Galaxies–intergalactic medium interaction calculation – I. Galaxy formation as a function of large-scale environment

Robert A. Crain,[1,2]★ Tom Theuns,[1,3] Claudio Dalla Vecchia,[4] Vincent R. Eke,[1] Carlos S. Frenk,[1] Adrian Jenkins,[1] Scott T. Kay,[5] John A. Peacock,[6] Frazer R. Pearce,[7] Joop Schaye,[4] Volker Springel,[8] Peter A. Thomas,[9] Simon D. M. White[8] and Robert P. C. Wiersma[4] (The Virgo Consortium)

[1] *Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE*
[2] *Centre for Astrophysics and Supercomputing, Mail No. H39, Swinburne University of Technology, PO Box 218, Hawthorn, Victoria 3122, Australia*
[3] *Department of Physics, University of Antwerp, Campus Groenenborger, Groenenborgerlaan 171, B-2020 Antwerp, Belgium*
[4] *Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, Netherlands*
[5] *Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, Alan Turing Building, University of Manchester, Manchester M13 9PL*
[6] *Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh EH9 3HJ*
[7] *Department of Physics and Astronomy, University of Nottingham, Nottingham NG7 2RD*
[8] *Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85740 Garching bei München, Germany*
[9] *Astronomy Centre, University of Sussex, Falmer, Brighton BN1 9QH*

ABSTRACT

We present the first results of hydrodynamical simulations that follow the formation of galaxies to the present day in nearly spherical regions of radius $\sim 20\,h^{-1}$ Mpc drawn from the Millennium Simulation (Springel et al.). The regions have mean overdensities that deviate by $(-2, -1, 0, +1, +2)\sigma$ from the cosmic mean, where σ is the rms mass fluctuation on a scale of $\sim 20\,h^{-1}$ Mpc at $z = 1.5$. The simulations have mass resolution of up to $\sim 10^6\,h^{-1}\,M_\odot$, cover the entire range of large-scale cosmological environments, including rare objects such as massive clusters and sparse voids, and allow extrapolation of statistics to the $(500\,h^{-1}\,\mathrm{Mpc})^3$ Millennium Simulation volume as a whole. They include gas cooling, photoheating from an imposed ionizing background, supernova feedback and galactic winds, but no AGN. In this paper, we focus on the star formation properties of the model. We find that the specific star formation rate density at $z \lesssim 10$ varies systematically from region to region by up to an order of magnitude, but the global value, averaged over all volumes, closely reproduces observational data. Massive, compact galaxies, similar to those observed in the GOODS fields (Wiklind et al.), form in the overdense regions as early as $z = 6$, but do not appear in the underdense regions until $z \sim 3$. These environmental variations are not caused by a dependence of the star formation properties on environment, but rather by a strong variation of the halo mass function from one environment to another, with more massive haloes forming preferentially in the denser regions. At all epochs, stars form most efficiently in haloes of circular velocity $v_c \sim 250\,\mathrm{km\,s^{-1}}$. However, the star formation history exhibits a form of 'downsizing' (even in the absence of AGN feedback): the stars comprising massive galaxies at $z = 0$ have mostly formed by $z = 1$–2, whilst those comprising smaller galaxies typically form at later times. However, additional feedback is required to limit star formation in massive galaxies at late times.

Key words: methods: *N*-body simulations – galaxies: abundances – galaxies: clusters: general – galaxies: formation – intergalactic medium.

1 INTRODUCTION

Numerical simulations have emerged, over the past two decades or so, as a useful technique for modelling the formation and

★E-mail: rcrain@astro.swin.edu.au

evolution of cosmic structures. In particular, they have yielded accurate predictions for the clustering and evolution of dark matter, a relatively simple problem in which the dynamics are determined purely by gravitational forces (e.g. Davis et al. 1985). Modelling the evolution of small-scale structures involving baryons is considerably more challenging because of the complexity of the physics involved and the greater uncertainties in the modelling techniques (e.g. Katz & Gunn 1991). In the regime where baryonic processes are non-negligible, the ability of numerical simulations to yield unique and robust predictions is therefore reduced. This is particularly true for systems in which radiative cooling is efficient.

The inherent difficulties and computational expense of modelling baryons in this regime has stimulated the development of semi-analytic techniques (White & Frenk 1991; Kauffmann, White & Guiderdoni 1993; Cole et al. 1994, 2000; Somerville & Primack 1999; Baugh 2006). These employ simplified prescriptions for the evolution of baryons in dark matter haloes. The formation histories of the haloes are followed using merger trees constructed analytically with Monte Carlo techniques, or extracted directly from *N*-body simulations. Notable examples of the latter procedure are the studies of Bower et al. (2006), Croton et al. (2006), De Lucia et al. (2006) and Font et al. (2008), based on merger histories derived from the Millennium Simulation (Springel et al. 2005). The scale and resolution of this dark matter simulation permits semi-analytical calculations of the evolution of galaxies more massive than the Small Magellanic Cloud within a comoving volume similar to that probed by the 2dF galaxy redshift survey (2dFGRS) (Colless et al. 2001) and Sloan Digital Sky Survey (SDSS) (York et al. 2000) at their median redshifts.

Simulations employing semi-analytic models have yielded critical insights into the effects of baryonic processes on the observable Universe. Nevertheless, the approximations inherent in this technique limit the kind of processes that can be followed reliably in a relatively simple manner. For example, while semi-analytic techniques can be extended to study the evolution of the intergalactic medium (IGM; e.g. Benson et al. 2002), they are not well suited for studying the interaction of gas with galactic ejecta, such as winds, and the associated dynamical evolution. Processes of this kind can only be studied in detail with simulations that explicitly follow the full hydrodynamic evolution of the baryons (Wiersma et al. 2009b).

The introduction and subsequent development of hydrodynamic numerical methods within the framework of cosmological simulations has led to important advances in our understanding of structure formation. For example, the relationship between the Lyman-α forest and the cosmic large-scale structure was first convincingly demonstrated using such simulations (e.g. Cen et al. 1994; Zhang, Anninos & Norman 1995; Hernquist et al. 1996; Miralda-Escudé et al. 1996; Zhang et al. 1997; Theuns et al. 1998; Davé et al. 1999). At the opposite end of the density scale, hydrodynamic simulations of the formation of individual galaxies have confirmed the generic requirement in hierarchical models (first identified using semi-analytic techniques) for energy feedback mechanisms to prevent the overcooling of gas in small haloes at early times (White & Rees 1978; Cole 1991; White & Frenk 1991), and the associated transfer of angular momentum from baryons to dark matter (e.g. Navarro & Benz 1991; Weil, Eke & Efstathiou 1998; Thacker & Couchman 2001; Abadi et al. 2003a,b; Sommer-Larsen, Götz & Portinari 2003; Governato et al. 2004; Robertson et al. 2004; Okamoto et al. 2005; Governato et al. 2007; Scannapieco et al. 2008; Zavala, Okamoto & Frenk 2008).

The traditional approach of hydrodynamical simulations has involved a trade-off between resolution and computational volume.

Such trade-offs, however, are not possible if an unbiased statistical sample of well-resolved simulated galaxies is desired. This presents a considerable computational challenge, and consequently previous simulations have been limited to tracing volumes that are relatively small in comparison with those used in large-scale structure calculations and those probed by observational surveys. The gulf in mass and spatial scale between individual galaxies and their large-scale environment compounds the more fundamental dynamic range problem faced by simulations of galaxy formation, that stems from the influence of individual stars (via stellar winds and supernovae) on gas.

Recent attempts to produce such samples include the smoothed particle hydrodynamics (SPH) simulations of Pearce et al. (1999, 2001) using the HYDRA code (Couchman, Thomas & Pearce 1995); those of Davé, Finlator & Oppenheimer (2006), Oppenheimer & Davé (2006), Davé & Oppenheimer (2007), Oppenheimer & Davé (2008a), Croft et al. (2008), Oppenheimer & Davé (2008b), using variants of the GADGET2 code (Springel 2005); and those of Brooks et al. (2007) using the GASOLINE code (Wadsley, Stadel & Quinn 2004). The adaptive mesh refinement (AMR) technique for hydrodynamics has also been employed for simulations of this type, e.g. by Nagamine et al. (2005, see references therein for code details), and by Ocvirk, Pichon & Teyssier (2008), using the RAMSES code (Teyssier 2002). To highlight an example, Croft et al. (2008) simulated (with 486^3 particles of gas and dark matter) a periodic volume in which haloes of Milky Way type galaxies were resolved with $\sim 10^5$ particles. Although ~ 1000 galaxies formed in this simulation, its relatively small size (a cube of side $L = 33.75\,h^{-1}$ Mpc) yielded few massive galaxies and sampled a relatively narrow range of cosmological environments. A further technical limitation of relatively small volumes is that fluctuations on scales comparable to the size of the box become non-negligible at a relatively high redshift (Bagla & Ray 2005; Sirko 2005), after which the simulation can no longer be considered a faithful representation of the underlying model. For this reason, Croft et al. (2008) halted their simulation at $z = 1$.

The need for detailed predictions of the properties of the low-redshift IGM and the evolution of galaxies over the interval $0 < z \lesssim 1$ is highlighted by a number of observational developments. The fate of the baryons seen in the Lyman-α forest at high redshift is still largely uncharted but measurements of the Lyman-α forest are beginning to probe the diffuse baryons associated with the filaments of the cosmic web (e.g. Pichon et al. 2001; Caucci et al. 2008), where a significant fraction may be in the form of a 'warm–hot intergalactic medium (WHIM)' phase (e.g. Cen & Ostriker 1999; Davé et al. 1999). The recent installation of the *Cosmic Origins Spectrograph* (Green, Wilkinson & Morse 2003) aboard the *Hubble Space Telescope* (*HST*) may yield significant advances in the understanding of the low-redshift IGM. Alongside studies of the evolution of cosmic gas, realistic modelling of the low-redshift evolution of galaxies is essential to connect the snapshots of cosmic evolution provided by galaxy surveys at different epochs, for example DEEP2 (Davis et al. 2003), VVDS (Le Fèvre et al. 2005) and COMBO-17 (Wolf et al. 2003) at $z \sim 1$ with the 2dFGRS and SDSS at $z \sim 0$.

Simulating large cosmological volumes ($L \gtrsim 100\,h^{-1}$ Mpc) at high resolution ($m_{gas} \lesssim 10^7\,h^{-1}\,M_\odot$) to $z = 0$ is infeasible with current computational resources. To circumvent this fundamental limitation, we have devised a new approach consisting of simulating at high resolution regions extracted from the dark matter Millennium Simulation whose overdensities at $z = 1.5$ represent $(-2, -1, 0, +1, +2)\sigma$ deviations from the cosmic mean, where σ is the rms mass fluctuation on a scale of $\sim 20\,h^{-1}$ Mpc. These
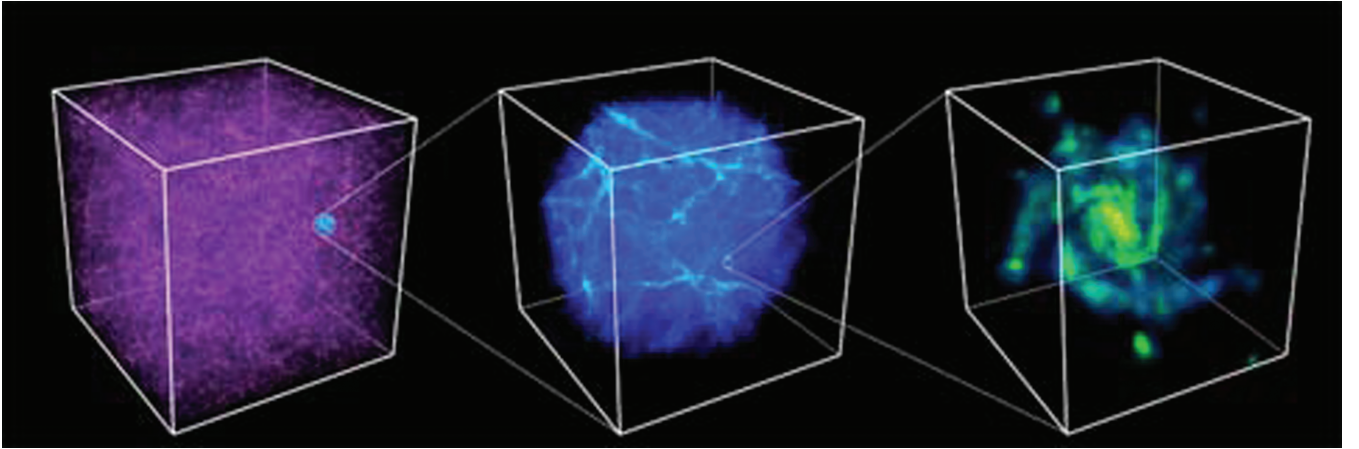
Figure 1. Schematic representation of GIMIC, illustrating the large dynamic range allowed by the 'zoomed' initial conditions technique. The left-hand panel shows the density distribution within the $L = 500\,h^{-1}$ Mpc periodic volume of the Millennium Simulation at $z = 1.5$. Shown embedded within this volume, to scale, is the 0σ GIMIC region. The centre panel shows the gas density of this region, zoomed down to $L \simeq 50\,h^{-1}$ Mpc. The right-hand panel, zoomed yet further to $L \simeq 50\,h^{-1}$ kpc, depicts the gas density within a disc galaxy of mass similar to that of the Milky Way.

simulations, the *Galaxies–Intergalactic Medium Interaction Calculation* (GIMIC), follow the evolution of five roughly spherical regions of radius $\sim 20\,h^{-1}$ Mpc from 'zoomed' initial conditions (Frenk et al. 1996; Power et al. 2003; Navarro et al. 2004) using full gas dynamics, while the rest of the Millennium volume is resimulated at lower resolution following dark matter only. We thus follow a very wide range of cosmological environments whilst only simulating 0.13 per cent of the Millennium volume at high resolution. The extent of the dynamic range allowed by our procedure is illustrated in Fig. 1 which, starting from the full Millennium Simulation volume ($L = 500\,h^{-1}$ Mpc), zooms in, first by a factor of 10 to show a full GIMIC sphere in the central panel and then by a further factor of 1000 to show a disc galaxy within this region.

In contrast to simulations of small periodic volumes, our adopted strategy allows us to follow the cosmic evolution to $z = 0$, since fluctuations on the scale of the Millennium Simulation are still well described by linear theory. Judicious selection of the five GIMIC regions across environmentally diverse regions yields a sample spanning the range of structures found within the Millennium Simulation. In our highest resolution realizations, the Jeans scale in the IGM after the epoch of reionization is marginally resolved. A further advantage of resimulating regions of the Millennium Simulation is that they complement the existing semi-analytic calculations implemented on it (Bower et al. 2006; Croton et al. 2006; De Lucia et al. 2006; Font et al. 2008).

The simulations were performed using GADGET3, a significantly upgraded variant of the GADGET2 code described by Springel (2005), and including new treatments of baryonic processes, such as radiative cooling, heating by the metagalactic ultraviolet (UV) background radiation, star formation, stellar feedback and chemodynamics. The simulations are computationally expensive, and so we have limited ourselves to running them with a single implementation of the code. Analyses of these simulations will therefore, in general, focus on the environmental effects that these simulations are able to explore; we do not examine the role of, for example, changes to the assumed initial mass function (IMF) or the implementation of feedback. Such tests will be explored in the closely related Overwhelmingly Large Simulations (OWLS; Schaye et al., in preparation) project. In this, the first paper in a series, we focus on one central aspect of these simulations, the difference in the star formation rate density (SFRD) amongst the five different environments, and on

how these combine to produce a 'cosmic' SFRD (e.g. Lilly et al. 1996; Madau et al. 1996).

The paper is laid out as follows. In Section 2, we describe the simulations, concentrating on new aspects of our code, and provide a brief overview of the generation of the initial conditions. Our suite of simulations includes runs with varying resolution to enable us to check for numerical convergence. In Section 3, we discuss the halo mass functions of the five GIMIC regions, which are very different from each other. In Section 4, we calculate the evolution of the star formation rate (SFR) of each region and present an estimate for the entire Millennium Simulation based on a weighted average of these measurements. The star formation properties of haloes and their galaxies are discussed in Section 5. In Section 6, we carry out a limited comparison with observational data. We summarize and conclude in Section 7.

The simulations adopt the same cosmological parameters as the Millennium Simulation: $\Omega_{\rm m} = 0.25$, $\Omega_{\Lambda} = 0.75$, $\Omega_{\rm b} = 0.045$, $n_{\rm s} = 1$, $\sigma_8 = 0.9$, $H_0 = 100\,h$ km s^{-1} Mpc^{-1}, $h = 0.73$. This work is part of the programme of the Virgo consortium for cosmological simulations.

2 METHODOLOGY

This section describes the three central aspects of our methodology: the code used to perform the simulations, the setup of the initial conditions and the identification of dark matter haloes and galaxies. Technical details of the generation of the initial conditions are deferred to an appendix.

2.1 Simulation code

For this study, we use GADGET3, an updated version of the GADGET2 code (Springel 2005), to which several physics modules have been added. The domain decomposition differs from that in GADGET2 in that it improves load balancing particularly for simulations (such as those described here) with strongly clustered particle distributions run on parallel supercomputers with a large number of cores. The hydrodynamics implementation, taken from GADGET2, is the entropy conserving formulation of SPH (Gingold & Monaghan 1977; Lucy 1977), as discussed in Springel & Hernquist (2002). SPH represents a fluid by a set of particles that carry along a number of internal

properties (e.g. mass, entropy), while other properties are computed by interpolation over neighbouring particles (e.g. the density or pressure gradient).

We have used new physics modules for star formation (Schaye & Dalla Vecchia 2008), stellar feedback (Dalla Vecchia & Schaye 2008), radiative cooling (Wiersma, Schaye & Smith 2009a) and chemodynamics (Wiersma et al. 2009b). Here, we only provide a brief overview of the main features of these modules.

(i) *Gas cooling and photoionization.* Radiative cooling was implemented as described in Wiersma et al. (2009a).[1] In brief, we assume the IGM is ionized and heated by a pervading uniform, but redshift-dependent, ionizing background from galaxies and quasars, as computed by Haardt & Madau (2001). We assume hydrogen reionizes at redshift $z = 9$ and He II at $z = 3.5$ (Schaye et al. 2000; Theuns et al. 2002). The cooling rate is computed as a function of redshift, gas density, temperature and composition on an element-by-element basis using interpolation tables computed with CLOUDY (as described in Ferland et al. 1998), assuming the gas to be optically thin and in ionization equilibrium. Owing to uncertainties in low-temperature metal-line cooling rates, we impose a minimum cooling temperature of 100 K; in practice, very few particles approach this limit after reionization because we impose an equation of state for cold, dense gas (see below) and because lower density gas is kept warm due to the inclusion of a photoionizing background. During the reionization of He II, we inject 2 eV per proton, smoothed with a Gaussian function, centred at $z = 3.5$ and with width $\sigma(z) = 0.5$, to mimic non-equilibrium and radiative-transfer effects (e.g. Abel & Haehnelt 1999). As demonstrated in Wiersma et al. (2009b), with this extra heat input the predicted temperature at the mean density is in good agreement with the measurements of Schaye et al. (2000).

(ii) *Quiescent star formation and feedback.* Cosmological simulations cannot, at present, resolve the multiphase structure of star-forming gas that arises from the combination of gas cooling and heating due to massive stars and supernovae (SNe). These processes make star-forming gas resistant to compression, since the energy output from SNe rapidly heats the interstellar medium (ISM), increasing its pressure. We mimic this by *imposing* an effective equation of state $P = \kappa\rho^{\gamma_{EOS}}$ on gas with density exceeding $n_H = 0.1\,\mathrm{cm}^{-3}$, the critical density for the onset of the thermogravitational instability (Schaye 2004). Schaye & Dalla Vecchia (2008) demonstrate that the exponent $\gamma_{EOS} = 4/3$ makes both the Jeans mass and the ratio of the SPH kernel and the Jeans length independent of the gas density, thereby preventing spurious fragmentation due to a lack of numerical resolution. We adopt this value here. The equation of state is normalized such that $P/k = 10^3\,\mathrm{cm}^{-3}$ K for atomic gas at the density threshold.

Observations of galaxies reveal a tight relationship between gas column density and SFR, the Kennicutt–Schmidt law (e.g. Kennicutt 1998). Schaye & Dalla Vecchia (2008) show that since gas column density in a self-gravitating system is directly related to pressure, it is possible to implement a local Kennicutt–Schmidt law as a pressure law, independent of the chosen value of γ_{EOS} and even for very low numerical resolution. The star formation algorithm therefore converts gas particles with densities above the threshold $n_H = 0.1\,\mathrm{cm}^{-3}$ into stars at a rate that depends on the gas pressure,

$$\dot{m}_\star = A(1\,\mathrm{M}_\odot\,\mathrm{pc}^{-2})^{-n} m_g \left(\frac{\gamma}{G} f_g P\right)^{(n-1)/2}, \tag{1}$$

where m_g is the mass of the gas particle for which we are computing $\dot{m}_\star$, $\gamma = 5/3$ for a monatomic gas (and should not be confused with γ_{EOS}), f_g is the mass fraction in gas (assumed here to be unity) and P is the gas pressure. The normalization, A, and slope, n, follow from the observed Kennicutt (1998) law,

$$\dot{\Sigma}_* = 1.5 \times 10^{-4}\,\mathrm{M}_\odot\,\mathrm{yr}^{-1}\,\mathrm{kpc}^{-2} \left(\frac{\Sigma_g}{1\,\mathrm{M}_\odot\,\mathrm{pc}^{-2}}\right)^{1.4}, \tag{2}$$

where we have divided Kennicutt's normalization by a factor of 1.65 to account for the fact that we assume the IMF takes the form proposed by Chabrier (2003) rather than the commonly adopted form due to Salpeter (1955). The conversion of gas particles into stars, based upon their associated SFR, is stochastic, and since the spawning of multiple star particles from a gas particle can lead to an artificial reduction in the efficiency of feedback prescriptions (for a discussion, see Schaye & Dalla Vecchia 2008), we convert entire gas particles into stars. Conveniently, this practice also conserves the particle number throughout our simulations.

(iii) *Kinetic feedback.* Feedback from star formation in the disc is represented in part by the imposed relation $P = \kappa\rho^{\gamma_{EOS}}$. However, observed starburst galaxies exhibit signs of galaxy-wide winds (Heckman, Armus & Miley 1990; Martin 1999; Pettini et al. 2002; Adelberger et al. 2003; Shapley et al. 2003; Wilman et al. 2005; Swinbank et al. 2007) that are thought to be responsible for enriching the IGM with metals. We model the generation of winds as follows. After a delay of 3×10^7 yr, corresponding to the maximum lifetime of stars that undergo core collapse SNe, newly formed star particles impart a randomly directed 600 km s^{-1} kick to, on average, $\eta \equiv \dot{m}_{wind}/\dot{m}_\star = 4$ of its neighbours. This mass loading value was chosen to scale the global SFRD, $\dot{\rho}_\star(z)$, such that it reasonably matches observational data.

Assuming that each star with initial mass in the range 6–100 M$_\odot$ injects a kinetic energy of 10^{51} erg, these parameter values imply that the total wind energy accounts for 80 per cent of the available kinetic energy for our chosen IMF. Note that contrary to the widely used kinetic feedback recipe of Springel & Hernquist (2003a) wind particles are injected *locally* and are *not* temporarily decoupled hydrodynamically. As discussed in Dalla Vecchia & Schaye (2008), these changes have a large effect on the structure of galaxy discs and outflows.

(iv) *Chemodynamics.* Each star particle represents a single stellar population, with metal abundances inherited from its parent gas particle. Given our assumed IMF (Chabrier 2003, stellar mass range 0.1–100 M$_\odot$) and stellar evolution tracks dependent on metal abundance (Portinari, Chiosi & Bressan 1998; Marigo 2001; Thielemann et al. 2003), we follow the delayed release of 11 elements (hydrogen, helium, carbon, nitrogen, oxygen, neon, magnesium, silicon, sulphur, calcium and iron) by Type Ia and II SNe, and AGB stars. Star particles distribute the synthesized elements and lost mass during each time-step to neighbouring gas particles using the SPH interpolation scheme. Throughout this study, we use the solar abundances of CLOUDY; solar metallicity is therefore taken to be $Z_\odot = 0.0127$.

2.2 Initial conditions and run details

For brevity, we present here only an outline of the generation of the initial conditions; a detailed technical description is deferred to the Appendix. There, we also discuss how we address the existence of an artificial boundary between the high-resolution region and the external, low-resolution pressureless region, and how we combine results from the five individual GIMIC regions to construct estimates of properties for the entire Millennium Simulation volume.

[1] We used their equation (3), rather than (4), and CLOUDY version 05.07.

We follow the evolution of a wide range of environments by simulating five spherical regions whose overdensities at $z = 1.5$ are $(-2, -1, 0, +1, +2)$ times the rms deviation, σ, from the mean on some spatial scale. Ideally, the simulations would resolve the Jeans mass at our imposed star formation threshold density, but this remains beyond the scope of our available computational resources. Therefore, for our highest resolution simulation, we aimed to resolve the Jeans mass in the post-reionization IGM, which was shown by Theuns et al. (1998) to be a prerequisite for attaining converged properties of the Lyα forest. This requires a gas particle mass of $\sim 10^6 \, h^{-1} \, M_\odot$, limiting the size of the regions to a radius of $18 \, h^{-1}$ Mpc. To ensure, we also followed a rich cluster; we imposed a precondition that the $+2\sigma$ region be centred on a massive dark matter halo at $z = 0$ and increased its radius to $25 \, h^{-1}$ Mpc.

In order to preserve the correct large-scale forces, the volume external to these five regions was simulated at lower resolution and without baryonic physics. This methodology is shown schematically in Fig. 1. The multiresolution particle distribution generated for each simulation is explained in the Appendix; the particle displacement fields were calculated using the techniques described in Power et al. (2003) and Navarro et al. (2004).

We created two realizations of the $18 \, h^{-1}$ Mpc spheres: high-resolution (marginally resolving the IGM Jeans mass) and intermediate-resolution (with a factor of 8 fewer particles); we reserve the term 'low-resolution' for the original Millennium Simulation realization. The $+2\sigma$ region was generated at intermediate resolution only because of its far greater computational demands. We have hence carried out nine simulations in total, with the characteristics listed in Table 1. All intermediate-resolution realizations have been run to $z = 0$. The least cpu-demanding region (-2σ) was also run to $z = 0$ at high resolution, whilst the three regions, $(-1, 0, +1)\sigma$, were run to $z = 2$ at high resolution. Simulations of the same region at different resolution are needed in order to check numerical convergence. In what follows, we perform such tests at $z \geq 2$ for the $(-1, 0, +1)\sigma$ and all the way to $z = 0$ for the -2σ region.

In all simulations, the gravitational forces of the baryonic and high-resolution dark matter particles were softened over the same length-scale. The softening length is initially fixed in comoving space, but becomes fixed in physical space at a predefined redshift, i.e. $\epsilon_{\rm com}(a)' = \min(\epsilon_{\rm com}, \epsilon_{\rm phys}^{\rm max}/a)$. The softenings were chosen such that at $z = 3$, they are fixed at $\epsilon_{\rm phys}^{\rm max} = (1.0, 0.5) \, h^{-1}$ kpc for the intermediate- and high-resolution runs, respectively. The gas particle masses are 1.16×10^7 and $1.45 \times 10^6 \, h^{-1} \, M_\odot$ at intermediate and high resolution, respectively. These are much smaller than the limit required to avoid artificial two-body heating effects (Steinmetz & White 1997) that dominate over radiative cooling.

The dark particle masses are a factor of $(\Omega_{\rm m} - \Omega_{\rm b})/\Omega_{\rm b} = 4.56$ larger than the gas particle masses.

2.3 Halo and galaxy identification

We identify haloes by applying the friends-of-friends (FoF) algorithm to dark matter particles, using the standard value of the linking length in units of the mean interparticle separation ($b = 0.2$; Davis et al. 1985). In order to identify the baryonic content of dark matter haloes, our group finding algorithm locates the nearest dark matter particle to each baryonic (i.e. gas or star) particle. If this dark matter particle has been grouped by FoF, the corresponding baryonic particle is also associated with the FoF group. The assignment of baryons to dark matter haloes is therefore unambiguous. This method occasionally introduces artefacts, mostly for groups with few particles. For example, we find that the baryon fraction of low-mass haloes varies widely, from haloes with almost no baryons to haloes with more than the cosmic mean value, because our scheme allows the baryon distribution around a halo to be more extended than the dark matter. Since these problems are mostly restricted to small haloes close to the resolution limit, we have not attempted to overcome them.

The FoF algorithm identifies isodensity contours of $\delta \simeq 3/(2\pi b^3) \simeq 60$ (Lacey & Cole 1994), but noise in the particle distribution leads to the detection of many small, artificial objects clustered around the true halo. Since these artefacts tend to be transient, we perform an unbinding calculation using the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2008), and omit from our analyses any FoF haloes that do not have at least one self-bound substructure. This additional procedure successfully removes artificial haloes from our sample. The version of SUBFIND used here is modified from the standard implementation, such that baryonic particles are also considered when identifying self-bound substructures (for more details, see Dolag et al. 2008). This provides an unambiguous definition of a *galaxy* within the simulations as the set of star particles bound to individual substructures. Thus, an individual halo may host more than one galaxy.

When quoting halo masses, we refer to the total mass of the FoF halo in all components [i.e. gas, stars and cold dark matter (CDM)]. In some cases, however, it is appropriate to modify this definition slightly. When comparing results with simulations that follow only dark matter (e.g. Section 3), we consider only the CDM component of a FoF halo and boost its mass by a factor of $\Omega_{\rm m}/(\Omega_{\rm m} - \Omega_{\rm b})$ to account for the baryonic component. Additionally, when measuring baryon fractions it is common to specify fractions within the well-defined volume of a spherical overdensity (SO) group (e.g. Lacey & Cole 1994); in this case, we quote the mass within a sphere, centred

Table 1. Parameters for the five GIMIC regions.

Region	x $(h^{-1}$ Mpc)	y $(h^{-1}$ Mpc)	z $(h^{-1}$ Mpc)	Comoving radius $(h^{-1}$ Mpc)	N (int. res)	N (high res)	$z_{\rm final}$ (int. res, high res)
-2σ	153.17	347.90	424.81	18	2.23×10^7	1.78×10^8	[0,0]
-1σ	387.91	316.48	113.46	18	2.80×10^7	2.24×10^8	[0,2]
0σ	271.94	108.29	107.45	18	3.44×10^7	2.75×10^8	[0,2]
$+1\sigma$	179.51	379.22	196.64	18	4.30×10^7	3.44×10^8	[0,2]
$+2\sigma$	233.10	139.30	387.38	25	1.24×10^8	N/A	[0,-]

Note. Columns 2–5 give the location (in Millennium Simulation coordinates) and the nominal comoving radius of the regions at $z = 1.5$. The following two columns show the number of gas (or, equivalently, dark matter) particles within the zoomed region of the simulation, whilst the final column gives the redshift at which each simulation was terminated. Note that the $+2\sigma$ region was not run at high resolution.

on the local minimum of the gravitational potential, whose radius encloses a mean density (in all components) of $200\rho_c$, where ρ_c is the critical density for a flat Universe. We also quote the circular velocity, v_c, of haloes at this radius, r_{200}.

3 EVOLUTION OF THE HALO AND STELLAR MASS FUNCTIONS

3.1 The halo mass function in different environments

As pointed out by Frenk et al. (1988), and discussed more recently by Cole (1997) and Sheth & Tormen (2002), a fundamental difference between regions of varying overdensity is the rate at which their dark matter halo populations evolve. This is reflected in the clustering bias of haloes as a function of mass (Mo & White 1996). As we will see below, many environmental effects revealed by our simulations, e.g. variations in the SFRD, can be traced back to differences in the halo mass function in different environments. We therefore preface the discussion of star formation with a brief overview of the evolution of the halo mass function in each region, including resolution tests.

In Fig. 2, we show the differential number density of haloes identified with the FoF algorithm, $dn(M)/dM$, multiplied by M^2 in order to reduce the dynamic range of the plot, and thus more clearly highlight the differences between regions. As noted in Section 2.3, the halo mass in this plot is taken to be $M_{FOF} = M_{FOF}^{DM} \Omega_m / (\Omega_m - \Omega_b)$, so as to enable the most direct comparison possible with the results of dark matter-only simulations. The regions simulated at intermediate resolution are shown as coloured curves, while the high-resolution realization of the -2σ region is shown as a grey curve. In the region of overlap, the agreement between the high- and intermediate-resolution realizations is excellent, suggesting that the mass function in the latter has converged down to $1 \times 10^{10} \, h^{-1} \, M_\odot$ (corresponding to ~ 200 particles). Extrapolating to the high-resolution simulation, we estimate that the halo mass

function is reliable down to $\sim 10^9 \, h^{-1} \, M_\odot$. For subsequent analyses of the baryonic properties of haloes, we therefore consider only FoF haloes with at least 200 particles.

The GIMIC mass functions are also compared with the global mass function fit of Reed et al. (2007), which extends to both higher and lower mass scales since it was derived from an ensemble of dark matter simulations covering a wide dynamic range in mass (including the Millennium Simulation; see also Jenkins et al. 2001). As expected, the mass functions of the underdense and overdense GIMIC regions bracket the global function, which is reasonably well matched by the mass functions of the 0σ and $+1\sigma$ regions. We have checked that the halo mass functions of the regions at $z = 1.5$ (when they were selected) are consistent with those of all possible candidate spheres in the Millennium Simulation. As an example, the inset shows the multiplicity function of the intermediate-resolution 0σ region, overplotted on the locus that encompasses the 16th and 84th percentiles of the multiplicity functions of 100 spheres, drawn from the Millennium Simulation, that qualified as candidates for the 0σ region (for details, see the Appendix).

Although the different halo mass functions appear as scaled versions of one another at high redshift, a large difference develops over time, as the amplitude and shape of the function evolves. By the present day, there are approximately three times more haloes of intermediate mass in the $+2\sigma$ region compared to the -2σ; the difference increases with mass, with large-mass haloes found only in the $+2\sigma$ region. To determine the efficiency of halo formation per unit mass, we show in the right-hand panel of Fig. 2 the halo mass function now normalized by the total enclosed mass, rather than by the total enclosed volume. Normalized this way, the differences between the GIMIC regions are smaller but there is still a large systematic variation from region to region at all masses. Haloes form more efficiently in the high-density regions because they are the most dynamically advanced and the most massive haloes form only in the most overdense regions. These results agree with previous numerical and analytical studies of dark matter
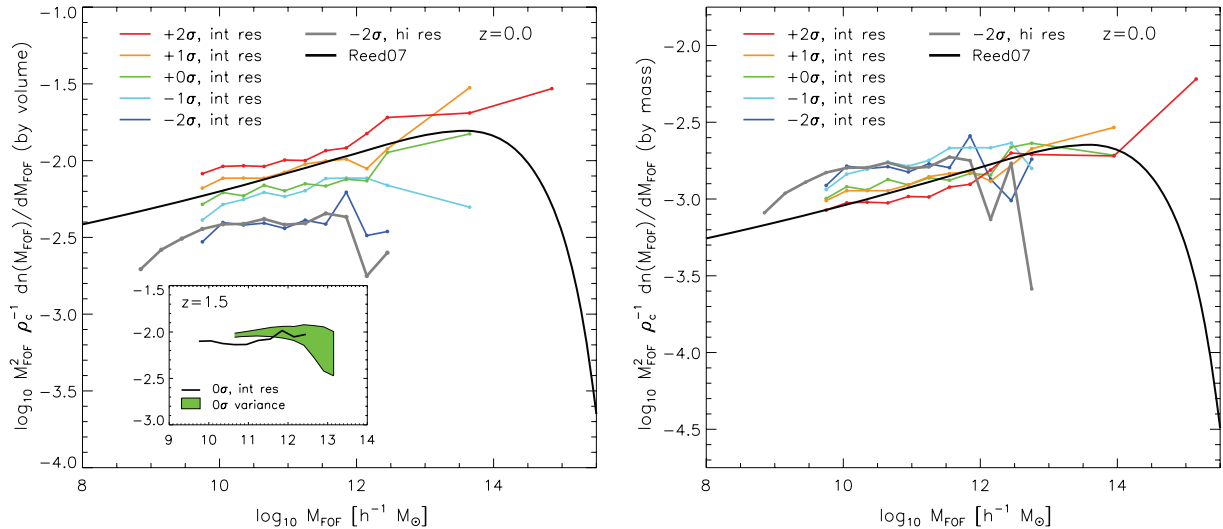
Figure 2. Left-hand panel: differential number density of haloes as a function of mass at $z = 0$, normalized by volume, and multiplied by M^2 in order to reduce the dynamic range of the plot. Coloured lines represent the intermediate-resolution GIMIC regions and the grey line represents the -2σ region at high resolution. The global halo mass function of Reed et al. (2007, black) is bracketed by the extreme GIMIC regions. Inset is the halo mass function at $z = 1.5$, plotted for the 0σ region. The green shaded area delineates the 16th and 84th percentiles of the mass functions of 100 candidate 0σ spheres drawn from the Millennium Simulation. Right-hand panel: as left-hand panel, but the differential number density is normalized by the total enclosed mass of each region, rather than by its volume. The difference between the regions is smaller when the regions are normalized by mass, but there is still a systematic variation from region to region at all masses. The most massive haloes form only in the most overdense, and hence most dynamically advanced, regions.

halo evolution (e.g. Frenk et al. 1988; Mo & White 1996; Sheth & Tormen 2002).

3.2 The galaxy stellar mass function

The galaxy stellar mass function as a function of redshift provides a useful check of the realism of our simulations. This function is shown in Fig. 3 and compared with observational data at $z = 2$ and 0. The stellar mass functions in each of the five GIMIC regions at intermediate resolution are shown as coloured curves. From these, we obtain an estimate of the stellar mass function for the entire Millennium Simulation volume (black curve), by a suitably weighted average; details of this procedure are given in the Appendix.

As in Fig. 2, we also plot results from the high-resolution realization of the -2σ region (grey curve) as a numerical test. This shows that at $z = 2$ the intermediate-resolution stellar mass function has approximately converged for $M_\star \geq 10^9\,h^{-1}\,M_\odot$. Below this value, there are roughly twice as many galaxies in the intermediate-resolution simulation. This excess almost certainly reflects a reduction in the efficiency of SNe feedback in poorly resolved galaxies. At $z = 0$, the convergence properties are similar: below $M_\star \sim 10^9\,h^{-1}\,M_\odot$, the intermediate-resolution simulation overpredicts the number of galaxies by about a factor of 2, and above this mass it underpredicts it by a similar factor.

The weighted average stellar mass function closely tracks that of the 0σ region over the mass range for which haloes are present in this region. At the high-mass end, only the overdense regions contribute to the stellar mass function. This highlights the need to follow a wide range of large-scale environments when attempting to compare simulations with large galaxy redshift surveys. For comparison, we overplot the stellar mass functions derived from the FORS and Great Observatories Origins Deep Survey (GOODS) deep fields at $z = 2$ (Drory et al. 2005) and from SDSS data at $z = 0$ (Li & White 2009). At $z = 2$, the data span a factor of 100 in stellar mass. Over this limited range, the model stellar mass function has a similar shape to the data but the amplitude of the weighted model function

is approximately 0.5 dex higher, which is a slightly less than the difference between this model function and that of the -2σ region.

At $z = 0$, the data span nearly four orders of magnitude in stellar mass. The simulations produce too many small galaxies, with $M_\star < 10^9\,M_\odot$, and too few in the mass range 10^9–$10^{10}\,M_\odot$, but inspection of the cumulative function shows that they produce about the right number density at $M_\star > 10^9\,M_\odot$ and then at $M_\star > 2.5 \times 10^{10}\,M_\odot$. The faint end slope of the weighted average stellar mass function evolves subtly with time; measured over the range $10^{7.5} < M_\star < 10^{9.5}\,h^{-1}\,M_\odot$, the function $dn/d\log_{10} M_\star$ varies as $M_\star^{-1.07}$ at $z = 2$ and $M_\star^{-0.96}$ at $z = 0$.

At larger masses, the two overdense GIMIC regions produce a population of massive galaxies that is not seen in the data. The overproduction of faint galaxies and the 'dip' at intermediate masses result from the simple wind model that we have adopted. The excess at large masses in the overdense regions results from our neglect of heating processes that can quench cooling flows in clusters. In semi-analytic models of galaxy formation (Bower et al. 2006; Croton et al. 2006; Somerville et al. 2008) and cosmological hydrodynamical simulations (Di Matteo et al. 2008; Sijacki et al. 2008; Booth & Schaye 2009), this is achieved by invoking feedback from AGN, which is not modelled in our simulations.

The sensitivity of the stellar mass function to variations in the physical assumptions and parameters in the simulations will be explored in the OWLS project (Haas et al., in preparation). For our purposes, it is sufficient to note that there is reasonably good agreement between the simulations and the data at high redshift while, at low redshift, the simulations produce roughly the right number of galaxies, albeit with an incorrect distribution of stellar masses.

4 THE EVOLUTION OF COSMIC STAR FORMATION

The cosmic SFRD is a key tracer of structure formation in the universe. After more than a decade of observational studies, it appears to be fairly well constrained at least out to intermediate redshifts
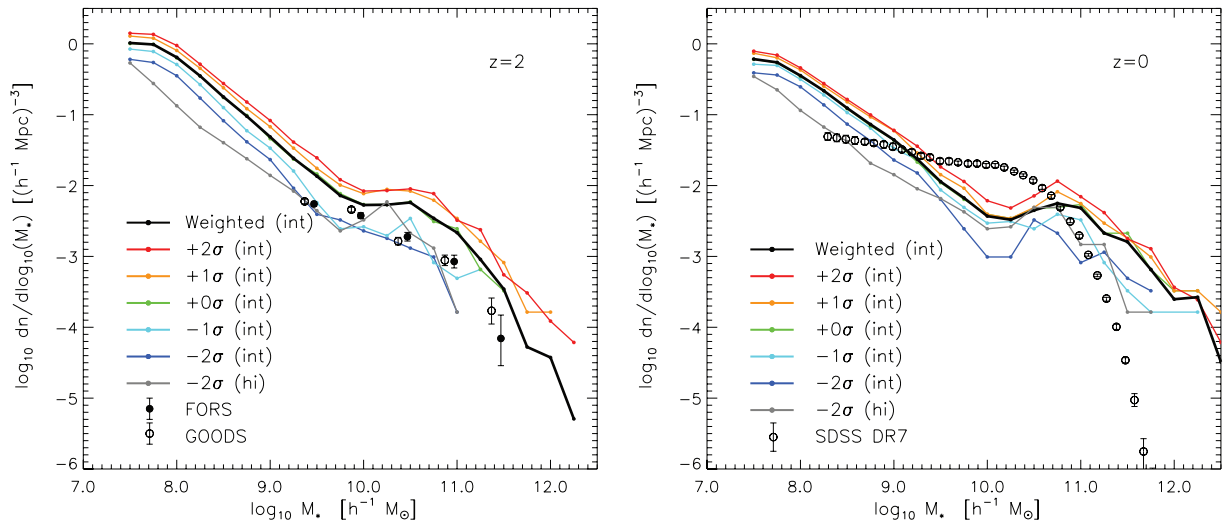
Figure 3. The stellar mass function of galaxies at $z = 2$ (left-hand panel) and $z = 0$ (right-hand panel). Results are shown for all five intermediate-resolution simulations (coloured curves) and their weighted average (black curve), which represents an estimate of the function over the entire Millennium Simulation volume. The stellar mass function of the -2σ region at high resolution is also shown (grey curve) to illustrate the degree of convergence. Observational data from the FORS and GOODS deep fields (Drory et al. 2005) are overplotted at $z = 2$, and from SDSS DR7 (Li & White 2009) at $z = 0$. The simulations are in rough agreement with the data at $z = 2$; at $z = 0$, the overall number of galaxies more massive than $\sim 10^9\,M_\odot$ is correctly reproduced by the simulations, but the number density of intermediate mass galaxies is too low, whilst that of the most massive galaxies is too high.

($z \lesssim 3$), given assumptions about the IMF and reddening (Hopkins 2007). Taking into account the mass fraction recycled during stellar evolution, an integral over $\dot{\rho}_\star(z)$ gives the present-day cosmic stellar mass density (e.g. Cole et al. 2001; Rudnick et al. 2003; Eke et al. 2005; Li & White 2009). These are all quantities that are calculated in our simulations.

In this section, we investigate the star formation history in our simulations and how it varies from one GIMIC region to another. We discuss the numerical convergence of our results and the nature of the dominant contributors to the star formation activity at high redshift. The averaging technique described in the Appendix is applied to construct an estimate of the global SFRD that can be compared both to observations and to the results of the semi-analytic calculations applied to the Millennium Simulation.

4.1 Large-scale environmental modulation

For each region, we consider all gas particles within the sampling sphere discussed in the Appendix, and use the volume of the sphere to compute the SFR per unit volume, $\dot{\rho}_\star(z)$, shown in Fig. 4. Comparison of the intermediate- and high-resolution simulations shows that our results have not converged at $z \gtrsim 6$; the high-resolution realizations exhibit significantly higher overall SFRs than their intermediate-resolution counterparts. At later times, however, convergence of $\dot{\rho}_\star(z)$ is achieved to within ~ 50 per cent, and inspection of the curves for the -2σ region realizations demonstrates that this conclusion extends to the present epoch.

The evolution of $\dot{\rho}_\star(z)$ may depend on resolution because the simulation fails to form small-mass haloes and/or because the SFR in the haloes that do form depends on resolution. The particle mass of the high-resolution realizations was chosen so that the Jeans mass in the photoionized IGM is marginally resolved. Hence, these simulations should be able to follow the formation of all haloes in which the gas cools by atomic line cooling after reionization. We show later that the SFR for resolved haloes is very similar at the two resolutions. We are therefore reasonably confident that our estimate of $\dot{\rho}_\star(z)$ is close to numerically converged in our high-

resolution simulations for $z < z_{\rm reion} = 9$, but clearly this is not the case in the intermediate-resolution realizations for $z \gtrsim 6$. The actual value of the SFR is, of course, sensitive to our subgrid modelling, in particular to the implementation of galactic winds. Note that the lack of high-redshift agreement in the value of $\dot{\rho}_\star$ between the two resolutions has only a small effect on the total amount of stars formed and on the distribution of stellar ages at low-redshift, because the duration of the high-redshift period before the values converge is small compared to the age of today's universe.

The main features of the SFRD in our simulations are:

(i) $\dot{\rho}_\star$ increases with decreasing redshift, peaks at $z \sim 2-3$, and then drops rapidly by a factor of ~ 6 to $z = 0$;

(ii) the shape of $\dot{\rho}_\star(z)$ is similar for all regions and

(iii) the amplitude of $\dot{\rho}_\star(z)$ varies by *an order of magnitude* between the most extreme regions, at all epochs.

The amplitude increases monotonically from the -2σ to the $+2\sigma$ region at all epochs, but the difference between the $+1\sigma$ and the $+2\sigma$ regions is far smaller than the difference between the -2σ and the -1σ regions. The large variations in the amplitude of $\dot{\rho}_\star(z)$ between the different regions is not simply caused by the greater mass contained in the higher density volumes. As shown in the right-hand panel of Fig. 4, higher σ regions also have a higher *specific* SFR (sSFR), $\dot{M}_{\rm star}/M_{\rm total}$ (where $M_{\rm total}$ is the total mass enclosed in each sphere). The SFR *per unit mass* still varies by ~ 0.5 dex between the regions. We show later in Section 5.2 that this is due to the different halo mass functions in regions of different overdensity, and the variation in star formation efficiency with halo mass.

A notable feature of the SFR in the high-resolution simulations (dashed lines) is the kink in $\dot{\rho}_\star$ at $z \simeq 9$, the epoch when the IGM is reionized and heated. The rapid increase in $\dot{\rho}_\star$ prior to this epoch is slowed by reionization in all regions. The onset of reionization has two main effects on the evolution of gas: (i) the photoheating boosts and maintains the gas temperature at a near-uniform level of $\sim 10^4$ K and (ii) the photoionization suppresses cooling due to line emission from hydrogen and helium (Efstathiou 1992) as well as heavy elements (Wiersma et al. 2009a). Small
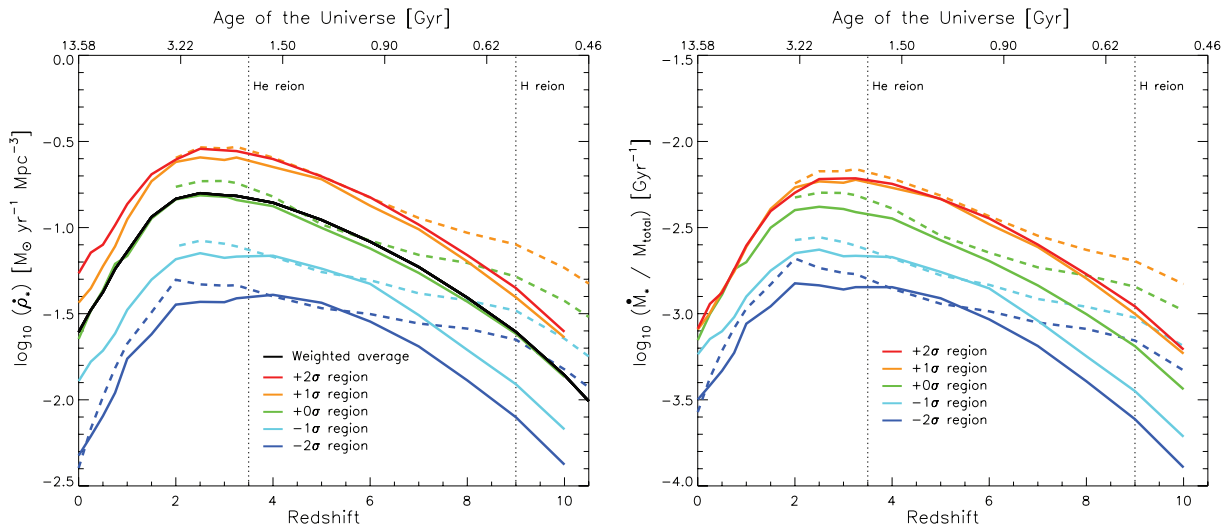
Figure 4. Left-hand panel: SFRD, $\dot{\rho}_\star(z)$, as a function of redshift in regions of different overdensity (colours), at intermediate (solid lines) and high resolution (dashed lines). We have not run the $+2\sigma$ simulation at high resolution. The vertical dotted lines denote the epochs of reionization of hydrogen (H I) and helium (He II), respectively. The variation in $\dot{\rho}_\star$ amongst the large-scale environments is pronounced; the most extreme regions differ by an order of magnitude at all epochs. Right-hand panel: the evolution of the total SFR *per unit mass*, in regions of different large-scale overdensity. The persistent offsets demonstrate that the increase of $\dot{\rho}_\star$ with large-scale overdensity does not arise purely because overdense regions enclose a greater mass.
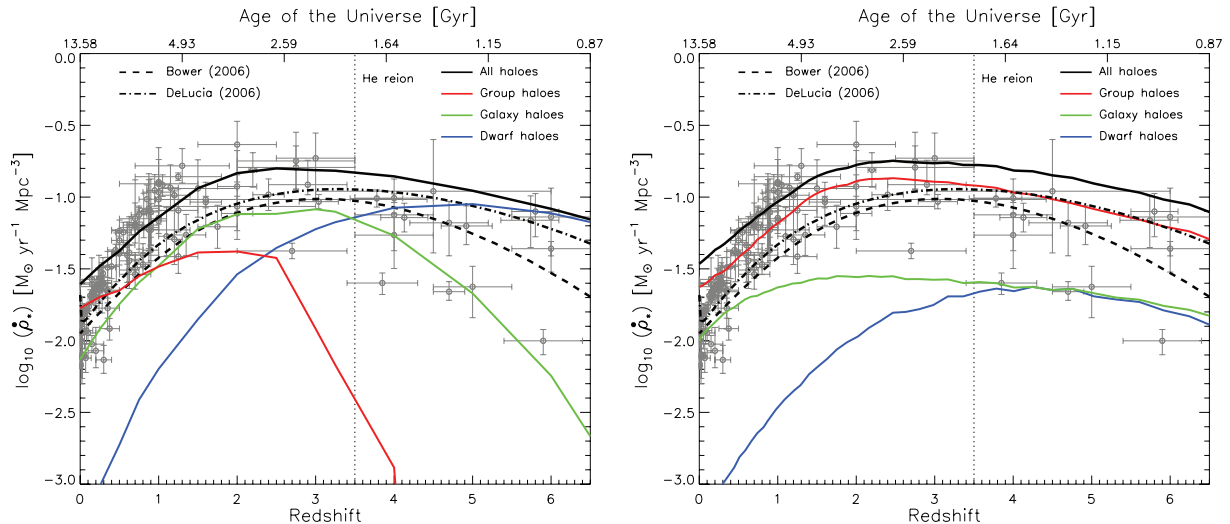
Figure 5. Left-hand panel: SFRD, $\dot{\rho}_\star(z)$, in the $L = 500\,h^{-1}$ Mpc Millennium Simulation volume (black solid line), derived from a weighted average of the intermediate-resolution GIMIC simulations. Coloured lines show the instantaneous contribution to the total from dwarf galaxy haloes (blue), regular galaxy haloes (green) and group/cluster haloes (red). The evolution of $\dot{\rho}_\star$ predicted by the semi-analytic calculations of Bower et al. (2006, black dashed line) and De Lucia et al. (2006, black dot–dashed line) are shown for comparison. Symbols with error bars are observational estimates from the compilation of Hopkins (2007), converted to the Chabrier IMF. Right-hand panel: the 'archaeological' description of the SFRD. Haloes are assigned to bins according to their mass at $z = 0$ and the coloured curves show the value of $\dot{\rho}_\star(z)$ *due to their progenitors*.

haloes, whose virial temperature is below the temperature of the photoionized IGM, cannot hold on to their photoheated gas at $z < z_{\rm reion}$, and star formation within them ceases (Pawlik & Schaye 2009). Both Hoeft et al. (2006) and Okamoto, Gao & Theuns (2008) used numerical simulations to construct a simple analytic model for how the characteristic mass, $M_{\rm c}(z)$, below which haloes lose most of their baryons, depends on the shape of the temperature–density relation for the IGM. This characteristic mass is not well resolved in the intermediate-resolution simulations (solid lines) for the ionizing background adopted here (Haardt & Madau 2001), which explains the absence of the kink in that simulation. In accord with the star formation history shown in the left-hand panel of Fig. 5, this also indicates that star formation at this epoch is dominated by low-mass haloes ($T_{\rm vir} \lesssim 10^4$ K); otherwise reionization would have had little effect on the *global* SFRD.

The maximum in the SFRD at $z \sim 2$–3 occurs significantly later than in the simulations of Springel & Hernquist (2003b), which have comparable numerical resolution to ours (and only slightly different cosmological parameters; they adopted $\Omega_{\rm m} = 0.3$, $\Omega_\Lambda = 0.7$, $\Omega_{\rm b} = 0.04$, $\sigma_8 = 0.9$, $h = 0.7$). We believe that this is only partly due to our inclusion of metal line cooling that affects the SFR particularly at lower redshift, as discussed and illustrated in the analytical model of Hernquist & Springel (2003). The largest simulations of Oppenheimer & Davé (2006) cover volumes of $(32\,h^{-1}\,{\rm Mpc})^3$ and have 256^3 particles, which results in a factor of $\sim$2 poorer mass resolution than our intermediate-resolution simulations and a factor of $\sim$15 poorer resolution than the high-resolution GIMIC runs. At this level of resolution, our simulations would certainly be far from converged. The SFR in most of their models also peaks at a higher redshift than ours. The simulations of both Springel & Hernquist (2003b) and Oppenheimer & Davé (2006) use the multiphase ISM implementation of Springel & Hernquist (2003a) for gas in galaxies, as well as a different prescription for the generation of galactic winds. The details of the ISM and wind implementations are probably the main cause of the formation history differences between these simulations and those presented here.

4.2 The Millennium Simulation star formation rate density: comparison with observations and semi-analytic models

The SFRs in the individual GIMIC regions can be combined in the manner discussed in the Appendix in order to estimate the net SFR of the whole Millennium Simulation volume. In order to obtain results down to $z = 0$, we use the intermediate-resolution simulations; we consider only epochs $z < 6$ where our results are approximately converged.

Our estimate of $\dot{\rho}_\star(z)$ for the Millennium Simulation volume is shown in Fig. 5. It increases with time to reach a broad plateau at $\approx 0.25\,{\rm M}_\odot\,{\rm yr}^{-1}\,{\rm Mpc}^{-3}$ over the interval $3 \gtrsim z \gtrsim 1$ with a maximum around $z \sim 2$, followed by a rapid decline by a factor of $\sim$6 to $z = 0$. The behaviour of the global $\dot{\rho}_\star(z)$ closely follows that of the mean density (0σ) region.

The coloured curves in the left-hand panel decompose $\dot{\rho}_\star(z)$ into contributions from haloes binned according to their mass at that epoch; bins are chosen to correspond approximately to the masses of haloes that host dwarf galaxies ('dwarf haloes'; $M < 5 \times 10^{11}\,h^{-1}\,{\rm M}_\odot$), typical galaxies ('galaxy haloes'; $5 \times 10^{11}\,h^{-1}\,{\rm M}_\odot < M < 5 \times 10^{12}\,h^{-1}\,{\rm M}_\odot$) and groups and clusters of galaxies ('group haloes'; $M > 5 \times 10^{12}\,h^{-1}\,{\rm M}_\odot$). As might be expected within the CDM hierarchical assembly model, the peak of star formation within low-mass haloes precedes that of more massive haloes. For our particular classification, these peaks are broad, occurring at $z \sim 4$–5 for dwarf haloes, $z \sim 2$–3 for galaxy haloes and $z \sim 1$–2 for group haloes.

The simulation outputs record the *formation time* of each star particle and this allows us to relate present-day systems to the properties of the progenitors in which their stars formed. This 'archeological' description of the star formation history is computed for haloes of a given mass at $z = 0$ by summing the initial mass of their stellar particles (i.e. the mass at the time of formation, prior to mass loss due to stellar evolution), binned by their formation redshift, and is shown in the right-hand panel of Fig. 5. Note that the solid black curve, integrating over haloes of all masses, differs slightly from

that of the left-hand plot at epochs other than $z = 0$, because the population of baryonic particles within the sphere used to define the volume-limited 'sample' for these plots (see the Appendix for details of how this is selected) evolves slightly over time. This plot demonstrates that star formation is dominated at all redshifts by the progenitors of galaxies that today reside in the most massive haloes (i.e. groups and clusters). This is perfectly compatible with a hierarchical buildup, because the star formation in massive haloes became dominant only recently, as shown in the left-hand panel of the figure.

To compare the GIMIC simulations to observations, we use the compilation of SFRs given by Hopkins (2007). For consistency, we have adjusted the published values to account for the cosmological parameters and IMF adopted in the simulations; to convert from the IMF of Salpeter (1955) to that of Chabrier (2003), $\dot{\rho}_\star$ is scaled down by a factor of 1.65. As noted in Section 2.1, the mass loading in the galactic winds was chosen by requiring small test simulations to produce an approximate match to the amplitude of the observed $\dot{\rho}_\star(z)$. As a consequence, the broad agreement with the data is largely a check that the weighting scheme works well.

The evolution of the SFRD shows a broad plateau of nearly constant SFR during $4 \gtrsim z \gtrsim 1$, followed by a rapid drop to $z = 0$. The reduction in $\dot{\rho}_\star$ towards the present day is smaller in the simulation than in the data. This is mostly due to the high SFRs in massive haloes in the simulations (red curve in the left-hand panel of Fig. 5), which is also the cause for the overproduction of bright galaxies seen in Fig. 3. Recent semi-analytic models (e.g. Bower et al. 2006; Croton et al. 2006; De Lucia et al. 2006; Font et al. 2008) and cosmological hydrodynamical simulations (Di Matteo et al. 2008; Sijacki et al. 2008; Booth & Schaye 2009) include AGN radio-mode feedback to suppress the cooling of gas in such haloes and quench their star formation. As a result, $\dot{\rho}_\star$ in these models falls faster from $z = 1$ to 0, in better agreement with the data (dashed and dot–dashed lines in Fig. 5). Kobayashi, Springel & White (2007) obtained a qualitatively similar star formation history to that of GIMIC in hydrodynamic simulations with very strong feedback from hypernovae, but no AGN feedback. Like us, they found star formation at late times to be dominated by massive blue galaxies.

The discrepancy in $\dot{\rho}_\star$ occurs at late times when the SFR in all objects is relatively low, and therefore has little effect on the global stellar mass density, $\Omega_\star$ (in units of the critical density). The simulation produces a *current* stellar mass density [obtained by integrating $\dot{\rho}_\star(z)$ and subtracting the mass recycled by stellar evolution] of $\Omega_\star \sim 3.2 \times 10^{-3}$ at $z = 0$. Observational estimates of this quantity give $\Omega_\star = (1.37 \pm 0.4) \times 10^{-3}$ from the 2dFGRS (Eke et al. 2005) and $\Omega_\star = (1.44 \pm 0.4) \times 10^{-3}$ from the SDSS (Li & White 2009), assuming the Kennicutt and Chabrier IMFs, respectively. These estimates have a systematic uncertainty of approximately a factor of 2 arising from the choice of IMF.

4.3 Cosmic downsizing

Observations indicate that star formation in the most massive galaxies observed today essentially concluded at $z \sim 1$ (e.g. Drory et al. 2003, 2005; Pozzetti et al. 2003; Fontana et al. 2004; Kodama et al. 2004), and that their SFRs were higher in the past than they are at present (e.g. Cowie et al. 1996; Juneau et al. 2005). These findings have been cited as a challenge to the Λ cold dark matter (ΛCDM) cosmogony, since a naive view of galaxy formation within a hierarchical assembly framework would have star formation shifting from smaller to larger objects as the hierarchy builds up. This apparent conflict has spawned terms such as 'downsizing' and 'antihierarchi-

cal galaxy formation'. Tentative evidence that these observations are compatible with ΛCDM was first presented by Pearce et al. (2001), using hydrodynamic simulations that featured star formation but no feedback mechanisms. Bower et al. (2006) and De Lucia et al. (2006) convincingly demonstrated that the observations are entirely compatible with their ΛCDM semi-analytic models. We now explore whether our hydrodynamic simulations also exhibit downsizing.

One measure of downsizing, first highlighted by Bower et al. (2006, their fig. 7), is the ratio of the past average SFR of galaxies of a given mass, $M_\star(z)/t_H(z)$, where t_H is the Hubble time, to their SFR at the epoch of observation $\dot{M}_\star(M_\star, z)$. This is shown in a volume-averaged sense in Fig. 6, for which galaxies have been drawn from all five intermediate-resolution regions. Galaxies enter the regime above the horizontal dotted line that marks $\rho_\star/t_H = \dot{\rho}_\star$ when they have assembled most of their stellar mass and their present star formation is adding only a minor contribution.

In common with results from the semi-analytic model of Bower et al. (2006), galaxies in GIMIC exhibit a segregation by stellar mass, such that the least massive galaxies continue to grow until the present epoch, whilst more massive galaxies have essentially concluded their star formation at $z > 1$–2. Although this effect is weaker in our simulations than in the semi-analytic model and in the observational data, this is a significant result because our model does not include AGN radio-mode feedback. Rather than causing the downsizing effect directly, AGN feedback appears merely to exacerbate it, pushing massive galaxies further into the regime where the past average SFR dominates the current SFR. AGN feedback does seem to be required, however, to prevent the formation of the overly bright and blue galaxies that form in large haloes in our simulations but which are not observed in nature. The sort of downsizing discussed in this section results not from AGN feedback but is likely to be a consequence of both the 'natural' downsizing of dark matter halo progenitor masses that was discussed by Neistein,
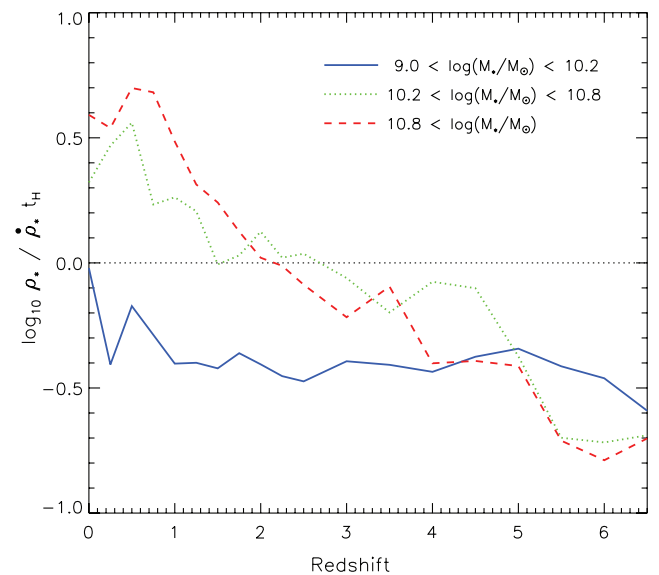
Figure 6. The ratio of past average SFR to the current SFR at the redshift of observation, for three mass bins. The binning, matched to fig. 7 of Bower et al. (2006), is done according to the stellar mass at the redshift of observation. The dotted line at $\rho_\star/t_H = \dot{\rho}_\star$ separates the regimes where (i) on-going star formation dominates the current stellar mass (below) and (ii) on-going star formation contributes a negligible fraction to the total stellar mass (above).

van den Bosch & Dekel (2006) and the complex interplay of gas cooling, star formation and feedback that develops as structure assembles hierarchically.

4.4 High-redshift star formation

The observational data for $z \gtrsim 6$ are uncertain and differ by up to a factor of 10 between different studies; compare, for example, the results of Bunker et al. (2006) and Bouwens et al. (2008). Our high-resolution simulations favour a modest drop of ~50 per cent in $\dot{\rho}_\star$ from its peak at $z \sim 2$ to $z = 6$, with most stars forming in faint galaxies, which have SFRs of $\lesssim 0.1\,\mathrm{M}_\odot\,\mathrm{yr}^{-1}$. Such SFRs are typically an order of magnitude below the sensitivity limits of most current surveys.

The sky coverage of each $r = 18\,h^{-1}\,\mathrm{Mpc}$ GIMIC region is $320\,\mathrm{arcmin}^2$ at $z = 6$. Whilst this is comparable to the sky coverage of the GOODS fields (Dickinson et al. 2003), it is much greater than the 5–20 arcmin^2 of the Hubble Ultra Deep Field (HUDF; Beckwith et al. 2006) in various passbands, which is an important source for the detection and characterization of galaxies at $z > 6$. Since the mean thickness of a sphere is $4r/3$, the comoving 'depth' of each region is $\sim 24\,h^{-1}\,\mathrm{Mpc}$. The expansion rate at $z = 6$ is $H_0 = 930\,h\,\mathrm{km\,s}^{-1}\,\mathrm{Mpc}^{-1}$, yielding a mean velocity range of $\Delta v \sim 3190\,\mathrm{km\,s}^{-1}$. This corresponds to $\Delta z \sim 0.07$ and therefore to a redshift 'window' of $z = 6 \pm 0.035$. Although this is much narrower than the $\Delta z = 1.5$ estimated by Bouwens et al. (2008) for their survey at $z = 7$, the actual volume probed with the pencil-beam geometry of the HUDF at these epochs is similar to that of each GIMIC region. Our simulations indicate that $\dot{\rho}_\star(z)$ varies by up to an order of magnitude between regions of this volume and so sample variance is a significant source of systematic uncertainty in measurements of $\dot{\rho}_\star$ at high redshift. This is further compounded by the dominant contribution to $\dot{\rho}_\star$ at these epochs of low mass, inefficiently star-forming galaxies that are below current detection limits.

The high-resolution simulations produce a slowly varying SFRD from $z = 6$ to 9. If the escape fraction of ionizing radiation from the small haloes that dominate $\dot{\rho}_\star$ at these epochs is high, say $\sim$(25–80) per cent as suggested by Wise & Cen (2009), then these galaxies will be the dominant contributors to the UV background at high redshift, and thus may have been the main sources of the radiation that led to reionization (Srbinovsky & Wyithe 2008). The faintness of each individual source would reconcile the apparent dearth of detected sources of UV photons with the inferred ionization state of the IGM at $z \sim 6$ (Bolton & Haehnelt 2007).

Although small galaxies dominate the SFRD, the wide range of overdensities represented in the GIMIC regions results in massive stellar systems ($M_\star \sim 10^{11}\,h^{-1}\,\mathrm{M}_\odot$) forming in the simulations as early as $z = 6$, as shown in Fig. 7. The dark matter haloes that host these massive galaxies are so strongly biased towards overdense regions that galaxies of comparable mass do not form in the void region (-2σ) until $z \sim 3$ (dark blue dashed line). These massive galaxies are embedded in large, nearly spherical coronas of hot ($T \sim 10^7$ K) gas of comoving radius $\sim 0.3\,h^{-1}\,\mathrm{Mpc}$, but the galaxies themselves are extremely compact (comoving radii of $\sim 3\,h^{-1}\,\mathrm{kpc}$). The stars in these galaxies have near-solar metallicity and are highly overabundant in the α elements produced in Type II SNe (such as oxygen), relative to those produced in Type Ia SNe (such as iron). They are relatively old (ages ~ 0.2 Gyr) at this redshift. Such galaxies are reminiscent of the $z \gtrsim 5$ massive and evolved galaxy candidates found in the GOODS fields by Wiklind et al. (2008).
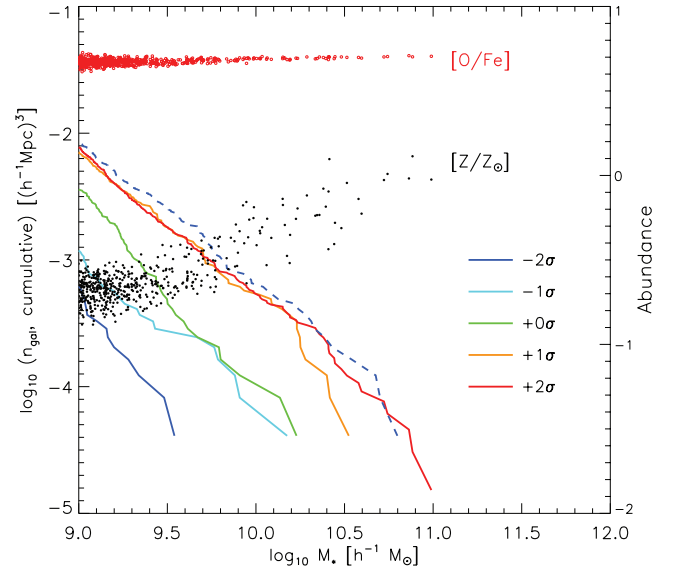
Figure 7. Cumulative number density of galaxies as a function of stellar mass at $z = 6$ for the five GIMIC regions (solid coloured lines) at intermediate resolution. The number density for the -2σ region is also shown at $z = 3$ (dark blue dashed line). The data plotted clearly illustrate the strong bias of massive stellar systems towards overdense regions. The most massive galaxies, drawn from the $+2\sigma$ region, have solar abundances (black points, right-hand y-axis) and are highly overabundant in α-process elements (red points, right-hand y-axis).

## 5 THE STAR FORMATION PROPERTIES OF HALOES

We now turn to an exploration of the physical basis for the environmental variation in the cosmic SFRD described in Section 4. We show that large-scale environment has no impact upon the evolution of the SFR in individual haloes. Instead, the strongly differing amplitudes of $\dot{\rho}_\star$ in the GIMIC regions can be traced back to their different mass functions.

We begin by considering the SFR in haloes of different circular velocity and how this depends on the 'subgrid' physics assumed in the simulations. We compare the star formation properties of the simulated haloes with analytic models for the regulation of star formation by gas cooling, and show that winds effectively regulate star formation up to a threshold value of the halo circular velocity. In common with conclusions drawn in Section 4, from comparison with a semi-analytic model that includes AGN feedback, we infer that efficient feedback in massive, hydrostatic haloes is necessary to regulate the growth of the most massive galaxies and to reconcile the cosmic SFRD at low redshift in the simulations with observations.

5.1 The specific star formation rate in haloes

To characterize the star formation efficiency of haloes, we define the *halo* sSFR $= \dot{M}_\star/M_{200}$, where M_{200} is the total mass (baryons plus dark matter) contained in a sphere of radius r_{200}. For clarity, we will retain the 'halo' prefix to avoid confusion with the more common definition of sSFR that applies strictly to *galaxies*, $\dot{M}_\star/M_\star$.

In common with previous sections, we start by discussing the degree of numerical convergence in the halo sSFR achieved by our simulations. The mean halo sSFR as a function of circular velocity, $v_c = (GM_{200}/r_{200})^{1/2}$, and M_{200} for haloes identified in the five intermediate-resolution regions are shown at various redshifts in Fig. 8. The small box and whisker elements denote one and two
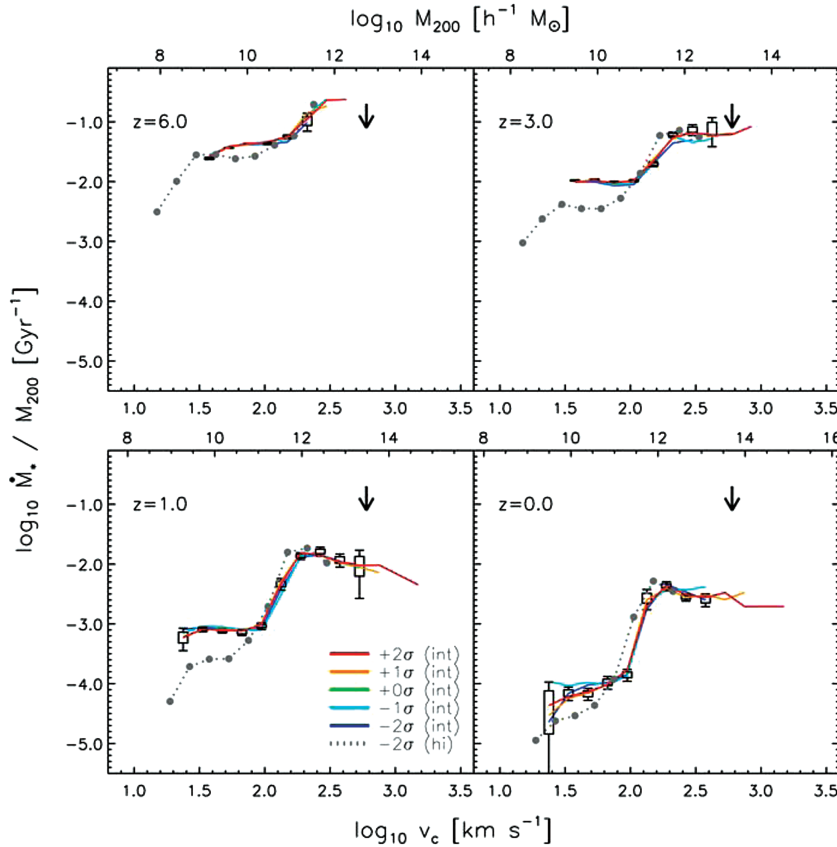
Figure 8. The 'halo sSFR', $\dot{M}_\star/M_{200}$, as a function of circular velocity (v_c, bottom axis) or halo mass (M_{200}, top axis) for various redshifts. Different colours are for each intermediate-resolution GIMIC region as indicated. Dotted grey lines are the results of the high-resolution -2σ simulations. Box and whisker elements indicate one and two standard errors on the mean for the 0σ region. The halo sSFR–v_c relations of the five GIMIC regions are essentially indistinguishable and show a break at $v_c \sim 250\,\mathrm{km\,s^{-1}}$, independent of redshift. The redshift evolution of the halo sSFR either side of this break reflects the differing mechanisms acting to regulate star formation at different halo mass scales. The arrow at $600\,\mathrm{km\,s^{-1}}$ indicates the launch speed of the winds.

standard errors on the mean. Since the relation for each region is virtually identical at each redshift, convergence can be assessed by comparison with the relation from the high-resolution realization of the -2σ region for all epochs (dark blue solid and *grey dotted* curves).

For $v_c \gtrsim 250\,\mathrm{km\,s^{-1}}$, the halo sSFR in the -2σ region is very similar for the two resolutions; for $100\,\mathrm{km\,s^{-1}} \lesssim v_c \lesssim 250\,\mathrm{km\,s^{-1}}$, the high-resolution simulation is only slightly higher than the intermediate-resolution simulation. However, below $100\,\mathrm{km\,s^{-1}}$, the intermediate-resolution simulation clearly overpredicts the halo sSFR (except at the highest redshift) and becomes approximately independent of v_c. Extrapolating from this behaviour to haloes resolved with the same number of particles in the high-resolution simulations, we conclude that these should give reliable results for $v_c \gtrsim 50\,\mathrm{km\,s^{-1}}$.

The halo sSFR–v_c relations for all GIMIC regions are virtually identical at every redshift. This similarity may be unexpected, given the large variation in the $\dot\rho_\star$ between the regions; we defer to Section 5.2 a detailed explanation of how these results are reconciled, in order to focus initially on the origin of the halo sSFR–v_c relation. The relation varies strongly with redshift, but drops rapidly below a characteristic circular velocity that is essentially independent of redshift. This drop occurs at $\sim 250\,\mathrm{km\,s^{-1}}$, in the regime where the high- and intermediate-resolution simulations agree well, and is caused by the action of galactic winds. This scale is significantly lower than the speed with which galactic winds are actually launched

($v_w = 600\,\mathrm{km\,s^{-1}}$, indicated by an arrow) because the *effective* velocity of the winds is suppressed by drag forces imparted by the dense gas of the ISM. Although the escape velocity of the halo, v_{esc}, can be considerably greater than v_c and therefore comparable to v_w, particularly in cases where the potential of the halo is steepened by the stellar mass at the centre (e.g. Smith et al. 2007), Dalla Vecchia & Schaye (2008) demonstrated that the efficacy of winds is determined by the ISM pressure: winds become ineffective when the kick velocity falls below a critical value that increases with the pressure, and hence also with v_c or v_{esc}. Previous studies such as Springel & Hernquist (2003b) did not find this because the wind particles were temporarily decoupled from the hydrodynamics and hence could escape freely without experiencing any drag from the ISM. The feature becomes more pronounced at later times when the halo sSFR drops faster with time for lower masses.

Haloes of virial mass $M_{200} \sim 10^9\,h^{-1}\,\mathrm{M_\odot}$ and circular velocity $\sim 50\,\mathrm{km\,s^{-1}}$ (that are well resolved in the -2σ high-resolution simulation) have low SFRs, $\dot{M}_\star \lesssim 0.05\,\mathrm{M_\odot\,yr^{-1}}$, at redshifts $z \sim 6$. Such objects are ubiquitous at that time and they dominate the total SFR. The halo sSFR drops rapidly at lower masses as haloes lose their baryons due to a combination of stellar feedback and heating of their gas by the UV background. The GADGET2 simulations of Springel & Hernquist (2003b) yield quantitatively similar results, as do the AMR simulations of idealized dwarfs described by Wise & Cen (2009). In the latter, blowout of baryons prevents continued star formation at masses below $10^8\,\mathrm{M_\odot}$. In our simulations, the halo

sSFR in these small haloes drops rapidly by more than two orders of magnitude towards $z = 0$.

Above a circular velocity of $\sim 250\,\mathrm{km\,s^{-1}}$, where haloes are less susceptible to galactic winds in our model, the halo sSFR is $\sim 10^{-1}\,\mathrm{Gyr^{-1}}$, almost independent of mass above $z \sim 3$. The level of this plateau drops by a factor of ~ 5 to $z = 1$. The simulations of Springel & Hernquist (2003b) yield an sSFR for a halo of mass $10^{12}\,h^{-1}\,\mathrm{M_\odot}$ of $10^{-1.7}\,\mathrm{Gyr^{-1}}$ at $z = 3$, a factor of 3 below our value. They also show the halo sSFR at $z = 1.5$, and the drop from $z = 3$ is a factor of 5 similar to the drop in our simulations over the interval $z = 1\text{–}3$. Whereas the halo sSFR is nearly constant with mass for massive objects in our simulations, it increases rapidly with halo mass in those of Springel & Hernquist (2003b).

The halo sSFR plummets by a factor ~ 50 for a $10^{12}\,h^{-1}\,\mathrm{M_\odot}$ halo between $z = 6$ and 0, and by an even greater factor at lower masses. The cause of this strong redshift evolution can be understood from simple physical arguments. White & Frenk (1991) present an analytic model in which the halo sSFR of a massive object of given circular velocity drops with redshift as star formation becomes limited by the gas cooling time rather than by its free-fall time. Defining the cooling rate, $\Lambda(u)$, via

$$\rho \dot{u} = -\Lambda n_{\mathrm{H}}^2, \qquad (3)$$

where $n_{\mathrm{H}} = X\rho/m_{\mathrm{H}}$ is the hydrogen number density for gas at density ρ, of which a mass fraction X is hydrogen, and u is the thermal energy per unit mass, this cooling time is then

$$\tau_{\mathrm{c}}(\rho) \equiv -\frac{u}{\dot{u}} = -\frac{u}{\Lambda n_{\mathrm{H}}^2/\rho} \propto -\frac{u}{\Lambda \rho}. \qquad (4)$$

The cooling function depends on temperature, T. The virial temperature of a halo is related to its circular velocity by

$$T_{\mathrm{vir}} = \frac{1}{3}\frac{\mu m_{\mathrm{p}}}{k_{\mathrm{B}}} v_{\mathrm{c}}^2 = 3.6 \times 10^5\,\mathrm{K}\left(\frac{\mu}{0.59}\right)\left(\frac{v_{\mathrm{c}}}{100\,\mathrm{km\,s^{-1}}}\right)^2, \qquad (5)$$

where μ is the mean molecular weight of the gas.

At high redshift, star formation is limited by the gas consumption time-scale since τ_{c} becomes very small. So long as inverse Compton cooling is unimportant, which is true at low redshifts, the cooling time (and hence the halo sSFR) of haloes with virial temperature $\sim 10^7\,\mathrm{K}$ (corresponding to $v_{\mathrm{c}} = 530\,\mathrm{km\,s^{-1}}$) scales $\propto \Lambda^{-1}\,\rho^{-1}$ which is approximately $\propto \rho^{-1} = (1+z)^{-3}$ if Λ does not evolve strongly.

The evolution of the normalization of the halo sSFR–v_{c} relation according to this model was explored analytically by Hernquist & Springel (2003). Starting from pseudo-isothermal initial gas profiles, their model yields sSFRs for haloes of $T_{\mathrm{vir}} \simeq 10^7\,\mathrm{K}$ that closely trace those in the simulations of Springel & Hernquist (2003b) at intermediate redshifts ($7 \gtrsim z \gtrsim 2$). This model is compared to the high-resolution GIMIC simulations in Fig. 9 where we have computed the cooling rate taking into account the effect of an evolving ionizing background as given by Haardt & Madau (2001) and described by Wiersma et al. (2009a). Note that this plot does not represent the *evolution* of the halo sSFR of a given object, but rather that of objects of a given circular velocity at *different* redshifts.

The effect of metallicity on the sSFR is pronounced in the model, but interestingly we find that there is, in fact, very little metallicity evolution for the hot gas in the simulations, $Z \sim 0.1\,\mathrm{Z_\odot}$ from $z = 9$ to the present. For this value of the metallicity, the model works very well for massive haloes, indicating that the cooling time controls star formation in these objects in the simulations, as expected from the White & Frenk (1991) calculation. (Note that the inclusion of an efficient feedback mechanism in massive haloes, such as energy injected by AGN, could alter the sSFR). The simulation results and
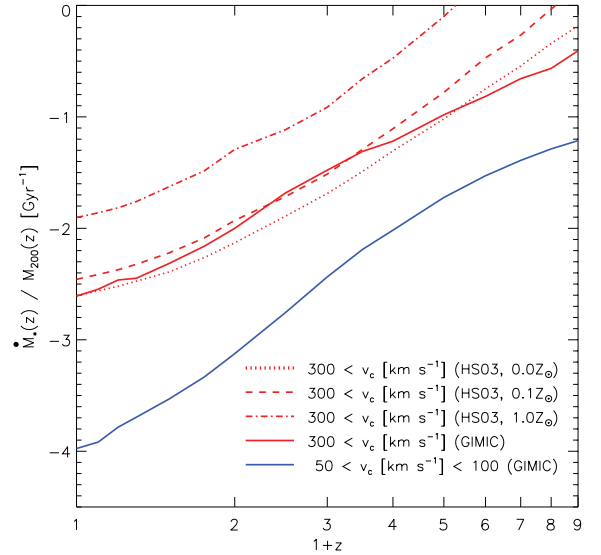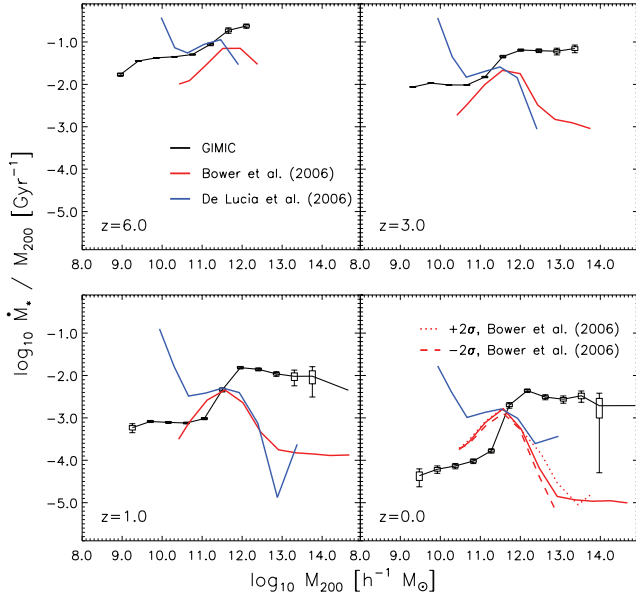

Figure 9. The mean halo sSFR, $\dot{M}_\star/M_{200}$, as a function of redshift for haloes of circular velocity $v_{\mathrm{c}} < 300\,\mathrm{km\,s^{-1}}$ (solid red line) and $v_{\mathrm{c}} = 50\,\mathrm{km\,s^{-1}}$ (solid blue line), drawn from the high-resolution -2σ GIMIC simulation. Results are compared with the analytic model of Hernquist & Springel (2003, labelled HS03) for the massive haloes, assuming three different hot gas metallicities (primordial, dotted red; $0.1\,\mathrm{Z_\odot}$, dashed red; $1.0\,\mathrm{Z_\odot}$ dot–dashed red). The metallicity of the hot gas in the simulations is nearly constant at $\sim 0.1\,\mathrm{Z_\odot}$, and the model predicts the result for the massive halo in the simulation very well. At lower redshifts, the drop in star formation follows that expected from the decrease in density. The analytic model is a poor description for low-mass haloes (blue), since it neglects the action of winds; the predicted halo sSFR is orders of magnitude higher than the simulated values, and so is not shown in the plot.

the Hernquist & Springel (2003) model diverge at high redshift because these authors chose not to include the gas consumption time-scale in this calculation; however, as White & Frenk (1991) show, taking account of this constraint would reduce the SFR and improve the match at high redshift.

Efficient feedback depresses the halo sSFR in low-mass objects far below the rate expected from the cooling time argument alone. The accelerated drop-off (relative to the trend for massive haloes) at $z \lesssim 3$ is due to feedback from galactic winds which depletes the baryon fraction in small haloes and increases the cooling time of the remaining gas, quenching star formation. Winds therefore introduce a second characteristic scale in galaxy formation in our model, that at which star formation is quenched due to *the ejection of baryons*. For our chosen wind speed of $v_{\mathrm{w}} = 600\,\mathrm{km\,s^{-1}}$, this scale is around $v_{\mathrm{c}} \sim 250\,\mathrm{km\,s^{-1}}$ and depends on the details of the wind implementation. This value is close to that advocated by White & Frenk (1991) as the scale separating massive haloes with long cooling times from smaller haloes within which gas cools efficiently.

5.1.1 Comparison with semi-analytic models

Although we defer a detailed comparison of the GIMIC and semi-analytic treatments of galaxy formation in the Millennium Simulation to a later paper, it is instructive to carry out a limited comparison here. We focus on the halo sSFR, $\dot{M}_\star/M_{200}$, and compare, in Fig. 10, the results obtained by Bower et al. (2006) and De Lucia et al. (2006) with those from GIMIC. Note that the GIMIC simulation extends to lower halo masses because of its higher mass resolution.

Figure 10. The halo sSFR, $\dot{M}_\star/M_{200}$, as a function of halo mass, M_{200} (as in Fig. 8) for all GIMIC regions combined, compared to the semi-analytic models of Bower et al. (2006, red lines) and De Lucia et al. (2006, blue lines). At $z = 0$, results of the semi-analytic model are shown, limited to the -2σ (red dotted curves) and $+2\sigma$ (dashed curves) regions, rather than the entire Millennium Simulation volume; no difference is apparent, as in the hydrodynamic simulations.

For halo masses $M_{200} \gtrsim 10^{11}\, h^{-1}\, \mathrm{M}_\odot$, the two semi-analytic models give very similar results, particularly for $z < 3$. At smaller masses, however, there are substantial differences between the semi-analytic models and also between these and the simulations. These differences are largely due to different treatments of feedback. The de Lucia et al. model has substantially higher sSFR than the Bower et al. model reflecting the weaker feedback assumed in the former; the feedback adopted in GIMIC is intermediate between the two semi-analytic prescriptions. For halo masses around $3 \times 10^{11}\, h^{-1}\, \mathrm{M}_\odot$, the GIMIC results agree well with the semi-analytic models. However, at higher masses GIMIC produces much higher SFRs. The reason for this discrepancy is simply that the semi-analytic models include feedback from the AGN radio mode. As explained in their respective papers (see also Benson et al. 2003; Croton et al. 2006), this 'radio-mode feedback' is required in order to prevent the formation of overly massive galaxies in large haloes.

Finally, dotted and dashed lines in the bottom-right panel of Fig. 10 compare the Bower et al. (2006) halo sSFRs in the $+2\sigma$ and -2σ GIMIC regions. As in the GIMIC simulations (see Fig. 8), the semi-analytic treatment finds little, if any, dependence of the halo sSFR on the large-scale environment.

5.2 The multiplicity function of star formation and its environmental dependence

We now return to an explanation of the physical reason for the strong dependence of the cosmic SFRD, $\dot{\rho}_\star$, on large-scale environment presented in Section 4.1. As shown in Fig. 4, $\dot{\rho}_\star$ varies by approximately an order of magnitude amongst the GIMIC regions, which may seem surprising given the similarity of the sSFR–v_c relation in each region.

The halo SFRD may be written in terms of the sSFR per halo as

$$\dot{\rho}_\star(z) = \int M_{200}\, \frac{\mathrm{d}N(M_{200}, z)}{\mathrm{d}\ln M_{200}}\, \frac{\dot{M}_\star}{M_{200}}\, \mathrm{d}\ln M_{200}$$

$$\equiv \int g(M_{200}, z)\, \mathrm{d}\ln M_{200}\,, \tag{6}$$

where $\mathrm{d}N(M_{200}, z)/\mathrm{d}\ln M_{200}$ is the mass function of dark matter haloes. The function

$$g(M_{200}, z) = M_{200}\, \frac{\mathrm{d}N(M_{200}, z)}{\mathrm{d}\ln M_{200}}\, \frac{\dot{M}_\star}{M_{200}} \tag{7}$$

is the *multiplicity function* of star formation, which describes the SFRD in haloes of a given mass. As we showed in Section 3, the halo mass function is a strong function of environment, with the largest haloes forming only in the densest regions. However, as we showed in Section 5.1 the star formation properties of a halo *at a given epoch*, including its sSFR $\dot{M}_\star/M_{200}$, depend on its mass, not on its environment (Fig. 8), i.e. at any given epoch $\dot{M}_\star/M_{200}$ is essentially independent of large-scale environment. The dependence of $\dot{\rho}_\star$ on large-scale environment therefore reflects the dependence of the halo mass function on large-scale environment.

The multiplicity function of star formation, $g(M_{200}, z)$, is shown in Fig. 11 for each GIMIC region at three different redshifts; the SFRD of each region is proportional to the shaded area. As highlighted by Springel & Hernquist (2003b), the history of $\dot{\rho}_\star(z)$ is shaped by the buildup of matter in haloes, as described by $\mathrm{d}N/\mathrm{d}\ln M_{200}$, and the evolution of the sSFR, $\dot{M}_\star/M_{200}$. On the one hand, advancing structure formation shifts an ever increasing fraction of bound mass into more massive haloes, thereby pushing more and more gas above the efficiency thresholds imposed by the photoionizing background and galactic winds. Conversely, the amplitude of $\dot{M}_\star/M_{200}$ is greatest at early times and drops steadily over time. The combination of these effects yields a broad plateau in $\dot{\rho}_\star$ at intermediate redshifts ($6 \gtrsim z \gtrsim 2$), from which it drops off towards lower and higher redshifts.

Massive haloes are underrepresented in low-density or 'void' regions and consequently star formation within voids is dominated by low-mass haloes at high redshift (i.e. $z = 6$), with the multiplicity function of star formation peaking at $\sim 10^9\, h^{-1}\, \mathrm{M}_\odot$. Only overdense regions have sufficient numbers of massive haloes at these early times for massive galaxies to contribute significantly to $\dot{\rho}_\star$; in these regions, a second peak develops at $\sim 10^{11}\, h^{-1}\, \mathrm{M}_\odot$. By $z = 1$, massive haloes begin to be found even in voids and so the first peak in g moves to haloes of mass $\sim 10^{12}\, h^{-1}\, \mathrm{M}_\odot$ in all regions with a second one developing around $\sim 10^{13}\, h^{-1}\, \mathrm{M}_\odot$ in the high-density regions. At the present time, star formation is dominated by galactic haloes ($M_{200} \sim 10^{12}\, h^{-1}\, \mathrm{M}_\odot$) in the void region, by groups ($M_{200} \sim 10^{13-14}\, h^{-1}\, \mathrm{M}_\odot$) in the intermediate-density regions and by clusters ($M_{200} \sim 10^{15}\, h^{-1}\, \mathrm{M}_\odot$) in the highest density region. It is the absence of massive, efficient star-forming haloes in voids that so severely inhibits their overall SFRD. This explains the much greater difference in $\dot{\rho}_\star$ between the underdense -2σ and -1σ regions, compared to that between the overdense $+1\sigma$ and $+2\sigma$ ones.

Our simulations overestimate the sSFRs in massive galaxies and haloes; hence, the contribution of clusters is likely to be exaggerated. Even so, the multiplicity function plot illustrates the dynamic range problem faced by simulations of galaxy formation: $\dot{\rho}_\star$ is dominated by $10^9\, h^{-1}\, \mathrm{M}_\odot$ haloes at $z = 6$ but (in our simulations at least) by $10^{15}\, h^{-1}\, \mathrm{M}_\odot$ haloes at $z = 0$. As an example, the Millennium Simulation does not resolve haloes below $2 \times 10^{10}\, h^{-1}\, \mathrm{M}_\odot$ and
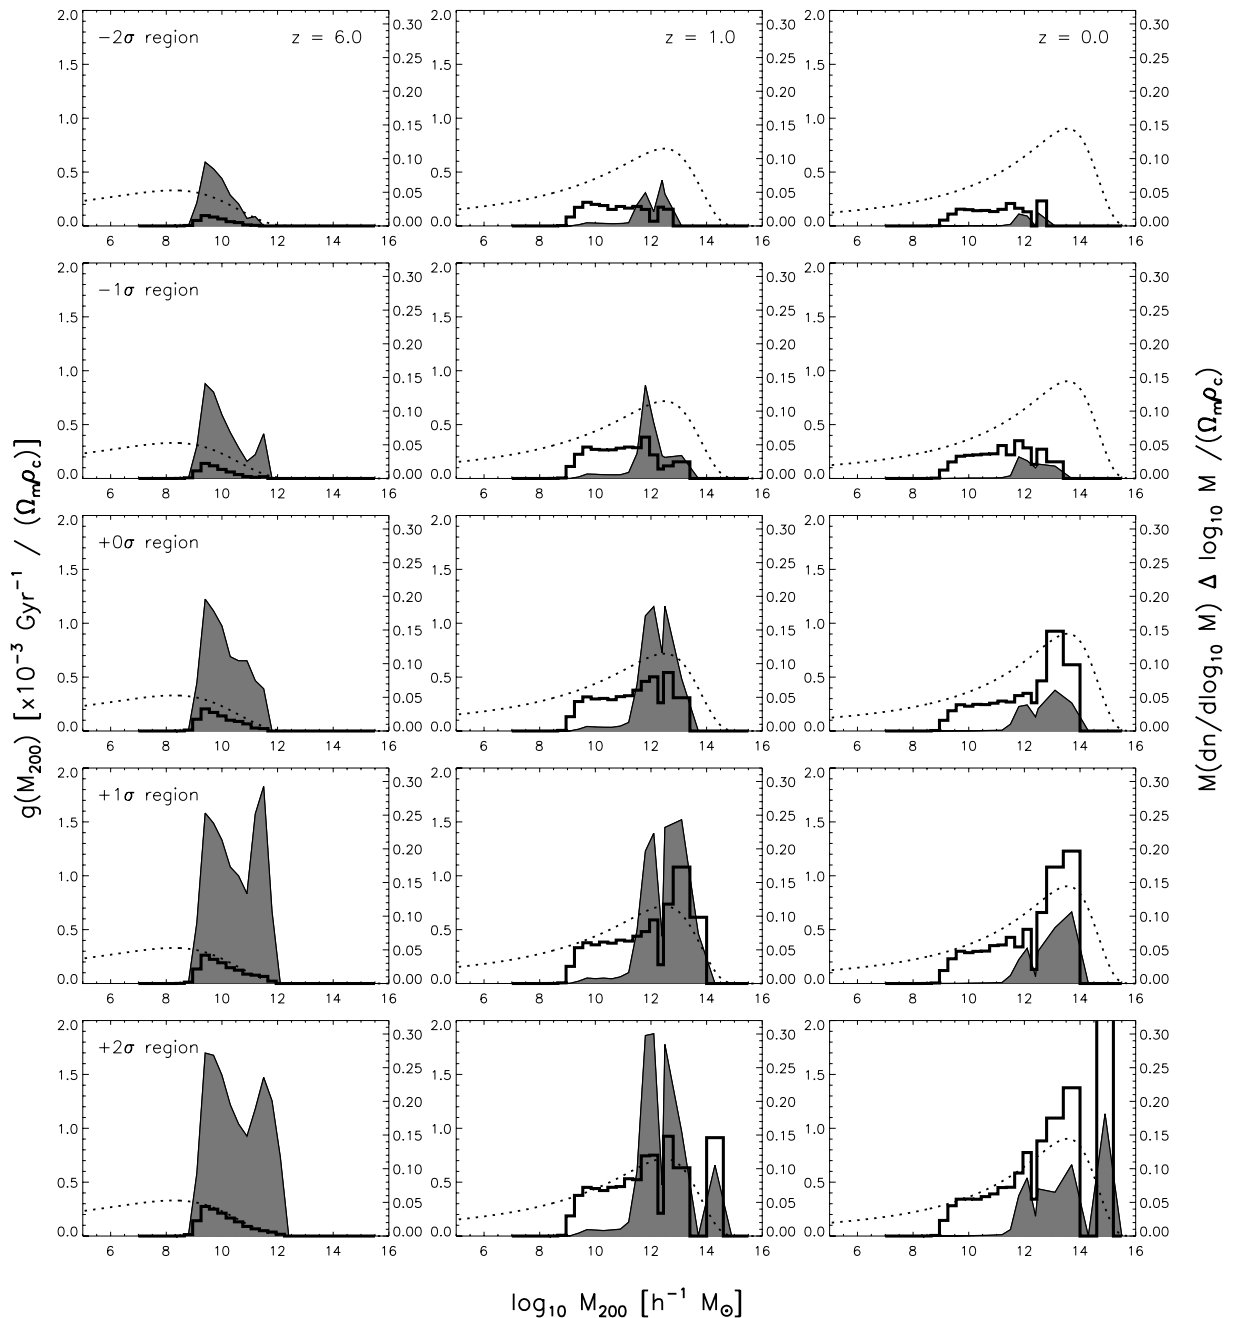
Figure 11. The multiplicity function, $g(M_{200}, z)$, of star formation (shaded) as a function of region (rows) and redshift (columns). This function is the product of the halo sSFR and the halo multiplicity function, both of which are functions of redshift and halo mass. For reference, the halo multiplicity function of each region is reproduced from Fig. 2 (black histogram, right-hand axis), as is the multiplicity function derived from the universal mass function fit of Reed et al. (2007, dotted curve, right-hand axis).

hence misses the objects that (in our simulations) are the dominant contributors to star formation at $z = 6$. On the other hand, periodic simulations of, for example, side length $L = 32\,h^{-1}$ Mpc enclose a total mass of $\sim 10^{15}\,h^{-1}\,\mathrm{M}_\odot$, and so cannot contain even a single cluster of the type that dominates $\dot{\rho}_\star$ in the $+2\sigma$ region. However, the fact that $\dot{M}_\star/M_{200}$ turns out to depend mostly on halo mass and not on overdensity makes it possible, in principle, to obtain accurate values for $\dot{\rho}_\star$ and its dependence on environment. This may be achieved by combining accurate estimates of $\dot{M}_\star/M_{200}$ as a function of halo mass with the dependence of the halo mass function on environment.

5.3 Halo baryon fractions

Stars comprise a relatively small fraction of the total baryonic content of the Universe. In this section, we investigate how the total baryonic content of haloes varies with halo circular velocity and with time. This allows us not only to generate an inventory of the baryons in our simulations but also to understand further the processes responsible for the star formation properties discussed in the previous section. As was the case for the sSFR, the baryon fractions of haloes are essentially independent of environment and so, when necessary, we combine haloes from all five GIMIC regions to
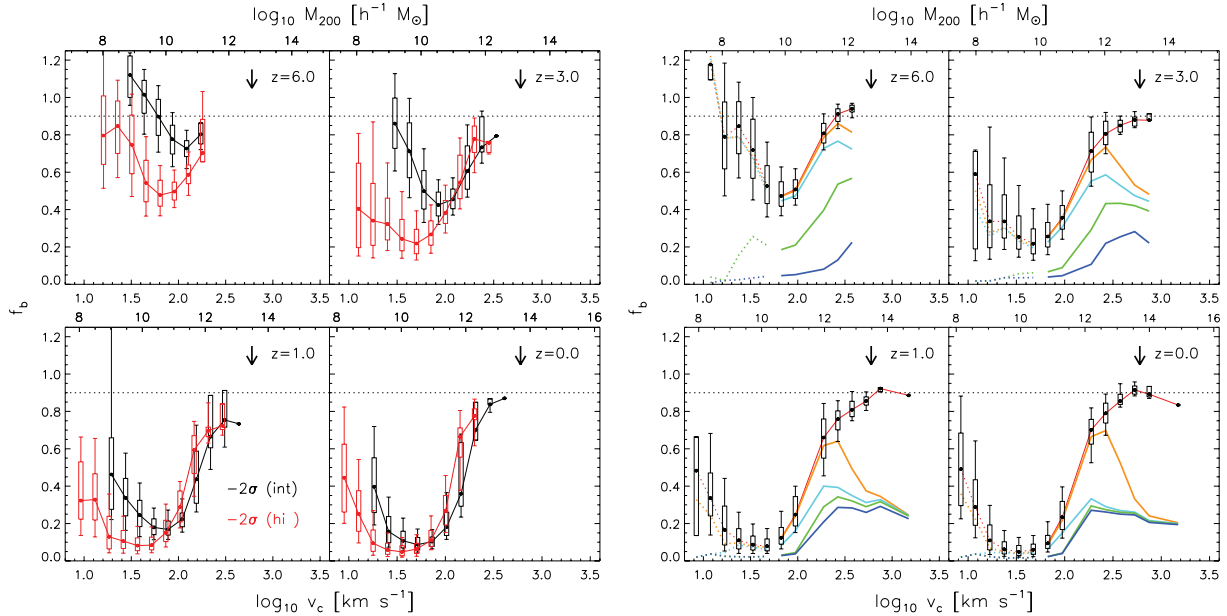
Figure 12. Median baryon fraction, $f_b \equiv (M_b/M_{200})/(\Omega_b/\Omega_m)$, of the simulated haloes. The dotted line indicates the 90 per cent level expected in the non-radiative regime; the arrows indicate the wind launch velocity. Box and whisker elements show the median and 10th, 25th, 75th and 90th percentiles of the total. Left-hand panel: resolution test, comparing the two realizations of the -2σ region. Agreement between the high- and intermediate-resolution runs for the total baryon fraction is achieved at a circular velocity that decreases with redshift; at $z = 0$, this scale is $\sim 50\,\mathrm{km\,s^{-1}}$. Right-hand panel: coloured curves decompose the baryon fraction into the *cumulative* contribution made by stars (dark blue), equation of state gas (green), cold gas (cyan), warm gas (orange) and hot gas (red), in that order. Below $160\,\mathrm{km\,s^{-1}}$, results are drawn from the high-resolution -2σ region; above this scale, they are drawn from all five intermediate-resolution regions. Dotted sections of the curves ($v_c < 50\,\mathrm{km\,s^{-1}}$) indicate the regime where we believe the high-resolution simulations to be reliable.

improve the statistics. As before, we begin this discussion with a convergence analysis.

The baryon fraction in units of the mean, $f_b \equiv (M_b/M_{200})/(\Omega_b/\Omega_m)$, is plotted as a function of circular velocity (and halo mass, top axis) at several redshifts in Fig. 12. To investigate convergence, we use the intermediate- and high-resolution versions of the -2σ region. The left-hand plot shows that approximate convergence is achieved at a circular velocity of $v_c \sim 100\,\mathrm{km\,s^{-1}}$, the same value at which, as shown in the preceding subsection, approximate convergence in sSFR is also achieved. As before, extrapolating from the behaviour seen in Fig. 12, we assume that baryon fractions in the high-resolution simulation are reliable down to $v_c \sim 50\,\mathrm{km\,s^{-1}}$ (that corresponds to the resolution limit scaling with the number of particles in the intermediate- and high-resolution simulations.) The rapid upturn in the baryon fraction at $v_c \lesssim 30\,\mathrm{km\,s^{-1}}$ is thus likely to be an artefact of the limited resolution.

The right-hand plot of Fig. 12 shows how the baryon content of a halo is made up from the following components:

 (i) stars
 (ii) gas with $n_H > 0.1\,\mathrm{cm^{-3}}$ (the value above which we impose an equation of state to mimic the multiphase ISM – see Section 2.2.)
 (iii) cold gas
 (iv) warm gas
 (v) hot gas.

The cold and warm gas components are separated at $T = 10^5\,\mathrm{K}$, whilst the warm and hot gas components are separated at $T = 10^{6.5}\,\mathrm{K}$. These limits are chosen so as roughly to distinguish the warm–hot medium from hotter X-ray emitting gas. The fractions of the different components are plotted *cumulatively* in the order listed

above, so that the curve showing the hot component is equivalent to the total baryon fraction. For $v_c < 160\,\mathrm{km\,s^{-1}}$, we consider only the (well-resolved) haloes from the high-resolution -2σ simulation; for larger values of v_c, we draw haloes from all five intermediate-resolution runs to obtain a large sample of well-resolved haloes.

Massive haloes with circular velocity comparable to the launch speed of winds or greater hold on to their baryons and, as shown in Fig. 12, have $f_b \sim 0.9$, as expected in the non-radiative regime (e.g. Eke, Navarro & Frenk 1998; Crain et al. 2007). The baryon fraction drops rapidly for haloes with $v_c \lesssim 200\,\mathrm{km\,s^{-1}}$. These haloes have a virial temperature that coincides with the peak in the cooling function. Efficient gas cooling leads to high rates of star formation, but the associated galactic winds are powerful enough to drive the gas out of the potential wells. As a result these haloes have, in fact, *lower* sSFRs (see Fig. 8) and baryon fractions than more massive haloes.

Whereas the total baryon fraction in haloes with $v_c \gtrsim 300\,\mathrm{km\,s^{-1}}$ is nearly independent of redshift, the fraction in the star-forming 'equation-of-state' gas phase drops dramatically from 35 per cent at $z = 6$ to a negligible fraction at $z = 0$, when ~ 75 per cent is in the form of hot gas. The overall baryon fraction of less massive haloes, $v_c \sim 100\,\mathrm{km\,s^{-1}}$, drops significantly with redshift. Their stellar fractions increase with time and become roughly constant for $z < 3$, at around 5 per cent of the cosmic mean. Because of their short gas cooling time, most of the gas in these low-mass haloes is initially in the cold phase. As star formation proceeds, winds typically eject over half of the baryons that are partially converted into warm and hot gas. Other mechanisms such as ram-pressure stripping of gas and tidal stripping of stars can contribute to the removal of baryons from these haloes.

6 SUMMARY AND CONCLUSIONS

We have presented the first results from GIMIC, a programme of hydrodynamical simulations conducted by the Virgo consortium that follow the formation of galaxies in a ΛCDM universe. The goal of the programme is to investigate the joint evolution of galaxies and the IGM. The simulations follow the evolution of nearly spherical regions of radius $\sim$20 h^{-1} Mpc, drawn from the Millennium Simulation (Springel et al. 2005). The regions have mean overdensities of $(-2, -1, 0, +1, +2)\sigma$, where σ is the rms overdensity on the scale of the regions at $z = 1.5$. The rest of the Millennium Simulation volume is simulated at much lower resolution using collisionless particles. This provides the correct tidal forces on the high-resolution regions.

The five GIMIC regions were followed to $z = 0$ at intermediate resolution ($m_{gas} = 1.16 \times 10^7 \, h^{-1} \, M_\odot$). Of the high-resolution simulations ($m_{gas} = 1.45 \times 10^6 \, h^{-1} \, M_\odot$), the -2σ region was continued to $z = 0$, whilst the $(-1, 0, +1)\sigma$ regions were stopped at $z = 2$. In the highest resolution simulations, the Jeans mass in the IGM after reionization is marginally resolved and thus these simulations can, in principle, track the formation of all galaxies whose cooling is dominated by H I.

Our simulation strategy has a number of advantages: (i) it samples the entire range of large-scale environments – including rare objects such as massive cluster haloes and voids – allowing us to examine the dependence of galaxy properties on large-scale overdensity; (ii) it allows accurate integration to $z = 0$ since fluctuations on the scale of the computational volume remain linear; (iii) by suitably averaging over the five regions, it allows an estimate of the statistical properties of representative cosmological volumes.

The simulations were carried out using the GADGET3 code, a substantial upgrade of GADGET2 (Springel 2005) that includes:

(i) a recipe for star formation designed to enforce a local Kennicutt–Schmidt law (Schaye & Dalla Vecchia 2008);

(ii) stellar evolution and the associated delayed release of 11 chemical elements (Wiersma et al. 2009b);

(iii) the contribution of metals to the cooling of gas, computed element by element, in the presence of an imposed UV background (Wiersma et al. 2009a) and

(iv) galactic winds that pollute the IGM with metals and can quench star formation in low-mass haloes (Dalla Vecchia & Schaye 2008).

For these simulations, we do not, however, follow the evolution of black holes or feedback effects associated with them.

The simulations produce the correct number density of galaxies with mass greater than about $10^9 \, M_\odot$. At $z = 2$, the stellar mass function is consistent with observations. At $z = 0$, there are too many galaxies of very low and very high mass and not enough of intermediate mass. The excess at the massive end reflects the absence of a heating mechanism, such as that often ascribed to AGN feedback, to prevent too much gas cooling. At small and intermediate masses, the discrepancy with observations reflects inadequacies in our treatment of galactic winds. For example, the paucity of galaxies with stellar mass $\sim 10^{10} \, M_\odot$ seems to be directly related to our simple wind model.

In this paper, we have concentrated on the star formation properties of the simulations. These are governed by an interplay between accretion, gas cooling, photoheating by the imposed ionising background and galactic winds. The sSFR in our simulations is converged for $z \lesssim 6$, and we estimate that the value of $\dot{\rho}_\star(z)$ in the high-resolution simulations is reliable for $z < z_{reion} = 9$. Indeed,

the reionization process is 'visible' in these simulations as a depression of the SFR in low-mass haloes. In all five GIMIC regions, $\dot{\rho}_\star(z)$ increases slowly with time, reaches a broad maximum around $z \sim 2$ and then falls rapidly by a factor of $\sim$6 to the present day. In general, the global value of $\dot{\rho}_\star(z)$, obtained by a weighted average over the individual regions, is in reasonable agreement with the observational compilation of Hopkins (2007). However, the decline in $\dot{\rho}_\star$ at low redshift in the simulation is not as pronounced as in the observations, largely because of excessive star formation in massive haloes.

One of the main results of this paper is the strong environmental dependence we find of the SFRD, $\dot{\rho}_\star(z)$, which is manifest as a large variation amongst the five GIMIC regions. Although the shapes of the $\dot{\rho}_\star(z)$ curves are similar for all regions, the offset in amplitude between the -2σ to the $+2\sigma$ region is approximately an order of magnitude. The 0σ region closely tracks the global value reconstructed from a weighted average of all five regions and is typically a factor of 2 below the $+2\sigma$ region. Even when normalized by total mass rather than volume, $\dot{\rho}_\star$ still varies by more than a factor of 3 between the two extreme regions. This environmental dependence makes it difficult to obtain reliable observational estimates, particularly at high-redshift. For example, at $z = 6$ the $18 \, h^{-1}$ Mpc GIMIC spheres correspond to 320 arcmin^2 in the sky, which is a considerably larger area than the (5–20 arcmin^2) covered by the HUDF.

At a given redshift, we find that the sSFR in dark matter haloes, $\dot{M}_\star/M_{200}$ (the 'halo sSFR'), depends only on halo mass, not on large-scale environment. At all redshifts, the specific star formation is most efficient in haloes of circular velocity $v_c = 200$–250 km s^{-1}. Above this value, the sSFR varies only little with v_c, but below this value it drops rapidly due to the action of galactic winds that deplete these haloes of baryons to below $\sim$10 per cent of the cosmic mean by $z = 0$. Above this critical value of v_c, the sSFR is well described by the spherically symmetric cooling model of White & Frenk (1991). Because of their generally higher baryon fractions and gas densities, haloes of a given mass tend to have higher sSFR at earlier times.

Since the halo sSFR is independent of environment, the reason for the strong dependence of the SFRD with environment is a strong dependence of the halo mass function on environment. There are always more haloes and galaxies in the $+2\sigma$ region than in the -2σ region. For example, the volume-normalized number density of galaxies with stellar mass $>10^{10} \, h^{-1} \, M_\odot$ at $z = 0$ is a factor of $\sim$3 greater in the $+2\sigma$ than the -2σ region. In addition, the most massive haloes are strongly biased towards regions of high overdensity. Since the sSFR in our simulations is always dominated by high-mass haloes with $v_c = 200$–250 km s^{-1}, even when normalized to the total mass in each region, the $+2\sigma$ region always has a higher *specific* SFRD than the -2σ region. The combined effect of halo number density and star formation efficiency results in star formation at the present time that is dominated by galactic haloes ($M_{200} \sim 10^{12} \, h^{-1} \, M_\odot$) in the void region, by groups ($M_{200} \sim 10^{13-14} \, h^{-1} \, M_\odot$) in the intermediate-density regions and by clusters, $M_{200} \sim 10^{15} \, h^{-1} \, M_\odot$, in the highest density region.

The plateau in the global $\dot{\rho}_\star(z)$ relation at $6 \gtrsim z \gtrsim 2$ is due to the balance between the advancing formation of dark matter haloes and the decline in the sSFR in haloes of a given mass (see Figs 8 and 11). Eventually, the latter dominates and $\dot{\rho}_\star$ plummets in all regions. The lack of feedback effects associated with the growth of central black holes may affect the quantitative, but perhaps not the qualitative, behaviour of these trends. This view is supported by the agreement, to within a factor of $\sim$2, between the overall SFRDs inferred for the Millennium Simulation from our simulations and

from the semi-analytic models of Bower et al. (2006), Croton et al. (2006) and De Lucia et al. (2006), which do include the effects of AGN feedback.

In common with the conclusions of recent studies employing semi-analytic modelling, we find that the growth of galaxies in our simulations is segregated by stellar mass: low-mass galaxies continue to grow at all epochs due to ongoing star-formation, whilst more massive galaxies have largely concluded their star formation at $z > 1$. This apparent 'downsizing' of the star formation activity occurs in spite of the absence of AGN feedback in the simulations.

The SFRD in the simulations is always dominated by the progenitors of galaxies that today reside in the most massive haloes (group/cluster haloes with $M \geq 5 \times 10^{12} \, h^{-1} \, M_\odot$). Yet, these stars were predominantly formed either in dwarf haloes ($M \leq 5 \times 10^{11} \, h^{-1} \, M_\odot$) before $z \sim 3.5$ or in galaxy-sized haloes after that. Galactic winds are an effective mechanism for quenching star formation to the observed levels in these smaller systems. The integrated stellar mass density at $z = 0$ agrees with observational estimates to within a factor of $\lesssim 2$.

Some properties of the galaxies in the simulations can be compared with observations. Galaxies form first in the high-density region. Between $z \simeq 9$ and 6, they are very compact (radii $\sim 3 \, h^{-1}$ kpc), massive ($M_\star \sim 10^{11} \, h^{-1} \, M_\odot$), contain relatively old stars (~ 0.2 Gyr old), have high metallicity ($Z \sim Z_\odot$) and are over-abundant in α-elements. These galaxies are reminiscent of those recently found in the GOODS fields by Wiklind et al. (2008).

In summary, the GIMIC simulations provide a useful resource to investigate the evolution of and interaction between galaxies and the intergalactic gas in different large-scale environments. In this first paper of a series, we have explored the star formation properties of haloes and some of the properties of the galaxies that form within them. In subsequent papers, we will examine further properties of both galaxies and intergalactic gas.

ACKNOWLEDGMENTS

REFERENCES

Abadi M. G., Navarro J. F., Steinmetz M., Eke V. R., 2003a, ApJ, 591, 499
Abadi M. G., Navarro J. F., Steinmetz M., Eke V. R., 2003b, ApJ, 597, 21
Abel T., Haehnelt M. G., 1999, ApJ, 520, L13
Adelberger K. L., Steidel C. C., Shapley A. E., Pettini M., 2003, ApJ, 584, 45
Bagla J. S., Ray S., 2005, MNRAS, 358, 1076
Baugh C. M., 2006, Reports Progress Phys., 69, 3101
Beckwith S. V. W. et al., 2006, AJ, 132, 1729
Benson A. J., Lacey C. G., Baugh C. M., Cole S., Frenk C. S., 2002, MNRAS, 333, 156
Benson A. J., Bower R. G., Frenk C. S., Lacey C. G., Baugh C. M., Cole S., 2003, ApJ, 599, 38

Bolton J. S., Haehnelt M. G., 2007, MNRAS, 382, 325
Booth C. M., Schaye J., 2009, preprint (arXiv:0904.2572)
Bouwens R. J., Illingworth G. D., Franx M., Ford H., 2008, ApJ, 686, 230
Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, MNRAS, 370, 645
Brooks A. M., Governato F., Booth C. M., Willman B., Gardner J. P., Wadsley J., Stinson G., Quinn T., 2007, ApJ, 655, L17
Bunker A., Stanway E., Ellis R., McMahon R., Eyles L., Lacy M., 2006, New Astron. Rev., 50, 94
Caucci S., Colombi S., Pichon C., Rollinde E., Petitjean P., Sousbie T., 2008, MNRAS, 386, 211
Cen R., Miralda-Escudé J., Ostriker J. P., Rauch M., 1994, ApJ, 437, L9
Cen R., Ostriker J. P., 1999, ApJ, 514, 1
Chabrier G., 2003, PASP, 115, 763
Cole S., 1991, ApJ, 367, 45
Cole S., 1997, MNRAS, 286, 38
Cole S., Aragon-Salamanca A., Frenk C. S., Navarro J. F., Zepf S. E., 1994, MNRAS, 271, 781
Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, MNRAS, 319, 168
Cole S. et al., 2001, MNRAS, 326, 255
Colless M. et al., 2001, MNRAS, 328, 1039
Couchman H. M. P., Thomas P. A., Pearce F. R., 1995, ApJ, 452, 797
Cowie L. L., Songaila A., Hu E. M., Cohen J. G., 1996, AJ, 112, 839
Crain R. A., Eke V. R., Frenk C. S., Jenkins A., McCarthy I. G., Navarro J. F., Pearce F. R., 2007, MNRAS, 377, 41
Croft R. A. C., Di Matteo T., Springel V., Hernquist L., 2008, MNRAS, submitted (arXiv:0803.4003)
Croton D. J. et al., 2006, MNRAS, 365, 11
Dalla Vecchia C., Schaye J., 2008, MNRAS, 387, 1431
Davé R., Oppenheimer B. D., 2007, MNRAS, 374, 427
Davé R., Hernquist L., Katz N., Weinberg D. H., 1999, ApJ, 511, 521
Davé R., Finlator K., Oppenheimer B. D., 2006, MNRAS, 370, 273
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
Davis M. et al., 2003, in Guhathakurta P., ed., Proc. SPIE Conf., Vol. 4834, Discoveries and Research Prospects from 6- to 10-Meter-Class Telescopes II. SPIE, Bellingham, p. 161
De Lucia G., Springel V., White S. D. M., Croton D., Kauffmann G., 2006, MNRAS, 366, 499
Di Matteo T., Colberg J., Springel V., Hernquist L., Sijacki D., 2008, ApJ, 676, 33
Dickinson M., Giavalisco M., The Goods Team, 2003, in Bender R., Renzini A., eds, The Mass of Galaxies at Low and High Redshift. p. 324
Dolag K., Borgani S., Murante G., Springel V., 2008, preprint (arXiv:0808.3401)
Drory N., Bender R., Feulner G., Hopp U., Maraston C., Snigula J., Hill G. J., 2003, ApJ, 595, 698
Drory N., Salvato M., Gabasch A., Bender R., Hopp U., Feulner G., Pannella M., 2005, ApJ, 619, L131
Efstathiou G., 1992, MNRAS, 256, 43P
Eke V. R., Navarro J. F., Frenk C. S., 1998, ApJ, 503, 569
Eke V. R., Baugh C. M., Cole S., Frenk C. S., King H. M., Peacock J. A., 2005, MNRAS, 362, 1233
Ferland G. J., Korista K. T., Verner D. A., Ferguson J. W., Kingdon J. B., Verner E. M., 1998, PASP, 110, 761
Font A. S. et al., 2008, MNRAS, 389, 1619
Fontana A. et al., 2004, A&A, 424, 23
Frenk C. S., Evrard A. E., White S. D. M., Summers F. J., 1996, ApJ, 472, 460
Frenk C. S., White S. D. M., Davis M., Efstathiou G., 1988, ApJ, 327, 507
Gingold R. A., Monaghan J. J., 1977, MNRAS, 181, 375
Governato F. et al., 2004, ApJ, 607, 688
Governato F., Willman B., Mayer L., Brooks A., Stinson G., Valenzuela O., Wadsley J., Quinn T., 2007, MNRAS, 374, 1479
Green J. C., Wilkinson E., Morse J. A., 2003, in Blades J. C., Siegmund O. H. W., eds, Proc. SPIE Conf., Vol. 4854, Future EUV/UV and Visible Space Astrophysics Missions and Instrumentation. SPIE, Bellingham, p. 72

Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, Clusters of Galaxies and the High Redshift Universe Observed in X-rays, preprint (astro-ph/0106018)

Heckman T. M., Armus L., Miley G. K., 1990, ApJS, 74, 833

Hernquist L., Katz N., Weinberg D. H., Miralda-Escudé J., 1996, ApJ, 457, L51

Hernquist L., Springel V., 2003, MNRAS, 341, 1253

Hoeft M., Yepes G., Gottlöber S., Springel V., 2006, MNRAS, 371, 401

Hopkins A. M., 2007, in Afonso J., Ferguson H. C., Mobasher B., Norris R., eds, ASP Conf. Ser. Vol. 380, Deepest Astronomical Surveys. Astron. Soc. Pac., San Francisco, p. 423

Jenkins A., Frenk C. S., White S. D. M., Colberg J. M., Cole S., Evrard A. E., Couchman H. M. P., Yoshida N., 2001, MNRAS, 321, 372

Juneau S. et al., 2005, ApJ, 619, L135

Katz N., Gunn J. E., 1991, ApJ, 377, 365

Kauffmann G., White S. D. M., Guiderdoni B., 1993, MNRAS, 264, 201

Kennicutt R. C. Jr., 1998, ARA&A, 36, 189

Kobayashi C., Springel V., White S. D. M., 2007, MNRAS, 376, 1465

Kodama T., Balogh M. L., Smail I., Bower R. G., Nakata F., 2004, MNRAS, 354, 1103

Lacey C., Cole S., 1994, MNRAS, 271, 676

Le Fèvre O. et al., 2005, A&A, 439, 877

Li C., White S. D. M., 2009, MNRAS, in press (arXiv:0901.0706)

Lilly S. J., Le Fevre O., Hammer F., Crampton D., 1996, ApJ, 460, L1

Lucy L. B., 1977, AJ, 82, 1013

Madau P., Ferguson H. C., Dickinson M. E., Giavalisco M., Steidel C. C., Fruchter A., 1996, MNRAS, 283, 1388

Marigo P., 2001, A&A, 370, 194

Martin C. L., 1999, ApJ, 513, 156

Miralda-Escudé J., Cen R., Ostriker J. P., Rauch M., 1996, ApJ, 471, 582

Mo H. J., White S. D. M., 1996, MNRAS, 282, 347

Nagamine K., Cen R., Hernquist L., Ostriker J. P., Springel V., 2005, ApJ, 618, 23

Navarro J. F., Benz W., 1991, ApJ, 380, 320

Navarro J. F. et al., 2004, MNRAS, 349, 1039

Neistein E., van den Bosch F. C., Dekel A., 2006, MNRAS, 372, 933

Ocvirk P., Pichon C., Teyssier R., 2008, MNRAS, 1131

Okamoto T., Eke V. R., Frenk C. S., Jenkins A., 2005, MNRAS, 363, 1299

Okamoto T., Gao L., Theuns T., 2008, MNRAS, 390, 920

Oppenheimer B. D., Davé R., 2006, MNRAS, 373, 1265

Oppenheimer B. D., Davé R., 2008a, MNRAS, 387, 577

Oppenheimer B. D., Davé R., 2008b, MNRAS, 387, 577

Pawlik A. H., Schaye J., 2009, MNRAS, 396, L46

Pearce F. R., Jenkins A., Frenk C. S., White S. D. M., Thomas P. A., Couchman H. M. P., Peacock J. A., Efstathiou G., 2001, MNRAS, 326, 649

Pearce F. R. et al., 1999, ApJ, 521, L99

Pettini M., Rix S. A., Steidel C. C., Adelberger K. L., Hunt M. P., Shapley A. E., 2002, ApJ, 569, 742

Pichon C., Vergely J. L., Rollinde E., Colombi S., Petitjean P., 2001, MNRAS, 326, 597

Portinari L., Chiosi C., Bressan A., 1998, A&A, 334, 505

Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, MNRAS, 338, 14

Pozzetti L. et al., 2003, A&A, 402, 837

Reed D. S., Bower R., Frenk C. S., Jenkins A., Theuns T., 2007, MNRAS, 374, 2

Robertson B., Yoshida N., Springel V., Hernquist L., 2004, ApJ, 606, 32

Rudnick G. et al., 2003, ApJ, 599, 847

Salpeter E. E., 1955, ApJ, 121, 161

Scannapieco C., Tissera P. B., White S. D. M., Springel V., 2008, MNRAS, 389, 1137

Schaye J., 2004, ApJ, 609, 667

Schaye J., Dalla Vecchia C., 2008, MNRAS, 383, 1210

Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, MNRAS, 318, 817

Shapley A. E., Steidel C. C., Pettini M., Adelberger K. L., 2003, ApJ, 588, 65

Sheth R. K., Tormen G., 2002, MNRAS, 329, 61

Sijacki D., Pfrommer C., Springel V., Enßlin T. A., 2008, MNRAS, 387, 1403

Sirko E., 2005, ApJ, 634, 728

Smith M. C. et al., 2007, MNRAS, 379, 755

Somerville R. S., Primack J. R., 1999, MNRAS, 310, 1087

Somerville R. S., Hopkins P. F., Cox T. J., Robertson B. E., Hernquist L., 2008, MNRAS, 391, 481

Sommer-Larsen J., Götz M., Portinari L., 2003, ApJ, 596, 47

Springel V., 2005, MNRAS, 364, 1105

Springel V., Hernquist L., 2002, MNRAS, 333, 649

Springel V., Hernquist L., 2003a, MNRAS, 339, 289

Springel V., Hernquist L., 2003b, MNRAS, 339, 312

Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726

Springel V. et al., 2005, Nat, 435, 629

Srbinovsky J., Wyithe S., 2008, preprint (arXiv:0807.4782)

Steinmetz M., White S. D. M., 1997, MNRAS, 288, 545

Swinbank A. M., Bower R. G., Smith G. P., Wilman R. J., Smail I., Ellis R. S., Morris S. L., Kneib J.-P., 2007, MNRAS, 376, 479

Teyssier R., 2002, A&A, 385, 337

Thacker R. J., Couchman H. M. P., 2001, ApJ, 555, L17

Theuns T., Leonard A., Efstathiou G., Pearce F. R., Thomas P. A., 1998, MNRAS, 301, 478

Theuns T., Bernardi M., Frieman J., Hewett P., Schaye J., Sheth R. K., Subbarao M., 2002, ApJ, 574, L111

Thielemann F.-K. et al., 2003, Nuclear Phys. A, 718, 139

Wadsley J. W., Stadel J., Quinn T., 2004, New Astron., 9, 137

Weil M. L., Eke V. R., Efstathiou G., 1998, MNRAS, 300, 773

White S. D. M., Frenk C. S., 1991, ApJ, 379, 52

White S. D. M., Rees M. J., 1978, MNRAS, 183, 341

Wiersma R. P. C., Schaye J., Smith B. D., 2009a, MNRAS, 393, 99

Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, MNRAS, in press (arXiv:0902.1535)

Wilman R. J., Gerssen J., Bower R. G., Morris S. L., Bacon R., de Zeeuw P. T., Davies R. L., 2005, Nat, 436, 227

Wiklind T., Dickinson M., Ferguson H. C., Giavalisco M., Mobasher B., Grogin N. A., Panagia N., 2008, ApJ, 676, 781

Wise J. H., Cen R., 2009, ApJ, 693, 984

Wolf C., Meisenheimer K., Rix H.-W., Borch A., Dye S., Kleinheinrich M., 2003, A&A, 401, 73

York D. G. et al., 2000, AJ, 120, 1579

Zavala J., Okamoto T., Frenk C. S., 2008, MNRAS, 387, 364

Zhang Y., Anninos P., Norman M. L., 1995, ApJ, 453, L57

Zhang Y., Anninos P., Norman M. L., Meiksin A., 1997, ApJ, 485, 496

APPENDIX A: METHODS IN DETAIL

In this appendix, we expand on the technical details of the generation of the initial conditions for the GIMIC simulations. Secondly, we discuss how we address the existence of an artificial boundary between the high-resolution, baryon-filled region of each simulation and the external low-resolution region that is simulated with collisionless dynamics. As part of this latter discussion, we describe our technique for combining results from the five GIMIC simulations to construct estimates of properties of the Millennium Simulation as a whole.

A1 Generation of the initial conditions

Here, we provide a more detailed description of the generation of the initial conditions, than the brief overview that was given in Section 2.2. As described there, we aim to trace a representative sample of the $(500\,h^{-1}\,\mathrm{Mpc})^3$ Millennium Simulation volume with five separate simulations, and also to follow the evolution of rare structures such as voids and clusters that are only found in such

large periodic volumes. Clearly, it is desirable to perform these simulations at high resolution, and we aimed to have a sufficiently low particle mass that the Jeans scale in the IGM (after the epoch of reionization) would be resolved in our 'high-resolution' runs (which have eight times more particles than our 'intermediate-resolution' runs). In practice, this limits the size of the regions to $\sim 18 \, h^{-1}$ Mpc. However, our additional constraint that the $+2\sigma$ region should be centred on a rich cluster at $z = 0$ required the size of this region to be increased to $25 \, h^{-1}$ Mpc; this is because the mass of a sphere of mean density and radius $18 \, h^{-1}$ Mpc at $z = 0$ is $1.7 \times 10^{15} \, h^{-1} \, M_{\odot}$, a mass comparable to that of the largest cluster in the Millennium Simulation. We describe the generation of the initial conditions in two parts: (i) the selection of the regions to be resimulated and (ii) the generation of the particle distribution and the displacement field.

A1.1 Selecting the five regions

The four $(-2, -1, 0, +1)\sigma$ spheres were selected from a target list of 10^5 randomly placed spheres of radius $18 \, h^{-1}$ Mpc in the Millennium Simulation volume at $z = 1.5$. Since by $z = 1.5$ the overdensity distribution has become significantly skewed (i.e. non-Gaussian), we defined the overdensities of the $(-2, -1, 0, +1)\sigma$ regions as those that correspond in terms of the rank ordering by overdensity, to a Gaussian distribution; the overdensities for the $(-2, -1, 0, +1)\sigma$ regions are $(-0.407, -0.243, -0.032, 0.241)$, respectively. From the complete list of 10^5 spheres, a candidate list of centres was generated for each of the regions by selecting all spheres with overdensities within 0.002 of the respective limits.

The procedure for the $+2\sigma$ sphere was different because of the requirement that a rich cluster should form at its centre by $z = 0$. The starting point was to identify the particles that form the 126 most massive groups (that have a characteristic separation of $\sim 100 \, h^{-1}$ Mpc) in the Millennium Simulation volume at $z = 0$, identified by a FoF group finder (Davis et al. 1985) with linking length $b = 0.2$. These particles were traced back to $z = 1.5$ and their centres of mass used as centres for spheres of radius $25 \, h^{-1}$ Mpc. The $+2\sigma$ sphere was then selected from a randomly sorted list of 12 candidates that had an overdensity close to the appropriate value. Since the first cluster had a mass at $z = 0$ of $M_{200} = 4 \times 10^{14} \, h^{-1} \, M_{\odot}$, which is rather low for a 'rich' cluster, we chose the second sphere on the list, with a mass of $M_{200} = 8 \times 10^{14} \, h^{-1} \, M_{\odot}$ at $z = 0$. This cluster is relatively isolated, with no material in the $25 \, h^{-1}$ Mpc belonging to another rich cluster.

Having selected the centre for the $+2\sigma$ sphere, and having made candidate lists of centres for the four other spheres, we selected our final five sphere centres from a list of all five centres generated as follows. We selected a sphere, at random, from each of the candidate lists of centres for the $(-2, -1, 0, +1)\sigma$ regions and added the chosen $+2\sigma$ sphere centre to form a quintet of candidate centres. Each quintet was considered eligible for the final list if it satisfied two criteria: (i) none of the five centres had one or more of their x, y or z coordinates within $50 \, h^{-1}$ Mpc of the periodic boundary of the simulation volume and (ii) each centre was at least $200 \, h^{-1}$ Mpc away from the other four sphere centres in the quintet. The first criterion was imposed to maintain the same coordinate system as the Millennium Simulation, and also allow the use of an intermediate-resolution particle-mesh (PM) force algorithm implemented within our simulation code, which requires that the isolated mesh used for this calculation should not cross the periodic boundaries of the simulation volume. The second criterion was chosen to prevent two or more spheres from being near neighbours, and to force the

five spheres to extend over a significant range of the Millennium Simulation volume, so as to be relatively independent of each other. It is possible to find many quintets satisfying these conditions; we chose the first quintet generated by the code that used a pseudo-random number generator to search for candidates. The centres and radii of all five spheres are presented in Table 1.

In order to maximize the load and memory efficiency of the intermediate-range PM algorithm in our code, which automatically resizes to encompass all high-resolution particles, we chose to make five separate sets of initial conditions, one for each of the regions chosen; this is because the inclusion of all five spheres within a single set of initial conditions would create an isolated mesh filling most of the simulation volume. Such a mesh would incur a significant memory cost and offer little speed up relative to the main PM algorithm.

A1.2 Generating the particle load and applying the displacement field

The procedure for making the initial conditions is identical for all five regions and consists of two main stages: (i) the creation of a multimass particle distribution to represent the uniform, unperturbed particle distribution and (ii) the recreation of the displacement field used by the Millennium Simulation, with the addition of extra short wavelength power in the resimulated regions, sampled from the same power spectrum, down to the appropriate Nyquist frequency. The multimass uniform particle distribution is based upon a cubic mesh of 320^3 cells that encloses the entire simulation volume. Each cell of the mesh is treated independently. Particles, which at this point are composite and represent both dark matter and gas, are placed in each cell so that the mean density of all cells is equal to the matter density of the Universe. A total of n^3 equal mass particles is placed in each cell, arranged as a cubic grid with a spacing of $1/n$ of the cell size and with the centre of mass of all the particles coinciding with the cell centre. The value of n varied from 32, for the high-resolution region encompassing the sphere itself, to unity for cells that are distant from the sphere and are needed only to provide the correct tidal forces on the region of interest.

The value of n, as a function of cell position, was determined by the following algorithm: all of the particles from the Millennium Simulation inside or within $0.5 \, h^{-1}$ Mpc ($1 \, h^{-1}$ Mpc for the $+2\sigma$ sphere) of a particular GIMIC sphere at $z = 1.5$ were traced back to the grid cells from which they originated at very high redshift. Any cell occupied by one or more of these particles is labelled as a high-resolution cell. A check is made to determine if the high-resolution cells form a single contiguous region, where contiguous cells are defined as those sharing a face. If there were non-contiguous cells present then any cell contiguous to a high-resolution cell is additionally defined as high-resolution cells, and this process is continued until all high-resolution cells are contiguous. A set of boundary layers is then defined as follows: level 1 boundary cells are defined as those cells adjacent (i.e. sharing a face, an edge or a vertex) to a high-resolution cell. Level 2 boundary cells are then adjacent to level 1 boundary cells, and so on until level 9, which in addition includes all cells that have not yet been assigned.

We then realize the zoomed initial conditions at two resolutions; since the original Millennium Simulation represents our low-resolution basis, we refer to the two GIMIC realizations as high and intermediate resolution. For the high-resolution realization, we used $n = 32$ in the high-resolution cells, while the neighbouring boundary cells had $n = (20, 15, 10, 8, 6, 4, 3, 2, 1)$ for the boundary levels 1–9, respectively. This yields a mass for the composite

(i.e. dark matter and gas) particles in the high-resolution region of $8.08 \times 10^6 \, h^{-1} \, \mathrm{M_\odot}$. For the intermediate-resolution realization, we effectively halve the value of n in each level (odd values are rounded up), yielding initial conditions with composite particle masses a factor of 8 greater than the high-resolution case, $6.46 \times 10^7 \, h^{-1} \, \mathrm{M_\odot}$. We therefore resolve haloes of mass $10^{12} \, h^{-1} \, \mathrm{M_\odot}$ (comparable to the Milky Way's halo) with $\sim$19 000 particles at intermediate-resolution and $\sim$150 000 at high-resolution. The total number of particles used in each simulation is presented in Table 1.

Having created the particle distribution, a displacement field was calculated for each particle using the techniques for making resimulation initial conditions described in Power et al. (2003) and Navarro et al. (2004). The displacement consists of two parts; on large scales, the same mode amplitudes and phases were used as in the Millennium Simulation, whilst short wavelength power down to the Nyquist frequency was added in a cubic region of dimension 50–$76 \, h^{-1} \, \mathrm{Mpc}$ containing all of the high-resolution cells. The break between the long-wavelength and short-wavelength power was chosen such that only wavenumbers $\boldsymbol{k} = (k_x, k_y, k_z)$ satisfying $\max(|k_x|, |k_y|, |k_z|) < \pi \, h \, \mathrm{Mpc^{-1}}$ were retained from the Millennium Simulation. This scale is sufficient to ensure that haloes resolved with at least a few hundred particles in the Millennium Simulation are present in the GIMIC resimulation. A large mesh of 2048^3 cells was used for all Fourier transforms to minimize interpolation errors.

Finally, the high-resolution composite particles were split into gas and dark matter particles, with masses appropriate to yield a cosmic baryon fraction of $\Omega_b/\Omega_m = 0.045/0.25 = 18$ per cent. The gas and dark matter particles were offset in each dimension from their common centre of mass by a mass-weighted fraction of one interparticle separation.

A2 Sampling and combining regions

The use of zoomed initial conditions complicates the interpretation of simulations because of the presence of an artificial boundary around the high-resolution region. Inside it, baryons are modelled explicitly with SPH particles; outside it, dark matter and baryons are represented by composite particles whose dynamics are collisionless. Clearly, the collisionless region should be excluded from any analysis, but the existence of the boundary has other indirect effects. For example, the absence of gas beyond the boundary artificially reduces the mass of gas that might accrete on to peripheral haloes. In addition, the lack of external pressure allows the easy escape of galactic outflows into the vacuum of the boundary region.

In the analyses presented in this paper that measure properties on a particle-by-particle basis, we consider only those particles that are within $18 \, h^{-1} \, \mathrm{Mpc}$ (comoving; $25 \, h^{-1} \, \mathrm{Mpc}$ for the $+2\sigma$ region) of the centre of mass of all baryonic and high-resolution dark matter particles; in the case of analyses where we consider the properties of haloes or galaxies, we consider those whose centre of mass is within this volume. For $z \gtrsim 1.5$, this automatically leaves a small boundary layer of SPH and high-resolution dark matter particles external to the sphere for 'safety', since a padding region was included in the initial conditions by design. However, at lower redshifts the morphology and comoving volume of each region deviates slightly from that of the sphere that defined them at $z = 1.5$. This is due to the varying overdensity of each region, since the underdense regions expand and the overdense regions contract in comoving space. However, we have explicitly checked that haloes towards the edge of the volume do not differ significantly from those at the centre, for instance in terms of their baryon fraction, SFR or metallicity, and so we do not consider this an issue for the analyses we present. In addition,
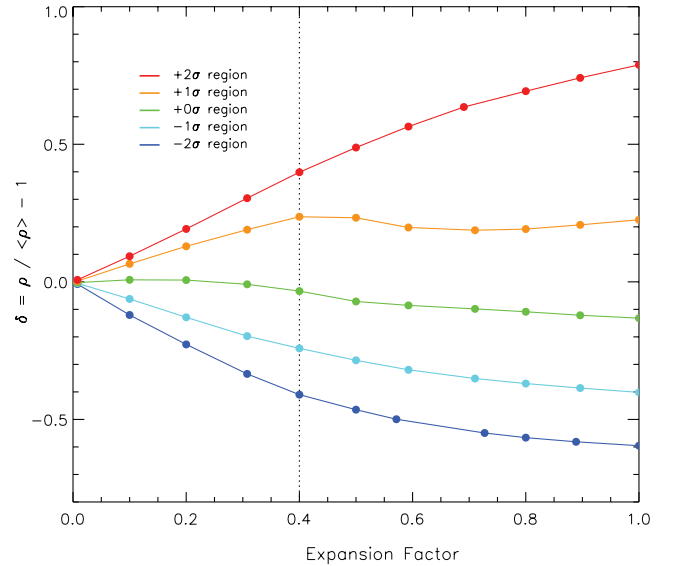
Figure A1. The overdensity evolution of the five spherical regions at intermediate resolution. The vertical dotted line denotes the epoch at which the regions were selected ($z = 1.5$). Note the drop, after $z = 1.5$, in overdensity of the $+1\sigma$ region; this results from a massive halo with high peculiar velocity leaving the spherical region.

the use of a simple and time-invariant sampling geometry (a sphere of constant comoving radius) allows various quantities, such as the SFRD $\dot{\rho}_\star$, to be normalized by volume in an unambiguous fashion.

In Fig. A1, we present the overdensity evolution of the spherical regions within each simulation. At early times, the regions necessarily have the mean density of the Universe, but their enclosed overdensities clearly diverge as they evolve. This figure highlights the advantage of zoomed initial conditions as a means to probe varying environments, since a periodic volume maintains $\delta = 0$ at all epochs. The dip in the overdensity of the $+1\sigma$ region at $a \simeq 0.5$ ($z \simeq 1$) results from the departure of a high-mass halo with a large peculiar velocity from the spherical 'analysis' region.

At $z = 1.5$, the distribution of overdensities over spherical regions of size comparable to the GIMIC regions is close to Gaussian. Ideally, to obtain the universal average of a quantity which is a function of overdensity, it is necessary to integrate over overdensity with a Gaussian weighting. Because the GIMIC simulation provides only five overdensities. $(-2, -1, 0, +1, +2)\sigma$, we approximate the integral as a sum and choose the weights for the five regions so that the summation correctly integrates any odd polynomial of overdensity

Table A1. Numerical weights applied to each GIMIC region when extrapolating statistics to the whole $(500 \, h^{-1} \, \mathrm{Mpc})^3$ Millennium Simulation volume.

Sphere	Summation weight	Volume weight	Net factor
-2σ	1/12	1	1/12
-1σ	1/6	1	1/6
0σ	1/2	1	1/2
$+1\sigma$	1/6	1	1/6
$+2\sigma$	1/12	$(18/25)^3$	486/15625

Note. The net factor (final column) is the product of the summation weight assigned to each region (first column) and the volume weighting (second column). The latter is applied to account for the larger volume of the $+2\sigma$ region.
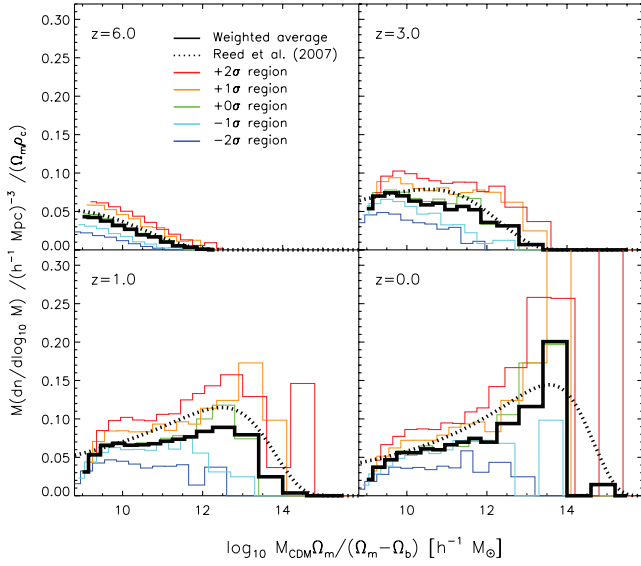
Figure A2. The halo multiplicity function in the five regions at intermediate resolution. The black histogram is the halo multiplicity function yielded by our weighting procedure used to generate estimates for the Millennium Simulation volume from the GIMIC regions. For comparison, we show the multiplicity function derived from the Reed et al. (2007) mass function (dotted line).

and even polynomials up to and including a quartic power in over-density. These weights are given in the second column of Table A1. An additional complication is that the $+2\sigma$ region has a different size, so an additional weight is needed to compensate for the larger volume (Table A1, third column). The actual weights are given in the fourth column Table A1.

These weights can only be approximate because the overdensity distributions are not exactly Gaussian, and strictly the regions only correspond to $(-2, -1, 0, +1, +2)\sigma$ deviations from the mean at $z = 1.5$. In Fig. A2, we show the *halo multiplicity function*, defined as the fraction of all mass bound into haloes of mass in the interval $(\log M, \log M + \mathrm{d} \log M)$. This description of the halo population is attractive since it has a smaller dynamic range than the halo mass function, and so more clearly highlights the *differences* between the halo populations of the GIMIC regions. Inspection of this plot shows that our weighted average (black histogram) constructed from the individual regions (coloured histograms) recovers the halo multiplicity function (and therefore the halo mass function also) of the Millennium Simulation (dotted line) reasonably well at all times. It is this fact that ultimately justifies our weighting scheme.

This paper has been typeset from a TEX/LATEX file prepared by the author.