

Feedback in Dwarf Galaxies

Sam Geen

Trinity College, University of Oxford



A thesis submitted for the degree of Doctor of Philosophy
at the University of Oxford.

Hilary Term 2012

To my mother, Josephine Geen

*Think when we talk of stars, that you see them
Printing their loud winds in the receiving sky;
For 'tis a machines' thoughts that now must make these things,
Carry them here and there; jumping o'er times,
Turning the accomplishment of many years
Into an hour-glass: for the which supply,
Admit me Chorus to this history;
Who prologue-like your humble patience pleases,
Gently to hear, kindly to judge, my thesis.*

Declaration

I declare that no part of this thesis has been accepted, or is currently being submitted, for any degree or diploma or certificate or any other qualification in this University or elsewhere. This thesis is the result of my own work unless otherwise stated. All the research was carried out in collaboration with my supervisor, Adrienne Slyz, and Julien Devriendt. The material in Chapter 5 is the result of collaboration with Jeremy Blaizot and Joakim Rosdahl at the Observatoire de Lyon, Saint-Genis-Laval, France. Results found in Chapter 4 have been published in The Monthly Notices of the Royal Astronomical Society. Some figures in chapter 3 have been reproduced, with kind permission, from Peirani et al. (2010).

Sam GEEN

(April 2012)

Acknowledgements

This thesis would not exist but for the encouragement and assistance of many people. I would thus like to thank some of those people. Any omission or error contained in these acknowledgements is completely unintended; I love you all.

I would first like to thank my supervisor Adrienne Slyz, with Julien Devriendt, for their help and encouragement throughout this DPhil, and for their continual tolerance for my inability to write a coherent sentence. I would also like to thank Jonathan Patterson for his patience and assistance as I continually crashed nodes and unwittingly wrote 1TB text output files on the supercomputers under his guardianship. I would like to thank Joakim Rosdahl for his help and advice regarding RAMSES-RT, and working with me on blowing bubbles. Thanks too goes to Jeremy Blaizot for his help and assistance in the final parts of this thesis, and for not actually firing me when I beat him at table football. Remerciements should also be directed to Romain Teyssier for the use of his code RAMSES, and for explaining how the hell the cooling module works on an early morning run around Saint-Genis-Laval. The proof-reading done by Adrienne Slyz, Julien Devriendt and Matt Geen was of vast help. I would like to also thank all the other staff of the University of Oxford Physics Department and Trinity College who made my stay at Oxford pleasant, well-administered and, God forbid, productive. To anyone in or outside of the department who offered advice, guidance or otherwise useful information that helped me to write this thesis: thank you.

I would like to thank my family for supporting me and encouraging me throughout my DPhil, in particular Matt and Emma Geen, for the strength they gave me in the more difficult moments, especially towards the end of this DPhil. Your love, conversations, posted cakes and fudge, financial support, long walks, allergens, transportation, postcards, birthday cards, music, stories, hugs and kindness have made all the difference. Thanks should also go to the students and post docs of Oxford Astrophysics, and especially

the denizens of “the Tower” who facilitated my caffeine habit, tolerated my presence in the various pubs around Oxford, fought for a working temperature above 17 degrees Celcius, saved my eyesight from the wrath of malevolent poultrygeists and were generally excellent friends. You are all too numerous and kind to risk omission through being named individually, so I shall just say this: you (yes, you!) are amazing. I would also like to acknowledge those who saw the start but not the end of this DPhil, and whose friendship I remember fondly: John Brown, Graham Dawes, Josephine Geen, Rita Hunter, Joseph Lawrence, Hugh Mowat, Vera Pickford, Steve Rawlings. Thank you.

This DPhil has been funded by a Science and Technology Facilities Council (STFC) studentship, and a Cosmocomp Fellowship. As well as these organisations, I would like to thank the Royal Astronomical Society for funding travel to present the work done during the course of my DPhil. The computational work discussed in this thesis was made possible by the Oxford Supercomputer Centre, the Centre Informatique National de l’Enseignement Supieur and a STFC DiRAC grant.

Abstract

Stellar feedback processes have been suggested as a mechanism for explaining various properties of galaxies, especially dwarf galaxies, which have weaker potentials and thus lower escape velocities for galactic winds. In this thesis, I present work done during my DPhil to better understand these processes. I begin by discussing the techniques used to simulate galaxies as collections of astrophysical fluids in a cosmological context, and present some methods for interpreting the results of such simulations. I then present two projects aimed at furthering our understanding of feedback in dwarf galaxies.

The first project is the investigation of a suite of simulations of satellites of a Milky Way-class halo. We discuss the formation of high-redshift dwarf galaxies and the effect that supernova feedback and reionisation have on the gas content and star formation history of these objects. We find that neither process has a dramatic effect on the star formation rates in high redshift dwarf galaxies that have already begun forming stars prior to reionisation. We do find, however, that the population of satellites is dramatically altered by the presence of cooled gas in the host halo, which increases the tidal stripping of satellites that pass close to the host. The second project concerns detailed simulations of a $15M_{\odot}$ star throughout its evolution, studying photoionisation, wind and supernova feedback from this star in various environments. Preliminary results are given for these simulations, which are compared to the results of previous authors.

¹ M_{\odot} is a solar mass, i.e. the mass of the Sun.

Contents

Declaration	vii
Acknowledgements	ix
Abstract	xi
1 Introduction	1
1.1 Cosmological Galaxy Formation	3
1.1.1 Missing Satellites, Dim Dwarfs	6
1.2 Thesis Summary	11
2 Numerical Methods with RAMSES	13
2.1 Numerical Hydrodynamics	13
2.1.1 Deriving the Fluid Equations	14
2.1.2 Discretising the Fluid	16
2.2 RAMSES	18
2.2.1 Adaptive Mesh Refinement	19
2.2.2 Solving the Poisson Equation	20
2.2.3 Eulerian Fluid Dynamics	23
2.2.4 Eulerian vs Lagrangian Techniques	25
2.3 Galaxy Formation Simulations	26
2.3.1 Cosmology & Initial Conditions	27
2.3.2 Radiative Cooling	27
2.3.3 Star Formation	31
2.3.4 Supernova Feedback and Reionisation	32
2.3.5 Radiative Transfer	34
2.4 Resimulation	35
2.4.1 Resimulating Dwarf Galaxies	37
2.4.2 Velocity Shift Method	39
2.4.3 Non-spherical resimulation region	39
2.4.4 Moving Region Refinement	40
2.4.5 Tracer Fluid	49
2.5 Instabilities, Odd-even Decoupling and Carbuncles	49
2.6 Shock-Induced Instabilities	61
2.7 Conclusions	64

3	Analysis & Visualisation	67
3.1	Halos, Merger Trees and Parallel Universes	67
3.1.1	Finding Halos with HALOMAKER	67
3.1.2	Merger Trees	70
3.1.3	Twinning Simulations	73
3.1.4	Twinning the Nut Suite	76
3.2	Visualisation and Data Handling	77
3.2.1	Simulation Data Handling	78
3.2.2	Alice: A Virtual Telescope	80
3.2.3	Case Study: Residual Star Formation in Mergers	83
4	Satellites of the Milky Way	91
4.1	Introduction	91
4.2	Simulating the Milky Way Satellites	92
4.2.1	Sugarplum Fairies: The Nut Suite	92
4.2.2	Mapping Halos in the Suite	95
4.2.3	Tracking high redshift galaxies down to $z = 0$	95
4.2.4	Resolution effects on galaxy formation	99
4.3	Comparison of Halo Properties	100
4.3.1	Tracking satellite galaxies down to $z = 0$	100
4.3.2	Dark Matter Satellite Halos	108
4.4	Feedback in Milky Way satellite progenitors	109
4.4.1	Supernova feedback	109
4.4.2	Reionisation Feedback	118
4.5	Stripping and Satellites at $z = 0$	119
4.5.1	The Effect of Baryonic Physics on Satellite Galaxy Survival	119
4.5.2	Properties of Satellite Galaxies at $z = 0$	129
4.6	Discussion and Conclusions	132
5	Single Supernovae and the ISM	137
5.1	Introduction	137
5.2	Model of a Massive Star	147
5.3	Numerical Simulations	148
5.4	Results	150
5.4.1	Wind-Free Supernovae	150
5.4.2	Stellar Winds and Supernovae	163
5.4.3	Stellar Photoionisation	176
5.5	Discussion and Future Work	180
6	Conclusions	183

Chapter 1

Introduction

One of the greatest surprises in 20th century cosmology was that the universe we observe – stars, planets, interstellar gas – makes up but a small fraction of the physical cosmos. Most of the mass in the universe does not interact with photons, and, upon discovery that the universe is accelerating in its expansion, most of the energy density is unaccounted for as well. But given these observations and their application to N-body simulations, astronomers are able to successfully replicate the general picture of galaxy formation that we see today. We can reproduce the number of Milky Way-like galaxies and their distribution in the universe, the formation of massive elliptical galaxies via collisions between other galaxies, using only these cosmological models of the universe to lay out the game board and a copy of Newton’s Principia as the rulebook.

Yet there remain unanswered questions. Solving one scientific mystery inevitably reveals more, skittering about in the darkness underneath. While purely gravitational models are highly successful, the observable universe is made up of stars, dust and gas – complex entities that obey the rules of fluid dynamics, chemistry and atomic physics. Various problems arise if one ignores these processes. The “missing satellite problem”, in which purely gravitational models of our galaxy over-predict the number of dwarf galaxies orbiting and spiralling into the Milky Way itself, is discussed at length in this thesis. Some theorists even invoke alternative cosmologies to explain discrepancy, giving dark matter a temperature or modifying gravity itself. Whether this is true, or whether we can explain these problems through the known interactions between gas particles alone, remains to be seen.

These interactions are not gentle things, either. The typical view of space as a cold, dark void is true if you compare the interstellar medium (hereafter ‘ISM’) to the air you are breathing in this room. But zoom out, and you find gas at a million degrees, winds travelling at thousands of kilometres per second, clouds cooling through collisions between their particles to tens of degrees above absolute zero or below, collapsing under gravity to form stars – and supernovae. The focus of this thesis, supernovae are phenomenally violent events, outshining entire galaxies for days, pumping vast quantities of energy into the galaxy, even blowing gas completely out of the galaxy as galactic winds. In some cases, this removes gas that might otherwise form stars. At the same time, supernovae can have a positive impact on star formation; the introduction of “metals” (elements heavier than helium) makes gas cooling more efficient, since they have more emission lines that can radiate away energy; improved cooling allows the gas to collapse faster, increasing the star formation rate. Blastwaves impacting gas clouds can either heat them, causing them to dissipate, or else compress them, taking them to the density required to form stars. Even the radiation from stars can affect galaxy formation; photoionisation by UV radiation heats gas, and could even prevent the formation of small haloes once the universe has been photoionised. Feedback cycles are set up, where gas forms stars which in turn affect the way gas forms the next generation of stars. Whether the feedback is positive and enhances star formation, or negative and regulates or suppresses star formation, is still unclear, and depends strongly on the properties of the galaxy in question.

The smallest galaxies, “dwarf” galaxies, are ideal test subjects for studying these processes. The weaker gravitational potential of these galaxies makes it easier for supernovae to create galactic winds, or for photoionisation to heat up the gas and prevent it from ever collapsing to form stars. The problem of feedback is a complex one, and requires vast computational power to study. Recently, computers have become powerful enough to allow us to probe down to parsec scales in simulations that form galaxies from cosmological initial conditions. This is important, because it allows us to form star-forming gas clouds without invoking “sub-grid” models for the gas in the ISM. Recent work has suggested that both star formation and supernova feedback are governed by small-scale processes rather than global galactic properties, and so high-resolution simulations of galaxy formation are central to our understanding of feedback as a complex hydrodynamic process. ? have successfully reproduced the star formation laws in galaxies by producing a relationship based on the amount of molecular gas in galaxies. The implication is that star formation rates are dominated

by small-scale properties of galaxies such as molecular cloud formation, rather than by global processes such as rotation or disk stability. Similarly, Slyz et al. (2005) noted that it is crucial to capture parsec-scale structure in the interstellar medium, such that a multiphase medium with hot, shocked supernova gas and cold, dense star-forming regions are present. While models at lower resolution are made that attempt to reproduce the observed effect of supernova feedback with “sub-grid” models (e.g. Springel & Hernquist (2003), who model supernova winds by injecting “wind” particles into their galaxies), the results of simulations of supernova feedback are still heavily implementation-dependent (Scannapieco et al., 2011). Complex, high-resolution hydrodynamic studies remain the only way to gain a proper theoretical understanding of the effect of supernovae on galaxy formation.

This thesis is concerned with studying feedback as a phenomenon affecting galaxies on a parsec scale or lower, both in isolated simulations of individual stars, and in full cosmological simulations including an array of physical mechanisms. In this way, I hope to shed some light on the precise effect of supernova feedback on galaxy formation, using dwarf galaxies as test subjects. In this chapter, I review the literature concerning dwarf galaxy formation, and the effect that supernovae have on the gas inside them. I then lay out the plan for this thesis, and the work I have done in attempting to resolve some of the outstanding issues concerning stellar feedback in dwarf galaxies.

1.1 Cosmological Galaxy Formation

The current paradigm of galaxy formation in a cosmological context is that of galaxies forming inside dark matter halos in a cold dark matter universe (White & Rees, 1978; Blumenthal et al., 1984) with a cosmological constant (Λ CDM) (Riess et al., 1998). The existence of matter that does not interact with light was posited by Zwicky (1933) as a way to explain the rotation curves of galaxies. Indeed, the latest measurements of the Cosmic Microwave Background (CMB) indicate that only $\sim 20\%$ of the mass in the universe, and only $\sim 2\%$ of the energy density of the universe today, is contributed by visible matter (Larson et al., 2010). The model of hierarchical galaxy formation proceeds as follows: dark matter fluctuations collapse under gravity, unimpeded by pressure from particles interacting via the electromagnetic force. This collapse forms 2D “sheets”, which collapse into filaments, which feed the growth of dark matter halos. Since visible matter and dark matter are coupled via gravity, baryons (i.e. atoms and nucleons) follow the dark

matter into these halos, where they are shock-heated and form gas halos. Collisions between the baryonic particles allows the gas to radiate and cool, forming cold, dense gas disks at the centre of these halos. Once the density reaches a certain threshold, stars are able to form.

There exists a bifurcation (called the “Hubble sequence”, or “Hubble tuning fork”) between disk galaxies and elliptical galaxies. Disk galaxies often exhibit spiral structures, which are attributed to density waves formed as the disk rotates (Lin & Shu, 1964), though a complete theory of spiral formation still eludes us. Stars often trace these spirals due to the higher concentration of dense, star-forming gas, though as we find in the NUT simulation discussed later in the thesis, this requires stable spiral structures in the disk gas, found more often at lower redshifts (i.e. closer to the present day). The disk also allows gas to condense further than it would if the same radial distribution of gas was distributed in a spherically-symmetric fashion. By contrast, we also find elliptical galaxies, which are usually more massive and exhibit much lower star formation rates than disk galaxies. A long-standing problem in galaxy formation was the explanation of the morphology (shape) of galaxies. (White & Rees, 1978) proposed that the two types of galaxy can be explained through hierarchical galaxy mergers, in which smaller galaxies collide and merge to form larger galaxies, which in turn collide and merge to form larger galaxies still (Kauffmann et al., 1993; Baugh et al., 1998). Major mergers (i.e. a merger in which the two progenitors are of roughly equal mass) will disrupt the existing structure of two colliding disk galaxies and creates the properties observed in elliptical galaxies (Holmberg, 1941; Toomre, 1977; Barnes & Hernquist, 1992). Meanwhile, the smooth accretion of matter onto a halo allows disk galaxies to form and grow, although minor merger events (in which one galaxy is much larger than the other) play an important role in the growth of spiral disk galaxies as well, since minor mergers will not significantly alter the structure of spirals (Steinmetz & Mueller, 1994; Navarro & White, 1994; Katz & Gunn, 1991). Dwarf elliptical galaxies also exist, although these tend not to have been formed by mergers, and are either too small to generate stable disks, or else have been stripped and moulded into spheroids by interactions with larger galaxies.

Local group satellite galaxies provide a useful catalogue of nearby dwarf galaxies that can be studied to build theoretical models of dwarf galaxy formation and evolution. In the hierarchical model of galaxy formation, smaller halos are captured by larger halos, and the galaxies they contain become satellite galaxies of the host galaxy until dynamical friction finally forces them to coalesce. The Milky Way and Andromeda

both have populations of satellite galaxies that are in the process of being disrupted and consumed by their hosts; the results of a recent survey are given in McConnachie (2012). In the process of merging, tidal and ram-pressure stripping disrupts these galaxies and their halos, producing tidal streams (Koposov et al., 2011) and dwarf spheroidals (Grebel et al., 2003). Tidal stripping is the effect of differential forces from a host's potential on a halo, causing stars and dark matter inside the galaxy to be pulled apart, forming long tidal streams (Toomre & Toomre, 1972). Ram-pressure stripping affects the gas content of galaxies and halos moving through clusters and host halos (Gunn & Gott, 1972); a clump of gas moving through a gaseous medium will be stripped as friction between the clump surface and the medium induces Kelvin-Helmholtz instabilities, disrupting the clump. Early attempts to reproduce the observed Milky Way satellite population using dark-matter-only simulations overproduced the number of low mass satellites by several orders of magnitude, if it is assumed that each simulated dark matter satellite halo contains a galaxy with a constant mass-to-light ratio (Moore et al., 1999). Current state-of-the-art dark matter simulations, such as the Aquarius (Springel et al., 2008) and the Via Lactea II (Diemand et al., 2008) reach the same conclusion.

Whilst the existing sample of Milky Way satellite galaxies is almost certainly incomplete (Koposov et al., 2008; Tollerud et al., 2008) and new satellite galaxies are continually being discovered (e.g. Belokurov et al. (2009, 2010)), it appears extremely unlikely that new observations will uncover galaxies populating every dark matter substructure predicted to exist around the Milky Way. As a result, various authors have attempted to explain this discrepancy by invoking physical mechanisms that reduce or prevent star formation in the majority of the smaller halos, making these dark matter substructures fainter or completely dark in the process. In particular, photoionisation and feedback from supernovae have been proposed as the most likely mechanisms to prevent gas from condensing to form stars, although other mechanisms, such as cosmic rays, have also been suggested (Wadepuhl & Springel, 2010a).

The suggestion that the Λ CDM model is incorrect has also been recently re-introduced as a solution to the missing satellite problem. Boylan-Kolchin et al. (2011); Lovell et al. (2012); Di Cintio et al. (2011) examine the relationship between the maximum circular velocity of the dwarf spheroidal satellites and the radius at which this velocity is found, and determine that the velocities found are higher than those found in either Λ CDM pure dark matter simulations or simulations with baryon physics using smoothed particle hydrodynamics (SPH). In fact, Di Cintio et al. (2011) determine that gas cooling makes the problem worse,

since the central density of the halo increases and the radius at which the maximum circular velocity is found decreases. Lovell et al. (2012) suggests that by using warm dark matter (WDM) instead of cold dark matter, the problem is resolved. They model dark matter particles as having pressure, hence slowing the collapse of halos and preventing the formation of halos below a certain mass threshold. However, Lovell et al. (2012) do not run simulations with both WDM and baryon physics, and thus the combined effect of baryon physics and WDM cannot be confirmed. In the following section, I detail some of the literature that discusses the effect that baryon physics and stellar feedback processes have on dwarf galaxy formation.

1.1.1 Missing Satellites, Dim Dwarfs

Since most of the mass in the universe is comprised of dark matter, dark matter-only simulations of the universe have enjoyed considerable success in explaining the structure of the universe and the properties of galaxies (Springel et al., 2005). However, if we assume that each dark matter halo hosts a galaxy of proportional mass, we find insufficient galaxy luminosity at both the low and high mass end of the luminosity function. For high mass galaxies, which are fainter than expected, active galactic nucleus (AGN) or quasar feedback has been invoked (Murray et al., 2005; Di Matteo et al., 2005; Silk & Rees, 1998). In this model, energy from infalling matter being accreted onto the central black hole is coupled to the gas in the rest of the galaxy, again removing it from the halo and suppressing star formation. A review by Veilleux et al. (2005) finds a wide array of both observed examples of and theoretical models for galactic winds, and the effect that stellar feedback and AGN have on star formation and galaxy evolution. This thesis, for now, ignores the effect of AGN, since it is not thought to contribute strongly to the properties of dwarf galaxies; Izotov & Thuan (2008) do find AGN within low-metallicity dwarfs, but they are extremely rare. Instead, we turn our attention to the feedback processes that are likely to have the biggest impact on dwarf galaxy formation.

The UV ionising background has been argued to be effective at halting or preventing star formation in low-mass halos in studies using analytic arguments, observations, N-body simulations and semi-analytic models (Bullock et al., 2000; Benson et al., 2002; Somerville, 2002; Kravtsov et al., 2004; Moore et al., 2006; Simon & Geha, 2007; Strigari et al., 2007; Madau et al., 2008; Macciò et al., 2010; Muñoz et al., 2009; Busha et al., 2010). Following the pioneering work of Quinn et al. (1996); Gnedin (2000b), hydrodynamic simulations of Milky Way-like galaxies have recently been employed to study the problem (Hoeft

et al., 2006; Libeskind et al., 2010; Nickerson et al., 2011; Okamoto et al., 2008; Okamoto & Frenk, 2009; Parry et al., 2011; Sawala et al., 2011; Scannapieco et al., 2011; Wadepuhl & Springel, 2010a). Many of these authors have found a best fit to the observed luminosity function by adopting instantaneous reionisation at $z \sim 11$, in agreement with the latest WMAP estimate ($z_{\text{reion}} = 10.5 \pm 1.2$ (Larson et al., 2010)). However, Hoefft et al. (2006); Wadepuhl & Springel (2010a); Guo et al. (2010) find reionisation to have limited effectiveness in completely suppressing star formation in low-mass halos that have already begun forming stars. Hoefft et al. (2006) determine that there is a characteristic mass of $6.5 \times 10^9 h^{-1} M_{\odot}$ below which halos become unable to retain baryons to $z=0$, and hence cannot form stars. Okamoto et al. (2008) find a similar result, corresponding to a maximum circular velocity v_{max} of 25km/s. Okamoto & Frenk (2009) refine this finding, stating that they find a sharp cut-off at around $v_{\text{max}} = 12\text{km/s}$ at reionisation ($z=9$ here) below which halos remain dark until $z=0$. To further obscure the picture, the extent of the epoch of reionisation itself is poorly constrained, with only a lower limit of $z \gtrsim 6$ on its completion provided by observations (e.g. Cen et al. (2009); Mesinger (2009)).

3.26 In isolated dwarf galaxies, Dekel & Silk (1986) proposed that the suppression of star formation by supernova feedback in dwarf galaxies embedded in dark matter halos could explain the observed scaling relations in luminosity, metallicity and radius; Dekel & Woo (2003) find that metallicity, surface brightness/radius and the “characteristic” velocity (rotation velocity or velocity dispersion) scale with stellar mass according to a power law, and attribute this to regulation by supernova feedback. This is corroborated by observational evidence of galactic winds in dwarf galaxies (e.g. Puche et al. (1992); Stewart et al. (2000); Schwartz & Martin (2004)). Benson et al. (2003) posit that supernovae can help explain the unexpectedly low dwarf galaxy luminosity function, along with photoionisation. By removing the gas from galaxies via the injection of thermal and kinetic energy into the interstellar medium (ISM), supernovae would reduce the number of stars formed inside dwarf halos by reducing the number of baryons available for star formation. The efficiency of transfer of kinetic energy from supernovae transferred to the ISM is typically $<3\%$ (Chevalier, 1977; Goodwin et al., 2000), or $\sim 10^{49}$ erg/supernova. Hydrodynamic simulations by authors such as Low & Ferrara (1999); Mashchenko et al. (2008); Ricotti et al. (2008); Ceverino & Klypin (2009) are able to convert this energy into massive supernova-driven galactic winds. However, it is still unclear as to precisely what effect supernovae have in turn on the star formation in low-mass galaxies. Tassis et al. (2008) claim that including supernova feedback

does not affect the properties of their simulated dwarf galaxies, despite the success of authors cited earlier in this chapter in reproducing observed dwarf galaxy properties by invoking stellar feedback. Instead, Tassis et al. (2008) attribute the scaling relations found in dwarf galaxies to low star formation efficiencies in weak potentials.

The contradictions between these authors find their origin in the different numerical recipes and spatial or mass resolutions adopted in the models, along with the variety of galaxy masses and merger histories. Scannapieco et al. (2011) presents a suite of simulations with the same initial conditions but run with different codes and techniques, and finds no consensus in the effect of supernova feedback on galaxies in hydrodynamic simulations. Indeed, supernovae are potentially able to drive either positive or negative feedback cycles. Outflows from supernovae can remove gas from the galaxy, preventing it from forming stars. However, they also release metals into the surrounding gas. Metal line cooling increases the efficiency of gas cooling and hence promotes the collapse of gas clouds into star-forming regions (Powell et al., 2011). Moreover, shock compression from stellar winds and supernovae has been observed to trigger the formation of star-forming regions (Assousa & Herbst, 1980; Oey et al., 2005).

As well as photoionisation and supernova feedback, stellar winds from massive stars are found to have an impact on the evolution of star-forming regions, as detailed in the seminal review by Weaver et al. (1977). This is because massive stars, which live only a few Myr, emit the strongest winds (Lejeune & Schaerer, 2001; Marigo & Girardi, 2007) closest to where they were formed, and can ionise and disrupt star-forming regions (Strömgren, 1939; Oort & Spitzer, 1955; Kahn, 1954), affecting the ability of these regions to form stars. We explore the issue of feedback from stellar winds and supernovae further in chapter 5. Evidence of all three types of stellar feedback are found in observations of local group dwarfs. The Large and Small Magellanic clouds (LMC and SMC respectively), the largest satellite galaxies of the Milky Way, are hosts to ongoing star formation. 30 Doradus in the LMC, for example, is the largest HII region in the local group (Brandl, 2005; Malumuth & Heap, 1994) and exhibits complex structures and high-velocity outflows (Chu & Kennicutt, 1994; Redman et al., 2003). On a galactic scale, the review by Veilleux et al. (2005) documents several instances of starburst events coinciding with galactic-scale winds, suggesting that the evolution of massive stars can indeed drive galactic-scale mass loss.

The problem of dwarf galaxy formation, and the importance of stellar feedback in producing the ob-

served properties of these galaxies, is thus still not well constrained. In order to better understand this problem, we are required to resolve the scales on which this feedback cycle operate, including star formation and the detailed hydrodynamics of the ISM. In this thesis I introduce some of the methods used to approach this problem from a numerical hydrodynamic perspective, and introduce two studies aimed at furthering our understanding of galaxy formation. The first study is designed to constrain satellite galaxy formation and evolution, and more specifically the role played by supernova feedback and reionisation in the process. To this end, we use the NUT suite of high resolution hydrodynamic cosmological simulations of a Milky Way-like galaxy (Powell et al., 2010; Kimm et al., 2011). At scales of 1-10pc (parsec, equal to 3.26 light years), these resolve large molecular clouds, and hence model the interstellar gas and stellar feedback in greater detail, which potentially affects star formation histories (Slyz et al., 2005). This approach differs from previous work investigating the Milky Way satellites using hydrodynamic simulations as these generally capture the ISM at lower resolution, and attempt to compensate for this by introducing analytic expressions to account for the multiphase ISM and outflows (Scannapieco et al., 2005, 2006; Murante et al., 2010). As it is currently too costly to simulate a Milky Way-sized halo and its substructures with parsec resolution and hydrodynamics throughout the lifetime of the Universe, most of our analysis is restricted to high redshift¹ ($z > 6$). However, observational studies conclude that a vast majority of satellite galaxies contain stars which formed prior to $z \sim 2$ and in many cases prior to even $z \sim 5$ (see recent review by Tolstoy et al. (2009) and also Kirby et al. (2011a)). Therefore, a high redshift study of these objects should be able to shed light on the problem, provided one is able to accurately predict their spatial distribution at $z = 0$.

This latter requirement is in itself a major challenge as it presumably requires hydrodynamics simulations which include (at least) radiative cooling. Indeed, since gas cooling can significantly increase the central density of dark matter halos (e.g. Blumenthal et al. (1984)), one expects physical processes like dynamical friction and tidal disruption of the satellites to be altered. The extent of these differences needs to be quantified because most of the studies mentioned earlier in this introduction rely on pure dark matter simulations to underpin analytic arguments or graft semi-analytic models of galaxy formation. Several groups have looked at differences between simulations of galactic halos containing baryons and their pure dark matter N-body counterparts. For instance Peirani (2010) found that identical simulations of a local

¹Redshift (or z) denotes the change in wavelength of light from distant galaxies reaching Earth. It can be used to determine the distance between a source of light and Earth, and hence (since we know the speed of light) the time since the light was emitted. It is used commonly as a measure of time in cosmology.

group-like volume with and without baryons matched well, but did not comment on satellites within halos. In their constrained simulations of the Local Group, Libeskind et al. (2010) find more satellite halos when baryons are included than in the identical pure dark matter run. Their radial distribution is also significantly more concentrated. By contrast, although they also follow a more concentrated radial distribution, satellites in the baryonic simulations of Romano-Díaz et al. (2010, 2009) survive for shorter times than their pure dark matter counterparts, which yields an overall lower number of satellites within R_{vir} when baryons are included. Schewtschenko & Macciò (2011) finds a similar result to Libeskind et al. (2010) in terms of number of halos but suggests that these results are not incompatible with those of Romano-Díaz et al. (2010). We re-examine this issue using Eulerian adaptive mesh refinement (AMR) grid hydrodynamics instead of Lagrangian smoothed particle hydrodynamics (SPH), and improving on both mass and force resolution for the dark matter.

The second study in this thesis involves a more detailed examination of the role of supernova feedback in the ISM. In this, I aim to revisit the work of Thornton et al. (1998), who perform 1D numerical simulations of supernova explosions in a uniform medium with various densities and metallicities. The goal of their work was to produce a set of empirical models for supernova remnants, which can be applied to lower-resolution galaxy simulations to create more realistic sub-grid models for supernova feedback. There are many limitations to this approach, however. In 1D, convection currents and other hydrodynamic instabilities are unable to form, though Gull (1973) attempts to quantify such processes using linear stability analysis. I review some of these processes in chapters 2 and 5. Secondly, non-radially-symmetric media such as power-law density fields or a turbulent ISM cannot be captured in 1D. Since the structure of the ISM can be highly non-linear, the use of 1D simulations can be limiting. Further, Thornton et al. (1998) do not consider physical processes such as stellar winds or radiative transfer. I introduce the preliminary results from a set of 3D simulations designed to rectify this. The first simulations are designed to reproduce the models of Thornton et al. (1998) in a 3D uniform field. I then introduce the first results of simulations including stellar winds and radiative transfer, and discuss the future work that will be done to introduce further physical complexity to these systems, in an effort to better constrain our understanding of the effect of supernova remnants on the evolution of the ISM in cosmological galaxy formation simulations.

1.2 Thesis Summary

This thesis proceeds as follows:

Chapter 2 I describe the algorithms used to describe astrophysical fluids in numerical simulations, including techniques for capturing an isolated dwarf galaxy moving through a cosmological volume. I then present a discussion of numerical instabilities in Riemann solvers and some possible solutions.

Chapter 3 I discuss methods for analysing and visualising raw simulation data, including halofinding and merger tree generation. I then present two visualisation tools written during this thesis; RAMTOVA-POR, a tool for allowing RAMSES adaptive mesh data to be read by VAPOR, a general-purpose fluid visualisation tool, and Alice, a specialised astrophysics visualisation tool. I finish by summarising the results of a study that uses Alice to predict the detectability of features of minor mergers with Hubble Space Telescope observations.

Chapter 4 I discuss the various feedback and gas dynamics processes thought to influence the survival of dwarf satellite galaxies to $z=0$ in the NUT suite, a set of cosmological hydrodynamic simulations of a Milky Way-class halo with up to 0.5pc resolution. I begin by determining the effect of supernova feedback and reionisation on the star formation in these halos. I then follow the merger tree of these objects to $z=0$ to determine which halos present at reionisation survive to $z=0$ when various physical processes are included in simulations. I end by comparing the satellites of our NUT galaxy to those of the Milky Way and M31. This chapter is based on material found in Geen et al. (2012).

Chapter 5 I present the first results of a set of fully-3D simulations of individual supernova explosions in media corresponding the various phases of the interstellar medium. I also discuss the effects of including stellar winds and radiative transfer in these models. I begin by introducing the model used to represent a $15M_{\odot}$ star throughout its lifetime. I then discuss and quantify the effect that environment, stellar winds and radiative transfer have on the evolution of supernova remnants. The aim of this work is to motivate future models of supernova feedback in galaxy formation simulations.

Chapter 6 I lay out the conclusions of this thesis, and discuss opportunities for future work on the subject

of feedback in dwarf galaxies using high-resolution hydrodynamic simulations.

Chapter 2

Numerical Methods with RAMSES

2.1 Numerical Hydrodynamics

A galaxy is a fluid made up of several components. The principal three with which we will concern ourselves in this thesis are dark matter, stars, and gas (while stars are made up of the same stuff as interstellar gas, we consider them as a separate phase for the purpose of our models). The first two can be thought of as collisionless fluids. Dark matter, at least the cold dark matter (CDM) with which we concern ourselves here, does not experience pressure, and interacts with itself and other fluids only through gravity. Stars are not collisionless, but their cross-section for interaction is so small that they rarely collide head-on, much less transmit any kind of energy through other stars by doing so. However, in order to create stars, we must simulate interstellar gas, which does experience pressure and viscosity via inter-particle collisions, even at densities found in the intergalactic medium (IGM).

Following the movement of every atom and molecule in a fluid is virtually impossible. A galaxy has $10^9 M_{\odot}$ or more baryons in it, as well as a large quantity of dark matter. Even following the motions of every one of the 10^{11} stars in our galaxy is beyond our current computational power. Thankfully, the mathematical model for treating this flow of particles as a continuous fluid is well established. In the following section I give a brief overview. For a more detailed treatment, readers are encouraged to consult Shu (1992), Toro et al. (1994) or, for a broader overview, Bodenheimer et al. (2007). In this section we discuss the equations of fluid mechanics when applied to this problem, and discuss the various discretisation techniques used to translate these equations into algorithms solvable by computer simulations.

2.1.1 Deriving the Fluid Equations

If we assume a fluid to be a collection of point particles (i.e. no rotations or vibrations) we can state that the number of particles with a given position and momentum is given by:

$$dN = f(\mathbf{x}, \mathbf{u}, t) d\mathbf{x} d\mathbf{u} \quad (2.1)$$

where \mathbf{x} and \mathbf{u} are the position of a given particle, and f is the *distribution function* at a given time t . One example of f is the Maxwellian distribution, in which collisions are elastic, and collisions involving more than two particles are assumed to be negligible.

Using the standard formula for deriving partial differential equations (PDEs), and noting that the rate of change in \mathbf{x} and \mathbf{u} with time are \mathbf{u} and \mathbf{F} (the force per unit mass) respectively, the rate of change in f is given by

$$\frac{\partial f}{\partial t} + u_i \frac{\partial f}{\partial x_i} + F_i \frac{\partial f}{\partial u_i} = \left[\frac{\partial f}{\partial t} \right]_{\text{coll}} \quad (2.2)$$

for $i=1,2,3$ (in 3D). The right-hand side of the equation accounts for changes to f by collisions; for collisionless fluids, this term will be (unsurprisingly) zero. This is known as *Boltzmann's transport equation* or simply “the Boltzmann equation”.

The density, bulk velocity and specific internal energy of the fluid at a position \mathbf{x} and time t are hence given by:

$$\rho(\mathbf{x}, t) \equiv \int m f(\mathbf{x}, \mathbf{u}, t) d\mathbf{u} \quad (2.3)$$

$$\mathbf{v}(\mathbf{x}, t) \equiv \frac{1}{\rho} \int \mathbf{u} m f(\mathbf{x}, \mathbf{u}, t) d\mathbf{u} \quad (2.4)$$

$$\varepsilon(\mathbf{x}, t) \equiv \frac{1}{\rho} \int m \tilde{u}^2 f(\mathbf{x}, \mathbf{u}, t) d\mathbf{u} \quad (2.5)$$

where \tilde{u} is the speed of a given particle with respect to the bulk velocity of the fluid and particles have a mass m . For a Maxwellian distribution:

$$f(\mathbf{x}, \mathbf{u}, t) d\mathbf{u} = n(\mathbf{x}, t) \left[\frac{m}{2\pi k_B T(\mathbf{x}, t)} \right]^{3/2} \exp \left[-\frac{m(\mathbf{u} - \mathbf{v})^2}{2k_B T(\mathbf{x}, t)} \right] d\mathbf{u} \quad (2.6)$$

we recover the usual formula for the specific energy of a monotonic gas $\varepsilon = \frac{3}{2} \frac{k_B}{m} T$ at temperature T , where

k_B is the Boltzmann constant.

The equations of hydrodynamics depend on expressions of continuity of mass, momentum and energy. These may be expressed as the 0th, 1st and 2nd order moments of the Boltzmann equation, i.e. multiplying equation 2.2 with \mathbf{u}^n and integrating over \mathbf{u} . Note that the total mass, momentum and energy of a fluid is conserved, and hence the collisional term on the right-hand side of the Boltzmann equation vanishes when integrated over particle velocity. Further, since the total mass, momentum and energy of a system must be finite, we can write $\lim_{u \rightarrow \infty} u^n f = 0$ for $n=0,1,2$. Hence, by integrating equations 2.3 to 2.5 by parts, we find

$$\int \frac{\partial f}{\partial u_i} d\mathbf{u} = 0 \quad (2.7)$$

$$\int u_j \frac{\partial f}{\partial u_i} d\mathbf{u} = -\delta_{ij} \frac{\rho}{m} \quad (2.8)$$

$$\int u^2 \frac{\partial f}{\partial u_i} d\mathbf{u} = -\frac{\rho}{m} v_i \quad (2.9)$$

where δ_{ij} is the Kronecker delta function, which is 1 when the indices are the same and 0 otherwise.

By integrating equation 2.2 with \mathbf{u} , multiplying by m , noting that equation 2.7 makes the third term disappear and noting that the collisional components disappear as stated above, we get

$$\frac{\partial}{\partial t} \int m f d\mathbf{u} + \frac{\partial}{\partial x_j} \int u_j m f d\mathbf{u} = 0 \quad (2.10)$$

Swapping with equations 2.3 and 2.4 we get:

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j} (\rho v_j) = 0 \quad (2.11)$$

which expresses mass conservation. Equations for momentum and energy conservation can similarly be found:

$$\frac{\partial \rho v_i}{\partial t} + \frac{\partial}{\partial x_j} (\rho v_i v_j) = -\frac{\partial P}{\partial x_i} + \rho F_i \quad (i = 1, 2, 3) \quad (2.12)$$

$$\frac{\partial e}{\partial t} + \frac{\partial}{\partial x_i} (e v_i) = -P \frac{\partial v_j}{\partial x_j} + \rho v_j F_j \quad (2.13)$$

where P is the fluid pressure. These equations can be parameterised as:

$$\frac{\partial U}{\partial t} + \frac{\partial \mathfrak{F}(U)}{\partial x} = 0 \quad (2.14)$$

in 1D, and if $\mathbf{F} = 0$, $U \equiv (\rho, \rho v, e)$ and $\mathfrak{F}(U) \equiv (\rho v, \rho v^2 + P, (e + P)v)$. We will return to this parameterisation when discussing the numerical solution of these equations in section 2.2.3.

Together these are the *continuity equations*, and are the basic equations of fluid mechanics. One can add terms to account for viscosity (as in the Navier-Stokes equations), radiative transfer, magnetic fields, etc, as well as derive new continuity equations to describe the flow of, say, chemicals or photon packets. However, this is beyond the scope of this thesis, and I direct the keen reader to the texts referenced at the beginning of this section.

Typically, the dynamics of a fluid may be calculated either by following each parcel of fluid as it moves, or by monitoring fixed positions and calculating the flow through these fixed positions. The former method is known as *Lagrangian*; the latter *Eulerian*. In theory both methods will give the same answer, provided they both solve the same underlying equations. In principle, since in numerical fluid dynamics most treatments are approximations (save for a few analytic cases), each method has its own advantages and disadvantages, which we will discuss briefly in the following sections. However, for the purposes of this thesis we will concentrate on Eulerian, grid-based techniques for non-collisionless fluids.

2.1.2 Discretising the Fluid

In order to translate these equations into forms solvable by a computer, we must discretise the equations and the underlying fluid; in other words, we must break the problem up into chunks that can be computed piecemeal by a processor. The simplest way to consider a fluid is to divide it up on a grid and compute the continuity equations at each grid point by computing the PDEs with respect to neighbouring grid cells. We do this using *finite difference equations*.

For a difference h between two points on the x -axis, we can represent the differential of the function $f(x)$ as:

$$\frac{df}{dx} = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} \quad (2.15)$$

For a PDE of multiple variables, we can write:

$$\frac{\partial f}{\partial x_i} = \lim_{h \rightarrow 0} \frac{f(x_1, \dots, x_i + h, \dots, x_N) - f(x_1, \dots, x_i, \dots, x_N)}{h} \quad (2.16)$$

This equation can be approximated by a finite difference where $h > 0$ if h is small; it is equivalent to taking the first term of the Taylor expansion of $f(x_1, \dots, x_i + h, \dots, x_N)$:

$$f(x_1, \dots, x_i + h, \dots, x_N) = \sum_{n=0}^{\infty} f^{(n)}(x_1, \dots, x_i, \dots, x_N) \frac{h^n}{n!} \quad (2.17)$$

Rearranging, we hence get

$$\frac{\partial f}{\partial x_i} = \frac{f(x_1, \dots, x_i + h, \dots, x_N) - f(x_1, \dots, x_i, \dots, x_N)}{h} + O(h^2) \quad (2.18)$$

Instead of the *forward difference* above, we can write this equation *backwards difference*, i.e.

$$\frac{\partial f}{\partial x_i} = \frac{f(x_1, \dots, x_i, \dots, x_N) - f(x_1, \dots, x_i - h, \dots, x_N)}{h} + O(h^2) \quad (2.19)$$

We can also combine these equations to give a *centred difference*:

$$\frac{\partial f}{\partial x_i} = \frac{f(x_1, \dots, x_i + h, \dots, x_N) - f(x_1, \dots, x_i - h, \dots, x_N)}{2h} + O(h^2) \quad (2.20)$$

For second order differentials, we must include a second term in the Taylor expansion. In 1D (for simplicity), and noting the above centred-difference equation, we get:

$$f''(x) = \frac{f(x+h) - 2f(x) + f(x-h)}{h^2} + O(h^3) \quad (2.21)$$

For small h , we can ignore the higher order terms in h given above in order to obtain approximate expressions for $f'(x+h)$ and $f''(x+h)$. We may also use the Taylor expansion to provide more accurate estimates of $f'(x)$ for a given value of h . Examples of this include the fourth-order Runge-Kutta equation (e.g. Press et al. (1986)).

To apply this to a simulation, we begin from some initial conditions and then step through in time with a

timestep Δt . The simplest approach is to multiply $f'(\mathbf{x}, t)$ at a given instant with Δt and add it to the previous value of $f(\mathbf{x}, t)$ to find $f(\mathbf{x}, t + \Delta t)$. Note that this can be done by using f' at t or $t + \Delta t$, or the average of both (forward, backward or centred time differencing). Forward time differencing is *conditionally stable*; that is, it does not exponentially diverge from an analytic solution if Δt is small enough. Backward and centred time differencing are *unconditionally stable*; that is, the difference from an analytic solution decreases regardless of the size of Δt (although a smaller value will give faster convergence). One approach to solving fluid equations is hence to use some combination of the forward and backward time differencing, as in the Crank-Nicholson method (Crank & Nicolson, 1996). For the time evolution of a fluid variable, solutions that use the differential at t are known as *explicit* and $t + \Delta t$ as *implicit*. As above, explicit solutions are conditionally stable, and implicit solutions unconditionally stable. Unconditionally stable solutions are desirable, since they prevent errors growing with time. Note that, in order for us to find the backward difference or perform any implicit time integration, we must solve the equations by locating the value at $t + \Delta t$ first. To do this, we must produce an estimate for the solution at $t + \Delta t$. This can be done by iterating until we locate a solution with the required accuracy, or by using a Taylor expansion about the solution at t . The latter method is known as *semi-implicit*, since it approximates an implicit solution.

In the following section I introduce the astrophysical numerical hydrodynamics code RAMSES (Teyssier, 2002), and discuss the application of these techniques to methods used to simulate galaxy formation numerically.

2.2 RAMSES

RAMSES is a general-purpose Eulerian AMR and N-body simulation code, which is specialised for the study of galaxy formation, particularly from a cosmological context. It was first presented in Teyssier (2002), and has received several improvements since then (e.g. those described in Dubois & Teyssier (2008a,b)). It is written in Fortran 90, and can be run in parallel across different computers using the Message-Passing Interface (MPI). In this section I briefly outline the methods used by RAMSES to apply the equations discussed in section 2.1 to the problem of simulating supernova feedback in dwarf galaxies.

2.2.1 Adaptive Mesh Refinement

The simplest way to solve the fluid equations on a grid is to use a uniform, rectilinear Eulerian grid. However, we may wish to provide better spatial resolution (and thus, as we will discuss later, time resolution). For example, if we are simulating a supernova propagating through the ISM then we may wish to resolve the shock boundary more highly than the relatively uniform unshocked medium. In cosmological simulations, the density variation between molecular clouds and the intergalactic voids is several orders of magnitude; hence, expending the same computational effort on a patch of void as a molecular cloud of the same volume is not feasible or desirable.

The solution that we employ in this case is to use an adaptive mesh. We start with a uniform grid of a given resolution. We then allow the grid to refine by subdividing itself by a power-of-two. Hence a cell with a level l will have a spatial length of 2^{-l} times the length of the entire box (RAMSES employs cubic – or square in 2D – regions). Eventually, the grid either exceeds the resolution that is required by the refinement criterion, or reaches some maximum threshold set by the user. Conversely, if a grid cell does not meet this refinement criterion then it will be destroyed and the next coarser level considered. The mesh is then smoothed about the refined regions in order to reduce the noise in the refinement map, where very small or thin regions such as shock boundaries activate the refinement criterion. This is important partly because it means that if the bulk velocity of the gas activating the refined region is non-zero, the nearby cells will also be refined when the high-resolution region moves.

The criteria used to determine whether a particular grid cell should refine depend on the application. As alluded to in the last paragraph, simulations of strong shocks benefit from refinement based on pressure or density gradients. Cosmological simulations, on the other hand, are defined by their density contrast, and hence the grids in these simulations will typically be refined based on the mass contained in a given cell.

A widely-used technique in astrophysical numerical simulations is to split a computational task across multiple processors, often with independent system memory. This is done using the *Message Passing Interface* (MPI). A simulation volume is split into multiple “domains”, which is a volume of adjacent grid cells and particles, as well as “ghost” regions surrounding the domain that allow flows to be computed. After a specified number of timesteps, the code communicates changes in adjacent domains between processors, as well as updating the location of the domains if significant refinement has been made in one region or if a




Figure 2.1: The Peano-Hilbert curve (Peano, 1890; Hilbert, 1891) for grids with 1 (red), 2 (black) and 3 (green) levels of refinement respectively. The curve is constructed so that each grid cell is adjacent to the next cell in the curve.

region of high refinement (e.g. a galaxy) has drifted between domains. RAMSES uses the Peano-Hilbert curve (Peano, 1890; Hilbert, 1891) to determine the location of these domains. The curve (see figure 2.1), is designed to link cells such that each cell is adjacent to the next cell on the curve. The cells are split along this curve so as to give an equal number of grid cells to each processor. This process is known as “load-balancing”. The ability of a code to optimally use multiple processors is known as “scaling” - if a particular simulation takes t seconds for 1 processor and t/N seconds for N processors, it scales optimally. If it takes closer to t seconds for N processors, it scales poorly. Processes that require significant communication between CPUs, or processes that must be done on one processor and communicated to all the other processors, typically reduce the quality of the scaling. RAMSES typically scales well for cosmological simulations, where there are multiple regions that only interact via gravity, whereas for problems with a single point source (e.g. a single supernova event as described in chapter 5), the initial scaling can be poor.

2.2.2 Solving the Poisson Equation

RAMSES uses N -body particles to compute the evolution of collisionless fluids. The force on the i th N -body particle (i.e. the force at a given point in a gravitational field made up of N massive particles) is given by:

$$\mathbf{F}_i(\mathbf{x}_i) = \sum_{j=1; j \neq i}^N \frac{Gm_i m_j (\mathbf{x}_i - \mathbf{x}_j)}{|\mathbf{x}_i - \mathbf{x}_j|^3} \quad (2.22)$$

where G is the gravitational constant. Solving this equation requires initial conditions in \mathbf{x}_i and \mathbf{v}_i , and the two equations:

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{v}_i \quad (2.23)$$

$$m_i \frac{d\mathbf{v}_i}{dt} = \mathbf{F}_i \quad (2.24)$$

It is common to express the gravitational field in terms of the potential $\Phi(\mathbf{x})$. Further, since we are modelling a continuous density field rather than a set of discrete particles¹, we can write the potential at \mathbf{x} as:

$$\Phi(\mathbf{x}) = -G \int \frac{\rho(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d^3\mathbf{x}' \quad (2.25)$$

Equations 2.23 and 2.24 can be expressed as the *Vlasov-Poisson* equations. Here, equation 2.24 becomes:

$$\frac{d\mathbf{v}_i}{dt} = -\nabla\Phi \quad (2.26)$$

while equation 2.25 can be written as a Poisson equation:

$$\nabla^2\Phi = 4\pi G\rho. \quad (2.27)$$

In order to solve this equation, RAMSES employs the following basic approach, known as a Particle Mesh (PM) method (Hockney & Eastwood, 1981; Klypin & Shandarin, 1983):

1. Compute the mass density on a regular mesh.
2. Solve the Poisson equation (equation 2.27) to find Φ .
3. Compute the particle acceleration at each grid point using the computed value of Φ .
4. Update the particle velocity using the particle acceleration as in equation 2.24.
5. Update the particle positions using the particle velocity as in equation 2.23.

(adapted from Teyssier (2002)). In principle, this process is fairly straightforward. Most of the subtleties of the methods employed by RAMSES stem from the use of the adaptive mesh rather than a fixed-resolution

¹It should be noted that the particles in the N-body method in this instance represent volumes of a continuous fluid, rather than discrete solid bodies. Even star particles in large simulations represent large groups of star. This is due to computational limits and the prohibitively large number of stars in a typical galaxy.

Eulerian grid.

For all the cosmological simulations discussed in this thesis, a periodic volume is used; in other words, a particle moving off one edge of the volume will appear at the other. The gravitational force also passes across the grid boundary in this fashion. This means that every point in the simulation volume is equal in terms of the accuracy of gravitational forces, and edge effects do not creep in over the Hubble time. However, care must be taken to use a box size such that the low-frequency modes of the density power function are not truncated, i.e. the larger clusters are retained and not prevented from forming due to the box size.

The density field is computed using the ‘‘cloud-in-cell’’ (CIC) technique (Hockney & Eastwood, 1981). In this technique, each particle is considered to be a cubic ‘‘cloud’’ with the physical size of its grid cell centred on the particle; mass is then distributed between neighbouring grid cells according to how much of the cube overlaps with each cell. For an adaptive mesh, we adapt this approach as follows. First, the density field from particles on a given level is computed. Next, particles in neighbouring cells with one coarser level than the current cell are considered in order to determine whether their clouds overlap with the given cell. We can thus compute the density in each grid cell for a mesh with non-uniform resolution.

Once this is done, the Poisson equation must be solved at every point. This is done in RAMSES using a *fast Fourier transform* (FFT) of the density field, and using it to calculate the potential at each point. The advantage of solving the Poisson equation in Fourier space can be demonstrated as follows. The Fourier transform of a function from 1D physical space to Fourier space is given as:

$$f(x) = \frac{1}{\sqrt{2\pi}} \int F(k) e^{ixk} dk \quad (2.28)$$

for all k , where k is the wavenumber of each mode summed to reproduce $f(x)$. Inserting this into the Poisson equation for $\Phi(x)$, $\rho(x)$, we get:

$$\nabla^2 \left(\frac{1}{\sqrt{2\pi}} \int \Phi(k) e^{ixk} dk \right) = -4\pi G \left(\frac{1}{\sqrt{2\pi}} \int \rho(k) e^{ixk} dk \right) \quad (2.29)$$

Noting that in 1D, $\nabla^2 \equiv \frac{\partial^2}{\partial x^2}$ and that the only dependence in x on the left hand side is in the term e^{ixk} , we

obtain:

$$\Phi(k) = -\frac{4\pi G}{k^2}\rho(k) \quad (2.30)$$

The advantage of this method is two-fold. Firstly, we are not required to solve a differential equation any more, and can simply compute a solution to the required accuracy. Secondly, we can ignore high-frequency (i.e. small-scale) modes for distant objects. We can perform a reverse Fourier transform to arrive back at the original solution. Examples of algorithms using discrete Fourier transforms to solve the Poisson equation are detailed in, e.g. Bodenheimer et al. (2007) and Binney & Tremaine (2008).

This method works well for collisionless fluids moving under gravity. However, the problem becomes complicated if we include collisional fluids, since now the fluid parcels must exchange physical quantities such as mass, momentum and energy with neighbouring parcels as well as interacting with the global gravitational field. To do this, RAMSES uses the same AMR structure used to build the particle mesh to store and compute the collisional fluid variables. This requires a very different method to compute the flow, which is discussed in the following section.

2.2.3 Eulerian Fluid Dynamics

In order to compute the flow of mass, momentum and energy in collisional fluids, RAMSES solves equations 2.11 to 2.13 in section 2.1.1, with an additional term to describe the gravitational force as in section 2.2.2. The gravitational source terms for mass, momentum and energy conservation are 0, $-\rho\nabla\Phi$ and $-\rho\mathbf{u}\cdot\nabla\Phi$ respectively. RAMSES also uses the equation of state:

$$p = (\gamma - 1)\rho\left(e - \frac{1}{2}u^2\right) \quad (2.31)$$

where e is the specific total energy, γ is the adiabatic index or heat capacity ratio, and u is the fluid velocity, with all other terms defined as before. Teyssier (2002) notes that RAMSES is not explicitly energy-conserving due to the inclusion of gravity as a source term, with errors of a few percent or less. However, the use of the continuity equations mean that this does not affect the bulk flow dynamics.

RAMSES computes fluxes across cell boundaries using a *time-centred* approximation - i.e. the value of the fluxes is considered most representative in the middle of the timestep. Note also that the fluid vari-

ables are sampled at the cell *centre*. The gravitational source terms are computed by averaging the values for the terms at the start and the end of the timestep. The fluxes are computed using a second-order Godunov/Piecewise Linear Method (PLM). The PLM is an approximation of the Godunov solution to the Riemann problem (Collela, 1990; Saltzman, 1994). The Riemann problem requires the solution of the advection of fluid properties across shocks given by the Rankine-Hugoniot conditions. In the rest frame of the shock front, these can be written as:

$$\rho_2 u_2 = \rho_1 u_1 \quad (2.32)$$

$$\rho_2 u_2^2 + P_2 = \rho_1 u_1^2 + P_1 \quad (2.33)$$

$$u_2 \left(\frac{1}{2} \rho_2 u_2^2 + e_2 + P_2 \right) = u_1 \left(\frac{1}{2} \rho_1 u_1^2 + e_1 + P_1 \right) \quad (2.34)$$

The conditions specify that mass, momentum and energy respectively are conserved across the shock front. Note that these conditions are also true for non-shock conditions across any boundary.

We now consider solutions for equation 2.14. Neglecting the source term for gravity, we can write the cell-centred solution to this equation:

$$U_{i+1/2}^{n+1} = \frac{1}{\Delta x_{i+1/2}} \left[\int_{x_i}^{x_{i+1/2}} \mathfrak{U}_i(x, t^{n+1}) dx + \int_{x_{i+1/2}}^{x_{i+1}} \mathfrak{U}_{i+1}(x, t^{n+1}) dx \right] \quad (2.35)$$

and hence it can be shown that:

$$U_{i+1/2}^{n+1} = U_{i+1/2}^n + \frac{\Delta t^{n+1/2}}{\Delta x_{i+1/2}} \left[\bar{\mathfrak{F}}(\mathfrak{U}_i(x_i)) - \bar{\mathfrak{F}}(\mathfrak{U}_{i+1}(x_{i+1})) \right] \quad (2.36)$$

where $\bar{\mathfrak{F}}$ is the flux averaged over $\Delta t^{n+1/2}$. \mathfrak{U}_i and \mathfrak{U}_{i+1} are the solutions to the Riemann problem between $U_{i-1/2}^n$ and $U_{i+1/2}^n$, $U_{i+1/2}^n$ and $U_{i+3/2}^n$ respectively. A linear fit to each value of U is made, which allows U to vary within grid cells. This has the effect of reducing numerical diffusion of flows, which can become an issue for structures with a constant, uniform bulk velocity through the grid.

The Riemann solver used by RAMSES hence computes the Godunov states using the wave equations given above, with transverse fluxes as well as those along the x,y,z axes of the grid. Various Riemann solvers are available, but the one used in the majority of the work in this thesis is the HLLC (Harten, Lax, van-Leer Contact) solver. This is based on the HLL (Harten, Lax, van-Leer) method, which treats the fluid across

the interface as being composed of two wave solutions moving between three constant states (Toro, 2006). These wave expressions are then inserted into the continuity equation to determine a solution to the flux at the grid interface. The problem with this method is that the assumption of two waves is only truly correct for systems containing two hyperbolic equations. With more complex flows, the HLL method will be unable to resolve features such as contact surfaces or shear waves. The HLLC method was proposed as a solution to this problem by Toro et al. (1994), and includes a term to account for contact surfaces and rarefaction waves.

One further consideration is the length of the timestep. A flow should not be allowed to pass information beyond the adjacent cells, since this will introduce errors and instabilities. To prevent this, RAMSES uses the *Courant-Friedrichs-Lewy* (CFL, or simply “Courant”) condition to determine the minimum timestep. Here, the timestep is limited to a fraction of the sound crossing time for a given cell (or the time taken for the bulk flow to travel across the grid cell, whichever is shorter). Typically, this fraction is around 0.5 of the crossing time. Further conditions are made to ensure that the free-fall time and the particle crossing time do not exceed some fraction of the timestep.

This is hence the basis for all fluid computations in RAMSES. In the following sections we briefly discuss the performance of these methods in comparison to other fluid solvers, and introduce extra models to account for processes in galaxies that occur below the resolution of the grid.

2.2.4 Eulerian vs Lagrangian Techniques

It is worth briefly noting that Eulerian grid solvers are not the only way to treat gas dynamics in astrophysics. One may also treat fluid parcels as Lagrangian, allowing them to move along the flow vector. There is no mathematical reason why Eulerian techniques are superior to Lagrangian techniques; indeed, many authors use a Smoothed Particle Hydrodynamics (SPH) formalism for solving the fluid equations in astrophysical hydrodynamics simulations (e.g. Springel (2005); Wadsley et al. (2004)). In this formalism, fluid variables are transferred between particles based on a smoothing kernel around each particle, where a specified number of nearest neighbours are sampled to determine the radius of this kernel.

There are many reasons why this approach is attractive. By modelling gas flow with particles, one may track their evolution without specifically adding tracer particles to follow the gas flow (tracer particles

have now been implemented in some versions of RAMSES). Further, if one refines based on mass, no adaptive refinement is needed since the particles naturally trace the flow of mass into high density regions. Eulerian techniques can experience numerical diffusion caused by numerical errors in tracing the movement of structures across the grid in bulk flows; however, Tasker et al. (2008) show that this effect is small.

The chief problem with SPH is that it is poor at capturing instabilities. Agertz et al. (2007) demonstrate that for dense flows moving past low-density flows tend to undersample the low-density fluid as a result of the smoothing kernel, and hence are unable to produce Kelvin-Helmholtz instabilities in this case (for flows of comparable density, the result is more favourable). Further, a dense sphere moving through a low-density wind does not, for similar reasons, experience shredding from turbulence and survives much longer than the Kelvin-Helmholtz timescale expected. This is worrying for simulations tracing the evolution of Milky Way satellite galaxies; if we accept the analogue of a satellite falling through the Milky Way halo in this example, the simulation will not capture ram-pressure stripping from the Milky Way's gas halo accurately. Further, Mitchell et al. (2009) demonstrate that SPH particles retain their entropy, causing artificial cusps to form in simulations of two gas clouds colliding (as in the merger of galaxy clusters).

There are a number of fixes to this issue discussed in the literature (Abel, 2011; Cha et al., 2010; Heß& Springel, 2010; Junk et al., 2010; Kawata et al., 2009; Murante et al., 2011; Price, 2008; Read et al., 2010), though many SPH codes remain unmodified and hence still experience this issue. Springel (2010) suggests a bolder approach in replacing the smoothing kernel in Lagrangian codes with a grid computed via Voronoi tessellation. This code, AREPO, is still in its early stages, but test cases demonstrate a favourable comparison with Eulerian AMR.

In the following sections we discuss techniques used to turn RAMSES into a galaxy formation simulation code.

2.3 Galaxy Formation Simulations

Once the basic equations of cosmological fluid dynamics are solvable in a numerical code, we still require some further steps before we are able to simulate galaxy formation in a cosmological context. First, cosmological initial conditions are required. Typically, we begin in the regime of linear perturbations to the density field after the Big Bang, and start the simulation just before the evolution of nonlinearities. Sec-

ond, we require certain sub-grid models that compensate for our inability to resolve processes such as star formation and stellar evolution using direct gas dynamics. We thus discuss models for the transfer of gas mass, momentum and energy to stars during star formation, and then back again if supernova feedback is implemented.

2.3.1 Cosmology & Initial Conditions

Our approach to modelling the initial conditions of the universe is to convolve the power spectrum of the universe at the end of the linear regime of density perturbations with a white noise field. This is because the equations of fluid dynamics are nonlinear, and hence we cannot recompute the precise initial conditions of a given observed galaxy such as the Milky Way (although some authors attempt to reproduce the broad properties of the Local Group, such as Libeskind et al. (2010)). Instead, we produce a universe that has the same properties as our own, and assume that a simulated galaxy in this universe will exhibit similar properties to galaxies of similar mass in a similar environment to galaxies in our own universe.

In order to generate initial conditions, we use the code MPGRAFIC (Bertschinger, 2001; Prunet et al., 2008). MPGRAFIC takes as its inputs the cosmological parameters Ω_Λ , Ω_M , h , Ω_b , n and σ_8 (in order, energy density parameter given by the cosmological constant, matter density parameter, baryon density parameter and the root-mean-square mass fluctuation amplitude in spheres of size $8h^{-1}\text{Mpc}$), as well as parameters for the box size in physical size units, the number of grid cells per spatial dimension, the white noise seed and the physical model to use for the power spectrum. The results are given as density and velocity fields for the dark matter and baryons, statistically distributed according to the selected power spectrum. We use the power spectrum given in Eisenstein & Hu (1998), and assume a WMAP cosmology (Larson et al., 2010). These results may then be read by RAMSES and used as the initial conditions of the particles and grid cells that are hence evolved by the solvers described in this chapter.

2.3.2 Radiative Cooling

In order to form stars inside halos, we first require the gas to cool. This may seem counter-intuitive at first glance, since stars are hot objects, but in order to allow the gas inside halos to collapse to densities that allow nuclear fusion to take place (i.e. the interior of stars), the gas must lose the pressure support that counteracts

the gravitational force of a halo. It is able to do this via *radiative cooling*. Early implementations of gas cooling in simulations are discussed in Sutherland & Dopita (1993) and Ferland et al. (1998).

The thermal energy density of the gas in a halo is given by:

$$u = (\gamma - 1)nk_{\text{B}}T = (\gamma - 1)\frac{\rho}{\mu m_{\text{H}}}k_{\text{B}}T \quad (2.37)$$

If we invoke a cooling rate $\Lambda(\rho, T)$ in energy emitted per unit volume per unit time, we can write an expression for the “cooling time” t_{c} , i.e. a scale time over which the gas loses thermal energy:

$$t_{\text{c}} \equiv \frac{u}{\dot{u}} = \frac{T}{\dot{T}} = \frac{u}{\Lambda(\rho, T)} \propto \frac{\rho T}{\Lambda(\rho, T)} \quad (2.38)$$

Note that the cooling rate is a function of the gas density rather than the total matter density, which includes dark matter. If we assume that cooling is predominantly a process caused by collisions (and that $n_{\text{e}} \propto n_{\text{H}}$), we find $\Lambda \propto \rho^2$. This gives us $t_{\text{c}} \propto 1/\rho$. In other words, high density gas cools faster than low density gas.

There are three regimes in which gas may cool:

1. $t_{\text{c}} > t_{\text{H}}$

Here, the cooling time is longer than t_{H} , the Hubble time $\equiv 1/H$, where H is the Hubble constant. In other words, the gas cannot cool, even on cosmological timescales. This is a good approximation to the conditions inside galaxy clusters. Here, ρ and T are constant.

2. $t_{\text{H}} > t_{\text{c}} > t_{\text{dyn}}$

Here, the gas cools but does so quasi-adiabatically, i.e. at each time the gas can return to a dynamic equilibrium and will maintain a constant Jeans mass M_{J} (i.e. the gas pressure forces and gravitational forces will balance). $M_{\text{J}} \propto \rho L_{\text{J}}^3$ where L_{J} (the Jeans length) $\propto c_{\text{s}}\rho^{-1/2}$ (e.g. Binney & Tremaine (2008)), and c_{s} (the sound speed) $\propto T$. We can thus obtain $\rho \propto T^3$.

3. $t_{\text{H}} > t_{\text{dyn}} > t_{\text{c}}$

In this case, the cooling occurs faster than the system can change dynamically, and hence the temperature will change faster than the density, and the system can become unstable. If $t_{\text{c}} \gg t_{\text{dyn}}$, the cloud will have no time to collapse, and the mass will exceed the Jeans mass. In this case, the system

will fragment until the mass of the fragments is of order of the Jeans mass and is able return to the previous, quasi-adiabatic state. This regime is important because fragmentation is necessary for stars to form.

There are processes that cause astrophysical particles to radiate energy. The chief processes that cause radiative cooling involve a combination of collisions between either electrons and ions, or ions and photons. In the latter case the cooling rate is proportional to density, since the number of particles that a photon may interact with increases linearly with density. In the former to density squared, since it describes the collision between two particles (the number of electrons is related to the number of ions, though multiply ionised particles also exist depending on the temperature of the gas). Cen (1992) produces a detailed description of the construction of such a cooling model, in which the change in thermal energy and number density of species of electrons and hydrogen and helium atoms and ions due to various cooling processes is considered. For completeness, these processes are listed briefly below:

- Collisional ionisation cooling: radiative losses from collisions between electrons and ions/atoms causing ionisation.
- Recombination cooling: radiative losses from collisions between electrons and ions recombining to form neutral / less ionised particles.
- Collisional excitation: radiative losses caused by electrons in atoms/ions radiating energy from excitation by collisions.
- Bremsstrahlung cooling: Radiation emitted by an electron being decelerated via electromagnetic interaction as it passes an ion.
- Compton cooling: Low-energy cosmic microwave background photons removing energy from particles via Compton scattering.

Heating processes include:

- Photoionisation: External photons exciting and removing electrons from atoms/ions
- Compton heating: Heating from background radiation if the gas temperature is lower than the cosmic microwave background (CMB) temperature (i.e. the reverse of Compton cooling).

- Shock heating: as the gas collapses, bulk flows collide, which induces collisional heating between particles as a result of viscosity which is naturally included in the fluid equations solved by numerical simulation codes.

Dalgarno & McCray (1972) and Rosen & Bregman (1995) also discuss the effect of metal-line cooling on the cooling function - since metals have more emission lines than hydrogen or helium, metal-enriched gas provides more efficient cooling than gas with a primordial metallicity.

If one assumes that the number density of each ion is roughly in equilibrium at each value of total gas density ρ , temperature T and metallicity Z , we may construct a table for $\Lambda(\rho, T, Z)$ and find the value of the cooling function for a given (ρ, T, Z) by interpolating across the table. This is the approach RAMSES takes, since it does not require memory-intensive tracking of the chemistry of the gas. Using parameterisations similar to those described above, we can use the value for \dot{u} demonstrated in equation 2.38 to compute the rate of energy loss (or gain, if heating exceeds cooling) for a gas cell at a given time. RAMSES has the option to use various models for cooling, but the one used in the majority of the simulations discussed is given in Courty & Teyssier, (in prep.), which includes the UV background as a separate parameter in Λ . Further, below 10^4K RAMSES then uses a semi-implicit integration method to find the change in thermal energy. As mentioned in section 2.1.2, the implicit solution requires knowing the value of temperature at the next timestep, T^{n+1} . Iterating to find this value using, say, a binary search over the whole parameter space can be computationally expensive. Instead, the semi-implicit method approximates T^{n+1} by using a Taylor expansion of $\Lambda(\rho, T)$ in T . If more accuracy is required, RAMSES sub-cycles this process inside a single timestep until we find the solution. This provides the stability and improved energy conservation of an implicit method without the computational cost of iteratively finding T^{n+1} .

Note that RAMSES decouples thermal energy from kinetic energy, since for cold flows the kinetic energy may dominate the thermal energy. Thus if we do not decouple the two, numerical errors may cause the code to assume that all of the energy is kinetic and erroneously return a value of 0K for the gas temperature. In the next section we describe how the cooled gas in RAMSES can form stars. Since we are in the business of galaxy formation, and galaxies contain stars, the next step is crucial if we are to recreate the grand structures that we observe in the universe around us.

2.3.3 Star Formation

Cosmological hydrodynamic galaxy formation simulations are unable to capture the length scales required to simulate star formation directly in the gas. Instead we must adopt a *sub-grid model* for star formation. This involves applying an analytic expression that converts a certain portion of the gas in a grid cell into stars when certain criteria are met.

We use the following equation for star formation rate, motivated by a Schmidt-Kennicutt law (Kennicutt, 1998):

$$\dot{\rho}_* = \frac{\rho}{t_*} \text{if } \rho > \rho_c \quad (2.39)$$

where ρ_c is the star formation density threshold and t_* is the timescale on which the stars form. t_* is a multiple of the local free-fall time t_{ff} ($\sim 1/\text{sqr}(G\rho)$). Typically t_* is of the order of 10-100 times t_{ff} ; the precise value is resolution dependent, since if we are able to resolve molecular clouds, we can set the density threshold to only form stars in these regions, and here the conversion of gas to stars will be more efficient than if we just consider a global law. Krumholz et al. (2012) find by comparing a variety of observations that star formation occurs according to a similar law:

$$\dot{\rho}_* = \varepsilon_{\text{ff}} f_{\text{H}_2} \frac{\rho}{t_{\text{ff}}} \quad (2.40)$$

where ε_{ff} is the efficiency of star formation compared to the free-fall timescale ($\sim 1\%$ in Krumholz & Tan (2007); Krumholz et al. (2012)) and f_{H_2} is the fraction of H_2 molecules in the gas. They determine that this is a local law rather than a global law, so stars form according to the local conditions in gas clouds rather than as a result of the global properties of their host galaxy. Depending on the resolution, the fraction of H_2 will vary (though not by more than an order of magnitude).

Hence, ρ_c can be used to constrain the places in which stars are able to form, and will typically be set according to the Jeans mass of a cell at the lowest level of refinement. For simulations that do not capture the spatial resolution of molecular clouds, this can be as low as 0.1 atoms/cm^3 ; i.e., molecular clouds are assumed to be distributed across the galaxy. If molecular clouds are included, we can set the threshold to be 10^5 atoms/cm^3 . However, we may need to tune the density threshold, since the assumption that a given cell collapses as a fixed gas cloud may not hold as we lose the ability to resolve the multiphase ISM. This is

discussed in chapter 4. We also need to ensure that the density threshold is not set so low as to allow stars to form in the filaments, which is unphysical.

Once we have formed stars, they are decoupled from the gas as star particles, which behave as N-body particles tracing a continuous distribution of stars with a given IMF. In the next section we consider stellar feedback, and how supernova events return mass, momentum and energy to the gas, and hence induce complex feedback cycles in the star formation rate.

2.3.4 Supernova Feedback and Reionisation

In order to model supernova feedback in cosmological hydrodynamic simulations in RAMSES, the code deposits a pre-evolved Sedov-Taylor blastwave (Sedov, 1946; Taylor, 1950) onto the grid over 2-3 grid cells (Dubois & Teyssier, 2008a). The reason for doing this is that typically the initial blastwave is poorly resolved. Previously, RAMSES implemented supernovae by creating a number of “debris” particles and ejecting them from the supernova, then recoupling them to the gas after a given delay. By pre-computing the Sedov blastwave in the given density field and placing it onto the grid, we are able to attain an accurate picture of the early supernova evolution. Previous authors had a problem with overcooling after supernova events. This was caused by poor time and spatial resolution, which made the supernova-heated gas radiatively cool dramatically, losing most of its energy before it could expand into the ISM. Gerritsen & Icke (1997) addressed this problem by allowing the gas to expand adiabatically shortly after the supernova, essentially turning off cooling during this phase. Springel & Hernquist (2003) simply removed an analytically-determined proportion of gas from galaxy disks as winds. However, for high-resolution simulations of dwarf galaxies with modern supercomputers, we are able to avoid this problem simply by resolving the length scales necessary to capture the multiphase ISM without dramatic loss of thermal energy from radiative cooling. In general, the Sedov-Taylor model for supernova blastwaves is good for early expansion (see chapter 5). However, the profile of expansion will depend on the precise density profile into which the supernova expands. In chapter 5, we discuss detailed simulations of isolated supernova events intended to address this issue.

RAMSES has various models for how the energy from supernovae is deposited in the gas, and is subject to ongoing development. This is not a trivial choice, since if we release all the supernova energy instantly

after the star is formed, it detonates in a high-density star-forming region; if it drifts, it can move into regions of lower density. Slyz et al. (2005) demonstrate that this causes a dramatically different star formation rate (SFR) and density distribution of the ISM.

For a simulation where the star particle mass is large, we can consider it to be a continuous stellar population. For this model, Cen & Ostriker (1992) propose depositing mass, energy and momentum in a distribution governed by the star formation timescale discussed in section 2.3.3, or the expected lifetime of a massive star ($\sim 10\text{Myr}$), whichever is larger. Ceverino & Klypin (2009) suggest a more complex model, in which they deposit energy from supernovae as well as stellar winds according to a Starburst99 model (Leitherer et al., 1999). They also suggest that adding a momentum “kick” to star particles, simulating the transfer of momentum from interactions between stars shortly after forming in molecular clouds, can provide results that converge when the resolution of simulations is reduced.

For a small stellar mass ($M \lesssim 100M_{\odot}$), this assumption breaks down, since the stellar initial mass function (IMF; e.g. Salpeter (1955)) produces approximately one supernova per $\sim 100\text{-}200M_{\odot}$. In chapter 4, I introduce $\overline{\text{NUT}}$, a suite of simulations that assumes a fixed time delay of 10Myr over which to deposit its energy, since its star particle mass reaches $80M_{\odot}$ for the most detailed simulation. This is based on the assumption that a star particle contains on average one Type II supernova with a lifetime of 10Myr (Powell et al., 2010; Kimm et al., 2011). In this suite, we assume that supernova events entrain 50% of the initial mass of the star particle ($\eta_{\text{W}} = 1$ in the notation of Dubois & Teyssier (2008b)) in a wind and have a metal yield of 0.1. The energy released is given by $\eta_{\text{SN}} \frac{m_{\star}}{m_{\text{SN}}} e_{\text{SN}}$, where m_{\star} is the mass of the star particle, m_{SN} and e_{SN} are typical values for a massive star going supernova and η_{SN} is the fraction of the total mass formed that is turned into supernova ejecta. We use $\eta_{\text{SN}} = 0.106$, $m_{\text{SN}} = 10M_{\odot}$ and $e_{\text{SN}} = 10^{51}\text{ergs}$.

One further process for gas heating is modelled - that of photon heating. Benson et al. (2003) suggested that invoking photoionisation as a heating process from the UV background is necessary if we are to explain the low end of the galaxy mass function. In RAMSES, we adopt a Haardt-Madau model of photon heating (Haardt & Madau, 1996), in which a heating term is added uniformly to the gas after a given redshift. This redshift is set according to the start of the epoch of reionisation; values from $z=11$ to 6 are adopted; in chapter 4, we use $z=8.5$. This approach ignores the effects of internal versus external heating of a given halo; a halo is able to “self-shield” some of the photon energy coming from outside of the halo. This is known to

occur around densities $n_{\text{H}} \sim 0.1 \text{ atoms/cm}^3$ (Susa & Umemura, 2004). The sources of cosmic UV photons are still uncertain. Quasars and massive stars at high redshift have previously been thought to be the source of high-energy photons, but recent work has suggested that shocks from supernovae and accretion onto halos may also play a role (Johnson & Khochfar, 2011; Wyithe et al., 2011). In the next section we discuss the methods recently implemented in RAMSES to compute direct radiative transfer inside hydrodynamic simulations, which can be used in future simulations of the photoionisation of gas halos.

2.3.5 Radiative Transfer

In the previous section on radiative cooling, the techniques discussed typically assume that the photons radiated are lost to the system. This is not an unreasonable assumption; the universe is generally optically thin beyond the epoch of *recombination*, when the mean free path of photons became comparable to the size of the universe. Photons emitted from a galaxy typically do not interact with another system, although they may be absorbed by the IGM in the process. However, this is not always a reasonable assumption. Dust inside a galaxy renders regions of the ISM optically thick. Self-shielding, the absorption of photons by the outer parts of a gas halo, can protect the inner parts of a halo from photoionisation heating (Susa & Umemura, 2004). On smaller scales, photoionisation from massive stars can disrupt molecular clouds and regulate further star formation (Rodriguez-Gaspar et al., 1995). A hydrodynamic model of radiative transfer is hence important for detailed simulations of the physical processes affecting galaxy formation.

Rosdahl & Blaizot (2012a) present the first publication using RAMSES-RT (Rosdahl & Blaizot, (in prep.)). I use this module in chapter 5 to investigate the evolution of supernova remnants and their impact on the ISM. RAMSES-RT uses a moment-based method as given in Aubert & Teyssier (2008). This models radiative transfer using packets of photons in discrete wave bands as a fluid on the grid, and follows their evolution from sources such as star particles. It also modifies the cooling module to follow the evolution of the ionisation fractions H_{II} , He_{II} and He_{III} , in order to self-consistently follow the evolution of UV absorption and emission in the gas.

The advantage of the moment-based method is that it allows the grid to be split into domains, as before, without requiring significant extra communication between processors, as ray tracing methods often require due to the high speed of light. Further, since every gas cell is considered a source, multiple scattering is

handled implicitly by the code. The results of RAMSES-RT compare favourably with the results of other radiative transfer modules described in Iliev et al. (2006, 2009). One issue is that the Courant condition also applies to photons, since they are considered to be fluids on the grid. Hence the high speed of light means that the timestep required is prohibitively short. However, a solution proposed by Gnedin & Abel (2001) involves reducing the speed of light used in the simulations. Rosdahl & Blaizot (2012a) note that reducing the speed of light by a factor of 1000 has no effect on their results.

The final ingredient we will discuss in simulations of dwarf galaxy formation is resimulation, a technique for simulating only specific objects inside a cosmological volume.

2.4 Resimulation

A simulation with an adaptive grid is able to simulate a cosmological volume containing filaments, halos, galaxies and voids such that we can resolve the processes occurring inside galaxies at the same time as capturing cosmological flows. However, we are still limited by computational resources in this instance. For example, if we wish to simulate dwarf galaxies in a cosmological context, we must also capture clusters containing massive ellipticals. Hence using a uniform refinement criterion on the AMR grid, the dwarf galaxies will be barely resolved unless we also simulate massive galaxies with the same resolution. This is a problem if we want to simulate the feedback processes in dwarf galaxies with sufficient resolution to capture the multiphase ISM.

One solution to this problem is to use a *resimulation*, or “zoom” simulation. These restrict refinement to the spatial co-moving volume inside which the dwarf galaxy moves. By doing this, we can use the cosmological volume in order to obtain a physically realistic accretion history and external gravitational field at the same time as resolving the ISM of the dwarf galaxy in great detail. In order to do this, we must first run a low-resolution simulation containing only dark matter in order to constrain the volume in which the matter travels throughout the Hubble time. If we increase the initial refinement of the grid in a given region at the “0th” timestep, RAMSES will also increase the resolution of the dark matter field and put $2^{n_{\text{dim}}\Delta L}$ extra particles of $2^{-n_{\text{dim}}\Delta L}$ lower mass, where n_{dim} is the number of dimensions ($= 3$ for cosmological simulations) and ΔL is the difference in refinement level between the refinement region defined in the resimulation and the coarse level of the grid outside that region. If we require $\Delta L > 1$, we will need to include several buffer




Figure 2.2: Projection of a $10h^{-1}\text{Mpc}$ on a side box showing the hierarchy of grids for a test simulation of a $10^{11}M_{\odot}$ galaxy at $z=0$. The colours represent the level of refinement. *Left panel:* Slice through the entire simulation volume, illustrating the spherical resimulation region (green (512^3 equivalent resolution), with red (256^3 equivalent resolution) buffer region) embedded in a 128^3 root grid (white). The blue and purple (1024^3 and 2048^3 equivalent resolution) represent additional AMR refinement applied only inside the resimulation region. *Right panel:* Magnification of the resimulation region, showing the extent of the refinement around the object being studied. The green arrow shows the location of the galaxy being studied.

regions that increase in refinement level by 1 until we reach the required resolution. This is to limit the mixing of particles at the coarse resolution and the highest-resolution particles.

Note that only using N-body dynamics is computationally less expensive than including gas dynamics, and we only require that we capture the bulk flows in order to carve out this region. Further, we will need to include a large buffer region around the selected high-resolution region to account for non-linearities in particle trajectories (although a halo's position may converge to a given value, individual particles in the N-body equation have been known to move chaotically since Poincaré). If a low-resolution (and hence high mass) particles enter the refinement region, they will disrupt the gravitational field of the dwarf galaxy. In a naive particle-particle force calculation, this will be equivalent to a large black hole moving through a stellar system. Even with a CIC N-body solver, this will still present problems.

We then force the resolution inside the refinement region to reach at least the specified value. We may then allow further refinement to occur inside this region according to normal adaptive mesh refinement

criteria. Using this technique it is now computationally feasible to simulate one (or more, if we use more than one refinement region) dwarf galaxies in a cosmological context while capturing the ISM at high resolution.

2.4.1 Resimulating Dwarf Galaxies

Typically, resimulations are performed inside a static co-moving volume (i.e. one whose position does not vary with time in co-moving space). The advantage of this is that it is simple to parameterise and implement (for a sphere, we require only the position and radius). Further, for galaxies with the mass of the Milky Way or larger, this is reasonable, since the galaxies does not move much in co-moving space over the Hubble time.

However, if we consider dwarf galaxies, a static spherical volume containing the trajectory of every parcel of fluid in it at $z=0$ may be around 10% of a $10h^{-1}\text{Mpc}$ volume. This is because a dwarf halo can travel $\sim 1\text{Mpc}$ over the Hubble time towards regions of higher density (see figure 2.3); simple classical mechanics states that for a given gravitational force, an object of low mass will accelerate faster for a given gravitational force. Since we require the dwarf galaxy to be resolved in high detail over its entire lifetime, resolving the entire volume it travels through during its lifetime can lead to unnecessary calculations over volumes of space that for large periods of time do not host the dwarf galaxy. Worse, other galaxies that we might not want to study may drift into the high resolution region and consume CPU time.

For these reasons, I present new techniques for reducing the volume of the resimulation region in dwarf galaxy simulations. In the test case presented in this section of a $10^{11}M_{\odot}$ dwarf halo, we reduce the refinement region's volume by 30% compared to a static resimulation region tracing the same fluid parcel trajectories. It may be possible to reduce the volume more for smaller dwarf galaxies that travel further in their lifetime.

I describe three methods for achieving this. The first is to give the initial conditions a velocity 'kick' corresponding to the dwarf galaxy's average peculiar velocity (where the average is taken over the lifetime of the galaxy) so that the galaxy drifts less and requires a smaller resimulation region. The second is to use a non-spherical geometry, thus capturing only the region in which the galaxy moves. The third is to move the refinement region with the velocity field of the dark matter particles inside the halo. The third method is




Figure 2.3: Top left to bottom right: Images taken with Alice (see section 3.2.2) of the dark matter particles in a dwarf halo from the initial conditions (top left image) at $z=500$, $z=4$ (top right), $z=1.5$ (bottom left) to $z=0$ (bottom right image), using co-moving coordinates. White particles are particles inside the virial radius of the dwarf halo at $z=0$; blue particles are the rest of the dark matter particles in the simulation. The pink dot shows the halo particles' mean position in the initial conditions. The halo moves ~ 1 Mpc from its starting point's co-moving position at $z=0$. By tracking the position of every particle in the halo at every time, we can carve out a region of high resolution inside which to resimulate the halo.

found to provide the most significant improvements, but the other methods are discussed for completeness. A fourth method, used in Hahn et al. (2010), is also presented and discussed briefly in comparison to the moving region technique.

2.4.2 Velocity Shift Method

One possible correction for the drift of dwarf galaxies in co-moving space is to modify the initial conditions such that the drift of the dwarf galaxy is minimised, while preserving Galilean invariance. To do this, we calculate the dwarf galaxy's average peculiar velocity over the Hubble time, and subtract that from the initial conditions. For small objects moving 1Mpc in co-moving coordinates over the age of the universe, this velocity can have a magnitude of $\sim 100\text{km/s}$.

The problem with this approach is that the velocity shift perturbs the velocity field most strongly early in the simulation. There are two reasons for this. Firstly, early in the simulation dark matter and gas velocities are just transitioning from the linear regime and are therefore low compared to the mean peculiar velocity of the dwarf galaxy over Hubble time. Secondly, early in the simulation the physical box size is smaller, so when kicked with its average peculiar velocity, the dwarf galaxy travels a larger physical distance before it slowly returns to its original position. Therefore the result of giving the initial conditions a velocity kick is that the dwarf progenitor nearly instantaneously undergoes a large displacement. Typically, this displacement is about the same size as the original drift, thus negating the value of this method.

2.4.3 Non-spherical resimulation region

One of the simplest ways of improving computation time is to remove from the refinement region any volume that does not contain objects of interest. When using a spherical region, this might be a large fraction of the volume of the region. Generally, isolated objects (i.e. not orbiting a larger halo) move in a roughly cylindrical volume, first undergoing an expansion phase due to Hubble expansion, before collapsing under gravity at late times. Thus two geometries are suggested – one, a cylinder, two, a pear-shaped region. The underlying mathematics for the two geometries is the same, with differing boundary conditions. We define the pear-shaped region as a cylinder with linearly varying radius (i.e. a conic section) truncated by half-spheres instead of flat planes. The benefit of the pear shape is that the spheres at either end capture

the geometry of the initial and final positions of the halo's particles. The early position of the particles is diffuse, while at $z=0$ they have collapsed under gravity to a smaller radius.

One of the problems with using a non-spherical resimulation region is that more calculations must be performed than a spherical region to determine if a given grid cell belongs to the region. Optimisation is thus desirable, even though the computation time saved in solving the gas flow at lower resolution more than compensates for the extra time spent performing the region selection calculations.

The algorithm here selects first for objects within the two spheres at either end of the pear, then for objects in the conic section between them.

If \mathbf{x} is the cell position being tested, and $(\mathbf{p}, r)_i$ for $i = \{1, 2\}$ are the two points and associated radii for the spheres at either end of the pear, then we have:

$$(\mathbf{x} - \mathbf{p}_1)^2 < r_1^2 \text{ or } (\mathbf{x} - \mathbf{p}_2)^2 < r_2^2 \quad (2.41)$$

for the cell to be flagged as belonging to the pear. In a similar vein, defining intermediate variables, \mathbf{q}, L and r' such that $\mathbf{q} \equiv \mathbf{p}_2 - \mathbf{p}_1$, $L \equiv \mathbf{q} \cdot \frac{(\mathbf{x} - \mathbf{p}_1)}{q^2}$, and $r' \equiv r_2 + (1 - L)r_1$, the cell will be flagged as part of the pear (and cylinder) if:

$$r'^2 < (\mathbf{p}_1 + L\mathbf{q} - \mathbf{x})^2 \text{ and } 0 < L < 1 \quad (2.42)$$

The first condition (2.41) selects cells within spheres centered on \mathbf{p}_1 and \mathbf{p}_2 , and $\mathbf{p}_1 + L\mathbf{q}$ is the line equation between these two points, normalised in such a way that $L = 0$ at \mathbf{p}_1 and $L = 1$ at \mathbf{p}_2 . Finally, r' is the scaled radius of the cylinder (actually a cone if $r_1 \neq r_2$). Hence the final condition (2.42) selects cells located within spheres of radius r' whose centers lie on the segment joining \mathbf{p}_1 and \mathbf{p}_2 .

This method works well if it is important that the entire region stays refined for the whole simulation. However, the method described in section 2.4.4 provides a larger decrease in resimulation volume, and thus I concentrate on the results from this method.

2.4.4 Moving Region Refinement

While we model interstellar gas with an Eulerian grid, dwarf galaxies move with the cosmic flow towards regions of higher density. It thus makes sense that the resimulation region follows the Lagrangian coordi-

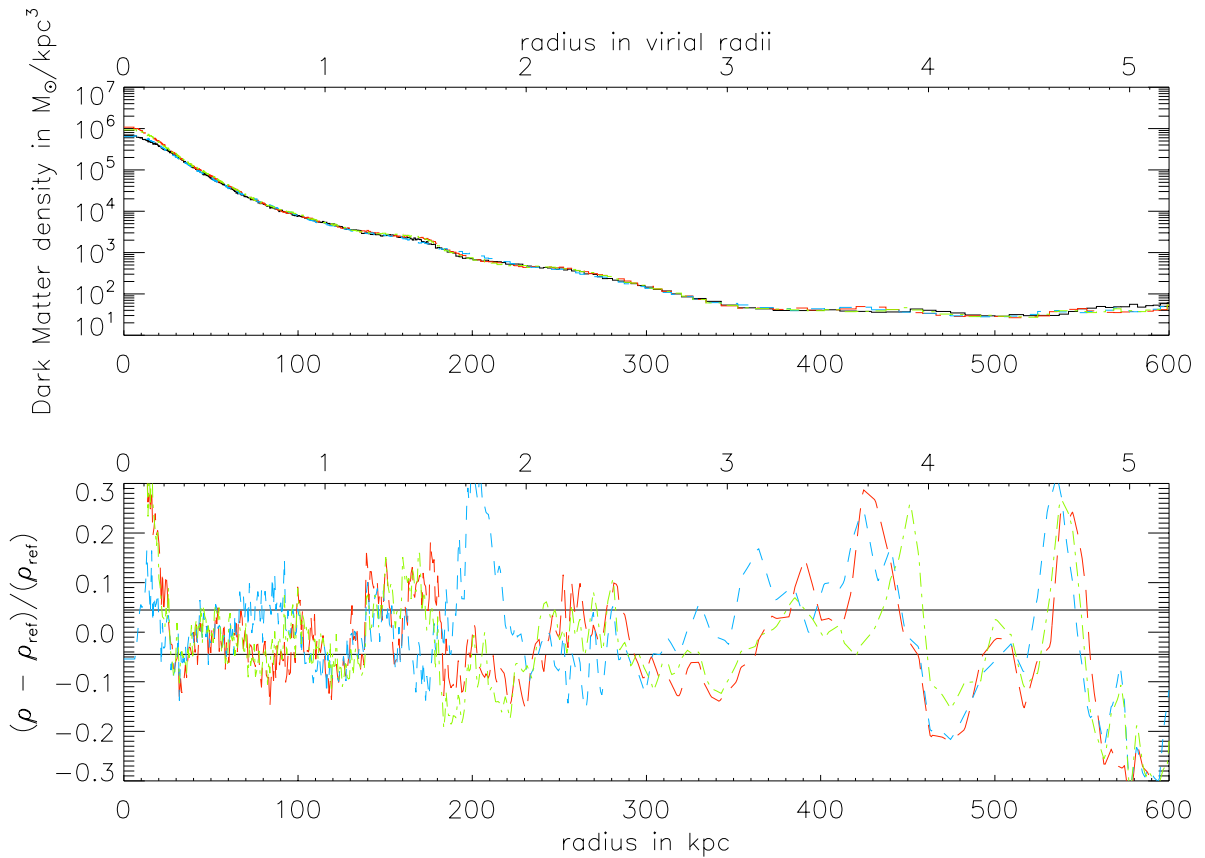


Figure 2.4: Moving Region refinement comparison of runs 1-4 (see section 2.4.4), showing dark matter density as a function of radius of each simulated halo. Profiles are binned radially with an equal number of samples in each bin. The black line shows the static resimulation (run 1). The green line shows a moving region (run 2), with the same initial conditions as the static resimulation. The red line shows a moving region (run 3) with a smaller radius than run 2. The blue line shows the smallest region (run 4), representing the limit of the volume necessary to capture the trajectory of the dark matter particles. The lower plot shows the error with respect to run 1, with the horizontal black lines representing 1 sampling error in each bin.

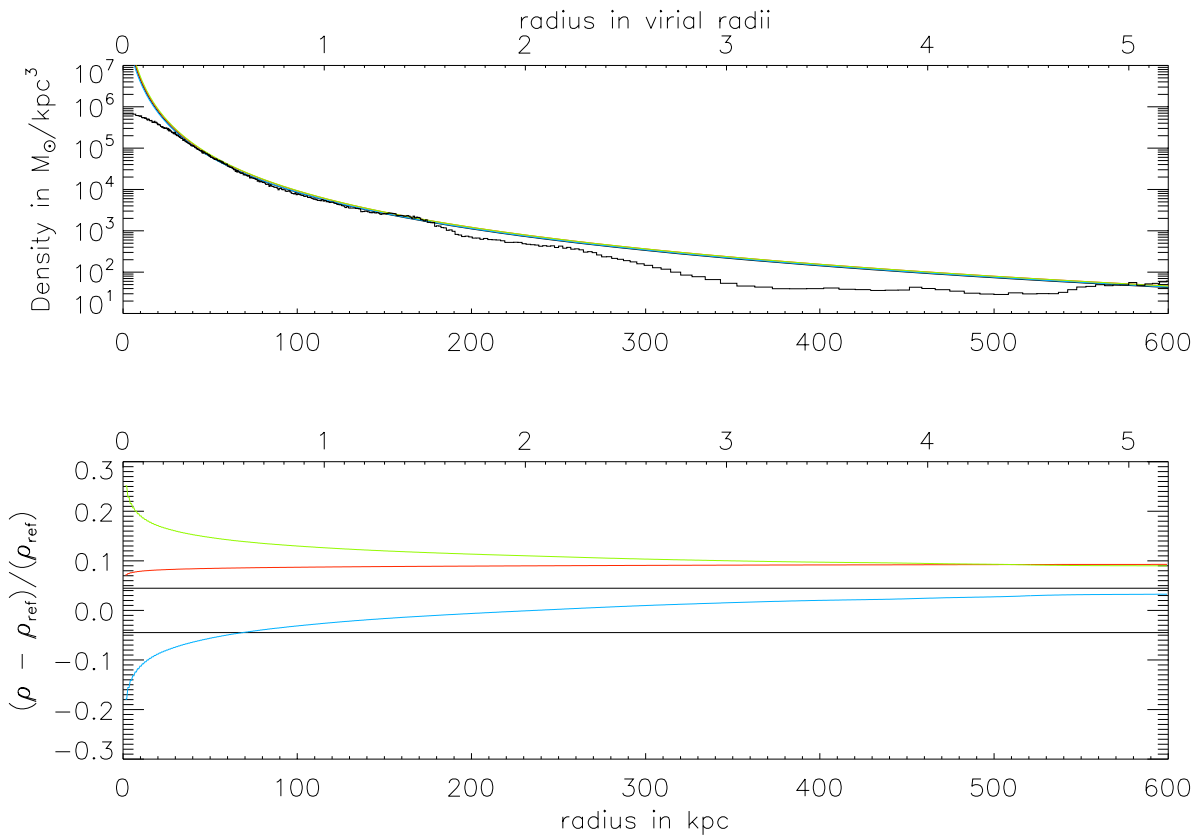


Figure 2.5: Navarro-Frenk-White (NFW) profiles fitted to the radial dark matter density distribution in figure 2.4 (equation 2.43). The profiles are fitted for radii between $[0.4, 1.5]r_{\text{vir}}$ using values from table 2.10. The run 1 dark matter density profile is overplotted as solid black line. See figure 2.4 for a description of the data.




Figure 2.6: Gas density profile as a function of radius. See figure 2.4 for a description of the data.

rates of the matter inside the dwarf galaxy’s halo from the initial conditions to $z=0$. To do this, we capture the dwarf halo’s position and size at each timestep, and define a track which we interpolate along to find the resimulation region’s coordinates at a given time. In the test cases described here, we sample the halo finely in time and use linear interpolation between these sampled outputs.

In this section, we compare “moving region” resimulations to a static resimulation in order to determine whether unacceptable numerical errors (e.g. a deviation of 3σ from the static-region control run) were encountered. Static resimulations have small numerical errors around the boundaries of the refinement region due to the discontinuities with the lower resolution grid that it is embedded in. These errors slowly creep into the high resolution region but if the resimulation region is large enough they do not interact with the object of interest. However, for fast-moving regions, the errors may sweep into the object, and depending on the size of these errors this may be problematic.

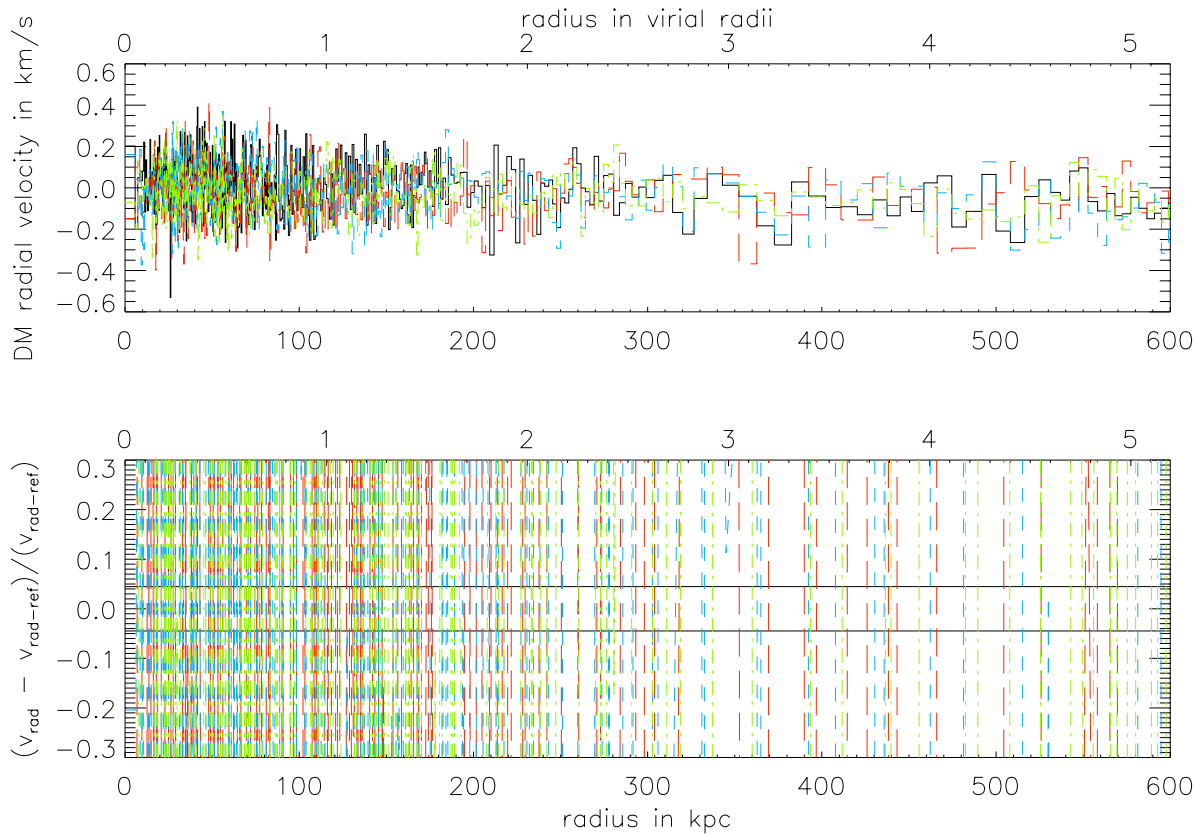


Figure 2.7: Dark matter radial velocity profile as a function of radius. Note that a positive value indicates outflow, and a negative value indicates inflow. See figure 2.4 for a description of the data.

In this study, we identify a halo of mass $10^{11}M_{\odot}$ with a virial radius of 113kpc at $z=0$ in a $10h^{-1}$ Mpc cubic volume. We then run four simulations with different resimulation methods:

- **Run 1** - A static resimulation capturing the whole trajectory of the galaxy and its host halo's dark matter particles.
- **Run 2** - A moving resimulation with the same initial refinement region as the static resimulation; this captures the region that the halo moves in with full resolution at all times as in Run 1, hence acting as a measure of the effect of nonlinearity on the halo.
- **Runs 3 & 4** - Two moving resimulations of decreasing refinement region diameter, designed to test the anticipated numerical artefact problem in the region edges described above. In both these cases, the position of the galaxy at $z=0$ is not inside the refinement region in the initial conditions.




Figure 2.8: Gas radial velocity profile as a function of radius. Note that a positive value indicates outflow, and a negative value indicates inflow. See figure 2.4 for a description of the data.

Each of these simulations are otherwise identical. The simulations were run with a 128^3 base resolution and a 512^3 -equivalent maximum resolution inside the resimulation region with a 256^3 -equivalent buffer region. The refinement was applied to both the dark matter particles and gas cell fields in the initial conditions. No AMR was used inside the refinement region, in order to make comparison clearer. Test cases with and without cooling were taken. Only the cooling runs are presented in this report as they represent the more realistic case. Equally, I did not include metals, UV background or star formation, again for ease of comparison, as well as lower computational expense. The simulations were run on the Oxford Supercomputer Centre's AMD x86-64 MPI Cluster using 20 processors and taking between 70 to 110 CPU hours per run.

I then take the dwarf galaxy's density profile, radial velocity profile and gas temperature profile, and compare them between the simulations. Each radial bin in the profile contains the same number of grid cells / dark matter particles to provide a constant sampling error. Binning is necessary as while individual




Figure 2.9: Gas temperature profile as a function of radius. The horizontal black line in the top plot is the minimum temperature in the cooling table. The virial temperature is 2.25×10^5 K. See figure 2.4 for a description of the data.

Run	$\rho_a(\text{M}_\odot/\text{kpc}^3)$	R(kpc)
1	1.034×10^8	4.480
2	1.866×10^8	3.784
3	1.056×10^8	4.583
4	5.063×10^7	5.752

Figure 2.10: Table of NFW profile parameters in equation 2.43, as fitted for each run. See section 2.4.4 for a description of the runs.

structures are convergent, the individual fluid elements are subject to the equations of fluid dynamics and gravity, which are inherently nonlinear (Thiébaud et al., 2008; Macciò et al., 2007).

The results for $z = 0$ are presented in figures 1-5. For the full suite of test runs (runs 1 through 4), the graphs show the radial variation of dark matter and gas densities (figures 2.4, 2.6), dark matter and gas radial velocities (figures 2.7, 2.8), and gas temperature (figure 2.9), with fractional differences in each bin compared against the static resimulation (run 1). The figures show good agreement in most cases, with the exception of the dark matter radial velocity, which are very noisy, since it tracks the random motion of particles in the halo.

The dark matter density plot (fig. 2.4) shows good agreement in the logarithmic curve plot, and the results are generally within a few sampling errors of the static resimulation. We fit a Navarro-Frenk-White (NFW) profile to the dark matter halo (Navarro et al., 1996):

$$\rho(r) = \frac{\rho_a}{\frac{r}{R}(1 + \frac{r}{R})} \quad (2.43)$$

where R is the scale radius and ρ_a is the (density) constant of proportionality. The results are tabulated in figure 2.10. The values are in reasonable agreement, but it should be noted that the agreement between the parameters depends heavily on the range over which the profile is fitted - note that in figure 2.4, the difference between a given run's density profile and the static run (run 1) varies with radius. The quality of the fit as seen on figure 2.5 is good, with the smallest moving region (run 4) providing (surprisingly) the best agreement with the static resimulation region (run 1), despite the values given in table 2.10 for run 4 providing the worst agreement with the run 1.

The error in the gas density profile when compared to the static resimulation (fig. 2.6) is much noisier than the dark matter profile's error, spiking intermittently. Looking at the gas density profile itself, all

the runs have some inherent noise and trace the same profile. Rebinning the gas with a larger number of samples per bin reveals some small systematic differences between the static run (run 1) and the other runs between 1.5 and $2 r_{\text{vir}}$. However, a larger number of samples per bin improves the agreement inside r_{vir} , and dramatically decreases the spiking in the residue (lower) plot. In the gas we see disc formation, and hence the density at the same radius can vary dramatically depending on whether we are sampling inside or outside the disk. Upon inspecting the galaxies formed, the orientation of the gas disks appears not to be convergent. This does not appear to materially affect the results above, however. Interestingly, the central dark matter density in the smallest moving region test (run 4) agrees best with the central dark matter density in the static resimulation test (run 1). However, this effect is not seen in the gas density profile, and is likely to be subject to some nonlinearity.

The dark matter radial velocity is shown in fig. 2.7. All the runs have completely virialised halos, in that there is no notable deviation from the 0km/s line. The profile appears noisy, but the noise amplitude is under 0.5 km/s , which is much lower than values found for typical cosmological velocities. Further, one would expect a self-gravitating particle cloud with negligible bulk flows to contain a statistical distribution of velocities that appears random when plotted radially. The gas radial velocity (fig. 2.8) shows considerable inflow of matter (note that a negative radial velocity denotes inflow, positive denotes outflow). For the gas radial velocity, there is good agreement, to within the sampling error in nearly all cases.

The gas temperature values (fig. 2.9) are in agreement as well. All values lying within two sampling errors of the static resimulation values for two virial radii despite the considerable scatter in the results. Again, this scatter is likely to be due to non-spherically-symmetric structures such as disks and cloud and/or filamentary accretion. Note that the peak of the temperature curve is where the efficiency of cooling balances the efficiency of gravitational heating of the gas. The temperature outside 2-3 virial radii is below the minimum cooling temperature – this is because the gas has not yet been gravitationally heated, and so does not need to cool to have a lower temperature. The virial temperature ($2.25 \times 10^5\text{K}$) is too high to be shown on this plot, as radiative cooling reduces the gas temperature in the centre of the galaxy to below this value.

Overall, the agreement between the curves is good, with variables within 1-2 virial radii agreeing to within 2-3 sampling errors or less (excluding the dark matter radial velocity, which is near-zero and therefore the residual is dominated by the scatter of the data). Given the non-linearity of both the N-body and Euler

equations, this is reasonable, and so this method appears to provide robust results that are not dominated by numerical error.

As mentioned earlier, this method means that we only have to resimulate, in the case of the smallest region, 30% of the volume required for a static resimulation. For a smaller halo that may move up to 1 Mpc in co-moving coordinates through its lifetime, this will be a much smaller value, and so the benefits to the computational time may be significant. The actual improvement in computation time depends on factors such as the resimulation region's resolution, and whether by decreasing the size of the region we can avoid wasting resources by resimulating objects that drift close to the galaxy being studied. In the test case studied, there was a factor of two improvement in total computation time between the static resimulation and the smallest moving resimulation.

2.4.5 Tracer Fluid

A final method is introduced in Hahn et al. (2010). Here, the authors refine only in AMR cells with a non-zero metallicity. Metals are only required inside the galaxy being resimulated; the rest of the cosmological volume is retained only to provide the correct gravitational field inside the refinement region. Hence the authors introduced metals into the initial conditions such that the metal distribution traces the movement of the halo to $z=0$, as described in the previous sections. One problem with this model is that if outflows are significant, metals will pollute the IGM, allowing refinement in these regions. Kirby et al. (2011b) estimate that the Milky Way dwarf spheroidal (dSph) satellites lose 96-99% of the metals they produce. However, this is only an issue if these metals pollute other halos.

2.5 Instabilities, Odd-even Decoupling and Carbuncles

One of the benefits of Godunov solvers is that they represent flows as a series of piecewise states. This means that discontinuities can be represented reasonably accurately (see section 2.2.3). However, they are also prone to numerical errors. Quirk (1994) discusses a number of issues with Riemann solvers. In particular, the author discusses a grid-induced artefact called the “carbuncle”, an unphysical bulge that appears in numerical simulations of curved shocks where the direction of the flow aligns with that of the shock, and “odd-even decoupling”, in which a significant axis-aligned instability develops from fluctuations such as

numerical round-off errors in the hydro variables. In lower resolution simulations, the large spacing of the grid cells acts to damp the instability. The instability manifests itself as a growing sawtooth instability in the shock boundary, growing into a finger-like structure. Quirk (1994) proposes switching to a more diffusive solver on shock fronts. This is because diffusive solvers suffer from extra numerical dissipation, which acts to damp these instabilities. To do this, the local pressure gradient $|p_1 - p_2|/\min(p_1, p_2)$ is measured, and if it falls above a certain value, the code switches to the more diffusive solver. Unfortunately, there is no rigorous method for finding this threshold, and it must be optimised for each problem. However, the fix is simple, since it requires only a switch between two solvers already existing in the code. This method is implemented in a simulation of a core-collapse supernova by Kifonidis et al. (2003). Since Quirk (1994), a number of other other solutions have been presented. A review of some of these is given by Pandolfi & D'Ambrosio (2001). Ren Y.-X. (2003) suggests a technique in which the axes in which the Riemann problem is solved are rotated, such that axis-aligned instabilities are avoided, without using a more diffusive solver such as the HLL method. The axis of rotation is chosen to be aligned to the vector defined by the difference in velocities across the grid cell interface. Another method given in Nishikawa & Kitamura (2008) is to switch between diffusive and non-diffusive solvers along this axis of rotation, without using the free parameter found in Quirk (1994). Huang et al. (2010) find that both these methods give good results, but that the hybrid method is more computationally efficient. One further approach is to use a method not based on solving the Riemann problem. For example, the Bhatnagar-Gross-Krook (BGK) scheme presented in Prendergast & Xu (1993), is based on solutions to the Boltzmann equation. This method has recently been implemented in RAMSES and is in the early stages of testing.

The following section details work to RAMSES to use the Quirk (1994) method for use in single supernova simulations, due to its simplicity and robustness (it requires no changes to the Riemann solvers already present in the code). The examples are from simulations run in preparation for chapter 5. We model a $15M_{\odot}$ star evolving in a uniform medium with density 100 atoms/cm^3 and temperature of 10K. The simulation volume is a cube of length 100pc with maximum grid resolution 0.1pc (corresponding to 10 levels of refinement), and includes metal line cooling. We chose this resolution as simulations with 8 levels of refinement show no sign of generating instabilities in the wind or supernova remnant. The grid is refined on density, pressure and velocity such that if the difference between the values of any variable either

side of a cell boundary is more than 20%, the grid is refined up to the maximum level. The total energy released into the ISM by the wind is 10^{48} ergs and 1.2×10^{51} ergs for the supernova event at 14.125 Myr. The wind is deposited into the medium as thermal energy (Weaver et al., 1977) in a sphere of radius 20 cells (2 pc at 0.1 pc resolution) according to Marigo et al. (2008). Full details of the stellar model used can be found in chapter 5. In these tests we calibrate the hybrid solver to change from HLLC to HLL when the pressure gradient is over 0.5% (calibrated on a simulation with uniform wind speed and no cooling, such that any kinks in the wind sphere caused by grid-aligned numerical instabilities were damped). We perform four simulations with identical physics but different solvers: pure HLLC, hybrid with a threshold of 50%, hybrid with a threshold of 0.5% and HLL. We call these simulations HLLC, Hybrid50, Hybrid0.5, and HLL respectively.

Figures 2.11 to 2.13 show the evolution of the stellar wind bubble prior to the supernova as a density slice through the centre of the box in each of the simulations, while figures 2.14 and 2.15 show similar images of the subsequent supernova explosion in two snapshots. During the evolution of the wind bubble, the HLLC and Hybrid50 simulations appear to be very similar. Both simulations exhibit instabilities on the surface of the bubble of shocked interstellar wind Weaver et al. (1977). Even the Hybrid50 simulation generates similar instabilities by 14 Myr, though the boundary between the shocked stellar wind and shocked interstellar gas is not well resolved due to the diffusivity of the HLL solver. Stellar wind shocks are typically unstable in certain regimes, which we discuss in section 2.6. All simulations exhibit some form of convective instabilities inside the wind bubble, though the HLL simulation remains oddly symmetrical in the grid throughout the simulation. The density profiles in figure 2.16 during the stellar wind phase appear similar in all simulations, with the density peak appearing to be a little sharper in the HLL simulation due to the greater radial symmetry of the wind bubble. The temperature profiles also appear to be roughly similar, though the temperature at the wind shock is lower in the HLL simulation than the other runs, with the Hybrid runs both exhibiting spikes in the temperature profile inside the wind bubble. The evolution of the radius of the wind bubble prior to the supernova event at 14.125 Myr in figure 2.18 is also very similar in all runs. A curious step function is visible here; this is due to the winds preferentially “inflating” the wind bubble in the direction of the grid axes, before the rest of the wind bubble inflates, returning it to a spherical shape. This occurs in all simulations, though the effect is weaker in the HLL simulation.

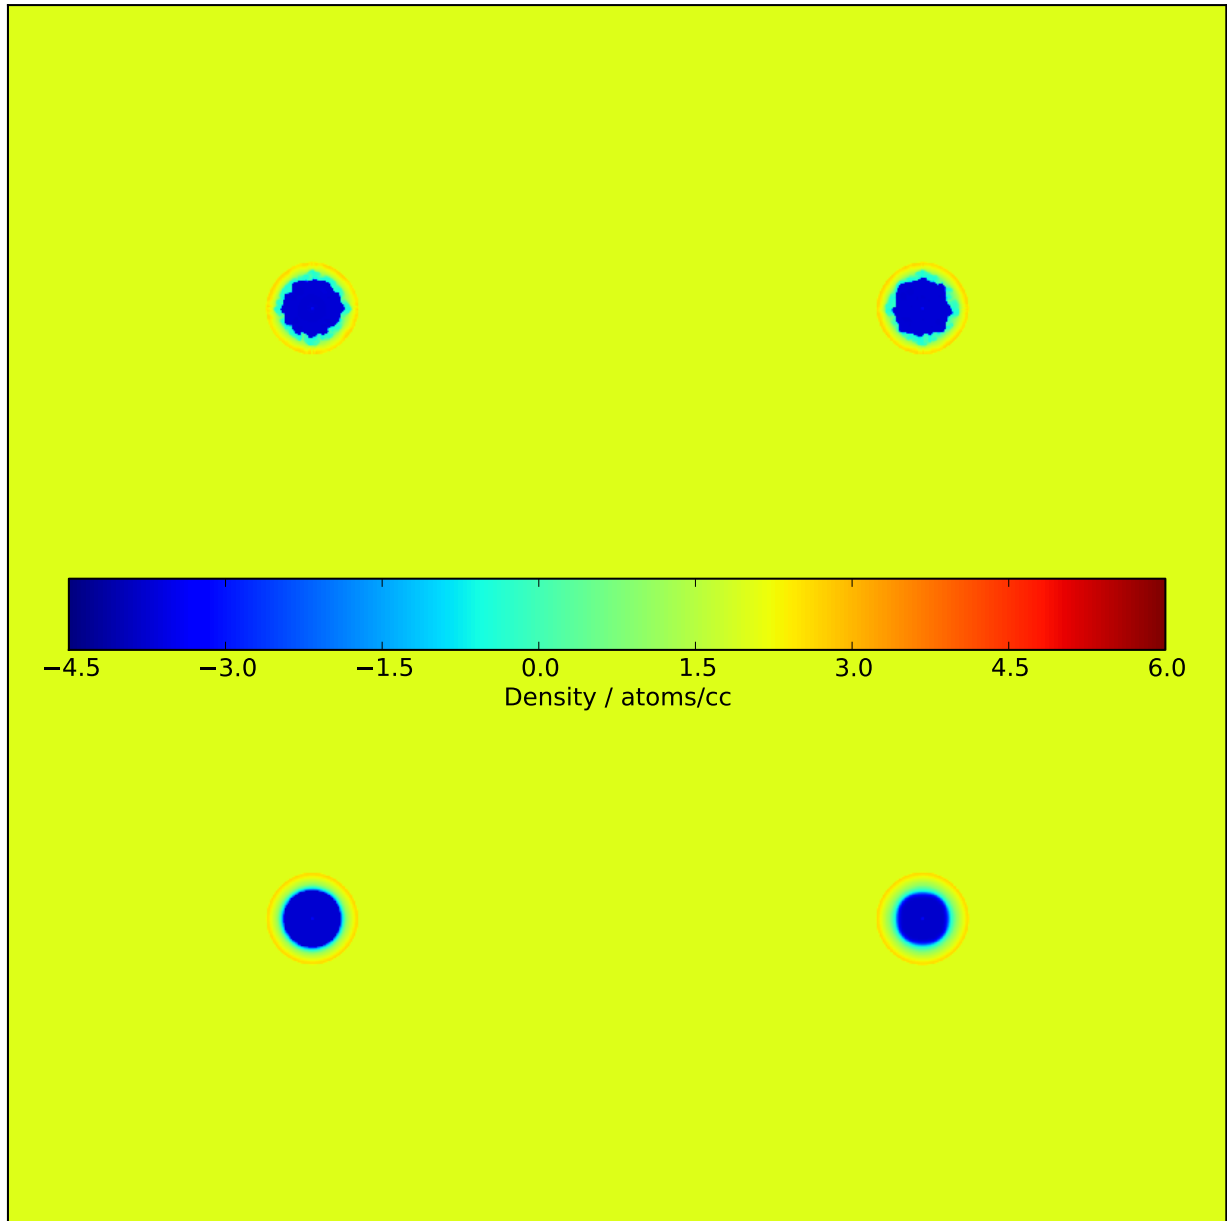


Figure 2.11: Four density slices at 5Myr through the simulation volume of a $15M_{\odot}$ star evolving in a uniform medium of density 100 atoms/cm^3 and temperature of 10K, demonstrating the effect of the hybrid solver on the evolution of the wind. Each of the images are 100pc by 100pc. From top left to bottom right, the images are of simulations HLLC, Hybrid50, Hybrid0.5, and HLL.




Figure 2.12: As in figure 2.11, but at 10Myr




Figure 2.13: As in figure 2.11, but at 14Myr, just prior to the star going supernova at 14.125Myr




Figure 2.14: As in figure 2.11, but at 15Myr, after the star goes supernova at 14.125Myr




Figure 2.15: As in figure 2.11, but at 20Myr




Figure 2.16: Radial gas density profiles of the CSM up to 50pc in each simulation, from 5Myr to 20Myr. The star goes supernova at 14.125Myr. From top left to bottom right, the profiles are of simulations HLLC, Hybrid50, Hybrid0.5, and HLL. The high densities in the centre of the hybrid figures at 15Myr (red) are the result of a reverse shock from the supernova bouncing off the wind bubble's shell.

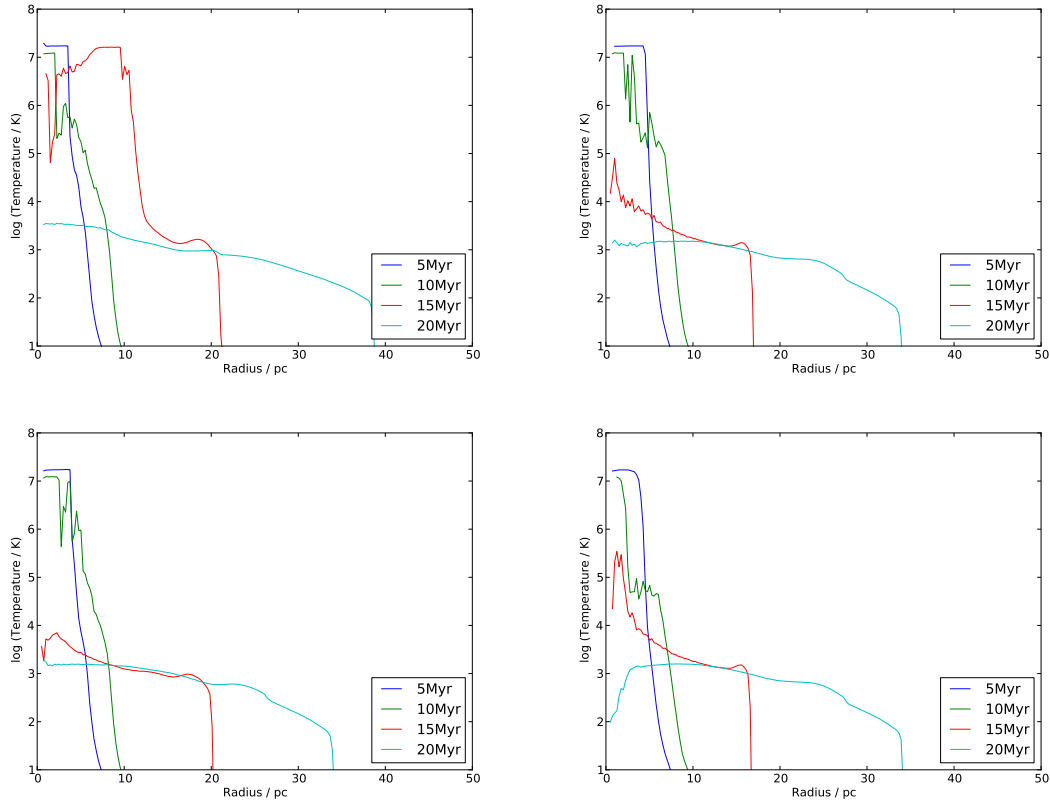


Figure 2.17: As in figure 2.16, but for gas temperature.

After the supernova event, the results seen in figures 2.14 and 2.15 diverge. Part of this is due to slight variations in output time; RAMSES output times vary slightly due to the varying timestep of the simulations. As a result, the fast movement of the shock results in small differences in the timestep causing large differences in the position of the shock, in particular the reverse shock filling the wind bubble. The HLLC simulation at 15Myr retains its hot, low-density wind bubble, while in the Hybrid and HLL simulations the wind bubble is filled by the reverse shock; in the Hybrid50 simulation, a high-density peak is formed. The slower reverse shock in the HLLC simulation can be achieved by triggering convective instabilities inside the shock, which removes energy from the radial kinetic energy of the shock and hence slows its growth until it can be converted back into thermal energy. Another possibility is that the higher thickness of the wind bubble shell in the HLL and Hybrid simulations seen in figure 2.13 may affect the evolution of the reverse shock as the supernova shock impacts the wind-driven shock. The HLLC simulation exhibits grid-aligned Rayleigh-Taylor instabilities manifesting themselves as bubbles on the shock front, similar to those seen in

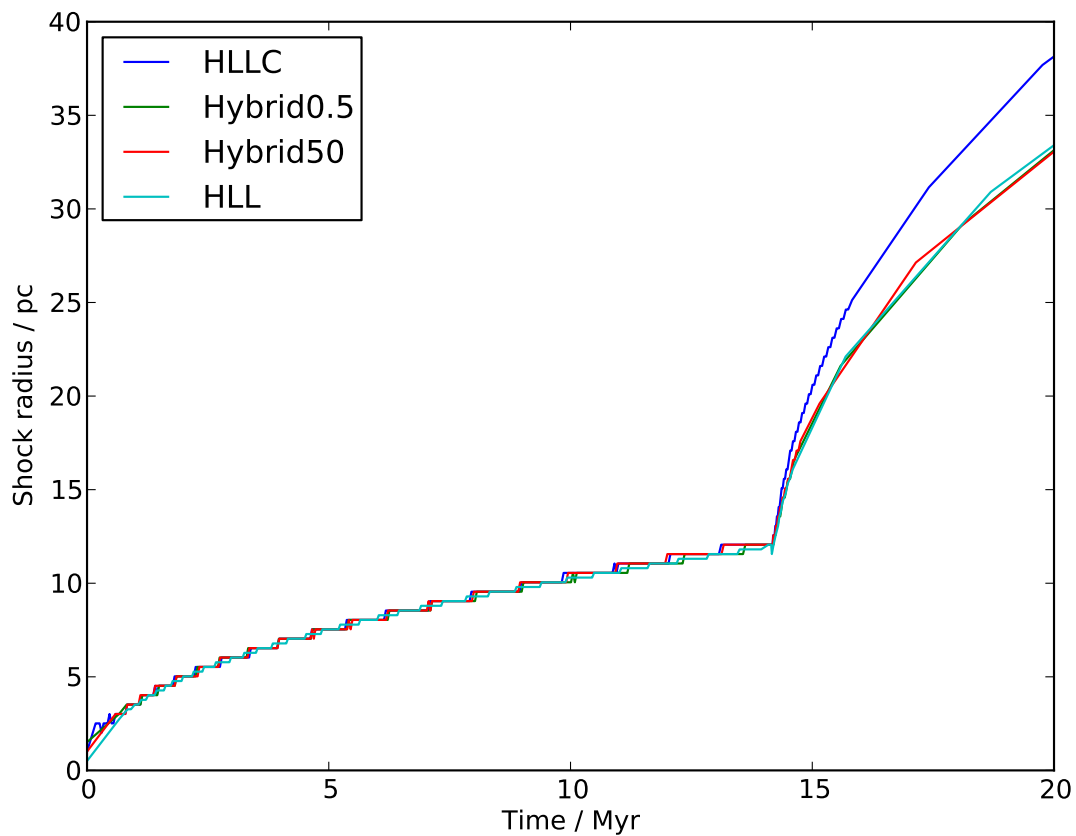


Figure 2.18: The evolution of the wind and subsequent supernova shock fronts in each simulation. The shockfront is measured by locating the point of highest pressure, ignoring reverse shocks; in one output in the HLL simulation, the reverse shock has a higher pressure than the outwards shock. The steps are due to the stellar winds preferentially “inflating” the wind bubble in the direction of the grid axes, before the rest of the wind bubble inflates, returning it to a spherical shape. The Hybrid and HLL simulations have a lower supernova shock radius than the HLLC simulation; this is because the shock front in all of these simulations is captured by the HLL solver, which is more diffusive than the HLLC solver.

Fraschetti et al. (2010), who also use a pure HLLC solver. This is a by-product of odd-even decoupling, and is not found in the Hybrid or HLL simulations.

One of the most worrying features of these plots is that the supernova blastwave in the Hybrid and HLL simulations has a shock with a smaller radius than the pure HLLC simulation, as seen in figure 2.18. At 20Myr after the formation of the star, the radii differ by approximately 10%. As Toro et al. (1994) demonstrates, the HLL method is more diffusive on shock boundaries than the HLLC method, resulting in a gradual smearing of the shock over time. The HLLC method, by contrast, captures a stationary shock almost perfectly, and exhibits significantly less smearing for a moving shock. One possible cause of this lies in the way gas cools radiatively, which is roughly proportional to density squared (see section 2.3.2). A loss of thermal energy from cooling will result in a much smaller supernova bubble. As noted above, the HLLC run retains its low-density wind bubble longer than the others, which fill it with a reverse shock (evidenced by higher central densities in figure 2.16 at 15Myr (the red line)). As a result, the outward expansion of the shock may be delayed long enough to allow the supernova remnant to preserve more of its radiative energy than the Hybrid or HLL simulations.

A higher resolution in the HLL run will resolve some of these issues, though this obviously incurs further computational cost; due to measuring the gradients in temperature, density and velocity, a large number of the cells in the volume will become refined, particularly on the surface of shocks. Increasing the resolution will trigger refinement in all these cells, so each time we add a level it will add $8(=2^3)$ cells in each of these zones, as well as halving the timestep due to the Courant condition.

It appears that the supernova shock, being both instantaneous and more energetic than the stellar wind, has a large pressure gradient and hence triggers the HLL solver in the Hybrid simulations with even a high threshold in pressure gradient. By contrast, the expanding wind bubble does not appear to trigger the hybrid solver with a threshold in pressure gradient of 50%. Hence it is difficult to tune the Quirk (1994) solver for this particular problem. Moreover, the effects of odd-even decoupling itself appear to be weak. The HLL solver introduces significant diffusivity into the simulation, while all of the simulations, including HLL, exhibit grid-aligned artefacts and a step-function caused by preferential inflation of the wind bubble along the grid axes. Indeed, the HLL simulation retains symmetry with the grid to a greater extent than any other of the simulations, since its dissipative nature damps the instabilities that break up these symmetries. As




Figure 2.19: Density slice through the centre of the wind bubble of a $15M_{\odot}$ star in a medium of $n_{\text{H}}=100$ atoms/cm³ and $T=10\text{K}$ at 5Myr. The image is 50pc in length on each side. I use the HLLC solver with a Courant number of 0.4, and seed the medium with density perturbations of $\pm 0.5\%$. The simulation includes metal-line cooling. The wind bubble exhibits grid-symmetric instabilities, which grow as the wind velocity and mass loss from the star increases. Eventually, Kelvin-Helmholtz instabilities on the shock boundary will break up the grid-symmetric structure. The supernova blast will travel preferentially along paths of low density, creating a non-spherical blast wave.

we discuss in the next section, many of these instabilities do indeed have a physical origin. Finally, the HLL and Hybrid solvers may require a higher resolution to capture the correct shock radius, which presents difficulties for simulations that have this purpose.

In the next section, we discuss other sources of numerical instabilities and attempts to separate physical from non-physical instabilities.

2.6 Shock-Induced Instabilities

In the previous section, I discussed the effect of odd-even decoupling on a stellar wind and subsequent supernova shock. I noted that the hybrid solver suggested by Quirk (1994) can remove the grid-aligned

instabilities seeded by odd-even decoupling. This method is attractive because it is simple to implement and requires few changes to existing solvers, since it is simply a switch between two methods. However, it also has a number of drawbacks. Firstly, it involves setting a free parameter that defines the pressure gradient threshold for which the code switches from the odd-even-decoupling-prone HLLC solver to the more diffusive but more stable HLL solver. There is no good way of setting this threshold other than manually iterating through test runs until a value is found. Even then a simulation with both stellar winds and supernovae may require different thresholds for different epochs in the star's evolution. Every time the simulation set-up is changed, this process must be repeated. Secondly, the shock radius is smaller in the runs containing HLL. This is because HLL is more diffusive, and is worse at capturing shocks than HLLC Toro et al. (1994). This is important for a study investigating the properties of supernova blastwaves. Thirdly, the increased diffusivity of the HLL method means that physical instabilities will be reduced or absent in regions where it is turned on. Fourthly, none of the test simulations we run avoid grid-aligned numerical artefacts; indeed, the HLL solver retains these to a greater extent than the HLLC-HLL hybrid or pure HLLC simulations. While in section 2.5 we listed a number of alternative methods that avoid these drawbacks, they are more difficult to implement, and hence we do not attempt this for the purposes of this thesis. Finally, certain instabilities are indeed physical - this is part of the motivation for running in 3D rather than using a spherically-symmetric 1D model.

There are many types of instabilities that occur in different regimes, but in this thesis we treat three principal types. The first type are Kelvin-Helmholtz instabilities, which are triggered by shear forces between two fluids moving in opposite directions. These can be seen in, for example, the counter-rotating zones and belts in Jupiter's atmosphere. The second are Rayleigh-Taylor instabilities, which can occur when a denser fluid lies above a lighter fluid (note that "above" in this context can refer to the local acceleration field, e.g. in a shockwave). A classic example of this is in a mushroom cloud. Thirdly, and relevant to the supernova and wind shocks in chapter 5, we have "Vischniac" instabilities. Vishniac (1983) describes a type of physical instability on a shock surface in which gravitational and/or thermal processes drive the growth of oscillations on the shock front. These types of instabilities are often called "overstabilities", since they are driven in part by restoring forces that are stronger in places than the force driving the oscillations. In this case, the restoring force is the pressure differential inside the shock and the reaction force from the

shock impacting the stationary external medium. These instabilities are discussed in more detail, along with formulae for growth rates, in the context of the stellar wind simulations in section 5.4.2.

These instabilities can all occur in the context of stellar wind and supernova shocks (Vishniac, 1995). Kahn (1980) determines that Kelvin-Helmholtz instabilities will grow on the interface between stellar winds and the interstellar with a sufficiently sharp interface. However, they also note that second-order effects should damp these instabilities in the absence of strong cooling. Mac Low & Norman (1993) confirmed the existence of such instabilities in 2D numerical simulations of an idealised shock. Vishniac (1994) also confirms the existence of such instabilities in the non-linear regime as well as the linear regime a discussed in the previous paper. More recently, Kushnir et al. (2005) and Ramachandran & Smith (2005) have created models that attempt to parameterise shock instabilities using perturbation and linear stability analysis.

Dgani et al. (1996) further find the existence of “transverse acceleration instabilities” in bow shocks from linear stability analysis and nonlinear numerical models. Smith & Rosen (2003) note that Gaetz et al. (1988); Innes et al. (1987) expect atomic cooling in fast shocks greater than 140km/s at 10^5 K should generate instabilities, the production of CO and H₂O in molecular gas on the shock front can also drive instabilities. Nishi & Kamaya (2000) note that there exists an interplay between the self-gravity of the shock versus the shock deceleration, such that a “decelerating shock instability” due to shock deceleration dominates on a small scale, while at the late stages of the shock evolution, self-gravity can drive larger-scale instabilities on timescales comparable to the freefall time of the external medium. Williams (1999) notes that the ionisation front from massive stars can drive instabilities as well, due to heating the CSM and driving the expansion of a high-temperature bubble. Shocks from supernovae have also been found to be unstable to convection: Gull (1973) determines using linear stability analysis of 1D numerical models that there should exist convective zones that trigger Rayleigh-Taylor instabilities behind the shock front in supernova shocks, as observed in models by, e.g. Dwarkadas (2007) and Fraschetti et al. (2010).

One problem with modelling these instabilities in numerical simulations on Cartesian grids is that the grid can imprint itself on the structure of the instabilities. The stellar wind bubble in Cartesian grid codes will sometimes form diagonal spike-shaped instabilities that are symmetric in the grid, as in figure 2.19. This is caused by the continual injection of wind into the CSM by the star; if a small instability is triggered on the grid, the subsequent winds will follow this, and push out into the medium around the shocked gas.

Krause et al (in prep.) determine that these instabilities are physical, even though they develop preferentially along grid axis when using a Cartesian grid, by measuring their growth against the growth rate of Vishniac instabilities given in Vishniac (1983). The authors note that these fluctuations grow similarly to Vishniac instabilities. The authors note that although the phenomena they observe appear to be aligned with the grid and hence appear to be a numerical artefact, they are triggered by genuine physical processes.

One possible way to reduce the impact of the grid on shocks is to seed the initial grid with random noise, such that small fluctuations in the shock imprinted by the grid are damped out while not destroying the sphericity of the shock. Again, this needs to be tuned to achieve the balance required. Another role of this kind of noise is to seed instabilities in the shock that would be formed physically if sufficient resolution was achieved. I find in one test simulation that adding a uniform distribution of random noise to the density field between $\pm 5\%$ of the average density of the medium begins to break up the grid-aligned shocks into more randomly oriented structures, though the imprint of the grid is still visible at this level. We do not use this noise in the simulations shown below. The reasons for this are twofold. Dwarkadas (2007) and Fraschetti et al. (2010) find that they do not require noise to seed instabilities, as they achieve sufficient resolution to capture their initial evolution. Secondly, as we have argued above the instabilities we observe are physical, even if they are symmetric along the grid axes. Thirdly, these simulations represent highly idealised conditions. In media dominated by turbulence or power-law fluctuations such as found in the ISM, the results are not expected to be spherically symmetric, and the high density contrast of the fluctuations will damp out any small grid-seeded instabilities.

A final, simple way to reduce (but not remove entirely) numerical instabilities is to reduce the Courant number, which in RAMSES is a free parameter. A typical value is 0.8, but a smaller value will create more stable flows, since the flow through an interface will be better captured if the timestep is reduced. The trade-off is that the simulation will take more computational resources to complete. In the following simulations we use a Courant number of 0.5.

2.7 Conclusions

In this chapter I have reviewed the techniques used in RAMSES to model astrophysical fluids, and in particular the formation of dwarf galaxies in cosmological hydrodynamic simulations. I also present techniques

for optimising the computation time for resimulations of dwarf galaxies. I then discuss issues with the Riemann solvers currently found in RAMSES and possible solutions, in order to illustrate potential sources of error in numerical simulations of shocks. Note that this is not a complete summary of the physical processes modelled by RAMSES. Models exist, for example, for magnetohydrodynamics (MHD) Dubois & Teyssier (2008a) and active galactic nuclei (AGN) (Dubois et al., 2012). We do not discuss either of these processes in this thesis. AGNs are thought to be important in larger galaxies, where the effects of quasars act to heat the gas inside a galaxy and quench star formation. In Milky Way-sized galaxies and below, AGNs play a smaller role in the energetics of a galaxy (see chapter 1).

The role of magnetic fields, however, can be important in galaxy formation. Hill et al. (2012) find that magnetic fields do not dramatically change the vertical structure of a disk galaxy, since magnetic pressure and magnetic tension force roughly cancel each other out. However, they do find a reduction of $\sim 8\%$ in the amount of cold gas in their simulations. In simulations of supernova remnants, Gull (1973) determine that $\sim 1\%$ of the kinetic energy from a supernova is funnelled into instabilities, and this generates a magnetic field of around 10^{-7}T . This is because the Maxwell equations state that the rate of change in a magnetic field is proportional to the curl of the velocity field of charged particles, and hence fluid turbulence will act to generate magnetic fields. Various reasons exist for not including magnetic fields in our simulations. The first is banal; merging versions of RAMSES (such as the radiative transfer and MHD modules) takes time, and both modules are fairly new and subject to continuing development. The second is that we often do not have good information about the expected initial magnetic fields in our simulations, and the evolution of a magnetic field is sensitive to its initial state. Thirdly, in order to determine the precise effect of magnetic fields in simulations, we must compare with control experiments containing no MHD. Hence the non-inclusion of MHD in initial studies is not a catastrophe, though future work should attempt to include it.

In the next chapter I discuss procedures used to analyse and visualise cosmological simulations of galaxy formation. I use these techniques to present a test case for using simulations to produce mock observations from the Wide Field Camera 3 (WFC3) on the Hubble Space Telescope.

Chapter 3

Analysis & Visualisation

The analysis of astrophysical data is an important part of the scientific process in astronomy. In this chapter I discuss the conversion of astrophysical simulations from raw data to figures in a finished publication. The first section of this chapter is devoted to discussing the techniques used to convert simulation outputs into halos and halo merger trees, and compare simulations with the same initial conditions but different physics. I discuss the application of these techniques on the NUT suite of cosmological hydrodynamic simulations of a Milky Way-class halo and its satellites. In the second section of this chapter, I discuss two software packages written during the course of this research to analyse and visualise astrophysical data and the design patterns that underpin them.

3.1 Halos, Merger Trees and Parallel Universes

3.1.1 Finding Halos with HALOMAKER

The first task is to convert the array of discrete fluid elements into objects corresponding to dark matter and gas halos, galaxies and star clusters. To do this, we identify gravitationally bound collections of particles. Dark matter particles and star particles are considered separately, since they cluster in different ways at different scales. We do not apply this algorithm to AMR gas cells, though in principle this would be possible. This process is achieved with HALOMAKER, a code written as part of the GalICS (“Galaxies in Cosmological Simulations”) project (Hatton et al., 2003). There are two principal methods for identifying




Figure 3.1: Schematic illustrating the MSM halofinder (Tweed et al., 2009) in 1D. We first identify peaks in the density field (a,b and c), then build a tree of peaks by identifying saddle points (1 and 2). We truncate the tree at some threshold density (3). The most massive peak (c) is determined to be the host halo, while a and b are subhalos. It is possible using this method to find subhalos of subhalos.

halos in HALOMAKER. The first, “friend-of-friend” (FOF), finds connections between particles using a tree structure (Davis et al., 1985; Huchra & Geller, 1982). The second, “most-massive subhalo method” (MSM) Tweed et al. (2009) is a variant of ADAPTAHOP (Aubert et al., 2004). This method identifies overdensities in the particle density field. In this section I briefly outline the operation of these two methods and their application to identifying structures in cosmological galaxy formation simulations.

The FOF method identifies groups of particles which are closer than $b \langle |\mathbf{x}_i - \mathbf{x}_j| \rangle_{ij}$, i.e. a scalar factor “b” times the mean global interparticle distance. Tweed et al. (2009) note that $b = 0.2$ matches an average halo over-density of $178\rho_c$ (where ρ_c is the critical density, the mean density of the universe assuming zero curvature). They further find that a grouping must contain over 20 particles to be sure of sampling a genuine gravitationally-bound structure rather than particles clustering via random noise, ignoring parameters such as particle velocity or the potential field. The problem with the FOF method is that it is very difficult to select the value of b such that all the required structures are identified. There is no mechanism, for example, to identify substructures inside halos (“subhalos”) in FOF; either these substructures are identified as independent structures or they are determined to be part of the host halo. This becomes a particular problem in chapter 4, when we look at the satellites of the Milky Way.

The ADAPTAHOP method resolves some of these issues by considering the density field at each point and identifying local maxima. The process for identifying halos and subhalos using MSM with ADAPTAHOP is summarised from Tweed et al. (2009) as follows. We illustrate the process in figure 3.1.

1. Find each particle's nearest neighbours; we typically select between 20 and 64 depending on which gives the best results for the data. This is typically done by building a particle mesh with an octree as in 2.2.2. Hence calculate the local density at the particle's position by building a spline using the nearest neighbours; this is equivalent to the "SPH smoothing kernel" (e.g. Springel (2005); Wadsley et al. (2004)).
2. Locate local maxima in the density field by stepping through the particle tree and comparing adjacent nodes. Once this is done, apply a density threshold to all particles using the density calculated previously, and attach all particles above this threshold to their nearest local maximum.
3. Identify saddle points in the density field. These saddle points can be used to join together neighbouring local maxima, creating a tree of maxima. We truncate the tree at the threshold density and sort the tree by saddle point density.
4. This tree now defines a hierarchy of local maxima in the density field that must be mapped to halos and subhalos. We identify every leaf node as a halo or subhalo. We then must attach the matter below the saddle point to a halo. We choose this halo to be the most massive leaf node (hence "most-massive subhalo"). We can hence build a tree of subhalos; this can occur in hierarchical galaxy formation when a halo captures a halo and is then captured by another subhalo. In chapter 4, we find one subhalo of a subhalo that contains a satellite galaxy and survives to $z=0$, while the host subhalo is stripped and destroyed by interactions with the host halo.

HALOMAKER then identifies various physical properties for each halo, such as virial radius, maximum circular velocity, triaxial shape and orientation, etc. We then visually inspect the results of the halofinding and can tweak the parameters used by HALOMAKER in order to remove spurious detections and boost the detection of small halos and peaks close to each other.

Note that in RAMSES, we can perform this process to star particles as well as dark matter particles (in theory, it is also possible to do the same thing to the gas field, though we do not perform this calculation

in the work discussed here). We decouple the two fluids before doing this, however, since stars are formed by cooled gas, and hence cluster on smaller scales than their dark matter host halos, which are supported by their velocity dispersion. In chapter 4 we locate satellite galaxies; here, we identify these as being the galaxies whose host halos are in turn subhalos of the main halo being studied. In most cases, the galaxies do not overlap in the star field. However, MSM can be used to identify star clusters inside a galaxy and study their evolution.

This, briefly, is the method used to identify structures in the particle fields. Halofinding is a large and evolving field; Knebe et al. (2011) lists comparisons between various methods found in the literature. One possible improvement to the above process is to create a density field in 6D (position and velocity) rather than simply 3D, since this helps identify subhalos that happen to be passing through the centre of their hosts.

In the next section we take the outputs of this process and join them together to build halo merger trees, which allows us to study the evolution of galaxies and halos through time.

3.1.2 Merger Trees

Each HALOMAKER output from a given simulation is taken and inserted into two post-processing tools: TREEMAKER and MAKE_TREE (Tweed et al., 2009) (the two tools have since merged into one package). This code produces a *halo merger tree*, which tracks halos through time. This process is not trivial; halos merge and fragment, and deciding which halo in output N corresponds to a given progenitor in output N-1 requires careful consideration. Take the pathological case where, between snapshots, a halo enters another halo as a subhalo and loses most of its mass to its host. If we simply take the halo which contains most of the progenitor's particles, we would choose the host halo and hence conclude that this subhalo has been destroyed by the host. In fact, the halo still exists, albeit in a diminished form. This problem will again be revisited in chapter 4, when we discuss the satellite galaxies of the Milky Way.

Tweed et al. (2009) summarises the process used to identify merger trees as follows. First, TREEMAKER is given a list of outputs 1 to N. It then identifies the particle ID numbers in halos in output I and I-1, where I has a value from 2 to N. For every halo in output I (hereafter H_i^I , for $i=1, \text{numhalos}^I$), it identifies each halo that contains its member particles in output I-1 and lists them according to the fraction of mass that they donate to H_i^I . It also lists the fraction of mass accreted smoothly; i.e. particles present in H_i^I that are not




Figure 3.2: Example of a halo merger tree. The branch on the left is the main halo in the NUT suite in the “Reference run” (see chapter 4). Each branch along the x-axis represents a halo that becomes a subhalo of the Milky Way at some point in its history, up to $z = 0$. Redshift decreases along the y-axis from $z = 15.6$ to $z = 0$. We only show the most massive branches of the tree on the figure, because in total, we find 6630 branches, of which 394 survive down to $z = 0$. White circles represent independent halos; grey squares represent subhalos. Symbol size is proportional to dark matter halo mass; see the legend for values. Solid diagonal lines represent final mergers with the host, and vertical lines represent the time evolution of an individual halo. Dashed diagonal lines indicate when a halo becomes the subhalo of another halo. To the right of the plot (branch numbers > 70) are the most massive halos that survive down to $z = 0$.

present in any halos in output I-1.

We then decide which of these progenitor (sub)halos is the main progenitor; this is the “son” or “child” (sub)halo (H_j^I is called the “father” or “parent” halo). In this process we ignore fragmentation; since halos merge under gravity, significant fragmentation is unusual for dark matter halos. If we consider the “ith” (sub)halo at step I, H_i^I , and the “jth” (sub)halo at step I-1, H_j^{I-1} , we take the ratio of masses of the two halos m_i/m_j and state that:

- Each halo can have only one son, and only be the son to one halo in the following output.
- H_i^I 's son is the H_j^{I-1} for which m_i/m_j is maximal.

In this way, we construct a “simplified” merger tree that can be traced backwards and forwards in time (the non-simplified merger tree produced by TREEMAKER does not enforce the rules listed above, and hence a single halo can have multiple progenitors, and vice versa).

With this simplified merger tree we can now track the hierarchical merging of galaxies, and follow the evolution of satellite galaxies. To do this, we run HALOMAKER on the star field, and attach the galaxies and clusters found to the dark matter halos in the same output (for clarity, we define galaxies as independent structures in the star field and clusters as substructures, analogous to halos and subhalos). For each galaxy (and cluster), we make a list of halos whose radii enclose the galaxy’s position. We then select from this list the halo whose centre is closest to the galaxy’s centre. In this way, we:

1. Match galaxies with their subhalos, rather than the host halo of the subhalo, but:
2. If no subhalo is found, a satellite galaxy is attached to the host halo, even if there is already a galaxy embedded in the centre of this halo.

Note that while it is not expected that a satellite galaxy will survive while its dark matter halo is destroyed by interactions with the host halo, this may occur if the numerical resolution of the star field is much higher than the dark matter field (we will return to this issue in chapter 4).

One final process that we employ to complete the suite of halo analysis tools is the comparison of two outputs of different simulations with the same initial conditions, and is discussed in the following section.

3.1.3 Twinning Simulations

One of the roles of hydrodynamic simulations is to determine the role of various physical mechanisms, such as supernova feedback or radiative transfer, in the formation of galaxies. This can be done by running two simulations with identical initial conditions, one simulation with the physical process in question included, and a control simulation without. For example, in chapter 4, we run simulations with and without supernovae that are otherwise identical. Since supernova feedback does not affect halo formation (we require that halos are formed first in order to produce stars that can become supernovae), we will thus have two simulations with the same halo population but whose galaxies have different properties that can be compared in order to empirically determine the effect of our chosen supernova feedback model on galaxy formation.

To do this, we employ a very similar technique to that described in section 3.1.2. This process is also employed in Libeskind et al. (2010); Peirani (2010); Romano-Díaz et al. (2010, 2009); Schewtschenko & Macciò (2011), and proceeds as follows. A list of particles and associated halo ID numbers is found for each halo in each simulation for each output in a “reference” run. Due to the differences in the timestepping in each simulation there is typically 0.5-1Myr difference between a given reference run output and a given output in the runs we are comparing (hereafter the “comparison” runs). The list of halos in the reference run is sorted in descending order of halo mass. Particles in the reference run list are removed if they are not in halos in the comparison runs. If a certain halo in the reference run has fewer than 50% of its particles in halos in the comparison runs, it is considered to have no identifiable counterpart in the other simulations and is hence ignored.

For each halo in the reference run we calculate the fraction of its particles that belong to a halo in the comparison runs. We then select the single halo from each of the comparison runs that has the largest fraction of its particles in the reference run halo and has not already been assigned to another reference run halo. This provides a 1:1 mapping between the halos of any two simulations.

In the next section we present the results of this halo identification and twinning process, and determine both visually and statistically whether this process adequately matches halos at the same redshift between simulations.

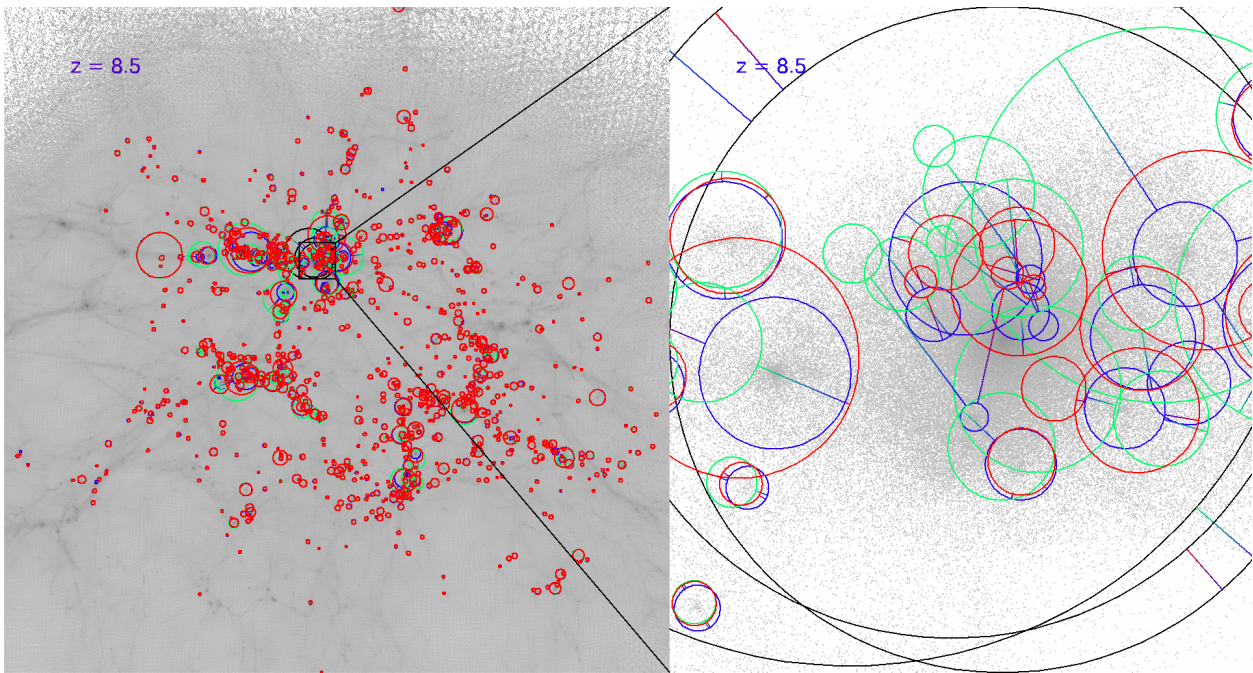


Figure 3.3: Visualisation of the halo twinning results at $z = 8.5$. The left panel shows the cubic volume (300kpc on a side) containing all the Milky Way satellite progenitor dark matter halos identified by $z = 0$. The right panel shows a zoom on the Milky Way progenitor halo outlined by the black square in the left image. The grey scale background represents the dark matter projected density distribution in the Reference run. Overlaid circles indicate the virial radii of halos identified as being Milky Way progenitors in the Reference run (blue) and their twins in the Cooling (green) and Feedback (red) runs, with colours overplotted in that order (hence halos with very similar positions and radii in all three runs appear as red circles). In the right-hand image we connect the halo in the Reference run with its twin in the other runs via straight lines. The Milky Way progenitor halo in each run is shown in black. Most of the twins are remarkably well matched in size and position, although unsurprisingly, the subhalos of the Milky Way progenitor show more pronounced discrepancies between runs, especially in the central region of the halo.




Figure 3.4: Histograms of percentage of halos in each mass bin in the Reference run with twins in the Cooling and Feedback run against virial mass in the Reference run in M_{\odot} . The error bars show the sampling error on the number of un-twinned halos in each bin. The colours correspond to the values at different redshifts (see legend). Twins in the Feedback Run are shown with a solid line, and twins in the Cooling Run are shown with a dashed line. For halos of mass greater than $10^7 M_{\odot}$, we find a success rate above 95% in the twinning procedure. For halos below $10^7 M_{\odot}$ (less than 200 DM particles) and at higher redshifts, this rate drops quite rapidly because of the lower spatial (force) resolution in the Reference run which reduces the overall number of collapsed objects.

3.1.4 Twinning the Nut Suite

In chapter 4 we present the results of a study of the effects of feedback processes on the Milky Way satellite galaxies and their progenitors at high redshift. In order to compare the effects of including supernova feedback against a control simulation without feedback, we must first demonstrate that our twinning procedure works so that we can compare halos between simulations with confidence.

To do this, we introduce the NUT suite of simulations (Powell et al., 2010; Kimm et al., 2011). NUT is a suite of simulations of a Milky Way-class halo with the same initial conditions but different physics and refinement criteria. The highest resolution simulations reach a resolution of 0.5pc. In this study we concentrate on three simulations:

1. The Reference run - A simulation containing gas cooling and star formation but no supernova feedback, with a resolution of 50pc run to $z=0$.
2. The Cooling run - As the Reference run, but with a resolution of 0.5pc run to $z=6-7$.
3. The Feedback run - As the Cooling run, but with a model of supernova feedback included.

These simulations are discussed in more detail in chapter 4, with descriptions of the physical models used. Every simulation has the same dark matter particle mass, initial conditions and cosmological parameters. The only difference is in the baryon physics implemented. For now, we discuss only the twinning of the dark matter distribution.

In order to visually confirm that the twinning procedure works, we plot a map of the Reference run's projected dark matter density field in Fig. 3.3. On the same figure, we overplot halos in the Reference, Cooling and Feedback runs as circles with radii proportional to their virial radius. We also link the halos in the Reference run to their twins in the Cooling and Feedback runs with straight lines connecting the corresponding circles. The figure clearly shows that in general, the twinning procedure yields excellent results for most halos (the vast majority of circles in the left panel of Fig 3.3 are red). For the region encompassed by the Milky Way progenitor (the right-hand zoomed-in panel), twinning results are still quite good, except in the very centre where positions and sizes of sub-halos become more discrepant as the non-linear nature of the system (shell-crossing) and the slight differences in output times between the runs begin to plague the comparison.

We now attempt to quantify the quality of the twinning procedure. The Reference run has 96.4% of its halos twinned with the halos in either the Feedback or Cooling run at $z=6.7$ (the final redshift for which all runs have data). If we relax the 1:1 mapping criterion and simply consider halos above the threshold where 50% of their particles in the Reference run also are found in halos in the other runs, 98.6% of halos have twins. In Fig. 3.4, we plot the success rate for twinning halos as a function of mass and redshift for the Feedback (solid lines) and Cooling runs (dashed lines). We find that for halos over $10^9 M_\odot$ there is a 100% success rate for the twinning procedure regardless of redshift. This success rate remains over 90% above $10^7 M_\odot$ (i.e. for halos containing $\gtrsim 200$ DM particles). For halos below this mass, the lower resolution of the Reference run causes the success rate of the twinning procedure to drop to between 70-90%. Note that only the *spatial* (or force) resolution in the Reference run is lower than in the Cooling or Feedback runs, the dark matter mass resolution is identical.

3.2 Visualisation and Data Handling

Visualisation is a vital aspect of astrophysical simulation post-processing. Visual inspection is the easiest and fastest way of gaining an intuitive feel for the results of a simulation. Bugs and gross errors in the algorithms used become obvious, the results of a halofinding algorithm can be compared against the brain's ability to locate patterns in images, and broad differences in simulations (such as the concentration of filaments with and without photoionisation) can be detected. This may then motivate further, more rigorous analytic work. If nothing else, visualisation is the is a good way of advertising a project to the general public and to non-scientists, since it requires no specialised knowledge to view an image of a simulated galaxy and appreciate its similarity to the real thing.

One of the key aspects of analysis is to compare simulation results to observations. Simulations take the role of repeatable experiments in astrophysics, but they are not of themselves laboratory experiments; numerical errors, resolution limits and incomplete physical models mean that we cannot take the results of simulations at face value. Instead, we must compare our results to observations to determine how well our simulations reproduce real galaxies. Observational data is typically noisy, and extracting physical quantities from it is a difficult problem. By comparison, we are able to extract simulation data with accuracy limited only by sampling error. Hence it is often desirable to produce synthetic observations from simulations and

compare these to real observations, since it is easier to add noise to an image than to remove it. In the following section we discuss a simple procedure for performing this task.

In order to visualise large amounts of astrophysical data, we must first extract that data from the simulation output and present it in a format that can be visualised on-screen. In this section, I present two tools developed during my DPhil for visualising data. The first, RAMTOVAPOR, is a tool for extracting RAMSES data and converting it to a form readable by the general-purpose 3D fluid dynamics visualisation package VAPOR (Clyne et al., 2007). The second, Alice, is a specialised astrophysical visualisation tool, which as well as allowing 3D visualisation contains routines for producing synthetic observations of the simulation data being visualised. I use the results of Peirani et al. (2010) as a case study, in which we produced simulated observations of the minor mergers with elliptical galaxies to determine whether the residual star formation induced by these mergers would be observable with Wide-Field Camera 3 (WFC3) on the Hubble Space Telescope (HST) (Rajan et al., 2011).

3.2.1 Simulation Data Handling

VAPOR is an interactive data visualisation and data analysis tool (Clyne et al., 2007). It is designed to visualise grid-based fluid dynamics simulations, and is available on a number of platforms with plug-ins to various scripting languages. Its native format is the .vdf file, which allows the storage of a number of fluid variables (e.g. density, pressure, velocity, passive scalars such as metallicity, etc). It also allows for data in octree-style adaptive meshes such as those used by RAMSES. In this section I describe a project to convert RAMSES data into the VAPOR .vdf format for visualisation.

The resulting program RAMTOVAPOR has three steps:

1. Load the data from RAMSES
2. Load the AMR tree from RAMSES into the .vdf file format
3. Load each hydro variable from RAMSES into .vdf file format

This must be done for each RAMSES output; currently, RAMTOVAPOR only reads each snapshot individually, though future versions could read multiple snapshots into the same .vdf file. VAPOR provides a number of C++ routines for creating .vdf files and loading tree data into those files. In order to read

the RAMSES data, we employ the Space Emperor library¹, which was written as part of this project. Space Emperor is a work-in-progress C++ library that provides a simple interface for loading-on-demand of RAMSES data, though the interface design is intended to allow generic access to any format of astrophysical data should readers be written for them. In this section I briefly discuss the design of Space Emperor and RAMTOVAPOR.

A simulation snapshot typically resides in a set of snapshots that represent the progress of a simulation. Often, multiple simulations will be packaged together into a suite of simulations that enable us to compare multiple realisations of the same system (e.g. whether simulations containing supernova feedback have different properties to those not containing supernova feedback). Inside each snapshot, fluid parcels represented by particles and/or grid cells are stored. In RAMSES, grid cells are separated into AMR files, which define the grid structure, and Hydro files, which define the hydrodynamic properties of the grid cells such as density, velocity and pressure. This is done since RAMSES uses the AMR grid in the Poisson solver (see chapter 2); hence a simulation containing only AMR and particle files can be run. In the latest version of RAMSES, a file containing information on the gravity field is also written. In certain cases (e.g. for calculating the density field around a halo), it may be helpful to consider every fluid parcel as a generic parcel of matter; in others (such as calculating the ratio of stars to gas in a halo), we will need to specify which type of fluid parcel (dark matter, stars, gas, etc) that we require.

Space Emperor is a C++ library that allows fast, load-on-demand access to the data elements inside RAMSES simulation snapshots. Each snapshot is modelled as an instance of a RAMSES class, which is a collection class which allows access to each of the files inside the snapshot via iterators (Gamma, 1995). The snapshot object, and each RAMSES file object, are “prepared” when their public methods (Deacon, 2005) are invoked. For the snapshot object, this involves locating each of the files stored in the snapshot. For the RAMSES file objects, this involves locating the start of each block of data in the file, by jumping through the file and reading the integer at the start of each Fortran file block. In this way, no data is processed until it is needed. Once this is done, the data can be read. To do this, the code accesses each data block via the name allocated to it (e.g. density, number of CPUs, stellar age). A series of public methods provide access to each variable in a user-readable form.

The library is written using *generic* programming on the lowest level by making use of templating

¹After postponing an after-work activity due to being stuck in the office, a friend noted “A space emperor’s work is never done.”

(Stroustrup, 1986) and the C++ Standard Template Library (STL) (Stepanov & Lee, 1995). The user interface of the library is *object-oriented* (e.g. Deacon (2005)). This provides both the flexibility and readability of object-oriented code with a speed comparable to C on the lowest level; templates in C++ are *statically polymorphic*, which means that the type of object being passed into a template class or method is resolved at compile time, making the run-time speed of a piece of templated code equivalent to that of a hand-coded block of C. Various routines exist for extracting information from Space Emperor and exporting it to ASCII/binary format, as well as IDL routines for importing this data. Python packages that interface with RAMSES such as YT (Turk et al., 2011) and PYMSES² currently provide more developed analysis tools, balancing the computational cost of Python with its rich environment of libraries and simpler syntax. Note that C++ is twice as complex as C by keyword count, and the ability to choose between pointers, references, various smart pointers and direct storage makes the language difficult to learn and debug in the absence of a good project coding standard (Schildt, 1991). It is also possible to interface C, C++, Java and Fortran with Python with a modicum of effort, and several Python libraries use this method to balance speed and ease-of-use.

RAMTOVAPOR is a wrapper around both the Space Emperor library and VAPOR's own AMR tree library. Space Emperor retrieves the grid positions and hydrodynamic variables from a RAMSES output folder. RAMTOVAPOR then passes these to a VAPOR .vdf file object, which writes this data to file for reading by VAPOR's visualisation software.

VAPOR is a well-developed general purpose fluid dynamics package. However, for some purposes, such as the creation of synthetic observations from simulation data, we require a specialised visualisation package. In the next section I discuss Alice, a dedicated astrophysical hydrodynamic simulation visualisation tool, and present a case study of its use to determine the detectability of formation of stars and dynamical structures formed during minor merger events.

3.2.2 Alice: A Virtual Telescope

Alice is a real-time 3D visualisation code for astrophysical simulations, written in C++ with OpenGL³. The aim of Alice is two-fold:

²<http://irfu.cea.fr/Projets/PYMSES/intro.html>

³www-astro.physics.ox.ac.uk/~samgeen/alice/index.html

1. Quick, easy visual inspection of a simulation at all outputs from any angle, with the ability to output images and videos.
2. Take any galaxy simulation at any camera angle, and reproduce the view from a given telescope instrument at a given redshift.

Upon being run, Alice reads a configuration file that contains simulation output file names, units information and a list of data fields (e.g. density, specific energy, stellar age, etc). It then pre-loads every simulation output and samples the data points to create a low-resolution version of each output, which are all loaded into memory to allow quick display when cycling between outputs. It then loads and displays the first output at full resolution. The user can then pan and zoom with the mouse and keyboard, cycle through the various data fields in the fluid in linear, log or squared form, enter a fly-through mode, save images and movies. It also includes a very simple halofinder and analysis tools (as in figure 2.3, where we tag all the particles in a halo at $z=0$ and trace them through previous outputs back to the initial conditions). It can also produce light echoes from supernova simulations (e.g. Sugerman et al. (2005)).

One of the problems with writing a code entirely in C++ is that it is not suited to modification by scripting. A more flexible approach, used in codes such as YT (Turk et al., 2011), PYMSES and PYNBODY⁴, is to write the user-facing code in Python and if necessary refactor objects or modules in C or C++ if extra speed is required. This combines the run-time speed of C/C++ with the ability to write scripts in Python. A further benefit is the ease with which Python modules can be combined; e.g. PYMSES can be used to import data into OpenGL easily. A prototype of Alice in Python has been written which performs this task.

Once the data in a simulation is loaded, and a projection axis chosen, we can generate a synthetic observation of our simulation data. The process is summarised visually in figure 3.5. The following procedure currently only works on particle data (Alice was originally written with GADGET data in mind), though the techniques discussed are broadly applicable to grid-based data as well.

First, the star and gas particles are projected onto a 2D array corresponding to the instrument's charge-couple device (CCD) pixels. This is done using a cosmological distance ladder, assuming a Λ CDM cosmology (Hogg, 1999), with default values of $H = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.3$ and $\Omega_\Lambda = 0.7$ and a flat universe ($k = 0$) (alternatively, users can import their own cosmological parameters). Next, we assign a

⁴<http://code.google.com/p/pynbody/>

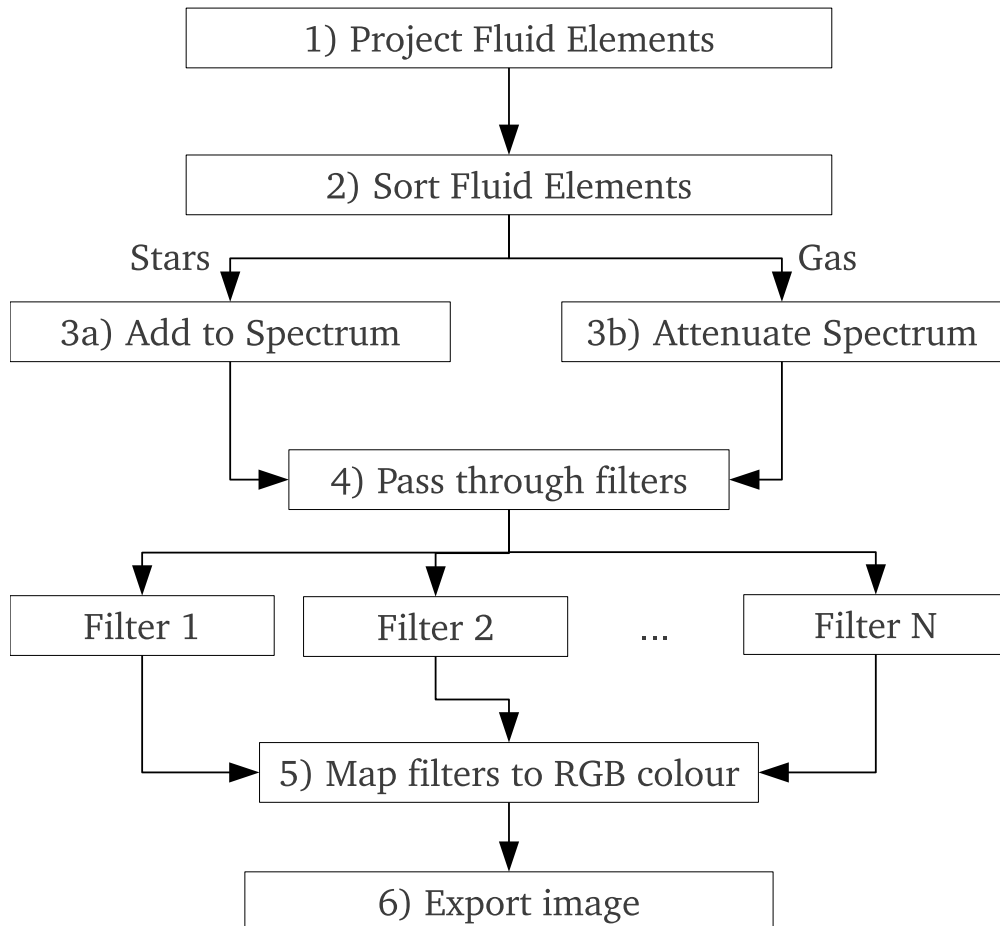


Figure 3.5: Flow diagram of the process Alice uses to produce synthetic observations from galaxy simulations. 1) Bin the fluid elements according to the instrument charge-couple device (CCD) cell they fall inside. Now, for each cell: 2) Sort the fluid elements by distance to the user, with furthest elements first. 3) Add the fluid element's contribution to the spectrum inside that cell (add mass-weighted SED spectrum for a star particle of a given mass and age; attenuate the spectrum according to an extinction law for the gas Calzetti (2001); Calzetti et al. (2000)); redshift according to the peculiar velocity of the particle and the cosmological redshift of the data. 4) Convolve the spectrum inside the CCD cell with each transmission filter in the instrument to produce a flux in each filter in each cell. 5) Map the logarithm of each filter's flux to its own RGB colour, and sum the RGB colours for each image. 6) Output the RGB image to .jpg or a .mpg movie. Alternatively, output the image from each filter as an array of fluxes.

stellar energy distribution (SED) and emission line strength for various emission lines to each star particle, following Bruzual & Charlot (2003). To do this, we use an empirically-determined age-metallicity relationship, derived from Prochaska et al. (2003). The spectra are frequency-shifted according to the Hubble and peculiar velocities relative to the simulated observer.

The code computes the emission from the star particles and the absorption from the gas particles. Each of a given pixel's spectra and emission line strengths are summed, and attenuated according to the gas column density (Calzetti et al., 2000; Calzetti, 2001). In some simulations, this approach will be inaccurate, since if we do not resolve the multiphase ISM, we will underestimate the effect of dust on extinction. Equally, we neglect emission from ionised gas; this should be included in future versions of the code. In the next step, these spectra are passed through a set of transmission filters. To do this, we linearly interpolate the spectra calculated with the given transmission curves. Finally, we produce flux maps corresponding to each transmission filter and emission line. We scale the fluxes to the minimum flux detectable by the instrument in each wavelength.

As examples, Alice contains instrument parameters taken from the WFC3 instrument. The UV filters are particularly good for identifying young stellar populations inside galaxies (see for instance, Hibbard et al. (2005)) while near infra-red (NIR) spectral window traces primarily the old stellar populations. This is because young stellar populations contain massive blue stars that have a lifetime on the order of 10Myr. Fluxes are calculated from these parameters by assuming a simulated exposure time.

In post-processing, it is necessary to add sky background, Poisson and readout noise. Further, images must be convolved with an appropriate Point-Spread Function (PSF). The HST PSF for WFC3⁵ can be derived using the TINYTIM⁶ software package (Krist & Hook, 2004).

3.2.3 Case Study: Residual Star Formation in Mergers

In this section we discuss the use of Alice to produce synthetic observations of minor mergers. The work in this section covers some of the results presented in Peirani et al. (2010). In this paper, we produce synthetic observations of minor mergers ($0.1 < M_1/M_2 < 0.25$) through the WFC3 instrument. We then determine whether the structures produced by these mergers, and the induced star formation, are detectable by the

⁵Winners don't do acronyms (WDDA).

⁶www.stsci.edu/software/tinytim




Figure 3.6: Projection of stellar component of the merging galaxies described in section 3.2.3 (experiment B_2 in Peirani et al. (2010)). Each image shows a 100kpc cubic volume from. Rows 1-2 show the stellar density, with black as most dense and white as least dense. Arrows indicate the trajectory and spin direction of the infalling satellite. Rows 3-4 show the ages of the stars formed since the beginning of the simulation. Reproduced, with permission, from Peirani et al. (2010).

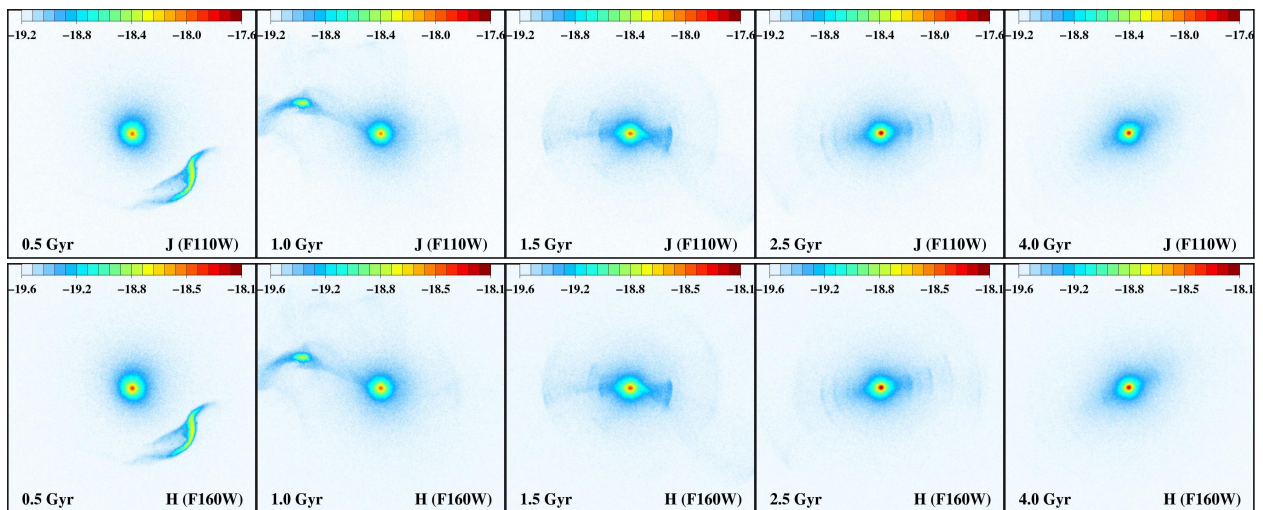


Figure 3.7: As in figure 3.6, but showing the emission through the J and H bands at a luminosity distance of 100Mpc ($z=0.023$). The first row shows the J band through the WFC3 F110W filter, corresponding to 0.9 - 1.4 μm . The second row shows the H band (F160W, at 1.4-1.7 μm). The images are Gaussian-smoothed with $\sigma = 2$ pixels. Reproduced, with permission, from Peirani et al. (2010).

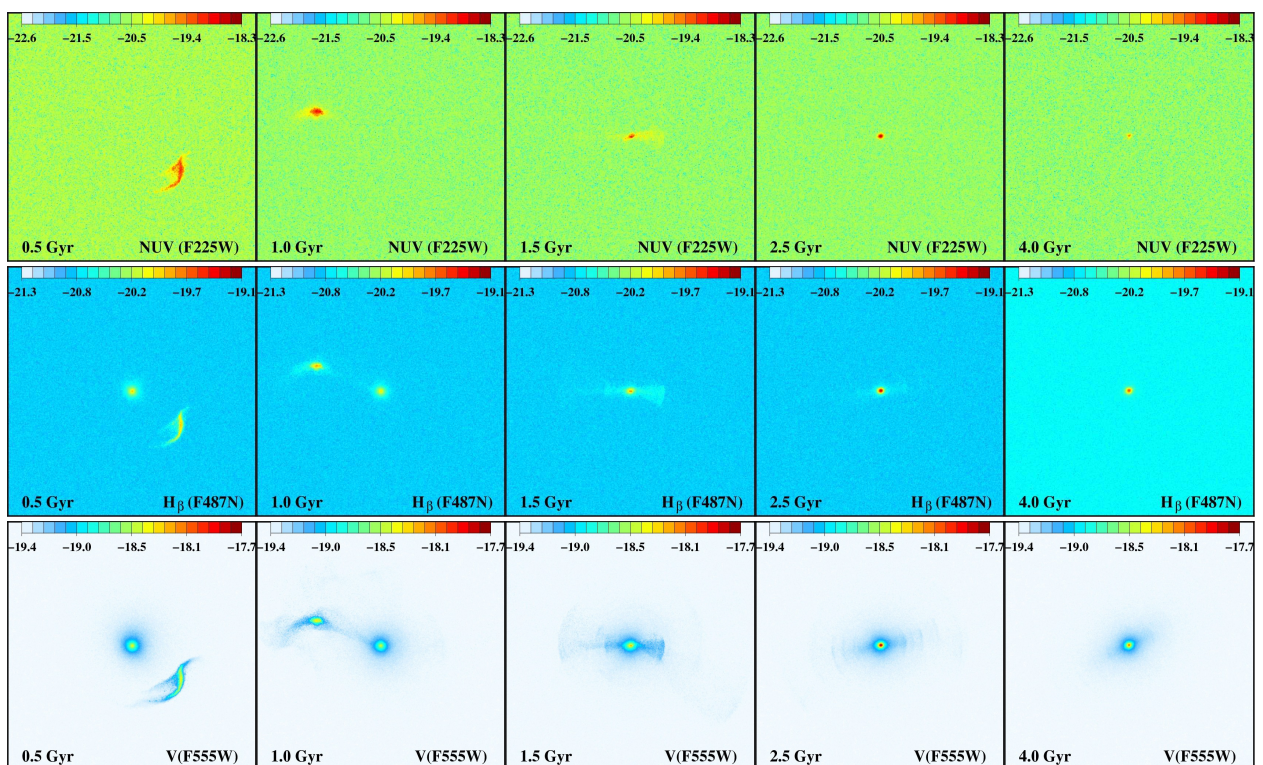


Figure 3.8: As in figure 3.7, but for the NUV, H_β and V bands. The first row shows the near-UV (NUV) filter (F225W, 0.2-0.275 μm). The second row shows the narrow-band H_β filter (F487N, 0.48-0.49 μm). The third row shows the V-band filter (F555W, 0.45-0.7 μm). Reproduced, with permission, from Peirani et al. (2010).




Figure 3.9: As in figure 3.7, but for a luminosity distance of 20Mpc ($z=0.05$). Note that the field-of-view (FOV) is smaller since the instrument has a fixed angular size. Reproduced, with permission, from Peirani et al. (2010).

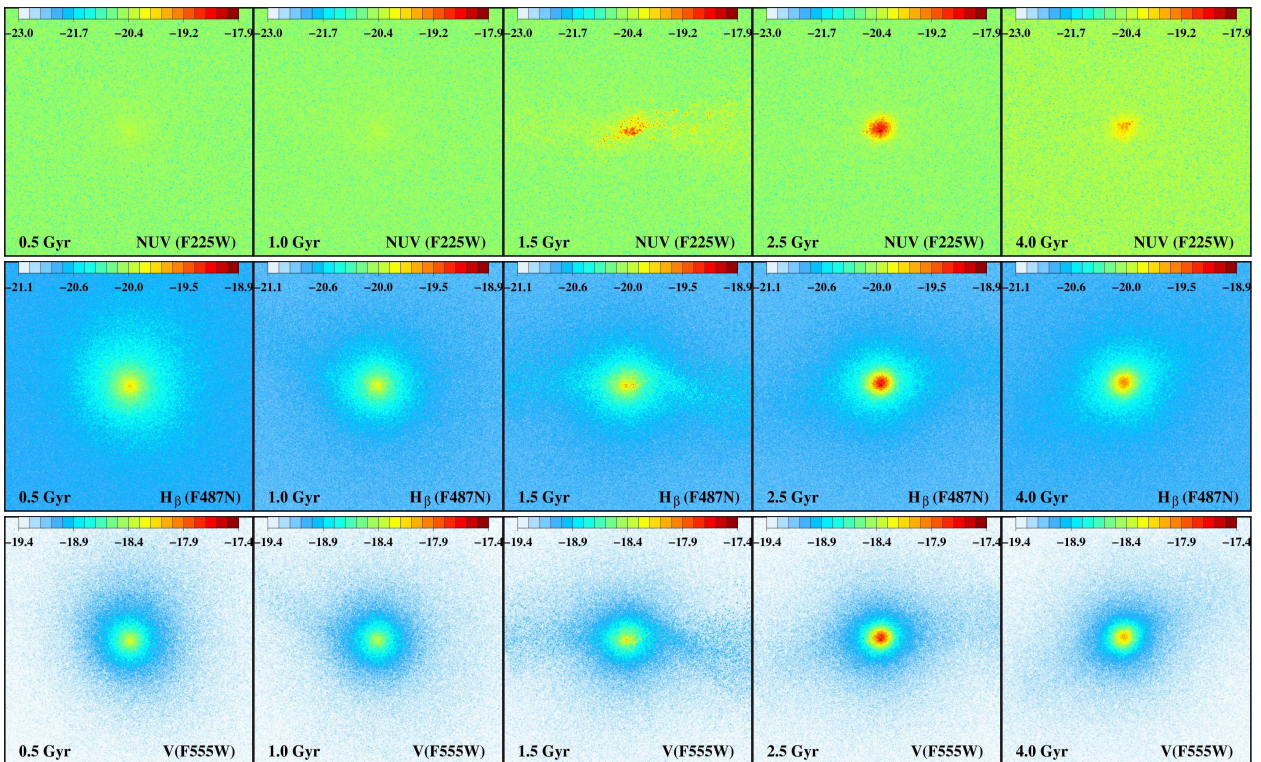


Figure 3.10: As in figure 3.8, but for a luminosity distance of 20Mpc ($z=0.05$). Reproduced, with permission, from Peirani et al. (2010).

WFC3.

We use idealised N-body and SPH gas dynamics simulations of a spiral galaxy merging with an elliptical galaxy (Springel et al., 2005). The elliptical galaxy has a mass of $10^{12}M_{\odot}$, with a spherical dark matter halo and stellar component with 4% of the total mass of the halo. The stellar component follows a Hernquist profile with effective radius 4.29kpc. For the spiral galaxy, we use a spherical dark matter halo and bulge-less stellar and gas disks. The spiral galaxy has a total mass of 0.25-0.1 of the mass of the elliptical, depending on the simulation run. The gas fraction in the disk is 20%, while the mass of the disk is 5% of the total mass. The satellite is placed in a prograde or retrograde parabolic orbit with an initial distance of 100kpc and a varying pericentric distance. The simulations are run in GADGET-2 (Springel, 2005), with additional prescriptions for cooling, star formation and supernova feedback. A gas particle may form stars if it has the properties: $T < 2 \times 10^4\text{K}$, $n > 0.1\text{cm}^{-3}$ and $\nabla \cdot v < 0$ according to the prescription described in section 2.3.3. Each simulation has 4M particles. The dark matter particle mass is $3.03 \times 10^5 M_{\odot}$. In the spiral galaxy, the gas and star particle mass is $4.5 \times 10^4 M_{\odot}$. In the elliptical galaxy, the gas and star particle mass is $1.35 \times 10^5 M_{\odot}$. Details of the simulations run and the parameters used in them are given in full in Peirani et al. (2010).

We focus here on a simulation from the paper with a galaxy mass ratio of 1:6 and a prograde orbit (experiment *B2* in Peirani et al. (2010)). Figure 3.6 shows the distribution of the particles in the simulation, along with the ages of the stars formed since the start of the merger. Shell-like structures are visible from 1.5Gyr to 4Gyr, with a fountain-like structure appearing between 1.5 and 2.5Gyr. The results indicate that merger-triggered star formation occurs mainly in the centre of the elliptical during this period, with older stars lying in the shells. This is due to the fact that the gas from the satellite falls into the centre of the elliptical's potential, fuelling star formation. For the first 1.5Gyr, star formation occurs more uniformly, since the gas from the satellite is still bound to the infalling object. Stars formed later are likely to have higher metallicity, and a metallicity gradient may be found that agrees with the observations of Clemens et al. (2009).

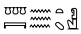
We take two sets of synthetic observations of the galaxy; one at $z = 0.023$ and one at $z = 0.005$ (corresponding to luminosity distances $D_L = 100\text{Mpc}$ and 20Mpc respectively). Since this is an idealised and not a cosmological simulation, redshift is not an implicit parameter and we are free to set the luminosity

distance by hand. We choose a low redshift since we attempt to reproduce the findings of, e.g. Malin & Carter (1983); van Dokkum (2005), who find numerous shell structures in elliptical galaxies in the local universe. These objects often exhibit recent star formation (Schweizer et al., 1990; Schweizer & Seitzer, 1992). Feldmann et al. (2008) find that these structures cannot be reproduced by equal-mass mergers but can be reproduced by minor mergers of existing elliptical galaxies and accreted disk galaxies. These kinds of mergers become more frequent than major mergers at lower redshift, as existing elliptical galaxies accrete their lower-mass neighbours (Khochfar & Silk, 2006; Genel et al., 2008; Khochfar & Silk, 2009).

In figures 3.7 and 3.8, we plot the infra-red (IR) and UV images of this merger at $z=0.023$ ($D_L=100\text{Mpc}$). Figure 3.7 shows that the shell structure is visible at 1.5Gyr and 2.5Gyr, though any observational sign of the shells disappears by 4.0Gyr. Both wavebands depicted appear to show very similar features. Figure 3.8 clearly traces the recent star formation in the UV filters, with only stars formed in the last 0.5Gyr visible in the near-UV (NUV) images. Hence using combined images of each filter in WFC3, we are able to locate features from merger events and separate the stellar population by age. Further, we can use the UV filters to constrain the gas mass of the infalling satellite by measuring the star formation rate.

We perform the same analysis for the same simulation at $z=0.005$ ($D_L=20\text{Mpc}$), which is shown in figures 3.9 and 3.10. In this case, the instrument FOV only covers the inner part of the galaxy. Plumes at 1.5Gyr show evidence of merger activity, and as in the previous study, star formation in the NUV band is clearly visible.

With synthetic observations, we are able to simulate detection of various phenomena, and bridge the gap between simulations and observations. Simulations predict a variety of physical features in galaxy formation, but it is often unclear as to whether these structures will be detectable by observations. By using this kind of analysis, we can determine whether these structures do indeed exist, or whether they are simply artefacts of our numerical recipes. Approaching the problem from another direction, we can ascertain the origins of structures detected in observations of galaxies by reproducing them in both idealised and cosmological simulations.

In the next chapter we turn our attention to simulations of accretion on a smaller scale; namely, the Milky Way satellites. We use the  (NUT) suite to model the formation of the satellites of a Milky Way-like galaxy from $z=499$ to $z=0$, with particular focus made on feedback processes inside these dwarf

galaxies at high redshift.

Chapter 4

Satellites of the Milky Way

4.1 Introduction

In this chapter, I present a study of feedback in the nearest dwarf galaxies to Earth; the Milky Way satellite galaxies. Surprisingly little is known about the origin and evolution of the Milky Way's satellite galaxy companions. UV photoionisation, supernova feedback and interactions with the larger host halo are all thought to play a role in shaping the population of satellites that we observe today, but the jury is still out on which of these effects, if any, dominates. In chapter 1 I review the earlier work done to examine the problem. I now revisit the issue by re-simulating a Milky Way class dark matter (DM) halo with unprecedented resolution. This work is also presented in Geen et al. (2012). In this paper, we employ a set of cosmological hydrodynamic AMR simulations, the NUT suite, to investigate the effect of supernova feedback and UV photoionisation at high redshift with sub-parsec resolution. We subsequently follow the effect of interactions with the Milky Way like halo using a lower spatial resolution (50pc) version of the simulation down to $z=0$. This latter produces a population of simulated satellites that we compare to the observed satellites of the Milky Way and M31. We find that supernova feedback reduces star formation in the least massive satellites but enhances it in the more massive ones. Photoionisation appears to play a very minor role in suppressing star and galaxy formation in all progenitors of satellite halos. By far the largest effect on the satellite population is found to be the mass of the host and whether gas cooling is included in the simulation or not. Indeed, inclusion of gas cooling dramatically reduces the number of satellites captured at high redshift which survive down to $z=0$.

Simulation	z_{\min}	m_{DM}	R_{\max} (level)	Gas Cooling	m_*	SNe	UV
Reference Run	0	$5.6 \times 10^4 M_{\odot}$	50pc (18)	✓	$3.5 \times 10^4 M_{\odot}$		✓
Cooling Run	6.7	$5.6 \times 10^4 M_{\odot}$	0.5pc (25)	✓	$167 M_{\odot}$		✓
Feedback Run	6.7	$5.6 \times 10^4 M_{\odot}$	0.5pc (25)	✓	$76 M_{\odot}$	✓	✓
Dark Matter Run	0	$6.7 \times 10^4 M_{\odot}$	50pc (18)		-		
Adiabatic Run	0	$5.6 \times 10^4 M_{\odot}$	50pc (18)		-		✓

Table 4.1: Table of properties of numerical simulations included in this chapter. The columns are, from left to right, the simulation name, the lowest redshift reached by the simulation, the minimum dark matter particle mass, the maximum spatial resolution of the AMR grid (with the associated level of refinement in brackets), whether the simulation includes gas cooling, the minimum star particle mass, whether the simulation includes supernova feedback, and whether the simulation includes a UV background. The bottom two simulations are considered only in section 4.5.1. For a complete description of the simulations, see section 4.2.1.

This chapter is split into four main parts. In section 4.2, we describe the simulations, our algorithms for comparing them and for tracking halos at high redshift down to $z = 0$. We then discuss the comparison of the dark matter properties of the satellites in section 4.3. Section 4.4 looks at the effect of feedback mechanisms on satellite galaxy formation at high redshift and ultra-high resolution. The fourth part, section 4.5, is devoted to the present epoch, and how our results at high redshift affect the satellite population we see today.

4.2 Simulating the Milky Way Satellites

4.2.1 Sugarplum Fairies: The Nut Suite

We analyse five simulations in the Nut Suite (NUT) suite of simulations (Nut Suite is the Ancient Egyptian goddess of the sky) (Powell et al., 2010; Kimm et al., 2011). Nut Suite (NUT) is a set of cosmological hydrodynamic resimulations of a Milky Way-like halo at $z = 0$ (throughout this chapter, this halo will be referred to as the “Milky Way” for simplicity). To run these simulations, we use the Adaptive Mesh Refinement (AMR) code RAMSES (Teyssier, 2002), presented in chapter 2. Each simulation starts from identical initial conditions, which are generated with MPGRAFIC (Bertschinger, 2001; Prunet et al., 2008) using cosmological parameters consistent with the WMAP 5 year measurements (Dunkley et al., 2009). The simulation volume is a periodic, cubic box of length $9h^{-1}\text{Mpc}$ with a minimum resolution of 128^3 dark matter particles and the same number of grid cells. Within this volume we carve out a spherical region of radius $1.44h^{-1}\text{Mpc}$,

centred on a halo that reaches a virial mass $M_{\text{vir}} = 5 \text{ times } 10^{11} M_{\odot}$ at $z=0$; a discussion of this mass is provided in section 4.3.2. We place three nested grids in this spherical region with effective resolutions of 256^3 , 512^3 and 1024^3 dark matter particles and grid cells. The minimum dark matter particle mass inside this region is equal to $5.6 \times 10^4 M_{\odot}$ (with the exception of the Dark Matter run, which we describe later in this section). We then allow the grid inside the refinement region to adaptively refine up to a given maximum level for each simulation. Our refinement strategy is quasi-Lagrangian: a grid cell is refined when the number of dark matter particles in the cell exceeds 8, or the baryonic mass of the cell reaches $8 m_{\text{SPH}}$, where $m_{\text{SPH}} = 9.4 \times 10^3 M_{\odot}$. The simulation parameters used are summarised in table 4.1, and in the text below.

The three main simulations that we consider in this chapter contain dark matter, gas cooling and a uniform UV background switched on at $z = 8.5$ to model reionisation (Haardt & Madau, 1996). Gas cooling is modelled as radiative energy loss from atomic processes including emission line cooling (below 10^4K), with a primordial metallicity of $0.001 Z_{\odot}$. Star formation in the simulation proceeds according to a Schmidt law on a local dynamical timescale (Cen & Ostriker, 1992) with an efficiency of 1%. The density threshold for star formation is set in each simulation to be comparable to the corresponding Jeans density of a cell on the highest level of refinement with a temperature of 100 K.

We first run a simulation that we call the ‘‘Reference run’’. The Reference run is allowed to refine adaptively to up to 8 times inside the fixed refinement region, such that the densest regions are allowed to reach a maximum physical resolution of 50pc at all times, between a few times and an order of magnitude higher than other cosmological hydrodynamics simulations of Milky Way satellites (Libeskind et al., 2010; Nickerson et al., 2011; Okamoto & Frenk, 2009; Parry et al., 2011; Ricotti & Gnedin, 2005; Romano-Díaz et al., 2010, 2009; Sawala et al., 2011; Scannapieco et al., 2011; Schewtschenko & Macciò, 2011; Wadepuhl & Springel, 2010a). The minimum star particle mass in this simulation is $3.5 \times 10^4 M_{\odot}$. We run this simulation to $z = 0$. Note that the main purpose of this run is to act as a lower spatial resolution ‘‘twin’’ of the two ‘‘high resolution’’ simulations in this study, allowing us to determine which of the galaxies formed at high redshift are progenitors of Milky Way satellite galaxies today. For this reason the DM mass resolution is kept identical in all runs. The threshold for star formation in the Reference run is 10 atoms/cm^3 .

We then run two high resolution simulations which are allowed to refine adaptively by up to 15 times so that its physical spatial resolution in the densest regions can reach a maximum of 0.5pc at all times. The

first of these high resolution simulations we call the ‘‘Cooling run’’. As with the Reference run, the Cooling run contains dark matter, gas cooling and a uniform UV background switched on at $z = 8.5$, but now the threshold density for star formation is 10^5 atoms/cm^3 . As a result, the minimum star particle mass formed in the Cooling run is $167M_{\odot}$. The second of the high resolution simulations is called the ‘‘Feedback’’ run. The Feedback run is identical to the Cooling run, except that it also includes supernova feedback. Following Dubois & Teyssier (2008b), supernovae are implemented as Sedov blast waves with a radius of 2 grid cells (1pc) around a star particle 10 Myr after it formed. Note that while we do not resolve individual stars, at this mass resolution and assuming a Salpeter initial mass function (Salpeter, 1955), we get one supernova per star particle. We assume supernova events entrain 50% of the initial mass of the star particle ($\eta_W = 1$ in the notation of Dubois & Teyssier (2008b)) in a wind and have a metal yield of 0.1. The energy released is given by $\eta_{\text{SN}} \frac{m_{\star}}{m_{\text{SN}}} e_{\text{SN}}$, where m_{\star} is the mass of the star particle, m_{SN} and e_{SN} are typical values for a massive star going supernova and η_{SN} is the fraction of the total mass formed that is turned into supernova ejecta. For this simulation, we use $\eta_{\text{SN}} = 0.106$, $m_{\text{SN}} = 10M_{\odot}$ and $e_{\text{SN}} = 10^{51} \text{ ergs}$ (Powell et al., 2010; Kimm et al., 2011). This translates into a minimum star particle final mass of $76M_{\odot}$ for the Feedback run.

Finally, we perform a further two simulations. These are called the Dark Matter run and the Adiabatic run. Both have the same initial conditions and refinement criteria as the Reference run. The Dark Matter run is a pure N-body dark matter simulation, in which the mass in baryons is replaced by mass in dark matter, giving it a minimum dark matter particle mass of $6.7 \times 10^4 M_{\odot}$ rather than $5.6 \times 10^4 M_{\odot}$ which is the value common to all the other runs. The Adiabatic run is identical to the Reference run, except that the gas is not permitted to radiate away its energy. As a result, no star formation takes place in the Adiabatic run though we still include the UV background for sake of comparison. These two simulations are used to determine the effect of including more physics on satellite galaxy evolution from $z \sim 6$ (redshift where the high resolution simulations stop) to $z = 0$. We discuss the results of this study in section 4.5.1 and compare the Dark Matter run to other pure N-body dark matter simulations of Milky Way satellites in section 4.3.2.

The Cooling and Feedback runs required on the order of a few million CPU hours, and are still running (with an end goal of $z = 3$). The Reference run took on the order of 200 000 CPU hours. The Adiabatic and Dark Matter runs required considerably less computation, both using 50 000 CPU hours or less.

4.2.2 Mapping Halos in the Suite

We use HaloMaker to identify dark matter halos and galaxies in each simulation output using the methods described in section 3.1.1 of chapter 3. We identify independent halos as dark matter overdensities not contained within another halo, and subhalos as halos that are identified as substructures of other halos. Similarly, we identify galaxies as overdensities in the star particles. For this procedure, we reject any identified halo that contains less than 40 particles, and any galaxy that contains less than 10 particles. The minimum total mass of a given dark matter (sub)halo is thus $2.2 \times 10^6 M_\odot$ and the minimum stellar mass for a galaxy is $3.5 \times 10^5 M_\odot$ in the Reference run, $1700 M_\odot$ in the Cooling run and $760 M_\odot$ in the Feedback run.

We perform this process in every simulation for both the dark matter and the stars whenever possible. We thus identify every dark matter halo and every galaxy above the mass limits given in the last paragraph. We visually inspect the results of the halo identification and tune the halofinder parameters such that we minimise spurious halo detection or halos that are visually identifiable but not detected by the halo finder.

We then identify the host halo of galaxies to determine which halos are luminous and which are dark. This sorting proceeds in two steps: (i) for each galaxy, we make a list of halos that enclose it within their virial radius (ii) we select the halo that lies closest to the galaxy centre. This final step is needed when the galaxy's host is a subhalo of a larger halo, and the galaxy lies within the virial radius of both the subhalo and its host. Once this is done, we adopt the halo twinning strategy described in section 3.1.3 of chapter 3. The results of this twinning process are described in the same section.

In this way, we can compare the properties of embedded galaxies in halos in each simulation. This enables us to study the effect of feedback on galaxies by comparing the galaxies formed in the Cooling run with the same galaxies formed in the Feedback run. However, if we want to study the Milky Way satellites, we must also track these galaxies to $z=0$. Since the Cooling and Feedback runs have not reached $z=0$, we must use the Reference simulation to follow their paths to $z=0$. We discuss this process in the following section.

4.2.3 Tracking high redshift galaxies down to $z = 0$

The ultimate goal of this project is to compare our simulated galaxies to observed Milky Way satellites. In order to achieve it, we need to evolve our simulated galaxies in the Cooling and Feedback run to $z = 0$. Since

it is computationally unfeasible at their nominal resolution, we instead track their evolution via their twin halos merger trees in the Reference run. This determines which galaxies at high redshift are the progenitors of Milky Way satellites today and allows us to quantify how advanced satellite galaxy formation is by the end the epoch of reionisation.

The fundamental assumption we make is that a halo which already contains stars at high redshift will still contain a galaxy at $z = 0$. This assumption is extremely plausible for two reasons. Firstly, even the lowest-mass halos are observed to be dark-matter dominated (Strigari et al., 2008), and thus we do not expect to find galaxies without dark matter halos. Secondly, galaxies are all predicted to be embedded within the inner part of the halo in which they form, so that the galaxy will be the last part of the halo to be destroyed, with tidal stripping affecting the outer regions of the halo first (e.g. Peñarrubia et al. (2010a)). We further comment on the validity of this assumption in section 4.5.2 where we identify galaxies in the Reference run at $z=0$ and locate their dark matter host halos. Finally, we also assume with this extrapolation technique that the dynamical friction and tidal stripping experienced by the satellite halos in the Reference run are similar to the Cooling and Feedback runs, i.e. that increased resolution and supernova feedback do not dramatically alter their efficiency. We discuss the validity of this assumption in more detail in section 4.5.1.

We track Reference run halos to $z=0$ using two techniques. The first is to build a halo merger tree as described in section 3.1.2 of chapter 3, where we enforce one child halo for every parent halo. One side-effect of this method is that if a subhalo loses more than 50% of its particles between two outputs, that subhalo is assumed to have been completely stripped by its host. To limit this occurrence, we use a large number of snapshots to build our merger tree (~ 100), so that our effective time resolution is roughly 150Myr. This method compares favourably to a second, more straightforward method we attempted, in which halos are directly matched between a single high-redshift output and the output at $z = 0$. In this alternative method, we chose as the 'child' the halo at $z = 0$ that was given the largest proportion of its particles by the high redshift halo. However the inevitable high mass loss rates between high and low z outputs led to a few pathological cases whereby the 'same' halo at two different high redshift outputs was attributed different children at $z = 0$. For this reason, we use the merger tree method for all further analysis.

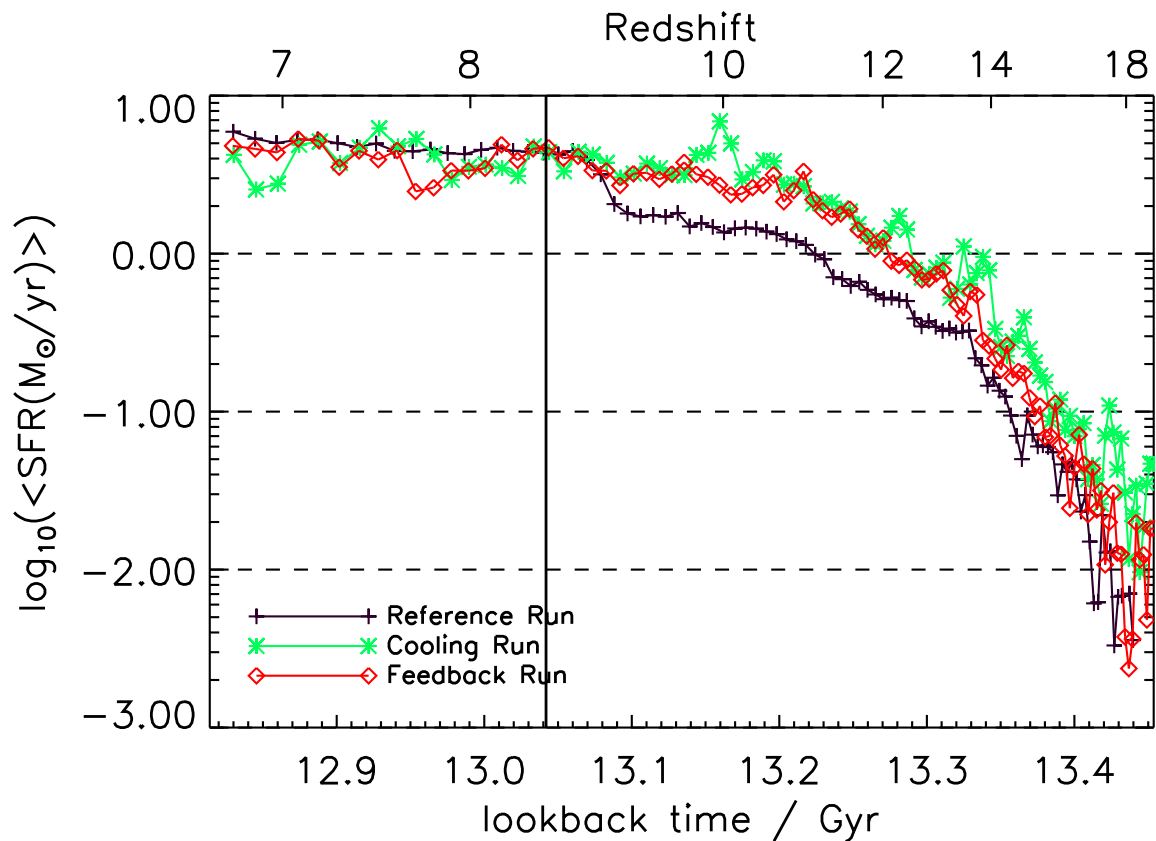


Figure 4.1: Star formation rate averaged over all halos in each simulation. The Reference run is shown in black, the Cooling run in green and the Feedback run in red. The vertical black line shows the time ($z = 8.5$) at which the universe is reionised in the simulation. Star formation in the Reference run is slightly delayed compared to the other runs, but catches up before reionisation. The jump in star formation rate at $z = 9$ is due to the triggering of a new level of refinement on the grid (from 14 to 15 levels) at this redshift, which allows the gas to collapse further and triggers star formation in all potentially star-forming halos.




Figure 4.2: Comparison of the mass-weighted stellar age for halos in the Reference run and their twins in the Feedback run. Halos that lie on the diagonal black line have the same mass-weighted stellar ages in both runs. Due to its lower resolution, as explained in section 4.2.4, the Reference run generally forms stars later than the Feedback run. However, the difference in star formation parameters between the runs allows a modest fraction of halos to form stars before, especially at lower redshifts. Similar results are found when comparing the Reference run to the Cooling run.

4.2.4 Resolution effects on galaxy formation

In order to assess how reliable the Reference run is to locate satellite galaxies at $z = 0$, we first compare it to the higher resolution Cooling and Feedback runs in the redshift range where all the runs overlap. In chapter 3, we used the twinning techniques described there to show that we are able to very successfully match halos more massive than $10^7 M_\odot$ between simulations. We now consider the effect that resolution has on star formation.

As previously mentioned in section 4.2.1, the Reference run has the same dark matter mass resolution as the Cooling and Feedback runs. However, the spatial resolution, which determines the accuracy of both the gravitational force and the properties of the gas is lower; 50pc in the Reference run instead of 0.5pc in the Cooling and Feedback runs. The density threshold for star formation is therefore lowered from 10^5atoms/cm^3 at high resolution down to 10atoms/cm^3 at low resolution, whilst the efficiency of star formation is preserved.

In figure 4.1, we compare the global star formation rate of halos in each of the runs. We find that before a lookback time of 13.1Gyr ($z = 9$), the Reference run's star formation rate is roughly half that of the Cooling and Feedback runs. However, after 13.1Gyr, all star formation rates agree within a few ten percent. This difference of behaviour before and after 13.1Gyr has nothing to do with reionisation, which occurs later on. Indeed this effect is purely numerical, and induced by the refinement criteria we choose to enforce. RAMSES refines the AMR grid using an octree, meaning that spatial resolution is a power-of-two fraction of the total box length (Teyssier, 2002). Since we specify a maximum spatial resolution for the grid in physical pc, we trigger a power-of-two increase in resolution each time the cosmological scale factor has increased enough that an extra level is necessary to achieve such a resolution. In the Reference run, such a jump in refinement level from 14 to 15 happens around $z = 9$. Rasera & Teyssier (2006) demonstrate that too low a maximum spatial resolution delays the collapse of low mass haloes/galaxy disks, preventing the ISM gas density in many of them from crossing the star formation threshold until a higher level of resolution is achieved. Such a delay eventually vanishes when the maximum spatial resolution becomes sufficient *at all times* as the lack of 'step' in the star formation histories of the Cooling and Feedback runs on Fig 4.1 clearly shows.

Note that the good agreement between the global star formation rates of all three runs for lookback times smaller than 13.1Gyr is expected by construction as we select the density threshold for the Reference

run to best match the star formation measured in the Cooling and Feedback runs. This has to obey two constraints: (i) the star formation density threshold should be smaller than the corresponding Jeans density ($\rho_J = (\pi c_s^2)/(\lambda_J^2 G) \sim 40 \text{ at/cm}^3$, Binney & Tremaine (2008)) on the highest level of refinement and (ii) stars should not form in smooth filaments, which yields a lower bound on the density threshold that we empirically determine to be $\sim 10 \text{ at/cm}^3$ (Powell et al., 2010). Alternative Reference runs with different values of the density threshold lead to the conclusion that the lowest one (10 at/cm^3) matches the global star formation histories in the Cooling and Feedback runs best, as shown on Fig 4.1.

We now consider the agreement between star formation histories of individual galaxies. In Fig. 4.2, we compare the mass weighted stellar ages of galaxies simulated at low (Reference run) and high (Feedback run) resolution and twinned at $z = 6.7$. We find a pattern similar to that of the global star formation histories presented in Fig 4.1; namely star formation is delayed in the Reference run but converges to values similar to the Feedback run at a lookback time comprised between 13.0 and 12.9 Gyr. After this epoch there is some inevitable scatter due to the nonlinear nature of star formation, but this scatter is centred around the line of equal age in Fig. 4.2. The lookback time at which the ages converge is later than the jump in star formation due to resolution because the mean age is skewed by the relative paucity of stars formed before 13.1 Gyr in the Reference run (Fig. 4.1). A similar result is found when comparing the Cooling run to the Reference run.

It is worth pointing out that the location of star formation within a galaxy is not guaranteed to match between simulations. Star formation in the Cooling and Feedback runs is confined to regions with densities similar to molecular cloud cores ($\rho > 10^5 \text{ atoms/cm}^3$), whereas in the Reference run star formation is allowed to occur in regions where the density is closer to that of typical diffuse clouds (10 at/cm^3). However, for analysing the bulk properties of satellite galaxies between reionisation and $z = 0$ this distinction is largely irrelevant.

4.3 Comparison of Halo Properties

4.3.1 Tracking satellite galaxies down to $z = 0$

In section 4.2.3, we discussed the techniques used to track the galaxies formed in the Cooling and Feedback run down to $z = 0$ using the Reference run. We now analyse the results of this tracking. The merger tree

for the Reference run is shown in Fig. 4.3 where we track halos (marked as circles) becoming subhalos (marked as squares), and being tidally disrupted (squares become smaller) before completely merging with their hosts. We resolve about 6630 such halos. Of these 6630, 394 survive as subhalos of the Milky Way halo at $z = 0$. In table 4.2, we list the number of galaxies formed by various redshifts from $z = 11$ to $z = 6.7$ in the Cooling and Feedback runs that survive to $z = 0$. It is apparent from the table that by this redshift, we have not formed enough satellite galaxies to match even the population of pre-SDSS satellites (Mateo, 1998). However, it is also clear that satellite galaxy formation continues after reionisation; three galaxies that end up as satellites of the Milky Way at $z = 0$ are formed between $z = 8$ and $z = 6.7$.

In fact, we find that the youngest Milky Way satellite galaxy in the Reference run begins forming stars at $z = 4.8$, or a lookback time of 12.4 Gyr. This is illustrated on Fig. 4.4 where we plot the stellar mass of each Milky Way satellite galaxy in the Reference run at $z = 0$ against the age of their oldest star particle. This figure shows that the youngest galaxies typically are the lowest mass galaxies. Since Fig. 4.2, further demonstrate that galaxies formed after a lookback time of 13.0 Gyr have ages that match well between the Reference and Feedback/Cooling runs, these results also are valid for the Feedback run. However, as we note in section 4.2.2, we are unable to identify galaxies with a stellar mass below $3.5 \times 10^5 M_{\odot}$ in the Reference run. Hence we cannot discount the possibility that the Cooling and Feedback runs might form more galaxies with lower masses that survive as Milky Way satellite galaxies. All we can conclude is that every galaxy above this mass threshold that survives as a Milky Way satellite in the Cooling and Feedback runs (through the twinning procedure) also survives as a satellite in the Reference run. However, it is possible that the trend of lower mass satellite galaxies forming at lower redshifts continues in the Cooling and Feedback runs, where stellar masses is better resolved.

Key to the survival process of satellite galaxies is the mass stripping they undergo as a function of time. We visualise this in Figs. 4.5 and 4.6 where we identify which halos containing stars at $z = 6.7$ survive to become Milky Way satellites at $z = 0$ and follow their dark matter and star particles through cosmic time. We locate these particles in outputs of the simulation at $z = 3, 1$ and 0 , and overplot them on top of their respective underlying density fields, colour coding them according to the redshift at which the halo they belong to is captured by that of the Milky Way. We find that halos captured before $z = 1$ experience significant disruption, with multiple tidal streams created despite the core of the halo surviving (halos ‘c’,




Figure 4.3: Halo merger tree for our Milky Way-like galaxy from $z = 15.6$ to $z = 0$ in the Reference run. We only show the most massive branches of the tree on the figure, because in total, we find 6630 branches, of which 394 survive down to $z = 0$. The left-most branch represents the mass evolution of the Milky Way halo itself. White circles represent independent halos; grey squares represent subhalos. Symbol size is proportional to dark matter halo mass; see the legend for values. Solid diagonal lines represent final mergers with the host, and vertical lines represent the time evolution of an individual halo. Dashed diagonal lines indicate when a halo becomes the subhalo of another halo. To the right of the plot (branch numbers > 70) are the most massive branches which reach $z = 0$.

‘e’, ‘f’ and ‘g’ in Fig. 4.5). By contrast, the only stars formed by $z = 6.7$ that are stripped from their galaxies are from satellites captured before $z = 3$ (Fig. 4.6). The surviving satellite that is captured just after $z = 3$ (halo ‘g’) arrives in the Milky Way halo as a subhalo of another halo, which is subsequently disrupted by the Milky Way. Hence this halo has already experienced some stripping by $z = 3$, as shown in Fig. 4.5. We discuss the effect that gas physics has on satellite survival in section 4.5.1.




Figure 4.4: Time of first star formation for satellite galaxies that survive to $z = 0$ (including the MW itself) in the Reference run against their stellar mass at $z = 0$. The region shaded in grey represents the lookback times which have not been simulated in the high resolution Cooling and Feedback runs. This illustrates that satellite galaxy formation is incomplete until $z = 4.8$ (i.e. after a lookback time of 12.4 Gyr). Note that this plot only shows galaxies in the Reference run.

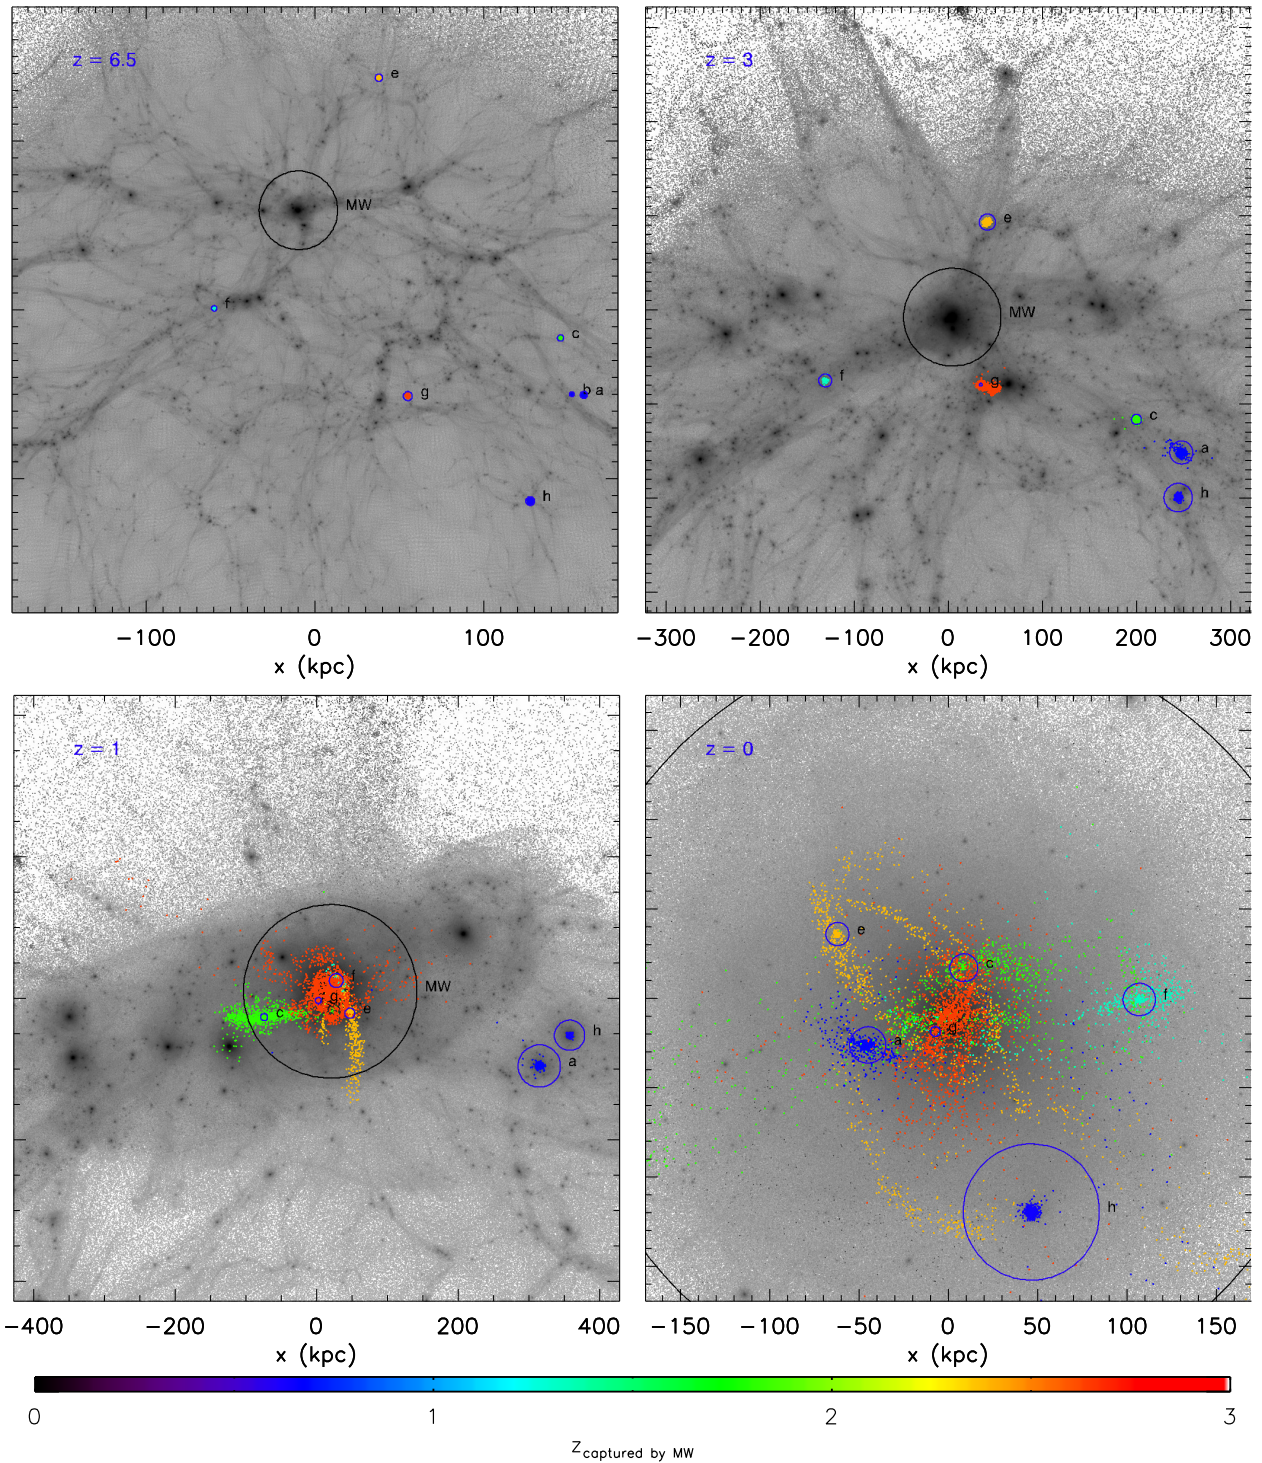


Figure 4.5: Tracking the location of dark matter particles in MW satellite progenitor halos from $z = 6.7$ to $z = 0$ in the Reference run. Each image is a projection of a cubic volume of length shown on the x-axis in physical kpc. Blue circles represent the virial radius of each halo tracked; the MW halo is shown as a black circle. The image at $z = 0$ lies largely inside the Milky Way virial radius. The colour of the particles in a halo represents the redshift at which the halo is captured and becomes a subhalo of the MW (see the colour bar). Halos captured before $z = 1$ exhibit significantly more stripping at $z = 0$ than halos captured after $z = 1$. The dark orange halo (halo ‘g’ in figure 4.8) is captured and partially stripped by another halo which, in turn, is captured and completely disrupted by the Milky Way between $z \leq 3$ and $z = 1$, while halo ‘g’ itself survives as a luminous satellite galaxy of the MW at $z = 0$.

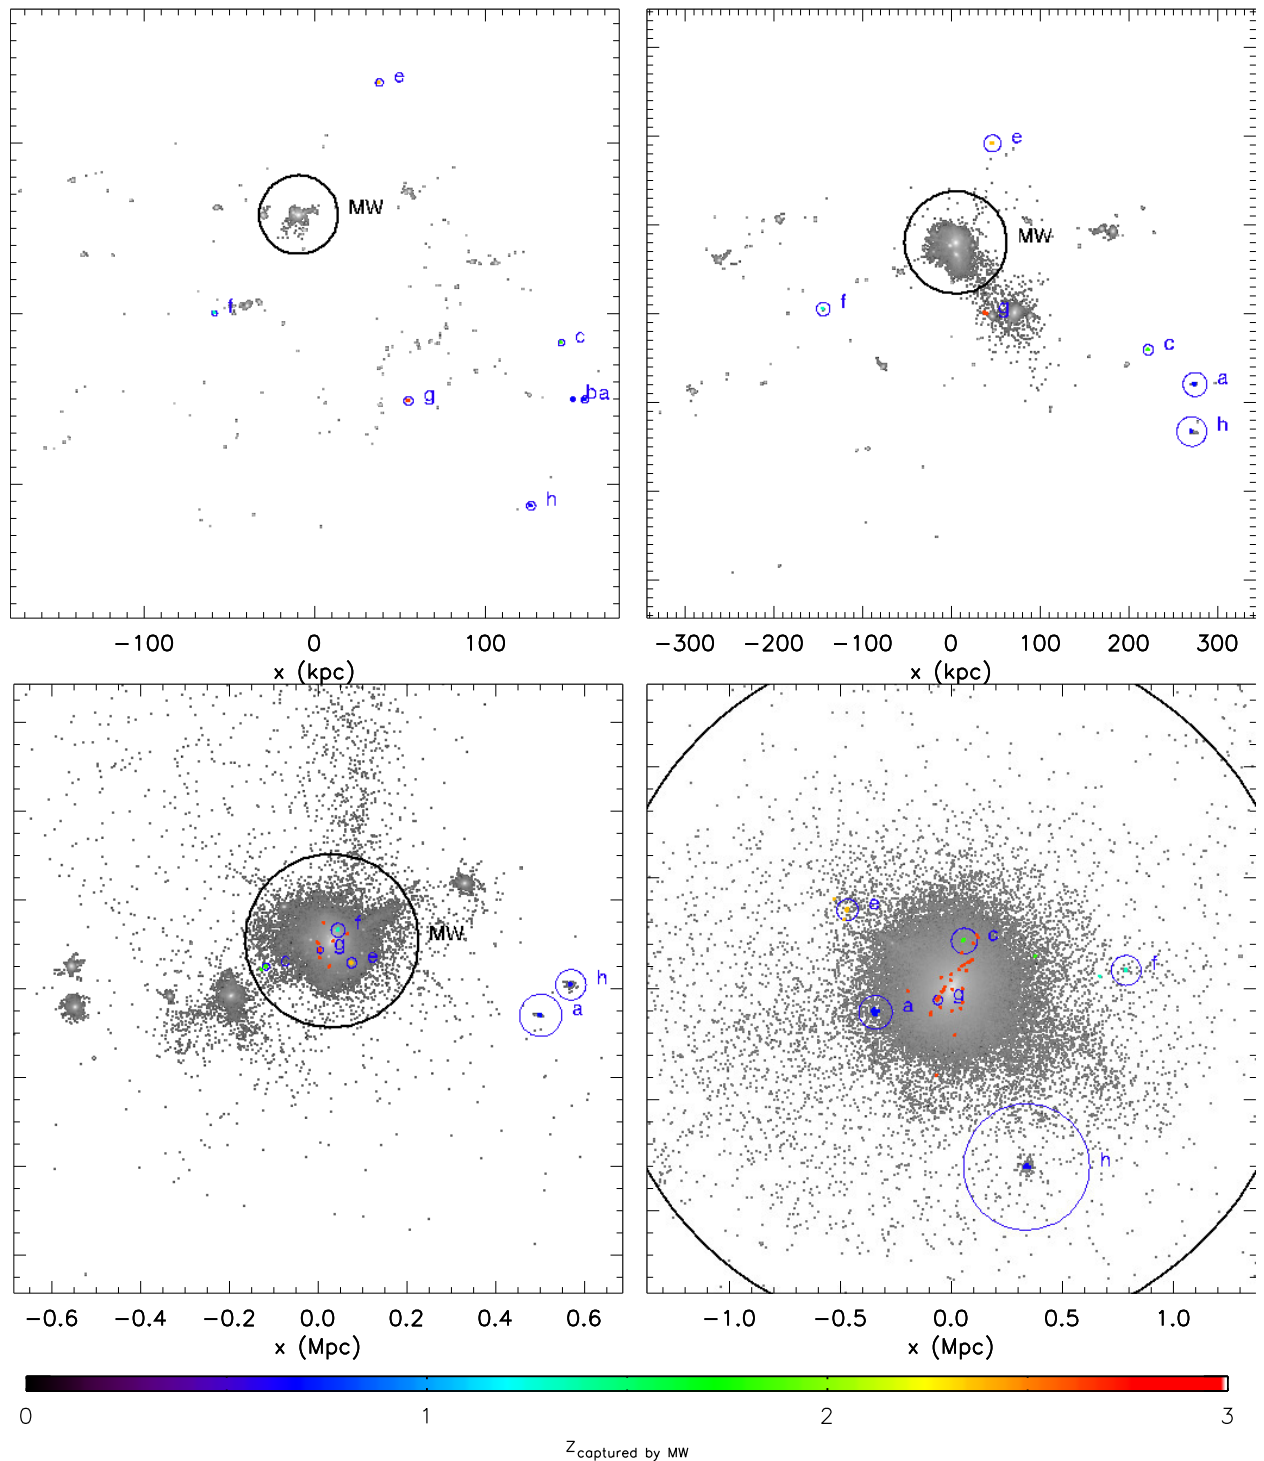


Figure 4.6: Tracking the location of star particles in MW satellite progenitor halos from $z = 6.7$ to $z = 0$ in the Reference run. As in Fig. 4.5, each image is a projection of a cubic volume of length shown on the x-axis in physical kpc and circles indicate virial radii. The background greyscale indicates the stellar column density. The colour of the particles in a galaxy represent the redshift at which the galaxy is captured and becomes a satellite of the MW (see the colour bar). With the exception of galaxies captured by the MW around $z \leq 3$ ('e', 'f' and especially 'g'), which form stellar streams around the MW centre, most stars formed before $z = 6.7$ remain part of their host galaxy.

High z	Simulation	Total Number high z (twinned)	Total Number $z = 0$	Satellite Mergers Destroyed (high $z \rightarrow z = 0$)	Merged with MW (high $z \rightarrow z=0$)
11	Cooling	66 (36)	1	0 (0 \rightarrow 0)	65
	Feedback	65 (35)	1	0 (0 \rightarrow 0)	64
	Reference	36 (36)	0	0 (0 \rightarrow 0)	36
9	Cooling	91 (61)	2	0 (0 \rightarrow 0)	89
	Feedback	89 (60)	2	0 (0 \rightarrow 0)	87
	Reference	63 (63)	1	0 (0 \rightarrow 0)	62
8.5	Cooling	89 (71)	3	0 (0 \rightarrow 0)	86
	Feedback	89 (71)	3	0 (0 \rightarrow 0)	86
	Reference	77 (77)	2	0 (0 \rightarrow 0)	75
8	Cooling	82 (68)	3	1 (2 \rightarrow 1)	78
	Feedback	90 (73)	3	1 (2 \rightarrow 1)	86
	Reference	85 (85)	4	0 (0 \rightarrow 0)	81
6.7	Cooling	85 (78)	6	1 (2 \rightarrow 1)	78
	Feedback	85 (79)	6	1 (2 \rightarrow 1)	78
	Reference	107 (107)	7	1 (2 \rightarrow 1)	99

Table 4.2: Table of the fate of galaxies formed between $z = 11$ and $z = 6.7$. The 6 columns are, from left to right: (1) the redshift at which the stellar population is sampled in the Reference, Cooling and Feedback runs ('high z '); (2) the simulation name; (3) the number of galaxies which become satellites of the MW between the sampled redshift and $z = 0$ not including the main MW progenitor halo (in brackets, the number of those galaxies whose twin halos in the Reference run also contain at least a galaxy); (4) the number of galaxies surviving as MW satellites at $z = 0$; (5) the satellite progenitors destroyed by mergers with other satellite progenitors - the figures in brackets show the number of objects taking part in mergers at high z , followed by the resulting number of objects after the satellite progenitor-satellite progenitor mergers at $z = 0$; (6) the number of galaxies merged with the MW and destroyed between high z and $z = 0$. See section 4.2.3 for a description of how these numbers are calculated. We find that the large majority of the halos containing galaxies captured by the MW by $z = 0$ merge with it and are destroyed, with two galaxies merging with each other before being captured by the MW and becoming a satellite galaxy. More mergers of MW progenitor galaxies are found, but these are all completely disrupted and destroyed after capture by the MW. We find that satellite galaxy formation in the Reference run is not complete by the lowest redshift reached by the high resolution runs ($z=6.7$). By looking at the ages of the star particles in satellite galaxies at $z = 0$ in the Reference run, we find that satellite galaxy formation continues until at least $z = 4.8$ (see Fig. 4.4).




Figure 4.7: Cumulative maximum circular velocity (V_{\max}) functions comparing our results to high resolution dark matter only simulations of MW-like objects at $z = 0$ published in the literature. The Reference run data is shown as a thick solid line, and the Dark Matter run as a thick dashed curve. Via Lactea II (VLII) data (Diemand et al., 2008) is overplotted as a thin dotted curve. A fit to the Aquarius data (Springel et al., 2008) is plotted as a thin dashed line. The empirical formula for $N(<V_{\max})$ given by Reed et al. (2005) is plotted as a thin solid line. Kinematically-derived values from observations given in Mateo (1998); Bekki & Chiba (2005); Bekki & Stanimirović (2009); Wolf et al. (2010) are shown as diamonds (see, e.g. Wolf et al. (2010) for an explanation of how these values are derived). In each case, V_{\max} is given by $\max(\sqrt{GM(<r)/r})$, where M is the mass inside r_{50} , the radius at which the density exceeds $50\rho_{\text{crit}}$. Our host halo has a V_{\max} of 126km/s in the Dark Matter run as opposed to ~ 200 km/s in the Aquarius or VLII simulations. Note that we plot all subhalos inside r_{50} (denoting the radius at which the density is above $50\rho_{\text{crit}}$) as opposed to r_{200} , which we use in the rest of this chapter. This is to allow comparison with Reed et al. (2005); Diemand et al. (2008); Springel et al. (2008). The Dark Matter run is reasonably well described by the Reed et al. (2005) formula indicating that the difference between our simulation and the Aquarius and VLII simulations are largely due to the MW host halo in the Aquarius and VLII having higher maximum circular velocities. By contrast, the Reference run lies below the empirical formula, which we comment on in section 4.5.1.

4.3.2 Dark Matter Satellite Halos

In this section, we consider the population of dark matter subhalos of the Milky Way in our runs at $z = 0$, comparing and contrasting their properties with similar dark matter simulations which exist in the literature.

For this purpose, we run a fourth simulation, the Dark Matter run, which is identical to the Cooling and Feedback runs except that the mass in baryons is replaced with dark matter mass, such that a dark matter particle is $1/(1 - f_b)$ times the mass of a dark matter particle in the Cooling and Feedback runs (see section 4.2.1). This simulation is described in more detail in section 4.5.1. We use it to compare our results directly with the Aquarius (Springel et al., 2008) and Via Lactea II (Diemand et al., 2008) simulations, which are the most resolved dark matter only simulations of MW-like objects available to date. The most basic comparison, a cumulative maximum circular velocity function, is presented in Fig. 4.7. In this figure, we also overplot the empirical prescription proposed by Reed et al. (2005) for the Milky Way halo in our Dark Matter run at $z=0$. We find that our Dark Matter run satellite halo data is well represented by this prescription, but that the Reference run predicts significantly fewer satellites at the low velocity end (between a factor 2 and 3 for $V_{\max} < 20\text{km/s}$), and more massive satellites ($V_{\max} > 30\text{km/s}$). We discuss the impact of simulation physics on the maximum circular velocity function in section (4.5.1).

In figure 4.7 we plot the data for observed Milky Way satellites given in Mateo (1998); Bekki & Chiba (2005); Bekki & Stanimirović (2009); Wolf et al. (2010). Above $V_{\max}=10\text{km/s}$, this trend matches our dark matter run well, though note that the normalisation of the curve depends on the host maximum circular velocity. Below 10km/s , the curve drops off. This may be down to low-mass satellite detection limits in the observations, which we do not correct for in this data. However, for the more massive satellites the data are likely to be complete.

From Fig.4.7, it is apparent that, whilst the cumulative maximum circular velocity function of our Dark Matter run has the same shape ($N(> V_{\max}) \propto V_{\max}^{-3}$) as that measured in both the Aquarius and Via Lactea II simulations, its normalisation is more than an order of magnitude lower. This large discrepancy can almost entirely be attributed to our choice for the mass of the MW host halo since our agreement with the Reed et al. (2005) prescription is quite reasonable (better than 20 %). Springel et al. (2008) report that their simulations overshoot the fitting formula of Reed et al. (2005) by a factor ~ 3 which they argue most likely arises from a systematic effect in the numerical technique used to perform the runs. On the other hand, Madau et al. (2008)

suggest it is due to the different (WMAP 1 instead of WMAP 3/5/7, i.e. different normalisation, σ_8 , and/or tilt, n_s , of the power spectrum) cosmology the Aquarius simulations employ. In any case, the remarkable conclusion that we draw from this comparison exercise, is that for a MW host halo with $V_{\text{max}}=126$ km/s in the Dark Matter run as opposed to ≈ 210 km/s in the Aquarius or VLII simulations, i.e. a difference of about 65 %, one gets a suppression in the number of satellites by about a factor of 10, which is enough to match the observed abundance of MW satellites with V_{max} between 10 and 30 km/s. Now, there is still an ongoing debate as to what the exact mass of the Milky Way halo is (e.g. Battaglia et al. (2005); Karachentsev & Kashibadze (2005); Watkins et al. (2010)). Our simulated halo admittedly lies at the very low end of the estimated range of values ($4.32 \times 10^{11} M_{\odot}$ within 195 kpc), and Aquarius and Via Lactea II somewhat on the high side ($1.85 \times 10^{12} M_{\odot}$ within 245 kpc). Note that in both instances the choice of mass is justified within the constraints given by observations. The low choice of mass was further motivated by considerations for other projects using the same data, e.g. measuring the properties of a “typical” galaxy at $z = 3$. It is not the purpose of the present work to create new constraints Milky Way halo mass or to argue which mass is more realistic, but only to illustrate how the uncertainty in the mass of the Milky Way halo and the inclusion of baryonic physics translate into an uncertainty in the number and properties of MW satellites one predicts. Thus we simply remark that the observed abundance of these satellites, taken at face value, seems to favour a light MW halo.

4.4 Feedback in Milky Way satellite progenitors

4.4.1 Supernova feedback

We now discuss the differences between the simulations with and without supernova feedback, in order to investigate the importance of the role of supernovae in high-redshift dwarf galaxy and Milky Way satellite formation. We use the twinning method described in chapter 3 to match halos in the Cooling and Feedback runs. We can thus determine whether the net effect of including supernovae in our sub-parsec resolution simulations enhances or suppresses star formation in halos of various masses.

Fig. 4.8 shows images of the progenitors of each satellite galaxy and the main Milky Way progenitor at $z = 6.7$. We find that every satellite progenitor has supernova outflows of different sizes, evidenced by metal-enriched outflows from the galaxy. Whether galaxies at $z = 6.7$ become satellite galaxies of the Milky




Figure 4.8: Progenitors of the Milky Way satellite galaxies formed before $z=6.7$ in the Feedback Run. The colour map shows projections of the density in red, the density-weighted temperature in yellow and the metallicity in green (highlighting the metal-enriched galactic outflows). The outer images are cubic projected images of satellite progenitors with a length of 10 dark matter virial radii, with the stellar column density overplotted as contours in black. The central image is a cubic projection of the full refinement region. White circles in the central image are galaxies, and red circles are halos containing galaxies that survive to become Milky Way satellites at $z = 0$ in the Feedback Run. The labels ‘a’-‘h’ indicate the position of the individual satellite images on the central image, with black boxes marking the size of the outer image on the central image. The Milky Way main progenitor is labelled ‘d’.

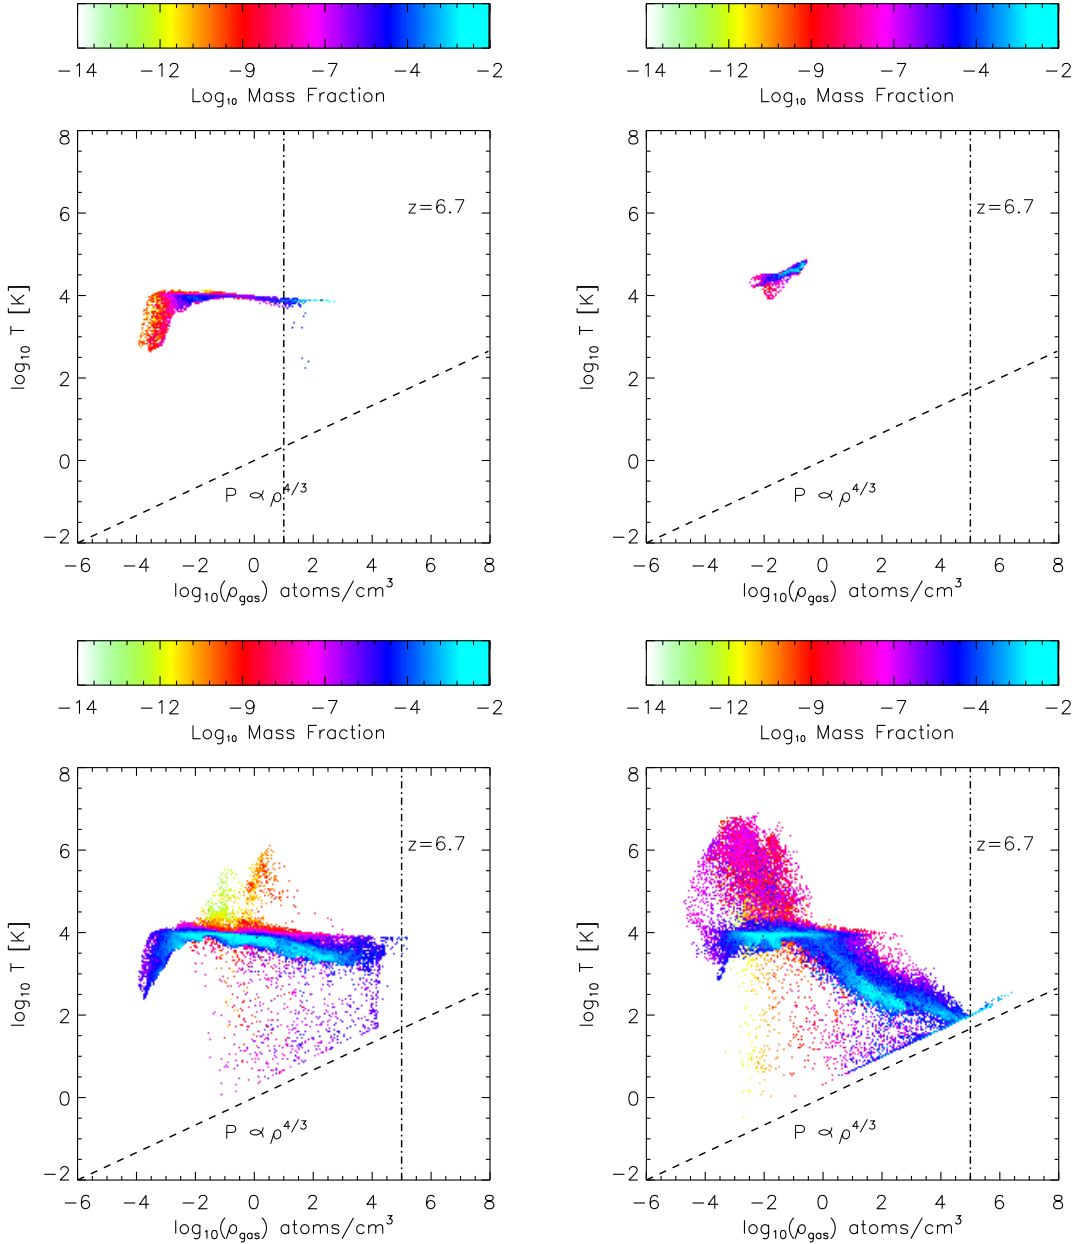


Figure 4.9: Phase diagrams of the gas inside r_{vir} of a single satellite progenitor at $z = 6.7$ (halo ‘g’ in Fig. 4.8). The run are, from top left to bottom right, Reference run, Adiabatic run, Cooling run, and Feedback run. The bulk of the gas in the halo is just below the T_{vir} ($\sim 10^4\text{K}$). The Feedback run exhibits high-temperature plumes corresponding to supernova-heated gas, while the Cooling run shows a tiny fraction of shock-heated infalling material. The Reference run does not contain any gas heated above the virial temperature. The Adiabatic run, by contrast, does contain gas above the virial temperature, but it exhibits neither dense gas nor gas significantly below the virial temperature, with most of the gas following the roughly slope of the polytrope. There is more cold, dense gas in the Feedback run than in the Reference and Cooling runs due to more efficient cooling from the injection of metals into the gas by supernovae. The dashed line shows the polytropic equation of state $P \propto \rho^{4/3}$, which is imposed between T and ρ to avoid numerical fragmentation. The vertical dot-dash line shows the density above which star formation can occur (10 at/cm^3 in the Reference run and 10^5 at/cm^3 in the Cooling and Feedback runs; no star formation can occur in the Adiabatic run).

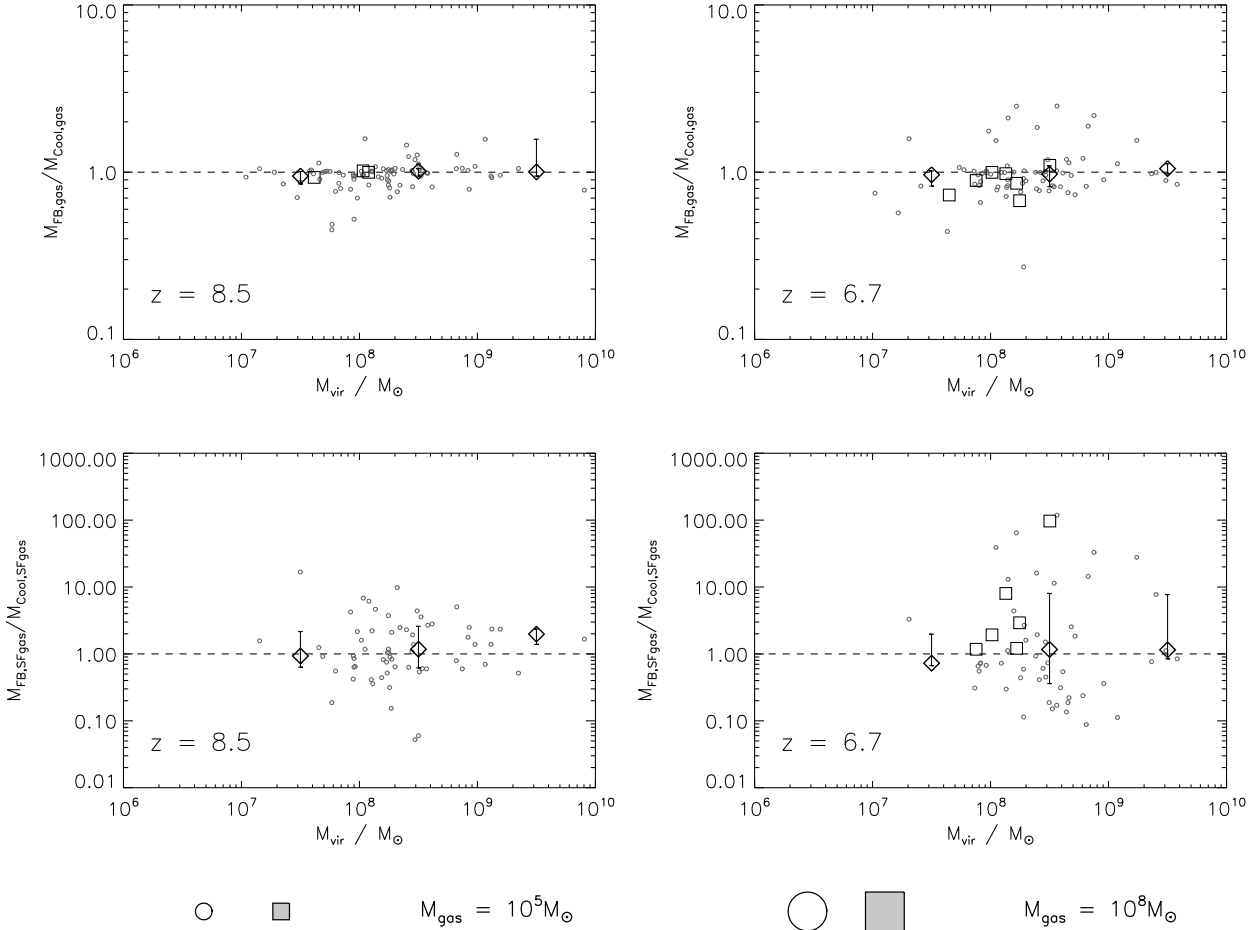


Figure 4.10: Ratio of gas mass in galaxies within r_{vir} in the Cooling and Feedback runs, plotted against the total halo mass of their twin halo in the Reference run. The figure shows galaxies at the redshift of reionization $z = 8.5$ on the left, and $z = 6.7$ on the right. The top plots show the ratio of total gas mass between runs, while the bottom plots show the ratio of the star-forming gas mass (i.e. $\rho > 10^5 \text{ atoms/cm}^3$). A black square indicates that the halo of that galaxy survives as a Milky Way satellite at $z = 0$; grey circles are halos that are completely disrupted by $z = 0$. We overplot as diamonds with error bars the median and interquartile range of the fractional differences in halo mass bins $10^7 M_{\odot} - 10^8 M_{\odot}$, $10^8 M_{\odot} - 10^9 M_{\odot}$ and $10^9 M_{\odot} - 10^{10} M_{\odot}$. The median values lie around the horizontal line marking an equal ratio, with the $10^7 M_{\odot}$ to $10^8 M_{\odot}$ bin having a ratio of ~ 0.95 and the higher mass bins having a 1:1 ratio or higher. There is a large amount of scatter in the relative amounts of star forming gas in halos in the two simulations. It is worth noting that some of the halos at each redshift sampled, including all of the Milky Way satellite progenitors at $z = 8.5$, do not contain star-forming gas. This is explained as star formation occurring in bursts, with the smaller galaxies containing no star-forming gas at certain instants in time.

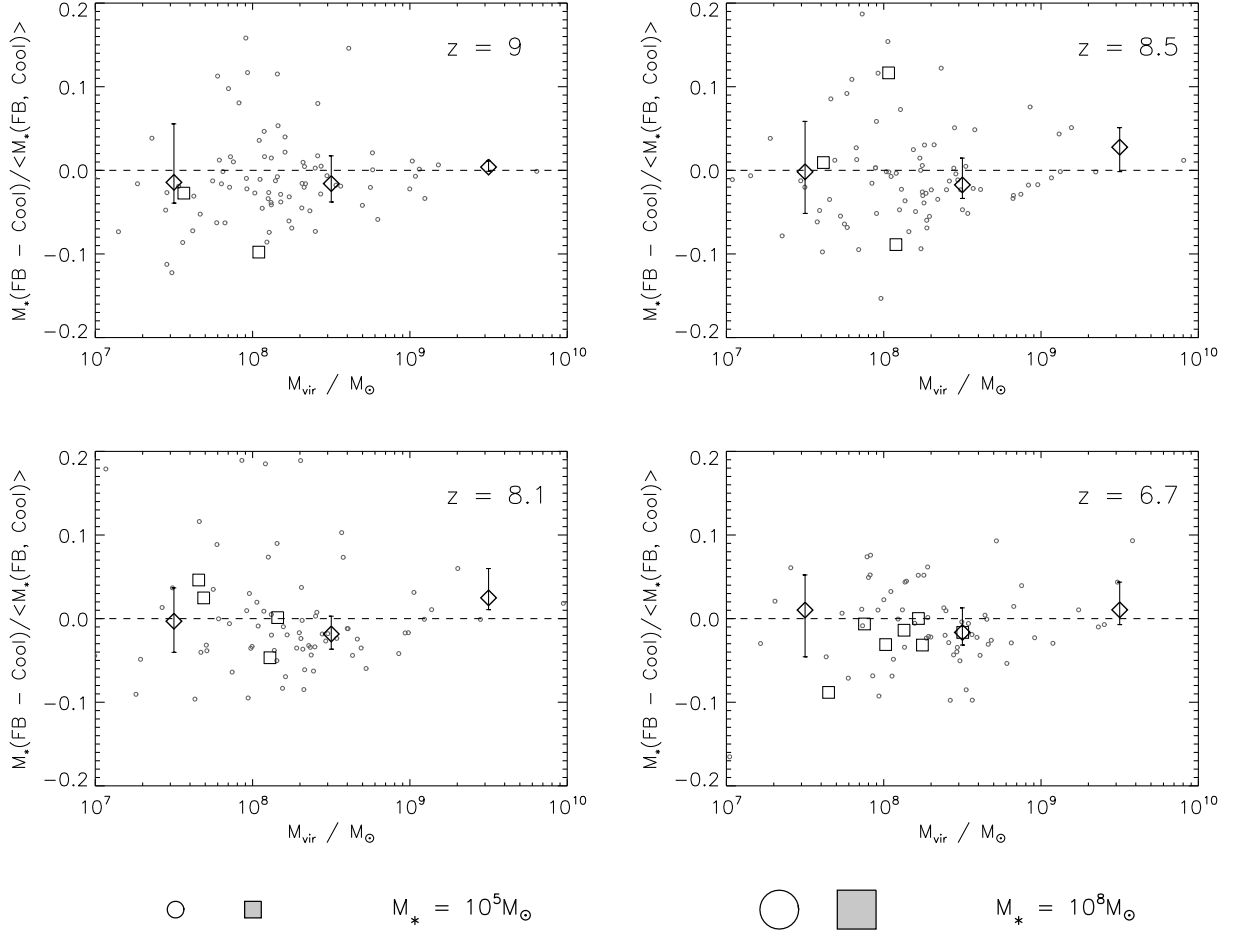


Figure 4.11: Fractional difference between stellar masses of galaxies in the Cooling and Feedback run ($M_{* \text{FB}} - M_{* \text{Cool}} / (0.5(M_{* \text{FB}} + M_{* \text{Cool}}))$) versus total twin halo mass in Reference Run, where $M_{* \text{FB}}$ is the stellar mass in the Feedback run and $M_{* \text{Cool}}$ is the stellar mass in the Cooling run). A positive value indicates that the inclusion of supernova feedback enhances star formation in the given galaxy, while a negative value means that supernova feedback suppresses star formation. Different panels show the values for different output redshifts, from $z = 9$ (top left) to $z = 6.7$ (bottom right). Black squares represent halos containing galaxies which survive as a Milky Way satellites at $z = 0$; a halo represented by a grey circle is completely disrupted by $z = 0$. We overplot as diamonds with error bars the median and interquartile range of the fractional differences in halo mass bins $10^7 M_{\odot} - 10^8 M_{\odot}$, $10^8 M_{\odot} - 10^9 M_{\odot}$ and $10^9 M_{\odot} - 10^{10} M_{\odot}$. Star formation is slightly suppressed in low mass halos with weaker gravitational potentials. The trend is reversed for high mass halos.

Way at $z = 0$ or are disrupted by interactions with the Milky Way halo is discussed in section 4.3.1.

Powell et al. (2010) discuss the differences in gas temperature and density when adding supernova feedback using phase diagrams of the Milky Way progenitor in the Cooling and Feedback runs. We now perform a similar study for a sample satellite galaxy progenitor. Fig. 4.9 shows typical phase diagrams for a star-forming halo at $z = 6.7$ (halo ‘g’ in Fig. 4.8) inside r_{vir} for each run. Following Powell et al. (2010), the phase diagram can be divided into distinct sections.

The first is the cold dense phase, containing the gas filaments. This gas, at a temperature $\sim 10^4\text{K}$ and densities comprised between $\sim 10^{-2}$ and 1 at/cm^3 accounts for a large fraction of the total gas in the halo, as shown by the colour coding in Fig. 4.9 where blue represents the largest mass fraction. Above $\sim 1 \text{ at/cm}^3$, the ISM begins and the temperature begins to drop below 10^4K especially in the presence of metals (Feedback Run). The vertical dash-dot lines in Fig 4.9 shows the densities at which we allow star formation to occur (10 at/cm^3 in the Reference run and 10^5 at/cm^3 in the Cooling and Feedback runs).

The final phase in Fig. 4.9 is the warm/hot diffuse phase, which represents gas that has been heated to or above the virial temperature of the halo. The Cooling and Feedback runs both produce gas in this phase. In both runs, a small quantity of gas falling into the halo diffusely, outside of the cold filaments, shock-heats. Due to its lower resolution, the Reference run does not exhibit this effect. In addition to this shock heated material, the Feedback run shows extra low-density hot gas originating from supernova explosions. These heat gas and potentially remove some from the centre of the halo, reducing star formation.

We now proceed to quantify to which extent the positive feedback processes (metal cooling, blastwave compression) or negative feedback processes (gas heating, outflows) dominate in halos of different masses. In Figs 4.10 and 4.11, we plot comparisons of the total gas mass, star-forming gas mass and total stellar mass in each halo in the Cooling and Feedback runs, using the halo twinning procedure described in chapter 2. Star-forming gas is defined as gas with a density above 10^5 at/cm^3 our density threshold for star formation. In each figure we overplot the median and interquartile range of the fractional differences in the mass bins $10^7 M_\odot - 10^8 M_\odot$, $10^8 M_\odot - 10^9 M_\odot$ and $10^9 M_\odot - 10^{10} M_\odot$. For the gas masses, since there exists a large scatter in the results we plot the ratio for each halo on a log scale to highlight both large and small differences. For the stellar masses, since differences are smaller, we plot the fractional difference between the runs. In other words, if we denote the ratio between stellar masses by R , we plot the quantity $(R - 1)/(0.5(R + 1))$.

Although Fig 4.10 shows a large scatter in ratios of gas masses of twinned halos in the Cooling and Feedback runs, the median values lie around an equal ratio, with the $10^7 M_{\odot}$ to $10^8 M_{\odot}$ mass bin showing $\sim 5\%$ less gas and the halos in the higher mass bins having a similar or slightly higher gas content in the Feedback run. This goes in the expected direction since supernovae eject gas back into the galaxy, causing the total gas mass to increase if this ejecta is unable to escape the halo. There is also considerably more scatter in the instantaneous star-forming gas mass results, with some halos in the Feedback run containing over 100 times the mass of star-forming gas than their Cooling run twins. This effect can be attributed to the enhanced metal cooling which takes place after the first supernovae explode in the Feedback run. However, on average we find a similar pattern to the total gas mass ratios, i.e. lower mass halos in the Feedback run contain very slightly less star-forming gas than their Cooling run counterparts, and more massive halos slightly more, although we note that at $z = 8.5$ the highest mass bin has a median star-forming gas mass that is twice as high. The large scatter in the amount of star-forming gas is expected; Stinson et al. (2007) also find that star formation in their dwarf galaxies is quite bursty, because the instantaneous mass of star-forming gas can strongly fluctuate on short timescales, driven by catastrophic non-linear events (instabilities, mergers).

This is why we plot in Fig. 4.11 a time integrated quantity, the fractional difference between the stellar mass of each twinned halo that has formed stars in the Cooling and Feedback runs. Unsurprisingly, there is a maximum 20% difference between values, much smaller than that for the star-forming gas mass. For similar reasons, we also find more scatter in stellar mass ratio of low mass halos than of high mass halos: the length of time these former have been forming stars is generally shorter and therefore they are more affected by temporary fluctuations in their star formation histories. As with the gas mass comparison, we find that the positive feedback processes outweigh the negative feedback processes in the highest mass bin, leading to a net increase in the median fractional difference in stellar mass of a few percent when supernova feedback is added.

In summary, we find that the effect of feedback on the gas mass and star formation in a halo is complex, with lower mass halos being on average more affected by negative feedback processes such as outflows and gas heating, and higher mass halos by positive feedback processes such as blastwave compression and metal cooling. We also find that stellar masses in individual twin halos can differ by up to 20%; however, the

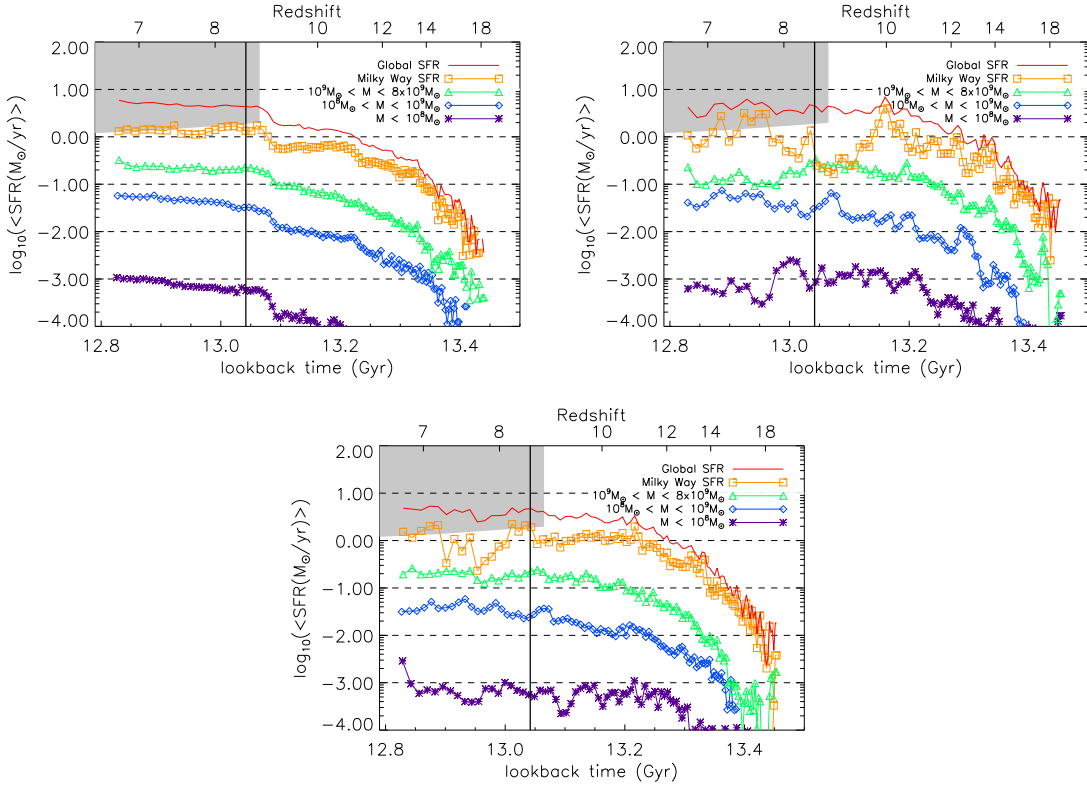


Figure 4.12: Mean star formation histories for different halo mass bins in the Reference run (top left) Cooling run (top right) and Feedback run (bottom). Mean star formation rates (SFR) in M_{\odot}/yr are given for halos of mass $M < 10^8 M_{\odot}$ (purple asterices), $10^8 M_{\odot} < M < 10^9 M_{\odot}$ (blue diamonds) and $10^9 M_{\odot} < M < 8 \times 10^9 M_{\odot}$ (green triangles), as well as the Milky Way (orange squares) and the global SFR for the entire high resolution region (red solid line). The vertical line at $z = 8.5$ (look-back time 13.042Gyr) shows the point at which reionisation is turned on in the simulations. The grey region shows the detectable star formation rates as determined by Wilkins et al. (2010). This suggests that the SFR for a Milky Way-like galaxy progenitor is almost detectable at $z \sim 8$ (lookback time ~ 13.0 Gyr). We find no sudden drop in star formation in any mass bin after reionisation for any of the simulations. In fact, some of the star formation rates increase by up to 0.5 dex at reionisation. The jumps in star formation rates in the Reference run are due to the triggering of a new level of refinement on the grid, which allows the gas to collapse to allow star formation in all potentially star-forming regions (Rasera & Teyssier, 2006).

median values only differ by a few percent in all mass bins, although values for low mass halos are more scattered. We therefore conclude that supernovae do not seem to have a significant effect on Milky Way satellite progenitors, at least at redshifts larger than 6.




Figure 4.13: Maximum circular velocity of satellite halos in the Reference run at $z = 0$ versus the maximum circular velocity of their main progenitor at, from top left to bottom right, $z = 8.5$ (just prior to reionisation), $z = 6.7$, $z = 3$, $z = 1$. In red are the satellites that contain stars at $z = 0$ in the Reference run; the blue satellites remain dark. The dotted horizontal line is at 12km/s and represents the threshold given in Okamoto & Frenk (2009) above which halos can form stars before reionisation. Like these authors, we find halos that form stars under this threshold to about 10km/s ; below this, no halos can form stars. Note that this threshold seems independent of redshift and that surviving Milky Way satellite galaxies begin to be captured by the Milky Way progenitor at $z = 3$. The halos with $v_{\text{max},z=0}$ around 80km/s are satellites labelled ‘a’ and ‘h’ (see Fig. 4.8).

4.4.2 Reionisation Feedback

At $z = 8.5$ in each simulation we include a simple instantaneous, uniform heating term that represents the UV background (Haardt & Madau, 1996). There are two ways such a background affects galaxy formation/evolution in our simulations: (i) it heats the ISM which could prevent the gas density within galaxies from crossing the star formation threshold and (ii) it heats the IGM which could cut off the gas accretion onto galaxies. Note that we very likely overestimate these effects in the simulation as we neglect self-shielding, which is known to occur around densities $n_{\text{H}} \gtrsim 0.1 \text{ at/cm}^3$ (Susa & Umemura, 2004). On the other hand, we do not account for the local UV radiation (from stars within the galaxy itself) which also photo-ionises the ISM/IGM. Nevertheless, by looking at the star formation rate in halos before and after $z = 8.5$, we should be able to estimate how efficient non-local ionisation is at halting star formation.

Fig. 4.12 shows that reionisation does not immediately stop star formation in halos already forming stars, in agreement with Kitayama et al. (2001); Machacek et al. (2001); Gnedin & Kravtsov (2006); Okamoto & Frenk (2009); Wadepuhl & Springel (2010a). Even for the lowest mass halos ($M_{\text{vir}} < 10^8 M_{\odot}$), we find that star formation continues after the uniform UV background is turned on. The free-fall time of a test particle falling from r_{vir} into one of the smallest galaxies formed at $z = 8.5$ is on the order of 50 Myr. Hence, we conclude that star formation is not stopped in these halos even after ~ 5 free-fall times (the amount of time elapsed between $z = 8.5$ and $z = 6.7$). In other words, if reionisation does halt star formation by heating up the ISM or cutting off the gas accretion in halos that have already formed stars, it does not do so abruptly, but rather over a significantly extended period of time.

Another potential effect of reionisation is to quench galaxy formation by preventing the collapse of gas within halos that have not yet formed stars (e.g. Gnedin (2000a); Somerville (2002); Benson et al. (2002)). However, the last $z = 0$ satellite galaxy to be formed in our Reference run begins forming stars at $z = 4.8$ in a halo with $M_{\text{vir}} = 1.4 \times 10^7 M_{\odot}$, and there are 9 other satellites galaxies hosted by halos with a similar mass which form their first star after $z = 8.5$. Hence, whilst it is still possible that UV photoionisation has a long-term role in preventing some galaxies from forming, it does not seem to be able to halt galaxy formation entirely.

In figure 4.13, we recast this statement in terms of minimal circular velocity for a star forming halo, v_{max} , below which halos are prevented from forming stars. This allows us to directly compare our results to

Redshift	Halo Type	Dark Matter	Adiabatic	Reference	Cooling	Feedback
0	Independent Halos	2084	1706	1731	-	-
	Subhalos	2730	2418	2209	-	-
7	Independent Halos	2405	1997	1997	2142	2104
	Subhalos	598	449	343	445	390

Table 4.3: Total number of independent halos and subhalos (dark and luminous) in the Dark Matter, Adiabatic and Reference runs at two sampled redshifts. Runs containing baryons compare favourably, whereas the Dark Matter run contains significantly more halos for a given redshift. We attribute this to the effect of pressure in baryons resisting gravitational infall. Dark matter particles do not experience this and are able to collapse more easily into halos under gravity. There are also fewer subhalos in the Reference run than either of the other two runs, as explained in section 4.5.1.

those presented by Okamoto & Frenk (2009). Like these authors, we cannot definitively conclude that this threshold arises entirely because of reionisation or the general inability of halos below a v_{\max} of 10km/s to cool and form stars by $z = 0$, since we do not run a simulation without reionisation. However, we note that reionisation in our simulation occurs instantaneously at $z = 8.5$ (close to the value of $z = 9$ of Okamoto & Frenk (2009)) and that, in stark contrast the threshold of $v_{\max} \approx 10$ km/s seems independent of redshift. Indeed, it remains quite constant both before and after reionisation has occurred, which leads us to argue that reionisation cannot play an important role in setting its level and only sustains it, in the best of cases.

4.5 Stripping and Satellites at $z = 0$

4.5.1 The Effect of Baryonic Physics on Satellite Galaxy Survival

In this section we discuss the effect of the baryonic physics included in the Reference run on the survival of satellite galaxies to $z=0$. Various authors have investigated the issue by running SPH galaxy simulations and comparing them to identical simulations in which the gas particles in the initial conditions are removed and their mass added to the dark matter particles, as we do in our Dark Matter run. For instance, Peirani (2010) adopts this strategy to study the Hubble flow on Mpc scales. He finds that the peculiar velocities of independent halos match well between simulations and that a cut in the maximum halo circular velocity of 30km/s in the pure dark matter run yields an excellent match with the halos containing galaxies in the baryonic run. Hence he concludes that for independent halos, predictions of which halos contain galaxies are robust. By contrast, we find that the baryonic physics included in our high resolution re-simulation of a




Figure 4.14: Density profiles of the main Milky Way-like halo up to r_{vir} at $z = 0$ for the Reference (thick solid line), Adiabatic (dotted line) and Dark Matter (dashed line) runs. We plot total density, i.e. dark matter, gas and stars except for the dot-dashed line which represents the DM density in the Reference run. We overplot a vertical dotted line at 200pc, equivalent to AMR level 16 at $z = 0$, i.e. 2 levels below the highest resolution reached by the runs to indicate the scale below which the gravitational force is underestimated, and lines at r_{200} , used in the bulk of the text and r_{50} , used for comparison purposes with figure 4.7. Power law profiles scaling like r^{-1} and r^{-2} and a NFW profile fit to the DM run (thin solid line) are also overplotted. The density profile in the Reference run departs from the others at 10kpc, and is several times higher at 2kpc. There is little noticeable difference between the Dark Matter and Adiabatic run density profiles. The high central density in the Reference run is due to gas cooling, and dramatically increases the stripping of sub-halos that pass in the central region of the halo. This effect is noticeable in figure 4.15, where the number of surviving satellite galaxies in the Reference run is dramatically lower than the number found in the Adiabatic and Dark Matter runs.




Figure 4.15: Number of satellite galaxies at $z = 0$ formed at or before a given redshift. Satellite galaxies are defined as sub-halos whose progenitors at high redshift have a twin in the Feedback run that contains one or more galaxies (see section 4.3.1). Each line represents a different run used to track halos down to $z = 0$. The solid line shows the Reference run, the dashed line shows the Dark Matter run and the dotted line shows the Adiabatic Run. The inclusion of gas cooling reduces the number of galaxies at high redshift that survive down to $z = 0$ by a factor ~ 3 to 10 depending on redshift. As discussed in the text, the central density profile of the Milky Way at $z = 0$ plays a key role in governing the survival of satellite galaxies.




Figure 4.16: Survival of halos as MW satellites at $z = 0$. Both panels show the redshift at which a halo becomes a MW sub-halo, z_{capture} , against the sub-halo mass when this capture happens, M_{capture} ; on the left is the Reference run, and on the right is the Adiabatic run. Black squares survive as MW satellite halos at $z = 0$. Note that they can be completely disrupted as long as they become part of another MW satellite halo by $z = 0$ in the disruption process. Grey circles are completely disrupted and simply become part of the diffuse MW halo by $z = 0$. Halos captured by the MW host before $z = 3$ do not survive to $z = 0$, with one exception in the Adiabatic run. Less than half of the low-mass halos ($M_{\text{vir}} \leq 10^9 M_{\odot}$) captured after $z = 3$ are able to survive to $z = 0$ in the Reference run, whereas the vast majority of them survive in the Adiabatic run.




Figure 4.17: Properties of a sampled satellite galaxy with time as it falls through the Milky Way halo. The plot on the right shows the halo Reference run (i.e. with gas cooling), and the left its twin in the Adiabatic run (i.e. without gas cooling). The blue line is the distance between the satellite and Milky Way centre. The red line is the mass of the satellite. The green line is 10% of the virial radius of the Milky Way halo, i.e. the approximate radius at which the density profile in the Reference run rises above that in the Adiabatic run.




Figure 4.18: r_{\max} vs v_{\max} for halos containing stars at $z=0$. Halos in the Reference run are shown as red diamonds. The blue shaded region represents the values Boylan-Kolchin et al. (2011) calculate for the Milky Way dwarf spheroidals to within 2σ . Note that the two points to the right correspond to halos ‘a’ and ‘h’, which are on their first infall and are unlikely to be dwarf spheroidals, matching more closely to the LMC and SMC, which are not included in Boylan-Kolchin et al. (2011). Further, some of our halos lie to the left of the shaded region; we cannot comment on whether these objects are outliers or whether they correspond to satellites that are too small to be included in the list of observed satellites in Boylan-Kolchin et al. (2011). Note that the Reference run does not contain supernova feedback.

MW class halo has a dramatic effect on the number of satellite galaxies it possesses at $z = 0$.

As described in section 4.2.1, we run three simulations to $z = 0$. The first one is the Dark Matter run. This is a pure N-body simulation, in which the mass in baryons corresponding to Ω_B is replaced by an equivalent increase in dark matter mass (note that in this case, RAMSES solves the Poisson equation using an AMR grid refined according to the same quasi-Lagrangian refinement criteria applied to dark matter alone (Teyssier, 2002)). The second simulation is the Adiabatic run. This simulation is identical to the Dark Matter run, except that it now contains the universal baryon mass fraction and solves the Euler-Poisson equations to follow gas dynamics. However, contrary to the third simulation, the Reference run, gas in the Adiabatic run is not permitted to cool radiatively. Star formation is also omitted in this run, though a Haardt & Madau (1996) UV background is included to minimise differences with other baryonic runs. Finally, to extend our analysis to runs with higher resolution (Cooling run) and including more detailed physics (Feedback run), we use the technique described in section 4.3.1 to predict the number of luminous halos at $z = 6.7$ that survive to $z = 0$.

The first thing to note is that the Dark Matter run produces more (sub-)halos ($\approx 20\%$ more) than any of the runs containing baryons; we present exact numbers in table 4.3. This difference is attributed to gas pressure delaying/suppressing halo collapse. Indeed, the collapse of purely collisionless dark matter halos is unimpeded by pressure forces. This effect is roughly independent of redshift. Secondly, we find fewer subhalos in the Reference run than in the two other runs. Figure 4.14 shows that the total (gas & DM) density profiles of the main MW halo at $z = 0$ are very similar between the DM and Adiabatic runs, with the Adiabatic run having a slightly lower density overall. However, gas in the Reference run is able to cool, and hence the density of the halo within a radius of 10kpc is significantly higher in this run, since cooled baryons are able to condense at the centre, pulling DM along with them (Flores & Primack, 1994). This higher central density plays an important role in destroying subhalos, and hence satellite galaxies: as mentioned in section 4.3.2, the Reference run subhalo cumulative maximum circular velocity function lies below the Dark Matter run's (figure 4.7).

In figure 4.15, we further quantify this effect by plotting the number of high redshift galaxies surviving to become satellite galaxies of the Milky Way at $z = 0$ (see sections 4.2.3 and 4.3.1 for details on the tracking of satellites from high to low redshift and table 4.2 for numbers specific to the Reference run). We find that

the Reference run contains far fewer (factor 3 to 7 depending on redshift) surviving satellite galaxy sub-halo hosts than the Dark Matter or Adiabatic runs. We attribute this quite dramatic difference to the stronger dynamical friction on satellites in the central region as this drag force scales proportionally to the host halo density (Binney & Tremaine, 2008). Finally, we note that none of the satellite galaxies at $z = 0$, in any of the simulations, were captured by the Milky Way progenitor halo at $z > 3$; all galaxies captured by the Milky Way progenitor halo at redshifts higher than $z = 3$ are destroyed. However, in the Adiabatic run, all halos captured at $z < 3$ survive as sub-halos down to $z = 0$, as opposed to the Reference run, where only a small fraction of them survives (see figure 4.16).

In Fig. 4.16, we investigate satellite survival to $z = 0$ in more detail. We compare the redshift at which a halo is captured by the Milky Way halo against its mass at capture in two simulations: the Reference run and the matching Adiabatic run. The only difference between these two runs is that an extra right hand side 'sink' term is included in the energy equation of the gas in the Reference run to model losses due to radiative cooling, as well as star formation (see section 4.2.1). We find that out of all the halos which survive to $z = 0$ in these two runs (6 in the Reference run, 20 in the Adiabatic run), only one is captured at $z > 3$ in the Adiabatic run and none in the Reference run. Moreover, higher mass halos ($M_{\text{vir}} > 10^9 \sim M_{\odot}$) only survive if they are captured later so that the highest mass satellites at $z = 0$ are systematically the ones that are captured last. This trend with mass is naturally expected since it is well known that dynamical friction makes massive satellites sink to the centre of their host halo faster than their lower mass counterparts (e.g. Binney & Tremaine (2008)). Whilst the Chandrasekhar formula described by these authors yields an incomplete picture of the interaction of satellites with their host (see e.g. Taylor & Babul 2001), one expects the trend in mass to persist in spite of tidal stripping and gravitational heating. However, what is perhaps more surprising is the relatively large impact that simply including radiative cooling has on the survival time of satellites. Indeed, we find that the number of satellites that survive is dramatically reduced in the Reference run (from 20 objects in the Adiabatic run down to 6), which follows a similar trend to that described by Romano-Diaz et al (2010) when comparing identical simulations with baryons (and cooling) to their pure dark matter counterparts. This is explained by two effects: (i) it is more difficult to strip mass from satellites as their central density increases and (ii) the cuspieness (and central density) of the host halo is increased. It is interesting to note that effect (i) could in principle lead to the opposite effect, i.e. an increase in the

lifetime of the satellites as it makes them more concentrated and thus more resistant to tidal disruption, but we find that reduction in dynamical friction time scales due to mass increase dominates. Of course these conclusions could be altered if a substantial mass of baryons was ejected out of the satellites, to the point where the trend that we measure could even be reversed. Pontzen & Governato (2012) argue that such a reversal is plausible, though it would require a significantly more efficient feedback mechanism than the one we observe in section 4.4.1.

An example of satellite stripping is given in figure 4.17. In this figure, a halo in the Reference run and its twin are tracked as they orbit the Milky Way. In the Reference run, the satellite is destroyed by $z=0$, whereas in the Adiabatic run the same halo survives to $z=0$. Both runs have very similar first approaches - the satellite falls on its first orbit directly into the centre of the Milky Way. We confirm this by measuring the fraction of the velocity of the satellite that is radial - in both cases, the halos that survive in the Adiabatic Run but not in the Reference run have radial velocity fractions that approach 1 at perihelion, i.e. the satellite passes almost directly through the centre. By contrast, the halos that survive typically have less radially-aligned orbits, and thus do not pass through the Milky Way's centre, or if they do, they do it close to $z=0$. Once the halo has passed through the centre of the Milky Way in figure 4.17, it is heavily stripped in the Reference run, while its orbit decays. By comparison, the same halo in the Adiabatic run exhibits less mass loss and its orbit decays much more slowly (though the precise distance at aphelion is somewhat erratic). Thus we can confirm that the dense centre of the halo plays a crucial role in satellite survival. We cannot comment on the effect of the structure of the satellites themselves, as we do not run controlled experiments as do, for example, Peñarrubia et al. (2010b).

A further consideration is that much of the cold gas in the Milky Way halo's centre in the Reference run is in a disk structure (see, e.g., Kimm et al. (2011)). This has the effect of concentrating the matter in the centre into a plane, rather than a diffuse cloud. We do not explore in detail the difference between this and a radially symmetric distribution of mass, though it may have an important effect on the satellites passing close to the centre of the galaxy.

Our finding that gas cooling reduces the number of surviving satellite galaxies seems to run counter to Libeskind et al. (2010), who claim that gas cooling allows satellite halos to survive longer. However, Romano-Díaz et al. (2010, 2009); Schewtschenko & Macciò (2011) (also running Smoothed Particle Hy-

drodynamics (SPH) simulations like Libeskind et al. (2010)) reach the opposite conclusion (in agreement with ours), arguing that sub-halos survive a shorter time when gas cooling is included. It is worth pointing out that these studies all use tree-SPH methods as opposed to the AMR technique we use to model the self-gravitating baryon and DM fluid. Beyond the documented inability of the standard implementation of SPH to properly capture hydrodynamical instabilities (Agertz et al., 2007; Mitchell et al., 2009, e.g.), which is likely to cause an underestimate of the amount of gas stripped from satellite galaxies as they plough through the hot diffuse medium of their host halo, the main difference between the two numerical approaches lies in their estimate of the gravitational force. Indeed, whilst this force in Libeskind et al. (2010) is computed using a tree method which involves direct particle summation and a constant softening on small scales (Springel, 2005, e.g.), Teyssier (2002) solves the Poisson equation on the AMR grid, i.e. with an adaptive softening directly linked to the grid cell size. Such a difference naturally leads to a better force resolution of the tree method in low density regions. However, the situation is reversed in high density regions typical of virialised halos, where the size of an AMR cell can become significantly smaller than the constant softening length used by tree methods. In our case, we reach a spatial resolution of $\leq 50\text{pc}$ in the central region of the Reference run MW halo, which is significantly smaller than the values reported in the aforementioned SPH studies (several hundred parsecs deduced from fig 4 of ?, $500\text{ h}^{-1}\text{pc}$ quoted in Schewtschenko & Macciò (2011) , and $150\text{ h}^{-1}\text{pc}$ in Libeskind et al. (2010)). This, in conjunction with a similar (or better) DM particle mass resolution as that reported in these studies, makes us quite confident that the force resolution achieved in the Reference run should not lead to a premature destruction of sub-halos, but allow them to better survive tidal disruption. Since we do not perform a tree-SPH counterpart to our AMR simulation using identical initial conditions, it is difficult to comment more quantitatively on the differences in number of satellites destroyed by these different techniques. However, we refer the interested reader to Schewtschenko & Macciò (2011) for a discussion of the possible reasons behind the discrepancy between their results and the results of Libeskind et al. (2010).

To conclude, even if our sub-halos had exactly the same mass in the DM and Reference runs, a higher density of the host halo would cause their orbits to decay faster, thus decreasing their survival time. This effect is exacerbated as the satellites in the Reference run are also more concentrated and therefore retain more mass on their way to the centre of the host as this renders them less prone to tidal disruption. We

believe that the main caveat of our Reference run is that it does not include a feedback model. Indeed, Pontzen & Governato (2012) have recently argued that the injection of energy from supernovae in the centre of dwarf galaxies at $4 > z > 2$ can dramatically alter their halo density profiles. Whilst we are not able to confirm this effect with our own set of simulations, should it prove to be able to lower the central density of the Milky Way halo as well, we predict that more satellites will survive to $z = 0$. However, whether this number will match that measured in pure DM simulations or will still reflect a significant suppression of satellites is likely to depend on the details of the numerical implementation of the feedback processes.

A further discrepancy between simulation and observation is discussed in Boylan-Kolchin et al. (2011); Lovell et al. (2012); Di Cintio et al. (2011), who find that simulations using Λ CDM models, either pure dark matter or with baryon physics modelled using SPH, create satellites whose v_{\max} is higher than that of observed dwarf spheroidal satellites of the Milky Way. In fact, baryon physics produce a worse fit to the observations (Di Cintio et al., 2011). Lovell et al. (2012) propose a warm dark matter model to explain this discrepancy. In figure 4.18, we reproduce the calculated values of r_{\max} and v_{\max} found in Boylan-Kolchin et al. (2011) and overplot the values for halos containing stars in the Reference run at $z=0$. We find good agreement between our results and the observations. Figure 4.19 shows that our values for the half-light radius and masses also overlap with the observed quantities for the Milky Way and Andromeda dwarf spheroidal satellites. Note that in both cases, halos ‘a’ and ‘h’ have high maximum circular velocities; this is attributed to the fact that they have only recently been captured by the Milky Way and have not lost their gas to stripping. The most likely candidates for these objects are the LMC and SMC, which are not dwarf spheroidals and are hence not included in the observational data. We also find objects that have maximum circular velocities that are too low to fit on the region in $r_{\max}-v_{\max}$ calculated for observed satellites. We cannot comment on whether this is caused by the numerical limits of the Reference run or whether the data for observed satellites is limited to the most massive satellites and hence does not include these objects. We posit that one explanation for the ability of satellites to retain a high v_{\max} in the SPH simulations of Di Cintio et al. (2011) is the inefficiency of SPH simulations at destroying spheres of gas falling through a gaseous medium (Agertz et al., 2007). Without modification, SPH will hence allow satellites to retain their cool gas for longer and hence the mass in the centre will be higher in SPH simulations than in AMR simulations, raising v_{\max} . As in figure 4.7, we note that the mass of satellites is heavily dependent on the mass of the

host, and hence the results found here will depend on the mass of the host chosen and how well it fits to the mass of the Milky Way. Again, note that the Reference run does not contain supernova feedback and hence we cannot comment on the effect this has on the kinematics of our satellites at $z=0$.

4.5.2 Properties of Satellite Galaxies at $z = 0$

In this section we compare the properties of galaxies found at $z = 0$ in our simulations to observations of known satellite galaxies of the Milky Way and M31. To do this, we make use of the Reference run results measured at $z = 0$ directly rather than relying on the satellite tracking algorithm described in section 4.3.1. In this run, we find fifteen satellite galaxies using a direct galaxy identification method described in section 4.2.2. Four of these (labelled i , l , r and q on Fig. 4.19) do not have associated dark matter host halos identified by our halo detection algorithm. Since these are the least massive satellite galaxies, this has two likely origins: (i) either the number of dark matter particles constituting the host sub-halo falls below the minimum number for detection as the outer regions are tidally stripped by the host, in which case the lack of dark matter host is simply the result of lack of mass resolution; (ii) or as pointed out by e.g. Knebe et al. (2011) since our halo finder does not consider particle velocities, sub-halos close to the centre of the host (r and q for instance in our case) can fail to be detected. In any case, including a DM sub-halo host with a mass just below our mass resolution threshold ($2.2 \times 10^6 M_{\odot}$) does not affect the results we present here.

Figure 4.19 shows a plot comparing the half stellar mass against half-light radius of our simulated satellite galaxies (grey squares), to the satellite galaxies of the Milky Way (from Wolf et al. (2010)) (white circles) and M31 (from Collins et al. (2011, 2010); Ferguson et al. (2000); Irwin et al. (2008); Irwin & Collins (2011, priv. comm.); Kalirai et al. (2010); Letarte et al. (2009); Martin et al. (2009); McConnachie et al. (2008); Morrison et al. (2003); Pustilnik et al. (2008); Richardson et al. (2011)) (white squares). The population of simulated galaxies lie in the same region as the observed satellites except for the two galaxies labelled a and h on the figure, which lie well above the observational data. These are the most massive galaxies in our sample, which have recently been captured by the Milky Way host halo and thus have not experienced significant stripping by $z = 0$. We note that beyond their larger stellar half mass, which can be somewhat overlooked since the data from Wolf et al. (2010) do not contain the LMC or SMC, these simulated galaxies seem too compact when compared to the rather well defined observational relation

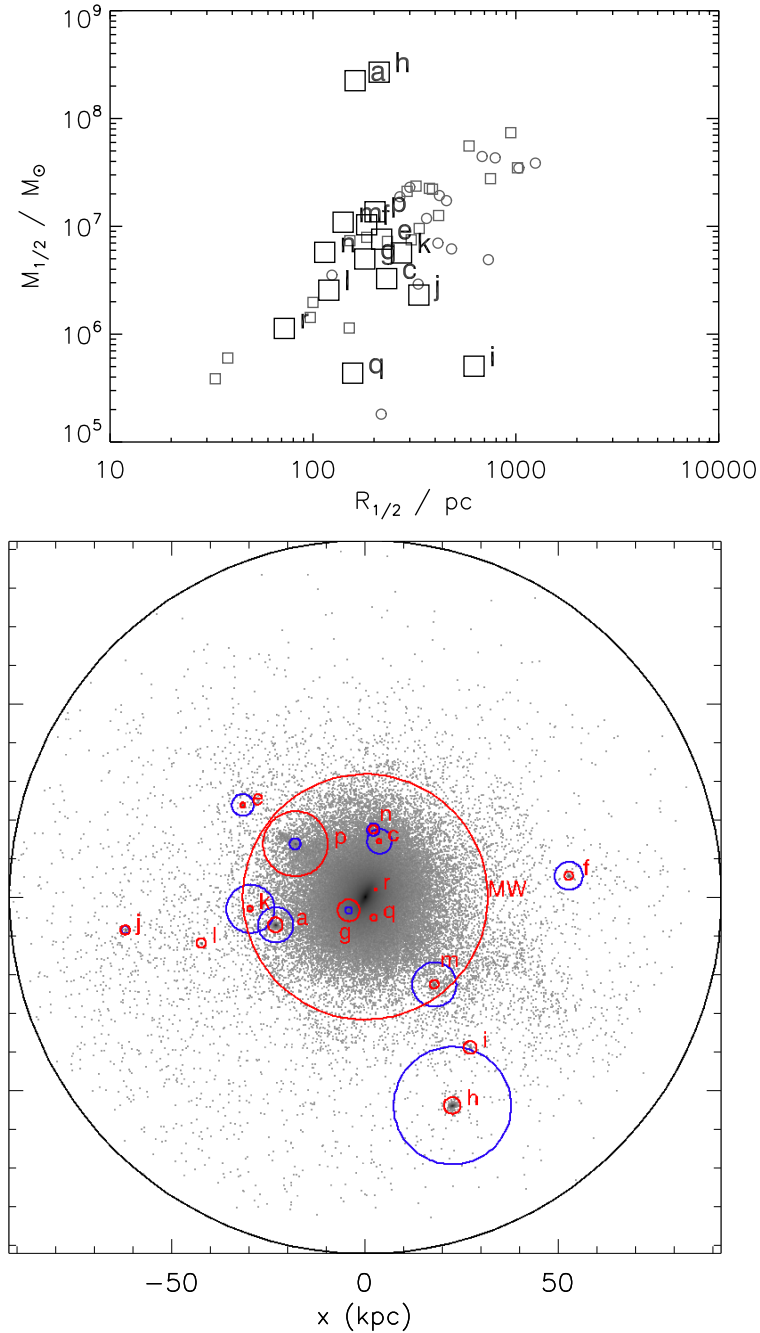


Figure 4.19: Properties of satellite galaxies surviving down to $z = 0$. The top panel shows half stellar mass against half-stellar-mass radius for satellite galaxies in the Reference run at $z = 0$ (large black squares), compared to half-light radii calculated for satellites of the Milky Way and M31 (small grey squares and small grey circles respectively). Data for the satellites of the Milky Way are taken from Wolf et al. (2010). Data for the satellites of M31 are derived from Collins et al. (2011, 2010); Ferguson et al. (2000); Irwin et al. (2008); Irwin & Collins (2011, priv. comm.); Kalirai et al. (2010); Letarte et al. (2009); Martin et al. (2009); McConnachie et al. (2008); Morrison et al. (2003); Pustilnik et al. (2008); Richardson et al. (2011). The bottom panel shows the spatial location of the satellite galaxies and their host DM sub-halos inside the Milky Way virial radius (red circles indicate galaxy radii, while blue circles represent sub-halo radii). The black circle shows the Milky Way virial radius. The background image is a cubic projection of the stellar density field. The two satellite galaxies (labelled *a* and *h*) which lie above the observed data in the top panel are galaxies that have been little affected by stripping due to their late capture by the Milky Way host. Note that the Milky Way data does not include the Magellanic clouds or ultrafaint satellites.


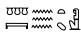


Figure 4.20: Star formation histories of satellite galaxies in the Reference run (black squares) compared to values deduced from observations of Milky Way satellites (Orban et al. (2008) and Wolf et al. (2010): grey circles). Each panel shows the fraction of stellar mass formed between $z = 0$ and the lookback time indicated on ordinate axis against the half stellar mass of the galaxy at $z = 0$. Stellar fractions and masses in each plot match well with the values estimated from observations, suggesting that the star formation histories of satellite galaxies in the Reference run do not wildly differ from those of observed Milky Way satellites.

linking satellites size and mass. While resolution undoubtedly plays an important role in getting accurate estimates of the sizes of simulated objects, it is nevertheless striking that satellite galaxies which lie on top of the observed results have all experienced tidal stripping to an advanced degree (see e.g. Figs 4.5 and 4.6). In other words, in our Reference run (which, once again, does not include feedback) every satellite galaxy which has experienced a large amount of tidal stripping falls remarkably close to the observed stellar mass - size relation.

Moreover, when we compare star formation histories of satellite galaxies in the Reference run with those derived from the analysis of colour-magnitude diagrams (CMDs) of Milky Way satellite galaxies observed by Orban et al. (2008) we find reasonable agreement. This is demonstrated in Fig. 4.20, where we plot the fraction of stars formed after lookback times of 1Gyr, 2Gyr, 5Gyr and 10Gyr both in simulated (grey squares) and observed (white circles) galaxies. It is somewhat reassuring that the star formation histories of our simulated satellite galaxies are broadly correct since we have previously shown (Fig 4.2) that they are fairly independent of resolution. Perhaps more importantly, this agreement also suggests that feedback, whatever its form and origin, cannot drastically alter the star formation histories of these galaxies: models where significant feedback at $z > 1$ completely quenches late star formation are clearly ruled out by the observational data. This makes it all the more challenging for stellar feedback to soften cusps of dark matter sub-halos and certainly favours a rapid, irreversible mechanism, very localised in time such as the one suggested by Pontzen & Governato (2012).

4.6 Discussion and Conclusions

The work discussed in this chapter has made use of the  (NUT) suite of high (a few tens of parsec) to ultra-high (sub-parsec) resolution cosmological hydrodynamic re-simulations to investigate the effect of baryonic physics on the evolution of a ‘Milky Way’ and its satellite galaxies. Whilst various other authors have simulated satellites of Milky Way-like galaxies down to $z = 0$, ours are the first to reach sub-parsec resolution down to the end of the epoch of reionisation. The motivation for this was to analyse in detail the evolution of satellite galaxies around the epoch of reionisation, which has been posited as a mechanism for suppressing star and galaxy formation in dwarf halos and hence for shaping the population of satellite galaxies that we observe around the Milky Way and M31. To the best of our knowledge, we are also the

first authors to use an adaptive mesh refinement (AMR) technique to study the evolution of these satellite galaxies down to $z = 0$ (other studies thus far either used smoothed particle hydrodynamics (SPH), or ended their simulations at higher redshifts).

In agreement with e.g. Guo et al. (2010); Ricotti & Gnedin (2005); Wadepuhl & Springel (2010b), we find that reionisation appears not to efficiently stop star, or even galaxy formation. Instead, we find that satellite galaxy formation continues down to at least $z = 4.8$ in our lowest (~ 50 pc) resolution simulation. The number of luminous satellite galaxies formed before reionisation ($z = 8.5$ in our case) is found to be far lower than the number of observed Milky Way satellites. These results are consistent with e.g. Okamoto & Frenk (2009); Hoefl et al. (2006): like these authors, we find that, down to the end of the reionisation era, there exists a threshold in v_{max} of about 10 km/s below which halos remain dark, never forming stars. This threshold persists at later redshifts, i.e. well after reionisation has ended. This is probably due to the fact that efficient atomic gas cooling increases halo central densities and hence v_{max} , separating halos that can cool gas and form stars from those that cannot. There are, however, at least two major limitations in our work that prevent us from commenting further. First, we do not run a simulation with self-consistent star formation and ionisation, and hence we cannot quantify the precise effect of UV photoionisation on star formation in galaxies already forming stars. Secondly, our model of reionisation consists of a uniform background which neglects the effect of gas self-shielding from external photoionisation sources. Future studies which include ionising photon radiative transfer and hence self-consistent re-ionisation, should allow us to determine whether proximity effects and/or self-shielding significantly alter our conclusions, but this does not seem likely.

The effect of supernova feedback on the gas and stellar mass of high redshift (both pre- and post-reionisation) galaxies seems to be quite stochastic. Indeed, our results exhibit a large scatter, even though one might argue we detect a slight systematic trend of star formation being enhanced by feedback as galaxy mass increases. A key feature of our model of supernova feedback is that it consists of star particles (with a standard Salpeter IMF) injecting mass, energy, momentum and metals into the surrounding gas according to a Sedov blast wave solution deposited onto the grid 10Myr after they have formed. Together with our sub-parsec resolution, this means that we should be able to track the effect of supernovae explosions fairly realistically on scales typical of small molecular clouds. Outflows and heating from these supernovae re-

duce the amount of gas available for star formation, but blast wave compression and an excess of radiative cooling due to the injection of metals into the ISM can potentially increase it. Our results show that, on average, adding supernovae feedback does reduce the gas and stellar mass of galaxies in halos below $10^9 M_\odot$ (negative feedback) but increases it for galaxies hosted by halos above $10^9 M_\odot$, where the deeper potential and extra metal injection negates the impact of outflows (positive feedback). In any case, the stellar masses of individual galaxies are only changed by maximum 10-20 percent either way by supernova feedback, and one has to invoke a much larger energetic input (for example a top-heavy IMF, or a large fraction of hypernovae) and/or one that is impervious to radiative losses (e.g. some kind of 'turbulent' energy) to overturn the situation in favour of negative feedback.

When we use lower resolution (50 pc) runs to track the descendants of galaxies in the high resolution (0.5 pc) Cooling and Feedback runs to $z = 0$, we find that very few of the galaxies which are captured by the Milky Way progenitor halo survive to $z = 0$. In fact, independent of the run used to perform the tracking (PureDM, Adiabatic or Reference), no satellite galaxy captured by the Milky Way progenitor before $z = 3$ survives as Milky Way satellite at $z = 0$. However, the main difference is that *all* galaxies (sub-halos) captured *after* $z = 3$ in the Pure DM or Adiabatic runs survive, while only a small fraction of these does so in the Reference run. This is caused by the much higher central density of the Milky Way halo (and sub-halos) in the Reference run (the only one to host a MW galaxy), which significantly shortens the dynamical friction timescales for satellites to spiral to the centre of the halo and experience disruption. We note that this result is still a matter of debate in the recent literature, with Libeskind et al. (2010) and Romano-Díaz et al. (2010, 2009); Schewtschenko & Macciò (2011), all using a very similar tree-SPH technique, reaching different conclusions. Whereas the former group of authors seems to disagree with us, the latter report a higher destruction rate of satellites when baryons and cooling are included in their simulations, consistent with our finding. We believe that the work presented here, which makes use of a completely different simulation technique and improves on the resolution of these previous studies, sheds a useful light on the issue. In particular, we emphasize that when we compare properties of the remaining satellite galaxies at $z = 0$ in the Reference run to their observational equivalents for the Milky Way or M31, we find that the star formation histories, stellar masses and half-light radii are in reasonable agreement. Only the more massive, recently captured satellites are found to be too compact when compared to Milky Way and M31 dwarf

spheroidal satellites. This puts rather tight constraints on the feedback mechanisms (amount of energy, duration, timing) required to soften the cusps of dark matter halos as they cannot have a major impact on these properties.

Finally, beyond the importance of the role played by baryonic physics in determining the number of satellites of MW class halos, our study also highlights the extreme sensitivity of this number to the circular velocity of the host halo. Whilst all simulated CDM halos in the literature match the shape of the $N_{sat} \propto V_{max}^3$ relation from Reed et al. (2005), not all of them agree on the constant of normalisation. Indeed, our Pure DM run agrees with the normalisation given by Reed et al. (2005) at the 10-20 % level whereas the Aquarius (Springel et al., 2008) simulations overshoot it by a factor ~ 3 , and by about 30 % more than the Via Lactea II (Diemand et al., 2008) simulation. Curiously enough, the cause of this significant discrepancy is still unresolved, even though the debate has received a lot of attention lately in papers such as Vera-Ciro et al. (2011); Boylan-Kolchin et al. (2012); Wang et al. (2012). Irrespective of this disagreement, our results, especially when baryonic physics is taken into account, argue in favour of a rather light ($5 \times 10^{11} M_{\odot} < M_{vir} < 10^{12} M_{\odot}$) MW halo. Note that if the likelihood for a satellite galaxy to survive down to $z = 0$ is suppressed by even half as much as we find in this work when feedback is properly incorporated, models relying on N-body dark matter-only simulations, such as semi-analytic models (SAMs) should be revised accordingly.

One possible weakness of this work is that the model for supernova feedback we implement, whilst possessing unprecedented resolution, is very simplistic, depositing a canonical quantity of energy and mass in a Sedov blastwave pattern onto the grid at a fixed time of 10Myr after the time of star formation. In the next chapter I introduce a study aimed to rectify this, by running a series of increasingly sophisticated 3D simulations of single supernovae and following the evolution of their remnants through an idealised model ISM.

Chapter 5

Single Supernovae and the ISM

In the previous chapters I discussed implementing supernova feedback in hydrodynamic galaxy formation simulations, and presented work studying its effect on the formation of the progenitors of the satellite galaxies of Milky Way-like galaxies. However, as I stated in previous chapters, the precise model used for supernova feedback and the resolution employed allows for hugely different results in the physical properties of the ISM, and subsequently the star formation rate (e.g. Slyz et al. (2005); Scannapieco et al. (2011)). It is thus important to understand the evolution of a single supernova in detail in order to determine the global effect of millions or even billions of supernova events on a single galaxy.

In the following chapter I discuss a project to simulate the evolution of a $15M_{\odot}$ star in a small patch of the ISM, including stellar winds prior to the supernova event and the supernova blastwave propagating into the wind-driven bubble once the star reaches the end of its lifetime. The aim of this study is to motivate future models for supernova feedback in galaxy-scale simulations.

5.1 Introduction

The effect of supernova feedback on the ISM and on star formation is complex and poorly understood. In particular, numerical simulations of galaxy formation are normally not able to resolve the hot and cold phase of the ISM, and hence cannot implicitly include the impact of supernova blastwaves on molecular clouds (Slyz et al., 2005). One approach is that of Murante et al. (2010), who produce a model for numerical simulations whose fluid element contains the hot and cold phases of the ISM at the same time, and introduces




Figure 5.1: Schematic of the pre-supernova evolution of the circumstellar medium (CSM), motivated by figure 1 in Weaver et al. (1977). In order of distance from the star, the phases of the evolution of the CSM proceed as follows: (i) The stellar wind escapes the star as a kinetic wind, but very quickly thermalises into (ii) the shocked stellar wind. This drives a shell of (iii) shocked photoionised gas, which is the interstellar medium as photoionised by the UV emissions from the star. A similar pattern is seen in the (iv) photoionised gas, where the expanding photoionised bubble terminates with a shell of (v) shocked interstellar gas, which in turn expands into the unshocked (vi) interstellar gas. This structure is revisited in sections 5.4.2 and 5.4.3.

analytic expressions to describe the transfer of matter between phases. However, this model requires parameter tuning in order to match observations. Meanwhile, Scannapieco et al. (2011) demonstrate that there is currently no convergence in the results of supernova feedback as implemented by grid and SPH codes. In order to be confident that our simulations accurately describe the physics of galaxy formation, it is desirable to obtain physically-motivated values for the parameters of any sub-grid models that I implement.

In order to create detailed models for supernova feedback on a parsec scale or lower, it is first necessary to understand the population of supernovae that exist in the universe and, to a lesser extent, the progenitor stars of these events. Complicating this process is the relatively limited sample of nearby stars that are identified before they go supernova. Minkowski (1940) created the label “Type I” and “Type II” to refer to supernovae without and with (respectively) hydrogen Balmer-lines in their light curves. Since then these classifications have been refined to create a zoo of supernovae with different spectral features in their light curves. For example, Wheeler & Harkness (1990) note that supernovae without hydrogen lines but with strong silicon features are fairly prevalent, and hence dubbed them Type Ia, which are believed to be the result of thermonuclear explosion of white dwarf stars. By contrast, Type II supernovae are largely thought to be the result of core-collapse, i.e. the collapse of the core of a massive star into a neutron star or black hole releasing a large quantity of gravitational energy into the rest of the star resulting in a supernova explosion (Woosley & Janka, 2005). Work by, e.g., Heger et al. (2003) and Nomoto et al. (2003) has attempted to catalogue the kinetic energy, chemistry and compact remnant mass and type for various initial single stellar masses; Nomoto et al. (2003) discuss highly energetic events called “hypernovae”, which Gal-Yam et al. (2006) attribute to “collapsar” events, which refers to the collapse of a fast-rotating massive star. However, there still exist large uncertainties on measurements and models of supernova properties and progenitors.

One aspect that provides further channels for supernova events is binary evolution. Observations by Wampler et al. (1990); Burrows et al. (1995) lead to the discovery of a triple-ring system around the source of SN1987A. White & Malin (1987); Podsiadlowski et al. (1991); Lloyd et al. (1995); Morris & Podsiadlowski (2009) concluded that the progenitor was a blue supergiant formed by the merger of two binary stars. This merger caused the blue supergiant to rotate quickly, causing the wind nebula to adopt an axi-symmetric form. The triple ring nebula was shaped from the disk of ejecta in the plane of rotation, and the two hemispherical lobes by ejecta along the axis of rotation of the disk. These were subsequently observed after

being illuminated by the supernova event. Type Ia are also thought to be caused by binary interaction. In one model, a white dwarf accretes matter from a binary companion; in another, two white dwarfs merge to create a supernova event (Hillebrandt & Niemeyer, 2000). Binary interaction is hence an important process in the creation of supernovae. In this study I focus on a single $15M_{\odot}$ star, but future work will require investigation of binary progenitors for supernova events in order to provide a more complete picture of the impact of supernovae on the ISM and hence on galaxy formation.

The internal mechanisms governing the conversion of gravitational potential energy from a core collapse event in a massive star to a supernova blastwave is still subject to active research by theorists (e.g. Müller et al. (2011)). Nonetheless, it is apparent that a portion of this energy is coupled to the gas in the stellar envelope. Once this occurs, the blastwave propagates through the stellar envelope and into the circumstellar medium (CSM). As the ejecta encounters the unshocked medium, reverse shocks are set up as the swept-up material lags behind the shock front. Rayleigh-Taylor (RT) instabilities then break up the ejected envelope into a series of clumps (Chevalier, 1977; Chevalier & Klein, 1978; Matzner & McKee, 1999).

Once the supernova has escaped the stellar envelope, it evolves into a strong shock that can be modelled simply as a Taylor-Sedov blastwave (Shu, 1992):

$$r_{\text{shock}}(t) = \xi_0 (Et^2/\rho_1)^{1/5} \quad (5.1)$$

where r_{shock} is the radius of the shock, E is the energy of the blast, t is the time since the supernova event and ρ_1 is the density of the ambient medium into which the shock passes, assuming a uniform density. This equation can be derived through simple dimensional analysis; using a more rigorous treatment, for a typical adiabatic index γ of $5/3$, $\xi_0 = 1.17$. Note that in the Taylor-Sedov phase, the shock is energy-conserving, i.e. it exhibits negligible radiative losses. As the blastwave expands, the internal pressure drops, and when the cooling time becomes comparable to the flow time (i.e. $r_{\text{shock}}/\dot{r}_{\text{shock}}$), radiative losses become significant. In this “shell-forming” phase, the supernova gathers a dense shell of swept-up gas just behind the shock front, which grows with time. Eventually this shell grows sufficiently dense that the blastwave enters the momentum-conserving “snow plough” phase, in which the dense shell sweeps up matter in the surrounding medium. In this phase, the expansion of the supernova remnant is composed of a cold, dense shell with a hot, low-density interior.

The evolution of a star prior to supernova can also have an impact on the ISM, and in particular on star-forming regions. A simple schematic of the various phases of the CSM around a star is given in figure 5.1. In many ways, the phases are similar to that described above for the supernova remnant. The supernova wind is initially kinetic, though it quickly thermalises due to shock interaction with the surrounding medium Weaver et al. (1977). There is a dense shell around the thermalised wind where the shocked external medium is swept up. Beyond that, a massive star creates a photoionised bubble where UV photons from the star heat the interstellar gas. This too can set up a shell of shocked interstellar medium. While the shocked wind and photoionised medium are typically low density and dominated by thermal energy, the shells are subject to rapid cooling due to their increased density from swept-up material in the external medium, leading to a momentum-driven phase as in the supernova remnant. See sections 5.4.2 and 5.4.3 for simulations of these various phases.

Strömberg (1939); Oort & Spitzer (1955); Kahn (1954) proposed that O stars can affect the evolution of star-forming clouds through radiation and kinetic feedback. As well as regulating star formation, these feedback from these stars leads to the formation of photoionised bubbles and disrupts star-forming clouds, creating channels through which further stellar winds and supernova blastwaves can escape dense clouds. The effect of this stellar feedback, e.g. in the HII region 30 Doradus in the LMC is captured in observations by Chu & Kennicutt (1994); Redman et al. (2003). Much work has been done in recent years on the effect of pre-supernova stellar evolution on the CSM, particularly inside star-forming clouds, as well as work on smaller stars and planetary nebula (PN) evolution (for reviews, see Hensler (2008); García-Segura (2009)). Dwarkadas & Owocki (2002); Chita et al. (2008) note that stellar rotation appears to alter the formation of nebulae around stars, creating bipolar nebulae as discussed earlier in this section. Garcia-Segura et al. (1999)

find that a star must be $> 1.3M_{\odot}$ to retain enough rotation to produce this bipolar shape. This non-sphericity can then set up shear forces in the CSM, creating Kelvin-Helmholtz instabilities (Dwarkadas & Balick, 1998). Somewhat conversely, Toalá & Arthur (2011) find that fast rotation can act to damp Vishniac instabilities forming in the wind bubble in massive ($40-60M_{\odot}$) stars. van Marle et al. (2008) also find that fast rotation in massive stars causes the density profile to deviate from the inverse square law found for non-rotating stars. Another effect to consider is the interaction between wind and photoionisation bubbles;

colliding wind bubbles can also generate complex structures in star-forming clouds (e.g. van Marle et al. (2011)). Further, as discussed in section 2.6 of chapter 2, physical instabilities do occur in stellar wind bubbles as well as supernova remnants, which act to modify the expansion of the shocks through the external medium.

The effect on the supernova blastwaves from the pre-supernova evolution of stars is varied. For example, massive stars undergoing core collapse will have ejected part of their envelope as stellar winds, altering the shape of the CSM. Dwarkadas (2011) note that this wind produces a power-law density field into which the supernova blastwave expands (Chevalier, 1982). For some stars, such as the progenitor of SN1987A, the pre-supernova ejecta is so dense that the supernova cannot expand as a Sedov blastwave and moves straight to the snowplough phase (Tenorio-Tagle et al., 1990); Tenorio-Tagle et al. (1991) find that this is the case for shocks initially travelling at $\lesssim 300\text{km/s}$ (i.e. in media of $n_{\text{H}} > 10^5\text{atoms/cm}^3$). This is in contrast to Type Ia supernovae, whose progenitors are typically low-mass stars that do not produce strong winds. The result is that the expansion of the supernova shock can be altered. Dwarkadas (2007) further notes that the interaction between the stellar wind bubble and subsequent supernova can explain observed Wolf-Rayet star termination shocks. The movement of stars as they emit wind can create tubes and bow-shocks in the ISM (e.g. Mohamed et al. (2011)). Rozyczka et al. (1993) attribute barrel-shaped supernovae to the creation of these tube-like structures, since the low-density tube will create a channel for the supernova shock to propagate along (Tenorio-Tagle et al., 1990). Equally, in very low-density environments such as galactic bulges or elliptical galaxies Tang & Wang (2005) find that cooling can be neglected in a low-density medium as the blastwave never sweeps up enough matter to begin shell formation. This allows a highly efficient transfer of energy from the supernova to the gas. The environment into which a supernova expands thus depends heavily on the properties of the CSM and ISM surrounding the progenitor star.

In our research, I focus on a single $15M_{\odot}$ star. I do this as stars of this mass are more common than stars of higher masses studied by the majority of the previously-mentioned authors according to the IMFs derived by, e.g. Salpeter (1955); Kroupa & Weidner (2003). The cost of this is that the lifetime of the star is much longer, and hence I must simulate our star for $\sim 14\text{Myr}$ as opposed to $\sim 5\text{Myr}$ for a $35M_{\odot}$ star. While the final study will cover a range of environments, I begin with a uniform medium in order to compare our results with those in the literature (Thornton et al., 1998). We also study a supernova exploding

in an environment typically found near molecular clouds, where the high density means the blast travels slower (Sedov, 1946), and hence the computational domain can be smaller, allowing a higher resolution; authors such as Garcia-Segura et al. (1996) employ a similar strategy. Further, unlike Dwarkadas (2007) or Frascchetti et al. (2010), I use a static computational domain, which allows more complex environments in future studies. The drawback of using a moving shock in a static medium, as I do, is that I simulate the early evolution of the shock with much lower resolution than these authors (in particular, Frascchetti et al. (2010) does not simulate the stellar wind prior to the supernova event and as such only models the first few parsecs of the supernova blast). Frascchetti et al. (2010) notes that another problem with using a moving shock in a static medium is that it is more susceptible to errors in capturing the shock due to Galilean invariance, which is an effect that causes structures to diffuse slightly as they move across the grid. Numerical artefacts can also be formed by imposing a spherical geometry on a Cartesian grid. Another drawback is the loss of early formation of instabilities that otherwise comes from resolving the early evolution of the wind or supernova shock. However, we do not implement a co-moving grid as Dwarkadas (2007) or Frascchetti et al. (2010) do. The main reason for this is that the setup, while effective for simulations of spherically symmetric systems in a uniform medium, cannot model multiple supernovae or non-uniform media, which we intend to do in future studies.

We hence have a picture of massive stars dramatically altering their environment prior to going supernovae; very massive stars ($\geq 20\text{-}40M_{\odot}$) produce photoionised and kinetic-wind-driven bubbles that can deviate wildly from spherical symmetry, both due to thermal and gravitational instabilities and external variations in the environment. These primarily affect star-forming regions, since these stars have shorter lifetimes and hence drift less from their place of birth. A large amount of research has been done on this phenomenon. However, controlled experiments of supernova events in various media are less well studied in recent times, particularly in 3D. In this study I aim to reproduce the results of Thornton et al. (1998), who perform 1D simulations of a single supernova, in full 3D. The goal is to capture the dynamics that cannot be captured in 1D, such as non-spherical supernova propagation and dynamic instabilities predicted but not followed by Gull (1973). The effect of turbulent instabilities is to remove radial kinetic energy and to convert it into thermal energy via the turbulence energy cascade; in our simulations, the resolution limit will dictate at what scale this happens.




Figure 5.2: Values for cumulative mass loss from a $15M_{\odot}$ star of metallicities Z_{\odot} and $0.1Z_{\odot}$, where Z_{\odot} is set to 0.02. The ionisation energy is the total energy in photons emitted from the star at frequencies that ionise HI, HeI and HeII. See section 5.2 for a full description of the values used.

I then introduce simulations including the effects of pre-supernova processes such as photoionisation of the CSM and stellar winds on the supernova remnant’s evolution. The purpose of this is to motivate future models of supernova feedback in galaxy formation simulations in a variety of settings. I begin by describing the model for our $15M_{\odot}$ star, and the numerical simulations employed. I then discuss methods for controlling numerical instabilities in the shock, and determining which instabilities are created by physical processes and which are solely the result of numerical artefacts. Finally I present the results of our simulations and discuss the possibilities for future work.




Figure 5.3: Values for mass loss rates from a $15M_{\odot}$ star of metallicities Z_{\odot} and $0.1Z_{\odot}$, where Z_{\odot} is set to 0.02. Note that this is derived from the cumulative mass loss tables plotted in figure 5.2.

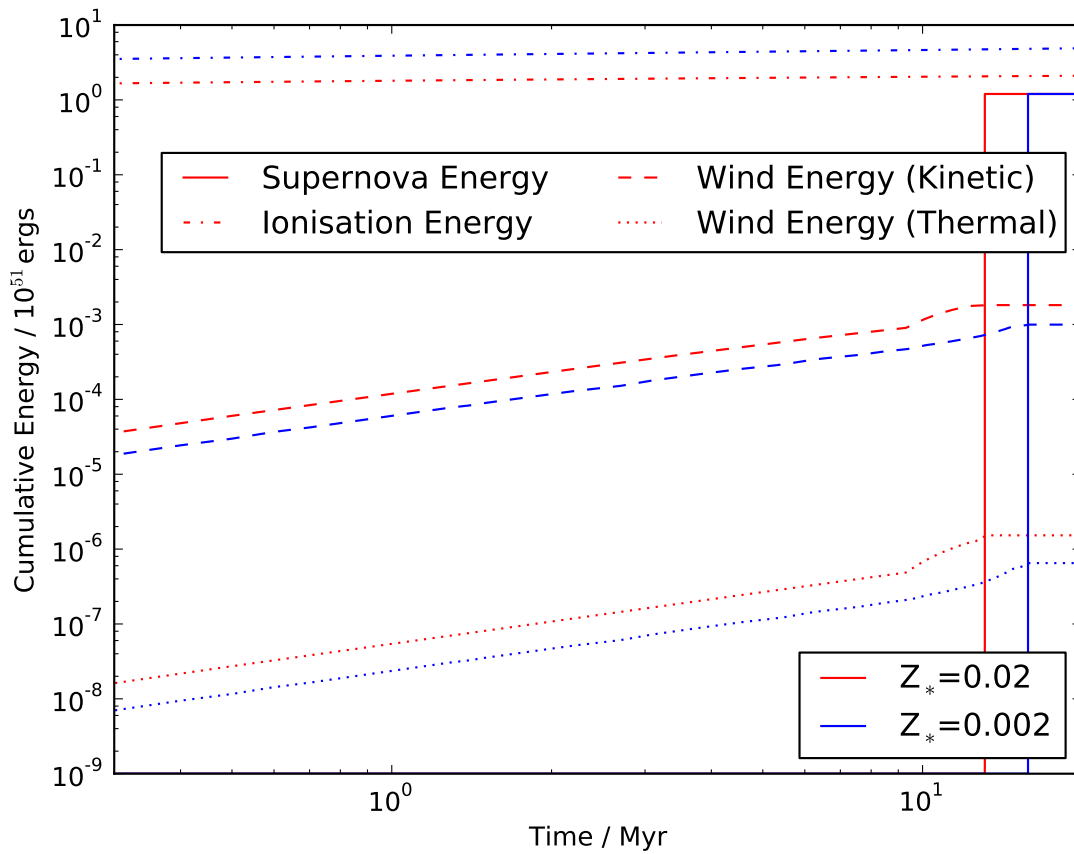


Figure 5.4: Values for cumulative energy output from a $15M_{\odot}$ star of metallicities Z_{\odot} and $0.1Z_{\odot}$, where Z_{\odot} is set to 0.02. The ionisation energy is the total energy in photons emitted from the star at frequencies that ionise HI, HeI and HeII.

5.2 Model of a Massive Star

In this section I discuss the construction of a model for the release of mass, momentum, energy and metal yield from a massive star, as well as the spectrum of flux from the star to use in the RHD module. The basic operation of the module proceeds as follows. I first extract a table of values from a model of stellar evolution for the parameters of interest. I then create a module in RAMSES that reads these tables and interpolates these values in stellar mass, age and metallicity. Finally, RAMSES queries this module with the current age, mass and metallicity of the star being modelled and receives values for the mass, momentum, energy, metal yield and intensity of the photon packets under consideration to deposit onto the grid.

I generate the single star wind models from the Padova stellar evolution models (Marigo et al., 2008). The wind velocity is taken as being the escape velocity from the surface of the star. This is reasonable for a star of this mass; Kudritzki & Puls (2000) find that for more massive stars, however, the wind velocity can reach three times the velocity of the star. The temperature of the wind is taken as the surface temperature of the star. The temperature of the wind is taken as being the surface temperature of the star. While Runacres & Owocki (2005) argue that the temperature decreases rapidly once it leaves the star, the kinetic energy of the wind dominates by roughly three orders of magnitude and thus the precise temperature value of the wind is unimportant. The mass loss is taken as being the cumulative difference between the initial mass of the star and the current mass of the star; the cumulative mass loss is used as it is more stable and useful for sampling between timesteps than the mass loss rate as a function of time, which can vary considerably as a function of time. I produce 3D tables (the dimensions being initial mass, time since star formation/stellar age and metallicity). I plot the values for the $15M_{\odot}$ star in figures 5.2 to 5.4. I also allow for momentum from the wind to be converted to thermal energy or vice-versa, since I do not resolve the surface of the star. Following Weaver et al. (1977), for the purposes of section 5.3 I consider all energy from the wind to be thermal. For runs including radiative transfer with RAMSES-RT (Rosdahl & Blaizot, (in prep.)), I produce a set of spectra for a $15M_{\odot}$ star using the Starburst99 web-based software and data package (Leitherer et al., 1999), which is used in the place of the Bruzual & Charlot (2003) model for the SEDs of stellar populations. Initial work has also been done to model a small stellar population emitting photons and winds; however, this is an extra level of complication that must be handled carefully, and is not explored in this thesis.

Once the star has reached the end of its lifetime, it goes supernova. I use Smartt et al. (2009) and Kovetz

et al. (2009) to determine values for energy, lifetime of the star and compact remnant mass, which I set to 1.2×10^{51} ergs, 14.125 Myr and $1.5M_{\odot}$. I also obtain a metal yield as computed from our Starburst99 data of 0.4. In the following section I discuss the application of these models to idealised numerical simulations of supernova remnants and stellar wind bubbles.

5.3 Numerical Simulations

Once the model for the star is complete, I am able to feed the hydrodynamic variables into RAMSES and allow it to evolve the stellar wind and supernova blast with or without atomic line cooling and radiative transfer. In this section I discuss the numerical model used to imprint the stellar wind and supernova blast onto the grid while minimising grid artefacts and ensuring that physical instabilities are allowed to evolve while numerical instabilities are suppressed as far as is possible in a Cartesian grid code.

The first consideration is the deposition of the stellar wind onto the grid. The volume around the star is first refined to the highest level allowed by the simulation in order to capture the initial evolution of the winds. I then cast 10^6 rays into a sphere of radius 20 cell lengths at the highest level of refinement around the position of the star, following Lemaster et al. (2007). I then create a cubic array that corresponds to the grid cells inside this radius, which stores the proportion of rays that fall inside each cell. This creates a Monte-Carlo-generated field of weights over which the wind variables are inserted onto the grid. The process naturally provides an inverse square fall-off in density as the wind propagates through its immediate surroundings. I deposit mass, energy and metals into the grid in this way; while the wind is initially kinetic as it leaves the star, it quickly thermalises the surrounding gas and hence at the resolution I am capturing the supernova energy is almost entirely thermal (Weaver et al., 1977). However, when running tests with various ratios of injected kinetic and thermal wind energy, I find very little difference in the resulting wind bubble as the partition of energy quickly reaches a natural equilibrium.

I present two sets of initial conditions in this section. In all simulations I reach a maximum 0.1 pc resolution; the grid is refined on density, pressure and velocity such that if the difference between the values of any variable either side of a cell boundary is more than 20%, the grid is refined up to the maximum level. I run each set of initial conditions with three separate realisations of the stellar model - one with just the supernova and no radiative cooling, one with just the supernova and radiative metal-line cooling, and one

including the stellar wind model as discussed above and radiative metal-line cooling. I unimaginatively dub these NoCool, SNCool and WindCool respectively.

The first set of initial conditions I name the “Cloud” simulation, which is motivated by the environmental conditions of a molecular cloud. I use a uniform medium with hydrogen number density n_{H} of 100 atoms/cm³, a temperature T of 20K and a metallicity Z of $0.3Z_{\odot}$, stationary with respect to the star. Note that in this study I do not consider the effects of turbulence, motion with respect to the star or feedback from other stars on the medium, which I will explore in future work. I situate the star in the centre of the grid, since the blast wave propagates more slowly in a dense medium and hence I can achieve the required spatial resolution even when simulating the whole volume rather than merely a segment with one corner on the star. This is also beneficial when implementing physical processes such as motion with respect to the external medium, a non-uniform medium and self-gravity, which will be included in future work.

The second, named “T1998” is motivated by the initial conditions featured in Thornton et al. (1998). I use a uniform medium with $n_{\text{H}} = 0.1$ atoms/cm³, $T = 1000\text{K}$ and $Z = 0.3Z_{\odot}$, stationary with respect to the star. Note that Thornton et al. (1998) focus on their simulation at Z_{\odot} , though they run a suite of models at various metallicities and densities. Further, Thornton et al. (1998) use a slightly lower supernova energy than us (10^{51} ergs as opposed to 1.2×10^{51} in our work; see section 5.2). The supernova is centred on the corner of the volume, such that I only simulate 1/8 of the shock (following, e.g. Dwarkadas (2007) in 2D and Fraschetti et al. (2010) in 3D). I do this as it enables my simulations to achieve twice the spatial resolution in each dimension for the same computational cost as simulating the full sphere. I simulate a cubic volume of length 200pc with a maximum resolution of 0.1pc, corresponding to 11 levels of refinement. This is done by creating a uniform grid of length 2^7 cells and then allowing the grid to adaptively refine for a further 4 levels. Each cell refines if velocity, pressure or density gradients rise above 0.2 and the maximum resolution is not already reached.

The computing time requirements for each simulation varied with the physics and initial conditions. Adding radiative transfer drastically increased the requirements, both because extra fluid elements were added to account for the photons, and because the ionisation front increased the volume of shocked gas considerably. The lower density of the T1998 simulations also led to a larger shock radius, and hence more computational time requirements. The time required for each simulation ranged from 10 000 to 100 000

CPU hours.

5.4 Results

5.4.1 Wind-Free Supernovae

In this section I present preliminary results from our two sets of initial conditions, T1998 and Cloud, for the simulation set-ups NoCool, SNCool, i.e. simulations of supernovae with and without cooling that do not consider pre-supernova winds. I discuss the results of adding wind (simulation set-up WindCool) in a subsequent section. I also present preliminary results from simulations with radiative transfer.

Figures 5.5 and 5.8 show radial distributions of various physical properties at different snapshots in time, from 10^4 to 10^6 yrs after the supernova event for the T1998 SNCool simulation. A sharp density peak is produced at the interface of the shock and the surrounding medium, where the interstellar medium is swept up by the shock. There is a weaker pressure peak, and this dies off almost entirely in later times, although this is partly due to the radial averaging of a shock that becomes increasingly non-spherical at late times; by contrast, small peaks in pressure and temperature are visible in Thornton et al. (1998), since these are 1D spherical models by construction. Note that this shell is also where the kinetic energy fraction exceeds the thermal energy fraction, since the velocity peaks at the surface of the shock, while cooling dominates at the surface where the (relatively) cold medium is swept up by the shock. Various reflection shocks are also seen in the figures, seen as peaks in velocity and density behind the main shock. Again, these are not readily visible in the pressure or temperature plots. A large underdensity is also observed at the centre of the supernova remnant, which is not visible in Thornton et al. (1998), though they do find some underdensity there. It is unclear what causes this, though it does not affect the broad agreement with the other radial distributions measured in their work. By 10^6 yr, large Vishniac instabilities (Vishniac, 1983) are visible at the shock interface. Large grid-axis-aligned spikes are also visible; these are likely to be the result of odd-even-decoupling, which is inevitable if we do not accept the constraints of the “fixes” to this problem discussed in chapter 2. Also discussed in 2, the instabilities exhibit symmetry with the grid, though this is likely due to the grid seeding the instabilities rather than their being unphysical. The energy in turbulent flows varies with time, usually corresponding to interfaces in density or pressure inside the remnant, though the fractional contribution to the total energy in a radial bin does not exceed 1% (note that this figure might

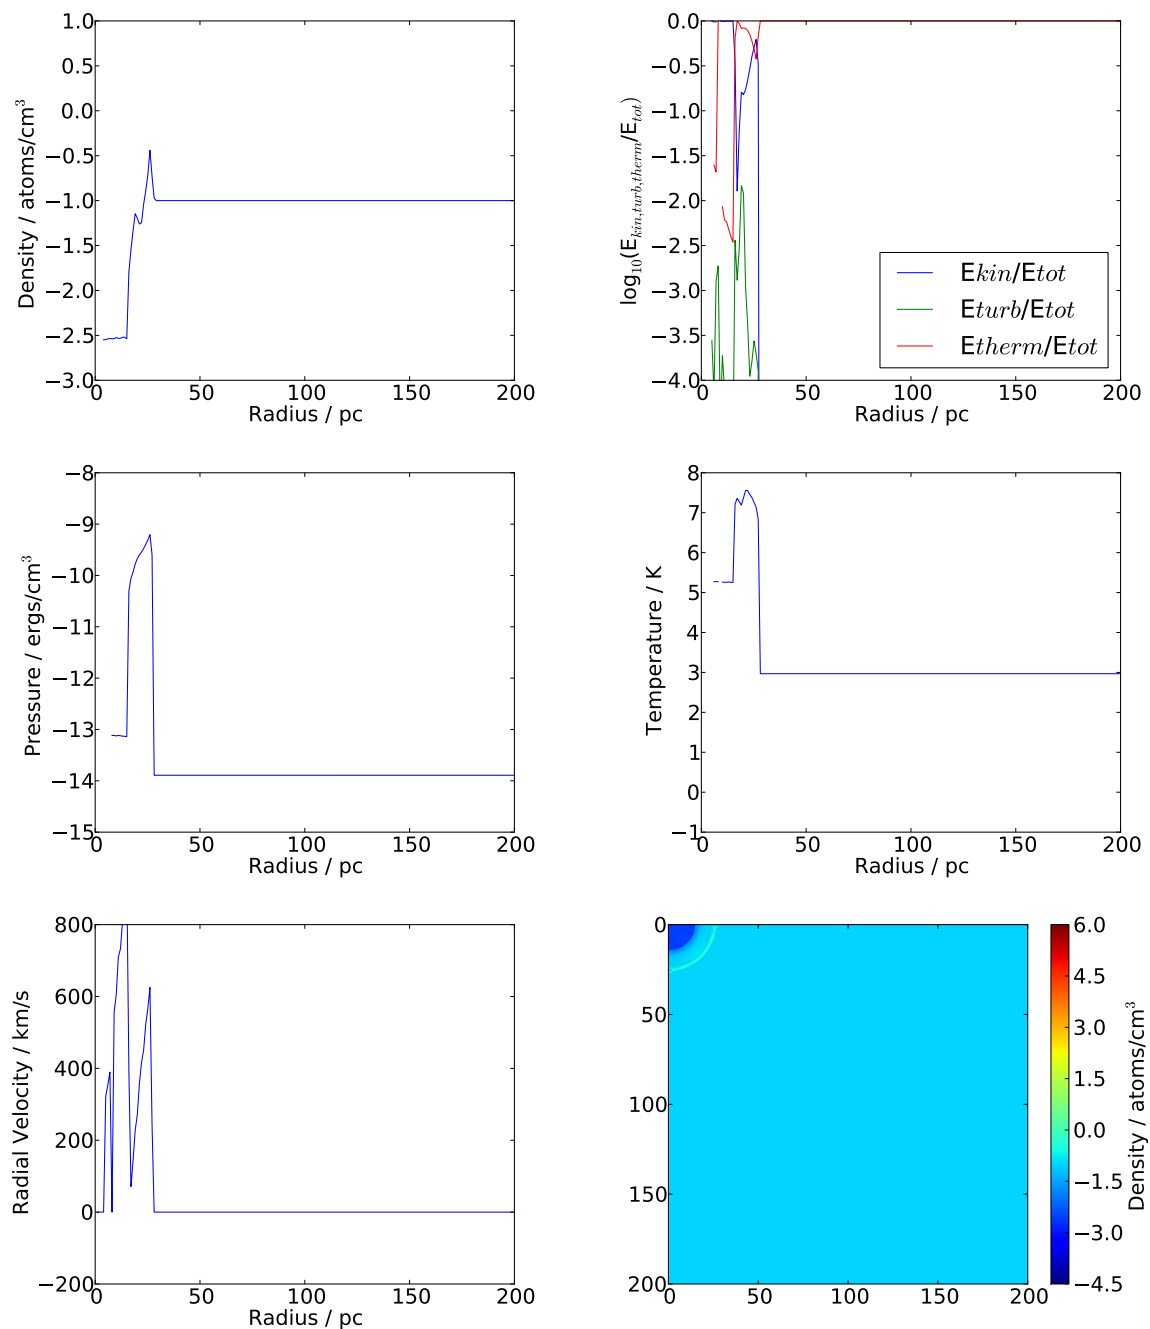


Figure 5.5: Physical properties of the supernova remnant by radius at 10^4 yr after the supernova event for the T1998 SNCool simulation. From top left to bottom right I plot density, energy partition, pressure, temperature, radial velocity against radius, and a 2D slice through the centre of the supernova in density. Energy partition is measured as the fraction of the total energy in a radial bin found in kinetic, thermal and turbulent energy. “Kinetic energy” is found by measuring all radially-aligned velocities. “Turbulent energy” is approximated by measuring all non-radial velocities. The simulation is still in the “quasi-adiabatic Sedov phase”, and has yet to form a kinetic energy-dominated shell. Multiple velocity peaks indicate multiple reflection shocks off the contact surface with the external medium.

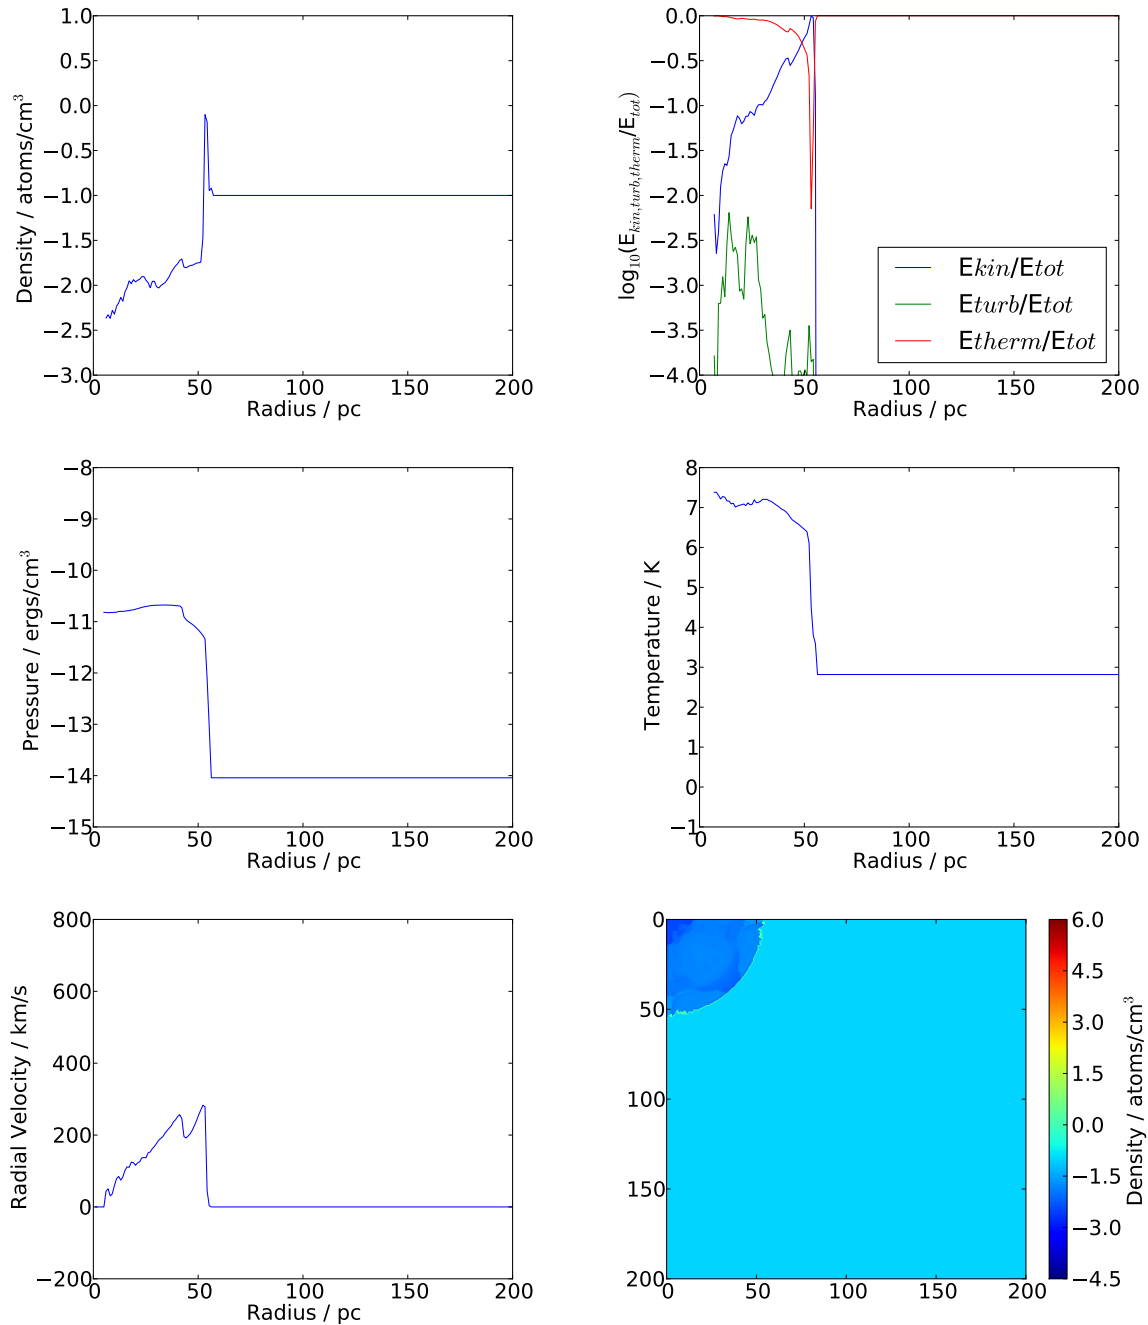


Figure 5.6: As in figure 5.5, but at 10^5 yr after the supernova event. A momentum-driven shell has begun to form, but the velocity profile is settling into a nearly-linear function with radius as per the Sedov solution. Any peak in pressure at the interface of the shock is smoothed out due to differences in the radius of the shock.

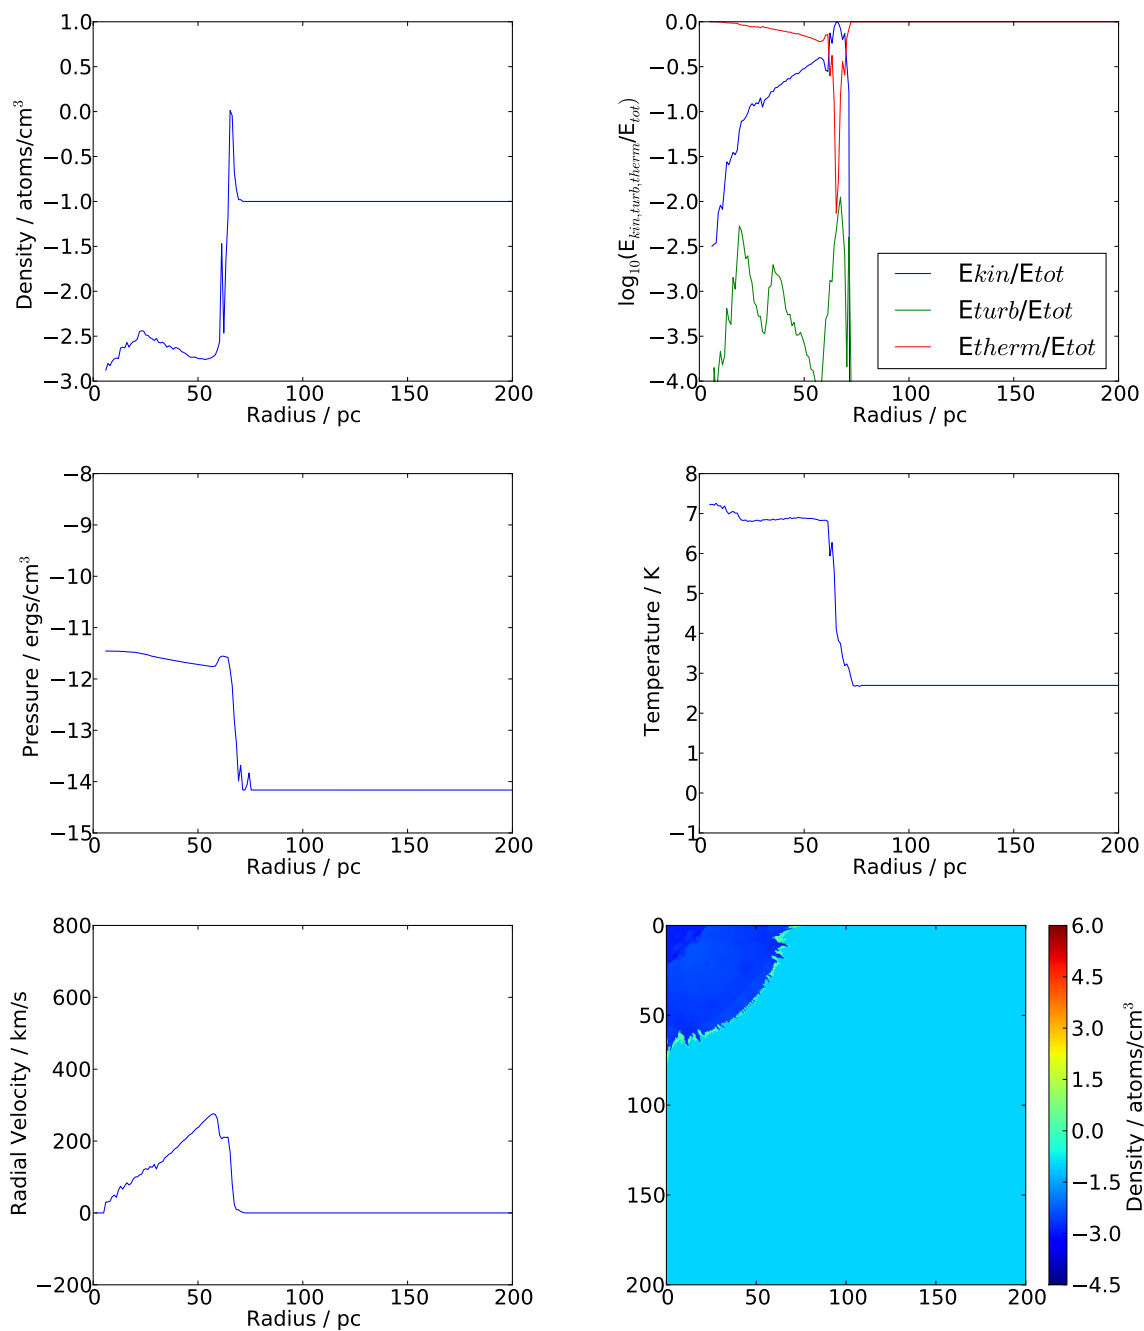


Figure 5.7: As in figure 5.5, but at 2×10^5 yr after the supernova event. Dense shell formation has begun, and finger-like Vishniac instabilities (Vishniac, 1983) have begun to form. A secondary density peak inside the remnant is also visible; this correlates with a convective zone in the density slice image, which contains relatively high energy contributions from non-radial velocities.

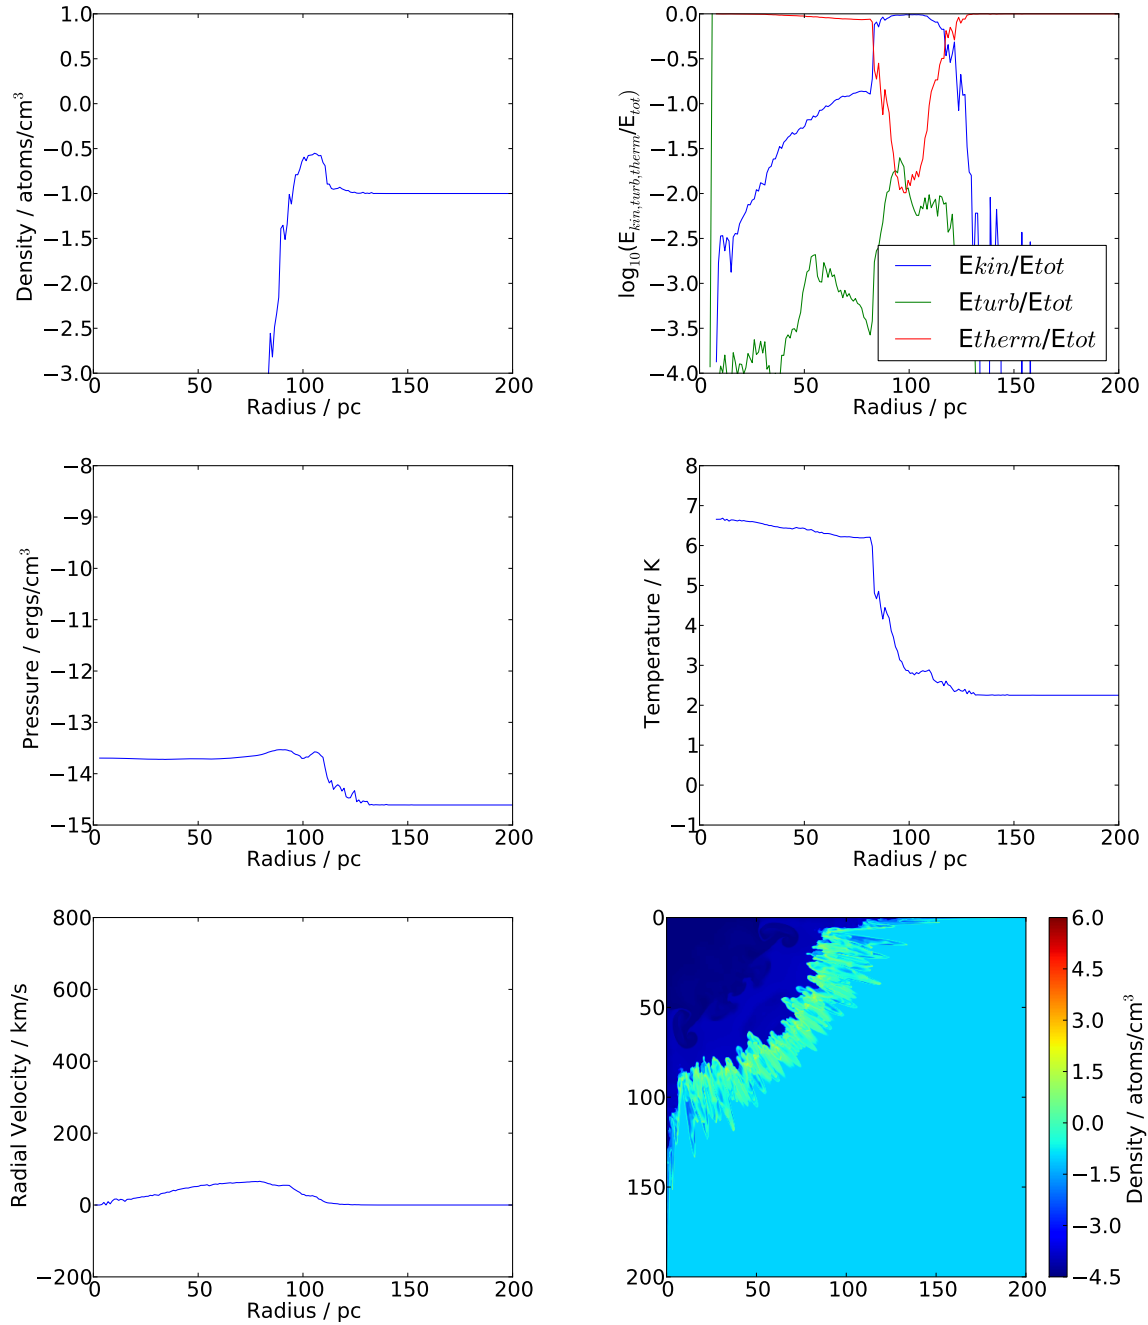


Figure 5.8: As in figure 5.5, but at 10^6 yr after the supernova event. The shock has begun to slow down, with large Vishniac instabilities now visible. The momentum-driven shell hence appears to be radially stretched due to this layer of instabilities on the contact surface.

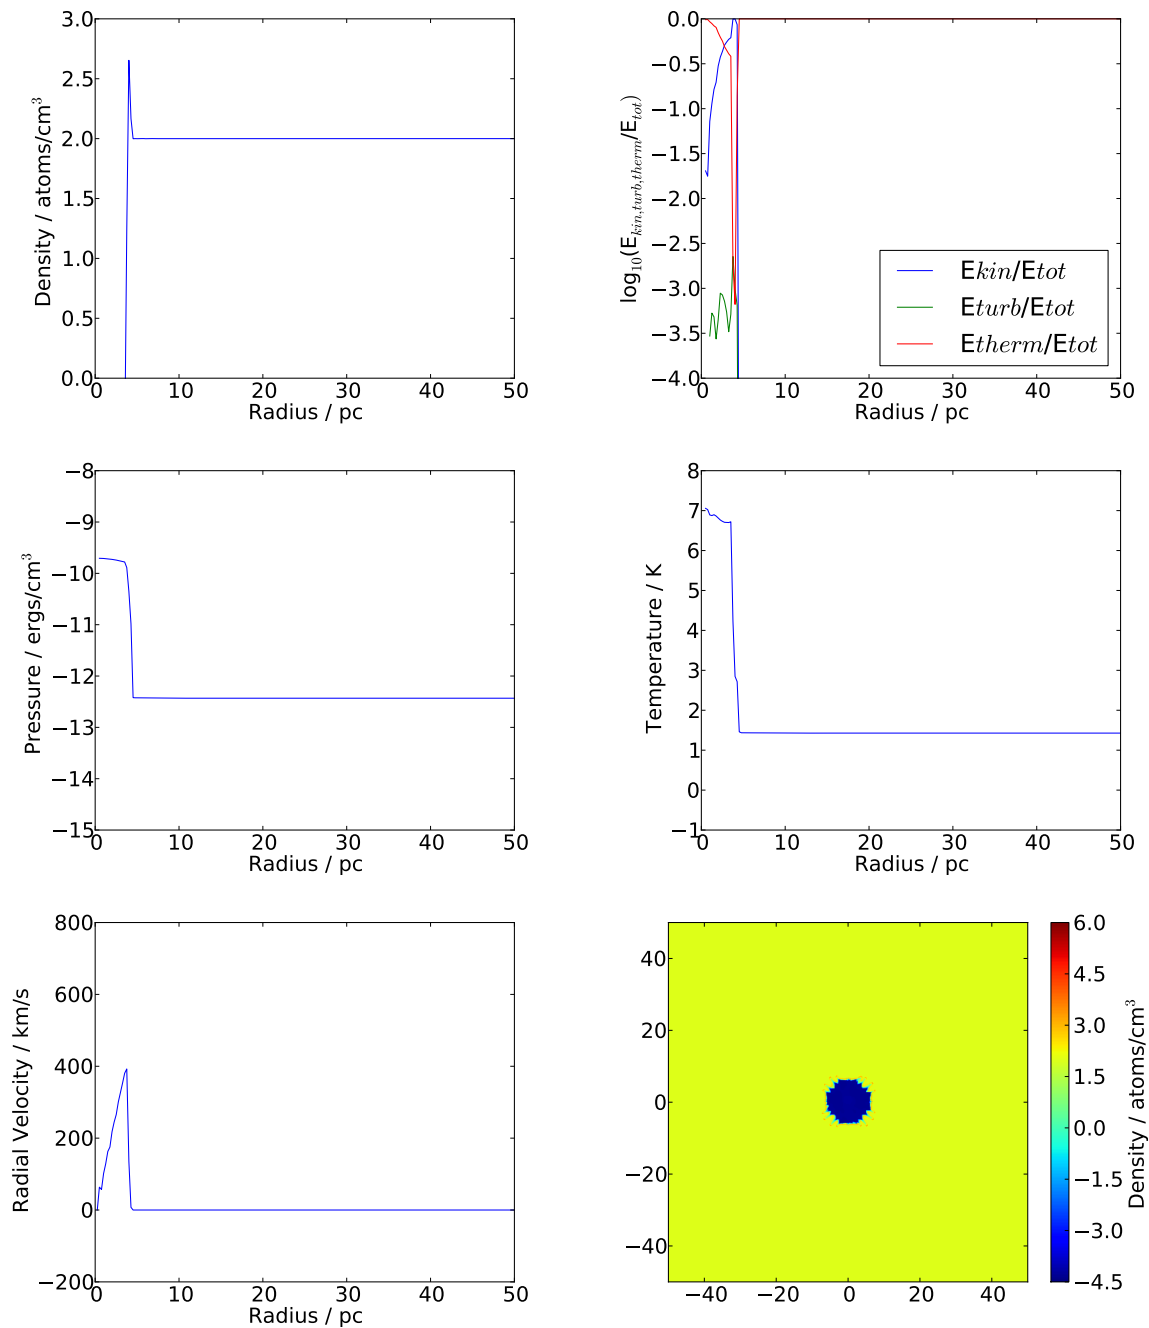


Figure 5.9: As in figure 5.5, but for the Cloud initial conditions. The formation of Vishniac instabilities is much faster in this run than in the T1998 initial conditions (see equation 5.4). The radial expansion is also slower as the external medium has a higher density.

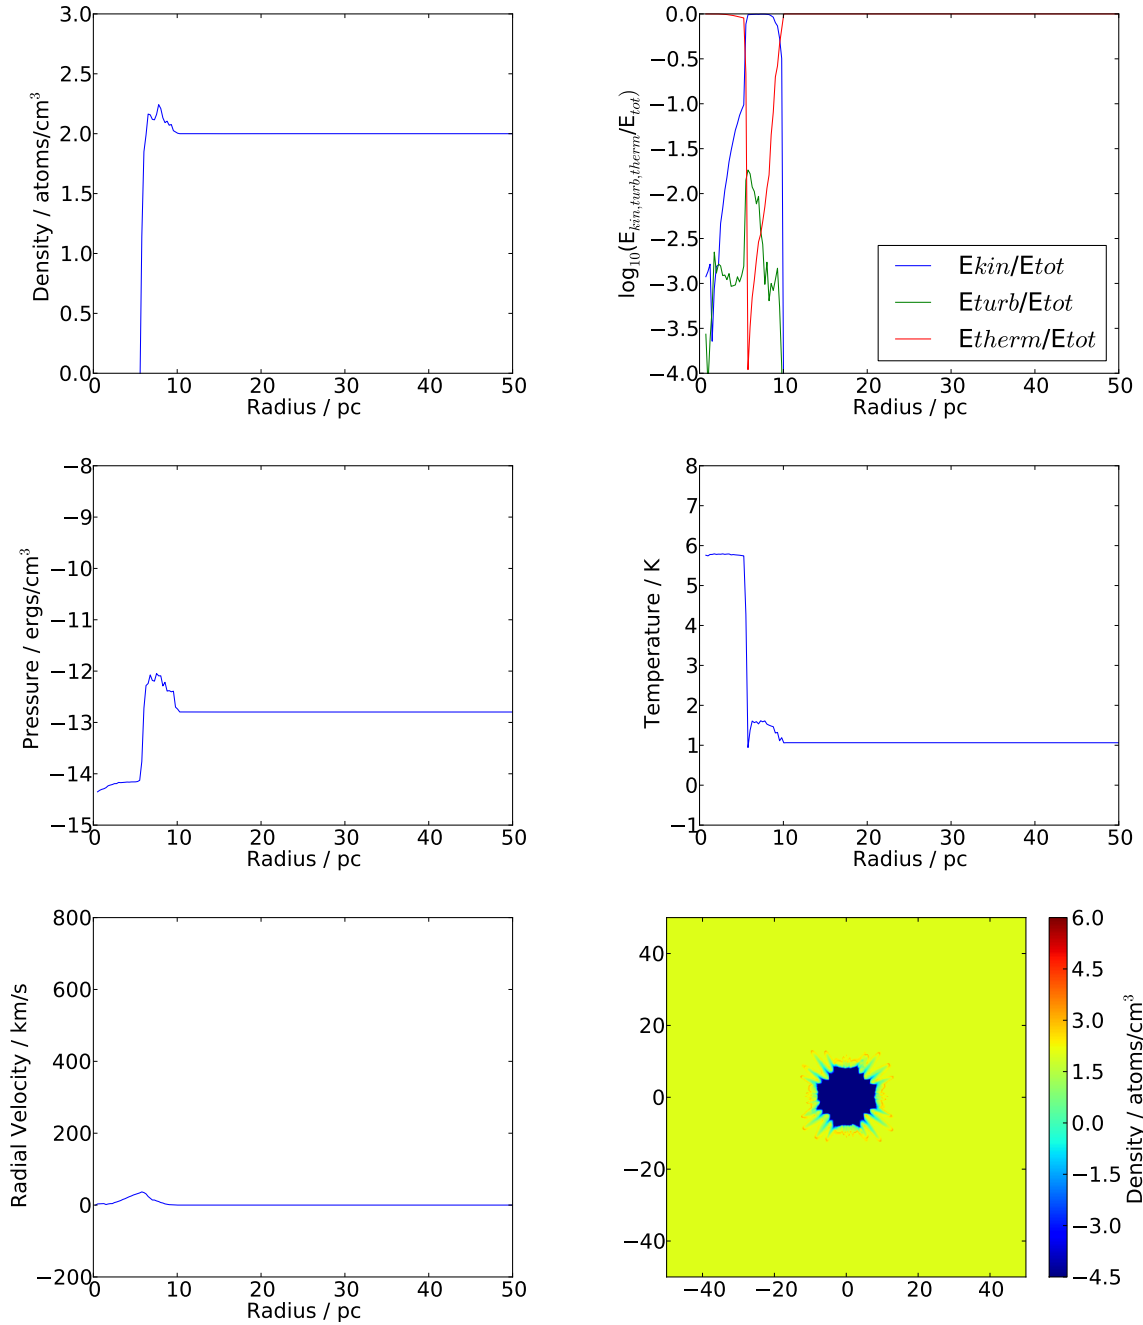


Figure 5.10: As in figure 5.9, but at 10^5 yr after the supernova event. The energy in turbulence on the shell/bubble boundary is growing.

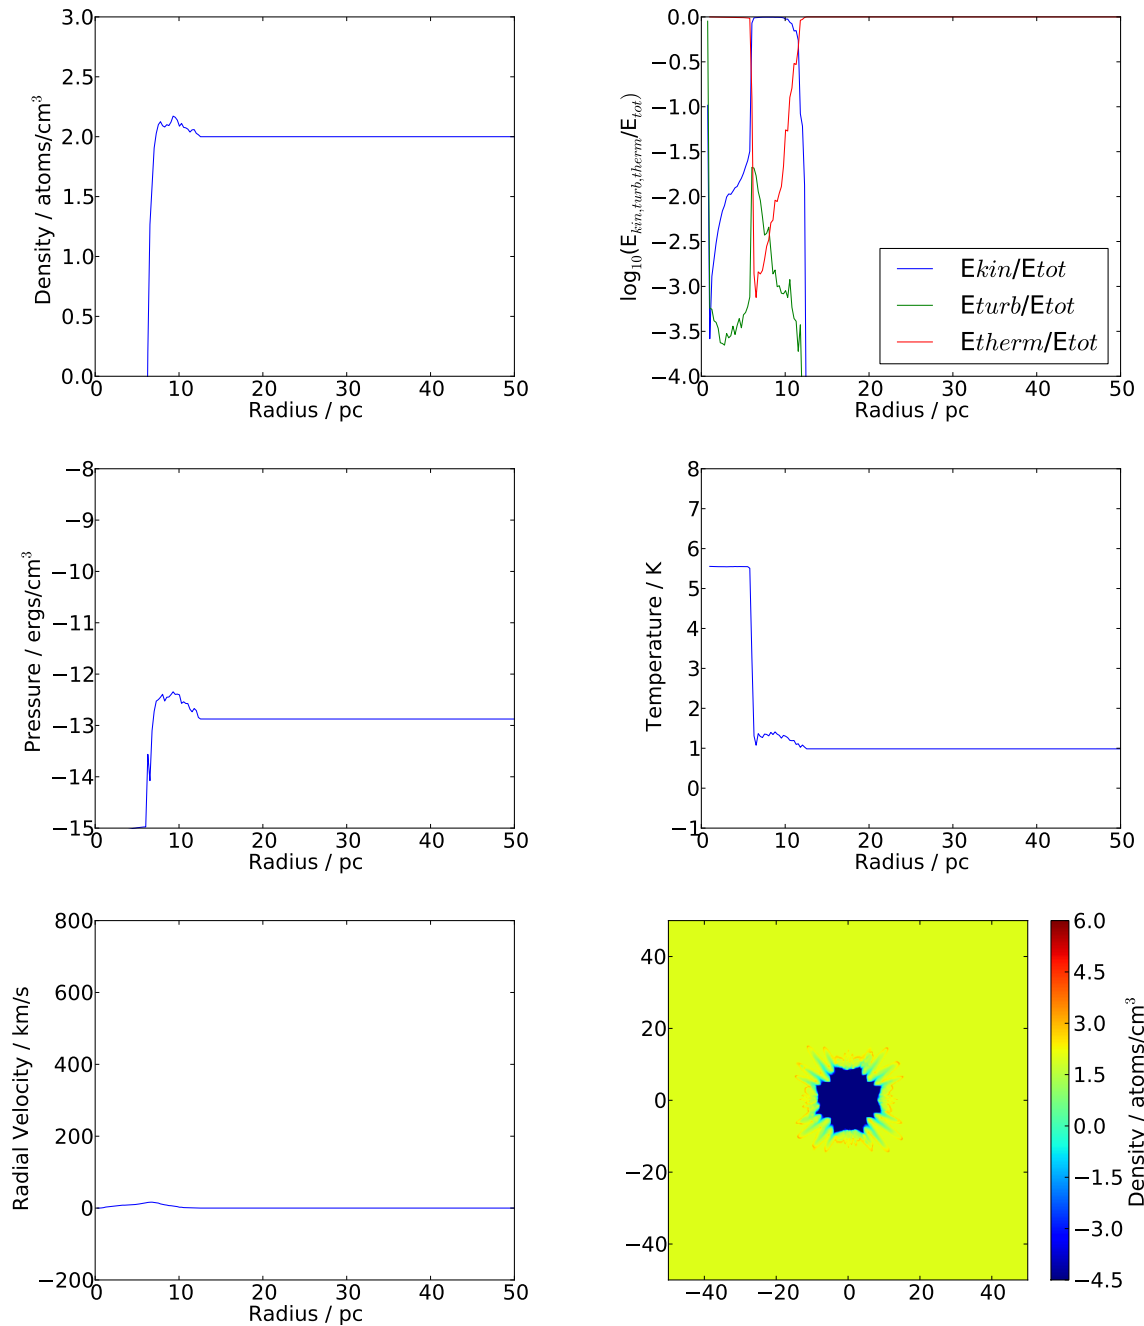


Figure 5.11: As in figure 5.9, but at 2×10^5 yr after the supernova event. The instabilities have grown, creating a wide shell as measured in radial coordinates. The radial velocity of the shock is reduced dramatically.

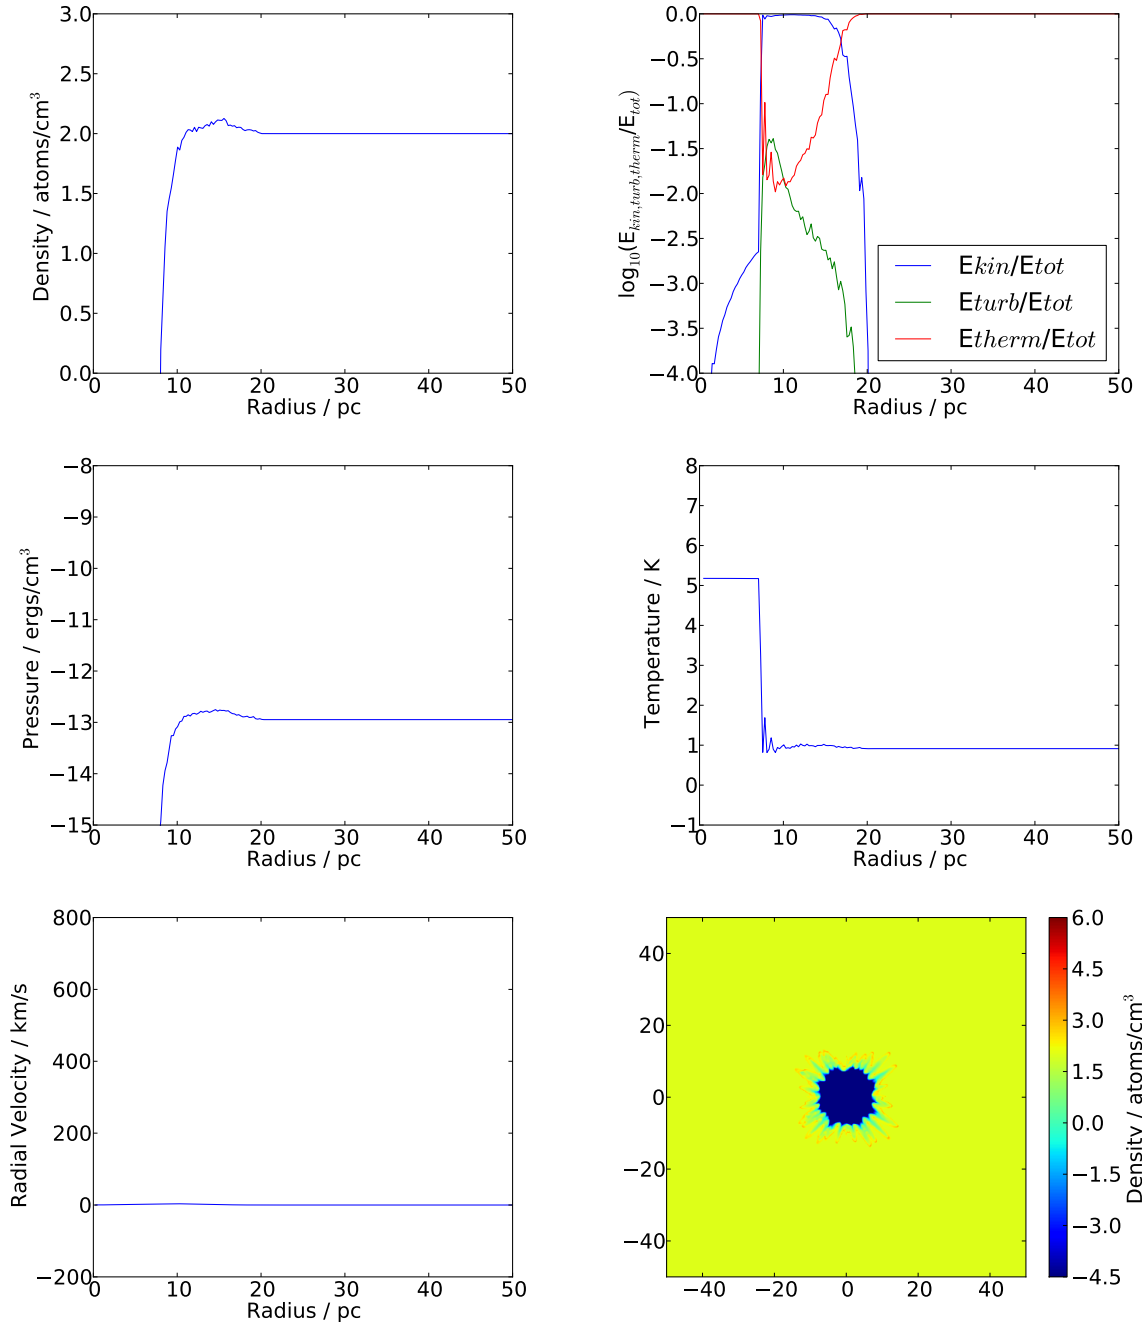


Figure 5.12: As in figure 5.9, but at 10^6 yr after the supernova event. By this point, the shock is virtually stationary, and the remnant dominated by the shell. The internal temperature of the bubble has dropped from 10^7 K to 10^5 K.




Figure 5.13: Blastwave energy as a function of time. The blue line shows thermal energy, the green line kinetic energy and the red line turbulent energy (see section 5.4 for definitions). The cyan line shows the total energy. The top plots show data for the Cloud run, and the bottom plots the T1998 run (see section 5.3). The left-hand plots show data from the non-radiating runs, and the right-hand plots the runs including metal-line cooling. Note that estimating the shock radius is difficult for highly non-spherical shocks and the estimate may exclude part of the shock or include part of the background medium. Note that the first $\sim 10^{4.5}$ yrs are poorly resolved in this plot and hence the energies are underestimated.




Figure 5.14: Mass contained behind the shock as a function of time for the Cloud and T1998 runs with and without gas cooling (see section 5.3). The supernova remnant mass is significantly higher than the star's original mass due to the amount of matter swept up from the interstellar medium in the blast. I give power-law fits to these lines in the text.




Figure 5.15: Shock radius as a function of time for the Cloud and T1998 runs with and without gas cooling (see section 5.3) The sampling of growing instabilities in the SNCool runs create uneven profiles. I give power-law fits to these lines in the text.

increase by 50% if we assume that there is also a radial component in turbulence that is not included here - this is considered in figure 5.13). It is worth mentioning that the Vishniac instabilities at 10^6 yrs will not contribute a great deal to this figure, since they are radially aligned and have not yet been broken up by Kelvin-Helmholtz instabilities; hence measurements of non-radial flows will not capture these.

Figures 5.9 and 5.12 contain similar plots for the Cloud initial conditions. The broad physical features of the radial distributions are very similar. However, radial instabilities develop much earlier, resulting in a much wider shell when measured radially. I also observe many fewer reverse shocks captured in these plots. The internal temperature of the remnant is also roughly an order of magnitude lower than in the T1998 simulation; the high density of the external medium allows more efficient cooling than in the T1998 simulations.

The Cloud and T1998 simulations also exhibit different energetics. As figure 5.13 shows, The NoCool simulations (i.e. no radiative cooling) are very similar in both sets of initial conditions, with roughly 7% kinetic energy, 93% thermal energy, and a negligible amount of turbulence; in fact, the flow becomes less turbulent as it evolves. However, when cooling is introduced, the picture changes. The T1998 simulation retains the Sedov partition of energy until 10^5 yr, with the kinetic energy finally dominating the thermal energy of the blast at 3×10^5 yr, whereas the Cloud simulation very quickly begins to lose energy as it sweeps up the dense medium into a dense shell and moves to the momentum-conserving, snowplough phase; this matches with the findings of Tenorio-Tagle et al. (1991). In this simulation, the kinetic energy dominates for all measured times, as the thermal energy is very quickly dissipated, though it appears to hit a floor around 10^5 yr where the cooling becomes less efficient. Note that kinks in these figures are possibly due to the sampling of the energies due to the non-sphericity of the remnant; it is difficult to measure the exact radius of the interface of the remnant and the surround medium, and hence some uncertainty in the total energy of the remnant is derived.

We measure the global properties of the supernova remnant by integrating over the radial profiles in figures 5.5 to 5.12. For the T1998 SNCool model at 10^6 Myr, I find a remnant mass of $1.64 \times 10^4 M_{\odot}$ and radius 138 pc. The total energy inside the remnant of 1.02×10^{50} ergs is partitioned between kinetic and thermal energy as 0.308:0.691. This gives us reasonable agreement with the results of Thornton et al. (1998), though have not run a full suite of models to compare our results fully. In a uniform medium, the

non-sphericity of the shock appears not to have a dramatic effect on the evolution of the supernova, at least before 10^6 yr. However, for a non-uniform medium, as discussed in section 5.1, density gradients have a dramatic effect on the evolution of the supernova shock. In a power law density gradient, for example, it is expected that the spokes of the shock will provide channels for the energy to stream into the low-density medium outside. In rough agreement with Gull (1973), I find that the turbulent energy is 1.3% of the total energy and 1.9% of the kinetic energy. Thornton et al. (1998) further determine the energy of the shell and bubble separately; however, a glance at the images in figure 5.12 and 5.8 reveal why this is a non-trivial problem in these simulations.

Figures 5.2 and 5.3 demonstrate a power law relationship between the remnant mass and radius, and time after the supernova. For the T1998 SNCool model, I derive the following linear relationship:

$$\log(M_{\text{remnant}}) = 0.9521 \log(t_{\text{SN}}) - 1.801 \quad (5.2)$$

$$\log(R_{\text{remnant}}) = 0.3121 \log(t_{\text{SN}}) + 0.1642 \quad (5.3)$$

where M_{remnant} is the remnant mass in M_{\odot} , R_{remnant} is the remnant radius in pc, and t_{SN} is the time after the supernova event in years, up to 1Myr. I do not explore the difference in slopes between the Sedov phase and the momentum-conserving, snowplough phase in this thesis; a more comprehensive analysis will be possible once I have explored the parameter space of initial conditions and physical models. As part of this process, I now turn your attention to the effect of stellar winds on the evolution of the supernova remnant.

5.4.2 Stellar Winds and Supernovae

In this section I introduce the WindCool set of simulation parameters, which I run for both the T1998 and Cloud initial conditions. The aim of this simulation set-up is to model the evolution of a $15M_{\odot}$ star's wind bubble and the effect that it has subsequent supernova remnant evolution in a uniform medium. This model is not particularly physically realistic for a number of reasons. First, a star will not typically be stationary with respect to the external medium due to the random motion of stars and interactions during its formation (e.g. Bate (2012)). Hence a bubble expanding into a medium that is stationary with respect to the source is an unlikely event. In a molecular cloud environment, one would also expect many other stars, in particular

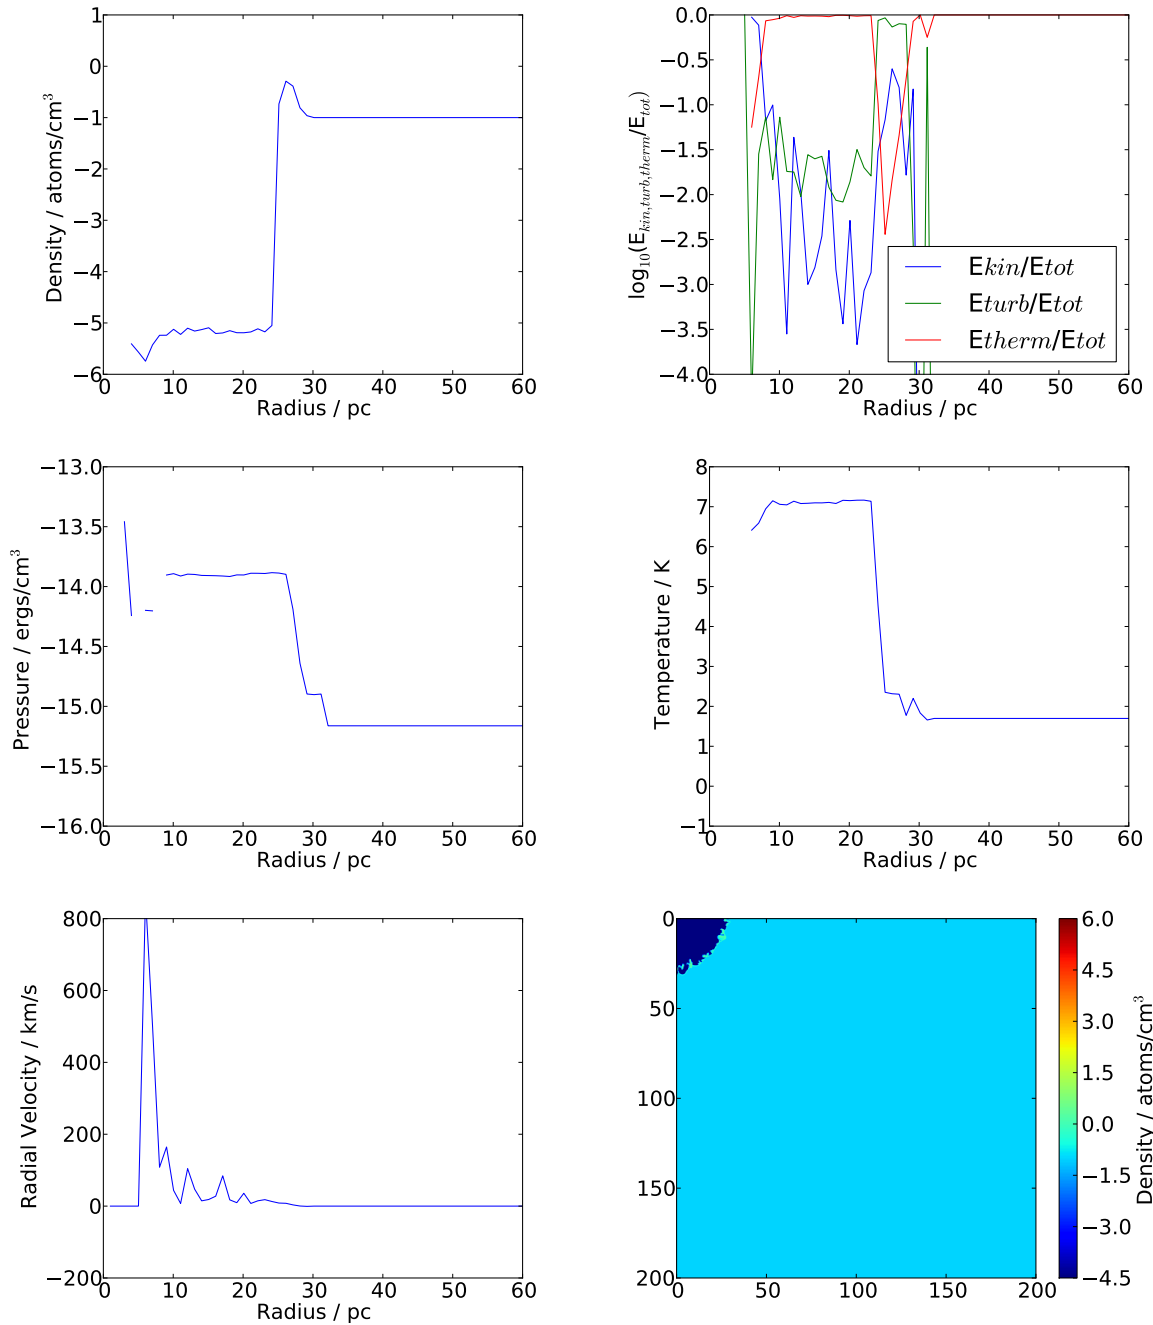


Figure 5.16: Physical properties of the wind bubble by radius at 5Myr after the formation of a $15M_{\odot}$ star for the T1998 WindCool simulation. From top left to bottom right I plot density, energy partition, pressure, temperature, radial velocity against radius, and a 2D slice through the centre of the supernova in density. Energy partition is measured as the fraction of the total energy in a radial bin found in kinetic, thermal and turbulent energy. “Kinetic energy” is found by measuring all radially-aligned velocities. “Turbulent energy” is approximated by measuring all non-radial velocities. Note that the internal profile is poorly sampled due to the continual injection of the wind energy and mass in a sphere of 2pc radius. Instabilities on the shock front are visible; see the text for an analysis of this phenomenon. Note also that the wind is highly turbulence-dominated.

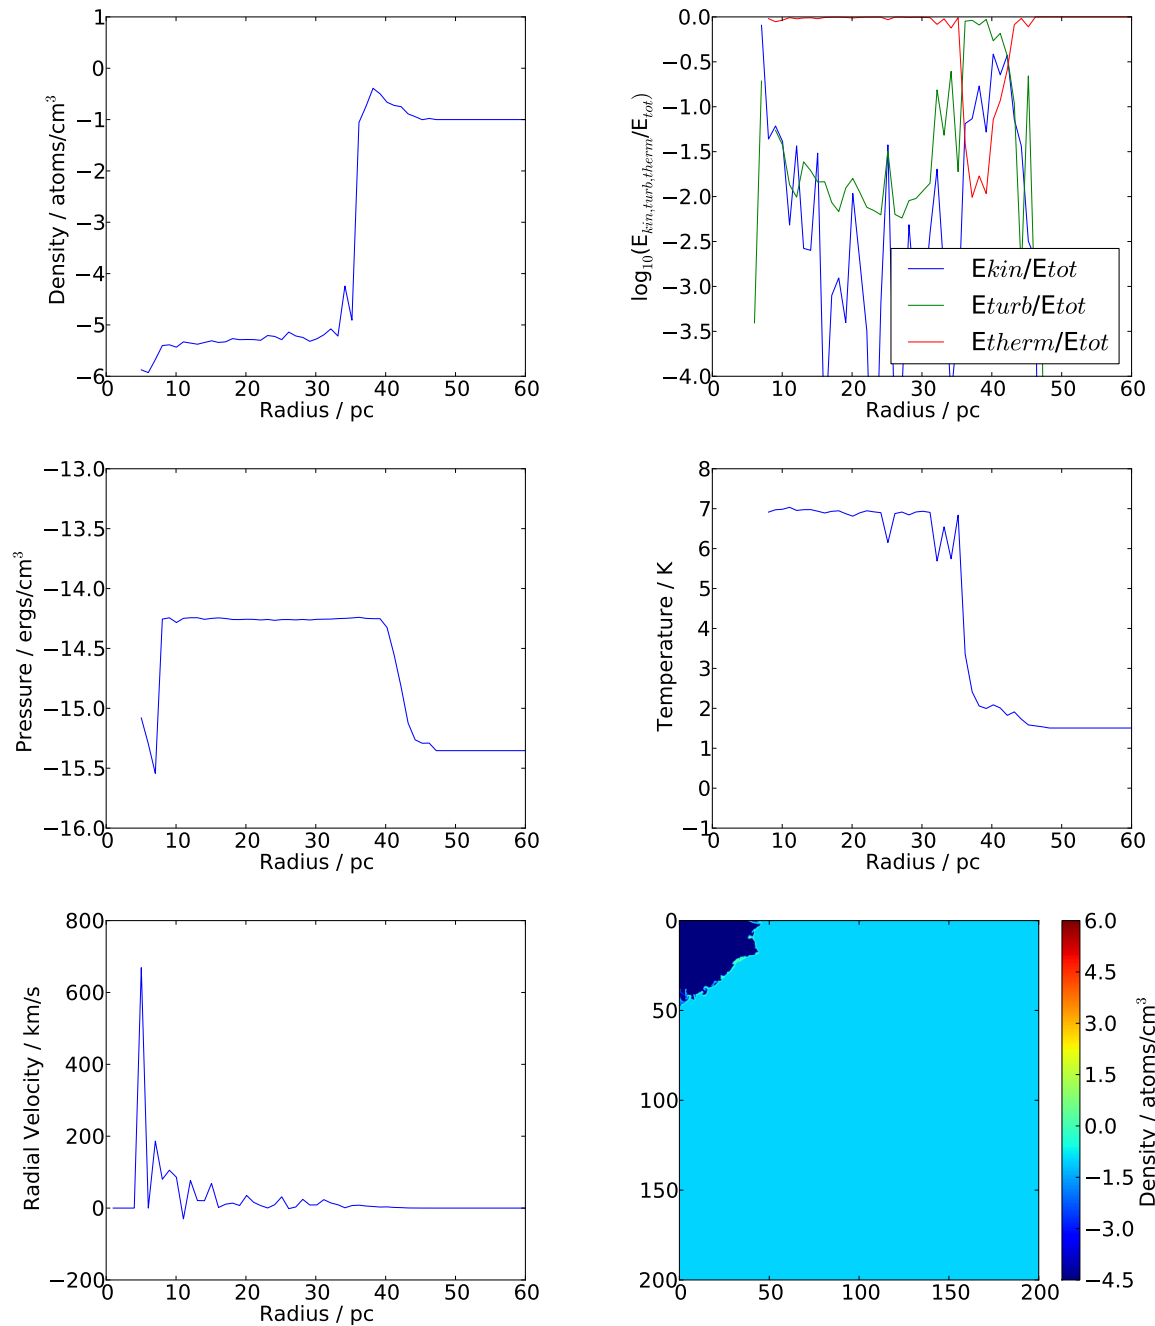


Figure 5.17: As in figure 5.17, but at 10Myr. A cooling-driven, turbulent shell begins to form as the instabilities at the shock front grow.

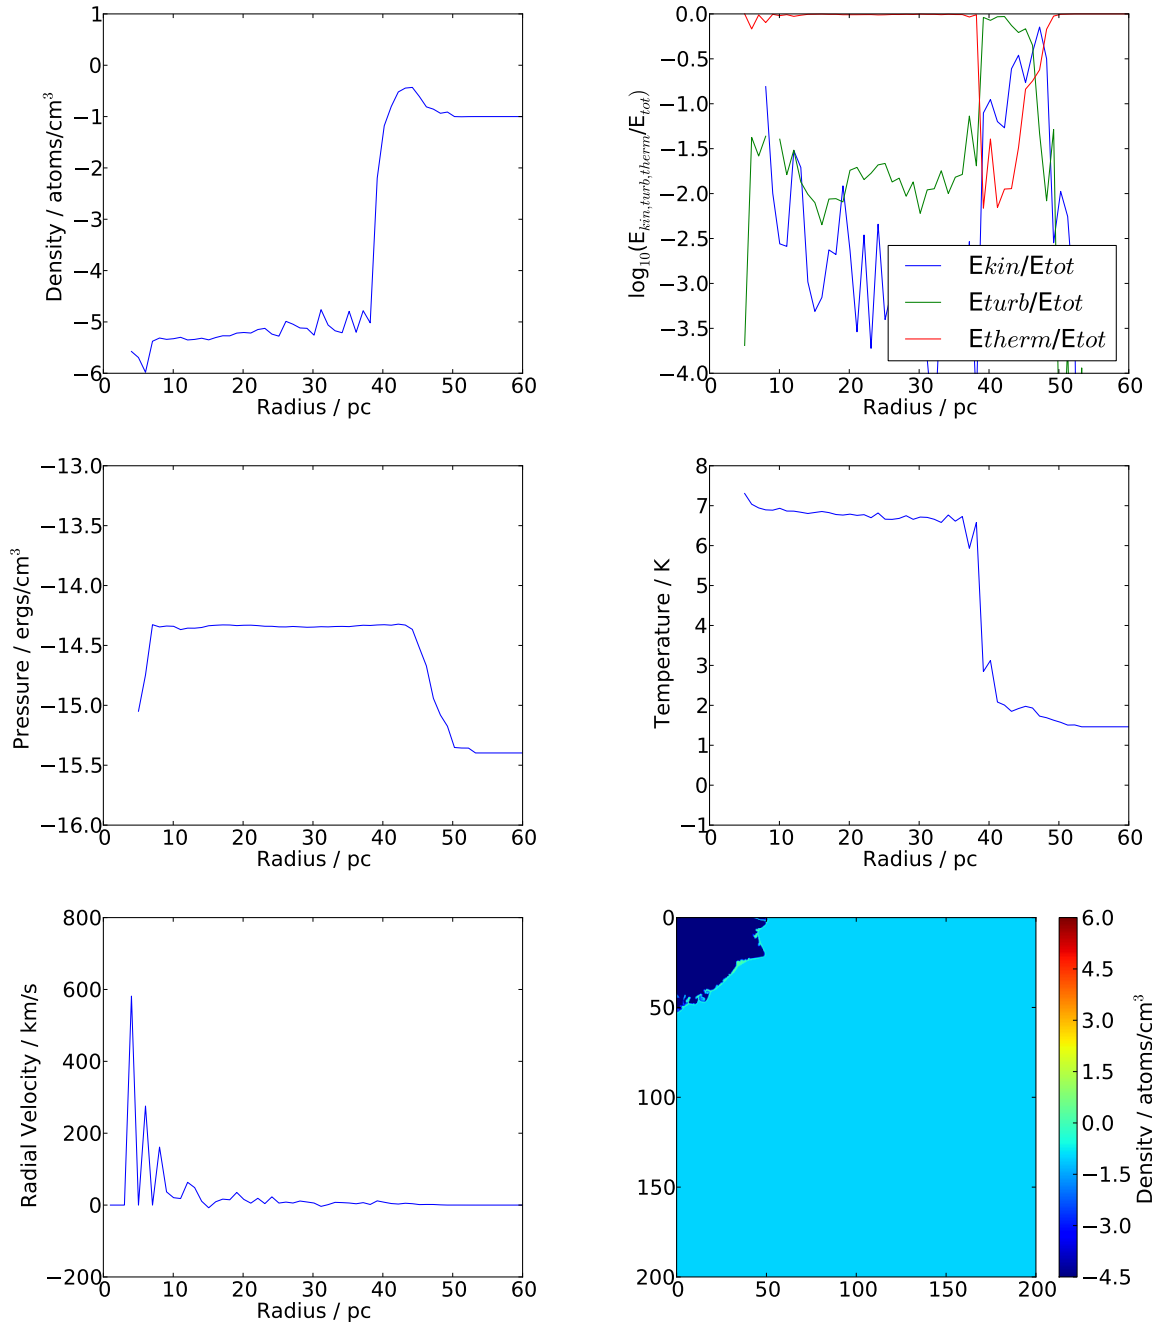


Figure 5.18: As in figure 5.18, but at 12Myr. Near the centre, the shocked wind gains some radial velocities, but this is quickly turned into turbulent motion towards the outer edge of the wind bubble.

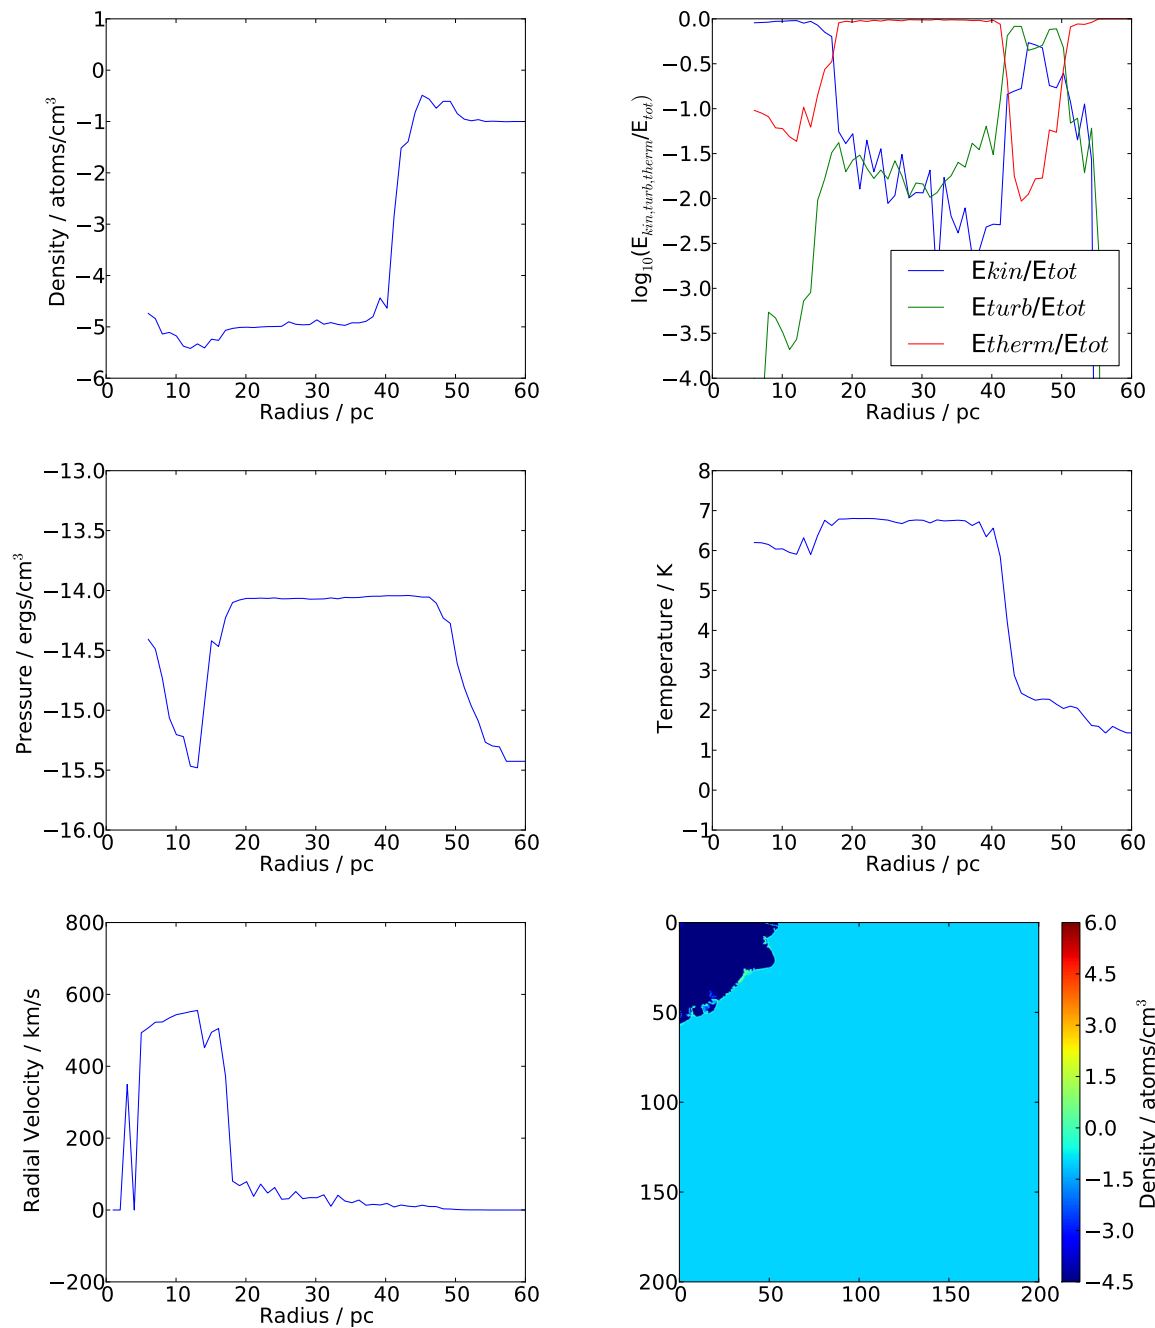


Figure 5.19: As in figure 5.19, but at 14 Myr. As the momentum and mass in the outflows from the star grows (see section 5.2, the bubble forms a kinetic energy-dominated core, which thermalises and becomes turbulent again at the interface of the fluid.

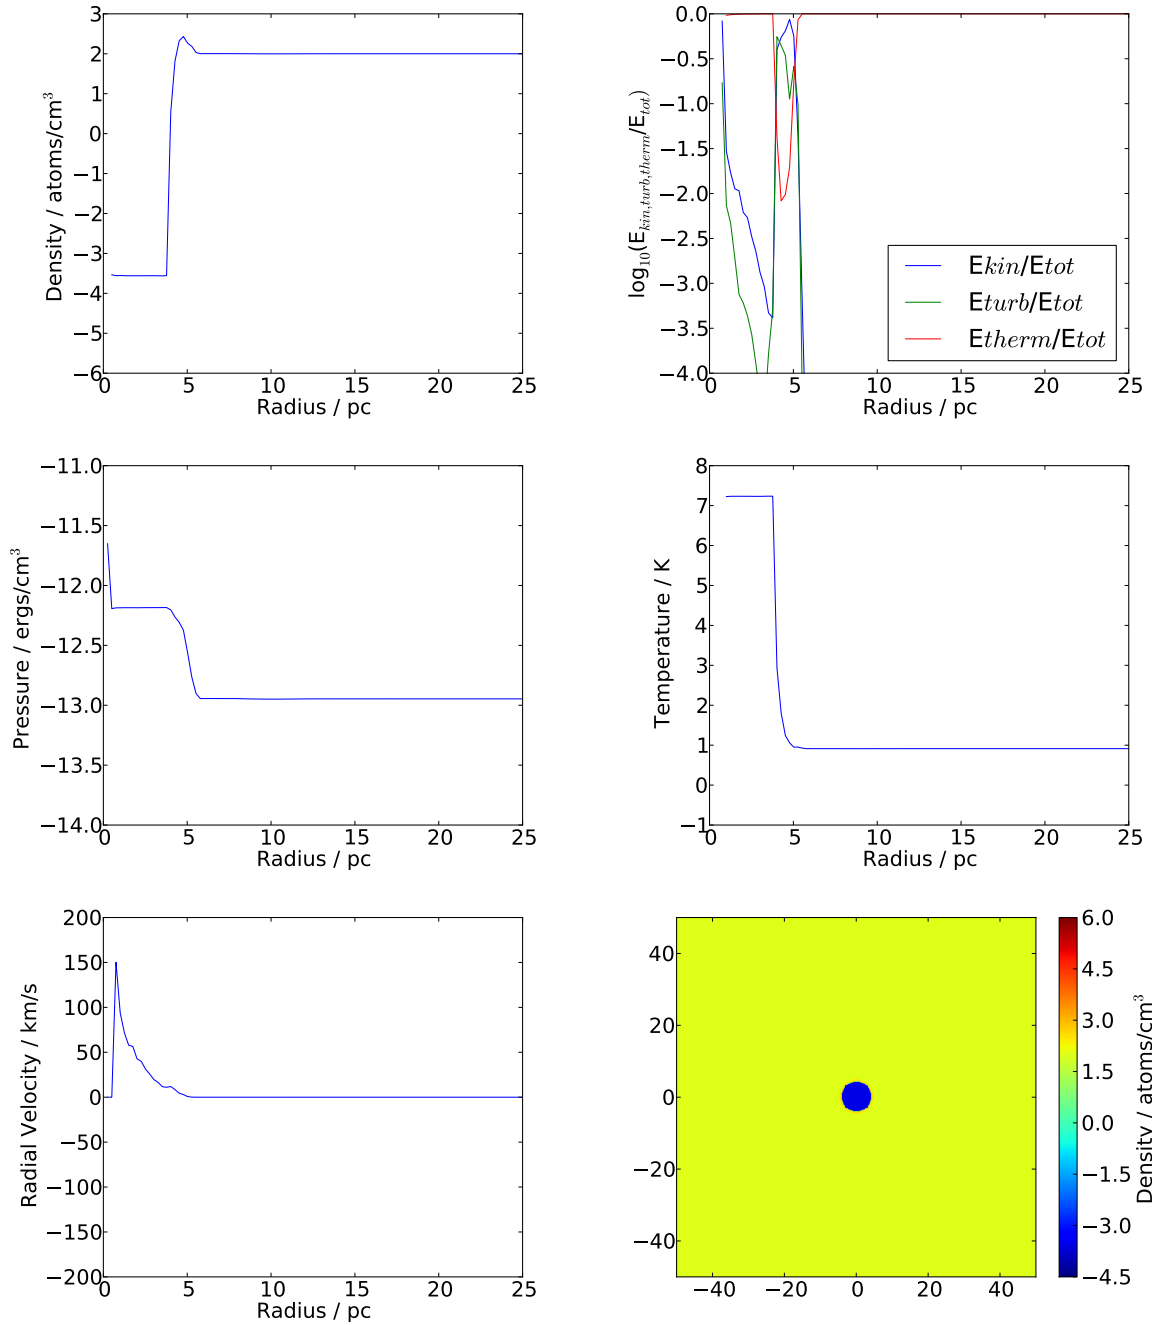


Figure 5.20: As in figure 5.20, but for the Cloud initial conditions at 5Myr. Though the flow still contains a lot of its kinetic energy in turbulent motions, there is no sign of instabilities on the wind bubble surface. I analyse this phenomenon in the text.

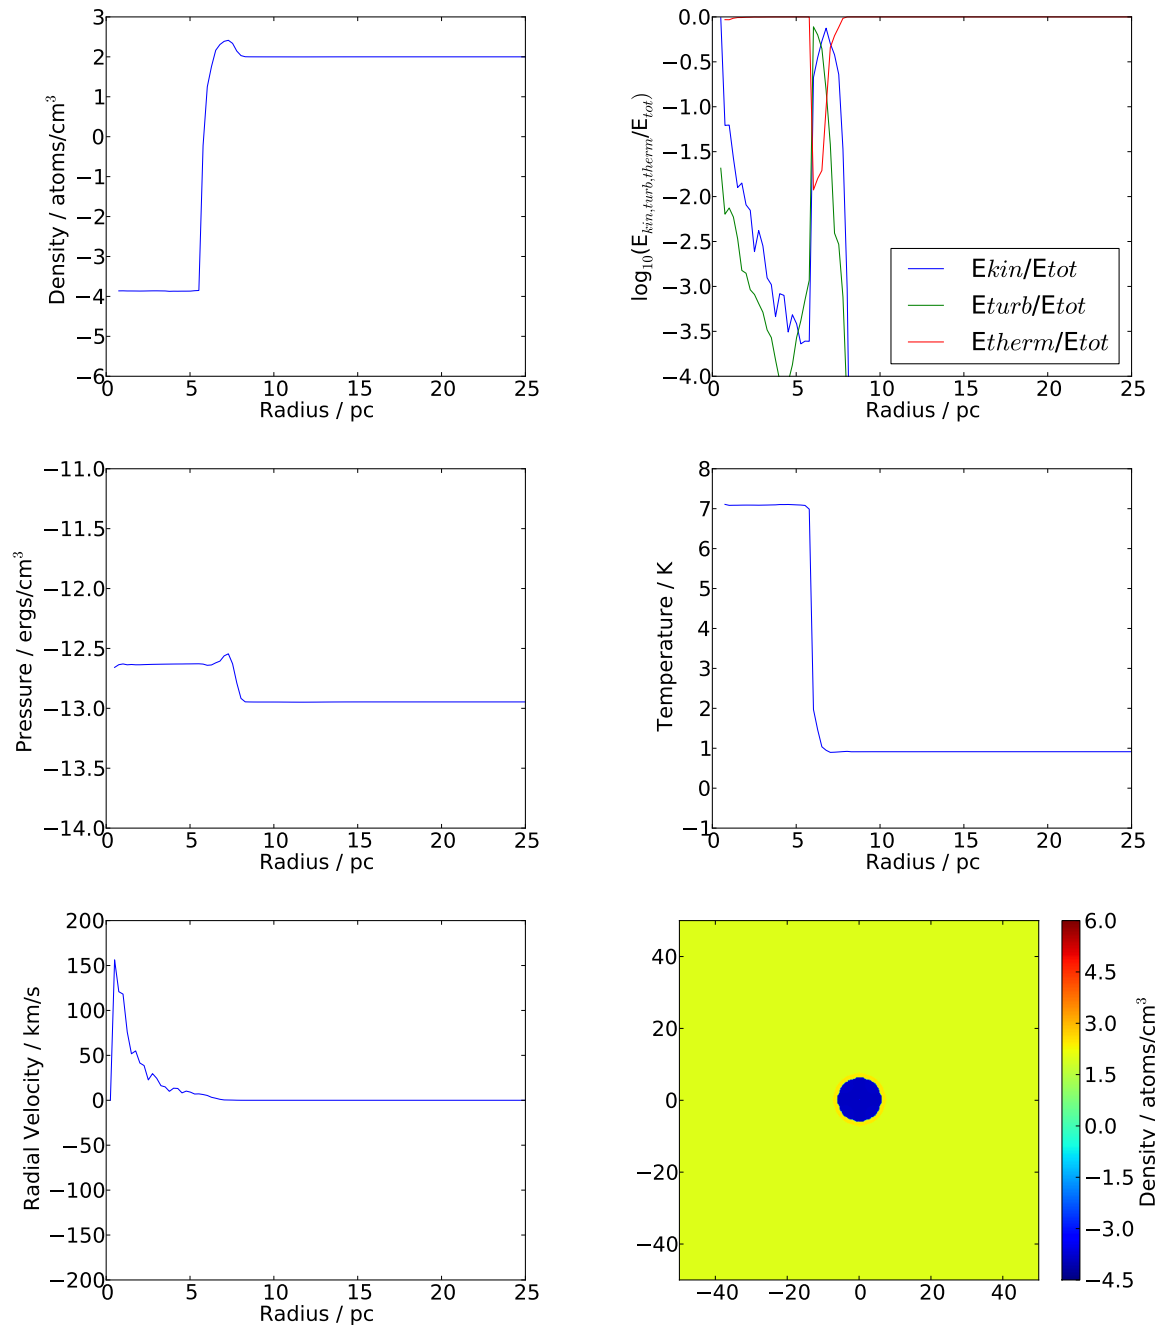


Figure 5.21: As in figure 5.21, but at 10 Myr. The bubble continues to grow much as in the first 5 Myr, though non-spherical structures are beginning to form on the shock surface.

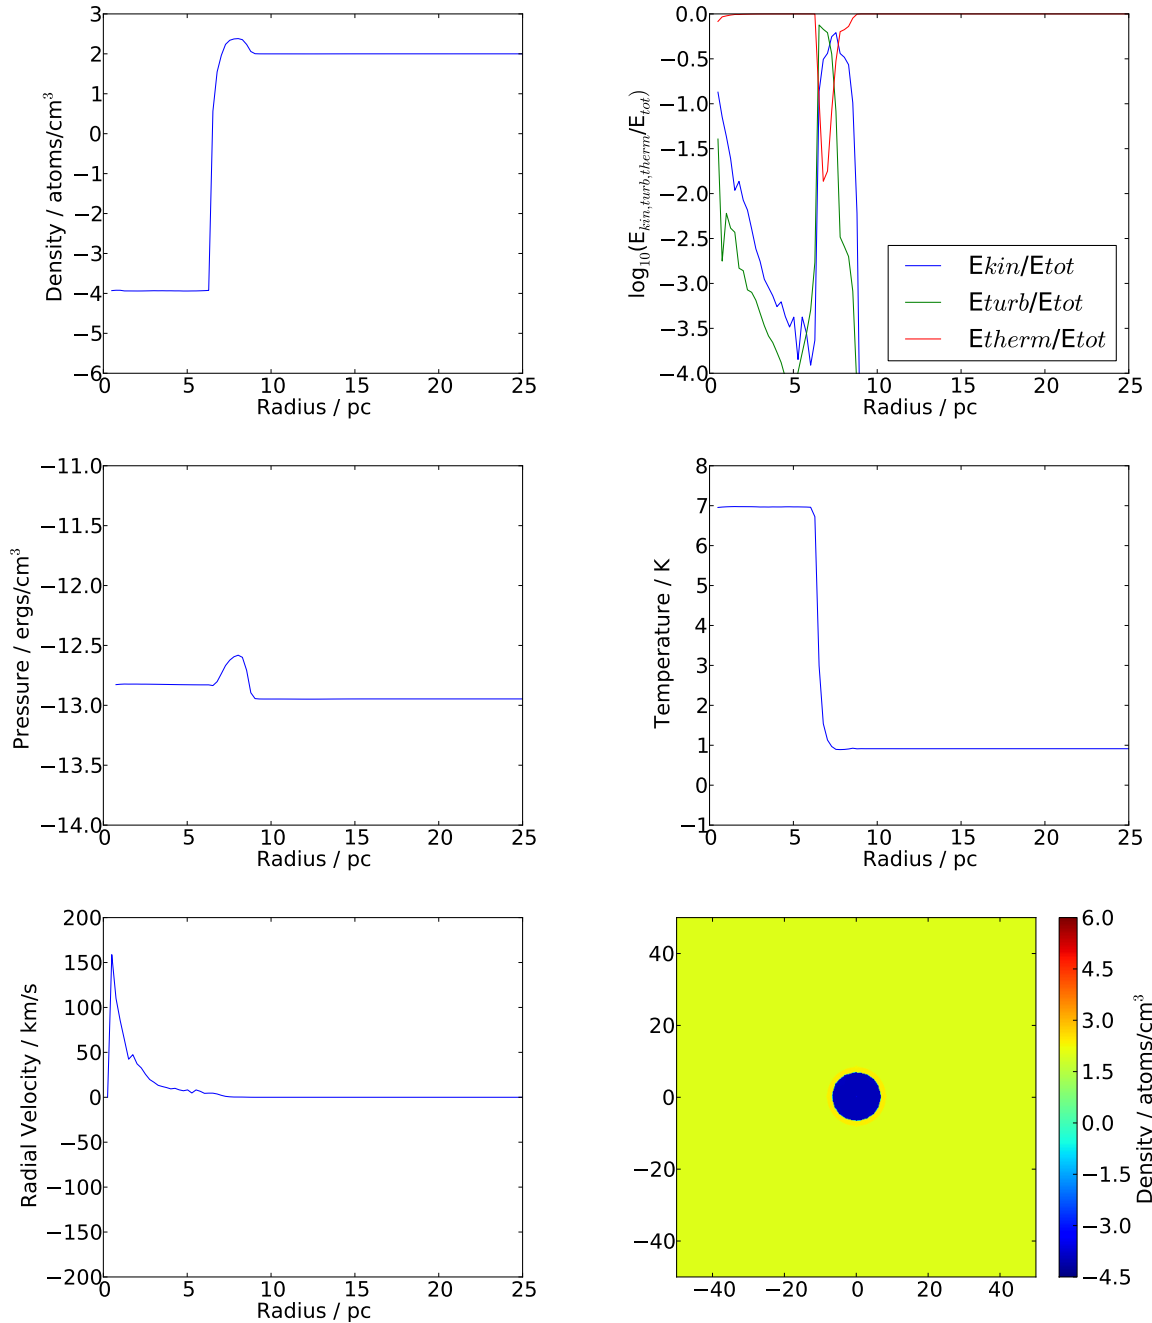


Figure 5.22: As in figure 5.22, but at 12 Myr. The profiles are again fairly similar, and the non-sphericity at the surface appears to have been damped.

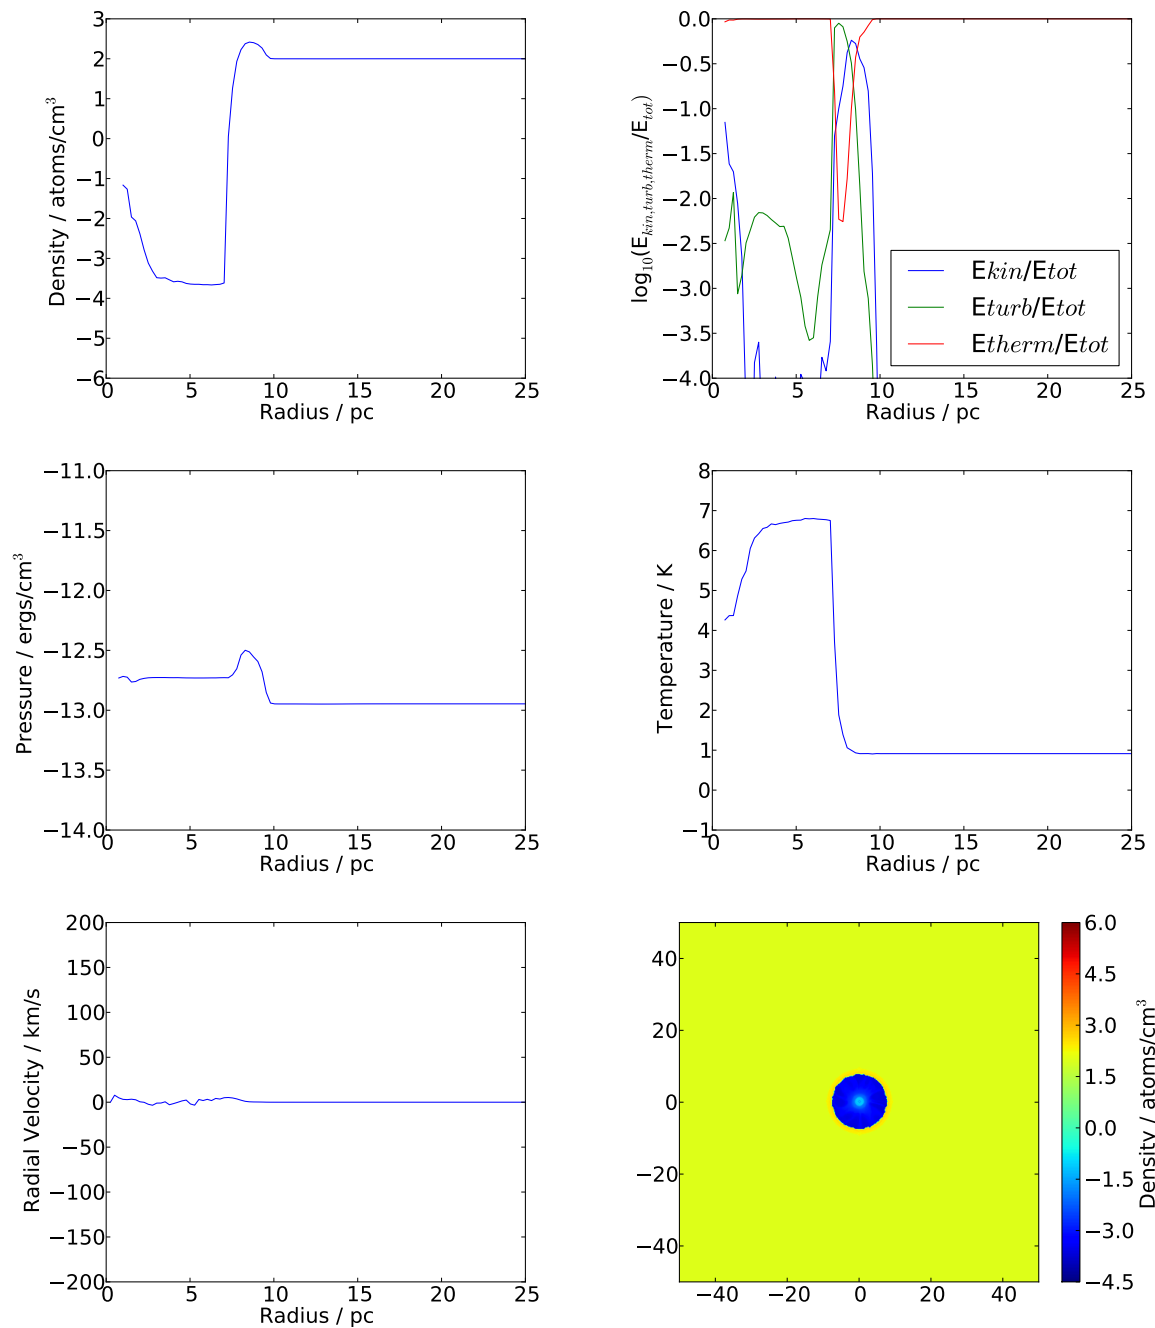


Figure 5.23: As in figure 5.23, but at 14Myr. As per figure 5.19 for the T1998 initial conditions, the increase in mass loss has disrupted the system; in this case, much of the kinetic energy is converted into turbulent motions. Small non-spherical structures are re-forming on the shock surface.

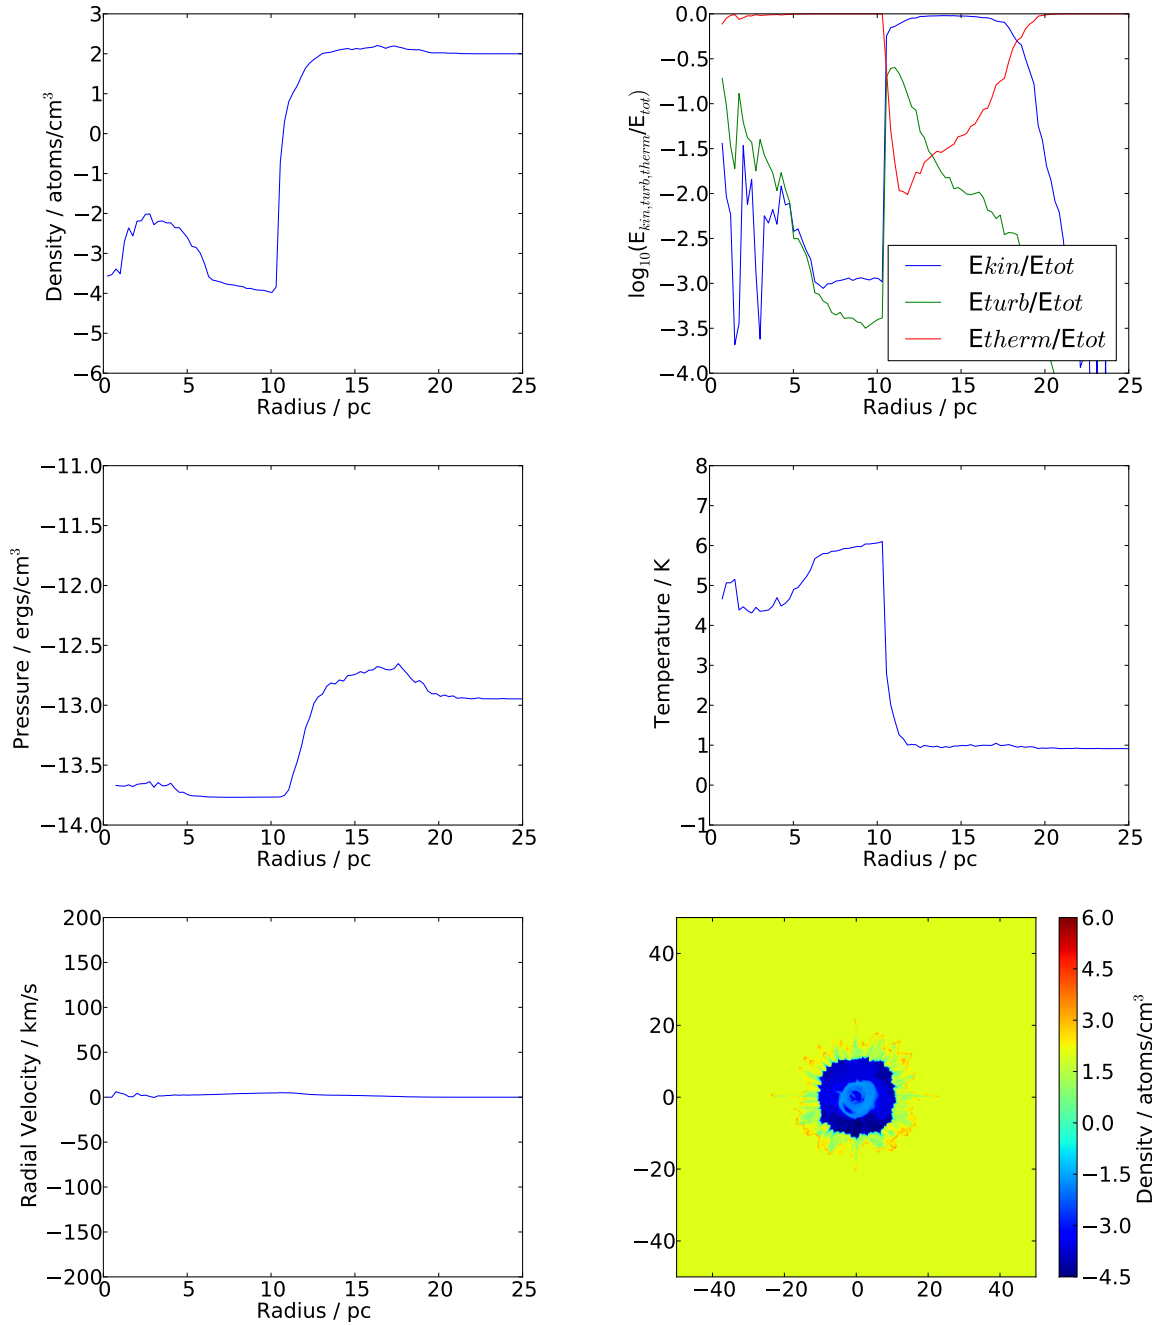


Figure 5.24: As in figure 5.24, but at 15.125 Myr, 1 Myr after the supernova event at 14.125 Myr. The radial structure of the supernova remnant is noticeably altered, with a radius that is $\sim 10\%$ than in the SNCool run - see the text for further analysis.

OB associations, or groups of more massive stars that dominate the photoionisation and kinetic wind output (see section 5.1). These stars will also go supernova first. In a future work I may study the evolution of a supernova blast into such a medium, though this will require careful simulation of the initial stellar population and cloud environment. For now, I concentrate on the simple example of the star's own wind, as discussed in section 5.2.

First, in figures 5.16 to 5.23 I consider the radial properties of the bubble itself in both sets of initial conditions before the supernova event. I also follow the evolution of the supernova remnant for the Cloud initial conditions in figure 5.24; the simulation for the T1998 initial conditions is still ongoing, as the larger simulation volume necessitated by the lower background density takes longer to simulate. Note that in all of these figures, the inner region of the bubble is poorly sampled by our profiles due to their small volume compared to the rest of the radial shells in the bubble. As discussed in section 5.3, I deposit the stellar wind onto a sphere of 2pc radius, meaning that the inner 2pc will usually reflect the wind solution imposed onto it by the code rather than the subsequent, natural evolution of the wind.

All of these simulations exhibit a low-density, few-parsec-wide momentum-driven shell and thermally-dominated shocked wind bubble, with very low internal densities and a 4-6 order of magnitude density contrast. The velocities inside the wind bubble are dominated by turbulent, non-radial flows, due to the continuous injection of mass and energy in increasing quantities into the bubble rather than a single blast as in the previous section. At the very centre, where the energy is injected as thermal energy, the wind quickly converts 10% of its energy into kinetic energy, but this drops with increasing radius as the turbulent gas thermalises, until it hits the cooling-dominated, momentum-driven shell.

I now wish to address the creation of instabilities on the surface of the wind bubble using quantitative analysis of the growth rates of various classes of instabilities. To do this, I refer to the values shown in the radial profiles given in figures 5.16 to 5.23 inside and outside the wind bubble. In all images of the Cloud simulation, the surface of the wind bubble appears to remain stable, exhibiting only small Vishniac instabilities. By comparison, the interface with the medium in the T1998 initial conditions very quickly becomes unstable. Following Vishniac (1983), the growth rate of the instabilities on the surface of the bubble may be approximated as:

$$\Gamma_v \simeq \left(\frac{P_{\text{internal}}}{\rho_{\text{external}} v_{\text{shock}}^2} \right)^{\frac{1}{4}} \frac{c_s}{h_{\text{shell}}} \quad (5.4)$$

where P_{internal} is the pressure inside the bubble, ρ_{external} is the external density, v_{shock} is the expansion velocity of the shock, c_s is the speed of sound inside the bubble and h_{shell} is the width of the shell. The external density in the T1998 initial conditions is three orders of magnitude lower than that found in the Cloud initial conditions, while the internal pressure is only 2 orders of magnitude lower on average. The shock velocities in both runs are within the same order of magnitude. The speed of sound squared is proportional to the internal temperature, which is roughly 10^7K in both simulations. It is difficult to judge the width of the shell due to the existing turbulence in the Thornton WindCool run. P_{internal} , ρ_{external} and v_{shock} hence all seem to balance each other to a reasonable degree, plus the power of $1/4$ will act to damp any differences. Hence the growth rate is dominated by the width of the shell, which is difficult to measure precisely in these figures, and the speed of sound, which is very similar in both simulations. I hence estimate that $\tau_v \equiv 1/\Gamma_v$ of $\sim 10^4\text{--}10^5\text{yr}$ in each simulation. Hence the differing growth rates of Vishniac instabilities are unlikely to be the cause of the larger instabilities observed on the surface of the wind bubble in the Thornton simulation, though they are likely seeds for further instability owing to their fast growth rate relative to the length of the simulation. Further, since Vishniac instabilities are in fact overstabilities that create finger-like structures (see, e.g. Mac Low & Norman (1993)), they do not cause in of themselves highly turbulent velocity fields as seen in the T1998 run figures. However, as suggested by Kahn (1980), the shear forces caused by these fingers can generate Kelvin-Helmholtz instabilities, which we turn our attention to now.

Following Agertz et al. (2007), the growth rate of Kelvin-Helmholtz instabilities may be approximated by:

$$\Gamma_{\text{KH}} \equiv \frac{1}{\tau_{\text{KH}}} \simeq \frac{v_{\text{shock}}}{\lambda} \frac{(\rho_1 \rho_2)^{1/2}}{\rho_1 + \rho_2} \simeq \frac{v_{\text{shock}}}{\lambda} \frac{1}{\chi^{1/2}} \quad (5.5)$$

where λ is the wavelength of the instability, ρ_1 , ρ_2 are the densities either side of the interface and χ is the density contrast. The density difference between the bubble and the external medium spans 4 orders of magnitude in the T1998 simulation and 6 in the Cloud simulation (ignoring the small overdensity in the shell), while the speed of the expansion of the bubble in the T1998 simulation is $\simeq 5$ times faster than in the Cloud simulation (which travels 30pc rather than 6pc over the first 5Myr). Hence the growth rate for Kelvin-Helmholtz instabilities should be fairly similar in both runs, though the Cloud simulation should exhibit slightly a larger growth rate. In the T1998 simulation, an instability of wavelength 1pc would take $\sim 16\text{Myr}$ to grow assuming that the bubble has travelled 30pc in 5Myr. Hence this is unlikely to be the cause

of instability growth in either simulation, though a more careful analysis will be needed to confirm this for all parts of the bubble.

Finally, following Sharp (1984), the growth rate of Rayleigh-Taylor instabilities is given by:

$$\Gamma_{\text{RT}} \equiv \frac{1}{\tau_{\text{RT}}} \simeq \sqrt{\frac{2\pi}{\lambda} g \frac{\rho_1 - \rho_2}{\rho_1 + \rho_2}} \quad (5.6)$$

where g is the local acceleration of the flow, i.e. the acceleration of the shock with respect to the uniform medium, or indeed any two layers in the wind bubble. The vertical height of the instability is also governed by this growth rate. Using the internal pressure, bubble mass and surface area of the bubble to estimate the acceleration of the wind with respect to the external medium, I find a value of $\sim 5 \times 10^{-7} \text{cm/s}^2$. For a perturbation of wavelength 1pc and $\frac{\rho_1 - \rho_2}{\rho_1 + \rho_2} \simeq 1$, I obtain a growth timescale of $10^4 - 10^5 \text{yr}$. By contrast, the Cloud simulation sweeps up more mass and hence the bubble mass is ~ 2 orders of magnitude higher. The internal pressure is also 2 orders of magnitude higher. However, the bubble radius is $\simeq 5$ times smaller, resulting in a lower shock acceleration and hence slower instability growth rate. This finding, as well as the geometry of the instabilities, is likely to be the cause of the difference in the growth rate and radial height of the instabilities observed in both simulations. It is clear, however, that more careful analysis than these rough order-of-magnitude calculations must be done in order to confirm these results.

In general, Vishniac instabilities dominate in dense environments where cooling is more efficient, while Rayleigh-Taylor instabilities dominate where the shock acceleration is greater. Note also that in the Kelvin-Helmholtz and Rayleigh-Taylor instabilities, the growth rate is higher for smaller wavelengths. However, the amplitude of the instabilities is smaller for smaller wavelengths. The precise function governing scales of instabilities in a given problem is complex and can depend on aspects such as fluid viscosity. Magnetic fields can also act to damp out small-scale instabilities, though I cannot comment on this using the current simulations, which do not include magnetic fields.

In figure 5.24 I plot the results for the Cloud simulation 1Myr after the supernova event occurs inside the wind bubble (as stated before, the T1998 simulation has not reached this epoch yet). I now compare this to figure 5.12, which shows the same data for the SNCool simulation (i.e. identical, except for the inclusion of the pre-supernova stellar wind). Both snapshots exhibit broadly similar features, though the internal structure in the WindCool simulation is more complex. Note that the stellar wind over the lifetime of the

star produces a total energy that is $\sim 0.1\%$ of the energy released by the supernova into the CSM (see figure ??). Despite this large difference in energetics, I do find a small difference in the structure of the supernova remnant. The maximum radius of the remnant in the wind bubble simulation, defined by the extent of the spokes of the Vishniac instabilities in the shock, is 23-24pc in the WindCool run as opposed to 20pc in the SNCool run (the radius of the hot bubble is 10-11pc as opposed to 7-8pc). Further, the temperature towards the edge of the hot bubble is 10^6K as opposed to 10^5K . In both cases I find negligible radial velocities, though as per the pre-supernova wind bubble simulations, the centre of the bubble contains highly turbulent flows. It is hence apparent that, at least in the higher density simulations, the wind bubble plays a crucial role in scouring out a low-density bubble for the supernova to expand into, and hence allows it to expand faster and lose less energy to radiative cooling. It is expected that in lower densities this effect is less marked, though confirmation of this, as well as a quantitative analysis of the problem, awaits analysis of the T1998 simulation outputs and a more comprehensive grid of simulations at various densities and metallicities.

5.4.3 Stellar Photoionisation

In the previous section, I considered the effect of stellar winds on the evolution of the supernova remnant. However, I did not discuss the effects of photoionisation from the star on either the evolution of the wind bubble or the supernova remnant. Photoionisation from stars acts to heat the gas around them, which in turn affects the evolution of star-forming regions, as first parameterised by Strömgren (1939); Oort & Spitzer (1955); Kahn (1954). The dominant source of feedback from stellar photoionisation inside galaxies is the effect of OB associations on star-forming regions. This is because their short lifetimes mean that they spend the majority of their time in or close to star-forming regions, and their high luminosity provides a significant source of energy for photoionisation (Lejeune & Schaerer, 2001; Marigo & Girardi, 2007). Again, I do not consider this in these simulations, though in future work I will include a model of photoionisation from a population of stars as well as from the $15M_{\odot}$ star on its own.

In this section I present first results from simulations containing winds and radiative transfer. These simulations are more computationally expensive than simulations containing simply radiative cooling, and hence they are not yet complete. Simulations containing just radiative transfer and supernovae are run, though these are of limited interest as they have not yet reached the supernova event. Further, the photoion-

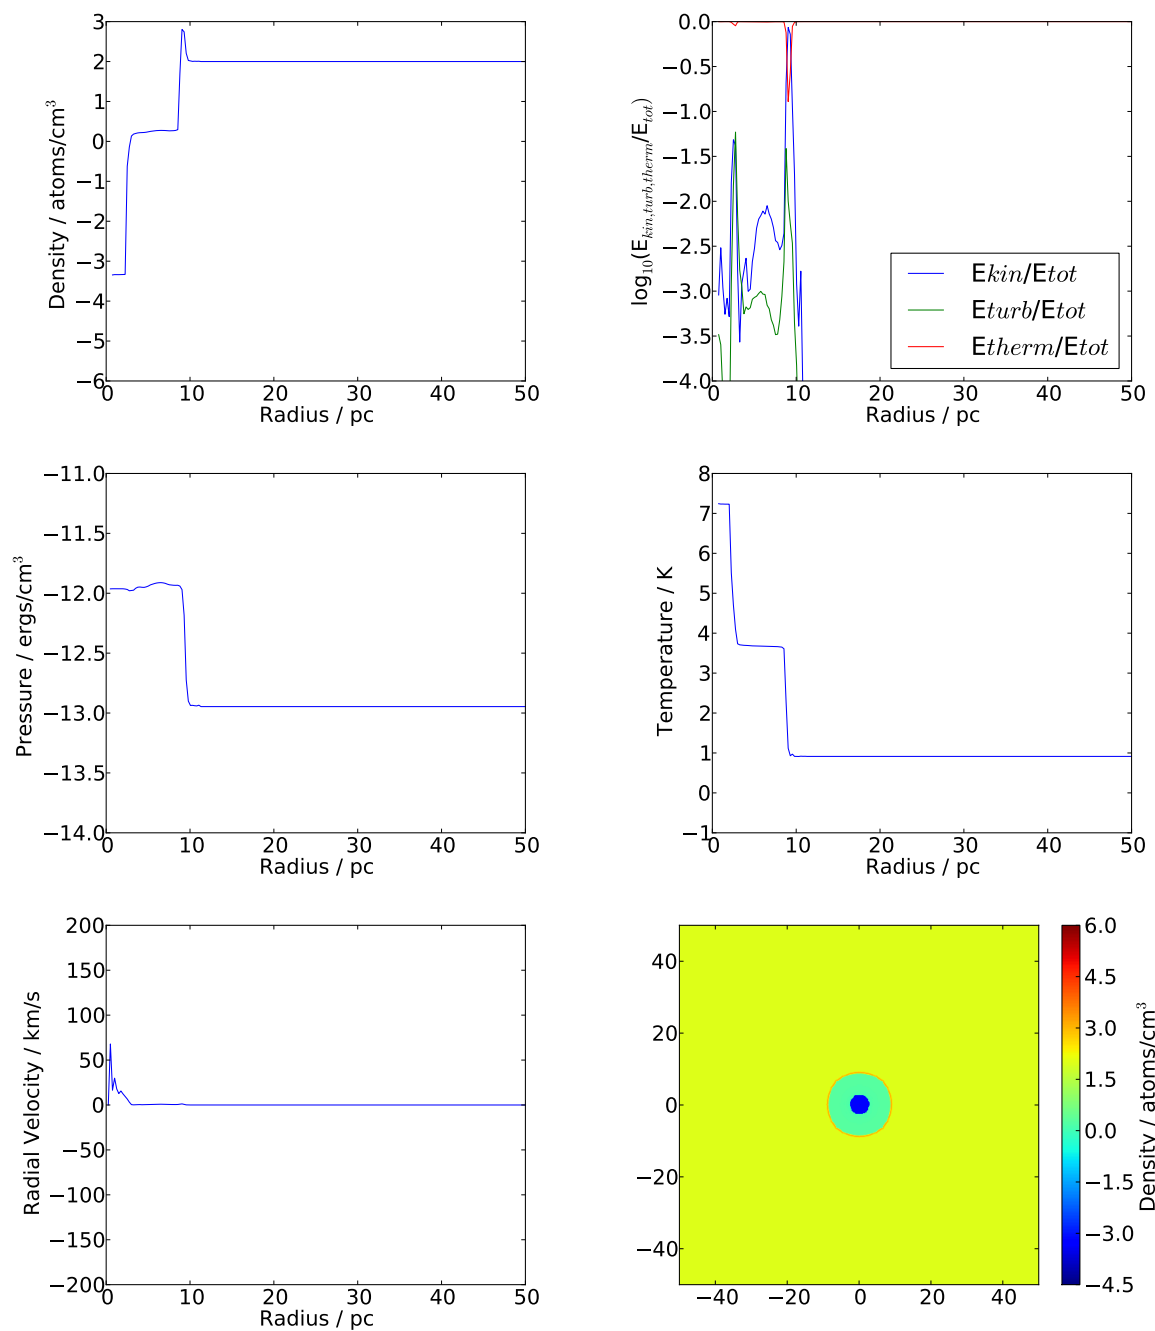


Figure 5.25: As in figure 5.26 but at 5.35 Myr in the Cloud WindRT simulation. At this stage, more instability in the contact surface between the shocked wind and the photoionised ISM is seen, as well as more turbulence inside the shocked wind bubble. The lower density and higher temperature in the photoionised gas relative to the external medium allows the wind bubble to evolve as if the external medium were at that density and temperature. The photoionised bubble is also forming a momentum-driven shell at its contact surface with the ISM.

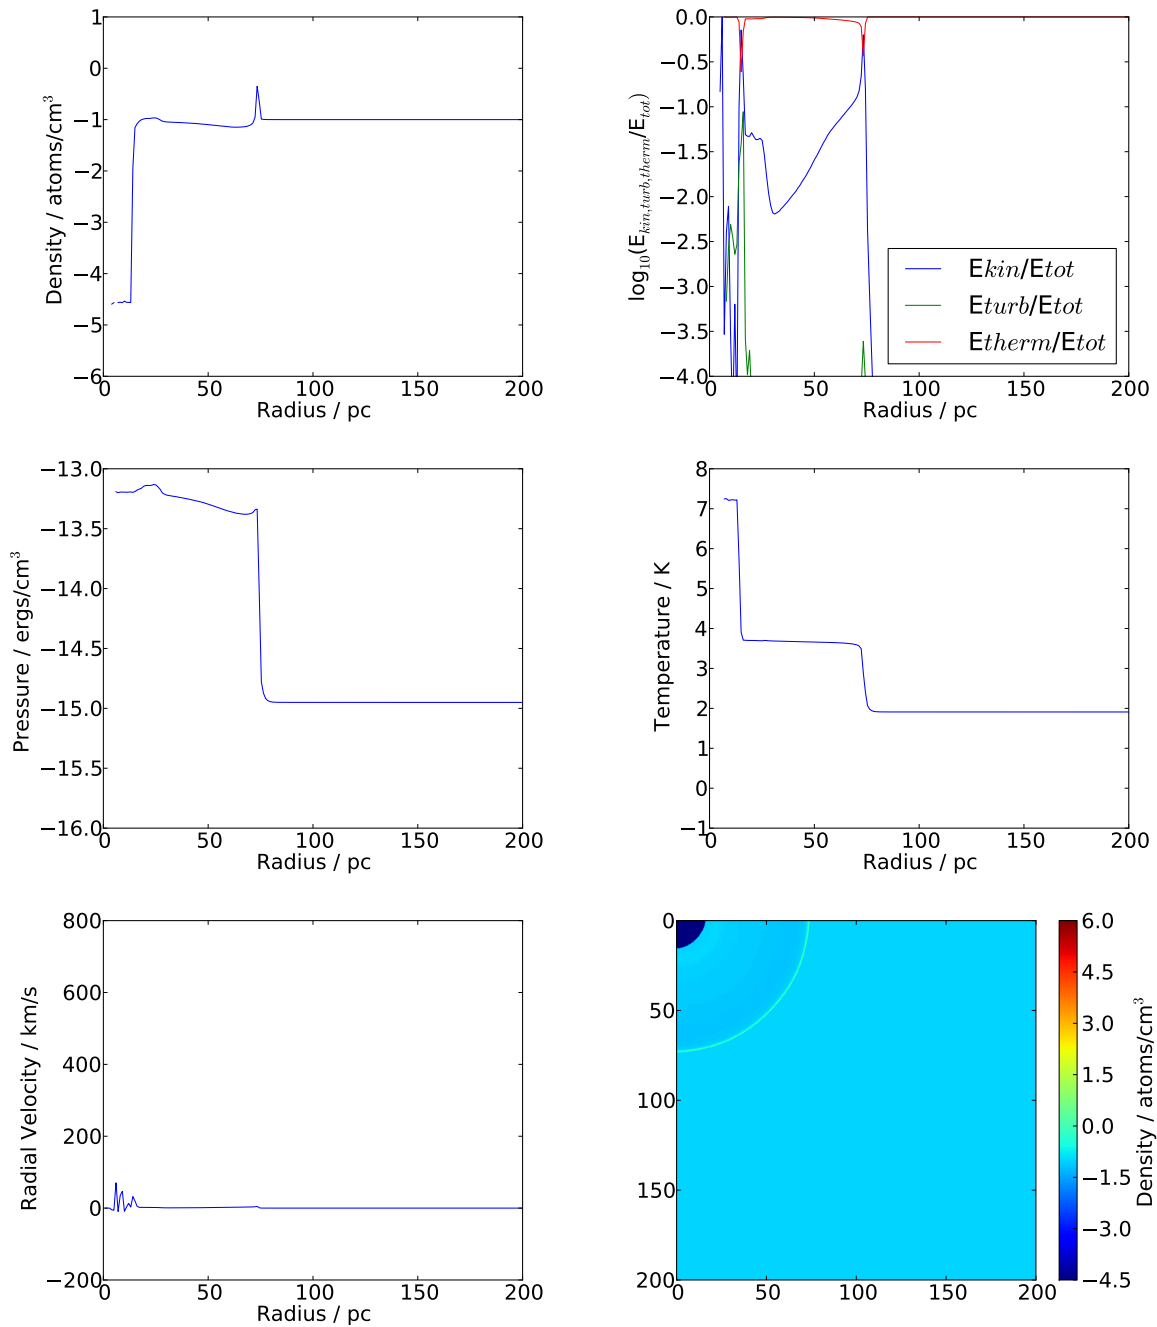


Figure 5.26: Physical properties of the wind bubble by radius at 2.55 Myr after the formation of a $15M_{\odot}$ star for the T1998 WindRT simulation. From top left to bottom right I plot density, energy partition, pressure, temperature, radial velocity against radius, and a 2D slice through the centre of the supernova in density. Energy partition is measured as the fraction of the total energy in a radial bin found in kinetic, thermal and turbulent energy. “Kinetic energy” is found by measuring all radially-aligned velocities. “Turbulent energy” is approximated by measuring all non-radial velocities. The extra photoionised shocked gas is seen, as described in figure 5.1, with the shocked gas still in the quasi-adiabatic thermally-dominated phase. Both the photoionisation and the wind shock appear to be roughly spherical.

isation bubble is very similar in simulations with and without stellar winds. I do not perform quantitative analysis to any great degree due to latest simulation outputs capturing only the early stage of the evolution of these bubbles. Rather, I simply present a snapshot of the progress of the simulations to illustrate the effect that early photoionisation from the star has on the simulations.

In figure 5.25 I plot the WindRT data for the same radial profiles as the WindCool data in figure 5.20, also at 5Myr, for the Cloud initial conditions. There are various quantitative and qualitative differences between the two figures. The most obvious is a photoionised bubble at just under 10^4K and 1 atom/cm^3 that has roughly twice the radius of the wind bubble. Note that the speed of light in this simulation is reduced (see chapter 2) in order to prevent the Courant condition from slowing down the simulation to a crawl; however, tests given in Iliev et al. (2006, 2009) have shown that this does not affect the expansion of the photoionised bubble if the medium is dense enough. This bubble then represents the extent of the star's influence on the external medium, and hence it is the edge of this bubble and not the wind bubble that hosts the momentum-driven shell, though a wide peak in kinetic energy is seen at the interface between the wind and the photoionised gas. This peak exhibits stronger Vishniac instabilities than in figure 5.20 due to the lower external density (see equation 5.4). These results indicate that even the photoionisation from a $15M_{\odot}$ star can have an impact on the conditions around a star-forming region by the time the star goes supernova 10Myr after this snapshot is taken.

The data in for the T1998 initial conditions shown in figure 5.26 is less comparable to figure 5.16 as it has only reached 2.55Myr, but there are some striking differences between the figures. Firstly, the photoionised bubble does create a thin overdensity at the shell of shocked interstellar gas, though inside the photoionised bubble the density does not change much from the background level of $n_{\text{H}} = 0.1\text{ atoms/cm}^3$. Secondly, there are now two shells where the kinetic energy partition peaks, one at the interface of the wind with the photoionised gas and one at the interface of the photoionised gas and the external medium. As in figure 5.25, this photoionised gas is heated to just under 10^4K , an order of magnitude above the background temperature. The radial velocity inside the bubble is also small, though this may change over time. In particular, no instabilities are observed on the surface of the wind bubble, and very little turbulence is seen, though again these may grow over time.

It is clear from these plots that photoionisation has a noticeable effect on the ISM prior to the launching

of stellar winds or supernova blastwaves, and hence is an important aspect to monitor as the simulations evolve.

5.5 Discussion and Future Work

In this chapter I have presented the preliminary results from a study of supernovae in uniform media, with the goal of improving on the work of Thornton et al. (1998), who explored empirical parameterisations of the evolution of supernova remnants in 1D simulations. I performed a similar study in 3D with the goal of exploring the effect of instabilities and turbulence, which do not exist in 1D, on the evolution of the SN remnant and its interaction with the ISM. Compared to these authors, I also model the pre-supernova phase of stellar evolution, accounting for stellar winds and ionising photon radiative transfer. I also assess their impact on the evolution of the subsequent supernova remnant.

In our most comparable model to the Thornton et al. (1998) work, I found good agreement in our results, though a full parameter study in the density, temperature and metallicity of the medium into which the supernova explodes has yet to be completed. I documented the creation of Vishniac instabilities (Vishniac, 1983) on the surface of the supernova remnants, which will play an important role in shaping the evolution of supernovae in more complex media when we study them in future simulations. I gave parameterisations for the evolution of the remnant mass, energy and radius for the initial conditions bearing closest resemblance to the test case given in Thornton et al. (1998).

I then presented results from experiments including stellar wind, and noted small but noticeable differences between supernova simulations that contain stellar winds from those that do not. I explored the creation of turbulence on the wind bubble surface and determined that the most likely cause of these is Rayleigh-Taylor instabilities, though they are likely to be influenced by Vishniac instabilities that occur on a similar timescale. More interesting is the role of photoionisation, since this noticeably affects the structure of the medium even in the first few Myr of the evolution of the star.

Further effects will also be considered as I expand the stable of simulations run as part of this study. The role of self-gravity has yet to be considered, though it is included as a term in the continuity equation for momentum in Thornton et al. (1998). Vishniac (1983) also notes that Vishniac instabilities can be seeded by gravitational as well as thermal fluctuations, and hence it may be important to consider the effect that

self-gravity has on the evolution of instabilities on the wind and supernova shocks.

The role of magnetic fields is more complicated, and more poorly constrained. Gull (1973) parameterises expected magnetic field strengths in his 1D simulations, though finding realistic initial conditions for magnetic fields in the ISM is still the subject of research. It is also worth noting that significant technical challenges exist in merging the MHD versions of RAMSES with the RAMSES-RT. Finally, the ISM is not a uniform medium with a single source of energy. Modelling the turbulent and non-uniform density field is important for understanding the evolution of supernova blastwaves into the ISM, as discussed in section 5.1. This will be the subject of a future series of experiments, once I have constrained the parameterisations for supernovae in a uniform density field.

Finally, it may be important to expand our study beyond a single $15M_{\odot}$ star. Many supernovae observed are Type Ia supernovae, which are thought to be formed not by core collapse but from thermonuclear detonation of a white dwarf. The assumption that winds escape a star in a spherically-symmetric pattern has been shown in studies presented in section 5.1 to be untrue in stars that exhibit significant rotation or non-spherical internal structures. It may be desirable to implement more sophisticated models of the stars themselves in order to better understand the full impact that supernovae of all shapes and sizes have on the interstellar medium once I have better characterised the evolution of individual 3D core-collapse supernovae remnants of $15M_{\odot}$ stars in a variety of homogeneous density fields.

Chapter 6

Conclusions

In this thesis, I have explored the subject of stellar feedback in dwarf galaxies using hydrodynamic simulations in RAMSES (Teyssier, 2002). In the first part of the thesis, I described the numerical and analytical techniques used to simulate galaxies and supernovae, and then interpreted these results. I presented some numerical, analysis and visualisation projects completed in the course of this thesis that aim to aid future work with numerical simulations of stellar feedback in galaxies.

I then presented in chapter 4 a study of the satellite galaxies of a suite of simulations of a Milky Way-like galaxy, as submitted for publication in MNRAS (Geen et al., 2012). I first discussed the effects of supernova feedback and reionisation at high redshift on these objects in sub-parsec resolution simulations. This is done because studies find that much of the stellar mass of observed dwarfs is in place at high redshift. I found that negative feedback from supernova feedback due to gas loss and positive feedback from metal enrichment and blast compression balance each other, with negative feedback being more important in lower mass halos where the weak potential allows for stronger mass loss, and vice-versa for larger halos. However, I also found that neither supernova feedback nor reionisation have a particularly significant effect on the star formation rates of high-redshift dwarf galaxies. As I followed the simulations to $z=0$, I found that the baryon physics included in the simulation has the biggest impact on the satellite population observed due to the effect that it has on the density profile of the host halo; a more centrally concentrated profile increases dynamical friction and tidal stripping in the centre of the halo. It thus appears that supernova feedback, while of consequence, is of secondary importance in understanding the evolution of satellite galaxies of

Milky Way-like galaxies. The models of feedback included in this simulation are something of a first pass, however - both the models of supernova feedback and photoionisation can be improved with the numerical spatial resolution now available to us, which is a problem I begin to tackle in the second main study in this thesis.

In chapter 5 I presented first results of a study that aims to improve our understanding of the effect of stellar feedback on the ISM by simulating isolated stars evolving in a uniform medium in 3D. This builds upon the work of Thornton et al. (1998), who use 1D simulations of supernovae to create a set of empirical models for the structure and energetics of supernova remnants in media of various densities and metallicities. I noted that my results, which aim to reproduce the case study in Thornton et al. (1998), agree with their findings, though strong instabilities are observed on the surface of the supernova shock that cannot otherwise form in 1D simulations and that affect the radial evolution of the remnant. I then presented a comparison of two simulations of stellar winds, and in one of these simulations the effect it has on the evolution of the subsequent supernova remnant. The simulation produces a low-density stellar wind bubble in a high-density medium, allowing the supernova shock to expand faster than it would otherwise do in a high-density medium. However, I cannot as yet say whether this result holds for lower-density media. Finally, I presented in-progress snapshots of a simulation with photoionisation from the star using the radiative transfer module in RAMSES (Rosdahl & Blaizot, 2012a,b), and found that even after a few Myr it has a strong impact on the temperature and density of the ISM surrounding the star. In this way, I this work encompasses a range of parameters that effect the expansion of supernova remnants into the ISM.

As numerical simulations of galaxy formation increase in complexity and resolution in line with improvements in the power of computer technology, so the sophistication of models of the physical processes inside them can and must increase. It is clear from this work and other works mentioned in this thesis that the effect of stellar feedback on the gas around them is complex and dramatic; even if, as this work suggests, the global impact on the star formation rates of dwarf galaxies is small, stellar winds, photoionisation and supernovae do have strong and varied effects on the properties of the ISM. If we are to reproduce the detailed observable characteristics of galaxies of all sizes, detailed simulations of stellar feedback in galaxies will be crucial in achieving this. As we have noted in earlier sections, dwarf galaxies remain ideal test candidates for this due to their relatively weak gravitational potentials and small spatial extent, allowing for simulations

with higher resolution at lower cost, and with more dramatic results.

This thesis does not aim to conclusively resolve the issues surrounding stellar feedback. Rather, it seeks to answer specific questions concerning dwarf galaxy formation and provides motivations for future studies of feedback. Stellar feedback is a complex, non-linear problem, and producing a reliable model for its impact on galaxies eludes the community (Scannapieco et al., 2011). In coming work, I hope to expand the results of this thesis to provide new stepping stones towards a global model for supernova feedback in the gas inside galaxies. The NUT suite is a vital step towards achieving this goal, providing a framework in which individual supernovae can be followed in a cosmological galaxy formation simulation with enough spatial resolution to generate a multiphase ISM without resorting to sub-grid prescriptions for the different phases of interstellar gas. Meanwhile, the parameter study of individual stars in idealised simulations introduced in this thesis will motivate better models for these individual supernovae in cosmological simulations, which will be able to determine more precisely the effect of stellar feedback on the interstellar gas, and eventually on star formation. The success of galaxy formation simulations of stellar feedback depends on a detailed understanding of how stars evolve into the interstellar gas, and this thesis presents work that I hope will help to advance this understanding.

*If these simulations haven't ended
Think but this, and all is mended,
Just sit back and slumber here
And wait for data to appear.*

Bibliography

- Abel T., 2011, *Monthly Notices of the Royal Astronomical Society*, 413, 271
- Agertz O., Moore B., Stadel J., Potter D., Miniati F., Read J., Mayer L., Gawryszczak A., Kravtsov A., Nordlund A. k., 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
- Assousa G. E., Herbst W., 1980, In: *Giant molecular clouds in the Galaxy; Proceedings of the Third Gregynog Astrophysics Workshop*, pp 275–284
- Aubert D., Pichon C., Colombi S., 2004, *Monthly Notices of the Royal Astronomical Society*, 352, 376
- Aubert D., Teyssier R., 2008, *Monthly Notices of the Royal Astronomical Society*, 387, 295
- Barnes J. E., Hernquist L. E., 1992, *Annual Review of Astronomy and Astrophysics*, 30, 705
- Bate M. R., 2012, *Monthly Notices of the Royal Astronomical Society*, 419, 3115
- Battaglia G., Helmi A., Morrison H., Harding P., Olszewski E. W., Mateo M., Freeman K. C., Norris J., Shectman S. A., 2005, *Monthly Notices of the Royal Astronomical Society*, 364, 433
- Baugh C. M., Cole S., Frenk C. S., Lacey C. G., 1998, *The Astrophysical Journal*, 498, 504
- Bekki K., Chiba M., 2005, *Monthly Notices of the Royal Astronomical Society*, 356, 680
- Bekki K., Stanimirović S., 2009, *Monthly Notices of the Royal Astronomical Society*, 395, 342
- Belokurov V., Walker M. G., Evans N. W., Gilmore G., Irwin M. J., Just D., Koposov S., Mateo M., Olszewski E., Watkins L., Wyrzykowski L., 2010, *The Astrophysical Journal Letters*, 712, L103
- Belokurov V., Walker M. G., Evans N. W., Gilmore G., Irwin M. J., Mateo M., Mayer L., Olszewski E., Bechtold J., Pickering T., 2009, *Monthly Notices of the Royal Astronomical Society*, 397, 1748
- Benson A. J., Bower R. G., Frenk C. S., Lacey C. G., Baugh C. M., Cole S., 2003, *The Astrophysical Journal*, 599, 38
- Benson A. J., Lacey C., Baugh C., Cole S., Frenk C., 2002, *Extragalactic Gas at Low Redshift*, 254, 354
- Bertschinger E., 2001, *The Astrophysical Journal Supplement Series*, 137, 1
- Binney J., Tremaine S., 2008, *Galactic Dynamics: Second Edition*
- Blumenthal G. R., Faber S. M., Primack J. R., Rees M. J., 1984, *Nature*, 311, 517

- Bodenheimer P., Laughlin G. P., Rózycka M., Yorke H. W., 2007, *Numerical Methods in Astrophysics: An Introduction*. Part of the Series in Astronomy and Astrophysics. Edited by Peter Bodenheimer
- Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2011, *Monthly Notices of the Royal Astronomical Society: Letters*, 415, L40
- Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2012, *Monthly Notices of the Royal Astronomical Society*, -1, no
- Brandl B. R., 2005, *Starbursts: From 30 Doradus to Lyman Break Galaxies*, 329, 49
- Bruzual G., Charlot S., 2003, *Monthly Notices of the Royal Astronomical Society*, 344, 1000
- Bullock J. S., Kravtsov A. V., Weinberg D. H., 2000, *The Astrophysical Journal*, 539, 517
- Burrows C. J., Krist J., Hester J. J., Sahai R., Trauger J. T., Stapelfeldt K. R., Gallagher J. S., Ballester G. E., Casertano S., Clarke J. T., Crisp D., Evans R. W., Griffiths R. E., Hoessel J. G., Holtzman J. A., Mould J. R., Scowen P. A., Watson A. M., Westphal J. A., 1995, *The Astrophysical Journal*, 452, 680
- Busha M. T., Alvarez M. A., Wechsler R. H., Abel T., Strigari L. E., 2010, *The Astrophysical Journal*, 710, 408
- Calzetti D., 2001, *Publications of the Astronomical Society of the Pacific*, 113, 1449
- Calzetti D., Armus L., Bohlin R. C., Kinney A. L., Koornneef J., Storchi-Bergmann T., 2000, *The Astrophysical Journal*, 533, 682
- Cen R., 1992, *The Astrophysical Journal Supplement Series*, 78, 341
- Cen R., McDonald P., Trac H., Loeb A., 2009, *The Astrophysical Journal Letters*, 706, L164
- Cen R., Ostriker J. P., 1992, *Astrophysical Journal*, 399, L113
- Ceverino D., Klypin A., 2009, *The Astrophysical Journal*, 695, 292
- Cha S.-H., Inutsuka S.-i., Nayakshin S., 2010, *Monthly Notices of the Royal Astronomical Society*, 403, 1165
- Chevalier R. A., 1977, In: *Annual review of astronomy and astrophysics*. Volume 15. (A78-16576 04-90) Palo Alto, 15, 175
- Chevalier R. A., 1982, *The Astrophysical Journal*, 258, 790
- Chevalier R. A., Klein R. I., 1978, *Astrophysical Journal*, 219, 994
- Chita S. M., Langer N., van Marle A. J., García-Segura G., Heger A., 2008, *Astronomy and Astrophysics*, 488, L37
- Chu Y.-H., Kennicutt R. C., 1994, *The Astrophysical Journal*, 425, 720
- Clemens M. S., Bressan A., Panuzzo P., Rampazzo R., Silva L., Buson L., Granato G. L., 2009, *Monthly Notices of the Royal Astronomical Society*, 392, 982

- Clyne J., Mininni P., Norton A., Rast M., 2007, *New Journal of Physics*, 9
- Collela P., 1990, *J. Chem. Phys.*, 87, 171
- Collins M. L. M., Chapman S. C., Irwin M. J., Martin N. F., Ibata R. A., Zucker D. B., Blain A., Ferguson A. M. N., Lewis G. F., McConnachie A. W., Peñarrubia J., 2010, *Monthly Notices of the Royal Astronomical Society*, 407, 2411
- Collins M. L. M., Chapman S. C., Rich R. M., Irwin M. J., Peñarrubia J., Ibata R. A., Arimoto N., Brooks A. M., Ferguson A. M. N., Lewis G. F., McConnachie A. W., Venn K., 2011, *Monthly Notices of the Royal Astronomical Society*, 417, 1170
- Courty S., Teyssier R., 2012
- Crank J., Nicolson P., 1996, *Advances in Computational Mathematics*, 6, 207
- Dalgarno A., McCray R. A., 1972, *Annual Review of Astronomy and Astrophysics*, 10, 375
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *The Astrophysical Journal*, 292, 371
- Deacon J., 2005, *Object-oriented analysis and design : a pragmatic approach*. Pearson Addison-Wesley, Harlow England [u.a.]
- Dekel A., Silk J., 1986, *Astrophysical Journal*, 303, 39
- Dekel A., Woo J., 2003, *Monthly Notices of the Royal Astronomical Society*, 344, 1131
- Dgani R., van Buren D., Noriega-Crespo A., 1996, *The Astrophysical Journal*, 461, 372
- Di Cintio A., Knebe A., Libeskind N. I., Yepes G., Gottlöber S., Hoffman Y., 2011, *Monthly Notices of the Royal Astronomical Society: Letters*, -1, no
- Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604
- Diemand J., Kuhlen M., Madau P., Zemp M., Moore B., Potter D., Stadel J., 2008, *Nature*, 454, 735
- Dubois Y., Devriendt J., Slyz A., Teyssier R., 2012, *Monthly Notices of the Royal Astronomical Society*, 420, 2662
- Dubois Y., Teyssier R., 2008a, *Astronomy and Astrophysics*, 482, L13
- Dubois Y., Teyssier R., 2008b, *Astronomy and Astrophysics*, 477, 79
- Dunkley J., Komatsu E., Nolte M. R., Spergel D. N., Larson D., Hinshaw G., Page L., Bennett C. L., Gold B., Jarosik N., Weiland J. L., Halpern M., Hill R. S., Kogut A., Limon M., Meyer S. S., Tucker G. S., Wollack E., Wright E. L., 2009, *The Astrophysical Journal Supplement*, 180, 306
- Dwarkadas V. V., 2007, *The Astrophysical Journal*, 667, 226
- Dwarkadas V. V., 2011, eprint arXiv:1108.3576
- Dwarkadas V. V., Balick B., 1998, *The Astrophysical Journal*, 497, 267
- Dwarkadas V. V., Owocki S. P., 2002, *The Astrophysical Journal*, 581, 1337

- Eisenstein D. J., Hu W., 1998, *The Astrophysical Journal*, 496, 605
- Feldmann R., Mayer L., Carollo C. M., 2008, *The Astrophysical Journal*, 684, 1062
- Ferguson A. M. N., Gallagher J. S., Wyse R. F. G., 2000, *The Astronomical Journal*, 120, 821
- Ferland G., Korista K., Verner D., Ferguson J., Kingdon J., Verner E., 1998, *Publications of the Astronomical Society of the Pacific*, 110, 761
- Flores R. A., Primack J. R., 1994, *The Astrophysical Journal*, 427, L1
- Fraschetti F., Teyssier R., Ballet J., Decourchelle A., 2010, *Astronomy and Astrophysics*, 515, A104
- Gaetz T. J., Edgar R. J., Chevalier R. A., 1988, *The Astrophysical Journal*, 329, 927
- Gal-Yam A., Fox D. B., Price P. A., Ofek E. O., Davis M. R., Leonard D. C., Soderberg A. M., Schmidt B. P., 2006, *Nature*, 444, 1053
- Gamma E., 1995, *Design patterns : elements of reusable object-oriented software*. Addison-Wesley, Reading Mass.
- García-Segura G., 2009, *Magnetic Fields in the Universe II: From Laboratory and Stars to the Primordial Universe* (Eds. A. Esquivel, 36, 67
- Garcia-Segura G., Langer N., Mac Low M.-M., 1996, *Astronomy and Astrophysics*
- Garcia-Segura G., Langer N., Rozyczka M., Franco J., 1999, *The Astrophysical Journal*, 517, 767
- Geen S., Slyz A., Devriendt J., 2012, *Monthly Notices of the Royal Astronomical Society*
- Genel S., Genzel R., Bouché N., Sternberg A., Naab T., Schreiber N. M. F., Shapiro K. L., Tacconi L. J., Lutz D., Cresci G., Buschkamp P., Davies R. I., Hicks E. K. S., 2008, *The Astrophysical Journal*, 688, 789
- Gerritsen J. P. E., Icke V., 1997, *Astronomy and Astrophysics*
- Gnedin N. Y., 2000a, *The Astrophysical Journal*, 542, 535
- Gnedin N. Y., 2000b, *The Astrophysical Journal*, 535
- Gnedin N. Y., Abel T., 2001, *New Astronomy*, 6, 437
- Gnedin N. Y., Kravtsov A. V., 2006, *The Astrophysical Journal*, 645, 1054
- Goodwin S., Pearce F., Thomas P., 2000, eprint arXiv, p. 1180
- Grebel E. K., Gallagher III J. S., Harbeck D., 2003, *The Astronomical Journal*, 125, 1926
- Gull S. F., 1973, *Monthly Notices of the Royal Astronomical Society*, 161, 47
- Gunn J. E., Gott J. R., 1972, *The Astrophysical Journal*, 176, 1
- Guo Q., White S., Li C., Boylan-Kolchin M., 2010, *Monthly Notices of the Royal Astronomical Society*, 404, 1111

- Haardt F., Madau P., 1996, *Astrophysical Journal* v.461, 461, 20
- Hahn O., Teyssier R., Carollo C. M., 2010, *Monthly Notices of the Royal Astronomical Society*, pp no–no
- Hatton S., Devriendt J. E. G., Ninin S., Bouchet F. R., Guiderdoni B., Vibert D., 2003, *Monthly Notices of the Royal Astronomical Society*, 343, 75
- Heger A., Fryer C. L., Woosley S. E., Langer N., Hartmann D. H., 2003, *The Astrophysical Journal*, 591, 288
- Henry M., 2010, BBC Food (http://www.bbc.co.uk/food/recipes/maryhenrysproperappl_67463)
- Hensler G., 2008, *Proceedings of the International Astronomical Union*, 4, 309
- HeßS., Springel V., 2010, *Monthly Notices of the Royal Astronomical Society*, 406, 2289
- Hibbard J. E., Bianchi L., Thilker D. A., Rich R. M., Schiminovich D., Xu C. K., Neff S. G., Seibert M., Lauger S., 2005, *The Astrophysical Journal*, 619, L87
- Hilbert D., 1891, *Mathematische Annalen*, 38, 459
- Hill A. S., Ryan Joung M., Mac Low M.-M., Benjamin R. A., Matthew Haffner L., Klingenberg C., Waagan K., 2012, *The Astrophysical Journal*, 750, 104
- Hillebrandt W., Niemeyer J. C., 2000, *Annual Review of Astronomy and Astrophysics*, 38, 191
- Hockney R. W., Eastwood J. W., 1981, *Computer Simulation Using Particles*
- Hoefl M., Yepes G., Gottlöber S., Springel V., 2006, *Monthly Notices of the Royal Astronomical Society*, 371, 401
- Hogg D. W., 1999
- Holmberg E., 1941, *The Astrophysical Journal*, 94, 385
- Huang K., Wu H., Yu H., Yan D., 2010, *International Journal for Numerical Methods in Fluids*, 9999
- Huchra J. P., Geller M. J., 1982, *The Astrophysical Journal*, 257, 423
- Iliev I. T., Ciardi B., Alvarez M. A., Maselli A., Ferrara A., Gnedin N. Y., Mellema G., Nakamoto T., Norman M. L., Razoumov A. O., Rijkhorst E.-J., Ritzerveld J., Shapiro P. R., Susa H., Umemura M., Whalen D. J., 2006, *Monthly Notices of the Royal Astronomical Society*, 371, 1057
- Iliev I. T., Whalen D., Mellema G., Ahn K., Baek S., Gnedin N. Y., Kravtsov A. V., Norman M., Raicevic M., Reynolds D. R., Sato D., Shapiro P. R., Semelin B., Smidt J., Susa H., Theuns T., Umemura M., 2009, *Monthly Notices of the Royal Astronomical Society*, 400, 1283
- Innes D. E., Giddings J. R., Falle S. A. E. G., 1987, *Monthly Notices of the Royal Astronomical Society (ISSN 0035-8711)*, 226, 67
- Irwin M. J., Collins M. L. M., 2011, private communication
- Irwin M. J., Ferguson A. M. N., Huxor A. P., Tanvir N. R., Ibata R. A., Lewis G. F., 2008, *The Astrophysical Journal*, 676, L17

- Izotov Y. I., Thuan T. X., 2008, *The Astrophysical Journal*, 687, 133
- Johnson J. L., Khochfar S., 2011
- Junk V., Walch S., Heitsch F., Burkert A., Wetzstein M., Schartmann M., Price D., 2010, eprint arXiv, 1004, 1957
- Kahn F. D., 1954, *Bulletin of the Astronomical Institutes of the Netherlands*, 12
- Kahn F. D., 1980, *Astronomy and Astrophysics*, 83, 303
- Kalirai J. S., Beaton R. L., Geha M. C., Gilbert K. M., Guhathakurta P., Kirby E. N., Majewski S. R., Ostheimer J. C., Patterson R. J., Wolf J., 2010, *The Astrophysical Journal*, 711, 671
- Karachentsev I. D., Kashibadze O. G., 2005, *Astrofizika*
- Katz N., Gunn J. E., 1991, *The Astrophysical Journal*, 377, 365
- Kauffmann G., White S. D. M., Guiderdoni B., 1993, *R.A.S. MONTHLY NOTICES* V.264
- Kawata D., Okamoto T., Cen R., Gibson B. K., 2009, eprint arXiv, 0902, 4002
- Kennicutt R. C., 1998, *Astrophysical Journal* v.498, 498, 541
- Khochfar S., Silk J., 2006, *Monthly Notices of the Royal Astronomical Society*, 370, 060620041143002
- Khochfar S., Silk J., 2009, *Monthly Notices of the Royal Astronomical Society*, 397, 506
- Kifonidis K., Plewa T., Janka H.-T., Müller E., 2003, *Astronomy and Astrophysics*, 408, 621
- Kimm T., Devriendt J., Slyz A., Pichon C., Kassir S. A., Dubois Y., 2011, eprint arXiv, 1106, 538
- Kirby E. N., Martin C. L., Finlator K., 2011a, *The Astrophysical Journal*, 742, L25
- Kirby E. N., Martin C. L., Finlator K., 2011b, p. 5
- Kitayama T., Susa H., Umemura M., Ikeuchi S., 2001, *Monthly Notices of the Royal Astronomical Society*, 326, 1353
- Klypin A. A., Shandarin S. F., 1983, *Royal Astronomical Society*, 204, 891
- Knebe A., Knollmann S. R., Muldrew S. I., Pearce F. R., Aragon-Calvo M. A., Ascasibar Y., Behroozi P. S., Ceverino D., Colombi S., Diemand J., Dolag K., Falck B. L., Fasel P., Gardner J., Gottlöber S., 2011, *Monthly Notices of the Royal Astronomical Society*
- Koposov S., Belokurov V., Evans N. W., Hewett P. C., Irwin M. J., Gilmore G., Zucker D. B., Rix H.-W., Fellhauer M., Bell E. F., Glushkova E. V., 2008, *The Astrophysical Journal*, 686, 279
- Koposov S. E., Belokurov V., Evans N. W., Gilmore G., Gieles M., Irwin M. J., Lewis G. F., Niederste-Ostholt M., Penarrubia J., Smith M. C., Bizyaev D., Malanushenko E., Malanushenko V., Schneider D. P., Wyse R. F. G., 2011
- Kovetz A., Yaron O., Prialnik D., 2009, *Monthly Notices of the Royal Astronomical Society*, 395, 1857

- Kravtsov A. V., Gnedin O. Y., Klypin A. A., 2004, *The Astrophysical Journal*, 609, 482
- Krist J., Hook R., 2004, STScI Documentation
- Kroupa P., Weidner C., 2003, *The Astrophysical Journal*, 598, 1076
- Krumholz M. R., Dekel A., McKee C. F., 2012, *The Astrophysical Journal*, 745, 69
- Krumholz M. R., Tan J. C., 2007, *The Astrophysical Journal*, 654, 304
- Kudritzki R.-P., Puls J., 2000, *Annual Review of Astronomy and Astrophysics*, 38, 613
- Kushnir D., Waxman E., Shvarts D., 2005, *The Astrophysical Journal*, 634, 407
- Larson D., Dunkley J., Hinshaw G., Komatsu E., Nolta M. R., Bennett C. L., Gold B., Halpern M., Hill R. S., Jarosik N., Kogut A., Limon M., Meyer S. S., Odegard N., Page L., Smith K. M., Spergel D. N., Tucker G. S., Weiland J. L., Wollack E., Wright E. L., 2010, eprint arXiv, 1001, 4635
- Leitherer C., Schaerer D., Goldader J. D., Delgado R. M. G., Robert C., Kune D. F., de Mello D. F., Devost D., Heckman T. M., 1999, *The Astrophysical Journal Supplement Series*, 123, 3
- Lejeune T., Schaerer D., 2001, *Astronomy and Astrophysics*, 366, 538
- Lemaster M. N., Stone J. M., Gardiner T. A., 2007, *The Astrophysical Journal*, 662, 582
- Letarte B., Chapman S. C., Collins M., Ibata R. A., Irwin M. J., Ferguson A. M. N., Lewis G. F., Martin N., McConnachie A., Tanvir N., 2009, *Monthly Notices of the Royal Astronomical Society*, 400, 1472
- Libeskind N. I., Yepes G., Knebe A., Gottlöber S., Hoffman Y., Knollmann S. R., 2010, *Monthly Notices of the Royal Astronomical Society*, 401, 1889
- Lin C. C., Shu F. H., 1964, *The Astrophysical Journal*, 140, 646
- Lloyd H. M., O'Brien T. J., Kahn F. D., 1995, *Monthly Notices of the Royal Astronomical Society*, 273, L19
- Lovell M. R., Eke V., Frenk C. S., Gao L., Jenkins A., Theuns T., Wang J., White S. D. M., Boyarsky A., Ruchayskiy O., 2012, *Monthly Notices of the Royal Astronomical Society*, 420, 2318
- Low M.-M. M., Ferrara A., 1999, *The Astrophysical Journal*, 513, 142
- Mac Low M.-M., Norman M. L., 1993, *The Astrophysical Journal*, 407, 207
- Macciò A. V., Kang X., Fontanot F., Somerville R. S., Kogosov S., Monaco P., 2010, *Monthly Notices of the Royal Astronomical Society*, 402, 1995
- Macciò A. V., Sideris I., Miranda M., Moore B., Jesseit R., 2007, eprint arXiv:0704.3078
- McConnachie A. W., 2012, eprint arXiv:1204.1562
- McConnachie A. W., Huxor A., Martin N. F., Irwin M. J., Chapman S. C., Fahlman G., Ferguson A. M. N., Ibata R. A., Lewis G. F., Richer H., Tanvir N. R., 2008, *The Astrophysical Journal*, 688, 1009
- Machacek M. E., Bryan G. L., Abel T., 2001, *The Astrophysical Journal*, 548, 509

- Madau P., Diemand J., Kuhlen M., 2008, *The Astrophysical Journal*, 679, 1260
- Madau P., Kuhlen M., Diemand J., Moore B., Zemp M., Potter D., Stadel J., 2008, *The Astrophysical Journal*, 689, L41
- Malin D. F., Carter D., 1983, *The Astrophysical Journal*, 274, 534
- Malumuth E. M., Heap S. R., 1994, *The Astronomical Journal*, 107, 1054
- Marigo P., Girardi L., 2007, *From Stars to Galaxies: Building the Pieces to Build Up the Universe. ASP Conference Series*, 374, 33
- Marigo P., Girardi L., Bressan A., Groenewegen M. A. T., Silva L., Granato G. L., 2008, *Astronomy and Astrophysics*, 482, 883
- Martin N. F., McConnachie A. W., Irwin M., Widrow L. M., Ferguson A. M. N., Ibata R. A., Dubinski J., Babul A., Chapman S., Fardal M., Lewis G. F., Navarro J., Rich R. M., 2009, *The Astrophysical Journal*, 705, 758
- Mashchenko S., Wadsley J., Couchman H. M. P., 2008, *Science*, 319, 174
- Mateo M. L., 1998, *Annual Review of Astronomy and Astrophysics*, 36, 435
- Matzner C. D., McKee C. F., 1999, *The Astrophysical Journal*, 510, 379
- Mesinger A., 2009, eprint arXiv, 0910, 4161
- Minkowski R., 1940, *Publications of the Astronomical Society of the Pacific*, 52, 206
- Mitchell N. L., McCarthy I. G., Bower R. G., Theuns T., Crain R. A., 2009, *Monthly Notices of the Royal Astronomical Society*, 395, 180
- Mohamed S., Mackey J., Langer N., 2011, eprint arXiv:1109.1555
- Moore B., Diemand J., Madau P., Zemp M., Stadel J., 2006, *Monthly Notices of the Royal Astronomical Society*, 368, 563
- Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, *The Astrophysical Journal*, 524, L19
- Morris T., Podsiadlowski P., 2009, *Monthly Notices of the Royal Astronomical Society*, 399, 515
- Morrison H. L., Harding P., Hurley-Keller D., Jacoby G., 2003, *The Astrophysical Journal*, 596, L183
- Muñoz J. A., Madau P., Loeb A., Diemand J., 2009, eprint arXiv, 0905, 4744
- Müller E., Janka H. T., Wongwathanarat A., 2011
- Murante G., Borgani S., Brunino R., Cha S.-H., 2011, eprint arXiv, 1105, 1344
- Murante G., Monaco P., Giovalli M., Borgani S., Diaferio A., 2010, *Monthly Notices of the Royal Astronomical Society*, 405, 1491
- Murray N., Quataert E., Thompson T. A., 2005, *The Astrophysical Journal*, 618, 569

- Navarro J. F., Frenk C. S., White S. D. M., 1996, *The Astrophysical Journal*, 462, 563
- Navarro J. F., White S. D. M., 1994, *Monthly Notices of the Royal Astronomical Society*, 267, 401
- Nickerson S., Stinson G., Couchman H. M. P., Bailin J., Wadsley J., 2011, *Monthly Notices of the Royal Astronomical Society*, p. 662
- Nishi R., Kamaya H., 2000, *The Astrophysical Journal*, 532, 1172
- Nishikawa H., Kitamura K., 2008, *Journal of Computational Physics*, 227, 2560
- Nomoto K., Maeda K., Mazzali P. A., Umeda H., Deng J., Iwamoto K., 2003
- Oey M. S., Watson A. M., Kern K., Walth G. L., 2005, *The Astronomical Journal*, 129, 393
- Okamoto T., Frenk C. S., 2009, *Monthly Notices of the Royal Astronomical Society: Letters*, 399, L174
- Okamoto T., Gao L., Theuns T., 2008, *Monthly Notices of the Royal Astronomical Society*, 390, 920
- Oort J. H., Spitzer L., 1955, *The Astrophysical Journal*, 121, 6
- Orban C., Gnedin O. Y., Weisz D. R., Skillman E. D., Dolphin A. E., Holtzman J. A., 2008, *The Astrophysical Journal*, 686, 1030
- Pandolfi M., D'Ambrosio D., 2001, *Journal of Computational Physics*, 166, 271
- Parry O. H., Eke V. R., Frenk C. S., Okamoto T., 2011, eprint arXiv, 1105, 3474
- Peñarrubia J., Benson A. J., Walker M. G., Gilmore G., McConnachie A. W., Mayer L., 2010a, *Monthly Notices of the Royal Astronomical Society*, 406, 1290
- Peñarrubia J., Benson A. J., Walker M. G., Gilmore G., McConnachie A. W., Mayer L., 2010b, *Monthly Notices of the Royal Astronomical Society*, 406, 1290
- Peano G., 1890, *Mathematische Annalen*, 36, 157
- Peirani S., 2010, *Monthly Notices of the Royal Astronomical Society*, 407, 1487
- Peirani S., Crockett R. M., Geen S., Khochfar S., Kaviraj S., Silk J., 2010, *Monthly Notices of the Royal Astronomical Society*, 405, 2327
- Podsiadlowski P., Fabian A. C., Stevens I. R., 1991, *Nature*, 354, 43
- Pontzen A., Governato F., 2012, *Monthly Notices of the Royal Astronomical Society*, 421, 3464
- Powell L. C., Slyz A., Devriendt J., 2010, eprint arXiv, 1012, 2839
- Powell L. C., Slyz A., Devriendt J., 2011, *Monthly Notices of the Royal Astronomical Society*, 414, 3671
- Prendergast K. H., Xu K., 1993, *Journal of Computational Physics*, 109, 53
- Press W. H., Flannery B. P., Teukolsky S. A., 1986, Cambridge: University Press
- Price D. J., 2008, *Journal of Computational Physics*, 227, 10040

- Prochaska J. X., Gawiser E., Wolfe A. M., Castro S., Djorgovski S. G., 2003, *The Astrophysical Journal*, 595, L9
- Prunet S., Pichon C., Aubert D., Pogosyan D., Teyssier R., Gottloeber S., 2008, *The Astrophysical Journal Supplement Series*, 178, 179
- Puche D., Westpfahl D., Brinks E., Roy J.-R., 1992, *The Astronomical Journal*, 103, 1841
- Pustilnik S. A., Tepliakova A. L., Kniazev A. Y., Burenkov A. N., 2008, *Astrophysical Bulletin*, 63, 102
- Quinn T., Katz N., Efstathiou G., 1996, *Monthly Notices of the Royal Astronomical Society*, 278
- Quirk J. J., 1994, *International Journal for Numerical Methods in Fluids*, 18, 555
- Rajan A., Soummer R., Hagan J. B., Gilliland R. L., Pueyo L., 2011, *WFC3 Instrument Science Report 2011-07*
- Ramachandran B., Smith M. D., 2005, *Monthly Notices of the Royal Astronomical Society*, 357, 707
- Rasera Y., Teyssier R., 2006, *Astronomy and Astrophysics*, 445, 1
- Read J. I., Hayfield T., Agertz O., 2010, *Monthly Notices of the Royal Astronomical Society*, 405, 1513
- Redman M. P., Al-Mostafa Z. A., Meaburn J., Bryce M., 2003, *Monthly Notices of the Royal Astronomical Society*, 344, 741
- Reed D., Governato F., Quinn T., Gardner J., Stadel J., Lake G., 2005, *Monthly Notices of the Royal Astronomical Society*, 359, 1537
- Ren Y.-X. 2003, *Computers and Fluids*, 32, 25
- Richardson J. C., Irwin M. J., McConnachie A. W., Martin N. F., Dotter A. L., Ferguson A. M. N., Ibata R. A., Chapman S. C., Lewis G. F., Tanvir N. R., Rich R. M., 2011, *The Astrophysical Journal*, 732, 76
- Ricotti M., Gnedin N. Y., 2005, *The Astrophysical Journal*, 629, 259
- Ricotti M., Gnedin N. Y., Shull J. M., 2008, *The Astrophysical Journal*, 685, 21
- Riess A. G., Filippenko A. V., Challis P., Clocchiatti A., Diercks A., Garnavich P. M., Gilliland R. L., Hogan C. J., Jha S., Kirshner R. P., Leibundgut B., Phillips M. M., Reiss D., Schmidt B. P., Schommer R. A., Smith R. C., Spyromilio J., Stubbs C., Suntzeff N. B., Tonry J., 1998, *The Astronomical Journal*, 116, 1009
- Rodriguez-Gaspar J. A., Tenorio-Tagle G., Franco J., 1995, *The Astrophysical Journal*, 451, 210
- Romano-Díaz E., Shlosman I., Heller C., Hoffman Y., 2009, *The Astrophysical Journal*, 702, 1250
- Romano-Díaz E., Shlosman I., Heller C., Hoffman Y., 2010, *The Astrophysical Journal*, 716, 1095
- Rosdahl J., Blaizot J., 2012a, *Monthly Notices of the Royal Astronomical Society*, 423, 344
- Rosdahl J., Blaizot J., 2012b
- Rosen A., Bregman J. N., 1995, *The Astrophysical Journal*, 440, 634

- Rozyczka M., Tenorio-Tagle G., Franco J., Bodenheimer P., 1993, *Royal Astronomical Society*, 261, 674
- Runacres M. C., Owocki S. P., 2005, *Astronomy and Astrophysics*, 429, 323
- Salpeter E. E., 1955, *Astrophysical Journal*, 121, 161
- Saltzman J., 1994, *Journal of Computational Physics*, 115, 153
- Sawala T., Scannapieco C., White S., 2011, eprint arXiv, 1103, 4562
- Scannapieco C., Tissera P. B., White S. D. M., Springel V., 2005, *Monthly Notices of the Royal Astronomical Society*, 364, 552
- 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., Al E., 2011, eprint arXiv:1112.0315
- Scannapieco C., White S. D. M., Springel V., Tissera P. B., 2011, eprint arXiv, 1105, 680
- Schewtschenko J. A., Macciò A. V., 2011, *Monthly Notices of the Royal Astronomical Society*, 413, 878
- Schildt H., 1991, *C++: the complete reference*. Osborne McGraw-Hill, Berkeley
- Schwartz C. M., Martin C. L., 2004, *The Astrophysical Journal*, 610, 201
- Schweizer F., Seitzer P., 1992, *The Astronomical Journal*, 104, 1039
- Schweizer F., Seitzer P., Faber S. M., Burstein D., Dalle Ore C. M., Gonzalez J. J., 1990, *The Astrophysical Journal*, 364, L33
- Sedov L. I., 1946, *Journal of Applied Mathematics and Mechanics*, 10, 241
- Sharp D., 1984, *Physica D: Nonlinear Phenomena*, 12, 3
- Shu F. H., 1992, *The physics of astrophysics. Volume II: Gas dynamics*.
- Silk J., Rees M. J., 1998, *Astronomy and Astrophysics*
- Simon J. D., Geha M., 2007, *The Astrophysical Journal*, 670, 313
- Slyz A. D., Devriendt J. E. G., Bryan G., Silk J., 2005, *Monthly Notices of the Royal Astronomical Society*, 356, 737
- Smartt S. J., Eldridge J. J., Crockett R. M., Maund J. R., 2009, *Monthly Notices of the Royal Astronomical Society*, 395, 1409
- Smith M. D., Rosen A., 2003, *Monthly Notices of the Royal Astronomical Society*, 339, 133
- Somerville R. S., 2002, *The Astrophysical Journal*, 572, L23
- Springel V., 2005, *Monthly Notices of the Royal Astronomical Society*, 364, 1105
- Springel V., 2010, *Monthly Notices of the Royal Astronomical Society*, 401, 791

- Springel V., Di Matteo T., Hernquist L., 2005, *Monthly Notices of the Royal Astronomical Society*, 361, 776
- Springel V., Hernquist L., 2003, *Monthly Notices of the Royal Astronomical Society*, 339, 289
- 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, *Nature*, 435, 629
- Steinmetz M., Mueller E., 1994, *Astronomy and Astrophysics (ISSN 0004-6361)*, 281
- Stepanov A., Lee M., 1995, Technical report, The Standard Template Library. HP Laboratories Technical Report 95-11(R.1)
- Stewart S. G., Fanelli M. N., Byrd G. G., Hill J. K., Westpfahl D. J., Cheng K., O'Connell R. W., Roberts M. S., Neff S. G., Smith A. M., Stecher T. P., 2000, *The Astrophysical Journal*, 529, 201
- Stinson G. S., Dalcanton J. J., Quinn T., Kaufmann T., Wadsley J., 2007, *The Astrophysical Journal*, 667, 170
- Strigari L. E., Bullock J. S., Kaplinghat M., Diemand J., Kuhlen M., Madau P., 2007, *The Astrophysical Journal*, 669, 676
- Strigari L. E., Bullock J. S., Kaplinghat M., Simon J. D., Geha M., Willman B., Walker M. G., 2008, *Nature*, 454, 1096
- Strömberg B., 1939, *The Astrophysical Journal*, 89, 526
- Stroustrup B., 1986, *The C++ programming language*. Addison-Wesley, Reading Mass.
- Sugerman B. E. K., Crofts A. P. S., Kunkel W. E., Heathcote S. R., Lawrence S. S., 2005, *The Astrophysical Journal*, 627, 888
- Susa H., Umemura M., 2004, *The Astrophysical Journal*, 610, L5
- Sutherland R. S., Dopita M. A., 1993, *The Astrophysical Journal Supplement Series*, 88, 253
- Tang S., Wang Q. D., 2005, *The Astrophysical Journal*, 628, 205
- Tasker E. J., Brunino R., Mitchell N. L., Michielsen D., Hopton S., Pearce F. R., Bryan G. L., Theuns T., 2008, *Monthly Notices of the Royal Astronomical Society*, 390, 1267
- Tassis K., Kravtsov A. V., Gnedin N. Y., 2008, *The Astrophysical Journal*, 672, 888
- Taylor G., 1950, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 201, 159
- Tenorio-Tagle G., Bodenheimer P., Franco J., Rozyczka M., 1990, *Royal Astronomical Society*, 244, 563
- Tenorio-Tagle G., Rozyczka M., Franco J., Bodenheimer P., 1991, *Royal Astronomical Society*, 251, 318

- Teyssier R., 2002, *Astronomy and Astrophysics*, 385, 337
- Thiébaud J., Pichon C., Sousbie T., Prunet S., Pogosyan D., 2008, *Monthly Notices of the Royal Astronomical Society*, 387, 397
- Thornton K., Gaudlitz M., Janka H.-T., Steinmetz M., 1998, *The Astrophysical Journal*, 500, 95
- Toalá J. A., Arthur S. J., 2011, *The Astrophysical Journal*, 737, 100
- Tollerud E. J., Bullock J. S., Strigari L. E., Willman B., 2008, *The Astrophysical Journal*, 688, 277
- Tolstoy E., Hill V., Tosi M., 2009, *Annual Review of Astronomy and Astrophysics*, 47, 371
- Toomre A., 1977, *Evolution of Galaxies and Stellar Populations*
- Toomre A., Toomre J., 1972, *The Astrophysical Journal*, 178, 623
- Toro E. F., 2006, *Riemann solvers and numerical methods for fluid dynamics - practical introduction - 2nd edition*. Springer
- Toro E. F., Spruce M., Speares W., 1994, *Shock Waves*, 4, 25
- Turk M. J., Smith B. D., Oishi J. S., Skory S., Skillman S. W., Abel T., Norman M. L., 2011, *The Astrophysical Journal Supplement Series*, 192, 9
- Tweed D., Devriendt J., Blaizot J., Colombi S., Slyz A., 2009, eprint arXiv, 0902, 679
- van Dokkum P. G., 2005, *The Astronomical Journal*, 130, 2647
- van Marle A. J., Keppens R., Meliani Z., 2011, *Astronomy & Astrophysics*, 527, A3
- van Marle A. J., Langer N., Yoon S.-C., García-Segura G., 2008, *Astronomy and Astrophysics*, 478, 769
- Veilleux S., Cecil G., Bland-Hawthorn J., 2005, *Annual Review of Astronomy and Astrophysics*, 43, 769
- Vera-Ciro C. A., Sales L. V., Helmi A., Frenk C. S., Navarro J. F., Springel V., Vogelsberger M., White S. D. M., 2011, *Monthly Notices of the Royal Astronomical Society*, 416, 1377
- Vishniac E. T., 1983, *The Astrophysical Journal*, 274, 152
- Vishniac E. T., 1994, *The Astrophysical Journal*, 428, 186
- Vishniac E. T., 1995, *Annals of the New York Academy of Sciences*, 773, 70
- Wadepuhl M., Springel V., 2010a, eprint arXiv, 1004, 3217
- Wadepuhl M., Springel V., 2010b, *Monthly Notices of the Royal Astronomical Society*, 410, no
- Wadsley J., Stadel J., Quinn T., 2004, *New Astronomy*, 9, 137
- Wampler E. J., Wang L., Baade D., Banse K., D'Odorico S., Gouiffes C., Tarenghi M., 1990, *The Astrophysical Journal*, 362, L13
- Wang J., Frenk C. S., Navarro J. F., Gao L., 2012, eprint arXiv:1203.4097

- Watkins L. L., Evans N. W., An J. H., 2010, *Monthly Notices of the Royal Astronomical Society*, 406, 264
- Weaver R., McCray R., Castor J., Shapiro P., Moore R., 1977, *The Astrophysical Journal*, 218, 377
- Wheeler J. C., Harkness R. P., 1990, *Reports on Progress in Physics*, 53, 1467
- White G. L., Malin D. F., 1987, *Nature*, 327, 36
- White S. D. M., Rees M. J., 1978, *Royal Astronomical Society*, 183, 341
- Wilkins S. M., Bunker A. J., Lorenzoni S., Caruana J., 2010, eprint arXiv, 1002, 4866
- Williams R. J. R., 1999, *Monthly Notices of the Royal Astronomical Society*, 310, 789
- Wolf J., Martinez G. D., Bullock J. S., Kaplinghat M., Geha M., Muñoz R. R., Simon J. D., Avedo F. F., 2010, *Monthly Notices of the Royal Astronomical Society*, p. 778
- Wosley S., Janka T., 2005, *Nature Physics*, 1, 147
- Wyithe S., Mould J., Loeb A., 2011
- Zwicky F., 1933, *Helvetica Physica Acta*, 6, 110