Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **457,** 1931–1943 (2016) doi:10.1093/mnras/stw145

The APOSTLE simulations: solutions to the Local Group's cosmic puzzles

Till Sawala,[1,2★] Carlos S. Frenk,[1] Azadeh Fattahi,[3] Julio F. Navarro,[3†] Richard G. Bower,[1] Robert A. Crain,[4] Claudio Dalla Vecchia,[5,6] Michelle Furlong,[1] John. C. Helly,[1] Adrian Jenkins,[1] Kyle A. Oman,[2] Matthieu Schaller,[1] Joop Schaye,[7] Tom Theuns,[1] James Trayford[1] and Simon D. M. White[8]

[1] *Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK*
[2] *Department of Physics, University of Helsinki, Gustaf Hällströmin katu 2a, FI-00014 Helsinki, Finland*
[3] *Department of Physics and Astronomy, University of Victoria, 3800 Finnerty Road, Victoria, British Columbia V8P 5C2, Canada*
[4] *Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK*
[5] *Instituto de Astrofísica de Canarias, C/ Vía Láctea s/n, E-38205 La Laguna, Tenerife, Spain*
[6] *Departamento de Astrofísica, Universidad de La Laguna, Av. del Astrofísico Francisco Sánchez s/n, E-38206 La Laguna, Tenerife, Spain*
[7] *Leiden Observatory, Leiden University, Postbus 9513, NL-2300 RA Leiden, the Netherlands*
[8] *Max-Planck Institute for Astrophysics, Karl-Schwarzschild-Strae 1, D-85741 Garching, Germany*

Accepted 2016 January 13. Received 2016 January 13; in original form 2015 November 3

ABSTRACT

The Local Group galaxies offer some of the most discriminating tests of models of cosmic structure formation. For example, observations of the Milky Way (MW) and Andromeda satellite populations appear to be in disagreement with *N*-body simulations of the 'lambda cold dark matter' (ΛCDM) model: there are far fewer satellite galaxies than substructures in CDM haloes (the 'missing satellites' problem); dwarf galaxies seem to avoid the most massive substructures (the 'too-big-to-fail' problem); and the brightest satellites appear to orbit their host galaxies on a thin plane (the 'planes of satellites' problem). Here we present results from APOSTLE (A Project Of Simulating The Local Environment), a suite of cosmological hydrodynamic simulations of 12 volumes selected to match the kinematics of the Local Group (LG) members. Applying the EAGLE code to the LG environment, we find that our simulations match the observed abundance of LG galaxies, including the satellite galaxies of the MW and Andromeda. Due to changes to the structure of haloes and the evolution in the LG environment, the simulations reproduce the observed relation between stellar mass and velocity dispersion of individual dwarf spheroidal galaxies without necessitating the formation of cores in their dark matter profiles. Satellite systems form with a range of spatial anisotropies, including one similar to the MWs, confirming that such a configuration is not unexpected in ΛCDM. Finally, based on the observed velocity dispersion, size, and stellar mass, we provide estimates of the maximum circular velocity for the haloes of nine MW dwarf spheroidals.

Key words: galaxies: evolution – galaxies: formation – cosmology: theory.

1 INTRODUCTION

The ability of the cold dark matter (CDM) model to predict observables on different scales and at different epochs lies at the root of its remarkable success. Predictions for the anisotropy of the microwave background radiation (Peebles 1987), and the large-scale distribution of galaxies (Davis et al. 1985), were made soon after the model was formulated, and have since been spectacularly validated by observations (Planck Collaboration XVI 2014). However, observations on scales currently testable only within the Local Group (LG) have yielded results that appear to be in conflict with CDM predictions.

The 'missing satellites' problem (Klypin et al. 1999; Moore et al. 1999) refers to the apparent paucity of luminous satellite galaxies compared to the large number of dark matter substructures predicted in ΛCDM.[1] That the number of observed dwarf galaxies does not

★ E-mail: till.sawala@helsinki.fi
† Senior CIfAR Fellow.

[1] Throughout this paper, we use ΛCDM to refer to the Λ Cold Dark Matter model, where Λ refers to the cosmological constant model for dark energy, which is constrained primarily by observations outside the LG. We also assume that the dark matter is collisionless.

directly mirror the number of substructures is perhaps no surprise: it has long been predicted that processes such as supernova feedback (e.g. Larson 1974), and cosmic re-ionization (e.g. Efstathiou 1992) should reduce the star formation efficiency in low-mass haloes, and even prevent the smallest ones from forming stars altogether.

The potential of these processes to bring the stellar mass function in ΛCDM into agreement with observations has already been demonstrated using semi-analytical models (e.g. Benson et al. 2002; Somerville 2002; Font et al. 2011; Guo et al. 2011). However, these do not investigate whether the resulting satellites would also match the observed kinematics, or indeed predict ratios between stellar mass and maximum circular velocity, v_{max}, that are far lower than observed. Extending the success of ΛCDM from the overall galaxy population down to the number of observed LG dwarf galaxies through direct simulations has remained a challenge.

Furthermore, the apparent excess of substructures in ΛCDM is not just limited to the lowest mass scales: simulations also predict the presence of subhaloes so massive that they should not be affected by re-ionization [and hence deemed 'too big to fail', (Boylan-Kolchin, Bullock & Kaplinghat 2011)], but whose internal structure seems incompatible with that of the brightest observed satellites. While the two main galaxies within the LG both have several satellite galaxies whose rotation curves are consistent with massive subhaloes, their number falls short of the number of such substructures predicted to surround Milky Way (MW) or M31 mass haloes in ΛCDM.

Finally, it has long been known that most of the 11 brightest, so called classical MW satellites, appear to orbit the Galaxy on a thin plane, and a similar (but distinct) plane has subsequently also been identified among satellites of Andromeda. While ΛCDM satellite systems are known to be anisotropic, the satellite systems of the MW and Andromeda are purported to be extremely unusual (Pawlowski & Kroupa 2013).

The problems, at times reported as fatal for ΛCDM, arise when observations are confronted with predictions from dark matter only (DMO) simulations that treat the cosmic matter content as a single collisionless fluid, a poor approximation on scales where baryonic processes are important. It has, of course, long been recognized that the distribution of light is not a precise tracer of dark matter, but simple models for populating dark matter structures with galaxies do not capture the complexity of galaxy formation physics.

On the other hand, hydrodynamic simulations which have confirmed the importance of baryonic effects have largely focused on individual dwarf galaxies, ignoring the LG environment (e.g. Stinson et al. 2009; Governato et al. 2010; Sawala et al. 2010, 2011; Nickerson et al. 2011; Shen et al. 2014; Wheeler et al. 2015). While earlier hydrodynamic simulations have been able to reproduce the observed large-scale structures in the Local Universe (e.g. Gottloeber, Hoffman & Yepes 2010; Yepes, Gottlöber & Hoffman 2014), they did not reproduce the observed LG dwarf galaxy population (Benítez-Llambay et al. 2015). The lack of a single model able to reconcile all of the LG observations with ΛCDM predictions has led to renewed interest in alternatives to CDM, such as warm (e.g. Lovell et al. 2012) or self-interacting dark matter (e.g. Rocha et al. 2013; Vogelsberger et al. 2014).

Here, we test the ΛCDM model with a new suite of cosmological hydrodynamic simulations, with initial conditions tailored to match the LG environment. In particular, we focus on pairs of haloes that match the separation, approach velocity, and relative tangential velocity of the MW and Andromeda (M31). From a large cosmological simulation, we have selected 12 pairs of haloes with combined virial masses of $\sim 2.3 \pm 0.6 \times 10^{12}\,M_\odot$, compatible with the most recent dynamical constraints (González, Kravtsov & Gnedin 2014; Peñarrubia et al. 2014). The selection and set-up of our initial conditions are described in more detail by Fattahi et al. (2015).

We have re-simulated each LG volume at several resolutions, both as DMO simulations, and as hydrodynamic simulations, with the code developed for the Evolution and Assembly of GaLaxies and their Environments project (EAGLE; Crain et al. 2015; Schaye et al. 2015).

In this paper, we give an introduction to the simulations, and compare our results to the observed LG galaxy population. We show that the abundance of galaxies in the LG can be reproduced within ΛCDM using a galaxy formation model calibrated on much larger scales. By comparing the DMO and hydrodynamic simulations, we show that the apparent discrepancies between observations and ΛCDM predictions can be resolved once baryonic effects are included.

This paper is organized as follows. In Section 2, we give a brief overview of the EAGLE code (Section 2.1), followed by a description of the APOSTLE volumes (Section 2.2). In Section 3, we present our key results, and implications for the questions outlined above: the stellar mass function (Section 3.3), and the satellites 'too big to fail' (Section 3.4). We demonstrate that our simulations also produce the right satellite galaxies in the right satellite haloes (Section 3.6), and provide estimates for v_{max} of observed satellites based on the measured velocity dispersion and stellar mass (Section 3.5). We revisit the apparent kinematical and spatial alignment of satellites in Section 3.7 and conclude with a summary in Section 4.

2 METHODS

2.1 The EAGLE galaxy formation model

The simulations presented in this paper were performed with the code developed for the EAGLE (Crain et al. 2015; Schaye et al. 2015) project. The EAGLE code is a substantially modified version of P-GADGET-3, which itself is an improved version of the publicly available GADGET-2 code (Springel 2005). Gravitational accelerations are computed using the standard Tree-PM scheme of P-GADGET-3, while hydrodynamic forces are computed in the smoothed particle hydrodynamics scheme of ANARCHY described in Dalla-Vecchia et al. (in preparation) and Schaller et al. (2015b), which uses the pressure–entropy formalism introduced by Hopkins (2013).

EAGLE is an evolution of the models used in the GIMIC (Crain et al. 2009) and OWLS (Schaye et al. 2010) projects and has been calibrated to reproduce the $z = 0.1$ stellar mass function and galaxy sizes from 10^8 to $10^{11}\,M_\odot$ in a cosmological volume of $100^3\,\mathrm{Mpc}^3$. In addition, the EAGLE simulations also successfully reproduce many other properties and scaling laws of observed galaxy populations, including the evolution of the stellar mass function (Furlong et al. 2015), and the luminosities and colours of galaxies (Trayford et al. 2015).

The subgrid physics model of EAGLE is described in detail by Schaye et al. (2015). It includes radiative cooling, star formation, stellar evolution and stellar mass-loss, and thermal feedback that captures the collective effects of stellar winds, radiation pressure and supernova explosions. It also includes black hole growth fuelled by gas accretion and mergers, and feedback from active galactic nuclei(AGN; Booth & Schaye 2009; Rosas-Guevara et al. 2015). Within the LG, AGN feedback is negligible, and the main processes regulating the formation of galaxies are gas cooling and heating by

the UV background, star formation, and supernova feedback, which are described in more detail below.

Following Wiersma, Schaye & Smith (2009a), net cooling rates are computed separately for 11 elements, assuming ionization equilibrium in the presence of uniform UV and X-ray backgrounds from quasars and galaxies (Haardt & Madau 2001), and the cosmic microwave background. Hydrogen is assumed to re-ionize instantaneously at $z = 11.5$, which is implemented by turning on the ionizing background. At higher redshifts, the background is truncated at 1 Ryd, limiting its effect to preventing the formation of molecular hydrogen. During re-ionization, an extra 2 eV per proton mass are injected to account for the increase in the photoheating rates of optically thick gas over the optically thin rates that are used otherwise. For hydrogen, this is done at $z = 11.5$, ensuring that the gas is quickly heated to 10^4 K, but for He II, the extra heat is distributed in time following a Gaussian centred at $z = 3.5$ with $\sigma(z) = 0.5$, which reproduces the observed thermal history (Schaye et al. 2000; Wiersma et al. 2009b). In order to prevent artificial fragmentation of the interstellar medium, a temperature floor is imposed on the gas through a polytropic equation of state with index $\gamma = 4/3$, normalized to $T = 8 \times 10^3$ K at a gas threshold density of $n_{\rm H} = 10^{-1}$ cm^{-3} (Schaye & Dalla Vecchia 2008).

The star formation rate is assumed to be pressure-dependent (Schaye & Dalla Vecchia 2008) and follows the observed Kennicutt–Schmidt star formation law with a metallicity-dependent density threshold (Schaye 2004). Energy feedback from star formation is implemented using the stochastic, thermal prescription of Dalla Vecchia & Schaye (2012). The expectation value for the energy injected per unit stellar mass formed decreases with the metallicity of the gas and increases with the gas density to account for unresolved radiative losses and to help prevent spurious numerical losses. The injected energy is calibrated to reproduce the observed, $z = 0$ galaxy stellar mass function and sizes (Crain et al. 2015). On average, it is close to the energy available from core collapse supernovae alone (Schaye et al. 2015). Galactic winds develop naturally, without imposing mass loading factors, velocities or directions.

In our highest resolution simulations, each of the main galaxies contains more than 20 million particles, comparable to the best simulations of individual MW-sized galaxies published to date. Nevertheless, they still barely resolve the scaleheight of the MW thin disc, and the effective resolution is also limited by the equation of state imposed on the gas. Furthermore, the resolution and the physics included in the EAGLE code do not resolve individual star-forming regions or supernova feedback events. We rely instead on a well-calibrated subgrid physics model to parametrize the star formation and feedback processes. For our study, we have used the same parameter values that were used in the 100^3 Mpc^3 L100N1504 EAGLE reference simulation (Schaye et al. 2015) independently of resolution.

While there is clearly scope for future improvements, the relevant properties discussed in this paper, such as the stellar mass function and the circular velocity function of substructures are well converged, as demonstrated in the appendix. This indicates that our numerical resolution is sufficient to capture the physical mechanisms of structure formation, gas accretion and outflows in the EAGLE model.

2.2 The APOSTLE simulations

Our 12 LG regions are zoom simulations selected from a DMO simulation of 100^3 Mpc^3 with 1620^3 particles in *Wilkinson Microwave Anisotropy Probe* (*WMAP*-7) cosmology (Komatsu 2011). The re-simulation volumes were selected to match the available dynamical constraints of the LG. Each volume contains a pair of haloes in the virial mass[2] range 5×10^{11}–2.5×10^{12} M$_\odot$, with median values of 1.4×10^{12} M$_\odot$ for the primary (more massive) halo and 0.9×10^{12} M$_\odot$ for the secondary (less massive) halo of each pair. The combined masses of the primary and secondary range from 1.6×10^{12} to 3.6×10^{12} M$_\odot$ with a median mass of 2.3×10^{12} M$_\odot$, consistent with recent estimates of $2.40^{+1.95}_{-1.05} \times 10^{12}$ M$_\odot$ from dynamical arguments and CDM simulations (González et al. 2014), or $2.3 \pm 0.7 \times 10^{12}$ based on equations of motions that take into account the observed velocities of galaxies in the local volume (Peñarrubia et al. 2014).

We further require that the two haloes be separated by 800 ± 200 kpc, approaching with radial velocity of (0–250) km s^{-1} and with tangential velocity below 100 km s^{-1}; to have no additional halo larger than the smaller of the pair within 2.5 Mpc from the mid-point of the pair, and to be in environments with a relatively unperturbed *Hubble* flow out to 4 Mpc. More details about the selection criteria, and implications of the different dynamical constraints on the total mass of the LG may be found in Fattahi et al. (2015).

The high-resolution initial conditions were created using second-order Lagrangian perturbation theory (Jenkins 2010). The cosmological parameters and the linear phases of the parent volume, which are based on the public multiscale Gaussian white noise field PANPHASIA, are given in tables 1 and 6 of Jenkins (2013), who also describes the method used to make the LG zoom initial conditions.

Each region sampled with baryons and at the highest resolution comprises a sphere of at least 2.5 Mpc radius from the LG barycentre at $z = 0$. Outside of these regions, dark matter particles of increasing mass are used to sample the large-scale environment of the 100^3 Mpc^3 parent simulation. To investigate the impact of baryons, we also repeated all our simulations as DMO, where the dark matter particle masses in the high-resolution region are larger by a factor of $(\Omega_{\rm b} + \Omega_{\rm DM})/\Omega_{\rm DM}$ than in the corresponding hydrodynamic simulations.

The three different resolution levels of the APOSTLE simulations labelled 'L1', 'L2' and 'L3' have primordial gas (DM) particle masses of approximately $1.0(5.0) \times 10^4$, $1.2(5.9) \times 10^5$ and $1.5(7.5) \times 10^6$ M$_\odot$, respectively, and maximum gravitational softening lengths of 134, 307 and 711 pc. L3 is close to the resolution of the EAGLE L100N1504 simulation. While the EAGLE simulations use the *Planck*-1 cosmology (Planck Collaboration XVI 2014), APOSTLE was performed in the slightly different *WMAP*-7 cosmology (Komatsu 2011), with density parameters at $z = 0$ for matter, baryons and dark energy of $\Omega_{\rm M} = 0.272$, $\Omega_{\rm b} = 0.0455$ and $\Omega_\lambda = 0.728$, respectively, a *Hubble* parameter of $\rm H_0 = 70.4\,km\,s^{-1}\,Mpc^{-1}$, a power spectrum of (linear) amplitude on the scale of $8\,h^{-1}$ Mpc of $\sigma_8 = 0.81$ and a power-law spectral index $n_{\rm s} = 0.967$. On LG scales, we expect the effect of cosmological parameters to be minimal.

3 RESULTS

3.1 Formation of LG galaxies

While the LG observations are all made at $z = 0$, and the focus of our paper is on the relation between the observable stellar component

[2] For halo masses, we generally quote m_{200}, the mass enclosed in a spherical volume whose mean overdensity is 200 times the critical density.

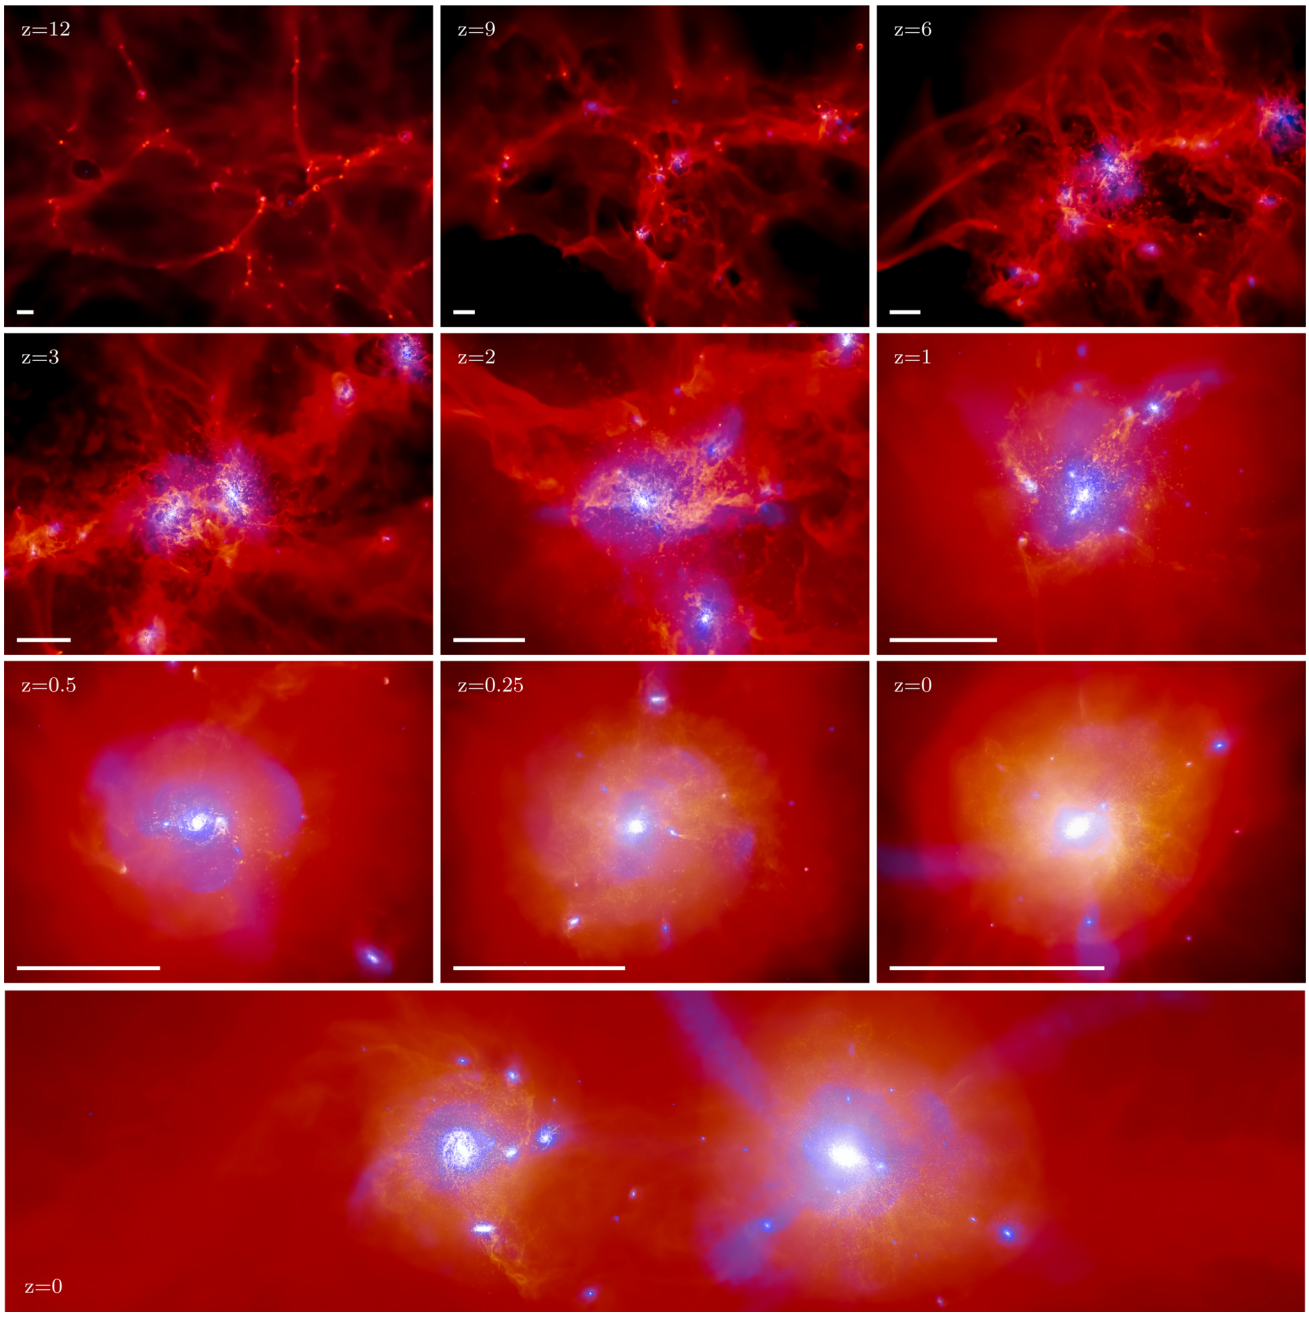

Figure 1. Top three rows: evolution of gas density (red colours) and stellar density (blue colours) in a comoving volume of side length $150 \times 200\ \mathrm{h}^{-1}$ kpc, centred on one of the main LG galaxies and its progenitors at resolution L1. For scale, a bar of length $100\,h^{-1}$ kpc is shown on each panel. At $z = 12$, before re-ionization, stars have already formed in some of the highest density regions, and feedback from supernovae has begun to blow bubbles into the gas. Over time stars form in many more haloes, and star-forming regions merge to form larger galaxies. By $z = 1$, the main galaxy has formed, and continues to accrete both gas and satellites, many of which lose their gas on infall and are also tidally disrupted. At $z = 0$, the central galaxy is surrounded by many satellite galaxies, and a complex stellar halo with visible shells and streams. Bottom row: slice through the entire LG in the same simulation.

and the underlying dark matter model, our simulations allow us to follow its evolution from before the formation of the first stars to the present day. Fig. 1 illustrates the evolution of gas and stars in a comoving region of side length $200 \times 150\ \mathrm{h}^{-1}$ kpc from $z = 12$–0, centred on the particles that become one of the central galaxies in one of the APOSTLE simulations at resolution level L1. At early times, the gas traces the filamentary structure, and stars begin to form in the highest density regions, often found in the nodes at the intersection of filaments. In this simulation, the first stars begin to form at $z \sim 17$, in the progenitors of what will become the pair of main LG galaxies, analogues to the MW and M31. Immediately after the first stars have formed, feedback associated with star formation begins to blow out gas from the then very low-mass dark matter and gas haloes. At $z = 11.5$, re-ionization heats the intergalactic gas and rarefies gas already collapsed in haloes, quenching further gas cooling and accretion into small haloes. As a result, the formation

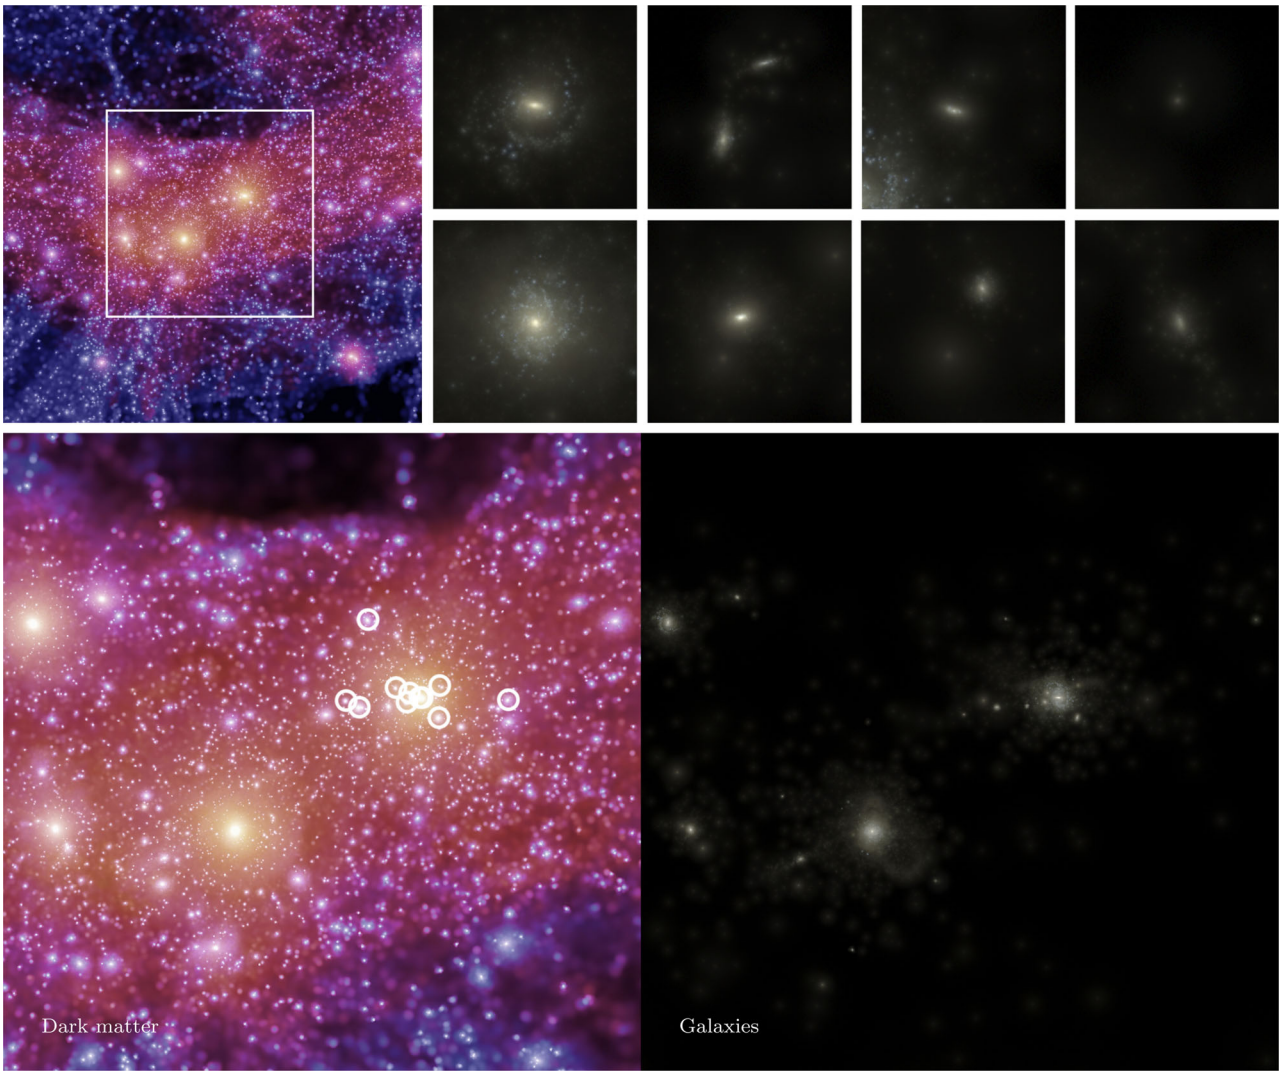

Figure 2. Top left: projected dark matter density in one of our re-simulations at resolution level L2 in a cube of side length 4 Mpc. The main panels contrast the vast number of dark matter substructures (left) with the stellar light distribution (right) in the 2 Mpc cube indicated by the square in the top left panel. Circles indicate the locations of the 11 brightest satellites of one of the main galaxies, whose spatial distribution is as anisotropic as that of the 11 brightest Milky Way satellites, and which align with the filament that contains most of the haloes and galaxies in the region. The small panels in the top row are of side length 125 kpc and reveal in more detail the stellar component of some of the different types of galaxies formed in this simulation; central galaxies (first and second columns) and satellite galaxies (third and fourth columns), which have realistic sizes, colours, and morphologies. Dark matter substructures are abundant in the APOSTLE simulations, but due to the complexity of galaxy formation, starlight paints a very different picture.

of new galaxies is disrupted, until sufficiently massive haloes begin to form. Over time, star formation begins anew in more and more haloes, while individual star-forming regions merge to assemble larger galaxies.

Shortly after $z \sim 3$, the proto-galaxy undergoes a final major merger, with minor mergers continuing up to $z = 0$. The progenitor continues to accrete new satellites that mostly lose their gas on infall due to ram-pressure stripping. A stellar halo also builds up, with shell-like and stream-like substructures originating from tidally disrupted satellite galaxies. By $z = 0$, a pair of large disc galaxies have formed, both surrounded by shells and streams, along with many dwarf galaxy satellites.

3.2 Galaxies that only scratch the surface

Fig. 2 illustrates the dark matter and starlight in another of our re-simulations at redshift $z = 0$. The top left panel shows the dark matter distribution in a cube of side length 4 Mpc, encompassing the spherical volume commonly considered as the LG. It reveals a cosmic filament that envelopes the two principal haloes and most of the galaxies in the region. The bottom row zooms in on a region of side length 2 Mpc around the simulated LG barycentre, contrasting the distribution of dark matter (left-hand panel) and star light (right-hand panel). While the simulations contains tens of thousands of dark matter substructures, galaxies appear as highly biased tracers of the dark matter, forming almost exclusively in the most massive haloes.

Also highlighted in the top left panel are the positions of the satellite haloes that host the 11 brightest satellites of one of the central galaxies. The alignment of the satellites is indicative of a thin plane seen in projection, that is also aligned with the orientation of the filament.

The small insets in Fig. 2 show the stellar structure of some of the many galaxies formed in this simulation. The images use multiband

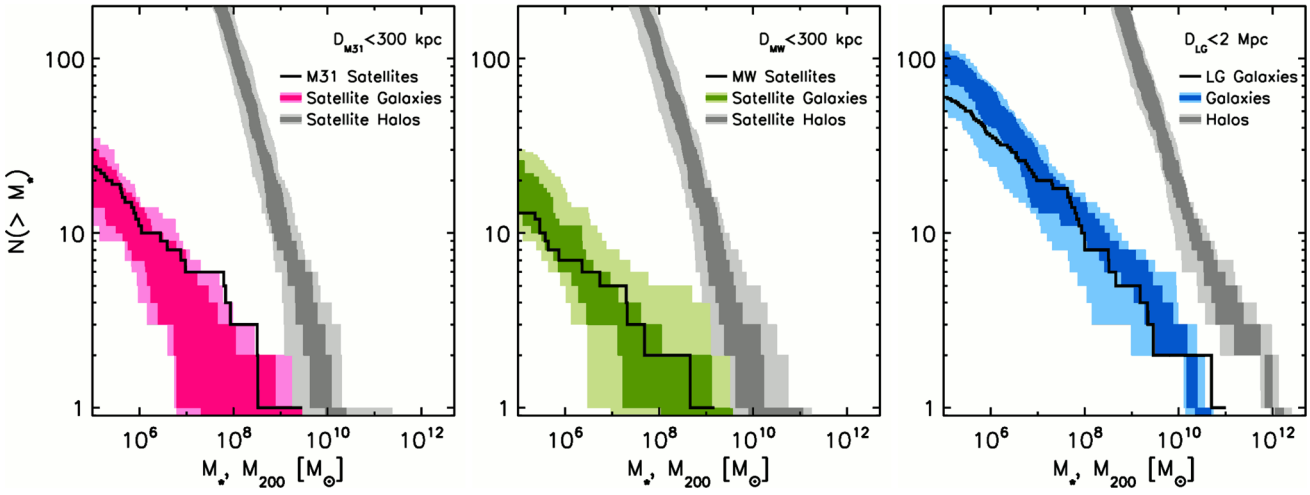

Figure 3. Stellar mass functions from 12 APOSTLE simulations at resolution L2 compared to observations. In the left and centre, shaded regions show the mass functions of satellites within 300 kpc of each of the primary (left) and secondary (centre) of the two main Local Group galaxies from each simulation volume, while lines show the observed stellar mass function within 300 kpc of M31 (left) and the MW (centre). In the right, the shaded region shows all galaxies within 2 Mpc of the Local Group barycentre in the simulations, while the line is the stellar mass function of all known galaxies within the same region. On each panel, the dark colour-shaded areas bound the 16th and 84th percentiles; light shaded areas indicate the full range among our 12 Local Group realizations. For comparison, the grey area on each panel corresponds to the mass function of all dark matter haloes. All observational data are taken from the latest compilation by McConnachie (2012). Note that while the M31 satellite count is likely to be complete to $10^5\,\mathrm{M}_\odot$, the count of satellites of the MW and the total count within 2 Mpc should be considered as lower limits to the true numbers due to the limited sky coverage of local galaxy surveys and the low surface brightness of dwarf galaxies. See Fig. A1 for numerical convergence.

colours rendered using a spectrophotometric model (Trayford et al. 2015). A variety of disc and spheroid morphologies, luminosities, colours, and sizes are clearly visible, reminiscent of the diversity of observed LG galaxies.

3.3 No missing satellites

Fig. 3 shows the galaxy stellar mass functions in the simulations, using data from all 12 of the APOSTLE volumes at resolution L2. Results are plotted both within 300 kpc from each of the two main galaxies per volume (labelled 'primary' and 'secondary' in order of halo mass), as well as within 2 Mpc from the LG barycentre, which includes both central and satellite galaxies.

The primary and secondary galaxies have 20^{+10}_{-6} and 18^{+8}_{-5} satellites more massive than $M_* = 10^5\,\mathrm{M}_\odot$ inside 300 kpc, respectively, where the errors indicate the scatter equivalent to 1σ about the median values. This is in good agreement with the observed number of MW and M31 satellites. Within 2 Mpc of the LG barycentre, there are $\sim$60 galaxies with $M_* > 10^5\,\mathrm{M}_\odot$ presently known; our simulations produce 90^{+20}_{-15}. The modest number of luminous galaxies is in stark contrast to the very large number of dark matter haloes found within the same volume, indicated by the grey shaded area in Fig. 3. While feedback from supernovae and stellar winds regulates star formation in those haloes where a dwarf galaxy has formed, re-ionization has left most of the low-mass haloes completely dark. The observed stellar mass function of the LG and those of the MW and M31 satellites are within the 1σ scatter of the average stellar mass function in our re-simulations over most of the stellar mass range. The relative scatter is larger for the satellite galaxies, reflecting the larger relative sampling error, and the fact that the relative variation in single-halo mass among the different APOSTLE volumes is larger than that of the total LG mass.

Excluding substructures, the stellar masses of the MW and M31 analogues in our simulations lie in the range $1.5\text{–}5.5 \times 10^{10}\,\mathrm{M}_\odot$, on the low end compared to the observational estimates for the MW [$5 \times 10^{10}\,\mathrm{M}_\odot$ (Flynn et al. 2006; Bovy & Rix 2013)] but lower than those for M31 [$10^{11}\,\mathrm{M}_\odot$ (Tamm et al. 2012)]. As noted by Schaye et al. (2015), the subgrid physics used in the Reference model of the EAGLE code, which we have adopted in this work, generally results in slightly low stellar masses in haloes of around $10^{12}\,\mathrm{M}_\odot$ compared to abundance matching expectations (e.g. Guo et al. 2010), while the MW and M31 both appear to lie above the average stellar-to-halo mass relation. While the predicted abundance of satellites and dwarf galaxies within the LG depends on its total mass, as discussed in Section 2.2, and in more detail by Fattahi et al. (2015), we have selected our LG analogues based on their dynamical properties in a pure dark matter simulation, and independently of the stellar mass in the primaries, which may be affected by the limitations of subgrid physics model.

That the simulations reproduce the stellar mass function of galaxies and satellites in the LG over all resolved mass scales is remarkable, given that these simulations use the very same EAGLE model that matches the shape and normalization of the galaxy stellar mass function in large cosmological volumes. Not only are our simulations free of the 'missing satellites' problem, but they indicate that the observed stellar mass functions of the LG volume and of the MW and M31 satellites are entirely consistent with ΛCDM.

3.4 The baryon bailout

We next consider the 'too-big-to-fail' problem (Boylan-Kolchin et al. 2011; Parry et al. 2012). As demonstrated by Strigari, Frenk & White (2010) using the AQUARIUS DMO simulations (Springel et al. 2008), a MW mass halo in ΛCDM typically contains at least one satellite substructure that matches the velocity dispersion profiles measured for each of the five MW dwarf spheroidal satellites for which high-quality kinematic data are available. However, that work addressed neither the question of whether those haloes which match the kinematics of a particular satellite would actually host a comparable galaxy, nor whether an observed satellite galaxy can

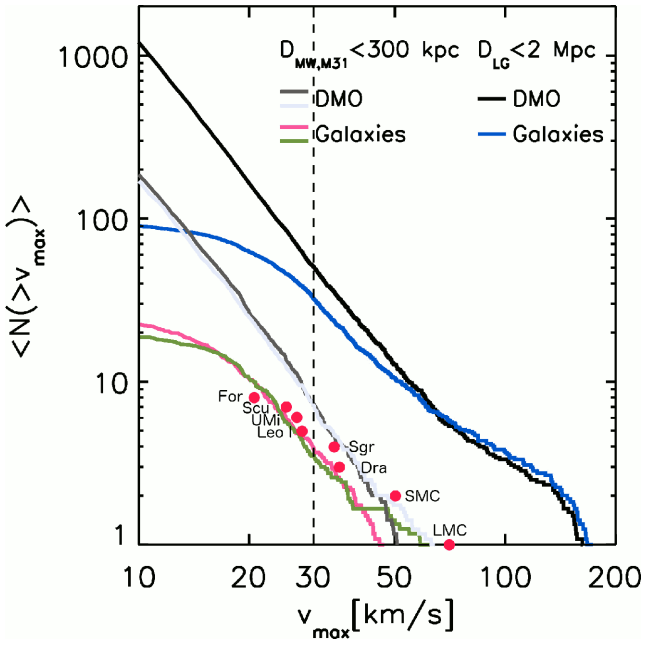

Figure 4. Cumulative number of haloes as a function of maximum circular velocity, v_{max}, averaged over 12 APOSTLE volumes at resolution level L2. The four bottom curves correspond to satellite haloes within 300 kpc of each of the two main galaxies; the top two curves to all systems within 2 Mpc from the LG barycentre. Grey/black curves are from dark matter only (DMO) simulations. Coloured curves are for systems that contain luminous galaxies in the hydrodynamic runs. Red circles show measurements of the most massive MW satellites by Peñarrubia et al. (2008). For guidance, the dashed line denotes a v_{max} value of 30 km s^{-1}. The abundance of satellites with $v_{max} > 30$ km s^{-1} is halved in the hydrodynamic simulations and matches the Milky Way observations. At lower values of v_{max}, the drop in the abundance relative to the DMO case increases as fewer and fewer subhaloes host an observable galaxy. See Fig. A1 for numerical convergence.

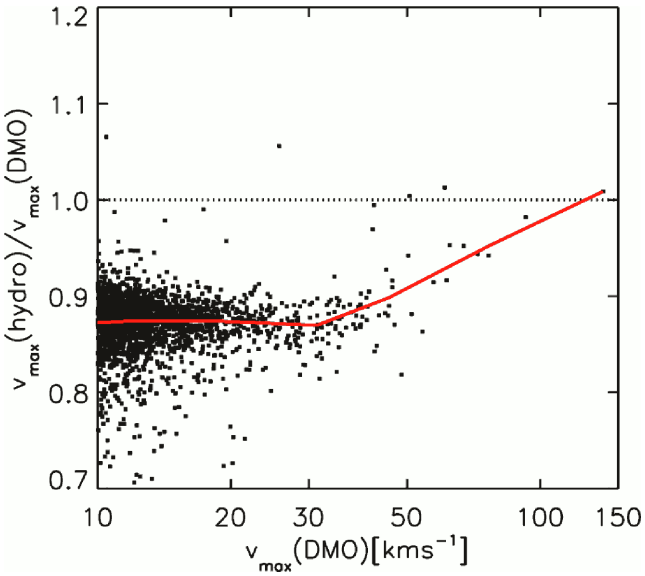

Figure 5. Ratio between the maximum circular velocity, v_{max}, of individual, isolated haloes in the hydrodynamic simulation and the DMO simulation of the same volume, as a function of v_{max}, at resolution L1. The red line shows the binned median ratio. The loss of baryons and the truncated growth leads to a reduction in v_{max} for haloes below $v_{max} \sim 100$ km s^{-1}.

be found to match each of the many predicted satellite haloes. Indeed, the identification in the same simulations of an excess of massive substructures with no observable counterparts, and the implication that the brightest satellites of the MW appear to shun the most massive CDM substructures, constitutes the 'too big-to-fail' problem (Boylan-Kolchin et al. 2011; Parry et al. 2012). A simple characterization of the too-big-to-fail problem is given by the number of satellite haloes with maximum circular velocities, $v_{max} = \max\left(\sqrt{GM(<r)/r}\right)$, above $\sim$30 km s^{-1}, where all satellite haloes are expected to be luminous (Okamoto, Gao & Theuns 2008; Sawala et al. 2016). Only three MW satellites are consistent with haloes more massive than this limit (the two Magellanic Clouds and the Sagittarius dwarf), whereas DMO ΛCDM simulations of MW-sized haloes produce two to three times this number.

Indeed, as shown in Fig. 4, when we consider the DMO counterparts of our LG simulations, the MW and M31 haloes each contain an average of 7–8 satellites with $v_{max} > 30$ km s^{-1} inside 300 kpc, more than twice the observed number of luminous satellites. This is despite the fact that, in order to match the most recent dynamical constraints (González et al. 2014; Peñarrubia et al. 2014), the average halo masses of M31 and the MW in the APOSTLE simulations are lower than those in which the problem was first identified (Boylan-Kolchin et al. 2011).

The situation changes, however, when we consider the hydrodynamic simulations: each main galaxy in our hydrodynamic simulation has on average only 3–4 luminous satellites with $v_{max} > 30$ km s^{-1}. Furthermore, the average velocity function of the most massive substructures across the APOSTLE simulations appears to be in excellent agreement with the MW estimates quoted by Peñarrubia, McConnachie & Navarro (2008) and overplotted as red circles in Fig. 4. It should be noted that the true v_{max} values of dwarf spheroidal galaxies cannot easily be measured, and their estimates rely on additional assumptions. For Fig. 4, we have chosen to use measurements obtained independently from our own simulations, but we revisit this topic in Section 3.5 where we provide ranges of the likely v_{max} values of nine MW dwarf spheroidals with measured stellar masses and velocity dispersions.

Several factors contribute to the reduction in the measured satellite v_{max} function in our hydrodynamic simulations compared to DMO simulations: (i) a reduction in the mass of each subhalo due to baryonic effects as discussed below, (ii) the failure of a fraction of subhaloes of $v_{max} < 30$ km s^{-1} to form any stars, and (iii) those haloes of $v_{max} < 30$ km s^{-1} today that actually contain observable dwarf galaxies having been affected by tidal stripping, even more strongly than typical satellite haloes of the same mass today.

In Fig. 5, we compare the maximum circular velocity of individual isolated haloes matched in our hydrodynamic and DMO simulations. In agreement with Sawala et al. (2013) and Schaller et al. (2015a), we find that while the more massive haloes of $v_{max} > 100$ km s^{-1}, such as those that host the MW and M31 are not significantly affected by baryonic effects, the haloes of dwarf galaxies end up being less massive than their DMO counterparts, because the loss of baryons due to re-ionization and supernova feedback, results in a reduced halo growth rate and leads to a $\sim$15 per cent reduction in v_{max}. The average reduction in mass is similar for the haloes of satellite and isolated galaxies prior to infall, but the more massive satellites in the hydrodynamic simulations experience a further mass-loss relative to their DMO counterparts due to the ram-pressure stripping of the remaining gas.

For haloes below 30 km s^{-1}, the intrinsic reduction in v_{max} due to baryonic effects is compounded by the fact that not all low-mass haloes host galaxies: at 10 km s^{-1}, the fraction of luminous systems

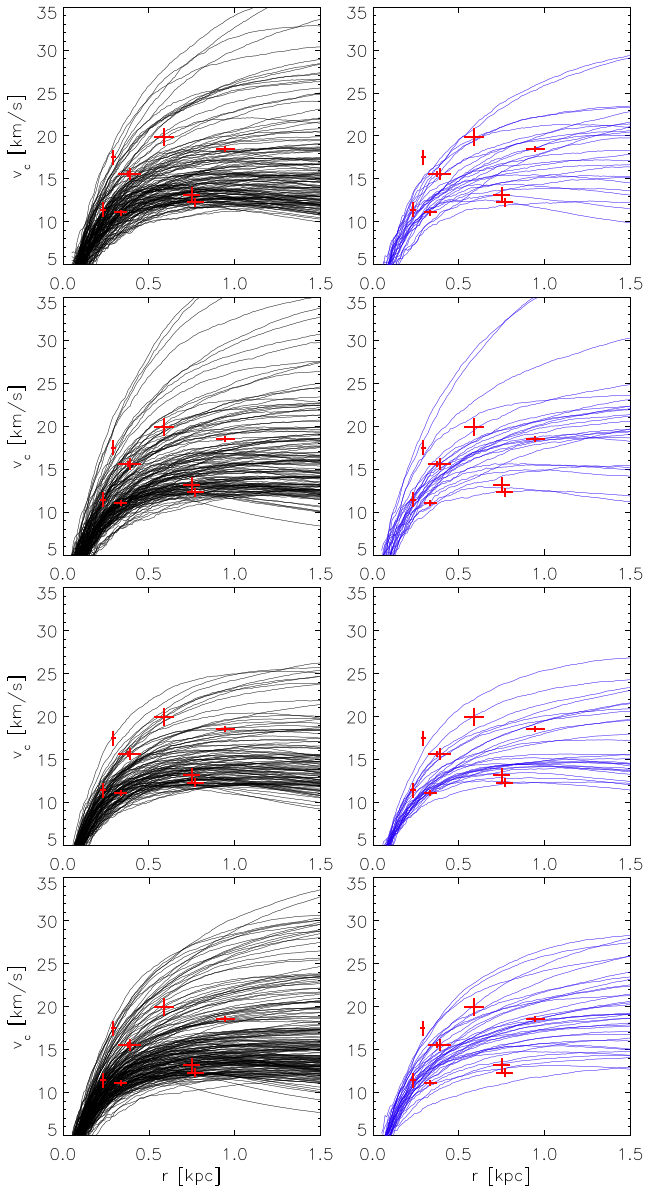

Figure 6. Left: circular velocity profiles of all satellites with $v_{max} > 12\,km\,s^{-1}$ and within 300 kpc from the halo centre, in four individual LG haloes from two APOSTLE simulations at resolution L1. Right: as above, for the satellite galaxies in the corresponding hydrodynamic simulations. Overplotted in red on each panel are the half-mass radii and corresponding circular velocities of nine observed MW dwarf spheroidal satellites adopted from Wolf et al. (2010). While three of the four haloes in the DMO simulations have multiple massive subhaloes without observable counterparts, the discrepancy is resolved in the hydrodynamic simulations. As the LMC, SMC and the Sagittarius dSph, all consistent with $v_{max} > 30\,km\,s^{-1}$, are excluded from this sample, we have also removed the three satellites with the highest v_{max} values from each panel.

is well below 10 per cent and decreases even further towards lower masses.

Fig. 6 shows a comparison of individual circular velocity curves of the most massive satellites within 300 kpc of four of the main LG galaxies, in the DMO simulations (black, left-hand column) and in the corresponding hydrodynamic (blue, right-hand column) simulations. While there is considerable scatter due to the fact that individual satellites can evolve differently in the two simulations, particularly after infall, it can be seen that the satellites in the hydrodynamic simulations have systematically lower circular velocity curves compared to the DMO counterparts. While three of the four haloes in the DMO simulations contain a number of satellites whose circular velocity curves cannot be matched by any of the observed satellite galaxies shown, the velocity curves of satellites in the hydrodynamic simulations are consistent with the observed stellar kinematics.

3.5 V_{max} estimates for dwarf spheroidal galaxies

Because the visible stellar components of dwarf spheroidal galaxies probe only the innermost part of their dark matter haloes, connecting the measured line-of-sight velocity dispersion, σ_{los}, and the size of a satellite to the maximum circular velocity, v_{max}, of its halo is not straightforward, and relies on several assumptions.

Peñarrubia et al. (2008) estimated v_{max} values of individual MW satellite galaxies, assuming that their stellar and dark matter components follow King profiles and NFW profiles, respectively. Boylan-Kolchin, Bullock & Kaplinghat (2012) also gave values for v_{max} for nine MW satellites with stellar masses above $10^5\,M_\odot$. Using the result of Wolf et al. (2010) that the uncertainty on the enclosed mass for an observed line-of-sight velocity dispersion is minimal at the stellar half-light radius, $r_{1/2}$, they used the circular velocity profiles of satellite haloes from the AQUARIUS DMO simulations to determine the most likely v_{max} value of satellite haloes in CDM that matched the measured values of $r_{1/2}$ and $\sigma_{los}(r_{1/2})$. Independently, using the same simulations, Strigari et al. (2010) also determined the v_{max} values of five dwarf spheroidals with resolved kinematics. Instead of relying on the velocity dispersion at the half-light radius, they determine the most likely value of v_{max} for a given observed satellite from the best-fitting velocity dispersion profile in the simulated haloes. The observed structural parameters of nine individual MW satellite galaxies compiled by Wolf et al. (2010), and the v_{max} estimates of Strigari et al. (2010), Peñarrubia et al. (2008) and Boylan-Kolchin et al. (2012) are reproduced in Table 1.

For those galaxies where multiple estimates are available, the v_{max} values of Boylan-Kolchin et al. (2012) tend to be lower than those obtained by the two other authors. Indeed, as pointed out by Pawlowski et al. (2015), these low values would still lead to an overprediction of the MW satellite v_{max} function in our simulations, even though the extrapolation down to $10\,km\,s^{-1}$ is somewhat misleading, because it only includes v_{max} values for a small fraction of the known MW satellites.

However, the most likely subhalo in a DMO simulation like AQUARIUS whose enclosed mass inside $r_{1/2}$ corresponds to the observed $\sigma_{los}(r_{1/2})$ may give a low estimate of the satellite's true v_{max}. Not only do baryons change the dark matter subhaloes, as discussed in Section 3.4; but not all subhaloes are expected to host satellite galaxies, and the probability for a low-mass subhalo to host a satellite galaxy depends on its v_{max}, and the typical stellar mass of a satellite galaxy also depends on v_{max}. Hence, not all satellite haloes have the same probability of matching an observed satellite galaxy of a known stellar mass.

Indeed, the sample of satellite galaxies analysed by Boylan-Kolchin et al. (2012) contains nine of the 12 most luminous known MW satellites, which are likely to be amongst the satellites whose haloes have the highest v_{max} values. ΛCDM predicts many more low v_{max} haloes than high v_{max} haloes. While Boylan-Kolchin et al. point out that this does not affect their analysis, it may amplify any potential bias caused by the implicit assumption that all haloes have an equal chance of hosting a satellite of a given stellar mass.

Table 1. Structural parameters and v_{max} estimates for Milky Way satellite galaxies with $M_* > 10^5\,M_\odot$.

Notes and references	M_* ($M_\odot$) (1)	$r_{1/2}$ (pc) (2)	$\sqrt{\langle\sigma^2_{los}(r_{1/2})\rangle}$ (km s^{-1}) (2)	v_{max} (km s^{-1}) (3)	v_{max} (km s^{-1}) (4)	v_{max} (km s^{-1}) (5)	v_{max} (DMO) (km s^{-1}) this work (6)	v_{max} (hydro) (km s^{-1}) this work (7)
Carina	$4.3^{+1.1}_{-0.9}\times10^5$	334 ± 37	6.4 ± 0.2	16	17	$11.4^{+1.1}_{-1.0}$	$13.7^{+4.8}_{-2.2}$	$14.6^{+6.4}_{-2.9}$
Draco	$2.2^{+0.7}_{-0.6}\times10^5$	291 ± 14	10.1 ± 0.5	–	35	$20.5^{+4.8}_{-3.9}$	$27.6^{+12}_{-4.1}$	23.7*
Fornax	$1.7^{+0.5}_{-0.4}\times10^7$	944 ± 53	10.7 ± 0.2	21	21	$17.8^{+0.7}_{-0.7}$	$19.6^{+1.4}_{-1.1}$	20.4*
Leo I	$5.0^{+1.8}_{-1.3}\times10^6$	388 ± 64	9.0 ± 0.4	22	30	$16.4^{+2.3}_{-2.0}$	$18.5^{+9.0}_{-3.4}$	$18.8^{+12.3}_{-4.0}$
Leo II	$7.8^{+2.5}_{-1.9}\times10^5$	233 ± 17	6.6 ± 0.5	–	18	$12.8^{+2.2}_{-1.9}$	$14.2^{+6.4}_{-2.6}$	$17.7^{+6.0}_{-3.2}$
Sculptor	$2.5^{+0.9}_{-0.7}\times10^6$	375 ± 54	9.0 ± 0.2	26	27	$17.3^{+2.2}_{-2.0}$	$20.9^{+9.6}_{-4.0}$	$20.1^{+7.2}_{-1.6}$
Sextans	$5.9^{+2.0}_{-1.4}\times10^5$	768 ± 47	7.1 ± 0.3	12	11	$11.8^{+1.0}_{-0.9}$	$12.4^{+1.6}_{-1.0}$	$13.2^{+2.5}_{-1.5}$
Ursa Minor	$3.9^{+1.7}_{-1.3}\times10^5$	588 ± 58	11.5 ± 0.6	–	29	$20.0^{+2.4}_{-2.2}$	$22.6^{+6.6}_{-4.0}$	$21.5^{+2.8}_{-0.4}$
CanVen I	$2.3^{+0.4}_{-0.3}\times10^5$	750 ± 48	7.6 ± 0.5	–	–	$11.8^{1.3}_{-1.2}$	$12.8^{+2.8}_{-1.6}$	$13.1^{+3.4}_{-1.4}$

Notes. (1): McConnachie (2012), (2): Wolf et al. (2010), (3): Strigari et al. (2010), (4) Peñarrubia et al. (2008), (5) Boylan-Kolchin et al. (2012), (6) satellite haloes from DMO APOSTLEsimulations that match the observed $v_{1/2}$ at the observed $r_{1/2}$, (7) satellite haloes from hydrodynamic APOSTLE simulations that match the observed $v_{1/2}$ at the observed $r_{1/2}$, and that host galaxies that match the observed stellar mass. *For Fornax and Draco, there are too few simulated counterparts to estimate the range reliably.

In Fig. 7, we examine the circular velocity profiles of individual satellite haloes and galaxies in APOSTLE, and compare them to individual observed MW satellites. On each panel, we show the circular velocity curves of satellites located within 300 kpc from the four central galaxies in two LG simulations at resolution L1, and compare them to the circular velocities inferred from $\sigma_{los}(r_{1/2})$ for nine observed MW dwarf spheroidal galaxies.

On the top left panel of Fig. 7 (black curves), we have selected satellite subhaloes from the DMO simulation, requiring that the circular velocity v_c be within three times the quoted observational uncertainty of $\sigma_{los}(r_{1/2})$ for a given satellite at its observed half-light radius, assuming the relation between v_c and $\sigma_{los}(r_{1/2})$ of Wolf et al. (2010).

For the remaining three panels of Fig. 7, we have selected satellite subhaloes from the corresponding hydrodynamic simulations. In the top right (grey curves), we use the same selection criterion as for the DMO case. For the bottom left panel (cyan curves), we have not applied any velocity criterion, but require the stellar mass in the simulation to be within 50 per cent of the observed stellar mass of the individual satellite. Finally, in the bottom right panel (dark blue curves), we combine the two previous criteria, and select satellites from the hydrodynamic simulation whose stellar mass and measured circular velocity are both compatible with those of the observed satellite.

Due to the range in concentration of CDM haloes, both the DMO and hydrodynamic simulations allow a large range of v_{max} values for most observed satellites, in particular for those where the half-light radius, $r_{1/2}$, is small compared to the radius where the circular velocity is maximal, r_{max}. It is also worth noting that when satellites are selected purely by stellar mass, we find that many of the simulated galaxies live in subhaloes that are consistent with the observed kinematics. As discussed in Section 3.6, a notable exception is the Draco dSph, which appears to have an unusually high halo mass for its stellar mass.

Combining the velocity and stellar mass criteria generally reduces the range of compatible v_{max} values for a particular satellite. We list the most likely v_{max} value for each of the nine observed satellites according to our simulations in the two rightmost columns of Table 1, for both the DMO simulation using only the velocity criterion, and for the hydrodynamic simulation, using both the velocity and stellar mass constraints.

It should be noted that while the ranges quoted for the values of Boylan-Kolchin et al. (2012) denote the uncertainty of the most likely value, for our own values we give the 16th and 84th percentiles of v_{max} for the compatible satellite haloes (equivalent to $\pm1\sigma$ for a normal distribution). The latter can be much larger, and reflect the diversity of haloes in ΛCDM and in our simulations. It is also worth pointing out again that any of the quoted v_{max} estimates from observed velocity dispersions are only valid in the ΛCDM context and, for the values that also use the stellar mass, under the additional assumptions made in our simulations.

3.6 The right satellites in the right haloes

We have shown in Figs 3 and 4 that the APOSTLE simulations reproduce the number of satellites as a function of stellar mass and of v_{max}, as inferred from the stellar velocity dispersion and sizes. However, this success does not necessarily imply that the simulations reproduce the stellar mass–v_{max} relation for individual satellites. Furthermore, the stellar-to-total mass ratios of individual dwarf spheroidals show a surprising amount of scatter: Fornax is roughly 100 times brighter than Draco, but appears to inhabit a halo of similar mass.

In our simulations, the key to understanding this puzzling result lies in the fact that, in the v_{max} range below 30 km s^{-1}, where some haloes remain dark, those satellite haloes that host galaxies similar to the MW dwarf spheroidals have typically experienced much more severe tidal stripping, with a resulting reduction in mass or v_{max} greater than expected for typical haloes of the same mass (Sawala et al. 2016). This phenomenon is illustrated in Fig. 8, which compares the stellar mass–v_{max} relation of both central and satellite galaxies in our simulations with data for the actual satellites of the MW. While isolated haloes in our simulations fall below the stellar mass–v_{max} relation of observed satellites, the strong tidal stripping experienced by the satellite galaxies and their haloes brings them into good agreement with the observed MW dwarf spheroidal data. The effect of stripping can also account for the large scatter in stellar mass of dwarf galaxies in a very narrow v_{max} range. In this scenario, the high stellar mass–v_{max} ratio of satellites like Fornax and Sextans would be explained in part by tidal stripping. By contrast, the low stellar-to-halo mass ratio of Draco is more typical of isolated

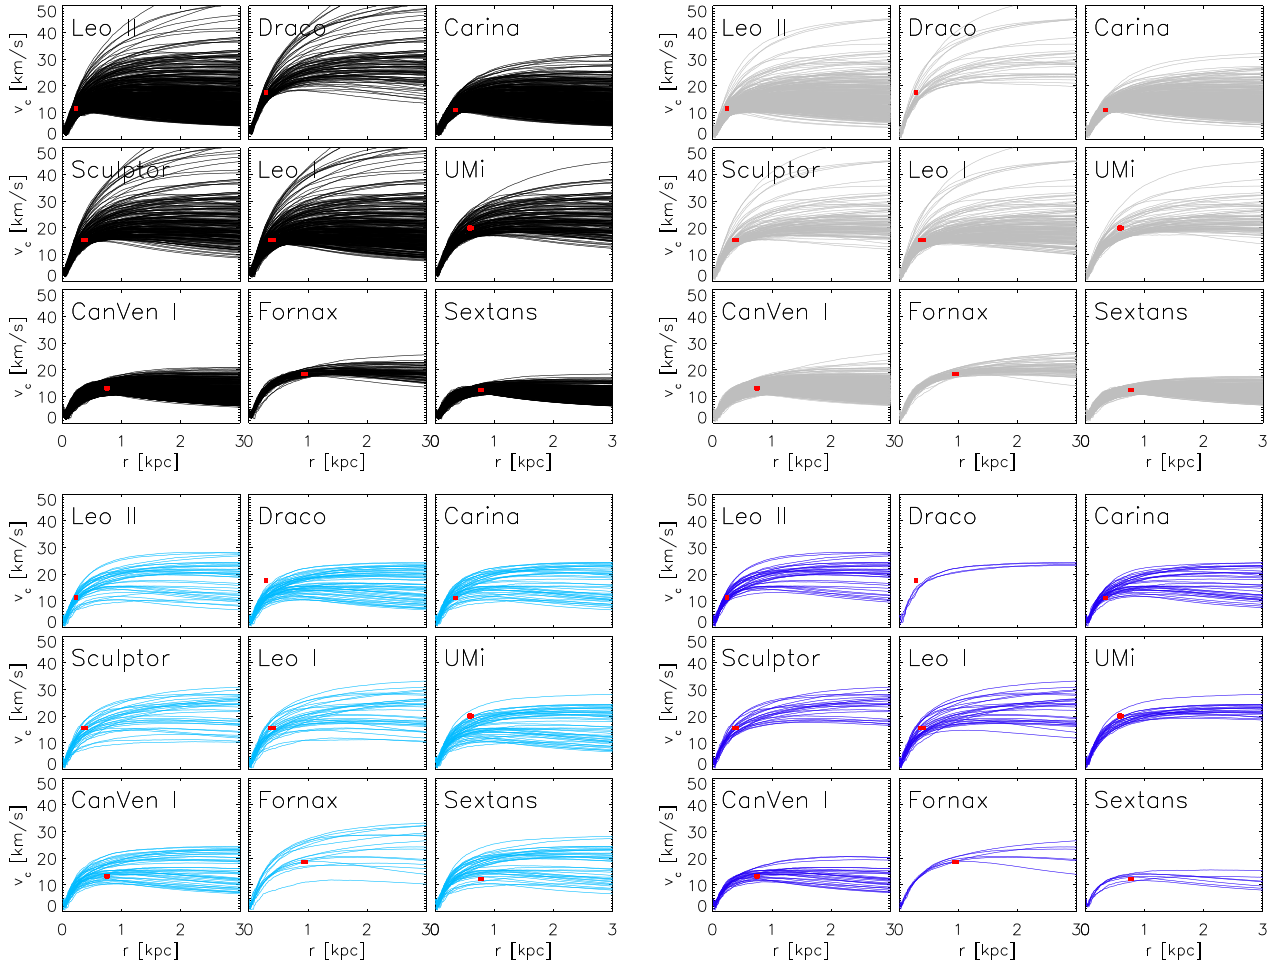

Figure 7. Circular velocity profiles for satellites from a total of four haloes from two of our LG simulations at resolution L1, which have been matched to nine observed dSph galaxies. On the top left, black lines show satellites from the DMO simulations whose circular velocities are within 3σ of the value corresponding to the observed velocity dispersion at the observed half-light radius. On the top right, grey lines show satellites from the corresponding hydrodynamic simulations selected with the same criteria. On the bottom left, cyan lines are for satellites from the hydrodynamic simulations, selected by stellar mass to be within a factor of 2 of the observed value, and on the bottom right, dark blue lines are for satellites from the hydrodynamic simulation that satisfy both the stellar mass and circular velocity criteria. On all panels, red symbols show measurements for observed dwarf spheroidals at the projected 3D half-light radius, adopted from Wolf et al. (2010).

dwarf galaxies in our simulations, suggesting that Draco has not yet experienced strong tidal effects.

In summary, not only do our hydrodynamic simulations reproduce both the observed satellite stellar mass function and satellite circular velocity function; but satellite galaxies of stellar masses comparable to observed dwarf spheroidals also live in haloes with compatible velocity profiles. The reduction in subhalo mass due to baryonic effects, and the strong stripping of haloes that host the luminous satellites, combine to give a satellite population that not only matches the MW and M31 satellite luminosity function, but also the total velocity function of the observed satellite population.

3.7 Unsurprisingly aligned

Finally, the anisotropy, first noticed by Lynden-Bell (1976), and the apparent orbital alignment of the 11 brightest, so-called classical MW satellites, have been regarded as highly improbable in ΛCDM (e.g. Pawlowski & Kroupa 2013). Fig. 9 compares the observed angular distribution and orbital kinematics of the 11 brightest MW satellites to the 11 brightest satellites around one of our simulated LG galaxies at $z = 0$, as identified in Fig. 2. Both the simulated and observed satellite populations show highly anisotropic distributions.

To characterize the anisotropy of each satellite system, we compute the ratio of the minimal to maximal eigenvalues, c and a, of the reduced inertia tensor (e.g. Bailin & Steinmetz 2005) defined by the 11 brightest satellites, $I_{\alpha,\beta} = \sum_{i=1}^{11} r_{i,\alpha} r_{i,\beta} / r_i^2$. From a total of 24 MW or M31-like haloes in our 12 LG simulations at resolution L2, we find values of $\sqrt{c/a}$ in the range 0.34–0.67, compared to 0.36 for the MW and 0.53 for M31.

Clearly, ΛCDM can produce satellite systems with a range of anisotropies, and consistent with measurements for both the MW and M31.

However, considering only the 11 brightest satellites, the satellite distribution of the MW appears to be more anisotropic than all but one of our 24 systems, which still leaves open the question of its statistical significance. It should be noted that, by the same criterion,

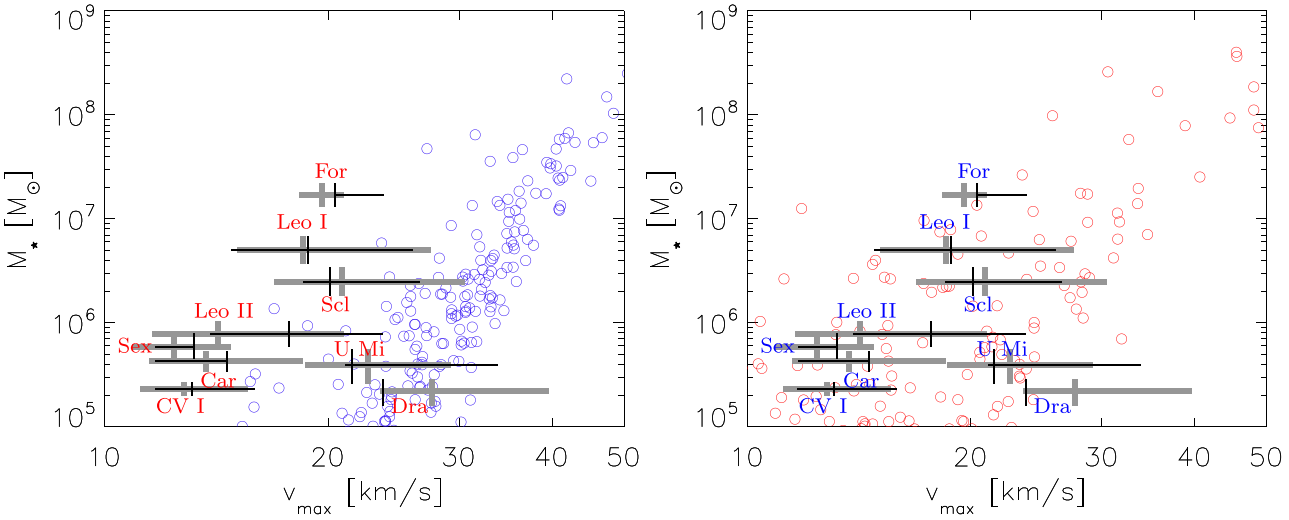

Figure 8. Stellar mass as a function of maximum circular velocity, v_{max}, for isolated galaxies (blue circles, left) and satellite galaxies (red circles, right) from two LG simulations at resolution L1. Also shown are measurements for the nine MW satellites listed in Table 1, with stellar masses adopted from Wolf et al. (2010) and the range of possible v_{max} values computed from the observed $r_{1/2}$ and $\sigma_{los}(r_{1/2})$ as discussed in Section 3.5, based on haloes from the DMO simulations (grey) and based on satellite galaxies from the hydrodynamic simulations (black). While isolated galaxies in the simulation fall below the MW satellites, the reduction in v_{max} particularly of low-mass satellites by tidal stripping, brings the simulated satellite galaxies into good agreement with observations.

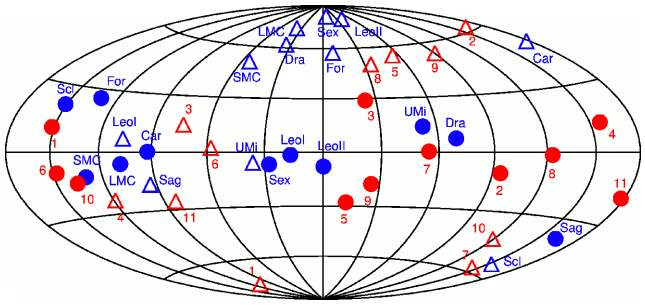

Figure 9. Angular distribution and kinematics of the 11 brightest satellites for a Milky Way like system in our simulation (red), and for the observed Milky Way satellites (blue). Circles show the location on the sky, while triangles indicate the orientation of the corresponding angular momentum vector. For both systems, the equator is chosen to align with the respective plane of satellites. The 11 brightest satellites in the simulated system are distributed on a plane just as flat as those of the Milky Way and several of them have a coherent rotation. Selecting the brightest satellites, systems as anisotropic as the MWs can be formed in ΛCDM, although they are not typical.

the satellite system of M31 is much less anisotropic than that of the MW, and indeed quite typical of our simulated LG galaxies. Given that the MW and M31 clearly formed within the same cosmology with different satellite anisotropies, it may be premature to consider the anisotropy of one a failure of the cosmological model.

The inclusion of baryonic physics in the APOSTLE simulations allow us to directly identify the brightest satellite galaxies out of a much larger number of satellite haloes. However, their spatial anisotropy is clearly not caused by baryonic effects; it is intrinsic to the assembly of ΛCDM satellite systems. As Cautun et al. (2015) have recently demonstrated, even when subsets of observed satellites are carefully chosen to exhibit maximal anisotropy (e.g. Ibata et al. 2014), rigorous statistical analysis shows that the observed spatial and kinematic anisotropies are not inconsistent with ΛCDM and that such apparently 'unusual' systems are in fact quite ubiquitous.

That satellite alignments are consistent with, and indeed expected to exist in ΛCDM does not render them uninteresting. Each individual satellite system can still contain information about the assembly history of the halo and its relation to the large-scale structure. Across our simulations, the satellite plane most like the MW's is aligned with a cosmic filament that envelopes the LG. This supports the scenario proposed by Libeskind et al. (2005), whereby the accretion of satellites from the 'cosmic web' imparts a degree of coherence to the timing and direction of satellite accretion.

4 SUMMARY

APOSTLE is a series of high-resolution zoom simulations of the LG within the ΛCDM cosmology, performed both as DMO, and including baryonic processes with the EAGLE code. We find that our simulations accurately reproduce the observed stellar mass function of the LG, and also result in satellite populations that are in good agreement with the observed relation between dwarf galaxies and their dark matter haloes.

We conclude that the 'problems' often cited as challenges to ΛCDM are resolved in simulations that reproduce the dynamical constraints of the LG environment and include a realistic galaxy formation model. Re-ionization and supernova feedback allow galaxy formation to proceed only in a small subset of dark matter haloes, eliminating the 'missing satellites' problem. It is notable that the very same galaxy formation model calibrated to reproduce the galaxy population in cosmologically representative volumes naturally produces an LG galaxy population in volumes consistent with the LG kinematics in ΛCDM.

The loss of baryons due to re-ionization and star formation feedback, and from satellites through ram pressure stripping, affects the growth of low-mass haloes, leading to a reduction in their maximum circular velocity compared to a DMO simulation. Combined with the effect of tidal stripping, strongly enhanced for luminous,

low-mass haloes, this not only resolves the 'too-big-to-fail' problem, but it also leads to a satellite circular velocity function that matches observations, at the same time as it matches the observed stellar mass function. Furthermore, kinematics of individual simulated satellites are consistent with the kinematics of observed satellites matched by stellar mass. We also find that individual galaxies formed in our simulations follow the observed relation between v_{max} and stellar mass, and use our DMO and hydrodynamic simulations to estimate the range of likely v_{max} values for nine dwarf spheroidals with measured stellar mass and $\sigma_{los}(r_{1/2})$ in the ΛCDM cosmology.

Finally, we find a diversity of satellite systems with spatial anisotropy similar to those of the M31, and one system that is similar to the 'plane of satellites' around the MW. We conclude that the observed anisotropies of these satellite systems do not falsify ΛCDM, but may reflect their assembly histories within the ΛCDM paradigm.

Another often-cited difficulty for ΛCDM is the inference of constant density cores in dark matter haloes (e.g. Walker & Peñarrubia 2011), whereas *N*-body simulations predict cusps. While the observed kinematics of LG dwarf spheroidals are, in fact, consistent with either cores or cusps (Strigari, Frenk & White 2014), it has also been argued that cores are required to solve the too-big-to-fail problem (Brooks et al. 2013). This is not the case: the star formation and feedback model in our simulations and the effect of tidal stripping give rise to a realistic LG galaxy population without cores. Hence, we conclude that cores are not necessary to solve the perceived small-scale problems of ΛCDM.

Our simulations predict that the relation between stellar mass and v_{max} should differ between satellites and isolated dwarf galaxies, as the observed satellites with high stellar mass–v_{max} ratios live in haloes that experienced particularly strong tidal stripping. This prediction is testable: alternative scenarios in which the dark matter haloes are modified independently of environment should give much more similar relations between the two.

Our simulations are still limited by resolution and the shortcomings of the subgrid physics model. Nevertheless, they result in galaxy populations compatible with many of the observations of the LG galaxy population and do not suffer from any of the problems often interpreted as a failure of ΛCDM. Using a model already shown to reproduce the galaxy population on much larger scales, they suggest that the success and predictive power of the ΛCDM cosmology extend far into the LG regime, once the effects of galaxy formation and of the particular LG environment are taken into account.

ACKNOWLEDGEMENTS

We are indebted to Dr Lydia Heck who looks after the supercomputers at the ICC. This work was supported by the Science and Technology Facilities Council [grant number ST/F001166/1 and RF040218], the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Grant agreement 278594 'GasAroundGalaxies', the National Science Foundation under Grant No. PHYS-1066293, the Interuniversity Attraction Poles Programme of the Belgian Science Policy Office [AP P7/08 CHARM]. TS acknowledges the Marie-Curie ITN CosmoComp and the support of the Academy of Finland grant 1274931. CSF acknowledges ERC Advanced Grant 267291 'COSMIWAY' and SW acknowledges ERC Advanced Grant 246797 'GALFORMOD'. RAC is a Royal Society University Research Fellow. This work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk), and resources provided by WestGrid (www.westgrid.ca) and Compute Canada / Calcul Canada (www.computecanada.ca). The DiRAC system is funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grant ST/H008519/1, STFC DiRAC Operations grant ST/K003267/1, and Durham University. DiRAC is part of the National E-Infrastructure.

REFERENCES

Bailin J., Steinmetz M., 2005, ApJ, 627, 647
Benítez-Llambay A., Navarro J. F., Abadi M. G., Gottlöber S., Yepes G., Hoffman Y., Steinmetz M., 2015, MNRAS, 450, 4207
Benson A. J., Lacey C. G., Baugh C. M., Cole S., Frenk C. S., 2002, MNRAS, 333, 156
Booth C. M., Schaye J., 2009, MNRAS, 398, 53
Bovy J., Rix H.-W., 2013, ApJ, 779, 115
Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2011, MNRAS, 415, L40
Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2012, MNRAS, 422, 1203
Brooks A. M., Kuhlen M., Zolotov A., Hooper D., 2013, ApJ, 765, 22
Cautun M., Bose S., Frenk C. S., Guo Q., Han J., Hellwing W. A., Sawala T., Wang W., 2015, MNRAS, 452, 3838
Crain R. A. et al., 2009, MNRAS, 399, 1773
Crain R. A. et al., 2015, MNRAS, 450, 1937
Dalla Vecchia C., Schaye J., 2012, MNRAS, 426, 140
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
Efstathiou G., 1992, MNRAS, 256, 43P
Fattahi A. et al., 2015, preprint (arXiv:1507.03643)
Flynn C., Holmberg J., Portinari L., Fuchs B., Jahreiß H., 2006, MNRAS, 372, 1149
Font A. S. et al., 2011, MNRAS, 417, 1260
Furlong M. et al., 2015, MNRAS, 450, 4486
González R. E., Kravtsov A. V., Gnedin N. Y., 2014, ApJ, 793, 91
Gottloeber S., Hoffman Y., Yepes G., 2010, preprint (arXiv:1005.2687)
Governato F. et al., 2010, Nature, 463, 203
Guo Q., White S., Li C., Boylan-Kolchin M., 2010, MNRAS, 404, 1111
Guo Q. et al., 2011, MNRAS, 413, 101
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. Editions Frontieres, Paris, p. 64
Hopkins P. F., 2013, MNRAS, 428, 2840
Ibata N. G., Ibata R. A., Famaey B., Lewis G. F., 2014, Nature, 511, 563
Jenkins A., 2010, MNRAS, 403, 1859
Jenkins A., 2013, MNRAS, 434, 2094
Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, ApJ, 522, 82
Komatsu E. et al., 2011, ApJS, 192, 18
Larson R. B., 1974, MNRAS, 169, 229
Libeskind N. I., Frenk C. S., Cole S., Helly J. C., Jenkins A., Navarro J. F., Power C., 2005, MNRAS, 363, 146
Lovell M. R. et al., 2012, MNRAS, 420, 2318
Lynden-Bell D., 1976, MNRAS, 174, 695
McConnachie A. W., 2012, AJ, 144, 4
Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, ApJ, 524, L19
Nickerson S., Stinson G., Couchman H. M. P., Bailin J., Wadsley J., 2011, MNRAS, 415, 257
Okamoto T., Gao L., Theuns T., 2008, MNRAS, 390, 920
Parry O. H., Eke V. R., Frenk C. S., Okamoto T., 2012, MNRAS, 419, 3304
Pawlowski M. S., Kroupa P., 2013, MNRAS, 435, 2116
Pawlowski M. S., Famaey B., Merritt D., Kroupa P., 2015, ApJ, 815, 19
Peebles P. J. E., 1987, ApJ, 315, L73
Peñarrubia J., McConnachie A. W., Navarro J. F., 2008, ApJ, 672, 904
Peñarrubia J., Ma Y.-Z., Walker M. G., McConnachie A., 2014, MNRAS, 443, 2204
Planck Collaboration XVI, 2014, A&A, 571, A16
Rocha M., Peter A. H. G., Bullock J. S., Kaplinghat M., Garrison-Kimmel S., Oñorbe J., Moustakas L. A., 2013, MNRAS, 430, 81

Rosas-Guevara Y. M. et al., 2015, MNRAS, 454, 1038
Sawala T., Scannapieco C., Maio U., White S., 2010, MNRAS, 402, 1599
Sawala T., Guo Q., Scannapieco C., Jenkins A., White S., 2011, MNRAS, 413, 659
Sawala T., Frenk C. S., Crain R. A., Jenkins A., Schaye J., Theuns T., Zavala J., 2013, MNRAS, 431, 1366
Sawala T. et al., 2016, MNRAS, 456, 85
Schaller M. et al., 2015a, MNRAS, 451, 1247
Schaller M., Dalla Vecchia C., Schaye J., Bower R. G., Theuns T., Crain R. A., Furlong M., McCarthy I. G., 2015b, MNRAS, 454, 2277
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
Schaye J. et al., 2010, MNRAS, 402, 1536
Schaye J. et al., 2015, MNRAS, 446, 521
Shen S., Madau P., Conroy C., Governato F., Mayer L., 2014, ApJ, 792, 99
Somerville R. S., 2002, ApJ, 572, L23
Springel V., 2005, MNRAS, 364, 1105
Springel V. et al., 2008, MNRAS, 391, 1685
Stinson G. S., Dalcanton J. J., Quinn T., Gogarten S. M., Kaufmann T., Wadsley J., 2009, MNRAS, 395, 1455
Strigari L. E., Frenk C. S., White S. D. M., 2010, MNRAS, 408, 2364
Strigari L. E., Frenk C. S., White S. D. M., 2014, preprint (arXiv:1406.6079)
Tamm A., Tempel E., Tenjes P., Tihhonova O., Tuvikene T., 2012, A&A, 546, A4
Trayford J. W. et al., 2015, MNRAS, 452, 2879
Vogelsberger M., Zavala J., Simpson C., Jenkins A., 2014, MNRAS, 444, 3684
Walker M. G., Peñarrubia J., 2011, ApJ, 742, 20
Wheeler C., Oñorbe J., Bullock J. S., Boylan-Kolchin M., Elbert O. D., Garrison-Kimmel S., Hopkins P. F., Kereš D., 2015, MNRAS, 453, 1305
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, 399, 574
Wolf J., Martinez G. D., Bullock J. S., Kaplinghat M., Geha M., Muñoz R. R., Simon J. D., Avedo F. F., 2010, MNRAS, 406, 1220
Yepes G., Gottlöber S., Hoffman Y., 2014, New Astron. Rev., 58, 1

APPENDIX A: NUMERICAL CONVERGENCE

As discussed in Section 2.2, we have performed our simulations at three resolution levels, varying by a factor of ∼144 in particle mass. Fig. A1 illustrates the numerical convergence of the mass functions (left-hand panel) and velocity function (right-hand panel) for one APOSTLE volume at the three resolution levels. For both the halo mass function and halo velocity function, results are shown from the DMO simulation. Here, convergence is excellent. For the galaxy mass function, and the velocity function of the haloes containing galaxies (coloured curves in both figures), results are shown from the hydrodynamic simulation.

Convergence of the cumulative stellar mass function remains very good. By comparison to Fig. 3, which includes results from all 12 APOSTLE volumes at resolution L2, it is clear that the difference in the stellar mass function due to numerical resolution is much smaller than the variation between the different APOSTLE volumes.

The velocity function of galaxies, which falls below the corresponding total velocity function of haloes in the DMO simulation due the effect of baryons on haloes at $v_{max} \sim 60\,km\,s^{-1}$, and due to the appearance of 'dark' haloes at $v_{max} \sim 30\,km\,s^{-1}$, is equally well converged. In particular, both the average reduction in v_{max} of subhaloes, and the fraction of dark haloes, are independent of resolution.

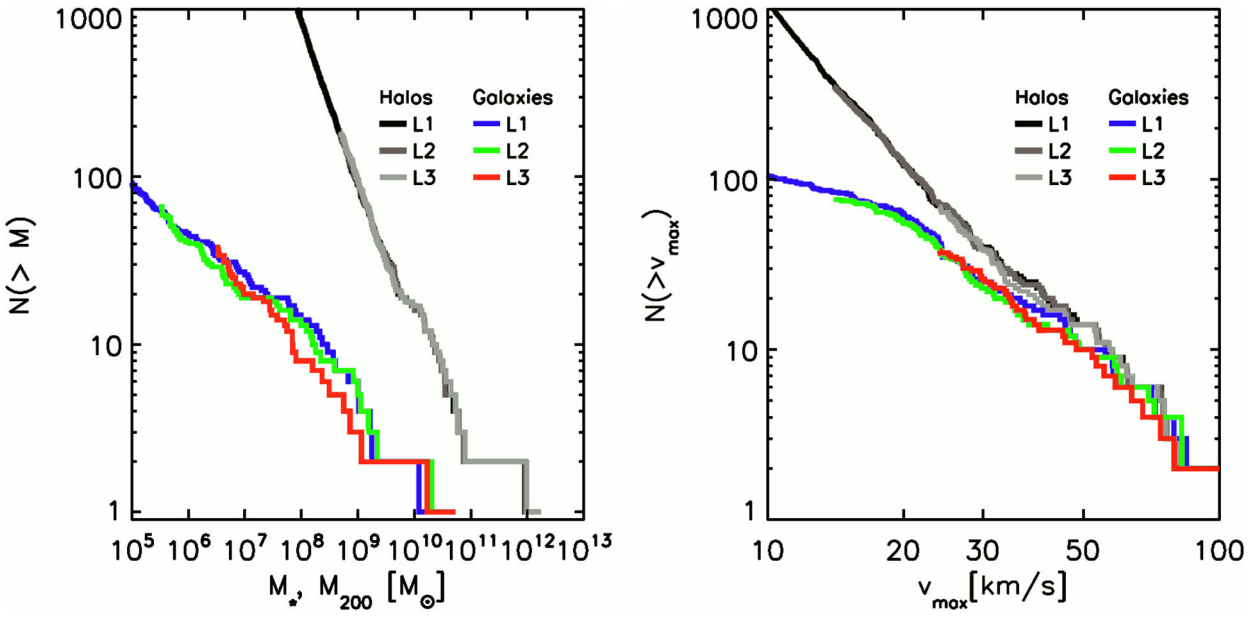

Figure A1. Stellar and halo mass functions (left-hand panel) and maximum circular velocity functions (right-hand panel) within 2 Mpc from the barycentre of one of the APOSTLE volumes at three different resolutions, for galaxies in the hydrodynamic simulations, and haloes in the corresponding DMO simulations. The stellar and halo mass functions, and the velocity functions of galaxies and haloes are well converged at resolution L2.

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