



LJMU Research Online

Fattahi, A, Navarro, JF, Sawala, T, Frenk, CS, Oman, KA, Crain, RA, Furlong, M, Schaller, M, Schaye, J, Theuns, T and Jenkins, A

The APOSTLE project: Local Group kinematic mass constraints and simulation candidate selection

<http://researchonline.ljmu.ac.uk/id/eprint/3669/>

Article

Citation (please note it is advisable to refer to the publisher's version if you intend to cite from this work)

Fattahi, A, Navarro, JF, Sawala, T, Frenk, CS, Oman, KA, Crain, RA, Furlong, M, Schaller, M, Schaye, J, Theuns, T and Jenkins, A (2016) The APOSTLE project: Local Group kinematic mass constraints and simulation candidate selection. MONTHLY NOTICES OF THE ROYAL ASTRONOMICAL SOCIETY.

LJMU has developed [LJMU Research Online](#) for users to access the research output of the University more effectively. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LJMU Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain.

The version presented here may differ from the published version or from the version of the record. Please see the repository URL above for details on accessing the published version and note that access may require a subscription.

For more information please contact researchonline@ljmu.ac.uk

<http://researchonline.ljmu.ac.uk/>

The APOSTLE project: Local Group kinematic mass constraints and simulation candidate selection

Azadeh Fattahi,¹★ Julio F. Navarro,¹† Till Sawala,² Carlos S. Frenk,³ Kyle A. Oman,¹ Robert A. Crain,⁴ Michelle Furlong,³ Matthieu Schaller,³ Joop Schaye,⁵ Tom Theuns³ and Adrian Jenkins³

¹Department of Physics and Astronomy, University of Victoria, PO Box 3055 STN CSC, Victoria, BC V8W 3P6, Canada

²Department of Physics, University of Helsinki, Gustaf Hällströmin katu 2a, FI-00560 Helsinki, Finland

³Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK

⁴Astrophysics Research Institute, Liverpool John Moores University, IC2, Liverpool Science Park, 146 Brownlow Hill, Liverpool L3 5RF, UK

⁵Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands

Accepted 2015 December 18. Received 2015 November 22; in original form 2015 July 10

ABSTRACT

We use a large sample of isolated dark matter halo pairs drawn from cosmological N -body simulations to identify candidate systems whose kinematics match that of the Local Group (LG) of galaxies. We find, in agreement with the ‘timing argument’ and earlier work, that the separation and approach velocity of the Milky Way (MW) and Andromeda (M31) galaxies favour a total mass for the pair of $\sim 5 \times 10^{12} M_{\odot}$. A mass this large, however, is difficult to reconcile with the small relative tangential velocity of the pair, as well as with the small deceleration from the Hubble flow observed for the most distant LG members. Halo pairs that match these three criteria have average masses a factor of ~ 2 times smaller than suggested by the timing argument, but with large dispersion. Guided by these results, we have selected 12 halo pairs with total mass in the range $1.6\text{--}3.6 \times 10^{12} M_{\odot}$ for the APOSTLE project (A Project Of Simulating The Local Environment), a suite of hydrodynamical resimulations at various numerical resolution levels (reaching up to $\sim 10^4 M_{\odot}$ per gas particle) that use the subgrid physics developed for the EAGLE project. These simulations reproduce, by construction, the main kinematics of the MW–M31 pair, and produce satellite populations whose overall number, luminosities, and kinematics are in good agreement with observations of the MW and M31 companions. The APOSTLE candidate systems thus provide an excellent testbed to confront directly many of the predictions of the Λ cold dark matter cosmology with observations of our local Universe.

Key words: methods: numerical – galaxies: dwarf – galaxies: haloes – Local Group – dark matter.

1 INTRODUCTION

The Local Group (LG) of galaxies, which denotes the association of the Milky Way (MW) and Andromeda (M31), their satellites, and galaxies in the surrounding volume out to a distance of ~ 3 Mpc, provides a unique environment for studies of the formation and evolution of galaxies. Their close vicinity implies that LG galaxies are readily resolved into individual stars, enabling detailed exploration of the star formation, enrichment history, structure, dark matter content, and kinematics of systems spanning a wide range of masses and morphologies, from the two giant spirals that dominate the LG gravitationally, to the faintest galaxies known.

This level of detail comes at a price, however. The LG volume is too small to be cosmologically representative, and the properties of its galaxy members may very well have been biased by the peculiar evolution that led to its particular present-day configuration, in which the MW and M31, a pair of luminous spirals ~ 800 kpc apart, are approaching each other with a radial velocity of ~ 120 km s^{−1}. This galaxy pair is surrounded by nearly 100 galaxies brighter than $M_V \sim -8$, about half of which cluster tightly around our Galaxy and M31 (see e.g. McConnachie 2012, for a recent review).

The LG is also a relatively isolated environment whose internal dynamics are dictated largely by the MW–M31 pair. Indeed, outside the satellite systems of MW and M31, there are no galaxies brighter than $M_B = -18$ (the luminosity of the Large Magellanic Cloud, hereafter LMC for short) within 3 Mpc from the MW. The nearest galaxies comparable in brightness to the MW or M31 are just beyond

* E-mail: azadehf@uvic.ca

† Senior ClifAR Fellow.

3.5 Mpc away (NGC 5128 is at 3.6 Mpc; M81 and NGC 253 are located 3.7 Mpc from the MW).

Understanding the biases that this particular environment may induce on the evolution of LG members is best accomplished through detailed numerical simulations that take these constraints directly into account. This has been recognized in a number of recent studies, which have followed small volumes tailored to resemble, in broad terms, the LG (see e.g. Gottloeber, Hoffman & Yepes 2010; Garrison-Kimmel et al. 2014). This typically means selecting ~ 3 -Mpc-radius regions where the mass budget is dominated by a pair of virialized haloes separated by the observed MW–M31 distance and whose masses are chosen to match various additional constraints (see e.g. Forero-Romero et al. 2013).

The mass constraints may include estimates of the individual virial¹ masses of both MW and M31, typically based on the kinematics of tracers such as satellite galaxies, halo stars, or tidal debris (see e.g. Battaglia et al. 2005; Sales et al. 2007; Smith et al. 2007; Xue et al. 2008; Watkins, Evans & An 2010; Deason et al. 2012; Boylan-Kolchin et al. 2013; Barber et al. 2014; Piffi et al. 2014, for some recent studies). However, these estimates are usually accurate only for the mass enclosed within the region that contains each of the tracers, so that virial mass estimates are subject to non-negligible, and potentially uncertain, extrapolation.

Alternatively, the MW and M31 stellar masses may be combined with ‘abundance matching’ techniques to derive virial masses (see e.g. Guo et al. 2010; Behroozi et al. 2013; Kravtsov, Vikhlinin & Meshcheryakov 2014, and references therein). In this procedure, galaxies of given stellar mass are assigned the virial mass of dark matter haloes of matching number density, computed in a given cosmological model. Shortcomings of this method include its reliance on the relative ranking of halo and galaxy mass in a particular cosmology, as well as the assumption that the MW and M31 are average tracers of the halo mass–galaxy mass relation.

A further alternative is to use the kinematics of LG members to estimate virial masses. One example is the ‘numerical action’ method developed by Peebles et al. (2001) to reconstruct the peculiar velocities of nearby galaxies which, when applied to the LG, predicts a fairly large circular velocity for the MW (Peebles, Tully & Shaya 2011). A simpler, but nonetheless useful, example is provided by the ‘timing argument’ (Kahn & Woltjer 1959), where the MW–M31 system is approximated as a pair of isolated point masses that expand radially away after the big bang but decelerate under their own gravity until they turn around and start approaching. Assuming that the age of the Universe is known, that the orbit is strictly radial, and that the pair is on first approach, this argument leads to a robust and unbiased estimate of the total mass of the system (Li & White 2008). Difficulties with this approach include the fact that the tangential velocity of the pair is neglected (see e.g. González, Kravtsov & Gnedin 2014), together with uncertainties relating the total mass of the point-mass pair to the virial masses of the individual systems.

Finally, one may use the kinematics of the outer LG members to estimate the total mass of the MW–M31 pair, since the higher the mass, the more strongly LG members should have been decelerated from the Hubble flow. This procedure is appealing because of its simplicity but suffers from the uncertain effects of nearby massive structures, as well as from difficulties in accounting for the directional dependence of the deceleration and, possibly, for the

gravitational torque/pull of even more distant large-scale structure (see e.g. Peñarrubia et al. 2014; Sorce et al. 2014, for some recent work on the topic).

A review of the literature cited above shows that these methods produce a range of estimates (spanning a factor of 2–3) of the individual masses of the MW and M31 and/or the total mass of the LG (see Wang et al. 2015, for a recent compilation). This severely conditions the selection of candidate LG environments that may be targeted for resimulation, and is a basic source of uncertainty in the predictive ability of such simulations. Indeed, varying the mass of the MW halo by a factor of 3, for example, would likely lead to variations of the same magnitude in the predicted number of satellites of such systems (Boylan-Kolchin, Besla & Hernquist 2011; Wang et al. 2012; Cautun et al. 2014), limiting the insight that may be gained from direct quantitative comparison between simulations and observations of the LG.

With these caveats in mind, this paper describes the selection procedure, from a simulation of a large cosmological volume, of 12 viable LG environment candidates for resimulation. These 12 candidate systems form the basis of the EAGLE–APOSTLE project, a suite of high-resolution cosmological hydrodynamical resimulations of the LG environment in the Λ cold dark matter (Λ CDM) cosmogony. The goal of this paper is to motivate the particular choices made for this selection whilst critically reviewing the constraints on the total mass of the LG placed by the kinematics of LG members. Preliminary results from the project (which we shall hereafter refer to as APOSTLE, a shorthand for ‘A Project Of Simulating The Local Environment’), which uses the same code developed for the EAGLE simulations (Crain et al. 2015; Schaye et al. 2015), have already been reported in Sawala et al. (2015, 2016) and Oman et al. (2015).

We begin by using the Millennium Simulations (Springel et al. 2005; Boylan-Kolchin et al. 2009) to select relatively isolated halo pairs separated by roughly the distance between the MW and M31 and derive the distribution of total masses of pairs that reproduce, respectively, the relative radial velocity of the MW–M31 pair, or its tangential velocity, or the Hubble flow deceleration of distant LG members.

Given the disparate preferred masses implied by each of these criteria when applied individually, we decided to select pairs within a narrow range of total mass that match loosely the LG kinematics rather than pairs that match strictly the kinematic criteria but that span the (very wide) allowed range of masses. This choice allows us to explore the ‘cosmic variance’ of our results given our choice of LG mass, whilst guiding how such results might be scaled to other possible choices. We end by assessing the viability of our candidate selection by comparing their satellite systems with those of the MW and M31 galaxies.

The plan for this paper is as follows. We begin by assessing in Section 2 the constraints on the LG mass placed by the kinematics of the MW–M31 pair and other LG members. We describe next, in Section 3, the choice of APOSTLE candidates and the numerical resimulation procedure. Section 4 analyses the properties of the satellite systems of the main galaxies of the LG resimulations and compares them with observed LG properties. We end with a brief summary of our main conclusions in Section 5.

2 THE MASS OF THE LOCAL GROUP

2.1 Observational data

We use below the positions, Galactocentric distances, line-of-sight velocities, and V-band magnitudes (converted to stellar masses

¹ We define the virial mass, M_{200} , as that enclosed by a sphere of mean density 200 times the critical density of the Universe, $\rho_{\text{crit}} = 3H^2/8\pi G$. Virial quantities are defined at that radius, and are identified by a ‘200’ subscript.

Table 1. The parameters of the cosmological simulations used in this paper.

Simulation	Cosmology	Ω_m	Ω_Λ	Ω_b	h	σ_8	n_s	Cube side (Mpc)	Particle number	$m_{p,DM}$ (M_\odot)
MS-I	WMAP-1	0.25	0.75	0.045	0.73	0.9	1	685	2160 ³	1.2×10^9
MS-II	WMAP-1	0.25	0.75	0.045	0.73	0.9	1	137	2160 ³	9.4×10^6
EAGLE (L100N1504)	Planck	0.307	0.693	0.04825	0.6777	0.8288	0.9611	100	1504 ³	9.7×10^6
DOVE	WMAP-7	0.272	0.728	0.0455	0.704	0.81	0.967	100	1620 ³	8.8×10^6

assuming a mass-to-light ratio of unity in solar units) of LG members as given in the compilation of McConnachie (2012). We also use the relative tangential velocity of the M31–MW pair derived from M31’s proper motion by van der Marel et al. (2012). When needed, we assume a local standard of rest (LSR) velocity of 220 km s^{-1} at a distance of 8.5 kpc from the Galactic Centre and that the Sun’s peculiar motion relative to the LSR is $U_\odot = 11.1 \text{ km s}^{-1}$, $V_\odot = 12.24 \text{ km s}^{-1}$, and $W_\odot = 7.25 \text{ km s}^{-1}$ (Schönrich, Binney & Dehnen 2010), to refer velocities and coordinates to a Galactocentric reference frame.

According to these data, the MW–M31 pair is 787 ± 25 kpc apart, and is approaching with a relative radial velocity of $123 \pm 4 \text{ km s}^{-1}$. In comparison, its tangential velocity is quite low: only 7 km s^{-1} with 1σ confidence region $\leq 22 \text{ km s}^{-1}$. We shall assume hereafter that these values are comparable to the relative velocity of the centres of mass of each member of the pairs selected from cosmological simulations. In other words, we shall ignore the possibility that the observed relative motion of the MW–M31 pair may be affected by the gravitational pull of their massive satellites, i.e. the Magellanic Clouds (in the case of the MW) and/or M33 (in the case of M31). This choice is borne out of simplicity; correcting for the possible displacement caused by these massive satellites requires detailed assumptions about their orbits and their masses, which are fairly poorly constrained (see e.g. Gómez et al. 2015).

We shall also consider the recession velocity of distant LG members, measured in the Galactocentric frame. This is also done for simplicity, since velocities in that frame are more straightforward to compare with velocities measured in simulations. Other work (see e.g. Garrison-Kimmel et al. 2014) has used velocities expressed in the LG-centric frame defined by Karachentsev & Makarov (1996). This transformation aims to take into account the apex of the Galactic motion relative to the nearby galaxies in order to minimize the dispersion in the local Hubble flow. This correction, however, is sensitive to the volume chosen to compute the apex, and difficult to replicate in simulations.

2.2 Halo pairs from cosmological simulations

We use the Millennium Simulations, Millennium-I Simulation (MS-I; Springel et al. 2005) and Millennium-II Simulation (MS-II; Boylan-Kolchin et al. 2009), to search for halo pairs with kinematic properties similar to the MW and M31. The MS-I run evolved 2160³ dark matter particles, each of $1.2 \times 10^9 M_\odot$, in a box 685 Mpc on a side adopting Λ CDM cosmological parameters consistent with the 1-year *Wilkinson Microwave Anisotropy Probe* (WMAP-1) measurements. The MS-II run evolved a smaller volume (137 Mpc on a side) using the same cosmology and number of particles as MS-I. Each MS-II particle has a mass of $9.4 \times 10^6 M_\odot$. We list in Table 1 the main cosmological and numerical parameters of the cosmological simulations used in our analysis.

At $z = 0$, dark matter haloes in both simulations were identified using a friends-of-friends (FoF; Davis et al. 1985) algorithm

run with a linking length equal to 0.2 times the mean interparticle separation. Each FoF halo was then searched iteratively for self-bound substructures (subhaloes) using the SUBFIND algorithm (Springel et al. 2001). Our search for halo pairs include all pairs of separate FoF haloes, as well as single FoF haloes with a pair of massive subhaloes satisfying the kinematic and mass conditions we list below. The latter is an important part of our search algorithm, since many LG candidates are close enough to be subsumed into a single FoF halo at $z = 0$.

The list of MS-I and MS-II haloes retained for analysis include all pairs separated by 400 kpc to 1.2 Mpc whose members have virial masses exceeding $10^{11} M_\odot$ each but whose combined mass does not exceed $10^{13} M_\odot$. These pairs are further required to satisfy a fiducial isolation criterion, namely that no other halo more massive than the less massive member of the pair be found within 2.5 Mpc from the centre of the pair (we refer to this as ‘medium isolation’ or ‘MedIso’, for short). We have also experimented with tighter/looser isolation criteria, enforcing the above criterion within 1 Mpc (‘loosely isolated’ pairs; or ‘LoIso’) or 5 Mpc (‘highly isolated’ pairs; ‘HiIso’). Since the nearest galaxy with mass comparable to the Milky Way is located at ~ 3.5 Mpc, the fiducial isolation approximates best the situation of the LG; the other two choices allow us to assess the sensitivity of our results to this particular choice.

2.3 Radial velocity constraint and the timing argument

The left-hand panel of Fig. 1 shows the radial velocity versus separation of all pairs in our MS-I samples, selected using our maximum isolation criterion. Each point in this panel is coloured by the total mass of the pairs, defined as the sum of the virial masses of each member.

The clear correlation seen between radial velocity and mass, at given separation, is the main prediction of the ‘timing argument’ discussed in Section 1. Timing argument predictions are shown by the dashed lines, which indicate the expected relation for pairs with total mass as stated in the legend. This panel shows clearly that low-mass pairs as distant as the MW–M31 pair are still, on average, expanding away from each other (positive radial velocity), in agreement with the timing argument prediction (top dashed curve). Indeed, for a total mass as low as $2 \times 10^{11} M_\odot$, the binding energy reaches zero at $r = 914$ kpc, $V_r = 43 \text{ km s}^{-1}$, where the dashed line ends.

It is also clear that predominantly massive pairs have approach speeds as large as the MW–M31 pair ($\sim 120 \text{ km s}^{-1}$). We illustrate this with the dotted line, which shows the evolution in the r – V_r plane of a point-mass pair of total mass $5 \times 10^{12} M_\odot$, selected to match the MW–M31 pair at the present time. The pair reached ‘turnaround’ (i.e. null radial velocity) about 5 Gyr ago and has since been approaching from a distance of 1.1 Mpc (in physical units) to reach the point labelled ‘M31’ by $z = 0$, in the left-hand panel of Fig. 1.




Figure 1. Left: the relative radial velocity versus distance of all halo pairs selected at highest isolation (‘HiIso’) from the Millennium Simulation. Colours denote the total mass of the pairs (i.e. the sum of the two virial masses), as indicated by the colour bar. The starred symbol indicates the position of the MW–M31 pair in this plane. The dotted line illustrates the evolution of a point-mass pair of total mass $\sim 5 \times 10^{12} M_{\odot}$ in this plane (in physical coordinates) which, according to the timing argument, ends at M31’s position. Dashed lines indicate the loci, at $z = 0$, of pairs of given total mass (as labelled) but different initial energies, according to the timing argument. The box surrounding the M31 point indicates the range of distances and velocities used to select pairs for further analysis. Right: same as left-hand panel, but for the relative tangential velocity. Dashed curves in this case indicate the mean distance–velocity relation for pairs of given total mass, as labelled.

An interesting corollary of this observation is that the present turnaround radius of the LG is expected to be well beyond 1.1 Mpc. Assuming, for guidance, that the turnaround radius grows roughly like $t^{8/9}$ (Bertschinger 1985), this would imply a turnaround radius today of roughly ~ 1.7 Mpc, so that all LG members just inside that radius should be on first approach. We shall return to this point when we discuss the kinematics of outer LG members in Section 2.5 below.

2.4 Tangential velocity constraints

The right-hand panel of Fig. 1 shows a similar exercise to that described in Section 2.3, but using the relative tangential velocity of the pairs. This is compared with that of the MW–M31 pair, which is measured to be only $\sim 7 \text{ km s}^{-1}$ by Sohn, Anderson & van der Marel (2012) and is shown by the starred symbol labelled ‘M31’.

This panel shows that, just like the radial velocity, the tangential velocity also scales, at a given separation, with the total mass of the pair. In general, higher mass pairs have higher speeds, as gleaned from the colours of the points and by the three dashed lines, which indicate the average velocities of pairs with total mass as labelled.

The average relative tangential velocity of a $5 \times 10^{12} M_{\odot}$ pair separated by ~ 800 kpc is about $\sim 100 \text{ km s}^{-1}$, and very few of such pairs (only ~ 6 per cent) have velocities as low as that of the MW–M31 pair. The orbit of a typical pair of such mass is thus quite different from the strictly radial orbit envisioned in timing argument estimates. Indeed, the low tangential velocity of the MW–M31 pair clearly favours a much lower mass for the pair than derived from the timing argument (see González et al. 2014, for a similar finding).

The kinematics of the MW–M31 pair is thus peculiar compared with that of halo pairs selected from cosmological simulations: its radial velocity is best matched with relatively large masses, whereas its tangential velocity suggests a much lower mass. We shall return

to this issue in Section 2.6, after considering next the kinematics of the outer LG members.

2.5 Hubble flow deceleration

The kinematics of the more distant LG members is also sensitive to the total mass of the MW–M31 pair. In particular, we expect that the larger the mass, the more the recession velocities of those members would be decelerated from the Hubble flow. We explore this in Fig. 2, where the symbols in each panel show the Galactocentric radial velocity of all galaxies in the McConnachie (2012) catalogue and in the Extragalactic Distance Database of Tully et al. (2009) found within 3 Mpc of the MW. These data illustrate two interesting points. One is that all galaxies beyond ~ 1.3 Mpc from the MW are receding; and the second is that the mean velocity of all receding galaxies is only slightly below the Hubble flow, $V_r = H_0 r$, indicated by the dot–dashed line for $H_0 = 70.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

The first result is intriguing, given our argument in Section 2.3 that the turnaround radius of the LG should be around ~ 1.7 Mpc at the present time. The second point is also interesting, since it suggests that the local Hubble flow around the MW (beyond ~ 1.3 Mpc) has been relatively undisturbed. The first point suggests that the M31 motion is somewhat peculiar relative to that of the rest of the LG members; the second that the total mass of the MW–M31 pair cannot be too large, for otherwise the recession velocities would have been decelerated more significantly.

We examine this in more detail in the four panels of Fig. 2, where the coloured lines show the mean recession speed as a function of distance for haloes and subhaloes surrounding candidate MS-II²

² We use here only MS-II pairs; the numerical resolution of the MS-I simulation is too coarse to identify enough haloes and subhaloes around the selected pairs to accurately measure the velocities of outer LG members.




Figure 2. Radial velocity versus distance for all LG members out to a distance of 3 Mpc from the Galactic Centre. Symbols repeat in each panel, and correspond to LG galaxies in the Galactocentric reference frame, taken from the compilations of McConnachie (2012) and Tully et al. (2009). Coloured bands are different in each panel and correspond to simulations. Solid symbols are used to show LG members that lie, in projection, within 30° from the direction to M31. The dotted line in each panel is the timing argument curve for M31, as in Fig. 1. Each panel corresponds to pairs of different mass, as given in the legends. The coloured lines and shaded regions correspond, in each panel, to the result of binning MedIso MS-II halo pairs separated by 600–1000 kpc and least-square fitting the recession velocities of the outer members (between 1.5 and 3 Mpc from either primary). The solid coloured line shows the median slope and zero-point of the individual fits. The shaded regions show the interquartile range in zero-point velocity. Note that the recession speeds of outer members decrease with increasing LG mass. The dot–dashed line shows the unperturbed Hubble flow, for reference.

pairs of different mass, binned as listed in the legend. Recession speeds are measured from each of the primaries for systems in the distance range 1.5–3 Mpc and are least-square fitted independently to derive a slope and velocity zero-point for each pair. The coloured lines in Fig. 2 are drawn using the median slope and zero-point for all pairs in each mass bin. The shaded areas indicate the interquartile range in the zero-point velocity. This figure shows clearly that the more massive the pair, the lower, on average, the recession velocities of surrounding systems.

This is also clear from Fig. 3, where we plot the individual values of the slopes and zero-point velocities (measured at $r = 2.5$ Mpc), compared with the values obtained from a least-squares fit to the LG data. The zero-point of the fit is most sensitive to the total mass of the pair; as a result, the relatively large recession velocity of the outer LG members clearly favours a low total mass for the MW–M31 pair.

2.6 Mass distributions

The histograms in Fig. 4 summarize the results of the previous three subsections. Each of the top three panels shows the distribution of the total masses of pairs with separations in the range (600, 1000) kpc selected to satisfy the following constraints: (i) relative radial velocity in the range $(-175, -75)$ km s^{-1} (dark blue histograms; see selection box in the left-hand panel of Fig. 1); (ii) tangential velocity in the range $(0, 50)$ km s^{-1} (light-blue histogram; see box in the right-hand panel of Fig. 1); and (iii) Hubble velocity fits in the observed range (lightest blue histograms; see box in Fig. 3). The histograms in red correspond to the few systems that satisfy all three criteria at once. (The scale of the red histograms has

been increased in order to make them visible; see legends.) In these three panels the isolation criteria for selecting pairs tightens from top to bottom, and, as a consequence, the total number of selected pairs decreases. As discussed before, the radial velocity constraint favours high-mass pairs, whereas the tangential velocity constraint favours low-mass ones. The Hubble velocity constraint gives results intermediate between the other two.

Fig. 4 shows that the main effect of relaxing the isolation criteria is to enable relatively low-mass pairs to pass the kinematic selection. The total mass of ‘HiIso’ pairs matching the radial velocity constraint, for example, clusters tightly about the timing argument prediction ($\sim 5 \times 10^{12} M_\odot$; see dark blue histogram in the ‘HiIso’ panel). However, a long tail of pairs with much lower masses appears in the other panels, indicating that neighbouring structures can have a non-negligible effect on the kinematics of a pair. In particular, ‘infall’ on to a massive structure may accelerate low-mass pairs to much higher relative velocities than they would be able to reach in isolation.

The magnitude of the effect is expected to scale with the mass and distance to the most massive neighbouring system, and it seems legitimate to question whether an object like the Virgo cluster might have a discernible effect on the kinematics of the MW–M31 pair. We have checked this explicitly by identifying the subsample of pairs in Fig. 4 that have a halo of mass $\gtrsim 5 \times 10^{13} M_\odot$ in the distance range³ 15–20 Mpc. We find no statistically significant difference between the pair mass distribution of this subsample and that of the MedIso and LoIso samples shown in Fig. 4. (This subsample is too

³ These numbers are chosen to match the Virgo cluster, a $\sim 4 \times 10^{14} M_\odot$ system situated 17 Mpc from the MW (McLaughlin 1999; Tonry et al. 2001).




Figure 3. Parameters of the least-squares fits to the recession speeds of outer LG members as a function of distance. The solid starred symbol labelled ‘LG’ indicates the result of fitting the distance versus Galactocentric radial velocity of all LG galaxies between 1.5 and 3 Mpc from the MW. The Hubble flow is shown by the solid circle. Each coloured point corresponds to a halo pair selected from the MS-II simulation assuming medium isolation, and uses all haloes in the 1.5–3 Mpc range resolved in MS-II with more than 100 particles (i.e. masses greater than $1 \times 10^9 M_\odot$). Velocities and distances are measured from either primary, and coloured according to the total mass of the pair. Note that recession velocities decrease steadily as the total mass of the pair increases. The square surrounding the ‘LG’ point indicates the error in the slope and zero-point, computed directly from the fit to the 33 LG members with distances between 1.5 and 3 Mpc. Open starred symbols correspond to the 12 pairs selected for the APOSTLE project (see Section 3.2).

small to draw a statistically sound conclusion for the HiIso case.) This suggests that only nearby massive systems can skew the pair mass distribution, and that a cluster as distant as Virgo is unlikely to play a substantial role in the kinematics of the M31–MW pair.

A further feature highlighted by Fig. 4 is that very few pairs satisfy all three constraints simultaneously. For the case of medium isolation (second panel from the top in Fig. 4) only 14 pairs satisfy the three criteria simultaneously in the $(137 \text{ Mpc})^3$ volume of the MS-II simulation. The bottom panel of Fig. 4 shows the mass distribution of pairs that satisfy all three criteria in the case of ‘medium isolation’, the closest to our observed LG configuration (empty red histogram). The mass distribution of pairs that satisfy the three criteria simultaneously peaks at $\sim 2 \times 10^{12} M_\odot$, a value significantly lower than the mass suggested by the timing argument. The distribution is quite broad, with an rms of 0.4 dex and a full range of values that stretches more than a decade; from 2.3×10^{11} to $6.1 \times 10^{12} M_\odot$ in the case of medium isolation.

This finding is one of the main reasons guiding our choice of mass for the halo pairs that we select for resimulation in the APOSTLE project; their distribution is shown by the solid black histogram in the bottom panel of Fig. 4. We describe the selection procedure of these 12 pairs in detail next.

3 THE APOSTLE SIMULATIONS

The APOSTLE project consists of a suite of high-resolution cosmological hydrodynamical simulations of 12 LG-like environments




Figure 4. Total mass distributions of all pairs in the MS-II with separation and relative radial velocity similar to the MW–M31 pair, selected assuming different isolation criteria. The dark blue histogram shows pairs that satisfy the distance–radial velocity criterion shown by the box in the left-hand panel of Fig. 1. The lighter blue histogram shows those pairs satisfying the constraint shown by the box in the distance–tangential velocity panel of Fig. 1. The lightest blue histogram identifies pairs satisfying the ‘Hubble flow’ criterion shown by the box in Fig. 3. The red histogram corresponds to the intersection of all three criteria; their scale has been increased to make them more easily visible. The axis on the right shows the scale of the red histogram. From top to bottom, the first three panels show results for different levels of isolation, as labelled. The numbers in brackets indicate the number of pairs in each histogram. Arrows indicate the median of each distribution. The bottom panel shows in red the results for medium isolation (empty histogram), and compares them with the mass distribution of the 12 candidate pairs selected from the DOVE simulation for the APOSTLE simulation project (in solid black).

selected from large cosmological volumes of a Λ CDM universe. Each of these volumes has been simulated at three different resolutions with the same code as was used for the EAGLE simulation project (Crain et al. 2015; Schaye et al. 2015).

Results from a subset of these simulations have already been presented in Sawala et al. (2015, 2016) and Oman et al. (2015). Here

we describe the selection procedure for the resimulation volumes together with the physical processes included in the simulations. We assess the viability of these candidates by analysing, in Section 4, their satellite populations and comparing them with observations of the satellites of the MW and M31 galaxies.

3.1 The code

The APOSTLE simulations were run using a highly modified version of the N -body/smoothed particle hydrodynamics (SPH) code, P-GADGET3 (Springel 2005). The code includes subgrid prescriptions for star formation, feedback from evolving stars and supernovae, metal enrichment, cosmic reionization, and the formation and energy output from supermassive black holes/active galactic nuclei (AGN). The details of the code and the subgrid physics model are described in detail in Schaye et al. (2015) and Crain et al. (2015). When used to evolve a cosmologically significant volume, this code gives a galaxy stellar mass function in good agreement with observations in the range of galaxy stellar masses between $\sim 10^8$ and $\sim 10^{11} M_{\odot}$ (Schaye et al. 2015). The APOSTLE runs use the same parameter choices as the ‘reference’ model in Schaye et al. (2015). As discussed below, we find good numerical convergence in the galaxy properties we investigate, without further calibration.

3.2 Candidate selection

The pairs selected for resimulation in the APOSTLE project were drawn from the DOVE⁴ cosmological N -body simulation described by Jenkins (2013). DOVE evolved a periodic box 100 Mpc on a side assuming cosmological parameters consistent with the WMAP-7 estimates and summarized in Table 1. DOVE has 134 times better mass resolution than MS-I, with a particle mass of $8.8 \times 10^6 M_{\odot}$ (comparable to MS-II). Halo pairs were identified in DOVE using the procedure described in Section 2.2.

Guided by the discussion in Section 2, we chose 12 different pairs from the MedIso sample that satisfied, at $z = 0$, the following conditions:

- (i) separation between 600 and 1000 kpc;
- (ii) relative radial velocity, V_r , in the range $(-250, 0) \text{ km s}^{-1}$;
- (iii) relative tangential velocity, V_t , less than 100 km s^{-1} ;
- (iv) recession velocities of outer LG members in the range defined by the box in Fig. 3;
- (v) total pair mass (i.e. the sum of the virial masses of the two primary haloes) in the range $\log(M_{\text{tot}}/M_{\odot}) = [12.2, 12.6]$.

The relative velocities and separations of the 12 pairs are shown in Fig. 5, where each pair is labelled with a number from 1 to 12 for future reference. The main properties of these pairs are listed in Table 2 and the histogram of their total masses is shown (in black) in the bottom panel of Fig. 4. The parameters of fits to the recession velocities of the outer members are shown by the open starred symbols in Fig. 3.

Note that the 12 pairs chosen span a relatively small range of masses compared to what is allowed according to the kinematic constraints described in the previous section. The lowest mass of the pairs is $1.6 \times 10^{12} M_{\odot}$ and the most massive pair is $3.6 \times 10^{12} M_{\odot}$, with an average mass of $2.3 \times 10^{12} M_{\odot}$ and an rms of only 30 percent. The small mass range of the pairs was chosen in order to explore the cosmic variance at fixed mass. Given the large




Figure 5. Relative radial (squares) and tangential (circles) velocities versus separation for all 12 candidates selected for resimulation in the APOSTLE project. The corresponding values for the MW–M31 pair are shown by the starred symbols. The shaded area indicates the region where we would expect the radial velocities of all 12 candidates to lie if the predictions of the timing argument held.

range of allowed masses it would have been impossible to cover the whole range with high-resolution resimulations.

It is interesting to compare the masses of the pairs with those estimated from the timing argument applied to each pair. This is shown by the shaded region of Fig. 5, which brackets the region in distance–radial velocity plane that should be occupied by the 12 candidates we selected, given their total masses. The timing argument does reasonably well in eight out of the 12 cases, but four pairs fall outside the predicted region. Three pairs, in particular, have approach velocities as high or higher than observed for the MW–M31 pair (pairs 7, 8, and 11), demonstrating that our mass choice does not preclude relatively high approach velocities in some cases.

3.3 Resimulation runs

Volumes around the 12 selected pairs were carefully configured in the initial conditions of the DOVE simulation so that the inner 2–3 Mpc spherical regions centred on the mid-point of the pairs contain no ‘boundary’ particles at $z = 0$ (for details of this ‘zoom-in’ technique, see Power et al. 2003; Jenkins 2013). For each candidate pair, initial conditions were constructed at three levels of resolution using second-order Lagrangian perturbation theory (Jenkins 2010). (Further details are provided in Appendix A.) Resolution levels are labelled L1, L2, and L3 (high, medium, and low resolution), where L3 is equivalent to DOVE, L2 is a factor of ~ 10 improvement in mass resolution over L3, and L1 is another factor of ~ 10 improvement over L2. The method used to make the zoom initial conditions has been described in Jenkins (2013). Particle masses, softening lengths, and other numerical parameters for each resolution level are summarized in Table 2.

Each run was performed twice, with one run neglecting the baryonic component (i.e. assigning all the matter to the dark matter component, hereafter referred to as dark-matter-only or ‘DMO’ runs) and the other using the EAGLE code described in Section 3.1. Note that the cosmology assumed for the APOSTLE runs is the same as

⁴ DOVE simulation is also known as COLOR.

Table 2. The parameters of the APOSTLE resimulations. The first two columns list labels identifying each run. The following columns list the virial masses of each of the primaries at $z = 0$; their relative separation, radial velocity, and tangential velocity in the DOVE simulation, as well as the initial baryonic mass per particle in the hydrodynamical runs. The dark matter particle mass is $m_{\text{DM}} = (1/f_{\text{bar}} - 1)m_{\text{gas}}$, where f_{bar} is the universal baryon fraction. (Dark matter-only runs have a particle mass equal to the sum of $m_{\text{gas}} + m_{\text{DM}}$.) The last column lists the value of the Plummer-equivalent gravitational softening, which is comoving at early times, but fixed at the listed value after $z = 3$.

Name	Run (resolution)	$M_{200}^{[1]}$ [$10^{12}(M_{\odot})$]	$M_{200}^{[2]}$ [$10^{12}(M_{\odot})$]	Separation (kpc)	V_r (km s^{-1})	V_t (km s^{-1})	m_{gas} [$10^4(M_{\odot})$]	ϵ_{max} (pc)
AP-1	L1/L2/L3	1.66	1.10	850	-51	35	0.99/12.0/147	134/307/711
AP-2	L2/L3	0.85	0.83	809	-39	97	12.5/147	307/711
AP-3	L2/L3	1.52	1.22	920	-35	84	12.5/147	307/711
AP-4	L1/L2/L3	1.38	1.35	790	-59	24	0.49/12.2/147	134/307/711
AP-5	L2/L3	0.93	0.87	828	-33	101	12.5/147	307/711
AP-6	L2/L3	2.36	1.21	950	-18	60	12.7/137	307/711
AP-7	L2/L3	1.88	1.09	664	-174	24	11.3/134	307/711
AP-8	L2/L3	1.72	0.65	817	-120	96	11.0/137	307/711
AP-9	L2/L3	0.96	0.68	814	-28	48	10.9/138	307/711
AP-10	L2/L3	1.46	0.87	721	-63	48	11.0/146	307/711
AP-11	L2/L3	0.99	0.80	770	-124	22	11.1/153	307/711
AP-12	L2/L3	1.11	0.58	635	-53	50	10.9/138	307/711

that of DOVE (see Table 1), which differs slightly from that of EAGLE. These differences, however, are very small, and are expected to have a negligible effect on our results. At the time of writing, only two volumes had been completed at the highest (L1) resolution: AP-1 and AP-4.

4 LOCAL GROUP SATELLITES

In this section we explore the properties of the satellite population that surrounds each of the two primary galaxies of the APOSTLE resimulations to assess whether our choice of total mass and subgrid physics gives results consistent with the broad properties of the satellite systems of the MW and M31. We examine consistency with properties that may be considered problematic, given that our choice for the total mass is significantly lower than indicated by the timing argument. In particular, we focus on (i) whether the total number of satellites brighter than $M_V = -8$ and -9.5 (or, equivalently, stellar mass,⁵ M_{gal} , greater than 1.4×10^5 or $5.6 \times 10^5 M_{\odot}$, respectively) is comparable to the observed numbers around the MW and M31; (ii) whether satellites as massive/bright as the Large Magellanic Cloud or M33 are present; (iii) whether the velocity dispersion of the satellites is consistent with observations; and, finally, (iv) whether the observed satellites within 300 kpc of either primary (i.e. the MW or M31) are gravitationally bound given our choice of total mass for the LG candidates.

4.1 Satellite masses/luminosities

4.1.1 Number of bright satellites

The total number of satellites above a certain mass is expected to be a sensitive function of the virial masses of the primary haloes (see e.g. Boylan-Kolchin et al. 2011; Wang et al. 2012), so we begin by considering the number of satellites whose stellar masses exceed

$1.4 \times 10^5 M_{\odot}$ or, equivalently,⁶ that are brighter than $M_V = -8$, as a function of the virial mass, M_{200} , of each host. This brightness limit corresponds roughly to the faintest of the ‘classical’ dwarf spheroidal companions of the MW, such as Draco or Ursa Minor. We choose this limit because there is widespread consensus that surveys of the surroundings of the MW and M31 are complete down to that limit (see e.g. Whiting et al. 2007; McConnachie et al. 2009; McConnachie 2012).

We compare in the top panel of Fig. 6 the number of simulated satellites within 300 kpc of each primary galaxy to the numbers observed around the MW and M31 (shown as horizontal line segments). Satellite numbers correlate strongly with virial mass, as expected, and it is encouraging that the observed number of satellites of the MW and M31 (shown by short horizontal line segments in Fig. 6) are well within the range spanned by our simulations. Since a lower mass limit of $M_{\text{gal}} = 1.4 \times 10^5 M_{\odot}$ corresponds to just a few particles at L2 resolution, we repeat the exercise in the bottom panel of Fig. 6, but for a higher mass limit, i.e. $M_{\text{gal}} > 5.6 \times 10^5 M_{\odot}$ ($M_V < -9.5$). The results in either panel are reassuringly similar.

The impact of numerical resolution may be seen by comparing the results for the medium resolution (L2; blue ‘+’ symbols) with those obtained for the L1 (high-resolution) runs in Fig. 6 (red crosses). On average, the number of satellites increases by only ~ 10 per cent when increasing the mass resolution by a factor of 10, indicating rather good convergence. One of the haloes in L1, however, hosts almost twice as many satellites as its counterpart in L2. The reason is that a relatively large group of satellites has just crossed inside the 300 kpc boundary of the halo in L1, but is still outside 300 kpc in L2. (We do not consider the low-resolution L3 runs in this plot because at that resolution the mass per particle is $1.5 \times 10^6 M_{\odot}$ and satellites fainter than $M_V \sim -13$ are not resolved, see Table 2.)

We also note that, had we chosen larger masses for our LG primaries, consistent with the timing argument and abundance matching, our simulations would have likely formed a much larger number of bright satellites than observed. For example, for a virial mass

⁵ Stellar masses are measured within a radius, r_{gal} , equal to 15 per cent of the virial radius, r_{200} , of the surrounding halo. For satellites we estimate r_{gal} from its peak circular velocity, V_{max} , and the relation between V_{max} and r_{200} for isolated galaxies.

⁶ We assume a constant mass-to-light ratio of 1 in solar units for converting stellar masses to V-band luminosities. This is done only for simplicity, since our primary aim is to assess consistency rather than to provide quantitative predictions. Later work will use spectrophotometric models and internal extinction as laid out in Trayford et al. (2015).




Figure 6. Top: number of satellites with stellar mass greater than $1.4 \times 10^5 M_{\odot}$ (or brighter than $M_V = -8$) found within 300 kpc of each of the primary galaxies in our APOSTLE runs plotted versus the primary’s virial mass. We show results for medium resolution (L2, blue plus symbols) and high resolution (L1, red crosses). The thin dashed lines connect the L1 haloes to their counterparts in L2. The solid, coloured lines indicate the best fit to L2 data with unit slope; the shaded areas indicate the rms range about this fit. Small horizontal lines indicate the observed numbers for the MW and M31. Arrows indicate virial masses estimated for the MW and M31 from abundance matching of Guo et al. (2010). The ‘error bar’ around each arrow spans different predictions by Behroozi et al. (2013) and Kravtsov et al. (2014). Bottom: same as the top panel but for satellites with stellar mass greater than $5.6 \times 10^5 M_{\odot}$ (or brighter than $M_V = -9.5$).

of $\sim 2 \times 10^{12} M_{\odot}$ (the value estimated for the MW by abundance-matching analyses; see the upward arrow in Fig. 6) our simulations give, on average, 25 satellites brighter than $M_V = -8$ and 15 satellites brighter than $M_V = -9.5$. These are well in excess of the 12 and eight satellites, respectively, found in the halo of the MW. The same conclusion applies to M31, where abundance matching suggests a halo mass of order $7 \times 10^{12} M_{\odot}$. It is clear from Fig. 6 that our simulations would have produced a number of satellites well in excess of that observed around M31 for a halo as massive as that.

The correlation between host virial mass and satellite number is, of course, sensitive to our choice of galaxy formation model, and it would be possible, in principle, to reduce the number of bright satellites by reducing the overall galaxy formation efficiency in low-mass haloes. It is interesting that the same galaxy formation




Figure 7. Stellar mass of the primary galaxies (circles) and their brightest satellites (crosses) for our 12 APOSTLE simulations, as a function of their maximum circular velocity. The large black solid symbols denote observational estimates for the two brightest members of the LG and their brightest satellites: MW, M31, LMC, and M33, respectively. Grey curves indicate ‘abundance matching’ predictions from Guo et al. (2010), Behroozi et al. (2013), and Kravtsov et al. (2014). Note that satellites as bright and massive as the LMC and M33 are not uncommon in our simulations. Note also that our primaries tend to be undermassive relative to observations and to predictions from abundance matching. The reverse applies at lower circular velocities. Black dashed lines connect matching systems at different resolution levels. See text for further discussion.

model able to reproduce the galaxy stellar mass function in large cosmological volumes (Schaye et al. 2015) can reproduce, without further tuning, the number of satellites in the APOSTLE resimulations.

4.1.2 Most massive satellites

Another concern when adopting a relatively low virial mass for the LG primaries (compared to timing argument and abundance-matching estimates) is that satellites as bright and massive as M33 in the case of M31, or the LMC in the case of the MW, may fail to form. We explore this in Fig. 7, where we plot the stellar mass of the primaries (open circles) and that of their most massive satellite found within 300 kpc (crosses, ‘plus’, and square symbols for L1, L2, and L3, respectively) as a function of their peak circular velocity. We also plot, for reference, the rotation velocity and stellar masses of the MW and M31, as well as those of their brightest satellites, the LMC and M33, respectively. For the LMC, the error bar indicates the velocity range spanned by two different estimates, 64 km s^{-1} from van der Marel et al. (2002), and 87 km s^{-1} from Olsen et al. (2011). For M33, we use a rotation speed of 110 km s^{-1} , with an error bar that indicates the range of velocities observed between 5 and 15 kpc by Corbelli et al. (2014): we use this range as an estimate of the uncertainty because the gaseous disc of M33 has a very strong warp in the outer regions which hinders a proper determination of its asymptotic circular speed.

Fig. 7 makes clear that there is no shortage of massive satellites in the APOSTLE simulations: six out of 24 primaries in the L2 runs have

a massive subhalo with V_{\max} exceeding 100 km s^{-1} (comparable to M33), and 12 out of 24 have $V_{\max} > 60 \text{ km s}^{-1}$ (comparable to the LMC). This result is robust to numerical resolution; the numbers above change only to 7 and 11, respectively, when considering the lower resolution L3 runs (see open squares in Fig. 7).

With hindsight, this result is not entirely surprising. According to Wang et al. (2012), the *average* number of satellites scales as $10.2(\nu/0.15)^{-3.11}$, where $\nu = V_{\max}/V_{200}$ is the ratio between the maximum circular velocity of a subhalo and the virial velocity of the primary halo. A halo of virial mass $10^{12} M_{\odot}$ has $V_{200} = 145 \text{ km s}^{-1}$, so we expect, on average, that one in 11 of such haloes should host a satellite as massive as M33 and one in 2.3 one like the LMC.⁷ Our results seem quite consistent with this expectation.

A second point to note from Fig. 7 is that the stellar mass of the primary galaxies is below recent estimates for either the MW ($5.2 \times 10^{10} M_{\odot}$, according to Bovy & Rix 2013) or M31 ($1.03 \times 10^{11} M_{\odot}$, according to Hammer et al. 2007). Indeed, none of our 24 primaries has masses exceeding $5 \times 10^{10} M_{\odot}$, a consequence of the low virial masses of our selected LG pairs coupled with the relatively low galaxy formation efficiency of the EAGLE ‘Ref’ model in $\sim 10^{12} M_{\odot}$ haloes. This may be seen in Fig. 7, where the open circles lie systematically below the grey lines, which show abundance matching predictions taken from three recent papers.

This issue has been discussed by Schaye et al. (2015), and is reflected in their fig. 4, which shows that the EAGLE ‘Ref’ model we use here underpredicts the number of galaxies with stellar masses a few times $10^{10} M_{\odot}$. That same figure shows that the opposite is true for smaller galaxies: the ‘Ref’ model actually overpredicts slightly but systematically the number of galaxies with stellar masses a few times $10^8 M_{\odot}$.

Galaxy formation efficiency in our runs thus seems slightly too low in MW-sized haloes, and slightly too high in systems of much lower mass. This slight mismatch in the galaxy mass–halo mass relation manifests itself more clearly on LG scales, making it difficult to match simultaneously the stellar masses of the MW and M31 as well as that of their luminous satellites. Indeed, increasing the halo masses of our LG candidates would lead to a better match to the stellar masses of the primaries, but at the expense of overpredicting the number of bright satellites (see e.g. Fig. 6), unless the star formation and feedback parameters are recalibrated.

The stellar masses of LG galaxies are therefore a sensitive probe of the galaxy formation efficiency on cosmological scales, and provide an important constraint on the ability of cosmological codes to reproduce the observed galaxy population.

4.2 Satellite kinematics

We can also use the kinematics of MW and M31 satellites to gauge the consistency of our results with the simulated satellite population. We begin by considering, in the top panel of Fig. 8, the radial velocity dispersion, σ_r , of all satellites more massive than $1.4 \times 10^5 M_{\odot}$ (brighter than $M_V = -8$) within 300 kpc from either primary. We expect the satellite velocity dispersion to scale as $\sigma_r \propto V_{200} \propto M_{200}^{1/3}$, so the solid line indicates the best fit with that slope to the data for our 24 systems. (Symbols are as in Fig. 6; blue + symbols indicate L2 resolution, red crosses correspond to L1 resolution.) This scaling




Figure 8. Top: radial velocity dispersion of luminous satellites within 300 kpc of each primary in the APOSTLE runs, as a function of the host’s virial mass. The red crosses and blue plus signs correspond to the resolution levels, L1 and L2, respectively. The thin dashed lines connect haloes of different resolutions. The solid coloured line indicates the best power-law fit $\sigma \propto M^{1/3}$; the shaded areas indicate the rms around the fit. Note the strong correlation between velocity dispersion and host mass; the relatively low radial velocity dispersion of the MW satellites is best accommodated with a fairly low virial mass, lower than expected from abundance matching (see upward pointing arrow, the same as in Fig. 6). Bottom: same as top, but for the line-of-sight velocity dispersion of one satellite system as seen from the other primary. Compared with the observed line-of-sight velocity dispersion of M31 satellites, shown by the horizontal line, this measure favours a lower virial mass than inferred from abundance-matching analyses, indicated by the arrow.

describes the correlation shown in Fig. 8 fairly well, and its best fit suggests that a $10^{12} M_{\odot}$ system should host a satellite system with a radial velocity dispersion of $98 \pm 17 \text{ km s}^{-1}$. (The shaded area in Fig. 8 indicates the rms scatter about the $M^{1/3}$ best fit.) The observed dispersion of the MW satellites (86 km s^{-1} ; shown by the dashed horizontal line) thus suggests a mass in the range $3.9 \times 10^{11} - 1.1 \times 10^{12} M_{\odot}$. The scatter, however, is large and these data alone can hardly be used to rule out larger virial masses.

The bottom panel of Fig. 8 shows a similar analysis, but applied to the satellites of M31. Since M31-centric radial velocities are not

⁷ These numbers assume a maximum circular velocity of 100 km s^{-1} for M33 and 60 km s^{-1} for the LMC.

directly available from observations,⁸ we repeat the analysis for our 24 satellite systems using *projected* velocities, measured along the line of sight from the other primary of each pair. (The analysis uses only satellites within 300 kpc from the galaxy’s centre.) The relation between the projected velocity dispersion, $\sigma_{v_{\text{los}}}$, and virial mass is again reasonably well described by the expected $M_{200}^{1/3}$ scaling. In this case, the observed projected dispersion of 99 km s^{-1} for the M31 satellites suggests a mass in the range $7.6 \times 10^{11} - 2.1 \times 10^{12} M_{\odot}$ although the scatter is again large.

The main conclusion to draw from Fig. 8 is that the observed velocity dispersions of both the MW and M31 satellites are well within the ranges found in our simulations. The relatively low mass of our pairs thus does not seem to pose any problems reproducing the kinematics of the LG satellite population. On the contrary, as was the case for satellite numbers, the satellite kinematics would be difficult to reconcile with much larger virial masses. For example, if the MW had a virial mass of $1.8 \times 10^{12} M_{\odot}$, as suggested by abundance matching (Guo et al. 2010), then its satellite systems would, on average, have a velocity dispersion of order 120 km s^{-1} , which would exceed the 86 km s^{-1} observed for the MW companions.

Finally, we check whether any of the MW or M31 satellites would be unbound given the mass of the primaries chosen for our sample. We explore this in Fig. 9, where we plot the observed velocity (radial in the case of the MW; line-of-sight in the case of M31) of satellites as a function of their distance to the primary’s centre (solid symbols).

The dotted lines delineate the escape velocity as a function of distance for haloes with virial masses of 6×10^{11} and $2 \times 10^{12} M_{\odot}$, corresponding to roughly the minimum and maximum virial masses of all primaries in our sample. The escape velocities assume a Navarro–Frenk–White (NFW) profile (Navarro, Frenk & White 1996, 1997), with a concentration $c = 10$. ‘Escapers’ (i.e. satellites with 3D velocities exceeding the escape speed for its primary) are shown by circled symbols. These are rare; only three of the $439 M_V < -8$ satellites examined are moving with velocities exceeding the nominal NFW escape velocity of their halo (see also Boylan-Kolchin et al. 2013).

It is clear from this figure that, for our choice of masses, none of the MW or M31 satellites would be unbound given their radial velocity. Indeed, even Leo I and And XIV (the least bound satellites of the MW and M31, respectively) are both within the bound region, and, furthermore, in a region of phase space shared with many satellites in our APOSTLE sample. The kinematics of the satellite populations of MW and M31 thus seems consistent with that of our simulated satellite populations.

5 SUMMARY AND CONCLUSIONS

We have analysed the constraints placed on the mass of the LG by the kinematics of the MW–M31 pair and of other LG members. We used these constraints to guide the selection, from a large cosmological simulation, of 12 candidate environments for the EAGLE–APOSTLE project, a suite of hydrodynamical resimulations run at various numerical resolution levels (reaching $\sim 10^4 M_{\odot}$ per gas particle at the highest level) and aimed at studying the formation of galaxies in the local Universe.

⁸ The M31-centric radial velocities can in principle be inferred from the data, but only by making further assumptions; see e.g. Karachentsev & Kashibadze (2006).




Figure 9. Top: radial velocity versus distance for all satellites in the medium-resolution (L2) APOSTLE runs (‘+’ symbols), compared with the observed radial velocities of the MW luminous satellites (filled squares). The positions of observed satellites in this phase space correspond to regions that are well populated in the simulations. Very few satellites are ‘unbound’ (marked by circles), as judged by escape velocity curves computed assuming a NFW profile of given virial mass (see dotted lines). Bottom: same as top but for line-of-sight velocities of satellites as seen from the other primary. This may be directly compared with data for M31 satellites (filled triangles).

APOSTLE uses the same code and star formation/feedback subgrid modules developed for the EAGLE project, which yield, in cosmologically representative volumes, a galaxy stellar mass function and average galaxy sizes in good agreement with observations. This ensures that any success of our simulations in reproducing LG-scale observations does not come at the expense of subgrid module choices that might fail to reproduce the galaxy population at large. We also compare the simulated satellite populations of the two main galaxies in the APOSTLE resimulations with the observed satellite systems of M31 and MW to assess consistency with observation.

Our main conclusions may be summarized as follows.

(i) The kinematics of the MW–M31 pair and of other LG members are consistent with a wide range of virial masses for the MW and M31. Compared with halo pairs selected from the Millennium Simulations, the relatively fast approach velocity of MW and M31 favours a fairly large total mass, of order $5 \times 10^{12} M_{\odot}$. On the

other hand, the small tangential velocity and the small deceleration from the Hubble flow of outer LG members argue for a significantly smaller mass, of order $6 \times 10^{11} M_{\odot}$. Systems that satisfy the three criteria are rare – only 14 are found in a $(137 \text{ Mpc})^3$ volume – and span a wide range of masses, from 2.3×10^{11} to $6.1 \times 10^{12} M_{\odot}$, with a median mass $\sim 2 \times 10^{12} M_{\odot}$.

(ii) Given the wide range of total masses allowed, the 12 candidate pairs selected for resimulation in the APOSTLE project were chosen to loosely match the LG kinematic criteria and to span a relatively narrow range of masses (from 1.6×10^{12} to $3.6 \times 10^{12} M_{\odot}$, with a median mass of $2.3 \times 10^{12} M_{\odot}$). This enables us to explore the cosmic variance of our results at fixed mass, and, potentially, to scale them to other mass choices, if needed.

(iii) Large satellites such as LMC and M33 are fairly common around our simulated galaxies, although their total virial mass is well below that estimated from the timing argument.

(iv) The overall abundance of simulated satellites brighter than $M_V = -8$ is a strong function of the virial mass assumed for the LG primary galaxies in our simulations. The relatively few (12) such satellites around the MW suggests a fairly low mass ($\sim 6 \times 10^{11} M_{\odot}$); the same argument suggests a mass for M31 about twice as large ($\sim 1.2 \times 10^{12} M_{\odot}$).

(v) The velocity dispersions of simulated satellites are consistent with those of the MW and M31. This diagnostic also suggests that virial masses much larger than those adopted for the APOSTLE project would be difficult to reconcile with the relatively low radial velocity dispersion observed for the MW satellite population, as well as with the projected velocity dispersion of the M31 satellite population.

(vi) The primary galaxies in the simulations are less massive than current estimates for the MW and M31. The most likely reason for this is an inaccuracy in the subgrid modelling of star formation and feedback and its dependence on halo mass.

Our overall conclusion is that, despite some shortcomings, the APOSTLE simulation suite should prove a wonderful tool to study the formation of the galaxies that populate our cosmic backyard.

ACKNOWLEDGEMENTS

We would like to thank Dr Lydia Heck for support through the computational facility at the ICC. We are grateful to Alan McConnachie for help in interpreting the observational data, and also Richard Bower for reading the manuscript and giving us useful feedback. We also would like to thank the anonymous referee for the constructive report. This research was supported by the National Science Foundation under Grant No. PHY11-25915 and the hospitality of the Kavli Institute for Theoretical Physics at UC Santa Barbara. This work was supported in part by the Science and Technology Facilities Council (grant number ST/F001166/1), European Research Council (grant numbers GA 267291 ‘Cosmiway’ and GA 278594 ‘Gas Around Galaxies’), the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013), the National Science Foundation under Grant No. PHYS-1066293, the Interuniversity Attraction Poles Programme of the Belgian Science Policy Office [AP P7/08 CHARM], and the hospitality of the Aspen Center for Physics. JFN acknowledges a Leverhulme Visiting Professor grant held at the Institute for Computational Cosmology, Durham University. TS acknowledges the Marie-Curie ITN CosmoComp. 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 grants ST/H008519/1 and ST/K00087X/1, STFC DiRAC Operations grant ST/K003267/1, and Durham University. DiRAC is part of the National E-Infrastructure.

REFERENCES

- Barber C., Starkeburg E., Navarro J. F., McConnachie A. W., Fattahi A., 2014, *MNRAS*, 437, 959
- Battaglia G. et al., 2005, *MNRAS*, 364, 433
- Behroozi P. S., Marchesini D., Wechsler R. H., Muzzin A., Papovich C., Stefanon M., 2013, *ApJ*, 777, L10
- Bertschinger E., 1985, *ApJS*, 58, 39
- Bovy J., Rix H.-W., 2013, *ApJ*, 779, 115
- Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, *MNRAS*, 398, 1150
- Boylan-Kolchin M., Besla G., Hernquist L., 2011, *MNRAS*, 414, 1560
- Boylan-Kolchin M., Bullock J. S., Sohn S. T., Besla G., van der Marel R. P., 2013, *ApJ*, 768, 140
- Cautun M., Frenk C. S., van de Weygaert R., Hellwing W. A., Jones B. J. T., 2014, *MNRAS*, 445, 2049
- Corbelli E., Thilker D., Zibetti S., Giovanardi C., Salucci P., 2014, *A&A*, 572, A23
- Crain R. A. et al., 2015, *MNRAS*, 450, 1937
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- Deason A. J. et al., 2012, *MNRAS*, 425, 2840
- Forero-Romero J. E., Hoffman Y., Bustamante S., Gottlöber S., Yepes G., 2013, *ApJ*, 767, L5
- Garrison-Kimmel S., Boylan-Kolchin M., Bullock J. S., Lee K., 2014, *MNRAS*, 438, 2578
- Gómez F. A., Besla G., Carpintero D. D., Villalobos Á., O’Shea B. W., Bell E. F., 2015, *ApJ*, 802, 128
- González R. E., Kravtsov A. V., Gnedin N. Y., 2014, *ApJ*, 793, 91
- Gottlöber S., Hoffman Y., Yepes G., 2010, preprint ([arXiv:1005.2687](https://arxiv.org/abs/1005.2687))
- Guo Q., White S., Li C., Boylan-Kolchin M., 2010, *MNRAS*, 404, 1111
- Hammer F., Puech M., Chemin L., Flores H., Lehnert M. D., 2007, *ApJ*, 662, 322
- Jenkins A., 2010, *MNRAS*, 403, 1859
- Jenkins A., 2013, *MNRAS*, 434, 2094
- Kahn F. D., Woltjer L., 1959, *ApJ*, 130, 705
- Karachentsev I. D., Kashibadze O. G., 2006, *Astrophysics*, 49, 3
- Karachentsev I. D., Makarov D. A., 1996, *AJ*, 111, 794
- Kravtsov A., Vikhlinin A., Meshcheryakov A., 2014, *ApJ*, preprint ([arXiv:1401.7329](https://arxiv.org/abs/1401.7329))
- Li Y.-S., White S. D. M., 2008, *MNRAS*, 384, 1459
- McConnachie A. W., 2012, *AJ*, 144, 4
- McConnachie A. et al., 2009, *Nature*, 461, 66
- McLaughlin D. E., 1999, *ApJ*, 512, L9
- Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493
- Olsen K. A. G., Zaritsky D., Blum R. D., Boyer M. L., Gordon K. D., 2011, *ApJ*, 737, 29
- Oman K. A. et al., 2015, *MNRAS*, 452, 3650
- Peebles P. J. E., Phelps S. D., Shaya E. J., Tully R. B., 2001, *ApJ*, 554, 104
- Peebles P. J. E., Tully R. B., Shaya E. J., 2011, preprint ([arXiv:1105.5596](https://arxiv.org/abs/1105.5596))
- Peñarrubia J., Ma Y.-Z., Walker M. G., McConnachie A., 2014, *MNRAS*, 443, 2204
- Piffi T. et al., 2014, *MNRAS*, 445, 3133
- Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, *MNRAS*, 338, 14
- Sales L. V., Navarro J. F., Abadi M. G., Steinmetz M., 2007, *MNRAS*, 379, 1464
- Sawala T. et al., 2015, *MNRAS*, 448, 2941
- Sawala T. et al., 2016, *MNRAS*, 456, 85

- Schaye J. et al., 2015, MNRAS, 446, 521
 Schönrich R., Binney J., Dehnen W., 2010, MNRAS, 403, 1829
 Smith M. C. et al., 2007, MNRAS, 379, 755
 Sohn S. T., Anderson J., van der Marel R. P., 2012, ApJ, 753, 7
 Sorce J. G., Courtois H. M., Gottlöber S., Hoffman Y., Tully R. B., 2014, MNRAS, 437, 3586
 Springel V., 2005, MNRAS, 364, 1105
 Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726
 Springel V. et al., 2005, Nature, 435, 629
 Tonry J. L., Dressler A., Blakeslee J. P., Ajhar E. A., Fletcher A. B., Luppino G. A., Metzger M. R., Moore C. B., 2001, ApJ, 546, 681
 Trayford J. W. et al., 2015, MNRAS, 452, 2879
 Tully R. B., Rizzi L., Shaya E. J., Courtois H. M., Makarov D. I., Jacobs B. A., 2009, AJ, 138, 323
 van der Marel R. P., Alves D. R., Hardy E., Suntzeff N. B., 2002, AJ, 124, 2639
 van der Marel R. P., Fardal M., Besla G., Beaton R. L., Sohn S. T., Anderson J., Brown T., Guhathakurta P., 2012, ApJ, 753, 8
 Wang J., Frenk C. S., Navarro J. F., Gao L., Sawala T., 2012, MNRAS, 424, 2715
 Wang W., Han J., Cooper A. P., Cole S., Frenk C., Lowing B., 2015, MNRAS, 453, 377
 Watkins L. L., Evans N. W., An J. H., 2010, MNRAS, 406, 264
 Whiting A. B., Hau G. K. T., Irwin M., Verdugo M., 2007, AJ, 133, 715
 Xue X. X. et al., 2008, ApJ, 684, 1143

APPENDIX A: PARAMETERS OF THE INITIAL CONDITIONS

The initial conditions for the DOVE and APOSTLE simulations were generated from the PANPHASIA white noise field (Jenkins 2013) using second-order Lagrangian perturbation theory (Jenkins 2010). The coordinates of the centres and the radii of the high-resolution Lagrangian regions in the initial conditions, as well as the positions of the MW and M31 analogues at $z = 0$, for the 12 APOSTLE volumes are given in Table A1.

Table A1. The positions of main haloes at $z = 0$ and parameters of the high-resolution Lagrangian regions of the APOSTLE volumes in the initial conditions. The first column labels each volume. The next columns list the (x,y,z) coordinates of each of the primaries at $z = 0$. The final four columns give the comoving coordinate centre and radius of a sphere that contains the high-resolution Lagrangian region in the initial conditions, for each of the zoom initial conditions. The phase descriptor for the APOSTLE runs is, in PANPHASIA format, [Panph1,L16,(31250,23438,39063),S12,CH1292987594,DOVE].

Name	X_1 (Mpc)	Y_1 (Mpc)	Z_1 (Mpc)	X_2 (Mpc)	Y_2 (Mpc)	Z_2 (Mpc)	X_1 (Mpc)	Y_1 (Mpc)	Z_1 (Mpc)	R_1 (Mpc)
AP-1	19.326	40.284	46.508	18.917	39.725	47.001	26.5	39.1	39.0	7.9
AP-2	28.798	65.944	17.153	28.366	65.981	16.470	28.1	60.2	18.4	14.3
AP-3	51.604	28.999	11.953	51.091	28.243	12.061	46.0	31.7	11.6	16.8
AP-4	63.668	19.537	72.411	63.158	20.137	72.467	57.1	20.6	74.9	8.4
AP-5	42.716	87.781	93.252	42.872	88.478	93.671	40.8	85.4	91.8	13.8
AP-6	35.968	9.980	43.782	36.171	9.223	43.251	32.9	13.1	45.2	9.9
AP-7	91.590	43.942	14.826	91.822	43.323	14.885	99.3	39.7	15.9	12.9
AP-8	4.619	22.762	85.535	4.604	23.508	85.203	4.9	20.4	89.9	9.9
AP-9	57.044	88.490	74.765	57.496	87.889	74.456	55.2	93.4	76.5	7.9
AP-10	61.949	24.232	98.305	61.867	24.925	98.124	62.5	24.5	93.5	11.7
AP-11	12.564	48.080	35.249	12.484	47.793	35.959	18.3	43.1	29.9	8.2
AP-12	97.553	89.587	72.093	97.351	90.100	72.407	98.5	91.9	81.9	7.9

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