

Parameter tests within cosmological simulations of galaxy formation

Scott T. Kay,¹★ F. R. Pearce,¹ A. Jenkins,¹ C. S. Frenk,¹ S. D. M. White,² P. A. Thomas³
and H. M. P. Couchman^{4,5}

¹*Physics Department, University of Durham, Science Laboratories, South Road, Durham DH1 3LE*

²*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild Straße, D-85740, Garching bei München, Germany*

³*Astronomy Centre, CPES, University of Sussex, Falmer, Brighton BN1 9QJ*

⁴*Department of Astronomy, Physics and Astronomy Building, University of Western Ontario, London, Ontario N6A 3K7, Canada*

⁵*Department of Physics and Astronomy, McMaster University, Hamilton, Ontario L8S 4M1, Canada*

Accepted 2000 February 25. Received 1999 December 8; in original form 1999 August 12

ABSTRACT

Numerical simulations of galaxy formation require a number of parameters. Some of these are intrinsic to the numerical integration scheme (e.g., the time-step), while others describe the physical model (e.g., the gas metallicity). In this paper we present results of a systematic exploration of the effects of varying a subset of these parameters on simulations of galaxy formation. We use N -body and ‘Smoothed Particle Hydrodynamics’ techniques to follow the evolution of cold dark matter and gas in a small volume. We compare a fiducial model with 24 different simulations, in which one parameter at a time is varied, focusing on properties such as the relative fraction of hot and cold gas, and the abundance and masses of galaxies. We find that for reasonable choices of numerical values, many parameters have relatively little effect on the galaxies, with the notable exception of the parameters that control the resolution of the simulation and the efficiency with which gas cools.

Key words: hydrodynamics – methods: numerical – galaxies: formation – cosmology: theory.

1 INTRODUCTION

Over the past two decades numerical simulations have become the most commonly used technique for investigating the formation of structure in the Universe. Originally, N -body simulations were used to calculate the clustering evolution of collisionless dark matter (e.g. Centrella & Melott 1983; Klypin & Shandarin 1983; Davis et al. 1985). Subsequently, countless N -body studies have addressed problems ranging from the large-scale distribution of dark matter to the inner structure of dark matter haloes in a programme whose results underpin much of modern cosmology.

The success of the N -body approach stimulated the extension of simulation techniques to processes involving gas. Eulerian, Lagrangian and deformable grid methods have been developed for this purpose. One of the simplest cosmological problems involving gas is the evolution of the intracluster medium in rich clusters (e.g. Evrard 1990; Thomas & Couchman 1992, hereafter TC92; Cen & Ostriker 1994; Navarro, Frenk & White 1995; Anninos & Norman 1996; Bryan & Norman 1998; Eke, Navarro & Frenk 1998). Since the cooling time of the gas is longer than the Hubble time in all but the central regions of clusters, the intracluster gas may be approximated as non-radiative. This is an important simplification that allows reliable and repeatable results to be obtained with a variety of techniques and numerical resolutions, as demonstrated by a recent systematic comparison

based on a simulation of a galaxy cluster in the cold dark matter (CDM) cosmogony (Frenk et al. 1999).

Another cosmological problem that can be addressed using gas dynamics is the evolution of cold gas clouds at high redshift. These are responsible for the Lyman- α forest lines seen in quasar spectra, and their low temperatures and moderate density contrasts ensure that cooling processes are relatively unimportant. Simulations of the Lyman- α forest in CDM models have been very successful, and have radically changed the prevailing view of the high-redshift gaseous Universe (e.g. Cen et al. 1994, 1998; Hernquist et al. 1996; Katz, Hernquist & Weinberg 1999). Recent work, however, has highlighted the need to take numerical effects carefully into account when making comparisons with observations (Theuns et al. 1998).

Simulating the formation of galaxies requires the incorporation of radiative cooling processes into the fluid equations, enabling the gas to achieve high enough density contrasts so that stars can plausibly form (e.g. Cen & Ostriker 1992; Katz, Hernquist & Weinberg 1992; Evrard, Summers & Davis 1994; Pearce et al. 1999). In general, the inclusion of such processes presents a formidable numerical challenge, because in hierarchical clustering theories, such as CDM, cooling is most efficient in small objects at high redshift where cooling times are very short. As a result, the outcome of a simulation depends on its resolution. Fortunately, the problem is less intransigent than it might appear. It is clear that only a small fraction of gas can have cooled into small pre-galactic objects at high redshift. This is normally attributed to the

★ E-mail: Scott.Kay@durham.ac.uk

effects of ‘feedback’, a term frequently used to denote a generic mechanism that prevents the gas from cooling catastrophically and turning into stars (White & Rees 1978; Cole 1991; White & Frenk 1991). Although feedback is poorly understood, it almost certainly involves stellar energy input (supernova explosions, galactic winds, etc.) to the intergalactic medium, and, perhaps, even energy input from active galactic nuclei.

At present, feedback effects need to be included in simulations using a phenomenological model. In some cases, a prescription to turn some of the gas into stars and to add thermal or kinetic energy to the rest is explicitly incorporated into the simulation (Katz 1992; Navarro & White 1993; Steinmetz & Müller 1995). In other cases (e.g. Pearce et al. 1999), the effect of resolution itself is used to cut off cooling; only gas in objects above the resolution limit of the simulation can cool efficiently.

Although resolution is the most important numerical concern in simulations of galaxy formation, several other numerical and physical parameters influence the outcome of a simulation. Examples of the former are the gravitational softening or the size of the time-step; examples of the latter are the metallicity of the gas or the baryon fraction. Often parameter values are set in the absence of rigorous criteria. In this paper we undertake a systematic study of a *limited* subset of the parameters required in simulations of galaxy formation. We consider the way in which such parameters affect, for example, the amount of gas that cools, the number and masses of the galaxies that form, etc. This is not intended as an exhaustive investigation, but rather as a guide to the sensitivity of simulation results to these particular parameters.

We restrict our study to one particular implementation of numerical hydrodynamics, the ‘Smoothed Particle Hydrodynamics’ (SPH) technique (Gingold & Monaghan 1977; Lucy 1977). This Lagrangian approach is well suited to studies of galaxy formation, because large dynamic variations in gas density and temperature are easily accommodated. In the implementation used in this paper, our code models processes such as adiabatic compression and expansion, shock heating and radiative cooling. For simplicity, we have chosen to ignore the effects of photoionization (studied previously, for example, by Navarro & Steinmetz 1997 and Weinberg, Hernquist & Katz 1997) and feedback. We intend to include these effects in a subsequent set of tests.

The rest of this paper is organized as follows. In Section 2 we outline the numerical methods used, particularly our implementation of gravity and hydrodynamics with radiative cooling. In Section 3 we define a ‘fiducial’ simulation and describe its evolution, focusing on the distribution of gas in the temperature-density plane and on the properties of ‘galaxies.’ In Section 4 we compare this fiducial simulation to a set of simulations in which one parameter at a time is varied. In particular, we consider the correspondence of galaxies in different simulations, as well as properties such as their mass distribution and the distribution of baryons within dark matter haloes. We summarize our results in Section 5 and conclude in Section 6.

2 NUMERICAL METHOD

The data analysed in this paper were generated using HYDRA¹ (Couchman, Thomas & Pearce 1995, hereafter CTP95), an

adaptive particle–particle/particle–mesh (AP³M) code incorporating SPH. The simulations were performed on single-processor workstations, although the same code has been implemented in parallel on a Cray T3D (Pearce & Couchman 1997). Below we briefly outline how HYDRA works.

2.1 Gravity

The AP³M method for calculating the gravitational forces involves three distinct parts: the particle–mesh (PM), particle–particle (PP) and refinement placing algorithms. The implementation used in HYDRA is detailed in Couchman (1991), although we outline the key concepts below.

The PM algorithm is an efficient method for calculating forces accurately down to scales of the order of the Nyquist wavelength of the mesh used. The computational speed scales $\sim O(L^3 \log L)$, where L^3 is the total number of mesh cells. The density field is discretized by smoothing the particles on to a regular cubic mesh using the Triangular-Shaped Cloud (TSC) kernel (Hockney & Eastwood 1981). The potential of this distribution is then obtained by performing a Fast Fourier Transform (FFT), multiplying by the appropriate Green’s function, and doing an inverse FFT. The potential at each mesh point is differenced to obtain the forces, which are then interpolated back to the particle positions using the same TSC kernel.

The PM method is augmented on small scales by the PP method, which calculates pairwise forces directly via a neighbour search. This is performed using a coarse mesh and searching only the 27 nearest cells for neighbours. The coarser this mesh is relative to the FFT mesh, the more accurate the force calculation becomes at the expense of extra PP work.

On very small scales, the gravitational force must be softened not only to avoid effects due to two-body interactions, but also to limit the need for very small time-steps. HYDRA uses a spline (S2) softening function which returns the Newtonian value of the force for $r \geq s$, where s is the softening length. For convenience, we express this value as the more common Plummer softening, ϵ (e.g. Binney & Tremaine 1987), using the equivalence $s = 2.34\epsilon$. (This produces asymptotically matching force laws at both large and small scales.) At late times we employ a fixed physical softening which is the value normally quoted as the softening for the simulation. In comoving coordinates this will grow backwards in time proportional to $(1+z)$. To maintain force accuracy, we do not allow the softening to grow indefinitely, introducing a maximum value, s_{\max} , which is typically set to equal 0.3 FFT cells. Above a certain redshift (typically around $z = 1$) the softening is constant in comoving coordinates.

Since gravity is an attractive force, the particle distribution evolves into a highly clustered state where large overdensities are inevitable. Hence the ratio of PP to PM work rapidly increases, leading to a substantial increase in the computational effort for each step. The refinement placing algorithm (Couchman 1991) eases this by locating highly clustered regions and recursively placing submeshes (or refinements) over these volumes. This technique drastically reduces the total amount of PP work per time-step, passing much of the effort over to the faster PM algorithm.

2.2 Hydrodynamics

Hydrodynamical forces are modelled using the SPH formalism, which, because of its Lagrangian nature, is well suited to problems

¹This code is in the public domain and can be obtained from the Hydra Consortium at the following URL: http://phobos.astro.uwo.ca/hydra_consort/

with a large dynamic range in density – a natural occurrence in cosmological simulations. Solving the hydrodynamic equations involves estimating various fields such as the density and pressure gradient, from the discrete particle distribution. SPH tackles this by smoothing the particles over a finite volume using a spherically symmetric smoothing kernel. To cope with large density fluctuations, the range of the kernel is set for each particle to be the radius of a sphere containing approximately N_{SPH} neighbouring particles (Wood 1981). The neighbour search is performed at the same time as the PP calculation and hence introduces no additional computational overhead. This procedure introduces a maximum SPH search length equivalent to the size of a PP chaining cell which leads to a minimum resolved density. For normal simulation parameters this limit is approximately the mean cosmological density.

Stable integration of the fluid equations requires the introduction of an artificial viscosity term (Monaghan & Gingold 1983). The implementation used in this study is based on the divergence of the flow, rather than the more usual relative pairwise velocity of particles.

Some excellent reviews of the SPH technique are available in the literature (e.g. Benz 1990; Monaghan 1992) the implementation in HYDRA is summarized in TC92; see also CTP95). The reader is also referred to recent work by Thacker et al. (1998), who have investigated different implementations of the SPH force calculation, specifically the force symmetrization and artificial viscosity assumptions. Part of that study (section 3.5) utilized the same initial conditions as we use here, and examined the effects of different SPH implementations. Our implementation is equivalent to version 1 of Thacker et al. (1998).

2.3 Radiative cooling

Radiative cooling requires adding a sink to the energy equation derived from the emissivity function,

$$\epsilon(n, T) = n^2 \Lambda(T) \quad (2.1)$$

where n is the number density of atoms/ions and electrons, and $\Lambda(T)$ is the (temperature-dependent) normalized cooling function. Our version of HYDRA uses a cooling function made up from a series of power-law fits to the optically thin radiative cooling code of Raymond, Cox & Smith (1976), as detailed by TC92. This fit contains contributions from bremsstrahlung and from line cooling due to hydrogen, helium and heavier elements with an assumed abundance Z times the solar value. The fit is relatively good above 10^4 K – below this temperature the cooling function drops precipitously, and we take it to be zero. Fig. 1 illustrates these curves for metallicities $Z/Z_{\odot} = 0.0, 0.5, 1.0$. The $Z = 0.5 Z_{\odot}$ cooling function of Sutherland & Dopita (1993) is overlaid.

2.4 Numerical integration

The numerical integration algorithm is described in CTP95, who evaluated several algorithms and settled for a simple PEC (Predict, Evaluate, Correct) scheme. This has the advantage of keeping storage to a minimum while also allowing arbitrary changes in the time-step, if the force changes abruptly in time (as it can do in the vicinity of shocks). HYDRA evaluates the time-step, Δt , such that $\Delta t = \kappa \min(0.4\Delta t_v, 0.25\Delta t_a, 0.0625\Delta t_H)$, where $\Delta t_v = \epsilon/v_{\text{max}}$, $\Delta t_a = (\epsilon/a_{\text{max}})^{1/2}$ and $\Delta t_H = 1/H$, with v_{max} being the current largest speed, a_{max} the current largest acceleration, and H the




Figure 1. The fits used in HYDRA for the normalized cooling function, $\Lambda(T)$. The three curves illustrate this for three values of the metallicity: $Z/Z_{\odot} = 0.0, 0.5$ and 1.0 . The function is plotted only for temperatures down to 10^4 K, since this is the limit where HYDRA switches off the cooling. (Physically, negligible thermal energy is lost below this value.) Overlaid on the $Z = 0.5 Z_{\odot}$ line is the interpolated curve from the tabulated values given by Sutherland & Dopita (1993).

Hubble parameter. The value of κ is a normalization constant typically set to unity for adiabatic simulations (TC92). A Courant-style condition is implicit in Δt_v , since the minimum SPH smoothing length for the gas is set by the softening. Note that a single time-step is used to advance all particles. In principle, a further gain in efficiency could be obtained by implementing individual time-steps for each particle.

Although the cooling time-scale is often shorter than the dynamical (or sound-crossing) time-scale, we do not shrink the time-step to follow cooling explicitly. Instead, we apply the cooling condition, assuming a constant density, at the end of the time-step. Thermal energy is removed such that

$$\int_{e_i}^{e_i - \Delta e_i} \frac{de_i}{\Lambda(T_i)} = - \frac{n_i^2}{\rho_i} \Delta t, \quad (2.2)$$

with subscript i denoting the i th particle, and e its specific thermal energy. This results in stable and physically reasonable behaviour, even for relatively long time-steps.

3 THE FIDUCIAL SIMULATION

3.1 Initial conditions

Throughout this paper we examine the effects of varying various numerical and physical parameters relative to a fiducial simulation. Our fiducial initial conditions were chosen to correspond to the standard cold dark matter model (SCDM). Thus we choose the cosmic density parameter, $\Omega = 1$, the cosmological constant, $\Lambda = 0$, and the Hubble constant, $h = 0.5$.² Although this model is now out of favour (for example, its power spectrum cannot simultaneously fit both the large-scale COBE normalization and the

² $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$.

amplitude of small-scale galaxy clustering), it is a well-studied model with moderately large power on all scales relevant to galaxy formation. In any case, the precise form of the initial conditions is not important for the tests carried out here.

The initial mass distribution was generated as described by Efstathiou et al. (1985), by perturbing a uniform distribution of particles using the Zel’dovich approximation. The amplitudes of the Gaussian random field were set using Bond & Efstathiou’s (1984) approximation to the CDM power spectrum. The normalization was set so that σ_8 , the rms of linear mass fluctuations in top-hat spheres of comoving radius, $8 h^{-1}$ Mpc, matches the value derived from the local abundance of rich clusters. We adopted $\sigma_8 = 0.6$ (White, Efstathiou & Frenk 1993; Eke, Cole & Frenk 1996; Vianna & Liddle 1996). The size of the simulation box (in comoving coordinates) was fixed at $10 h^{-1}$ Mpc, with a mesh resolution of $L = 64$ (implying a cell-length of $156.25 h^{-1}$ hpc). Although the box is not big enough to follow the evolution of large-scale structure accurately (waves on the scale of the box are not evolving linearly at the final time), this is not important since we are not attempting to compare our results directly to observations – the volume is far too small for such purposes.

Table 1 summarizes the fiducial values of all the parameters that we study in this paper. We represent the mass using 32^3 gas and 32^3 dark matter particles. The initial softening was set to 0.3 FFT cells (equivalent to a comoving Plummer softening of $\sim 20 h^{-1}$ hpc), but after $z \sim 1$ it switched to a fixed physical softening of $10 h^{-1}$ hpc. The number of neighbours over which the SPH algorithm smooths was set to $N_{\text{SPH}} = 32$. The simulation was run from an initial epoch, $z_i = 24$, to the present ($z = 0$), with the data output at regular intervals.

Regarding the physical parameters, we chose to set the baryon fraction equal to $\Omega_b = 0.06$, consistent with the results from nucleosynthesis calculations (Copi, Schramm & Turner 1995). The global gas metallicity was set to $Z = 0.5 Z_\odot$, similar to the value obtained from spectral observations of intracluster gas (e.g. Mushotzky et al. 1996). The gas was given a cold start, with the initial temperature set to 100 K.

3.2 The baryon phase diagram at $z = 0$

The two objects that we study in this paper are the dark matter haloes and the cold, dense clumps of gas that we define as galaxies. All our comparisons are carried out at $z = 0$. The

Table 1. The values of the parameters chosen for the fiducial simulation that were subsequently varied. N is the total number of particles, N_{SPH} is the number of gas particles over which the SPH smoothing is carried out, L is the size of the FFT mesh, s_{max} is the maximum comoving size of the softening in units of FFT cells, ϵ_0 is the effective Plummer softening, κ is the time-step normalization, Ω_b is the baryon fraction, Z is the gas metallicity, z_i is the initial redshift and T_i is the initial gas temperature. The choice of these parameters is discussed in more detail in the text.

Parameter	Value
N	2×32^3
N_{SPH}	32
L	64
s_{max}	0.3
ϵ_0	$10 h^{-1}$ hpc
κ	1
Ω_b	0.06
Z	$0.5 Z_\odot$
z_i	24
T_i	100 K

selection criteria used to catalogue the haloes and the galaxies are described in detail in Section 3.4, but as motivation for the galaxy selection procedure, we initially develop an understanding of the correspondence between the spatial distribution and the thermal evolution of the baryons. Fig. 2 illustrates the temperature–density distribution of the gas at $z = 0$, spanning 7 orders of magnitude in density and 9 orders of magnitude in temperature. We have split the distribution into three *phases*, which we label as the *uncollapsed* phase, the *shocked* phase and the *galaxy* phase. Accompanying this figure are three dot-plots, showing the corresponding spatial distribution of each phase.

3.2.1 The uncollapsed phase

We define the uncollapsed phase as the regime in which gas particles have $T \leq 1$ K and $\rho < 10^3 \langle \rho \rangle$. The boundaries for this phase are chosen such that they enclose the particles that lie on the power-law relation, which is an adiabat. These particles are still largely in free expansion and are, on average, cooling adiabatically, with $T \propto (1+z)^2$. Spatially, the particles are homogeneously distributed, as a result of occupying a small range in density. The hard lower limit on this quantity (visible for all temperatures) is artificial, imposed by the algorithm due to the maximum distance it can search for N_{SPH} neighbours. For the fiducial simulation, $\rho_{\text{min}} \sim 1.9 \langle \rho \rangle$, clearly an undesirably large value. As a result, the low-density gas is forced to unphysically high values (the gas should have $\rho < \langle \rho \rangle$). The latest version of HYDRA now properly accounts for low-density gas, and we have checked that this makes no significant difference to the properties of the high-density material.

3.2.2 The shocked phase

The shocked gas phase contains particles with $\rho < 10^3 \langle \rho \rangle$ and $T > 1$ K, as well as all particles with $T > 10^5$ K. Gas entered this phase due to the collapse and violent relaxation of gravitationally unstable regions, in which the bulk kinetic energy of the gas was transferred, via shocking, into heat. As a result, the spatial distribution of this phase exhibits large fluctuations in density: halo atmospheres lie in this regime, and the shocked gas is largely absent from the voids.

3.2.3 The galaxy phase

The particles in the galaxy phase lie in the regime $\rho \geq 10^3 \langle \rho \rangle$, and $T \leq 10^5$ K. The strong horizontal line at 10^4 K is due to the fact that the cooling rate drops to zero at this temperature. Particles above a temperature of 10^5 K must lose a significant fraction of their thermal energy before reaching the main locus, a transition made possible by the large contributions from hydrogen and helium line cooling at these temperatures. The plot of their positions highlights the fact that cold, dense gas forms very tight clumps, with spatial extents of the order of the gravitational softening. To select ‘galaxies’, we look at the particles that lie in this region of the phase diagram.

3.3 Evolution of the baryon phase diagram

To illustrate the global evolution that led to the baryon phase diagram in Fig. 2, we plot the mass fraction in each phase in Fig. 3, as a function of the expansion factor (normalized to the present

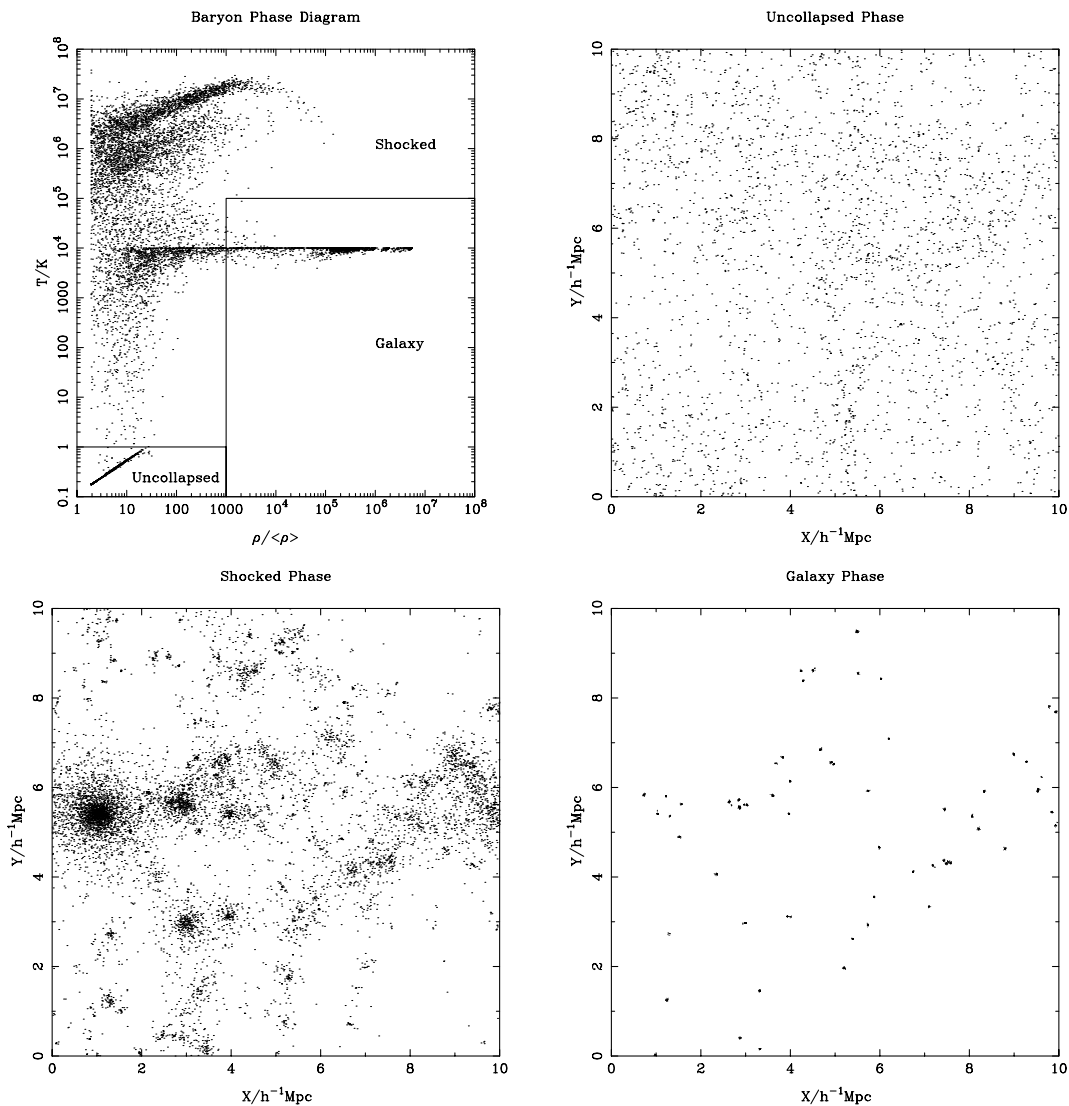


Figure 2. (top-left) The temperature–density phase diagram of all baryons in the fiducial simulation at the final output time ($z = 0$). The density is displayed in units of the mean gas density, and the temperature in Kelvin. Positions of the gas particles obeying the various selection criteria defined in the text are shown in the other three plots, specifically the uncollapsed, shocked and galaxy phases. It is evident that each phase shows a considerable difference in structure: the uncollapsed phase particles are homogeneously distributed, the shocked phase particles show large fluctuations in density, and the galaxy phase is exclusively made up of tight clumps of particles.

value). All phase boundaries are fixed, with the exception of the uncollapsed phase – the temperature threshold is modified to $T < (1 + z)^2 \text{ K}$, accounting for the adiabatic expansion of the uncollapsed gas. Initially, all the gas is in this phase, with a temperature of 100 K. As the simulation progresses, some of the gas transfers to the shocked phase, as the first virialized haloes become resolved ($z \sim 10$). Gas does not start to enter the galaxy phase (forming the first objects) until later ($z \sim 3$). Both the shocked and galaxy phases continue to accrete more and more mass as the uncollapsed phase becomes progressively depleted. By the present day, ~ 18 per cent of the gas is still in the uncollapsed phase, whilst the shocked gas phase contains about 64 per cent of the gas mass and the remaining 18 per cent has cooled into the galaxy phase.

Another important question we have addressed is to determine the precise trajectories in the ρ – T plane of the particles that finally end up in the galaxy phase. Fig. 4 illustrates the two main paths

that the galaxy particles take, for the whole duration of the fiducial simulation. The marked points indicate the redshifts $z = (10, 5, 3, 2, 1, 0.5 \text{ and } 0.1)$ for each trajectory.

The first example (solid line) shows the expected evolution of a galaxy particle – it becomes initially shock heated to over 10^6 K before cooling back down on to the 10^4 K boundary and, on average, progressively increasing in density until the final time. The second example (dotted line) takes a different path – again, the particle is shock-heated, but only as far as 10^4 K . Both particles undergo a series of further heating events and, notably, the second one reaches its largest temperature at some later stage when a major galaxy merger takes place.

Quantitatively, we find that only ~ 11 per cent of the particles which end up in the galaxy phase were initially heated above 10^5 K for a minimum of one time-step (~ 32 per cent initially reached temperatures above 12 500 K). Most of the galaxy mass reaches high density in the smallest haloes close to the resolution

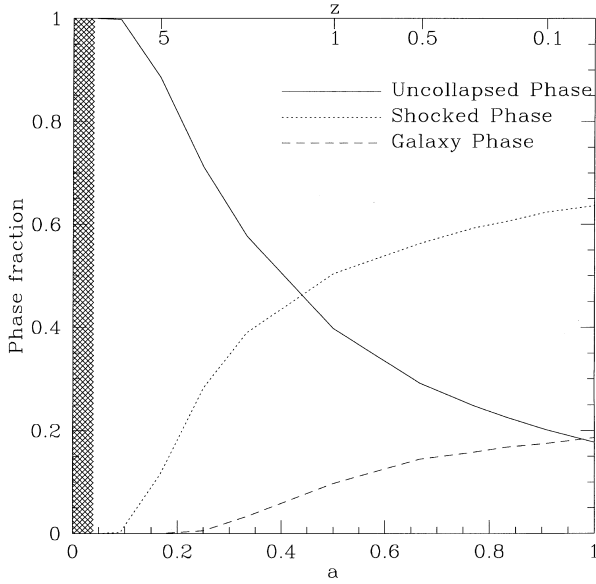


Figure 3. The evolution of the uncollapsed, shocked and galaxy phases, plotted as the fraction of gas particles in each phase (with respect to the total mass of gas particles), as a function of the expansion factor, $a = (1+z)^{-1}$. The shaded region marks the time before the start of the simulation ($a = 0.04$). Redshift is indicated along the top of the figure.

limit of the simulation. Consequently, their very short cooling times cause them to be heated above the 10^4 K locus and subsequently cool back down within one time-step.

3.4 Selection criteria

To select the ‘galaxies’, we initially extracted the subset of the gas particles that lie in the defined galaxy phase (Fig. 2). This was then run through a friends-of-friends group finder (Davis et al. 1985), which links particles together to form groups until no unlinked particles remain that are closer to a group than the linking length, $l = b\bar{n}^{-1/3}$, where \bar{n} is the global mean number density. We used the value $b = 0.1$ ($l \sim 30 h^{-1}$ hpc), which gives a linking length of the order of the final S2 softening length. Since the clumps are very tight and are well separated, the output of the group finder is insensitive to variations in the choice of b : changes of the order of a few per cent in mass were noted when varying b by a factor of 2 in each direction.

Once this catalogue was generated, it was truncated by throwing away objects containing fewer than N_{SPH} particles. Objects below this mass are poorly defined, because the SPH algorithm does not sample their properties well. For the fiducial simulation, 16 groups out of 69 were discarded in this manner, making up ~ 4 per cent of the mass in the galaxy phase. This constitutes most of the residual material in the galaxy phase that is not part of the final galaxy catalogue.




Figure 4. Two examples of particle trajectories in the ρ - T plane, from the initial epoch ($z = 24$) to the present day. The legends depict their maximum temperatures. The solid path illustrates the particle that has the highest temperature on initial heating (i.e., the first time it crosses the 10^4 K boundary), whereas the dotted path illustrates the evolution of a particle which is only heated to 10^4 K. Marked points indicate redshifts $z = (10, 5, 3, 2, 1, 0.5, 0.1)$.

We have also checked the effect of varying our defined limits of the galaxy phase. Since we make use of the total mass of gas in this phase, it is important to ensure that we have selected a region within which most of the particles are part of objects in the galaxy catalogue. The most important boundary is the lower limit to the gas density ($\rho/\langle\rho\rangle = 10^3$). Since the 10^4 K feature is visible down to the lowest densities, there are particles that have gone through the cooling process that lie outside our galaxy phase. Lowering the density limit by an order of magnitude caused only a 3 per cent increase in the total number of particles grouped into galaxies. On the other hand, the fraction of the galaxy phase present in the final galaxy catalogue dropped from 96 to 80 per cent. Reducing the density by the same amount again caused no further change in the former quantity, but the completeness again decreased to 56 per cent, indicating that at these low densities there is contamination from the diffuse halo gas. The temperature boundary for the galaxy phase includes all relevant particles for $\rho > 10^3\langle\rho\rangle$, since the hotter particles are still cooling down. As a check, however, we lowered the upper limit on the temperature from 10^5 K to 12000 K. This produced a change of less than 1 per cent to the completeness and lowered the number of galaxies in the catalogue from 53 to 52.

The positions of the fiducial galaxies are illustrated in Fig. 5. Here we plot the projected positions of the particles linked together by the group finder. The circles illustrate the size of the objects, centred on the median position of the linked particles in each group, with radii proportional to the cube root of the number of particles linked.

The dark matter haloes were identified by using a spherical overdensity algorithm (Lacey & Cole 1994) on the positions of the dark matter particles in the fiducial simulation endstate. The algorithm makes an estimate of the local density, which is then used to find the centre of overdense regions. Successive spheres are placed around this centre with growing radii, until the mean internal density equals some value – we assume the usual value of 178 times the mean density, as predicted by the spherical collapse model for the virial radii of haloes in an $\Omega = 1$ universe. Iterations are performed, with the centre redefined as the centre of mass of the object, until the difference becomes small. Substructure within haloes was ignored for the purposes of this paper. The halo finder produces a catalogue of all haloes with more than $N_h = 100$ particles in the fiducial case, corresponding to a mass of $\sim 8 \times 10^{11} h^{-1} M_\odot$. Since the dark matter dominates the dynamical evolution, a lower limit in halo mass (rather than particle number) was applied when comparing different simulations. Hence, for simulations with poorer mass resolution, the lower limit on the number of particles was reduced in order to keep the halo mass threshold constant. This limit never went below 10 particles, roughly the lowest number that can be trusted (Efstathiou et al. 1988).

The output of the halo finder for the fiducial simulation is illustrated in Fig. 6. The circles are centred on the halo centre of mass, and have radii equal to the virial radius of each halo. The dots represent the positions of all the dark matter particles within the virial radius of each halo. The largest halo has a total mass of $\sim 5 \times 10^{13} h^{-1} M_\odot$, which corresponds to a virial temperature of $\sim 10^7$ K.

The ρ – T distribution of the baryons located within the halo virial radii for the fiducial simulation are shown in Fig. 7. The range of temperatures for the majority of the hot phase gas agrees with the corresponding range of virial temperatures of the haloes for the mass range of the simulation (from $\sim 5 \times 10^5$ to 10^7 K).




Figure 5. Projected positions of the fiducial set of 53 galaxies. The dots show the particles that are linked together by the group finder, and the circles illustrate the size of the galaxies, centred on the median position of each galaxy with radii proportional to the cube root of the galaxy mass.

Note that only a negligible amount of cold gas falls within the halo boundaries.

4 SIMULATION COMPARISONS

4.1 Global features

The complete set of simulations that we analysed are listed in Table 2, split into three sections. The first section includes




Figure 6. Projected positions of the dark matter particles in the fiducial simulation selected by the halo finder. Only haloes with more than 100 particles are shown. The circles are centred on each halo's centre of mass, and their radii are set equal to the distance from the centre to the most outlying particle of the halo. There are 31 haloes in total in the fiducial catalogue.




Figure 7. Phase diagram of all the gas particles within the selected halo virial radii in the fiducial simulation. Overlaid are the boundaries which we have defined in order to split the gas into separate phases.

simulations in which features other than simple parameter values were varied. The second section contains simulations in which numerical parameters (i.e., those that affect the algorithm) were varied. Finally, the third section contains simulations in which the physical parameters were altered. For most of the parameters,

there is a set of three simulations: the fiducial case and two additional simulations with the appropriate parameter varied about the fiducial value. The columns indicate the comparison number (13 in total); the alteration made to the fiducial parameter set; the fraction of particles identified in the uncollapsed, shocked and galaxy phases respectively; the number of galaxies in the final catalogue; the number of haloes in the final catalogue; and the *completeness*, which we define as the fraction of particles in the galaxy phase that appear in the final object set. For discussion, we label the sections *general*, *numerical* and *physical* respectively.

4.1.1 General comparisons

The first comparison was designed to probe effects that might be present due to the fact that runs were performed on different computer architectures. The fiducial simulation was run on a Sun Sparcstation, whereas most of the comparison runs were performed on a Dec Alpha. Hence, we have analysed an identical simulation to the fiducial case that was executed on the latter platform. Differences arise from the system-specific changes made to the code by the compiler. Clearly, this change can introduce variations in the phase fractions at the level of ~ 1 –2 per cent. Furthermore, the number of galaxies in this comparison differs by 2, and the number of haloes differs by 1.

Comparison 2 was performed in order to investigate the numerical effects introduced by perturbing the particles slightly from their initial positions. We set up two simulations to test this. The first also tested the periodicity of the algorithm, by translating the particles by a fixed amount (in this case, 1 FFT cell) before

Table 2. Details of the set of simulations analysed in this paper, following the various definitions and selection procedures discussed in Section 3. The first column is the comparison number given to the simulation, with the second column indicating the alteration made to the fiducial initial conditions. The next three columns give the fraction of baryons in each of the three phases defined in the text. Columns six and seven give the number of galaxies and haloes respectively that remain in the catalogue after the various selection procedures have been applied. The final column is the *completeness*, i.e., the fraction of baryons in the galaxy phase that are present in the final galaxy catalogues. Bold type has been used to denote values that deviate significantly from the fiducial case (2 per cent for the phase fractions, 5 percent for the galaxies, haloes and completeness).

No.	Alteration	Uncollapsed Phase	Shocked Phase	Galaxy Phase	Galaxies	Haloes	Completeness
0	Fiducial Simulation	0.18	0.64	0.18	53	31	0.96
1	Architecture	0.16	0.65	0.19	51	32	0.96
2	Position Offset (1 cell)	0.17	0.65	0.18	50	31	0.96
2	Position Offset (random)	0.18	0.64	0.18	55	30	0.97
3	New Cooling Function	0.18	0.62	0.20	55	31	0.98
4	Cold gas decoupled	0.18	0.67	0.15	51	31	0.95
5	$s_{\max} = 0.6$	0.17	0.64	0.19	53	30	0.96
6	$\kappa = 0.5$	0.18	0.63	0.19	52	31	0.97
6	$\kappa = 2.0$	0.06	0.78	0.18	53	31	0.95
7	$\epsilon_0 = 5 h^{-1} \text{ hpc}$	0.16	0.69	0.15	45	31	0.95
7	$\epsilon_0 = 20 h^{-1} \text{ hpc}$	0.18	0.63	0.19	53	30	0.97
8	$N_{\text{SPH}} = 16$	0.18	0.56	0.26	97	31	0.98
8	$N_{\text{SPH}} = 64$	0.17	0.70	0.13	28	31	0.89
9	$N = 2 \times 16^3$	0.20	0.80	0	0	29	0
9	$N = 2 \times 24^3$	0.17	0.74	0.09	14	30	0.71
9	$N = 2 \times 48^3$	0.17	0.57	0.26	141	34	0.98
10	$1 + z_i = 10$	0.17	0.65	0.18	55	32	0.97
10	$1 + z_i = 50$	0.19	0.63	0.18	51	29	0.96
11	$\Omega_b = 0.03$	0.18	0.69	0.13	49	31	0.94
11	$\Omega_b = 0.12$	0.18	0.58	0.24	53	31	0.97
12	$Z = 0.0 Z_{\odot}$	0.18	0.78	0.04	9	30	0.44
12	$Z = 1.0 Z_{\odot}$	0.18	0.60	0.22	54	31	0.97
13	$T_i = 0 \text{ K}$	0.18	0.64	0.18	51	31	0.95
13	$T_i = 10^7 \text{ K}$	0.22	0.61	0.17	50	30	0.96

starting the simulation, then moving them back at the final time. The second was performed in order to examine how small inaccuracies in the initial positions affect the results, by adding a random perturbation (tiny compared to the initial displacements) to each particle position in the initial conditions. Variations in the phase fractions and number of objects are similar to the first comparison; hence the two comparisons set a baseline for the size of variations that can be introduced solely by noise. Simulations that only differ by this amount should be viewed as being indistinguishable from the fiducial case.

For comparison 3, we have added a run that uses a cooling function in tabulated form, as given by Sutherland & Dopita (1993), rather than a series of power-law fits. The two cooling functions used are shown in Fig. 1 for the fiducial metallicity, $Z/Z_{\odot} = 0.5$. The power law fits the tabulated function relatively well for the range of relevant temperatures (10^4 – 10^8 K), although the latter shows a significant enhancement in $\Lambda(T)$ around 10^5 K, (as much as a factor of 3) due to helium line cooling. However, as Table 2 shows, there appears to be only a slight enhancement in the amount of gas in the galaxy phase.

The final comparison that was made in this section was between the fiducial simulation and the version of HYDRA used by Pearce et al. (1999). The latter contains a change in the SPH algorithm to reduce cooling in large haloes (see also Thacker et al. 1998). It is based on a *multiphase* approach, in which the cold, dense particles are decoupled from the hot medium: gas above 10^5 K does not see gas below 12000 K for the purposes of calculating the gas density and forces; for all other gas particles, they are calculated in the standard way. This procedure ensures that infalling cold diffuse particles still feel the halo’s outer accretion shock, and allows hot haloes to cool at a rate determined by only the hot gas and galaxies to dissipate energy, merge and feel drag as they move around the halo environment. The effects of this alteration on the phase distribution are evident in the table: the fraction of gas particles in the galaxy phase is reduced by 3 per cent (with ~ 2 per cent due to the largest object), while the fraction in the shocked phase increases to compensate.

4.1.2 Varying numerical parameters

The first numerical parameter we varied (comparison 5) was the initial (maximum) softening, s_{\max} . Increasing s_{\max} moves the changeover from a comoving to a physical softening to higher redshift. We doubled the value of the maximum softening to $s_{\max} = 0.6$, so that for the same final softening length, the changeover now occurred at $z \sim 3$ rather than at $z \sim 1$. More importantly, the softening now has a larger comoving value (equivalent to a Plummer softening of $\sim 40 h^{-1}$ hpc) before the changeover. The galaxies that form before this epoch are more loosely bound (increasing their effective volume), and therefore are susceptible to greater variations in their mass, due to ram-pressure stripping, tidal disruptions or accretion as they move through the hot halo gas. No significant changes in the final phase fractions are present, however.

The time-step normalization value κ (comparison 6) was varied by a factor of 2 about the fiducial value of 1. This effectively doubles or halves the length of each time-step, causing the simulation to take roughly twice or half as many time-steps to run to $z = 0$. Doubling the fiducial value introduces inaccuracies in both the positions and the velocities of the particles. The phase fractions show that there is significantly more shocked gas in the

$\kappa = 2$ simulation than in the fiducial run, at the expense of the uncollapsed gas. The increased efficiency at which the gas is heated can be explained by the fact that the larger time-steps in the $\kappa = 2$ run increase the error in particle positions and velocities; this leads to larger values in the velocity divergence, increasing the efficiency of gas heating from viscous interactions.

As was hoped, the value $\kappa = 0.5$ had no significant effect on the endstate properties of the galaxies and haloes, vindicating CTP95’s choice of $\kappa = 1$.

Comparison 7 consisted of halving and doubling the value of the final softening length about the fiducial value, $\epsilon_0 = 10 h^{-1}$ hpc. Changing the value of ϵ_0 has an effect only when the physical softening length has dropped below the maximum defined by s_{\max} . For $\epsilon_0 = 5 h^{-1}$ hpc this occurs at $z = 3$, whilst for $\epsilon_0 = 20 h^{-1}$ hpc the softening is always comoving. Doubling the fiducial value of ϵ_0 has no significant effect on the final number of galaxies and the fraction of gas in each phase. However, halving the final softening length systematically decreases the number of galaxies by ~ 15 per cent, with the residual material still in the shocked phase. This might be due to the effects of two-body relaxation artificially heating gas particles (Steinmetz & White 1997). For a dark matter halo with 50 particles that forms at $z = 0$, the two-body relaxation times are approximately 10.4, 7.5 and 5.9 Gyr for $\epsilon_0 = 20, 10$ and $5 h^{-1}$ hpc respectively. For $\epsilon_0 = 5 h^{-1}$ hpc, the relaxation time is less than half the age of the Universe.

The value of N_{SPH} was varied in comparison 8. Since this parameter sets the resolution at which the density field is evaluated, large overdensities are progressively smoothed out as the value of N_{SPH} increases. It then becomes increasingly difficult to form galaxies, since the emissivity is proportional to the square of the gas density. On the other hand, using a smaller value of N_{SPH} increases the noise arising from the discrete nature of the mass distribution. Thus a compromise in the value is sought. It is common to assume $N_{\text{SPH}} = 32$, which is our fiducial value, and this comparison looks at the effects introduced by changing this number by a factor of 2. The number of galaxies in each simulation clearly shows the expected trend, with more galaxy phase gas and more galaxies if $N_{\text{SPH}} = 16$ than if $N_{\text{SPH}} = 64$. This material comes from the shocked phase, as the uncollapsed gas fraction is more or less constant in all cases. The significant drop in completeness for the $N_{\text{SPH}} = 64$ simulation is due to 11 groups being discarded because they are below the imposed resolution threshold (i.e., for these galaxies, $N_g < N_{\text{SPH}}$).

Comparison 9 consisted of changing the mass resolution by varying the total number of particles. The initial power spectrum was truncated appropriately to account for the varying Nyquist frequency of the particle mesh. The four runs have $N = 2 \times (16, 24, 32, 48)^3$ particles respectively, where the 2 indicates the number of species. The masses of the two particle species in each of these runs are given in Table 3. The three runs with the largest N were selected to quantify the result of increasing the resolution of the simulation. We also added an $N = 2 \times 16^3$ run to probe the

Table 3. The masses of each particle species for the runs performed with varying values of N .

N	$M_{\text{dark}}(10^{10} h^{-1} M_{\odot})$	$M_{\text{gas}}(10^{10} h^{-1} M_{\odot})$
2×16^3	6.38	0.41
2×24^3	1.89	0.12
2×32^3	0.80	0.05
2×48^3	0.24	0.015

regime where artificial two-body heating of the gas particles by the dark matter particles should significantly affect the cooling rate of the gas. Fig. 8 shows the critical dark matter particle mass (defined as the mass where the cooling time of the gas equals the two-body heating time) as a function of gas temperature, for the metallicities and baryon fractions studied in this paper (Steinmetz & White 1997). We have assumed the gas fraction to be equal to the appropriate value of Ω_b , although the values can be scaled, if desired. The Coulomb logarithm is assumed to be 5, appropriate for the largest haloes in the simulations studied in this paper (and therefore demonstrating the worst-case scenario for this effect). Overlaid as dotted horizontal lines are the mass of the dark matter particles for the simulations in this comparison. For $\Omega_b = 0.06$ and $Z = 0.5 Z_\odot$, both the fiducial simulation and the $N = 2 \times 48^3$ simulation are acceptable for the range of temperatures relevant to this work, and the $N = 2 \times 24^3$ simulation is borderline. Since the $N = 2 \times 16^3$ run lies well above the critical line, the amount of gas cooled in this simulation should be significantly affected by two-body heating. The effect is indeed severe: only 24 particles make it into the galaxy phase, and no galaxies form at all. (Note that since the cooling time at the resolution limit of this simulation is $\sim 10^8$ yr – a small fraction of the age of the Universe – in the absence of artificial heating effects, we should have been able to form a population of cooled objects.)

For the remaining three simulations, the amount of uncollapsed gas remains roughly constant, and the biggest change is in the increasing amount of gas that cools from the shocked phase to the

galaxy phase with larger N . This is reflected both in the galaxy fractions and in the number of galaxies present in the final catalogues. This effect is expected when varying the mass resolution of the simulations. When N is increased, the sampling of the density field improves, and hence smaller mass haloes are able to be resolved. The smallest haloes in the simulation will, on average, form first, and therefore have higher density contrasts than the haloes that form later. As discussed in Section 3.3, most of the galaxy mass is accumulated in these objects.

Finally, we investigate the effect of changing the initial redshift of the simulation by considering runs with $1 + z_i = 10, 25$ and 50 respectively (comparison 10). The main potential pitfall when choosing z_i is that choosing too small a value will suppress the amount of small-scale power, causing objects on small scales to form too late. Hence the initial redshift should be high enough that scales of the order of the initial softening should not have already become non-linear (i.e., $\delta \ll 1$). For the $1 + z_i = 50$ run, the fiducial boundaries of the uncollapsed phase are sufficient to segregate the uncollapsed gas from the rest, while for the $1 + z_i = 10$ run, the temperature boundary is rescaled by a factor of 6.25, i.e., $(25/10)^2$. Evidently, there are no significant changes in any of the gas phases or in the number of galaxies, when comparing all three runs.

4.1.3 Physical comparisons

The mass fraction of baryons was varied in comparison 11,




Figure 8. The critical dark matter particle mass for spurious two-body heating to exceed cooling, as a function of temperature (Steinmetz & White 1997). The three solid curves in each plot correspond to different gas metallicities (from top to bottom, $Z/Z_\odot = 1.0, 0.5$ and 0.0 , with the fiducial metallicity, 0.5 , in bold). The different panels assume a local gas fraction, f , corresponding to the values of the global baryon fraction analysed in this paper, $\Omega_b = 0.03, 0.06$ and 0.12 . The dotted horizontal lines give the masses of the dark matter particles for $N = 2 \times X^3$, where X is the label on each line. A value of 5 was assumed for the Coulomb logarithm.

forming the set of values $\Omega_b = (0.03, 0.06, 0.12)$. Varying the baryon fraction alters the number density of ions and free electrons per simulation gas particle (and therefore the masses of both the gas and dark matter particle), and consequently changes the cooling rate, since the emissivity scales as n^2 . Again (see Fig. 8), it is important that the baryon fraction be consistent with the critical mass of the dark matter particles, given both the metallicity and the value of N . Clearly, for our choice of parameters, the value $\Omega_b = 0.03$ ($Z = 0.5, N = 2 \times 32^3$) is borderline, and this will contribute to the reduction in the number of galaxies formed. Increasing the baryon fraction has little effect upon the number of galaxies, but they are systematically heavier.

Comparison 12 was performed in order to examine the effect of changing the metallicity of the gas. Each particle is assigned a constant amount of *metals* (i.e., of nuclei that have an atomic number greater than 2). The metallicity affects the shape of the cooling function, as illustrated in Fig. 1. Also affected is the critical mass of the dark matter particles (which depends on the cooling function), as shown in Fig. 8. The simulation with $Z = 0$ is clearly on the wrong side of the line, and the galaxy phase fractions are consequently affected: slightly more material cools with the higher metallicity, but the run with no metals shows a catastrophic drop in the amount of material in the galaxy phase, the number of galaxies and the completeness.

Finally, in comparison 13, we created a set of runs with different initial gas temperatures in order to measure how the dynamical and thermal evolution of the gas depends on its initial thermal state. We included two runs, one with $T_i = 0$ K and the other with $T_i = 10^7$ K. Although these parameter choices may seem hard to justify in physical terms, we decided to stick with our approach of varying a single parameter at a time. The fiducial boundaries of the uncollapsed phase are sufficient for the run with $T_i = 0$ K, however the boundaries are redrawn for the run with $T_i = 10^7$ K, such that $s < 2.5$ and $\rho/\langle\rho\rangle < 50$, where $s = \log_{10}(T/(\rho/\langle\rho\rangle)^{2/3})$ measures the specific entropy of the gas. (Using these constraints makes it easier to pick out the uncollapsed gas, since it is isentropic.) For the run with $T_i = 0$ K, there are no significant differences in the measured quantities; however, the run with $T_i = 10^7$ K shows a small increase in the uncollapsed gas phase at the expense of the shocked gas. For this run, the gas cools as expected from adiabatic expansion until $z \sim 5$. However, by $z = 3$ the uncollapsed phase is at temperatures below 10^4 K, over an order of magnitude below the adiabatic temperature. At these redshifts, the density and the cooling rate of the gas are high enough to allow it to radiate a significant fraction of its energy away. Consequently, only a small fraction of the gas is too hot to undergo gravitational collapse.

4.2 The mass distribution of galaxies

In Fig. 9 we plot the cumulative mass functions of galaxies (i.e., the number density of galaxies greater than a specified baryon mass, M_g), with the fiducial simulation always plotted as the solid line. The dynamic range for these simulations is about 2 orders of magnitude in abundance, and about 1.5 orders of magnitude in mass. The functions all have a characteristic shape: a shallow slope at low masses, turning over to a steeper slope at higher masses, and a large tail at the high-mass end. The last feature is simply due to the existence of a single massive galaxy that has a baryonic mass of $\sim 5 \times 10^{11} h^{-1} M_\odot$. Galaxies this massive are extremely rare in the Universe; we shall discuss this object below.

The biggest differences in the mass functions occur in the physical comparisons, in which the gas cooling rate varies. The simulation with $Z = 0$ fails to produce as many galaxies as the fiducial simulation – the abundance is down by a factor of 5, and the galaxies that do form are much less massive. The large differences seen in the runs with varying baryon fraction are mainly due to the mass of the individual gas particles being different by a factor of 2. However, when the galaxy masses are scaled to take out this effect, there is still a systematic difference between each function. This is due to the effect the baryon fraction has on the cooling rate of the gas.

Varying the number of SPH neighbours shows a progression in the smallest resolved galaxy mass due to the selection criterion that the number of constituent particles $N_g > N_{\text{SPH}}$. There are fewer galaxies for larger N_{SPH} , and they are systematically less massive. This is a reflection of the fact that the algorithm has to smooth over a larger range, which consequently leads to a poorer resolution of the density field. Since emissivity scales as n^2 , the lower cooling rate inhibits the galaxy formation process – an entirely numerical effect. Varying the value of N shows similar results for the same reason, i.e., a larger number of particles enables the SPH algorithm to resolve higher gas densities.

The remaining differences are relatively small. Changing the cooling function to the tabulated form of Sutherland & Dopita (1993) produces more cooling around 10^5 K, boosting the cooling rate and systematically increasing the mass of the galaxies by ~ 10 per cent. However, there is no significant difference in the shape of the mass function.

Comparison 4 (in which the cold gas is decoupled from the hot gas) also shows a small difference in the mass functions. This difference is dominated by the largest object in the simulation volume, which is reduced in mass by a factor of 3 when the cold, dense gas is decoupled from the hot gas. Notably, this alteration to the algorithm has a negligible effect on the smaller objects.

The comparison set in which the final softening length is varied shows an offset for $\epsilon_0 = 5 h^{-1}$ hpc. This simulation produces galaxies that are systematically lighter than the simulations with larger values of ϵ_0 , as expected if they are being affected by two-body heating from the dark matter particles.

4.3 Comparison of matched galaxies and haloes

To compare the individual properties of the galaxies and haloes, we performed object-to-object matching between each simulation and the fiducial run. Since the number of objects in each catalogue is different in nearly every case, there is not a straightforward one-to-one correspondence between the sets. To circumvent this, we always took the smallest object set first and compared the larger one to it. When more than one object in the larger set was matched with the same object in the smaller set, the closest pair was selected. Generally, this was not a problem, since the discarded objects are usually at considerably larger distances than their matched object. However, this method will inevitably produce outliers in both separation and mass differences due to the fact that, on occasion, an object in one simulation is fragmented into smaller objects in another. This prompted us to use the median and semi-interquartile ranges as statistical measures of the offset and scatter of any relative quantity we measured. Positions of the galaxies were defined as the median position of the linked particles and, for the haloes, their centre of mass. Below, we detail the individual comparisons we made.




Figure 9. Cumulative galaxy mass functions for the full set of simulation comparisons. The plot gives the comoving number density of galaxies with baryon mass greater or equal to M_g . The fiducial simulation is always plotted with a solid line-type.

4.3.1 Galaxy–galaxy displacements

The first property we examined was the scatter in the positions of the matched galaxies. The results are illustrated in Figs 10 and 11, which show the displacement between each matched pair (in h^{-1} hpc) plotted against the average object mass. The solid line illustrates the median separation for each comparison. The number of galaxies matched and the median plus semi-interquartile ranges of the displacements are presented in Table 4, columns 3, 4 and 5. It is clear from the comparison between the simulations run on

different computer architectures that there is an intrinsic offset in matched positions at the level of a few softening lengths. Furthermore, a significant number of matches exceeded this level by as much as a factor of 4. As a consequence of this, we regard as significant only those comparisons that have median offsets in excess of the level seen for these two simulations.

The first set to show a significant change is the simulations with different values of N . Varying the mass resolution causes differences in the number of objects formed and also affects the background dark matter distribution, so it is of no surprise that the




Figure 10. Separations between matched pairs of galaxies in the fiducial simulation and all other simulations, plotted against the average galaxy mass. The labels indicate the parameter that varies in each plot. The median scatter is shown as a solid horizontal line. Matches above $100 h^{-1}$ hpc are not plotted.

displacements show a significantly larger amount of scatter in these comparisons. Varying the value of N_{SPH} also produces very different numbers of galaxies, but this does not affect the median separation at all, demonstrating that it is the difference in the dark matter that primarily drives the increased scatter in the galaxy positions rather than mergers or interactions between the different numbers of galaxies.

A systematic difference in the matched positions is also introduced when the initial redshift of the simulation is changed:

the median displacement increases by a factor of 2 in both comparisons. Variations in the initial positions and velocities of the particles leads to asynchrony in the subsequent spatial trajectories of the galaxies.

Regarding the physical comparisons, substantially increasing the initial temperature of the gas leads to a significant effect. Particles that have been supplied with this much thermal energy naturally have greater pressure support, which directly affects the equation of motion, and therefore the particle trajectories.




Figure 11. Separations between matched pairs of galaxies (continued).

4.3.2 Galaxy–galaxy masses

We have also analysed the scatter in the masses of the matched galaxy pairs. Figs 12 and 13 illustrate the results by showing the logarithm of the mass ratio of each matched pair plotted against their average baryonic mass, always defined as the variant relative to the fiducial case. The solid line illustrates a ratio of unity (i.e., equal masses), and the dashed line illustrates the median value of the mass ratio. The median and semi-interquartile ranges are given in columns 6 and 7 of Table 4.

Adopting the tabulated cooling function causes an increase in the mass of each galaxy, because of an overall enhancement in the cooling rate around the temperatures most appropriate for these simulations ($\sim 10^5$ – 10^7 K). Decoupling the hot gas from the cold gas causes the masses of the galaxies to decrease slightly with only a few exceptions, notably the largest object. Decreasing the softening length from $\epsilon_0 = 10$ to $5 h^{-1}$ hpc also has the effect of decreasing the masses of the galaxies. Similarly, a larger value of N_{SPH} causes the masses of the galaxies to be systematically lower, due to its effect on the density field and hence the cooling rate.

Table 4. Statistical measures of the scatter present in all simulation comparisons of the matched galaxies and haloes. The first two columns detail the particular comparison, and the third column lists the number of galaxies in each catalogue that were matched with the fiducial simulation, using the method detailed in the text. The rest of the columns list the median and semi-interquartile range (siqr) for the measured scatter in galaxy separation (in h^{-1} hpc), galaxy mass, halo dark matter mass and halo baryon mass respectively.

No.	Alteration	Matches	Δr (galaxies)		$\Delta \log M$ (galaxies)		$\Delta \log M$ (dark matter)		$\Delta \log M$ (baryons)	
			median	siqr	median	siqr	median	siqr	median	siqr
1	Architecture	50	17.	6.0	0.00	0.01	0.003	0.01	-0.002	0.01
2	Position Offset (1 cell)	49	21.	5.4	0.00	0.02	0.007	0.02	0.0	0.02
2	Position Offset (random)	52	15.	5.5	0.00	0.02	0.004	0.009	0.002	0.01
3	New Cooling Function	53	14.	5.8	0.04	0.01	0.0	0.01	0.01	0.02
4	Cold gas decoupled	49	19.	4.0	-0.03	0.03	-0.008	0.01	-0.02	0.03
5	$s_{\max} = 0.6$	51	24.	7.3	0.01	0.02	0.002	0.02	0.02	0.02
6	$\kappa = 0.5$	51	19.	7.1	0.00	0.01	-0.002	0.01	0.004	0.01
6	$\kappa = 2.0$	50	22.	7.4	0.00	0.03	0.004	0.01	-0.002	0.02
7	$\epsilon_0 = 5 h^{-1}$ hpc	44	23.	8.7	-0.08	0.03	-0.007	0.02	-0.07	0.02
7	$\epsilon_0 = 20 h^{-1}$ hpc	50	15.	4.9	0.01	0.02	0.0	0.01	0.01	0.01
8	$N_{\text{SPH}} = 16$	53	18.	5.9	0.09	0.04	-0.003	0.02	-0.04	0.02
8	$N_{\text{SPH}} = 64$	28	16.	4.3	-0.06	0.04	-0.003	0.02	-0.02	0.03
9	$N = 2 \times 24^3$	14	95.	62.	-0.13	0.06	-0.01	0.05	0.07	0.05
9	$N = 2 \times 48^3$	51	93.	55.	0.02	0.06	0.005	0.03	0.01	0.05
10	$1 + z_i = 10$	50	46.	16.	0.01	0.03	0.0	0.04	0.01	0.03
10	$1 + z_i = 50$	48	39.	23.	0.00	0.03	-0.001	0.03	-0.004	0.03
11	$\Omega_b = 0.03$	48	26.	7.5	-0.41	0.03	-0.006	0.01	0.34	0.03
11	$\Omega_b = 0.12$	51	18.	4.9	0.37	0.04	-0.02	0.01	0.40	0.02
12	$Z = 0.0 Z_{\odot}$	9	19.	1.7	-0.56	0.04	0.004	0.01	0.10	0.04
12	$Z = 1.0 Z_{\odot}$	51	14.	4.7	0.05	0.03	0.0	0.02	0.03	0.01
13	$T_i = 0$ K	51	18.	6.7	0.00	0.02	0.0	0.01	0.003	0.01
13	$T_i = 10^7$ K	49	41.	18.	-0.02	0.02	-0.01	0.03	0.002	0.03

The dispersion in the mass ratio for the comparisons with varying N is large because of the large scatter in object positions for these runs. This makes spurious matches much more likely. For $N = 2 \times 24^3$ the objects are significantly less massive, whilst for $N = 2 \times 48^3$ the objects are only a little more massive than the fiducial simulation.

The largest differences in galaxy mass result from altering the cooling rate via the physical parameters, Ω_b and Z . For the comparisons with varying baryon fraction, the galaxies are more than a factor of 2 heavier for the larger value of Ω_b . Furthermore, there appears to be a slight trend for increasing mass excess for heavier galaxies, but this is difficult to measure due to the small dynamic range in mass. Varying the metallicity produces a similar effect to varying Ω_b .

4.3.3 Dark matter halo–halo masses

The method used for comparing masses of matched galaxy pairs was also performed on the dark matter haloes. Columns 8 and 9 in Table 4 quantify the differences. Overall, very tight correlations are seen for most of the simulations, as the dark matter is the dominant source of the gravitational potential, and so it is not significantly affected by changes in the gas dynamics. The simulation with $\Omega_b = 0.12$ shows the largest difference – here the contribution from the gas starts to make a difference to the properties of the dark matter haloes.

4.4 The distribution of baryons in haloes

The final item we investigate in this study is how parameter variations affect the final distribution of the baryons within haloes, as a function of the halo virial mass. Specifically, we looked at the mass fraction of baryons present in the galaxy phase, the ratio of

baryon masses for matched haloes and the total halo baryon fraction normalized to the global value.

4.4.1 Fraction of baryons in galaxies

Fig. 14 illustrates the fraction of baryons in the galaxy phase as a function of the halo virial mass, for each comparison. Because of our small samples, we have binned the data as histograms; the last bin corresponds to the single largest object. This statistic is remarkably stable; the first two comparisons show virtually no difference. The trend is for higher mass haloes to have less of their mass in galaxies. The mass fraction varies from ~ 0.65 in the smallest haloes down to ~ 0.2 in the highest mass halo. The comparison between the fiducial simulation and the case in which the galaxy phase is decoupled from the hot gas reflects the reduction in the total amount of galaxy phase baryons in the latter – about a factor of 3 for the galaxy in the largest halo ($\sim 5 \times 10^{13} h^{-1} M_{\odot}$), but only ~ 10 per cent for haloes with masses of $\sim 10^{11} h^{-1} M_{\odot}$. Other comparisons that give significant changes in the mass fractions show nothing unexpected.

4.4.2 Local halo baryon fraction

The variations in the total mass of baryons within the virial radius of matched haloes are quantified in columns 10 and 11 of Table 4. The median differences are slightly higher than for the dark matter, but smaller than the galaxies alone. Interestingly, the sign of the median difference is not always the same as the equivalent comparison for the galaxies, indicating that fluctuations in the total baryon content are not due to the galaxies alone. Significant differences are most evident in the simulations in which parameter changes have affected the cooling rate. Inducing a higher cooling rate causes the gas to lose some of its pressure support, and




Figure 12. The logarithm of the ratio of the masses of matched pairs of galaxies plotted against the logarithm of the mass of the fiducial galaxy. Labels indicate the parameter corresponding to each comparison (with the relative difference defined as the second minus the first value). The solid line corresponds to equal masses, and the dashed line illustrates the median value for each comparison.

therefore to become more compressed within its halo. The simulation with $\epsilon_0 = 5 h^{-1} \text{ hpc}$ shows a significant change in the same direction as the change in the galaxy masses, supporting the idea that the gas is being more strongly heated from two-body encounters with the dark matter particles.

We also analysed the total baryon fraction in haloes as a function of the halo virial mass. The results are displayed in Fig. 15. The fraction is normalized to the global value of Ω_b for each simulation. The fiducial simulation has halo baryon fractions

ranging from ~ 93 per cent of the global value in the smallest haloes to ~ 82 per cent in the largest. All simulations agree with the fiducial values to within ~ 20 per cent. Notably, not all simulations show a monotonically decreasing baryon fraction with halo mass. However, the scatter is considerable because of the small number of objects in each bin (with only one in the highest mass bin). The differences reflect the varying amounts of material that has cooled. Note, for example, the drop in the first bin for $Z = 0.0$, for which galaxies do not form in small haloes, and the




Figure 13. The logarithm of the ratio of the masses of matched pairs of galaxies (continued).

rise in the first bin for $N_{\text{SPH}} = 16$. For almost all of the simulations and across the entire range of halo masses, the baryon fraction is less than the global value. The baryon fraction in the largest halo in this simulation (0.82 for the fiducial case) is somewhat smaller than the values, 0.9–1.0, found in the rich galaxy cluster simulations presented by Frenk et al. (1999). The value for our largest cluster, however, is noisy, as indicated by the first comparison in the figure, and it depends on metallicity, as indicated also in the figure. Furthermore, this cluster is approximately 10 times less massive than the cluster in Frenk

et al. and, as shown by Pearce et al. (in preparation), there is a weak trend of increasing baryon fraction with mass for clusters in the SCDM cosmology.

5 SUMMARY OF THE COMPARISON RESULTS

We now summarize the main effects seen in each comparison, in turn.

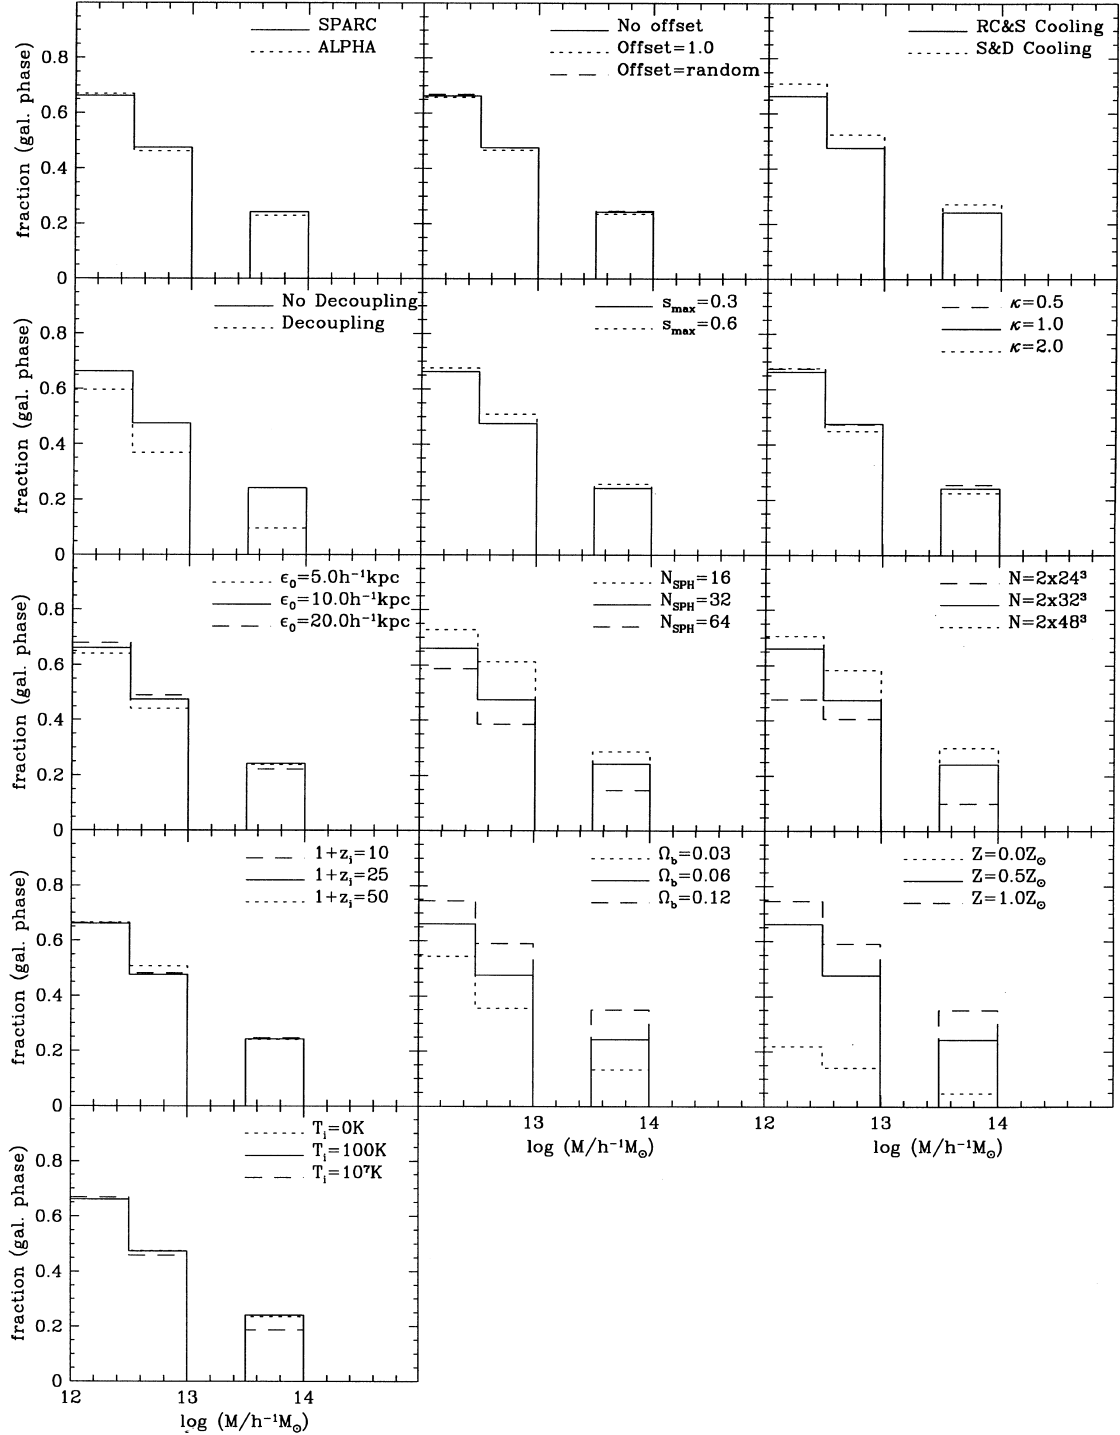


Figure 14. The mass fraction of baryons in the galaxy phase plotted as a function of the total halo mass. The fiducial simulation is the solid histogram.

In comparison 1, we looked at the effects introduced by running the simulation on different computer architectures. No significant differences were found between the two simulations. However, the displacements between matched pairs of galaxies at $z=0$ were found to be of the order of an S2 softening length. It is therefore incorrect to trust either the positions of individual galaxies or their trajectories on scales smaller than this.

Comparison 2 was intended to look at the effects of small displacements in the positions of the particles at the start of the simulation. We ran one simulation with the particles

initially displaced by 1 FFT cell, to test the periodic nature of the algorithm. We also ran a simulation with tiny random perturbations applied to the initial particle positions, to test the level at which differences can grow over the course of a few thousand time-steps (the typical length of each simulation). Again, no significant changes were detected amongst the three simulations.

We next changed the cooling function from the series of power-law fits, used in the fiducial simulation, to a set of tabulated values from Sutherland & Dopita (1993). A slight enhancement in the

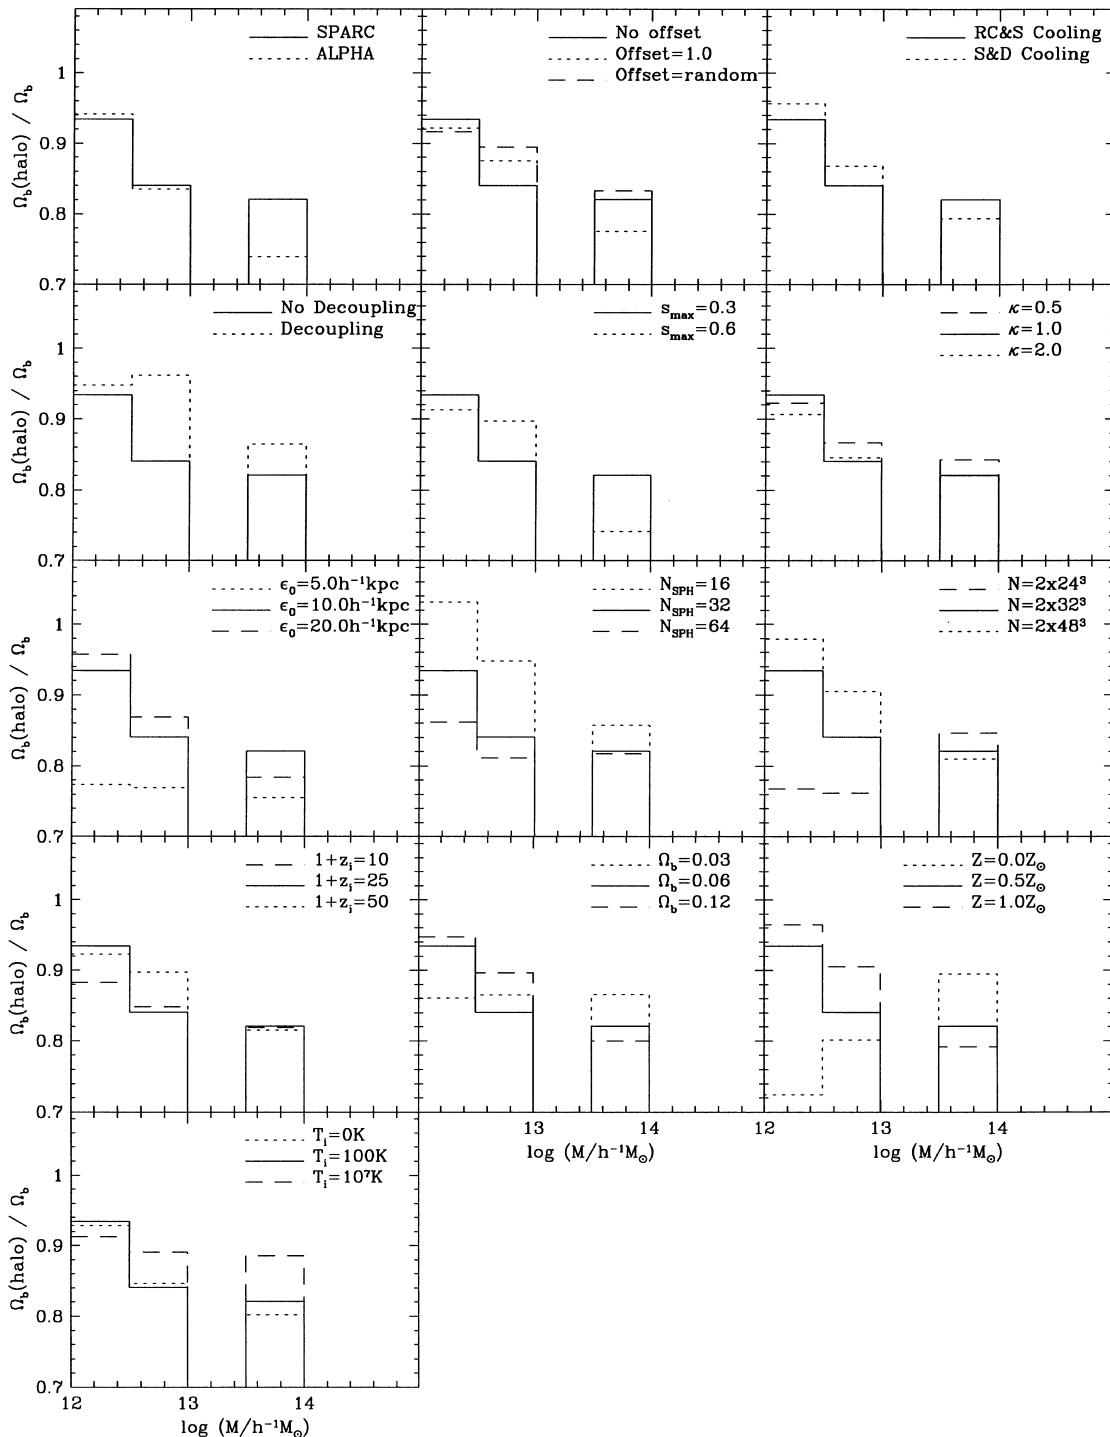


Figure 15. The baryon fraction in haloes as a function of the total halo mass, in units of the global baryon fraction. The fiducial simulation is the solid histogram.

galaxy masses, at the 10 per cent level, was produced in the latter case, although no significant changes were produced in the spatial distribution of galaxies.

In comparison 4, we compared the fiducial simulation to one in which the galaxy phase gas is partially decoupled from hot gas in the SPH algorithm (see Pearce et al. 1999). This modification provides a better estimate of the hot gas density near galaxies; specifically, it prevents an overestimate which leads to an artificial enhancement in the cooling rate. Only the largest galaxy in the

simulation was significantly affected, with its mass being reduced by a factor of 3.

Comparison 5 consisted of doubling the initial (maximum) softening length. This results in a period of time in which the softening is larger than in the fiducial simulation. This led to slight differences in the spatial distribution of the galaxies, although it left the mass distribution unaffected.

We altered the size of the time-steps by a factor of 2 each way, in comparison 6, by changing the constant κ . Increasing the

time-step resulted in much more hot gas at the expense of the cold phase. This illustrates the effect of the time-step on the dynamics of the particles: introducing significant errors into the positions and velocities leads to enhanced heating of the gas. The galaxies themselves remained largely unaffected. No significant change was evident when halving the value of κ , thus justifying the value $\kappa = 1$, adopted as fiducial.

Comparison 7 consisted of changing the final value of the softening length, ϵ_0 . In the simulation with $\epsilon_0 = 5 h^{-1}$ hpc, the galaxies were about 20 per cent less massive than in the fiducial simulation ($\epsilon_0 = 10 h^{-1}$ hpc), possibly due to discreteness effects caused by harder two-body encounters. A larger value of $\epsilon_0 = 20 h^{-1}$ hpc, which implies a constant comoving softening, produced no significant differences.

In comparison 8, we varied the value of N_{SPH} , making it larger and smaller by a factor of 2, thus looking at the effects produced when the SPH algorithm uses a different number of neighbours when calculating the smoothing length for each gas particle. These changes had a dramatic effect on the simulations. The value of N_{SPH} determines the minimum mass for which the gas density can be resolved: smaller values will resolve smaller masses at the expense of larger random fluctuations. We found that the number of galaxies increases strongly with smaller N_{SPH} , as did the masses of the matched galaxy pairs, though without significantly affecting their positions.

Comparison 9 consisted of changing the number of particles in the simulation, thus altering the mass resolution. To demonstrate the effect of spurious heating by dark matter particles, we ran a simulation with a dark matter particle mass above the critical value calculated by Steinmetz & White (1997). In this case, gas is unable to cool and form galaxies. For the main set of comparisons, we had a suite of runs with $N = 2 \times [24, 32, 48]^3$. The number and mass of galaxies increased with N . Improving the mass resolution also affected the final distribution of dark matter. The two simulations with the highest resolution were in better agreement than the two with the lowest resolution, although the discrepancies in the galaxy phase material (e.g., the fraction of cooled gas) were still large.

In comparison 10, the initial redshift of the simulations was changed ($1 + z_i = 10, 25, 50$). The masses of the objects were not significantly altered by this change, but small changes in the spatial distribution of the galaxies were produced.

Comparison 11 involved changing the value of the baryon fraction ($\Omega_b = 0.03, 0.06, 0.12$). The higher value of Ω_b causes the cooling rate to increase (since emissivity is proportional to Ω_b^2). The mass of each galaxy also increases, since m is proportional to Ω_b . The run with $\Omega_b = 0.03$ has dark matter masses of the order of the critical mass for two-body heating, although the reduction in the cooling rate produced by this effect is confused by the fact that the cooling rate is significantly less anyway. Overall, large differences were noted in all analyses in which the baryon fraction was varied.

In comparison 12, we changed the value of the (fixed) gas metallicity ($Z = 0.0, 0.5$ and 1.0 in solar units). The run with no metals differed dramatically from the fiducial simulation, producing very few objects by $z = 0$. This simulation had a dark matter mass well above the critical value for two-body heating. The run with $Z = 1.0 Z_\odot$ showed an enhancement in galaxy masses relative to the fiducial simulation, although the change was far less severe.

Finally, in comparison 13, the initial gas temperature was changed from 100 K to 0 K and 10^7 K respectively. No significant changes were measured in the former case. Although the uncollapsed gas in the latter simulation is on a higher adiabat

($T \sim 10^3$ K), this had very little effect on the final state of the gas in galaxies and haloes, with the exception of introducing significant displacements between matched galaxy pairs.

6 CONCLUSIONS

In this paper we have explored how variations in simulation parameters and physical assumptions affect the population of galaxies which form in cosmological simulations. We used the AP^3M algorithm to calculate the gravitational potential field, and SPH with radiative cooling to model gas dynamics. We started with a carefully chosen fiducial simulation and calculated a number of variants, each differing in only one parameter. Galaxy and halo catalogues were selected in a consistent manner before analysing the effects that the parameter changes had on the thermodynamic state of the gas, on the mass and spatial distribution of the galaxies, and on the baryon content of haloes.

By carefully tracking the phase evolution of gas particles we have shown that most galaxy material (around 90 per cent in the fiducial simulation) originally cools in small haloes where the virial temperature is $\sim 10^5$ K. As Fig. 14 shows, the percentage of halo baryons that are cold rises as the halo mass falls. Only a small fraction of the baryons cool at later stages of the evolution. Mergers of these small fragments can reheat the gas briefly as massive galaxies are built up. This implies that the most dynamically interesting part of the process occurs at the resolution threshold of the simulation, and it is therefore sensitive to many numerical and physical parameters.

Numerical noise leads to a scatter in the final galaxy positions that is of the order of the (spline) softening length for most galaxies, but can be much higher in some cases. Individual galaxy trajectories should therefore be treated with caution.

The length of time-step recommended by CTP95 is adequate. Halving it caused no significant change in the galaxy catalogues, but choosing a larger value led to a significant increase in the shocked gas relative to the uncollapsed gas. This can be attributed to errors in the particle positions and velocities, resulting in an artificial enhancement in the thermal heating of the gas.

We confirm that simulations with dark matter particle masses greater than the Steinmetz & White (1997) critical mass for two-body heating do indeed show a catastrophic reduction in the final amount of cooled gas and in the number of galaxies, even when the shortest cooling time is a small fraction of the age of the Universe. The choice of dark matter particle mass should be comfortably below this critical value. The mass in galaxies is also strongly affected by the choice of baryon fraction and metallicity.

Pre-heating of the gas to a high initial temperature (10^7 K at $z = 24$) produces small changes in the final positions of the galaxies but has little effect on their final mass. A significant fraction of the thermal energy of the gas is radiated away around $z = 3-5$, allowing the majority of the baryons to undergo gravitational collapse.

Increasing the mass resolution by using a larger number of particles causes both the number of galaxies and their masses to increase. Improved resolution of the density field enhances the cooling rate within haloes, particularly at high redshift when the first objects form. As the mass resolution of the simulation increases, modelling feedback processes correctly becomes critical in order to avoid overcooling of baryonic material.

Lowering the number of neighbours used by the SPH algorithm to calculate hydrodynamical quantities has a similar effect on the

amount of cooled material as increasing the total number of particles in the simulation. Smaller values of N_{SPH} enable higher overdensities to be resolved (therefore enhancing the cooling rate in such regions), at the expense of more contamination from sampling noise.

The fraction of the baryonic material that cools within a halo and forms galaxies is a very stable quantity which varies inversely with halo mass. Reducing the total amount of gas affects all scales in a similar way. The ratio of baryonic to dark material within the virial radius of haloes is found to be around 0.8 ± 0.1 of the global value, for the range of halo masses studied here.

In this study we have used four different prescriptions for the gravitational softening. We have varied the initial comoving softening, from $20 h^{-1}$ hpc to $40 h^{-1}$ hpc, and we have also used a range of final physical softenings, from $5 h^{-1}$ hpc to $10 h^{-1}$ hpc and $20 h^{-1}$ hpc. With the exception of the $5 h^{-1}$ hpc case, this variation had very little effect on the galaxy population. The smaller final softening systematically reduced both the mass of galaxies and the mass of baryons within haloes by ~ 20 per cent. This effect is probably due to the artificial heating of the gas from two-body encounters with the dark matter particles discussed above (Steinmetz & White 1997). The fiducial softening is borderline for this resolution.

Finally, in almost every simulation we obtained an overly massive galaxy at the centre of the largest halo. With a baryonic mass of $\sim 5 \times 10^{11} h^{-1} M_{\odot}$, such an object is not expected within such a small volume of the Universe. Decoupling the central dense gas from the hot halo material effectively reduces the mass of this object to a more reasonable value (see also Pearce et al. 1999).

We conclude that varying most of the numerical or physical parameters required in simulations of galaxy formation within reasonable bounds produces only relatively small changes in the abundance, mass and spatial distribution of the galaxies that form. The most important numerical parameter is mass resolution, which has an important effect because most of the gas cools in low-mass haloes (in which the cooling times are shortest). The most important physical parameters are those that control the cooling time of the gas, the density of baryons and its metallicity. Varying these within plausible ranges can substantially change the abundance and mass of galaxies. Careful modelling of physical processes occurring in low-mass haloes where the thermodynamic state of the gas is likely to be affected, for example, by energy injected by supernovae and stellar winds is the next great challenge facing simulators of galaxy formation.

ACKNOWLEDGMENTS

We thank the referee, Matthias Steinmetz, for his constructive suggestions that have improved the clarity of the manuscript. The work presented in this paper was carried out as part of the programme of the Virgo Supercomputing Consortium (<http://star-www.dur.ac.uk/~frazierp/virgo/>). We acknowledge a NATO collaborative research grant (CRG 970081). This work was supported by the EC network for ‘Galaxy formation and evolution’. STK acknowledges the support of a PPARC postgraduate studentship; CSF acknowledges a PPARC Senior Research Fellowship and a Leverhulme Research Fellowship.

REFERENCES

Anninos P., Norman M. L., 1996, *ApJ*, 459, 12

- Benz W., Buchler J. R., 1990, *Numerical Modelling of Stellar Pulsations: Problems and Prospects*, Kluwer, Dordrecht p. 269
- Binney J., Tremaine S., 1987, *Galactic Dynamics*. Princeton Series in Astrophysics
- Bond J. R., Efstathiou G., 1984, *ApJ*, 285, L45
- Bryan G. R., Norman M. L., 1998, *ApJ*, 495, 80
- Cen R., Ostriker J., 1992, *ApJ*, 393, 22
- Cen R., Ostriker J., 1994, *ApJ*, 429, 4
- Cen R., Miralda-Escudé J., Ostriker J. P., Rauch M., 1994, *ApJ*, 437, L9
- Cen R., Phelps S., Miralda-Escudé J., Ostriker J. P., 1998, *ApJ*, 496, 577
- Centrella J., Melott A. L., 1983, *Nat*, 305, 196
- Cole S., 1991, *ApJ*, 367, 45
- Copi C. J., Schramm D. N., Turner M. S., 1995, *ApJ*, 455, 95
- Couchman H. M. P., 1991, *ApJ*, 368, L23
- Couchman H. M. P., Thomas P. A., Pearce F. R., 1995, *ApJ*, 452, 797 (CTP95)
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- Efstathiou G., Davis M., Frenk C. S., White S. D. M., 1985, *ApJS*, 57, 241
- Efstathiou G., Davis M., Frenk C. S., White S. D. M., 1988, *MNRAS*, 235, 715
- Eke V. R., Cole S. M., Frenk C. S., 1996, *MNRAS*, 282, 263
- Eke V. R., Navarro J. F., Frenk C. S., 1998, *ApJ*, 503, 569
- Evrard A. E., 1990, *ApJ*, 363, 349
- Evrard A. E., Summers F. J., Davis M., 1994, *ApJ*, 422, 11
- Frenk C. S. et al., 1999, *ApJ*, 525, 554
- Gingold R. A., Monaghan J. J., 1977, *MNRAS*, 181, 375
- Hernquist L., Katz N., Weinberg D. H., Miralda-Escudé J., 1996, *ApJ*, 457, L57
- Hockney R. W., Eastwood J. W., 1981, *Computer Simulation Using Particles*. McGraw-Hill
- Katz N., 1992, *ApJ*, 391, 502
- Katz N., Hernquist L., Weinberg D. H., 1992, *ApJ*, 399, L109
- Katz N., Hernquist L., Weinberg D. H., 1999, in Carilli C. L., Radford S. J. E., Menten K. M., Langston G. I., eds, *ASP Conf. Series*, Vol. 156, *Highly Redshifted Radio Lines*. Astron. Soc. Pac., San Francisco, p. 1
- Klypin A., Shandarin S. F., 1983, *MNRAS*, 204, 891
- Lacey C., Cole S., 1994, *MNRAS*, 271, 676
- Lucy L., 1977, *AJ*, 82, 1013
- Monaghan J. J., 1992, *ARA&A*, 30, 543
- Monaghan J. J., Gingold R. A., 1983, *J. Comput. Phys.*, 52, 374
- Mushotzky R., Loewenstein M., Arnaud K. A., Tamura T., Fukazawa Y., Matsushita K., Kikuchi K., Hatsukade I., 1996, *ApJ*, 466, 686
- Navarro J. F., Steinmetz M., 1997, *ApJ*, 478, 13
- Navarro J. F., White S. D. M., 1993, *MNRAS*, 265, 271
- Navarro J. F., Frenk C. S., White S. D. M., 1995, *MNRAS*, 275, 720
- Pearce F. R., Couchman H. M. P., 1997, *New Astronomy*, 2, 411
- Pearce F. R. et al., 1999, *ApJ*, 521, L99
- Raymond J. C., Cox D. P., Smith B. W., 1976, *ApJ*, 204, 290
- Steinmetz M., Müller E., 1995, *MNRAS*, 276, 549
- Steinmetz M., White S. D. M., 1997, *MNRAS*, 288, 545
- Sutherland R. S., Dopita M. A., 1993, *ApJS*, 88, 253
- Thacker R. J., Tittley E. R., Pearce F. R., Couchman H. M. P., Thomas P. A., 1998, submitted *MNRAS* (astro-ph/9809221)
- Theuns T., Leonard A., Efstathiou G., Pearce F. R., Thomas P. A., 1998, *MNRAS*, 301, 478
- Thomas P. A., Couchman H. M. P., 1992, *MNRAS*, 257, 11 (TC92)
- Viana P. T. P., Liddle A. R., 1996, *MNRAS*, 281, 323
- Weinberg D. H., Hernquist L., Katz N., 1997, *ApJ*, 477, 8
- White S. D. M., Frenk C. S., 1991, *ApJ*, 379, 52
- White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341
- White S. D. M., Efstathiou G., Frenk C. S., 1993, *MNRAS*, 262, 1023
- Wood D., 1981, *MNRAS*, 194, 201

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.