Simulations of mass accretion onto dark matter haloes and angular momentum transfer to a Milky Way disk at high redshift

Henry William Tillson

The Department of Physics and Lincoln College, Oxford

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

Trinity Term 2012

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. Except where explicit reference is made to the work of others, the work contained in this thesis is my own, and is not the outcome of work done in collaboration.

Section 3.2 of Chapter 3 and the majority of Chapter 4 have been published in the Monthly Notices of the Royal Astronomical Society:
Tillson, H., Miller, L. and Devriendt, J., 'The environment and redshift dependence of accretion on to dark matter haloes and subhaloes', MNRAS, 417, 666-680, 2011

The majority of Chapters 5 and 6 are about to be submitted to the Monthly Notices of the Royal Astronomical Society:
Tillson, H., Devriendt, J., Slyz, A., Miller, L. and Pichon, C., 'Angular momentum transfer to a Milky Way disk at high redshift'

Acknowledgments

I wish to thank Julien Devriendt, Lance Miller and Adrianne Slyz for academic guidance during my DPhil years. I feel that my ability at tackling complex problems has greatly improved under their supervision.

I also acknowledge several friends studying physics with whom I have had countless fruitful academic conversations and have enjoyed many social occasions. These include Olaf Davis, Philip Jones, Gunnar Langfahl-Klabes, Richard Masters, Sarah Miller and Tim Wang. I must also thank Sam Geen and Taysun Kimm, fellow students within the Oxford simulation group. Richard Booth, Roger Deane and Thibaut Louis merit particular recognition for engaging in numerous invigorating academic discussions and contributing towards a balanced sense of perspective during my graduate years. I also wish to highlight the positive impact friends outside my field have had, both in Oxford—Bipana Bantawa, Nahal Khabbazbashi, Joseph Rogel, Richard Simmonds—and in Cambridge—Antony Chan, Jan Guthrie, Thomas Owen and Peter Sarkies.

The people to whom I am truly indebted, however, and to whom I dedicate this thesis, are my parents Gary and Judith, my sister Charlotte, and my very close friend Janeth Peña-Heredia. These people have been a gigantic source of emotional support and motivation, and I celebrate the completion of this thesis with them.

Simulations of mass accretion onto dark matter haloes and angular momentum transfer to a Milky Way disk at high redshift

A thesis submitted for the degree of Doctor of Philosophy

Henry William Tillson, Lincoln College

Trinity Term 2012

Abstract

This thesis presents results from two simulation studies of galaxy formation. In the first project, a dark-matter-only HORIZON simulation is used to investigate the environment and redshift dependence of mass accretion onto haloes and subhaloes. It is found that the halo accretion rate varies less strongly with redshift than predicted by the Extended Press–Schechter formalism, and that low accretion events may drive the radio-mode feedback hypothesized for recent galaxy formation models. The subhaloes at $z < 0.5$ in the simulation accrete at higher rates than haloes, on average, and it is argued that this is due to their enhanced clustering at small scales. There is no dependence of accretion rate on environment at $z \sim 2$, but a weak correlation emerges at $z \leq 0.5$. The results further support previous suggestions that at $z > 1$, dark matter haloes and their associated black holes grew coevally, but imply that haloes could be accreting at fractional rates that are up to a factor of 3–4 higher than their associated black holes by the present day.

In the second project, outputs from one of the Adaptive Mesh Refinement NUT simulations are analyzed in order to test whether filamentary flows of cold gas are responsible for the build-up of angular momentum within a Milky Way type disk at $z \geq 3$. A set of algorithms are presented that use the resolved physical scale of 12 pc to identify: (i) the central gas disk and its plane of orientation; (ii) the complex individual filament trajectories that connect to the disk, and; (iii) the infalling satellites. The results suggest that two filaments at $z \gtrsim 5.5$, which later merge to form a single filament at $z \lesssim 4$, drive the angular momentum and mass budget of the disk between $3 \lesssim z \lesssim 8$, whereas luminous satellite mergers make negligible fractional contributions. These findings hence provide strong quantitative evidence that the growth of thin disks in low mass haloes at high redshift is supported via inflowing streams of cold gas.

Contents

Chapter 1

Introduction

This chapter focusses primarily on describing the standard cosmological model encoded within galaxy simulations.

1.1 The origin of structure

1.1.1 The Friedmann–Lemaître–Robinson–Walker formalism

The Universe is thought to have undergone a period of very rapid inflation at very early times (Guth, 1981). In the standard inflationary paradigm there is a single scalar field (Guth & Pi, 1982; Linde, 1982) and fluctuations in this field give rise to fluctuations in the space-time metric, as described by the Einstein equations, which, conforming to the notations adopted by Coles & Lucchin 2002, are expressed as follows:

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = \frac{8\pi G}{c^4} \tilde{T}_{\mu\nu}. \tag{1.1}$$

Here G refers to Newton's gravitational constant, c is the speed of light, R is the Ricci scalar ($= \sum_{\mu\nu} g_{\mu\nu} R^{\mu\nu}$), and the terms $R_{\mu\nu}, g_{\mu\nu}$ and $\tilde{T}_{\mu\nu}$ represent the Ricci tensor, the space-time metric tensor and the effective energy-momentum tensor respectively. The expression for $\tilde{T}_{\mu\nu}$ includes the cosmological constant Λ term:

$$\begin{aligned} \tilde{T}_{\mu\nu} &= T_{\mu\nu} + \frac{\Lambda c^4}{8\pi G} g_{\mu\nu} & (1.2) \\ &= -\tilde{p} g_{\mu\nu} + (\tilde{p} + \tilde{\rho} c^2) U_\mu U_\nu. & (1.3) \end{aligned}$$

The latter step holds if the fluid within the universe is assumed to be a perfect fluid with a four velocity U_μ and an effective pressure $\tilde{p}$ and an effective density $\tilde{\rho}$ given by the following expressions:

$$\tilde{p} = p - \frac{\Lambda c^4}{8\pi G} \tag{1.4}$$

$$\tilde{\rho} = \rho + \frac{\Lambda c^2}{8\pi G}, \tag{1.5}$$

where p and ρ represent the pressure and density of the fluid. Solutions to equation (1.1) hence depend on the geometry of the universe, encapsulated by the space-time metric. Perhaps the simplest description of a universe is one that is statistically homogeneous and isotropic that contains no special spatial locations and directions, in concordance with the Cosmological Principle (homogeneity and isotropy are more formally defined later in this section). The popular Friedmann–Lemaître–Robinson–Walker (hereafter FLRW) metric describes an expanding universe with these properties, which for a 3D surface with curvature is of the form:

$$g_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & \frac{a^2(t)}{1-kr^2} & 0 & 0 \\ 0 & 0 & a^2(t)r^2 & 0 \\ 0 & 0 & 0 & a^2(t)r^2\sin^2\theta \end{pmatrix}, \tag{1.6}$$

where $a(t)$ is the dimensionless scale factor (whose explicit time dependence is now dropped for brevity) normalized with respect to its present day value, and k is the curvature parameter (with dimensions of $1/r^2$). This choice of metric, which is written in terms of dimensionless comoving distances, follows the conventions adopted by several textbooks on the subject (see, for example, Peacock 1999 and Liddle 2003). By using equation (1.6) and making the assumption that the universe consists of a perfect fluid, it follows that the solutions to equation (1.1) are:

$$\ddot{a} = -\frac{4\pi}{3} G \left(\tilde{\rho} + \frac{3\tilde{p}}{c^2} \right) a \tag{1.7}$$

$$\left(\frac{\dot{a}}{a} \right)^2 = \frac{8\pi G}{3} \tilde{\rho} - \frac{kc^2}{a^2}. \tag{1.8}$$

Equations (1.7) and (1.8) hence completely determine the evolution of a homogeneous and isotropic universe of a chosen cosmology. Defining the Hubble factor as $H(a) \equiv \dot{a}/a$, the critical density as

$\rho_c(a) \equiv 3H^2(a)/(8\pi G)$, and $\Omega(a) \equiv \tilde{\rho}(a)/\rho_c(a)$, equation (1.8) becomes:

$$\Omega(a) - 1 = \frac{kc^2}{a^2 H^2(a)}. \tag{1.9}$$

In a universe with $\Omega_\Lambda(a) = 0$, the expansion is undone by collapse to a singularity if $k > 0$, whereas the expansion is eternal if $k \leq 0$.

1.1.2 Statistical properties of our Universe

Scalar field perturbations induce perturbations in the space-time metric: density perturbations around the mean mass density in the Universe $\bar{\rho}$ are hence born and evolve with time as the Universe expands (e.g. Peebles 1993). A density perturbation around $\bar{\rho}$ at a given spatial location $\vec{x}$ is often parametrized by the dimensionless density contrast $\delta(\vec{x})$:

$$\delta(\vec{x}) = \frac{\rho(\vec{x}) - \bar{\rho}}{\bar{\rho}}, \tag{1.10}$$

whose time index has been dropped for brevity. If the Universe is modelled as a periodic box of length L and volume V $(= L^3)$, then the real-space density contrast at a given location $\delta(\vec{x})$ can be decomposed into a finite sum of Fourier modes $\delta_{\vec{k}}$, where $|\vec{k}| = 2\pi p/L$ and p is an integer. It is hence useful to transform to Fourier-space, and all of the derivations and discussions in this subsection follow the approaches adopted in several of the popular reviews of cosmology (e.g. Peebles 1993, Peacock 1999, Coles & Lucchin 2002, Zentner 2007). If the box volume V approaches infinity, $\delta(\vec{x})$ becomes:

$$\delta(\vec{x}) = \frac{V}{(2\pi)^3} \int \delta_{\vec{k}} \exp(-i\vec{k} \cdot \vec{x}) \, \mathrm{d}^3 k. \tag{1.11}$$

Similarly, each Fourier mode $\delta_{\vec{k}}$ can be expressed in terms of $\delta(\vec{x})$ by taking the inverse of equation (1.11):

$$\delta_{\vec{k}} = \frac{1}{V} \int \delta(\vec{x}) \exp(i\vec{k} \cdot \vec{x}) \, \mathrm{d}^3 x. \tag{1.12}$$

L is sufficiently large (much greater than the size of large-scale structure) that the Universe is accurately sampled and $\int |\delta(\vec{x})| \, \mathrm{d}^3 x$ is finite.

The standard model of cosmology, which is discussed at length later in this chapter, makes two important assumptions about the nature of the initial Fourier modes $\delta_{\vec{k}}$ in the linear regime of perturbation growth:

- their phases are random and independent of each other, and;
- their distribution is statistically isotropic and homogeneous, in concordance with the Cosmological Principle.

As a consequence of the central limit theorem, in the limit of a large number of $\delta_{\vec{k}}$, the probability density function for a set of perturbations $\{\delta(\vec{x}_i)\}$ of size N is described by an N-variate Gaussian distribution. For a homogeneous density field, the mean density $\langle\delta(\vec{x})\rangle$ (which is zero for a random density field) and the correlation between two density samples at two different arbitrary spatial locations $\xi(\vec{x}_i, \vec{x}_j) \equiv \langle\delta(\vec{x}_i)\delta(\vec{x}_j)\rangle$, are translationally invariant. The averages are performed over independent pairs of points whose spatial separation is $|\vec{r}_{ij}| = |\vec{x}_i - \vec{x}_j|$. Isotropy further enforces rotational invariance, such that $\xi(\vec{x}_i, \vec{x}_j) = \xi(|\vec{x}_i - \vec{x}_j|)$. The correlation between two perturbations located at $\vec{x}$ and $\vec{x} + \vec{r}$ can hence be found by using this property of rotational invariance and equation (1.11), which is now represented in its discrete form for brevity (following Peebles 1993):

$$\xi(\vec{r}) = \left\langle \sum_{\vec{k}} \sum_{\vec{k}'} \delta_{\vec{k}} \delta_{\vec{k}'} \exp\left[-i(\vec{k} + \vec{k}') \cdot \vec{x}\right] \exp(-i\vec{k} \cdot \vec{r}) \right\rangle . \tag{1.13}$$

According to equation (1.12), $\delta_{-\vec{k}} = \delta^*_{\vec{k}}$, and by invoking the property that $\delta(\vec{x})$ is real, equation (1.13) becomes:

$$\begin{aligned} \xi(\vec{r}) &= \left\langle \sum_{\vec{k}} \sum_{\vec{k}'} \delta_{\vec{k}} \delta^*_{\vec{k}'} \exp\left[i(\vec{k}' - \vec{k}) \cdot \vec{x}\right] \exp(-i\vec{k} \cdot \vec{r}) \right\rangle & (1.14) \\ &= \sum_{\vec{k}} \sum_{\vec{k}'} \left\langle \delta_{\vec{k}} \delta^*_{\vec{k}'} \right\rangle \exp\left[i(\vec{k}' - \vec{k}) \cdot \vec{x}\right] \exp(-i\vec{k} \cdot \vec{r}). & (1.15) \end{aligned}$$

The Fourier modes obey the relation:

$$\left\langle \delta_{\vec{k}} \delta^*_{\vec{k}'} \right\rangle = \delta^D_{\vec{k}\vec{k}'} \left\langle \left|\delta_{\vec{k}}\right|^2 \right\rangle \left(\frac{2\pi}{L}\right)^3, \tag{1.16}$$

where $\delta^D_{\vec{k}\vec{k}'}$ refers to the 3-dimensional Dirac–delta function (and not the density contrast), and so the expression for $\xi(r)$ in equation (1.15) simplifies to:

$$\xi(r) = \frac{V}{(2\pi)^3} \int \left\langle \left|\delta_{\vec{k}}\right|^2 \right\rangle \exp(-i\vec{k} \cdot \vec{r})\, \mathrm{d}^3k. \tag{1.17}$$

If the Fourier transform of the power spectrum is defined as:

$$P(k) \equiv V \left\langle \left|\delta_{\vec{k}}\right|^2 \right\rangle, \tag{1.18}$$

where $k = |\vec{k}|$, equation (1.17) becomes:

$$\begin{aligned}
\xi(r) &= \frac{1}{(2\pi)^3} \int P(k) \exp(-ikr\cos\theta)\, \mathrm{d}^3k & (1.19)\\
&= \frac{1}{(2\pi)^3} \int_\phi \int_\theta \int_k P(k) \exp(-ikr\cos\theta) k^2 \sin\theta\, \mathrm{d}k\, \mathrm{d}\theta\, \mathrm{d}\phi & (1.20)\\
&= -\frac{1}{(2\pi)^2} \int_\theta \int_k P(k) \exp(-ikr\cos\theta) k^2\, \mathrm{d}k\, \mathrm{d}(\cos\theta) & (1.21)\\
&= \frac{1}{2\pi^2} \int_{-\infty}^{\infty} k^3 P(k) \frac{\sin(kr)}{kr}\, \mathrm{d}(\ln \mathrm{k}). & (1.22)
\end{aligned}$$

Since $P(k)$ has dimensions of volume, it is useful to define the parameter $\Delta^2(k)$:

$$\Delta^2(k) \equiv \frac{k^3 P(k)}{2\pi^2}, \tag{1.23}$$

which can be inserted into equation (1.22) to yield:

$$\xi(r) = \int_{-\infty}^{\infty} \Delta^2(k) \frac{\sin(kr)}{kr}\, \mathrm{d}(\ln \mathrm{k}). \tag{1.24}$$

It follows that the two-point correlation function is the Fourier transform of the power spectrum (equation 1.19) and corresponds to the mass variance (equation 1.24), which can be thought of as the sum in logarithmic intervals of each $\Delta^2(k)$ mode modulated by a factor $\sin(kr)/(kr)$.

The k dependence of $P(k)$

In the single field standard model of inflation, the mean squared fluctuations in the gravitational potential Φ are independent of scale as they enter the horizon, such that Δ^2_Φ is constant:

$$\Delta^2_\Phi \propto k^3 \left\langle \left|\delta\Phi_{\vec{k}}\right|^2 \right\rangle \propto \frac{k^3 \left\langle \left|\delta_{\vec{k}}\right|^2 \right\rangle}{k^4} \propto \frac{P(k)}{k} \propto \text{const.} \tag{1.25}$$

The second step has used the relation $\delta\Phi_{\vec{k}} \propto \delta_{\vec{k}}/k^2$, which is found by taking the Fourier transform of the Poisson equation. This primordial scale invariance in Δ^2_Φ hence corresponds to a primordial power spectrum in the density fluctuations that grows linearly with wavenumber k, i.e. $P(k) \propto k$,

yielding $\Delta^2(k) \propto k^4$ (according to equation 1.23). At later times, as the horizon scale of the Universe $l_{\rm H} = 1/[aH(a)]$ expands after inflation, this scale invariance is not preserved because perturbations on different scales cross the horizon at different epochs. For example, at the time of matter-radiation density equality $z_{\rm eq}$, perturbations on scales larger than the horizon scale at this epoch, $l_{\rm H}(z_{\rm eq})$, collapse under their own gravity and do not enter the horizon until later times when the Universe is matter-dominated. By contrast, perturbations on small scales that enter the horizon shortly after inflation are driven by the radiation-dominated expansion, which provides a pressure support against gravitational collapse. These perturbations are hence frozen until $z_{\rm eq}$, where non-baryonic matter (i.e. dark matter) decouples from the radiation field and becomes unstable to gravitational collapse. The linear power spectrum on small scales at later epochs is therefore a modified version of its primordial form, with a dependence on k and z that is characterized by the transfer function $T(k,z)$ and the evolution of the linear growth factor $D(z)$: $P(k,z) \propto P(k,z\rightarrow\infty)T^2(k,z)D^2(z)$. Bardeen et al. (1986) demonstrated that $T(k,z)$ becomes approximately independent of redshift for $z < 100$ and found that a function of the following form for a flat universe dominated by cold dark matter (this model universe is discussed in more detail in Section 1.3) provides a good fit to the observed transfer function:

$$T(k) = \frac{\ln(1+2.34q)}{2.34q}\left[1+3.89q+(16.1q)^2+(5.46q)^3+(6.71q)^4\right]^{-0.25}, \tag{1.26}$$

where q depends on k and the fraction of the total present day energy density in adiabatic cold dark matter, $\Omega_{\rm M}$, and is defined as:

$$q \equiv \frac{k}{h^2\Omega_{\rm M}\,{\rm Mpc}^{-1}}. \tag{1.27}$$

1.2 From density perturbations to virialized haloes of dark matter

1.2.1 The spherical top-hat model of halo formation

The previous section has introduced the FLRW model that describes the evolution of the Universe across cosmic time, and has demonstrated that the real-space density perturbations can be represented as a superposition of Fourier modes whose amplitude and time evolution are characterized by the power spectrum $P(k,z)$. This section explains how haloes of dark matter—the fundamental building block of galaxies—form from these real-space density fluctuations as they subsequently

evolve in an expanding FLRW universe. The spherical top-hat model of halo formation is presented and is divided into three phases: expansion, turnaround and virialization upon collapse. The discussions that accompany each of these stages follow the arguments presented in Peacock (1999) and for simplicity, refer to a matter-dominated, closed universe with $\Omega_\Lambda = 0$. There are, however, several additional works that complement the following (Gunn & Gott, 1972; Lacey & Cole, 1993) and apply the spherical top-hat model to a universe with $\Omega_\Lambda \neq 0$ (Eke et al., 1996).

1. **Expansion.** One prediction of Birkhoff's theorem is that a spherically symmetric overdense region containing a mass M grows at the same rate as a sphere of equal mass evolving with the background expansion. The radius $R(t)$ of a given density perturbation, modelled as a uniformly overdense sphere, therefore has the following time dependence:

$$\dot{R}^2(t) = \frac{8\pi G\rho(t)R^2(t)}{3} - kc^2. \tag{1.28}$$

Assuming that the perturbation grows in a matter-dominated universe, the density $\rho(t)$ evolves as $\rho(t) = \rho_0 R_0^3/R^3(t)$, where the subscript '0' refers to initial quantities. Since collapse to a singularity occurs at time $t_{\rm coll}$ in the spherical top-hat model, the overdense region can be thought of as behaving like a closed universe with $k = +1$, provided that the scale factor is rescaled by the curvature parameter: $a(t) \rightarrow a(t)/\sqrt{k}$. Equation (1.28) hence simplifies to:

$$\dot{R}(t)^2 = \frac{2G}{R(t)}\frac{4\pi R_0^3\rho_0}{3} - c^2 \tag{1.29}$$

$$= \frac{2GM}{R(t)} - c^2, \tag{1.30}$$

where M represents the mass within the sphere:

$$M = \frac{4}{3}\pi R^3(t)\rho(t) = \frac{4}{3}\pi R_0^3\rho_0 = M_0, \tag{1.31}$$

which is constant at all times in the halo formation process. As the perturbation becomes non-linear ($\delta > 1$), this assumption will no longer hold because local matter will gravitationally sink inwards, and so the perturbation will acquire a radial density profile (for an estimate of the functional form of this profile, see Penston 1969 and Peebles 1993). Nevertheless, the assumption of constant mass is preserved in what follows since the aim is to make simple

model predictions. By defining conformal time η as:

$$\eta = c \int_0^t \frac{\mathrm{d}t'}{R(t')}, \tag{1.32}$$

equation (1.30) can be further simplified:

$$\left(\frac{\mathrm{d}R}{\mathrm{d}\eta}\right)^2 = \frac{2GMR}{c^2} - R^2 \tag{1.33}$$

$$\left[\left(\frac{c^2}{GM}\right)\frac{\mathrm{d}R}{\mathrm{d}\eta}\right]^2 = \frac{2GMR}{c^2}\left(\frac{c^2}{GM}\right)^2 - R^2\left(\frac{c^2}{GM}\right)^2 \tag{1.34}$$

$$\left[\frac{\mathrm{d}}{\mathrm{d}\eta}\left(\frac{R}{R_\star}\right)\right]^2 = 2\left(\frac{R}{R_\star}\right) - \left(\frac{R}{R_\star}\right)^2, \tag{1.35}$$

where $R_\star$ is a constant defined as $R_\star \equiv GM/c^2$, and the explicit η dependence of R has been dropped for brevity. The solution to equation (1.35) can be found by completing the square and changing variables, and is of the form:

$$R(\eta) = R_\star(1 - \cos\eta). \tag{1.36}$$

It follows from equation (1.32) that t can also be written as a function of η:

$$t(\eta) = \frac{1}{c}\int_0^\eta R(\eta')\,\mathrm{d}\eta' = \frac{R_\star}{c}(\eta - \sin\eta). \tag{1.37}$$

For early times, $\eta << 1$, and so by performing a Taylor series expansion in $\cos\eta$ and $\sin\eta$ around $\eta \sim 0$ up to the η^5 term and subsequently eliminating η, it can be shown that:

$$\lim_{\eta\to 0} R(t) \simeq \frac{R_\star}{2}\left(\frac{6ct}{R_\star}\right)^{2/3}\left[1 - \frac{1}{20}\left(\frac{6ct}{R_\star}\right)^{2/3}\right] = R_0\left[1 + \frac{\delta R_0}{R_0}\right]. \tag{1.38}$$

By invoking the relation $\rho(R) \propto 1/R^3$, the linear theory prediction for the density contrast at early times can be found:

$$\delta \equiv \frac{\delta\rho}{\rho} = \frac{\delta\rho}{\delta R}\frac{\delta R}{\rho} = -3\frac{\delta R}{R} \simeq \frac{3}{20}\left(\frac{6ct}{R_\star}\right)^{2/3}. \tag{1.39}$$

The density perturbation therefore initially grows linearly with the scale factor ($a(t) \propto t^{2/3}$ in a matter-dominated universe).

2. **Turnaround**. The perturbation expands with the background until a point of maximum expansion at time t_t, and it can be seen that R in equation (1.36) is equal to its maximum value of $2R_\star$ when $\eta = \pi$, corresponding to a turnaround epoch $t_t = \pi R_\star/c$ in equation (1.37). The density contrast of the perturbation ρ_p with respect to the background density ρ_b is:

$$\delta = \frac{\rho_p}{\rho_b} - 1, \tag{1.40}$$

which at $t = t_t$ becomes:

$$\delta = \frac{\left[(R_\star/2)(6ct_t/R_\star)^{2/3}\right]^3}{(2R_\star)^3} - 1 = \frac{9\pi^2}{16} - 1 \sim 4.55, \tag{1.41}$$

according to the spherical top-hat model in equation (1.38). This prediction is roughly a factor of 4–5 larger than the linear theory estimate of equation (1.39):

$$\delta_l \simeq \frac{3}{20}\left(\frac{6ct_t}{R_\star}\right)^{2/3} = \frac{3}{20}(6\pi)^{2/3} \sim 1. \tag{1.42}$$

It is hence clear that the growth of the perturbation is non-linear at turnaround.

3. **Collapse and virialization.** The perturbation is instantaneously stationary at $t = t_t$ and so all of its energy is gravitational potential energy V. The kinetic energy T of the subsequent collapse is converted into random motions. When $R = R_\star$, i.e. half the maximum turnaround radius, the kinetic and potential energy of the perturbation obey the virial relation $2T+V = 0$. Equation (1.36) shows that $R = R_\star$ when $\eta = 3\pi/2$, which yields a virial overdensity of $\delta_{\rm vir} \sim 146$ upon substitution into equation (1.40). A more common assumption, however, is that the halo is virialized at the collapse epoch where $R = 0$, which occurs when $\eta = 2\pi$, yielding a collapse time $t_{\rm coll} = 2\pi R_\star/c$. The virial density contrast estimate at $t = t_{\rm coll}$ predicted by the spherical top-hat model is therefore:

$$\delta_{\rm vir} = \frac{(R_\star/2)^3(6ct_{\rm coll}/R_\star)^2}{R_\star^3} - 1 = \frac{1}{8}\left(\frac{6c}{R_\star}\frac{2\pi R_\star}{c}\right)^2 - 1 = 18\pi^2 - 1 \sim 177, \tag{1.43}$$

which is significantly greater than the linear theory estimate:

$$\delta_l \simeq \frac{3}{20}\left(\frac{6ct_{\rm coll}}{R_\star}\right)^{2/3} = \frac{3}{20}(12\pi)^{2/3} \sim 1.69. \tag{1.44}$$

So far it has been assumed that the density perturbations are sufficiently dense that they will collapse to form haloes. Lower density perturbations will form haloes at later times, however, and so several authors have introduced the notion of a 'critical' density contrast for halo formation at a given epoch (Gunn & Gott, 1972; Press & Schechter, 1974; Lacey & Cole, 1993). This features in the theory of Press & Schechter (1974), which ultimately leads to a prediction of the number density of collapsed virialized haloes in a given mass range at a given redshift, and is discussed in detail in Chapter 3.

1.2.2 Virial properties of dark matter haloes

Many of the arguments presented in this thesis make reference to a halo's virial radius $r_{\rm vir}$ and virial mass $M_{\rm vir}$, and so these quantities are now defined. The halo-finders used in this study (which are introduced in Chapter 2) compute the virial radius by dividing each resolved halo into ellipsoidal bins and recording the distance from the halo's centre to the outermost bin whose 'virial accuracy parameter' (i.e. ratio of virial energy to total energy) is at least 20%. The virial radius is then defined as:

$$r_{\rm vir} = \sqrt[3]{abc}, \tag{1.45}$$

where a, b and c refer to the semi-principal axes of the outermost virial ellipsoid. Equation (1.45) hence ensures that the volume of the spherical virial region and the outermost virial ellipsoid are equal. It follows that the virial mass of a spherical halo is given by $M_{\rm vir} = \frac{4}{3}\pi r_{\rm vir}^3 \Delta_{\rm vir} \bar{\rho}$, where $\bar{\rho}$ is the mean background mass density and $\Delta_{\rm vir}$ is the virial density contrast. Bryan & Norman (1998) provide functional fits for $\Delta_{\rm vir}$ that depend on redshift and cosmology: for an Einstein–de Sitter and ΛCDM universe at $z = 0$, $\Delta_{\rm vir} \sim 180$, 340 respectively.

1.2.3 Observational evidence for the existence of dark matter

Since reference has been made to haloes of dark matter throughout this section, it is worth briefly summarizing the observational evidence in support of the existence of this dark component, before proceeding to the ΛCDM model.

The first prediction for the existence of dark matter was made by Zwicky (1933), who applied the virial theorem to the Coma cluster and estimated its mass-to-light ratio as $\Upsilon \sim 300\, h\, {\rm M}_\odot/{\rm L}_\odot$, where h is a dimensionless normalization of the Hubble constant: $h = H_0/(100\,{\rm km\,s^{-1}\,Mpc^{-1}})$. It was assumed that light traces galaxy mass, and so this excess of mass in the Coma cluster

was attributed to an unseen, dark component. This claim was later supported by White et al. (1993), who used X-ray data from hot intracluster gas to measure the mass of Coma and estimated that baryons account for only ∼10% of its mass. Similar studies have further confirmed that early predictions of constant Υ hold to a good approximation (Merritt, 1987; Carlberg et al., 1997) within the central regions ($r < 1\,h^{-1}\,\mathrm{Mpc}$) of X-ray clusters (typically enclosing a total mass of $\sim 10^{15}\,\mathrm{M}_{\odot}$), provided that the cluster gas is in hydrostatic equilibrium. Weak-lensing measurements of cluster masses that probe to scales ∼1 Mpc from cluster centres also yield mass-to-light ratios that are in concordance with X-ray estimates (e.g. Hoekstra et al. 1998).

Other evidence for dark matter comes from the velocity rotation curves of spiral galaxies, denoted by $V(r)$, which flatten beyond ∼5–10 kpc from halo centres (Rubin et al., 1980; van Albada et al., 1985). This flattening is in tension with the $V(r) \propto 1/r$ Keplerian prediction and implies that $M(r) \propto r$ in the outer regions of these rotationally supported systems. Dark matter is thought to constitute the unobserved mass at large r.

1.3 The ΛCDM model of the Universe

To date, almost every galaxy simulation models the formation of dark matter haloes and luminous galaxies in a flat universe whose energy density is dominated by a cosmological constant and a collisionless cold dark matter (CDM) component that acts via the gravitational force but perhaps also via the weak force. In this ΛCDM framework, the initial density fluctuations are close to scale-invariant and are Gaussian distributed, growing with the expansion of the background until their associated gravity induces collapse and subsequent virialization, leading to the formation of haloes of dark matter, as described in the previous section. Luminous galaxies form as baryonic matter cools and condenses at halo centres (White & Rees, 1978; Fall & Efstathiou, 1980), and grow with their host haloes by accreting mass via mergers in a hierarchical assembly of structure formation, giving rise to rich clusters of galaxies in the local Universe ($z < 1$) with cluster masses of the order $10^{15}\,\mathrm{M}_{\odot}$ (Marriage et al., 2011; Maughan et al., 2012).

There are three independent, concrete pieces of evidence that support the ΛCDM paradigm:

- **The Cosmic Microwave Background (CMB) angular power spectrum of temperature fluctuations.** A large collection of ground- and space- based experiments (COBE: Smoot et al. 1992; WMAP1: Spergel et al. 2003; CBI: Readhead et al. 2004; VSA: Dickinson

et al. 2004; ACBAR: Kuo et al. 2004; BOOMERANG: Jones et al. 2006; WMAP3: Spergel et al. 2007; WMAP5: Dunkley et al. 2009) have measured the autocorrelation function $C(\theta)$ as a function of angular separation θ:

$$\begin{aligned} C(\theta) &= \left\langle \frac{\delta T}{T}(\vec{n}_1)\frac{\delta T}{T}(\vec{n}_2) \right\rangle \\ \theta &= \cos^{-1}(\vec{n}_1 \cdot \vec{n}_2). \end{aligned} \tag{1.46}$$

The unit vectors $\vec{n}_1$ and $\vec{n}_2$ represent random directions in the sky, T refers to the mean temperature of the CMB measured as $\sim$2.73 K (Smoot et al., 1992), and the average is taken over all independent pairs of temperature fluctuations separated by an angle θ. The ΛCDM model is fundamentally characterized by six cosmological parameters, $\{\Omega_b, \Omega_c, \Omega_\Lambda, \Delta^2_{\mathcal{R}}, n_s, \tau\}$, which are respectively the fraction of the total present day energy density in baryons, cold dark matter, and the cosmological constant fluid phase, the amplitude of the curvature perturbations, the spectral index of the primordial power spectrum, and the optical depth of the epoch of cosmic reionization. Unless explicitly stated otherwise, it should be assumed throughout this thesis that any cosmological parameter refers to its present day value.

Table 1.1 (copied from Dunkley et al. 2009) shows the mean values and associated 1σ errors of both these standard ΛCDM parameters and the other cosmological parameters that have been fit to the angular temperature power spectrum, which now extends across a range of multipole moments from $\ell = 2$ to $\ell \sim 10\,000$ (Dunkley et al., 2009; Das et al., 2011). Table 1.1 indicates that the CMB data is fully consistent with the ΛCDM model: the matter content is dominated by cold dark matter, the energy density is mostly attributed to the cosmological constant, and the initial power spectrum is close to scale-invariant.

- **The power spectrum of density fluctuations as measured by the SDSS and the 2dFGRS sky surveys.** The Two Degree Field Galaxy Redshift Survey (2dFGRS) is a map with an area of $\sim$2000 deg^2, divided into two strips that span across the northern and southern galactic pole, whose entire data sample contains the photometry, spectroscopy and redshifts of around 250 000 galaxies (Colless et al., 2001). Several authors have computed the galaxy-galaxy clustering correlation function of these sample members, and in particular, have detected large-scale baryon oscillations in the power spectrum for wavenumbers k in the range $0.01 < k/(h\,\mathrm{Mpc}^{-1}) < 1$ (Percival et al., 2001; Tegmark et al., 2002; Cole et al.,

ΛCDM Model Parameters and 68% Confidence Intervals From the Five-Year *WMAP* Data Alone

Parameter	Three-Year Mean	Five-Year Mean	Five-Year Max Like
$100\Omega_b h^2$	2.229 ± 0.073	2.273 ± 0.062	2.27
$\Omega_c h^2$	0.1054 ± 0.0078	0.1099 ± 0.0062	0.108
Ω_Λ	0.759 ± 0.034	0.742 ± 0.030	0.751
n_s	0.958 ± 0.016	$0.963^{+0.014}_{-0.015}$	0.961
τ	0.089 ± 0.030	0.087 ± 0.017	0.089
$\Delta^2_{\mathcal{R}}$	$(2.35 \pm 0.13) \times 10^{-9}$	$(2.41 \pm 0.11) \times 10^{-9}$	2.41×10^{-9}
σ_8	0.761 ± 0.049	0.796 ± 0.036	0.787
Ω_m	0.241 ± 0.034	0.258 ± 0.030	0.249
$\Omega_m h^2$	0.128 ± 0.008	0.1326 ± 0.0063	0.131
H_0	$73.2^{+3.1}_{-3.2}$	$71.9^{+2.6}_{-2.7}$	72.4
$z_{\rm reion}$	11.0 ± 2.6	11.0 ± 1.4	11.2
t_0	13.73 ± 0.16	13.69 ± 0.13	13.7

Table 1.1: The best fit values of the cosmological parameters as measured by WMAP (this table has been copied from Dunkley et al. 2009).

2005; Percival et al., 2007). The complementary Sloan Digital Sky Survey (SDSS) is an even larger probe of large-scale structure (York et al., 2000). Eisenstein et al. (2005), for example, measured the comoving (s) two-point correlation function clustering $\xi(s)$ in redshift-space of $\sim$45 000 SDSS galaxies that span over 3000 deg^2, covering a volume of $\sim 0.7\,h^{-1}\,\mathrm{Gpc}^{-3}$ and a redshift range $0.16 < z < 0.47$. They reported the detection of a prominent baryonic acoustic peak at $s \sim 100\,h^{-1}\,\mathrm{Mpc}$, in agreement with the predicted local scale of the recombination-generated acoustic peaks as measured by WMAP. Recent work on galaxy clustering in the SDSS has also provided tight constraints on Ω_m, Ω_b and Ω_k (Tegmark et al., 2006; Percival et al., 2010), and the inferred values of these parameters are close to the values listed in Table 1.1. The combination of smaller scale WMAP measurements with the large-scale 2dFGRS and SDSS measurements of the density power spectrum yields excellent agreement with a ΛCDM cosmology over two decades in comoving wavenumber from $k \sim 0.001\,h\,\mathrm{Mpc}^{-1}$ to $k \sim 0.1\,h\,\mathrm{Mpc}^{-1}$ (Sánchez et al., 2006, 2009).

- **The magnitude-redshift Hubble relation of Type Ia supernovae.** Astronomers seek to measure the light curves of Type Ia supernovae (hereafter SNe Ia) as they are known to be standardizable candles, which when scaled by the width-luminosity relation (Phillips, 1993; Riess et al., 1996, 2004), yield estimates of H_0 (Hoeflich & Khokhlov, 1996; Riess et al., 1998, 2007, 2011). By probing high redshift SNe Ia one is hence able to measure the time evolution of the expansion of the Universe and infer values of $\Omega_{\rm M}$ and Ω_Λ.

Hamuy et al. (1996) were one of the first groups that attempted to measure the Hubble relation for SNe Ia at low redshift. Their sample consisted of ∼30 SNe Ia at $0.01 < z < 0.1$, taken from the Calán/Tololo supernovae survey (Hamuy et al., 1993). The pioneering work conducted by Perlmutter et al. (1999) then combined the Hamuy et al. (1996) sample with a sample of 42 SNe Ia in the redshift range $0.18 < z < 0.83$, detected using the Supernova Cosmology Project survey (Perlmutter et al., 1995, 1997). They found that the best-fit values of $\{\Omega_{\rm M}, \Omega_\Lambda\}$ to the B-band Hubble relation for a flat universe were roughly $\{0.28, 0.72\}$, in good agreement with both the corresponding ΛCDM values in Table 1.1 and earlier claims by Riess et al. (1998), who reported that their SNe Ia detections strongly disfavoured a flat or a closed universe with $\Omega_\Lambda = 0$ at the 7σ confidence level. Recent results released by the Carnegie Supernova Project (Freedman et al., 2009) and Supernova Legacy Survey (Conley et al., 2011; Sullivan et al., 2011) teams further constrain Ω_m, Ω_k, Ω_Λ and w (the dark energy equation of state parameter) and lend additional support to the ΛCDM model of our Universe.

With a standard model in place, recent years have witnessed a huge number of observational and theoretical experiments that test ΛCDM predictions for the origin and evolution of galaxy properties. A review of the most important ΛCDM 'anomalies' is now provided, mostly from the perspective of galaxy simulations, which forms the general subject heading of this thesis.

1.4 Establishing observed galaxy properties using simulations in a ΛCDM context

1.4.1 The galaxy luminosity function

Perhaps the most fundamental property of the galaxy population at a given epoch is the galaxy luminosity function, which represents the differential number density of galaxies per unit luminosity L at a given redshift z. Observations in the near-infrared K-band ($\lambda_{\rm eff}^{\rm K} = 2190\,{\rm nm}$) have revealed that the present day galaxy luminosity function in a ΛCDM cosmology obeys a Schechter function of the form $\phi(L/L_\star) \propto (L/L_\star)^\alpha \exp(-L/L_\star)$ (Cole et al., 2001; Kochanek et al., 2001), where $L_\star$ corresponds to an absolute K-band magnitude $M_\star(K) \sim -24$, and $\alpha \sim -1.37$ (Huang et al., 2003). Yet if one assumes a constant mass-to-light ratio for all galaxies and transforms the present day halo mass function (which is discussed in more detail in Chapter 3) into a predicted luminosity

function, there are too many low luminosity and high luminosity galaxies compared with the number observed, and agreement is found at just a single luminosity, corresponding to $L \sim L_\star$ (Benson et al., 2003). In order to bring the observed and mass-function-predicted luminosity functions into agreement, the ΛCDM picture requires a mass-to-light ratio with a luminosity index n_L in the optical B-band ($\lambda_\mathrm{eff}^\mathrm{B} = 445\,\mathrm{nm}$) whose value changes from $n_\mathrm{L}(B) \sim -1/2$ for $L < L_\star$ to $n_\mathrm{L}(B) \sim 1/2$ for $L > L_\star$ (Marinoni & Hudson, 2002). This relationship indicates a non-trivial suppression of star formation at both the low and high luminosity ends of the luminosity function, seemingly in conflict with the standard model of galaxy formation.

1.4.2 Cosmic downsizing

Another process that appears to be in tension with the hierarchical build-up of halo mass via mergers is 'downsizing' (originally coined by Cowie et al. 1996), which broadly falls into two categories: downsizing in the galaxy star formation rate and downsizing in the black hole accretion rate. The observational evidence for the former process is very strong. Brinchmann & Ellis (2000) analyzed a sample of ∼300 near-infrared K-band selected galaxies detected with the Hubble Space Telescope, and found that their dwarf galaxies experienced much higher specific star formation rates between $0 < z < 1$ than their ellipticals. Bauer et al. (2005) computed the specific star formation rates $\dot{n}_\star$ of the near-infrared K- and I-band ($\lambda_\mathrm{eff}^\mathrm{I} = 806\,\mathrm{nm}$) selected galaxies in the redshift range $0 < z < 1.5$. They found that low stellar mass galaxies had a higher $\dot{n}_\star$ than high stellar mass galaxies at every redshift bin in this interval, and further reported that these two star forming regimes are separated by a threshold mass that decreases towards the present day. Bundy et al. (2006) confirmed this behaviour and argued that this threshold mass M_t scales with redshift as $M_t \propto (1+z)^{3.5}$. A clear message from these studies is that the star formation in massive haloes appears to occur at high redshifts, and much of the present day stellar mass in massive haloes is already established by $z \sim 1$. This is further corroborated by the observation that most massive galaxies at low redshift ($z < 0.1$) are 'red and dead', outnumbering their blue star-forming counterparts by at least an order of magnitude for $M_\star \gtrsim 10^{11}\,\mathrm{M}_\odot$ (Baldry et al., 2004).

By contrast, AGN downsizing is far less well understood, although there appears to be strong evidence for a luminosity downsizing at X-ray wavelengths. Hasinger et al. (2005) selected ∼1000 soft X-ray (0.5–2 keV) Type-I AGN from the Chandra, ROSAT and XMM-Newton surveys, probing the redshift interval $0 < z < 4$, and found that the peak in both AGN number density and black hole

emissivity shifts to lower redshift for lower luminosity AGN. Barger et al. (2005) examined several hundred hard X-ray (2–8 keV) sources from the Chandra survey and argued that their reported drop in integrated hard X-ray luminosity density around $z \sim 1$ was driven by a downsizing in X-ray luminosity. Similar luminosity dependent density evolution in other wavebands, in particular radio wavebands, has proved somewhat less conclusive. Rigby et al. (2011), however, claim to have detected a downsizing in radio luminosity with a sample of ∼135 radio-loud AGN that span a similar redshift range to the X-ray detected downsizing signals ($0 < z < 4$), but the trend is far less convincing than at X-ray wavelengths, and only tentatively exists for the lower luminosity end of the radio-loud population in their sample ($P \leq 10^{26}$ W) at low redshift ($z < 1$).

1.4.3 The absence of cooling flows

It has now been convincingly demonstrated that the plasma in X-ray clusters cools via thermal bremsstrahlung radiation and line emission, and that the plasma density is higher and the temperature is lower in the core of clusters than in the outer regions (Cowie & Binney, 1977; Fabian & Nulsen, 1977). Cooling times of plasma clouds in X-ray cluster cores are known to be short, on timescales $\lesssim 1$ Gyr (White et al., 1997; Fabian et al., 2000; McNamara et al., 2000), hence in the absence of heating, a large mass of very low temperature material is expected to reside in the central region (few kpc) of local clusters ($z < 0.1$). Yet only a very small fraction of the total plasma gas is found to radiate at even modest temperature, and recent observations of clusters detected by XMM-Newton have demonstrated that there is a significant emission deficit and about an order of magnitude mass deficit of centrally concentrated low temperature plasma compared with simple theoretical predictions (Peterson et al., 2003). Although this is not a direct issue with the ΛCDM paradigm, any successful galaxy formation model incorporated within a ΛCDM context should account for the apparent absence of cooling flows in X-ray clusters.

1.4.4 The missing satellites problem

Klypin et al. (1999b) and Moore et al. (1999a) were among the first groups to show that simulations of Milky Way type galaxies, designed to test the ΛCDM model on galaxy scales, predict the existence of ∼300 present day dark matter subhaloes in the Local Group, roughly an order of magnitude more than the observed luminous satellite population (Mateo, 1998). This discrepancy is referred to as the 'missing satellites problem' (Moore et al., 1999a; Kamionkowski & Liddle, 2000).

1.4.5 The cusp-core problem

Numerous dark-matter-only N-body simulation studies performed in a ΛCDM cosmogony have demonstrated that the density of dark matter within the inner cores of dwarf galaxies ($M_\star \lesssim 10^9\,\mathrm{M}_\odot$) typically spanning $5 \times 10^{-4} \leq r/r_{\mathrm{vir}} \leq 5 \times 10^{-3}$ has a power-law scaling with radius of the form $\rho(r) \propto r^\alpha$, where $-1.5 \lesssim \alpha \lesssim -0.8$ (Moore et al., 1999b; Stadel et al., 2009; Navarro et al., 2010). These 'cuspy' distributions of dark matter are in tension with observations of the central kiloparsec of low mass galaxies, which predict a far shallower 'cored' index $-1 \lesssim \alpha \lesssim 0$ (Swaters et al., 2003; Simon et al., 2005; Donato et al., 2009). This discrepancy in the allowed range of α values is commonly referred to as the 'cusp-core' problem.

1.4.6 Modelling the physics of baryons within the ΛCDM paradigm

Given the baryonic nature of the processes described in Sections 1.4.1–1.4.5, it is clear that there are missing ingredients in a standard model that just describes the physics of dark matter. This subsection hence introduces the recipes that model the physics of luminous galaxy components.

At first glance, both forms of downsizing appear to be in conflict with ΛCDM predictions. If the star formation rate of luminous galaxies is positively correlated with their dark host's mass, then the largest star formation rates should be found in the most massive galaxies at a given epoch, and the rate of star formation in these massive galaxies should increase towards the present day as haloes become more massive, the opposite of what is observed in the optical band. Similarly if AGN mass scales with AGN luminosity, then the findings in Section 1.4.2 imply that at high redshift the most massive black holes are most active and their lower mass counterparts are most active at low redshifts, which is again the opposite to a bottom-up, merger-driven, hierarchical prediction of structure formation.

During the past twenty years, the proposed solutions to reconciling the behaviour of the galaxy luminosity function and downsizing in a ΛCDM framework have all been borne from various implementations of the technique of semi-analytic modelling. The semi-analytic recipe is as follows: construct a physically motivated model that is characterized by as few free parameters as possible, and tune these parameters to reproduce known galaxy observables (Baugh, 2006; Benson, 2010). These include the present day galaxy luminosity functions, the colour-magnitude bimodality distribution, the I-band Tully–Fisher relation (Tully & Fisher, 1977), the M_{BH}-M_{BUL} relation, and the cosmic star formation history of the Universe (Hatton et al., 2003; Cattaneo et al., 2006; Dekel

& Birnboim, 2006; De Lucia et al., 2006; Somerville et al., 2008). Most authors rely on executing their models on the Millennium Simulation run (Springel et al., 2005), which provides them with the necessary statistics. The challenge for these models is to provide physical mechanisms that are capable of heating or expelling gas in order to suppress star formation at both ends of the luminosity distribution and hence account for the observed deficit of low and high luminosity galaxies.

Cole et al. (1994) were one of the first groups to demonstrate that supernovae feedback was able to fit the flat, low luminosity end for each of K- and $b_{\rm J}$- bands (the term 'luminosity function' is henceforth made in reference to one or both of these bands at the present day). Benson et al. (2003) later developed these ideas and considered three supernovae-driven physical mechanisms capable of reproducing the luminosity function: disk reheating, thermal conduction of energy from ionized gas to the central region, and superwinds. Reheating of the disk provides a poor fit across the whole luminosity range unless a high disk heating efficiency of $\epsilon \sim 0.4$ is assumed, in which case the flat low luminosity behaviour is captured. Several orders of magnitude discrepancy still exist at the high luminosity end, however. Thermal conduction cannot recover the shape unless both an unphysically high conduction efficiency and a high disk heating efficiency are assumed. Superwinds are capable of reproducing the low luminosity distribution, but an energy injection that is greater than the amount of energy released by supernovae has to be assumed. More recently, Bower et al. (2006), Croton et al. (2006) and De Lucia et al. (2006) have found that feedback via AGN is able to sufficiently quench star formation in high luminosity galaxy systems, and predicts a number density in concordance with the observed exponential tail of the luminosity function. Yet these studies propose different models for the feedback; Croton et al. (2006) and De Lucia et al. (2006) describe a model that strongly correlates with halo mass and that requires black holes to accrete hot gas, whilst Bower et al. (2006) propose a self-regulated, Eddington-limited feedback loop with weaker dependence on halo mass and no dependence on the temperature of the accreted gas. While observational evidence in support of the nature of these accretion modes is lacking, evidence that favours AGN as the source of feedback has recently been mounting. Schawinski et al. (2007, 2009) examined low redshift samples ($z < 0.1$) from the SDSS at millimetre wavelengths and reported that molecular gas local to the accreting black holes disappeared on $\sim$100 Myr timescales, which these authors interpreted as an indication of a low level of feedback from a local AGN population.

The emerging picture, then, is that weak AGN feedback is able to heat molecular clouds and prevent star formation, which reduces the number of high luminosity objects and explains why

massive local ellipticals are not forming stars, therefore providing a solution to star formation rate downsizing in a ΛCDM cosmogony as well as yielding agreement with the exponential tail of the luminosity function. Accounting for observed AGN downsizing, however, still remains somewhat elusive and is a two-fold degenerate phenomenon, driven either by low mass black holes accreting at near-Eddington rates (Heckman et al., 2004) or by supermassive black holes accreting at low rates (Babić et al., 2007). Whether AGN heating is able to suppress cooling flows in clusters is also still questionable from the semi-analytic perspective, given the lack of observational evidence supporting the Bower et al. (2006) and Croton et al. (2006) models. A galaxy formation model that is able to both convincingly account for the absence of cooling flows in clusters and reproduce known galaxy properties (e.g. the luminosity function, the colour distribution of present day galaxies and the star formation history of the Universe) within the ΛCDM framework is hence lacking.

Two of the most outstanding issues facing the ΛCDM model at present, however, are the missing satellites problem and the cusp-core problem. Although these phenomena are not investigated in this thesis, their importance warrant separate attention. The challenge for any model that aims to resolve the missing satellites problem is not just one of luminosity abundance matching: the physical properties of the Local Group satellites also have known correlations, encapsulated by the local scaling relations between surface brightness, halo circular velocity, metallicity and stellar mass. Dekel & Woo (2003) examined the low mass dwarf satellites of the Local Group with observed stellar masses in the range $6 \times 10^5 \leq M_\star/\mathrm{M}_\odot \leq 3 \times 10^{10}$, and found that their prescription of supernovae feedback was able to reproduce the scaling relations, arguing that this mechanism was efficient at heating gas bound to haloes with circular velocity $v_c \leq 100\,\mathrm{km\,s^{-1}}$. Benson et al. (2002) investigated the importance of heating from the UV background and claimed fairly good luminosity abundance agreement at absolute optical V-band ($\lambda_{\mathrm{eff}}^{\mathrm{V}} = 551\,\mathrm{nm}$) magnitudes below -10, but had to assume that all of the UV ionizing photons were injected into the intergalactic medium in order to yield a realistic reionization epoch $z_{\mathrm{re}} \sim 8$. The most fundamental problem these models face, however, is that luminous satellites residing in haloes with dynamical mass as low as $M_{\mathrm{dyn}} \sim 10^7\,\mathrm{M}_\odot$ have been detected (Mateo, 1998), and the proposed mechanisms above make no such prediction. Kravtsov et al. (2004), whose work was motivated by this tension, has perhaps provided the most promising explanation of the missing satellites problem and the existence of low mass satellites. They traced the mass and circular velocity evolution of the subhaloes belonging to three present day Milky Way type objects and found that ∼10% were much more massive at their formation

redshifts ($z \geq 2$), and that the present day low mass dwarf spheroidals in particular experienced heavy tidal stripping during merger episodes and hence suffered the largest loss of mass and circular velocity. Upon implementing their model of galaxy formation, they found that the circular velocity functions and cumulative radial distributions of their satellites agreed far better with observations. These authors hence suggested that the initially massive spheroidals represent sites of intense star formation at high redshift, but are subsequently tidally stripped, explaining why their low mass present day counterparts still host luminous galaxies despite their weaker associated gravity.

From the perspective of the cusp-core problem, it has been argued that simply invoking warm, as opposed to cold, dark matter does not reduce the cuspiness of the density profile over the observed scales within the central kiloparsec region of low mass haloes (Knebe et al., 2002; Strigari et al., 2007; Colín et al., 2008). (The observational evidence disfavouring the Warm Dark Matter paradigm is presented later in this section.) It is generally agreed, however, that if dark matter particles are heated then their central concentration is reduced, and there are two popular heating sources discussed in the literature: dynamical friction (Tonini et al., 2006; Goerdt et al., 2010b) and supernovae feedback (Mo & Mao, 2004; Mashchenko et al., 2006). Focussing on the latter mechanism, Pontzen & Governato (2012) developed an analytical model of supernovae feedback whereby repetitive supernovae explosions induce local perturbations in the gravitational potential experienced by a dark matter particle, boosting its energy. Using this model, they found that a core of $\sim$1 kpc was generated in their resimulated halo (whose present day virial mass was $M_{\rm H} \sim 10^{10}\,{\rm M}_{\odot}$) between $2 \lesssim z \lesssim 4$. This work hence provides a promising solution to the long-standing problem of the distribution of dark matter within the central regions of low mass haloes.

1.5 Forming galaxy disks in a ΛCDM framework

Disks of gas residing at the centre of luminous galaxies have proved particularly difficult to simulate in recent years. This section hence shifts focus to understanding how disks acquire their angular momentum and their observed present day properties.

1.5.1 The origin of galaxy angular momentum

The understanding of the origin of angular momentum within galaxies dates back to the pioneering works of Hoyle (1951) and Peebles (1969), who modelled protogalaxies as spherical Eulerian patches and argued that torques are exerted on a given patch due to tidal interactions with neighbouring

patches. Doroshkevich (1970) removed the assumption of spherical symmetry and demonstrated that as dark halo perturbations follow the expansion of the background, their total angular momentum evolves with the scale factor as $J(a) \propto a^{3/2}$ in an Einstein–de Sitter universe ($\Omega_{\rm M} = 1$). An N-body simulation study by White (1984) later confirmed that this growth of angular momentum continues until the perturbation, modelled in a Lagrangian framework, breaks away from the background expansion and collapses to form a halo, whereupon its angular momentum remains approximately constant. White (1984) hence argued that the magnitude of the angular momentum acquired for a given halo is set at the epoch of maximum expansion. These calculations, however, were performed using linear perturbation theory and therefore only hold while the density contrast is linear ($\delta < 1$). Barnes & Efstathiou (1987) demonstrated that haloes containing embedded substructure often have more associated angular momentum than haloes without dense substructures, due to a more efficient coupling to the external tidal field. They further argued that the evolution of the subhalo angular momentum modulus is highly non-linear, and varying amounts of angular momentum are transported to the outer regions of the host as subhaloes sink to the central host region. N-body simulations provide a natural framework for non-linear experimentation, and the papers by White (1984) and Barnes & Efstathiou (1987) were amongst the first to confirm the linear tidal torque theory prediction of $J(a) \propto a^{3/2}$ for a linear density perturbation in an Einstein–de Sitter universe.

The theories of Rees & Ostriker (1977), White & Rees (1978) and Fall & Efstathiou (1980) expanded beyond the early predictions of tidal torque theory introduced by Peebles (1969), and consequently painted a simple picture of galaxy disk formation. Gas residing in the potential wells of collapsed density perturbations (whose depth is dominated by dark matter) is shock heated to the virial temperature of the halo and the inner gas regions subsequently cool and lose their pressure, sinking to the halo centre on a free-fall timescale. Galaxy disks hence form from the 'inside-out', as the infalling gas retains its angular momentum. Several authors have since reported that the mean stellar age decreases (de Jong, 1996; MacArthur et al., 2004; Gogarten et al., 2010) and specific star formation rates increase (Muñoz-Mateos et al., 2007) as a function of radial position from the disk centre, thereby lending support to the inside-out paradigm of disk growth. This model makes two important assumptions:

(i) The initial specific angular momentum distributions of gas and dark matter are equal.

(ii) The specific angular momentum of gas is conserved upon collapse to the central region.

The first assumption holds if gas and dark matter mix and are subject to the same torques. The second assumption is a natural consequence of the standard theory of galaxy formation, which assumes that infalling gas crosses an adiabatic shock at the halo virial radius and then cools, with no further shock heating episodes.

This model has two appealing features. Firstly, by making a prediction for the relationship between the disk scalelength R_d and the host halo virial radius $R_{\rm vir}$, it is able to recover locally observed distributions of spiral disk scalelengths and the I-band Tully–Fisher relation (Dalcanton et al., 1997; Mo et al., 1998; de Jong & Lacey, 2000). A simple calculation using the dimensionless dark matter spin parameter λ (Peebles, 1969) best illustrates this:

$$\lambda = \frac{J|E|^{1/2}}{GM^{5/2}} = \frac{j|E|^{1/2}}{GM^{3/2}} \propto \frac{jV_c}{M}, \tag{1.47}$$

where J, j, E and M are respectively the angular momentum, the specific angular momentum, the total energy and mass of a dark halo. It has been assumed that at the halo virial radius $R_{\rm vir}$, the tail of the velocity curve has converged to the circular velocity V_c, and that the density profile $\rho(r)$ of dark matter follows an isothermal distribution ($\rho(r) \propto 1/r^2$) such that $E \propto MV_c^2$. If assumptions (i) and (ii) above hold, then $j \propto j_{\rm gas}$ and $j_{\rm gas} = j_d$, where $j_d = R_d V_c$ is the specific angular momentum of the disk, whose scalelength is R_d (it is also assumed that the velocity at the disk's scalelength is equal to V_c). Equation (1.47) hence simplifies to:

$$\lambda \propto \frac{R_d V_c^2}{M} \propto \frac{R_d}{R_{\rm vir}} \Rightarrow R_d \propto \lambda R_{\rm vir}. \tag{1.48}$$

The distribution of disk scalelengths hence depends on the distribution of halo spin parameters, which is approximately log-normal with median $\lambda \sim 0.05$ (Barnes & Efstathiou, 1987), and the distribution of $R_{\rm vir}$. The second success for this model is the discovery of the remarkable property that the cumulative fraction of mass within a halo that has specific angular momentum less than j, denoted by $M(< j)$, is well fit by a universal function that depends on a single free parameter, which holds for haloes of all masses (Bullock et al., 2001).

1.5.2 Reproducing observed disk properties

Both the disk scalelength distribution and the fitting function for $M(< j)$ rely on the validity of the assumption that gas conserves its specific angular momentum upon infall. Earlier work

by Navarro & Benz (1991) challenged this assumption as these authors found that upon cooling, the gas fragments into a clumpy distribution of gas clouds and dynamical friction between these clumps transports some of their angular momentum to the host halo. Subsequent simulation studies by Navarro & White (1994) and Navarro et al. (1995b) therefore removed the requirement that gas conserves its specific angular momentum and consequently experienced enormous difficultly in replicating observed disk scalelengths in a ΛCDM context, predicting disks at the present day whose scalelength was an order of magnitude too small. Evidently the gas in these simulations is losing too much angular momentum to the dark matter host. During the past twenty years, several groups have sought various physical remedies to this 'disk angular momentum problem' and the general issue of replicating known galaxy properties. The basic aim is to simultaneously reproduce observed baryonic spin parameter distributions of dwarf and disk galaxies, exponential surface brightness profiles, the slope and normalization of local spiral Tully–Fisher I-band relations, and local distributions of disk scalelengths (Abadi et al., 2003; Sommer-Larsen et al., 2003; Governato et al., 2004; Okamoto et al., 2005; Scannapieco et al., 2009). To date, there have been essentially four proposed explanations that could in principle satisfy these constraints, all under the assumption that the problem lies within the simulations:

- **Unrealistic treatment of the gas density distribution**

  Dynamical friction between merging gas clumps is very efficient at transporting gas angular momentum to the host haloes, and the stars that form within the clumps merge to form a strong bulge component (Navarro & Benz, 1991; Navarro & Steinmetz, 2000). With a smoother gas density distribution, star formation will be suppressed at high redshift and the gas will retain more of its angular momentum, leading to a more extensive disk.

- **Incorrect standard model of cosmology**

  One way to reduce the excessive transfer of angular momentum to the host is to smooth the gas density via heating. Sommer-Larsen & Dolgov (2001) were able to reproduce the observed local I-band Tully–Fisher relation by invoking warm dark matter (hereafter WDM). This result is a little unsurprising because WDM only differs from CDM on small scales due to relativistic free streaming, which smoothes the density and quenches small scale power (Bardeen et al. 1986 showed that the ratio of transfer functions is of the form: $T_{\rm WDM}/T_{\rm CDM} \propto \exp\left(-x-x^2\right)$, where x depends on the mass of the warm dark matter particle $m_{\rm WDM}$). The corollary of

this is that the number density of low mass haloes is quenched compared with a CDM model, and the extent of the quenching is ultimately determined by $m_{\rm WDM}$.

WDM models are generally disfavoured, however. Recent observations conducted by Viel et al. (2008) demonstrate that the Lyman-α flux power spectrum at multiple redshifts in the range $2.5 < z < 5.5$ is consistent with a warm dark matter particle mass of $\sim$8 keV, which despite being significantly warmer than the cold dark matter particle mass of $\sim$1 GeV, rules out any sterile neutrino WDM candidate (Seljak et al., 2006) by being above the allowed upper limit of $\lesssim$8 keV imposed by the observed deficit of X-rays emitted from neutrino decays (Abazajian et al., 2001; Boyarsky et al., 2006). The epoch of reionization is also delayed in the WDM paradigm and is largely inconsistent with the $z_{\rm re} \sim 11$ value that best fits the WMAP CMB temperature power spectrum (shown in Table 1.1).

- **Errors in the numerical methods**

  There are currently two schemes that are used to follow the formation of luminous galaxies within haloes. The first is a Lagrangian, particle-based Smoothed Particle Hydrodynamics treatment (Monaghan & Lattanzio, 1985), hereafter SPH, and the second is an Eulerian adaptive grid-based treatment (e.g. Kravtsov et al. 1997). Both schemes have their relative merits, which are summarized in Table 1.2 and are now discussed in turn.

  Galilean invariance. Galaxies often travel with large velocities of the order several hundred kilometres per second with respect to the background. Invariance of the solutions to the equations of motion (which are discussed in Chapter 2) in the presence of bulk flows is hence an important constraint that any given galaxy evolution code should satisfy in order to reliably capture the time evolution of systems with large bulk velocities. A notable flaw of grid codes is that they are not Galilean invariant (e.g. Tasker et al. 2008; Wadsley et al. 2008), as their numerical errors depend on the relative velocity of the fluid with respect to the grid. SPH codes, on the other hand, are Galilean invariant by construction (e.g. Springel 2010). The importance of Galilean invariance was recently highlighted by Tasker et al. (2008), who gave their simulated cluster a bulk velocity that returned the cluster back to its original starting position after 1 Gyr. They found that in the grid simulations the inner core region ($r \lesssim 1$ kpc) of the cluster was offset by $\sim$100–200 kpc from its starting location after 1 Gyr, whereas this effect was absent in the particle simulations.

Table 1.2: A brief summary of the important relative advantages and disadvantages of particle-based Lagrangian schemes versus Eulerian grid-based schemes in simulating galaxy formation. The letters 'B' and 'W' correspond to 'better' and 'worse' than the alternative scheme.

Property	Particle	Adaptive Grid
Galilean invariance	B	W
Resolving fluid instabilities	W	B
Gravitational softening	W	B
Artificial viscosity	W	B
Artificial angular momentum transfer	W	B

Resolving fluid instabilities. There are two common instabilities that feature in galaxy simulations: the Kelvin-Helmholtz instability and the Rayleigh-Taylor instability. The former arises wherever there is velocity shear at the boundary separating two fluids. This occurs when satellites travel through their host clusters, for example, with Kelvin-Helmholtz instabilities forming at the boundary between satellite gas and intracluster gas (e.g. Mori & Burkert 2000). Rayleigh-Taylor instabilities arise when a low density fluid accelerates a high density fluid. An example of this effect is the propagation of a high density supernova remnant through the interstellar medium (e.g. Warren et al. 2005; Fraschetti et al. 2010). Agertz et al. (2007) have demonstrated that fluids in grid simulations naturally mix and that grid codes are able to resolve both of these instabilities, unlike standard implementations of SPH, which suffer from the infamous pressure 'blip' problem (e.g. Price 2012). Price (2008), however, has argued that models of artificial conductivity (combined with artificial viscosity, which is discussed in more detail below), lead to a more efficient mixing between fluid layers that largely damps the pressure blip. This scheme is hence a promising mechanism for resolving instabilities in SPH simulations.

Gravitational softening. A common process that plagues any particle-based representation of dark matter, stars and gas, is two-body heating (Hernquist & Barnes, 1990; Pfenniger & Friedli, 1993): particles travelling along orbits that pass sufficiently close to one another undergo large angle scattering, thereby inducing perturbations in the gravitational potential, which causes particle heating. The energy of a particle is therefore no longer a constant of motion. This random heating of particles can induce artificial accelerations, which themselves induce spurious increases in the random velocities of the particles affected. Such random velocities inevitably reduce the angular momentum transported by the particle and alter the temperature dependent cooling function. Two-body heating can hence result in gas angular

momentum loss (transferred to the dark matter halo) that is purely an artefact of the discrete nature of particle sampling, and only vanishes in the limit of an infinite number of particles. Modern day techniques strive to suppress two-body relaxation interactions by smoothing each particle's mass over a softening scale, in attempt to add continuity to the particle density field and prevent gravitational divergence in close encounters. This scale therefore represents the spatial resolution limit, and the choice of fixed value that best achieves the trade-off between spatial resolution and noise has been the subject of much debate (Merritt, 1996; Athanassoula et al., 2000; Power et al., 2003), with recent studies opting for a spatially adaptive scale that depends on local environment density (Iannuzzi & Dolag, 2011). Note that this softening length (for the stellar and dark matter particles) is smaller in grid-based treatments and is set by the minimum cell width.

Artificial viscosity. Particle codes include an ad-hoc dissipative artificial viscosity term in the equation of motion in order to correctly capture shocks (e.g. Agertz et al. 2007; Mayer et al. 2008). This is not required in grid codes, but artificial viscosity still features here too due to the errors associated with interpolating fluid properties from cell centres to cell boundaries.

Artificial angular momentum transfer. One source of the loss of angular momentum of gas in particle simulations is two-body heating, which is largely circumvented by the introduction of gravitational softening techniques that suppress encounters below a fixed scale, as described above. However, dark matter particles are typically more massive than gas particles and so artificial angular momentum loss can still occur, even with softening: it seems that the particle graininess is one of the main sources of loss in angular momentum modulus of the gas component. Governato et al. (2004) and Kaufmann et al. (2007) examined isolated disk galaxies and found that at least 10^5 dark matter particles, 10^6 gas particles and softening scales below 1 kpc are required to prevent severe disk angular momentum losses. Yet even with their best model (sub kpc softening, 10^5 dark matter particles and 5×10^6 gas particles), Kaufmann et al. (2007) report a loss in the initial angular momentum of disk particles of more than 10% over a 5 Gyr period.

Another source of angular momentum loss of the gas component in SPH simulations is artificial viscosity, which also arises in grid simulations (as described above), although the magnitude of this effect is much lower than in standard SPH simulations where ad-hoc inser-

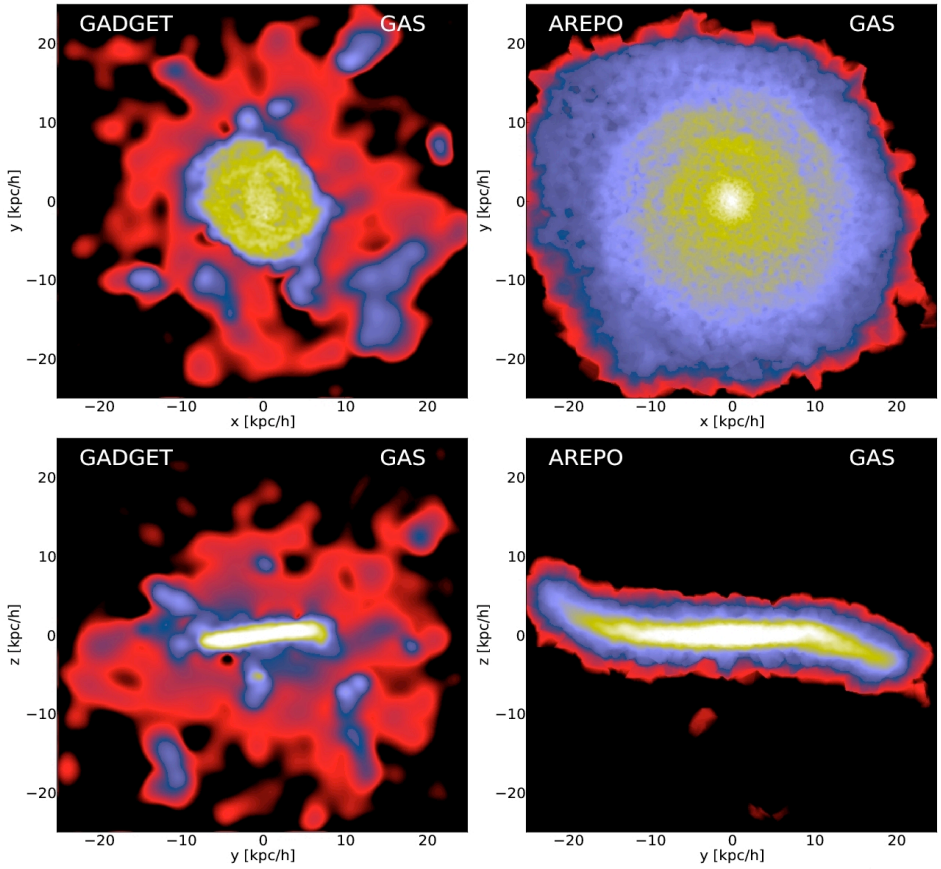

Figure 1.1: These images have been taken from Keres et al. (2011) and illustrate face-on (z projections, first row) and edge-on (x projections, second row) views of the gas disk at $z = 0$ from the same galaxy resimulated using the particle code GADGET (left column) and the grid tessellation code AREPO (right column). The disk scalelength resulting from the grid treatment is approximately a factor of two more extended than found by using the particle treatment.

tion of artificial viscosity introduces dissipation at all spatial locations. This is demonstrated in Fig. 1.1, taken from Keres et al. (2011), which compares the performance of AREPO (an adaptive Lagrangian tessellation grid code) and GADGET (an SPH code) at producing extended disks. Both simulations resimulate the same 24 galaxies in a $20\,h^{-1}\,\mathrm{Mpc}$ comoving periodic box with identical initial conditions but differing schemes for solving the Euler equations (which are discussed in Chapter 2). Clearly the adaptive grid code AREPO yields more extended disks, owing to the reduced amounts of artificial viscosity. Similar results were found by Commerçon et al. (2008), who simulated the collapse of a protostellar cloud using the adaptive mesh code RAMSES (Teyssier, 2002) and the SPH code DRAGON (Goodwin et al., 2004).

It should be noted, however, that some sophisticated models of artificial viscosity have recently

been proposed (e.g. Cullen & Dehnen 2010) that are particularly effective at only introducing dissipation at the locations of shocks. These viscosity 'switches' (see, for example, Price 2008) may yield disks with similar scalelengths to those found in grid simulations.

- **Missing physics in standard cosmological simulations**

  As argued by Agertz et al. (2011), the angular momentum problem implies that star formation in molecular clouds is too efficient at high redshift. Physical mechanisms (alternatives to WDM) that can heat the interstellar medium, quench star formation and smooth the gas density have therefore been the subject of much debate in recent years. The candidates that have received the most attention are: (i) heating from a UV background established at the epoch of reionization; (ii) supernovae feedback, and; (iii) feedback from quasars.

  The UV background has been studied in two redshift regimes corresponding to $z \lesssim 3$ and $z \sim 10$. For the former regime, it has been argued that UV photoheating reduces the fraction of gas that collapses within haloes whose circular velocities approach $v_c \sim 75\,\mathrm{km\,s^{-1}}$ (Thoul & Weinberg, 1996). Gnedin (2000) later reported that this collapsed gas fraction decreases towards a characteristic mass scale that decreases at higher redshift, implying that UV heating becomes less effective at gas removal at early epochs. Dijkstra et al. (2004) supported these claims by demonstrating that the velocity threshold is reduced to $v_c \sim 10\,\mathrm{km\,s^{-1}}$ at $z \sim 10$. It hence appears that UV heating, by itself, is not the solution to accurate disk formation.

  Feedback from supernovae (quasar feedback is not discussed in the following) is probably the most popular missing ingredient that is thought to help yield realistic disks, but due to the sub-parsec nature of star formation and the kpc scale of galaxies, several issues have arisen with its implementation. State of the art numerical simulations that have started to probe resolution scales of a few parsecs have resorted to 'sub-grid' phenomenology in order to model star formation and supernovae feedback (discussed later in Section 2.1.5), and it is the difference in the chosen prescriptions that leads to the conflicting conclusions found upon implementing these models. An early SPH study of disk formation by Katz (1992), for example, argued that energy from supernovae was quickly radiated away on a cooling timescale far below the dynamical timescale for redshifts $1 < z < 4$. These claims were then contested by Maller & Dekel (2002), whose supernovae feedback model parametrized the size of the satellite halo region occupied by baryons and the fraction of gas that remains in the

halo post-supernova. They found that their prescription sufficiently expelled the angular momentum of baryonic material at the low end of the distribution, yielding agreement with the observed dwarf baryonic spin distributions. More recently, Dutton & van den Bosch (2009) introduced an energy- and momentum- driven mass ejection scheme that varies with distance to the gas disk centre. Their models were characterized by two efficiency parameters $\epsilon_{\rm E}$ and $\epsilon_{\rm M}$ that account for the fraction of released (kinetic) energy and momentum that couples to the supernovae winds, as well as an additional parameter that indicates the number of supernovae that are produced for every solar mass of stars formed. They found that including the energy- and momentum- driven supernovae winds produced relations between disk scale sizes, stellar masses and rotation velocities in concordance with those observed, but their models required unphysical values of either 0.25 or 1 for $\epsilon_{\rm E}$ and $\epsilon_{\rm M}$ respectively to reproduce this agreement.

It is hence clear that the hypothesis that supernovae feedback can shape disk properties is highly sensitive to the model adopted to parametrize the process. Selective gas heating seems unavoidable too, as bulges form from low angular momentum gas. Indeed Scannapieco et al. (2009) conducted an analysis on eight simulated Milky Way type galaxies and found that each system consisted of a disk that was deficient in mass, and that disks were easily perturbed by major mergers, which induced mass transfer to the bulge component.

In summary, the solution to the angular momentum problem in particle and grid simulations probably ultimately depends on spatial and particle resolution. Increasing the spatial resolution will help abolish the need for semi-analytic prescriptions of star formation and supernovae feedback, and will reduce the numerical errors associated with time-evolving averaged physical quantities over a cell volume in grid simulations, while increasing the number of particles in particle simulations will remove numerical artefacts that contribute to angular momentum loss from the baryons.

1.6 The aims and layout of this thesis

Results from two galaxy simulation projects are presented in this thesis. The aim of the first project is to monitor the growth of dark matter haloes as they are accreted onto dense groups and clusters in a bottom-up hierarchical context within an N-body, dark-matter-only simulation, and to account for any possible accretion rate dependence on environment. By not simulating the luminous component, the issues in Section 1.4 are ignored and the fundamental component of

galaxies, which is likely to shape their luminous properties, is examined. The aim of the second project is to determine how a Milky Way type disk at high redshift acquires its mass and angular momentum. The analysis is conducted on a single high redshift system traced down to $z = 3$, which, when run to $z = 0$ at coarser physical resolution, boasts a disk scalelength and height that are very similar to the measured present day Milky Way values. This holds for the runs that include and exclude a prescription for supernovae feedback.

A brief scientific motivation for both of these projects is now presented.

1.6.1 Project 1: Accretion onto structures in dense environments

Several authors have argued that dense dark haloes often contain embedded subhaloes and have demonstrated that low mass subhaloes can survive in their hosts for several billion years (Tormen, 1997; Tormen et al., 1998; Moore et al., 1999a). One challenge for cosmological N-body simulations is to link dark matter haloes and subhaloes with luminous galaxies (Bower et al., 2006; Conroy et al., 2006; Vale & Ostriker, 2006). Understanding this relationship has proved difficult (Diemand et al., 2004; Gao et al., 2004; Nagai & Kravtsov, 2005) and most explanations are provided by semi-analytic models (White & Frenk 1991; Somerville & Primack 1999; Hatton et al. 2003; Bower et al. 2006; Cattaneo et al. 2006; Croton et al. 2006). Nonetheless, a vital ingredient in explaining luminous galaxy growth in large groups and clusters is an understanding of how dark matter haloes and subhaloes accrete mass in dense environments. The growth and survival of dark matter haloes—the most fundamental components of galaxies—across their accretion history will probably ultimately help shape the galaxy properties highlighted in Section 1.4.

In this study, outputs from a high resolution dark-matter-only N-body simulation have been used and a new robust method for defining accretion onto haloes and subhaloes is provided, building on previous simulation studies and moving beyond the analytic Extended Press–Schechter calculation. The primary aim is to investigate exactly how accretion onto haloes and subhaloes behaves as a function of redshift, mass and environment. A derivation and explanation of the analytic tool that can be used to predict the mass growth of a halo at a given time—the Extended Press–Schechter formalism—is presented in Chapter 3. The halo and subhalo accretion algorithm, which extends beyond both this analytic calculation and the algorithms from previous simulation studies, is also introduced in Chapter 3. Chapter 4 then presents the results found by implementing this algorithm

and includes a detailed discussion of their implications for: a) subhalo growth in dense environments; b) the Extended Press–Schechter formalism of halo growth, and; c) black hole growth and possible accretion rate downsizing.

1.6.2 Project 2: Transporting angular momentum to a Milky Way disk

The past decade has witnessed a resurgence in tackling one of the defining characteristics of the standard theory of disk formation: that gas crossing a halo's virial sphere is shock heated to the virial temperature of the halo, which typically corresponds to X-ray temperatures of a few 10^6 K. Binney (1977) was the first to claim that, contrary to the standard paradigm, high density gas in low mass haloes (i.e. most haloes at high redshift) probably crosses an isothermal rather than an adiabatic shock front at the host halo virial radius, and hence argued that the vast majority of gas is expected to stream in towards the disk in cold flows with temperatures of a few 10^4 K. Birnboim & Dekel (2003) reignited this idea by demonstrating that for haloes less massive than $\sim 10^{11.6}\,\mathrm{M}_\odot$, gas cools efficiently, loses some of its pressure, and is unable to support a virial shock. It was hence concluded that for haloes of mass $M_\mathrm{H} \lesssim 10^{11.6}\,\mathrm{M}_\odot$, gas flows to the disk in a cold mode as opposed to a hot mode. The lack of redshift evolution in this threshold mass and the penetration of cold streams to the disk boundary in most galaxies was later confirmed by several numerical simulation studies (Kereš et al., 2005; Ocvirk et al., 2008). The existence of a bimodal gas accretion phase has enormous implications for star formation, and may, as the above authors have speculated, account for the observed lack of soft X-ray flux and copious Lyman-α emission in high redshift galaxies (Birnboim & Dekel, 2003), the bimodality in galaxy colours and its relation to galaxy morphology, including the existence of red galaxies at $z > 1$ and massive bursts of star formation at $z \sim 2$–4 (Dekel & Birnboim, 2006), and possibly star formation rate downsizing (Kereš et al., 2005; Ocvirk et al., 2008). These are hence very exciting times for galaxy disk theory.

Having established that gas inflow in most haloes at high redshift proceeds via cold flows that often terminate in the outskirts of gaseous disks (Brooks et al., 2009; Powell et al., 2011), one naturally wonders about the processes governing the transport of angular momentum onto disks within these overdense regions. The second project presented in this thesis attacks this question, and by resolving the multiphase nature of the interstellar medium, it joins a series of papers (Kimm et al., 2011; Pichon et al., 2011; Danovich et al., 2012) that have helped to construct a picture of angular momentum transport in high redshift galaxies, from supergalactic scales of the order Gpc

to halo virial regions of the order kpc. Chapter 5 describes the methods that have been employed to: a) compute the orientation plane of a resimulated Milky Way type disk; b) resolve the individual gas filaments and satellites that merge onto it, and; c) quantify the angular momentum locked-up in each of these galaxy components. A discussion of the implications of the results obtained upon application of these routines is then provided in Chapter 6. Chapter 7 follows with a short summary of the overall conclusions of this thesis and a presentation of several possible future projects, including an introduction to an algorithm that has recently been developed to monitor the time evolution of the angular momentum transported along filaments from large-scale virial regions to the point of disk contact.

Before proceeding to the analysis chapters, however, an explanation of the numerical methods and techniques encoded in the simulations, which ultimately form the source of the results of this thesis, is essential, and is the subject of the following chapter.

Chapter 2

The simulations

The ultimate aims of this thesis are to answer the following two questions: how do the accretion rates of haloes and subhaloes vary as a function of environment, and how does a Milky Way type disk acquire its angular momentum at high redshift? These questions have been tackled by analyzing outputs from one of the HORIZON[1] and one of the NUT simulations (Powell et al., 2011).

The full HORIZON suite consists of three separate projects. The first is called the HORIZON-4II run (Teyssier et al., 2009), which refers to a collection of outputs from a pure dark matter simulation containing 4096^3 particles in a $2\,h^{-1}\,\mathrm{Gpc}$ periodic box, generated with the RAMSES code (Teyssier, 2002). The second contains outputs from smaller pure dark matter simulations with fewer particles but enhanced spatial resolution, run with both the RAMSES and GADGET-II (Springel, 2005) codes. The third is called the HORIZON-MARENOSTRUM project (details given in Ocvirk et al. 2008, Dekel et al. 2009 and Devriendt et al. 2010), a hydrodynamic simulation that uses the RAMSES code. In this thesis, one of the pure dark matter GADGET-II simulations is analyzed and 'the HORIZON simulation' is henceforth made in reference to this run alone. The NUT suite also consists of an array of different runs, but unlike the HORIZON suite each set of simulation outputs is generated using the same code—RAMSES.

The first section of this chapter introduces the basic physics encoded within the simulations used in this study, and the second discusses the algorithms that have been implemented to detect haloes and subhaloes. The third section concludes the chapter with a brief summary of the technical specifications of each of the simulation runs.

[1] http://www.projet-horizon.fr

2.1 Simulating the formation of galaxies with the GADGET-II and RAMSES codes

There are currently two different types of adaptive mesh refinement (hereafter AMR): tree-based methods (e.g. ART: Kravtsov et al. 1997) that subdivide parent cells into child cells and introduce the concept of a threaded tree structure, and less efficient patch-based methods (e.g. ENZO: Bryan & Norman 1997; Bryan et al. 2001) that refine cells and their neighbours in rectangular patch regions. RAMSES is a tree-based code, and throughout this section there are several references to the refinement levels of its tree-grid data structure. This section hence starts by defining this adaptive grid, which is very similar to the grids used in other tree-based AMR codes.

2.1.1 The adaptive grid in RAMSES

In its simplest form, a galaxy system has three main constituents: dark matter, stars and gas. RAMSES simulates the gas component by defining a Cartesian grid-based tree structure with varying levels of refinement ℓ, as described in Teyssier (2002). The coarsest grid is given by $\ell = 0$, the base of the tree structure, and higher ℓ correspond to progressively smaller physical scales and hence higher spatial resolutions. The cells in the grid are arranged into octs, with a single oct containing 2^{N_d} cells, where N_d is the number of grid dimensions. A grid at level ℓ contains a maximum of 2^{ℓ} cells (and hence $2^{\ell-1}$ octs) along each dimension. An oct at a given ℓ is linked to its single parent host cell at $\ell - 1$, and to its 2^{N_d} child octs at $\ell + 1$. The 'end nodes' of the tree are defined as child octs which do not contain further child octs, signifying the highest level of refinement. By using the above scheme, the tree structure is defined at a given time-step. In order to compute the time evolution of physical properties like gas density and gas pressure (examples of 'flow variables'), RAMSES adopts an adaptive ℓ-based time sequencing algorithm: the higher the ℓ of the oct in the tree, the higher the assigned integer multiple of time-steps for that ℓ, per coarse time-step interval. For example, if the integer time-step scaling between two successive levels is 2, then for every coarse time-step there will be twice as many time-steps for an $\ell = 3$ oct compared with an $\ell = 2$ oct.

Having devised a scheme that links octs at different ℓ, RAMSES next computes a map of the cells that are to be refined. Any cell within an oct containing a child cell flagged for refinement is also flagged. However, RAMSES forces the flagged cells to contain $3^{N_d} - 1$ parent cell neighbours, so as to ensure a steady increase in resolution at successively higher ℓ. The refinement levels

are triggered once a flow variable (e.g. density) or a non-flow variable (e.g. mass), or their first and/or second derivatives, exceeds some pre-determined threshold (a free parameter). Random local fluctuations about the mean values of the flow variables can, however, exceed the threshold value, and so RAMSES employs a mesh smoothing technique to eliminate this random error. The refinement grid is hence modified. An 'attaching and snipping' method then decides which octs in the tree are to be refined: child octs belonging to cells marked for refinement are attached to the refinement grid, and those belonging to unflagged cells are cut from the tree.

With the grid and refinement map now defined, the remainder of this section examines each of the galaxy components in turn, starting with the dark matter and stellar particles and their mutual gravitational forces.

2.1.2 The N-body approach to solving the equations of motion of particle flow

Particle-based simulations consist of a finite number of particles N_p, each with mass M_p. The HORIZON simulation used in this study has 512^3 dark matter particles, and each particle roughly corresponds to a dwarf galaxy in mass ($M_p = 6.8 \times 10^8\,\mathrm{M}_\odot$). Since the number density of dark matter (stellar) particles in an N-body simulation is always far less than the true number density of haloes (stars) in the Universe (which itself is low), the frequency of two- and three- body interactions is negligible, and so a given particle's trajectory is well described by an effective gravitational potential due to all the other particles. It is hence assumed that all the particles in these simulations are collisionless (e.g. Springel 2005). Following Binney & Tremaine (2008), these collisionless particles with position and velocity vectors $\vec{x}$ and $\vec{v}$ are often considered in a phase-space $\vec{\omega} \equiv \{\vec{x}, \vec{v}\}$, with time evolution $\dot{\vec{\omega}} = \{\vec{v}, \dot{\vec{v}}\}$. Each particle in three dimensions is therefore represented by a point that flows with time in this six dimensional space. The flow described by $\dot{\vec{\omega}}$ conserves the number of particles present in the system and so the phase-space distribution function of particles $f(\vec{x}, \vec{v}, t)$ as a function of position, velocity and time obeys a continuity equation, which in the collisionless limit simplifies to the collisionless Boltzmann equation:

$$\frac{\mathrm{d}f}{\mathrm{d}t} = \frac{\partial f}{\partial t} + \vec{v} \cdot \vec{\nabla} f - \vec{\nabla}\Phi \cdot \vec{\nabla}_{\vec{v}} f = 0. \tag{2.1}$$

The effective gravitational potential $\Phi(\vec{x}, t)$ of the system is a solution of the Poisson equation:

$$\nabla^2 \Phi(\vec{x}, t) = 4\pi G \int f(\vec{x}, \vec{v}, t)\, \mathrm{d}^3 v. \tag{2.2}$$

The variables and operators in the two equations above, and in all of the equations throughout this section, are represented in a coordinate system that accounts for the expansion of the Universe. GADGET-II and RAMSES adopt the standard comoving and the 'supercomoving' coordinates (Martel & Shapiro, 1998), respectively. The latter choice is particularly attractive because all of the variables can be expressed in dimensionless units, and the fundamental equations—the collisionless Boltzmann equation and Euler equations (discussed later in this section)—are unchanged for a monatomic ideal gas with adiabatic index $\gamma = 5/3$ (the sole exception is the Poisson equation, however, which does acquire a slightly different form compared to its non-cosmological representation).

The coupling between equations (2.1) and (2.2) and the high dimensionality of f render the computation of Φ and f using standard finite difference schemes non-trivial (the explicit dependences of Φ and f have been dropped for brevity). As discussed in Springel et al. (2001b), the N-body solution to this problem, which is an approximation and is incorporated within the GADGET-II and RAMSES codes, involves a discrete sampling of f via use of particles, yielding an approximation to Φ. Equation (2.1) can be converted to an ordinary differential equation using the method of characteristics, and the characteristics are given by the time evolution of the particle positions ($\vec{x}_i$) and velocities ($\vec{v}_i$). The overall problem is hence reduced to solving Newton's equations for a system of particles whose dynamics are governed by their mutual self gravity (Hockney & Eastwood, 1981; Sellwood, 1987; Quinn et al., 1997; Springel et al., 2001b):

$$\frac{\mathrm{d}\vec{x}_i}{\mathrm{d}t} = \vec{v}_i \tag{2.3}$$

$$\frac{\mathrm{d}\vec{v}_i}{\mathrm{d}t} = -\vec{\nabla}\phi \tag{2.4}$$

$$\nabla^2 \phi = 4\pi G \rho, \tag{2.5}$$

where ϕ is the gravitational potential, G is Newton's gravitational constant and ρ is the density. These equations could be solved by direct summation, but this is an N_p^2 process (where N_p is the number of particles) and is hence inefficient. The following subsection hence discusses a more refined approach.

2.1.3 Particle Mesh and Tree-Particle Mesh techniques

The GADGET-II code adopts the Particle Mesh technique (Hockney & Eastwood, 1981; Klypin & Shandarin, 1983; Barnes & Hut, 1986) to solve equations (2.3), (2.4) and (2.5) when the separation r between particles is larger than a characteristic scale r_s, and uses the Tree-Particle Mesh method (Barnes & Hut, 1986; Hernquist, 1987) when $r < r_s$, for a more accurate force computation. By contrast, the RAMSES code uses the Particle Mesh technique for all separation scales.

Following Teyssier (2002), the Particle Mesh prescription for solving equations (2.3), (2.4) and (2.5) proceeds as follows:

1. Find the density on the mesh.
2. Solve the Poisson equation on the mesh.
3. Compute the acceleration on the mesh and interpolate to particle positions.
4. Use a numerical time integrating leapfrog-based scheme to solve for the particle velocity and position as a function of time.

In what follows, all of the comments made in reference to RAMSES and GADGET-II can be found in Teyssier (2002) and Springel (2005) respectively.

Step 1: Computing the mesh density

The density of a given cell $\{i, j, k\}$ in the mesh grid is computed using the Cloud-in-Cell interpolation technique originally developed by Hockney & Eastwood (1981):

$$\rho_{ijk} = \frac{1}{\Delta x \Delta y \Delta z} \sum_p m_p W(\vec{r}_p - \vec{r}_{ijk}). \tag{2.6}$$

The weighting function $W(\vec{r}_p - \vec{r}_{ijk})$ represents the fraction of the p^{th} particle's mass m_p, located at $\vec{r}_p$, that contributes to the density of cell $\{i, j, k\}$. In 3D, the weighting function is given by:

$$W(\vec{r}_p - \vec{r}_{ijk}) = \int_{z-\Delta z/2}^{z+\Delta z/2} \int_{y-\Delta y/2}^{y+\Delta y/2} \int_{x-\Delta x/2}^{x+\Delta x/2} S(x_p - x') S(y_p - y') S(z_p - z') \, \mathrm{d}x' \, \mathrm{d}y' \, \mathrm{d}z', \tag{2.7}$$

where $\{x, y, z\}$ denotes the three Cartesian components of cell $\{i, j, k\}$ and $\Delta x, \Delta y$ and Δz refer to the size of cell $\{i, j, k\}$ along each dimension. $S(x)$ is the particle shape function, which controls

the fraction of particle p's mass that is assigned to location $\vec{r}_{ijk}$. In the Cloud-in-Cell scheme, the shape function has the following form:

$$S(x) = \frac{1}{\Delta x} \begin{cases} 1 & \text{if } |x| < \frac{1}{2}\Delta x \\ 0 & \text{otherwise.} \end{cases} \tag{2.8}$$

According to equation (2.8), each particle can be thought of having an associated 'cloud' volume of $\Delta x \Delta y \Delta z$. If the centre of cell $\{i, j, k\}$ lies within particle p's cloud, then p's mass is averaged over the cell. The sum in equation (2.6) in GADGET-II is hence performed over all particle clouds containing the centre of cell $\{i, j, k\}$, whereas in RAMSES it is performed by looping over particles within the same oct at the same tree node. It is also possible that particle clouds at the previous coarser tree node cross into the finer oct, and so these particle contributions are included by RAMSES in the density computation too.

Step 2: Estimating the gravitational potential on the mesh

With an estimate of the density in each grid cell, the gravitational potential can now be found. Both RAMSES and GADGET-II (for large scales) adopt the discrete Fast Fourier Transform technique to compute the potential, as described by Hockney & Eastwood (1981). The potential in real-space $\phi(\vec{x})$ is expressed as a convolution of the real-space mesh density field $\rho(\vec{x})$ and a real-space Green's function for the potential $\mathcal{G}(\vec{x})$. By taking Fourier transforms and applying the convolution theorem, it follows that $\phi(\vec{k}) = \mathcal{G}(\vec{k})\rho(\vec{k})$. The computation of the potential is hence reduced to a simple multiplication and is performed in Fourier-space, and the real-space potential $\phi(\vec{x})$ is subsequently obtained by taking the inverse Fourier transform of $\phi(\vec{k})$.

In order to achieve a higher level of accuracy, RAMSES replaces the Laplacian by its $2N_d + 1$ point finite difference approximation, where N_d is the number of dimensions. In 1D, equation (2.5) becomes:

$$\phi_{i+1} - 2\phi_i + \phi_{i-1} = 4\pi G \rho_i (\Delta x)^2, \tag{2.9}$$

where Δx denotes the grid spacing. The Fourier transform of Equation (2.9) yields:

$$\left[\exp(ik\Delta x) - 2 + \exp(-ik\Delta x)\right] \phi(k) = -4\sin^2\left(\frac{k\Delta x}{2}\right) \phi(k) = 4\pi G \rho(k) (\Delta x)^2, \tag{2.10}$$

and so it follows that the expression for $\mathcal{G}(k)$ in 1D is:

$$\mathcal{G}(k) = -\frac{(\Delta x)^2 \pi G}{\sin^2(k\Delta x/2)}, \tag{2.11}$$

where $k = |\vec{k}|$. The Fast Fourier Transform technique cannot be used for the more refined levels of the grid because the grid is no longer regular (octs can contain child octs which themselves can contain further child octs etc). Equation (2.5) is hence solved using the Gauss–Seidel successive over-relaxation method for cells at higher levels of refinement than the coarse grid. Each Cartesian component of the $\vec{\nabla}^2$ operator, $\left[\frac{\partial^2}{\partial x^2}, \frac{\partial^2}{\partial y^2}, \frac{\partial^2}{\partial z^2}\right]$, is represented by its second-order central difference Taylor series approximation, and so the solution to equation (2.5) for the potential at site $\{i, j, k\}$ in a cube reduces to:

$$\phi^{n+1}_{ijk} = \frac{1}{6}\left(\phi^n_{i+1,j,k} + \phi^n_{i-1,j,k} + \phi^n_{i,j+1,k} + \phi^n_{i,j-1,k} + \phi^n_{i,j,k+1} + \phi^n_{i,j,k-1} - (\Delta x)^2 4\pi G \rho_{ijk}\right), \tag{2.12}$$

where the index n is an iterative index and an equal grid spacing of Δx has been assumed (note that the subscript k in equation 2.12 should not be mistaken for a Fourier mode). The value of ϕ^{n+1}_{ijk} is subsequently corrected to a new value by use of the over-relaxation method:

$$\phi^{n+1}_{ijk} = \alpha\phi^n_{ijk} + (1-\alpha)\phi^{n+1}_{ijk}. \tag{2.13}$$

The parameter α is approximated by the Press et al. (1992) expression:

$$\alpha = \frac{2}{1 + \beta\frac{\pi}{N}}, \tag{2.14}$$

where N is the number of cells along each dimension of the grid and β is a parameter whose optimum value depends on the choice of boundary conditions. The rate of convergence of equation (2.13) is sensitive to the initial choice of ϕ at a given site, ϕ^0_{ijk}. RAMSES approximates ϕ^0_{ijk} at ℓ to the converged value at $\ell - 1$.

The major disadvantage of the Particle Mesh method is that the accuracy of the force computation (described in the next subsection) is determined by the scale of the mesh, which is kept suitably large for quick application of the Fast Fourier Transform, rendering it unsuitable for small scale force calculations (Bouchet & Kandrup, 1985; Efstathiou et al., 1985; Bagla & Padmanabhan, 1997). GADGET-II therefore implements a scale-dependent scheme for computing $\phi(\vec{k})$: on large

scales the Particle Mesh technique is employed, but on smaller scales a tree code is used to resolve the particle distribution (Quinn et al., 1997). The gravitational potential in Fourier-space is hence decomposed into a long range $\phi_l(\vec{k})$ and a short range $\phi_s(\vec{k})$ component (e.g. Bode et al. 2000, Bagla 2002) using $\mathcal{G}(\vec{k}) = -4\pi G/(\vec{k}\cdot\vec{k})$ (this Green's function is also sometimes used in RAMSES for the coarse grid), which is found by applying the mapping $\vec{\nabla} \rightarrow -i\vec{k}$ to equation (2.5):

$$\begin{aligned}\phi(\vec{k}) &= \frac{-4\pi G\rho(\vec{k})}{k^2} \\ &= \frac{-4\pi G\rho(\vec{k})}{k^2}\exp\left(-k^2r_s^2\right) - \frac{4\pi G\rho(\vec{k})}{k^2}\left[1-\exp\left(-k^2r_s^2\right)\right] \\ &= \phi_l(\vec{k}) + \phi_s(\vec{k}),\end{aligned} \tag{2.15}$$

where:

$$\phi_l(\vec{k}) = -\frac{4\pi G\rho(\vec{k})}{k^2}\exp\left(-k^2r_s^2\right) = \phi(\vec{k})\exp\left(-k^2r_s^2\right) \tag{2.16}$$

$$\phi_s(\vec{k}) = -\frac{4\pi G\rho(\vec{k})}{k^2}\left[1-\exp\left(-k^2r_s^2\right)\right] = \phi(\vec{k})\left[1-\exp\left(-k^2r_s^2\right)\right], \tag{2.17}$$

and $k = |\vec{k}|$. A scalelength r_s that separates the short and long scale regimes is therefore introduced. GADGET-II uses the Particle Mesh technique (which is adopted for all scales in the RAMSES code) to compute the force on scales $r > r_s$. As discussed in Springel (2005), the long range component $\phi_l(\vec{k})$ for each particle is found by solving the Poisson equation for $\phi(\vec{k})$ and using the expression in equation (2.16). By contrast, the short range force $F_s(r)$ as a function of distance r from a given particle at location $\vec{r}_p$ has a direct analytic solution in real-space, and is computed using the particle tree approximation. By plotting both $F_s(r)$ and the total force $F_t(r)$ as a function of r, Bagla & Ray (2003) showed that masses further than $\simeq 5r_s$ from $\vec{r}_p$ contribute less than 1% of $F_t(r \gtrsim 5r_s)$. The short range force computation in GADGET-II is hence performed over neighbouring cells within a sphere of radius $\simeq 5r_s$ centred at $\vec{r}_p$. Neighbouring cells of size L_j within $5r_s$ at distance r_j from $\vec{r}_p$ are considered, and the following cell acceptance criterion (CAC) is enforced for the j^{th} neighbour:

$$r_j > \frac{L_j}{\theta}, \tag{2.18}$$

where θ is free parameter that determines the desired accuracy of the force computation. If equation (2.18) is satisfied, then cell j, which may contain several particles, is sufficiently far away that all

of its mass is assumed to act at a point (cell centre of mass), whereas it is further subdivided into a child oct if equation (2.18) is not satisfied. The cell subdivision process can, in principle, continue until each particle occupies its own cell, marking the end nodes of the tree mesh. The choice of the cell acceptance criterion varies in the literature (Barnes & Hut, 1986; Salmon & Warren, 1994; Springel et al., 2001a), but in all cases the number of particle contributions is far lower than the number of particles in the simulation.

The tree hence reduces the error ϵ in the total force computation for a given particle to $\epsilon < 2\%$ (Springel, 2005), and for the smallest r where direct summation over individual particles is performed, the error is zero (e.g. Bagla 2002). The disadvantage of the tree algorithm is that the normalization of its $N_p \log_{10}(N_p)$ scaling, where N_p is the number of particles, is a factor of 10–50 larger than that of the Particle Mesh code (e.g. Bode & Ostriker 2003), and so the tree method is a factor of 10–50 slower (Hernquist, 1987).

Step 3: Computing forces on the mesh and interpolating them to particle positions

The force computation on short scales in GADGET-II is performed in real-space, and so no interpolation to real-space particle positions is required. The acceleration $\vec{a}_{ijk}$ (and hence the force) at a given site on the mesh in RAMSES and on large scales in GADGET-II is computed via the relation $\vec{a}_{ijk} = -\vec{\nabla}\phi_{ijk}$. RAMSES uses a five-point finite difference representation of $\vec{\nabla}$ with an error of $\mathcal{O}(\Delta x)^5$, whereas GADGET-II adopts a four-point scheme with an error of $\mathcal{O}(\Delta x)^4$. The force vector $\vec{F}(\vec{r}_p)$ on a particle at position $\vec{r}_p$ is then computed using an inverse Cloud-in-Cell scheme with the same interpolation kernel used for the density computation in equation (2.6):

$$\vec{F}(\vec{r}_p) = \sum_{ijk} W(\vec{r}_p - \vec{r}_{ijk})\vec{F}_{ijk}. \tag{2.19}$$

Step 4: Following the particles across time

RAMSES uses a three stage 'kick-drift-kick' leapfrog integrator (Quinn et al., 1997) of the form $\mathrm{K}(\Delta t/2)\mathrm{D}(\Delta t)\mathrm{K}(\Delta t/2)$ to solve equations (2.3) and (2.4). The 'kick' step updates the particle velocity over a given time interval while holding the positions constant, and the 'drift' step does

the reverse. The K($\Delta t/2$)D(Δt)K($\Delta t/2$) method is a shorthand for the following set of equations:

$$\vec{v}_i^{n+1/2} = \vec{v}_i^n - \vec{\nabla}\phi^n \Delta t^n/2 \quad (2.20)$$

$$\vec{x}_i^{n+1} = \vec{x}_i^n + \vec{v}_i^{n+1/2}\Delta t^n \quad (2.21)$$

$$\vec{v}_i^{n+1} = \vec{v}_i^{n+1/2} - \vec{\nabla}\phi^{n+1}\Delta t^n/2, \quad (2.22)$$

where the time index for the i^{th} particle is given by the superscript n. It hence appears that the Poisson equation has to be integrated twice: once at t^n in equation (2.20) and once at t^{n+1} in equation (2.22). RAMSES avoids the Poisson-solver call at t^{n+1} by waiting until t^{n+1} before updating $\vec{v}_i^{n+1}$. As is further presented in Teyssier (2002), RAMSES also offers two schemes for choosing the time interval Δt. The first uses the same time Δt for all particles, with the advantage that Δt can be chosen to be the smallest required interval for all levels of refinement. Although accurate, this leapfrog integration is a slow computation and is unnecessarily inefficient in regions of low refinement level. The second scheme therefore adopts an individual time interval for each refinement level. This ℓ-based adaptive time sequencing algorithm (discussed in Section 2.1.1) is clearly less accurate than the single interval approach, but is quicker. The NUT simulation in this study uses both time sampling routines: the fixed time interval method is used for $7 \leq \ell \leq 10$ and the quicker ℓ-based method is used for $\ell > 10$.

GADGET-II adopts the KDK leapfrog integrator for large-scale force computations, and a series of KDK operations with varying time intervals for the small scale forces. It also offers an adaptive time-sequencing scheme for short range force computations that yields individual time intervals Δt for each particle depending on their accelerations $\vec{a}$:

$$\Delta t = \min\left[\Delta t_{\max}, \left(\frac{2\eta\epsilon}{|\vec{a}|}\right)^{1/2}\right], \quad (2.23)$$

where η represents a chosen level of accuracy, ϵ is the gravitational softening scale and $\Delta t_{\max}$ is a maximum upper time-step limit imposed for large-scale force computations. Each particle Δt from equation (2.23) is then grouped into discrete time intervals Δt_j according to a power-of-two integer scaling between the number N_j of particle time intervals and the coarsest time interval Δt_c: $\Delta t_j = \Delta t_c/2^{N_j}$ (Hernquist & Katz, 1989). This time sequencing system, which is used in the HORIZON simulation in this thesis, hence assigns several small time intervals to particles that experience large accelerations on small scales per larger time interval for particles with low

accelerations on large scales.

The particle trajectories within the HORIZON and NUT simulations can hence be tracked as a function of time. The next task is to compute the properties of gas over a range of scales. All of the hydrodynamics in this thesis have been modelled using the RAMSES code, and so no further discussion of GADGET-II beyond the N-body schemes presented in this section is provided.

2.1.4 Solving the time evolution of the gas flow variables

RAMSES solves the Euler equations in order to estimate the time evolution of the gas density ρ, momentum density $\rho\vec{u}$ and energy density ρe within a given cell:

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \vec{\nabla} \cdot \rho\vec{u} &= 0 \\ \frac{\partial (\rho\vec{u})}{\partial t} + \vec{\nabla} \cdot (\rho\vec{u} \otimes \vec{u}) + \vec{\nabla} p &= -\rho\vec{\nabla}\phi \\ \frac{\partial (\rho e)}{\partial t} + \vec{\nabla} \cdot [\vec{u}(\rho e + p)] &= -\rho\vec{u} \cdot \vec{\nabla}\phi. \end{aligned} \tag{2.24}$$

The variables u and e refer to the velocity and total specific energy (thermal plus kinetic) of the fluid respectively. In the RAMSES code, the gas pressure p and density ρ are related by the following equation of state for an ideal gas with adiabatic index γ, specific internal energy $e_{\rm int}$ and total specific energy e:

$$p = (\gamma - 1)\rho e_{\rm int}, \tag{2.25}$$

where:

$$e_{\rm int} = e - \frac{1}{2}\left|\vec{u}\right|^2. \tag{2.26}$$

Note that equations (2.24) cannot be expressed in conservative form in the presence of a gravitational potential ϕ as spatial fluctuations in ϕ act as source gravity terms, and so momentum and energy density are not conserved at the level of floating point precision (but errors are only at the percent level). Discontinuities at cell boundaries in the solutions to equations (2.24) can exist (e.g. the Sod shock tube problem), and such discontinuities for conservative fluids fall under the class of Riemann problems. There are many Riemann solvers that are able to solve problems of this nature (e.g. Godunov & Ryabenki 1964, Roe 1981) and in the NUT simulation the HLLC solver is used (e.g. Toro 2009).

Equations (2.24) can be written in the form of an advection equation with a source term $\vec{S}$:

$$\frac{\partial \vec{U}}{\partial t} + \vec{\nabla} \cdot \vec{F} = \vec{S}, \qquad \text{where:} \tag{2.27}$$

$$\vec{U} = \begin{pmatrix} \rho \\ \rho\vec{u} \\ \rho e \end{pmatrix}, \qquad \vec{F} = \begin{pmatrix} \rho\vec{u} \\ [\rho\vec{u} \otimes \vec{u} + p\mathbb{1}] \\ \vec{u}(\rho e + p) \end{pmatrix}, \qquad \vec{S} = \begin{pmatrix} 0 \\ -\rho\vec{\nabla}\phi \\ -\rho\vec{u} \cdot \vec{\nabla}\phi \end{pmatrix}, \tag{2.28}$$

and $\mathbb{1}$ is the 3×3 identity matrix. The solution to equation (2.27) along a given dimension is:

$$U_i^{n+1} = U_i^n - \frac{\Delta t}{\Delta x}\left(F_{i+1/2}^{n+1/2} - F_{i-1/2}^{n+1/2}\right) + S_i^{n+1/2}\Delta t, \tag{2.29}$$

where n represents the time index, $i \pm 1/2$ refers to the boundary either side of the centre of cell i that is shared with neighbouring cells $i \pm 1$, and:

$$F_{i+1/2}^{n+1/2} = \frac{1}{\Delta t}\int_{t^n}^{t^{n+1}} F(x_{i+1/2}, t)\,\mathrm{d}t \tag{2.30}$$

$$S_i^{n+1/2} = -\left(0, \frac{\rho_i^n\vec{\nabla}\phi_i^n + \rho_i^{n+1}\vec{\nabla}\phi_i^{n+1}}{2}, \frac{(\rho\vec{u})_i^n \cdot \vec{\nabla}\phi_i^n + (\rho\vec{u})_i^{n+1} \cdot \vec{\nabla}\phi_i^{n+1}}{2}\right). \tag{2.31}$$

The mid-interval components $S_i^{n+1/2}$ are hence arithmetic averages of the source terms at t^n and t^{n+1}. The solution to equation (2.30) is computed in RAMSES by using a second-order Godunov scheme as described in Toro (2009).

2.1.5 Modelling stellar processes

The remainder of the first section of this chapter introduces the high resolution NUT suite run that has been analyzed in this thesis to track the evolution of gas flow onto a Milky Way type disk. The NUT simulation includes stars and so this subsection provides a discussion of the prescriptions used for modelling star formation, which are described in detail by Rasera & Teyssier (2006) and Dubois & Teyssier (2008), and are briefly summarized by Powell et al. (2011).

Star formation

The gas in a cell can cool to temperatures $T \sim 10^4\,\mathrm{K}$ via bremsstrahlung radiation (effective until $T \sim 10^6\,\mathrm{K}$), and via collisional and ionization excitation followed by recombination (dominant for

$10^4 \leq T/\mathrm{K} \leq 10^6$). Further cooling to molecular cloud temperatures of $\sim$1 K occurs via metal line emission, and once the gas is sufficiently dense (i.e. $\rho > \rho_0$, where ρ_0 is a density threshold and is chosen to represent the density of the interstellar medium, hereafter ISM) star formation naturally ensues, a process which is modelled by a Schmidt law (Schmidt, 1959):

$$\frac{\mathrm{d}\rho_\star}{\mathrm{d}t} = -\frac{\rho}{t_\star}. \tag{2.32}$$

The star formation timescale $t_\star$ is given by:

$$t_\star = t_0 \left(\frac{\rho}{\rho_0}\right)^{-\frac{1}{2}}, \tag{2.33}$$

and is hence proportional to the local free-fall timescale. The two free parameters in this model—ρ_0 and t_0—are very poorly constrained by observations and are both scale-dependent (e.g. Krumholz & Tan 2007). In the NUT suite, the choice of value assigned to t_0 depends on the star formation efficiency ϵ (i.e. the fraction of gas mass of a given star-forming gas cloud that ends up in the stellar phase per free-fall time), and ρ_0. For the NUT outputs analyzed in this thesis, which correspond to a resolved physical scale of 12 pc, the efficiency parameter ϵ is fixed at 1% per free-fall time in concordance with observations (Krumholz & Tan, 2007), and ρ_0 is chosen to be equal to 400 hydrogen atoms per cubic centimetre. This choice of parameters yields t_0 of a few Myrs. If the Jean's scale λ_J is close to or below the minimum resolved physical scale Δx_min of the simulation, it is possible that gas clouds artificially fragment, leading to artificial star formation. Therefore, when in excess of the ISM density ρ_0, the gas is forced to follow a polytropic equation of state of the form:

$$T = T_0 \left(\frac{\rho}{\rho_0}\right)^{\gamma-1}, \tag{2.34}$$

where $T_0 = 100\,\mathrm{K}$. This choice of parameter values ensures that $\lambda_\mathrm{J} > \Delta x_\mathrm{min}$, and setting $\gamma = 4/3$ provides a sufficiently high source of outward pressure support to prevent the gas in the cell from collapsing under its own gravity at scales of the order Δx_min. The number of stellar particles that form in a cell whose gas is sufficiently cool and dense is a randomly chosen sample from a Poisson distribution:

$$P(N|\lambda) = \frac{\lambda^N \exp(-\lambda)}{N!}, \tag{2.35}$$

where the mean of the distribution λ is given by:

$$\lambda = \left[\frac{\rho(\Delta x)^3}{m_{\star,\min}}\right]\frac{\Delta t}{t_\star}, \tag{2.36}$$

and Δx and Δt are respectively the cubic cell length and time-step interval at the relevant level of refinement. The minimum stellar mass $m_{\star,\min}$ depends on parameters that account for mass loss via supernovae explosions, and so a more detailed discussion is postponed until the following subsection. In order to prevent an overproduction of stellar particles in a given cell, all of the N newly formed stellar particles are merged into a single particle of mass $m_\star = N m_{\star,\min}$. A maximum mass of $0.9M_g$ (where M_g is the gas mass) in a cell is allowed to form stars during a given time-step interval, as fractions in excess of this lead to numerical instabilities.

Supernovae feedback

Even though the NUT suite runs that include the local heating and injection of metals into the ISM via stellar explosions are not examined in this thesis, the value of $m_{\star,\min}$ in equation (2.36) is set by the parameters that define the supernovae feedback model and hence a brief description of this model is now provided. Each singly fused stellar particle in the NUT suite is guaranteed to undergo a supernova phase 10 Myr after its creation, corresponding to the typical main sequence lifetime of a Type-II supernova progenitor of mass $10\,\mathrm{M}_\odot$ (e.g. Ostlie & Carroll 2007). In so doing, a certain fraction of the stellar mass, characterized by η_{SN}, is recycled via supernovae explosions into the ISM. The resultant supernova Sedov shock also transports some of the mass surrounding the stellar particle, parametrized by η_{W}, into the ISM on scales below the minimum cell width $\Delta x_{\min}$ (which corresponds to the resolution limit). A gas cloud of mass M_g forming stars is hence assumed to distribute its mass as follows:

$$M_g = m_\star(1 + \eta_{\mathrm{SN}} + \eta_{\mathrm{W}}), \tag{2.37}$$

where $m_\star$ represents the mass that permanently belongs to the stars. It follows that $m_{\star,\min}$ in equation (2.36) is given by:

$$m_{\star,\min} = \frac{\rho_0(\Delta x_{\min})^3}{1 + \eta_{\mathrm{SN}} + \eta_{\mathrm{W}}}, \tag{2.38}$$

yielding $m_{\star,\mathrm{min}} \sim 2 \times 10^4\,\mathrm{M}_\odot$ for the $\Delta x_{\mathrm{min}} = 12\,\mathrm{pc}$ CO run, where η_{SN} and η_{W} are both zero. Stellar particles in the $\Delta x_{\mathrm{min}} = 12\,\mathrm{pc}$ NUT suite are therefore stellar clumps of up to several thousand stars. The value of η_{SN} in the supernovae feedback runs is found by truncating the Salpeter initial mass function between $0.1\,\mathrm{M}_\odot$ and $100\,\mathrm{M}_\odot$ and computing the fraction of stars whose mass lies above the minimum progenitor mass threshold for a Type-II supernova to occur ($\sim 8\,\mathrm{M}_\odot$), corresponding to $\eta_{\mathrm{SN}} \sim 0.1$. The free parameter η_{W} is set by imposing the condition that angular momentum is conserved pre- and post- the Sedov blast-wave phase, yielding $\eta_{\mathrm{W}} \sim 1$. This choice of η_{W} hence implies that the amount of gas surrounding the stellar particle that is eventually entrained by the blast wave, is equal to the mass permanently locked up within the stellar particle, and that this entrained mass originally belonged to the stellar particle before the supernovae phase.

The total energy E_{T} released by the supernova of the stellar particle that couples to the ISM with efficiency η_{SN} is:

$$E_{\mathrm{T}} = \eta_{\mathrm{SN}} \frac{m_\star}{M_{\mathrm{pro}}} E_{\mathrm{SN}}, \tag{2.39}$$

where E_{SN} is the energy expelled by a single stellar supernova and M_{pro} is the mass of the progenitor of each star. For a stellar particle progenitor of mass $M_{\mathrm{pro}} = 10\,\mathrm{M}_\odot$ that undergoes a Type-II supernova, $E_{\mathrm{SN}} \sim 10^{51}\,\mathrm{erg}$ (e.g. Ostlie & Carroll 2007). Half of the ejecta energy E_{T} in equation (2.39) is thermal and heats the local ISM and the other half is kinetic and drives the expansion of the Sedov shock through the ISM.

2.1.6 Resimulation techniques

The NUT suite simulations have adopted the 'resimulation' technique, also known as the 'zoom' technique. Following Navarro et al. (1995a), the aim of this method is to locate a region of interest within a simulation volume and harness most of the available computer power into that region. Usually a dark-matter-only simulation is run with initial conditions generated using a standard package (e.g. MPgrafic, Prunet et al. 2008) to aid selection of the region of interest. Haloes and subhaloes within the simulation grow with time until the simulation is ended at some truncation redshift z_t, and for the purpose of illustration, it is assumed in the following that $z_t = 0$. A target at z_t is then chosen—perhaps a Milky Way type halo H as in the NUT suite (Kimm et al., 2011; Powell et al., 2011). All of the member particles within H at z_t are then tagged and tracked back to the initial time-step of the simulation. The challenge is to define a fixed comoving box to be used in the full resimulation run that encloses all of the low mass particles belonging to H across

its evolution, and that excludes the higher mass 'contamination' particles. In order to achieve this, the centre of mass of the member particles (identified at z_t) is computed for a collection of time-steps, whose interval spacing is chosen by the user. Note that this is more reliable than choosing a box defined at just a single time-step, as this approach is not sensitive to possible halo migration outside the box as the system evolves. The distance of the furthest z_t-identified member particle from the new centre of mass of the z_t-identified member particles is then used to define a candidate box for each sample time-step. The minimum lower left corner and maximum upper right corner of all the boxes are then found, and these coordinates define a preliminary comoving resimulation box. The sidelength of this box is then multiplied by a 'buffer' value N (Tormen et al., 1997) in order to damp the discontinuities in mass between the resimulation particles and the contamination particles that are not included in the resimulation box. In the NUT simulations, the scaling of this region is $N \sim 3$. The adaptive mesh is then imprinted on the resimulation volume (which fully encloses the object of interest) and the entire simulation cube (containing both the contaminated and resimulated regions) is simulated with a much coarser sampling than used in the resimulation volume, so as to correctly account for the effects of gravity on large scales.

2.2 Dark matter halo and subhalo finding algorithms

Outputs from galaxy simulations contain information on the distribution of dark matter particles at any given epoch, which is completely determined by solving equations (2.3), (2.4) and (2.5). This distribution exists in a six dimensional phase-space, with three spatial and three velocity components per dark matter particle. It is hence the goal of any halo-finder to resolve this distribution into a collection of gravitationally bound objects called haloes and subhaloes, at each epoch (e.g. Knebe et al. 2011). Once these objects have been defined, it is possible to track their constituent particle positions back to earlier epochs, and these particle trajectories trace out a merger tree for every halo and subhalo in the simulation.

2.2.1 Resolving dark matter haloes

Despite recent claims that 3D halo-finders systematically underestimate the mass of detected haloes compared with 6D halo-finders (Maciejewski et al., 2009), the latter, more sophisticated algorithms are few in number (6DFOF: Diemand et al. 2006; HSF:Maciejewski et al. 2009; ROCKSTAR: Behroozi et al. 2011) and are very much in their infancy, and are hence avoided in this thesis.

The friends-of-friends algorithm

Perhaps the most popular 3D (spatial) halo-finder is the friends-of-friends (hereafter FOF) method (Davis et al., 1985), which groups together particles that are spatially separated by a distance that is a fraction b of the mean inter-particle separation (b is usually 0.2). Unfortunately, the FOF algorithm has several flaws:

- It is prone to halo-loss for haloes whose mass is near the widely accepted $20M_p$ threshold.

  Tweed et al. (2009) have shown that a non-negligible fraction of FOF haloes whose mass satisfies $M \leq 30M_p$ pass below the threshold at the subsequent time-step, and concluded that at least 40 particles are needed to reduce the halo-loss rate to the percent level.

- It cannot consistently resolve substructure.

  One might think that subhaloes could be found by reducing the linking length b (Klypin et al., 1999a). However, as argued by Tweed et al. (2009), there is no obvious choice of fixed value for this free parameter that is guaranteed to distinguish all levels of structure. The value of b required to resolve the sub-subhalo of halo H_1, for example, may be completely different from that required to resolve the sub-subhalo of H_2, simply because H_1 and H_2 could have different merger histories and formation epochs.

- It significantly overestimates the halo merger rate for haloes that are about to merge (Genel et al., 2009; Hopkins et al., 2010).

A more sophisticated algorithm is hence required.

The AdaptaHOP method

Throughout this work the halo-finder of choice is the AdaptaHOP algorithm, hereafter AHOP (Aubert et al., 2004). The basic idea behind AHOP is to identify peaks in the density field and saddle points between surfaces joining these peaks, and to use the density at the saddle points to create a 'node-tree'. Haloes and subhaloes represent certain trajectories from the base node to the leaves of this tree. Following Aubert et al. (2004), there are three core steps to the AHOP method:

- Assign a density estimate to each dark matter particle.

To compute the density of particle j, its $N_{\rm SPH}$ neighbours are found, with each neighbour donating a fraction of its mass to j according to the SPH cubic spline kernel (Monaghan & Lattanzio, 1985):

$$W_{\rm C}(x;R) = \frac{8}{\pi R^3}\begin{cases} 1-6\left(x/R\right)^2+6\left(x/R\right)^3 & \text{if } 0 \leq x/R \leq \frac{1}{2} \\ 2\left(1-x/R\right)^3 & \text{if } \frac{1}{2} < x/R \leq 1 \\ 0 & \text{if } x/R > 1, \end{cases} \tag{2.40}$$

where x represents the distance of the neighbour from j, and R is the smoothing scale of the window function. The number of neighbours $N_{\rm SPH}$ that contribute to the density estimate at a given location is a free parameter and is usually fixed to a few tens (Eisenstein & Hut, 1998). The distance to the furthest of the $N_{\rm SPH}$ neighbours determines the value of R.

- Impose a minimum density threshold ρ_t on all of the SPH particles and find local density maxima throughout the entire density field.

  The value assigned to ρ_t at each epoch is $\sim 80\rho_c$ (where ρ_c is the critical density of the Universe) as this roughly corresponds to a FOF linking length of $b = 0.2$ (Eisenstein & Hut, 1998). Each particle is then associated with its nearest density maximum.

- Compute the node-tree.

  Fig. 2.1, taken from Tweed et al. (2009), illustrates the concept of the node-tree. A node is defined as a region containing particles whose density is between a lower and upper limit, and it can contain several local density maxima. Each density maximum has an associated surface, and saddle points between the overlapping surfaces of the density maxima are found. Each node can therefore contain several density maxima connected by several saddle points. The densities at these saddle points then set the subsequent upper density threshold barriers. For example, node 1 in Fig. 2.1 contains particles with densities above ρ_t and below ρ_{23}, where ρ_{23} is the density of the lowest saddle point above ρ_t. Particles with $\rho_t \leq \rho \leq \rho_{23}$ are then distributed to their nearest density maximum, and in this example, there are two maxima. The ρ_{23} barrier therefore splits node 1 into two separates nodes, 2 and 3, with the density at the saddle point of node 3 being higher than that at node 2 (hence the number ordering). The process is then repeated for nodes 2 and 3 (i.e. the higher density saddle points are found)

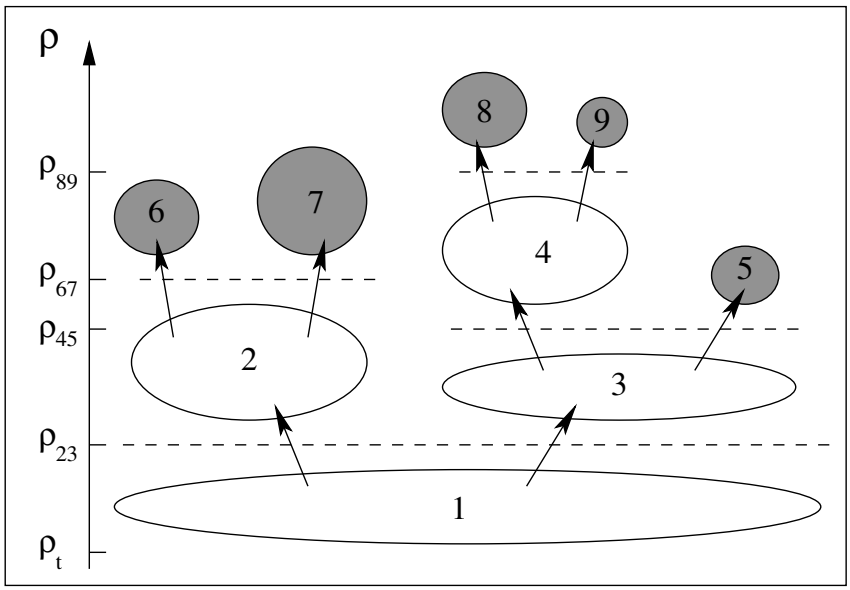

Figure 2.1: A schematic illustrating the node-tree structure, taken from Tweed et al. (2009). Each node is numbered according to the density at its saddle point: node 1 has the lowest saddle point density, whilst 9 has the highest. Only nodes $5, 6, 7, 8$ and 9 have physical meaning as they represent the location of the density peaks, and form the 'leaves' of the tree. The bigger the leaf circle, the more massive the leaf.

until the highest density peaks with no more saddle points have been reached, defining the leaves of the node-tree.

2.2.2 Resolving the substructure

With the leaves of the node-tree resolved for a given base node, the task now is to associate each leaf with a certain level of structure, i.e. decide which leaf is the host halo, which leaves are the subhaloes of the host, and whether there are any sub-subhaloes etc. The algorithm used in this thesis to address these issues is the 'Most massive Sub-node Method' (Tweed et al., 2009), hereafter MSM. For each base node the chain of nodes that leads to the most massive leaf is found, and this chain denotes the host halo. In Fig. 2.1, the host is given by the chain $1 \rightarrow 2 \rightarrow 7$ because node 7 is the most massive leaf. A wall is then drawn around nodes 1, 2 and 7, and the connections from nodes 1 and 2 are cut. The second most massive leaf, leaf 8, is next considered. Its chain is given by $1 \rightarrow 3 \rightarrow 4 \rightarrow 8$ but because connections from node 1 have been cut, this reduces to $3 \rightarrow 4 \rightarrow 8$, forming a subhalo of halo $1 \rightarrow 2 \rightarrow 7$. A wall is then drawn around nodes 3, 4 and 8 and the connections from nodes 3 and 4 are cut. Leaf 6 is also a subhalo of $1 \rightarrow 2 \rightarrow 7$, whereas leaves 5 and 9 are subhaloes of subhalo $3 \rightarrow 4 \rightarrow 8$, since links to nodes 3 and 4 have been cut. The various

levels of structure are hence defined. Note that the MSM method is not the only possible way of resolving structure from the node-tree. Suppose the order of selection of the leaves is performed according to density rather than mass. The densest leaf, leaf 8, would be considered first, defining the main halo $1 \rightarrow 3 \rightarrow 4 \rightarrow 8$. Following the above procedure, 9, 5 and $2 \rightarrow 7$ would be subhaloes of $1 \rightarrow 3 \rightarrow 4 \rightarrow 8$, and 6 would be a sub-subhalo of $1 \rightarrow 3 \rightarrow 4 \rightarrow 8$. This 'Densest Profile Method' therefore yields a different ordering of structure compared with the MSM ordering (3 subhaloes and 1 sub-subhalo for the former vs 2 subhaloes and 2 sub-subhaloes for the latter, and different host haloes). This example illustrates that the most massive leaf is not necessarily the most dense, and if this is the case the two methods will always resolve different structure hierarchies. The MSM algorithm is chosen in this thesis, however, as Tweed et al. (2009) have found that the Densest Profile Method, unlike the MSM routine, can predict subhaloes to have higher masses than their host haloes, which is unphysical.

At various stages throughout this thesis, reference will be made to the AHOP halo: this refers to the combined mass of all the leaves of each resolved base node (i.e. $\sum_{i=5}^{i=9} m_i$ in Fig. 2.1). The MSM halo, by contrast, refers to the host halo $1 \rightarrow 2 \rightarrow 7$ and by construction its mass is lower than the mass of its AHOP counterpart, with the mass difference shared amongst higher levels of MSM-resolved substructure.

2.2.3 Constructing the merger tree

Once a halo or a subhalo at a given time-step t_n has been identified, its constituent particles are known. The next step is to link haloes and subhaloes at different time-steps and follow a given (sub)halo's merger tree. The TreeMaker code (Tweed et al., 2009) is used in this thesis to follow the trajectories of each particle across time. Suppose that a halo H at time-step t_n contains 40 dark matter particles. Each of these particles is traced back to the earlier time-step t_{n-1}. Every unique (sub)halo containing at least one of these 40 particles is recorded as a 'father' progenitor of H and so it is possible for H to have several fathers. Note that it is also possible that a particle is not traced to a halo or a subhalo at t_{n-1}: instead it may belong to the set of 'background' particles, which are not bound to any resolved structure (e.g. tidally-stripped material from a merger event). The fathers of all the fathers of H are then examined and the same process is repeated. The particles belonging to H can also be traced forward in time to time-step t_{n+1}, and their owners are tagged as descendent 'sons'. Halo H may therefore have several sons. This tracking of particle

trajectories across time in the direction of both the sons and the fathers is a key feature of the accretion algorithm presented in the following chapter as it guarantees mass conservation.

2.3 A simulation specification overview

Having explained the basic models and techniques encoded into the simulations, this chapter concludes with a summary of the technical specifications of the simulation runs analyzed in this thesis.

2.3.1 The HORIZON simulation

The dark-matter-only HORIZON simulation used to examine accretion onto haloes and subhaloes contains 512^3 dark matter particles within a box of comoving side length $100\,h^{-1}\,\mathrm{Mpc}$ in a ΛCDM universe. Haloes and subhaloes form in a universe described by the Eisenstein & Hu (1999) matter transfer function and a WMAP3 cosmology (Spergel et al., 2007), where $\Omega_\mathrm{M} = 0.24, \Omega_\Lambda = 0.76$, $\Omega_b = 0.042, n = 0.958, h = 0.73$ and $\sigma_8 = 0.77$. The time-steps are separated by 0.01 in scale factor from $z = 99$ to the present day, but analysis is restricted to haloes and subhaloes in the redshift range $0 \leq z \leq 9$. The mass of each particle, M_p, is $6.8 \times 10^8\,\mathrm{M}_\odot$ and haloes and subhaloes with a recorded accretion value contain at least 40 particles. The mass of a (sub)halo used in this study corresponds to the total mass M_T detected by the halo-finder, which is found by summing the mass of all the tagged (sub)halo member particles. The total mass hence satisfies $M_\mathrm{T} \geq M_\mathrm{vir}$, where M_vir is the (sub)halo's virial mass. For reference, the MSM algorithm resolves 223 781 objects at $z = 0$ and ∼20% of these objects are subhaloes.

2.3.2 The NUT cooling run

One of the NUT simulations has been examined in a quest to understand how a Milky Way type disk acquires its angular momentum. Following Powell et al. (2011), the NUT suite is a set of ultra-high resolution simulations that employ a zoom technique (see Section 2.1.6) to resimulate a Milky Way type galaxy and follow its evolution across redshift in a ΛCDM cosmology. The highest resolution runs have a physical resolved scale Δx_min of 0.5 pc but terminate at high redshifts around $z = 6$ due to computational time restrictions, and so in this thesis the lower resolution $\Delta x_\mathrm{min} = 12\,\mathrm{pc}$ outputs are analyzed as they probe redshifts down to $z = 3$. The NUT simulations have been run with the adaptive mesh refinement code RAMSES, which incorporates the following physics:

(i) Epoch of reionization—a spatially uniform, redshift-dependent ultraviolet (UV) background, instantaneously switched on at $z = 8.5$ (following the Haardt & Madau 1996 UV model).

(ii) Cooling (a modified version of the Sutherland & Dopita 1993 cooling function) and star formation.

(iii) Supernovae feedback and hence metal enrichment, which in turn alters the metal cooling function.

Outputs from the $\Delta x_{\rm min} = 12\,{\rm pc}$ cooling run (hereafter CO run) are examined in this thesis, which include (i) and (ii) with a fixed metallicity of $10^{-3}\,{\rm Z}_{\odot}$ (the NUT supernovae feedback runs also include (iii), but are not considered in this thesis).

The chosen resimulated Milky Way type halo satisfies two conditions. Firstly it is not embedded in a cluster environment at $z = 0$, and secondly its present day mass ($M_{\rm H} \simeq 5 \times 10^{11}\,{\rm M}_{\odot}$) is comparable to the threshold mass ($M_{\rm H} \sim 4 \times 10^{11}\,{\rm M}_{\odot}$) reported by Ocvirk et al. (2008) below which gas streams along filamentary structures at low temperatures of ${\rm T} \lesssim 2 \times 10^4\,{\rm K}$. The entire simulation box has a comoving side length of $9\,h^{-1}\,{\rm Mpc}$ and starts at $z = 499$, with the initial conditions generated using the package MPgrafic (Prunet et al., 2008). The resimulation box around the main halo has a comoving side length of $\sim 2.7\,h^{-1}\,{\rm Mpc}$, and the simulation evolves in a universe of WMAP5 cosmology (Dunkley et al., 2009) with $\Omega_{\rm M} = 0.258, \Omega_{\Lambda} = 0.742$, $\Omega_b = 0.045, \sigma_8 = 0.8$ and $h = 0.72$. The coarse root grid for the entire simulation has 128 cells along each dimension of the $9\,h^{-1}\,{\rm Mpc}$ box, whereas the resimulated region contains three higher resolution nested grids yielding an equivalent particle resolution of 1024^3 dark matter particles, each with mass $M_{\rm DM} = 5.4 \times 10^4\,{\rm M}_{\odot}$. A quasi-Lagrangian refinement scheme is used to maintain $\Delta x_{\rm min} = 12\,{\rm pc}$ in physical coordinates as the simulation evolves. Higher levels of refinement are spawned once either: a) the baryonic mass in the cell exceeds $8 m_{\rm SPH}$ ($m_{\rm SPH} = 9.4 \times 10^3\,{\rm M}_{\odot}$) or; b) the number of dark matter particles in the cell exceeds eight. This scheme therefore strives to achieve a roughly equal gas mass per cell.

The outputs from the NUT CO run were then fed into the MSM algorithm (with an identical cosmology) in order to detect the haloes, subhaloes, stellar clumps and stellar subclumps at each time-step. Each resolved halo (stellar clump) object was forced to contain at least 40 (100) particles to ensure reliable detections. The stellar clump resolution limit was set above the dark halo limit because the clump-finder identified groups of stars within the disk as being separate from the disk when lower thresholds were imposed.

Armed with an understanding of the processes modelled in the HORIZON and NUT simulations, the following chapters introduce algorithms that have been developed to tackle both of the projects presented at the end of Chapter 1, and discuss results acquired from their implementation.

Chapter 3

Techniques for measuring mass accretion onto haloes

The analytical and computational tools that have been used in this thesis to compute the mass accretion onto haloes and subhaloes across time are now discussed. This serves as the methods precursor to the following chapter, which presents results obtained from employing these accretion schemes and addresses the possible dependence of halo and subhalo accretion rates on environment.

The first section of this chapter begins with a discussion of the theory of halo evolution, which ultimately leads to a prediction of the mass accretion rate onto dark haloes as a function of mass and redshift. This theory suffers from several limitations, however, with its most significant shortcoming being that it cannot account for accretion onto substructures embedded in dense environments. The second section therefore appeals to N-body simulations to resolve this issue, and moves beyonds the standard theory by presenting an algorithm that computes accretion onto all levels of structure, be it haloes, subhaloes, sub-subhaloes etc.

3.1 The theory of dark matter halo evolution and accretion

The spherical top-hat model of halo formation was introduced in Section 1.2.1, and it was demonstrated that this model yields simple estimates of the density contrast of a collapsed spherical perturbation at its virialization epoch. This section moves beyond these predictions by accounting for the time evolution of the density perturbations.

3.1.1 Estimating halo masses

By smoothing the initial density field on a scale R, it is possible to compute the mass enclosed within R. This smoothing process is equivalent to convolving the density field with a window function $W(\vec{x}; R)$ (e.g. Coles & Lucchin 2002; Zentner 2007):

$$\delta(\vec{x}; R) = \int W(|\vec{x} - \vec{x}'|; R)\delta(\vec{x}')\,\mathrm{d}^3x', \tag{3.1}$$

with Fourier modes:

$$\delta(\vec{k}; R) = W(\vec{k}; R)\delta(\vec{k}). \tag{3.2}$$

By using equation (1.11), the mass variance $\sigma^2(R) \equiv \left\langle \delta^2(\vec{x}; R) \right\rangle$ on scale R reduces to:

$$\sigma^2(R) = \int_{-\infty}^{\infty} \Delta^2(k)|W(k; R)|^2\,\mathrm{d}(\ln \mathrm{k}). \tag{3.3}$$

Choice of window function

There are three filters that are frequently used to smooth densities, and their real-space and Fourier-space expressions are given below, along with their corresponding real-space volumes:

1. Spherical top-hat.

$$\begin{aligned} W(\vec{x}; R) &= \left(\frac{4\pi}{3}R^3\right)^{-1} \Theta(1 - x/R) \\ W(\vec{k}; R) &= 3\left[\frac{\sin(kR) - (kR)\cos(kR)}{(kR)^3}\right] \\ V &= \frac{4\pi}{3}R^3 \end{aligned} \tag{3.4}$$

2. Gaussian.

$$\begin{aligned} W(\vec{x}; R) &= \left[(2\pi)^{3/2}R^3\right]^{-1} \exp\left(-\frac{x^2}{2R^2}\right) \\ W(\vec{k}; R) &= \exp\left(-\frac{k^2R^2}{2}\right) \\ V &= (2\pi)^{3/2}R^3 \end{aligned} \tag{3.5}$$

3. Sharp k-space.

$$\begin{aligned} W(\vec{x};R) &= \left(6\pi^2 R^3\right)^{-1} 3\left[\frac{\sin(x/R)-(x/R)\cos(x/R)}{(x/R)^3}\right] \\ W(\vec{k};R) &= \Theta(1-kR) \\ V &= 6\pi^2 R^3, \end{aligned} \tag{3.6}$$

where $x = |\vec{x}|$, $k = |\vec{k}|$ and $\Theta(y)$ is the Heaviside step function, defined as:

$$\Theta(y) = \begin{cases} 1 & \text{if } y \geq 0 \\ 0 & \text{if } y < 0. \end{cases} \tag{3.7}$$

As argued in several recent reviews (Peacock, 1999; Coles & Lucchin, 2002; Zentner, 2007), it is useful to define a volume for each window function V as this provides an estimate of the enclosed mass ($M \propto \bar{\rho}V$, where $\bar{\rho}$ is the mean density of the Universe). Since the window functions above have units of inverse volume, W is normalized to a dimensionless quantity W' whose maximum value is unity, and V is hence defined as:

$$V \equiv \int W'(\vec{x};R)\,\mathrm{d}^3x \tag{3.8}$$

$$= \int \frac{W(\vec{x};R)}{W(0;R)}\,\mathrm{d}^3x. \tag{3.9}$$

Each filter has its relative advantages and disadvantages. Following Bond et al. (1991), the spherical top-hat function has a well-defined mass ($M = \frac{4}{3}\pi R^3\bar{\rho}$) but yields a smoothed density field in equation (3.1) that cannot be differentiated with respect to position $\vec{x}$, and also suffers from sharp transitions in real-space translating to power on all scales in Fourier-space. The Gaussian filter has a variance which decreases monotonically with R, but has a less well-defined volume. The sharp k-space filter has the major advantage that a decrease in filter scale in real-space introduces uncorrelated modes in Fourier-space, a property that is exploited in computing the probability for a region to collapse under a change in filter scale (which is discussed in the next section). Its disadvantage is that it is not physical in real-space: $W(\vec{x};R)$ in equation (3.6) is not finite.

3.1.2 The Press–Schechter (PS) formalism

In the ΛCDM model, $\sigma^2(M) \to 0$ as $R \to \infty$. As the window scale is decreased, one could imagine that $\delta(\vec{x}; R)$ becomes sufficiently dense that gravitational collapse into virialized structures takes place. In Section 1.2.1 it was demonstrated that the spherical top-hat model of halo formation predicts that $\delta \sim 1$ at turnaround, implying that a density contrast of order unity will break from the background expansion and subsequently form a virialized dark matter halo. Press & Schechter (1974) used the above ideas to make a prediction for the number density of haloes in a certain mass range at a given epoch. As highlighted by Zentner (2007), the standard theory makes the following important assumptions:

- the density field $\delta(\vec{x}; R)$ is Gaussian distributed (Fourier modes have random phases), homogeneous and isotropic (the $\vec{x}$ dependence of δ is henceforth dropped for brevity);
- density perturbations evolve linearly, as described by the linear growth factor;
- the collapse of overdense regions at large scales is independent of smaller scale collapse, despite the non-linearities in the density field that arise on these smaller scales once a region of the Universe has exceeded the critical density contrast for collapse $\delta_c(z)$ at that redshift and formed a virialized halo;
- $\delta_c(z)$ is independent of mass (although Sheth et al. 2001 later introduced a mass dependence);
- the initial density field can be smoothed on various scales to deduce the mass function of virialized haloes that collapse in that density field.

The obvious weakness associated with this latter assumption is that Press–Schechter framework does not model halo collapse as a dynamical, non-linear process, but a more formal discussion of the flaws of the theory is postponed until the end of this subsection.

Press & Schechter (1974) argued that in order to find the fraction of mass $F(> M)$ at a given epoch that is locked-up in virialized objects whose mass is greater than M, one needs to sum all the regions with a density contrast above the critical threshold:

$$F(> M) = \int_{\delta_c(z)}^{\infty} P(\delta; R)\, \mathrm{d}\delta. \tag{3.10}$$

An expression for the probability $P(\delta; R)$ is hence required to further evaluate equation (3.10). Since the density field is assumed to be a random Gaussian field, the smoothed field in equation (3.1) is also Gaussian, and so the probability of finding a density perturbation between δ and $\delta + \mathrm{d}\delta$ when smoothed on a scale R is:

$$P(\delta; R)\,\mathrm{d}\delta = \frac{1}{\sqrt{2\pi\sigma^2(R)}} \exp\left[-\frac{\delta^2}{2\sigma^2(R)}\right] \mathrm{d}\delta. \tag{3.11}$$

Hence $F(> M)$ in equation (3.10) simplifies to:

$$F(> M) = \frac{1}{2}\mathrm{erfc}\left(\frac{\nu}{\sqrt{2}}\right), \tag{3.12}$$

where $\nu \equiv \delta_c/\sigma(M)$ and $\mathrm{erfc}(x)$ denotes the complementary error function. As the smoothing scale is reduced to zero, all of the $\delta(R)$ will be above δ_c and so $F(> M)$ should approach unity, yet equation (3.12) predicts that only half the mass will be bound to virialized haloes. Press & Schechter (1974) speculated that this 'missing' mass, a feature at all epochs, arises from not including the underdense regions that accrete onto the overdense collapsed objects at later times, and hence a correction factor of 2 was inserted into equation (3.12). As argued by several authors (e.g. Peacock & Heavens 1990, Bond et al. 1991), that this exclusion should precisely correspond to a factor of 2 is highly dubious, but the convention is preserved here nonetheless. It follows that the comoving differential number density of haloes per unit mass at a time where the Universe has a mean comoving background density $\bar{\rho}$, is given by:

$$\frac{\mathrm{d}n}{\mathrm{d}M}\,\mathrm{d}M = 2\frac{\bar{\rho}}{M}\left|\frac{\mathrm{d}F(> M)}{\mathrm{d}M}\right| \mathrm{d}M \tag{3.13}$$

$$= \sqrt{\frac{2}{\pi}}\frac{\bar{\rho}}{M^2}\nu\frac{\mathrm{d}(\ln\nu)}{\mathrm{d}(\ln M)} \exp\left(-\frac{\nu^2}{2}\right) \mathrm{d}M. \tag{3.14}$$

3.1.3 Extended Press–Schechter (EPS) theory

The standard implementation of the Press–Schechter formalism assumes that if $\delta(R) < \delta_c$ at some scale R, then collapse can only proceed at some smaller scale $R' < R$. It is, however, possible that $\delta > \delta_c$ at a larger smoothing scale, and accounting for this would increase the fraction of mass in haloes $F(> M)$. This issue is commonly referred to as the 'cloud-in-cloud' problem and solving it requires computing the maximum R, or correspondingly the smallest value of $\sigma^2(R)$, at which

δ first surpasses the critical density barrier. Bond et al. (1991) showed that the EPS formalism solves the cloud-in-cloud problem and alleviates the ad-hoc insertion of the factor 2 in equation (3.13). The remainder of this subsection is dedicated to demonstrating this result, and follows the approach taken by Zentner (2007).

It is common convention to represent the evolution of δ with scale by transforming to $[S, \delta(S)]$ space, where $S \equiv \sigma^2(R)$. The EPS formalism in this new space seeks to find the smallest value of S where δ first surpasses the δ_c threshold. As R is successively decreased, $S(R)$ is increased, and $\delta(S_i)$ traces out a trajectory. If the $\delta(S_i)$ along the trajectory are correlated, modelling this trajectory is highly non-trivial as the correlations between all neighbouring $\delta(S_i)$ have to be computed. This problem is bypassed when smoothing the density field with the sharp k-space filter, however, since a transition of ΔS in real-space adds a set of uncorrelated Fourier modes $\delta_{\vec{k}}$ with random phases to the density field, and so the trajectory $\delta(S_i)$ follows a Markovian random walk. The probability of transitioning from δ_1 to δ_2 ($= \delta_1 + \Delta\delta$) after increasing the mass variance from S_1 to S_2 ($= S_1 + \Delta S$) is hence described by a Gaussian:

$$\Pi(\delta_2, S_2)\,\mathrm{d}\delta_2 = \frac{1}{\sqrt{2\pi\Delta S}} \exp\left[-\frac{(\Delta\delta)^2}{2\Delta S}\right]\,\mathrm{d}(\Delta\delta). \tag{3.15}$$

When the step size in S is further reduced so that trajectories during the interval ΔS that surpass the density barrier are not missed, the probability distribution that describes the transition from δ_1 to δ_2 using the sharp k-space window function is a solution to a diffusion equation:

$$\frac{\partial\Pi}{\partial S} = \frac{1}{2}\frac{\partial^2\Pi}{\partial\delta^2}. \tag{3.16}$$

The simplest way to find the cumulative probability $F(> M)$ that the first surpassing of the critical density barrier has occurred at smaller mass variance than some $S(R)$, is to compute the probability that it has not. This imposes the necessary boundary conditions to solve equation (3.16), namely: a) $\Pi(\delta, S)$ is finite as $\delta \to -\infty$ and; b) $\Pi(\delta_c, S) = 0$, i.e. the trajectory is absorbed at the barrier. Taking the Fourier transform of equation (3.16) and using these boundary conditions yields the solution:

$$\Pi(\delta_2, S_2) = \frac{1}{\sqrt{2\pi\Delta S}}\left(\exp\left[-\frac{(\Delta\delta)^2}{2\Delta S}\right] - \exp\left[-\frac{(2(\delta_c - \delta_1) - \Delta\delta)^2}{2\Delta S}\right]\right), \tag{3.17}$$

where $\Delta\delta = \delta_2 - \delta_1$ and $\Delta S = S_2 - S_1$. The total fraction of all trajectories in $[S, \delta(S)]$ space that

have already surpassed the barrier is hence given by:

$$F(S) = 1 - \int_{-\infty}^{\delta_c} \Pi(\delta(S), S)\,\mathrm{d}\delta = \mathrm{erfc}\left(\frac{\delta_c - \delta_1}{\sqrt{2}\Delta S}\right). \tag{3.18}$$

It follows that the probability an upcrossing occurs between S and $S + \mathrm{d}S$ is (Bond et al. 1991; Lacey & Cole 1993, 1994):

$$\begin{aligned} f(S_2\,|\,\delta_1, S_1)\,\mathrm{d}S &= \left|\frac{\mathrm{d}F}{\mathrm{d}S}\right|\,\mathrm{d}S \\ &= \left(\int_{-\infty}^{\delta_c} \frac{\partial\Pi}{\partial S}\mathrm{d}\delta\right)\,\mathrm{d}S \\ &= \frac{1}{2}\left(\int_{-\infty}^{\delta_c} \frac{\partial^2\Pi}{\partial\delta^2}\mathrm{d}\delta\right)\,\mathrm{d}S \\ &= \left[\frac{1}{2}\frac{\partial\Pi}{\partial\delta}\right]_{-\infty}^{\delta_c}\,\mathrm{d}S \\ &= \frac{\delta_c - \delta_1}{\sqrt{2\pi}(\Delta S)^{3/2}} \exp\left[-\frac{(\delta_c - \delta_1)^2}{2\Delta S}\right]\,\mathrm{d}S, \end{aligned} \tag{3.19}$$

where equation (3.16) has been used in transitioning from the second to the third line. If the initial coordinates of the random walk are $S_1 = 0$ and $\delta_1 = 0$, the EPS mass function with sharp k-space filtering reproduces equation (3.14) without the ad-hoc insertion of the factor 2:

$$\frac{\mathrm{d}n}{\mathrm{d}M}\,\mathrm{d}M = \frac{\bar{\rho}}{M}\left|\frac{\mathrm{d}F}{\mathrm{d}S}\frac{\mathrm{d}S}{\mathrm{d}M}\right|\,\mathrm{d}M \tag{3.20}$$

$$= \sqrt{\frac{2}{\pi}}\frac{\bar{\rho}}{M^2}\frac{\delta_c}{\sigma}\left|\frac{\mathrm{d}(\ln\sigma)}{\mathrm{d}(\ln M)}\right| \exp\left(-\frac{\delta_c^2}{2\sigma^2}\right)\,\mathrm{d}M \tag{3.21}$$

$$= \sqrt{\frac{2}{\pi}}\frac{\bar{\rho}}{M^2}\nu\frac{\mathrm{d}(\ln\nu)}{\mathrm{d}(\ln M)} \exp\left(-\frac{\nu^2}{2}\right)\,\mathrm{d}M. \tag{3.22}$$

3.1.4 EPS predictions of halo accretion

Following equation (3.19), the probability that a particle belonged to a halo of mass M_2 at a time t_2 given that it is part of a halo of larger mass M_1 at a later time t_1 is (Lacey & Cole, 1993, 1994):

$$\frac{\mathrm{d}P}{\mathrm{d}M_2}(M_2, t_2\,|\,M_1, t_1) = \frac{\Delta\delta_c}{\sqrt{2\pi}(\Delta\sigma^2)^{3/2}}\left|\frac{\mathrm{d}\sigma^2}{\mathrm{d}M_2}\right| \exp\left[-\frac{(\Delta\delta_c)^2}{2\Delta\sigma^2}\right], \tag{3.23}$$

where $\Delta\delta_c = \delta_c(t_2) - \delta_c(t_1)$ and $\Delta\sigma^2 = \sigma^2(M_2) - \sigma^2(M_1)$. Miller et al. (2006) used this expression to derive a formula for the mass accretion rate onto a halo of mass M_H as a function of time. Their

formalism starts by considering the expected mass accreted ΔM onto a halo of mass M_1 at time t_1 during the time interval $\Delta t = t_1 - t_2$, which according to equation (3.23) is:

$$\langle \Delta M \rangle = \int_0^{M_1} (M_1 - M_2) \frac{\Delta \delta_c}{\sqrt{2\pi}(\Delta\sigma^2)^{3/2}} \left| \frac{\mathrm{d}\sigma^2}{\mathrm{d}M_2} \right| \exp\left[-\frac{(\Delta\delta_c)^2}{2\Delta\sigma^2} \right] \mathrm{d}M_2. \tag{3.24}$$

This can be simplified by making the substitution $y^2 \equiv 1/\Delta\sigma^2$:

$$\langle \Delta M \rangle = M_1 \Delta\delta_c \sqrt{\frac{2}{\pi}} \int_0^{\infty} \left(1 - \frac{M_2}{M_1} \right) \exp\left[-\frac{1}{2} y^2 (\Delta\delta_c)^2 \right] \mathrm{d}y. \tag{3.25}$$

As $\Delta t \to 0$, $\Delta\delta_c \to 0$ and so equation (3.25) becomes:

$$\lim_{\Delta\delta_c \to 0} \langle \Delta M \rangle \to M_1 \Delta\delta_c \sqrt{\frac{2}{\pi}} \int_0^{\infty} \left(1 - \frac{M_2}{M_1} \right) \mathrm{d}y, \tag{3.26}$$

or in terms of $\Delta\sigma^2$:

$$\langle \Delta M \rangle \to M_1 \frac{\Delta\delta_c}{\sqrt{2\pi}} \int_0^{\infty} \left(1 - \frac{M_2}{M_1} \right) \frac{\mathrm{d}(\Delta\sigma^2)}{(\Delta\sigma^2)^{3/2}}. \tag{3.27}$$

It follows that in the limit $\Delta\delta_c \to 0$, the average mass accretion rate onto a halo of mass M_H at time t is:

$$\left\langle \dot{M}_\mathrm{H} \right\rangle = \lim_{\Delta\delta_c \to 0} \frac{\langle \Delta M \rangle}{\Delta\delta_c} \left| \frac{\mathrm{d}\delta_c(t)}{\mathrm{d}t} \right| = M_\mathrm{H} \left| \frac{\mathrm{d}\delta_c(t)}{\mathrm{d}t} \right| f(M_\mathrm{H}), \tag{3.28}$$

where:

$$\left| \frac{\mathrm{d}\delta_c(t)}{\mathrm{d}t} \right| = \left| \frac{\mathrm{d}\delta_c(t)}{\mathrm{d}D(a)} \frac{\mathrm{d}D(a)}{\mathrm{d}a(t)} \frac{\mathrm{d}a(t)}{\mathrm{d}t} \right|, \tag{3.29}$$

and:

$$f(M_\mathrm{H}) = \sqrt{\frac{1}{2\pi}} \int_0^{\infty} \left(1 - \frac{M_2}{M_\mathrm{H}} \right) \frac{\mathrm{d}(\Delta\sigma^2)}{(\Delta\sigma^2)^{3/2}}. \tag{3.30}$$

These expressions for $\mathrm{d}\delta_c(t)/\mathrm{d}t$ and $f(M_\mathrm{H})$ can be further simplified:

- $\mathrm{d}\delta_c(t)/\mathrm{d}t$

  EPS theory assumes that $\delta_c(a)$ can be normalized by its present day value:

$$\delta_c(a) = \frac{\delta_c D}{D(a)}, \tag{3.31}$$

where $D(a)$ is the linear growth factor (Heath, 1977):

$$D(a) = \frac{5}{2}\Omega_{\rm M} H_0^2 H(a) \int_0^a \frac{1}{[aH(a)]^3}\,{\rm d}a. \tag{3.32}$$

Equation (3.31) is an approximation: Percival (2005) showed that $\delta_c(a)$ has a dependence on the adopted cosmology, and cannot therefore be described solely by $D(a)$. However, Miller et al. (2006) later demonstrated that this dependence is weak, and argued that the errors associated with this approximation are below the 2% level. Hence by assuming that $\delta_c = 1.686$, the value of the present day critical density contrast in an Einstein–de Sitter universe (Gunn & Gott, 1972), equation (3.29) becomes:

$$\left|\frac{{\rm d}\delta_c(t)}{{\rm d}t}\right| \simeq 1.686 \frac{D}{D^2(a)} \frac{{\rm d}D(a)}{{\rm d}a} \frac{{\rm d}a(t)}{{\rm d}t}. \tag{3.33}$$

Note that the evolution of $\delta_c(t)$ is independent of mass.

- $f(M_{\rm H})$

  Equation (3.30) can be simplified by integrating by parts:

$$\begin{aligned} f(M_{\rm H}) &= \sqrt{\frac{2}{\pi}} \left[\left(1 - \frac{M_2}{M_{\rm H}}\right) (\Delta\sigma^2)^{-1/2} \right]_{M_2=0}^{M_2=M_{\rm H}} + \sqrt{\frac{2}{\pi}} \int_0^{M_{\rm H}} (\Delta\sigma^2)^{-1/2} \frac{1}{M_{\rm H}}\,{\rm d}M_2 \\ &= \sqrt{\frac{2}{\pi}} \frac{1}{M_{\rm H}} \int_0^{M_{\rm H}} (\Delta\sigma^2)^{-1/2}\,{\rm d}M_2, \end{aligned} \tag{3.34}$$

  and Miller et al. (2006) demonstrated that this mass dependence is weak.

Equation (3.28) can hence be evaluated by using equations (3.33) and (3.34), as was originally shown by Miller et al. (2006). These expressions, however, are not unique: consideration of the forward probability in equation (3.23) yields an alternative expression for $\left\langle \dot{M}_{\rm H} \right\rangle$. This arises due to the lack of self consistency within the EPS formalism as highlighted by Benson et al. (2005), who considered EPS-derived merger rates $R(M_i, M_j\,|\,t)$ between haloes of mass M_i and M_j at time t, and demonstrated that the condition $R(M_i, M_j\,|\,t) \equiv R(M_j, M_i\,|\,t)$ is not satisfied for all M_i/M_j.

3.1.5 The shortcomings of the EPS framework

EPS theory is only approximate and has several limitations:

1. The assumption of spherical collapse.

It is standard convention in EPS theory to perform a linear extrapolation of the density contrast at some epoch z to a normalization epoch z_n, usually chosen to be the present day. The linearly extrapolated $\delta(z)$ is denoted by $\delta^{\rm L}(z;z_n)$. Any $\delta(z)$ whose extrapolated value satisfies $\delta^{\rm L}(z;z_n) > \delta_c^{\rm L}(z;z_n)$ is assumed to have already collapsed into a virialized halo at z. By extrapolating the density contrast of an initial spherical perturbation that just breaks away from the Hubble expansion of a spherical background region of equal mass, it is possible to derive the critical density contrast as a function of collapse epoch (Gunn & Gott, 1972; Eke et al., 1996). For an Einstein–de Sitter universe, Gunn & Gott (1972) showed that:

$$\delta_c(z_{\rm coll}) \simeq 1.686(1+z_{\rm coll}), \tag{3.35}$$

where $z_{\rm coll}$ is the redshift at which the perturbation collapses. This hence predicts that a $\delta(z)$ with $\delta^{\rm L}(z;z_0) = 1.686$ will collapse at z_0 (the present day) in a universe with $\Omega_{\rm M}(z) = 1$, whereas a $\delta(z)$ with $\delta^{\rm L}(z;z_0) = 6\delta_c(z_0)$ will collapse at $z = 5$.

Dark-matter-only N-body simulations challenge the assumption of spherical symmetry and hence the applicability of the spherical collapse model described above and in Chapter 1, as haloes are found to be triaxial (Jing & Suto, 2002; Bailin & Steinmetz, 2005).

2. The lack of dynamical information.

   This is a serious problem because merging is a highly dynamical process, and even N-body simulations suffer from noisy 'fake' accretion events that are an artefact of discretizing the merger process (these fake events can be identified, however, as described in Section 3.2.3). Unlike simulations, however, EPS theory cannot account for rudimentary dynamical phenomena, like mass being stripped from one halo and then being accreted onto another, or halo merger rates being self-consistent (Benson et al., 2005).

3. Averaging over halo environment.

   EPS theory approximates the mass accretion rate onto an object of mass M but does not resolve the substructure that could be contained within M. The EPS formalism cannot, therefore, measure the mass accreted onto the individual substructures representing luminous satellite galaxies in a group or cluster. Clusters like the Local Group are modelled as a single object in EPS theory, and so a lot of physics that describes accretion onto resolved structures is lost (Bower, 1991; Sandvik et al., 2007; Zentner, 2007; Desjacques, 2008).

3.2 The simulation perspective of halo accretion

In the remainder of this chapter, the laboratory for measuring accretion onto haloes and subhaloes—dark matter simulations—is considered. In the following chapter, the rates of accretion onto these structures will be examined and directly compared with EPS predictions. This section begins with a summary of the results from previous simulation studies, and concludes by presenting a new robust method for estimating accretion onto simulated haloes and subhaloes, and an alternative halo accretion algorithm.

3.2.1 Previous simulation analyses

To date, several authors have defined prescriptions for computing accretion onto haloes using dark-matter-only simulations:

- Wechsler et al. (2002)—henceforth W02—identified the mass accretion history of ∼14 000 haloes at $z = 0$ using the ART code (Kravtsov et al., 1997) in a WMAP1 cosmology. Using their algorithm, W02 found that the accretion histories of their present day haloes were, on average, well fitted by:

$$M_{\rm H}(z) = M_0 \exp(-\alpha z), \tag{3.36}$$

  where M_0 is the present day mass of a halo and $\alpha(z_f)$ is a parameter that describes its formation epoch. According to equation (3.36) the predicted average accretion rate onto a halo of mass $M_{\rm H}$ is:

$$\left\langle \dot{M}_{\rm H} \right\rangle = \langle \alpha(z_f) M_{\rm H} \rangle \left| \frac{{\rm d}z}{{\rm d}t} \right|, \tag{3.37}$$

  which can be compared to the EPS prediction in equation (3.28):

$$\left\langle \dot{M}_{\rm H} \right\rangle = f(M_{\rm H}) M_{\rm H} \left| \frac{{\rm d}\delta_c(z)}{{\rm d}z} \frac{{\rm d}z}{{\rm d}t} \right|. \tag{3.38}$$

  Ignoring the slight mass dependencies of the $\alpha(z_f(M_{\rm H}))$ and $f(M_{\rm H})$ terms, it can be seen that equation (3.36) is a sensible fit for W02 to have chosen because in the case of an Einstein–de Sitter universe, where ${\rm d}\delta_c/{\rm d}z = 1.686$ (equation 3.35), their $\left\langle \dot{M}_{\rm H} \right\rangle$ has the same $M_{\rm H}\dot{z}$ dependence as equation (3.38), differing only in normalization.

- van den Bosch (2002) used the N-branch merger tree algorithm of Somerville & Kolatt (1999)

and found that a two parameter fit better described the mass accretion histories of his haloes, although M06 demonstrated that this two parameter fit becomes unphysical locally as it predicts that present day haloes are not accreting mass. van den Bosch (2002) also provided a relation for α and z_f that can be used in equation (3.36):

$$\alpha = \left(\frac{z_f}{1.43}\right)^{-1.05}, \tag{3.39}$$

but it is more common to define z_f as the epoch at which the present day halo of interest had half of its present day mass:

$$z_f = \frac{\ln 2}{\alpha}. \tag{3.40}$$

- More recently McBride et al. (2009) investigated the mass accretion histories of ∼500 000 haloes from the Millennium Simulation with $M_{\rm H} > 10^{12}\,{\rm M}_{\odot}$ and $0 \leq z \leq 6$ and found that only ∼25% were well described by equation (3.36). They introduced a second parameter, β, and showed that a function of the following form provided a better fit to the halo accretion trajectories:

$$M_{\rm H}(z) = M_0(1+z)^{\beta}\exp(-\gamma z). \tag{3.41}$$

- Fakhouri et al. (2010) used a joint dataset from the Millennium I and II Simulations and found that equation (3.41) held across five decades in mass up to $z = 15$.

These listed accretion fits only apply when averaged across all environments. In order to understand accretion in dense regions such as galaxy groups and clusters, one must resolve substructure and design an accretion algorithm that can account for accretion onto haloes and all levels of substructure. The difficulties in devising such an accretion algorithm are two-fold: firstly, it should define a single progenitor for each and every (sub)halo which accurately represents that object at earlier epochs, and secondly, it must conserve mass (which becomes harder to do when one introduces subhaloes). An algorithm satisfying these conditions is now presented. The mass of a halo or subhalo henceforth corresponds to the total mass detected by the AHOP and MSM halo-finders (i.e. the sum of all of the mass of the member particles belonging to a given halo or subhalo).

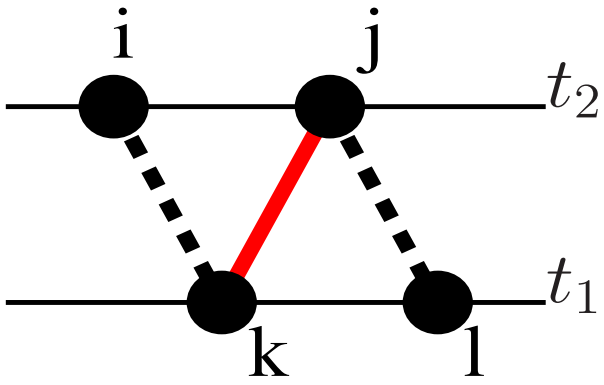

Figure 3.1: Halo i and j at time-step t_2 merge and form two different haloes, k and l, at the later time-step t_1. Halo k's main father is j and j's main son is k, and their main branch connection is shown by the solid red line.

3.2.2 A simple merger

A mock merger scenario is presented in Fig. 3.1. In order to compute the mass accretion rate onto k, one must define k's 'main branch' and various authors have adopted different prescriptions for identifying this channel (Springel et al. 2001a; W02). W02, for example, define halo j as the main father of k if k receives most of its mass from j, but they also require j's most bound particle to be included within k if the mass of j is not at least half of k's mass. These rules hence force each halo to have a single main son and a single main father. If the main son of j—halo k—has j as its main father, then the main branch for k is defined (shown as the solid red line in Fig. 3.1).

There is freedom to choose the main father (son) of k as either the physically most massive father (son), or the father (son) that contributes (receives) the most mass. The latter definition is adopted throughout this thesis as it has the advantage that it explicitly follows the exchange of mass across time in a dynamical manner, and therefore conserves mass by construction and better traces mergers. The former definition, however, is also briefly considered in an alternative accretion algorithm presented in Section 3.2.5, and the predicted halo accretion trajectories across time found using both of these main branch schemes are compared in the following chapter.

Both of the accretion algorithms in this thesis aim to measure the positive growth of bound structures, and do not include prescriptions for measuring mass loss via stripping. The motivation for making this choice is justified in Sections 3.2.5 and 3.2.6.

3.2.3 Anomalous events

Anomalous events describe haloes that spatially coincide at one time-step and then separate at later time-steps. These haloes might take several more time-steps to form a bound merger halo or they might never coincide again. One must hence be careful that their accretion estimator accounts for accretion onto bound objects only. To illustrate this point further, one would naïvely expect the mass accretion rate of halo k in Fig. 3.1 at time-step t_1 to be:

$$\dot{M}_k = \frac{M_k - M_j}{t_1 - t_2}, \tag{3.42}$$

but there were a large number of negative accretion events when this was applied to all of the haloes at t_1 (which was chosen to be the $z = 0$ output from the HORIZON simulation). This prediction is problematic as it is seemingly in tension with the hierarchical paradigm of halo growth, which argues that haloes, on average, should not be losing mass. The large mass loss signal arises because artificial anomalous merger events are included in equation (3.42)[1]. These unbound objects contaminate the 'real' mass loss signal, which is driven by material stripped from bound objects during mergers. As explained later in this section, the accretion estimators used in this study only measure positive mass accretion, hence it is crucial to wait until a system of objects undergoing a complex merger has formed relaxed, bound objects at later times, and to pinpoint the time interval during which mass is accreted onto these relaxed objects.

Identifying anomalous events

Testing to see whether an object is bound is one definitive way of excluding such fake events and it is common practise to sum the kinetic and potential energies of each object and disregard those objects whose total energy is positive (Maciejewski et al., 2009). These erroneous haloes and subhaloes arise because the AHOP and MSM finders resolve objects according to the peaks in the density field (Section 2.2), with no consideration of the total energy of particles in the halo and subhalo resolving process. In this thesis, this binding energy technique is combined with an independent anomalous detection method to identify unbound objects at each redshift.

[1]It is also likely that anomalous events contribute positive accretion signals at other time-steps: the sign of their contribution is determined by the stage of their evolution (e.g. Fig. 3.2). In either case, anomalous events represent false signals and must be excluded from measured accretion rates. It should also be noted that even if the amount of mass lost along the main branch over a single interval were exactly equal to the mass accreted onto it over the next interval, so that the total change in mass were zero over the combined intervals, it is unlikely that the total mass accretion *rate* along the main branch would be zero as the time intervals for each of these events would be different.

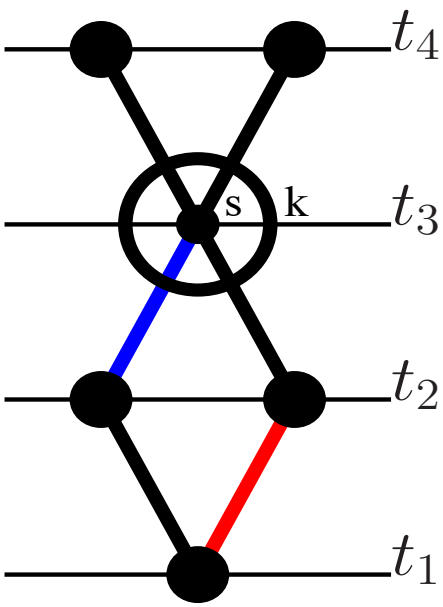

Figure 3.2: Two haloes at t_4 ($< t_3 < t_2 < t_1$) merge and appear to form a MSM halo (k) hosting a subhalo (s) at t_3. At the subsequent time-step t_2, however, s donates its mass to a different halo (blue line) compared with the son of its host, and the system returns to its original configuration. The objects at t_3 are hence recorded as participating in an anomalous event. The true mass accretion event eventually takes place during the $(t_1 - t_2)$ interval, and the main branch is shown by the solid red line.

The prescription adopted for identifying objects participating in anomalous events is illustrated in Fig. 3.2. The fathers of an object k at time-step t_3 are found and if object k has two or more halo fathers that each donate a mass $M_{\rm D} \geq 20M_p$, then object k is flagged as a possible fake merger candidate. A minimum threshold of $20M_p$ is chosen here rather than $40M_p$ used in later sections, because $20M_p$ is a common mass resolution limit used in other simulation studies (e.g. Springel et al. 2005) and it also maximizes the number of possible anomalous events, thereby ensuring that the measured accretion signals correspond to mass accretion events onto bound objects. The sons of k are then found and if k donates a mass $M_{\rm D} \geq 20M_p$ to two or more haloes, then it has fragmented and it is identified as an anomalous event candidate. In the case of AHOP haloes from the HORIZON simulation in this study, which average over their environment and whose substructure is not resolved, this is the sole anomalous event criterion and the same criterion is then imposed on the next halo at time-step t_3.

For subhaloes an additional condition is imposed. Imagine that two haloes at time-step t_4 merge to form a halo, k, which hosts a subhalo, s, at the subsequent time-step t_3, and that halo k and its subhalo are then detected as separate haloes at the following time-step t_2 ($t_2 > t_3 > t_4$). This system has transitioned over two time intervals from two haloes, to a halo and a subhalo, and

back to two haloes again, and is hence an anomalous event, as no merger has taken place. The substructures of a given host halo are therefore also examined if the host does not fragment. If a subhalo at t_3 donates a mass $M_{\rm D} \geq 20M_p$ to a halo at t_2 that is a different halo to the halo son of its host (blue line in Fig. 3.2), then it is identified as part of an anomalous event, as are its subhaloes (if it has any) and its host. The key ideas of this anomalous event detection method are therefore:

- to search for channels that receive/donate at least $20M_p$ from/to two or more different haloes;
- to ensure that the host and all associated substructures are flagged in the case of any one of these objects being classified as participating in an anomalous event.

Identifying unbound objects

Table 3.1 illustrates the relative importance of unbound MSM objects above the mass threshold in the HORIZON simulation ($M \geq 40M_p$) for each of the redshifts shown in column 1 (these redshifts have been chosen because the number of subhaloes increases with decreasing redshift in the simulation, as galaxy groups and clusters form). The percentages in Table 3.1 express the number of objects above the threshold mass satisfying the condition in each column as a fraction of the total number of objects above the threshold mass at the redshift in question, with the exception of the bracketed values in column 3, which show the fraction of anomalous events that are unbound.

There is a positive correlation between the independently identified anomalous events and unbound objects, with a large fraction of the anomalous events being unbound (henceforth 'unbound' refers either to an object with total energy $E_{\rm T} \geq 0$ or an object participating in an anomalous event, or both). Not all objects in column 3 have $E_{\rm T} \geq 0$, however, and so there is a small population of unbound objects at each redshift that would be missed if just a requirement of $E_{\rm T} \geq 0$ were imposed on every object.

Only bound objects above the mass threshold can have a recorded accretion value in this study, despite ∼38% of all the objects at each of the redshifts shown in Table 3.1 having a mass below the chosen threshold limit. Bound objects below threshold, however, are not removed from the sample and so it is possible for a bound object with $M < 40M_p$ to be a main father. This method hence avoids possible biasing of the accretion events in the simulation, whilst ensuring that only well-resolved objects have an accretion value.

Table 3.1: The relative importance of unbound MSM haloes and subhaloes above the threshold mass ($M \geq 40M_p$) in the HORIZON simulation.

Redshift	$E_{\rm T} \geq 0$ %	Anomalous ($E_{\rm T} \geq 0$) %	Objects with recorded accretion %
0.49	23.1	7.85 (84.8)	73.5
0.23	23.7	8.57 (87.3)	73.2
0.01	24.1	9.14 (89.7)	73.4

Column 4 shows the fraction of objects above the mass threshold with a recorded accretion value: ∼27% of the objects above threshold at each of the redshifts shown do not have a measured accretion rate because they are either unbound, or do not satisfy some additional criteria imposed by the accretion algorithm, which are explained in the following section.

3.2.4 The accretion algorithm

In detecting substructure, Springel et al. (2001a) required that several of the most bound particles of the main father were included in the main son—this was more robust than tracking the evolution of the single most bound particle, which essentially performs a random walk across time. The main son in this thesis is defined as the son that receives the most mass from the object of interest, thereby preserving consistency with the main father definition.

The algorithm that is used to compute accretion onto haloes and associated substructures is henceforth referred to as the 'halosub' method and is illustrated in Fig. 3.3. The routine begins by identifying the main son k (solid blue line) for object i at time-step t_2. Using the main son definition this means that most of i's mass goes to k and the remainder goes to m and p. The father that contributes the most mass to k is then found and in this example j is the main father (solid red line). The mass accretion onto k is therefore $(1 - f_j)M_k$, where f_j is the fraction of k's mass that comes from object j. Object k is then flagged and the accretion onto the other sons of i, m and p, is considered. Since m is not the main son of i and m does not have any other fathers, an accretion value for m is not recorded and it is flagged as an orphan. On the other hand, if one of the other sons, p, of the object of interest does experience mass accretion, its main father, q, is identified and the mass it accretes is recorded as $(1 - f_q)M_p$. Object p would then also be flagged. The principal features of the halosub method are summarized below:

- the measured mass accretion onto an object represents the sum of diffuse accretion (material

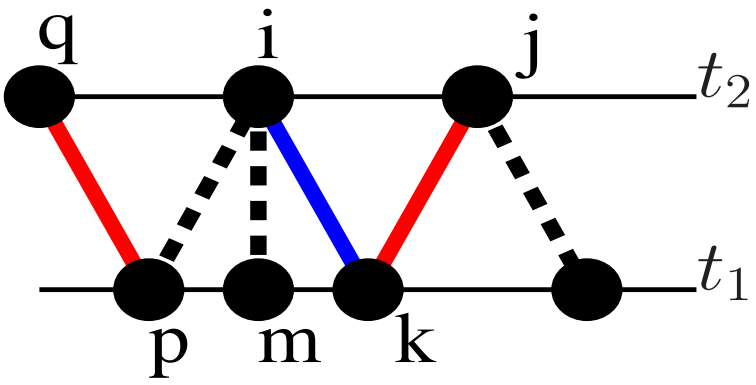

Figure 3.3: A schematic illustrating the halosub accretion algorithm that accounts for accretion onto haloes and subhaloes. In this example object k, whose main father is object j (solid red line), has been identified as the main son of object i (solid blue line). The mass accretion onto k using the halosub method is therefore $(1 - f_j)M_k$, where f_j is the fraction of k's mass that comes from j. Object m is not the main son of object i and because it does not have any other fathers, it is skipped. If q is the main father of p, then the mass accretion onto p is $(1 - f_q)M_p$. The halosub method therefore only ever records zero or positive accretion rates.

not bound to any resolved structure) and merger-driven growth;

- mass loss events are considered to be zero accretion events—measured accretion signals in this study are never negative;
- all objects with a recorded accretion value are bound and have a mass $M \geq 40M_p$, and;
- no distinction is made between haloes and different levels of substructure.

Since no attempt is made to measure the mass lost from an object during a given time interval, the accretion rate in this study can be thought of as an upper limit. Note that objects which only lose mass and have a recorded accretion rate of zero are identified as systems where the bound main son of the object of interest has only one bound father (which *is* the object of interest). If the halosub method encounters a flagged object it means that either the accretion onto that object has already been accounted for or that object has been identified as an orphan.

The halosub method has been used throughout this thesis, but before its limitations are discussed, a second accretion algorithm is considered.

3.2.5 An alternative accretion algorithm

In this subsection, another accretion algorithm that uses different definitions for the main son and main father to those adopted in the halosub method, is presented. The motivation for introducing

this alternative 'halo method' scheme, which is used to compute accretion onto haloes only, is that with two different accretion algorithms it is possible to directly test whether the measured halo accretion rates are sensitive to the prescription chosen to define the main branch. The halo method has the following three properties:

- it records mass accretion onto unflagged, bound haloes above the 40 particle mass threshold;
- it only measures positive accretion, which can involve tracing back several time-steps, and;
- it ensures that each halo accretor along the main branch is the most massive son of its most massive father.

Halo accretors identified by the halo method hence satisfy similar constraints to those identified using the halosub method, which allows a meaningful comparison to be made between the two algorithms. The conditions above also highlight the underlying principle of the halo method, which is to link the most massive father of a halo to its most massive son and to measure the positive accretion onto this main son.

These criteria are now explained in more detail with reference to a simple example shown in Fig. 3.4. The halo method routine starts by considering a given halo i at the penultimate time-step t_1 from the simulation, corresponding to $z \sim 0$ (the final $z = 0$ output cannot be the starting time-step because it is not possible to use the anomalous method to test for unbound haloes, as discussed in Section 3.2.3). If i is unbound or below the 40 particle mass threshold, it is skipped, otherwise its bound main father is found, which is defined as the most massive halo of the halo fathers. There are two possible scenarios: either i has just a single father, or it has several. If in the former case the father belongs to the background, then the mass accretion rate onto i is recorded as $\dot{M}_i = M_i/(t_1 - t_2)$ because i was born during the time interval. In the latter case, i's most massive halo father j at the earlier time-step t_2 is found.

Assuming that j is bound, the condition that i is the most massive son of j is also imposed. If this criterion is not satisfied, then the halo method is unable to define a main branch for i (shown as red lines in Fig. 3.4), and moves to the next halo at time-step t_1. If the condition is met, however, then i is the labelled as the main son of j and $\Delta M_{ij} = M_i - M_j$ is computed. If $\Delta M_{ij} \geq 0$, the mass accretion rate of i is recorded as $\dot{M}_i = (M_i - M_j)/(t_1 - t_2)$, whereas if $\Delta M_{ij} < 0$, j is flagged and its bound main father is subsequently found, halo l. The mass difference $\Delta M_{il} = M_i - M_l$ is then computed, provided that j is the main son of l, and the process above is repeated. If a main

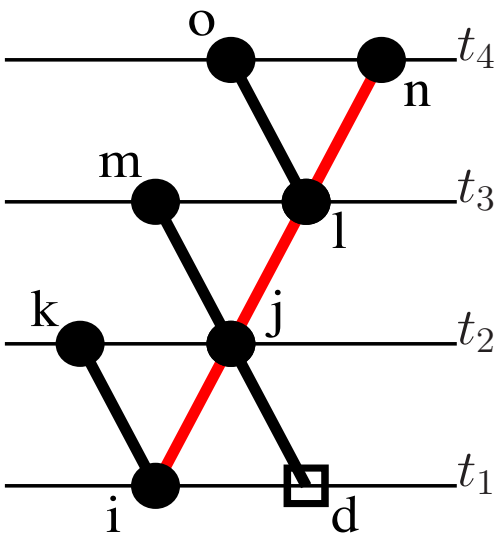

Figure 3.4: An example of a simple merger scenario that illustrates the halo method for computing accretion onto bound haloes above the mass threshold. The most massive father of halo i at t_1 ($z \sim 0$) is j, and i is the most massive son of j. A main branch connecting the main father and main son is hence defined and is shown as a red line. If the mass difference between i and j is not positive, the main father is traced back to earlier time-steps. Negative mass loss events along the main branch can occur if mass is stripped during a merger, contributing to the diffuse background component (as shown by the open box labelled 'd'). In the example above, the net positive mass accretion onto i occurs during time-steps t_1 and t_4, because $M_i > M_n$. The mass accretion rate onto i is hence recorded as $\dot{M}_i = [M_i - M_n]/[t_1 - t_4]$ (note that the halo circles are not of equal mass, despite their equal radii, and that the time intervals are not equally spaced). The bound haloes j and l are flagged because they are more massive than i, and are skipped when haloes at time-steps t_2 and t_3 are examined because they experience a net loss of mass along the main branch with respect to i. By contrast, there is a net positive accumulation of mass onto n as it undergoes mergers to form i and hence n is an accretion candidate at t_4.

branch cannot be defined for halo l then it is flagged, i is skipped, and the next halo at time-step t_1 is considered.

While this method is able to yield reliable estimates of halo accretion rates, it is somewhat limited compared with the halosub method for the following two reasons:

- It defines the main father and main son according to mass and not the mass they donate or receive (as in the halosub method), and hence suffers from a mass bias that penalizes subhaloes.

  Fig. 3.5 shows two haloes M_{10} and M_8 that merge at t_2 and form a halo M_{12} and a subhalo M_6 at the subsequent time-step t_1, with the numbers corresponding to masses in arbitrary units. Objects M_{12} and M_6 share the same halo fathers, M_{10} and M_8. The main branch, as defined in the halo method, will connect M_{12} to its most massive father M_{10}. Halo M_{12} is the main son of M_{10} as it is more massive than M_6, and so the main branch links M_{10} with

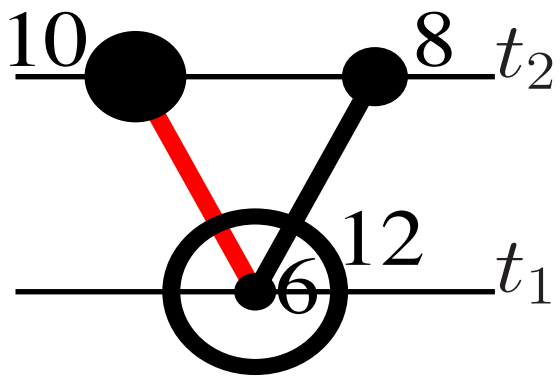

Figure 3.5: This diagram illustrates the mass bias of an algorithm that defines the main father and main son according to mass, with the numbers corresponding to masses in arbitrary units. Two haloes merge at t_2, forming a host halo and a subhalo at the following time-step t_1. If the main branch (red line) connects the most massive father M_{10} with its most massive son M_{12}, the accretion onto subhalo M_6 at t_1 cannot be accounted for because the most massive son is M_{12} at t_1. If instead the main branch connects a given object to its most massive father and removes the main son condition, the accretion scheme no longer conserves mass (see text).

M_{12}, leaving M_6 without a recorded accretion rate. By contrast, the halosub method allows for the possibility that the main branch for M_6 connects to M_8.

- It appears to require a non-trivial variant in order to account for mass accretion onto subhaloes in a mass conserving manner.

  If the requirement that an object has to be the most massive son of its most massive father were relaxed in an attempt to remove the mass bias described above, the measured mass accretion is no longer guaranteed to conserve mass, a condition that must hold at each time-step because the simulation is a closed system with no external sources or sinks of particles. Under this new scheme, objects M_{12} and M_6 would share the same main father M_{10} and have recorded mass accretions of $+2$ and -4 units respectively, yielding a total mass loss rate of 2 units. If the main father is instead defined as the father that donates the most mass, then it is possible that M_8 is the main father of M_6. In this scenario, subhalo M_6 would lose two units while M_{12} would gain two units, and so the total change in mass would be zero, as required by mass conservation. It is also possible, however, that M_6 receives more mass from M_{10} than M_8, and that the main progenitor of both M_6 and M_{12} is once again M_{10}, which violates mass conservation. The advantage of the halosub method is that it circumvents the difficult task of computing the total change in mass for a given halo or subhalo (or sub-subhalo etc) under detailed mass conservation, by measuring only the positively accreted component over

a single time interval in a mass conserving manner (as argued in the previous section).

It should be noted that if subhalo M_6 were not resolved and just a single halo of mass 18 units were to form at t_1, the halo method would predict an accretion rate of 8 units, which would perfectly conserve mass (the total accretion rate onto existing objects need no longer sum to zero in this case, since the number of objects has decreased). The halo method presented in this section is hence only suitable for computing accretion onto haloes, and not higher levels of structure.

3.2.6 Limitations

Other than finite mass and time resolutions, which are shortcomings of any simulation, the growth of haloes and subhaloes in a ΛCDM universe is modelled without a prescription for the gas physics. The dark-matter-only simulation satisfies the objective of this study, however: to determine whether halo and subhalo accretion is dependent on environment. Both accretion algorithms exclude tidal stripping from the measured accretion rate, but this is justified when drawing comparisons with EPS theory (which ignores mass loss) and when probing the link between halo mass accretion and AGN activity, a process that is driven by positive, not negative, accretion events. At any rate, objects are stripped of mass in the simulation as they undergo mergers, and this reduces their mass.

Having acquired an algorithm capable of measuring positive accretion onto all levels of structure, one is suitably equipped to test its predictions for accretion onto haloes and subhaloes as a function of mass, time, and environment, drawing comparisons with previous simulations studies and the EPS formalism. As alluded to above, it is also interesting to speculate on whether the growth of haloes can drive the activity of their associated black holes. These processes are examined and discussed at length in the following chapter.

Chapter 4

The environment and redshift dependence of dark matter halo and subhalo accretion

When a halo merges with a group of galaxies and changes its status to a subhalo, assuming it survives, does the rate at which it accretes mass differ significantly compared with the accretion rate onto a subhalo of similar mass in a much less dense environment? What exactly is meant by 'environment', and how is it best quantified? Do the measured halo accretion rates in dark-matter-only simulations show any discrepancy with the corresponding halo accretion rates derived using standard EPS theory? Is there evidence of a slowdown in accretion onto more massive haloes compared with their low mass counterparts as the simulation evolves towards the present day, thereby implying that the AGN downsizing discussed in Chapter 1 could be imprinted in the underlying halo mass accretion history? By measuring the rate of accretion onto haloes as a function of time, is it possible to reproduce the integrated cosmic black hole accretion history required by various semi-analytic prescriptions in order to yield agreement with the present day galaxy luminosity function?

All of these questions, among others, are now addressed. The forepart of this chapter presents the results found by applying the halosub accretion algorithm described in Chapter 3 to the haloes and subhaloes from the HORIZON simulation. A detailed discussion of the scientific implications of the results follows. Throughout this chapter:

1. 'object' refers to haloes and/or subhaloes, which are detected using the AHOP and the MSM methods described in Section 2.2;
2. the mass of an object corresponds to the total mass, $M_{\rm T}$, detected by the halo-finder;
3. accretion onto bound objects above the mass threshold, $M \geq 40M_p$, is examined;
4. the measured mass accretion is the sum of diffuse- and merger-driven accretion: mass loss has not been measured;
5. $\mu \equiv \dot{M}/M$ denotes the specific accretion rate in units of Gyr^{-1} onto an object of mass M;
6. $\delta \equiv \delta M_{\rm H}/M_{\rm H}$ (and hence does not refer to the density contrast), where $M_{\rm H}$ represents the mass of a halo.

4.1 Accretion onto all levels of halo structure

4.1.1 Comparing halo accretion with EPS

Sheth & Tormen (1999) demonstrated that standard EPS theory is not able to accurately reproduce the mass function of dark matter haloes formed in simulations, yielding too many haloes of low mass and too few haloes of high mass. Since then, several authors have devised parameter fits to better reproduce the N-body halo mass function (Sheth et al., 2001; Jenkins et al., 2001), and the most recent calibrated expression is accurate to $\leq 5\%$ (Tinker et al., 2008). The aim of this subsection is to further these findings by probing any possible discrepancy between EPS and N-body predictions for halo accretion rates.

Fig. 4.1 shows the average accretion rate onto the AHOP haloes from the simulation computed using the halo (dotted lines) and halosub (solid lines) algorithms, as a function of redshift and halo mass. Haloes with recorded accretion values are binned in mass at each redshift and the average accretion rate for each mass bin is computed. Averages of the corresponding mass bins over redshift then yield constant $\langle M_{\rm H} \rangle$ values, which for the halo and halosub method are shown in the top left and bottom right corners respectively (for an alternative $\langle M_{\rm H} \rangle$ computation, see W02, who binned the $z = 0$ haloes in mass and then averaged over all the accretion trajectories in each mass bin at each redshift). The error bars indicate the 1σ errors on the mean accretion rate, and the EPS predictions for each of the halosub $\langle M_{\rm H} \rangle$ bins have been computed using equation (3.28) and are shown as the dashed lines.

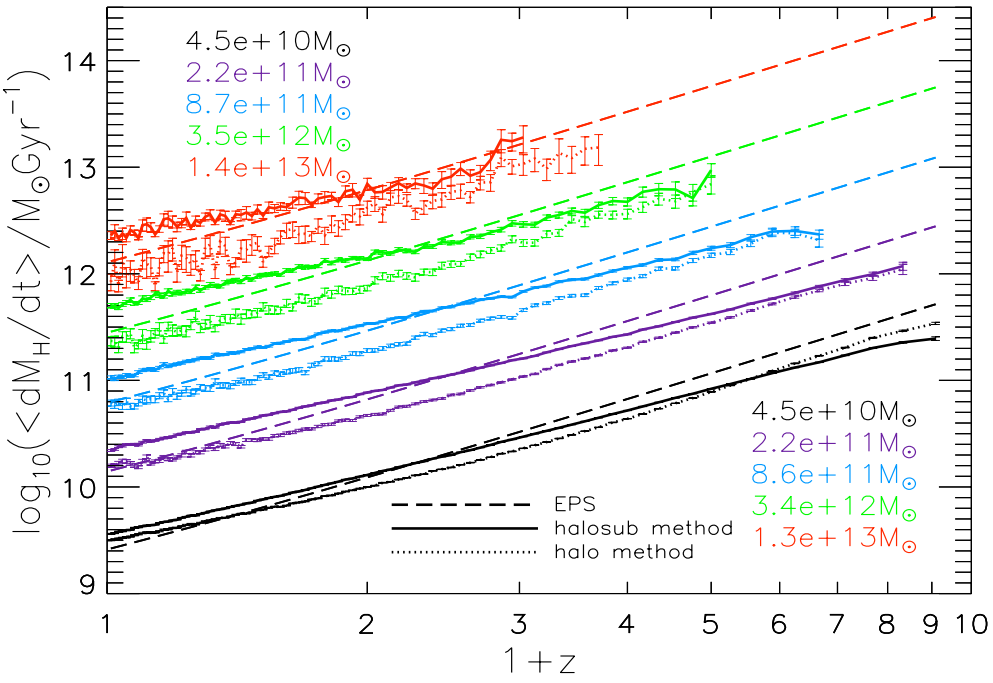

Figure 4.1: The average halo accretion rate as a function of redshift and halo mass. The accretion values onto the AHOP haloes using the halosub and halo methods are shown as the solid and dotted lines respectively for each of the five halo (top left) and halosub (bottom right) $\langle M_{\rm H} \rangle$ bins, with the errors corresponding to the 1σ errors on the mean accretion rate. The dashed EPS curves have been computed by inserting the values of the halosub mass bins into equation (3.28).

It can be seen from Fig. 4.1 that the halosub mass trajectories have a lower gradient across redshift than the EPS curves, which overestimate the accretion rate onto the lowest mass haloes in the simulation at high redshift by a factor of $\sim$2, and underestimate it by a factor of $\sim$1.6–1.8 at $z = 0$. It is tempting to think that the enhanced accretion onto haloes with respect to EPS theory at low redshift results from the exclusion of mass loss in the halosub algorithm. However, EPS does not account for mass loss from haloes either: haloes only grow with time, by construction. The offset with EPS should therefore be regarded as an offset in gradient, and Fig. 4.1 shows that the EPS curves are a factor of $\sim$1.3 steeper than the corresponding halosub trajectories.

The halo method yields accretion gradients that are in better agreement with EPS, but there is roughly a factor of 2 normalization offset for $z \gtrsim 1$, and slow convergence to EPS predictions towards the present day. By contrast, the gradient offset between the halo and the corresponding halosub curves increases towards $z = 0$. Given the similarity in the average mass bins, it is likely that both methods are measuring accretion onto similar haloes, and that the discrepancy does not therefore arise from differing main branch identification, but rather a difference in the time interval considered for the accretion. The offset is probably an environment-dependent effect: the

halo method has to trace back to earlier time-steps at low redshift in order to measure positive mass accretion onto haloes, and this time interval is likely to increase with decreasing redshift as haloes merge into denser group- and cluster- environments, where they experience more tidal stripping than overall positive mass growth. An increasing discrepancy between the halo and halosub methods with halo mass at low redshift lends further support to this view, as more massive haloes are more biased and hence populate more massive environments (Sheth & Tormen, 1999).

Evidently the level of agreement with EPS appears to be somewhat dependent on the adopted prescription of halo accretion. It can be argued that both methods are valid, but the halo method is henceforth abandoned due to its probable time biasing mentioned above, and the fact that it is insensitive to positive mass accretion onto haloes whose absolute mass decreases across time (for further comparisons between the two algorithms, see Section 3.2.5). A discussion regarding the possible sources of the EPS discrepancy and the implications of Fig. 4.1 for the relationship between black hole and halo growth, is reserved for Sections 4.2.1 and 4.2.3 respectively.

4.1.2 The different modes of accretion

The mass accreted onto the AHOP haloes in Fig. 4.1 is the summed contribution of diffuse accretion events and minor- and major-merger events, and so these accretion modes are decoupled in this subsection in an attempt to gauge their relative importance in driving halo growth as a function of redshift. It is argued that the cosmological evolution of the integrated low accretion events drives the 'radio-mode' of black hole activity (Croton et al., 2006).

In the upper panel of Fig. 4.2 the dimensionless quantity δ ($\equiv \delta M_{\rm H}/M_{\rm H}$) is computed for each accretion event at each redshift: the dashed lines and the thin solid lines respectively show haloes with $\delta \leq 0.02$ (minor-mergers & diffuse accretion) and $\delta \geq 0.08$ (likely to be dominated by major-mergers, but also includes diffuse accretion events). The total mass accretion rate per comoving cubic Mpc for haloes in a given mass bin and of a given δ at each redshift is then computed. The thick solid lines show the total mass accretion rate per comoving cubic Mpc integrated over all δ. For a given linestyle, the lower mass curves shift to higher redshifts.

At high redshift, all haloes are found to accrete mass in high fractional events with the peak in activity shifting to lower redshifts for more massive haloes. As the mass accreted onto the lowest mass haloes via minor-mergers and diffuse accretion starts to plateau at low redshift, minor-merger and diffuse accretion activity onto the more massive haloes starts to rapidly accelerate:

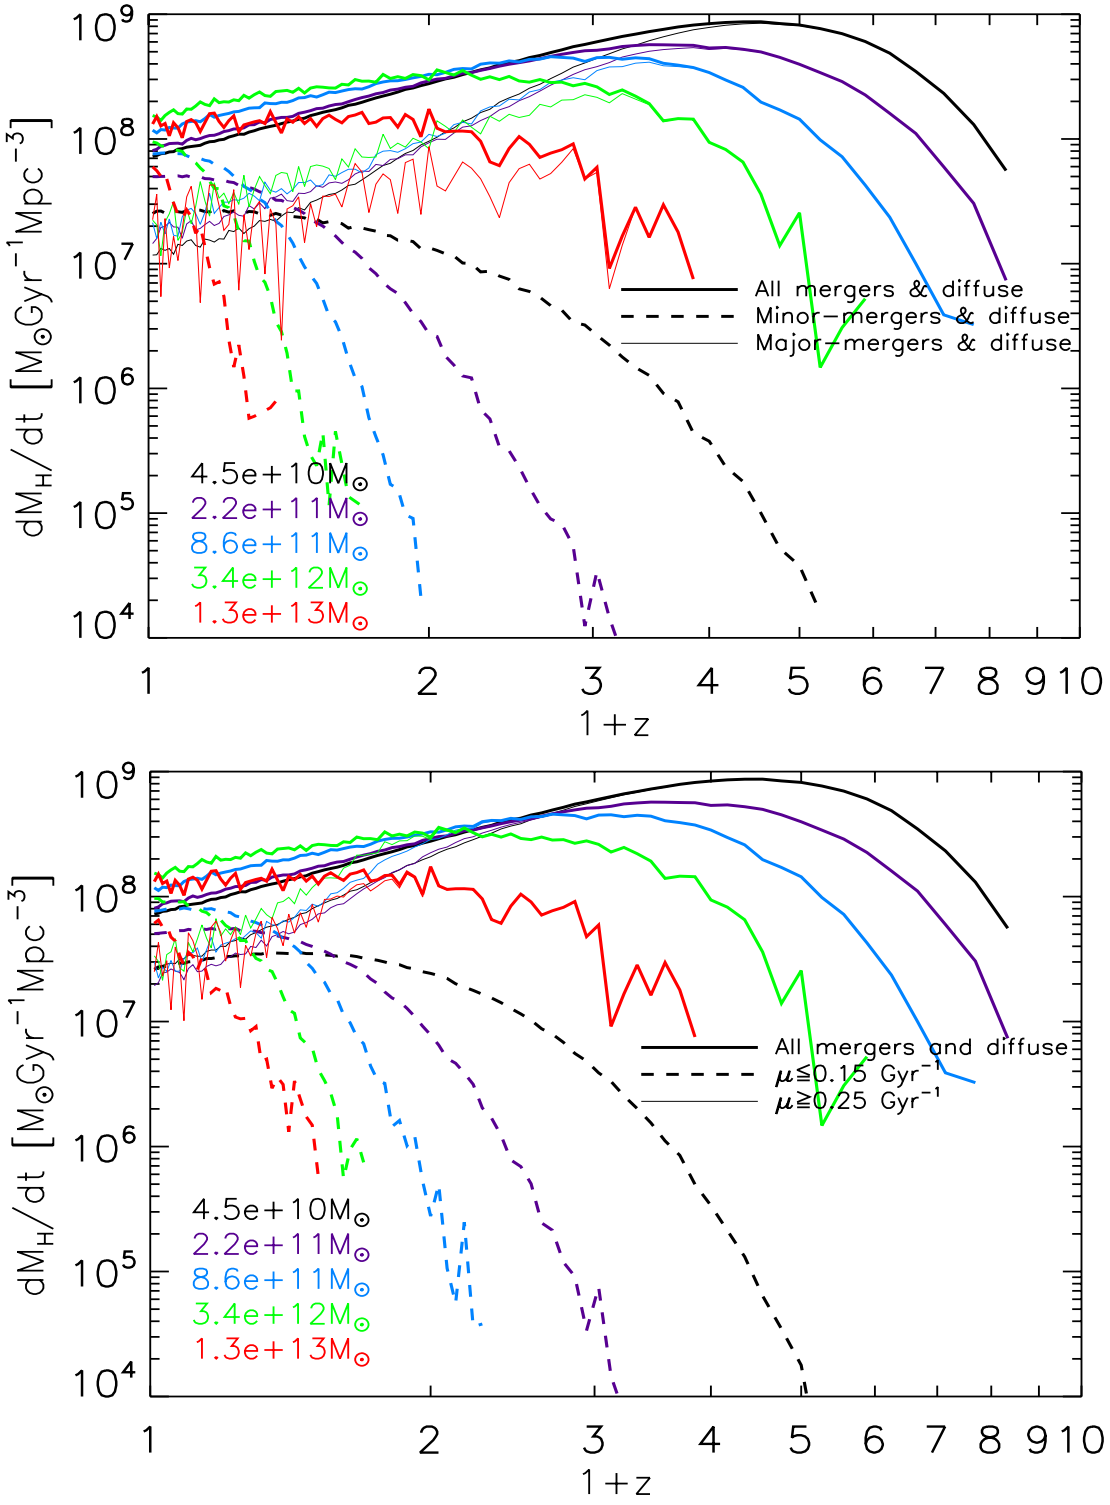

Figure 4.2: The total mass accretion rate onto the AHOP haloes per comoving cubic Mpc as a function of redshift, halo mass, δ ($\equiv \delta M_{\rm H}/M_{\rm H}$) and μ ($\equiv \dot{M}_{\rm H}/M_{\rm H}$). Both sets of mass bins correspond to the $\langle M_{\rm H}\rangle$ halosub bins in Fig. 4.1, with the lower mass curves shifting to higher z. In the top (bottom) panels, the dashed and thin solid lines show each mass bin decomposed into haloes with $\delta \leq 0.02$ ($\mu \leq 0.15\,{\rm Gyr}^{-1}$) and $\delta \geq 0.08$ ($\mu \geq 0.25\,{\rm Gyr}^{-1}$) respectively, and the thick solid lines show the accretion trajectories integrated over all δ (μ).

low mass haloes and non-halo material are being accreted onto larger structures. By $z = 0$, the combined minor-merger and diffuse accretion signals dominate the growth of all haloes. The dashed curves also have a strikingly similar cosmological evolution to the radio-mode integrated black hole accretion rate density curves found by Bower et al. (2006) and Croton et al. (2006), but a more detailed discussion is postponed until Section 4.2.4.

Fig. 4.3 shows the shift from major-merger and diffuse- dominated growth at high redshift to minor-merger and diffuse- dominated growth at low redshift, more clearly. The linestyles have the same meaning as in Fig. 4.2, except haloes with $0.02 < \delta < 0.08$ are also included, shown by the dotted lines. It can be seen that minor-mergers and diffuse accretion events start to significantly contribute to growth for $z < 0.5$, and by $z = 0$ drive accretion onto all halo masses.

Very similar results to those shown in the top panels of Figs 4.2 and 4.3 are found when haloes are binned in μ ($\equiv \dot{M}/M$) instead of δ, but the major-merger and diffuse curves decouple from the integrated curves at later epochs for all masses, as is illustrated in the bottom panels of Fig. 4.2 and 4.3. This is probably because in transitioning from δ to μ, one must divide δ by the time interval during which mass is accreted, and at higher redshifts this time interval is smaller (time is not a linear function of redshift) and μ is hence larger than it is for a given δ onto a halo of fixed mass at lower redshift.

So far it has been assumed that the imposed cuts in δ are capable of separating minor- and major-merger accretion events, and so this assumption is tested in Fig. 4.4 by adopting the more classical progenitor mass ratio definition (e.g. Fakhouri et al. 2010). Each progenitor j of accretor k is assumed to merge in turn with k's main father i, with progenitor mass ratio $\chi \equiv M_i/M_j$, donating $f_j M$ to accretor k at the following time-step, where f_j denotes the fraction of k's mass that comes from j. Events with $\chi \leq 3$ ($\chi > 3$) are recorded as major (minor) mergers. As an example, imagine that accretor k in Fig. 4.5 has three halo fathers a, b and c, and also accretes diffusely from the background d (shown as a filled box). Progenitor b is the main father of k and the main branch is shown by the red line. Two merger events (and hence two mass ratios with respect to b) and one diffuse event are recorded for k, with three corresponding δ measurements along each accretion channel. The dashed lines in Fig. 4.4 combine together diffuse events with $\delta \leq 0.02$ and minor-merger events with both $\chi > 3$ and $\delta \leq 0.02$ (note that diffuse events do not have a recorded mass ratio χ). The thin solid lines group together the diffuse events with $\delta \geq 0.08$ and the major-merger events with $\chi \leq 3$ and $\delta \geq 0.08$. It can be seen that the thin solid

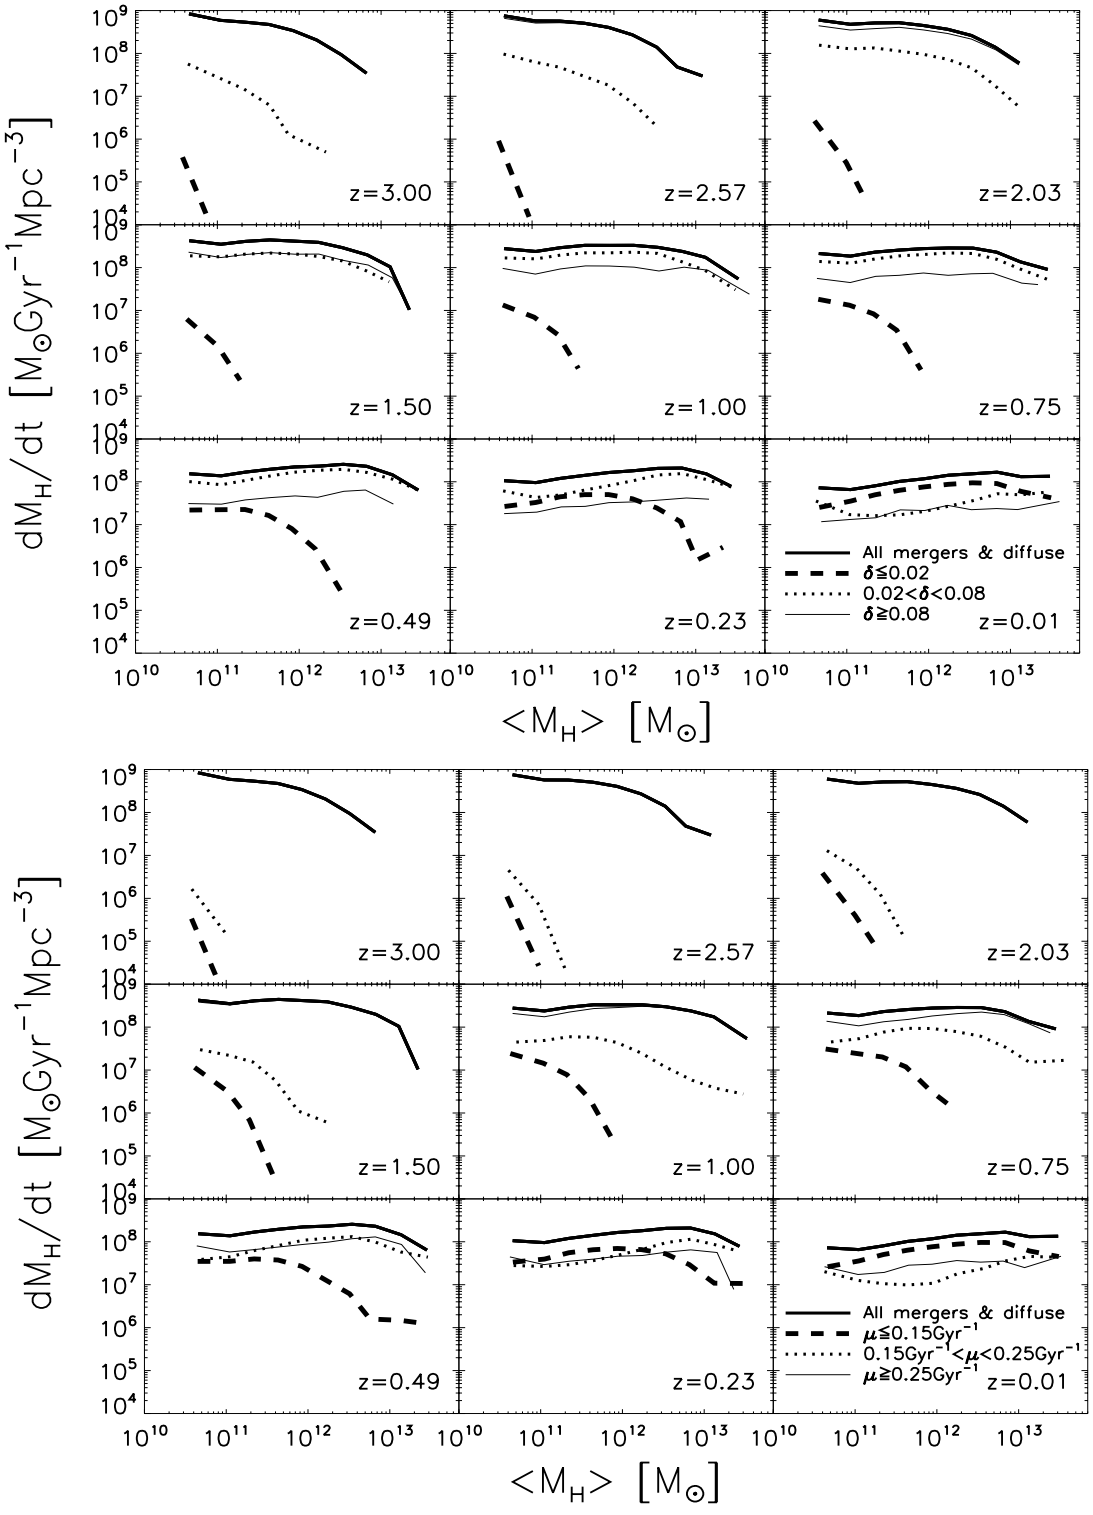

Figure 4.3: The total mass accretion rate onto the AHOP haloes per comoving cubic Mpc as a function of halo mass and accretion mode (denoted by δ and μ), shown for several redshifts. The linestyles have the same meaning as in Fig. 4.2, except that the haloes with $0.02 < \delta < 0.08$ (top) and $0.15\,\mathrm{Gyr}^{-1} < \mu < 0.25\,\mathrm{Gyr}^{-1}$ (bottom) are also included, and are shown by the dotted lines.

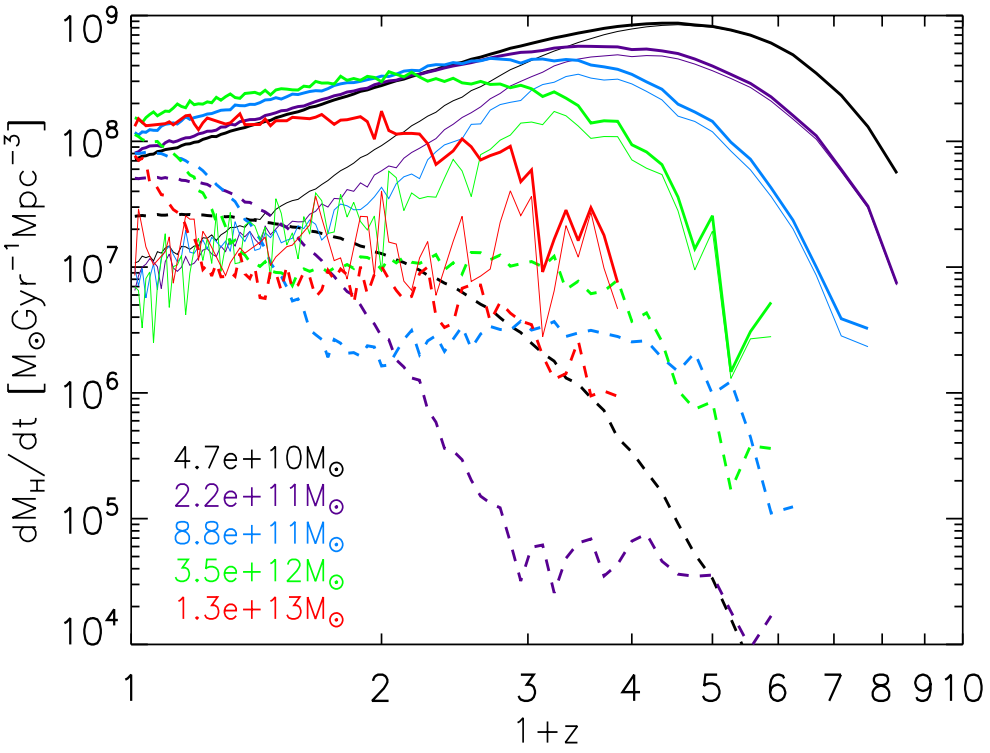

Figure 4.4: In this plot, the ratio χ between the mass of the main father and the mass of each of the other progenitors j is recorded for a given accretor k, with each progenitor j donating a mass $f_j M_k$ to k. The dashed lines represent the combination of $\chi > 3$ and $\delta \leq 0.02$ events (minor-mergers) with diffuse $\delta \leq 0.02$ events (which do not have a recorded mass ratio), whereas the thin solid lines represent the combination of $\chi \leq 3$ and $\delta \geq 0.08$ events (major-mergers) with diffuse $\delta \geq 0.08$ events. As in previous figures, the thick solid lines show the trajectories integrated over all δ.

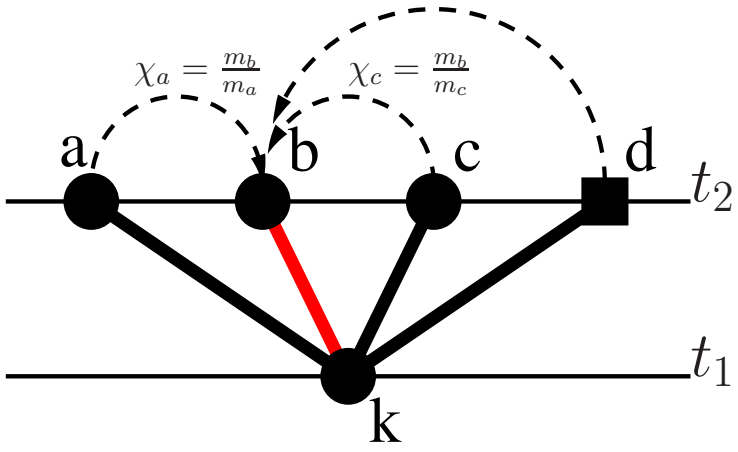

Figure 4.5: A schematic diagram illustrating the method used to test whether δ is able to distinguish between merger-type. In this example, halo k at t_1 accretes from three halo fathers a, b and c and the background d (shown as a box) at t_2, where $t_2 < t_1$. Two merger mass ratios χ_a and χ_c measured with respect to the main father b (red line) are recorded for k, along with three δ values equal to f_a, f_c and f_d, where f_j represents the fraction of k's mass received from progenitor j.

lines in Fig. 4.4 have a very similar cosmological evolution to the $\delta \geq 0.08$ curves in Fig. 4.2. The dashed lines in Fig. 4.4 also show a similar behaviour to the $\delta \leq 0.02$ curves in Fig. 4.2, except there are more minor-mergers at higher redshift for all but the lowest halo mass bin. It will be demonstrated in Section 4.2.4, however, that these features do not affect the main conclusions. Fig. 4.4 hence provides confirmation that the imposed δ cuts are able to accurately distinguish between merger-type.

4.1.3 Accretion onto subhaloes

In this section the AHOP haloes are resolved into constituent MSM haloes and subhaloes and the halosub method is applied to these resolved structures to account for accretion onto objects in groups and clusters. The section begins by comparing the AHOP halo and MSM halo and subhalo specific accretion rates with the results found in the W02 (Wechsler et al., 2002) simulation study. The mass of a halo or subhalo is henceforth denoted by M, in contrast with the previous section, which only recorded accretion onto haloes with mass $M_{\rm H}$.

Comparing the halosub accretion algorithm with W02

Fig. 4.6 plots the average specific accretion rate for all bound objects from the simulation as a function of average object mass for redshifts corresponding to $z = 0.49$ (triple-dot-dashed lines), $z = 0.23$ (dashed lines) and $z = 0.01$ (solid lines). These redshifts have been chosen because the epoch of large group and cluster formation is $z \lesssim 0.5$ (e.g. Mantz et al. 2010). The blue and red lines of a given linestyle represent the accretion onto the AHOP haloes and MSM haloes and subhaloes respectively. The thick black line shows the W02 result at $z = 0.01$, found using equation (3.37) (strictly, equation 3.37 applies to the accretion histories of haloes at $z = 0$ but the anomalous detection method of Section 3.2.3 cannot be used at this redshift). The W02 result is calculated by binning in mass each $z = 0.01$ bound AHOP halo accretor and computing the corresponding average W02 α parameter in equation (3.40) for each mass bin (α is inversely proportional to halo formation redshift).

The specific accretion rate onto the MSM objects is systematically larger than the AHOP specific accretion rates at every mass when considering a given redshift. The MSM method resolves the substructure that has been averaged out in the AHOP halo, so the main MSM host halo and subhaloes are individually less massive than the AHOP counterpart. The offset with MSM is probably caused

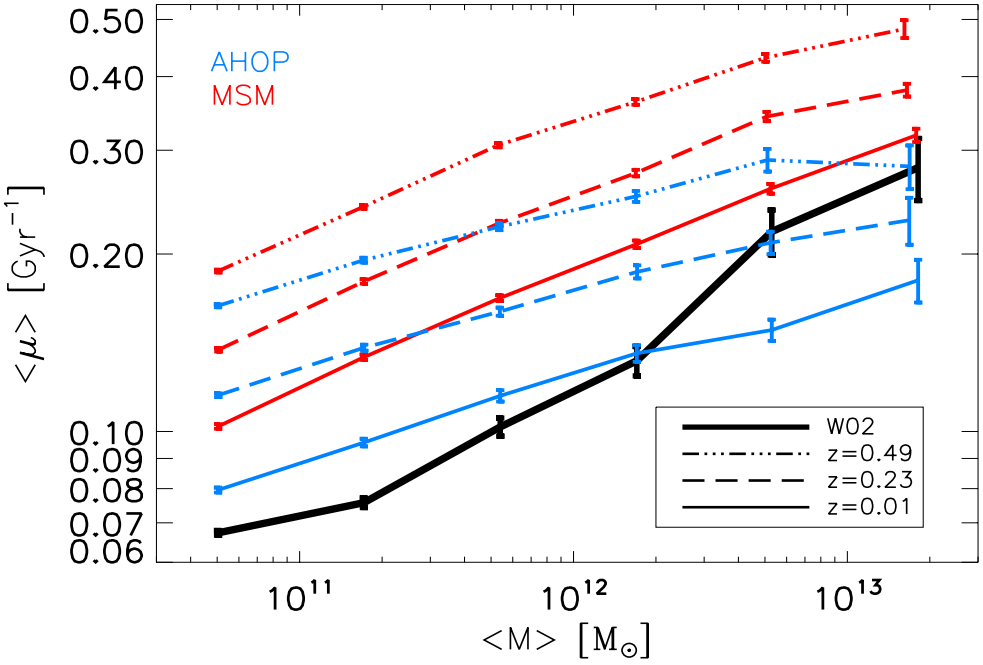

Figure 4.6: The mean specific accretion rate onto haloes and subhaloes using the halosub method, plotted as a function of mass for three redshifts corresponding to $z = 0.49$ (triple-dot-dashed lines), $z = 0.23$ (dashed lines) and $z = 0.01$ (solid lines). The blue and red lines of a given linestyle represent the accretion signals onto the AHOP haloes and the MSM haloes and subhaloes respectively. The thick black line shows the W02 result, obtained by using equation (3.37) and equation (3.40) at $z = 0.01$.

by dividing by the larger AHOP mass, and increases with increasing mass because at larger masses MSM subhaloes occupy a larger fraction of the total AHOP mass (Tweed et al., 2009). The mass difference between AHOP and the main host MSM halo therefore increases with increasing AHOP mass (and there are more detected haloes than subhaloes at a given redshift in the simulation, so the haloes dominate the MSM halo and subhalo accretion signal).

W02 fitted the accretion trajectories of their $z = 0$ haloes averaged over environment in a WMAP1 cosmology and so their result can be directly tested against the AHOP curve at $z = 0.01$ which also averages over environment, but in a universe with a WMAP3 cosmology (W02 argue that their fitting formula does not depend on the chosen cosmology). The W02 specific accretion rate has a stronger mass dependence than is found for the AHOP haloes in this study, and so for the large galaxy- and group- sized dark haloes, overpredicts the specific accretion rate by a factor of $\sim$1.5. Recent studies have shown that some halo-finding algorithms can lead to large uncertainties in the halo accretion rate (Genel et al., 2009; Hopkins et al., 2010). The disagreement across mass with W02 in Fig. 4.6, however, does not result from differences in halo-finder: the AHOP algorithm

is very similar to the modified bound density maxima technique of Bullock et al. (2001) used in W02 (Knebe et al., 2011). The disagreement most likely arises because W02 impose different criteria to identify the main son and main father. They adopt a policy, in some cases, of tracking the single most bound particle, which is misleading as it can hop between haloes undergoing a merger and perform a random walk across several time intervals, before eventually settling onto the relaxed object. By contrast, the anomalous method used in this work rigorously identifies false merger candidates and the halosub accretion algorithm tracks channels which donate/receive the most mass (and recall that allowing a bound object below the mass threshold to be a main father removes potential accretion event bias).

The comparison with W02 has therefore revealed that the predicted accretion rate is somewhat sensitive to imposed prescriptions. The advantage of the halosub method is that it avoids using ad-hoc criteria to define the main branch.

Accretion onto the MSM haloes and subhaloes

Fig. 4.7 shows the specific accretion rate from bottom to top of the MSM haloes, MSM haloes and subhaloes, and MSM subhaloes with the linestyles having the same meaning as in Fig. 4.6. The average specific accretion rates onto haloes ($\mu_{\rm H}$) and subhaloes ($\mu_{\rm S}$) have weak mass dependencies for each of the redshifts shown: $\langle\mu_{\rm H}\rangle \propto M^{0.2}$ and $\langle\mu_{\rm S}\rangle \propto M^{0.1}$ at $z = 0.01$. Each of the halo, halo and subhalo, and subhalo curves shift downwards with decreasing redshift: the average specific accretion rate onto a subhalo at $z = 0.49$ is a factor of 1.3–1.4 greater than at $z = 0.01$. Major-mergers and diffuse events at higher redshifts, when the Universe was denser, are more prominent.

Fig. 4.7 also reveals that the subhalo accretors (and this includes the subhaloes with a zero accretion rate) accrete at a larger rate, on average, than the halo accretors at $z < 0.5$ for the mass scales shown. This, however, only causes a modest shift from the halo curve to the halo and subhalo curve at each redshift, because there are more halo accretors than subhalo accretors in the simulation, indicating that the enhanced subhalo accretion rates are not responsible for the AHOP to MSM shift in accretion at each redshift in Fig. 4.6. The enhanced accretion onto subhaloes can be understood by examining their mutual clustering and the relative velocity of their progenitors compared to their internal velocity, and both of these processes are now discussed.

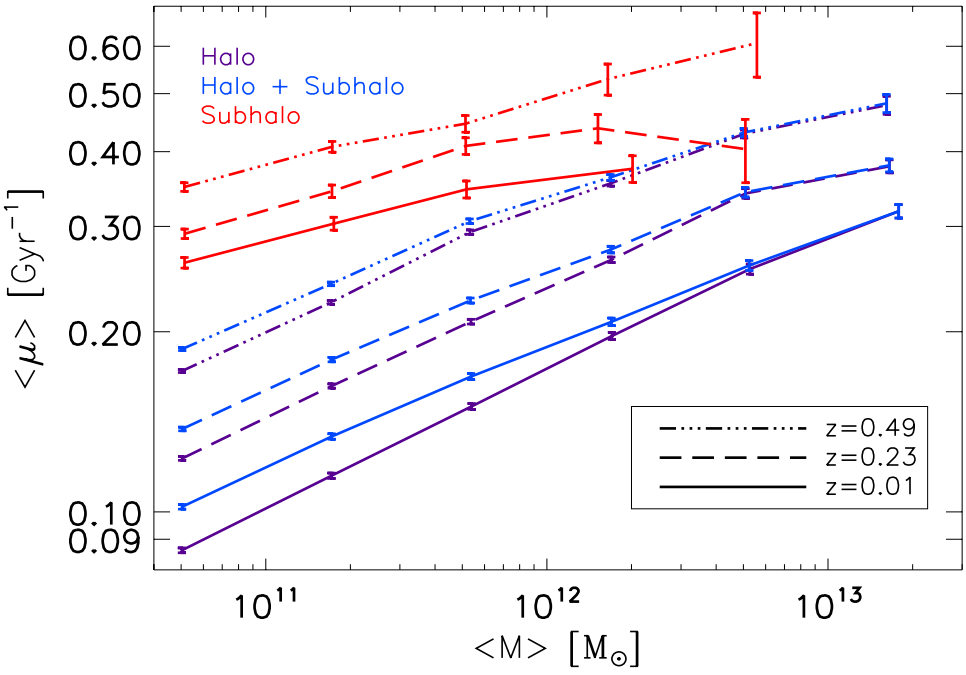

Figure 4.7: The mean specific accretion rate as a function of mass shown for $z = 0.49$ (triple-dot-dashed lines), $z = 0.23$ (dashed lines) and $z = 0.01$ (solid lines) using the halosub accretion method. For a given linestyle, the bottom purple line shows the MSM haloes, the middle blue line shows the MSM haloes and subhaloes and the top red line shows the MSM subhaloes.

The clustering of haloes and subhaloes

One of the main aims of this study is to investigate whether there is a relationship between the rate at which objects accrete mass and their environment and so in this section the clustering properties of haloes and subhaloes at different redshifts are examined. In the following section, accretors in different cluster-scale environments are specifically targeted.

Fig. 4.8 shows the two-point correlation function, ξ, for the MSM accretors from the simulation as a function of the physical separation distance r, at the same three redshifts shown in Figs 4.6 and 4.7 and at a much higher redshift of $\sim$2. The Landy & Szalay (1993) $\hat{w}_4$ estimator is used to compute ξ:

$$\hat{w}_4 = d - 2x + 1 \tag{4.1}$$

$$d = \frac{DD}{GN(N-1)/2} \tag{4.2}$$

$$x = \frac{DR}{GNN_r} \tag{4.3}$$

$$G = \frac{RR}{N_r(N_r-1)/2}. \tag{4.4}$$

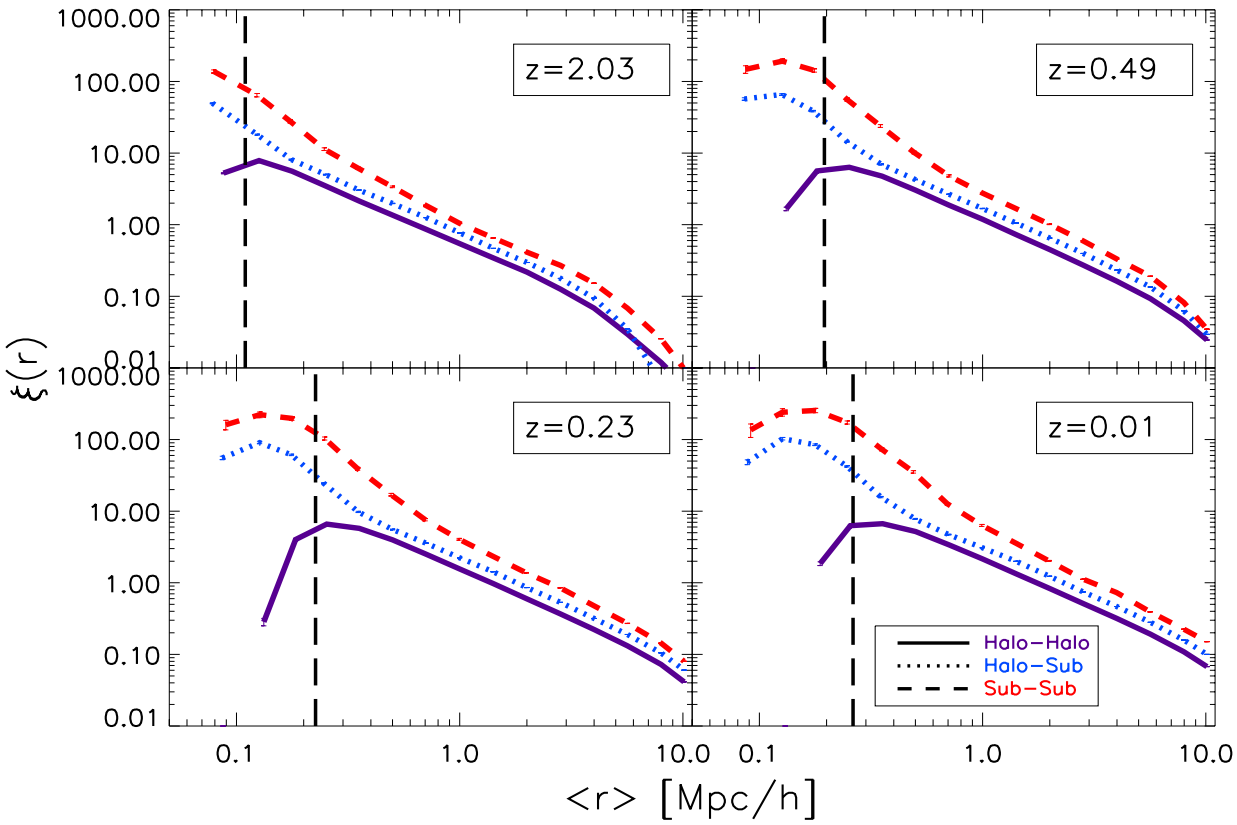

Figure 4.8: The two-point correlation function as a function of the mean inter-object separation (in physical coordinates) for four redshifts in the range $0 \lesssim z \lesssim 2$. The purple, blue and red lines in each panel show the bound MSM halo-halo, halo-subhalo and subhalo-subhalo accretor signals respectively, with the error bars corresponding to Poisson errors. The vertical long-dashed black lines represent an estimate of the resolution limit in r at each redshift.

Here DD, DR and RR represent the $N(N-1)/2, NN_r$ and $N_r(N_r-1)/2$ unique combinations of data-data, data-random and random-random pairs of points, where N and N_r refer to the total number of points in the real data set and the random data set respectively. Random catalogues are therefore required for each redshift. The catalogues used in this work sample 300 000 objects at each redshift and are hence larger than the corresponding total number of detected haloes and subhaloes (156 120, 211 537, 216 232 and 223 781 at $z = 2.03, 0.49, 0.23, 0$ respectively). For each panel in Fig. 4.8, the solid purple lines represent the halo-halo pairs, the dotted blue lines represent the halo-subhalo pairs, the dashed red lines represent the subhalo-subhalo pairs and the errors correspond to Poisson errors and are negligible. Only the clustering of bound accretors is measured: halo-subhalo pairs therefore correspond to the clustering of all bound halo accretors with all bound subhalo accretors. The vertical dashed black lines show the average total diameter of all the objects at the redshift in question, and represent an estimate of the resolution limit in r.

Fig. 4.8 demonstrates that subhalo-subhalo pairings are a factor of ∼2 more clustered than halo-halo pairings at large physical scales at low redshift. This factor increases to ∼10–15 at lower separation scales: subhaloes, by definition, reside within haloes and so cluster more strongly at small scales. The drop-off in clustering amplitude at the lowest scales should be ignored as this occurs at scales that are below the estimated resolution limit.

The subhalo-subhalo correlation function is the sum of two terms: the first describes the clustering of subhaloes within the same host and the second describes the clustering of subhaloes that belong to different hosts. For small separations, the subhalo-subhalo correlation function has a strong contribution from pairs of subhaloes in the same host. The clustering of halo-halo pairings is lower at these scales because these scales approach the size of haloes, and so it is less common to find two haloes close to each other without one or both member(s) of the pair being a subhalo. At larger scales, subhaloes belonging to different hosts contribute strongly to the subhalo-subhalo clustering strength.

The clustering amplitudes of the three curves also evolve with redshift, as shown by the evolution of the correlation length ($\xi(r_{\rm corr}) = 1$) of the subhalo-subhalo curve, which increases by a factor ∼3 towards $z = 0$. This is probably because at lower redshift there are more dense clusters and more subhaloes within a given host in the simulation, hence there is a stronger contribution to the subhalo-subhalo clustering amplitude than at higher redshift at the separation scales shown.

Measuring the relative velocities between the accretors' progenitors

Having established that subhaloes at sub-cluster scales are more clustered than haloes, especially at small scales, the distributions of $\Delta v/v_c$ are now examined, where Δv represents the relative velocity between an accretor's main father and one of its other progenitors, and v_c is the accretor's circular velocity. If $\Delta v/v_c$ tends to be smaller, on average, for subhalo accretors than halo accretors, then accretion onto haloes will tend to be more suppressed than accretion onto subhaloes. Fig. 4.9 shows the distributions of this ratio for haloes (thick lines) and subhaloes (thin lines) at the same redshifts shown in Fig. 4.7. The $\Delta v/v_c$ ratio is computed for each progenitor k (not equal to the main father j) of a given accretor: each particle accreted from the background is counted as an individual relative velocity event, as is each halo/subhalo progenitor. So if an accretor has a main father j, a father k, and also accretes two particles from the background, m and n, then three separate relative velocities with respect to j are computed for that accretor. The accretors are

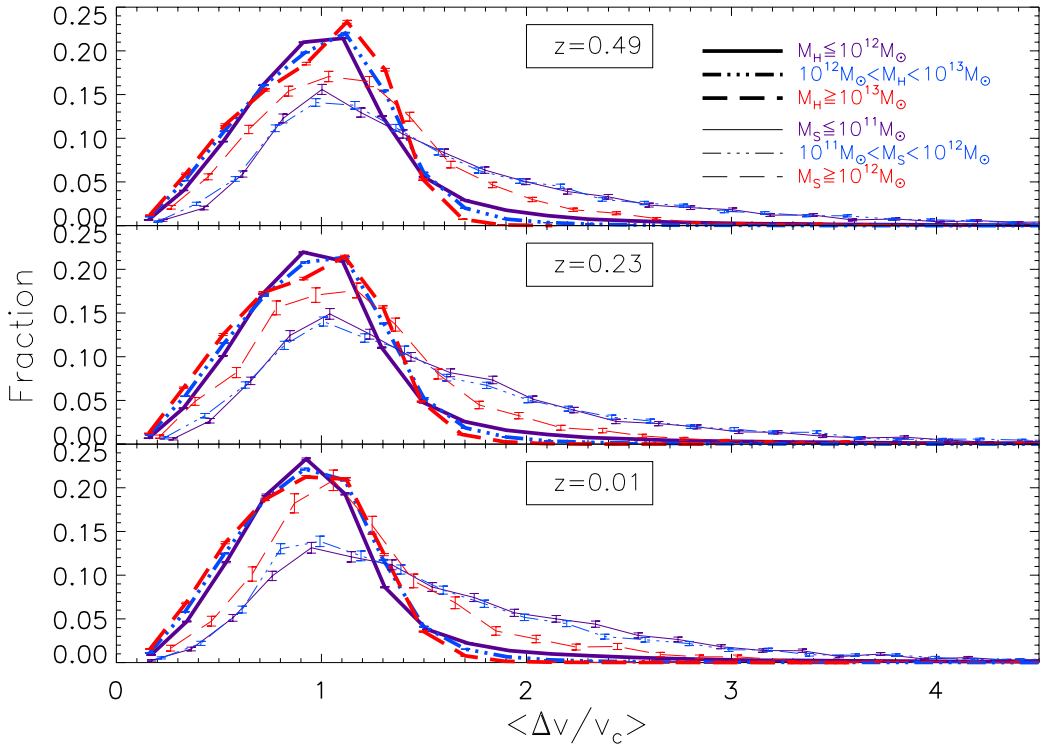

Figure 4.9: Distributions of $\Delta v/v_c$ for the halo (thick lines) and subhalo (thin lines) accretors, where Δv represents the relative velocity between an accretor's main father and one of its other progenitors, and v_c is the accretor's circular velocity. The halo and subhalo accretors from Fig. 4.7 have been divided into different mass bins, shown by the ranges of $M_{\rm H}$ and $M_{\rm S}$ respectively.

binned in mass and the different halo and subhalo mass bins are shown in Fig. 4.9 by the ranges of $M_{\rm H}$ and $M_{\rm S}$ respectively.

It can be seen from Fig. 4.9 that the distributions of $\Delta v/v_c$ for the halo and subhalo accretors are similar: they depend quite weakly on mass and their peaks coincide.

Revisiting the enhanced accretion onto subhaloes in Fig. 4.7

It is well established that in simulations, after infall, subhaloes experience mass loss via tidal stripping, tidal heating and disk shocking (Gnedin et al., 1999; Dekel et al., 2003; Taylor & Babul, 2004; D'Onghia et al., 2010), and have a large velocity dispersion that scales with their host's mass. Mass stripping events in this dark-matter-only study are recorded as zero accretion events, and so one would perhaps expect subhaloes to be accreting at low rates, on average. Fig. 4.10 therefore examines the importance of mass stripping by considering the spatial distribution of the subhalo accretors within their hosts at the same three redshifts considered in Fig. 4.7. The distance between the centre of every subhalo accretor and its host is recorded in units of its host virial radius. A

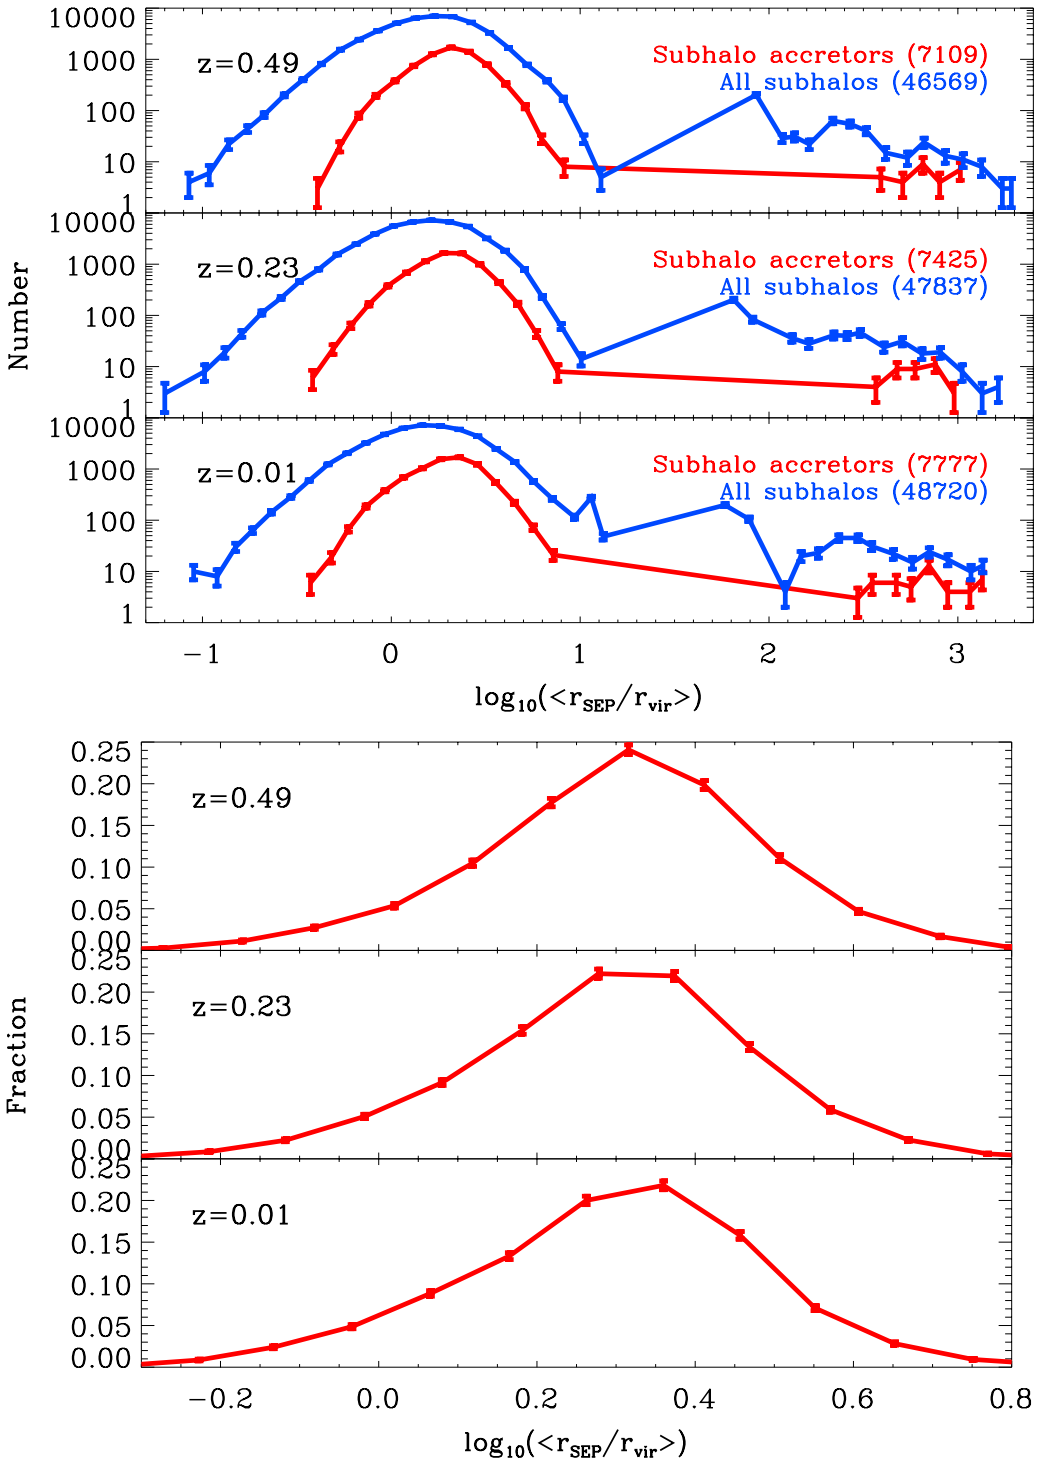

Figure 4.10: The upper panel shows the number of subhaloes as a function of distance from their host, measured in units of their host's virial radius. A distinction is made between the entire subhalo sample of both accretors and non-accretors ($M \geq 20M_p$, blue lines), and the subhalo accretor sample ($M \geq 40M_p$, red lines), with the size of each population shown in brackets. The number of subhalo accretors at each separation scale is then normalized by the subhalo accretor sample size in the lower panel. The subhaloes residing at the largest distances beyond their host's virial radius in the lower panel probably represent satellites that have just been accreted onto their host and that are undergoing the preliminary merger phase.

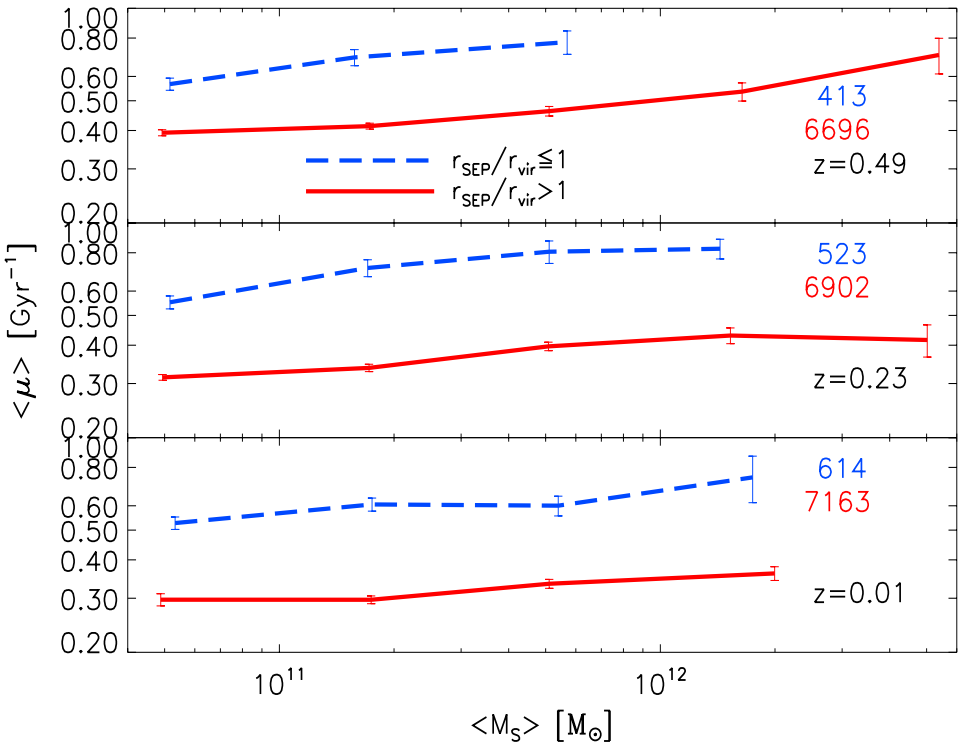

Figure 4.11: The specific accretion rate of the subhalo accretors from Fig. 4.10 as a function of mass for subhaloes within (blue) and beyond (red) their host's virial radius, with the corresponding numbers shown to the right of each panel.

given halo's virial radius (defined in Section 1.2.2) is similar to r_{500}, the spherical region within which the halo density is at least 500 times the critical density of the Universe. It can be shown that $r_{200} \sim 1.6 r_{500}$ for an isothermal density profile, and so for a good estimate of the distributions of $\log_{10}(\langle r_{\rm sep}/r_{200} \rangle)$ (the more common definition of the halo virial radius), one can simply shift the distributions of $\log_{10}(\langle r_{\rm sep}/r_{\rm vir} \rangle)$ in Fig. 4.10 in the negative x direction by $\log_{10}(1.6) \sim 0.2$ dex. The top (bottom) panel of Fig. 4.10 indicates the number of subhalo accretors (fraction of all subhalo accretors) in each radial bin. Most of the subhalo accretors in the simulation at $z < 0.5$ reside in the outer regions of their host with $\sim$70% located beyond r_{200}. Most of these subhaloes have therefore probably not been significantly stripped of their mass.

Fig. 4.11 further separates the subhalo accretors at each redshift into populations within (blue) and beyond (red) their host's virial radius, to test whether the accretion onto subhaloes depends on the proximity to their host. A dependence is found: subhaloes of a given mass within their host's virial radius accrete at a higher rate than those further away, on average. At first glance, this appears to be in conflict with recent findings which report that stripping is a strong function of distance to the host centre (Hester & Tasitsiomi, 2010). Diffuse accretion is almost certainly

responsible for resolving this apparent discrepancy: high rates of stripping within the virial region of the host at low redshift means that there is a large supply of diffuse material for subhaloes to accrete from, and the halosub accretion estimator measures this and does not measure the mass lost via stripping. The fraction of mass accreted via diffuse accretion also increases towards the present day (Figs 4.2, 4.3 and 4.4), which serves to enhance this effect.

That subhaloes of a given mass in the simulation have a larger rate of accretion, on average, than haloes of the same mass can now be understood. The majority of the mass-selected subhalo accretors reside beyond their host's virial radius (where they are unlikely to be experiencing significant mass stripping) and there is no significant difference between the halo and subhalo accretor distributions of $\Delta v/v_{\rm c}$. The very frequent interactions between subhaloes of the same host at small scales (Fig. 4.8) is therefore probably responsible for the enhanced subhalo accretion rates. Haloes are less clustered at small scales and so accrete at lower rates, on average.

4.1.4 Halo and subhalo environment

In this section the effect an object's environment at cluster scales has on the rate at which it accretes mass is investigated. There are two popular, independent measures of environment in the literature: the overdensity $\delta_{\rm R}(\vec{x})$ in a sphere of radius R (Lemson & Kauffmann, 1999; Wang et al., 2007) and halo bias (Sheth & Tormen, 2004; Gao & White, 2007). This work adopts two similar measures of an object's environment: the first defines an environment mass within a cluster-sized sphere and the second uses the two-point correlation function.

Environment mass

The environment of a halo and a subhalo is defined as the total mass, $M_{\rm E}$, contained within a sphere of radius R centred on the object of interest. $M_{\rm E}$ includes the mass of all those objects whose centres lie within the sphere as well as the mass of the object the sphere is centred on. Spheres of radii $R = 1.46\,h^{-1}\,$Mpc and $R = 3.65\,h^{-1}\,$Mpc are considered because: a) these scales represent both typical clusters and much larger clusters, and; b) various authors have found that the dependence of some halo properties on environment, such as halo formation redshift, are sensitive to the choice of sphere radius (Lemson & Kauffmann, 1999; Harker et al., 2006; Hahn et al., 2009). Both of these environment mass definitions are applied to each bound accretor at the redshift under consideration, with only bound accretors having a recorded $M_{\rm E}$ value. Unbound objects and

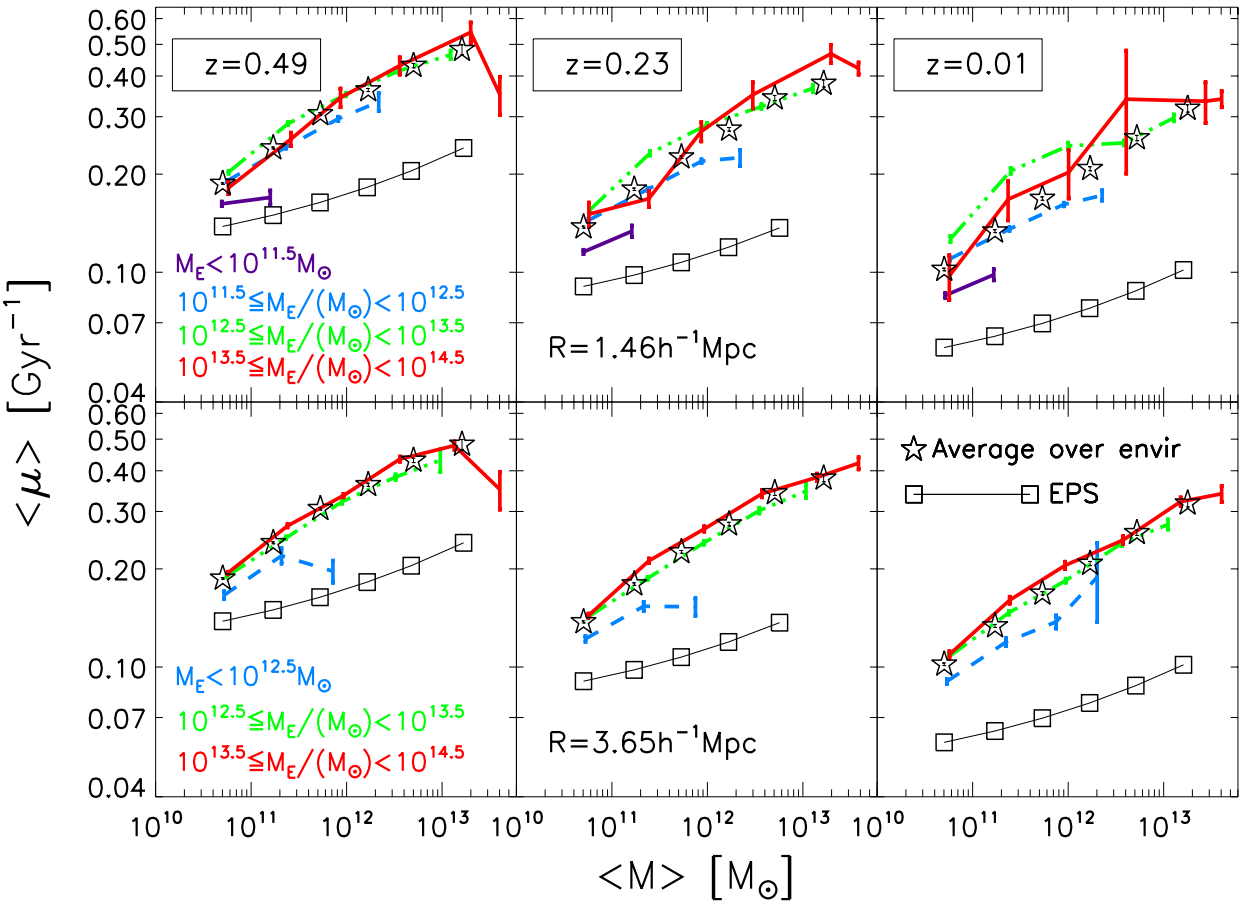

Figure 4.12: $\langle\mu\rangle$ plotted as a function of object mass (M) and environment mass ($M_{\rm E}$) for MSM-detected objects. The purple, blue, green and red lines in row one represent $M_{\rm E} < 10^{11.5}\,{\rm M}_\odot$, $10^{11.5}\,{\rm M}_\odot \leq M_{\rm E} < 10^{12.5}\,{\rm M}_\odot$, $10^{12.5}\,{\rm M}_\odot \leq M_{\rm E} < 10^{13.5}\,{\rm M}_\odot$ and $10^{13.5}\,{\rm M}_\odot \leq M_{\rm E} < 10^{14.5}\,{\rm M}_\odot$ respectively, using a sphere radius (R) of $1.46\,h^{-1}$ Mpc. The second row shows the results for a sphere radius of $3.65\,h^{-1}$ Mpc with environment mass bins $M_{\rm E} < 10^{12.5}\,{\rm M}_\odot$, $10^{12.5}\,{\rm M}_\odot \leq M_{\rm E} < 10^{13.5}\,{\rm M}_\odot$ and $10^{13.5}\,{\rm M}_\odot \leq M_{\rm E} < 10^{14.5}\,{\rm M}_\odot$, given by the blue, green and red lines respectively. The open squares joined by solid lines illustrate the EPS result using equation (3.28), and the open stars show accretion onto the MSM haloes and subhaloes independent of their environment. Columns one, two and three correspond to $z = 0.49$, $z = 0.23$ and $z = 0.01$.

resolved objects with $M \leq 40M_p$ are not, however, excluded from the sample as these objects could be part of a bound object's environment.

The first row of Fig. 4.12 plots the specific accretion rate onto haloes and subhaloes as a function of average object mass (M) and average environment mass ($M_{\rm E}$) for $z = 0.49$ (first column), $z = 0.23$ (second column) and $z = 0.01$ (third column) using a sphere radius of $1.46\,h^{-1}$ Mpc. For each panel in the first row, the purple, blue, green and red lines correspond to the environment mass bins $M_{\rm E} < 10^{11.5}\,{\rm M}_\odot$, $10^{11.5}\,{\rm M}_\odot \leq M_{\rm E} < 10^{12.5}\,{\rm M}_\odot$, $10^{12.5}\,{\rm M}_\odot \leq M_{\rm E} < 10^{13.5}\,{\rm M}_\odot$ and $10^{13.5}\,{\rm M}_\odot \leq M_{\rm E} < 10^{14.5}\,{\rm M}_\odot$ respectively. The second row of Fig. 4.12 shows the results us-

ing a larger sphere radius of $3.65\,h^{-1}\,\mathrm{Mpc}$ at the same three redshifts, divided into the following environment mass bins: $M_\mathrm{E} < 10^{12.5}\,\mathrm{M}_\odot$ (blue), $10^{12.5}\,\mathrm{M}_\odot \leq M_\mathrm{E} < 10^{13.5}\,\mathrm{M}_\odot$ (green) and $10^{13.5}\,\mathrm{M}_\odot \leq M_\mathrm{E} < 10^{14.5}\,\mathrm{M}_\odot$ (red). The stars in each panel represent the accretion signals onto the MSM haloes and subhaloes independent of their environment, and the squares joined by solid lines show the EPS results.

The relationships found in the previous sections are preserved in Fig. 4.12: the specific accretion rate increases with object mass for objects in most environments and decreases towards $z = 0$ (as is shown in Fig. 4.7), and EPS consistently underestimates the mass accreted onto all object masses (as is shown for haloes at $z < 1$ in Fig. 4.1). The most striking feature of Fig. 4.12, however, is that objects of a given mass residing in more massive environments do not accrete at a particularly enhanced rate compared with objects of the same mass in much lower mass environments[1]. This suggests that the specific accretion rate onto haloes and subhaloes does not depend strongly on environment at cluster scales. Objects in cluster-mass environments shown in the first row (red lines) mostly accrete less mass than in lower mass environments, but the number of objects in cluster-mass surroundings is limited by the choice of sphere radius. This effect is not seen for the larger scale environments shown in the second row, where merging between subhaloes on the outskirts of the host halo (shown by rightward tail of the $r_\mathrm{sep}/r_\mathrm{vir}$ distribution in the upper panel of Fig. 4.10) is probably driving accretion (but only at a slightly higher overall rate). The second row in Fig. 4.12 also shows that the specific accretion rate only depends weakly on environment, even for larger scales that probe the outermost regions of clusters. This weak environment dependence in rows one and two is discussed in more detail in Section 4.2.2.

Other authors have quantified environment by computing the overdensity δ_R in a sphere of radius R, rather than the mass (Lemson & Kauffmann, 1999; Harker et al., 2006; Hahn et al., 2009; Fakhouri & Ma, 2009, 2010). The advantage of a distance-weighted density measure is that the sensitivity to objects whose centre is just within the spherical boundary is reduced, as their contribution to the overall density has an associated low weighting factor. An alternative weighted

[1]Haloes dominate the accretion signals in Fig. 4.12, but the same trends were found at each of the chosen redshifts when just subhaloes were plotted as a function of their mass and environment mass (the subhalo curves have been omitted in Fig. 4.12 for clarity). There were two differences, however. The subhaloes were found to: a) accrete at higher rates, and; b) reside only in larger mass environments.

environment density measure for each halo and subhalo accretor has therefore also been computed:

$$\rho_i(\leq R) = \sum_j W_{\rm C}(|\vec{r}_i - \vec{r}_j|\,; R) M_j, \tag{4.5}$$

where the index j denotes a neighbouring object of mass M_j whose centre at $\vec{r}_j$ is within the environment sphere of radius R (which itself is centred on object i at $\vec{r}_i$). Here $W_{\rm C}$ represents the weighting function normalized by the volume of the spherical region, and is chosen to be the standard SPH cubic spline kernel given by equation (2.40). The mass contributions from objects close to the centre of the sphere are hence weighted more strongly than those further away. Fakhouri & Ma (2009) showed that for haloes more massive than $10^{14}\,{\rm M}_\odot$, the density of the object the sphere is centred on starts to dominate the contributions to $\delta_{\rm R}$, and so the central object's contribution is included and excluded in two separate weighted environment density measures, shown in Fig. 4.13.

When binned in environment density, the same weak environment dependence as in Fig. 4.12 is found in both cases. The error bars in Fig. 4.13 also show similar behaviour to Fig. 4.12. Increasing the sphere radius R: a) increases the number of objects (and hence accretors) included in the mass-weighted computation and therefore reduces the error bars on $\langle\mu\rangle$ for each of the $\rho_{\rm E}$ bins, and; b) reduces the environment densities, which has the effect of increasing the population in the lower $\rho_{\rm E}$ bins. By comparing the two environment definitions in Fig. 4.13 at a given R scale, it can be seen that the large mass haloes densely populate the high $\rho_{\rm E}$ bins if their mass is included in the sum in equation (4.5). If it is excluded, however, the large mass haloes have lower associated environment densities that populate the lower $\rho_{\rm E}$ bins, and the higher $\rho_{\rm E}$ bins become more noisy, as is seen.

Clustering in differing accretion schemes

The correlation function is now used as an alternative to the environment mass and weighted environment density measures of the previous section, and the analysis is not restricted to just cluster scales of a few Mpc. Samples of objects with very similar masses at different redshifts are considered in an attempt to examine whether objects of a given mass that accrete at larger rates have a larger clustering amplitude. This also tests the work by Percival et al. (2003), who found that at $z = 2$, haloes of a given mass accreting at differing rates do not cluster differently.

The $z = 2.03$ panel in Fig. 4.14 shows the correlation function for all those objects whose

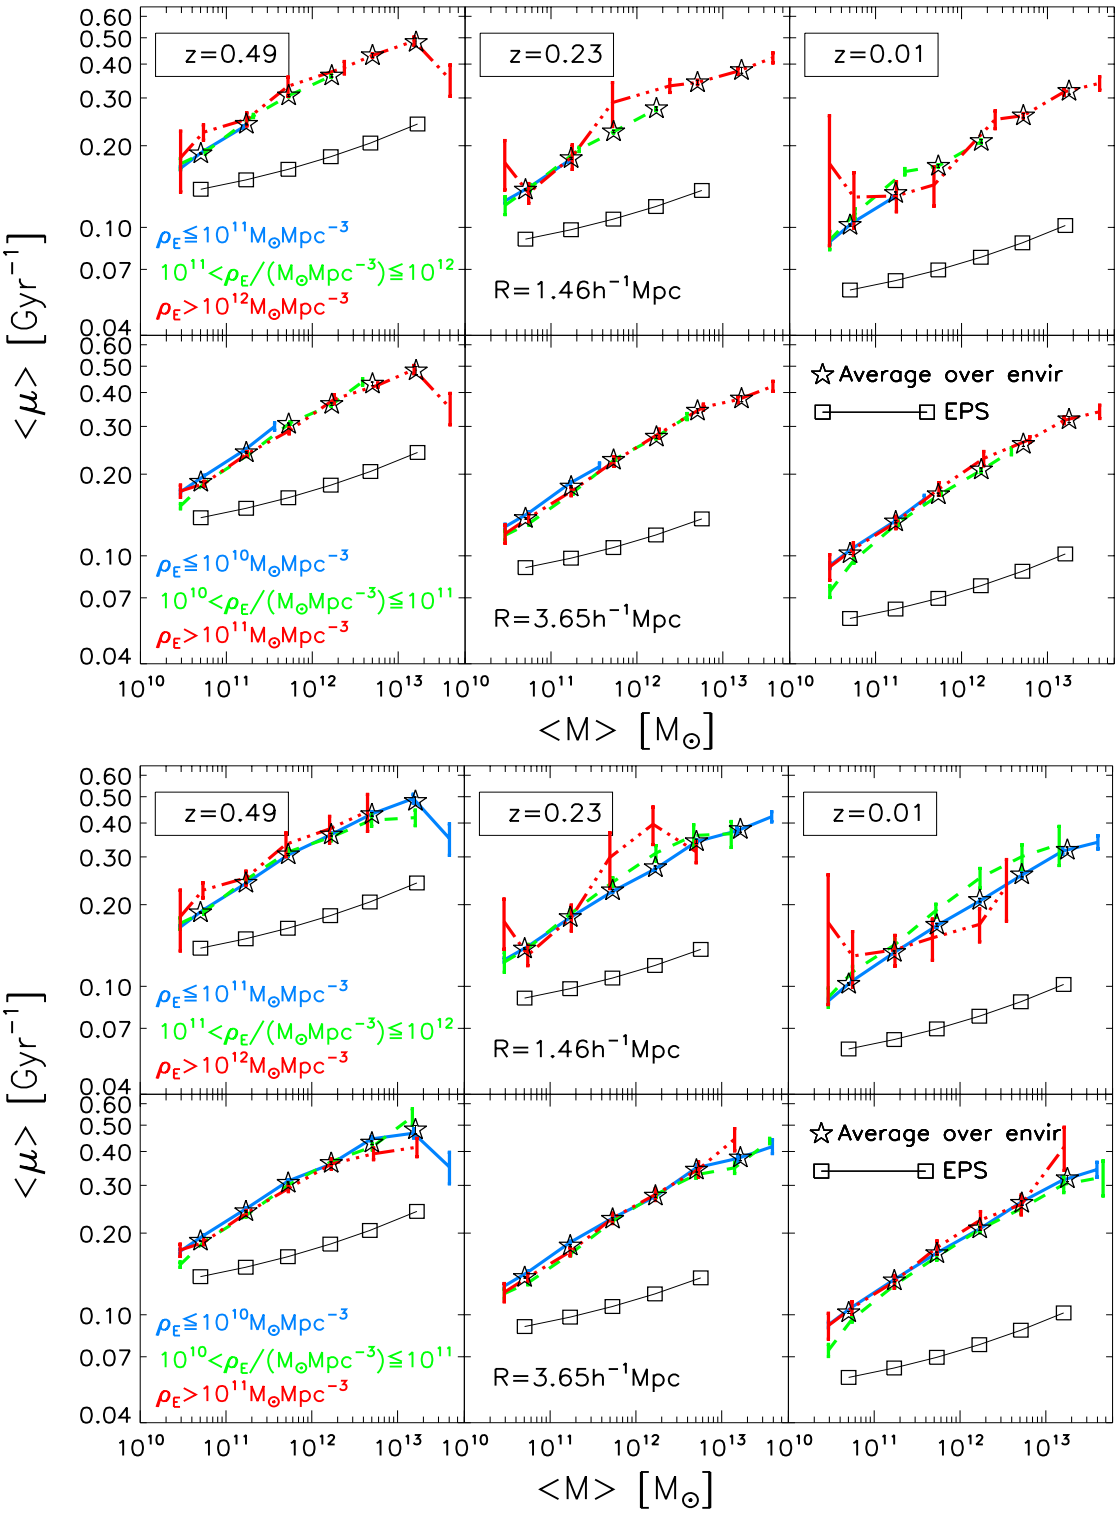

Figure 4.13: Same as Fig. 4.12 except that the environment measure $\rho_{\rm E}(\leq R)$ for each accretor is now a weighted sum of the mass of all the objects within a sphere of radius R centred on each accretor, normalized with respect to the sphere volume, computed using equation (4.5). The top and bottom panels respectively include and exclude the mass of each accretor in the weighted environment density measure. For a discussion of the size of the error bars, see the text.

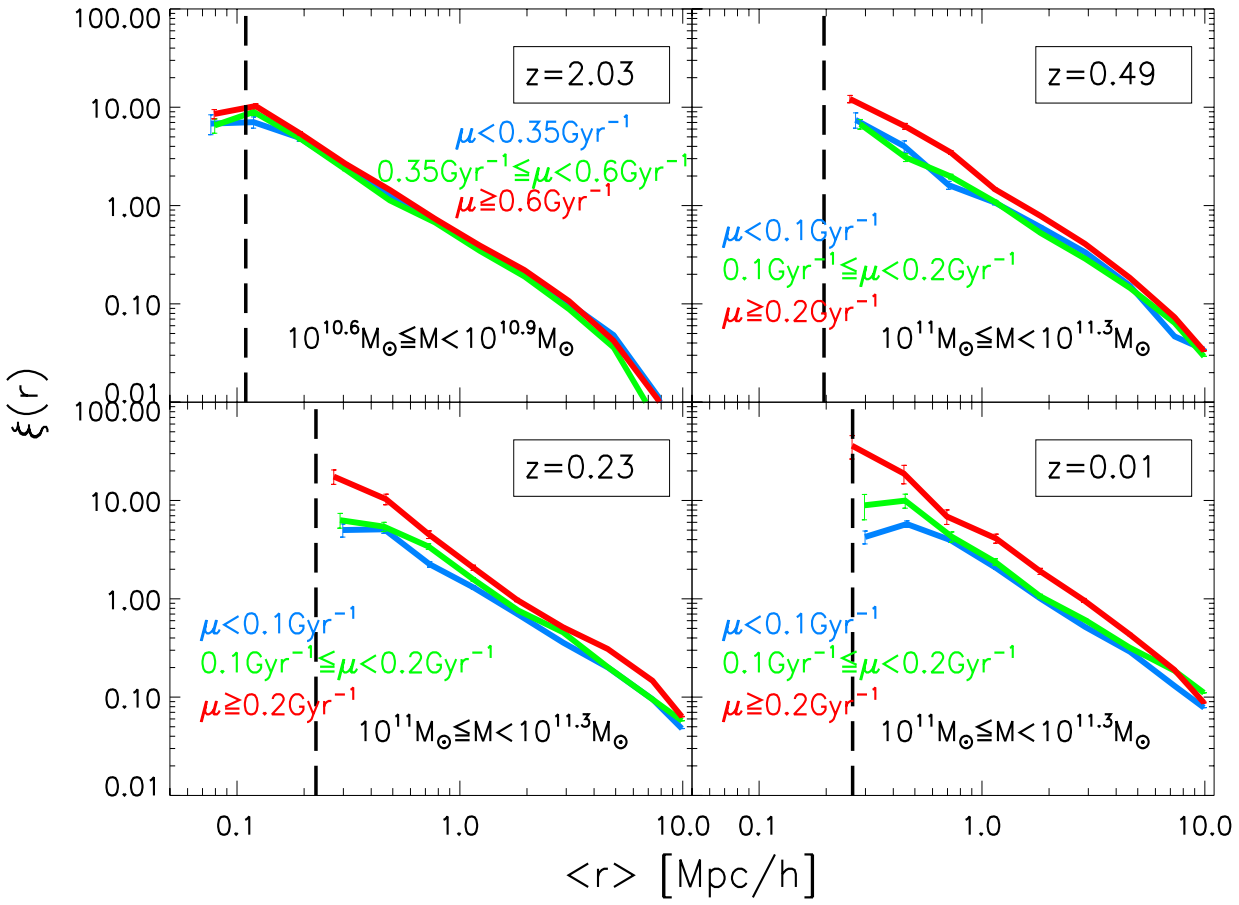

Figure 4.14: The two-point correlation function plotted as a function of the mean inter-object separation (in physical coordinates) for four redshifts in the range $0 \lesssim z \lesssim 2$. The $z = 2.03$ panel shows the MSM objects whose masses satisfy $10^{10.6}\,\mathrm{M}_\odot \leq M < 10^{10.9}\,\mathrm{M}_\odot$ with $\mu < 0.35\,\mathrm{Gyr}^{-1}$ (blue), $0.35\,\mathrm{Gyr}^{-1} \leq \mu < 0.6\,\mathrm{Gyr}^{-1}$ (green) and $\mu \geq 0.6\,\mathrm{Gyr}^{-1}$ (red). The lower redshift panels show the MSM objects whose masses satisfy $10^{11}\,\mathrm{M}_\odot \leq M < 10^{11.3}\,\mathrm{M}_\odot$ with $\mu < 0.1\,\mathrm{Gyr}^{-1}$ (blue), $0.1\,\mathrm{Gyr}^{-1} \leq \mu < 0.2\,\mathrm{Gyr}^{-1}$ (green) and $\mu \geq 0.2\,\mathrm{Gyr}^{-1}$ (red). The vertical long-dashed black lines represent an estimate of the resolution limit in the separation scale at each redshift.

mass satisfies $10^{10.6}\,\mathrm{M}_\odot \leq M < 10^{10.9}\,\mathrm{M}_\odot$ with $\mu < 0.35\,\mathrm{Gyr}^{-1}$ (blue), $0.35\,\mathrm{Gyr}^{-1} \leq \mu < 0.6\,\mathrm{Gyr}^{-1}$ (green) and $\mu \geq 0.6\,\mathrm{Gyr}^{-1}$ (red). The lower redshift panels show the correlation function for objects with $\mu < 0.1\,\mathrm{Gyr}^{-1}$ (blue), $0.1\,\mathrm{Gyr}^{-1} \leq \mu < 0.2\,\mathrm{Gyr}^{-1}$(green) and $\mu \geq 0.2\,\mathrm{Gyr}^{-1}$ (red) whose mass satisfies $10^{11}\,\mathrm{M}_\odot \leq M < 10^{11.3}\,\mathrm{M}_\odot$. The mass interval for $z < 0.5$ has been chosen because it lies below the break mass, $M_\star$, in the mass function at these redshifts and so μ is not biased. For comparison, the mass interval in the $z \sim 2$ panel lies closer to $M_\star$. The vertical dashed lines represent an estimate of the resolution limit in the separation scale (same as the vertical dashed lines in Fig. 4.8).

At well-resolved non-linear small scales for $z < 0.5$, objects with high specific accretion rates are up to a factor of ∼3 more clustered than the lower accreting objects, whereas at larger linear

scales the difference in clustering between different accretors is much smaller. For the cluster-scale environments of the first row of Fig. 4.12, corresponding to a maximum r value of $2.92\,h^{-1}\,\mathrm{Mpc}$, there is a weak environment dependence, with objects of larger μ being slightly more clustered. Fig. 4.14 therefore provides further evidence that the mass accreted onto haloes and subhaloes of a given mass weakly depends on their environment at cluster scales.

In contrast to the $z < 0.5$ behaviour, there is very little difference in clustering between different accretors with $10^{10.6}\,\mathrm{M}_\odot \leq M < 10^{10.9}\,\mathrm{M}_\odot$ at $z \sim 2$ and this holds for both the linear and non-linear scales shown. The results in Fig. 4.14 therefore agree with the conclusions of Percival et al. (2003) at $z \sim 2$ but show that they break down at $z < 0.5$, where there is a larger difference in clustering between high accretors and low accretors of a given mass at all scales.

4.2 Discussion

4.2.1 Disagreement with EPS theory

The analytic EPS calculation shows significant departures from the halo accretion rates found in the simulation using the halosub method at both low and high redshift (Fig. 4.1), and yields accretion curves that are a factor of ∼1.3 steeper than the corresponding simulation curves. This simulation study, however, is not the first to report disagreement with EPS theory at high redshift: Cohn & White (2008) examined the accretion onto haloes of mass $M_\mathrm{H} = 5\text{–}8 \times 10^8\,h^{-1}\,\mathrm{M}_\odot$ at $z = 10$ and found that EPS overestimated the halo accretion rate by a factor ∼1.5 (using a lookback time of 50 Myrs). Fig. 4.1 shows a similar behaviour, with EPS overpredicting the accretion rate onto haloes of mass $M_\mathrm{H} \sim 10^{10.7}\,\mathrm{M}_\odot$ by a factor of ∼2 at $z = 8$. One might expect EPS to overestimate accretion onto haloes at high redshift because it assumes that collapse is spherical (see Section 3.1.5) and that the density barrier is fixed in height (Lacey & Cole, 1993), whereas it has been shown that allowing for ellipsoidal collapse and treating the critical density contrast for collapse as a free parameter better reproduces the N-body halo mass function (Sheth et al., 2001; Sheth & Tormen, 2002). This modification reduces the critical density contrast for collapse by a factor of $\sqrt{0.7}$ (M06) which reduces $f(M_\mathrm{H})$ in equation (3.28) by the same factor, causing a slight shift in the dashed EPS curves in Fig. 4.1 but otherwise having no effect on the redshift or mass dependence.

The disagreement might arise because EPS theory is only approximate, as highlighted in Section 3.1.5. A simplified assumption of the EPS framework inherent in equation (3.28) is that haloes

accrete at rates that do not depend on their environment, which is particularly problematic in the non-linear regime at $z < 1$ when accretion onto structures embedded within clusters is of interest (Fig. 4.12). This restrictive assumption, however, is not a prediction of the theory and so various authors have recently relaxed it. Sandvik et al. (2007) implemented a multidimensional generalization of the EPS formalism and used an ellipsoidal collapse model where collapse depended both on the overdensity and the shape of the initial density field. They found only a weak dependence between halo formation redshift and halo clustering which was stronger for more massive haloes, in disagreement with the reported halo assembly bias in numerical simulations (Gao et al., 2005; Gao & White, 2007; Maulbetsch et al., 2007). Zentner (2007) modified the EPS formalism by using a Gaussian smoothing window function, and Desjacques (2008) allowed the density threshold to have an environment dependence, but both authors found that dense large-scale environments preferentially contain haloes that form later. An EPS model that is able to account for halo assembly bias and predict a modified analytic version of equation (3.28) for the halo accretion rate is hence lacking. A modification of this kind might produce better agreement with the simulation results for $z < 1$ in Fig. 4.12.

Yet perhaps the biggest weakness with EPS was highlighted by Benson et al. (2005), which could ultimately account for the offset in Figs 4.1 and 4.12. They showed that the Lacey & Cole (1993) EPS formula yields merger rates that are not symmetric under exchange of halo masses, and that do not predict the correct time evolution of the Press–Schechter mass distribution, indicating that constructed EPS merger trees are fundamentally flawed. Recent work by Neistein & Dekel (2008), however, has demonstrated that these discrepancies can be reconciled if the EPS formalism relaxes the binary merger condition (assumed in the Benson et al. 2005 paper), and allows haloes to have multiple progenitors. However, this approach requires the mergers between multiple progenitors to take place over a very small timescale compared with the Hubble time, so that the limit $\Delta t \to 0$ is justified, and the validity of the assumption that mergers between more than two progenitors occur on small timescales is questionable.

An alternative viewpoint is that the level of disagreement between the measured accretion onto simulated haloes and the corresponding EPS predictions is dependent on the chosen accretion model, as Fig. 4.1 has revealed that the halo method yields better agreement with the EPS accretion evolution than the halosub method. It should be noted, however, that the halo method prescription suffers from a time-biasing that probably increases towards the present day in more

massive environments, as the routine has to trace back further in time to achieve a positive mass difference between a given object and its most massive progenitor. The halo method also ignores the accreted component when mass loss dominates, and so may be dividing by an artificially large time interval. Yet it is unlikely that these limitations fully account for the discrepancies between the halo and halosub methods in Fig. 4.1, and so the general conclusion that the level of agreement with the EPS trajectories depends on the chosen prescription is expected to hold.

4.2.2 Understanding the weak relationship between local accretion rate and environment at cluster scales

By quantifying accretion onto substructures embedded in groups and clusters, this study has moved beyond the limited predictive power of the EPS formalism. Fig. 4.7 demonstrates that subhaloes in the simulation accrete at larger rates than haloes of the same mass, on average (by a factor of $\sim$3 for the lowest mass subhaloes at $z = 0$). At first glance this appears to contradict recent claims: Angulo et al. (2009) and Hester & Tasitsiomi (2010) have shown that subhalo-subhalo mergers are rare and that subhaloes are severely stripped of mass, which probably means that their accretion rates are likely to be low. The low redshift subhalo accretors in this study, however, form a subsample of subhaloes that are chosen to be safely above the mass resolution limit and that are mostly located at large distances from their host's centre, with $\sim$70% residing beyond r_{200} of their host (Fig. 4.10). These subhaloes are probably not therefore being significantly stripped of mass, unlike those in recent studies. The mass-selected nature of the subhalo accretors and the different spatial distribution within their host are hence the most likely causes of the apparent accretion rate discrepancy with these aforementioned studies. It should be noted, however, that the mass and spatial distributions of resolved subhaloes within their hosts are affected by the bias of the chosen halo finder, and so may not correspond to the 'true' distributions. Knebe et al. (2011) performed a comprehensive comparison between 18 different finders for the predicted velocity, mass, particle and spatial distributions of known mock haloes and subhaloes. They found that configuration-based finders (based on spatial, not velocity components) like the MSM algorithm employed in this work and the SUBFIND routine used in the Angulo et al. (2009) and Hester & Tasitsiomi (2010) studies, were not able to resolve M_{200} and $V_{\rm max}$ of the velocity profile of central subhaloes within 20–30% of R_{100} of their host, requiring the subhaloes to contain at least 30–40 particles for reliable detections (hence justifying the 40 particle threshold enforced throughout this study). The Knebe

et al. (2011) paper therefore further highlights that both the MSM and SUBFIND finders are biased towards large mass subhaloes that are not in close proximity to their host's centre, and so the requirement that subhalo accretors in this study must contain 40 particles probably accounts for the paucity of central subhalo accretors in Fig. 4.10. The absence of this mass threshold condition in the Angulo et al. (2009) and Hester & Tasitsiomi (2010) studies liberates the central subhalo population, but the predictions for M_{200} and $V_{\rm max}$ of these objects are not reliable.

Fig. 4.8 demonstrates that the subhalo accretors in this work are more clustered than the halo accretors at small scales and Fig. 4.9 shows that there is no significant difference between the halo and subhalo distributions of $\Delta v/v_c$, where Δv is the relative velocity between an accretor's main father and one of its other progenitors, and v_c is the accretor's circular velocity. The high subhalo accretion rates are therefore likely to be driven by the very frequent interactions at small scales with other subhaloes of the same host.

One might expect the accretion rate onto haloes and subhaloes to depend strongly on environment at larger cluster scales given the increased rate of interactions in dense environments, but only a weak dependence is found (Figs 4.12, 4.13 and 4.14). The subhalo accretors reside in only the most massive environments and probably accrete mostly locally from their nearby subhalo neighbours rather than their host, and so this is a possible explanation for their weak relationship between accretion rate and environment. One likely explanation for haloes is that the increased interaction rates of haloes in group- and cluster- mass environments are not large enough to significantly overcome the large halo relative velocities, resulting in only a modest net increase in accretion at cluster scales.

Fakhouri & Ma (2010) examined the environment dependence of accretion onto high mass haloes ($M_{\rm H} > 10^{12}\,{\rm M}_{\odot}$) from the Millennium Simulation and found a weak, negative correlation for galaxy-mass haloes. By contrast, Fig. 4.12 reveals a weak, positive dependence for all objects. The analysis in this study extends theirs by accounting for accretion onto substructures, and adopts a different expression for the mass accretion rate, which they define as $(M_i - M_j)/(t_i - t_j)$, where j is the most massive father of i (note that this definition is similar to the measured accretion rate presented in Section 3.2.5). Their adopted expression hence includes mass loss, unlike the halosub method (or the halo method). Another possible source of the discrepancy is the method used to identify anomalies. Genel et al. (2009) highlighted some fundamental problems with the 'stitching' algorithm used by Fakhouri & Ma (2010) to remove anomalous events. They demonstrated that it

can lead to a double counting of mergers and to a false counting of anomalous events as mergers, and they showed that this overestimation of the merger rate was particularly problematic for minor-mergers. Predicting the effects that overestimating the merger rate have on the accretion rate and how they vary as a function of environment is not trivial, but differences between the anomalous event detection methods could explain the difference in the sign of the trend between accretion rate and environment.

The $z = 2.03$ panel in Fig. 4.14 shows that at higher redshift, when haloes far outnumber subhaloes in the simulation, the rate of accretion onto haloes is independent of environment, confirming the Percival et al. (2003) result. The Percival et al. (2003) study examined the difference in clustering at $z = 2$ between haloes of a given mass accreting at different rates. They considered several mass intervals ranging from $10^{10.3}\,\mathrm{M}_{\odot} \leq M_{\mathrm{H}} \leq 10^{10.4}\,\mathrm{M}_{\odot}$ to $10^{13.3}\,\mathrm{M}_{\odot} \leq M_{\mathrm{H}} \leq 10^{13.6}\,\mathrm{M}_{\odot}$ and concluded that for each mass interval halo accretion rates do not depend on environment at this redshift. This apparent lack of environment dependence probably arises because the haloes in the Percival et al. (2003) study, and to a lesser extent the haloes considered in the $z \sim 2$ panel in Fig. 4.14, represent some of the most massive objects at $z \sim 2$ and hence have bias factors $b >> 1$ (Sheth & Tormen, 1999). These structures are located at the highest peaks in the density field and so by computing the clustering amplitude of these objects, one is essentially measuring the clustering pattern of the highest density peaks at this redshift. It is therefore unlikely that the highest mass haloes experiencing different instantaneous accretion rates differ in their clustering. By contrast, the lower mass haloes and subhaloes in the $z < 1$ panels are less biased and so more closely track the clustering of the underlying mass distribution.

4.2.3 Comparing dark halo growth with black hole growth

Under the assumption that, on average, black hole growth traces dark halo growth (so-called 'pure coeval evolution'), M06 tested the predictions of equation (3.28) for the evolution of the integrated AGN luminosity density for $z \leq 3$. The coeval evolution model tests the hypothesis that the fractional mass accretion rate onto black holes and onto haloes are equal (i.e. $\dot{M}/M$ is the same for both black holes and haloes), and is consistent with the tight relation inferred between black hole mass and galaxy bulge mass (Tremaine et al., 2002, but see Batcheldor 2010 for an alternative interpretation), and is easy to test. M06 found the predicted integrated AGN luminosity density to be in remarkable agreement with the bolometric AGN luminosity density measured using hard

X-ray data. They also found that for $z > 0.5$ average black hole growth was well approximated by pure coeval evolution, but for $z < 0.5$ the black hole luminosity density tailed off more quickly than dark halo growth, and by $z = 0$ was lower by a factor of $\sim$2. They suggested that this slowdown in black hole accretion could be related to cosmic downsizing (e.g. Barger et al., 2005).

Their predictions for dark halo growth were, however, based on EPS theory. The simulation trajectories in Fig. 4.1 show that the EPS formalism underestimates halo accretion for $z < 1$, and at $z = 0$ is a factor of $\sim$1.5–2 lower for all halo masses. This implies that present day dark haloes could be accreting at fractional rates that are up to $\sim$3–4 times higher than their associated black holes. However, for $1 < z < 3$, the simulated dark halo accretion trajectories in Fig. 4.1 are reasonably well approximated by EPS. It follows that a possible accretion scenario is that for $1 < z < 3$, black holes grow coevally with their dark hosts but for $z < 1$, the epoch of cluster formation, their growth significantly decouples from that of their hosts.

It is still plausible that this decoupling is linked to the inference that high mass black holes preferentially turn off at low redshifts, leaving the remaining accretion activity dominated by low mass black holes (Heckman et al., 2004). The cause of such downsizing is often assumed to be connected to the physics of the baryon component. This study reinforces this assumption: if downsizing were a whole-halo phenomenon, it would be manifest in the dark-matter-only simulation, and its absence in the results confirms that one should seek an explanation in the baryons.

4.2.4 Are low halo accretion events the cause of radio-mode feedback?

A number of authors have developed semi-analytic models of galaxy formation that are tuned to reproduce the galaxy luminosity function at low redshift (e.g. Bower et al. 2006; Croton et al. 2006; De Lucia et al. 2006). A key ingredient of these models is a low level of feedback from black hole accretion that arises in all galaxies and which increases in importance towards low redshifts. The feedback mechanism has still not been identified: luminous, high accretion-rate AGN only form a small subset of the galaxy population at low redshift and seem unlikely to provide the required feedback in all galaxies. Bower et al. (2006) require black holes to have relatively high accretion Eddington ratios, which may be inconsistent with observations: it seems that the accretion and an associated outflow need to be hidden from view in a so-called 'radio-mode'. Croton et al. (2006) have assumed that such a mode could be fuelled by Bondi accretion from the hot gas phase of their model, but the observational evidence for such a mechanism has not been demonstrated either.

The survey of Ho et al. (1997) revealed that a high fraction—over 40% of nearby galaxies rising to 50–75% of bulge systems—host low luminosity AGN (LLAGN), with the majority of LLAGN accreting at highly sub-Eddington rates in the range $10^{-5} < L_{\rm bol}/L_{\rm Edd} < 10^{-3}$. Ho (2005) argued that these are systems where accretion occurs via a radiatively-inefficient advection-dominated accretion flow (ADAF). The accretion flow puffs up the inner disk and material is advected towards the black hole (Narayan, 2002; Ho, 2002, 2008), with outflow being channelled along kinetic-energy-dominated jets (Collin et al., 2003; Ho, 2005, 2008). It is possible, then, that LLAGN, fuelled by low accretion rate ADAFs, may provide the radio-mode feedback.

In this dark-matter-only study, the integrated minor-merger and diffuse halo accretion rate density curves in Figs 4.2 and 4.3 increase in importance towards the present day for all halo masses. This qualitatively agrees with the cosmological evolution of the black hole radio-mode integrated accretion signal found for each of the different semi-analytic models (Bower et al. 2006; Croton et al. 2006). If the periods when galaxy halo growth is dominated by low accretion rate minor-mergers and diffuse accretion events, are mirrored by low accretion rates onto their associated black holes, these in turn could produce the LLAGN that may be the radio-mode required for the feedback models.

The predicted integrated accretion rate density onto black holes residing in galaxy-mass haloes that are accreting diffusely and via minor-mergers at $z = 0$ is also very similar to the integrated accretion rate density onto black holes residing in similar sized haloes found by Croton et al. (2006), who argue that radio-mode feedback is more effective in more massive systems. The following estimate for the total black hole accretion rate density tests the hypothesis that for black holes with mass $M_{\rm BH}$ residing in haloes with mass $M_{\rm H}$:

$$\sum_i \dot{M}_{i,{\rm BH}}(z=0) \sim \alpha \frac{M_{\rm BH}}{M_{\rm H}} \sum_i \dot{M}_{i,{\rm H}}(z=0), \tag{4.6}$$

where α is the power-law index in the black hole–dark halo mass relation, and the index i sums over all galaxy-mass dark haloes and all black holes residing in these haloes at $z = 0$. Equation (4.6) hence assumes that black hole growth positively traces dark halo growth, on average (recent claims by Kormendy & Bender 2011, however, argue that for bulgeless galaxies there is no such correlation between black holes and their dark hosts, but the interpretation of this as meaning that there is no such relation for *all* galaxies has been clearly refuted by Volonteri et al. 2011. In what

follows the reliability of the assumption in equation (4.6) is not considered, but rather its prediction for black hole growth is tested). Ferrarese (2002) found that $\alpha = 1.65$ and that galaxy-mass haloes with $M_{\rm H} \sim 10^{12}\,{\rm M}_{\odot}$ have a black hole–dark halo mass ratio of $\sim 10^{-5}$. According to Fig. 4.2 these haloes with $\delta M_{\rm H}/M_{\rm H} \leq 0.02$ have a total accretion rate density of $\sim 7.6 \times 10^{7}\,{\rm M}_{\odot}\,{\rm Gyr}^{-1}\,{\rm Mpc}^{-3}$ at $z = 0$, which when substituted into equation (4.6) yields a total black hole accretion rate density of $\sim 10^{-5.9}\,{\rm M}_{\odot}\,{\rm yr}^{-1}\,{\rm Mpc}^{-3}$. This is very similar to the integrated accretion rate density of $\sim 10^{-5.8}\,{\rm M}_{\odot}\,{\rm yr}^{-1}\,{\rm Mpc}^{-3}$ onto supermassive black holes at $z = 0$ reported by Croton et al. (2006).

The δ parameter ($\equiv \delta M_{\rm H}/M_{\rm H}$) is a free parameter in the halosub model, but it is found that adopting the more classical progenitor mass ratio, χ, to distinguish between merger-type (Fig. 4.4) yields almost identical results to Fig. 4.2. This provides confirmation that the chosen δ cuts are indeed capable of separating minor- and major- merger channels. The δ parameter is therefore probably no more unconstrained than χ.

To summarize, the low rates of accretion onto dark haloes, driven by minor-mergers and diffuse accretion, may provide an alternative explanation to that proposed by Croton et al. (2006) for the radio-mode feedback needed to reproduce the observed galaxy luminosity function. The low redshift feedback phenomenon and its cosmological evolution may be driven by the cosmological evolution of halo minor-mergers and diffuse accretion rather than requiring accretion out of a hot gas phase.

4.3 Summary

Outputs from one of the high resolution dark-matter-only HORIZON simulations have been used to investigate the environment and redshift dependence of accretion onto both haloes and subhaloes. A method that computes the combined merger- and diffuse- driven accretion onto haloes and all levels of substructure has been developed and it is found that:

- Halo accretion rates vary less strongly with redshift than predicted by the EPS formalism when a prescription for tracing the exchange of mass across time is adopted. The level of disagreement with EPS is somewhat dependent on the chosen accretion prescription, however.

- Comparison with an observational study of black hole growth suggests that dark haloes at $z = 0$ could be accreting at fractional rates that are up to 3–4 times higher than their black holes.

- Halo growth is driven by minor-mergers and diffuse accretion at low redshift. These latter accretion modes have both the correct cosmological evolution and inferred integrated black hole accretion rate density at $z = 0$ to drive radio-mode feedback, which has been hypothesized in recent semi-analytic galaxy formation models as the feedback required to reproduce the galaxy luminosity function at low redshift. Radio-mode feedback may therefore be driven by dark halo minor-mergers and diffuse accretion, rather than accretion of hot gas onto black holes, as has been recently argued.

- The low redshift subhalo accretors in the simulation form a mass-selected subsample safely above the mass resolution limit and mostly reside in the outer regions of their host, with $\sim$70% beyond their host's virial radius, and are probably not therefore being significantly stripped of mass. These subhaloes accrete at higher rates than haloes, on average, at low redshifts. This is shown to be due to their enhanced mutual clustering at small scales: there is no significant difference between the halo and subhalo accretor distributions of $\Delta v/v_c$, where Δv represents the relative velocity between an accretor's main father and one of its other progenitors, and v_c is the accretor's circular velocity. The very frequent interactions with other subhaloes of the same host drive the high subhalo accretion rates.

- Accretion rates onto haloes and subhaloes depend only weakly on environment at cluster scales. It appears that the increased interaction rates in group- and cluster- mass environments are not large enough to significantly overcome the large halo relative velocities, resulting in only a modest net increase in halo accretion at cluster scales. The subhalo accretors only reside in the densest environments and they are likely to be accreting mostly from their nearby subhalo neighbours, rather than from their host. It is further demonstrated that haloes accrete independently of their environment at $z \sim 2$, as has been found by other authors, but this behaviour results from examining the clustering of the most massive haloes with large bias factors. When less massive haloes below $M_\star$ at low redshift are considered, a weak dependence between accretion rate and environment at cluster scales arises.

Having examined accretion onto host haloes and their embedded substructures, attention is now turned towards the second project in this thesis, which investigates the transport of angular momentum in the vicinity of a Milky Way type gas disk.

Chapter 5

Angular momentum transfer to a Milky Way disk: computational techniques

Towards the end of Chapter 1 it was argued that several simulation studies have reported that gas disks residing at the centre of most galaxies at high redshift grow via a cold mode of gas accretion (Birnboim & Dekel, 2003; Kereš et al., 2005; Ocvirk et al., 2008; Brooks et al., 2009; Powell et al., 2011). This 'methods' chapter aims to resolve the various galaxy components that channel into the central region of a resimulated Milky Way galaxy, and to provide a framework for measuring the amount of angular momentum they transport to its disk.

5.1 Scientific motivation

Understanding how disks acquire their angular momentum is undoubtedly a complex, multiscale process. On the largest scales of several Gpc, the Universe is thought to be arranged in a complex web of cosmic structure (Bond et al., 1996; Pogosyan et al., 1998), and simulations with differing spatial resolution have helped to decompose the intricate patterns of this web (e.g. Sousbie 2011). It appears that large sheets surrounding Gpc-scale voids intersect to form filaments that funnel gas from the voids to the intersection nodes of the web, where haloes of dark matter form. The amount of mass shared between these separate phases has recently been the subject of debate and involves examining the tidal field tensor that describes the second order derivatives of the

gravitational field, in both the linear (Doroshkevich, 1970) and non-linear (Shen et al., 2006; Hahn et al., 2007; Aragón-Calvo et al., 2010) regimes of perturbation growth. Yet the general conclusions of these studies are in agreement: filaments dominate the mass budget. It is hence likely that this component also carries large amounts of angular momentum to the halo nodes of the cosmic web.

Several recent studies have shed some light on the subject of angular momentum transport by filaments, on both supergalactic and galactic scales. Danovich et al. (2012) performed a statistical study of 350 Milky Way dark matter haloes selected from the HORIZON-MARENOSTRUM simulation (details given in Ocvirk et al. 2008, Dekel et al. 2009 and Devriendt et al. 2010) at $z = 2.5$ with mass $M_{\rm H} \simeq 10^{12}\,{\rm M}_{\odot}$, whose luminous components were simulated with a physical resolution of 1 kpc. They found that the streams of cold gas flowing towards the disk were oriented in a narrow plane, and that the angular momentum transported by this infalling component was highly misaligned with respect to the angular momentum direction of the disk, until the approximate disk boundary, whereupon it dramatically swung into close alignment. Pichon et al. (2011) analyzed outputs from a similar run of the HORIZON-MARENOSTRUM suite and examined the nature of filament trajectories on kpc scales, in an attempt to understand the existence of thin gas disks at high redshift that are thought to have formed from the inside-out. They demonstrated that material accreted at the virial sphere carried more angular momentum at later times, and attributed this phenomenon to a 'lever' mechanism: recently accreted gas has travelled from a further distance and moves with greater velocity owing to the velocity sway of filaments on large scales. It was further speculated that these large-scale drift velocities arise from the asymmetric cancellation of motions of gas pumped out of voids. These authors hence concluded that the angular momentum transported along cold gas flows into halo virial regions originates from large-scale voids. Kimm et al. (2011) investigated these claims by probing to a lower physical resolution of $\sim$50 pc in the NUT simulation (Powell et al., 2011). They found that at the time of crossing the virial sphere of their resimulated NUT halo, the specific angular momentum of cooling gas was systematically larger than that of dark matter. They argued that this was a manifestation of a dark matter angular momentum cancellation effect caused by the mixing of dark matter particles carrying different amounts of angular momentum at the time of accretion, and justified this claim by demonstrating that there was a large distribution in the age of the dark matter particles at a given radius within the NUT halo virial region. The gas accreting onto low mass galaxies like the NUT halo, on the other hand, does not experience the same angular momentum cancellation effect, as it is

now generally accepted that gas within these systems cools too quickly to support a shock at the virial boundary, and hence subsequently acquires a preferential flow direction by streaming inward towards the central region in a cold phase with temperature $T \lesssim 10^4\,\mathrm{K}$ (Birnboim & Dekel, 2003; Kereš et al., 2005; Ocvirk et al., 2008). Kimm et al. (2011) then further argued that the observation within the simulation that freshly accreted gas carries more angular momentum as a function of time lends support to the lever effect reported by Pichon et al. (2011). Perhaps the most exciting result from the Kimm et al. (2011) study, however, was the discovery that the amount of specific angular momentum transported by gas was constant with radius at both low ($0 \leq z \leq 3$) and high redshifts ($z > 3$), except for $r \lesssim 0.1 r_{\mathrm{vir}}$, whereupon it fell dramatically in both of these redshift regimes. This hints at unresolved complex dynamics within the 'disk' regions (e.g. Bett et al. 2010; Book et al. 2011) of high redshift low mass galaxies, and perhaps arises due to the mixing of gas.

Yet due to physical resolution constraints, the Kimm et al. (2011), Pichon et al. (2011) and Danovich et al. (2012) studies were neither able to accurately resolve the disk scaleheight and scalewidth at high redshift, nor the dynamics within the central region. By analyzing outputs from one of the high resolution NUT simulations, which boasts a resolution scale of 12 pc in physical coordinates, this thesis rises to the ultimate challenge of attempting to account for the amount of angular momentum locked-up in a Milky Way type disk that resides in a halo fed by streams of cold gas, and complements the papers listed above.

5.2 Does a single Milky Way system suffice?

Numerous groups have invested a lot of effort into simulating the evolution of a single (Agertz et al., 2009; Powell et al., 2011) or a small number (Governato et al., 2004; Agertz et al., 2011; House et al., 2011; Scannapieco et al., 2011, 2012) of Milky Way systems, and the NUT simulation analyzed in this thesis falls into the former category. One hence naturally wonders:

- why there is an emphasis on these systems being representative of the Milky Way, and;
- whether it is possible to make statistically meaningful statements regarding the general growth of galaxy disks by resimulating just a single isolated disk galaxy.

The answer to the first question is that the mass of a Milky Way halo is below the Birnboim & Dekel (2003) cold mode halo mass threshold, M_{T}, throughout the cold gas filamentary phase at $z \gtrsim 3$, and so a resimulation of this system's evolution at high redshift provides a natural framework

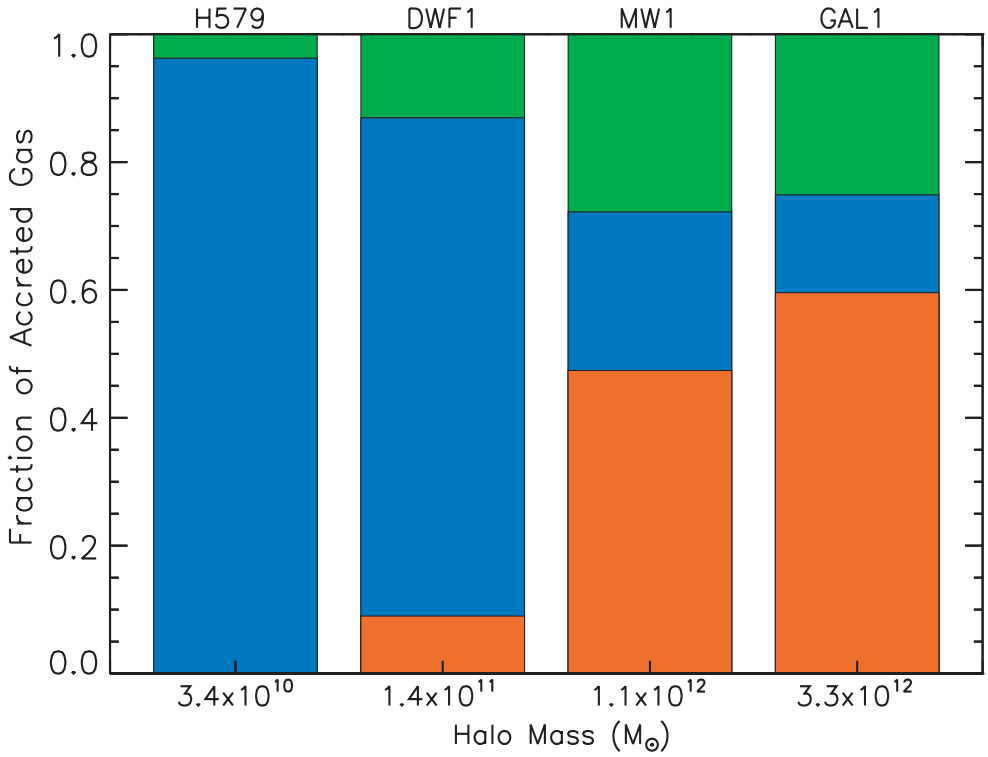

Figure 5.1: This plot has been taken from Brooks et al. (2009) and demonstrates the importance of the cold mode in driving the growth of luminous galaxies. Four resimulated haloes spanning two orders of magnitude in total halo virial mass (dark matter, stars and gas) are shown. The green, blue and orange regions correspond to the fractional contributions of clumpy (satellites and stripped material), unshocked and shocked gas to the total mass budget of all the gas that has accreted onto the virial regions of these haloes between $0 \leq z \leq 6$.

for examining the cold mode. Concerns of the latter type are therefore addressed in the remainder of this section, and it is argued that both the high mass nature and the lack of redshift evolution of $M_{\rm T}$ implies that the cold mode drives the growth of *most* galaxies at high redshift.

The importance of the high mass nature of $M_{\rm T}$ is best illustrated with reference to Fig. 5.1, taken from Brooks et al. (2009). These authors tracked the evolution of 7 galaxies with present day halo masses spanning $3.4 \times 10^{10} \leq M_{\rm H}/{\rm M}_{\odot} \leq 3.3 \times 10^{12}$, resimulated using the GASOLINE code (Wadsley et al., 2004). Their findings in Fig. 5.1 show that over 90% of the total gas mass that has accreted onto their lowest mass halo since $z = 6$ belongs to the unshocked gas phase. What is particularly intriguing is that even for their highest mass halo, whose mass is nearly an order of magnitude above the cold mode halo mass threshold of $M_{\rm T} \sim 4 \times 10^{11}\,{\rm M}_{\odot}$ (Ocvirk et al., 2008), this fraction is only reduced to ∼15%. These findings hence suggest that cold gas accretion plays an important role in the formation and growth of disks in higher mass systems as well: haloes need to be significantly more massive than $M_{\rm T}$ for the cold mode to have contributed negligibly to the

growth of their luminous galaxies, and haloes of this kind sample the exponential tail of the halo mass function and hence constitute only a small fraction of the halo population at a given redshift. Most haloes are not this massive at high redshift, and hence their galaxy growth is probably largely supported by cold gas accretion.

The second property that drives the importance of the cold mode—the lack of redshift evolution of $M_{\rm T}$—is best highlighted by consideration of the evolution of the break mass $M_\star$ of the stellar mass function. Bell et al. (2003) have demonstrated that $M_\star \sim 7.5 \times 10^{10}\,{\rm M}_\odot$ in the K- and g-bands at the present day, which is only a factor of $\sim$3 larger than the Kereš et al. (2005) cold mode stellar mass threshold estimate of $M_{\rm T} \sim 2.5 \times 10^{10}\,{\rm M}_\odot$. Hence the fraction of galaxies at the present day that are expected to have grown via inflows of cold gas at high redshift is large. The decrease in $M_\star$ with increasing redshift due to the evolution of the stellar mass function (e.g. Mortlock et al. 2011), combined with the lack of redshift evolution in $M_{\rm T}$, also means that a larger fraction of galaxies are expected to be growing by accretion of cold gas at higher redshifts ($z \gtrsim 3$).

Hence even though the isolated NUT halo examined in this study, which formed at the intersection of three filaments in the simulation, is probably representative of the high redshift evolution of most low mass galaxies in the local Universe, it does not necessarily correspond to *every* low mass galaxy. Pichon et al. (2010) have demonstrated that while the average number of filaments connecting at halo nodes in the cosmic web peaks at 3, there is a minor population of haloes in both lower and higher density environments at high redshift that have a fewer and greater number of associated filaments than 3, respectively. Milky Way haloes at high redshift that are members of galaxy groups or that are embedded in galaxy clusters are not considered in this thesis, and it is not obvious whether these systems would further highlight the importance of the cold mode in driving galaxy growth. The Pichon et al. (2010) study suggests that these systems are connected by a greater number of filaments than the three that connect to the NUT halo, yet their massive host haloes are probably above $M_{\rm T}$, which may mean that these galaxies accrete a large fraction of their mass in the hot gas phase. At any rate, while this possible environment dependence is an interesting effect that deserves attention, the fraction of low mass galaxies at high redshift in high density environments is expected to be lower than the fraction in the field, as massive hosts form the exponential tail of the halo mass function.

Hence, in summary, despite representing just a single galaxy sample from the myriad of galaxies in existence at the present day, the NUT simulation examined in this thesis should be regarded as a

laboratory for conducting analysis on a mode of collimated gas inflow that is likely to dominate the growth of a large fraction of galaxies at high redshift ($z \gtrsim 3$). The aim of this chapter is to quantify each galaxy component's contribution to the overall angular momentum budget of the simulated NUT CO central disk as a function of time. Fulfilling this objective within an Eulerian grid-based framework inevitably requires a suite of computational tools, and so before analyzing the results, the algorithms that have been used in this study are presented.

5.3 Methods I: Computing flow properties on user-defined grids

There are numerous references to grids of various dimensions throughout this chapter, and so before delving into the array of techniques that have been developed to identify the separate galaxy components (Section 5.4) and compute their angular momentum (Section 5.5), a short summary that explains how the most fundamental feature of an arbitrary grid—its refinement level ℓ—has been determined, is provided. This is followed by an explanation of how these grids have been manipulated to compute: a) the density, temperature and velocity (hereafter 'flow variables') at different locations within the galaxy system, and; b) the centre of galactic rotation $\vec{r}_c$, and its velocity $\vec{v}_c$. The following terms—$r_{\rm vir}$, disk, satellites and virial sphere—are henceforth made in reference to the NUT host halo.

5.3.1 Grid notations

This study analyzes cubic and spherical grids $\mathcal{G}_x$ centred at $\vec{r}_c$ with physical half-lengths and radii equal to $xr_{\rm vir}$, where $x \in \{0.1, 0.15, 0.5, 1.0, 2.0\}$. Hence the spherical grid $\mathcal{G}_{1.0}$ corresponds to a grid that encompasses the spherical virial region. The numerical subscript x will be explicitly stated when discussing a certain grid type, but for comparisons between grid types (e.g. the gas grid $\mathcal{G}_g$ or the stellar grid $\mathcal{G}_\star$) it is replaced by a symbol denoting the type, and the extent of the grid is understood. All of the spherical (cubic) grids in this chapter with physical diameter (full-length) L are subdivided into n cubic cells of equal physical sidelength Δx and are hence 'fixed', obeying the relation:

$$\Delta x = \frac{L}{n}. \tag{5.1}$$

Since there are 2^ℓ cells along each dimension of the physical simulation grid length $L_s(z)$ at level ℓ (Section 2.1.1), it follows that the refinement level of $\mathcal{G}_x$ is given by:

$$\ell = \frac{\ln\left(L_s/\Delta x\right)}{\ln 2}, \tag{5.2}$$

where $L_s(z) = 9\,h^{-1}\,\mathrm{Mpc}/(1+z)$ and ℓ is rounded to its nearest integer value. It will be demonstrated later in this section that the routine responsible for generating fixed grids at a given location requires L and ℓ as inputs at runtime, and so the preliminary codes, whose purpose was to choose these parameter values, always strived to match the n and ℓ values that corresponded to $\Delta x = \Delta x_{\mathrm{min}}$, the minimum allowed value of Δx determined by the universal physical resolution limit of 12 pc in the cooling run (which is hereafter abbreviated to the CO run as in Section 2.3.2). Limited memory availability meant that n could not exceed 1024, however, and so ℓ was successively decreased in integer units in equation (5.2) until this constraint was satisfied. Hence Δx for the grids enclosing the r_{vir} and $2r_{\mathrm{vir}}$ regions tended to increase in deviation from Δx_{min} with time, as the virial radius of the host (and hence L) expanded. This time bias did not affect the smaller grids $\mathcal{G}_{0.1}$ and $\mathcal{G}_{0.15}$ used for performing the angular momentum computations, however, as their cell width was always maintained at the resolution limit.

5.3.2 Obtaining $\mathcal{G}_g, \mathcal{G}_\star$ and $\mathcal{G}_{\mathrm{DM}}$

The amr2cube routine from the RAMSES package[1] has been used to compute the flow variables within gas grids $\mathcal{G}_g$. This code uses the grid centre (corresponding to the centre of the clump of stars belonging to the main host), and the values of L and ℓ_{max} inserted at runtime, to perform two simple tasks. It begins by constructing $\mathcal{G}_g$ and locating the cells within $\mathcal{G}_g$ that do not contain further subcells. Each of these nodes is then assigned the gas flow variables that have been computed for its counterpart on the simulation grid over the same region. $\mathcal{G}_g$ hence contains multiple cells i with $\ell_{\mathrm{min}} \leq \ell_i \leq \ell_{\mathrm{max}}$, where $\ell_{\mathrm{min}} = 7$ and ℓ_{max} is given by fixing Δx in equation (5.2) to Δx_{min} (but recall that the inserted value of ℓ_{max} for $\mathcal{G}_{1.0}$ and $\mathcal{G}_{2.0}$ is not necessarily guaranteed to match this upper limit for late times, as discussed in Section 5.3.1). The second stage of the algorithm involves transforming the multi-level grid to a fixed grid of equal cell sizes, and so $\mathcal{G}_g$ is divided in such a way that each cell is refined to the same level (ℓ_{max}). The flow variables are then interpolated to the fixed grid.

[1] http://web.me.com/romain.teyssier/Site/RAMSES.html

Fig. 5.2 further illustrates the gas cube construction process within the amr2cube code by making reference to a density grid, which for simplicity is shown in 2D. The left panel represents the dimensions of the grid inserted at runtime and corresponds to the simulation grid $\mathcal{G}_{\rm sim}$ refined up to the inserted value of $\ell_{\rm max}$ over the same region. The cell labelled j in the purple oct represents a grid node as it does not contain further child octs, and so the density it is assigned is directly read from the corresponding cell on the simulation grid. Cell i within the red oct (which itself is a child of the blue oct) is also a node, and in this mock scenario it has been refined to the inserted value of $\ell_{\rm max}$, hence its density value is also read from the simulation grid. Once every node with the specified region has been assigned a density, amr2cube refines *every* cell to $\ell_{\rm max}$, as shown in the right panel. Cell j has therefore been divided into a higher level oct, with its density assigned to each cell member within this newly created oct. Cell i, by contrast, is already at $\ell_{\rm max}$ and so no interpolation is required. The appropriate cell volumes between the left and right panels ensure that mass (momentum) is conserved upon interpolation of density (velocity components).

Having computed $\mathcal{G}_g$, it may seem misleading to introduce the notion of dark matter grids $\mathcal{G}_{\rm DM}$ and stellar grids $\mathcal{G}_\star$, as these components are represented by particles in the simulation and hence their properties should not need to be computed on a grid. $\mathcal{G}_\star$ and $\mathcal{G}_{\rm DM}$ are required in the densest cell computation described later in Section 5.3.3, however, and so the part2cube routine from the RAMSES suite has also been used in order to generate particle grids with exactly the same dimensions as the gas grids computed by amr2cube (see the right panel of Fig. 5.2). The part2cube algorithm performs a Cloud-in-Cell smoothing technique of each particle's mass using periodic boundary conditions. The box centred on the yellow particle in Fig. 5.2 corresponds to the cloud volume, and each neighbouring grid cell (e.g. the green cell) that intercepts this cloud is assigned a portion of the particle's mass according to the fractional volume shared between the cloud and the neighbouring cell (shaded area). Mass is hence also conserved by construction in the part2cube routine.

With $\mathcal{G}_g$, $\mathcal{G}_\star$ and $\mathcal{G}_{\rm DM}$ now defined over a given region, the subscript in reference to the inserted value of the maximum refinement level is henceforth dropped, and it is understood that $\ell(\mathcal{G}_x) = \ell_{\rm max}(\mathcal{G}_x)$.

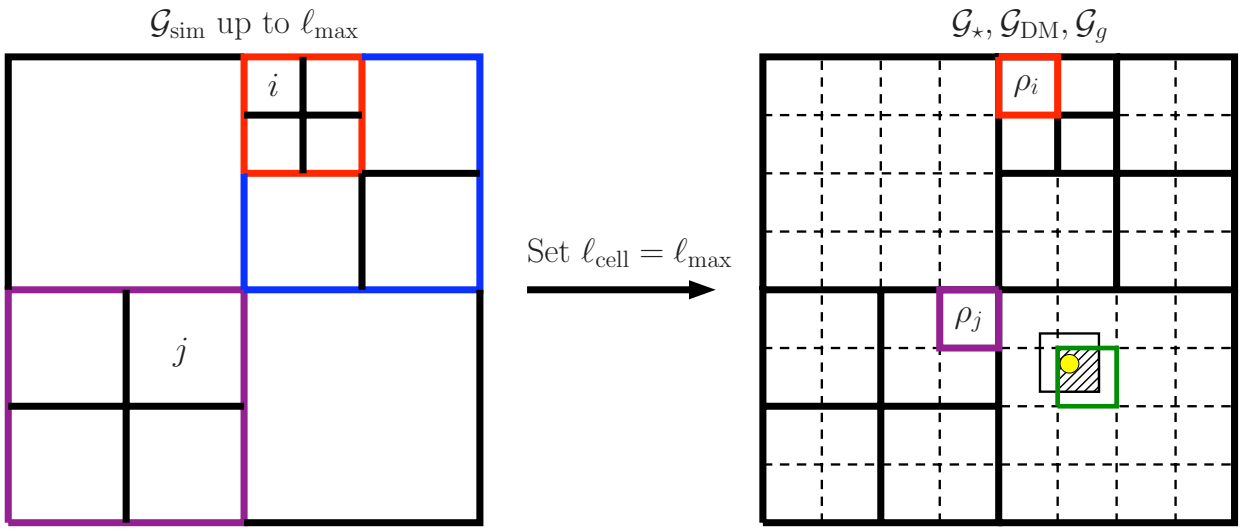

Figure 5.2: A simple example illustrating the principles of the amr2cube and part2cube routines that have been used to compute the flow variables within fixed grids. A density grid is considered here. The amr2cube code starts by reading the densities of all the cell nodes (e.g. i and j) from their counterparts on the simulation grid $\mathcal{G}_{\rm sim}$ (left panel), and these values are then interpolated to cells on a fixed gas grid $\mathcal{G}_g$ whose cell width is determined by the value of $\ell_{\rm max}$ inserted at runtime, as shown by the thin black dashed divisions in the right panel. The part2cube routine locates the star and dark matter particles within the fixed grid and performs a Cloud-in-Cell smoothing of their mass over a fixed cell volume. This scheme is shown for the yellow particle in the right panel, and the shaded area represents the fraction of its mass assigned to the neighbouring green cell. $\mathcal{G}_{\rm sim}$ is superimposed on the fixed grid in the right panel (thick black lines) to help illustrate the cellular interpolation process, and cells i and j are coloured according to their member octs on $\mathcal{G}_{\rm sim}$.

5.3.3 Estimating the location and velocity of the centre of galactic rotation

The grids from the previous section have been computed over regions centred on the densest location within the host galaxy's group of stars. As argued below, the location and motion of this point are not always the best representations of the centre of rotation $\vec{r}_c$ and velocity $\vec{v}_c$ of the host disk galaxy. More accurate estimates of these quantities, which are of central importance to most of the algorithms in this chapter, are hence required.

There are two obvious candidates for $\vec{r}_c$ within the central galaxy: 1) the densest cell, which is likely to correspond to the disk centre, and; 2) the centre of mass. Even though both of these definitions have been found to yield very similar estimates of $\vec{r}_c$, it is worth briefly highlighting their relative advantages and disadvantages in the general case. The densest cell has the desirable property that it represents a local minimum in the gravitational potential that is likely to often correspond to the centre of the gas disk. Its drawback is that it is potentially subject to noise, and may, for example, coincide with a massive clump of stars born in a region slightly offset from the

disk's centre, at an epoch where the stars dominate the mass budget of the densest cell. While this was not found to be a problem in the CO run, it could become more significant when supernovae feedback is included in future extensions. The centre of mass, on the other hand, has the advantage that it is well-defined for a gravitationally bound object, but at high redshifts where the main host undergoes major-mergers, it could be located in a region relatively devoid of mass, which is unlikely to correspond to the centre of the disk. This study therefore adopts the densest cell as the best estimate of $\vec{r}_c$.

Fig. 5.2 has shown that cells within a fixed grid have equal volume, and so the procedure for identifying the densest cell began by constructing grids $\mathcal{G}_{0.1}$ centred on the group of stars belonging to the host galaxy. A grid radius of $0.1r_{\rm vir}$ has been chosen for three reasons: (i) it represents the region where the amount of specific angular momentum transported by gas starts to decline (Kimm et al., 2011); (ii) it fully encloses both the gas disk at each time-step (as has been confirmed by visual inspection of the final output at $z = 3$) and the centre of the host dark matter halo identified by the halo-finder, and; (iii) it mostly includes material bound to the central galaxy, unlike more extended regions, which probably contain larger amounts of diffuse gas. The total density of the $i^{\rm th}$ member cell within $\mathcal{G}_{0.1}$ is given by:

$$\rho_i = \rho_{\rm DM} + \rho_\star + \rho_g, \tag{5.3}$$

and so the densest cell $\rho_{\rm max}$ was computed by simply summing $\mathcal{G}_g$, $\mathcal{G}_\star$ and $\mathcal{G}_{\rm DM}$ over the $0.1r_{\rm vir}$ region and identifying the cell with the maximum density. It was found that the gas dominated the mass budget of $\mathcal{G}_{0.1}$ at high redshift, whereas the stars dominated at lower redshifts. At $z = 3$ ($z \sim 9$), for example, the fractional contributions to $\rho_{\rm max}$ were $\sim$0.97 (0.17), $\sim$0.02 (0.01) and $\sim$0.01 (0.82) for the stars, dark matter and gas respectively.

The velocity of the centre of disk rotation $\vec{v}_c$ has been approximated by the centre of mass velocity $\vec{v}_{\rm com}$ of $\mathcal{G}_{0.1}$:

$$\vec{v}_c \simeq \vec{v}_{\rm com} = \frac{\sum_j m_j \vec{v}_j}{\sum_j m_j} \tag{5.4}$$

$$\vec{v}_j = \vec{u}_j + \vec{w}_j, \tag{5.5}$$

where the sum is performed over all of the gas cells and stellar and dark matter particles within $0.1r_{\rm vir}$. The velocity of the disk centre can be thought of as consisting of three components: a bulk

streaming part $\vec{u}$, an internal random component $\vec{w}$ and a rotational motion about its centre. In the centre of mass frame of disk rotation, this latter component is zero by definition. The disk is also likely to have a negligible net random velocity in this frame, as the non-zero $\vec{w}_j$ largely cancel upon mass-weighted summation. It is therefore important that the estimate for $\vec{v}_c$ is dominated by the disk's bulk motion. However, this constraint is not necessarily satisfied when one considers the velocity of the densest cell because the physical cell size of 12 pc limits the amount of gas within the cell, potentially yielding a non-negligible random velocity component. It is also possible that the densest cell rotates about the disk's centre of mass, especially if there are multiple dense stellar clumps in the central region. By contrast, there are a large number of cells and particles that contribute to the centre of mass computation in equation (5.4), and so there is a higher probability that the net random motion is zero and hence closer to the bulk velocity with this choice of $\vec{v}_c$.

Despite the above differences in definition, these approximations for $\vec{r}_c$ and $\vec{v}_c$ have proved to be robust—when just the stars were used in determining the densest cell or centre of mass, similar results were obtained.

5.4 Methods II: Resolving the galaxy components

This section describes the algorithms that have been used to identify the likely sources of the disk's angular momentum within the virial region, and is divided into three main subsections. The details of the halo- and galaxy- finder are briefly reviewed in part one, the filament identification routine is presented in part two, and the satellite-finding algorithm is discussed in part three.

5.4.1 Resolving galaxies and following their merger histories

The MSM subhalo-finding algorithm (Section 2.2.2) has been used to resolve dark matter haloes and stellar clumps, and has been run with the same cosmological parameters as the NUT CO run (Section 2.3.2). Each resolved dark matter object and stellar clump was forced to contain at least 40 and 100 particles respectively, to ensure reliable detections. The stellar clump resolution limit was set above the dark halo limit because the clump-finder fragmented the disk into several separate groups of stars when lower thresholds were imposed. The TreeMaker code (Section 2.2.3) then provided a list of all the fathers and sons of every halo and subhalo, and every stellar clump and subclump.

5.4.2 Disentangling the filaments

In order to ascertain whether cold filament flows control the mass and angular momentum budget of the disk at high redshift, it is necessary to first define gas belonging to the filamentary phase and then resolve the individual trajectories. The techniques adopted to perform these tasks are explained in this section.

The tracer particle colouring algorithm

Powell et al. (2011) have shown that parcels of gas whose number density of hydrogen atoms n and temperature T simultaneously satisfy $0.1 \leq n/\mathrm{cm}^{-3} \leq 10$ and $0 \leq T/\mathrm{K} \leq 2 \times 10^4$, are representative of gas belonging to filaments in the NUT simulations at high redshift ($z \geq 9$). To tackle the possible time evolution in the filament density cuts, Kimm et al. (2011) introduced a prescription between the lower threshold and the background density, arguing that large-scale filament density is gravitationally coupled to the expansion of the Universe. They defined the lower number density bound n_{L} as:

$$n_{\mathrm{L}} \equiv \frac{\delta_f \bar{\rho} f_b X_{\mathrm{H}}}{m_{\mathrm{H}}}, \tag{5.6}$$

where $\bar{\rho}, \delta_f, f_b, X_{\mathrm{H}}$ and m_{H} are respectively the mean density of the background, the density contrast of the filament (measured with respect to $\bar{\rho}$), the fraction of the total mass in baryons ($\equiv \Omega_b/\Omega_{\mathrm{M}}$), the primordial relative mass abundance of hydrogen (76%), and the mass of a hydrogen atom. A typical value of n_{L} from the NUT CO run is $\sim 0.01\,\mathrm{cm}^{-3}$, which is an order of magnitude lower than the fixed Powell et al. (2011) estimate. This difference is not surprising, however, as Powell et al. (2011) only examined filament gas flow rates onto the central NUT galaxy down to $z = 9$ and were therefore somewhat justified in ignoring the evolution of the density of the filament phase in equation (5.6). By contrast, the filaments in this thesis are resolved down to $z = 3$ and hence evolution in their density is more significant.

It was found that an upper temperature limit of $T_u \leq 2 \times 10^5\,\mathrm{K}$ detected filaments to a higher level of accuracy for $z \geq 3$ in the CO run than the classical cold mode temperature of $T_u \leq 2 \times 10^4\,\mathrm{K}$ (e.g. Kay et al. 2000; Birnboim & Dekel 2003; Kereš et al. 2005) at which gas is thought to cool via Lyman-α emission (e.g. Fardal et al. 2001). There are two likely reasons for this. Firstly, the UV background at the epoch of reionization ($z_{\mathrm{re}} = 8.5$) heats the ISM and hence the filament gas. Secondly, the absence of feedback from supernovae in the CO run causes the filaments to remain

hotter for longer, as cooling is not so efficient in low metallicity environments (e.g. Sutherland & Dopita 1993). Therefore, in order to identify the neutral hydrogen and helium gas in the filamentary phase in this study, the density and temperature (hereafter nT) criteria reported by Powell et al. (2011) have been used for $z \geq 3$, but the lower density threshold has been replaced by the value given by equation (5.6) and the upper temperature limit has been increased to $T_u = 2 \times 10^5\,\mathrm{K}$. This particular choice of T_u is also consistent with the definition adopted in other similar studies (e.g. Kereš et al. 2009; Faucher-Giguère et al. 2011).

The left panel in Fig. 5.3 shows an image of the three filaments within $2r_{\mathrm{vir}}$ of the host centre at $z \sim 10$ in the CO run, found by imposing the nT criteria. This image has been generated using the VAPOR software[2]. The rightward red, upward green and forward blue arrows correspond to the x, y and z directions respectively, and the colour scale indicates the corresponding $\log_{10}(\mathrm{n/cm^{-3}})$ densities. It will be demonstrated later in this section that the nT criteria can include gas gravitationally bound to satellites as part of the filaments, but a discussion of how gas that is shared by filaments and satellites is apportioned between these two phases is preserved for Section 5.4.3.

The filaments in the left panel of Fig. 5.3 occupy three distinct regions at $z \sim 10$ on large scales approaching $2r_{\mathrm{vir}}$, and the flow appears to be ordered and mostly radial. The dynamics of the gaseous motions in the central region are far more complex, however, as demonstrated in the right panel of Fig. 5.3, which shows the velocity flow vectors of filament gas within the virial region at a later time (corresponding to $z \sim 8$) when the system has evolved to a different configuration. VAPOR computes these velocity trajectories by performing a numerical flow integration of the velocity field within the volume-rendered region over a small time interval (see, for example, Clyne & Rast 2005). The starting points of the integration were randomly distributed over a small box that was placed along each filament outside the central region, and the flow lines generated from the subsequent integration represent an estimate of the individual filament trajectories, which bend, twist and mix with each other in the central region, yet often remain intact. In order to test the hypothesis that the disk acquires its angular momentum from filamentary motions, it is critical to follow these separate individual filament trajectories. A suitable tool for performing this Lagrangian exercise is the use of tracer particles, which have been demonstrated to preferentially trace filament flows in galaxy simulations (Pichon et al., 2011). The tracer particles in the NUT suite have an identical spatial distribution to dark matter at the beginning of each simulation, and are assigned

[2] www.vapor.ucar.edu

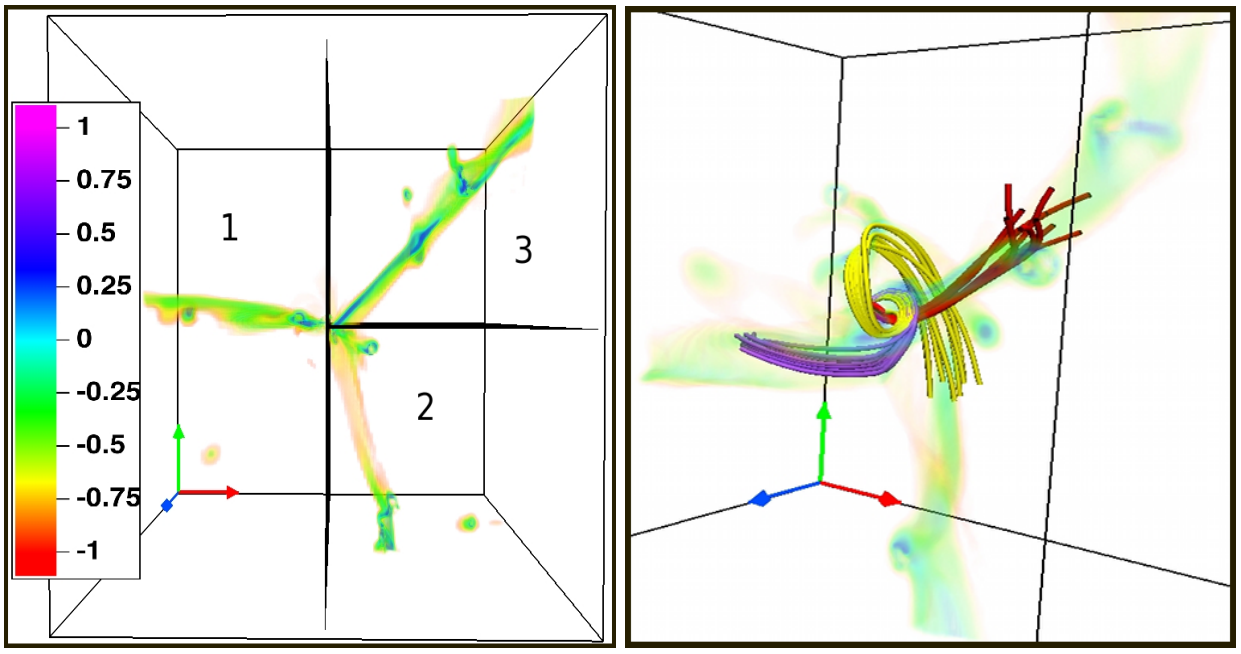

Figure 5.3: Images of the filament gas within a box centred on the main host at $z \sim 10$ (left) and $z \sim 8$ (right) in the CO run, generated using the VAPOR visualization tool. The sidelength of the box in the left panel corresponds to $4r_{\rm vir}$ ($\sim$17 kpc) and this face-on view illustrates the clear separation of the filaments into three distinct regions (labelled accordingly) at large scales and early times. The image in the right panel shows the inner virial region ($\sim$7.5 kpc) at lower redshift when the dynamics of the system have changed, and the velocity flow tubes demonstrate the complex trajectories of gas flow within the gas disk region, colour coded according to the filament that they are likely to trace. The density of the gas has been dimmed and the viewing angle rotated to best highlight each trajectory. Both panels have been refined to $\ell = 16$, corresponding to a physical cell width of $\sim$17 pc (left) and $\sim$22 pc (right), and the colour bar indicates the gas density in dimensionless units of $\log_{10}({\rm n/cm^{-3}})$, where n is the number density of hydrogen atoms per cubic centimetre. Filament gas has been detected by applying the nT criteria, but gas with density $\log_{10}({\rm n/cm^{-3}}) < -1$ has been excluded from both figures for clarity. The rightward red, upward green and forward blue axes correspond to the x, y and z directions, and the vertical and horizontal black planes in the left panel show the spatial cuts in the x and y directions that have been used to start the colouring of the tracer particles, which is a technique that has been developed to separate each of the filaments at high redshift (see text).

zero mass so as to follow the velocity field within the host virial region.

The algorithm that resolves each filament at high redshift is henceforth referred to as the 'tracer propagation method', and consists of three principal stages:

1. **Colouring the tracer particles across time**

   At some initial early time t_1, planar cuts are used to divide the filaments into distinct regions. Each region is assigned a unique colour, and tracer particles are initially coloured according to the region in which they reside.

For later times, the positions of coloured tracers $\vec{r}_i(t-1)$ from the previous time-step $t-1$ that spatially coincide with a filament at the current time-step t, are updated to their new locations $\vec{r}_i(t)$. Each uncoloured tracer j that has just moved onto a filament trajectory within the region $\mathcal{R}$ covered by the fixed spherical grid (which grows with time) is assigned the majority colour of its nearest N coloured tracer neighbours at t (N was at least 100).

2. **Colouring the grid**

   Once all of the filament tracer particles within $\mathcal{R}$ are coloured, the overall colours of the cells within the grid spanning $\mathcal{R}$ are determined. The prescription used in this study enforces the condition that each coloured tracer cell is 'pure', consisting entirely of particles with identical colour.

3. **Individual filament growth**

   Pure cell colours are then propagated to neighbouring uncoloured filament cells within $\mathcal{R}$.

Each of these steps is now explained in more detail.

Step 1: Colouring tracers at each time-step

The horizontal and vertical black planes in the left panel of Fig. 5.3 show the cuts that have been used to separate the filaments at the initial time-step in the CO run, corresponding to $z \sim 10$. These planar cuts have only been applied in the x and y directions and intersect at the densest cell location within the virial sphere, corresponding to the point of rotation of the system (Section 5.3.3). Each separate region is subsequently assigned a unique colour and all of the tracer particles within $r_{\rm vir}$ that coincide with a filament cell (as identified using the nT criteria) are given one of these colours depending on the region they occupy. This method therefore initially resolves a single filament per region, by construction, and throughout this chapter reference will be made to the purple, yellow and red filaments in regions $1, 2$, and 3 respectively.

The Lagrangian nature of the system poses two computational difficulties when assigning colour:

- filaments mix in the vicinity of the disk (as can be seen in the right panel of Fig. 5.3), and;
- filaments acquire a drift velocity on scales $r \gtrsim r_{\rm vir}$ (Pichon et al., 2011), which can lead to them switching between regions within $r_{\rm vir}$ as the host grows.

Filament trajectories are hence not necessarily confined to a single region, and so several tracer particles are incorrectly coloured when using the planar cuts. To tackle the first issue of mixing,

the filament trajectories are not resolved until the following time-step t_2, by which time the vast majority of the incorrectly coloured particles within $0.1r_{\rm vir}$ at t_1 have accreted onto the disk during the $\Delta t = t_2 - t_1$ interval, and are therefore no longer associated with a filament (the upper panel of Fig. 5.4, which is discussed in more detail later in this section, provides a visual justification of this assumption). To circumvent the second issue of filaments acquiring drift velocities on large scales, previously flagged tracer particles i at locations $\vec{r}_i(t_1)$ that have moved during the Δt interval, are updated according to their new instantaneous locations $\vec{r}_i(t_2)$ at the following time t_2, provided that they coincide with filament cells at t_2. All of the newly accreted tracer particles that are co-spatial with filaments within the (larger) host spherical region at t_2 that have not been assigned a colour, are given a colour according to their host region. At subsequent time-steps t_j (t_3 onwards), the newly accreted tracer particles are instead assigned the majority colour of their N nearest coloured tracer neighbours that sample the filaments at $\vec{r}_i(t_j)$. Rather than being directly sorted, which is computationally expensive for this type of problem, the distances between each new tracer particle and each of the coloured neighbours have been binned, and the particles in ascending bins have been extracted until N is at least 100 (the size of the bins are chosen such that the variation in the number of neighbours considered is small). The new tracer particles are then assigned the majority colour of their N neighbours. Note that the above neighbouring scheme cannot be used for t_2 because the constructed filament trajectories are not known until the end of the time-step (see steps 2 and 3 below).

The colouring of tracer particles moving along filaments at a given time-step t_j takes place over a spherical grid $\mathcal{G}_{2.0}$ of radius $2r_{\rm vir}$ centred on the densest cell within the central region. If a smaller colouring radius of $r_{\rm vir}$ is chosen, most of the tracer particles at the subsequent time-step t_{j+1} move into the $0.1r_{\rm vir}$ disk region. Filaments can then appear to mix at large scales or swap positions when using a nearest coloured neighbours scheme at t_{j+1}. To illustrate this point further, it is possible that a group of yellow tracer particles reside just outside the disk region and act as the nearest neighbours to an uncoloured tracer particle moving along the purple filament's trajectory. The new particle is hence coloured yellow rather than purple, and this colour risks being propagated along the length of the purple filament at later time-steps. One solution to this problem is to sample regularly along a greater length of each filament at every time-step, which involves colouring over a larger radius, and this is the approach undertaken in this study. A sphere radius of $2r_{\rm vir}$ is used because coloured tracers between $r_{\rm vir}(t_j)$ and $2r_{\rm vir}(t_j)$ were found to populate the full length of the

$r_{\rm vir}(t_{j+1})$ region with this scale choice, largely nullifying any risk of incorrect colour assignment in the central region where the angular momentum computations are performed.

The upper panel of Fig. 5.4 shows tracer particles colour coded according to the filament they sample at $z \sim 9.5$, updated to their new filament positions at $z \sim 9$. The displaced tracers extend across the full virial region at $z \sim 9$ (given by the inner solid circle), and so new tracers entering the $2r_{\rm vir}$ region at this epoch (denoted by the outer solid circle) are appropriately coloured. Since colour is propagated across time, one is able to test the overall performance of the method by considering the final output. Coloured filament tracers within $2r_{\rm vir}$ at $z = 3$ are hence shown in the lower panel of Fig. 5.4 (only 5% of all the coloured filament tracers are plotted for clarity, but this subsample is representative of the whole distribution). The purple and yellow filaments undergo a merger around $z \simeq 5.5$ in the NUT simulation (this is discussed in more detail Chapter 6), and the algorithm clearly captures the aftermath of this merger as only the purple and red tracer particles remain at $z = 3$. The filaments are highly perturbed within the inner virial region and trace complex trajectories, no longer displaying the ordered gas inflow seen at higher redshift in the upper panel. The dashed circles enclose the $2r_{\rm vir}$ regions at the previous time-steps, and are centred at the densest cell positions at these times. The offset in centres between the outer solid and dashed circles indicates the motion of the central point of rotation, which is more stable at late times when major-merger activity is reduced.

Step 2: Determining the colours of the grid cells

By this stage, every tracer particle in a filament cell within $r_{\rm vir}$ at time-step t_2 has a colour. It is, of course, possible that several tracer particles with different colour belong to the same cell using the above scheme, and so the next task is to assign an overall colour to each tracer cell (particles located in non-filament cells within $r_{\rm vir}$ are ignored, so that only filaments are resolved). A fourth 'contaminated' colour was therefore introduced, corresponding to cells that contained at least one tracer particle whose colour was different to that of the other particles. By contrast, 'pure' cells consisted entirely of tracer particles of identical colour.

In this study, the contaminated cells are ignored and only the pure cells are considered reliable filament tracers. It may seem that filaments are unnecessarily under-sampled with this approach (an extreme example being a cell containing N tracers, where N is large, with $N - 1$ of the same colour and only one of different colour), especially since the cell width of $\mathcal{G}_{2.0}$ increases with time (as discussed in Section 5.3.1) to $\Delta x \sim 190\,{\rm pc}$ (physical) at $z = 3$, hence increasing the likelihood

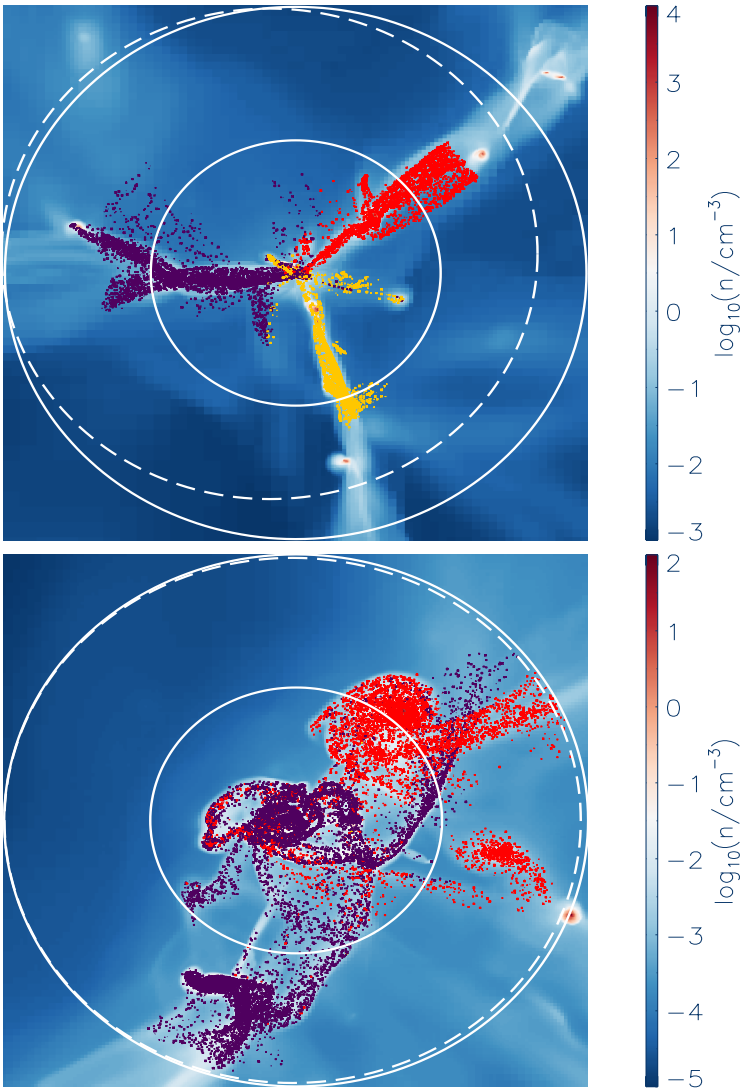

Figure 5.4: These gas density images are z projections of the large-scale region surrounding the central disk in the CO run, and illustrate the spatial distribution of particles tracing the filaments at $z \sim 9$ (upper panel) and $z = 3$ (lower panel). The inner and outer solid circles have radii equal to $r_{\rm vir}$ and $2r_{\rm vir}$ respectively, where $r_{\rm vir} \sim 5.3\,{\rm kpc}$ (upper) and $r_{\rm vir} \sim 31\,{\rm kpc}$ (lower). The dashed circles represent the $2r_{\rm vir}$ regions at the previous time-step, corresponding to $z \sim 9.5$ ($r_{\rm vir} \sim 4.7\,{\rm kpc}$) and $z \sim 3.04$ ($r_{\rm vir} \sim 30\,{\rm kpc}$). All of the circles are centred at the location of the densest cell in the disk region at the relevant time-step. Upper panel: Those tracer particles that have been identified with each of the purple, yellow and red filaments at $z \sim 9.5$ and that coincide with filament cells refined to $\ell = 15$ ($\Delta x \sim 38\,{\rm pc}$) at $z \sim 9$, have been updated to their new positions. The offset in centres between the outer solid and dashed circles arises due to the motion of the centre of disk rotation, and the gap between each set of coloured tracers and the dashed circle highlights the displacement of the tracers over the time interval. Uncoloured new tracer particles that have crossed the $2r_{\rm vir}$ sphere over the interval and that move along filament trajectories at $z \sim 9$ have been omitted. Lower panel: The filament configuration at $\ell = 14$ ($\Delta x \sim 190\,{\rm pc}$) for the final time-step at $z = 3$. This image includes both updated coloured tracers and new tracers that have been assigned a colour using the nearest neighbours method (see text). In Chapter 6, it will be demonstrated that the purple and yellow filaments start to merge around $z \sim 5.5$, with the former filament surviving, and this is confirmed by the absence of yellow tracers in this panel. Only 5% of all the coloured filament tracers are shown at this time-step, for clarity.

of finding a contaminated cell. The decision to only assign filament colour to pure cells is justified for two reasons, however: (i) there are sufficient numbers of pure coloured tracer cells to reproduce the apparent trajectories of each filament at a given time-step (as demonstrated in step 3 below), and; (ii) the use of ad-hoc criteria to help decide the overall colour of a cell that contains a certain ratio of particles of different colours is avoided.

Step 3: Constructing the individual filaments

The above pure tracer cell colour assignment yields far fewer coloured cells than the total number of filament cells found using the nT selection criteria (for example, $\mathcal{G}_{1.0}$ at $z \sim 8$ contained $\sim 2.6 \times 10^4$ pure tracer cells and $\sim 1.6 \times 10^7$ filament cells). One might naïvely think that the remaining uncoloured filament cells should be assigned the colour of their nearest tracer cells, in concordance with the tracer particle colouring method. However, given the sparse sampling of tracer cells with pure colour, there is a danger with this approach of incorrectly colouring a cell that belongs to a purple filament, red (for example), in regions where the two filaments come into proximity, because there just so happens to be a red tracer cell in its vicinity. This issue is not so important for the uncoloured tracer particles entering filaments within the $2r_{\rm vir}$ region (step 2), because: a) the distribution of coloured tracers along each filament is not as sparse as the corresponding number of pure tracer cells over the same region, and; b) most of the tracer colouring takes place on large scales, beyond the central mixing region.

The final step of the filament identification algorithm therefore includes a colouring scheme that grows each filament around the pure colour cells, illustrated in 2D in Fig. 5.5. The filaments are simultaneously grown around the filled squares, which correspond to the tracer cells with pure colour. Their filament neighbours (i.e. cells which satisfy the nT criteria) are shown as the hollow squares, and are coloured according to their filled neighbour's colour. Therefore, after the first round of growth, all of the hollow squares in Fig. 5.5 become filled squares and are flagged. In the second round of colouring, the uncoloured unflagged filament neighbours of these newly filled squares are considered. The procedure terminates once all neighbours of neighbours etc, are coloured and the only remaining neighbouring cells are non-filament or previously-coloured cells. Note that only pure colours are propagated: the contaminated tracer cells are skipped in order to avoid mixed regions extending into regions where filaments do not physically mix.

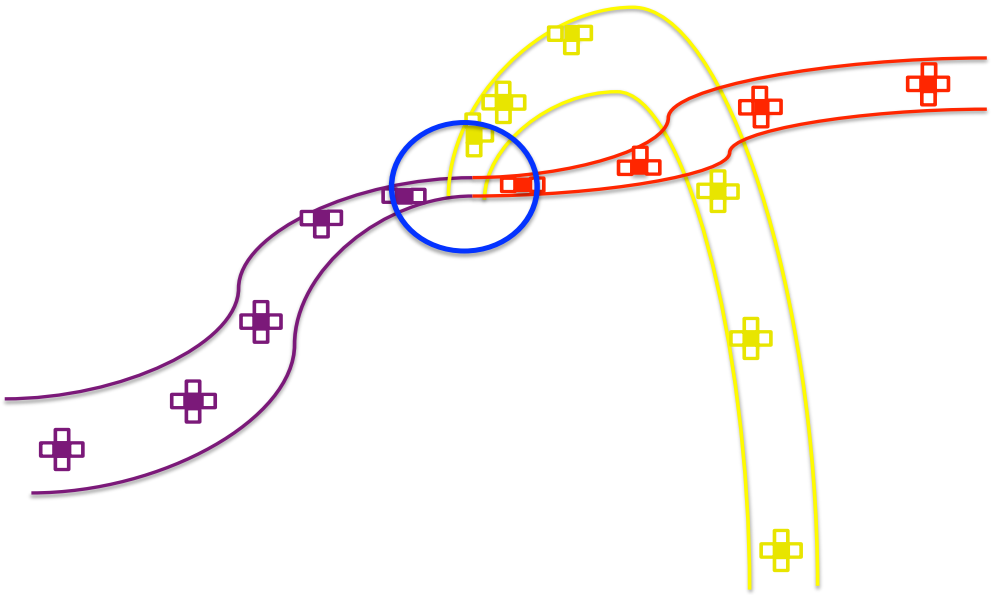

Figure 5.5: A diagram illustrating the colour propagation technique on $\mathcal{G}_{1.0}$ for a mock filament configuration. Each pure tracer cell is shown as a filled square and filament neighbours are drawn as hollow squares. The three filaments are grown simultaneously, promoting the hollow squares to filled squares, and the filament neighbour cells are flagged upon colour assignment. Neighbours of these cells are then examined and the colouring method is repeated until there are no more uncoloured filament neighbours remaining. The disk region is shown by the blue circle.

Errors

In order to resolve individual filament trajectories, one requires a prescription for determining gas in the filamentary phase, and in this study the nT cuts have been used. While these cuts are only approximate, Powell et al. (2011) have argued that their associated errors are likely to be negligible because these authors found that these cuts were able to successfully resolve filaments (as confirmed by visual verification) around the disk in the NUT suite runs with $\Delta x_{\rm min} = 0.5\,{\rm pc}$. Hence the errors associated with the nT criteria are ignored in what follows, and the constructed filaments are calibrated against their nT counterparts.

It is possible to estimate the significance of the errors in the filament colour propagation scheme by performing two simple tests:

1. compute the fraction of filament cells recovered, and;

2. determine whether the apparent behaviour of the individual filament trajectories is captured.

The colouring algorithm does not guarantee that every filament cell is assigned a colour because it is not able to propagate colour between physically unconnected filament clumps. Fig. 5.6 examines

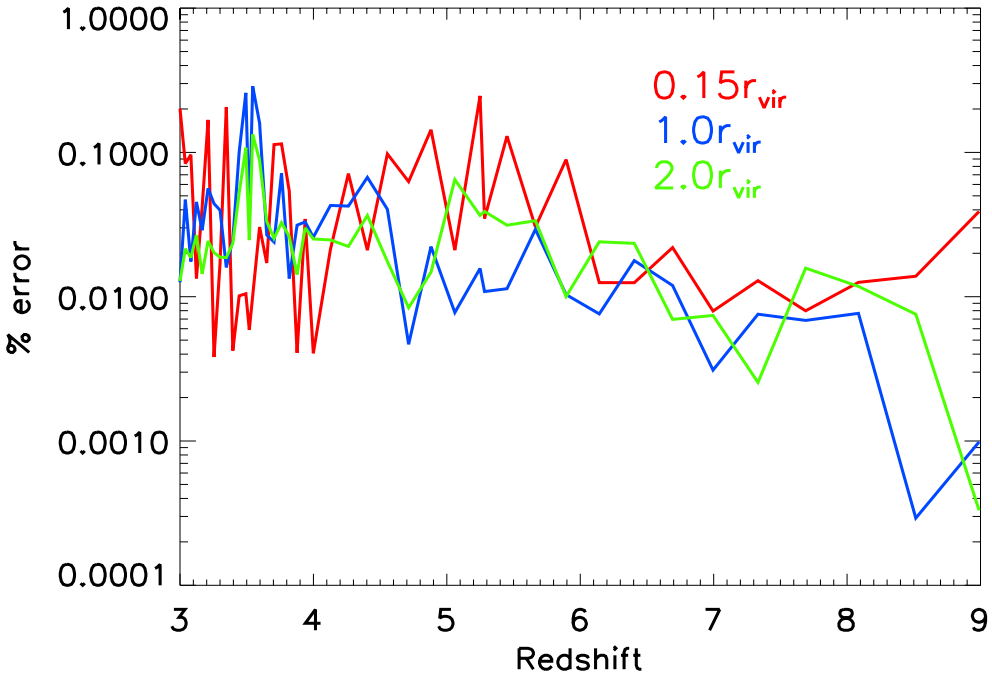

Figure 5.6: Filament recovery errors associated with the tracer colour propagation scheme, shown as a function of redshift and grid radius. The errors at each epoch correspond to the number of filament cells that are not assigned a pure colour by the algorithm, expressed as a fraction of the total number of filament cells found using the nT criteria. The grid covering the $0.15r_{\rm vir}$ region has the highest spatial resolution of the three grids and is used for the angular momentum computations.

the relative importance of these pathological scenarios by showing the number of filament cells left without a colour upon application of the propagation method, as a fraction of the total number of filament cells found by applying the nT criteria, for different grid sizes at $z \geq 3$ in the CO run. It can be seen that errors of type 1 mentioned above are negligible at all times and are never above the percent level.

Errors of type 2 are now examined. Since the aim of this project is to measure the amount of angular momentum transferred to the disk by filaments, the trajectories of these cold flows in the disk's vicinity are required to a high level of precision. The filament reconstruction stage (step 3 of the colour propagation algorithm) was hence repeated over the $0.15r_{\rm vir}$ region to ensure that the angular momentum calculations around the disk were performed at the resolution limit (this particular choice of grid radius is justified in Section 5.5.3). Once constructed, it is natural to wonder whether the filament trajectories are representative of their 'true' counterparts in the simulation. A rigorous quantification of the relative discrepancies appears to be a non-trivial task, and so a visual check was adopted instead. Fig. 5.7 shows 3D images of the density of each filament

from the right panel of Fig. 5.3, as found by implementing the colour propagation technique on $\mathcal{G}_{0.5}$ with spacing $\Delta x \sim 12\,\mathrm{pc}$ at $z \sim 8$. By zooming into the larger $0.5r_{\mathrm{vir}}$ region at the resolution limit, it is possible to test the algorithm's ability to follow each filament onto the disk. The viewing angle in each panel of Fig. 5.7 has been rotated to demonstrate the nature of the individual trajectories, and varies from filament to filament. By comparison with the right panel of Fig. 5.3, it appears that the separate paths are correctly identified to the point of disk contact, and that the curvature about the disk region is particularly well-resolved, which is most evident for the yellow filament in region 2. Note that the sudden upturn of the red filament's inner trajectory in Fig. 5.3 after mixing in the central region is captured by the colour propagation routine. Fig. 5.7 hence suggests that the algorithm is successful in constructing each filament's trajectory.

5.4.3 Locating the satellites

There are two types of satellite in the simulation: satellites with a dark matter halo and satellites whose dark halo component has been stripped by tidal forces and is hence not detected by the halo-finder. One property all satellites share in common, however, is that they are substructures of the main host. As shown in the upper panel of Fig. 5.8 (which is discussed in more detail later in this subsection), luminous satellites in the simulation are often found to stream along filaments, eventually merging onto the central disk, perhaps transporting large amounts of angular momentum in the process. One is therefore faced with the following dilemma: should the angular momentum of gas gravitationally bound to a satellite, which itself is drifting along a filament, be solely attributed to the satellite? The answer to this question in this study is chosen to be 'no': only gaseous material within the virial region of a satellite that does not satisfy the nT criteria and that lies above the upper filament density cut, is apportioned to the satellite phase. 'Satellite' is therefore made in reference to the central dense galaxy component that most probably dominates the 'true' satellite angular momentum signal. Since this definition neglects gas that is gravitationally bound to the satellite beyond its central region, it is tempting to think that an alternative scheme that attributes all the gas within the satellite dark halo to the satellite phase would provide a more reliable estimate of the satellite angular momentum. The issue with this approach is that large spherical regions would be masked out of the filaments, leaving behind unphysical filament trajectories. The advantages of the satellite definition adopted in this work are hence two-fold:

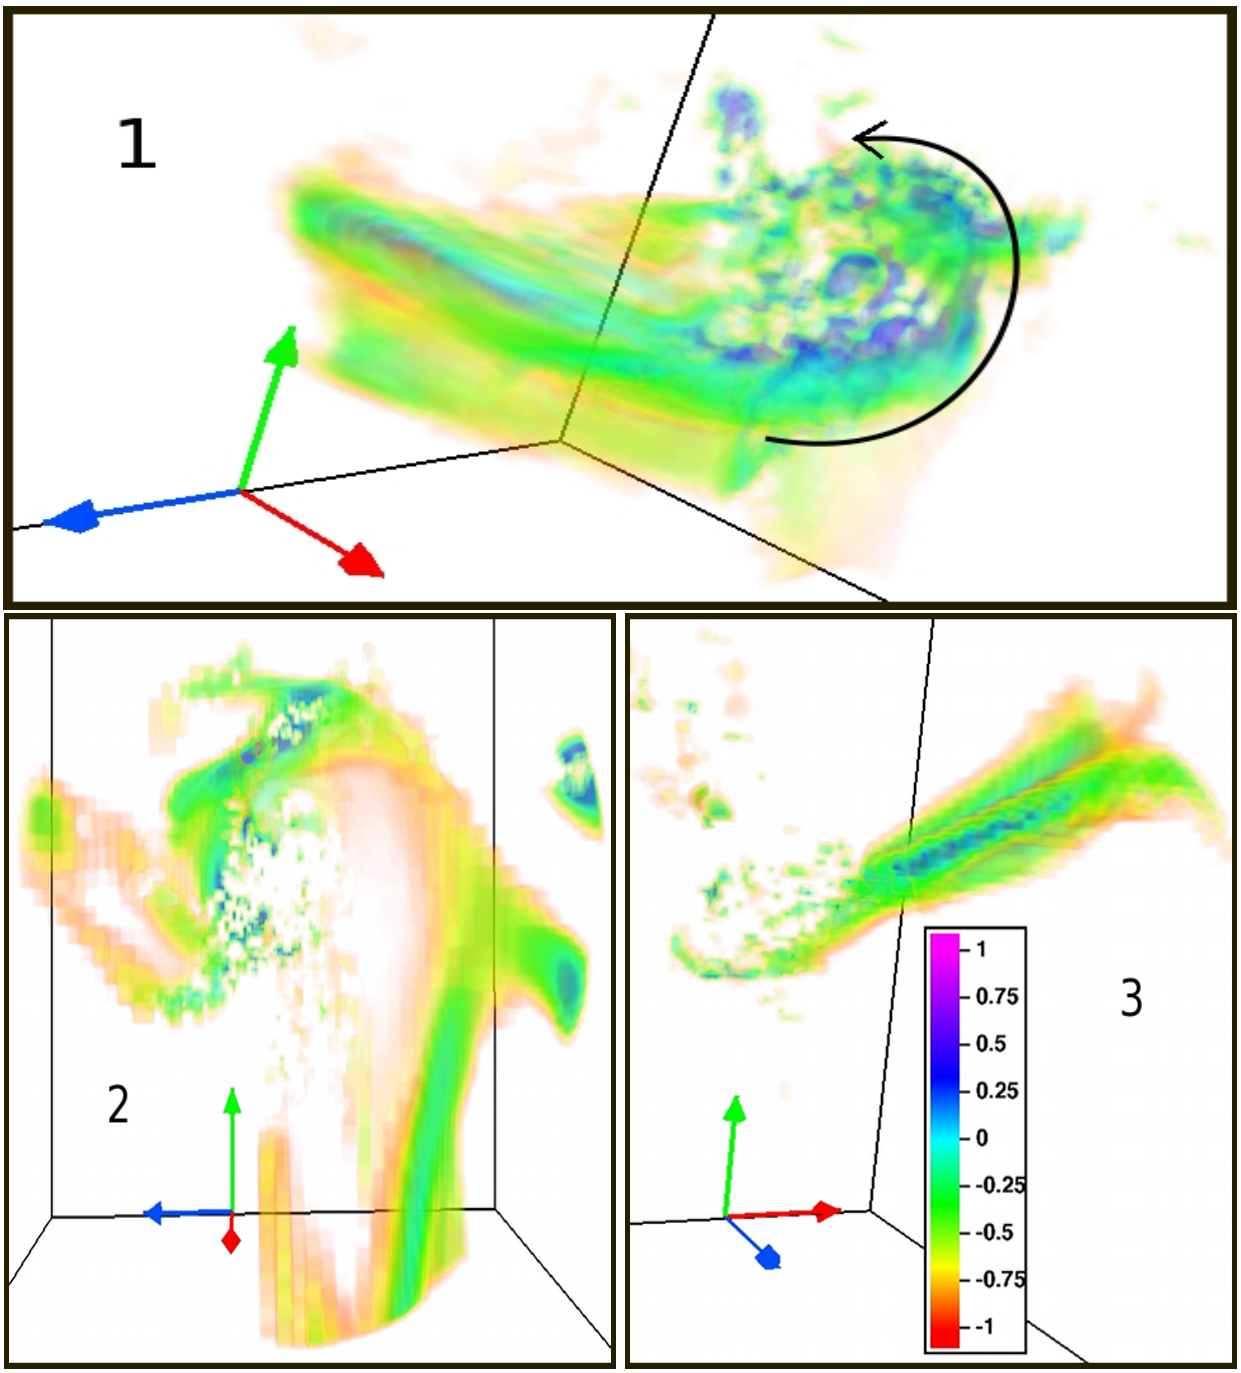

Figure 5.7: These VAPOR images illustrate the trajectories of each gas filament on $\mathcal{G}_{0.5}$ ($r_{\rm vir} \sim 7.5\,{\rm kpc}$) at $z \sim 8$ in the CO run with $\ell = 17$ ($\Delta x \sim 12\,{\rm pc}$), as identified by the colour propagation technique. Each panel corresponds to the updated versions of the purple, yellow and red filaments associated with regions 1, 2 and 3 respectively in the left panel of Fig. 5.3, and follows the same density colour scaling as in previous figures. The black curved arrow in the top panel indicates the gas flow direction of the purple filament.

- only small regions are masked from the density field, leaving the filament trajectories largely unperturbed, and;
- the highest density material that is likely to dominate the satellite angular momentum budget is accounted for.

The remainder of this subsection is now devoted to describing the technique that has been adopted to identify both the stellar and the dark matter satellites. For simplicity, reference is made to dark matter satellites in what follows, but the same method applies to stellar clumps detected by the clump-finder.

A substructure i that is classed as a satellite of the host at a given time-step t_i, satisfies both of the following criteria:

- it partially or fully infringes the host's virial sphere at t_i, and;
- it has previously been a halo.

The former condition ensures that only substructures within the host's virial region are considered, and is hence satisfied should any region within satellite i's virial radius also be contained by the host's virial sphere. If it is, then i is flagged as a satellite candidate, otherwise it is skipped. Assuming i is a satellite candidate, the latter condition involves computing its merger tree, and if at any time-step along its main branch one of the main progenitors becomes a halo, then i is flagged as a satellite.

In order to compute the main branch along i's merger tree, the progenitor that donates the most mass to i is identified: this progenitor, j, is called the 'main progenitor', and the link from i to j is referred to as the 'main branch' (following Section 3.2.2). Object j may be a halo (and it would then be concluded that i is a satellite), or it may be a subhalo (in which case the main progenitor of j, k, is found and the same set of criteria is applied to k). Another possibility is that j is diffuse background material, not bound to any halo or subhalo: in this scenario, i never originated from a halo and so would not be flagged as a satellite. A further possibility is that each of the main progenitors are bound objects but were never haloes, and that i was born as a subhalo—i would not be identified as a satellite if this were true either.

To summarize, the satellite-finding algorithm adopted in this study imposes the condition that in order for a subhalo to be a satellite at time t, it must have been a field halo at some point earlier

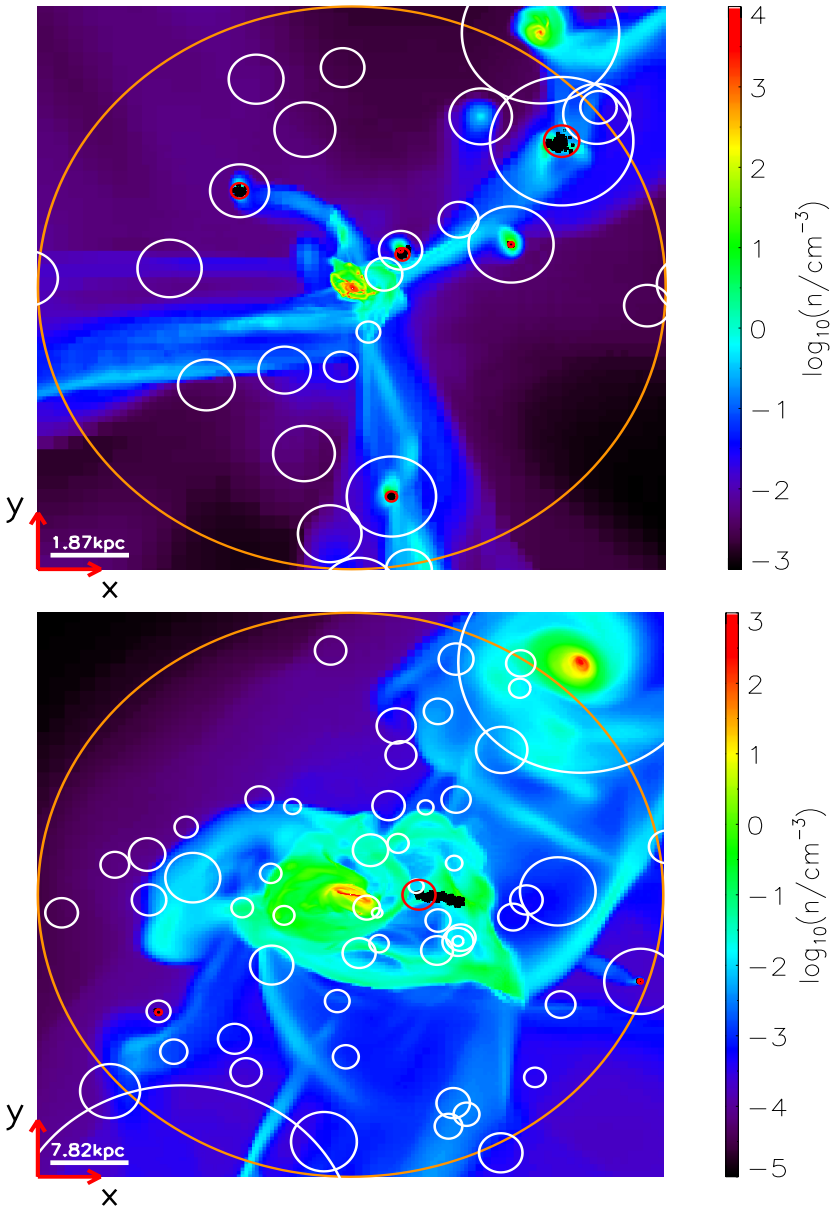

Figure 5.8: Projections along the z direction of all the gas within a box of sidelength equal to $2r_{\rm vir}$, corresponding to the diameter of the orange circles. The centre of each image coincides with the centre of the stellar clump of the host galaxy at $z \sim 8$ (top) and $z = 3$ (bottom), with $\ell = 16$ ($\Delta x \sim 22\,$pc) and $\ell = 15$ ($\Delta x \sim 95\,$pc). The virial regions of the stellar and dark halo satellites of the main host, as identified by the satellite-finding algorithm, are shown by the red and white circles respectively, with each circle centred at the densest member star identified as part of the satellite. The constituent satellite stars have been detected by the clump-finder, and are shown as small filled black circles. The white horizontal bars and vertical colour bars indicate the length scale (in physical kpc) and the density colour scale (in dimensionless units of $\log_{10}$ number of hydrogen atoms per cubic centimetre) respectively. Note that the large stellar satellites in the upper right corners of both panels have not been identified as luminous satellites because their virial regions do not infringe the host's virial sphere at these epochs.

in its merger history. Note that this method allows objects to share the same main progenitor and hence maximizes the number of possible satellites, resulting in the purest filamentary angular momentum signals once the central satellite galaxy regions have been masked. Fig. 5.8 shows z projections of the gas density within the host virial region (orange circle) at $z \sim 8$ (upper panel) and $z = 3$ (lower panel), and includes all the dark matter satellites (white circles) and stellar satellites (red circles) that have been found by applying the technique described in this subsection. Each of these circles is centred on the densest point of the object it represents, as identified by the halo and stellar clump-finder. The filled black circles correspond to the constituent stellar particles of each stellar satellite, which are not necessarily constrained to their host's virial region, as illustrated for the elongated satellite in the lower panel. This occurs because some of the member particles reside in the outer ellipsoidal shells that violate the virial accuracy condition (Section 1.2.2).

Fig. 5.8 hence confirms that all of the apparent luminous satellites are captured by the satellite-finding algorithm at the two extrema of the range in epochs examined in this study. Clearly the number of stellar satellites is reduced by $z = 3$, and those that are detected often have depleted gas reservoirs and stretched stellar morphologies, indicative of objects experiencing tidal stripping.

5.5 Methods III: Angular momentum measurement techniques

With the satellites and filaments now resolved at each time-step across the full host virial region, the framework for performing angular momentum measurements of material as it crosses the inner disk region is in place.

5.5.1 Calculating angular momentum in the general case

Unless stated otherwise, the total angular momentum $\vec{J}$ has been computed by using the following expression:

$$\vec{J} = \sum_i m_i(\vec{r}_i - \vec{r}_c) \times (\vec{v}_i - \vec{v}_c), \tag{5.7}$$

where i refers to either a stellar particle, a dark matter particle or a gas cell of mass m_i at location $\vec{r}_i$ with velocity $\vec{v}_i$. Angular momentum has hence been measured in the rest frame of the object of interest in this study—the gas disk—whose position and velocity vectors $\vec{r}_c$ and $\vec{v}_c$ are defined in Section 5.3.3.

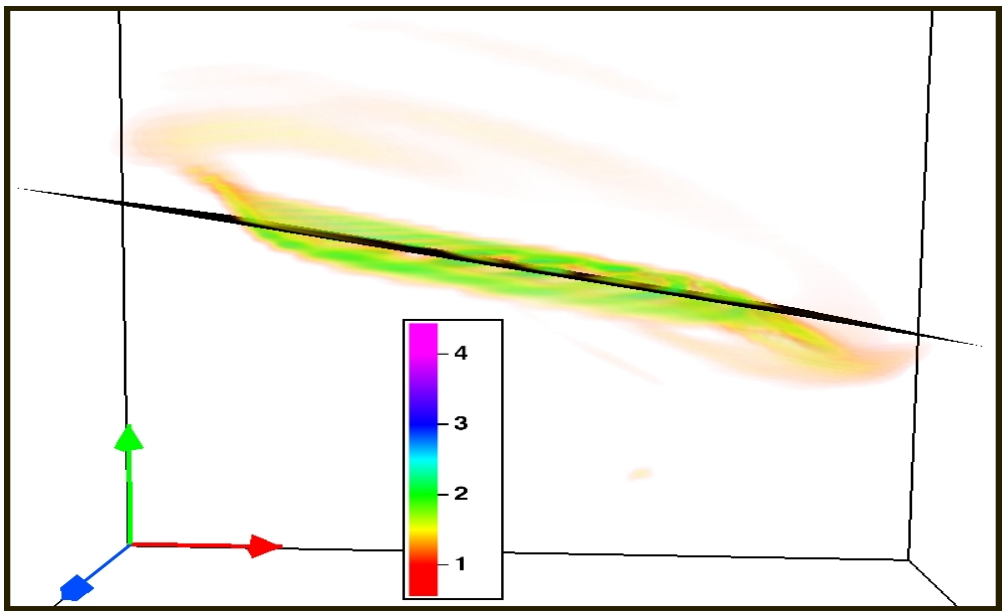

Figure 5.9: A VAPOR visualization in density of the gas disk refined to $\ell = 18$ ($\Delta x \sim 12\,\mathrm{pc}$) at $z = 3$. All of the gas within $0.1r_{\mathrm{vir}}$ that does not belong to filaments or the dense central components of luminous satellites is assumed to be part of the disk. The black slice shows the plane of disk rotation found by computing the disk angular momentum $\vec{J}_d$ within $\mathcal{G}_{0.1}$ (equation 5.7). The scaled sidelength of the box in this image is approximately 6 kpc in physical coordinates, and the units of the colour bar and the set of orthogonal axes are the same as in Figs 5.3 and 5.7.

5.5.2 Obtaining the equation of the gas disk's plane of rotation

It is assumed that the rotation of the gas disk is confined to a single plane centred on $\vec{r}_c$ at all times, and an estimate for the equation of this plane has been made by using equation (5.7) to compute the disk angular momentum $\vec{J}_d$. All of the gas within $\mathcal{G}_{0.1}$ associated with the luminous satellites (identified using the satellite-finding technique) and the filaments (identified using the nT criteria) has been flagged, and the remaining gas in this region is assigned to the disk. Fig. 5.9 shows a density image of the gas disk at $z = 3$ in the CO run, and overplots the plane found by adopting the above disk definition. The excellent agreement between the computed disk plane and its apparent orientation in Fig. 5.9 provides confirmation that both the estimates of $\vec{r}_c$ and $\vec{v}_c$ in equation (5.7) and the adopted disk definition, are reliable.

5.5.3 Quantifying the angular momentum transported to the disk

The material channelled inward along filaments f in the simulation gets compressed once it settles into the disk d, and star formation ensues when the gas is sufficiently dense. Mass and angular momentum transported by these flows is hence assumed to be shared amongst the central baryons

b, which include both the gas in the disk and the stars that form from the disk's fragmentation:

$$m_f \rightarrow m_b = m_d + m_\star \tag{5.8}$$

$$\vec{J}_f \rightarrow \vec{J}_b = \vec{J}_d + \vec{J}_\star. \tag{5.9}$$

At a given time-step, a list of identifiers of every particle that has ever been resolved as a member constituent of a satellite, as detected by the satellite-finding algorithm over the Milky Way host's virial region, is updated. Hence $\vec{J}_\star$ has been computed at each epoch by using equation (5.7) and including each particle within $0.1r_{\rm vir}$ that does not belong to this list. These non-satellite stars correspond to bulge and disk stars, but the former component has a weak angular momentum modulus (e.g. van den Bosch et al. 2002) and so the stellar disk signal is likely to dominate $\vec{J}_\star$. Hence m_b and $\vec{J}_b$ are made in reference to the mass and angular momentum of the 'baryonic disk' in what follows, which will often be abbreviated to 'the disk'. By excluding contributions from satellite stars, equations (5.8) and (5.9) focus on in-situ disk formation.

It is instructive to compare $\vec{J}_b$ at each time-step with the accumulated angular momentum of the satellite gas and cold filament gas within $0.1r_{\rm vir}$, projected along the direction of the disk's axis of rotation. After all, these components are thought to fuel the disk's angular momentum budget. It is also interesting to consider the fraction of $\vec{J}_b$ acquired via a hot mode of gas accretion, as this enables direct comparisons with the cold mode signals. The hot gas phase has hence been resolved by selecting all of the gas cells with a temperature above the imposed upper temperature limit of the cold phase (i.e. $T > 2 \times 10^5\,{\rm K}$). Hot gas should not be regarded as a separate component, however, because these cells are included in the satellite and disk signals by construction (see Section 5.5.2).

The hot and cold phase contributions to m_b and $\vec{J}_b$

Estimates of the angular momentum accumulated onto the disk via the accretion of hot and cold gas have been made by time-integrating the instantaneous flow rates of these components across a thin spherical shell located on the periphery of the disk region. This is a suitable method to employ because the inflowing gas in the simulation crosses the shell before settling onto the disk. The flow rates have been computed by using the technique described in Powell et al. (2011), and for simplicity, reference is made to the filament gas phase in the descriptions below (but the same

equations apply to the hot mode):

$$\frac{\mathrm{d}\vec{J}}{\mathrm{d}t} = \frac{\sum_i \rho_i v_{r,i} \left[(\vec{r}_i - \vec{r}_c) \times (\vec{v}_i - \vec{v}_c)\right] (\Delta x)^3}{\sum_j (\Delta x)^3} 4\pi r_m^2 \tag{5.10}$$

$$\frac{\mathrm{d}M}{\mathrm{d}t} = \frac{\sum_i \rho_i v_{r,i} (\Delta x)^3}{\sum_j (\Delta x)^3} 4\pi r_m^2. \tag{5.11}$$

The sum indexed by j is performed over all of the cells in the shell, irrespective of their gas phase, and $\vec{r}_c$ and $\vec{v}_c$ are the same as in equation (5.7). The variables $\vec{r}_i, \vec{v}_i$ and ρ_i correspond respectively to the position, velocity and density of the i^{th} cell satisfying the relevant imposed condition, which in this case is the requirement that i is a filament cell. Equations (5.10) and (5.11) hence yield the net radial flux transported across the spherical shell towards (negative sign) or away from (positive sign) the disk, where the radial velocity of each cell is given by:

$$v_{r,i} = (\vec{v}_i - \vec{v}_c) \cdot \frac{\vec{r}_i - \vec{r}_c}{|\vec{r}_i - \vec{r}_c|}. \tag{5.12}$$

It follows that the total angular momentum (and mass) transferred towards the disk by each filament over the time interval $\Delta t = t - t_i$ is:

$$\Delta\vec{J}_{\mathrm{flux}}(t) = \int_{t_i}^{t} \frac{\mathrm{d}\vec{J}}{\mathrm{d}t}\,\mathrm{d}t \tag{5.13}$$

$$\simeq \left(\left.\frac{\mathrm{d}\vec{J}}{\mathrm{d}t}\right|_t + \left.\frac{\mathrm{d}\vec{J}}{\mathrm{d}t}\right|_{t_i} \right) \frac{\Delta t}{2}, \tag{5.14}$$

where the latter step has performed a simple numerical trapezium integration, approximating the form of the unresolved flow rates between time-steps by a linear interpolation. The signal found using the 'spherical flux technique' of equation (5.14) at t is then accumulated with the corresponding signals at earlier times.

The summations indexed by i in equations (5.10) and (5.11) include all of the filament cells whose centres lie within a shell of inner, outer and mid radius $r_{\mathrm{in}} = r_m - \Delta r/2$, $r_{\mathrm{out}} = r_m + \Delta r/2$ and $r_m = 0.1 r_{\mathrm{vir}}$, respectively. The thickness of the shell Δr is a free parameter and has been fixed to 1% of r_{vir} (i.e. one tenth of the $0.1 r_{\mathrm{vir}}$ radius) at each time-step, as this choice strikes a good compromise between negligible Poisson error and the validity of the thin shell approximation (the shells typically included $\sim 2 \times 10^5$ cells at the highest redshifts around $z \sim 9$ and $\sim 2 \times 10^7$ cells at the lower redshifts around $z \sim 3$, where the shell was physically larger). Similar flux

measurements were found for shell thicknesses in the range ∼0.1–2.5% of $r_{\rm vir}$, which implies that there is some flexibility in the value assigned to Δr. A larger grid than $\mathcal{G}_{0.1}$ is hence required in order to compensate for the half of the shell that spills beyond the $0.1r_{\rm vir}$ region. The filaments have therefore been constructed over the $0.15r_{\rm vir}$ region according to step 3 of the colour propagation scheme described in Section 5.4.2. This particular grid length is chosen because the physical width of the cells within $\mathcal{G}_{0.15}$ is still at the 12 pc limit for every time-step, and so the extra grid length makes no difference to the accuracy of the mass and angular momentum computations.

It is clear from the above that a decision has been made to represent the quantities r_m and Δr as fractions of the host's virial radius, which grows with time in the simulation. The spherical flux method hence sweeps up gas by construction, and a simple trapezium integration of the flux signals across the shell could miss this component. An alternative way of thinking about this is to imagine that the filament gas circulates around the disk region over a given time interval in the θ and ϕ directions, with zero velocity in the radial direction. Equation (5.10) predicts that the angular momentum transported towards the disk is zero in this scenario, yet this filament gas is included within the $0.1r_{\rm vir}$ sphere as the sphere grows over the interval, and will probably accrete onto the disk at some later time. The amount of swept up gas will scale with the growth rate of the host's virial radius, which is likely to be quite slow between successive time intervals. Nonetheless, an accurate method would measure both the gas whose inward trajectory causes it to traverse the central region and the gas that is swept up due to the growth of the central region, and Fig. 5.10 illustrates one such method. The routine operates on a simple set of principles:

1. fix the sphere radius $r(t_1)$ over a given time interval $\Delta t_{21} (\equiv t_2 - t_1)$;

2. apply equations (5.10) and (5.11) to these identical spheres at the two time-steps, and perform the trapezium integration in equation (5.14), and;

3. combine the above spherical flux signals with the corresponding instantaneous angular momentum and mass of all the filament cells measured at the later time-step t_2 within the small region (shaded area at t_2) that is traced out due to the growth of the disk component.

By performing an instantaneous computation at the later time-step, this approach assumes that the swept up gas is static and that its measured mass and angular momentum is quickly deposited onto the disk. Despite not being strictly in concordance with the Lagrangian nature of the system, this is likely to be robust given that the associated errors are due to the time resolution of the

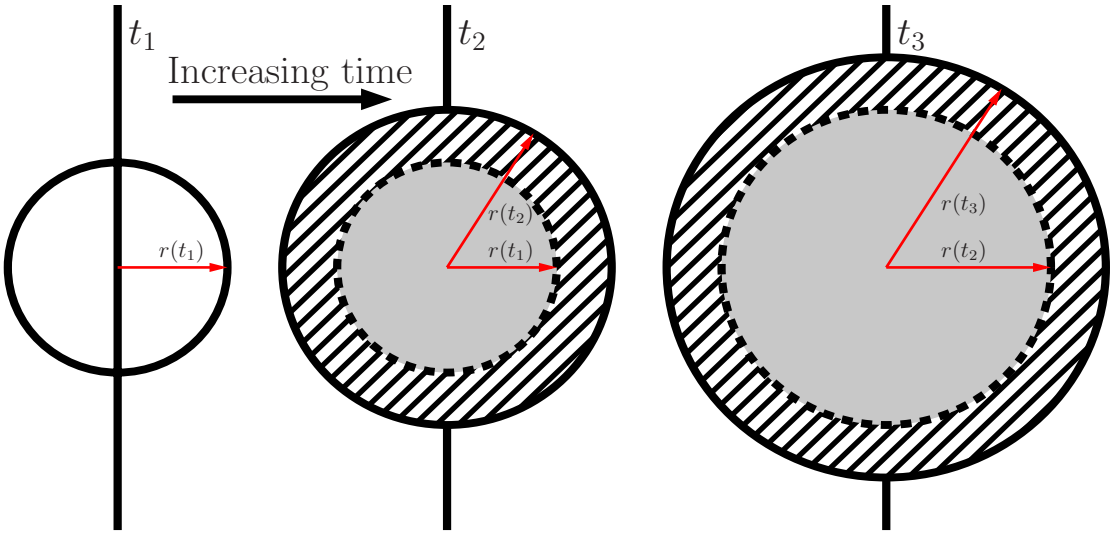

Figure 5.10: This diagram illustrates the technique that has been used to estimate the angular momentum and mass of filament gas and hot gas deposited onto the disk. Three consecutive time-steps are shown, with $t_3 > t_2 > t_1$. The radius of the central $0.1r_{\rm vir}$ region at each time-step t_j is abbreviated to $r(t_j)$, and the shaded spherical rim at t_2 and t_3 indicates the growth of this central region over the Δt_{21} and Δt_{32} time intervals. There are two angular momentum and mass recordings for $t_j > t_1$. Using $t_j = t_2$ as an example, the first is due to inflowing gas, which is measured using equations (5.10) and (5.11) on two stationary spheres of equal radius $r(t_1)$. The second signal arises from the material that is swept up by the outwardly propagating sphere (shaded spherical region at t_2), and is calculated instantaneously at t_2. For t_3, the radius at t_2 is fixed and the same technique is adopted. Note that contrary to the simplified scenario above, the spheres are centred at different spatial locations from time-step to time-step in the simulation due to the bulk motion of the centre of disk rotation (see Fig. 5.4).

simulation (an estimate of the importance of possible time lags between material crossing $0.1r_{\rm vir}$ and being accreted onto the disk is presented later in this section). The integrated quantities at time t hence correspond to:

$$M_{\rm T}(t) = \sum_{t_i}^{t} [\Delta M_{\rm flux}(t'_j) + \Delta M_{\rm stat}(t'_j)] \quad (5.15)$$

$$\vec{J}_{\rm T}(t) = \sum_{t_i}^{t} [\Delta \vec{J}_{\rm flux}(t'_j) + \Delta \vec{J}_{\rm stat}(t'_j)], \quad (5.16)$$

where the first and second terms refer to the signals measured using the spherical flux method and the 'static shell' approximation, respectively. The ratio of these terms in equation (5.15) satisfies $\Delta M_{\rm flux}/\Delta M_{\rm stat} \gtrsim 10$ at every time-step in the simulation, which indicates that the amount of gas swept up over a given time interval is somewhat negligible. Nevertheless, the static contributions are included for completeness. The initial time-step in equations (5.15) and (5.16) is denoted by t_i

and the sum is performed over all discrete time-steps $t_i \leq t'_j \leq t$, yielding disk projections:

$$J_p(t) = \vec{J}_{\rm T}(t) \cdot \hat{J}_b(t). \tag{5.17}$$

An estimate of the angular momentum transported by the hot gas and filament gas that is transferred to the disk can now hence be computed.

The satellite contributions to m_b and $\vec{J}_b$

Luminous satellites in the simulation merge with the host galaxy across its lifetime, depositing their stars and gas into the central region. Satellite stars are ignored in this study following the baryonic disk definition in equation (5.8), and so the angular momentum (and mass) contributions from satellites represent the total gas component within their virial regions above the upper filament density threshold of 10 hydrogen atoms per cubic centimetre.

The following simple prescription, which seeks to pinpoint mergers with the disk, has been used to measure the instantaneous angular momentum of satellites before they are accreted. The algorithm begins by finding all the luminous satellites (according to the method described in Section 5.4.3) that infringe the $0.1r_{\rm vir}$ and $r_{\rm vir}$ regions at t and $t + \Delta t$ respectively. Each satellite identification number at t is then compared with the list of satellite identifiers at $t + \Delta t$. Without a match, a given satellite has either fully merged with the disk or has been stripped of mass upon crossing the central region and has consequently passed below the minimum resolution limit of 100 particles imposed by the clump-finder. The merger has hence already taken place or is due to take place very shortly, and so the satellite signals at t are measured. If, however, the satellite is resolved at $t + \Delta t$ (it may have moved closer to the disk or outside the central region during the merger process), then it has evidently not yet been accreted, and so no signal is recorded. The gas residing within the satellite virial region above the upper filament density threshold is hence included in the cumulative signal at t provided the satellite is:

1. encroaching the $0.1r_{\rm vir}$ region at t, and;

2. unresolved at $t + \Delta t$.

The satellite gas mass M_s and angular momentum $\vec{J}_s$ deposited within $0.1r_{\rm vir}$ at t are subsequently computed using direct summation and equation (5.7) respectively, and are coupled with the totals

from the previous epochs in an analogous fashion to equations (5.15) and (5.16):

$$M^s_{\rm T}(\leq 0.1r_{\rm vir}(t),t) = \sum_{t_i}^{t} M_s(\leq 0.1r_{\rm vir}(t'_j),t'_j) \quad (5.18)$$

$$\vec{J}^s_{\rm T}(\leq 0.1r_{\rm vir}(t),t) = \sum_{t_i}^{t} \vec{J}_s(\leq 0.1r_{\rm vir}(t'_j),t'_j), \quad (5.19)$$

yielding a disk-projected component:

$$J^s_p(t) = \vec{J}^s_{\rm T}(t)\cdot\hat{J}_b(t). \quad (5.20)$$

Note that M_s and $\vec{J}_s$ are calculated by integrating over the $0.1r_{\rm vir}$ region rather than time-integrating the net instantaneous accretion rate across an outwardly propagating spherical shell, because unlike the filaments that continuously stream radially inward towards the disk, the satellite mergers are discrete burst events that can easily pass undetected through the shell over a given time interval. Applying equation (5.11) to the satellites could hence yield zero mass accretion events.

The explicit time dependences of $J_p, \vec{J}_{\rm T}$ and $\vec{J}_b$ are henceforth dropped for brevity.

Assumptions

By comparing the projected angular momentum of gas measured using equations (5.17) and (5.20) with the disk signal found using equation (5.9) at each epoch, two implicit assumptions are being made:

- the angular momentum of the material passing through the spherical shell at $0.1r_{\rm vir}$ as measured at t reflects the amount deposited onto the disk at the time of contact $t+\delta t$ (i.e. δt is assumed to be smaller than a given time interval in the simulation Δt), and;
- the time integrated filament and hot gas flux signals are well approximated by a trapezium integration.

The first condition draws attention to possible time lags between material crossing the shell and merging onto the disk, and its importance can be estimated by comparing the free-fall time $t_{\rm ff}$ at $0.1r_{\rm vir}$ with the separation between time intervals in the simulation Δt (for a defense of the view that filaments are indeed well approximated by radial free-fall, see Rosdahl & Blaizot 2012). If $t_{\rm ff} >> \Delta t$, then it is likely that mass piles up or follows a more complex trajectory before merging

with the disk, which would require its angular momentum to be recalculated at the time of accretion onto the disk. In the CO run, it was found that the range in ratio of these two timescales satisfied $0.1 \lesssim t_{\rm ff}/\Delta t \lesssim 1$. Taking the final output at $z = 3$ as an example, where $\Delta t \sim 32\,\mathrm{Myrs}$:

$$t_{\rm ff} = \sqrt{\frac{3\pi}{32G\rho}} = \sqrt{\frac{4\pi^2(0.1r_{\rm vir})^3}{32GM_{\rm T}(\leq 0.1r_{\rm vir})}} \sim 16\,\mathrm{Myrs}. \tag{5.21}$$

The virial radius at this epoch is $r_{\rm vir} \sim 31\,\mathrm{kpc}$ and the total mass (stars, dark matter and gas) enclosed in the disk region is $M_{\rm T}(\leq 0.1r_{\rm vir}) \sim 3 \times 10^{10}\,\mathrm{M}_{\odot}$. Given that $t_{\rm ff} < \Delta t$ for every time-step, it is likely that the change in system configuration during the interval does not yield a significantly different angular momentum signal transferred to the disk compared with the signal measured at $0.1r_{\rm vir}$, thereby justifying the first assumption.

The second condition questions the reliability of the integration step in the flux method (equation 5.14) and examines whether a linear interpolation provides an accurate representation of the unresolved flux signals across the time intervals. This issue is addressed in the following chapter, where the predictions of equation (5.15) for the integrated disk mass are directly compared with the measured disk mass from equation (5.8) at each epoch.

5.6 Summary

This chapter has tackled the computational aspect to monitoring the flow of angular momentum onto a Milky Way type disk at $z \geq 3$. A colour propagation scheme using Lagrangian tracer particles (Section 5.4.2) has been developed to identify particles within individual filaments spanning a region of length of $2r_{\rm vir}$ from the disk's centre at each epoch. The method assigns one of three colours to the tracer particles by using a nearest neighbours scheme, and subsequently locates the 'pure' cells within $\mathcal{G}_{2.0}$ that contain tracer particles of identical colour. Individual filament trajectories are then grown around these pure sites. A satellite finding algorithm has also been presented, which operates under the principle that a given satellite substructure of the main host at time t has at some earlier time been a separate field galaxy of its own. With the filaments and satellites resolved, the aim of the chapter has been to provide a technique capable of measuring their contributions to the host's disk angular momentum budget at each epoch. For the filament and hot gas components, the mass and angular momentum transported to the central disk has been computed by recording both the net inward flux across an outwardly propagating spherical shell on the disk's edge at $0.1r_{\rm vir}$, and

the contribution from the material swept up by the shell's outward radial motion. Satellite signals have been measured by adopting an alternative technique that pinpoints mergers with the disk and records their angular momentum at the final pre-accretion stage.

These angular momentum measurement schemes are now ready to be applied to each time output from the NUT CO simulation.

Chapter 6

Angular momentum transfer to a Milky Way disk: results

The techniques described in the previous chapter are now implemented in an attempt to examine whether: (i) filaments dominate $\vec{J}_b$ at high redshift, and; (ii) there are any special episodes of disk angular momentum acquisition.

Fig. 6.1 plots the redshift evolution of the integrated net inflow of mass towards the disk in the upper panel, and the projection of the integrated angular momentum along the disk's direction in the lower panel. All of the measurements have been performed at or within the $0.1r_{\rm vir}$ boundary as described in Section 5.5. Baryonic disk signals correspond to the solid black lines, while the satellite gas (equations 5.18 and 5.20) and hot gas (equations 5.15 and 5.17) signals are given by the green and light blue lines respectively, with the filled circles in the lower panel indicating negative cumulative angular momentum projections. Each filament is distinguished by its unique colour, and the dashed dark blue line represents their sum at each epoch (equations 5.15 and 5.17). The solid dark blue line shows the signals of all of the filament cells that have been found by applying the nT criteria at each redshift, and serves as a consistency check for the individual filament computations. The dashed and solid dark blue lines are indistinguishable from one another in both panels, which implies that the vast majority of the filaments cells have been accounted for, and reinforces the high level of accuracy of the tracer colouring technique (see Fig. 5.6). Finally, it is expected that the total gas angular momentum and mass contributions from the filament and satellite gas phases are equal to the baryonic disk signals at each epoch. The dashed black lines in Fig. 6.1 test this hypothesis and should be compared with the corresponding solid black lines.

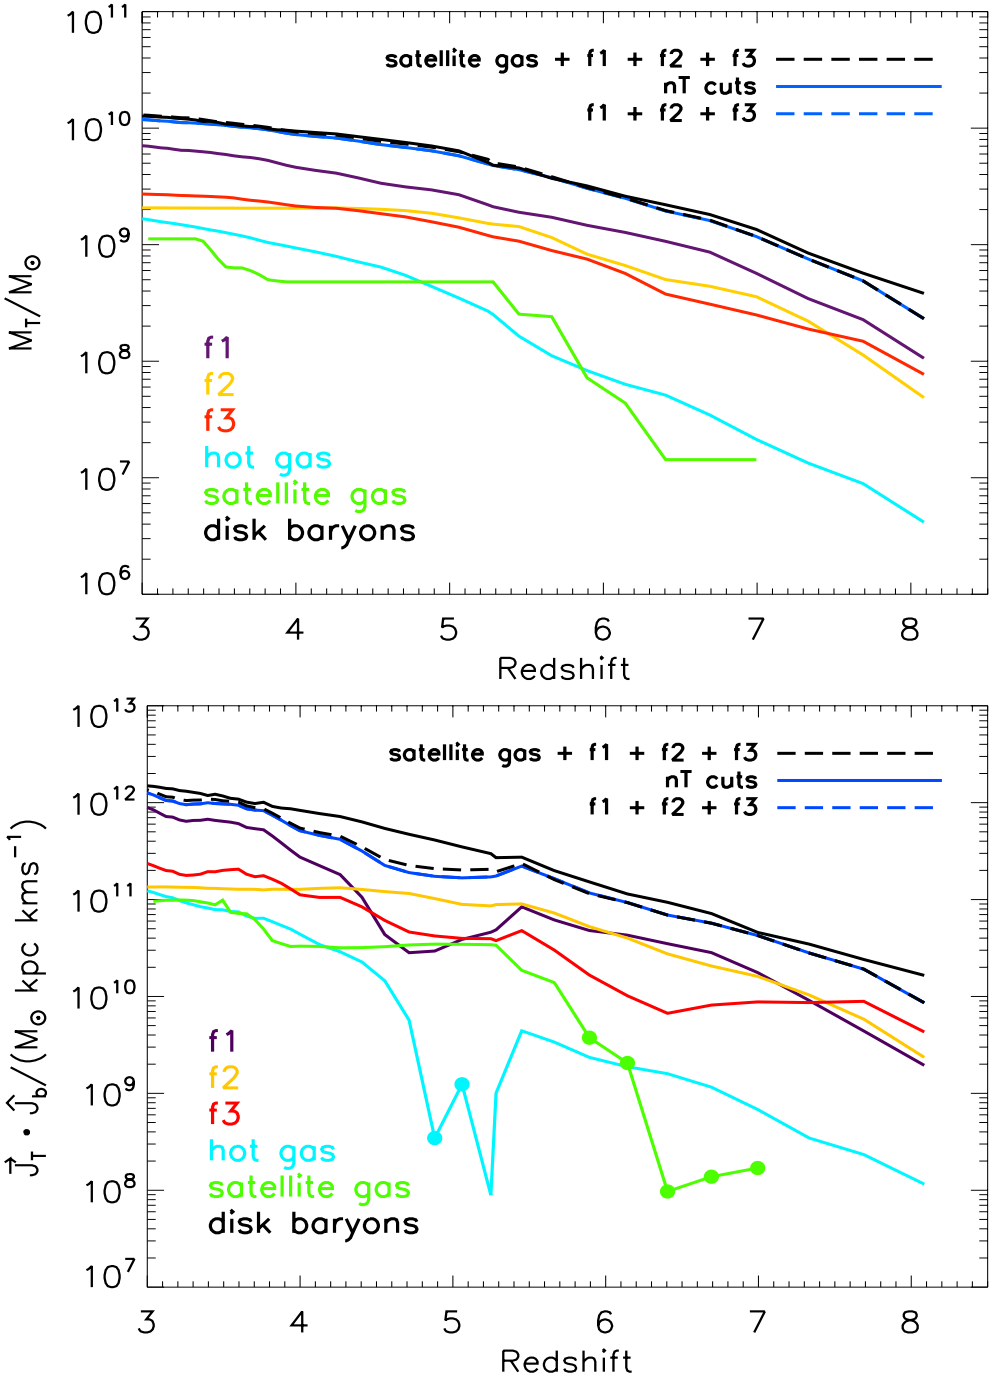

Figure 6.1: The relative importance of hot gas ($T > 2 \times 10^5$ K), cold filament gas and satellite gas in establishing the mass and angular momentum budget of the central disk at high redshift. $M_{\rm T}$ (upper panel) and $\vec{J}_{\rm T}$ (lower panel) have been accumulated with their corresponding signals from all prior epochs. The projection of $\vec{J}_{\rm T}$ along the unit disk direction $\hat{J}_b$ then yields $J_p(z)$ in equation (5.17) for the filament and hot gas phases. A similar computation has been performed for the satellites, except that equations (5.18) and (5.19) have been used in place of equations (5.15) and (5.16), with the filled circles corresponding to negative $J_p(z)$ contributions. The dashed dark blue lines sum over the individual filament signals at each epoch, and cannot be distinguished from the solid dark blue lines, which show the accumulated signals for all of the filament gas found by applying the nT cuts. The mass and angular momentum of the disk baryons defined by equations (5.8) and (5.9) have been computed at every time-step and are given by the solid black lines.

6.1 Redefining the system's initial state

The mass and angular momentum transported by the various galaxy components are first recorded at $z \sim 8$ in Fig. 6.1. It should be noted that this epoch does not correspond to the initial state of the galaxy system: the MSM halo-finder first resolved the Milky Way host's dark matter halo with 40 particles and the central luminous galaxy's stellar clump with 100 particles, at $z \sim 13$. The system was highly unsettled at this epoch, however, and so application of the tracer colouring algorithm in Fig. 6.1 was postponed until $z_{\rm in} \sim 10$, which was found to be the earliest time-step that displayed a smooth, distinct filament configuration in the CO run (shown in the left panel of Fig. 5.3). For reasons highlighted in Section 5.4.2, there is a two time-step delay between a given starting epoch and the epoch at which the individual filaments can first be resolved. This therefore pushes the starting point forward to $z'_{\rm in} \sim 9$, and the subsequent trapezium integration of equations (5.10) and (5.11) over the next time interval yields their first reliable estimates at $z \sim 8.5$, three time intervals after $z_{\rm in} \sim 10$. The first change in the measured angular momentum and mass signals with respect to the starting signals at $z \sim 8.5$ can hence be found for each filament by performing another trapezium over the following interval, which explains why the curves in Fig. 6.1 do not appear until $z \sim 8$. Hence in order to preserve consistency with the individual filaments, the initial state of the entire system is chosen to be $z'_{\rm in} \sim 8.5$, and the disk, filament, hot gas and satellite signals at $z'_{\rm in}$ are subtracted from their counterparts at $z < z'_{\rm in}$.

In summary, the $0.1r_{\rm vir}$ sphere is modelled as being empty at $z'_{\rm in}$ in order to test mass conservation within the inner regions of the luminous galaxy and the significance of cold flows in the build-up of disk angular momentum at high redshift. Even though the disk has probably existed since the first epoch of star formation in the central region around $z_1 \sim 20$, the approach adopted above of shifting the origin forward to $z'_{\rm in}$ is justified given: a) the negligible amount of time elapsed between z_1 and $z'_{\rm in}$ compared with z_1 and the final output at $z = 3$, and; b) the desire to maintain consistency with the starting point of the tracer colouring algorithm.

6.2 The mass growth of the central disk component

It appears that the technical details surrounding the choice of starting point are somewhat negligible because the upper panel of Fig. 6.1 indicates that the baryonic disk experiences rapid growth at high redshift, where its mass increases by a factor of $\sim$8 between $z \sim 8$ and $z \sim 6$. This growth

is slower at later times, with a factor of ∼2 increase in disk mass between $z \sim 5$ and $z = 3$. The periods of upturn in the green curves correspond to discrete luminous satellite merger events, which deposit ∼1% of the disk's mass in the form of gas at early times ($z \sim 7$) and ∼10% at late times ($z = 3$). Satellites do not therefore appreciably affect the mass budget of the disk in the redshift range $3 \lesssim z \lesssim 8$. The amount of hot gas accreted onto the disk increases at a very stable rate owing to the long cooling times associated with a metal-poor ISM in the CO run, but by $z = 3$ the hot phase only accounts for about one tenth of the disk's accumulated mass. This phase hence plays a similarly subdominant role to the luminous satellites. Evidently the mass budget of the disk is largely controlled by the behaviour of filaments, and Fig. 6.1 suggests that it is the purple filament corresponding to region 1 in the left panel of Fig. 5.3 that carries the most mass to the disk, typically transporting at least twice as much mass at a given epoch compared with the yellow and red filaments, whose contributions flow in approximately equal measure until $z \sim 5.5$, whereupon the purple and yellow filaments merge, with the former surviving (Fig. 5.4). The disappearance of the yellow filament is evident from the plateau in its accumulated mass evolution for $z \lesssim 5.5$.

By comparing the dashed black line with the solid black line in the upper panel of Fig. 6.1, it can be seen that the spherical flux method (which dominates the signal in equation 5.15), designed to predict the net radial inflow of gas accreted onto the disk, conserves mass to high precision across the disk's evolution. This is remarkable given that the predicted flux rates across the shell could, in principle, vary quite dramatically between time intervals and hence not be well approximated by a trapezium integration. For example, it is possible that fast moving clumps of material that are accreted onto the disk pass through the thin shell undetected, which might lead one into thinking that the flux method has a systematic tendency to underestimate the accumulated mass signal. The forepart of this argument appears to hold for the very first recorded mass signal at $z \sim 8$, where the inflow of gas is just under a factor of 2 below the disk signal in place at that epoch. Clearly this is only a phenomenon that occurs at high redshift, as it quickly becomes unimportant when the disk experiences rapid smooth growth of cold gas over the subsequent time interval, which is captured by the flux method. Even so, for completeness, it is interesting to speculate on this high redshift discrepancy. One might propose that by increasing the thickness of the spherical shell from $\Delta r = 1\%$ of $r_{\rm vir}$ to a larger value, there is a higher probability of capturing this apparent 'missing' filament mass. Yet it has been found that even when $\Delta r \sim 10\%$ of $r_{\rm vir}$ is used at the first few time-steps, there is a negligible difference in the filament (and hence filament plus satellite

gas) signals. The issue with increasing the shell thickness is that although more filament gas is included, the number of cells within the shell also increases, and it seems that these two competing effects approximately offset one another in equation (5.11). Although relatively large increases in Δr appear to be incapable of accounting for the discrepancy with the disk signal at $z \sim 8$, the insensitivity of the flux rate to these changes at high redshift symbolizes the robustness of the spherical flux method and implies that the offset is not caused by an underestimation of the amount of mass transported radially inward by the filaments. Given the large 100 particle threshold imposed by the clump-finder and the technique adopted by the satellite-finding algorithm in Section 5.5.3 of identifying satellites within $0.1r_{\rm vir}$, there does appear to be some scope for suggesting that the luminous satellite component is underestimated at this epoch: less massive, darting satellites could be excluded from the green signals. After all, the satellite trajectories in Fig. 6.1 may appear somewhat artificial in the sense that: a) it is unlikely that there are no satellite merger events beyond $z \sim 7$, and; b) their trajectories are somewhat irregular. Both of these traits are surely associated with the large particle threshold. Although motivated by the desire to record reliable satellite mass and angular momentum detections, it is possible that this issue could be resolved by simply reducing the particle threshold of the clump-finder, but then the disk fragments into several subclumps that may be identified as satellites, despite being rotationally supported in a disk. The central disk would hence have to be recomputed and its mass is likely to be underestimated at each epoch compared with its 'true' value, as fewer stars in equation (5.8) would be associated with the non-satellite phase. The level of agreement with the integrated filament signal at $z \sim 8$ in Fig. 6.1 is also likely to be sensitive to the chosen particle threshold. Therefore, given the above difficulties with tackling the offset, and the fact that this lack of mass conservation occurs for a single high redshift interval in Fig. 6.1, it is suggested that an underestimation of the luminous satellite signal at this epoch is probably responsible for the offset, but further analysis probing this discrepancy is considered beyond the scope of this project, and deemed unnecessary. This underestimate will, of course, not only have the effect of increasing the cumulative satellite gas signal at the higher redshifts but at the lower redshifts as well, and so the resulting level of mass conservation may not quite be as strong as is currently suggested by the dashed black line in Fig 6.1. This shift is expected to have a minor effect for $z \lesssim 7.5$, however, given the rapid rate of disk growth at early times.

6.3 Which components dominate the disk's angular momentum budget?

The general trends of satellites and hot gas contributing a weak mass signal are also imprinted in their angular momentum evolution, shown in the lower panel of Fig. 6.1. At $z = 3$, the hot gas signal (like the satellite signal) accounts for roughly one tenth of the disk's accumulated angular momentum, and with the exception of its evolution between $4.5 \lesssim z \lesssim 5.5$ (discussed in more detail below), increases at a steady, uniform rate towards late times. In terms of angular momentum modulus, it therefore appears that the relative contributions of the satellite and hot gas components to the disk angular momentum budget generally trace their relative disk mass contributions. In terms of angular momentum direction, the satellite material deposited onto the disk has a tendency to be misaligned with respect to the disk's direction at high redshift, but because the signal strength is relatively weak, this mismatch may only have a negligible effect on the disk's orientation. Certainly at the lower redshifts below $z \sim 4.5$, the mass that is headed towards the disk in the satellite (and hot gas) phase appears to be more closely aligned with the disk's direction.

When integrated across all three components (solid and dashed dark blue lines), the filament angular momentum signals show a similarly strong correlation with their corresponding mass trajectories. Clearly for the redshift ranges $3 \leq z \lesssim 4$ and $5.5 \lesssim z \lesssim 7.5$, the disk's budget is driven by the transport of angular momentum from filaments. When decomposed into its individual separate trajectories, however, the cold phase shows noticeably different behaviour: (i) the yellow and red filaments transport unequal amounts of angular momentum along the disk's direction for $z \gtrsim 4.5$, in contrast with their relative mass contributions, and; (ii) the purple and yellow filaments are equally dominant for $z \gtrsim 5.5$, unlike in the mass plot where the purple filament clearly dominates at all epochs. It should also be noted that the absence of a correlation between negative satellite signals and negative filament signals implies that satellites are not necessarily sensitive to the dynamics of the filament flow, despite often streaming along these flows (e.g. Fig. 5.8).

Perhaps the most interesting feature of Fig. 6.1 is the lack of response the disk has to the sudden dip in the projected component of the purple filament's cumulative angular momentum during the filament merger phase between $4.5 \lesssim z \lesssim 5.5$. It was found that the change in the direction of the disk's axis of rotation during successive intervals across this period was negligible (of the order a few degrees), and so the sharp drop-off implies that the purple filament transports misaligned

angular momentum towards the disk during the early phase of the merger (equation 5.17). The red filament's disk-projected angular momentum contribution also abruptly changes around $z \sim 5.5$, which when combined with the purple filament's signal, yields a flattened integrated dark blue signal that reflects the net small amounts of angular momentum being transported inwards by the filaments during this merger period. Given that the baryonic disk angular momentum in stars dominates the gas contributions at lower redshifts ($|\vec{J}_\star|/|\vec{J}_{\rm gas}| \sim 2.6$ at $z \sim 5$ for example), one might expect the new population of disk stars that form from this low angular momentum filament material to yield a slower disk growth rate than is observed in the lower panel of Fig. 6.1. Remarkably, the purple filament's angular momentum direction is quite quickly restored along the disk's direction around $z \lesssim 4$, whereupon it becomes the dominant filament contributor by at least a factor of $\sim$3. Hence one naturally wonders about the physical mechanisms at play that allow the filament gas phase during the merger epoch to continue to supply the disk with all of its mass (as can be seen from the upper panel of Fig. 6.1), but not all of its angular momentum. The above results suggest that the violent merger between the purple and yellow filaments that begins at $z \sim 5.5$ induces a change in the velocity and trajectory direction of both the red and purple filaments. The purple filament then probably experiences a torque from the disk that arises due to the misalignment between its angular momentum direction and that of the disk, and this torque swings its angular momentum vector back into alignment at later times. Meanwhile, the angular momentum that the spherical flux method measures from filaments at $0.1r_{\rm vir}$ appears to differ from the filament angular momentum deposited onto the disk during this merger period, most probably because the inner dynamics within $0.1r_{\rm vir}$ are severely perturbed by the filament merger (the resultant complex inner gas trajectories appear to be imprinted in the hot gas signal too, which becomes misaligned around $z \sim 5$). Perhaps a shock on the disk boundary is responsible for removing the misaligned filament velocity components and hence filtering aligned angular momentum to the disk, driving its angular momentum growth. This mechanism is probably not so important at earlier and later epochs, as the net inflowing filament angular momentum is already well aligned with the disk's angular momentum direction.

6.4 Discussion

Perhaps the most defining feature of disks in Milky Way type galaxies with $M_\star \lesssim 3 \times 10^{10}\,{\rm M}_\odot$ is the amount of angular momentum they possess, as it is this property that establishes many of their

known correlations, encapsulated by the Tully–Fisher relation (Governato et al., 2007) and the scaling relations between velocity, metallicity, surface brightness and stellar mass (Dekel & Woo, 2003; Dekel & Birnboim, 2006). The Tully–Fisher relation, in particular, is an important constraint that many semi-analytic models of galaxy formation strive to satisfy (Hatton et al., 2003; Cattaneo et al., 2006; Croton et al., 2006; De Lucia et al., 2006). Understanding how the disk acquires its angular momentum is hence of central importance, and has been the subject of this chapter.

6.4.1 The dominance of filaments at high redshift

The evidence favouring the growth of disks in low mass systems via a mode of cold gas accretion appears to be mounting. Kereš et al. (2005) performed SPH simulations of several hundred galaxies between $0 \leq z \leq 3$ using a comoving gravitational softening scale of $5\,h^{-1}\,\mathrm{kpc}$, and found evidence of a clear shift in the total fraction of gas accreted with temperatures $T < 2.5 \times 10^5\,\mathrm{K}$ around a galaxy baryonic mass (stars and cold gas) of $M_{\mathrm{gal}} \sim 2.5 \times 10^{10}\,\mathrm{M}_{\odot}$, with at most a factor of $\sim$1.6 deviation in M_{gal} across this redshift range. Ocvirk et al. (2008) conducted a similar statistical analysis on haloes from the HORIZON-MARENOSTRUM simulation with $10^{10} < M_{\mathrm{H}}/\mathrm{M}_{\odot} < 10^{13}$ at higher redshifts between $2 \leq z \leq 5$, in an attempt to measure the temperature and density of the gas deposited onto the central galaxies of these systems, which were simulated using the RAMSES code at a physical resolution scale of 1 kpc. Their results demonstrated that the fraction of gas accreted at temperatures $T < 2.5 \times 10^5\,\mathrm{K}$ sharply increases for $M_{\mathrm{H}} \lesssim 4 \times 10^{11}\,\mathrm{M}_{\odot}$, in concordance with the estimates provided by Birnboim & Dekel (2003) and Kereš et al. (2005). By resolving and examining the various components of the filamentary gas phase across a wide range in redshift, the results from Fig. 6.1 in this thesis provide quantitive support for the cold mode paradigm of gas accretion onto a Milky Way type disk, advocated by the above studies. They also show that a single filament is responsible for driving the mass budget of the baryonic disk, although this does not necessarily map to a dominance of the disk's angular momentum budget, as two filaments appear to closely share priority at the higher redshifts ($z \gtrsim 5.5$), despite a clear difference in their mass contributions. The flow of cold gas onto the central disk is not necessarily a smooth, continuous process either. While satellites probably do not significantly perturb the filament trajectories onto the disk, large-scale motions of filaments can lead to mergers between these cold streams (Pichon et al., 2011), and this process is able to change the orientation of angular momentum advected along these flows, probably even if gas channelled along each of the filament mergers in question is

co-planar with the disk's rotation before the merger.

A quantitative result of this kind has remained elusive to many previous studies mostly due to the parsec-scale physical resolution required to resolve the satellites, filaments and disk components within the inner tenth of the virial region at $z > 3$. Some authors have recently started to probe this resolution barrier, however. Kimm et al. (2011) measured the angular momentum of gas in radial bins as a function of redshift, for the same resimulated NUT Milky Way type halo analyzed in this thesis. Upon stacking the radial profiles in two separate redshift regimes ($z \leq 3$ and $z > 3$), they found evidence for a sudden loss in the amount of specific angular momentum transported by gas in the $0.1r_{\rm vir}$ region, but with a physical resolution scale of only $\sim$50 pc, were unable to speculate on the cause of this loss. Although the analysis in this chapter does not address this issue *per se*, it does show that the filaments supply enough angular momentum to the boundary of the disk region to be able to account for the disk's signal across most epochs at $z > 3$, a result that has been found by using a filament tracer colouring technique that can be used in future experiments wishing to probe the morphology of the filament trajectories and hence the physical mechanisms that control the redistribution of angular momentum in the disk's vicinity.

Several studies have also demonstrated that the angular momentum vectors of dark matter and gas are not necessarily uniform within the virial region of galaxies at both low ($z = 0$) and high ($z \leq 3$) redshift (Bett et al., 2010; Roškar et al., 2010). Danovich et al. (2012) reported a weak correlation between the angular momentum direction of gas in the inner ($\sim 0.1r_{\rm vir}$) and outer ($\sim 1r_{\rm vir}$) galaxy regions at $z = 2.5$, and a large body of evidence has now been presented that confirms the existence of mismatches in gas angular momentum directions over scales of a few disk scalelengths, giving rise to the population of 'warped' disks (García-Ruiz et al., 2002; Shen & Sellwood, 2006). Despite being observed as a low redshift phenomena, the recent simulation study by Roškar et al. (2010) has argued that warps may exist around $z \sim 2$–3. In Fig. 6.1 of this thesis, the amount of angular momentum transported along the disk's direction by the cumulative filament signal is very comparable to that locked-up in the disk, which would not be the case if large-scale misaligned signals reported above were preserved at the disk boundary. If the claim that the directions of the gas angular momentum vectors vary with distance from the central galaxy holds across the galaxy's lifetime, it is conceivable that some physical mechanism is responsible for aligning the infalling angular momentum at the disk's edge, and this can, in principle, be examined by using the filament tracer colouring algorithm (discussed in more detail in Section 7.2.3). For

example, Roškar et al. (2010) have argued that freshly accreted gas at the virial radius is strongly torqued by the hot halo gas component, and cited this as a possible cause of the warps between inner and outer disk structure. It is hence possible that shocks and/or torques are responsible for removing the components of the purple filament's velocity that are misaligned with respect to the disk's plane of rotation, thereby explaining why this filament is able to maintain a relatively stable transfer of aligned angular momentum to the disk between $3 \lesssim z \lesssim 8$.

Despite the evidence from both Fig. 6.1 and the aforementioned simulation studies favouring the cold gas paradigm, not all authors are convinced that this phase is quite so dominant in growing the disk components of low mass galaxies. Murante et al. (2012) monitored the fractional accretion rate onto two Milky Way type haloes and found that ∼50% of the accreted gas onto the central galaxy between $3 \leq z \leq 6$ was in a warm phase, with a temperature in the range $2.5 \times 10^5 < T/\mathrm{K} < 10^6$. They attributed this apparent discrepancy to a supernovae feedback prescription that only modelled thermal heating, as opposed to thermal and kinetic feedback. The effects of supernovae feedback on filament structure were also examined by van de Voort et al. (2011), who suggested that it reduces the growth rate of low mass central galaxies residing in haloes with mass $M_\mathrm{H} \lesssim 10^{12}\,\mathrm{M}_\odot$, implying that the filaments streaming cold gas at large inflow rates in Fig. 6.1 do not survive when supernovae feedback is included. Powell et al. (2011), however, refuted these claims and argued that this result is an artefact of giving particles an artificial kick velocity, because when the individual Sedov–Taylor blasts from the ultra-high resolution ($\Delta x_\mathrm{min} \sim 0.5\,\mathrm{pc}$) NUT feedback run simulation were resolved, the net mass inflow rates of gas in the filament phase were found to be an order of magnitude larger than the supernovae-driven mass outflow rates. The Powell et al. (2011) study hence implies that whilst the spatial distribution of filaments local to the disk is likely to be perturbed around sites of stellar explosions (e.g. infalling host satellite galaxies), the amount of filament gas flowing towards the disk is largely unaltered upon implementing a reliable model of supernovae feedback. Hence the amount of angular momentum aligned with the disk that is transported by filaments is probably unaltered too, given the general correlation between the filament mass and disk-projected angular momentum signals in Fig. 6.1. It is therefore expected that the difference between the supernovae feedback version of Fig. 6.1 and its counterpart presented in this thesis is negligible. Confirmation of the dominance of the cold mode in driving the angular momentum content of the disk when supernovae feedback is included is left as a future exercise, and is discussed in more detail in Section 7.2.2.

6.4.2 Understanding the impact of the merger signal

Fig. 6.1 shows that luminous satellite mergers constitute only a minor fraction of the disk's mass and angular momentum modulus between $3 \lesssim z \lesssim 8$. This result does not necessarily imply, however, that satellites as a general population have little effect on the disk's evolution. Bett & Frenk (2012) analyzed present day Milky Way haloes of mass $M_{\rm H} \sim 10^{12}$–$10^{12.5}\,h^{-1}\,{\rm M}_{\odot}$ from one of the Millennium Simulation runs, and examined the importance of 'spin flips', which are defined as abrupt changes of more than 45° in the orientation of a component's angular momentum vector. They argued that over 90% of their detected host halo spin flips were caused by minor-mergers, and demonstrated that the number of flip events increases in the inner halo where the central galaxy resides. They further speculated that these spin flips could destroy the host's stellar disk (which is included in the baryonic disk signal of Fig. 6.1), or torque it (see, for example, Ostriker & Binney 1989). In both of these scenarios one would expect a change in the angular momentum direction of the baryonic disk component. The validity of this hypothesis remains an open question, however, because the existence of a simple correlation between dark halo spin flips and disk spin flips is yet to be confirmed: several recent studies have in fact hinted at a lack of correlation between the two (Scannapieco et al., 2009; Stinson et al., 2010; Sales et al., 2011). It will hence be interesting to test this claim by Bett & Frenk (2012) for the NUT host halo in this study, which experiences multiple minor mergers across its accretion history (as indicated in Fig. 5.8).

6.4.3 High redshift contributions to the present day NUT CO disk signals

It is also informative to estimate the fraction of the disk's mass and angular momentum at $z = 0$ that is already in place at $z = 3$, as large high redshift contributions would further highlight the importance of understanding the primordial phase of Milky Way type galaxy growth. Crude estimates are hence provided in this section, but it should be noted that more rigorous measurements could be made by analyzing all of the lower resolution outputs from the NUT CO run between $0 \leq z \leq 3$ and extending Fig. 6.1 to $z = 0$.

Kimm et al. (2011) have demonstrated that the total amount of baryonic mass (i.e. gas plus stars) within $0.1r_{\rm vir}$ from the $\Delta x_{\rm min} = 48\,{\rm pc}$ NUT CO run at $z = 0$ is $M_b(z = 0) \sim 8 \times 10^{10}\,{\rm M}_{\odot}$. This measurement includes the central contributions from all of the gas phases (satellite, hot mode, cold mode and disk) and all of the stars (satellite, disk and bulge), and is a factor of ∼6 greater than the baryonic disk signal at $z = 3$ from Fig. 6.1, which is $M_b(z = 3) \sim 1.3 \times 10^{10}\,{\rm M}_{\odot}$. Multiplying the

modulus of the total baryonic specific angular momentum contributions within the $0.1r_{\rm vir}$ region at $z = 0$ from the Kimm et al. (2011) study by the relevant component masses yields an estimate of the total baryonic angular momentum signal: $J_b(z = 0) \sim 4.4 \times 10^{13}\,{\rm M}_\odot\,{\rm kpc\,km\,s^{-1}}$. This is a factor of ∼30 higher than the modulus of the baryonic disk angular momentum signal at $z = 3$ shown in Fig. 6.1: $J_b(z = 3) \sim 1.6 \times 10^{12}\,{\rm M}_\odot\,{\rm kpc\,km\,s^{-1}}$. Note that the mass (and most probably the angular momentum) ratios are upper limit estimates because the Kimm et al. (2011) measurements include the satellite star and non-disk gas contributions, whereas Fig. 6.1 just shows the baryonic disk signals defined according to equations (5.8) and (5.9).

In summary, at least ∼16% (∼4%) of the NUT disk's final mass (angular momentum) is probably in place by $z = 3$, suggesting that certainly in terms of mass, the high redshift epochs studied in this thesis represent an important period of the disk's accretion history.

6.5 A brief theoretical and observational outlook for the cold gas paradigm

Unfortunately there is no compelling observational evidence at present that confirms detections of filaments streaming cold gas to the centres of low mass galaxies, and given the high redshift nature of this phenomenon, securing a reliable detection of this kind seems a daunting task. Over the past decade, however, various authors have started to grasp the observational dynamics of these systems and the field has now become an active area of research.

With temperatures in the range $10^4 \lesssim T/{\rm K} \lesssim 2 \times 10^5$, cooling emission from filaments is expected to correspond to Lyman-α wavelengths, and several studies have argued that cold flows fuel the Lyman-α 'blobs' at $z \sim 3$ (Goerdt et al., 2010a; Matsuda et al., 2011). Faucher-Giguère et al. (2010) have contested these claims by showing that when the high density gas phase is allowed to form stars, the emitted Lyman-α luminosity is at least an order of magnitude below the observed blob luminosities of $L_{{\rm Ly}\alpha} \sim 10^{43}$–$10^{44}\,{\rm erg\,s^{-1}}$ (Steidel et al., 2000). An alternative approach is to try to detect filaments in absorption, but Steidel et al. (2010) have demonstrated that filaments do not account for a significant fraction of the bolometric luminosity and hence have low covering factors that become ever lower in an expanding background, rendering this a very difficult task. Nonetheless, Fumagalli et al. (2011) have attempted to measure these covering factors by performing a radiative transfer post-processing analysis on 7 resimulated Milky Way

type haloes between $1 \lesssim z \leq 4$. They found that the simulated Lyman-α equivalent width radial profile in absorption at $z \sim 3$ matched its observed counterpart, and attributed these 'Lyman-limit systems', whose metallicities were of the order $\sim 0.01\,\mathrm{Z}_{\odot}$, to gas in the cold filament phase.

Perhaps the most pressing issue facing the simulation community at present, however, is the inability to convincingly reproduce the stellar and gas components of a Milky Way galaxy across its lifetime. With a bulge-to-disk ratio of B/D ~ 0.75 at $z = 0$, a factor of $\sim$3 higher than the measured Milky Way value (e.g. Drory & Fisher 2007), the lower resolution ($\Delta x_{\mathrm{min}} \sim 48\,\mathrm{pc}$) NUT CO run is no exception to this problem. As argued in Section 1.5.2, and by several other authors (Governato et al., 2004; Mayer et al., 2008), it is likely that both a correct model for supernovae feedback and sub-parsec spatial resolution are required in order to yield accurate Milky Way disk properties. The recent work by Agertz et al. (2011) reinforces the importance of these ingredients. They demonstrated that by using the NUT star formation recipe presented in Section 2.1.5, and a low value for the star formation efficiency parameter ($\epsilon \sim 1\%$), it is possible to produce realistic present day disk candidates with scalelengths $r_d \sim 4$–$5\,\mathrm{kpc}$ and bulge-to-disk ratios B/D ~ 0.25. Achieving these results for multiple simulated disk galaxies is the next grand challenge.

Having now presented a suite of techniques developed for attacking the scientific objectives of both research projects examined in this thesis, and having discussed in detail the results found upon implementing these methods, the following final chapter provides a concise thesis summary and proposes several related projects for future research.

Chapter 7

Conclusions and future projects

7.1 The main results from this thesis

Two accretion algorithms—the halo and halosub methods—have been presented in Section 3.2. Both schemes measure the positive mass accretion onto a given bound AHOP halo i in the simulation using different prescriptions. The former examines the mass difference ΔM_{ij} between i and its main progenitors j along the main branch, possibly tracing back several time intervals before recording positive ΔM_{ij}, and is restricted to haloes, while the latter measures the fraction of i's mass acquired from mergers and accretion of diffuse non-halo material over a single interval, and can be applied to all levels of structure, conserving mass by construction. Section 3.2 also highlights the importance of quantifying accretion onto bound objects only. Unbound objects are classed as structures that either have a positive total energy or that participate in unrelaxed 'anomalous' merger events, and an algorithm is presented that identifies objects of this category and excludes them from the accretion signals. In Chapter 4, the results obtained from employing these methods are presented, and the main findings are as follows:

- EPS theory predicts a stronger evolution in the halo accretion rate compared with the measured rates in the simulation, but the level of relative discrepancy in the accretion trajectory gradients varies between the halo and halosub methods.

- Radio-mode feedback in semi-analytic models of galaxy formation could be caused by episodes of low mass accretion onto haloes, which yield both the correct redshift evolution and present day integrated black hole accretion rate, provided black hole growth correlates with halo growth.

- Present day haloes may be accreting at fractional rates that are up to a factor 3–4 higher than their associated black holes.

- The local subhalo accretors in the simulation experience enhanced clustering at small scales, which is likely to account for their higher average rates of accretion compared with haloes.

- Only a weak relationship is detected between local accretion rate and environment at scales typical of a large group or cluster. The subhalo accretion rates are probably largely insensitive to environment because the imposed minimum mass threshold of 40 particles for accretors in this study restricts analysis to a small range of massive subhalo environments, and the mutual interactions at small scales that drive accretion onto subhaloes in the outskirts of their hosts are unlikely to significantly differ between these environments.

Chapter 5 then switches focus to the second project and begins by introducing a Lagrangian filament construction technique and a satellite finding algorithm, which have both been applied to the CO run from the NUT suite with $\Delta x_{\rm min} = 12\,{\rm pc}$, in order to address whether the angular momentum acquired by the baryonic disk of a Milky Way type galaxy at high redshift is driven by filaments. Figs 5.6 and 5.7 highlight the ability of the tracer colouring filament method at reproducing the apparent individual filament trajectories that channel material onto the disk, and Fig. 5.8 confirms that the satellite-finding algorithm performs reliable detections of luminous satellite merger candidates. The angular momentum and mass transported by these components are then measured using the techniques described in Section 5.5.3, and the results obtained from implementing these algorithms in Chapter 6 suggest that:

- The cumulative effect of the cold filament gas phase dominates the mass and angular momentum budgets of the disk as a function of time, hence providing quantitative support for the cold gas paradigm of disk growth in low mass galaxies at high redshift.

- For $5.5 \lesssim z \lesssim 7.5$, the largest portion of the filament angular momentum signal is transported by two filaments, which start to merge around $z \sim 5.5$. This merger appears to induce a change in the orientation of the angular momentum transported towards the disk, but the surviving filament remnant quickly realigns itself along the disk direction and by $z \sim 4$, dominates the filament angular momentum budget by at least a factor of 3.

- The luminous satellites account for at most one tenth of the disk mass and angular momentum modulus at any given epoch.

As argued in Section 5.2, recent studies comparing local stellar mass functions with the redshift-independent cold mode halo mass threshold suggest that low temperature gas is likely to drive the growth of most galaxies at high redshift. Simulation studies following the accretion history of present day haloes that are up to an order of magnitude more massive than pure cold mode galaxy haloes have also revealed that a sizeable fraction of their gas was accreted cold. The high resolution NUT Milky Way type halo analyzed in this thesis has provided a framework for quantifying the importance of this cold mode paradigm of galaxy growth.

7.2 Current research and possible future extensions

The remainder of this chapter proposes several further projects that serve to shed more light on how angular momentum is transferred to the central galaxy disk component across time. It begins by introducing a method that has recently been developed at the time of writing this thesis for monitoring angular momentum signals along filaments, but whose predictions are yet to be tested.

7.2.1 Walking along the separate filaments

The tracer colouring technique of Section 5.4.2 is able to construct the separate filaments at a given time-step, but by itself is incapable of following their trajectories from the point they cross the virial sphere to the point they join with the disk. Defining these trajectories would enable one to monitor the orientation and magnitude of the angular momentum transported along the flow direction as a function of scale, and possibly locate regions where the vector changes, perhaps due to shocks or torquing from a galaxy component. This may also explain why the modulus of the specific angular momentum of gas $|\vec{j}_{\rm gas}|$ decreases within the $0.1r_{\rm vir}$ region (Kimm et al., 2011). The main filament trajectory is therefore likely to reveal information about the physics of the flow, and so the algorithm that has been designed to walk along filament flow directions is now described.

Using the skeleton to define each filament's main trajectory

The approach adopted in this study for performing walks along each filament is to follow their ridges of maximum density, as material local to the dense component is likely to contribute most

to the overall local filament angular momentum budget. The peaks in the density field have been computed by using the 'skeleton' code (Pogosyan et al., 2009; Sousbie, 2011). The purpose of the skeleton is to connect ridges of maximum density with local saddle points, providing a framework that joins critical locations $\vec{x}_{\rm crit}$ of the density field $\rho(\vec{x})$ at each epoch ($\vec{\nabla}\rho(\vec{x}_{\rm crit}) = 0$). This task reduces to computing the Hessian of $\rho(\vec{x})$:

$$\mathcal{H}(\rho)_{ij}(\vec{x}) = \frac{\partial^2 \rho}{\partial x_i \partial x_j}(\vec{x}), \tag{7.1}$$

and bridging links to points in the direction of the Hessian eigenvector with the largest eigenvalue, corresponding to the direction of the path with the least local density curvature (Pichon et al., 2010; Sousbie, 2011). The skeleton requires two dimensionless parameters at run-time, $\epsilon_{\rm smo}$ and $\epsilon_{\rm per}$: the former represents the number of iterations that are used to smooth the positions of the non-critical points with the positions of their nearest neighbours, while the latter corresponds to the allowed maximum difference in density between a pair of connected critical points (the so-called 'level of persistence'), normalized with respect to mean density of the inserted grid. The code then outputs a web of conjoined paths, where a single path is defined as a series of connected links that may or may not be of similar length. Each link has a start and end node, and so two connected links share at least one node (both nodes would be shared in the case of a closed loop). Fig. 7.1 shows all of the paths at $z \sim 8$ as black lines that were detected by the skeleton with $\epsilon_{\rm smo} = 200$ and $\epsilon_{\rm per} = 200$ over a box of sidelength equal to $2r_{\rm vir}$ ($\sim$15 kpc) centred on the centre of the host stellar clump in the NUT CO run. The skeleton has been computed on a coarse grid with a physical cell size of $\sim$44 pc ($\ell = 15$), and this choice of 'skeleton grid' is discussed in more detail at the end of this subsection.

Having found all of the paths along the maximum density ridges within the host virial region, the next task is to select a path that walks from the starting point of a given filament at the virial radius, to its terminus, which is most likely to be the outer regions of the gas disk. Fig. 7.2 shows an example of one such possible skeleton path configuration and serves as a visual guide to the 'walking' algorithm used in this study. This algorithm examines each filament in isolation and consists of five main stages:

1. **Choice of starting location**

   All of the filament links in a thin spherical shell spanning $0.98 \leq r/r_{\rm vir} \leq 1$ are found, with

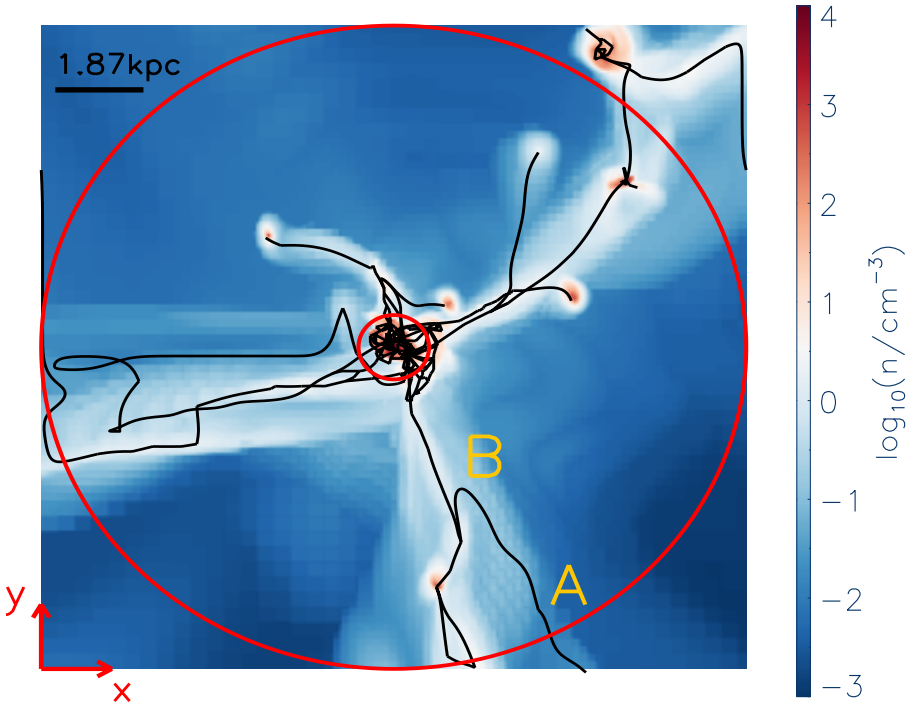

Figure 7.1: This figure shows the paths of maximum local density detected by the skeleton (black lines) when the gas is refined to $\ell = 15$ ($\Delta x \sim 44\,\mathrm{pc}$) over the host's virial region (outer red circle) at $z \sim 8$ in the NUT CO run with $\epsilon_{\mathrm{per}} = \epsilon_{\mathrm{smo}} = 200$. The skeleton identifies several paths along each filament that cross and mix in the $0.1r_{\mathrm{vir}}$ disk region (inner red circle), and the labels along one such path serve to highlight the difficulties with defining individual main trajectories (see text).

each of these links possibly representing a trajectory that leads to the disk. An example of a starting link in this radial bin is given in Fig. 7.2.

2. **Neighbouring link identification**

   Each link consists of two nodes, which are shown as filled black circles in Fig. 7.2, and is assigned a number according to the number of unflagged neighbouring filament links that connect to it along the local walking direction. The starting link in Fig. 7.2 has only one node connected to other filament links, and so this defines the local walking direction. It is of course possible that both of the starting link's nodes are joined to separate filament links, which would yield two local walking directions. In this scenario, however, it would not be long before one of the walks traverses outside the virial region (because the starting point is so close to the virial radius) or along a known path, and so a single walking direction towards the disk is considered in what follows. In Fig. 7.2, the algorithm assigns link i a '2' because it has two unflagged filament links, j and k, that join to it along the walking direction. Link

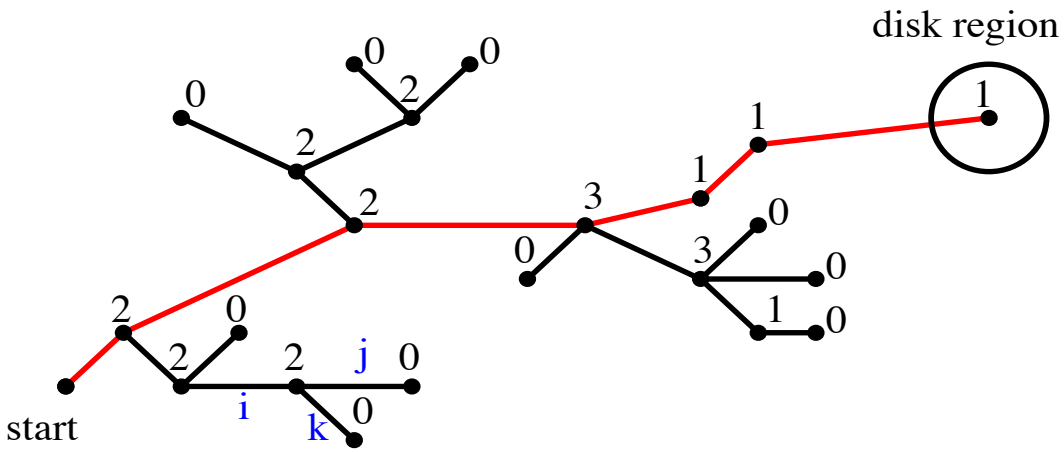

Figure 7.2: This diagram illustrates how the main skeleton path (red line) is selected from a set of connected filament links, given a starting position near the host's virial radius. Each link is assigned a value that represents the number of unflagged neighbouring links connected to it, and is subsequently flagged. 'Dead-ends' (e.g. j and k), defined as links outside the disk region with a '0', are then decayed backwards along their parent links (i) in such a way that after the whole decay process a chain of links with number '1' remain, revealing the main path (see text for details).

i is then flagged. The links that mark the end of the walk have terminal status (assignment of a '0'), and there are two ways such status can be achieved using this method:

- if the link in question is not connected to any unflagged neighbouring filament links, and/or;
- neither node of the link corresponds to a filament cell on the skeleton grid.

Link j and k are terminals because they are only connected to link i, which has already been flagged in reaching them, and hence violate the first condition. An exception to the rules above is made if a terminal link lies within the disk region, in which case its assigned value of '0' is changed to '1' (as shown in Fig. 7.2), distinguishing it from the 'dead-end' terminals beyond the disk region.

3. **The decay of dead-ends and subsequent collapse to the main path**

Once every filament link's nodes have been examined on the skeleton grid and each link has been assigned an appropriate number, the decay process begins. The algorithm starts at each dead-end that has been assigned a '0' and traces backwards, subtracting '1' from the parent link number for every mutual dead-end connection. Link i in Fig. 7.2 is the parent of links j and k, which themselves are dead-ends, and hence two subtractions from the number '2' assigned to i occur (one from j and the other from k), demoting it to a dead-end. Once this

decay process is iteratively performed over all dead-end links detected along the path, a single chain of links remain, each with the number '1', highlighting the main path. This is shown as the red line in Fig. 7.2.

4. **Main trajectory selection**

   The penultimate step of the algorithm makes a distinction between a main path and a main trajectory, as there are often several main paths like the one in Fig. 7.2 that reach the disk from different starting points near the virial radius, and the goal is to select the one that best follows the flow direction onto the disk. Once the main trajectory has been identified by the walking algorithm, it can be tested against visualizations of the density field within the virial region, which help reveal the 'true' local flow direction for each filament.

   In this study, the main trajectory that penetrates through to the disk region is defined as either the shortest mass-weighted path or the path with the largest integrated mass. Both prescriptions have their relative advantages and disadvantages. The former has the desirable property that it penalizes paths that follow unclosed loops (e.g. loop 'B' in Fig. 7.1) that have a larger integrated mass due to their additional length, but it has the undesirable property that it does not necessarily trace the curvature of the filament bends around the disk, as longer trajectories are disfavoured. The latter definition is likely to account for the dominant component of the angular momentum budget of each filament and gives preference to the longer paths that twist and turn onto the disk, but it is susceptible to choosing skeleton links that do not appear to best coincide with each filament's known flow direction (e.g. the path starting at 'A' in Fig. 7.1). Therefore the performance of both of these main trajectory definitions is tested by visual inspection for each individual filament at a given time-step in this study, and the appropriate path is subsequently chosen.

   The integrated mass of a given main path is computed by combining the mass of all the cells on the fine grid ($\Delta x \sim 12\,\mathrm{pc}$) within spheres that are centred on the midpoint of each member link. This technique is henceforth referred to as the 'spheres method' and is defined by a single free parameter r_s, corresponding to the sphere radius. Assigning each sphere a fixed value of r_s suffers from two problems: (i) it ignores the density distribution of filament gas along the main trajectory and could therefore yield a large distribution in enclosed mass, which would bias the angular momentum signal; (ii) it is insensitive to the physical nature

of a given filament's trajectory and could combine gas at different angular positions in the central region where the filament bends, which is analogous to averaging over the filament's curvature. A minimum density limit of $1000\,\mathrm{cm}^{-3}$ is hence imposed instead, in an attempt to remove the potential mass bias. This choice of parameter value approximately bisects the range in the nT filament density criteria for 1000 filament cells of equal volume, and so the Poisson errors are likely to be low. In order to avoid premature path truncation, the walking algorithm identifies the satellite cells on the coarse skeleton grid and permits the walk to continue when these cells are encountered, but ensures that their counterparts on the fine grid do not contribute to the mass budget along the path because satellites are not defined as part of filaments in this study.

Fig. 7.3 shows the main trajectories projected along the z and x axes for each of the filaments in the upper (purple), lower left (yellow) and lower right (red) panels of Fig. 5.7. The algorithm follows each filament towards the disk region and mostly captures the individual twists and turns, as is best illustrated for the yellow filament's trajectory. It appears that the red filament is truncated prematurely, but the mixing with the purple filament destroys most of its inner structure and the surviving inner red filament remnant cannot be resolved at a refinement level of $\ell = 15$. It can be seen in the lower right panel of Fig. 5.7, however, that this extra inner structure of the red filament that is resolved by increasing the refinement level to 17 at this epoch is minimal, and so the amount of angular momentum it carries is likely to be minimal too, especially since the trajectory appears to be mostly radial.

5. **Angular momentum transport along the flow**

   With the main trajectory of each filament identified at a given epoch, it is possible to compute the total angular momentum of all the cells on the fine grid within the spheres along the trajectory by using equation (5.7). This yields the angular momentum profile of the individual filaments.

Computational challenges

The main difficulties are associated with the skeleton web, which is sensitive to the values of ϵ_{per} and ϵ_{smo}. If both parameters are too small, the skeleton structure is noisy and can follow directions that locally do not conform to the direction of the bulk inward radial flow. If on the other hand

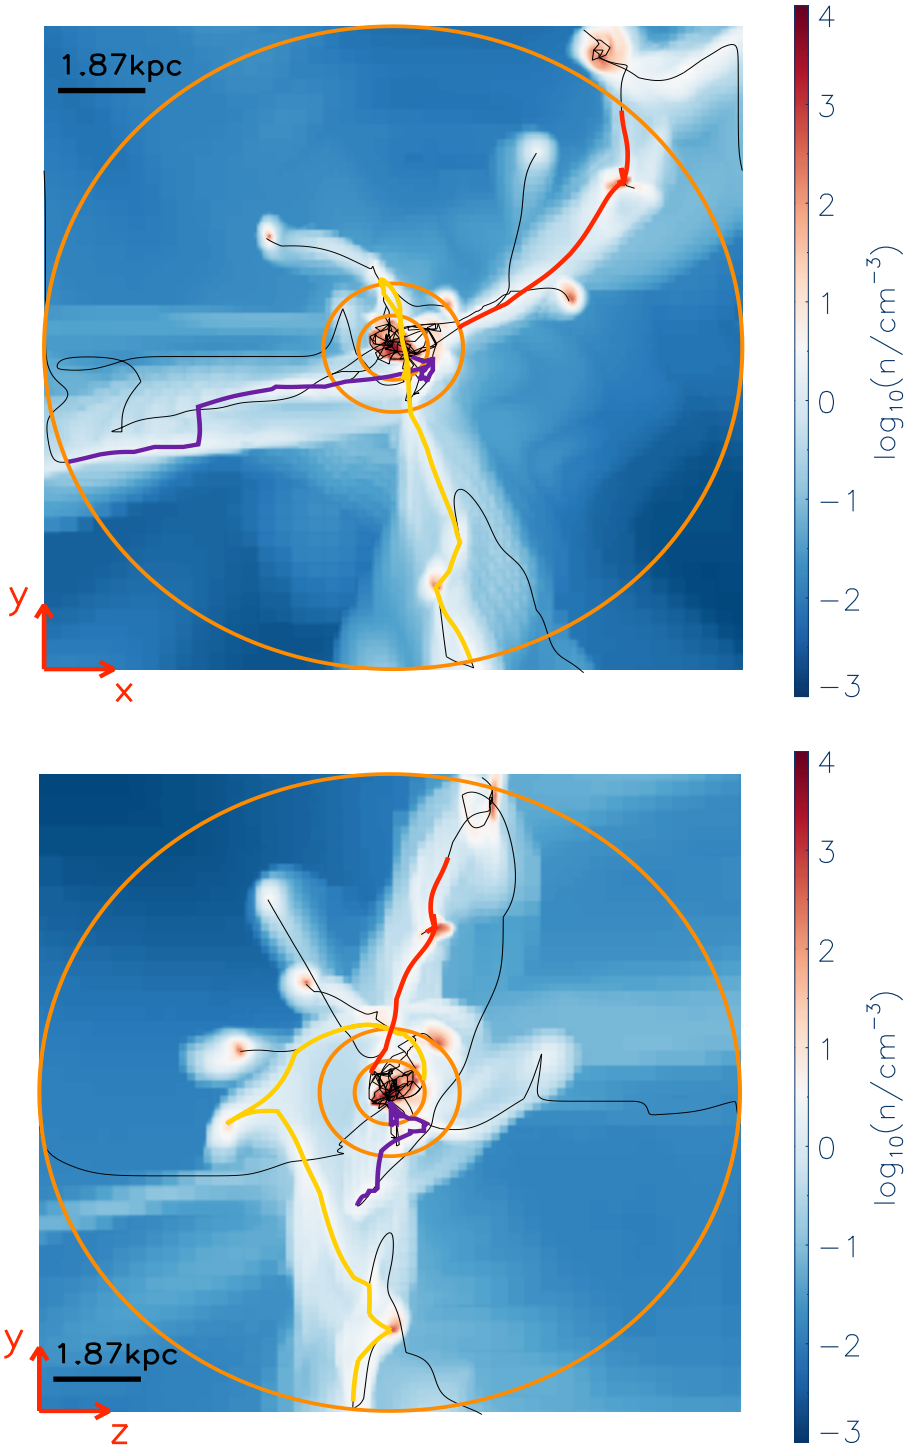

Figure 7.3: These images show the main filament trajectories at $z \sim 8$ selected from the skeleton paths (thin black lines) by the walking algorithm described in Section 7.2.1. The grid dimensions are identical to the skeleton grid shown in Fig. 7.1, and the purple, yellow and red filaments correspond to the top, lower left and lower right panels of Fig. 5.7 respectively. The upper and lower visualizations in gas density represent projections along the z and x axes, with the $0.1r_{\rm vir}$, $0.2r_{\rm vir}$ and $r_{\rm vir}$ regions shown by the inner, middle and outer orange circles respectively. Notice that the lower panel particularly highlights the success of the walking algorithm in following the curvature of the yellow filament about the disk.

both parameters are too large, connections are made between density maxima that have large spatial separations (which the walking algorithm may flag as part of the main trajectory), and the details of a given filament's curvature (e.g. sharp bends) are lost in excessive smoothing. Both of these extreme cases can yield trajectories that do not necessarily correspond to the apparent trajectory of each filament over the range of scales within the virial region. The ratio between the time it takes to compute the skeleton web on two grids refined to different levels also approximately scales with the grid volume ratio, which justifies the choice of coarse grid in Fig. 7.1. Coarser grids are further favoured when deciding on the best values of $\epsilon_{\rm per}$ and $\epsilon_{\rm smo}$, which could prove a very laborious exercise on finer grids. It is hence likely that only a small subset of time-steps suitably spaced between $z \sim 8$ and $z = 3$ will be considered when probing the time evolution of the system.

7.2.2 Analyzing the outputs from the more complex NUT runs

Once the walking algorithm and spheres method have been successfully implemented on the CO run, it will be interesting to conduct a similar analysis on the NUT outputs that include feedback from supernovae (FB run) according to the prescriptions described in Section 2.1.5. The inclusion of this additional heating source will probably perturb the density field local to filament regions that host satellites and that penetrate the disk region, which could yield a more clumpy density distribution. Powell et al. (2011) have demonstrated that the filaments in the NUT simulations survive upon inclusion of supernovae feedback, but that the gas density distribution becomes far less smooth, and so the walking algorithm presented in the previous section may require modification to avoid premature path truncation outside the disk region. It will be interesting to test these predictions on the ultra-high resolution NUT FB run with a physical resolved scale of 0.5 pc.

Supernovae also inject metals into their local surroundings, and so filaments are likely to cool more quickly than in the CO run. This increase in the metallicity of the ISM will therefore enable a further investigation of the role of cold gas in establishing the disk's angular momentum, as the fraction of gas at late times with $T \leq 2 \times 10^4$ K will increase. It should hence be possible to reduce the filament's upper temperature limit from $T_u = 2 \times 10^5$ K to $T_u = 2 \times 10^4$ K. Metallicity maps may also prove instructive when assessing the role stripped satellite gas has in growing disks. The gas bound to the outer regions of satellites streaming along filaments that is likely to largely constitute the stripped component has been attributed to the filament phase in this thesis for reasons discussed in Section 5.4.3, and Fig. 5.8 indicates that this could form a non-negligible signal. Decoupling this

from the filament phase and identifying it as a separate, metal-rich component may hence prove instructive. Metallicity cuts would inevitably be introduced to carry out this task.

7.2.3 Identifying gaseous shocks and measuring torques

It can be seen in Fig. 5.7 that filaments often follow trajectories that bend in the vicinity of the disk before eventually merging with it, which raises the question of why these filaments bend. Pichon et al. (2011) have speculated that the filaments encounter a radial isothermal shock at a certain angle of incidence on the periphery of the disk, which radiates away the component of the velocity perpendicular to the shock front as thermal energy, leaving behind the velocity component in the plane of the shock. They argue that this could cause the trajectory to bend, as the velocity vector will point in a different direction. It is, in principle, possible to test for radial shocks in both the CO and FB runs by measuring the radial and angular components of the linear momentum of filaments using the spheres method in combination with the walking algorithm, and by locating sudden drops in the radial component. Detecting oblique shocks would be far less trivial, however.

It is also conceivable that filament trajectories bend due to torquing from a galaxy component. Roškar et al. (2010) have measured the gravitational and hydrodynamical torques on hot and cold gas as a function of scale within the virial region and claim that at around $0.1r_{\rm vir}$ from the halo centre, the torque from dark matter is roughly one tenth of the torque from gas. By measuring the modulus and direction of the torques on the filament gas cells due to dark matter and gas, and by recording the angular positions of the filament segments, it may be possible to associate swings in the torque vector with swings in filament trajectories.

7.2.4 Resolving the inner disk structure

With a minimum physical scale of $\sim$0.5 pc, the angular momentum and mass profile of the disk can be resolved at each time-step, which could reveal evidence for an inner and outer disk component, in concordance with the lower redshift disk warps reported in the literature (e.g. García-Ruiz et al. 2002). It may also be interesting to separate the combined stellar and gas disk signals in Fig. 6.1 and study their interplay as a function of time. Given the distribution in stellar ages and the velocity kicks experienced after the supernova phase, perhaps there exists an old and a young stellar disk with some level of misalignment.

7.2.5 Lagrangian dynamics using SPH

There is a considerable amount of literature that debates the relative merits of SPH codes versus adaptive grid codes (see Table 1.2 for a concise summary). The NUT CO run analyzed in this thesis boasts a redshift-independent physical gravitational softening length of 12 pc, which invites studies similar to that presented in Chapter 6 to probe the redistribution of angular momentum in the disk's vicinity by providing the high resolution required to resolve the disk and the filaments. Most SPH codes, by contrast, are forced to adopt a model characterized by a parameter $z_{\rm fix}$ that switches the softening scale at some epoch from being fixed in comoving coordinates at high redshift to being fixed in physical coordinates at low redshift, and another parameter ϵ_g that determines the softening scale at the present day, which is typically between 200–700 pc (e.g. Springel et al. 2008; Okamoto et al. 2010). Yet despite having poorer spatial resolution, the SPH approach does offer an attractive framework that is, by definition, suited to the Lagrangian task of following accretion of gas onto the central disk component. It would hence be interesting to tag all of the particles that reside within the disk at the final output of a resimulation SPH run, record their identification numbers and subsequently rewind back to the start of the simulation, following each particle as soon as it crosses the virial sphere. Each filament could be resolved in a similar fashion to the tracer particle colouring scheme, except a nearest neighbours scheme would be used in place of a neighbour cell colour propagation method (Section 5.4.2). The aim would then be to follow the angular momentum modulus and direction of each coloured SPH filament gas particle as a function of time, before it is accreted onto the disk. The SPH approach would hence offer unique trajectories in the angular momentum versus time plane for collections of coloured filament particles accreted at the virial radius at similar times, which could prove instructive from the point of view of identifying physical mechanisms along the filaments, such as shocks. It would also be interesting to make a comparison with the SPH version of Fig. 6.1.

In summary, this research project would enable a thorough quantitative comparison between two alternative schemes for performing radiative gas hydrodynamics. Perhaps the results would support the recent study by Scannapieco et al. (2012), who argued that the disagreement between grid-based and SPH-based predictions for the present day stellar mass, size, rotation curve and morphology of their resimulated galaxy from the Aquarius Project simulation (Springel et al., 2008), was driven by differing models of star formation and feedback, rather than fundamental numerical differences between the particle and grid treatments themselves.

Bibliography

Abadi, M. G., Navarro, J. F., Steinmetz, M. & Eke, V. R., 2003. Simulations of Galaxy Formation in a Λ Cold Dark Matter Universe. I. Dynamical and Photometric Properties of a Simulated Disk Galaxy. ApJ, **591**, 499–514.

Abazajian, K., Fuller, G. M. & Tucker, W. H., 2001. Direct Detection of Warm Dark Matter in the X-Ray. ApJ, **562**, 593–604.

Agertz, O., Moore, B., Stadel, J., Potter, D., Miniati, F. et al., 2007. Fundamental differences between SPH and grid methods. MNRAS, **380**, 963–978.

Agertz, O., Teyssier, R. & Moore, B., 2009. Disc formation and the origin of clumpy galaxies at high redshift. MNRAS, **397**, L64–L68.

Agertz, O., Teyssier, R. & Moore, B., 2011. The formation of disc galaxies in a ΛCDM universe. MNRAS, **410**, 1391–1408.

Angulo, R. E., Lacey, C. G., Baugh, C. M. & Frenk, C. S., 2009. The fate of substructures in cold dark matter haloes. MNRAS, **399**, 983–995.

Aragón-Calvo, M. A., van de Weygaert, R. & Jones, B. J. T., 2010. Multiscale phenomenology of the cosmic web. MNRAS, **408**, 2163–2187.

Athanassoula, E., Fady, E., Lambert, J. C. & Bosma, A., 2000. Optimal softening for force calculations in collisionless N-body simulations. MNRAS, **314**, 475–488.

Aubert, D., Pichon, C. & Colombi, S., 2004. The origin and implications of dark matter anisotropic cosmic infall on $L_\star$ haloes. MNRAS, **352**, 376–398.

Babić, A., Miller, L., Jarvis, M. J., Turner, T. J., Alexander, D. M. & Croom, S. M., 2007. Low accretion rates at the AGN cosmic downsizing epoch. A&A, **474**, 755–762.

Bagla, J. S., 2002. TreePM: A Code for Cosmological N-body Simulations. JA&A, **23**, 185–196.

Bagla, J. S. & Padmanabhan, T., 1997. Cosmological N-body simulations. Pramana, **49**, 161.

Bagla, J. S. & Ray, S., 2003. Performance characteristics of TreePM codes. New Astronomy, **8**, 665–677.

Bailin, J. & Steinmetz, M., 2005. Internal and External Alignment of the Shapes and Angular Momenta of ΛCDM Halos. ApJ, **627**, 647–665.

Baldry, I. K., Glazebrook, K., Brinkmann, J., Ivezić, Ž., Lupton, R. H., Nichol, R. C. & Szalay, A. S., 2004. Quantifying the Bimodal Color-Magnitude Distribution of Galaxies. ApJ, **600**, 681–694.

Bardeen, J. M., Bond, J. R., Kaiser, N. & Szalay, A. S., 1986. The statistics of peaks of Gaussian random fields. ApJ, **304**, 15–61.

Barger, A. J., Cowie, L. L., Mushotzky, R. F., Yang, Y., Wang, W., Steffen, A. T. & Capak, P., 2005. The Cosmic Evolution of Hard X-Ray-selected Active Galactic Nuclei. ApJ, **129**, 578–609.

Barnes, J. & Efstathiou, G., 1987. Angular momentum from tidal torques. ApJ, **319**, 575–600.

Barnes, J. & Hut, P., 1986. A hierarchical $\mathcal{O}(N \log N)$ force-calculation algorithm. Nature, **324**, 446–449.

Batcheldor, D., 2010. The $M_\bullet$–$\sigma_\star$ Relation Derived from Sphere of Influence Arguments. ApJL, **711**, L108–L111.

Bauer, A. E., Drory, N., Hill, G. J. & Feulner, G., 2005. Specific Star Formation Rates to Redshift 1.5. ApJL, **621**, L89–L92.

Baugh, C. M., 2006. A primer on hierarchical galaxy formation: the semi-analytical approach. Reports on Progress in Physics, **69**, 3101–3156.

Behroozi, P. S., Wechsler, R. H. & Wu, H.-Y., 2011. The Rockstar Phase-Space Temporal Halo Finder and the Velocity Offsets of Cluster Cores. Preprint (arXiv:1110.4372).

Bell, E. F., McIntosh, D. H., Katz, N. & Weinberg, M. D., 2003. The Optical and Near-Infrared Properties of Galaxies. I. Luminosity and Stellar Mass Functions. ApJS, **149**, 289–312.

Benson, A. J., 2010. Galaxy formation theory. Physics Reports, **495**, 33–86.

Benson, A. J., Bower, R. G., Frenk, C. S., Lacey, C. G., Baugh, C. M. & Cole, S., 2003. What Shapes the Luminosity Function of Galaxies? ApJ, **599**, 38–49.

Benson, A. J., Frenk, C. S., Lacey, C. G., Baugh, C. M. & Cole, S., 2002. The effects of photoionization on galaxy formation - II. Satellite galaxies in the Local Group. MNRAS, **333**, 177–190.

Benson, A. J., Kamionkowski, M. & Hassani, S. H., 2005. Self-consistent theory of halo mergers. MNRAS, **357**, 847–858.

Bett, P., Eke, V., Frenk, C. S., Jenkins, A. & Okamoto, T., 2010. The angular momentum of cold dark matter haloes with and without baryons. MNRAS, **404**, 1137–1156.

Bett, P. E. & Frenk, C. S., 2012. Spin flips - I. Evolution of the angular momentum orientation of Milky Way-mass dark matter haloes. MNRAS, **420**, 3324–3333.

Binney, J., 1977. The physics of dissipational galaxy formation. ApJ, **215**, 483–491.

Binney, J. & Tremaine, S., 2008. *Galactic Dynamics*. Princeton University Press, 2nd ed.

Birnboim, Y. & Dekel, A., 2003. Virial shocks in galactic haloes? MNRAS, **345**, 349–364.

Bode, P. & Ostriker, J. P., 2003. Tree Particle-Mesh: An Adaptive, Efficient, and Parallel Code for Collisionless Cosmological Simulation. ApJS, **145**, 1–13.

Bode, P., Ostriker, J. P. & Xu, G., 2000. The Tree Particle-Mesh N-body Gravity Solver. ApJS, **128**, 561–569.

Bond, J. R., Cole, S., Efstathiou, G. & Kaiser, N., 1991. Excursion set mass functions for hierarchical Gaussian fluctuations. ApJ, **379**, 440–460.

Bond, J. R., Kofman, L. & Pogosyan, D., 1996. How filaments of galaxies are woven into the cosmic web. Nature, **380**, 603–606.

Book, L. G., Brooks, A., Peter, A. H. G., Benson, A. J. & Governato, F., 2011. Angular momentum evolution in dark-matter haloes. MNRAS, **411**, 1963–1976.

Bouchet, F. R. & Kandrup, H. E., 1985. Particle-mesh simulations of clustering in cosmology. ApJ, **299**, 1–4.

Bower, R. G., 1991. The evolution of groups of galaxies in the Press–Schechter formalism. MNRAS, **248**, 332–352.

Bower, R. G., Benson, A. J., Malbon, R., Helly, J. C., Frenk, C. S., Baugh, C. M., Cole, S. & Lacey, C. G., 2006. Breaking the hierarchy of galaxy formation. MNRAS, **370**, 645–655.

Boyarsky, A., Neronov, A., Ruchayskiy, O. & Shaposhnikov, M., 2006. Constraints on sterile neutrinos as dark matter candidates from the diffuse X-ray background. MNRAS, **370**, 213–218.

Brinchmann, J. & Ellis, R. S., 2000. The Mass Assembly and Star Formation Characteristics of Field Galaxies of Known Morphology. ApJL, **536**, L77–L80.

Brooks, A. M., Governato, F., Quinn, T., Brook, C. B. & Wadsley, J., 2009. The Role of Cold Flows in the Assembly of Galaxy Disks. ApJ, **694**, 396–410.

Bryan, G. L., Abel, T. & Norman, M. L., 2001. Achieving Extreme Resolution in Numerical Cosmology Using Adaptive Mesh Refinement: Resolving Primordial Star Formation. Preprint (arXiv:astro-ph/0112089).

Bryan, G. L. & Norman, M. L., 1997. Simulating X-Ray Clusters with Adaptive Mesh Refinement. In D. A. Clarke & M. J. West, eds., *Computational Astrophysics; 12th Kingston Meeting on Theoretical Astrophysics*, vol. 123 of *Astronomical Society of the Pacific Conference Series*, 363.

Bryan, G. L. & Norman, M. L., 1998. Statistical Properties of X-Ray Clusters: Analytic and Numerical Comparisons. ApJ, **495**, 80.

Bullock, J. S., Kolatt, T. S., Sigad, Y., Somerville, R. S., Kravtsov, A. V., Klypin, A. A., Primack, J. R. & Dekel, A., 2001. Profiles of dark haloes: evolution, scatter and environment. MNRAS, **321**, 559–575.

Bundy, K., Ellis, R. S., Conselice, C. J., Taylor, J. E., Cooper, M. C., Willmer, C. N. A., Weiner, B. J., Coil, A. L., Noeske, K. G. & Eisenhardt, P. R. M., 2006. The Mass Assembly History of Field Galaxies: Detection of an Evolving Mass Limit for Star-Forming Galaxies. ApJ, **651**, 120–141.

Carlberg, R. G., Yee, H. K. C. & Ellingson, E., 1997. The Average Mass and Light Profiles of Galaxy Clusters. ApJ, **478**, 462.

Cattaneo, A., Dekel, A., Devriendt, J., Guiderdoni, B. & Blaizot, J., 2006. Modelling the galaxy bimodality: shutdown above a critical halo mass. MNRAS, **370**, 1651–1665.

Clyne, J. & Rast, M., 2005. A prototype discovery environment for analyzing and visualizing terascale turbulent fluid flow simulations. In *Proceedings of Visualization and Data Analysis 2005*.

Cohn, J. D. & White, M., 2008. Dark matter halo abundances, clustering and assembly histories at high redshift. MNRAS, **385**, 2025–2033.

Cole, S., Aragon-Salamanca, A., Frenk, C. S., Navarro, J. F. & Zepf, S. E., 1994. A Recipe for Galaxy Formation. MNRAS, **271**, 781.

Cole, S., Norberg, P., Baugh, C. M., Frenk, C. S., Bland-Hawthorn, J. et al., 2001. The 2dF galaxy redshift survey: near-infrared galaxy luminosity functions. MNRAS, **326**, 255–273.

Cole, S., Percival, W. J., Peacock, J. A., Norberg, P., Baugh, C. M. et al., 2005. The 2dF Galaxy Redshift Survey: power-spectrum analysis of the final data set and cosmological implications. MNRAS, **362**, 505–534.

Coles, P. & Lucchin, F., 2002. *Cosmology: The Origin and Evolution of Cosmic Structure*. John Wiley & Sons, 2nd ed.

Colín, P., Valenzuela, O. & Avila-Reese, V., 2008. On the Structure of Dark Matter Halos at the Damping Scale of the Power Spectrum with and without Relict Velocities. ApJ, **673**, 203–214.

Colless, M., Dalton, G., Maddox, S., Sutherland, W., Norberg, P. et al., 2001. The 2dF Galaxy Redshift Survey: spectra and redshifts. MNRAS, **328**, 1039–1063.

Collin, S., Combes, F. & Shlosman, I., eds., 2003. *Active galactic nuclei : from the central engine to host galaxy : proceeding of a conference held in Paris-Meudon Observatory, Meudon, France, 23-27 July 2002*, vol. 290 of *Astronomical Society of the Pacific Conference Series*.

Commerçon, B., Hennebelle, P., Audit, E., Chabrier, G. & Teyssier, R., 2008. Protostellar collapse: a comparison between smoothed particle hydrodynamics and adaptative mesh refinement calculations. A&A, **482**, 371–385.

Conley, A., Guy, J., Sullivan, M., Regnault, N., Astier, P. et al., 2011. Supernova Constraints and Systematic Uncertainties from the First Three Years of the Supernova Legacy Survey. ApJS, **192**, 1.

Conroy, C., Wechsler, R. H. & Kravtsov, A. V., 2006. Modeling Luminosity-dependent Galaxy Clustering through Cosmic Time. ApJ, **647**, 201–214.

Cowie, L. L. & Binney, J., 1977. Radiative regulation of gas flow within clusters of galaxies - A model for cluster X-ray sources. ApJ, **215**, 723–732.

Cowie, L. L., Songaila, A., Hu, E. M. & Cohen, J. G., 1996. New Insight on Galaxy Formation and Evolution From Keck Spectroscopy of the Hawaii Deep Fields. ApJ, **112**, 839.

Croton, D. J., Springel, V., White, S. D. M., De Lucia, G., Frenk, C. S., Gao, L., Jenkins, A., Kauffmann, G., Navarro, J. F. & Yoshida, N., 2006. The many lives of active galactic nuclei: cooling flows, black holes and the luminosities and colours of galaxies. MNRAS, **365**, 11–28.

Cullen, L. & Dehnen, W., 2010. Inviscid smoothed particle hydrodynamics. MNRAS, **408**, 669–683.

Dalcanton, J. J., Spergel, D. N. & Summers, F. J., 1997. The Formation of Disk Galaxies. ApJ, **482**, 659.

Danovich, M., Dekel, A., Hahn, O. & Teyssier, R., 2012. Coplanar streams, pancakes and angular-momentum exchange in high-z disc galaxies. MNRAS, 2777.

Das, S., Marriage, T. A., Ade, P. A. R., Aguirre, P., Amiri, M. et al., 2011. The Atacama Cosmology Telescope: A Measurement of the Cosmic Microwave Background Power Spectrum at 148 and 218 GHz from the 2008 Southern Survey. ApJ, **729**, 62.

Davis, M., Efstathiou, G., Frenk, C. S. & White, S. D. M., 1985. The evolution of large-scale structure in a universe dominated by cold dark matter. ApJ, **292**, 371–394.

de Jong, R. S., 1996. Near-infrared and optical broadband surface photometry of 86 face-on disk dominated galaxies. IV. Using color profiles to study stellar and dust content of galaxies. A&A, **313**, 377–395.

de Jong, R. S. & Lacey, C., 2000. The Local Space Density of SB-SDM Galaxies as Function of Their Scale Size, Surface Brightness, and Luminosity. ApJ, **545**, 781–797.

De Lucia, G., Springel, V., White, S. D. M., Croton, D. & Kauffmann, G., 2006. The formation history of elliptical galaxies. MNRAS, **366**, 499–509.

Dekel, A. & Birnboim, Y., 2006. Galaxy bimodality due to cold flows and shock heating. MNRAS, **368**, 2–20.

Dekel, A., Birnboim, Y., Engel, G., Freundlich, J., Goerdt, T., Mumcuoglu, M., Neistein, E., Pichon, C., Teyssier, R. & Zinger, E., 2009. Cold streams in early massive hot haloes as the main mode of galaxy formation. Nature, **457**, 451–454.

Dekel, A., Devor, J. & Hetzroni, G., 2003. Galactic halo cusp-core: tidal compression in mergers. MNRAS, **341**, 326–342.

Dekel, A. & Woo, J., 2003. Feedback and the fundamental line of low-luminosity low-surface-brightness/dwarf galaxies. MNRAS, **344**, 1131–1144.

Desjacques, V., 2008. Environmental dependence in the ellipsoidal collapse model. MNRAS, **388**, 638–658.

Devriendt, J., Rimes, C., Pichon, C., Teyssier, R., Le Borgne, D., Aubert, D., Audit, E., Colombi, S., Courty, S., Dubois, Y., Prunet, S., Rasera, Y., Slyz, A. & Tweed, D., 2010. The dusty, albeit ultraviolet bright, infancy of galaxies. MNRAS, **403**, L84–L88.

Dickinson, C., Jodrell Bank Observatory Team, Cavendish Astrophysics Team & IAC Team, 2004. Latest CMB results from the Very Small Array, out to $\ell = 1500$. In *American Astronomical Society Meeting Abstracts*, vol. 36 of *Bulletin of the American Astronomical Society*, 1461.

Diemand, J., Kuhlen, M. & Madau, P., 2006. Early Supersymmetric Cold Dark Matter Substructure. ApJ, **649**, 1–13.

Diemand, J., Moore, B. & Stadel, J., 2004. Velocity and spatial biases in cold dark matter subhalo distributions. MNRAS, **352**, 535–546.

Dijkstra, M., Haiman, Z., Rees, M. J. & Weinberg, D. H., 2004. Photoionization Feedback in Low-Mass Galaxies at High Redshift. ApJ, **601**, 666–675.

Donato, F., Gentile, G., Salucci, P., Frigerio Martins, C., Wilkinson, M. I., Gilmore, G., Grebel, E. K., Koch, A. & Wyse, R., 2009. A constant dark matter halo surface density in galaxies. MNRAS, **397**, 1169–1176.

D'Onghia, E., Springel, V., Hernquist, L. & Keres, D., 2010. Substructure Depletion in the Milky Way Halo by the Disk. ApJ, **709**, 1138–1147.

Doroshkevich, A. G., 1970. Spatial structure of perturbations and origin of galactic rotation in fluctuation theory. Astrophysics, **6**, 320–330.

Drory, N. & Fisher, D. B., 2007. A Connection between Bulge Properties and the Bimodality of Galaxies. ApJ, **664**, 640–649.

Dubois, Y. & Teyssier, R., 2008. On the onset of galactic winds in quiescent star forming galaxies. A&A, **477**, 79–94.

Dunkley, J., Komatsu, E., Nolta, M. R., Spergel, D. N., Larson, D. et al., 2009. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Likelihoods and Parameters from the WMAP Data. ApJS, **180**, 306–329.

Dutton, A. A. & van den Bosch, F. C., 2009. The impact of feedback on disc galaxy scaling relations. MNRAS, **396**, 141–164.

Efstathiou, G., Davis, M., White, S. D. M. & Frenk, C. S., 1985. Numerical techniques for large cosmological N-body simulations. ApJS, **57**, 241–260.

Eisenstein, D. J. & Hu, W., 1999. Power Spectra for Cold Dark Matter and Its Variants. ApJ, **511**, 5–15.

Eisenstein, D. J. & Hut, P., 1998. HOP: A New Group-Finding Algorithm for N-body Simulations. ApJ, **498**, 137.

Eisenstein, D. J., Zehavi, I., Hogg, D. W., Scoccimarro, R., Blanton, M. R. et al., 2005. Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. ApJ, **633**, 560–574.

Eke, V. R., Cole, S. & Frenk, C. S., 1996. Cluster evolution as a diagnostic for Omega. MNRAS, **282**, 263–280.

Fabian, A. C. & Nulsen, P. E. J., 1977. Subsonic accretion of cooling gas in clusters of galaxies. MNRAS, **180**, 479–484.

Fabian, A. C., Sanders, J. S., Ettori, S., Taylor, G. B., Allen, S. W., Crawford, C. S., Iwasawa, K., Johnstone, R. M. & Ogle, P. M., 2000. Chandra imaging of the complex X-ray core of the Perseus cluster. MNRAS, **318**, L65–L68.

Fakhouri, O. & Ma, C., 2009. Environmental dependence of dark matter halo growth - I. Halo merger rates. MNRAS, **394**, 1825–1840.

Fakhouri, O. & Ma, C., 2010. Dark matter halo growth - II. Diffuse accretion and its environmental dependence. MNRAS, **401**, 2245–2256.

Fakhouri, O., Ma, C. & Boylan-Kolchin, M., 2010. The merger rates and mass assembly histories of dark matter haloes in the two Millennium simulations. MNRAS, 857.

Fall, S. M. & Efstathiou, G., 1980. Formation and rotation of disc galaxies with haloes. MNRAS, **193**, 189–206.

Fardal, M. A., Katz, N., Gardner, J. P., Hernquist, L., Weinberg, D. H. & Davé, R., 2001. Cooling Radiation and the Lyα Luminosity of Forming Galaxies. ApJ, **562**, 605–617.

Faucher-Giguère, C.-A., Kereš, D., Dijkstra, M., Hernquist, L. & Zaldarriaga, M., 2010. Lyα Cooling Emission from Galaxy Formation. ApJ, **725**, 633–657.

Faucher-Giguère, C.-A., Kereš, D. & Ma, C.-P., 2011. The baryonic assembly of dark matter haloes. MNRAS, **417**, 2982–2999.

Ferrarese, L., 2002. Beyond the Bulge: A Fundamental Relation between Supermassive Black Holes and Dark Matter Halos. ApJ, **578**, 90–97.

Fraschetti, F., Teyssier, R., Ballet, J. & Decourchelle, A., 2010. Simulation of the growth of the 3D Rayleigh-Taylor instability in supernova remnants using an expanding reference frame. A&A, **515**, A104.

Freedman, W. L., Burns, C. R., Phillips, M. M., Wyatt, P., Persson, S. E. et al., 2009. The Carnegie Supernova Project: First Near-Infrared Hubble Diagram to $z \sim 0.7$. ApJ, **704**, 1036–1058.

Fumagalli, M., Prochaska, J. X., Kasen, D., Dekel, A., Ceverino, D. & Primack, J. R., 2011. Absorption-line systems in simulated galaxies fed by cold streams. MNRAS, **418**, 1796–1821.

Gao, L., De Lucia, G., White, S. D. M. & Jenkins, A., 2004. Galaxies and subhaloes in ΛCDM galaxy clusters. MNRAS, **352**, L1–L5.

Gao, L., Springel, V. & White, S. D. M., 2005. The age dependence of halo clustering. MNRAS, **363**, L66–L70.

Gao, L. & White, S. D. M., 2007. Assembly bias in the clustering of dark matter haloes. MNRAS, **377**, L5–L9.

García-Ruiz, I., Sancisi, R. & Kuijken, K., 2002. Neutral hydrogen and optical observations of edge-on galaxies: Hunting for warps. A&A, **394**, 769–789.

Genel, S., Genzel, R., Bouché, N., Naab, T. & Sternberg, A., 2009. The Halo Merger Rate in the Millennium Simulation and Implications for Observed Galaxy Merger Fractions. ApJ, **701**, 2002–2018.

Gnedin, N. Y., 2000. Effect of Reionization on Structure Formation in the Universe. ApJ, **542**, 535–541.

Gnedin, O. Y., Hernquist, L. & Ostriker, J. P., 1999. Tidal Shocking by Extended Mass Distributions. ApJ, **514**, 109–118.

Godunov, S. K. & Ryabenki, V. S., 1964. *Theory of difference schemes - an introduction*. North-Holland Pub. Co.

Goerdt, T., Dekel, A., Sternberg, A., Ceverino, D., Teyssier, R. & Primack, J. R., 2010a. Gravity-driven Lyα blobs from cold streams into galaxies. MNRAS, **407**, 613–631.

Goerdt, T., Moore, B., Read, J. I. & Stadel, J., 2010b. Core Creation in Galaxies and Halos Via Sinking Massive Objects. ApJ, **725**, 1707–1716.

Gogarten, S. M., Dalcanton, J. J., Williams, B. F., Roškar, R., Holtzman, J., Seth, A. C., Dolphin, A., Weisz, D., Cole, A., Debattista, V. P., Gilbert, K. M., Olsen, K., Skillman, E., de Jong, R. S., Karachentsev, I. D. & Quinn, T. R., 2010. The Advanced Camera for Surveys Nearby Galaxy Survey Treasury. V. Radial Star Formation History of NGC 300. ApJ, **712**, 858–874.

Goodwin, S. P., Whitworth, A. P. & Ward-Thompson, D., 2004. Simulating star formation in molecular cloud cores. I. The influence of low levels of turbulence on fragmentation and multiplicity. A&A, **414**, 633–650.

Governato, F., Mayer, L., Wadsley, J., Gardner, J. P., Willman, B., Hayashi, E., Quinn, T., Stadel, J. & Lake, G., 2004. The Formation of a Realistic Disk Galaxy in Λ-dominated Cosmologies. ApJ, **607**, 688–696.

Governato, F., Willman, B., Mayer, L., Brooks, A., Stinson, G., Valenzuela, O., Wadsley, J. & Quinn, T., 2007. Forming disc galaxies in ΛCDM simulations. MNRAS, **374**, 1479–1494.

Gunn, J. E. & Gott, III, J. R., 1972. On the Infall of Matter Into Clusters of Galaxies and Some Effects on Their Evolution. ApJ, **176**, 1.

Guth, A. H., 1981. Inflationary universe: A possible solution to the horizon and flatness problems. PRD, **23**, 347–356.

Guth, A. H. & Pi, S.-Y., 1982. Fluctuations in the new inflationary universe. PRL, **49**, 1110–1113.

Haardt, F. & Madau, P., 1996. Radiative Transfer in a Clumpy Universe. II. The Ultraviolet Extragalactic Background. ApJ, **461**, 20.

Hahn, O., Porciani, C., Carollo, C. M. & Dekel, A., 2007. Properties of dark matter haloes in clusters, filaments, sheets and voids. MNRAS, **375**, 489–499.

Hahn, O., Porciani, C., Dekel, A. & Carollo, C. M., 2009. Tidal effects and the environment dependence of halo assembly. MNRAS, **398**, 1742–1756.

Hamuy, M., Maza, J., Phillips, M. M., Suntzeff, N. B., Wischnjewsky, M. et al., 1993. The 1990 Calan/Tololo Supernova Search. AJ, **106**, 2392–2407.

Hamuy, M., Phillips, M. M., Suntzeff, N. B., Schommer, R. A., Maza, J. & Aviles, R., 1996. The Hubble Diagram of the Calan/Tololo Type IA Supernovae and the Value of HO. AJ, **112**, 2398.

Harker, G., Cole, S., Helly, J., Frenk, C. & Jenkins, A., 2006. A marked correlation function analysis of halo formation times in the Millennium Simulation. MNRAS, **367**, 1039–1049.

Hasinger, G., Miyaji, T. & Schmidt, M., 2005. Luminosity-dependent evolution of soft X-ray selected AGN. New Chandra and XMM-Newton surveys. A&A, **441**, 417–434.

Hatton, S., Devriendt, J. E. G., Ninin, S., Bouchet, F. R., Guiderdoni, B. & Vibert, D., 2003. GALICS- I. A hybrid N-body/semi-analytic model of hierarchical galaxy formation. MNRAS, **343**, 75–106.

Heath, D. J., 1977. The growth of density perturbations in zero pressure Friedmann-Lemaitre universes. MNRAS, **179**, 351–358.

Heckman, T. M., Kauffmann, G., Brinchmann, J., Charlot, S., Tremonti, C. & White, S. D. M., 2004. Present-Day Growth of Black Holes and Bulges: The Sloan Digital Sky Survey Perspective. ApJ, **613**, 109–118.

Hernquist, L., 1987. Performance characteristics of tree codes. ApJS, **64**, 715–734.

Hernquist, L. & Barnes, J. E., 1990. Are some N-body algorithms intrinsically less collisional than others? ApJ, **349**, 562–569.

Hernquist, L. & Katz, N., 1989. TREESPH - A unification of SPH with the hierarchical tree method. ApJS, **70**, 419–446.

Hester, J. A. & Tasitsiomi, A., 2010. Dark Matter Halo Mergers: Dependence on Environment. ApJ, **715**, 342–354.

Ho, L. C., 2002. On the Relationship between Radio Emission and Black Hole Mass in Galactic Nuclei. ApJ, **564**, 120–132.

Ho, L. C., 2005. "Low-State" Black Hole Accretion in Nearby Galaxies. APSS, **300**, 219–225.

Ho, L. C., 2008. Nuclear Activity in Nearby Galaxies. ARAA, **46**, 475–539.

Ho, L. C., Filippenko, A. V. & Sargent, W. L. W., 1997. A Search for "Dwarf" Seyfert Nuclei. V. Demographics of Nuclear Activity in Nearby Galaxies. ApJ, **487**, 568.

Hockney, R. W. & Eastwood, J. W., 1981. *Computer Simulation Using Particles*. McGraw-Hill International Book Co.

Hoeflich, P. & Khokhlov, A., 1996. Explosion Models for Type IA Supernovae: A Comparison with Observed Light Curves, Distances, H_0, and q_0. ApJ, **457**, 500.

Hoekstra, H., Franx, M., Kuijken, K. & Squires, G., 1998. Weak Lensing Analysis of CL 1358+62 Using Hubble Space Telescope Observations. ApJ, **504**, 636.

Hopkins, P. F., Croton, D., Bundy, K., Khochfar, S., van den Bosch, F., Somerville, R. S., Wetzel, A., Keres, D., Hernquist, L., Stewart, K., Younger, J. D., Genel, S. & Ma, C., 2010. Mergers in ΛCDM: Uncertainties in Theoretical Predictions and Interpretations of the Merger Rate. ApJ, **724**, 915–945.

House, E. L., Brook, C. B., Gibson, B. K., Sánchez-Blázquez, P., Courty, S., Few, C. G., Governato, F., Kawata, D., Roškar, R., Steinmetz, M., Stinson, G. S. & Teyssier, R., 2011. Disc heating: comparing the Milky Way with cosmological simulations. MNRAS, **415**, 2652–2664.

Hoyle, F., 1951. The Origin of the Rotations of the Galaxies. In *Problems of Cosmical Aerodynamics*, 195.

Huang, J.-S., Glazebrook, K., Cowie, L. L. & Tinney, C., 2003. The Hawaii+Anglo-Australian Observatory K-Band Galaxy Redshift Survey. I. The Local K-Band Luminosity Function. ApJ, **584**, 203–209.

Iannuzzi, F. & Dolag, K., 2011. Adaptive gravitational softening in GADGET. MNRAS, 1346.

Jenkins, A., Frenk, C. S., White, S. D. M., Colberg, J. M., Cole, S., Evrard, A. E., Couchman, H. M. P. & Yoshida, N., 2001. The mass function of dark matter haloes. MNRAS, **321**, 372–384.

Jing, Y. P. & Suto, Y., 2002. Triaxial Modeling of Halo Density Profiles with High-Resolution N-body Simulations. ApJ, **574**, 538–553.

Jones, W. C., Ade, P. A. R., Bock, J. J., Bond, J. R., Borrill, J. et al., 2006. A Measurement of the Angular Power Spectrum of the CMB Temperature Anisotropy from the 2003 Flight of BOOMERANG. ApJ, **647**, 823–832.

Kamionkowski, M. & Liddle, A. R., 2000. The Dearth of Halo Dwarf Galaxies: Is There Power on Short Scales? PRL, **84**, 4525–4528.

Katz, N., 1992. Dissipational galaxy formation. II - Effects of star formation. ApJ, **391**, 502–517.

Kaufmann, T., Mayer, L., Wadsley, J., Stadel, J. & Moore, B., 2007. Angular momentum transport and disc morphology in smoothed particle hydrodynamics simulations of galaxy formation. MNRAS, **375**, 53–67.

Kay, S. T., Pearce, F. R., Jenkins, A., Frenk, C. S., White, S. D. M., Thomas, P. A. & Couchman, H. M. P., 2000. Parameter tests within cosmological simulations of galaxy formation. MNRAS, **316**, 374–394.

Keres, D., Vogelsberger, M., Sijacki, D., Springel, V. & Hernquist, L., 2011. Moving mesh cosmology: characteristics of galaxies and haloes. Preprint (arXiv:1109.4638).

Kereš, D., Katz, N., Fardal, M., Davé, R. & Weinberg, D. H., 2009. Galaxies in a simulated ΛCDM Universe - I. Cold mode and hot cores. MNRAS, **395**, 160–179.

Kereš, D., Katz, N., Weinberg, D. H. & Davé, R., 2005. How do galaxies get their gas? MNRAS, **363**, 2–28.

Kimm, T., Devriendt, J., Slyz, A., Pichon, C., Kassin, S. A. & Dubois, Y., 2011. The angular momentum of baryons and dark matter halos revisited. Preprint (arXiv:1106.0538).

Klypin, A., Gottlöber, S., Kravtsov, A. V. & Khokhlov, A. M., 1999a. Galaxies in N-body Simulations: Overcoming the Overmerging Problem. ApJ, **516**, 530–551.

Klypin, A., Kravtsov, A. V., Valenzuela, O. & Prada, F., 1999b. Where Are the Missing Galactic Satellites? ApJ, **522**, 82–92.

Klypin, A. A. & Shandarin, S. F., 1983. Three-dimensional numerical model of the formation of large-scale structure in the Universe. MNRAS, **204**, 891–907.

Knebe, A., Devriendt, J. E. G., Mahmood, A. & Silk, J., 2002. Merger histories in warm dark matter structure formation scenarios. MNRAS, **329**, 813–828.

Knebe, A., Knollmann, S. R., Muldrew, S. I., Pearce, F. R., Aragon-Calvo, M. A. et al., 2011. Haloes gone MAD: The Halo-Finder Comparison Project. MNRAS, **415**, 2293–2318.

Kochanek, C. S., Pahre, M. A., Falco, E. E., Huchra, J. P., Mader, J., Jarrett, T. H., Chester, T., Cutri, R. & Schneider, S. E., 2001. The K-Band Galaxy Luminosity Function. ApJ, **560**, 566–579.

Kormendy, J. & Bender, R., 2011. Supermassive black holes do not correlate with dark matter haloes of galaxies. Nature, **469**, 377–380.

Kravtsov, A. V., Gnedin, O. Y. & Klypin, A. A., 2004. The Tumultuous Lives of Galactic Dwarfs and the Missing Satellites Problem. ApJ, **609**, 482–497.

Kravtsov, A. V., Klypin, A. A. & Khokhlov, A. M., 1997. Adaptive Refinement Tree: A New High-Resolution N-body Code for Cosmological Simulations. ApJS, **111**, 73.

Krumholz, M. R. & Tan, J. C., 2007. Slow Star Formation in Dense Gas: Evidence and Implications. ApJ, **654**, 304–315.

Kuo, C. L., Ade, P. A. R., Bock, J. J., Cantalupo, C., Daub, M. D., Goldstein, J., Holzapfel, W. L., Lange, A. E., Lueker, M., Newcomb, M., Peterson, J. B., Ruhl, J., Runyan, M. C. & Torbet, E., 2004. High-Resolution Observations of the Cosmic Microwave Background Power Spectrum with ACBAR. ApJ, **600**, 32–51.

Lacey, C. & Cole, S., 1993. Merger rates in hierarchical models of galaxy formation. MNRAS, **262**, 627–649.

Lacey, C. & Cole, S., 1994. Merger Rates in Hierarchical Models of Galaxy Formation - Part Two - Comparison with N-body Simulations. MNRAS, **271**, 676.

Landy, S. D. & Szalay, A. S., 1993. Bias and variance of angular correlation functions. ApJ, **412**, 64–71.

Lemson, G. & Kauffmann, G., 1999. Environmental influences on dark matter haloes and consequences for the galaxies within them. MNRAS, **302**, 111–117.

Liddle, A., 2003. *An Introduction to Modern Cosmology*. Wiley, 2nd ed.

Linde, A. D., 1982. Scalar field fluctuations in the expanding universe and the new inflationary universe scenario. PLB, **116**, 335–339.

MacArthur, L. A., Courteau, S., Bell, E. & Holtzman, J. A., 2004. Structure of Disk-dominated Galaxies. II. Color Gradients and Stellar Population Models. ApJS, **152**, 175–199.

Maciejewski, M., Colombi, S., Springel, V., Alard, C. & Bouchet, F. R., 2009. Phase-space structures - II. Hierarchical Structure Finder. MNRAS, **396**, 1329–1348.

Maller, A. H. & Dekel, A., 2002. Towards a resolution of the galactic spin crisis: mergers, feedback and spin segregation. MNRAS, **335**, 487–498.

Mantz, A., Allen, S. W., Rapetti, D. & Ebeling, H., 2010. The observed growth of massive galaxy clusters - I. Statistical methods and cosmological constraints. MNRAS, **406**, 1759–1772.

Marinoni, C. & Hudson, M. J., 2002. The Mass-to-Light Function of Virialized Systems and the Relationship between Their Optical and X-Ray Properties. ApJ, **569**, 101–111.

Marriage, T. A., Acquaviva, V., Ade, P. A. R., Aguirre, P., Amiri, M. et al., 2011. The Atacama Cosmology Telescope: Sunyaev-Zel'dovich-Selected Galaxy Clusters at 148 GHz in the 2008 Survey. ApJ, **737**, 61.

Martel, H. & Shapiro, P. R., 1998. A convenient set of comoving cosmological variables and their application. MNRAS, **297**, 467–485.

Mashchenko, S., Couchman, H. M. P. & Wadsley, J., 2006. The removal of cusps from galaxy centres by stellar feedback in the early Universe. Nature, **442**, 539–542.

Mateo, M. L., 1998. Dwarf Galaxies of the Local Group. ARAA, **36**, 435–506.

Matsuda, Y., Yamada, T., Hayashino, T., Yamauchi, R., Nakamura, Y., Morimoto, N., Ouchi, M., Ono, Y., Kousai, K., Nakamura, E., Horie, M., Fujii, T., Umemura, M. & Mori, M., 2011. The Subaru Lyα blob survey: a sample of 100-kpc Lyα blobs at $z = 3$. MNRAS, **410**, L13–L17.

Maughan, B. J., Giles, P. A., Randall, S. W., Jones, C. & Forman, W. R., 2012. Self-similar scaling and evolution in the galaxy cluster X-ray luminosity-temperature relation. MNRAS, **421**, 1583–1602.

Maulbetsch, C., Avila-Reese, V., Colín, P., Gottlöber, S., Khalatyan, A. & Steinmetz, M., 2007. The Dependence of the Mass Assembly History of Cold Dark Matter Halos on Environment. ApJ, **654**, 53–65.

Mayer, L., Governato, F. & Kaufmann, T., 2008. The formation of disk galaxies in computer simulations. ASL, **1**, 7–27.

McBride, J., Fakhouri, O. & Ma, C., 2009. Mass accretion rates and histories of dark matter haloes. MNRAS, **398**, 1858–1868.

McNamara, B. R., Wise, M., Nulsen, P. E. J., David, L. P., Sarazin, C. L., Bautz, M., Markevitch, M., Vikhlinin, A., Forman, W. R., Jones, C. & Harris, D. E., 2000. Chandra X-Ray Observations of the Hydra A Cluster: An Interaction between the Radio Source and the X-Ray-emitting Gas. ApJL, **534**, L135–L138.

Merritt, D., 1987. The distribution of dark matter in the coma cluster. ApJ, **313**, 121–135.

Merritt, D., 1996. Optimal Smoothing for N-body Codes. AJ, **111**, 2462.

Miller, L., Percival, W. J., Croom, S. M. & Babić, A., 2006. The cosmological history of accretion onto dark halos and supermassive black holes. A&A, **459**, 43–54.

Mo, H. J. & Mao, S., 2004. Galaxy formation in pre-processed dark haloes. MNRAS, **353**, 829–840.

Mo, H. J., Mao, S. & White, S. D. M., 1998. The formation of galactic discs. MNRAS, **295**, 319–336.

Monaghan, J. J. & Lattanzio, J. C., 1985. A refined particle method for astrophysical problems. A&A, **149**, 135–143.

Moore, B., Ghigna, S., Governato, F., Lake, G., Quinn, T., Stadel, J. & Tozzi, P., 1999a. Dark Matter Substructure within Galactic Halos. ApJL, **524**, L19–L22.

Moore, B., Quinn, T., Governato, F., Stadel, J. & Lake, G., 1999b. Cold collapse and the core catastrophe. MNRAS, **310**, 1147–1152.

Mori, M. & Burkert, A., 2000. Gas Stripping of Dwarf Galaxies in Clusters of Galaxies. ApJ, **538**, 559–568.

Mortlock, A., Conselice, C. J., Bluck, A. F. L., Bauer, A. E., Grützbauch, R., Buitrago, F. & Ownsworth, J., 2011. A deep probe of the galaxy stellar mass functions at $z \sim 1 - 3$ with the GOODS NICMOS Survey. MNRAS, **413**, 2845–2859.

Muñoz-Mateos, J. C., Gil de Paz, A., Boissier, S., Zamorano, J., Jarrett, T., Gallego, J. & Madore, B. F., 2007. Specific Star Formation Rate Profiles in Nearby Spiral Galaxies: Quantifying the Inside-Out Formation of Disks. ApJJ, **658**, 1006–1026.

Murante, G., Calabrese, M., De Lucia, G., Monaco, P., Borgani, S. & Dolag, K., 2012. A Warm Mode of Gas Accretion on Forming Galaxies. ApJL, **749**, L34.

Nagai, D. & Kravtsov, A. V., 2005. The Radial Distribution of Galaxies in Λ Cold Dark Matter Clusters. ApJ, **618**, 557–568.

Narayan, R., 2002. Why Do AGN Lighthouses Switch Off? In M. Gilfanov, R. Sunyeav, & E. Churazov, ed., *Lighthouses of the Universe: The Most Luminous Celestial Objects and Their Use for Cosmology*, 405.

Navarro, J. F. & Benz, W., 1991. Dynamics of cooling gas in galactic dark halos. ApJ, **380**, 320–329.

Navarro, J. F., Frenk, C. S. & White, S. D. M., 1995a. Simulations of X-ray clusters. MNRAS, **275**, 720–740.

Navarro, J. F., Frenk, C. S. & White, S. D. M., 1995b. The assembly of galaxies in a hierarchically clustering universe. MNRAS, **275**, 56–66.

Navarro, J. F., Ludlow, A., Springel, V., Wang, J., Vogelsberger, M., White, S. D. M., Jenkins, A., Frenk, C. S. & Helmi, A., 2010. The diversity and similarity of simulated cold dark matter haloes. MNRAS, **402**, 21–34.

Navarro, J. F. & Steinmetz, M., 2000. Dark Halo and Disk Galaxy Scaling Laws in Hierarchical Universes. ApJ, **538**, 477–488.

Navarro, J. F. & White, S. D. M., 1994. Simulations of dissipative galaxy formation in hierarchically clustering universes - 2. Dynamics of the baryonic component in galactic haloes. MNRAS, **267**, 401–412.

Neistein, E. & Dekel, A., 2008. Merger rates of dark matter haloes. MNRAS, **388**, 1792–1802.

Ocvirk, P., Pichon, C. & Teyssier, R., 2008. Bimodal gas accretion in the Horizon-MareNostrum galaxy formation simulation. MNRAS, **390**, 1326–1338.

Okamoto, T., Eke, V. R., Frenk, C. S. & Jenkins, A., 2005. Effects of feedback on the morphology of galaxy discs. MNRAS, **363**, 1299–1314.

Okamoto, T., Frenk, C. S., Jenkins, A. & Theuns, T., 2010. The properties of satellite galaxies in simulations of galaxy formation. MNRAS, **406**, 208–222.

Ostlie, D. A. & Carroll, B. W., 2007. *An Introduction to Modern Stellar Astrophysics*. Pearson Addison-Wesley, 2nd ed.

Ostriker, E. C. & Binney, J. J., 1989. Warped and tilted galactic discs. MNRAS, **237**, 785–798.

Peacock, J. A., 1999. *Cosmological Physics*. Cambridge University Press.

Peacock, J. A. & Heavens, A. F., 1990. Alternatives to the Press–Schechter cosmological mass function. MNRAS, **243**, 133–143.

Peebles, P. J. E., 1969. Origin of the Angular Momentum of Galaxies. ApJ, **155**, 393.

Peebles, P. J. E., 1993. *Principles of Physical Cosmology*. Princeton University Press.

Penston, M. V., 1969. Dynamics of self-gravitating gaseous spheres - III. Analytical results in the free-fall of isothermal cases. MNRAS, **144**, 425.

Percival, W. J., 2005. Cosmological structure formation in a homogeneous dark energy background. A&A, **443**, 819–830.

Percival, W. J., Baugh, C. M., Bland-Hawthorn, J., Bridges, T., Cannon, R. et al., 2001. The 2dF Galaxy Redshift Survey: the power spectrum and the matter content of the Universe. MNRAS, **327**, 1297–1306.

Percival, W. J., Cole, S., Eisenstein, D. J., Nichol, R. C., Peacock, J. A., Pope, A. C. & Szalay, A. S., 2007. Measuring the Baryon Acoustic Oscillation scale using the Sloan Digital Sky Survey and 2dF Galaxy Redshift Survey. MNRAS, **381**, 1053–1066.

Percival, W. J., Reid, B. A., Eisenstein, D. J., Bahcall, N. A., Budavari, T. et al., 2010. Baryon acoustic oscillations in the Sloan Digital Sky Survey Data Release 7 galaxy sample. MNRAS, **401**, 2148–2168.

Percival, W. J., Scott, D., Peacock, J. A. & Dunlop, J. S., 2003. The clustering of halo mergers. MNRAS, **338**, L31–L35.

Perlmutter, S., Aldering, G., Goldhaber, G., Knop, R. A., Nugent, P. et al., 1999. Measurements of Omega and Lambda from 42 High-Redshift Supernovae. ApJ, **517**, 565–586.

Perlmutter, S., Gabi, S., Goldhaber, G., Goobar, A., Groom, D. E. et al., 1997. Measurements of the Cosmological Parameters Omega and Lambda from the First Seven Supernovae at $z \geq 0.35$. ApJ, **483**, 565.

Perlmutter, S., Pennypacker, C. R., Goldhaber, G., Goobar, A., Muller, R. A. et al., 1995. A supernova at $z = 0.458$ and implications for measuring the cosmological deceleration. ApJL, **440**, L41–L44.

Peterson, J. R., Kahn, S. M., Paerels, F. B. S., Kaastra, J. S., Tamura, T., Bleeker, J. A. M., Ferrigno, C. & Jernigan, J. G., 2003. High-Resolution X-Ray Spectroscopic Constraints on Cooling-Flow Models for Clusters of Galaxies. ApJ, **590**, 207–224.

Pfenniger, D. & Friedli, D., 1993. Computational issues connected with 3D N-body simulations. A&A, **270**, 561–572.

Phillips, M. M., 1993. The absolute magnitudes of Type IA supernovae. ApJL, **413**, L105–L108.

Pichon, C., Gay, C., Pogosyan, D., Prunet, S., Sousbie, T., Colombi, S., Slyz, A. & Devriendt, J., 2010. The Skeleton: Connecting Large Scale Structures to Galaxy Formation. In J.-M. Alimi & A. Fuözfa, ed., *American Institute of Physics Conference Series*, vol. 1241 of *American Institute of Physics Conference Series*, 1108–1117.

Pichon, C., Pogosyan, D., Kimm, T., Slyz, A., Devriendt, J. & Dubois, Y., 2011. Rigging dark haloes: why is hierarchical galaxy formation consistent with the inside-out build-up of thin discs? MNRAS, 1739.

Pogosyan, D., Bond, J. R. & Kofman, L., 1998. Origin and observables of the cosmic web. JRASC, **92**, 313.

Pogosyan, D., Pichon, C., Gay, C., Prunet, S., Cardoso, J. F., Sousbie, T. & Colombi, S., 2009. The local theory of the cosmic skeleton. MNRAS, **396**, 635–667.

Pontzen, A. & Governato, F., 2012. How supernova feedback turns dark matter cusps into cores. MNRAS, 2641.

Powell, L. C., Slyz, A. & Devriendt, J., 2011. The impact of supernova-driven winds on stream-fed protogalaxies. MNRAS, **414**, 3671–3689.

Power, C., Navarro, J. F., Jenkins, A., Frenk, C. S., White, S. D. M., Springel, V., Stadel, J. & Quinn, T., 2003. The inner structure of ΛCDM haloes - I. A numerical convergence study. MNRAS, **338**, 14–34.

Press, W. H. & Schechter, P., 1974. Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. ApJ, **187**, 425–438.

Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P., 1992. *Numerical recipes in C. The art of scientific computing*. Cambridge University Press, 2nd ed.

Price, D. J., 2008. Modelling discontinuities and Kelvin Helmholtz instabilities in SPH. Journal of Computational Physics, **227**, 10040–10057.

Price, D. J., 2012. Smoothed particle hydrodynamics and magnetohydrodynamics. Journal of Computational Physics, **231**, 759–794.

Prunet, S., Pichon, C., Aubert, D., Pogosyan, D., Teyssier, R. & Gottloeber, S., 2008. Initial Conditions For Large Cosmological Simulations. ApJS, **178**, 179–188.

Quinn, T., Katz, N., Stadel, J. & Lake, G., 1997. Time stepping N-body simulations. Preprint (arXiv:astro-ph/9710043).

Rasera, Y. & Teyssier, R., 2006. The history of the baryon budget. Cosmic logistics in a hierarchical universe. A&A, **445**, 1–27.

Readhead, A. C. S., Mason, B. S., Contaldi, C. R., Pearson, T. J., Bond, J. R. et al., 2004. Extended Mosaic Observations with the Cosmic Background Imager. ApJ, **609**, 498–512.

Rees, M. J. & Ostriker, J. P., 1977. Cooling, dynamics and fragmentation of massive gas clouds - Clues to the masses and radii of galaxies and clusters. MNRAS, **179**, 541–559.

Riess, A. G., Filippenko, A. V., Challis, P., Clocchiatti, A., Diercks, A. et al., 1998. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. AJ, **116**, 1009–1038.

Riess, A. G., Macri, L., Casertano, S., Lampeitl, H., Ferguson, H. C., Filippenko, A. V., Jha, S. W., Li, W. & Chornock, R., 2011. A 3% Solution: Determination of the Hubble Constant with the Hubble Space Telescope and Wide Field Camera 3. ApJ, **730**, 119.

Riess, A. G., Press, W. H. & Kirshner, R. P., 1996. A Precise Distance Indicator: Type IA Supernova Multicolor Light-Curve Shapes. ApJ, **473**, 88.

Riess, A. G., Strolger, L.-G., Casertano, S., Ferguson, H. C., Mobasher, B. et al., 2007. New Hubble Space Telescope Discoveries of Type Ia Supernovae at $z \geq 1$: Narrowing Constraints on the Early Behavior of Dark Energy. ApJ, **659**, 98–121.

Riess, A. G., Strolger, L.-G., Tonry, J., Casertano, S., Ferguson, H. C. et al., 2004. Type Ia Supernova Discoveries at $z > 1$ from the Hubble Space Telescope: Evidence for Past Deceleration and Constraints on Dark Energy Evolution. ApJ, **607**, 665–687.

Rigby, E. E., Best, P. N., Brookes, M. H., Peacock, J. A., Dunlop, J. S., Röttgering, H. J. A., Wall, J. V. & Ker, L., 2011. The luminosity-dependent high-redshift turnover in the steep spectrum radio luminosity function: clear evidence for downsizing in the radio-AGN population. MNRAS, **416**, 1900–1915.

Roe, P. L., 1981. Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes. Journal of Computational Physics, **43**, 357.

Rosdahl, J. & Blaizot, J., 2012. Extended Lyα emission from cold accretion streams. MNRAS, 2837.

Roškar, R., Debattista, V. P., Brooks, A. M., Quinn, T. R., Brook, C. B., Governato, F., Dalcanton, J. J. & Wadsley, J., 2010. Misaligned angular momentum in hydrodynamic cosmological simulations: warps, outer discs and thick discs. MNRAS, **408**, 783–796.

Rubin, V. C., Ford, W. K. J. & . Thonnard, N., 1980. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 $R = 4\,\mathrm{kpc}$ to UGC 2885 $R = 122\,\mathrm{kpc}$. ApJ, **238**, 471–487.

Sales, L. V., Navarro, J. F., Cooper, A. P., White, S. D. M., Frenk, C. S. & Helmi, A., 2011. Clues to the 'Magellanic Galaxy' from cosmological simulations. MNRAS, **418**, 648–658.

Salmon, J. K. & Warren, M. S., 1994. Skeletons from the treecode closet. Journal of Computational Physics, **111**, 136–155.

Sánchez, A. G., Baugh, C. M., Percival, W. J., Peacock, J. A., Padilla, N. D., Cole, S., Frenk, C. S. & Norberg, P., 2006. Cosmological parameters from cosmic microwave background measurements and the final 2dF Galaxy Redshift Survey power spectrum. MNRAS, **366**, 189–207.

Sánchez, A. G., Crocce, M., Cabré, A., Baugh, C. M. & Gaztañaga, E., 2009. Cosmological parameter constraints from SDSS luminous red galaxies: a new treatment of large-scale clustering. MNRAS, **400**, 1643–1664.

Sandvik, H. B., Möller, O., Lee, J. & White, S. D. M., 2007. Why does the clustering of haloes depend on their formation history? MNRAS, **377**, 234–244.

Scannapieco, C., Wadepuhl, M., Parry, O. H., Navarro, J. F., Jenkins, A. et al., 2012. The Aquila comparison project: the effects of feedback and numerical methods on simulations of galaxy formation. MNRAS, **423**, 1726–1749.

Scannapieco, C., White, S. D. M., Springel, V. & Tissera, P. B., 2009. The formation and survival of discs in a ΛCDM universe. MNRAS, **396**, 696–708.

Scannapieco, C., White, S. D. M., Springel, V. & Tissera, P. B., 2011. Formation history, structure and dynamics of discs and spheroids in simulated Milky Way mass galaxies. MNRAS, **417**, 154–171.

Schawinski, K., Lintott, C. J., Thomas, D., Kaviraj, S., Viti, S., Silk, J., Maraston, C., Sarzi, M., Yi, S. K., Joo, S., Daddi, E., Bayet, E., Bell, T. & Zuntz, J., 2009. Destruction of Molecular Gas Reservoirs in Early-Type Galaxies by Active Galactic Nucleus Feedback. ApJ, **690**, 1672–1680.

Schawinski, K., Thomas, D., Sarzi, M., Maraston, C., Kaviraj, S., Joo, S., Yi, S. K. & Silk, J., 2007. Observational evidence for AGN feedback in early-type galaxies. MNRAS, **382**, 1415–1431.

Schmidt, M., 1959. The Rate of Star Formation. ApJ, **129**, 243.

Seljak, U., Makarov, A., McDonald, P. & Trac, H., 2006. Can Sterile Neutrinos Be the Dark Matter? PRL, **97**(19), 191303.

Sellwood, J. A., 1987. The art of N-body building. ARAA, **25**, 151–186.

Shen, J., Abel, T., Mo, H. J. & Sheth, R. K., 2006. An Excursion Set Model of the Cosmic Web: The Abundance of Sheets, Filaments, and Halos. ApJ, **645**, 783–791.

Shen, J. & Sellwood, J. A., 2006. Galactic warps induced by cosmic infall. MNRAS, **370**, 2–14.

Sheth, R. K., Mo, H. J. & Tormen, G., 2001. Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes. MNRAS, **323**, 1–12.

Sheth, R. K. & Tormen, G., 1999. Large-scale bias and the peak background split. MNRAS, **308**, 119–126.

Sheth, R. K. & Tormen, G., 2002. An excursion set model of hierarchical clustering: ellipsoidal collapse and the moving barrier. MNRAS, **329**, 61–75.

Sheth, R. K. & Tormen, G., 2004. On the environmental dependence of halo formation. MNRAS, **350**, 1385–1390.

Simon, J. D., Bolatto, A. D., Leroy, A., Blitz, L. & Gates, E. L., 2005. High-Resolution Measurements of the Halos of Four Dark Matter-Dominated Galaxies: Deviations from a Universal Density Profile. ApJ, **621**, 757–776.

Smoot, G. F., Bennett, C. L., Kogut, A., Wright, E. L., Aymon, J. et al., 1992. Structure in the COBE differential microwave radiometer first-year maps. ApJL, **396**, L1–L5.

Somerville, R. S., Hopkins, P. F., Cox, T. J., Robertson, B. E. & Hernquist, L., 2008. A semi-analytic model for the co-evolution of galaxies, black holes and active galactic nuclei. MNRAS, **391**, 481–506.

Somerville, R. S. & Kolatt, T. S., 1999. How to plant a merger tree. MNRAS, **305**, 1–14.

Somerville, R. S. & Primack, J. R., 1999. Semi-analytic modelling of galaxy formation: the local Universe. MNRAS, **310**, 1087–1110.

Sommer-Larsen, J. & Dolgov, A., 2001. Formation of Disk Galaxies: Warm Dark Matter and the Angular Momentum Problem. ApJ, **551**, 608–623.

Sommer-Larsen, J., Götz, M. & Portinari, L., 2003. Galaxy Formation: Cold Dark Matter, Feedback, and the Hubble Sequence. ApJ, **596**, 47–66.

Sousbie, T., 2011. The persistent cosmic web and its filamentary structure - I. Theory and implementation. MNRAS, **414**, 350–383.

Spergel, D. N., Bean, R., Doré, O., Nolta, M. R., Bennett, C. L. et al., 2007. Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications for Cosmology. ApJS, **170**, 377–408.

Spergel, D. N., Verde, L., Peiris, H. V., Komatsu, E., Nolta, M. R., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Tucker, G. S., Weiland, J. L., Wollack, E. & Wright, E. L., 2003. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters. ApJS, **148**, 175–194.

Springel, V., 2005. The cosmological simulation code GADGET-2. MNRAS, **364**, 1105–1134.

Springel, V., 2010. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. MNRAS, **401**, 791–851.

Springel, V., Wang, J., Vogelsberger, M., Ludlow, A., Jenkins, A., Helmi, A., Navarro, J. F., Frenk, C. S. & White, S. D. M., 2008. The Aquarius Project: the subhaloes of galactic haloes. MNRAS, **391**, 1685–1711.

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. Simulations of the formation, evolution and clustering of galaxies and quasars. Nature, **435**, 629–636.

Springel, V., White, S. D. M., Tormen, G. & Kauffmann, G., 2001a. Populating a cluster of galaxies - I. Results at $z = 0$. MNRAS, **328**, 726–750.

Springel, V., Yoshida, N. & White, S. D. M., 2001b. GADGET: a code for collisionless and gasdynamical cosmological simulations. New Astronomy, **6**, 79–117.

Stadel, J., Potter, D., Moore, B., Diemand, J., Madau, P., Zemp, M., Kuhlen, M. & Quilis, V., 2009. Quantifying the heart of darkness with GHALO - a multibillion particle simulation of a galactic halo. MNRAS, **398**, L21–L25.

Steidel, C. C., Adelberger, K. L., Shapley, A. E., Pettini, M., Dickinson, M. & Giavalisco, M., 2000. Lyα Imaging of a Proto-Cluster Region at $\langle z \rangle = 3.09$. ApJ, **532**, 170–182.

Steidel, C. C., Erb, D. K., Shapley, A. E., Pettini, M., Reddy, N., Bogosavljević, M., Rudie, G. C. & Rakic, O., 2010. The Structure and Kinematics of the Circumgalactic Medium from Far-ultraviolet Spectra of $z \simeq 2 - 3$ Galaxies. ApJ, **717**, 289–322.

Stinson, G. S., Bailin, J., Couchman, H., Wadsley, J., Shen, S., Nickerson, S., Brook, C. & Quinn, T., 2010. Cosmological galaxy formation simulations using smoothed particle hydrodynamics. MNRAS, **408**, 812–826.

Strigari, L. E., Kaplinghat, M. & Bullock, J. S., 2007. Dark matter halos with cores from hierarchical structure formation. PRD, **75**(6), 061303.

Sullivan, M., Guy, J., Conley, A., Regnault, N., Astier, P. et al., 2011. SNLS3: Constraints on Dark Energy Combining the Supernova Legacy Survey Three-year Data with Other Probes. ApJ, **737**, 102.

Sutherland, R. S. & Dopita, M. A., 1993. Cooling functions for low-density astrophysical plasmas. ApJS, **88**, 253–327.

Swaters, R. A., Madore, B. F., van den Bosch, F. C. & Balcells, M., 2003. The Central Mass Distribution in Dwarf and Low Surface Brightness Galaxies. ApJ, **583**, 732–751.

Tasker, E. J., Brunino, R., Mitchell, N. L., Michielsen, D., Hopton, S., Pearce, F. R., Bryan, G. L. & Theuns, T., 2008. A test suite for quantitative comparison of hydrodynamic codes in astrophysics. MNRAS, **390**, 1267–1281.

Taylor, J. E. & Babul, A., 2004. The evolution of substructure in galaxy, group and cluster haloes - I. Basic dynamics. MNRAS, **348**, 811–830.

Tegmark, M., Eisenstein, D. J., Strauss, M. A., Weinberg, D. H., Blanton, M. R. et al., 2006. Cosmological constraints from the SDSS luminous red galaxies. PRD, **74**(12), 123507.

Tegmark, M., Hamilton, A. J. S. & Xu, Y., 2002. The power spectrum of galaxies in the 2dF 100k redshift survey. MNRAS, **335**, 887–908.

Teyssier, R., 2002. Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES. A&A, **385**, 337–364.

Teyssier, R., Pires, S., Prunet, S., Aubert, D., Pichon, C., Amara, A., Benabed, K., Colombi, S., Refregier, A. & Starck, J.-L., 2009. Full-sky weak-lensing simulation with 70 billion particles. A&A, **497**, 335–341.

Thoul, A. A. & Weinberg, D. H., 1996. Hydrodynamic Simulations of Galaxy Formation. II. Photoionization and the Formation of Low-Mass Galaxies. ApJ, **465**, 608.

Tinker, J., Kravtsov, A. V., Klypin, A., Abazajian, K., Warren, M., Yepes, G., Gottlöber, S. & Holz, D. E., 2008. Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality. ApJ, **688**, 709–728.

Tonini, C., Lapi, A. & Salucci, P., 2006. Angular Momentum Transfer in Dark Matter Halos: Erasing the Cusp. ApJ, **649**, 591–598.

Tormen, G., 1997. The rise and fall of satellites in galaxy clusters. MNRAS, **290**, 411–421.

Tormen, G., Bouchet, F. R. & White, S. D. M., 1997. The structure and dynamical evolution of dark matter haloes. MNRAS, **286**, 865–884.

Tormen, G., Diaferio, A. & Syer, D., 1998. Survival of substructure within dark matter haloes. MNRAS, **299**, 728–742.

Toro, E. F., 2009. *Riemann Solvers and Numerical Methods for Fluid Dynamics*. Springer, 3rd ed.

Tremaine, S., Gebhardt, K., Bender, R., Bower, G., Dressler, A., Faber, S. M., Filippenko, A. V., Green, R., Grillmair, C., Ho, L. C., Kormendy, J., Lauer, T. R., Magorrian, J., Pinkney, J. & Richstone, D., 2002. The Slope of the Black Hole Mass versus Velocity Dispersion Correlation. ApJ, **574**, 740–753.

Tully, R. B. & Fisher, J. R., 1977. A new method of determining distances to galaxies. A&A, **54**, 661–673.

Tweed, D., Devriendt, J., Blaizot, J., Colombi, S. & Slyz, A., 2009. Building merger trees from cosmological N-body simulations. Towards improving galaxy formation models using subhaloes. A&A, **506**, 647–660.

Vale, A. & Ostriker, J. P., 2006. The non-parametric model for linking galaxy luminosity with halo/subhalo mass. MNRAS, **371**, 1173–1187.

van Albada, T. S., Bahcall, J. N., Begeman, K. & Sancisi, R., 1985. Distribution of dark matter in the spiral galaxy NGC 3198. ApJ, **295**, 305–313.

van de Voort, F., Schaye, J., Booth, C. M., Haas, M. R. & Dalla Vecchia, C., 2011. The rates and modes of gas accretion on to galaxies and their gaseous haloes. MNRAS, **414**, 2458–2478.

van den Bosch, F. C., 2002. The universal mass accretion history of cold dark matter haloes. MNRAS, **331**, 98–110.

van den Bosch, F. C., Abel, T., Croft, R. A. C., Hernquist, L. & White, S. D. M., 2002. The Angular Momentum of Gas in Protogalaxies. I. Implications for the Formation of Disk Galaxies. ApJ, **576**, 21–35.

Viel, M., Becker, G. D., Bolton, J. S., Haehnelt, M. G., Rauch, M. & Sargent, W. L. W., 2008. How Cold Is Cold Dark Matter? Small-Scales Constraints from the Flux Power Spectrum of the High-Redshift Lyman-α Forest. PRL, **100**(4), 041304.

Volonteri, M., Natarajan, P. & Gültekin, K., 2011. How Important is the Dark Matter Halo for Black Hole Growth? ApJ, **737**, 50.

Wadsley, J. W., Stadel, J. & Quinn, T., 2004. Gasoline: a flexible, parallel implementation of TreeSPH. New Astronomy, **9**, 137–158.

Wadsley, J. W., Veeravalli, G. & Couchman, H. M. P., 2008. On the treatment of entropy mixing in numerical cosmology. MNRAS, **387**, 427–438.

Wang, H. Y., Mo, H. J. & Jing, Y. P., 2007. Environmental dependence of cold dark matter halo formation. MNRAS, **375**, 633–639.

Warren, J. S., Hughes, J. P., Badenes, C., Ghavamian, P., McKee, C. F., Moffett, D., Plucinsky, P. P., Rakowski, C., Reynoso, E. & Slane, P., 2005. Cosmic-Ray Acceleration at the Forward Shock in Tycho's Supernova Remnant: Evidence from Chandra X-Ray Observations. ApJ, **634**, 376–389.

Wechsler, R. H., Bullock, J. S., Primack, J. R., Kravtsov, A. V. & Dekel, A., 2002. Concentrations of Dark Halos from Their Assembly Histories. ApJ, **568**, 52–70.

White, D. A., Jones, C. & Forman, W., 1997. An investigation of cooling flows and general cluster properties from an X-ray image deprojection analysis of 207 clusters of galaxies. MNRAS, **292**, 419.

White, S. D. M., 1984. Angular momentum growth in protogalaxies. ApJ, **286**, 38–41.

White, S. D. M. & Frenk, C. S., 1991. Galaxy formation through hierarchical clustering. ApJ, **379**, 52–79.

White, S. D. M., Navarro, J. F., Evrard, A. E. & Frenk, C. S., 1993. The baryon content of galaxy clusters: a challenge to cosmological orthodoxy. Nature, **366**, 429–433.

White, S. D. M. & Rees, M. J., 1978. Core condensation in heavy halos - A two-stage theory for galaxy formation and clustering. MNRAS, **183**, 341–358.

York, D. G., Adelman, J., Anderson, Jr., J. E., Anderson, S. F., Annis, J. et al., 2000. The Sloan Digital Sky Survey: Technical Summary. AJ, **120**, 1579–1587.

Zentner, A. R., 2007. The Excursion Set Theory of Halo Mass Functions, Halo Clustering, and Halo Growth. International Journal of Modern Physics D, **16**, 763–815.

Zwicky, F., 1933. Die Rotverschiebung von extragalaktischen Nebeln. Helvetica Physica Acta, **6**, 110–127.