Galaxy properties and the cosmic web in simulations

Ofer Metuki,[1][*] Noam I. Libeskind,[2] Yehuda Hoffman,[1] Robert A. Crain[3] and Tom Theuns[4,5]

[1]*Racah Institute of Physics, Hebrew University, Jerusalem 91904, Israel*
[2]*Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, D-14482 Potsdam, Germany*
[3]*Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands*
[4]*Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK*
[5]*Department of Physics, University of Antwerp, Campus Groenenborger, Groenenborgerlaan 171, B-2020 Antwerp, Belgium*

ABSTRACT

We seek to understand the relationship between galaxy properties and their local environment, which calls for a proper formulation of the notion of environment. We analyse the Galaxies-Intergalactic Medium Interaction Calculation suite of cosmological hydrodynamical simulations within the framework of the cosmic web as formulated by Hoffman et al., focusing on properties of simulated dark matter haloes and luminous galaxies with respect to voids, sheets, filaments, and knots – the four elements of the cosmic web. We find that the mass functions of haloes depend on environment, which drives other environmental dependence of galaxy formation. The web shapes the halo mass function, and through the strong dependence of the galaxy properties on the mass of their host haloes, it also shapes the galaxy-(web) environment dependence.

Key words: galaxies: evolution – galaxies: haloes – large-scale structure of Universe.

1 INTRODUCTION

Galaxy formation is one of the main research foci in modern cosmology and one of the most challenging. According to the standard cosmological paradigm, the initial density field is nearly perfectly homogeneous, peppered with small perturbations that cause structures to form in the Universe. These perturbations grow via gravitational instability, creating a filamentary network known as the cosmic web (Bond, Kofman & Pogosyan 1996), delineated by geometrically distinct components termed voids, sheets, filaments, and knots. The cosmic web is the most striking manifestation of gravitational collapse on the scales of a few megaparsecs and above, and has been known since the advent of numerical simulations (Davis et al. 1985; White et al. 1987) and observed in large sky surveys, such as the CfA Redshift Survey (Geller & Huchra 1989).

On the submegaparsec scale, dark matter (DM) collapses into virialized structures called haloes. If a DM halo's potential well is deep enough, it may act as a site of galaxy formation, enabling the efficient cooling of gas and the formation of stars. That certain galaxy properties (such as luminosity, colour, morphology, etc.) correlate with the mass of their host halo is a fairly well-established feature of galaxy formation (Peacock & Smith 2000; Berlind & Weinberg 2002; Moster et al. 2010). However, it is not clear what role environment plays in the process of galaxy formation. Dressler (1980) has shown that galaxy morphology correlates with the galaxy density in their neighbourhood. More recent studies have shown correlations between the environment and both galactic properties (e.g. luminosities, surface brightnesses, colours, and profile shapes; see Blanton et al. 2005) as well as host halo properties (e.g. mass, spin parameter, shape, and mass assembly histories; see Avila-Reese et al. 2005; Gao, Springel & White 2005; Maulbetsch et al. 2007). Bahé et al. (2012) have also shown a correlation for satellite galaxies. These studies suggest that although galaxy formation occurs deep within a halo's potential well, the process may be significantly influenced by larger scale phenomena. However, these studies adopt different definitions of environment, and it is therefore necessary to scrutinize this definition. In observational studies, projected galactic number density is often used to define environment (a good overview of such methods can be found in Muldrew et al. 2012). In simulations, one may use additional measures such as the DM density.

Definitions of environment (e.g. Aragón-Calvo et al. 2007; Sousbie et al. 2008; Bond, Strauss & Cen 2010a,b) often do not employ kinematical data, since these are challenging if not impossible to directly measure in observational samples. However, within the past few years, a number of dynamical techniques have been developed for the analysis of simulations (e.g. Hahn et al. 2007; Forero-Romero et al. 2009; Hoffman et al. 2012; Cautun, van de Weygaert & Jones 2013). Hahn et al. (2007) and Forero-Romero et al. (2009) employed the tidal tensor, defined as the Hessian of the potential, to identify the cosmic web in cosmological

[*] E-mail: metuki@phys.huji.ac.il

simulations. More recently, Hoffman et al. (2012) used the velocity shear tensor, as an extension to the tidal tensor, to define the cosmic web. This so-called V-web is able to probe the non-linear regime deep into the submegaparsec scale. A growing body of literature (Altay, Colberg & Croft 2006; Aragón-Calvo et al. 2007; Wang et al. 2011; Codis et al. 2012; Libeskind et al. 2012, 2013; Cautun et al. 2014; Neyrinck et al. 2014) has employed these 'web finders' to unveil properties of the DM haloes that populate the cosmos.

Many models of galaxy formation implicitly limit the effect the environment may have. For example, halo occupation models (e.g. Conroy, Wechsler & Kravtsov 2006; van den Bosch et al. 2007) assume that galaxy properties are a function solely of the host halo's mass. More sophisticated semi-analytical models (SAMs; e.g. Kauffmann, White & Guiderdoni 1993; Cole et al. 1994, 2000) often assume that galaxy properties are determined uniquely by a halo's merger history. Cosmological hydrodynamical simulations implicitly include the environmental effects on galaxy formation but progress along this avenue is hindered by their high computational cost and uncertainties in the subgrid physics. Yet, progress has been made with the emergence of some recent simulations in which the limited dynamical range has been shifted to scales large enough so as to probe the environmental dependence of galaxy formation. This is the case of the Galaxies-Intergalactic Medium Interaction Calculation (GIMIC; Crain et al. 2009) suite of simulations, which are used here.

In this paper, we use the GIMIC suite of simulations to investigate how properties of galaxies differ across the cosmic web. The GIMIC suite of simulations has been used to study the properties of the hot, X-ray luminous circumgalactic medium (Crain et al. 2010a,b), the orientation of satellite galaxies (Deason et al. 2011), stellar haloes (Font et al. 2011; McCarthy et al. 2012a), star formation rates (SFRs; McCarthy et al. 2012b), the origin of discs (Sales et al. 2012), and metal abundance in the circumgalactic medium (Crain et al. 2013), and is uniquely suited for our purposes as it combines a variety of environments (in the sense of the mean density) of moderate volume (which can be combined to statistically represent a much larger volume) with high resolution down to galactic scales, and (for the intermediate resolution) has been run all the way to redshift zero. The V-web algorithm (Hoffman et al. 2012) is used to define the cosmic web of the GIMIC simulations.

The paper is organized as follows: Section 2 describes the simulations and the cosmic web formalism and the results are presented in Section 3. Section 4 concludes the paper with a discussion of the results and their implications.

2 METHODS

2.1 Simulations

The GIMIC is a gas-dynamical re-simulation of five roughly spherical regions drawn from the Millennium simulation (Springel et al. 2005 – a DM only cosmological simulation in a periodic box of side length $L_{\mathrm{box}} = 500 \, h^{-1}$ Mpc). We give a short description of the simulations, focusing only on the relevant aspects, and refer the interested reader to the thorough explanations provided in Crain et al. (2009).

The five re-simulated regions were picked to sample different environmental densities: each has a mean overdensity at $z = 1.5$ that deviates from the cosmic mean by $[-2, -1, 0, 1, 2]\sigma$, where σ is the rms mass fluctuation in a sphere of radius $\sim 20 \, h^{-1}$ Mpc.

'Zoomed' initial conditions are constructed and tested to ensure that few low-resolution ('tidal') particles end up within the high-resolution region. Gas particles are added to the initial conditions and the hydrodynamics are computed using smoothed particle hydrodynamics. This is done at various resolution levels for the purpose of convergence testing: at the highest resolution, the computational cost prevented the entire suite from being run to $z = 0$. As we are interested in the present epoch, we use the intermediate resolution simulations that were run to $z = 0$. The mass resolution for these runs is $M_{\mathrm{DM}} \sim 5.3 \times 10^7$ and $M_{\mathrm{gas}} \sim 1.16 \times 10^7 \, h^{-1} \, \mathrm{M}_{\odot}$. A fixed comoving softening length is used for $z > 3$. Below this, the softening is fixed in physical space at $1 \, h^{-1}$ kpc. Convergence tests on the -2σ simulation (which has been run to $z = 0$ at a higher resolution) demonstrate that for haloes of $M_{\mathrm{vir}} \geq 10^{10} \, \mathrm{M}_{\odot}$ results are converged (Crain et al. 2009), and for that reason we limit ourselves to haloes in that mass range.

The simulations were performed using a modified version of the GADGET3 simulation code (last described in Springel 2005), with modules to follow radiative gas cooling and reionization, star formation, kinetic supernova and stellar winds feedback, and chemical abundances. No active galactic nucleus (AGN) feedback is included. The same cosmology as employed by the Millennium simulation (which is consistent with *Wilkinson Microwave Anisotropy Probe* 1 parameters) is assumed: $\Omega_{\mathrm{m}} = 0.25$, $\Omega_{\Lambda} = 0.75$, $\Omega_{\mathrm{b}} = 0.0045$, $n_{\mathrm{S}} = 1$, $\sigma_8 = 0.9$, $H_0 = 100 \, h$ km s^{-1} Mpc^{-1}, where $h = 0.73$. Only the $z = 0$ snapshot of the five simulations is used in this paper.

2.1.1 Gas-dynamics, star formation, and feedback

The baryonic physics that GIMIC models includes the following processes.

(i) Gas cooling and photoionization were implemented as described in Wiersma, Schaye & Smith (2009) – a redshift dependent, spatially uniform ionizing background is assumed; the cooling rate is computed as a function of redshift, gas density, temperature, and composition.

(ii) Quiescent star formation and feedback were implemented by imposing a polytropic equation of state-to-gas particles with density $n_{\mathrm{H}} > 10^{-1}$ cm^{-3}. These particles are eligible for star formation, and are stochastically converted into star particles with a probability determined by the local gas pressure (Schaye & Dalla Vecchia 2008).

(iii) Kinetic feedback was implemented by stochastically imparting randomly oriented velocity kicks to, on average, $\eta = \dot{m}_{\mathrm{wind}}/\dot{m}_{\star} = 4$ neighbouring gas particles of newly formed stars. Kicks of 600 km s^{-1} were adopted, and implemented 3×10^7 yr after the formation of the star particle. This implementation is described by Dalla Vecchia & Schaye (2008).

(iv) Chemodynamics – star particles inherit their abundances in 11 elements (hydrogen, helium, carbon, nitrogen, oxygen, neon, magnesium, silicon, sulphur, calcium, and iron) from their parent gas particles. An initial mass function (Chabrier 2003, with a stellar mass range of 0.1–100 M$_{\odot}$) and stellar evolution tracks (dependant on metal abundance; Portinari, Chiosi & Bressan 1998; Marigo 2001; Thielemann et al. 2003) are assumed, and the delayed release of these elements (Wiersma et al. 2009) from supernovae and asymptotic giant branch stars is then followed as they are being distributed to neighbouring gas particles. Solar abundances are based on CLOUDY.

The lack of AGN feedback and the simplified wind model results in a luminosity function which does not properly match observations at high masses (see Crain et al 2009). However, this mismatch does

not affect our main point of interest, namely the relative properties across different environments.

2.1.2 DM halo and galaxy identification

DM haloes are identified using the publicly available Amiga Halo Finder (AHF; Knollmann & Knebe 2009), which locates local overdensities in an adaptively smoothed density field as prospective halo centres [the global mass function of haloes has been verified to be consistent with the that derived by Crain et al. (2009), using the SUBFIND algorithm; the 'halo finder comparison project' provides a thorough discussion on the comparisons of structure finders of haloes (Knebe et al. 2011, 2013b), subhaloes (Onions et al. 2012), and galaxies (Knebe et al. 2013a)] . The potential of each density peak is then calculated and bound particles are retained as halo members. AHF implicitly locates both parent and subhaloes, simultaneously ('parents' are haloes whose centre does not fall within the virial radius of any other halo; subhaloes are haloes contained within the virial radius of another halo). The virial radius is defined in the standard way (namely, the distance at which the mean interior density falls below $200\rho_{back}$, where ρ_{back} is the mean density in the simulation). The radius of a subhalo cannot be defined in this manner since the overdensity may only drop below the given threshold outside the parent halo, and so its extent is defined as the point at which the density profile starts rising, effectively signalling the transition from the subhalo to the main halo.

Following Sales et al. (2012), we identify galaxies by examining the particles within 15 per cent of r_{halo}, where r_{halo} is the parent or subhalo radius. This value is used to ensure that most of the stellar mass in the halo is inside the galaxy (see Kravtsov 2013). Within a given halo, the fraction of star particles that reside within $0.15r_{halo}$ depends on the host halo mass. For example, for high-mass haloes ($M_{halo} > 10^{12} \, h^{-1} \, M_{\odot}$), around 90 per cent of all halo star particles reside within $0.15r_{halo}$; this drops to around 50 per cent for low-mass haloes ($M_{halo} \sim 10^{10} \, h^{-1} \, M_{\odot}$, which is our resolution limit for haloes). In order to ensure that our results are not dependent on this choice, we have run our analysis using 0.1, 0.15, and 0.2 of r_{halo}, and found no differences in our main results, explained in Section 3. All baryonic particles within this distance define the 'galaxy' if the number of star particles exceeds 10. Such a galaxy finder yields 19 081 central galaxies (with $M_{vir} > 10^{10} \, h^{-1} \, M_{\odot}$), 733 of which contain satellites. The total number of satellite galaxies in the simulations is 2876. Haloes containing galaxies are defined here as 'luminous haloes'. Groups of galaxies are defined as all galaxies hosted within the same parent halo.

Care needs to be exercised when obtaining global averages from the various co-added GIMIC simulations, owing to the prevalence and size of each region. Thus, when combining results from the different simulations, each region is weighted according to appendix A2 of Crain et al. (2009); specifically haloes and galaxies are added using the weights that appear in table A1 of that appendix.

2.2 The web classification algorithm

We use the V-web method (Hoffman et al. 2012) to identify the four elements of the cosmic web in the GIMIC simulations, and through this associate haloes and galaxies with their web environment. A cubic subvolume of side length twice the radius of the re-simulated sphere at $z = 1.5$ ($2 \times 25 \, h^{-1}$ Mpc for the $+2\sigma$ simulation and $2 \times 18 \, h^{-1}$ Mpc for the rest) is identified. The density and velocity fields in this cubic volume are calculated on a 128^3 grid (see the

appendix for a discussion of grid sizes and convergence tests) using the clouds-in-cells (CIC) algorithm, resulting in a cell size of $390 \, h^{-1}$ kpc for the $\sigma = +2$ and $280 \, h^{-1}$ kpc for the rest of the simulations. The CIC has an inherent smoothing on the scale of two cells. However, an additional Gaussian smoothing with a kernel size of $(1, 2, 3, 4) \times$ the cell size of the $\sigma = +2$ simulation is also applied, so as to suppress the anisotropy induced by the Cartesian grid (see the appendix for a discussion of smoothing kernel choice). The reason we apply four different smoothing kernels to the same CIC is explained below. The (dimensionless) velocity shear is defined at each grid cell by

$$\Sigma_{\alpha\beta} = -\frac{1}{2H_0} \left(\frac{\partial v_\alpha}{\partial r_\beta} + \frac{\partial v_\beta}{\partial r_\alpha} \right), \tag{1}$$

where $\alpha, \beta = x, y, z$ and H_0 is the Hubble constant. It is calculated by means of fast Fourier transform (FFT) over the cubic regions. The effect of the periodic boundary conditions employed by the FFT has been tested against a variety of box sizes and found to be negligible, as is explained in the appendix. The shear tensor is diagonalized at each grid cell. Note a minus sign has been introduced in equation (1), so as to have a positive (negative) eigenvalue correspond to a collapse (expansion) along the direction of the corresponding eigenvector.

The cosmic web classification is done by counting the number of eigenvalues of the shear tensor at a particular point in the space (cell) that are above a given threshold. Following Hoffman et al. (2012) and Libeskind et al. (2012), this threshold is taken to be $\lambda_{th} = 0.44$, which best matched the density distribution visually. Thus, a cell is tagged as a void, sheet, filament, or knot if it has 0, 1, 2, or 3 eigenvalues greater than λ_{th}, respectively.

Haloes (and their galaxies) inherit the web classification from the cell within which their centres are embedded. The smoothing scale (and corresponding velocity field and shear tensor) assigned to each halo is chosen to be the smallest smoothing scale (of the four used) that is greater than a halo's virial radius.

3 RESULTS

The results presented in this section combine properties obtained for each simulation separately, as explained in Section 2.1. Care has been taken to ensure that for each web element, these properties do not differ significantly between the five simulations, and so these results are robust across different density environments.

In Fig. 1, we show a $0.5 \, h^{-1}$ Mpc thick slice through the centre of the $[-2, 0, +2]\sigma$ GIMIC regions, smoothed with a $390 \, h^{-1}$ kpc smoothing kernel. On the left, we show the logarithm of the density, with DM haloes overplotted (blue circles; circle radius scales with the radius of the halo), while on the right, we show the cosmic web, decomposed into voids, sheets, filaments, and knots with the galaxy distribution overplotted (purple–pink circles; the radius of the circles is proportional to the logarithm of the galaxy's stellar mass). That haloes and galaxies trace the matter distribution is expected and clearly seen in these plots.

The statistics which characterize the cosmic web are presented in Table 1. It includes the volume fractions, the DM and stellar mass fractions, and the fractions of haloes and galaxies in each web type, as well as the fractions of haloes above our resolution limit that contain galaxies in each web type (defined here as luminous haloes, as opposed to dark haloes that contain no luminous galaxies; for an examination of dark haloes in the GIMIC simulations see Sawala et al. 2013). The following conclusions follow. Haloes, galaxies, and the DM reside predominantly in filaments and sheets. Voids dominate
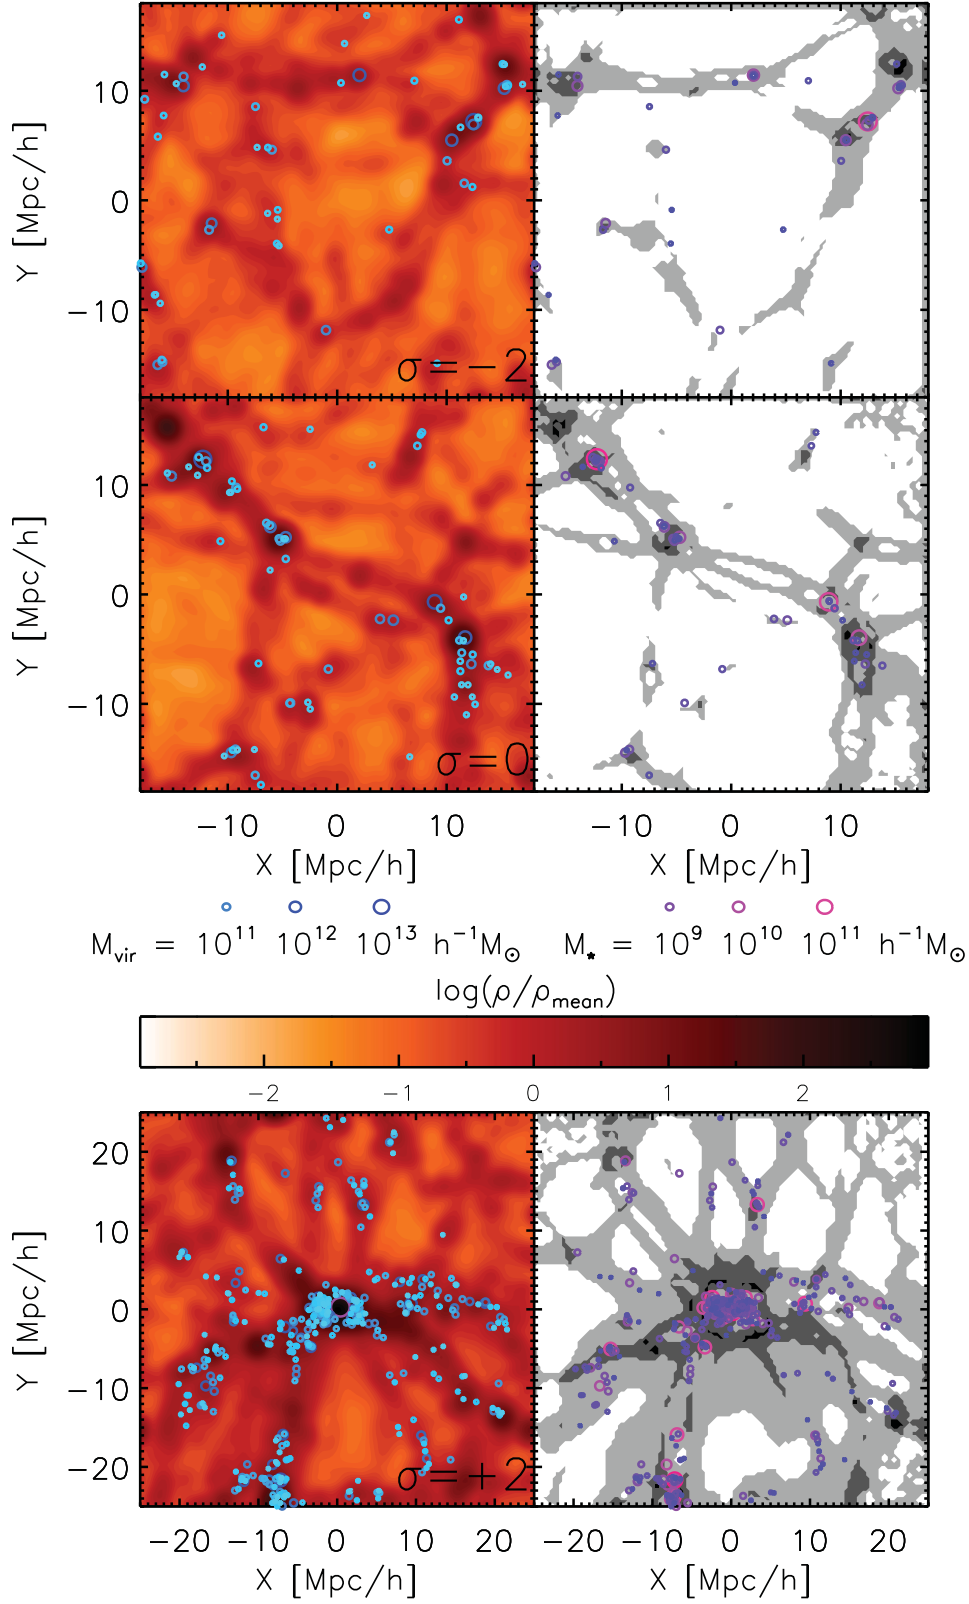
Figure 1. Slices of thickness $0.5\,h^{-1}$ Mpc of the $\sigma = [-2, 0, +2]$ simulations, smoothed with a $390\,h^{-1}$ kpc kernel. Left-hand panels – DM parent haloes (blue circles; symbol size is proportional to $\log(M_{\rm vir})$, as indicated by the legend) plotted on top of the logarithm of the DM density. The colour bar describes the logarithm of the density; all densities are scaled by the universal mean density. Right-hand panels – galaxies in the slice (purple–pink circles; symbol size is proportional $\log(M_\star)$) plotted on top of the cosmic web. The contours are coloured by web element (white for voids, light grey for sheets, dark grey for filaments, and black for knots).

Table 1. Statistics of different web elements – their volume filling fractions, the number of haloes and galaxies they host, and the DM and stellar masses they contain, each given as a fraction of the entire simulation. The bottom line gives the fraction of the total haloes (above our resolution limit of $10^{10}\,\mathrm{M_\odot}$) within each web element that contain galaxies (i.e. luminous haloes).

	Voids	Sheets	Filaments	Knots
Volume fraction	63.7 per cent	31.7 per cent	4.3 per cent	0.3 per cent
Halo no. fraction	15.9 per cent	51.5 per cent	28.8 per cent	3.8 per cent
DM mass fraction	16.2 per cent	35.4 per cent	33.0 per cent	15.4 per cent
Galaxy no. fraction	13.4 per cent	46.5 per cent	30.8 per cent	9.3 per cent
Stellar mass fraction	0.8 per cent	11.3 per cent	36.4 per cent	51.5 per cent
Fraction of haloes that are luminous	49.0 per cent	63.2 per cent	75.0 per cent	80.2 per cent

by volume, with a volume fraction of ∼64 per cent, yet contain less than 1 per cent of the stellar mass. These results agree with works using different web classification schemes, such as Aragón-Calvo et al. (2007) and Cautun et al. (2014). The number fraction of galaxies distributed by their web classification constitutes a fairly close proxy to that of haloes and to the DM mass fraction. This is not the case for the stellar mass whose web distribution varies significantly from that of the DM.

The fact that knots contain roughly half of all stellar mass available is a reflection of the following two properties. The first is inferred from the bottom row of Table 1, which shows the fraction of DM haloes above our resolution limit in each web type that host galaxies ('luminous haloes'). Approximately 80 per cent of haloes in knots have luminous galaxies, whereas less than a half of haloes in voids have luminous galaxies. The second is the different mass functions[1] (see Fig. 3).

The left- (right-)hand panel of Fig. 2 manifests this fact, by showing the fraction of haloes (galaxies) in each web element as a function of M_{vir} ($M_\star$). At low masses, most haloes and galaxies reside in sheets, at intermediate masses most of them are in filaments, and at high masses practically all of them reside in knots. This is in good agreement with the bottom panel of fig. 18 of Cautun et al. (2014).

It is important to understand how the different web classification selects haloes of different DM mass, galaxies of different stellar mass or luminosity, and so on. A number of these properties are shown in Figs 3(a)–(d).

In Fig. 3(a), we show the DM mass functions of parent haloes. The mass functions vary across the cosmic web, both in amplitude, shape, and (importantly) range. Sheets contain most of the lowest mass parent haloes, up to $3 \times 10^{11}\,h^{-1}\,\mathrm{M_\odot}$. Most haloes of mass between 3×10^{11} and $10^{13}\,h^{-1}\,\mathrm{M_\odot}$ are found in filaments, while the most massive haloes in the simulation ($M_{\mathrm{vir}} > 10^{13}\,h^{-1}\,\mathrm{M_\odot}$) are found almost exclusively in knots. Voids contain the smallest fraction of haloes with DM masses above $2 \times 10^{11}\,h^{-1}\,\mathrm{M_\odot}$. The general trend is similar to the ones seen in Aragón-Calvo et al. (2007) and Cautun et al. (2014), despite the fact that the web classification schemes used in those papers are different; this is not entirely surprising, as the most massive haloes are expected to reside in knots, while voids are expected to be mostly empty and definitely not contain any massive galaxies, and so on.

The subhalo mass function is shown in Fig. 3(b). From this figure, we infer that most subhaloes are embedded in knot-parents and the fewest subhaloes are embedded in void haloes (irrespective of subhalo mass). This is likely due to the fact that the parent knot halo mass function extends to higher masses and thus, above the given resolution limit, contain more substructures than all of the other web elements.

The right-hand panels of Fig. 3 show various galaxy stellar mass functions. In Fig. 3(c), the full galaxy stellar mass function is shown. It is similar to the halo mass function in the fact that knot environments host the most massive galaxies (where $M_\star \gtrsim 10^{12}\,h^{-1}\,\mathrm{M_\odot}$), with filaments and sheets dominating the environments of the lower mass bins ($10^{10} \lesssim M_\star/h^{-1}\,\mathrm{M_\odot} \lesssim 10^{12}$ and $M_\star \lesssim 10^{10}\,h^{-1}\,\mathrm{M_\odot}$, respectively). Note that the apparent 'bump' at stellar masses of $\gtrsim 10^{11}\,\mathrm{M_\odot}$ is driven by the lack of AGN feedback and the simplified wind model in the GIMIC subgrid physics. It is also interesting to note that in the lower mass end, all total mass functions (both DM and stellar) follow a power law with the same slope of -1.1, at least up to the point where the 'bump' in the stellar mass function begins.

The satellite stellar mass function, Fig. 3(d), is also broadly similar to its DM counterpart in showing a clear hierarchical trend for the stellar mass functions of the different web elements across the entire range of stellar masses. However, practically no satellite galaxies above our resolution limit inhabit void regions – this is likely because DM haloes in voids are on average less massive and thus have fewer resolvable substructures. This limits the number of both substructures and satellites.

This comparison between the left-hand column of Fig. 3, which describes the mass function of the DM haloes, with the right-hand column, which shows the various stellar mass functions, and their dependence on the web environment, leads to the inference that galaxy-environment dependence (right-hand column) is strongly driven by the halo mass-environment dependence (left-hand column). This seems even likelier when considering the fact that the trends do not change whether one plots the stellar mass functions of all haloes, groups of galaxies, or just the centrals.

In Fig. 4, we show the average number of subhaloes per host halo (a) and satellites per host galaxy (b) as a function of host mass, divided into web types. The shaded regions depict the standard deviation about the mean for all haloes/galaxies, while the error bars depict the standard error of the mean (i.e. the error involved in computing the mean; these errors were computed using a bootstrapping algorithm) about the mean for each web element. The lower panels show the fractional difference from the mean, $\Delta(N_{\mathrm{sub}})/\overline{N}_{\mathrm{sub}}$, for the different web elements. As can be seen, while there is some

---

[1] Throughout this paper, we use the term 'mass functions', plural, to denote the division into the four web elements.
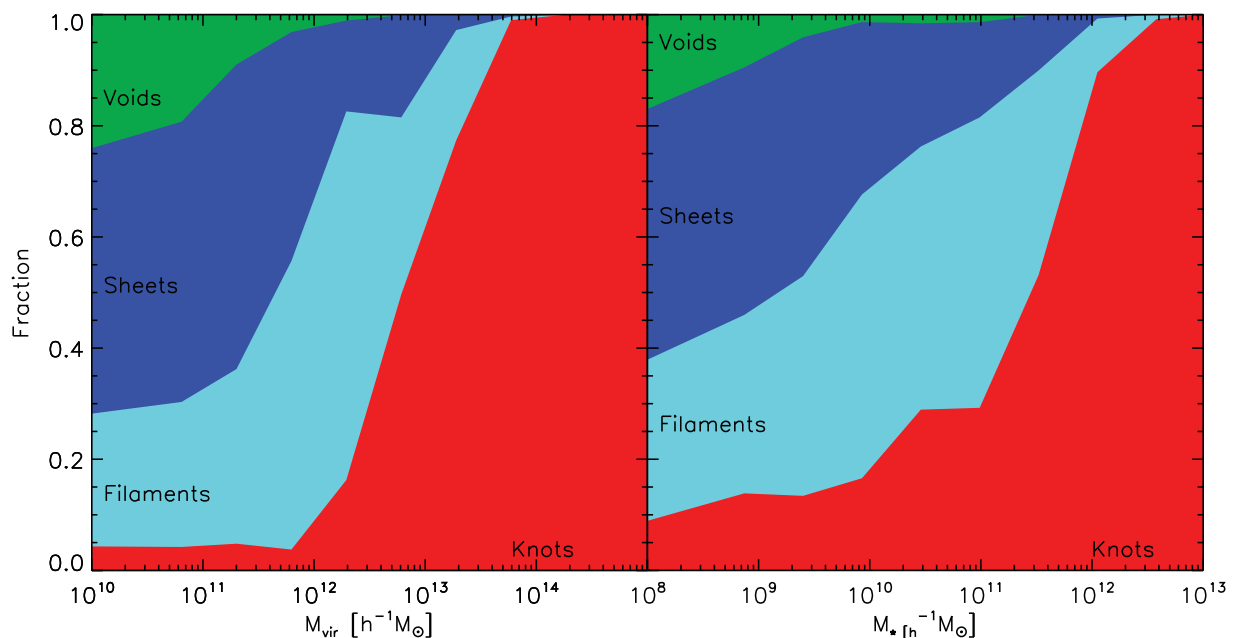
Figure 2. Left: division of the haloes between the different web elements as a function of halo mass. Right: division of the galaxies between the different web elements as a function of stellar mass. The coloured regions depict the different elements – red for knots, cyan for filaments, blue for sheets, and green for voids. The fractions are cumulative for each mass; i.e. the upper line of the knots gives the actual fraction of haloes/galaxies in knots at that mass, the upper line of the sheets gives the fraction in knots plus the fraction in sheets, etc.

difference in the number of satellite galaxies at a given galaxy stellar mass, it is small and within the standard error of the mean for almost all masses and web elements. For the number of subhaloes at a given halo mass, the difference is even smaller. The curves are almost entirely within the standard deviation around the mean.

In Fig. 5, we examine three properties of the baryonic component in luminous haloes as a function of the host halo's virial mass. The shaded regions depict the standard deviation around the mean for all haloes, while the error bars show the standard error of the mean for each web element.

In Fig. 5(a), we show the mean baryon mass fraction in haloes ($f_{bar} = M_{bar}/M_{vir}$), divided by the average baryon fraction in the simulation. The baryon mass fraction increases with host halo mass, in agreement with previous works (Somerville et al. 2008; McGaugh et al. 2010), but the differences between web elements for a given halo mass are very small. The baryon mass fraction never reaches the universal mean, in agreement with Crain et al. (2007).

In Fig. 5(b), we show f_{gas} – the gas baryon mass fraction (i.e. M_{gas}/M_{bar}). The gas fraction increases with halo mass for haloes in all web types. The increase with halo mass is in agreement with the recent work of Noh & McQuinn (2014). At low halo masses ($<10^{11} \, h^{-1} \, \mathrm{M_\odot}$), the gas fractions of the different web elements differ visibly, though still within a factor of a few. All curves lie within the standard deviation around the mean. In this mass range, knot haloes have a smaller gas fraction than filament haloes, which in turn have a smaller gas fraction than sheet and void haloes. This may be due to the low baryon masses at these host halo masses – as can be seen from Fig. 5(a), at these halo masses the baryon fraction is much lower than for higher halo masses, and so a slight change in star formation can mean a significant change in gas fraction.

In Fig. 5(c), we show $M_{\star, group}/M_{vir}$, the fraction of the virial mass that is locked up in galaxies. For low-mass haloes, the stellar mass component increases with host halo mass, and there is some difference between the different web elements, though within a factor of

2. A transitional mass between 10^{12} and $10^{13} \, h^{-1} \, \mathrm{M_\odot}$ exists, above which the fraction of halo mass in stars starts to decrease with host halo mass. This behaviour is in agreement with previous works (Yang, Mo & van den Bosch 2008; Bouché et al. 2010; McGaugh et al. 2010). These trends are irrespective of web type, and all curves are within the standard deviation around the mean.

Similar trends are visible across web types in all three plots, indicating again that, to first order, haloes behave similarly irrespective of web environment. Although the overall shapes of these plots agree with previous works, we do not stress the shapes of these plots too much, as they may be affected by the physics of the simulation. However, this should not change the underlying result – the lack of difference between the web elements for a given halo mass (with the exception of a small difference at small halo masses for some properties).

Is the ability to turn baryons into stars affected by the web environment? We investigate this possibility in Fig. 6, which shows the SFR divided by host halo mass as a function of lookback time for three host halo mass bins, and divided into web elements. Effectively, we calculate a (differential) histogram of ages for star particles that are bound to galaxies at $z = 0$. The SFR is the total mass of all star particles formed during a finite time interval divided by that time interval. The shaded regions depict the standard deviation around the mean for all galaxies, while the error bars show the standard error of the mean for each web element.

As can be seen, haloes of different masses have different star formation histories, the main difference being that high-mass haloes have much higher star formation at later times than intermediate- and low-mass haloes, with the peak of star formation between redshift 1 and 2. We do not stress the shapes of these plots, as they have been shown to be difficult to reconcile with observations (Weinmann et al. 2012). However, it is also seen that for high-mass haloes, the difference between the different web elements is very small at all times. Intermediate-mass haloes show slightly more difference for
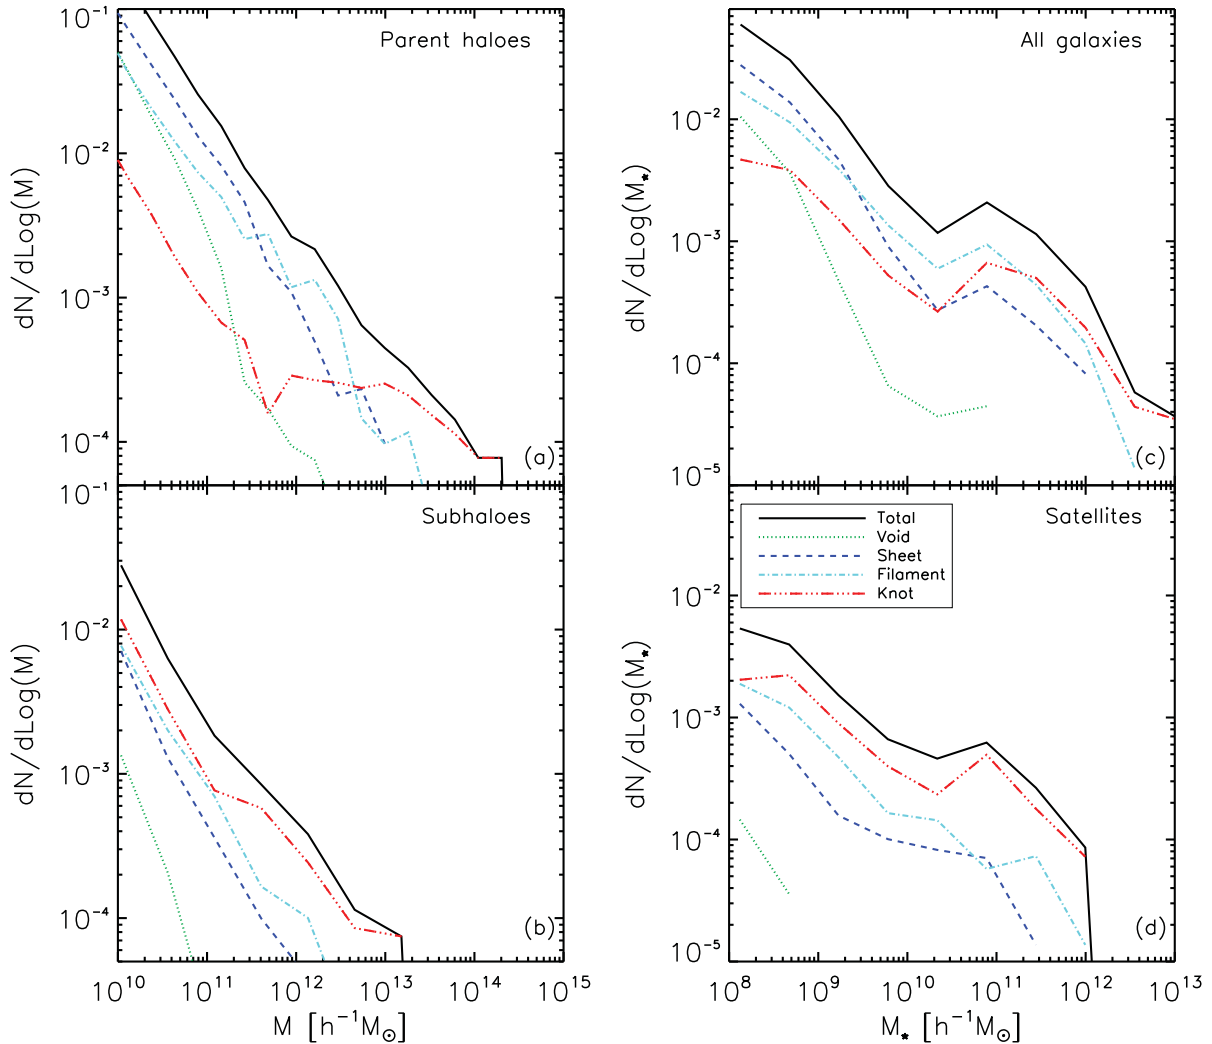
Figure 3. Left: the DM mass functions of all (top) and sub (bottom) haloes. Right: the stellar mass functions of all (top) and satellite (bottom) galaxies. Different line styles denote the total (solid black line), voids (green dotted line), sheets (blue dashed line), filaments (single dot–dashed cyan line), and knots (triple dot–dashed red line).

later epochs, but still within a factor of a few at most. Low-mass haloes show distinct differences at early times, and especially at the peak, but these are small (within a factor of 2). At all masses, the curves lie well within the standard deviation around the mean. The fact that the differences are small appears to be robust and independent of the shapes of the plots. Therefore, in agreement with Crain et al. (2009), we find that the main driver of the different SFRs is driven by its dependence on host halo mass, and through it an environmental dependence is induced.

4 DISCUSSION

One of the very basic components of any theory of galaxy formation is the dependence of at least some of the properties of galaxies on their ambient environment. Arguably, the morphology–density relation of Dressler (1980) is the best studied manifestation of such a dependence. This has led us to analyse the GIMIC set of simulations (Crain et al. 2009), a suite of state-of-the-art gas-dynamical simulations that include a myriad of physical phenomena. It is convenient to define the environment within which galaxies reside by the cosmic web, which consists of voids, sheets, filaments, and knots (Bond et al. 1996). The velocity shear tensor formulation (Hoffman

et al. 2012; Libeskind et al. 2012) has been used here to quantify the cosmic web. Our analysis consists of studying the properties and distribution of simulated DM haloes and luminous galaxies with respect to the cosmic web.

The main results of the paper are as follows.

(i) The DM halo mass function varies significantly with the web classification, namely the more massive a halo is the more likely it is to reside higher in the web sequence (voids, sheets, filaments, and knots). A similar trend is found when the sample of haloes is limited only to luminous ones, namely the ones who contain a galaxy.

(ii) The stellar mass functions of all galaxies vary significantly with the web elements, in a manner similar to the halo mass functions. Voids are populated mostly by faint galaxies, and moving up along the web sequence, sheets, filaments, and knots are progressively populated by more massive galaxies. A similar trend is found for the stellar mass functions of groups of galaxies and of central galaxies, where the group is defined by all galaxies embedded in a given parent halo.

(iii) Subhaloes show a more distinct trend with the web sequence, showing dominance according to web sequence in all subhalo masses. Satellite galaxies have a similar trend with stellar mass.
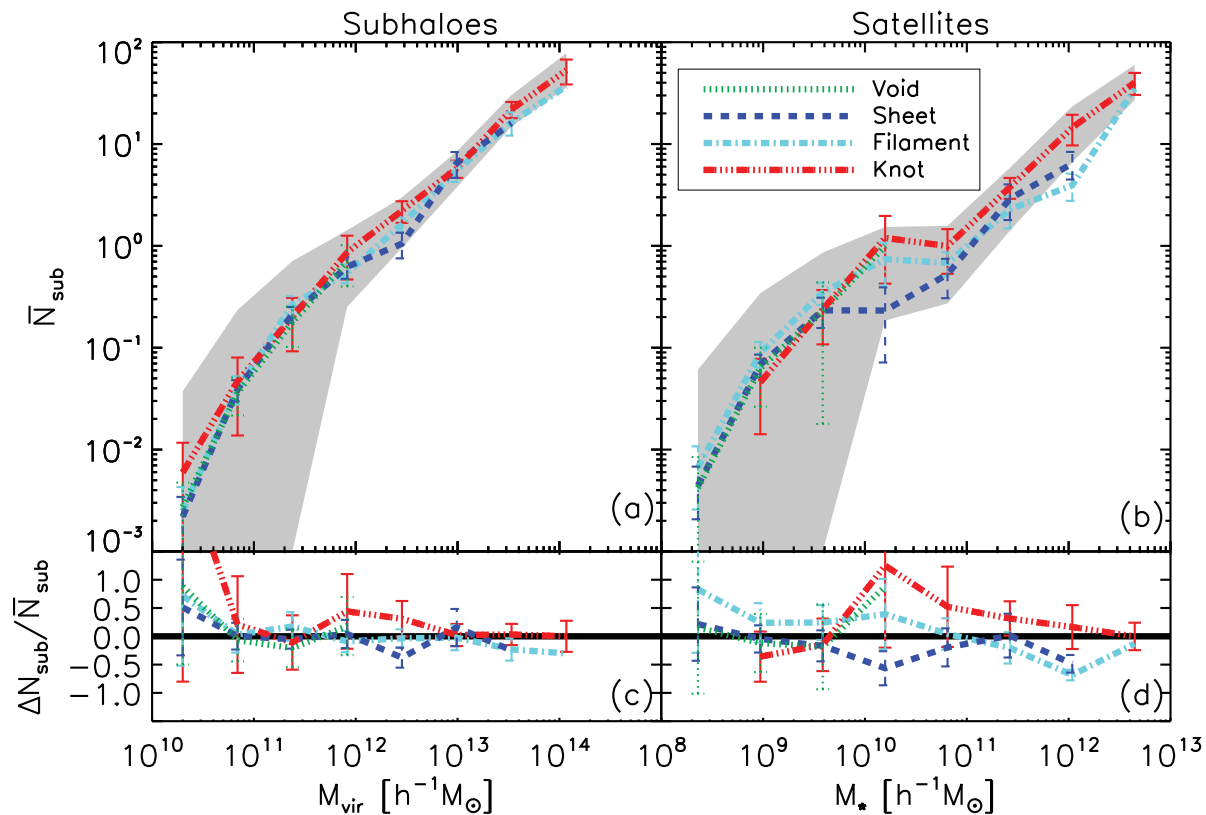
Figure 4. (a) Average number of subhaloes as a function of host halo virial mass. The shaded region depicts the standard deviation of the mean, the error bars depict the standard error of each web element. (b) Average number of satellite galaxies as a function of central galaxy stellar mass. The shaded region and error bars are as in (a). (c) The fractional difference from the mean number of subhaloes for each web element. (d) The fractional difference from the mean number of satellites for each web element. The average substructure numbers do not change significantly between web elements.

(iv) The number of satellites per galaxy of a given mass extends over more than three orders of magnitude for the range of haloes studied here, with a relatively small dependence on the web environment. The number of subhaloes per halo of a given mass shows an even smaller dependence on web environment.

(v) Considering the properties of baryons residing in DM haloes at a given mass range, these properties are broadly independent of the web classification. The baryon mass fraction, gas fraction, and stellar mass fraction in luminous haloes depend strongly on the virial mass of the halo, and show very little direct dependence on the web environment for most of the mass-range studies. Gas fractions at low halo masses do seem to differ noticeably between web elements; however, this is still within a factor of a few, and may be a result of the low baryon masses in these haloes. All curves lie within the standard deviation around the mean.

(vi) The SFR of central galaxies, divided by host halo mass, shows slight variation with respect to web type. For high-mass haloes, the different web types are practically indistinguishable within the errors; intermediate-mass haloes show some variation, but still within the standard errors; low-mass haloes show the most significant changes, but these are still small, a factor of 2 at the most. All curves lie well within the standard deviation around the mean. In agreement with the other 'baryonic' properties of galaxies residing in DM haloes, the difference seen between different mass bins is induced by the strong dependence of the star formation on the host halo virial mass.

The following understanding of the galaxy–environment relation emerges. There is a clear and strong dependence of the halo

mass function on the web classification. It has been long known that denser environments host more massive haloes, but this dependence has been quantified now in terms of the cosmic web classification that is adopted here. Galaxies are characterized here by a few global properties, such as gas and stellar mass fraction, stellar mass, and star formation. The distribution of haloes and galaxies with respect to these properties depends on the web type and this is best manifested by the dependence of the galaxy stellar mass function on the cosmic web attributes. Again, a very well known fact but quantified here within the framework of the cosmic web. Looking deeper into the web dependence, the distribution of the above properties of haloes and the galaxies that they host, within a given mass range, has been studied, again in the context of the cosmic web.

Our main finding is that as a first approximation the dependence of baryonic properties of galaxies on their ambient web environment is driven by the different mass function of the different web environments. One should note that it is not claimed here that environmental effects do not shape the properties of galaxies. Rather, to the extent that environment shapes the halo mass function, it also shapes the galaxy properties, but indirectly. The end product of such effects is that the halo mass is the main parameter through which the web environment affects processes related to galaxy formation.

The major conclusion of this paper is that the strong dependence of galaxy properties on the mass of their host halo is the leading effect that shapes the dependence of the galaxy properties on the cosmic web. For the most part, the baryonic parameters are indistinguishable within the standard errors, with the exception of gas fractions, stellar masses, and SFR/M_{vir} of low-mass haloes (10^{10}–$10^{11} h^{-1} M_\odot$). Even these are within a factor of a few, and we can
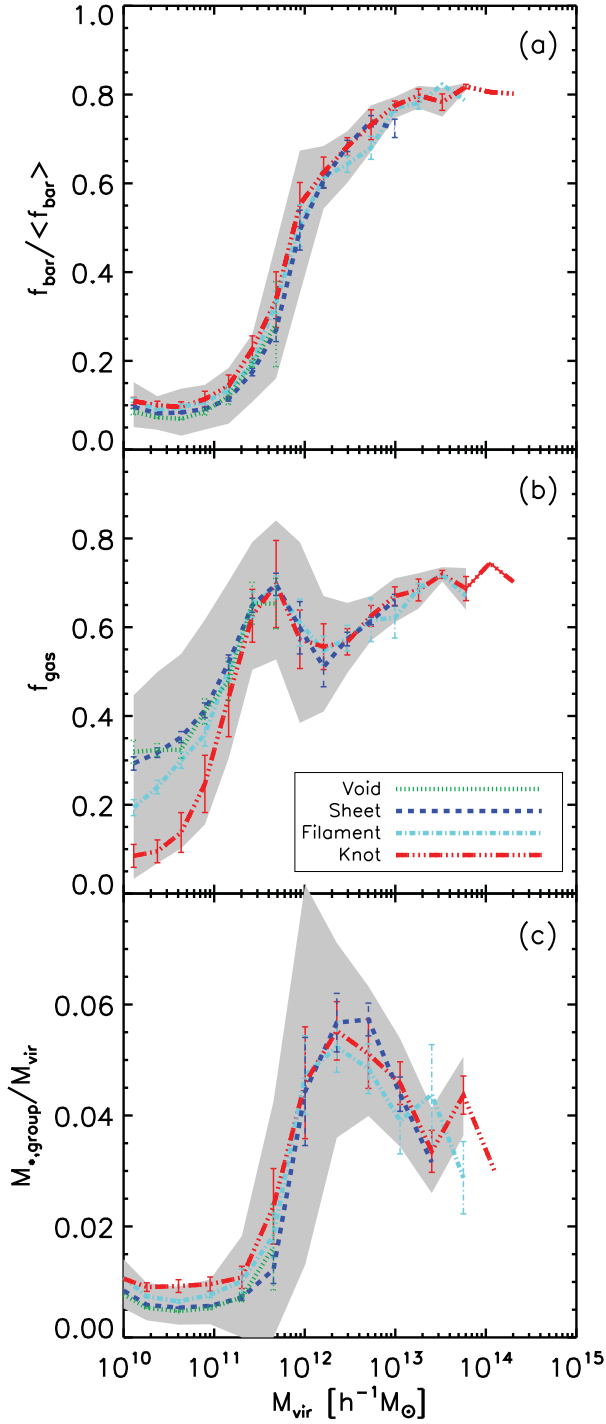
Figure 5. Three properties of the baryonic component of luminous haloes as a function of the halo's virial mass. Panel (a) shows the baryon mass fraction (normalized by the mean baryon mass fraction in the simulation); panel (b) presents the gas fraction of the baryonic component; panel (c) gives the total stellar mass in galaxies within the halo divided by the halo's virial mass. The shaded area depicts the standard deviation around the mean for all haloes, and the error bars are the standard error of the mean for each web element. All three plots show a distinctly small difference between the different web elements, with the exception of gas fractions at small host halo masses. All of the curves are within the 1σ range around the mean.
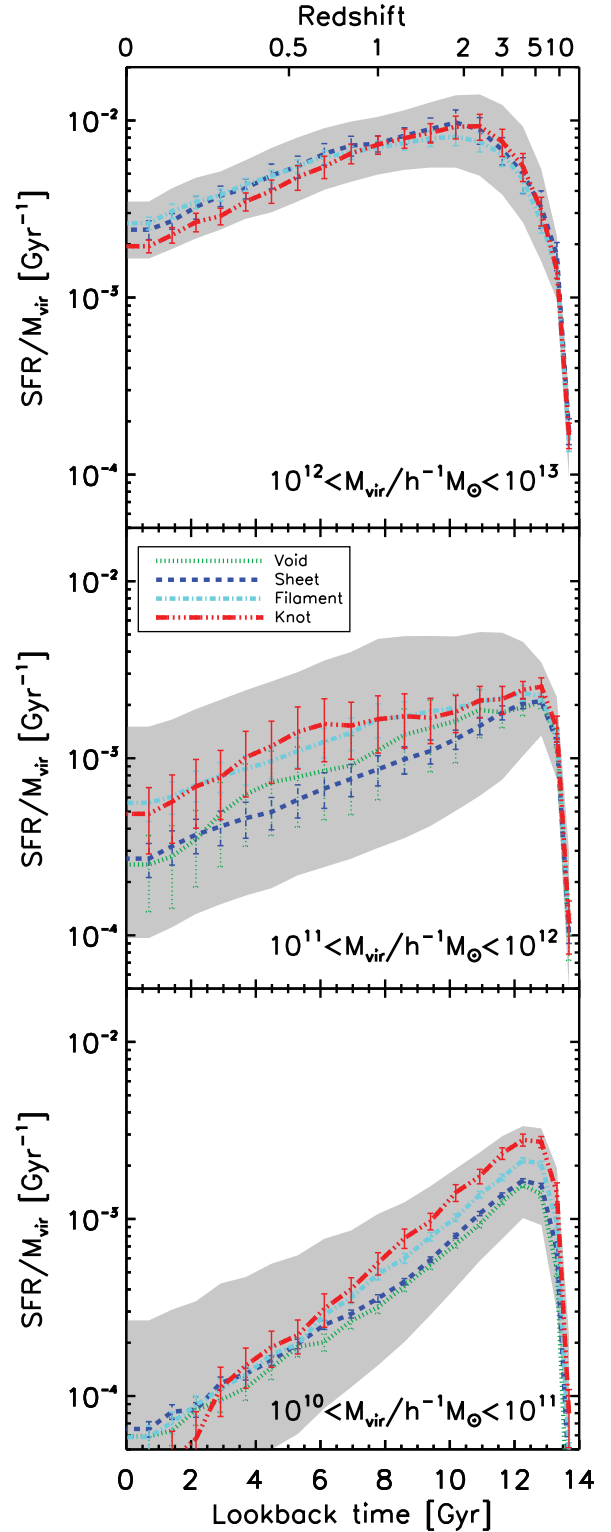
Figure 6. The SFR in central galaxies normalized by host halo mass as a function of lookback time in three different host halo mass bins (in units of $h^{-1}\,M_\odot$), divided into web elements. Shaded regions depict the standard deviation around the mean for all galaxies, while error bars show the standard error of the mean for each web element. While the behaviour changes between different mass bins, there is very little change between web elements in each mass bin; even at small masses, the plots lie outside of the standard errors of each other at early times, the change is within a factor of 2. All curves lie well within the standard deviation around the mean.

therefore conclude that for the GIMIC simulation the web environment affects the properties of galaxies only to a small degree, and that only for low-mass haloes.

This lends support to a wide range of SAMs (e.g. Cole et al. 2000; Bower et al. 2006) and statistical schemes of populating haloes with galaxies (e.g. CLF; van den Bosch et al. 2007), all of which assume that the galaxy properties are determined by the mass of a halo and its mass aggregation history.

ACKNOWLEDGEMENTS

REFERENCES

Altay G., Colberg J. M., Croft R. A. C., 2006, MNRAS, 370, 1422
Aragón-Calvo M. A., Jones B. J. T., van de Weygaert R., van der Hulst J. M., 2007, A&A, 474, 315
Avila-Reese V., Colín P., Gottlöber S., Firmani C., Maulbetsch C., 2005, ApJ, 634, 51
Bahé Y. M., McCarthy I. G., Crain R. A., Theuns T., 2012, MNRAS, 424, 1179
Berlind A. A., Weinberg D. H., 2002, ApJ, 575, 587
Blanton M. R., Eisenstein D., Hogg D. W., Schlegel D. J., Brinkmann J., 2005, ApJ, 629, 143
Bond J. R., Kofman L., Pogosyan D., 1996, Nature, 380, 603
Bond N. A., Strauss M. A., Cen R., 2010a, MNRAS, 406, 1609
Bond N. A., Strauss M. A., Cen R., 2010b, MNRAS, 409, 156
Bouché N. et al., 2010, ApJ, 718, 1001
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
Cautun M., van de Weygaert R., Jones B. J. T., 2013, MNRAS, 429, 1286
Cautun M., van de Weygaert R., Jones B. J. T., Frenk C. S., 2014, MNRAS, 441, 2923
Chabrier G., 2003, PASP, 115, 763
Codis S., Pichon C., Devriendt J., Slyz A., Pogosyan D., Dubois Y., Sousbie T., 2012, MNRAS, 427, 3320
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
Conroy C., Wechsler R. H., Kravtsov A. V., 2006, ApJ, 647, 201
Crain R. A., Eke V. R., Frenk C. S., Jenkins A., McCarthy I. G., Navarro J. F., Pearce F. R., 2007, MNRAS, 377, 41
Crain R. A. et al., 2009, MNRAS, 399, 1773
Crain R. A., McCarthy I. G., Frenk C. S., Theuns T., Schaye J., 2010a, MNRAS, 407, 1403
Crain R. A., McCarthy I. G., Schaye J., Frenk C. S., Theuns T., 2010b, preprint (arXiv:1011.1906)
Crain R. A., McCarthy I. G., Schaye J., Theuns T., Frenk C. S., 2013, MNRAS, 432, 3005
Dalla Vecchia C., Schaye J., 2008, MNRAS, 387, 1431
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
Deason A. J. et al., 2011, MNRAS, 415, 2607
Dressler A., 1980, ApJ, 236, 351
Font A. S., McCarthy I. G., Crain R. A., Theuns T., Schaye J., Wiersma R. P. C., Dalla Vecchia C., 2011, MNRAS, 416, 2802
Forero-Romero J. E., Hoffman Y., Gottlöber S., Klypin A., Yepes G., 2009, MNRAS, 396, 1815
Gao L., Springel V., White S. D. M., 2005, MNRAS, 363, L66
Geller M. J., Huchra J. P., 1989, Science, 246, 897
Hahn O., Porciani C., Carollo C. M., Dekel A., 2007, MNRAS, 375, 489
Hoffman Y., Metuki O., Yepes G., Gottlöber S., Forero-Romero J. E., Libeskind N. I., Knebe A., 2012, MNRAS, 425, 2049
Kauffmann G., White S. D. M., Guiderdoni B., 1993, MNRAS, 264, 201
Knebe A. et al., 2011, MNRAS, 415, 2293
Knebe A. et al., 2013a, MNRAS, 428, 2039
Knebe A. et al., 2013b, MNRAS, 435, 1618
Knollmann S. R., Knebe A., 2009, ApJS, 182, 608
Kravtsov A. V., 2013, ApJ, 764, L31
Libeskind N. I., Hoffman Y., Knebe A., Steinmetz M., Gottlöber S., Metuki O., Yepes G., 2012, MNRAS, 421, L137
Libeskind N. I., Hoffman Y., Forero-Romero J., Gottlöber S., Knebe A., Steinmetz M., Klypin A., 2013, MNRAS, 428, 2489
McGaugh S. S., Schombert J. M., de Blok W. J. G., Zagursky M. J., 2010, ApJ, 708, L14
McCarthy I. G., Font A. S., Crain R. A., Deason A. J., Schaye J., Theuns T., 2012a, MNRAS, 420, 2245
McCarthy I. G., Schaye J., Font A. S., Theuns T., Frenk C. S., Crain R. A., Dalla Vecchia C., 2012b, MNRAS, 427, 379
Marigo P., 2001, A&A, 370, 194
Maulbetsch C., Avila-Reese V., Colín P., Gottlöber S., Khalatyan A., Steinmetz M., 2007, ApJ, 654, 53
Moster B. P., Somerville R. S., Maulbetsch C., van den Bosch F. C., Macciò A. V., Naab T., Oser L., 2010, ApJ, 710, 903
Muldrew S. I. et al., 2012, MNRAS, 419, 2670
Neyrinck M. C., Aragón-Calvo M. A., Jeong D., Wang X., 2014, MNRAS, 441, 646
Noh Y., McQuinn M., 2014, MNRAS, 444, 503
Onions J. et al., 2012, MNRAS, 423, 1200
Peacock J. A., Smith R. E., 2000, MNRAS, 318, 1144
Portinari L., Chiosi C., Bressan A., 1998, A&A, 334, 505
Sales L. V., Navarro J. F., Theuns T., Schaye J., White S. D. M., Frenk C. S., Crain R. A., Dalla Vecchia C., 2012, MNRAS, 423, 1544
Sawala T., Frenk C. S., Crain R. A., Jenkins A., Schaye J., Theuns T., Zavala J., 2013, MNRAS, 431, 1366
Schaye J., Dalla Vecchia C., 2008, MNRAS, 383, 1210
Somerville R. S., Hopkins P. F., Cox T. J., Robertson B. E., Hernquist L., 2008, MNRAS, 391, 481
Sousbie T., Pichon C., Colombi S., Novikov D., Pogosyan D., 2008, MNRAS, 383, 1655
Springel V., 2005, MNRAS, 364, 1105
Springel V. et al., 2005, Nature, 435, 629
Thielemann F.-K. et al., 2003, Nucl. Phys. A, 718, 139
van den Bosch F. C. et al., 2007, MNRAS, 376, 841
Wang H., Mo H. J., Jing Y. P., Yang X., Wang Y., 2011, MNRAS, 413, 1973
Weinmann S. M., Pasquali A., Oppenheimer B. D., Finlator K., Mendel J. T., Crain R. A., Macciò A. V., 2012, MNRAS, 426, 2797
White S. D. M., Frenk C. S., Davis M., Efstathiou G., 1987, ApJ, 313, 505
Wiersma R. P. C., Schaye J., Smith B. D., 2009, MNRAS, 393, 99
Yang X., Mo H. J., van den Bosch F. C., 2008, ApJ, 676, 248

APPENDIX A: NUMERICAL CONVERGENCE TESTS

The web formalism includes a number of steps that could potentially alter our results. We have taken care to ensure that this is not the case.
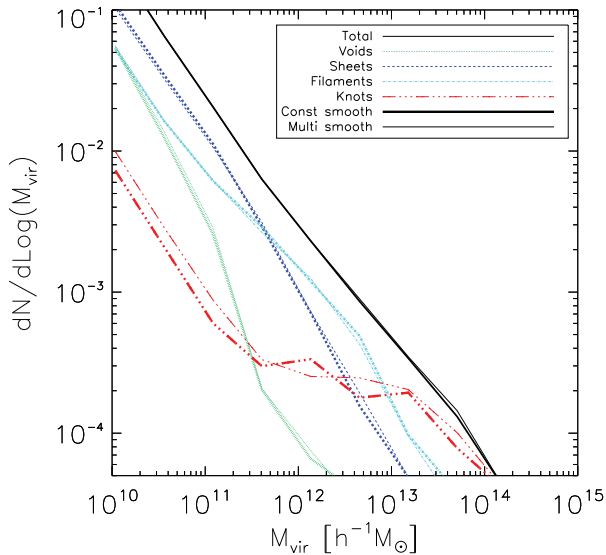
Figure A1. The mass functions of parent haloes, comparing the case of a constant smoothing kernel size (thick lines; solid lines), as was used in this work, and of a smoothing kernel size that changes according to the size of the computational box (thin lines; dashed lines). As can be seen, the two smoothing kernel sizes produce very similar curves.

Periodic boundary conditions

The V-web algorithm uses the FFT to calculate the spatial derivatives of the velocity field. The application of the FFT assumes periodic boundary conditions for the computational box. This is not the case in the present analysis of the GIMIC zoom subboxes and this might lead to some numerical artefacts in a 'boundary layer' that is a few cells thick (as was shown in Hoffman et al. 2012). However, this has virtually no effect in the present case where the CIC box is chosen to enclose the almost spherical zoom region, and haloes are chosen only from that spherical region. Hence, only very few high-resolution cells are affected.

Smoothing kernel size

That the cell sizes are different between the different simulations is unavoidable since the volumes of their area of interest (the high-resolution spheres) are different, and the FFT algorithm is fastest

Table A1. Convergence test for the grid size. Shown are the volume fractions occupied by each web element in the -2σ simulation, for three different griddings – a 128^3 grid from the MidRes simulation, a 128^3 grid from the HiRes simulation, and a 256^3 grid from the HiRes simulation.

	Voids	Sheets	Filaments	Knots
128^3 MidRes	76.9 per cent	20.5 per cent	2.5 per cent	0.1 per cent
128^3 HiRes	78.4 per cent	19.3 per cent	2.2 per cent	0.1 per cent
256^3 HiRes	76.6 per cent	20.9 per cent	2.4 per cent	0.1 per cent

when the grid size is a factor of 2. In order to be able to compare the different simulations, the same smoothing kernel sizes (1–4 cell sizes of the $+2\sigma$ simulation) were used in all the simulations. To ensure that this choice does not affect our results, we have made a similar analysis with smoothing kernel sizes that are 1–4 cell lengths for each simulation, rather than of the $+2\sigma$. All our results are reproduced in that case; as an example we bring a comparison of the mass functions of parent haloes, shown in Fig. A1.

Grid size

In this work, we use the intermediate resolution set (MidRes) of GIMIC simulations since not all simulations in the high-resolution set (HiRes) have been run to redshift zero. This forces us to use a 128^3 grid in order to ensure that there is (on average) at least one particle per cell. However, in order to ensure that the choice of grid size does not affect our results, we have performed resolution tests on different sized grids; since for the MidRes simulation a 256^3 grid is too sparse, we have used the HiRes version of the -2σ simulation, which has been run to redshift zero, and compared it also to the MidRes simulation. This was done by comparing the volume fractions of the different web elements from the webs computed on three grids – a 128^3 grid in the MidRes simulation, a 128^3 grid in the HiRes simulation, and a 256^3 grid in the HiRes simulation, all with the same smoothing kernel size. The results are summarized in Table A1, and as can be seen, they are extremely similar for all grids.

This paper has been typeset from a TeX/LaTeX file prepared by the author.