accretion history for a galaxy including the different modes of its accretion is thought to play a crucial role in determining its final dynamical state (e.g. Kereš et al.

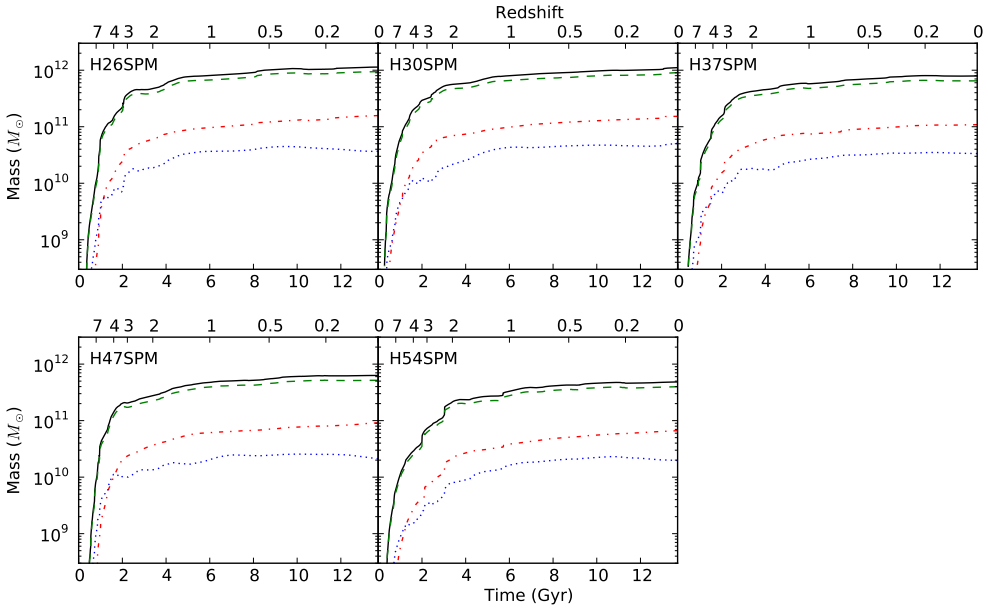

Figure 3.2 We plot the mass of our five canonical halos as a function of time. Different components of each halo are labeled differently: dark matter (green dashed line), stars (red dot-dashed line), gas (blue dotted line), and total mass (black solid line).

2005). In order to analyze the simulation in high time resolution, we record outputs from the simulation every 10 Myr. For each output we run the HOP algorithm (Eisenstein & Hut 1998) on the dark matter particles in order to identify halos. Given the particles in each of our five halos at $z = 0$, we identify and track these halos back to early times. Each halo is tracked backwards in time by identifying the local progenitor which shares the largest number of tightly-bound dark matter particles. The resulting mass-accretion history for each halo is shown in Figure 3.2. The halo masses are computed inside of r_{200}, the radius within which the mean density is 200 times the critical density of the universe at that redshift. At each time, we determine the center of the halo using an iterated center-

of-mass technique, which starts with the center of mass within r_{200}, and then successively recomputes the center of mass in smaller spherical volumes, decreasing the radius by 5% on each iteration and using the center of mass of the previous volume. This is necessary in order to make an accurate determination of the halo center (we found that simply choosing either the densest point or the center of mass within r_{200} did not produce a good estimate of the center in many cases). All masses calculated are masses contained within r_{200}, and are shown in Table 3.1.

In Figure 3.2, the mass histories of dark matter (green), gas (blue), stars (red) and total mass (black) are shown. The halos are arranged in decreasing z=0 total mass. All of the halos lack any major mergers over the last 10 Gyr, providing them with quiescent growth, in order to maximize the chance of producing disk systems. The lowest-mass system, halo 54, undergoes mergers at $z \sim 3.5$ and $z \sim 2.5$ causing discrete jumps in the mass of all its components at those times.

3.3.2 Star Formation History

Much of the angular momentum problem stems from an overproduction of stars and a buildup of the oversized stellar bulge, revealed by the galactic star formation rate history. If one can reduce star formation early in a halo's evolution it will moderate the amount of material in the inner region of the galaxy. While it is true that the bulk of the modeled bulge mass is stellar in nature, it remains unclear as to where these stars were created. Were they formed in clumps in the disk before plummeting to the center of the system, or

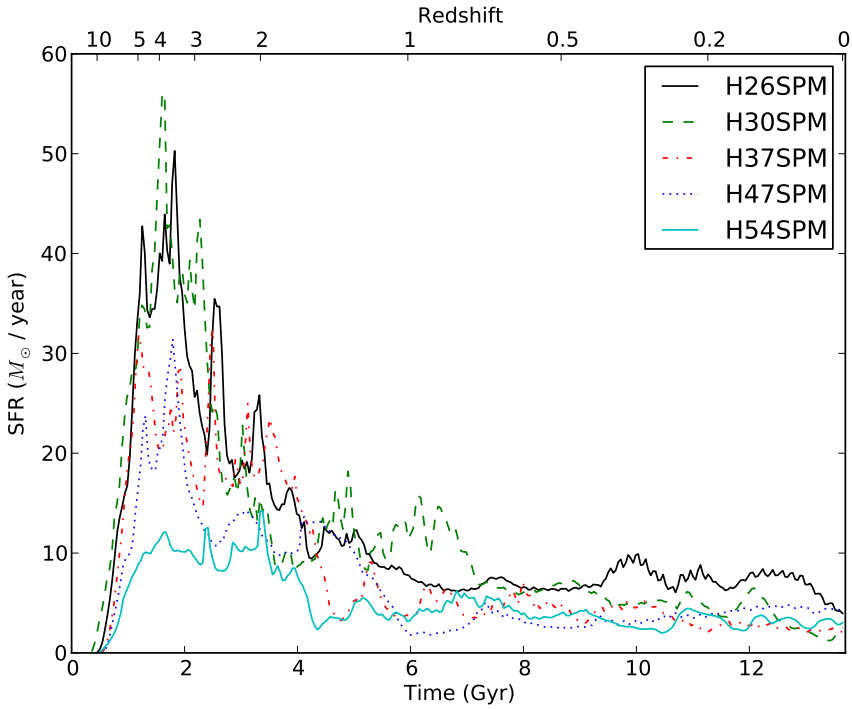

Figure 3.3 We plot the star formation histories for the five canonical halos. The bulk of the star formation occurs in the first 2-3 gigayears of evolution for each halo before settling into a low level of continuous star formation lasting to present.

did dense knots of gas spiral into the center of the galaxy where they ultimately collected and formed stars? This question will be examined more closely in Section 3.3.5.3.

In Figure 3.3, we show the star formation histories for our halos. These are computed by selecting the particles within r_{200} at $z = 0$ and using the stellar age of each particle to compute the implied star formation rate. This means that the rate shown includes star formation in all progenitor halos (not just the most massive progenitor shown in Figure 3.2). The star formation history shows an early burst, as dense, cold gas is rapidly

converted into stars, followed by a decline to a steady continuous level of SFR $\sim 3 - 10$ $M_\odot$ yr^{-1}. This low-level of star formation reflects both the decreasing gas supply and the decreasing amount of cold gas accretion (e.g. Kereš et al. 2005). Its value is slightly higher than but roughly consistent with observed levels in the Milky Way of ~ 1 $M_\odot$ yr^{-1} (Robitaille & Whitney 2010). Interestingly, our lowest mass halo, H54SPM, undergoes a much less pronounced early burst of star-formation, never forming more than 15 $M_\odot yr^{-1}$, which may be a reflection of a more extended merging period in the first few gigayears of evolution, as evidenced by its accretion history in Figure 3.2.

3.3.3 Disk Images

3.3.3.1 Gas

Each of the five halos produce a gaseous disk at $z = 0$. We use the analysis suite *yt* (Turk et al. 2011) in order to visualize the gaseous and stellar components of these halos, as shown in Figure 3.4. We determine the disk normal vector by computing the net angular momentum of all cold, dense gas (defined as $T < 2 \times 10^4$ K and $\rho > 10^{14}$ $M_\odot$ Mpc^{-3}) within 0.2 r_{200}. We then generate two images for each disk, one side-on and one face-on but each with the same scale of 25 kpc on a side. These two images are generated in slightly different ways. For faceon projections, we simply show the gas surface density in units of particles per square centimeter. For the edgeon image, we carry out a volumetric rendering to show transparent isodensity contours (in detail, we use the gas density to assign a transfer function at each point that consists of a set of narrow Gaussians, each

separated by a factor of 10 in density, and each given a different color). This allows us to show both the disk but also bring out structure in the halo.

These images display gas disks with radii of a few 10's of kpc, slightly smaller than, but approximately typical of present-day late-type systems. These cold disks are present from early times and are rotationally supported – we will examine their rotational velocities in more detail in the next section (Section 3.3.4). None of the disks display significant spiral structure at this particular timestep, although all of them seem to possess it at some point in their evolution. Interestingly, the most massive of the halos, halo 26 does not display as large of a disk as halo 47 despite having nearly two times as much total mass; however Table 3.1 reveals that the latter has significantly more cold gas, which is what is plotted here. The details of satellites and gas accretion tend to dominate the observed behavior of the disk at any point in time, and halo 26 recently accreted some small halos.

3.3.3.2 Stars

In the bottom part of Figure 3.4, we render the stellar particle distribution for each halo. Using the same scale and camera angles as we did for the gaseous components above, we generate stellar surface density plots. We represent each star particle as a gaussian with a sigma of 200 pc, then step through the volume depthwise, coadding a random sampling of 10% of all star particles. The color of a star is determined by its age, where we use a continuous color function of blue through red to represent the log of the age of the star as

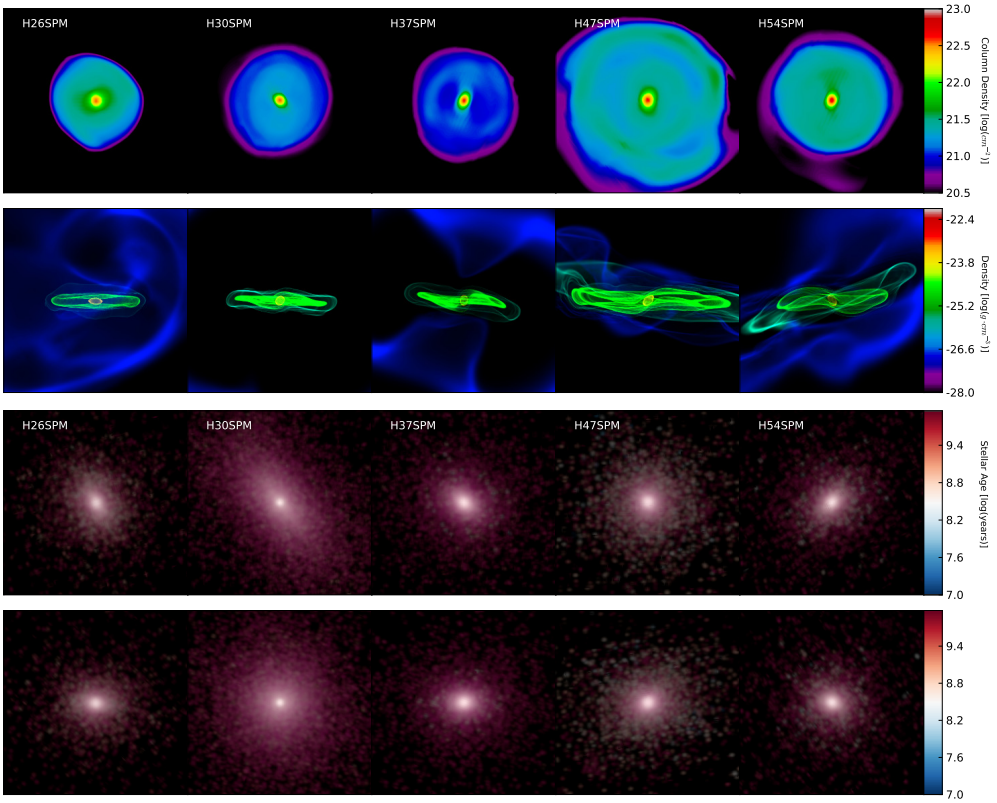

Figure 3.4 We render faceon and edgeon views for each of our five canonical halos at $z = 0$. The top two rows show the gas component of each galaxy, whereas the bottom two rows display each galaxy's stellar component. Each postage stamp has a width of 25 kpc. For the gas, edgeon views show a volume rendering with isocontours of the gas density, whereas faceon views are column density bricks. For the stars, we use the column density of stars in the faceon and edgeon views respectively. The color of a star particle is representative of its age where blue represents young stars and red represents old stars according to the color bars on the right.

shown in the colorbar. This represents stars just 10 Myrs old as pure blue and stars 13 Gyrs old as pure red. Just as in the gas surface density, the intensity of our stellar renderings is logarithmic with the bulk of the material residing in the inner core of each galaxy.

All of the stellar halos lack the disk structure we might expect from disk galaxies.

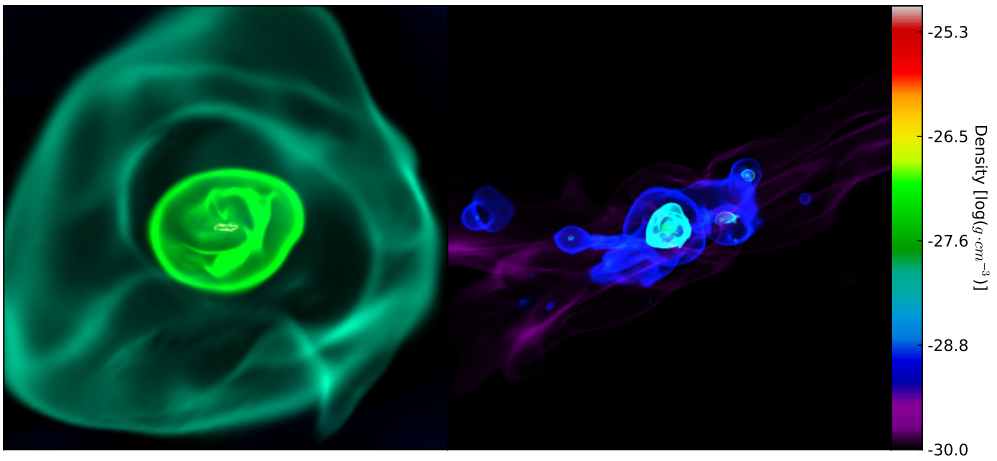

Figure 3.5 We display edgeon volume renderings of Halo 26 at $z = 0$. Like the edgeon views of the galaxies in Figure 3.4 these images show structure through the use of isodensity contours; however, unlike Figure 3.4, these images show the larger environment around the galaxy with 250 kpc and 2.5 Mpc on a side respectively.

Instead, they are completely dominated by a bulge component. This happens because star formation occurs most efficiency in the small, dense clumps that merge and lose angular momentum. Without feedback, the dense gas clumps efficiently deposit gas and stars in the center of the halo, leaving only a gas-poor disk which does not efficiently form stars. In addition, we do not see significant age gradients for most of the halos. We note that halo 47 does display the youngest, bluest overall stellar population (particularly relative to halo 30's aging stars), consistent with the star formation histories of Figure 3.3.

3.3.3.3 Larger gaseous environment

In addition, we generate large-scale volumetric renderings of halo 26 in order to show its extended halo and immediate environment. Figure 3.5 is generated in the same way as the

edgeon renderings described in Section 3.3.3.1. It shows H26SPM over a region with 250 kpc (approximately the virial radius) and 2.5 Mpc on a side respectively. These large-scale volume renderings (of the side-on disk) show that there is also a gaseous halo, extending out to at least $r_{200} \sim 200$ kpc. This hot halo gas is approximately spherical at $z = 0$, but contains a significant amount of substructure due to ongoing infall and asymmetric accretion. We do not see gas clumps cooling and condensing out of the smooth, hot gas halo (see also Binney et al. 2009; Joung et al. 2012a). The larger-scale image shows that the halo is embedded in a set of filaments, along which gas accrete and is typical of the other halos.

3.3.4 Rotation Curves

In this section, we focus on rotation curves at $z = 0$, as these are both directly comparable to observations and also immediately show the mass distribution of the halos. In Figure 3.6, we plot the the rotation curves for each of our five systems. The curves in each graph represent the equivalent circular velocity for each mass component: dark matter, gas and stars:

$$v_{\rm circ} = \sqrt{\frac{GM(\leq r)}{r}} \tag{3.3}$$

This figure shows that the rotation curves for each halo are highly peaked in their inner 5 kpc, primarily due to an overabundance of stars in their cores. This is clearly

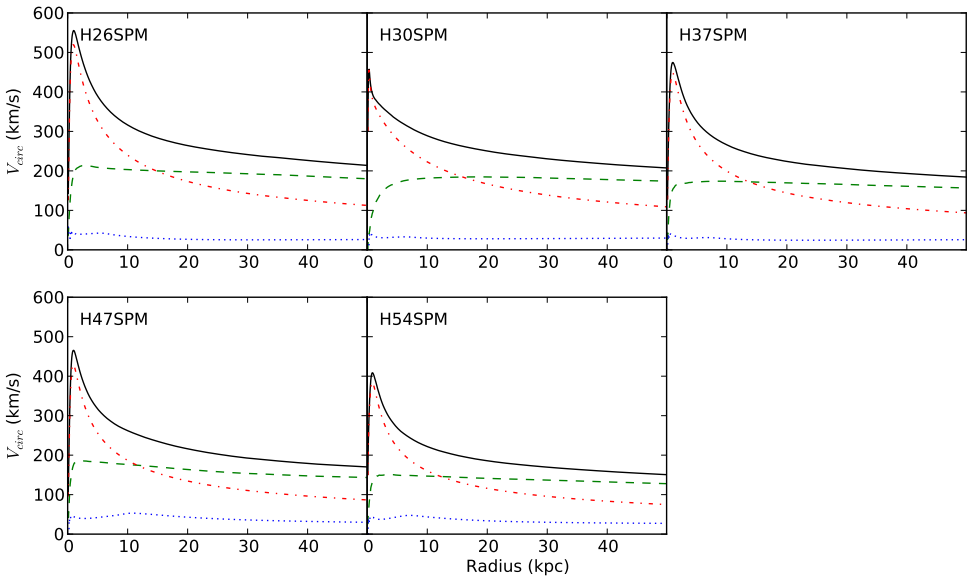

Figure 3.6 We plot the rotation curves for the five simulated halos. We represent the circular velocity of the gas due to the dark-matter component (green dashed line), the stellar component (red dot-dashed line), the gas component (blue dotted line), and the total of all halo components (black solid line). The stellar component seems to be driving the cusped rotation curve in the cores of each galaxy, whereas the gas and dark matter profiles appear consistent with observational expectations. In each system, the total-mass rotation curve is highly cuspy and unlike any observed galaxy, confirming that our simulations reproduce the angular momentum problem.

inconsistent with the nearly flat rotation curves observed in disk systems (e.g., Courteau 1997). The gas contributes negligibly, while the dark matter curve is steeper than expected for an NFW-profile because of contraction driven by the deep potential well of the stellar component, and remains nearly flat to the core, but is secondary to the stellar distribution.

3.3.5 Modifications to the Canonical Runs

In addition to our five canonical runs presented above, we conducted a series of additional simulations in which we systematically varied numerical parameters involving resolution, star formation, supernovae feedback, and gas physics. We used the initial conditions from halo H26SPM for each of these simulations, so that we could directly compare these results against one of our canonical models. In the following subsections, we present the results of these various modified runs and for each simulation, examine its star-formation history as well as its rotation curve. The plots for all of these runs are presented side-by-side in Figures 3.9, 3.10, 3.12 & 3.13. The bulk characteristics of these halos are presented in Table 3.1. We note that some runs of this halo have a slightly higher virial mass at $z = 0$ due to the presence of a fairly large satellite (10% of the main halo mass), which sits very close to the virial radius – slight shifts in its position in the various runs can lead to its inclusion in the total mass.

3.3.5.1 Resolution study

Since spatial resolution has been raised as an important issue affecting the angular momentum of the gas (e.g. Mayer 2004), we conducted three different modified runs, each with a factor of 2 change in the spatial resolution (done by systematically changing the maximum allowed refinement level). Runs D7H26SP, D8H26SP & D10H26SP follow the initial conditions of H26SPM (which itself has a resolution of 425 comoving pc) with 1700, 850, and 212 comoving parsecs spatial resolution respectively. There is no change to the

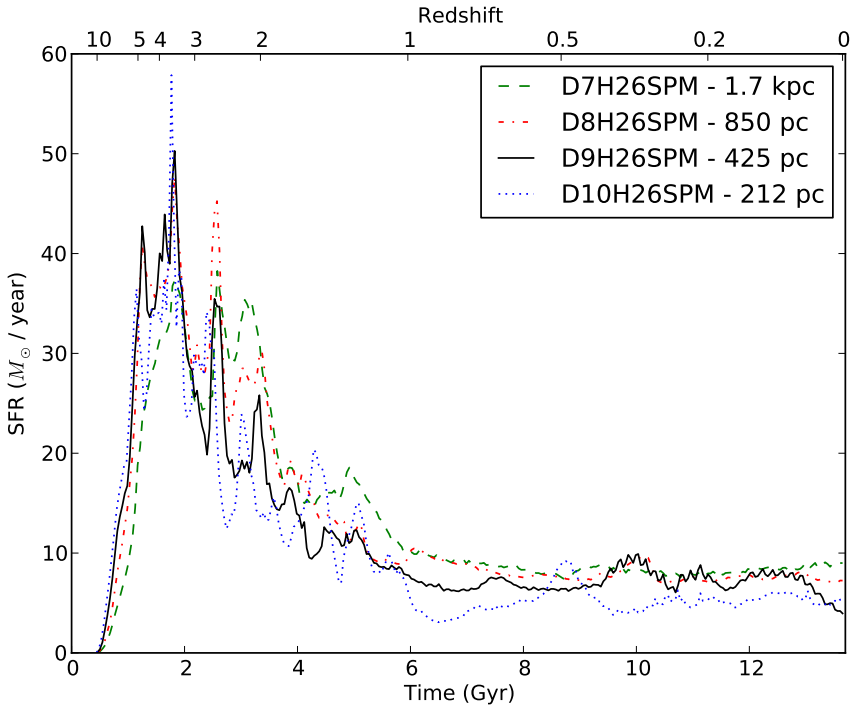

Figure 3.7 We plot the star formation histories for our resolution study. These runs all use identical initial conditions, those of canonical run H26SPM, which is also presented here as D9H26SPM. The only difference in each run is the maximum refinement level, that is, the maximum level of spatial resolution achieved which ranges from 212 comoving parsecs to 1.7 comoving kiloparsecs.

mass resolution in these modified runs.

Figures 3.7 and 3.8 show the star formation histories and $z = 0$ rotation curves for these runs. Up to the range we can probe, resolution alone doesn't appear to play a significant role in determining the overall mass distribution in a halo or the conversion of gas into stars. There is an indication from Figure 3.7 that lower resolution lowers the initial burst (because of the decreased central gas densities) and therefore shifts the peak of

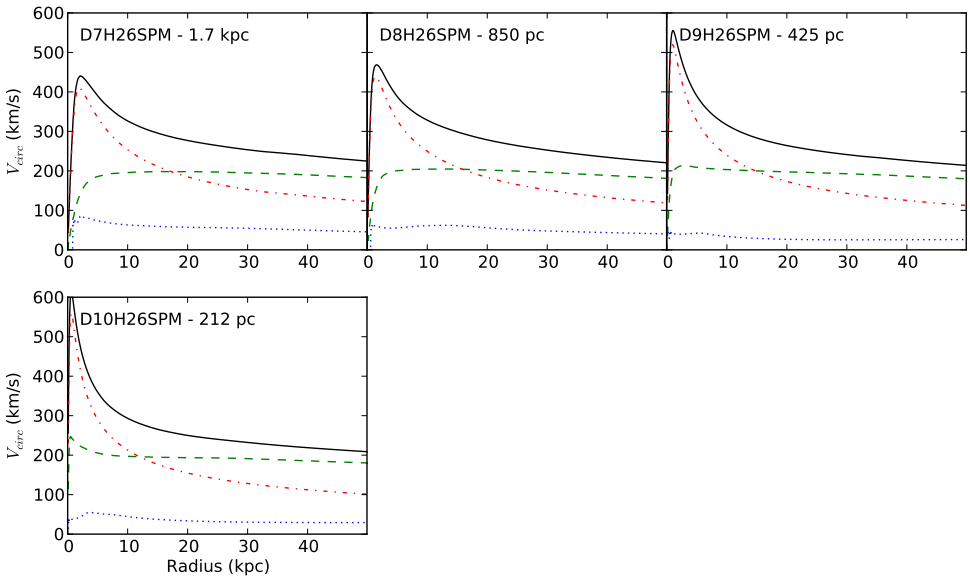

Figure 3.8 We plot the rotation curves for the resolution study. As in Figure 3.6, we represent the circular velocity of the gas due to the dark-matter component (green dashed line), the stellar component (red dot-dashed line), the gas component (blue dotted line), and the total of all halo components (black solid line). There is a clear trend in these runs that as one increases the resolution of the simulation, one allows a denser cusp of material in the core (primarily in the form of stars). In turn, the core circular velocity is increasingly driven upward.

star formation to later times; however, the effect is small. Increasing the spatial resolution actually allows material to condense to an even smaller volume in the cores of galaxies, which increases the inner cusp in the rotation curves of these systems, as seen in Figure 3.8. Thus, resolution by itself cannot solve the angular momentum problem. Perhaps increased resolution when coupled with a more sophisticated star-formation or feedback prescription will produce more realistic galaxies when one begins to resolve star-forming regions on parsec scales, since this is the scale over which star formation actually occurs.

3.3.5.2 Minimum Pressure Support

These simulations cannot resolve star-forming events on parsec scales, so our canonical runs employed an artificial minimum pressure to prevent Jeans fragmentation on smaller scales than we could resolve. By including this minimum pressure described in Machacek et al. (2001), we assured that the Jeans length is always refined by at least 8 cells, and therefore the Truelove criteria was met (see also Truelove et al. 1997; Robertson & Kravtsov 2008; Ceverino & Klypin 2009). It has been suggested that artificial disk fragmentation leads to large gas clumps that lose angular momentum via dynamical friction, and hence result in angular momentum loss. For modified run H26S, we turned off this minimum pressure support. The results of that run are presented in Figure 3.9 and 3.10. The results are nearly identical for both measures, indicating that the minimum pressure support did not have a significant effect on the star formation rate or the distribution of gas and stars in our simulated system.

3.3.5.3 Star Formation Efficiency

The star formation model we adopt in this work is based on the Kennicutt-Schmidt relation, but has an efficiency parameter that is not well-constrained. It has recently been suggested by Agertz et al. (2011) that star formation efficiency is a key parameter controlling the distribution of stars in the disk, and hence the rotation curve. Our efficiency is already fairly low; however in order to examine this suggestion in more detail, and also to probe the impact of decreasing the efficiency in general, we adopt $\epsilon_{\rm SF} = 10^{-3}$ in run

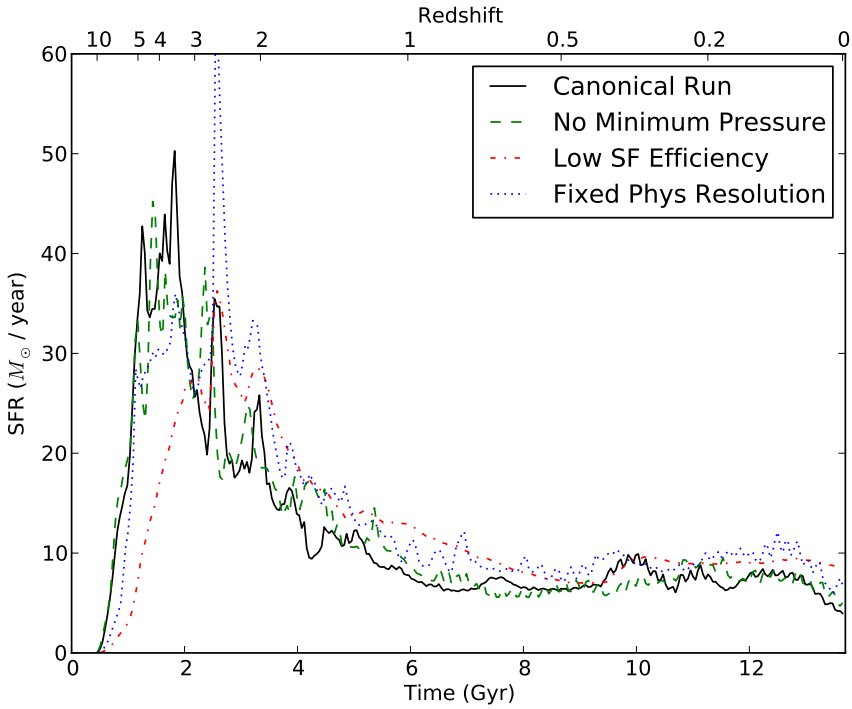

Figure 3.9 We plot the star formation histories for the first three of our modified runs. These runs all use identical initial conditions, those of canonical run H26SPM, which is also presented here. It appears that these three test runs have little effect on the bulk star formation history of the galaxy, although the simulation with the depressed star formation efficiency delays massive star formation for a gigayear or so.

H26SPML. The star formation history, in Figure 3.9, shows a significant delay to later time, as would be expected, although the net amount of stars produced is nearly identical. However, this does not translate into a flatter rotation curve, as can be seen in Figure 3.10.

To understand this result a little more, we look at the rotation curve for this run, at early times. Figure 3.11 shows the rotation curves at $z = 4.5$ of this low star formation efficiency run H26SPML compared against our canonical run H26SPM. During this epoch,

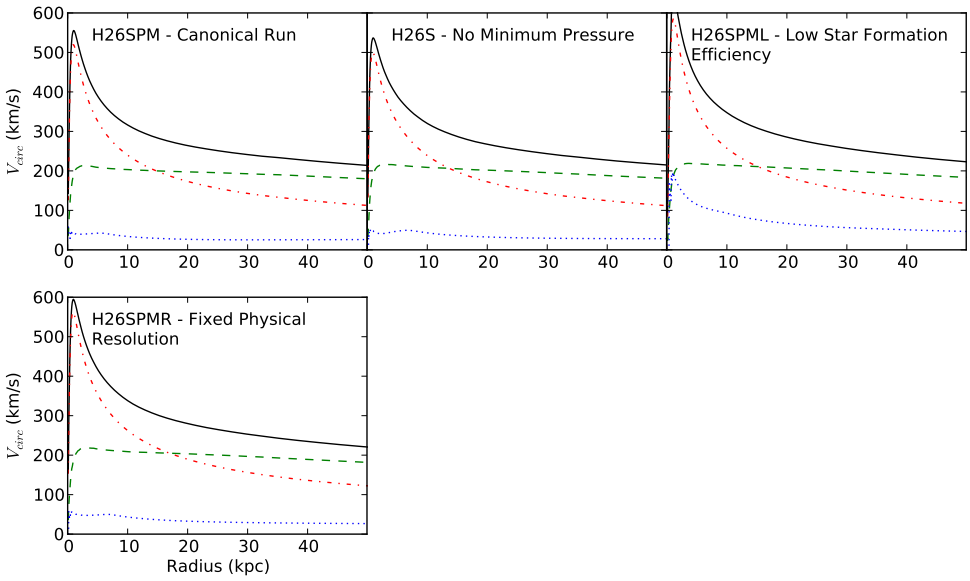

Figure 3.10 We plot the rotation curves for the first three of our modified runs. Just as in Figures 3.6 & 3.8, we represent the circular velocity of the gas due to the dark-matter component (green dashed line), the stellar component (red dot-dashed line), the gas component (blue dotted line) and the total of all halo components (black solid line). Like Figure 3.9, there is little improvement between these modified runs and their control run H26SPM. In fact, decreasing the star formation efficiency as in H26SPML actually drives more material into the core by $z = 0$.

much of the halo has formed, but star formation has not yet converted most of the gas into stars. This plot demonstrates that gas dominates the rotation curve at early times, and yet despite this, the gas clumps have already lost their angular momentum and formed a central cusp of compressed gas. In runs with higher star formation rates (e.g. the canonical run H26SPM), the gas clumps are partially converted to stars before accreting so the cusp is primarily composed of stars, but the net result is the same – the clumps lose angular momentum and form a centrally peaked rotation curve.

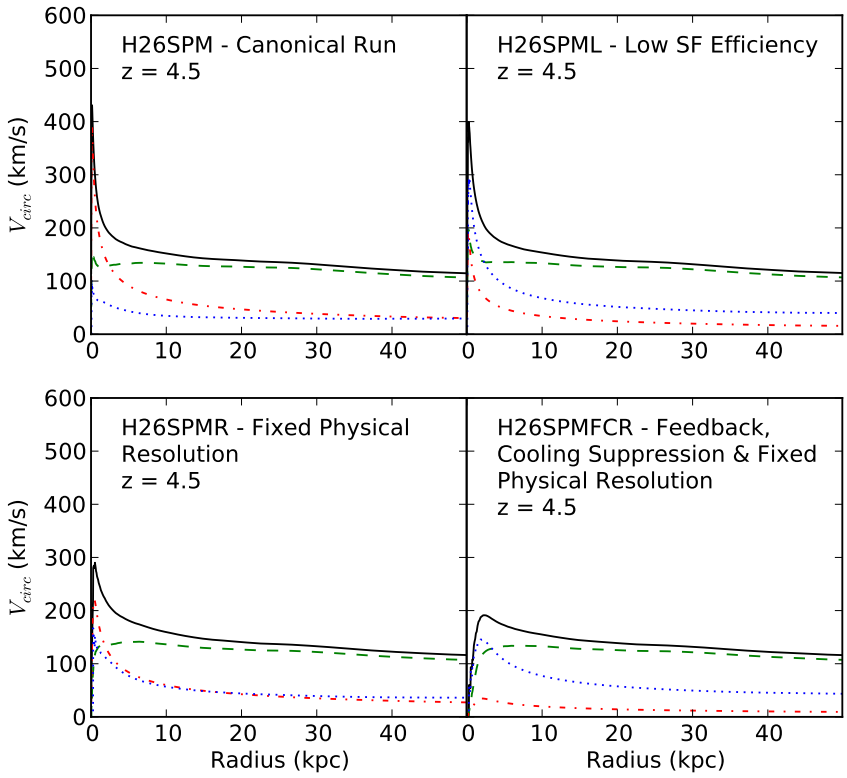

Figure 3.11 We plot the rotation curves for our canonical run (top-left) and the low star-formation run (top-right) at $z = 4.5$. Both exhibit a cusp of material in the core at this early epoch; however, in the canonical run the cusp is due to stars, whereas the low star-formation run's cusp is almost entirely compressed gas. This demonstrates that gas is first funneled into the core of a halo even in the absence of significant star formation. As seen in Figure 3.10 that in the end, the stellar component ends up dominating the cusp by $z = 0$, regardless of how things look at $z = 4.5$. In the bottom two panels we also show rotation curves at $z = 4.5$ for the fixed physical resolution run and the 'everything' run, for comparison.

Note that it is still possible that star formation efficiency coupled with feedback may be important – in particular these results are consistent with the idea that low star formation rates combined with sufficient feedback to puff up gas-dominated clumps would suppress angular momentum loss (indeed, the runs in Agertz et al. (2011) generally

include feedback and always include a stiff equation of state, $P \sim \rho^2$). However, we see that by itself, a low star formation efficiency does not change the distribution of matter in the simulated galaxies.

Also shown in Figure 3.11 are the rotation curves at $z = 4.5$ for the fixed physical resolution run (H26SPMR), as well as the simulation including feedback, cooling suppression, and fixed resolution (H26SPMFCR). As can be seen, the cusp already appears in the fixed resolution run (although not quite as high as in the canonical run), while gas dominates in the H26SPMFCR simulation. For that run, the additional pressure from feedback/cooling suppression has allowed the infalling clumps to be incorporated in the disk before losing a significant amount of angular momentum. For the rest of the modified physics runs (not shown), only the cooling suppression simulations differ significantly from the canonical run, and they are similar to the H26SPMFCR simulation, but with a somewhat more pronounced cusp.

3.3.5.4 Fixed Physical Resolution

Our canonical set of simulations used a fixed maximum refinement level, which translates to a best resolution achievable in comoving spatial units. Another common prescription is to fix the highest resolution as a fixed physical length scale (e.g. Agertz et al. 2011). We carry out a simulation of Halo 26 in which we vary the maximum allowed refinement level so as to keep as close as possible to a physical resolution of 425 pc. Note that because of the factor-of-two refinement in AMR, this implies discrete resolution changes at various

times during the simulation. The star formation history for this simulation is shown in Figure 3.9. The lower resolution at early times (for example, at $z = 5$, the spatial resolution is four times worse than in the canonical run) results in lower densities and hence a shift of the bulk of the star formation to later times (as in the lowered star formation efficiency run H26SPML). The rotation curve is shown in Figure 3.10, and again, there is not a significant change in the distribution of mass in the core.

3.3.5.5 Supernovae Feedback

As described in the methodology section (Section 3.2), we also performed one run with thermal feedback from Type II SN, using a moderate value of $\epsilon_{\rm FB}$, as given in Table 3.1. Figure 3.12 shows the star formation history for this run, and demonstrates that the inclusion of thermal feedback does have some effect. There is a reduction in the overall burst of star formation at high redshift relative to the canonical run with the same initial conditions. At late times the star formation rates are similar. However, in Figure 3.13 the end result for the rotation curve at $z = 0$ is the same as when feedback is not included. This occurs because the feedback is not strong enough to destroy the infalling clumps and prevent their loss of angular momentum. It is possible that stronger feedback will change this; however, we delay a systematic examination of feedback prescriptions for a later paper.

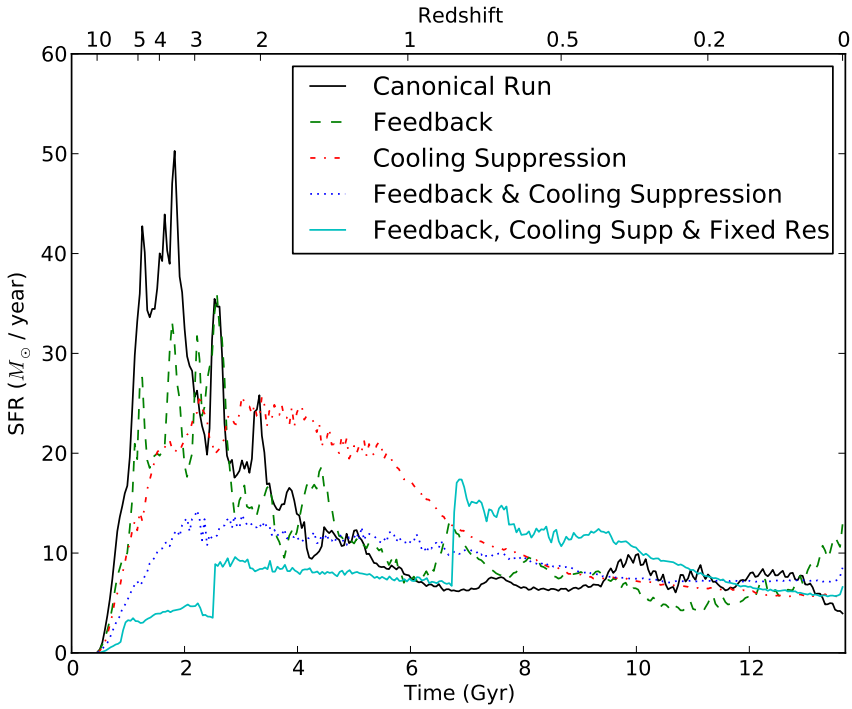

Figure 3.12 We plot the star formation histories for the last four of our modified runs. These runs all use identical initial conditions, those of canonical run H26SPM, which is also presented here. While feedback and cooling suppression each individually have some effect on lowering overall star formation, together their effects are amplified to keep star formation low throughout the galaxy's lifetime. The addition of a fixed physical resolution lowers that star formation even further. Unfortunately, because of the way we have implemented a "fixed" physical resolution in *Enzo*, there are discrete jumps in the star formation rate at times when we increase the comoving resolution (e.g. $t \sim 2.5$ & $t \sim 7$ gigayears).

3.3.5.6 Cooling Suppression Models

Many current studies employ a cooling suppression scheme in order to prevent cooling for a short period after a newly formed star is born, allowing the thermal feedback to efficiently operate. We performed three simulations which integrated cooling suppression:

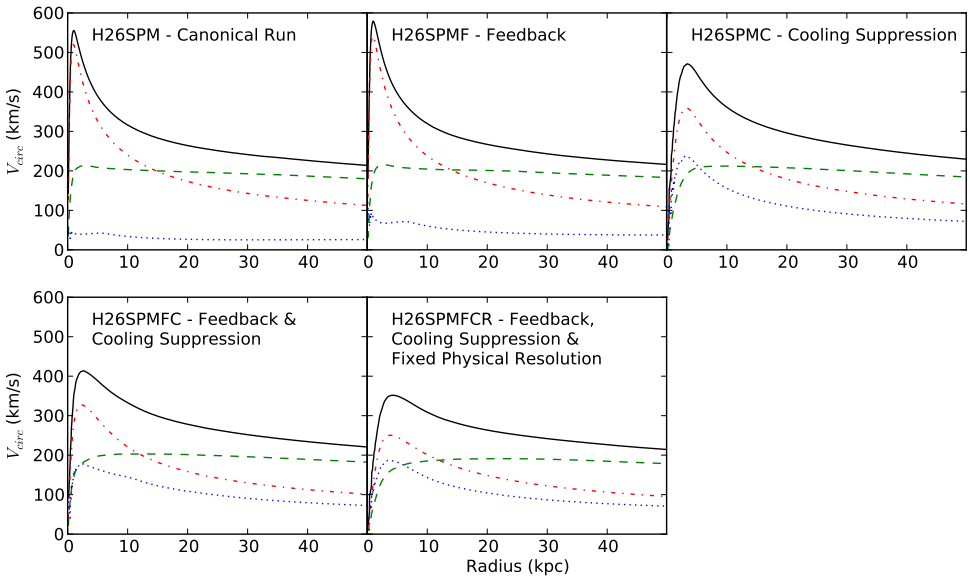

Figure 3.13 We plot the rotation curves for the last four of our modified runs. Just as in Figures 3.6 and 3.8, we represent the circular velocity of the gas due to the dark-matter component (green dashed line), the stellar component (red dot-dashed line), the gas component (blue dotted line) and the total of all halo components (black solid line). In agreement with Figure 3.12, we see that when working in concert the effects of feedback and cooling suppression are intensified to dampen the dense stellar cusp in the interior of our halo. The addition of the fixed physical resolution further decreases its rotation curve cusp and makes it look much more akin to something we might find from observations.

one with cooling suppression alone, one with cooling suppression in addition to thermal feedback, and one with cooling suppression, thermal feedback and fixed physical resolution. Figure 3.12 shows the star formation history of these three simulations. Interestingly, with just the lone addition of cooling suppression (red dot-dash line), our star formation is highly suppressed at early times, and extends to much later time, also being less bursty compared to the canonical run. The addition of thermal feedback on top of that further intensifies the effect of cooling suppression and the star formation rate becomes almost

constant throughout the simulation at about $10 - 12$ $M_{\odot}$/year. Finally, the synthesis of feedback, cooling suppression and fixed resolution results in an even lower overall star formation rate. In this last run, we see a step function in the SFR at two redshifts which correspond to the epochs when we change the allowed maximum refinement level, in order to preserve fixed resolution. These transitions show up particularly clearly in this case because the cooling suppression operates only in the local cell, and since most of the disk is refined to the maximum level, when this changes, it drastically lowers the efficacy of the cooling suppression, leading to more star formation. This demonstrates the sensitivity of our star formation results to the chosen parameters of the cooling suppression model.

In Figure 3.13, we show the resulting rotation velocity curves. We finally have an effect that has a significant impact on the rotation curves. Even cooling suppression by itself results in a peak rotation rate which is 100 km/s lower than the canonical run. Adding thermal feedback on top of this decreases the peak even more, to just over 400 km/s, and adding a fixed physical resolution brings it down to about 350 km/s. This occurs because the cooling suppression model (and the effective feedback that it permits) acts to decrease the density of infalling clumps. These clumps are then disrupted *before* they can lose a significant amount of angular momentum and so end up rotating at larger radius then they otherwise would have. We speculate that cooling suppression operates even in the absence of feedback because other forms of heating, such as shock heating and adiabatic compression, can play an important role. We note that the resulting rotation curves, while much more realistic, are still somewhat too strongly peaked at small radius,

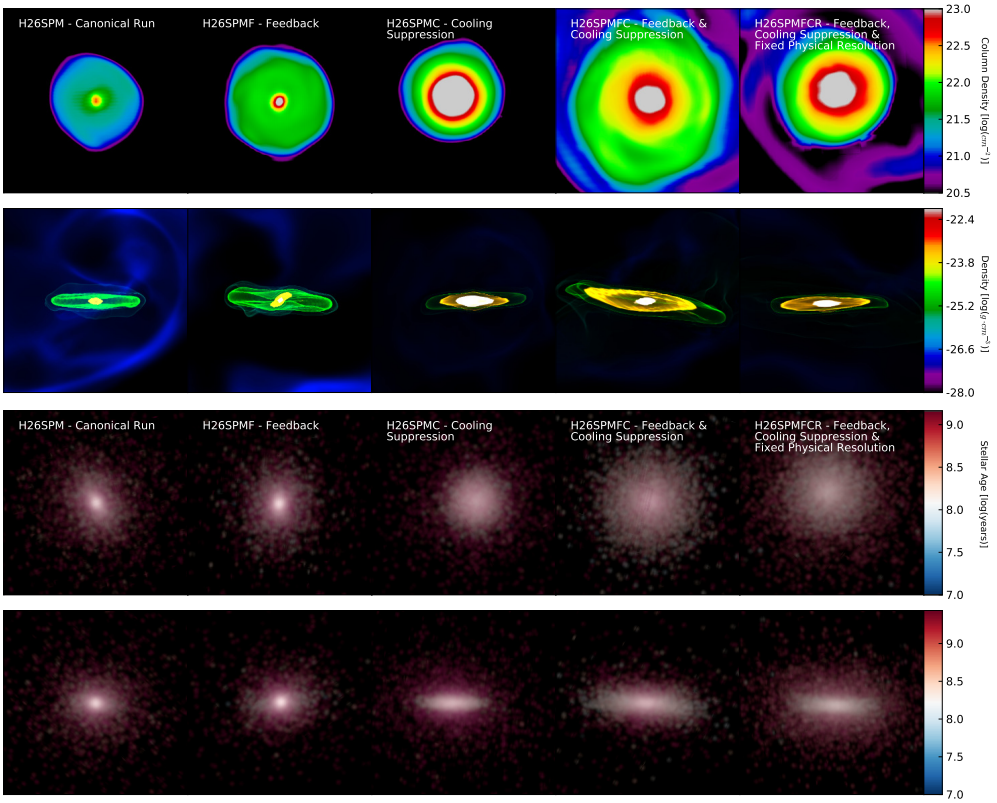

Figure 3.14 We render faceon and edgeon views for the last four of our modified halos at $z = 0$. The top two rows show each galaxy in gas, whereas the bottom two rows display each galaxy's stellar component. Each postage stamp has a width of 25 kpc. For the gas, edgeon views show a volume rendering with isocontours of the gas density, whereas faceon views are column density bricks. For the stars, we use the column density of stars in the faceon and edgeon views respectively. The color of a star particle is representative of its age where blue represents young stars and red represents old stars according to the color bars on the right. Notably, these modified runs produce young stellar disks and have a much higher density of gas in their cores than the canonical runs.

a characteristic shared (to a larger or smaller degree) by essentially all simulations that include cooling suppression.

Finally, Figure 3.14 displays visual renderings of the gaseous and stellar components of the halos, similar to those of the canonical runs in Figure 3.4. The simulations which

successfully reduced the buildup of material in the core of the galaxies have larger cold gas densities. There *is* still a lot of material in the core, but much of it is in compressed cold gas instead of stars. In the star projections, young stellar disks are present and embedded in halos of older stars. These disks have no significant bars, but are coaligned with their gas disks, as we might predict for systems of this type.

3.4 Discussion

These results demonstrate that it is challenging to generate disk systems with the correct mass distribution. Without effective feedback, the default outcome is for dense clumps to lose angular momentum and result in centrally-cusped rotation curves. In particular, the simulation with a very low star formation efficiency nicely demonstrates that this result is fundamentally a dynamical one, and does not depend on whether the clumps are primarily gas, or mostly stars. As long as they are concentrated, they will lose angular momentum and hence rotational support. Although this result is not new, and there is a long history of SPH simulations which found this result much earlier, we show it here clearly and systematically using a completely different numerical method (AMR). Therefore, the result is quite general.

In addition, we went on to systematically vary our numerical parameters and investigate a range of resolution and feedback methods. We confirmed that only cooling-suppression feedback models are capable of significantly changing the mass distribution

and hence the rotation curve.

Cooling suppression models do effectively enhance feedback, although it is unclear how physically meaningful this technique is (see Section 3.2.1.7), and a better approach might be to use high-resolution local models to generate subgrid models (see Yepes et al. 1997; Tasker & Bryan 2006; Ceverino & Klypin 2009; Joung et al. 2009b, for some attempts in this direction).

Another constraint is the baryon content of galactic halos: a variety of techniques have been used to infer that the baryon-to-dark-matter ratio in galaxies is much smaller than the cosmic mean (e.g. Moster et al. 2010; Behroozi et al. 2010), implying that a significant amount of mass has been ejected from galactic systems (or never accreted in the first place). For example, Milky-Way massed halos only appear to host 20% of their baryons, with the fraction decreasing rapidly for smaller-mass systems (Behroozi et al. 2010). We find that all of our simulations result in very high disk baryon fractions; even the run with feedback and cooling suppression has a baryon fraction of about 63% of the cosmic mean. This appears to be a general issue with cosmological galaxy simulations (see also the discussion in Avila-Reese et al. 2011).

We demonstrate that simulations must be run to $z \sim 0$ in order to gauge the efficacy of model parameters at minimizing the effects of the angular momentum problem. Many past studies (e.g. Ceverino & Klypin 2009; Joung et al. 2009a) have traded off simulation run time for increased resolution in their simulations. While there are some early indicators of the angular momentum problem like bursts of star formation and peaked rotation

curves at redshifts as early as $z = 5$, successful results during this epoch do not guarantee successful results at $z \sim 0$. We specifically demonstrate this in the case of the low star formation efficiency run H26SPML, which staved off early bursts of star formation but eventually succumbed to the same fate as runs with a normal star formation efficiency.

3.4.1 Comparison to Previous Work

In agreement with previous work using SPH, we find that unless we include efficient feedback, the resulting systems are dominated by a too-large spheroidal component, and the resulting rotation curve is peaked in the center (e.g. Navarro & Benz 1991; Navarro & Steinmetz 1997; Weil et al. 1998; Abadi et al. 2003; D'onghia et al. 2006; Zavala et al. 2008; Piontek & Steinmetz 2011).

A number of reasons have been suggested for this in the past, including purely numerical causes, such as angular momentum transfer between the disk and the hot halo (e.g. Okamoto et al. 2005), or between cold gas clumps and the hot halo (e.g. Thacker & Couchman 2000). The concern was particularly that SPH simulations might be susceptible to this issue because of smoothing between hot and cold phases. However, the fact that AMR simulations – which use a completely different numerical method to solve the fluid equations – find the same result, is an indication that these effects do not dominate.

Another suggested source of numerical angular momentum loss is in the form of gravitational instabilities which arise from inadequate resolution of the Jeans length in the disk (Truelove et al. 1997; Robertson et al. 2004). We tested this idea by running with

and without an additional numerical ("Jeans") pressure designed specifically to ensure that the Jeans length was adequately resolved, finding no difference in our results.

A third numerical reason is the lack of resolution (Governato et al. 2004; Kaufmann et al. 2007); however, we specifically test this over the computational range available for us, and find no significant difference from 200-1700 pc. This is in agreement with the SPH simulations of Piontek & Steinmetz (2011), who also varied their numerical resolution over a wide range, and found no difference.

There have been a number of recent cosmological AMR simulations which we can compare to. Colín et al. (2010) used the *ART* code (Kravtsov et al. 1997) to simulate a halo which is smaller by an order of magnitude (about 10^{11} $M_{\odot}$), finding peaked rotation curves (decreasing as the star formation density criterion was reduced). Although the halo masses are quite different, the essential result seems to be in agreement. Agertz et al. (2011) used the *RAMSES* AMR code (Teyssier 2002) to simulate a Milky-Way mass galaxy with similar resolution to that found here. They argued that a low star formation efficiency by itself was enough to produce nearly flat rotation curves (and that feedback was only efficient if extreme amounts of energy were injected). We have not been able to confirm the first suggestion – using a range of low efficiencies for star formation, we find that clumps lose angular momentum at high-redshift, and generate steep rotation curves, whether they are in gas or stellar form. The efficiency controls when the gas is converted to stars, but has little impact on the distribution of the material. In this, we are in agreement with previous SPH work (e.g. Weil et al. 1998; D'onghia et al. 2006; Piontek

& Steinmetz 2011).

Finally, we have found that the only way to significantly decrease the peak of the rotation curve was to introduce a sub-grid model which enhanced the efficiency of stellar feedback. We briefly explored the cooling suppression model and found this to be effective. This agrees with a substantial number of SPH simulations which adopt this mechanism (Gerritsen 1997; Thacker & Couchman 2000; Sommer-Larsen et al. 2003; Stinson et al. 2006; Governato et al. 2007; Colín et al. 2010; Agertz et al. 2011; Piontek & Steinmetz 2011; Guedes et al. 2011). In addition, Avila-Reese et al. (2011) used AMR simulations (with the *ART* code) and also found cooling suppression to be effective in obtaining approximately flat rotation curves. Ceverino & Klypin (2009) also used the *ART* code, but with a different sub-grid model, arguing that a model in which stars are born with significant velocities relative to the nascent gas will feed energy into low-density regions, producing efficient feedback and flat rotation curves, although the simulation is only run to $z \sim 3$.

3.5 Summary

In this work, we have carried out adaptive mesh refinement simulations of a sample of Milky-Way sized halos in order to better understand how numerical methodology impacts the content and structure predicted by such models. We selected a series of five halos ranging in mass from $0.5 - 1.1 \times 10^{12}$ $M_{\odot}$, each picked to have a quiescent mass accretion history over the last 10 Gyr, with no major mergers, in order to focus on halos that have

a high chance of hosting disk-like galaxies. We simulated the halos with high mass and spatial resolution, including gas, radiative cooling, star formation, and in some runs, feedback. We also took one halo and resimulated it a number of times, modifying the simulation parameters for each run. Our primary results are presented below.

- We find that, without any sort of feedback, all five halos produce rotationally supported gas disks. They are all, however, dominated by a massive and concentrated stellar spheroid, resulting in rotation curves that peaks at about 500 km/s in the central few kpc. Therefore, we confirm previous SPH simulation work that also found that dissipation allowed dense clumps to lose angular momentum and produce halos which are too cuspy.

- For one halo, we vary the comoving spatial resolution from 1700 to 212 pc and find the resulting cuspy halo to be completely robust against resolution change. Using constant physical resolution, rather than constant comoving resolution also has no effect in reducing the central concentration of mass for halos in the absence of other modified physics.

- Adding a model of local thermal feedback from Type II SN also has little effect, somewhat suppressing star formation, but with no impact on the $z = 0$ rotation curve. The use of an artificial pressure to ensure that the Jeans length was always resolved also resulted in little change to the star formation history or to the mass distribution.

- We find that reducing the star formation efficiency (by itself) has little impact on preventing the build up of a dense cusp of material in the core of the galaxy. In particular, we show that at $z \sim 4$, the cusp is already in place; however for the run with low star formation efficiency, it is primarily in the form of dense gas rather than stars, as in our other simulations. This demonstrates that the angular momentum problem is a dynamical one. Modifying the star formation prescription to slow the rate of stellar production is unlikely to solve the angular momentum problem in the absence of other efficient feedback (e.g. Agertz et al. 2011; Guedes et al. 2011).

- The only modification that did have a significant impact on the overall mass distribution in our halos was to suppress cooling in the vicinity of young stars (more precisely we suppress cooling in the local cell for 50 Myr after a star formation event). This, combined with the addition of thermal feedback and resolving to a fixed physical resolution, led to a large decrease in the peak of the rotation curve (although still larger than observations indicate in the inner few kpc).

To date, the most effective means for staving off the angular momentum problem is to employ cooling suppression in the vicinity of star formation events in order to intensify the effects of feedback. Most of the recent researchers in this field employ it by default, while at the same time they investigate the effects of other simulation parameters on preventing the angular momentum problem (e.g. Agertz et al. 2011; Guedes et al. 2011). Our results indicate that the other parameters are secondary to cooling suppression, and they do not seem to work effectively in the absence of it. We realize that the use of cooling

suppression is mostly motivated by its effectiveness in this regard, but we continue to search for other feedback parameterizations which are more physically derived.

There are several potential alternatives to cooling suppression as a form of feedback in cosmological hydrodynamics solutions of galaxy formation. Radiative feedback shows some promising results (Kim et al. 2011), although scaling its effects to a large number of particle sources is a current computational challenge. Another option is that the cosmic rays produced by supernovae could be used as a means for transporting energy and momentum to the surrounding medium as an additional fluid in a simulation (e.g., Miniati 2001; Jubelgas et al. 2008). Including magnetic fields may provide an additional support mechanism to stave off collapse and overcooling, although we lack significant constraints on the magnetic fields of observed galaxies. Alternatively, modifying the gas in dense regions to have a stiff equation of state (Agertz et al. 2011) may help "puff" up early collapsing pockets of gas. Perhaps there is even a redshift-dependent feedback prescription, similar to those used by Sommer-Larsen et al. (2003) and Okamoto et al. (2005), where it was assumed that the IMF was top-heavy in the distant past resulting in more supernovae and a higher feedback efficiency in that epoch. We are currently investigating some of these options to be presented in a forthcoming paper.

Chapter 4

A Comparison of Simulated and Observed Circumgalactic Media for Constraining Subgrid Models

4.1 Introduction

Following the flow of baryons in and out of galaxies is necessary to understand the star formation rate, stellar content and overall structure of galaxies of all sizes. Significant progress has been made in our theoretical comprehension of the inflow side of this equation, with the improved understanding of cold and hot-mode accretion (e.g. Binney 1977; Birnboim & Dekel 2003; Kereš et al. 2005; Ocvirk et al. 2008; Kereš et al. 2009). In this picture, halos with masses below about 3×10^{11} $M_{\odot}$ accrete most of their gas in the form

of ("cold") $\sim 10^4$ K filamentary flows, while those that are more massive develop an accretion shock and tend to accrete gas in a smooth, hot form. This implies that Milky Way progenitor systems, at high redshift, received baryons in the cold mode, but that present-day large spirals do not.

Although the existence of these cold streams is robustly predicted by ΛCDM simulations, observational confirmation remains challenging. Kimm et al. (2011) found that the predicted covering fraction of dense gas $N_{\rm HI} > 10^{20}$ cm^{-2} at $z \sim 2-3$ that would give rise to observable C II absorption in background quasi-stellar objects (QSOs) is only about 5% for Milky-Way-mass halos. Despite this low covering fraction, Fumagalli et al. (2011) found that the cold streams at high-redshift contribute significantly to the observed population of Lyman-limit systems (see also Ribaudo et al. 2011), and that the transition from low to high mass should be accompanied by a disappearance of Lyman-limit absorbers at moderate to high redshift (Stewart et al. 2011). At low redshift, the inflow picture is even less clear, as accretion is predicted to transition out of the cold-mode. Detailed predictions are currently lacking, although see Joung et al. (2012b); Fernández et al. (2012) for a detailed comparison of H I gas properties between simulations and observations at $z \sim 0$.

Galactic winds are the other side of the baryon flow equation (Davé et al. 2012), and evidence from the current stellar content of galaxies indicates that these must carry a substantial amount of mass and metals (Mandelbaum et al. 2006; Oppenheimer & Davé 2008; Behroozi et al. 2010). Numerical simulations have struggled to reproduce these large mass outflows (Katz 1992; Abadi et al. 2003; Hummels & Bryan 2012), but the

addition of subgrid models can produce high mass winds (e.g. Springel & Hernquist 2003a; Stinson et al. 2006; Oppenheimer & Davé 2006; Scannapieco et al. 2006; Governato et al. 2012). However, a detailed understanding of what drives winds is still lacking, with even the basic physical mechanism (e.g. ram pressure from supernova bubbles or radiation pressure) in dispute (e.g. Murray et al. 2011; Sharma & Nath 2012; Hopkins et al. 2012).

One important way to probe both the inflow and outflow of baryons, and an important test for any cosmological galaxy formation model, is the observed correlation between absorption lines in background quasars and the projected distance (and velocity offset) to the associated galaxy. A substantial body of such data has been built up at high redshift (Simcoe et al. 2004; Steidel et al. 2010). These data are beginning to be compared against simulations at high redshift (e.g. Altay et al. 2011; Tescari et al. 2011; Fumagalli et al. 2011; Barnes et al. 2011; Goerdt et al. 2012; Shen et al. 2012). At low redshift, where we focus, there exists a large body of data combining galaxy imaging and QSO spectroscopy (Chen et al. 2001; Chen & Tinker 2008; Thom & Chen 2008; Yao et al. 2008; Chen et al. 2010; Prochaska et al. 2011; Thom et al. 2011; Tumlinson et al. 2011), but only a few predictions from numerical simulations (Stinson et al. 2011; Fernández et al. 2012).

In this paper, we use high-resolution cosmological simulations of a galaxy run to $z = 0$ to make predictions for a wide variety of absorption lines as a function of impact parameter and compare them against observational samples. We include tracers of cold gas ($T \sim 10^4 - 5 \times 10^4$ K) such as H I, Mg II, Si II, Si III, and Si IV, probes of warm-hot

gas (C IV, N V, and O VI, for $T \sim 10^5 - 3 \times 10^5$ K), as well as gas at the virial temperature of the halo (O VII; $T \sim 10^6$ K). We will show that these comparisons make for a sensitive (and challenging!) probe of the phase of the inflowing and outflowing gas. These runs are conducted on identical initial conditions, systematically varying the stellar feedback prescription to demonstrate their effect on the column density distribution of absorbers in the galactic halo. In Section 4.2, we describe the details of our simulations and our post-processing steps on both the numerical and observational datasets. Section 4.3 presents the results of these simulations and the consequent analyses. We discuss these results and their implications for the nature of outflows, feedback and for use as a probe of different feedback models in Section 4.4. Lastly, we summarize these conclusions in Section 4.5 and discuss future directions for subgrid models to match these added observational constraints.

4.2 Methodology

4.2.1 Simulation Code

The numerical methodology and initial conditions used in this paper are very similar to those employed in Hummels & Bryan (2012, hereafter Paper I), and so we only briefly describe them here. We perform all of our simulations using the adaptive mesh refinement (AMR) hydrodynamics code, Enzo[1] (Bryan & Norman 1997; O'Shea et al. 2004). Enzo

[1]http://enzo-project.org

uses AMR to dynamically restructure the grid over the course of the simulation to achieve higher resolution in regions where predetermined refinement criteria are met. It treats dark matter and stellar populations as collisionless particles modeled with an N-body particle-mesh gravity solver (Efstathiou et al. 1985). The equations of gas dynamics are solved using the ZEUS hydrodynamics engine (Stone & Norman 1992).

A minimum pressure floor is included in these simulations in order to keep the Jeans length resolved across all grid cells (Truelove et al. 1997). We follow the implementation of Machacek et al. (2001), wherein the ratio of the Jeans length to cell size is maintained at $J = L_{\rm J}/\Delta x \geq 8$ by the addition of an artificial pressure.

We model the process of star formation following a prescription similar to that in Cen & Ostriker (1992). A cell will spawn a star particle if the following criteria are met: (1) the cell's gas overdensity exceeds $\delta_{\rm SF}$, (2) the local velocity divergence is negative, and (3) the cooling time is less than the dynamical time for the gas cell. When these conditions exist, a star particle with mass m_* is created according to the relation:

$$m_* = \epsilon_{\rm SF} \frac{\Delta t}{t_{\rm dyn}} \rho_{\rm gas} \Delta x^3 \tag{4.1}$$

where $\epsilon_{\rm SF}$ is the star formation efficiency per dynamical time; Δt is the local timestep; $t_{\rm dyn}$ is the local dynamical time, defined as $t_{\rm dyn} = (3\pi/32G\rho)^{1/2}$; $\rho_{\rm gas}$ is the gas density; and Δx is the grid cell width. To prevent a large number of low mass particles, we adopt a minimum star particle mass of $m_{*,\rm min}$. If $m_* < m_{*,\rm min}$, the star is formed with a probability

of $m_*/m_{*,\mathrm{min}}$, and its resulting mass is the minimum of $\mathrm{m}_{*,\mathrm{min}}$ and 90% of the mass in the gas grid element (e.g. Tasker & Bryan 2006). For these simulations, we set $\mathrm{m}_{*,\mathrm{min}} = 10^5\ \mathrm{M}_\odot$. We find that $\epsilon_{\mathrm{SF}} = 10^{-2}$ and $\delta_{\mathrm{SF}} = 10^3$ approximately reproduces the Kennicutt-Schmidt Law (see Paper I). This value of δ_{SF} is equivalent to 0.1 cm^{-3} at $z = 3$.

4.2.1.1 Supernova Feedback

Our implementation of stellar feedback includes a prescription for prompt injection of thermal energy from type II supernovae to the surrounding medium. Individual star formation events produce star particles of $m_* \sim 10^{4-5}\ \mathrm{M}_\odot$, each representing an entire stellar population. The high-mass stars within this stellar population will become supernovae and send energy, mass and metals into the local medium. Because a stellar population is rarely created instantaneously in nature, we "activate" a small component of the star particle's mass each timestep to represent the stars that were effectively created during that period. We parameterize the extended timescale over which this occurs with:

$$m_{\mathrm{FB}} = m_* \frac{\Delta t}{\tau} \frac{t - t_0}{\tau} e^{-(t-t_0)/\tau} \tag{4.2}$$

where m_{FB} is the mass of newly activated stars for a given timestep, t_0 is the star particle creation time, and $\tau = \max(t_{\mathrm{dyn}}, 10\ \mathrm{Myr})$. In order to prevent unrealistically short dynamical times, we apply a floor of 10 Myr to τ. We normalize our activated star mass so that the total mass released over the lifetime of the star particle is one-fourth the star's original

mass (i.e. $\sum m_{\rm FB} = 0.25 m_*$).

The returned energy $e_{\rm FB}$ and returned metals $m_{\rm metals}$ due to stellar feedback are directly proportional to the mass of these new stars:

$$e_{\rm FB} = m_{\rm FB} c^2 \epsilon_{\rm FB} \tag{4.3}$$

$$m_{\rm metals} = m_{\rm FB}\, y \tag{4.4}$$

where c is the speed of light; $\epsilon_{\rm FB}$ is a feedback efficiency parameter; and y is the fraction of metals returned relative to the activated stellar mass. Thus at any timestep Δt after the creation of a star particle, it returns $m_{\rm FB}$ mass, $e_{\rm FB}$ energy, and $m_{\rm metals}$ metals to the cell containing the star particle. This thermal energy should heat the surrounding gas, increasing the Jeans mass, and prevent further immediate star formation.

In the present investigation, we use a value of $y = 0.02$ consistent with previous work (Smith et al. 2011), and we include metal and mass feedback in all of our simulations. In four of our simulations, we vary the thermal energy efficiency value: $\epsilon_{\rm FB} = 3 \times 10^{-6}$ (low), 1×10^{-5} (medium), and 3×10^{-5} (high). These different energy efficiencies correspond to one 10^{51} erg SN for every 180, 60, and 18 $M_\odot$ stars formed respectively, effectively sampling the range of feedback efficiencies used by other studies on this topic (see discussions in Section 4.4.3 and Paper I).

4.2.1.2 Cooling Suppression

One of the shortcomings of the thermal feedback model is that thermal energy dumped into an underresolved volume will fail to bring that gas up to a high enough temperatures to prevent it from immediately radiating away this energy. To address this, many researchers have adopted an *ad hoc* solution, to turn off radiative cooling in the gas for a period after a star formation event (Gerritsen 1997; Thacker & Couchman 2000; Sommer-Larsen et al. 2003; Stinson et al. 2006; Governato et al. 2007) This model is justified as an application of the Sedov-Taylor solution for a supernova blast wave applied at scales below the spatial resolution of the simulation, although these arguments are not entirely physically consistent (see Paper I). However, cooling suppression seems to effectively quench star formation and slow the overcooling problem particularly when combined with an additional feedback prescription (Colín et al. 2010; Agertz et al. 2011; Guedes et al. 2011; Piontek & Steinmetz 2011, Paper I).

Our implementation of cooling suppression is relatively simple. After a star-formation event, we shut off cooling for a period $t_{\rm supp}$ in the gas cell containing our newly-formed star particle. We allow this effect to track the star particle until $t_{\rm supp}$ has passed, just as our thermal prescription for feedback tracks the star particle for a time after its creation. In this study, we set $t_{\rm supp} = 50$ Myr, a value consistent with previous studies (Stinson et al. 2006; Colín et al. 2010, Paper I).

4.2.1.3 Radiative Heating and Cooling

In modeling the thermal state of the gas in these simulations, Enzo accounts for photoelectric heating from a metagalactic ultraviolet background and radiative cooling from metals present in the gas. Because directly modeling the fractional components in each ionic state of every element is too computationally costly for simulations, we adopt a different tactic. We employ the method pioneered by Smith et al. (2008) and subsequently used by other groups in a variety of contexts (Booth & Schaye 2009), wherein all of the phases of H and He are followed individually throughout the simulation, but for more massive elements, a multi-dimensional lookup table is consulted to determine their aggregate contribution to heating and cooling using a single all-encompassing metallicity field. All phases of H and He (H I, H II, He I, He II, He III, and e^-) are calculated by a non-equilibrium, primordial chemistry network, which uses a backwards-differencing formula to account for collisional ionization, recombination, brehmstrallung, Compton cooling, photoexcitation and photoionization rates in these species (Abel et al. 1997; Anninos et al. 1997). The metal heating/cooling lookup table is compiled using version 07.02.01 of the spectral synthesis code, Cloudy[2] (Ferland et al. 1998). The table consists of five dimensions, one each for density, temperature, metallicity, electron fraction, and redshift. The redshift index is included to determine the intensity and spectrum of the metagalactic ultraviolet background at that time, which we activate at $z = 7$ and henceforth follow the prescription set in Haardt & Madau (1996) as updated in this version of Cloudy. Prior to reionization

[2]http://www.nublado.org

at $z = 7$, metal heating and cooling are only due to collisional ionization.

4.2.2 Initial Conditions

In this study, we adopt the WMAP 5-year results (Komatsu et al. 2009), specifically adopting $\Omega_0 = 0.258$, $\Omega_\Lambda = 0.742$, $\Omega_{\rm baryon} = 0.044$, $\sigma_8 = 0.796$, $H_{\rm o} = 71.9$ km s^{-1}. Using the *inits* program included in the Enzo distribution, we generate a 128^3 particle grid over our $20h^{-1}$ Mpc volume and carry out a low-resolution simulation to identify a dark matter halo with mass $M_{200} \sim 10^{12}$ M$_\odot$ (see Paper I for more details – we use halo 26 of that paper). Once the halo is identified, we select all particles within 2 virial radii at $z = 0$, and trace them back to their original positions at $z = 99$. The region containing all of these particles at $z = 99$ is further refined with two nested refinement regions, each with 2-times better spatial resolution and 8-times lower mass dark matter particles, resulting in $M_{\rm DM} = 4.9 \times 10^6$ M$_\odot$ dark matter particles in the vicinity of our target halo. Additional (adaptive) refinement occurs in a cell when its gas mass (or dark matter mass) exceeds 4×10^6 M$_\odot$ (or 1.9×10^7 M$_\odot$ for the dark matter), resulting in a further increase of the spatial resolution by a factor of two, up to nine levels of refinement, achieving a maximum resolution of 425 co-moving pc. The relaxation timescale for our particles of different mass, stellar and dark matter, is a on the order of a Hubble time, assuring that relaxation effects are avoided.

4.2.3 Description of the Simulations

We performed five different simulations using identical initial conditions and simulation parameters as described above. The way in which these simulations differed was simply in their supernova feedback and cooling suppression parameters. These runs include: a control run with no energetic feedback or cooling suppression ["Vanilla"]; three runs with no cooling suppression but varying feedback efficiencies of $\epsilon_{\rm FB} = 3 \times 10^{-6}, 1 \times 10^{-6}$, and 3×10^{-5} ["Low Feedback", "Medium Feedback", and "High Feedback", respectively]; and a run including cooling suppression and feedback with $\epsilon_{\rm FB} = 1 \times 10^{-6}$ ["Medium Feedback and Cooling Suppression"]. These runs and some of the properties of their halos at $z = 0$ are listed in Table 4.1. The simulations produced in this study have identical initial conditions and physical parameters to those associated with run H26SPM (i.e. the canonical run) in Paper I save for the treatment of feedback and the presence of metals. See that work for more detail. It is worth mentioning that we do not tune our models in any way to produce the results herein, all parameters being either deduced from theoretical calculations or values in the literature. Unless otherwise noted, all units herein are in comoving units, where h^{-1} is excluded.

4.2.4 Analysis Code

We applied a series of post-processing analyses to each simulation so as to render synthetic observations comparable to observational datasets for these simulated galaxies. Each simulation outputs a snapshot to disk every 10 Myr, representing the state of the gas,

stars and dark matter at that time. Following the method described in (Smith et al. 2011), we processed each datafile to compute the spatial distribution of various different ionic species. The density of the *i*th ionic species of a given element, X, is given by

$$n_{X_i} = n_{\rm H} \times (n_X/n_{\rm H}) \times (n_{X_i}/n_X) \tag{4.5}$$

where $n_{\rm H}$ is the density of hydrogen, $(n_X/n_{\rm H})$ is the elemental abundance of X and (n_{X_i}/n_X) is the ion fraction in the *i*th energized state. We internally track the neutral and excited states of hydrogen and helium internally throughout the simulation, but we represent all metals as a single metallicity field. Thus, we assume the metallicity of any given metal X tracks with this overall metallicity Z consistent with solar abundance ratios: $n_{\rm X}/n_{\rm H} = Z \times (n_{\rm X}/n_{\rm H})_\odot$. To calculate (n_{X_i}/n_X), we create a lookup table using the spectral synthesis code, Cloudy (version 07.02.01), which gives (n_{X_i}/n_X) uniquely determined by the gas temperature, gas density, and the photoionizing spectrum and intensity at that time. We use the redshift-dependent ultraviolet background included in Cloudy as the photoionizing source (a 2005 updated version of Haardt & Madau 1996), which takes into account reionization by QSOs and galaxies. The Cloudy lookup table accounts for collisional ionization as well as photoionization by the metagalactic UV field (although we do not track local photoionizing sources). This infrastructure coupled with the yt simulation analysis suite[3] (Turk et al. 2011), allows us to add any arbitrary ion and element to our simulation as an additional data field, as defined by the local baryon

[3]http://yt-project.org

density, metallicity, temperature, and redshift. We process all of our data outputs to include fields for the astrophysically important ions of H I, Mg II, Si II, Si III, Si IV, C IV, N V, O VI, and O VII.

For each data output of each simulation, we locate the galaxy of interest, and create a series of ten column-density projections centered on the galaxy from randomly oriented vectors passing through the galaxy in each of the selected fields. Each projection consists of two 128^2-pixel images, one with a field of view of 2 Mpc and the other with 100 kpc, allowing us to achieve a large dynamical range in resolution (1 Mpc to 1 kpc resolution in the central region of the galaxy). These images provide a map of the column density of each species and can be compared directly against observations. Figure 4.1 displays a few such projections from the MFCS run in H I, C IV, O VI and total gas density at $z = 0.5$. For illustration's sake, these images are 500 kpc on a side, with a circle overplotted representing r_{200}, the radius at which the galactic density drops to 200 times the mean density of the universe. The column density limits chosen are those used later in the radial profiles for each species.

For the image plane, pixels are identified by their comoving distance from the central pixel. These pixel radii are then binned, and the bins are populated by the column density maps of a series of images for a given galaxy. The median column density value in each radius bin is computed to produce a radial profile for that galaxy, representing the typical column density one might expect to detect at a given radius were one to observe it from a random angle. We found early on that there were significant temporal peculiarities (e.g.

infalling satellites) which made it difficult to compare one run at a given redshift against another, so we made radial profiles consisting of data outputs over a range of redshifts to wash out this effect. To give some insight into the range of column densities at each radius, we also include some quartile column density ranges (i.e. 25% and 75%) in these radial profiles.

For the physical quantities of metallicity and temperature, we take a mass-weighted average projection for each quantity as opposed to a column density. All further processing steps on these projections follow the same prescription as for the atomic species projections.

4.2.5 Processing the Observational Datasets

In order to say something meaningful about our numerical models, we compare them directly against observations taken in a similar manner. Typically, gas in these species is not observed in emission, but rather in absorption due to its low column density (emission $\propto N^2$, whereas absorption $\propto N$). Thus, observational investigations of the circumgalactic medium consist of finding quasars sightlines passing close to intervening galaxies, taking spectra of these targets, and measuring atomic absorption lines in these spectra.

We identified observational studies which collate the proximity (i.e. $\rho \equiv$ impact parameter) of absorbers to nearby host galaxies for a subsample of our chosen atomic species. The observational studies against which we compare our data are: H I (Lyα): Prochaska et al. (2011) (utilizing data from Penton et al. 2004; Danforth & Shull 2008;

Tripp et al. 2008); Mg II: Chen et al. (2010); C IV: Chen et al. (2001); O VI: Tumlinson et al. (2011); and O VII: Yao et al. (2008). In some of these studies (e.g. Chen et al. 2001, 2010), the absorption features were not fully resolved into individual components, and so only equivalent widths of the absorbing gas were provided, whereas in all of the other studies both equivalent width and column densities of the absorber exist. In order to directly compare our results against all of these studies, we needed to convert any equivalent widths to column densities.

Converting between equivalent width and column density is not a trivial proposition, as it requires knowledge of the temperature and velocity distribution of the gas containing the absorbing species. To make an approximate translation, we used a "curve of growth" analysis. We applied a fitting function relating equivalent width and column density as found in Draine (2011):

$$W \approx \begin{cases} \sqrt{\pi}\frac{b}{c}\frac{\tau_0}{1+\tau_0/(2\sqrt{2})} & \text{for } \tau_0 < 1.25393 \\ \left[\left(\frac{2b}{c}\right)^2 \ln\left(\frac{\tau_0}{\ln 2}\right) + \frac{b}{c}\frac{\gamma_{lu}\lambda_{lu}}{c}\frac{(\tau_0 - 1.25393)}{\sqrt{\pi}}\right]^{1/2} & \text{for } \tau_0 > 1.25393 \end{cases} \tag{4.6}$$

$$\tau_0 \approx \sqrt{\pi}\frac{e^2}{m_e c}\frac{N_l f_{lu}\lambda_{lu}}{b} \tag{4.7}$$

where W is the equivalent width of a given absorption line for a transition from state l to state u; $b \equiv \sqrt{2}\sigma_V$ is the velocity dispersion of the absorbing gas; c is the speed of light; τ_0 is the optical depth at absorption line center; γ_{lu} is the intrinsic width of the absorption line; λ_{lu} is the wavelength of the transition; N_l is the column density of material in the l

state; and f_{lu} is the oscillator strength of the transition.

This relation converts equivalent width to column density, but we want the inverse, so we numerically inverted the function applying the bounds-bisection method (Press et al. 1992). We used Verner et al. (1996) for accurate oscillator strength, wavelength, and γ values for the various transitions.

To deduce the appropriate value of b for our absorber populations, we utilized the equivalent width and column density data in Tumlinson et al. (2011) and Prochaska et al. (2011). Applying equation 4.6 to calculate column density from these equivalent widths for a variety of b-values, we then made a χ-squared fit with the listed column density values to ascertain the best choice of b for O VI and H I respectively. Our results indicated $b_{\rm OVI} = 58$ km s^{-1} and $b_{\rm HI} = 29$ km s^{-1}. Since the gas responsible for absorbing C IV is in a similar temperature regime (slightly lower) to that of O VI absorbing gas, we used $b_{\rm CIV} \approx b_{\rm OVI} = 58$ km s^{-1}. Additionally, since Mg II has a similar ionization energy as H I, and Mg II is thought to trace H I at $N_{\rm HI} > 10^{16}$ cm^{-2} (Putman et al. 2012), we let $b_{\rm MgII} \approx b_{\rm HI} = 29$ km s^{-1}. Since the lines of many species have multiple components, these b values should not be take too literally, but for the purposes of this current comparison, they suffice. We thus convert the observational datasets to approximate column densities (and proper units to comoving units) for comparison against the numerical results.

Name	Description	M_{200} (10^{10} M$_\odot$)	$M_{\rm DM}$ (10^{10} M$_\odot$)	$M_{\rm star}$ (10^{10} M$_\odot$)	$M_{\rm gas}$ [hot, cold][1] (10^{10} M$_\odot$)	r_{200} (kpc)	$v_{\rm circ,max}$ (km s^{-1})	$\epsilon_{\rm FB}$	$t_{\rm supp}$ (Myr)
V	Vanilla	110	93	17	2.5 [1.8, 0.6]	210	590	0	0
LF	Low Feedback	110	92	15	3.6 [2.8, 0.8]	210	560	3×10^{-6}	0
MF	Medium Feedback	110	91	11	5.1 [4.4, 0.8]	210	520	1×10^{-5}	0
HF	High Feedback	90	84	2	4.3 [3.0, 1.3]	200	200	3×10^{-5}	0
MFCS	Medium Feedback & Cooling Suppression	100	90	7	6.6 [5.7, 0.9]	200	390	1×10^{-5}	50

Table 4.1 We present each of our simulations, their simulation parameters, and the resulting halo characteristics at $z = 0$.

Species	$E_{\rm ion}$	Redshift	Reference
H I	13.6 eV	0.05 - 0.2	Prochaska et al. (2011)
Mg II	15.0	0.1 - 0.5	Chen et al. (2010)
Si II	16.3	—	—
Si III	33.5	—	—
Si IV	45.1	—	—
C IV	64.5	0.08 - 0.9	Chen et al. (2001)
N V	97.9	—	—
O VI	138.	0.1 - 0.35	Tumlinson et al. (2011)
O VII	739.	0	Yao et al. (2008)

Table 4.2 The observational datasets used to compare against our numerical radial profiles. We also include ionization energies for our selected ionic species.

4.3 Results

Our suite of simulations yields halos with similar bulk properties, consistent with the fact that they all share the same initial conditions. We catalog the mass of each component (dark matter, gas, and stars) and the maximum circular velocity for the primary halo in each of our simulations at $z = 0$ in Table 4.1. The mass in each component is measured interior to r_{200} at $z = 0$. There is little variation in the dark matter and total mass across our simulations except in the case of the high feedback run where large baryonic outflows have softened the potential and diminished the overall mass of the system.

In Paper I, we examined the rotation curves and star formation histories of a set of

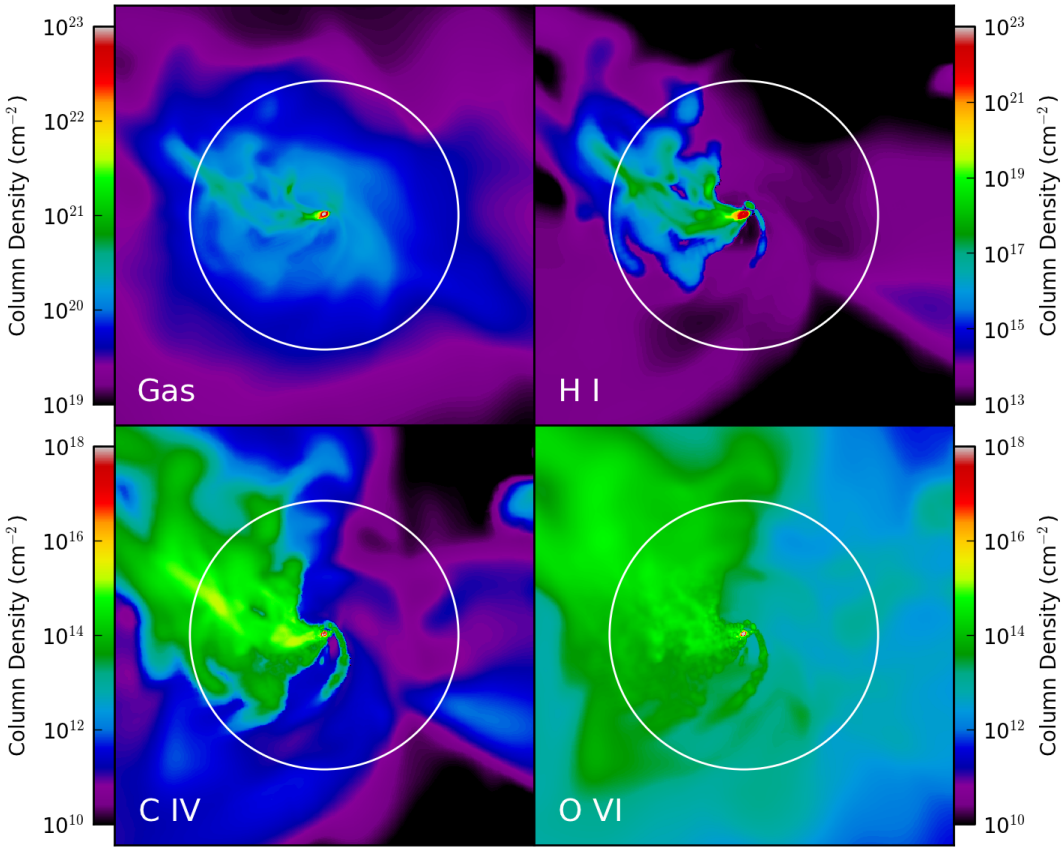

Figure 4.1 Sample projections from the medium feedback and cooling suppression run at $z = 0.5$. Each image represent a column-density map of a region with 500 comoving kpc on a side. The white circle denotes $r_{200} = 160$ kpc.

simulations very similar to those used in this paper. Here, we focus on the observable properties of the circumgalactic medium (CGM) at low redshift – in particular, we report on the column density distributions of a wide variety of atomic species, observable in absorption against background quasars. We initially restrict ourselves to low redshift, as there exists a substantial body of observations which link absorption strength to projected distance from the galaxy. In the following sections, we first examine the column density

distributions, then explore the physical properties of the gas giving rise to this absorption, and finally investigate the redshift evolution of the predicted absorption strength.

4.3.1 Radial Profiles for Atomic Species

In this study, we examine the radial column density distribution for a number of absorption-line-generating species, well-sampling the range of ionization energies (and thus temperature/density regimes) of the CGM as shown in Table 4.2. At the two ends of the spectrum, the H I traces the cold, dense clouds, whereas the O VII probes the hot coronal gas. Between these regimes is a continuum of layers of material acting as the interface between these two phases. This layering effect takes place on many different spatial scales from small star-forming clumps of cold gas up to the galactic scale. Figure 4.1 demonstrates this at the latter scale, where the cold H I of the inner halo is enshrouded in warmer C IV-bearing gas, which is further encased and extended by gas exhibiting O VI. This provides us with some intuition for understanding the following galactic profiles for these species.

We produce radial profiles of the median gas column density in nine different ionic species as a function of projected radius extending from 2 kpc to 1 Mpc for our simulated galaxies. We probe this gas from the disk-dominated region ($r < r_{\rm inner} \sim 20$ kpc) through the CGM ($r_{\rm inner} < r < r_{200}$) into the intergalactic medium ($r > r_{200}$). Figures 4.2, 4.3, and 4.4 present these results (Figure 4.5 shows total gas column densities for comparison). The left and right sides of each figure show the radial profiles of the same quantity but for different simulations. The left sides display a comparison of simulations with and

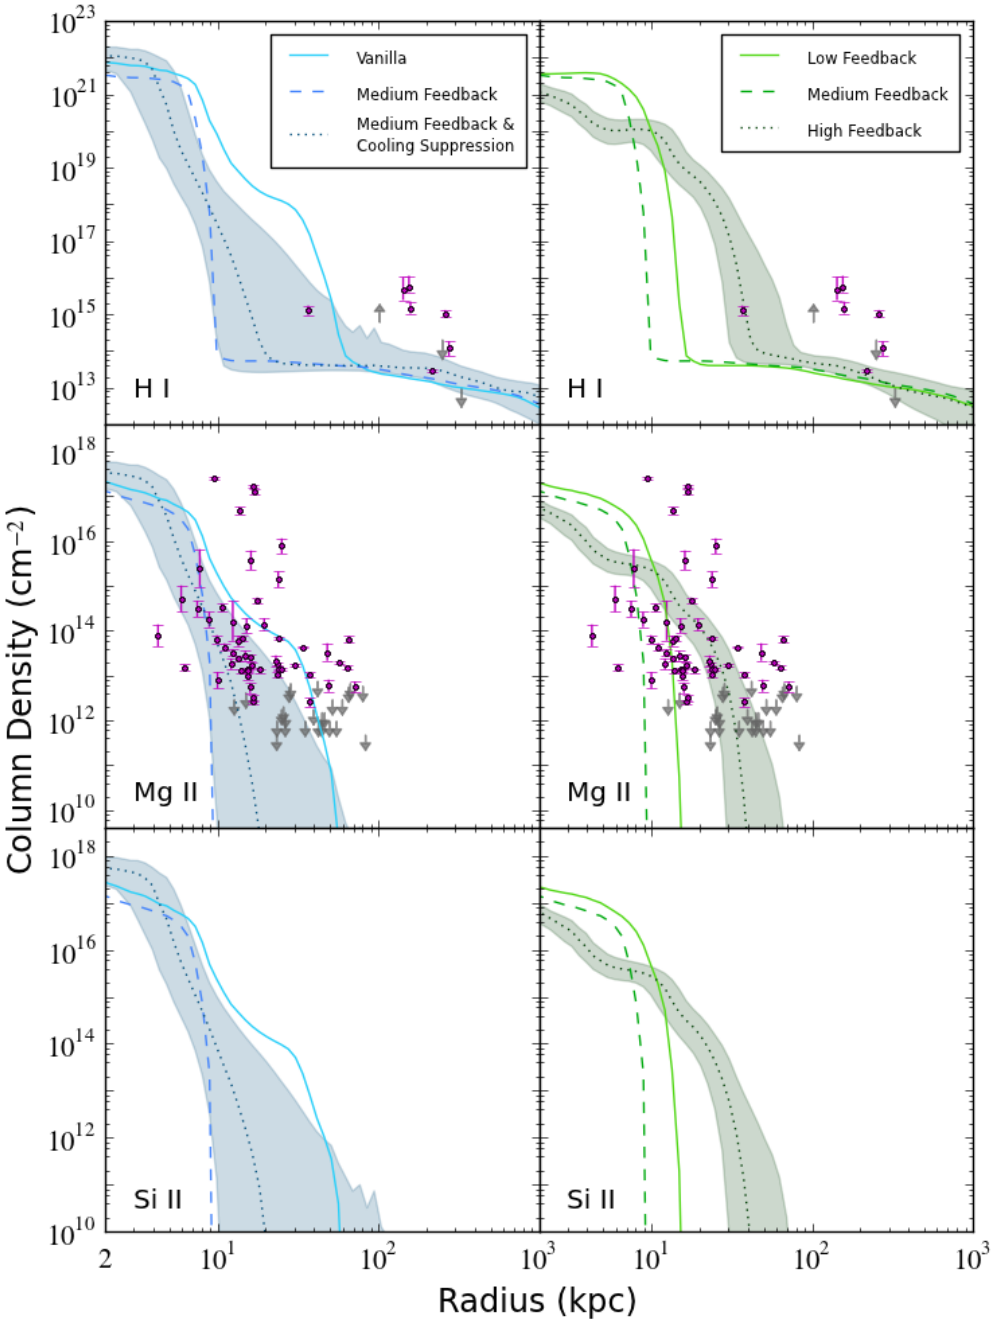

Figure 4.2 Radial profiles for column densities of H I, Mg II, and Si II in comoving units covering the period $z = 0.1 - 0.5$. The color bands represent the quartile values (±25%) overplotted on the median for two of the models: MFCS and HF.

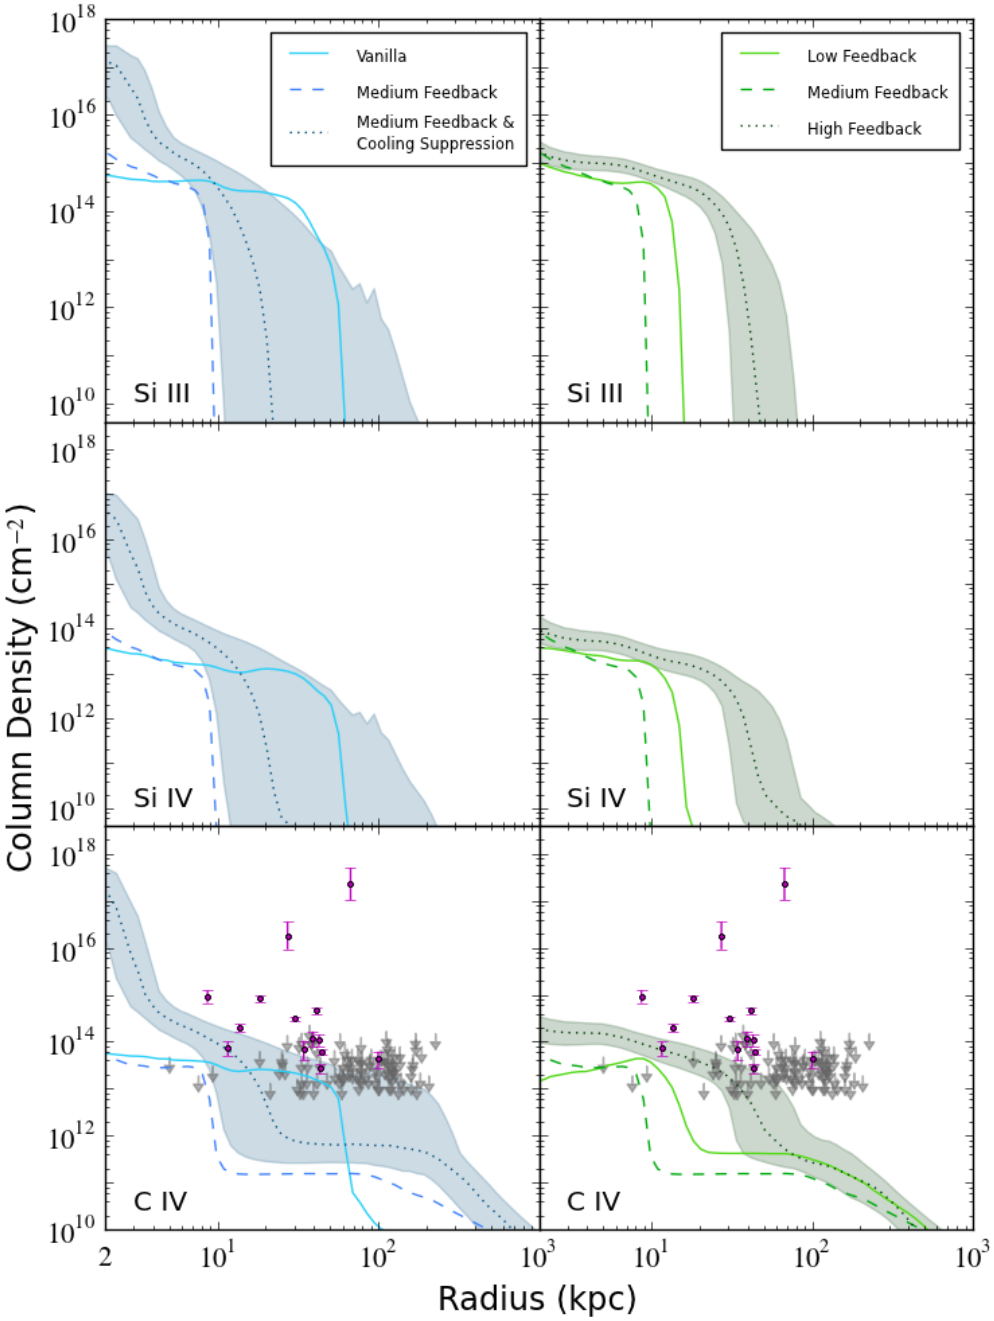

Figure 4.3 Similar to Figure 4.2 but for Si III, Si IV, and C IV.

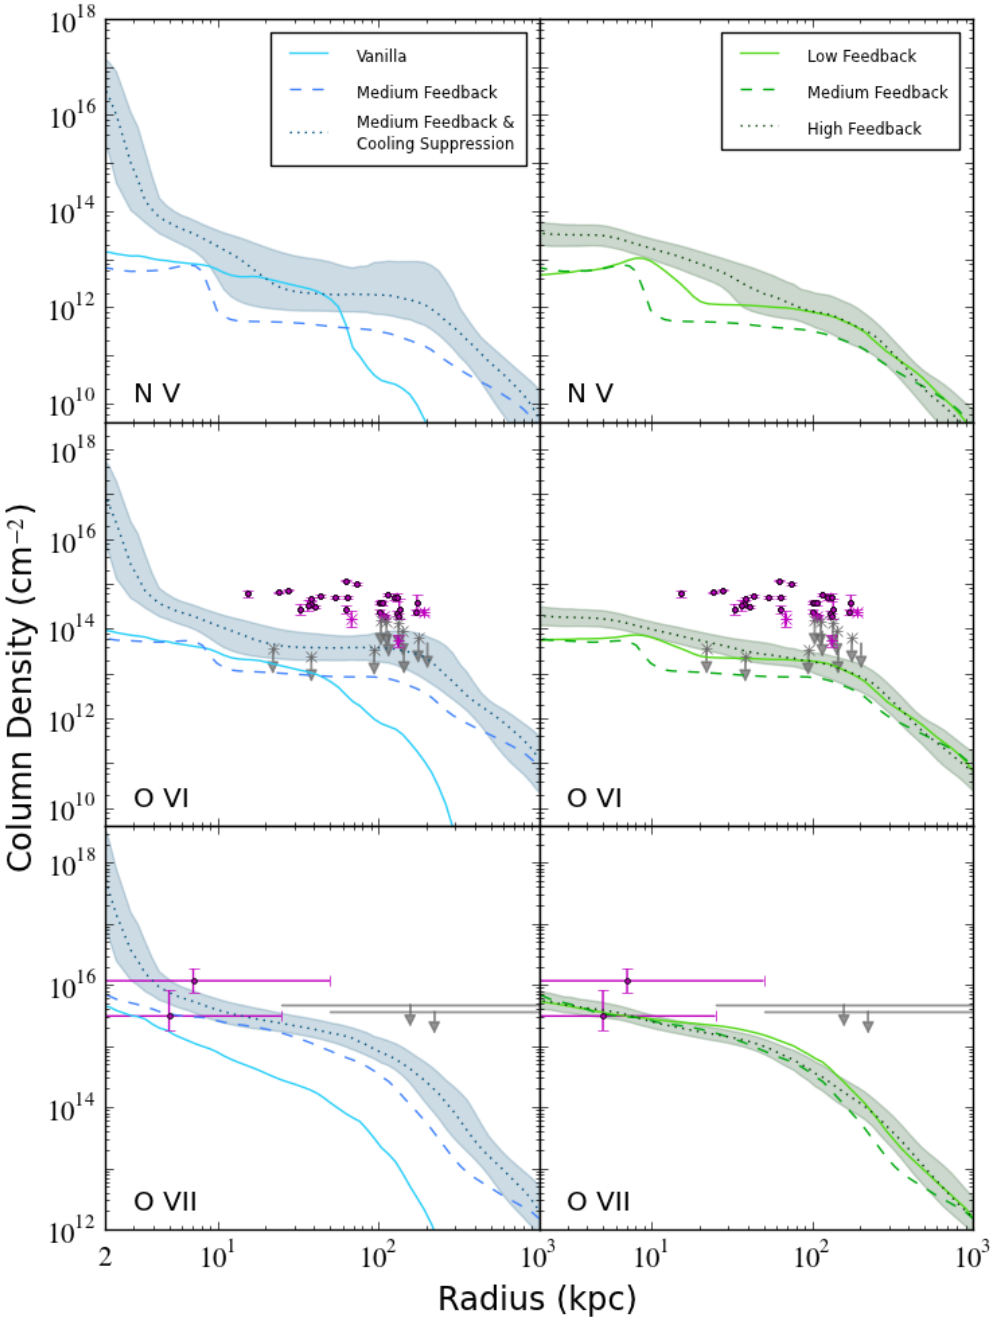

Figure 4.4 Similar to Figure 4.2 for N V, O VI, and O VII. The O VI dataset (Tumlinson et al. 2011) is split into star-forming galaxies (circle symbols and plain arrows) and passive galaxies ("asterisk" symbols and "asterisk" arrows).

without different forms of feedback: V, MF and MFCS, whereas the right sides exhibit a progression of thermal feedback intensities: LF, MF and HF (the MF simulation is repeated on both sides to ease comparison). Each profile gives the median value expected in an observation of a galaxy at a given radius over the redshift range $z = 0.1 - 0.5$ (selected to approximately match the redshift ranges of the observations). To give some idea of the variation around the median absorber strength, in each panel we overplot a color band representing the quartiles of the column density distribution (i.e. 25% to 75%) for one model: MFCS (left) and HF (right). The quartile distribution associated with the HF model is most indicative of the spread of all the other models, whereas MFCS produces a particularly spread-out distribution suggesting more multiphase substructure. In Figure 4.2 we present radial profiles for our lowest ionization energy species, H I, Mg II, and Si II; Figure 4.3 shows medium energy ions, Si III, Si IV, and C IV; and Figure 4.4 displays the highest energy species N V, O VI, and O VII. Overplotted on several profiles are low-redshift observational datasets, where magenta points represent detections and grey arrows mark limits. These datasets, their sources, and their redshift ranges are listed in Table 4.2.

The H I, Mg II, and Si II datasets (the H I group) in Figure 4.2 demonstrate similar behavior across all of our models, which is consistent with the fact that their ionization energies are very close to each other (see Table 4.2). All of the simulations show a high-column density core in these species out to $r \sim 10$ kpc, at which point the profiles steeply drop to trace levels by 10-50 kpc. This transition radius, $r_{\rm inner}$, marks the fall off of the

disk of the galaxy (and associated inflowing cold gas). Each feedback model predicts a different $r_{\rm inner}$, sometimes differing over 40 kpc in radial location, identifying $r_{\rm inner}$ as a useful differentiator between feedback models. Our V and HF simulations can match the bulk of the Mg II observations; however none of the models reproduce the most extended high column density observations in H I and Mg II. Note that we have subsampled the H I observations (Prochaska et al. 2011) including only those detections associated with galaxies in the $0.5L^* < L < 2L^*$ range since these best match the characteristics of our simulated galaxy.

In Figure 4.3, the medium-ionization species Si III, Si IV, and C IV (the "C IV group") display somewhat similar behavior to those of the H I group, each exhibiting a core and precipitous drop in material around 10-40 kpc. The floor level of the C IV absorption strength (beyond $r_{\rm inner}$) is larger than either of the other two lines, presumably due to the fact that it is tracing warmer gas. We also compare to the observed C IV dataset, but note that this comparison is difficult due to the somewhat contradictory limits and detections in the outer halo; however, all of the simulations appear to be an order magnitude too low in column density in the probed radial regime ($r \sim 10 - 200$ kpc).

The tracers of warm and hot gas, N V, O VI, and O VII, have more extended and flatter distributions than the H I and C IV groups. The O VII distribution, which probes gas with temperature close to the virial temperature of the halo is particularly flat (note the different vertical scale). Our models that best fit these observational datasets are the HF and MFCS simulations; however even those profiles are an order of magnitude below

the observations of O VI in the $r \sim 10 - 100$ kpc range. These same models reasonably fit the O VII observations, although this dataset has relatively weak constraints. The vanilla run conspicuously falls off in column density at $r \sim 50$ kpc in all of the species (we will address the reason for this in more detail below). The O VI observations are separated into "passive" galaxies (plotted with an asterisk) and star-forming galaxies (plotted without an asterisk), the latter being potentially more comparable to our simulated galaxy (see Tumlinson et al. 2011).

4.3.2 Radial Profiles for Physical Quantities

In order to better understand the physical basis for the simulation predictions, we next turn to the physical properties of the CGM gas itself. We present radial profiles for median gas column density, median metallicity, and median temperature in Figure 4.5 in a manner similar to Figures 4.2, 4.3, and 4.4. In order to make a fairer comparison to the projected column density profiles, we compute these profiles using the same machinery, first carrying out several projections for each quantitiy (mass-weighting for the metallicity and temperature), and then computing the median value for each projected radius over the same $z = 0.1$ to 0.5 redshift range used earlier. The profiles computed in this way are similar to the three-dimensional quantities, but are more straightforward to compare at a given projected radius.

For the gas column density profile, there is not a great deal of variation (within a factor of 2) from one simulation to another, suggesting that different feedback models

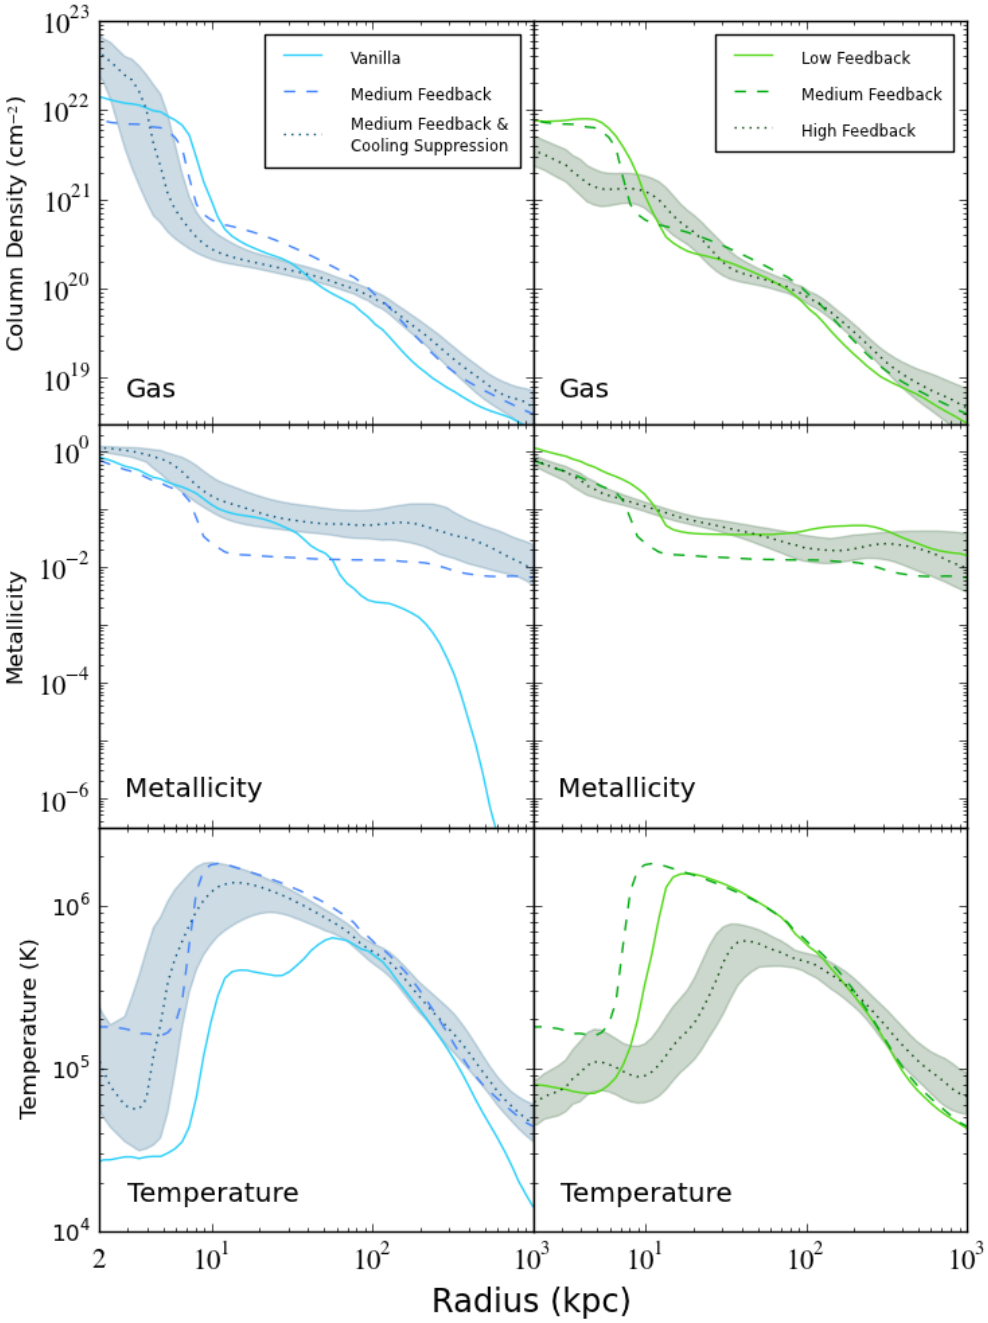

Figure 4.5 Radial profiles for physical quantities: column density of gas, mass-weighted metallicity, and mass-weighted temperature covering the period $z = 0.1 - 0.5$ for two of the models: MFCS and HF.

play only a modest role in rearranging the radial distributions of the gas mass within a galaxy.

The metallicity profiles demonstrate a much more dramatic difference between those runs incorporating energetic feedback (LF, MF, MFCS, HF) and the run which does not (V). While all of the models share roughly the same median metallicity in their inner 50 kpc (within 1 dex of Solar), the median metallicity of the vanilla run plunges to nearly primordial values beyond that distance, explaining why the vanilla run drops off in there in C IV, N V, O VI and O VII. All of our simulations share the ability to produce and return metals to the surrounding medium from star-formation events, but only the "feedback" runs are capable of depositing energy along with those metals. It is clear that the inclusion of some form of energetic feedback is necessary to mix and disperse high-metallicity gas beyond 50 kpc. Lacking energetic feedback, the V simulation effectively has no outflow of material, and the metals get mixed up by gas motions in the inner halo region. This fact makes it very useful for designating the Vanilla simulation as an *inflow-only* model, thus cleanly isolating the characteristics of infalling material versus models generating outflows. Because our inflow-only simulation shows similar concentrations of gas in the outer halo, we can conclude that most of the gas in the halo must have been deposited there by inflows. Evidently, the main effect of energetic feedback is to mix in higher-metallicity material to large halo distances, but not to enhance the density there dramatically.

All of the simulations share the same general behavior in the projected temperature profiles: low in the core (the cold gas of the disk) out to $\sim r_{\rm inner}$, higher values extending

to around the virial radius, and then decreasing beyond that. Note that the "cold" gas temperatures ($T \sim 10^5$ K) seen in the core in these profiles reflect the presence of multi-phase gas. The LF, MF and MFCS runs exhibit a high peak temperature of $T \sim 20$ million Kelvin at $r_{\rm inner}$, where the high density of the galactic disk drops off. Conversely, the V and HG runs peak farther out ($r \sim 40$ kpc) at a lower temperature ($T \sim 7 \times 10^5$ K). The lack of clear trends in the temperature profiles indicate that low values in the center reflect simply how much cold gas survives in each model, while the temperature in the outer parts of the halo probably depends largely on the halo's virial temperature. These temperature profiles, however, are clearly related to the species profiles shown earlier – the low ionization species (particularly the H I group) only appear where the mass-weighted gas temperature is low. Even the C IV group is correlated strongly with the extent of cold gas in the temperatures, probably indicating that it is formed in multi-phase gas. Finally, the column density profiles of the higher ionization species are generally flatter than the gas profile, indicating that they are primarily coming from the outer halo and their slight overdensity in the core is mostly a projection effect. This is particularly true of the O VII profile, which is very flat.

4.3.3 Redshift Evolution of Radial Profiles

There exists a significant body of observations as well as a growing number of computational studies investigating the radial profiles of metal-bearing gas in the CGM of galaxies at high redshift (e.g. Altay et al. 2011; Tescari et al. 2011; Fumagalli et al. 2011; Barnes et al.

2011; Shen et al. 2012). While the focus of the current work is low redshift, it is worth examining how the CGM evolves over time in our simulations. Figure 4.6 presents radial profiles of gas column density, mass-weighted metallicity, and O VI column density for two of our models (V and MFCS) as a function of redshift. These profiles are constructed using the same method of median-sampling many data outputs over a range of redshifts so as to wipe out any temporal peculiarities in the profile. The figure demonstrates trends in the behavior of certain quantities with time sampling redshifts from $z = 4$ to $z = 0$.

The gas column density behaves similarly in both of our models, consistent with an expanding universe. From an average redshift of 3 to 0, we see a change of the scale factor of four, so that the column density of a uniform medium integrated over a comoving path length should decrease by a factor $N^2 = 16$, which is approximately consistent with these profiles, particularly at large projected radii. On the other hand, metallicity appears to increase over time in both models, especially for the inflow-only simulation, although the effect is quite weak. In the MFCS run, the metallicity grows with time at all radii, but by only a few percent. Much of the star formation occurs at early times and is spatially distributed through the galaxy quickly enriching the gas at all radii, so subsequent star formation seems to do little to further increase the metallicity throughout the halo.

Finally, the O VI radial profile in this figure exhibits behavior similar to the gas profile, in that the column densities at all radii drop over time. While the high-redshift version of the MFCS model matches the O VI observational data relatively well, once it has reached low redshift (the redshift of the actual observational dataset), it has fallen

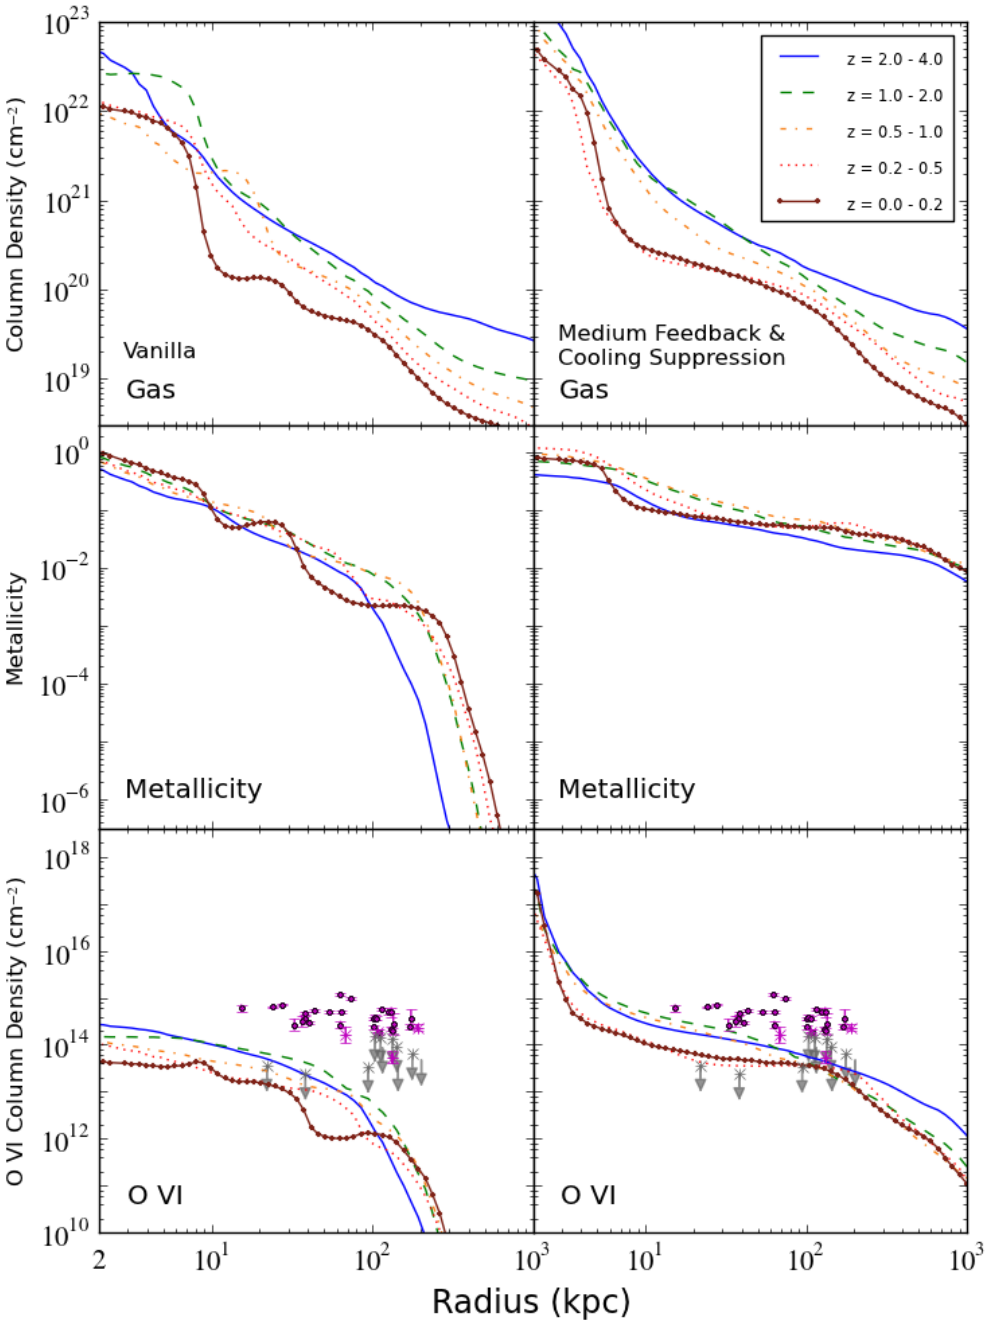

Figure 4.6 Radial profiles for gas, metallicity and O VI as they change over time. The left and right sides display the evolution of radial profile for the V and MFCS runs respectively.

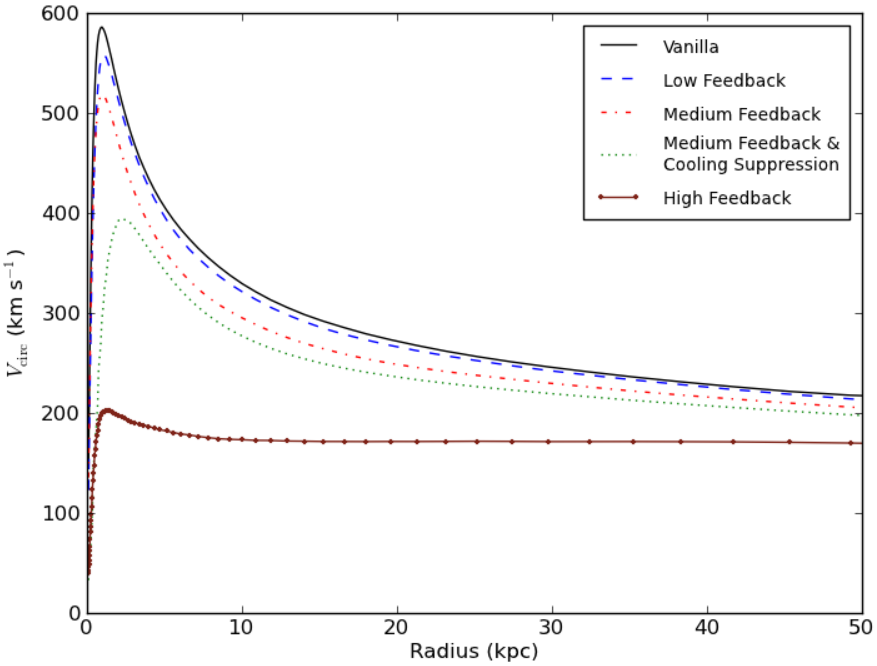

Figure 4.7 Rotation curves for each of the simulations at $z = 0$.

almost an order of magnitude. This behavior of ionic column densities to decline by an order of magnitude at all radii is somewhat generic across all of our species, although some of the line/simulation combinations have a more subtle time-dependence belying the complex nature of the Saha equation.

4.3.4 Rotation Curves

In Paper I, we analyzed the efficacy of different subgrid models (including feedback prescriptions) using the rotation curves of the resulting galaxies as a metric of success.

We revisit this analysis with the current models in Figure 4.7. Consistent with Paper I, simulations V (canonical run), LF and MF produce galaxies with very peaked rotation curves ($v_{\rm circ} \sim 600$ km s^{-1}), clearly suffering from the angular momentum problem. The simulation including cooling suppression (MFCS) was successful at reducing this peaked value significantly to $v_{\rm circ} \sim 400$ km s^{-1}. The high feedback run resulted in a system with a very flat rotation curve and about 20% less total mass than the other halos. While our MFCS and HF simulations may appear as successful models (i.e. realistic) based on this metric alone, the preceding analysis comparing ionic column densities makes it clearer that they do not predict significantly different absorber distributions from the lower feedback runs (and fail to match observations). This fact alone reveals that QSO absorption line observations are an important constraint for feedback models to match in addition to traditional measures of success (e.g. rotation curves, star formation rates, Tully-Fisher Relation, Kennicutt-Schmidt Relation).

4.4 Discussion

For all of the simulations, we see the same general behavior over each of the ionic radial profiles. The low-energy species of H I, Mg II, Si II, Si III, Si IV, and C IV display a core column density level out to a characteristic radius, $r_{\rm inner}$, beyond which they fall steeply to a floor of trace levels indicating that these ions tracing cool gas are mostly associated with the disk. As we increase the ionization energy of the species, the core column

density diminishes and the floor column density level increases indicating warmer, more distributed gas being traced. This trend continues to the highest energy ions of O VI and O VII, which display hot, extended gas in mostly flat radial profiles.

It is clear from our results that these simulations do not produce column densities distributions of various ionic species matching QSO absorption observations. In general, our models predict column densities which are consistently less extended and of lower concentration than the observational studies, save for the predictions of the hottest gas in O VII absorption. Furthermore, there is only a hint of a trend in our data suggesting a path to resolve this discrepancy (e.g. by increasing feedback efficiency). So why do our models fail to agree with observations?

One obvious problem with our method is that we only consider column densities of material within 1 Mpc of our galaxy. In reality, quasar sightlines pass through a much larger volume, and it is possible that additional intervening absorbers with velocities similar to our galaxy exist, which would increase the column densities in subsequent radial projections. However, this effect is likely to be small, so we seek additional explanations for these column density discrepancies.

4.4.1 Matching O VI

In order to better understand the shortcomings of our models, we attempt to reconstruct the conditions necessary to reproduce the observational datasets. In particular, we try to match the O VI observations, since of all the observational datasets compared, they

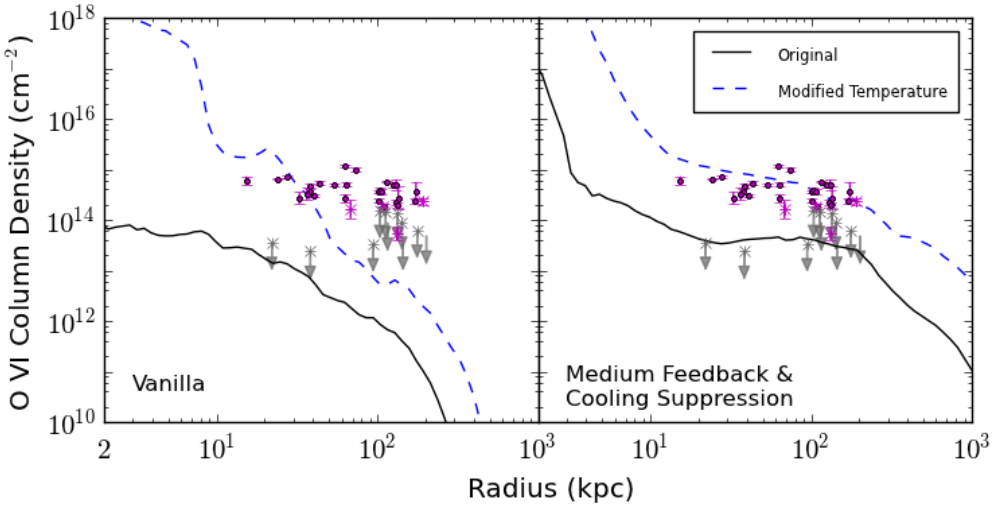

Figure 4.8 Radial profiles of O VI column density for a single data output at $z = 0.2$. The left and right sides show the V and MFCS runs respectively. The O VI column density of the original model is displayed in black, whereas the maximum O VI column density when we leave temperature as a free parameter is displayed in blue.

have the smallest errors and are the least likely to be susceptible to inaccuracies from photoionization (see section 4.4.2 below). As noted previously, our radial profiles of O VI column density are 1-2 orders of magnitude lower than the observations in the $r \sim 10-100$ kpc region.

Broadly speaking, one could imagine two explanations for this deficit. First, there could simply not be enough metal-enriched gas at the projected radii in question, or second, the gas could be there in sufficient quantities, but not in the correct ionization state, perhaps because the temperature is too high or too low. O VI is particularly susceptible to the thermal state of the gas, as it only exists in sufficient quantities in a relatively narrow temperature range. To explore which of these two explanations is most consistent with

our simulations, we take the radial profiles of physical quantities generated in Figure 4.5, and systematically vary either the density, temperature, or the metallicity at each projected radius in order to maximize the resulting O VI column densities, holding all other quantities constant. Under the assumption that collisional ionization of O VI dominates, we fix the radiation intensity (see Section 4.4.2).

We test this out on a single data output at $z = 0.2$ in the MFCS run although the result holds at other times. First, we find that for a fixed metallicity and temperature, changing the density has little effect – this might be somewhat counterintuitive, but is expected because we are keeping the column density constant, and O VI collisional ionization depends very little on density. The metallicity has a much larger effect, but we find that there is no solar or sub-solar metallicity values which will bring our column densities up to the observed levels (we limit the metallicity to solar on plausibility grounds). On the other hand, modifying the temperature (i.e. by lowering it to 3×10^5 K for most radii) while holding everything else fixed succeeds in producing sufficient column densities to match (or even exceed) the observational data, as seen in the right panel of Figure 4.8.

This suggests that our MFCS model possesses sufficient quantities of enriched gas at large radii, but that insufficient amounts of cooler gas exist (necessary for O VI). Note that we are not free to reduce the gas temperature everywhere (this would decrease the pressure and hence change the gas distribution); however, one solution would be to introduce multi-phase gas, preserving the temperature but increasing the amount of gas in the 3×10^5 K temperature range. This result appears to apply to the other feedback

simulations as well, that is, the median temperature is generally too high in our CGM, leading to depleted concentrations of these ions in preference to higher ionization states, but modifying the temperature allows us to match the observations.

Interestingly, the results for the Vanilla run (shown in the left panel of Figure 4.8) are quite different – fine-tuning the temperature in this model cannot bring it in line with observations, since the metallicity falls off far too quickly at $r \sim 50$ kpc to allow for the modified temperature to compensate for it, further confirming that feedback and outflows are necessary in order to produce an observationally consistent simulation.

4.4.2 Limitations

To better understand the limitations of our models, we look at uncertainties related to photo-ionization and collisional ionization. To help focus this discussion, we separate our chosen species based on their primary ionization source. Generally the 'high' ions (N V, O VI, O VII) are thought to be associated with collisional ionization, whereas the 'low' ions (Mg II, Si II, Si III, Si IV, C IV) can be the result of both collisional and photoionization (Danforth & Shull 2008). Oppenheimer & Davé (2009) suggest that many of the weaker O VI absorbers may be due to photoionization, however even they agree that absorbers with equivalent widths greater than 100 mÅ (most of our observational comparison dataset) are predominantly due to collisions. With that division made, we can discuss the the ionization sources separately.

4.4.2.1 Photoionization

Our low-energy species are partially ionized by absorption of energetic photons, so any problems in our modeling of photoionization will affect their column density profiles. Our study calculates photoionization for a metagalactic ionizing background in the optically thin regime, but it does not account for local photoionizing sources, such as energetic stars and galaxies within the volume itself. Inclusion of these sources will increase the ionization state of the various metals present, thus potentially increasing the column density of these ions. Its dominant effect will be felt close to the disk of the galaxy ($r < r_{\rm inner}$), where the bulk of the star formation is occurring and the local radiation field is high. However, the discrepancy between observations and the models occurs well outside of the disk, so while the lack of local photoionization may have some effect on our profiles it will not explain away the observed differences in column density profiles for H I, Mg II and possibly C IV beyond $r_{\rm inner}$.

Cool gas at high densities will act to shield itself from photoionizing radiation. In this paper, we have treated all gas in the optically thin regime for absorption of the ionizing background; however, gas with column densities $N_{\rm HI} > 10^{17.2}$ cm^{-2} should be partially or completely shielded from radiation more energetic than 13.6 eV (e.g., Faucher-Giguère & Kereš 2011). This will have the effect of lowering the ionization state of the various species present, increasing the column density of the lowest ionization species preferentially. Examination of the H I radial profile reveals that column densities of this magnitude only occur in the near-disk region of the galaxy ($r < r_{\rm inner}$). Interestingly, our

lack of self-shielding may in some cases be balanced by our lack of local photoionizing sources, since they both dominate in the disk and can have counteracting effects depending on the ion. Either way, neither seem likely to have significant repercussions at large radii, in the region of the greatest observational discrepancies.

Like many simulation groups, we use a Haardt-Madau photoionizing background (Haardt & Madau 1996, updated to the 2005 version), but there have been some advances in constraining the nature of this background and its sources in the universe since this time. A few improved approximations of the background have appeared recently (e.g. Faucher-Giguère et al. 2009; Haardt & Madau 2012). Employing a different background has the potential to modify the ionization balance, so as to change the column density of gas in these low-energy ions, although it is unlikely that it will elevate our results by the order of magnitude difference required to align them with observations.

Current work is ongoing to include the aforementioned physics into our simulations. However, while these effects might make some differences in the simulated profiles, the changes will be confined to the disk region, failing to bring our models substantially closer to agreement with observations.

4.4.2.2 Collisional Ionization Equilibrium

The use of Cloudy in calculating the ionization state of our various species assumes collisional ionization equilibrium (CIE), which is an ideal condition requiring $\tau_{\rm cool} >> \tau_{\rm CIE}$ (Dopita & Sutherland 2003). The CIE assumption holds in high-density, low-temperature

regions where the cooling time is short, but it can break down in the low-density regions of the IGM (typically $r > r_{\rm vir}$) (Cen & Fang 2006). Cen & Fang (2006) investigated the differences in O VI and O VII absorber strengths with and without the assumption of CIE and found that CIE overpredicts equivalent widths for these two ions by a small factor (less than 50%). Notably, properly accounting for non-equilibrium effects would decrease our column densities of O VI and O VII, pulling them farther away from the observational sample.

The high-energy species associated with warm and hot gas are primarily ionized due to collisions with ambient electrons. The rate of these collisions increases with electron density and density of the medium, so that clumpy regions of high density should have a large amount of collisional ionization. Thus, resolution becomes a key element in order to assure one can resolve clumps sufficiently for proper modeling of their ionization state.

4.4.2.3 Spatial Resolution

To investigate the impact of resolution and its potential nonlinear effects on the collisional ionization equilibrium in our species, we analyzed several outputs in an attempt to identify the characteristic size of absorbers in various different ionic species. We found that clumps of material at moderate to high density in our fields came in a variety of sizes, ranging from over 10 kpc down to 425 comoving pc (our spatial resolution limit), but there was no preferred size scale of absorbers. This continuum of sizes indicates that there may be more gas clumps forming on subgrid scales not resolved in this study.

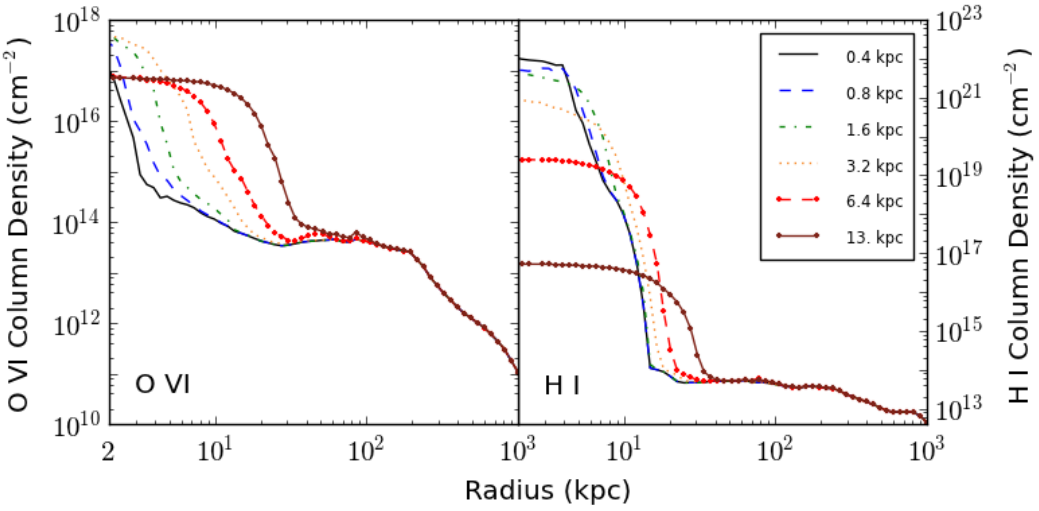

Figure 4.9 Radial profiles of O VI and H I column density as a function of downgraded spatial resolution for a single data output of MFCS at $z = 0.2$.

To further investigate this issue, we modified outputs from the simulations by re-sampling them at decreased maximum resolution (0.8, 1.6, 3.2, 6.4 and 13. kpc maximum spatial resolution instead of the 0.4 kpc used in the simulations) prior to post-processing them for the atomic species. This step is done on all fields (mass-weighted for the temperature and metallicity fields) in order to determine whether or not our atomic post-processing step reaches resolution convergence. This downsampling will show us where the multiphase substructure exists in the simulation, since smoothing two multiphase structures together (e.g. cold clumps in a hot medium) will result in gas with different ion populations.

Figure 4.9 displays the outcome of this study for O VI and H I column densities of an MFCS output at $z = 0.2$. Neutral hydrogen, which is very cuspy in the high-resolution

simulation (seen as the solid black line) is spread out over an increasingly large region as the spatial resolution diminishes. The core is particularly sensitive to this effect, so as the cold core is artificially mixed with the hotter surrounding medium, it raises the core mean temperature, leading to hydrogen ionization and less H I. On the other hand, the O VI column density actually increases in this region because the artificial warm component favors O VI over other oxygen ionization states. In each case, we observe that our change in resolution in post-processing has very little effect on the CGM at large radius (except for a small increase for O VI) suggesting that most of our multiphase substructure exists within $r \sim 10$ kpc. By this measure, we are converging on behavior outside of $r \sim 10$ kpc; however, this result could change if the underlying simulation had higher resolution in the halo, or if it employed a different prescription for generating outflows.

4.4.3 Comparison with Previous Work

A number of recent studies have investigated the state of the CGM in hydrodynamic galaxy simulations both in the high-redshift and low-redshift universe. Many such studies focus on neutral hydrogen in the CGM of simulated galaxies at high z. Altay et al. (2011) post-process an OWLS (Schaye et al. 2010) galaxy at $z = 3$ to include a self-shielding correction and the formation of molecular hydrogen, demonstrating that the inclusion of these features causes a flattening and steepening of the H I column density distribution function at $N_{\rm HI} \sim 10^{18}$ cm^{-2} and $N_{\rm HI} \sim 10^{21}$ cm^{-2} respectively. Faucher-Giguère & Kereš (2011) examine an SPH simulation of a galaxy at $z = 2 - 4$ by ray-tracing the ionizing

background to correct for self-shielding, discovering low covering factors of high-density H I in the CGM consistent with those we produce in our models. Barnes et al. (2011) investigate how modifying the intensity of galactic winds affects the H I distribution and subsequent Lyman-α emission from such systems at $z = 3$. Their momentum-driven wind feedback prescription yields H I distributions similar to our own MFCS run at that redshift.

Other researchers have explored using metal ions as possible probes of the CGM by post-processing their simulations to include additional species, as we do in the current study. Kacprzak et al. (2010) added a Mg II field to an AMR galaxy simulation at $z = 1$ (similar to our MF run), finding a Mg II-absorber size distribution and Mg II $r_{\rm inner} = 20$ kpc consistent with our own work. Fumagalli et al. (2011) examined outputs at $z \sim 2 - 3$ from an AMR simulation like our MF model, processing them to include self-shielding due to dust, radiative transfer from local ionizing sources, and H I and Si II fields. In agreement with our results, they determined that the concentration of Si II absorbers is an order of magnitude lower than observed values, suggesting they too may lack significant multiphase structures in the CGM. A suite of simulations modifying the feedback prescription to include AGN and wind-driven feedback was performed by Tescari et al. (2011), demonstrating that C IV statistics were considerably more sensitive to feedback implementation than H I in the intergalactic medium (IGM) at $z = 2 - 3$. While their conclusion agrees with our results in the outer parts of our halos ($r > 200$ kpc), we find that both H I and C IV are sensitive to feedback prescription in the CGM ($20 < r < 200$

kpc) as exhibited in Figure 4.2. Goerdt et al. (2012) post-process an AMR simulation of $z \sim 2.5$ galaxy output to include a self-shielding correction and H I, O I, C II, C IV, Si II, Si IV, Mg II, and Fe II, observing the distribution of these metals versus galactocentric radius using a feedback model similar to our MF simulation. In general our results at that redshift align well with their study, except our C IV column densities are an order of magnitude greater than theirs, which might be a due to their self-shielding correction suppressing higher levels of ionization in the CGM. Lastly, Shen et al. (2012) analyze a $z = 2.8$ output from the Eris Simulation (Guedes et al. 2011) akin to the MFCS model, including a spherically-symmetric correction for local photoionizing sources and adding H I, C II, Si II, Si IV, C IV, and O VI in post-processing. They show the radial dependence of O VI column density, generally agreeing (ours is flatter) with our MFCS results at $z \sim 3$ (see Figure 4.6).

At low redshifts, there have only been a few studies probing the nature of the CGM in numerical simulations. Van De Voort & Schaye (2011) demonstrated that inflowing, cold gas lacks significant metals in the outer CGM consistent with our vanilla run presented in Figure 4.5. Looking at H I, Fernández et al. (2012) examined a $z = 0$ galaxy (similar to our MF model) and investigated the source of H I feeding its CGM. They showed H I column densities and covering factors matching our own study, exhibiting a steep falloff of $N_{\rm HI}$ at $r \sim 10$ kpc but an extended floor of material beyond $r_{\rm vir}$. Stinson et al. (2011) produced a suite of SPH simulations run to $z = 0$ including two feedback intensities. Their "low" feedback run was similar to our LF run but included cooling suppression, whereas

their "high" feedback run (intensified in order to integrate the effects of radiation from massive, young stars) was $\sim 3\times$ stronger than our HF run while still including a form of cooling suppression. The resulting outputs were processed to include O VI data, and then radial profiles were produced for both H I and O VI. The column density profiles for the "low" feedback run are consistent with our LF simulation; however, the profiles of their "high" feedback run have column densities exceeding any of our simulations. These profiles align well with the observational datasets in H I and O VI, suggesting that their success may be a result of using extremely high feedback intensities (i.e. 10^{51} ergs released per 7.5 $M_\odot$ of stars formed). We see a possible trend in our models, potentially indicating that intensifying feedback could increase column densities of various species in the CGM region (notably O VI) but this prediction is not clear cut.

4.5 Conclusions

We perform the first study to probe the CGM at low redshifts comparing AMR simulations against a large host of observational probes. We carry out cosmological simulations of a Milky-Way-massed galactic halo using a variety of stellar feedback prescriptions, post-processing the data to compute the column density distributions of a wide range of atomic species, effectively probing different phases of the CGM. We median-average these distributions over many orientations and times (from $z = 0.1$ to $z = 0.5$) both in order to make a cleaner comparison with data, and to decrease the impact of sub-structure

on our resulting radial profiles. Lastly, we compare these results against a number of different low-z observational datasets correlating absorber intensity and impact parameter to nearby galaxies. Our main results are as follows.

- It is difficult for any of our feedback models to reproduce the column density distribution of some of the observed absorbers – in particular, we cannot match observations of Mg II beyond $r > 50$ kpc, and O VI at any radii. On the other hand, we find that the Mg II distribution at smaller radii ($r < 50$ kpc) can be reproduced in some of the models. In addition, the predicted O VII absorption (which probes hot gas) in simulations including moderate or strong feedback is in agreement with current observations.

- We find that some form of feedback is required to get significant amounts of metals into the CGM beyond ~ 50 kpc, but the total amount of gas in this region of the CGM is not strongly affected by feedback. In models which include some minimal amount of feedback, we find sufficient metals in our simulated CGM to reproduce observations, but that our gas falls in the wrong phase. This is clearest in the case of O VI, where the observations generally do not match any of our models because of this mismatched phase (i.e. temperature).

- Turning thermal feedback up to very high values (even without cooling suppression) is successful at producing a galaxy with a flat rotation curve; however, it still fails to match the observational datasets from QSO absorption line studies (notably O

VI) demonstrating that QSO observations provide additional constraints on subgrid models beyond traditional metrics of success.

- There is significant redshift evolution of the column density distribution for our atomic species. In general, the column densities for most species drop by almost an order of magnitude across all radii in evolving from $z \sim 3$ to $z = 0$.

Generally, we find that QSO observations place tight constraints on feedback models. Perhaps not surprisingly, the CGM which acts as both stellar fuel and the medium into which supernovae explode, contains a great deal of information regarding the specific nature of star formation and stellar feedback. Probing this medium through QSO absorption line studies can tap into this information and reveal clear ways of differentiating between subgrid models. Significantly sampled observations in any of H I, Si II, Si III, Si IV, or C IV identifying $r_{\rm inner}$, the transition radius where column densities plummet in the CGM for these low-ionization ions, should potentially favor a specific type and intensity of stellar feedback. Furthermore, even subgrid models which can meet other criteria for success (e.g. flat rotation curve, matching the Kennicutt-Schmidt Law) need not match QSO absorption line statistics, as is demonstrated in the case of our HF run producing a flat rotation curve but failing to reproduce the the Mg II, H I and O VI observational datasets. Thus, comparisons against QSO absorption lines place a constraint on galaxy simulations orthogonal to standard comparison modes.

Future work on this topic includes addressing some of the limitations considered in Section 4.4.2. As was discussed, it appears that we lack multiphase substructure

in the halos of our galaxies, which is likely responsible for our inability to produce sufficient column densities of certain atomic species required to match observations (e.g. H I, Mg II, O VI). It may be that the subgrid models of thermal feedback and cooling suppression which we employ result in hot outflows of material suppressing the creation of a multiphase medium in the CGM. By adopting a different subgrid model, such as one utilizing radiation pressure or kinetic feedback to parameterize energy ejection, we may cause bulk outflows of material in different phases, potentially more conducive to producing a multiphase CGM. Additional steps to ensure higher resolution in the halo may also result in smaller clumps forming there; however, one potentially must add new refinement criteria since a density-based criterion will fail to substantially increase spatial resolution in the low-density halo.

Chapter 5

Conclusions

The comparison studies in the previous chapters have shown that significant disagreements exist between observations and simulations of galaxies, indicating that our numerical models lack adequate treatment for certain physical processes. While the basic physics governing the structure and bulk characteristics of galaxies appear to be in place, details of the distribution of matter and its phase are far from accurate. Perhaps this is not surprising, given the resolutions we are able to achieve in our simulations and assuming simple recipes for behavior on scales below tens or even hundreds of parsecs. That said, the comparison methods implemented here will remain as well-defined constraints on future studies of galaxy evolution.

Based on the results of this dissertation, I believe the direction for forward progress in the field of numerical galaxy evolution lies in better hydrodynamical algorithms for increasing resolution selectively, so as to minimize the need for subgrid models, as well as

more realistic prescriptions for the treatment of stellar feedback. Our experiments indicate that the thermal feedback model, even with the aid of unphysical cooling suppression, is incapable of producing galaxies consistent with reality. Tuning the feedback parameters can produce galaxies able to meet some observational parameters such as a flat rotation curve or the Kennicutt-Schmidt relation, but ultimately these models fail to match the tightest observational constraint to date, the circumgalactic medium.

Future success in this field rests on employing kinetic, radiative, and cosmic ray prescriptions for energy injection following star formation. We know that all of these mechanisms are at work in the universe. Young, massive stars radiate copious amount of energy into their environments in a short time, heating gas and driving massive outflows by exerting radiation pressure. Supernovae blast waves sweep up large quantities of material as they travel outward from their hosts generating winds. Cosmic rays are created as supernovae shocks entrain and accelerate ions outward, depositing significant energy into this less-understood medium. In all of these cases, energy injection drives material outward and injects energy into a medium which cannot efficiently lose that energy (e.g. through radiative losses as in the thermal model). This may potentially reduce the frequent creation of dense knots of material, a manifestation of the angular momentum problem. Furthermore, pressure exerted on all phases of gas will likely encourage the multiphase substructure we see in the galaxy today, forcing simulated circumgalactic medium measurements more in line with observations.

Work has already begun in implementing kinetic feedback in the enzo code with

collaborators Greg Bryan and Jeremiah Ostriker. We have devised an efficient algorithm for its use, avoiding some of the pitfalls to which other momentum-driven feedback models have succumbed. By distributing the kinetic energy deposition over many gas elements, we will not have to resort to unphysical behaviors like decoupling "kicked" gas from the local medium. It is my prediction that the incorporation of a realistic energy injection mechanism in stellar feedback subgrid models, such as this kinetic model, will yield simulated galaxies which meet the observational tests provided in this dissertation, potentially opening up new problems in understanding the nature of star formation and the resultant stellar and galactic populations.

Bibliography

Abadi, M. G., Navarro, J. F., Steinmetz, M., & Eke, V. R. 2003, The Astrophysical Journal, 591, 499

Abel, T., Anninos, P., Zhang, Y., & Norman, M. L. 1997, New Astronomy, 2, 181

Agertz, O., Moore, B., Stadel, J., Potter, D., Miniati, F., Read, J., Mayer, L., Gawryszczak, A., Kravtsov, A., Nordlund, Å., Pearce, F., Quilis, V., Rudd, D., Springel, V., Stone, J., Tasker, E., Teyssier, R., Wadsley, J., & Walder, R. 2007, Monthly Notices of the Royal Astronomical Society, 380, 963

Agertz, O., Teyssier, R., & Moore, B. 2011, Monthly Notices of the Royal Astronomical Society, 410, 1391

Alongi, M., Bertelli, G., Bressan, A., Chiosi, C., Fagotto, F., Greggio, L., & Nasi, E. 1993, Astronomy and Astrophysics Supplement Series (ISSN 0365-0138), 97, 851

Altay, G., Theuns, T., Schaye, J., Crighton, N. H. M., & Dalla Vecchia, C. 2011, The Astrophysical Journal Letters, 737, L37

Anninos, P., Zhang, Y., Abel, T., & Norman, M. L. 1997, New Astronomy, 2, 209

Avila-Reese, V., Colín, P., González-Samaniego, A., Valenzuela, O., Firmani, C., Velázquez, H., & Ceverino, D. 2011, The Astrophysical Journal, 736, 134

Balogh, M. L., Pearce, F. R., Bower, R. G., & Kay, S. T. 2001, Monthly Notices of the Royal Astronomical Society, 326, 1228

Barnes, J. & Hut, P. 1986, Nature (ISSN 0028-0836), 324, 446

Barnes, L. A., Haehnelt, M. G., Tescari, E., & Viel, M. 2011, Monthly Notices of the Royal Astronomical Society, 416, 1723

Baum, W. A. 1959, Publications of the Astronomy Society of the Pacific, 71, 106

Behroozi, P. S., Conroy, C., & Wechsler, R. H. 2010, The Astrophysical Journal, 717, 379

Bertin, E. & Arnouts, S. 1996, Astronomy and Astrophysics Supplement, 117, 393

Bertschinger, E. 1998, Annu. Rev. Astro. Astrophys., 36, 599

Bigiel, F., Leroy, A., Walter, F., Brinks, E., de Blok, W. J. G., Madore, B., & Thornley, M. D. 2008, The Astronomical Journal, 136, 2846

Binney, J. 1977, Astrophysical Journal, 215, 483

Binney, J. & Merrifield, M. 1998, Galactic Astronomy, Princeton Series in Astrophysics (Princeton, NJ: Princeton University Press)

Binney, J., Nipoti, C., & Fraternali, F. 2009, Monthly Notices of the Royal Astronomical Society, 397, 1804

Birnboim, Y. & Dekel, A. 2003, Monthly Notices of the Royal Astronomical Society, 345, 349

Blanton, M. R., Hogg, D. W., Bahcall, N. A., Baldry, I. K., Brinkmann, J., Csabai, I., Eisenstein, D. J., Fukugita, M., Gunn, J. E., Ivezic, Z., Lamb, D. Q., Lupton, R. H., Loveday, J., Munn, J. A., Nichol, R. C., Okamura, S., Schlegel, D. J., Shimasaku, K., Strauss, M. A., Vogeley, M. S., & Weinberg, D. H. 2002, arXiv, astro-ph, 186

Bohlin, R. C., Savage, B. D., & Drake, J. F. 1978, Astrophysical Journal, 224, 132

Booth, C. M. & Schaye, J. 2009, Monthly Notices of the Royal Astronomical Society, 398, 53

Bressan, A., Fagotto, F., Bertelli, G., & Chiosi, C. 1993, Astronomy and Astrophysics Supplement Series (ISSN 0365-0138), 100, 647

Brook, C. B., Governato, F., Roskar, R., Stinson, G., Brooks, A. M., Wadsley, J., Quinn, T., Gibson, B. K., Snaith, O., Pilkington, K., House, E., & Pontzen, A. 2011, Monthly Notices of the Royal Astronomical Society, 415, 1051

Bruzual, G. & Charlot, S. 1993, Astrophysical Journal, 405, 538

—. 2003, Monthly Notices of the Royal Astronomical Society, 344, 1000

Bryan, G. L. & Norman, M. L. 1997, arXiv, astro-ph, 363

Cen, R. 2011, The Astrophysical Journal, 741, 99

Cen, R. & Fang, T. 2006, The Astrophysical Journal, 650, 573

Cen, R. & Ostriker, J. P. 1992, The Astrophysical Journal, 399, L113

Ceverino, D. & Klypin, A. 2009, The Astrophysical Journal, 695, 292

Chabrier, G. 2003, The Astrophysical Journal, 115, 763

Charlot, S. & Bruzual, G. 1991, Astrophysical Journal, 367, 126

Charlot, S. & Fall, S. M. 2000, The Astrophysical Journal, 539, 718

Chen, H.-W., Helsby, J. E., Gauthier, J.-R., Shectman, S. A., Thompson, I. B., & Tinker, J. L. 2010, The Astrophysical Journal, 714, 1521

Chen, H.-W., Lanzetta, K. M., Webb, J. K., & Barcons, X. 2001, The Astrophysical Journal, 559, 654

Chen, H.-W. & Tinker, J. L. 2008, The Astrophysical Journal, 687, 745

Colín, P., Avila-Reese, V., Vázquez-Semadeni, E., Valenzuela, O., & Ceverino, D. 2010, The Astrophysical Journal, 713, 535

Courteau, S. 1997, The Astronomical Journal, 114, 2402

Danforth, C. W. & Shull, J. M. 2008, The Astrophysical Journal, 679, 194

Davé, R., Finlator, K., & Oppenheimer, B. D. 2012, Monthly Notices of the Royal Astronomical Society, 421, 98

Diemand, J., Kuhlen, M., & Madau, P. 2006, arXiv, astro-ph, 262

Djorgovski, S. & Davis, M. 1987, Astrophysical Journal, 313, 59

D'onghia, E., Burkert, A., Murante, G., & Khochfar, S. 2006, Monthly Notices of the Royal Astronomical Society, 372, 1525

Dopita, M. A. & Sutherland, R. S. 2003, Astrophysics of the Diffuse Universe, Astronomy and Astrophysics Library (Berlin, New York: Springer)

Draine, B. T. 2011, Physics of the Interstellar and Intergalactic Medium by Bruce T. Draine. Princeton University Press, 2011. ISBN: 978-0-691-12214-4, -1

Dressler, A. 1980, Astrophysical Journal, 236, 351

Efstathiou, G., Davis, M., White, S. D. M., & Frenk, C. S. 1985, Astrophysical Journal Supplement Series (ISSN 0067-0049), 57, 241

Efstathiou, G. & Eastwood, J. W. 1981, Monthly Notices of the Royal Astronomical Society, 194, 503

Eisenstein, D. J. & Hut, P. 1998, The Astrophysical Journal, 498, 137

Faber, S. M. & Jackson, R. E. 1976, Astrophysical Journal, 204, 668

Fagotto, F., Bressan, A., Bertelli, G., & Chiosi, C. 1994a, Astronomy and Astrophysics Suppl. 104, 104, 365

—. 1994b, Astronomy and Astrophysics Suppl. 105, 105, 29

Fall, S. M. & Efstathiou, G. 1980, Monthly Notices of the Royal Astronomical Society, 193, 189

Faucher-Giguère, C.-A. & Kereš, D. 2011, Monthly Notices of the Royal Astronomical Society: Letters, 412, L118

Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2009, The Astrophysical Journal, 703, 1416

Ferland, G. J., Korista, K. T., Verner, D. A., Ferguson, J. W., Kingdon, J. B., & Verner, E. M. 1998, The Publications of the Astronomical Society of the Pacific, 110, 761

Fernández, X., Joung, M. R., & Putman, M. E. 2012, The Astrophysical Journal, 749, 181

Fryxell, B., Olson, K., Ricker, P., Timmes, F. X., Zingale, M., Lamb, D. Q., MacNeice, P., Rosner, R., Truran, J. W., & Tufo, H. 2000, ASTROPHYS J SUPPL S, 131, 273

Fumagalli, M., Prochaska, J. X., Kasen, D., Dekel, A., Ceverino, D., & Primack, J. R. 2011, Monthly Notices of the Royal Astronomical Society, 418, 1796

Gerritsen, J. 1997, Ph.D. thesis, Groningen University, the Netherlands, (1997)

Girardi, L., Bressan, A., Chiosi, C., Bertelli, G., & Nasi, E. 1996, Astronomy and Astrophysics Supplement, 117, 113

Gnedin, N. Y. 1995, Astrophysical Journal Supplement Series (ISSN 0067-0049), 97, 231

Gnedin, N. Y. & Kravtsov, A. V. 2011, The Astrophysical Journal, 728, 88

Goerdt, T., Dekel, A., Sternberg, A., Gnat, O., & Ceverino, D. 2012, arXiv, astro-ph.CO

Governato, F., Mayer, L., Wadsley, J., Gardner, J. P., Willman, B., Hayashi, E., Quinn, T., Stadel, J., & Lake, G. 2004, The Astrophysical Journal, 607, 688

Governato, F., Willman, B., Mayer, L., Brooks, A., Stinson, G., Valenzuela, O., Wadsley, J., & Quinn, T. 2007, Monthly Notices of the Royal Astronomical Society, 374, 1479

Governato, F., Zolotov, A., Pontzen, A., Christensen, C., Oh, S. H., Brooks, A. M., Quinn, T., Shen, S., & Wadsley, J. 2012, Monthly Notices of the Royal Astronomical Society, 422, 1231

Guedes, J., Callegari, S., Madau, P., & Mayer, L. 2011, The Astrophysical Journal, 742, 76

Haardt, F. & Madau, P. 1996, The Astrophysical Journal, 461, 20

—. 2012, The Astrophysical Journal, 746, 125

Holmberg, E. 1941, Astrophysical Journal, 94, 385

Hopkins, P. F., Quataert, E., & Murray, N. 2011, Monthly Notices of the Royal Astronomical Society, 417, 950

—. 2012, Monthly Notices of the Royal Astronomical Society, 421, 3522

Hubble, E. P. 1929, Astrophysical Journal, 69, 103

Hummels, C. B. & Bryan, G. L. 2012, The Astrophysical Journal, 749, 140

Joung, M. K. R. & Mac Low, M.-M. 2006, The Astrophysical Journal, 653, 1266

Joung, M. R., Bryan, G. L., & Putman, M. E. 2012a, The Astrophysical Journal, 745, 148

Joung, M. R., Cen, R., & Bryan, G. L. 2009a, The Astrophysical Journal Letters, 692, L1

Joung, M. R., Mac Low, M.-M., & Bryan, G. L. 2009b, The Astrophysical Journal, 704, 137

Joung, M. R., Putman, M. E., Bryan, G. L., Fernández, X., & Peek, J. 2012b, arXiv, astro-ph.CO

Jubelgas, M., Springel, V., Enßlin, T., & Pfrommer, C. 2008, Astronomy and Astrophysics, 481, 33

Kacprzak, G. G., Churchill, C. W., Ceverino, D., Steidel, C. C., Klypin, A., & Murphy, M. T. 2010, The Astrophysical Journal, 711, 533

Katz, N. 1992, The Astrophysical Journal, 391, 502

Katz, N., Weinberg, D. H., & Hernquist, L. 1996, The Astrophysical Journal Supplement, 105, 19

Kaufmann, T., Mayer, L., Wadsley, J., Stadel, J., & Moore, B. 2007, Monthly Notices of the Royal Astronomical Society, 375, 53

Kennicutt, R. C. J. 1989, Astrophysical Journal, 344, 685

Kent, S. M., Dame, T. M., & Fazio, G. 1991, Astrophysical Journal, 378, 131

Kereš, D., Katz, N., Fardal, M., Davé, R., & Weinberg, D. H. 2009, Monthly Notices of the Royal Astronomical Society, 395, 160

Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, Monthly Notices of the Royal Astronomical Society, 363, 2

Kim, J.-h., Wise, J. H., Alvarez, M. A., & Abel, T. 2011, The Astrophysical Journal, 738, 54

Kimm, T., Slyz, A., Devriendt, J., & Pichon, C. 2011, Monthly Notices of the Royal Astronomical Society: Letters, 413, L51

Komatsu, E., Dunkley, J., Nolta, M. R., Bennett, C. L., Gold, B., Hinshaw, G., Jarosik, N., Larson, D., Limon, M., Page, L., Spergel, D. N., Halpern, M., Hill, R. S., Kogut, A., Meyer, S. S., Tucker, G. S., Weiland, J. L., Wollack, E., & Wright, E. L. 2009, The Astrophysical Journal Supplement, 180, 330

Kravtsov, A. V., Klypin, A. A., & Khokhlov, A. M. 1997, The Astrophysical Journal Supplement, 111, 73

Krumholz, M. R. & Tan, J. C. 2007, The Astrophysical Journal, 654, 304

Larson, R. B. 1969, Monthly Notices of the Royal Astronomical Society, 145, 405

Lucy, L. B. 1977, Astronomical Journal, 82, 1013

Machacek, M. E., Bryan, G. L., & Abel, T. 2001, The Astrophysical Journal, 548, 509

Mandelbaum, R., Seljak, U., Kauffmann, G., Hirata, C. M., & Brinkmann, J. 2006, Monthly Notices of the Royal Astronomical Society, 368, 715

Mayer, L. 2004, arXiv, 11476

McKee, C. F. & Ostriker, J. P. 1977, Astrophysical Journal, 218, 148

Miller, R. H. & Prendergast, K. H. 1968, Astrophysical Journal, 151, 699

Miniati, F. 2001, Computer Physics Communications, 141, 17

Mo, H. J., Mao, S., & White, S. D. M. 1998, Monthly Notices of the Royal Astronomical Society, 295, 319

Moster, B. P., Somerville, R. S., Maulbetsch, C., van den Bosch, F. C., Macciò, A. V., Naab, T., & Oser, L. 2010, The Astrophysical Journal, 710, 903

Murray, N., Ménard, B., & Thompson, T. A. 2011, The Astrophysical Journal, 735, 66

Navarro, J. F. & Benz, W. 1991, The Astrophysical Journal, 380, 320

Navarro, J. F. & Steinmetz, M. 1997, The Astrophysical Journal, 478, 13

Navarro, J. F. & White, S. D. M. 1994, Monthly Notices of the Royal Astronomical Society, 267, 401

Norman, M. L. & Bryan, G. L. 1999, in Numerical Astrophysics : Proceedings of the International Conference on Numerical Astrophysics 1998 (NAP98), 19

O'Connell, R. W. 1999, arXiv, astro-ph, 603

Ocvirk, P., Pichon, C., & Teyssier, R. 2008, Monthly Notices of the Royal Astronomical Society, 390, 1326

Oemler, A. J. 1974, Astrophysical Journal, 194, 1

Okamoto, T., Eke, V. R., Frenk, C. S., & Jenkins, A. 2005, Monthly Notices of the Royal Astronomical Society, 363, 1299

Oppenheimer, B. D. & Davé, R. 2006, Monthly Notices of the Royal Astronomical Society, 373, 1265

—. 2008, Monthly Notices of the Royal Astronomical Society, 387, 577

—. 2009, Monthly Notices of the Royal Astronomical Society, 395, 1875

O'Shea, B. W., Bryan, G., Bordner, J., Norman, M. L., Abel, T., Harkness, R., & Kritsuk, A. 2004, arXiv, astro-ph

O'Shea, B. W., Nagamine, K., Springel, V., Hernquist, L., & Norman, M. L. 2005, ASTROPHYS J SUPPL S, 160, 1

Osterbrock, D. & Ferland, G. 2006, Astrophysics of Gaseous Nebulae and Active Galactic Nuclei, 2nd edn. (Sausolito, CA: University Science Books)

Penton, S. V., Stocke, J. T., & Shull, J. M. 2004, The Astrophysical Journal, 152, 29

Piontek, F. & Steinmetz, M. 2011, Monthly Notices of the Royal Astronomical Society, 410, 2625

Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1992, Numerical Recipes in FORTRAN: The Art of Scientific Computing, 2nd edn. (Cambridge, England: Cambridge University Press)

Press, W. H. & Schechter, P. 1974, Astrophysical Journal, 187, 425

Prochaska, J. X., Weiner, B., Chen, H.-W., Mulchaey, J., & Cooksey, K. 2011, The Astrophysical Journal, 740, 91

Putman, M., Peek, J., & Joung, M. 2012, Annual Reviews in Astrophysics, 1

Ribaudo, J., Lehner, N., & Howk, J. C. 2011, The Astrophysical Journal, 736, 42

Robertson, B., Yoshida, N., Springel, V., & Hernquist, L. 2004, The Astrophysical Journal, 606, 32

Robertson, B. E. & Kravtsov, A. V. 2008, The Astrophysical Journal, 680, 1083

Robitaille, T. P. & Whitney, B. A. 2010, The Astrophysical Journal Letters, 710, L11

Rubin, V. C., Burstein, D., Ford, W. K. J., & Thonnard, N. 1985, The Astrophysical Journal, 289, 81

Sarazin, C. L. & White, R. E. I. 1987, Astrophysical Journal, 320, 32

Scannapieco, C., Tissera, P. B., White, S. D. M., & Springel, V. 2006, Monthly Notices of the Royal Astronomical Society, 371, 1125

Scannapieco, C., Wadepuhl, M., Parry, O. H., Jenkins, A., Springel, V., Teyssier, R., Carlson, E., Couchman, H. M. P., Crain, R. A., Frenk, C. S., Kobayashi, C., Monaco, P., Okamoto, T., Quinn, T., Stinson, G. S., Theuns, T., Wadsley, J., White, S. D. M., & Woods, R. 2011, eprint arXiv:1112.0315

Schaye, J. & Dalla Vecchia, C. 2008, Monthly Notices of the Royal Astronomical Society, 383, 1210

Schaye, J., Dalla Vecchia, C., Booth, C. M., Wiersma, R. P. C., Theuns, T., Haas, M. R., Bertone, S., Duffy, A. R., McCarthy, I. G., & Van De Voort, F. 2010, Monthly Notices of the Royal Astronomical Society, 402, 1536

Schultz, G. V. & Wiemer, W. 1975, Astronomy and Astrophysics, 43, 133

Sedov, L. 1959, Similarity and Dimensional Methods in Mechanics, New York: Academic Press, 1959, -1,

Sharma, M. & Nath, B. B. 2012, The Astrophysical Journal, 750, 55

Shen, S., Madau, P., Guedes, J., Mayer, L., & Prochaska, J. X. 2012, arXiv, astro-ph.CO

Simcoe, R. A., Sargent, W. L. W., & Rauch, M. 2004, The Astrophysical Journal, 606, 92

Slipher, V. M. 1913, Lowell Observatory Bulletin, 2, 56

Smith, B., Sigurdsson, S., & Abel, T. 2008, Monthly Notices of the Royal Astronomical Society, 385, 1443

Smith, B. D., Hallman, E. J., Shull, J. M., & O'Shea, B. W. 2011, The Astrophysical Journal, 731, 6

Sneden, C., Gehrz, R. D., Hackwell, J. A., York, D. G., & Snow, T. P. 1978, Astrophysical Journal, 223, 168

Sommer-Larsen, J., Götz, M., & Portinari, L. 2003, The Astrophysical Journal, 596, 47

Springel, V. 2010, Monthly Notices of the Royal Astronomical Society, 401, 791

Springel, V. & Hernquist, L. 2003a, Monthly Notice of the Royal Astronomical Society, 339, 289

—. 2003b, Monthly Notice of the Royal Astronomical Society, 339, 312

Springel, V., Wang, J., Vogelsberger, M., Ludlow, A., Jenkins, A., Helmi, A., Navarro, J. F., Frenk, C. S., & White, S. D. M. 2008, Monthly Notices of the Royal Astronomical Society, 391, 1685

Springel, V., White, S. D. M., Jenkins, A., Frenk, C. S., Yoshida, N., Gao, L., Navarro, J., Thacker, R., Croton, D., Helly, J., Peacock, J. A., Cole, S., Thomas, P., Couchman, H., Evrard, A., Colberg, J., & Pearce, F. 2005, arXiv, astro-ph, 629

Springel, V., Yoshida, N., & White, S. D. M. 2001, New Astronomy, 6, 79

Steidel, C. C., Erb, D. K., Shapley, A. E., Pettini, M., Reddy, N., Bogosavljević, M., Rudie, G. C., & Rakic, O. 2010, The Astrophysical Journal, 717, 289

Steinmetz, M. & Muller, E. 1995, Monthly Notices of the Royal Astronomical Society, 276, 549

Steinmetz, M. & Navarro, J. F. 1999, The Astrophysical Journal, 513, 555

—. 2002, New Astronomy, 7, 155

Stewart, K. R., Kaufmann, T., Bullock, J. S., Barton, E. J., Maller, A. H., Diemand, J., & Wadsley, J. 2011, The Astrophysical Journal Letters, 735, L1

Stinson, G., Brook, C., Prochaska, J. X., Hennawi, J., Pontzen, A., Shen, S., Wadsley, J., Couchman, H., Quinn, T., Macciò, A. V., & Gibson, B. K. 2011, arXiv, astro-ph.CO

Stinson, G., Seth, A., Katz, N., Wadsley, J., Governato, F., & Quinn, T. 2006, Monthly Notices of the Royal Astronomical Society, 373, 1074

Stone, J. M. & Norman, M. L. 1992, The Astrophysical Journal Supplement, 80, 753

Strateva, I., Ivezić, Ž., Knapp, G. R., Narayanan, V. K., Strauss, M. A., Gunn, J. E., Lupton, R. H., Schlegel, D., Bahcall, N. A., Brinkmann, J., Brunner, R. J., Budavári, T., Csabai, I., Castander, F. J., Doi, M., Fukugita, M., Győry, Z., Hamabe, M., Hennessy, G., Ichikawa, T., Kunszt, P. Z., Lamb, D. Q., McKay, T. A., Okamura, S., Racusin, J., Sekiguchi, M., Schneider, D. P., Shimasaku, K., & York, D. 2001, The Astrophysical Journal, 122, 1861

Tasker, E. J. & Bryan, G. L. 2006, The Astrophysical Journal, 641, 878

Taylor, G. 1950, Royal Society of London Proceedings Series A, 201, 159

Tescari, E., Viel, M., D'Odorico, V., Cristiani, S., Calura, F., Borgani, S., & Tornatore, L. 2011, Monthly Notices of the Royal Astronomical Society, 411, 826

Teyssier, R. 2002, Astronomy and Astrophysics, 385, 337

Teyssier, R., Chapon, D., & Bournaud, F. 2010, The Astrophysical Journal Letters, 720, L149

Thacker, R. J. & Couchman, H. M. P. 2000, The Astrophysical Journal, 545, 728

Thom, C. & Chen, H.-W. 2008, The Astrophysical Journal, 683, 22

Thom, C., Werk, J. K., Tumlinson, J., Prochaska, J. X., Meiring, J. D., Tripp, T. M., & Sembach, K. R. 2011, The Astrophysical Journal, 736, 1

Toomre, A. & Toomre, J. 1972, Astrophysical Journal, 178, 623

Tripp, T. M., Sembach, K. R., Bowen, D. V., Savage, B. D., Jenkins, E. B., Lehner, N., & Richter, P. 2008, ASTROPHYS J SUPPL S, 177, 39

Truelove, J. K., Klein, R. I., McKee, C. F., Holliman, J. H., Howell, L. H., & Greenough, J. A. 1997, The Astrophysical Journal Letters, 489, L179

Tully, R. B. & Fisher, J. R. 1977, Astronomy and Astrophysics, 54, 661

Tumlinson, J., Thom, C., Werk, J. K., Prochaska, J. X., Tripp, T. M., Weinberg, D. H., Peeples, M. S., O'Meara, J. M., Oppenheimer, B. D., Meiring, J. D., Katz, N. S., Davé, R., Ford, A. B., & Sembach, K. R. 2011, arXiv, astro-ph.CO

Turk, M. J., Smith, B. D., Oishi, J. S., Skory, S., Skillman, S. W., Abel, T., & Norman, M. L. 2011, The Astrophysical Journal Supplement, 192, 9

Van De Voort, F. & Schaye, J. 2011, eprint arXiv:1111.5039

Verner, D. A., Verner, E. M., & Ferland, G. J. 1996, Atomic Data and Nuclear Data Tables, 64, 1

von Hoerner, S. 1960, Z. Astrophys. 50, 50, 184

Wadsley, J., Stadel, J., & Quinn, T. 2003, arXiv, astro-ph

Weil, M. L., Eke, V. R., & Efstathiou, G. 1998, Monthly Notices of the Royal Astronomical Society, 300, 773

Wetzstein, M., Nelson, A. F., Naab, T., & Burkert, A. 2009, The Astrophysical Journal Supplement, 184, 298

Wise, J. H. & Abel, T. 2011, Monthly Notices of the Royal Astronomical Society, 414, 3458

Yao, Y., Nowak, M. A., Wang, Q. D., Schulz, N. S., & Canizares, C. R. 2008, The Astrophysical Journal, 672, L21

Yepes, G., Kates, R., Khokhlov, A., & Klypin, A. 1997, Monthly Notices of the Royal Astronomical Society, 284, 235

Zavala, J., Okamoto, T., & Frenk, C. S. 2008, Monthly Notices of the Royal Astronomical Society, 387, 364

Zel'dovich, Y. B. 1970, Astronomy and Astrophysics, 5, 84