Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **446,** 2274–2290 (2015) doi:10.1093/mnras/stu2257

Creating mock catalogues of stellar haloes from cosmological simulations

Ben Lowing,[1] Wenting Wang,[1⋆] Andrew Cooper,[1,2] Rachel Kennedy,[1] John Helly,[1] Shaun Cole[1] and Carlos Frenk[1]

[1]*Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK*
[2]*National Astronomical Observatories, Chinese Academy of Sciences, 20A Datun Road, Chaoyang, Beijing 10012, People's Republic of China*

Accepted 2014 October 27. Received 2014 October 18; in original form 2014 May 31

ABSTRACT

We present a new technique for creating mock catalogues of the individual stars that make up the accreted component of stellar haloes in cosmological simulations and show how the catalogues can be used to test and interpret observational data. The catalogues are constructed from a combination of methods. A semi-analytic galaxy formation model is used to calculate the star formation history in haloes in an *N*-body simulation and dark matter particles are tagged with this stellar mass. The tags are converted into individual stars using a stellar population synthesis model to obtain the number density and evolutionary stage of the stars, together with a phase-space sampling method that distributes the stars while ensuring that the phase-space structure of the original *N*-body simulation is maintained. A set of catalogues based on the Λ cold dark matter Aquarius simulations of Milky Way mass haloes have been created and made publicly available on a website. Two example applications are discussed that demonstrate the power and flexibility of the mock catalogues. We show how the rich stellar substructure that survives in the stellar halo precludes a simple measurement of its density profile and demonstrate explicitly how pencil-beam surveys can return almost any value for the slope of the profile. We also show that localized variations in the abundance of particular types of stars, a signature of differences in the composition of stellar populations, allow streams to be easily identified.

Key words: methods: numerical – galaxies: haloes.

1 INTRODUCTION

Extensive diffuse stellar haloes are known to surround both the Milky Way and M31 (e.g. Deason, Belokurov & Evans 2011; Ibata et al. 2014). The Λ cold dark matter (ΛCDM) model of hierarchical structure formation predicts that these haloes should be a ubiquitous outcome of galaxy formation resulting from the accretion of satellite galaxies that fall into the potential of the larger host galaxy and are subsequently disrupted (Bullock & Johnston 2005; Cooper et al. 2010). The dynamical time of the stellar halo is very long compared to the age of the host galaxy and so it preserves a memory of its initial conditions. From the identification of distinct populations within haloes, and the study of properties such as kinematics or chemical composition, a vast amount of information about the assembly history of the galaxy can be inferred.

The outer regions of stellar haloes are dominated by multiple, extensive tidal streams, while the inner regions, where dynamical time-scales are shorter, experience a more complete destruction of substructure and therefore have a smoother stellar distribution. In addition to the accreted component, stellar haloes are thought to contain a second component made up of stars that formed *in situ*. Recent observational claims of a two-component halo in the Milky Way have focused attention on this possibility (e.g. Carollo et al. 2010; Beers et al. 2012; Kafle et al. 2013), while the evidence for two components is disputed by Schönrich, Asplund & Casagrande (2014). From a theoretical perspective, there is a wide variety of possible formation mechanisms for *in situ* halo stars, many of which overlap with possible origins of the Galactic thick disc. These include the scattering of stars from the thin disc, the collapse of the disc at high redshift, star formation in gaseous tidal streams and thermal instabilities in a hydrostatic gas halo (e.g. Abadi, Navarro & Steinmetz 2006; Zolotov et al. 2010; Font et al. 2011a; Tissera et al. 2013; Parry et al., in preparation). Most of these mechanisms result in *in situ* components that dominate in the inner few kiloparsecs of the galaxy and remain important out to ∼30 kpc.

Over the last few years there has been much effort dedicated to modelling stellar haloes (e.g. Bullock & Johnston 2005; Cooper et al. 2010; Zolotov et al. 2010; Font et al. 2011a; Tissera, White & Scannapieco 2012) motivated by improvement in observations. Recent surveys such as SDSS SEGUE (Yanny et al. 2009) and PanSTARRS (Kaiser et al. 2010) have started to probe deeper into

⋆E-mail: bilinxing.wenting@gmail.com

the Milky Way halo and are starting to build up a picture of its complex structure. The *Gaia* mission will lead to further major improvements in the Milky Way data (Jordi et al. 2010; de Bruijne 2012) and new photometric surveys are extending these studies to other nearby galaxies (McConnachie et al. 2009; Martínez-Delgado et al. 2010). We are now at the point where careful comparisons between the haloes predicted by theory and those observed in the Universe are becoming possible.

In order to make a direct comparison of theoretical predictions with observations of stellar haloes, a model of stellar halo formation is required, along with a way to convert the output into observable quantities. Furthermore, to obtain realistic accretion histories, a complete model of galaxy formation in a cosmological context is essential. Various studies have focused on supplementing *N*-body simulations with analytical modelling of star formation and chemical enrichment (e.g. Bullock, Kravtsov & Weinberg 2001; Robertson et al. 2005; Font et al. 2006; De Lucia & Helmi 2008; Johnston et al. 2008; Tumlinson 2010). Alternatively, cosmological hydrodynamical simulations have also been used to study the overall structure and assembly of stellar haloes (e.g. Crain et al. 2009; Zolotov et al. 2010; Libeskind et al. 2011; Tissera et al. 2012; Pillepich et al. 2014), although these do not yet have the capability to resolve the detailed kinematic and spatial structure of individual halo stellar populations. Even 'zoom' simulations of individual galaxies represent the stellar halo with a relatively small number of star particles, since this component amounts to only a few per cent of the total stellar mass.

In this paper, we employ the particle tagging method developed by Cooper et al. (2010, hereafter C10) to construct stellar halo models. The method is based on high-resolution dark matter simulations in which the star formation occurring in haloes is calculated using a semi-analytic galaxy formation model. At every output time in the simulation, an appropriate set of dark matter particles is tagged with the stellar mass that formed since the previous output time. The advantages of this technique are that it can create models of much higher resolution and at a much lower computational cost than full hydrodynamical simulations. C10 demonstrated that the method produces complex stellar haloes containing tidal streams, shells and other substructure, whose ages and metallicities are in broad agreement with recent observations.

A limitation of the C10 tagging method is that 'star particles' are not individual stars but rather represent entire stellar populations which can contain thousands of solar masses of stars. The challenge is to compare these sparsely sampled haloes to actual observational surveys which often focus on a particular type of star, such as blue horizontal branch (BHB) stars. These may not trace the overall stellar mass distribution in an unbiased way. To overcome this limitation, we use a method similar to that employed in GALAXIA (Sharma et al. 2011) to construct customizable mock catalogues of halo stars. Each star particle is treated as a separate simple stellar population using theoretical isochrones to determine the abundance of each type of star expected in a population of the corresponding age and metallicity. To create positions and velocities of individual stars, the star particle distribution is oversampled using a phase-space kernel method based on the ENBID code (Sharma & Steinmetz 2006) which maintains the underlying phase-space structure.

Our mock catalogues allow us to compare the structure of simulated stellar haloes to observational data in a more realistic way. For example, obscuration by the Milky Way disc, survey limits imposed by the high cost of deep photometry and spectroscopy and the low surface density of halo tracers mean that the structure of the Galactic stellar halo is only well sampled to large distances in small patches of the sky. Nevertheless, attempts are often made to infer the profile of the Milky Way stellar halo from these limited-area observations (e.g. Watkins et al. 2009; Sesar, Jurić & Ivezić 2011; Akhter et al. 2012). In this paper, we show how the complicated uneven structure of the outer accreted stellar halo results in large fluctuations among measurements of its global properties based on observations of small sky patches. We also show how different satellites that build up the halo create observable variations in the stellar populations across the sky, for example in the ratio of stellar types. This variation can potentially help to detect new halo substructures in the Milky Way.

In Section 2, we outline the methods used to create stellar halo models. The final output of the process, a set of mock catalogues of individual stars and their properties, is described in further detail in Section 3. Sections 4 and 5 present two simple example applications of the mock catalogues to study the overall structure of stellar haloes in our simulations and the distribution of different stellar types over the sky. Finally, in Section 6 we discuss some potential advanced applications of the mock catalogues and how they can be used to interpret and test observational strategies.

2 METHOD

Creating a model of stellar haloes based on a consistent theory of galaxy formation in a cosmological context requires a combination of various techniques. The foundation of our models is a set of *N*-body dark matter only 'zoom' cosmological simulations of Milky Way mass haloes on to which we have grafted the Durham semi-analytic model of galaxy formation, GALFORM (Cole et al. 2000). GALFORM predicts the amount of stars that form in every halo and subhalo at each output time, along with an estimate of the metallicity of each stellar population. This stellar mass is then assigned to dark particles within the *N*-body simulations using the C10 particle tagging technique, which allows us to track the accretion and disruption of satellites and follow the fate of the stripped stars. Finally, to convert the massive simulation particles into mock catalogues (mocks for short) of individual stars, we use a method based on stellar population modelling and phase-space sampling. Each of these steps is briefly outlined in the following subsections.

2.1 *N*-body simulations

For this work, we use the haloes from the Aquarius project, which simulates six dark matter haloes of mass $\sim 10^{12}\,\mathrm{M}_{\odot}$ at multiple resolution levels (Springel et al. 2008a,b; Navarro et al. 2010). The simulations assume the standard ΛCDM cosmology with parameters chosen to be consistent with the results from the *WMAP* 1-year data (Spergel et al. 2003) and the 2dF Galaxy Redshift Survey data (Colless et al. 2001): matter density parameter, $\Omega_{\mathrm{M}} = 0.25$; cosmological constant, $\Omega_{\Lambda} = 0.75$; power spectrum normalization, $\sigma_8 = 0.9$; spectral index, $n_{\mathrm{s}} = 1$; and Hubble parameter $h = 0.73$. The six haloes were selected randomly from a set of isolated $\sim 10^{12}\,\mathrm{M}_{\odot}$ haloes identified in a lower resolution (900^3-particle) parent simulation of a $100h^{-1}$ Mpc cubic volume (Gao et al. 2008). The isolation criterion requires a halo to have no neighbours exceeding half its mass within $1\,h^{-1}$ Mpc. This weak selection criterion ensures that the haloes are not members of any massive groups or clusters.

The Aquarius simulations are labelled Aq-A to Aq-E; in each case we use the second highest resolution ('level-2' in the Aquarius notation) simulation, with a particle mass of at most $10^4\,h^{-1}\,\mathrm{M}_{\odot}$. We have used only five of the six Aquarius haloes, omitting halo

Aq-F which undergoes two major mergers at $z \sim 0.6$ and is thus highly unlikely to host a disc galaxy at $z = 0$. For each simulation, we have 128 output times, with a constant spacing of 155 Myr after $z \approx 2.5$.

2.2 Galaxy formation model

Dark matter *N*-body simulations are ideal for following the non-linear growth of cosmic structure, including the detailed accretion history of a halo. However, including a separate, realistic treatment of baryons in a cosmological simulation at sufficient resolution for stellar halo studies is still computationally prohibitive. Semi-analytic galaxy formation modelling permits baryonic physics to be included in much higher resolution *N*-body simulations at low computational cost, and is thus ideal for this purpose. Based on simple theoretical treatments and empirical prescriptions, semi-analytic models can provide a good match to a remarkable number of observations of the local and distant galaxy populations (Baugh 2006). These models, however, are not designed to follow the internal structure of galaxies in detail, hence the need for the particle tagging extension described below.

The Durham semi-analytic model, GALFORM, is used in this work to post-process the Aquarius *N*-body simulations. GALFORM computes the mass of stars that forms in each halo between every two successive simulation outputs, along with the properties of these stars, such as their metallicity. Rather than using the Bower et al. (2006) model on which C10 was based, we use the more recent Font et al. (2011b) model, which includes a modified treatment of physics relevant to satellites. The only significant changes compared to the C10 model are the use of a higher yield, a modified feedback model for supernovae in which the mass ejection efficiency saturates in haloes with circular velocity $V_{\rm circ} \leq 65$ km s^{-1}, and an earlier epoch of hydrogen reionization. The Font et al. (2011b) model produces a satellite luminosity function that matches that of Milky Way and also generates stellar populations in satellites which match the observed luminosity–metallicity relation. This eliminates a shortcoming of C10 where the overall metallicity of the stellar halo was found to be significantly lower than that of the Milky Way's stellar halo.

2.3 Particle tagging

The C10 particle tagging method is a means to associate the stellar mass calculated by GALFORM with six-dimensional phase-space volumes, defined by carefully chosen sets of representative particles in high-resolution dark matter only simulations. This allows the 3D spatial distribution and kinematics of stellar populations to be followed, albeit approximately, without including gas physics in the original simulation.

After applying GALFORM to the outputs of the Aquarius simulations, C10 associated each newly formed stellar population (discretized in age by simulation output times) with the 1 per cent most tightly bound dark matter particles in their corresponding dark matter haloes. These sets of 'tagged' dark matter particles are chosen at the time when the population forms. Their diffusion in configuration and velocity space can then be tracked to the present day. This method is ideal for studying the formation of the outer accreted component of stellar haloes, which are thought to be created by the tidal disruption of satellites. The fraction of most-bound particles tagged, the free parameter introduced by the method, has an effect on the scale radii of the resulting galaxies. C10 fixed this value at 1 per cent in order to reproduce the distribution of observed half-mass radii for surviving satellites. Having fixed this scale, satellite profiles and velocity dispersions also match well to observations (see C10).

Most surviving satellites are highly dark matter dominated (e.g. Walker et al. 2007). Since the *N*-body potential is not altered by the gravity of the 'stars' represented by the tags, the technique is likely to be much less accurate in baryon-dominated systems, such as the centre of the Milky Way analogue (see C10, Cooper et al. 2013 and Cooper et al. 2014 for further details and discussion). Satellite growth rates, survival times and orbits may also be affected by a more self-consistent treatment baryons (e.g. Sawala et al. 2013). The relative importance of these effects is uncertain, however. For example, supernova feedback can create dark matter cores in haloes (e.g. Navarro, Eke & Frenk 1996; Pontzen & Governato 2013). This reduction in central density may reduce their survival time when they are subjected to tidal stripping (e.g. Peñarrubia et al. 2010). On the other hand, the central concentration of stars may also bind these galaxies more tightly in a self-consistent simulation – in this context we note that numerical softening of the gravitational force between particles creates artificial cores in our simulated haloes on scales of $\sim$100 pc.

To explore the consequences of such limitations, Bailin et al. (2014) compared tagged dark matter particles in hydrodynamical simulations to star particles formed self-consistently in the same simulations. They found moderate disagreement in the distributions of halo stars represented by these two sets of particles, and concluded that tightly bound dark matter is not a reliable proxy for the dynamics of stars in satellite galaxies. However, there is a fundamental difference between the tagging method tested by Bailin et al. (2014) and that used by C10. Bailin et al. (2014) carried out a single tagging procedure for each satellite at the time of its infall into the main halo, whereas C10 tag each stellar population at the time of its formation. Star particles undergo considerable diffusion in energy over time, particularly if their host halo is disturbed, and hence would not be expected to occupy the same region of phase space as tightly bound dark matter chosen long after they form. Dark matter particles with similar binding energy chosen at the time of star formation, on the other hand, will undergo similar diffusion to star particles. This diffusion is therefore captured in the C10 approach, but not in the method used by Bailin et al. (2014). Two forthcoming papers will test the tagging technique of C10 in a similar way (Le Bret et al., in preparation; Cooper et al., in preparation). Within the limitations and approximations of the method, mentioned above and in C10, these new studies find good correspondence between stars and tagged dark matter particles in hydrodynamical simulations.

2.4 Mock catalogue construction

The final step is to turn the output of the tagged simulations into full mock catalogues of individual stars. The output of the model is a collection of *N*-body particles, each tagged with some amount of stellar mass, and having both an age for when the stars were born and a metallicity at that time. When a single dark matter particle is tagged multiple times, multiple 'tags' are created with the same phase-space trajectory. Turning these tags into a full mock catalogue requires two steps. The first is to determine the total number of stars of different types that each tagged particle represents by considering each as a single-age stellar population, and the second step is to distribute the stars over the phase-space volume represented by the corresponding dark matter particle. The following subsections discuss each step in detail.

2.4.1 *Stellar population synthesis*

Unfortunately, even with the high-resolution Aquarius simulations we are a long way from being able to represent each individual star by a single *N*-body particle. Instead, in the C10 simulations, each stellar particle can represent up to $\sim 10^3\,\mathrm{M}_\odot$. In order to convert these massive stellar particles into individual stars we consider each tagged particle as a stellar population (SSP), in which all the stars formed at the same time from gas with the same metallicity; the theoretical justification of this is discussed in detail in Pasetto, Chiosi & Kawata (2012). The abundances of the different star types that make up the population can then be obtained using stellar population synthesis modelling.

The initial mass function (IMF) determines the number of stars forming within each stellar population as a function of initial stellar mass. Here we adopt a Kennicutt IMF with no correction for brown dwarfs (Kennicutt 1983), just as was used in the semi-analytical model, to ensure consistency. Theoretical stellar isochrones then specify whether each type of star still survives at the final time, along with the current properties of these stars.

Stellar isochrones allow us to assign properties such as temperature, magnitude and colour to stars of a given initial mass, age and metallicity. We employ the PARSEC isochrones (Bressan et al. 2012), as these are recent, up-to-date isochrones that extend over a wide range of metallicities ($0.0001 \leq Z \leq 0.06$). We create a grid of isochrones spanning the range of ages between $6.63 < \log_{10} t/\mathrm{yr} < 10.13$, with a step size ($\Delta \log(t) = 0.0125$), and a range of metallicities between $0.0001 \leq Z \leq 0.005$. Interpolating between ages and metallicities would require matching different evolutionary points along the isochrones, so instead we simply identify the nearest isochrone in the grid and associate that with a tag. A finer grid could be used if required, but the adopted resolution was found to be adequate for our current purposes.

The method can be summarized as follows: for each tag, split the population into a range of initial stellar mass bins. Use the IMF to determine the fractional weighting of each bin. Then find the closest isochrone in age and metallicity and use that to assign properties to the stars in each bin. Finally, multiply the mass of the tag by the weight of the bin to obtain the total mass and, from that, the number of stars in that bin. Since the lifetimes of massive stars are much shorter than the age of the Universe, only a fraction of the initial stellar mass of the tag will be converted to stars in the mock catalogue, the balance being returned to the interstellar medium or locked in stellar remnants.

2.4.2 *Phase-space sampling*

The next step is to split the massive stellar particles representing SSPs into stars with individual positions and velocities. Although the massive stellar particles provide a coarse sampling of the phase-space distribution of the stellar halo, it is desirable to smooth the distribution and increase the sampling. This is done by distributing stars over the entire phase-space volume that each tagged particle represents. Using the phase-space volume rather than just the real-space volume preserves coherent structures such as tidal streams. It must be noted that increasing the sampling does not increase the resolution and cannot reveal any more detailed structure than was in the original simulation.

To estimate the phase-space volume each stellar mass particle occupies, we have used the entropy-based binary decomposition (ENBID) code of Sharma & Steinmetz (2006). ENBID numerically estimates the densities of discretely sampled data based on a binary space partitioning tree. ENBID uses an entropy-based node splitting criterion to decide which subspace (real or velocity space) to split next. This makes the method metric free, as there is no need to specify in advance how velocity and configuration spaces relate to each other, and ensures the subspace with greater variation is subdivided more.

Once the binary tree has been built, the leaf nodes can be used to estimate the phase-space volume of each particle. The volume of the hypercube of the leaf nodes can vary significantly even between close neighbours due to Poisson sampling noise of the original particle distribution. Simply using the volumes of raw nodes would result in a large dispersion, but this can be reduced by using a smoothing scheme. We have modified ENBID's anisotropic kernel density estimation scheme to return a volume as well as a density. The scheme works by first calculating a local metric about each particle based on the shape of the hypercube of each leaf node and using the length of the six sides to scale the phase space in each direction. A covariance matrix is then calculated based on the nearest 64 neighbours in the scaled space. Diagonalizing the covariance matrix gives a set of eigenvalues and eigenvectors that can be used to perform a further transformation to a coordinate system in which the eigenvectors define the principal axes. In this new rotated and scaled space, the local particle distribution will have unit covariance in each direction. Finally, the phase-space volume of the particle is defined as a hypersphere with 1/40 of the volume of a hypersphere based on the distance to the 40th nearest neighbour, R_{40}. It therefore has a radius R_1 of

$$R_1 = \left(\frac{1}{40}\right)^{1/6} R_{40}. \quad (1)$$

The modified version of ENBID returns a 6×6 matrix that describes the phase-space volume of each particle.

The actual sampling is performed by choosing points from an 6D isotropic Gaussian distribution in which each component has zero mean and variance $\sigma^2 = \gamma R_1^2$. Each generated point is then transformed by applying the inverse operations to convert from the sampled space back to positions and velocities in the original simulation. Some care has to be taken with the choice of the scale parameter γ which relates the scale of the Gaussian kernel to the scale of the hypersphere returned by ENBID. Choosing too large a value of γ will result in the phase-space structure of each component of the halo being made hotter and too diffuse. Choosing too small a value of γ will result in an artificially clumpy phase-space distribution. As a compromise we have chosen $\gamma = 1/48$. In a separate study (Wang et al., in preparation), we have developed a maximum likelihood method for modelling the distribution function of the stellar halo in order to estimate the mass of the host dark matter halo. Using this approach, we have found that the recovered halo mass is biased high if γ is too large. With $\gamma = 1/48$ the level of the bias is controlled to be less than about 10 per cent, which is much smaller than typical measurement errors. Smaller values of γ would remove this bias entirely, but would negate the benefit of distributing stars over a phase-space kernel, because in the limit $\gamma \to 0$ one is merely placing all the stars at the same phase-space location as the tagged particle from which they arise.

In order to test our ENBID sampling method and demonstrate that we have made a suitable choice for the parameter γ, we attempt to reconstruct a debris stream from one of our simulations after artificially degrading its resolution by sampling only a fraction of its particles. We extracted from the simulation a satellite in the process of undergoing significant tidal stripping, such that its particles have an extended distribution in both configuration and velocity

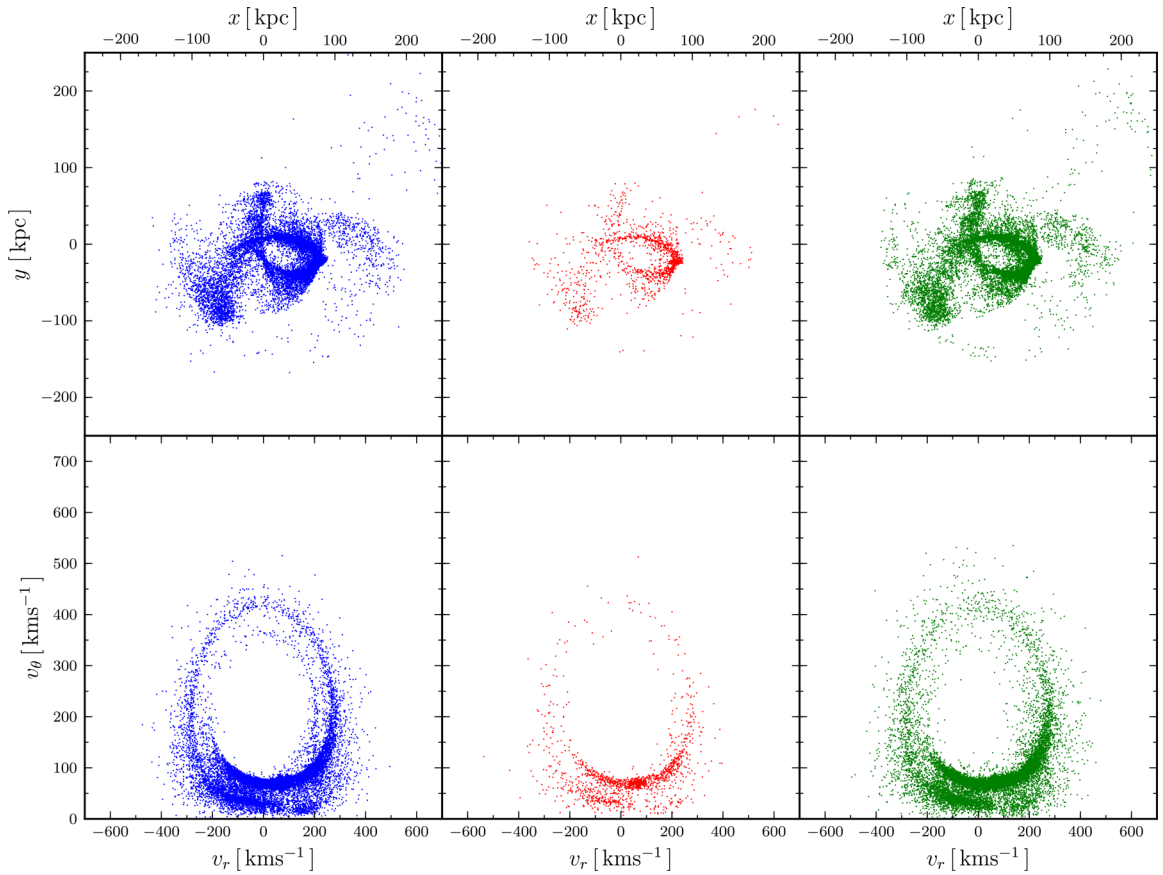

Figure 1. Real-space projection of an example satellite (upper row) and velocity-space projection (bottom row). The left-hand column shows the original set of satellite particles, while middle column is the down-sampled version with just 10 per cent of the number of particles. The right-hand column shows the result of resampling 10 per cent of the particles back to the original resolution using phase-space kernels based on ENBID density estimates.

space. This object comprises 15 250 particles. A lower resolution version was created by randomly selecting 10 per cent of the particles. Finally, the low-resolution version was then resampled using our ENBID method, with each particle being converted back into 10 particles, thus creating an object with the same resolution as the original satellite. A comparison between the two then demonstrates the success of our phase-space sampling method.

Fig. 1 shows projections of phase space for the different versions of the satellite. The top row is a real-space projection and the bottom row a velocity projection. The original satellite (left-hand column) has a large amount of structure in both real and velocity space, with several complete wraps of the tidal tails visible. In the down-sampled version (the second column from the left), the primary features are still present but the subtle ones are missing, while the remnant of the satellite core is more prominent. The ENBID up-sampled version (the third column from the left) is remarkably similar to the original satellite, though the subtle features tend to be smoothed out. In particular, the arc in velocity space corresponding to an orbital apocentre at $v_\theta \approx 400\,\mathrm{km\ s^{-1}}$, while still present, is less well resolved. This is expected, as information is lost during down-sampling and no resampling scheme will be able to recover it perfectly.

A further test of how well the original and ENBID versions match is a comparison between the volume distribution function of phase-space density, $v(f)$, as shown in the upper panel of Fig. 2. We have used the technique of Arad, Dekel & Klypin (2004), based on Delaunay tessellation, as an independent measurement of the phase-space density. The phase-space volume distribution function is found to be similar in all three cases, suggesting that it is robustly defined by even a small subset of the particles and that the ENBID resampling process does not significantly alter it. The differences can be seen more clearly in the bottom panel of Fig. 2. The relative difference between the down-sampled volume distribution function and the original is shown by the red curve, while the relative difference between the ENBID up-sampled version and the original is show by the green curve. We have also tested how much noise is introduced to $v(f)$ simply by randomly sampling the phase-space distribution defined by the original particles (as estimated by ENBID) at a ratio of 1 : 1 (grey line). The relative difference in all these cases is mostly within 10–50 per cent, as indicated by the two horizontal black dashed lines. There is no systematic trend in the relative difference across nine orders of magnitude in phase-space density. A few spikes correspond to larger differences in particular regions, which appear to be almost entirely due to the shot noise introduced by the discreteness of the particle distribution.

We also notice that at the low-density end of Fig. 2, the red dashed curve based on the down-sampled particles starts to plunge at $f \sim 10^{-14.5}$, whereas the blue curve based on the original sample continues to increase down to $f \sim 10^{-15.5}$. The down-sampled particles are more sparsely distributed in phase space and thus sample low density regions poorly, as well as increasing the Poisson noise. The green dot–dashed curve, which is based on the particles re-sampled with ENBID, recovers the same small-scale behaviour of the original particle distribution. The grey curve, where each particle is

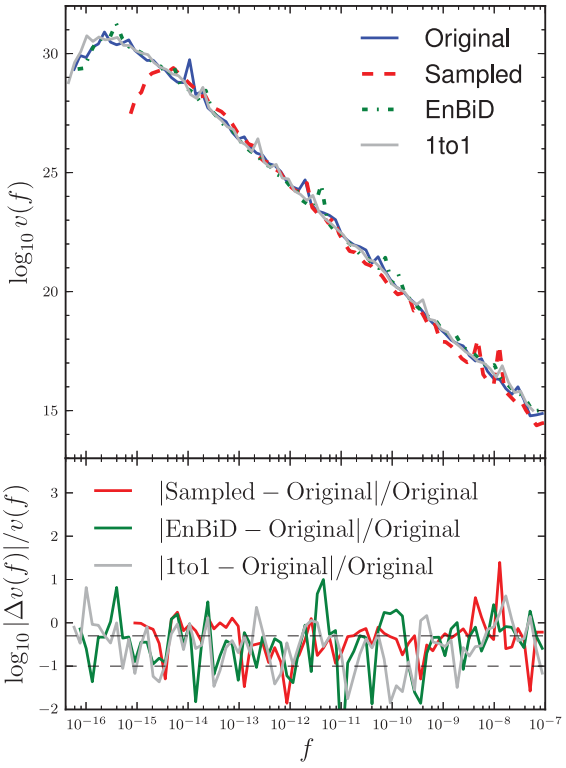

Figure 2. Top: comparison of the volume distribution function of phase-space density, $v(f)$, for different phase-space sampling schemes applied to a single satellite. Both axes have arbitrary units. Blue (labelled 'original'): the original high-resolution particle distribution. Red ('sampled'): a sample of 1/10 of the original distribution. Green ('ENBID'): the same number of particles as the original distribution, drawn from the phase-space kernels defined by the down-sampled (red) distribution. Grey('1-to-1'): each particle resampled once from the phase-space kernels defined by the original distribution. Bottom: relative difference between the volume distribution functions of the 'sampled', 'ENBID' and '1-to-1' results to the 'original' volume distribution function. Black dashed lines mark the 10 and 50 per cent level of difference.

resampled once from the phase-space kernel of the original sample, have almost the same small scale behaviour of the original sample as well.

When applying this resampling technique to the tagged particles an important consideration is how the particles are divided into subsets on which to run ENBID. The particular subset that is chosen defines which particles are considered as phase-space neighbours. We choose to apply the ENBID volume estimation to all the tagged particles that form along a given branch of the halo merger tree. This separates the tagged particles into sets based upon the halo they formed in.

It is also possible to split the tagged particles by individual assignments, i.e. only particles from an individual single-age stellar population are considered as potential phase-space neighbours. This has the advantage that the phase-space volumes of particular populations are not diluted by stars forming in the same halo at later times. The problem with this method is that sets of particles from individual assignments can be very small in many cases, such that their phase-space volumes are too poorly sampled to accurately constrain the structure in six dimensions. While our implementation allows either method to be used, we have chosen to use the first method in the rest of this paper.

ENBID requires a minimum of 64 particles for the covariance matrix calculation. We therefore exclude assignments with smaller numbers of particles than this. We have checked that this eliminates a negligible amount of stellar mass.

3 MOCK CATALOGUES

In this section, we describe the range of information provided by our mock catalogues, and present a set of example mocks generated from the Aquarius simulations. The combination of the various techniques used results in a rich data set.

3.1 Public catalogues

Five complete mock catalogues for the Aquarius haloes A–E are publicly available at the following website http://galaxy-catalogue.dur.ac.uk:8080/StellarHalo. Each mock is stored in a separate table within the StellarHalo data base and can be accessed via the web interface using SQL queries. The data base will contain all stars with absolute magnitude $M_g < 7$. This cut is intended to select bright stars most useful to studies of the distant halo, including main-sequence turn-off (MSTO) stars, red giant branch stars and BHB stars (see Section 3.4). It excludes a very large number of faint main sequence stars which are not accessible to surveys of the outer halo of the Milky Way.

Table 1 lists the information that is provided. Positions, velocities and apparent magnitudes[1] are given relative to both the centre of the halo and to an observer placed in a position similar to that of the Sun. In order to do this, we first define a plane that a galactic disc would likely lie in and then choose a position for the Sun within that plane. We choose to orient the disc plane perpendicular to the minor axis of the moment of inertia of the dark matter in the inner 10 kpc, as the minor axis is a stable orbital axis. While the inclusion of baryons might be expected to change the shape of the centre of the halo due to adiabatic contraction (Kazantzidis et al. 2004), their presence is unlikely to alter the halo's orientation. We have taken the Sun to be 8 kpc from the centre of the galaxy and moving with a velocity $(U, V, W)_\odot = (11, 232, 7)$ km s^{-1} (Schönrich, Binney & Dehnen 2010). The observer has been placed on the x-axis, though any orientation on the solar circle would be an equally valid choice.

In addition to the magnitudes, we provide a few intrinsic properties of stars, such as their age, metallicity, stellar mass, surface gravity and effective temperature. The semi-analytical galaxy formation model GALFORM uses physically motivated prescriptions to follow the evolution of the galaxies within each halo. While GALFORM outputs many properties for each galaxy, it is only the star formation and chemical history that are used to build the mocks. These two model predictions for the baryonic components of satellites are essential quantities required to populate the halo with stars. The final input are the theoretical isochrones that provide the detailed properties of the stars. For a given age and metallicity, together with an IMF, the isochrones provide a breakdown of these stellar properties. The type column is the original 'stage' flag in the PARSEC isochrones that denotes different phase of stellar evolution (see Table 1).

Beyond the spatial and kinematical information for individual stellar populations, the underlying dark matter N-body simulations provide the history of halo mass assembly. Each star is associated with a dark matter particle and the subhalo to which that particle is currently bound, through the dark matter ID and Subhalo ID. This information allows the dynamical history of the particle to be

[1] The magnitudes are based on SDSS u, g, r, i, z filters.

Table 1. Information provided in our online mock catalogues that may be accessed at http://galaxy-catalogue.dur.ac.uk:8080/StellarHalo. SQL can be used to query any of these fields.

Field	Data type	Unit	Description
ID	long		ID of the star, unique within the mock catalogue
Position (x, y, z)	float	kpc	Star position, relative to halo centre
Velocity (v_x, v_y, v_z)	float	km s^{-1}	Stars velocity, relative to halo rest frame
Galactic coordinates (l, b)	float	$^\circ$	Galactic longitude and latitude relative to solar observer
Radial distance (r)	float	kpc	Distance from solar observer
Relative velocity (v_r, v_θ)	float	km s^{-1}	Radial and tangential velocity relative to solar observer
$\log_{10}$(Age)	float	$\log_{10}$(yr)	Age of star
Metallicity	float		Fraction of star mass in metals
Mass	float	$M_\odot$	Current mass of star
$\log_{10}(g)$	float	log (cgs)	Surface gravity of the star
$\log_{10}(T_{\rm eff})$	float	$\log_{10}$(K)	Effective temperature of the star
Type	int		The original PARSEC evolutionary stage flag that labels the type of star (0=Pre-MS, 1=Main sequence, 2=Subgiant branch, 3=Red giant branch, 4,5,6=stages of core He burning, 7,8=stages of asymptotic giant branch)
$M_{(u, g, r, i, z)}$	float		Absolute rest-frame (u, g, r, i, z) band (SDSS) magnitude of the star
$m_{(u, g, r, i, z)}$	float		Apparent rest-frame (u, g, r, i, z) band (SDSS) magnitude of the star as seen by the solar observer
Dark matter ID	long		The tagged dark matter particle this star was spawned from
Subhalo ID	long		The current subhalo to which the associated dark matter particle belongs to. -1 if not currently a member of any subhalo, 0 if part of the main halo
Tree ID	long		The subhalo tree that this star came from. All stars that arrive in the main halo via the same satellite have the same ID
Infall redshift	float		Accretion redshift of the subhalo this star once belonged to. Defined as the redshift at which the satellite reaches its maximum mass (see the text).

traced back through the simulation. The satellite which brought the star into the main halo can be identified, along with the time of accretion of that satellite and the time at which the particle was stripped from the satellite. It is possible to find stellar structures and streams easily by grouping stars according to their satellite of origin, using the Tree ID field. The 'infall' redshift of the parent satellite has been provided, defined as be the redshift when the satellite reached its maximum mass (for most objects this correlates closely with the actual time of entry into the main FOF group and hence with crossing the virial radius of the central halo).

3.2 Errors and extinction

For generality and simplicity, the observables presented in our catalogue have not been convolved with sources of observational error. In order to simulate a particular survey, the user should therefore convolve physical quantities such as magnitude, velocity and metallicity with appropriate error functions. The user may also wish to impose a foreground dust extinction model, and superpose additional source distributions to account for contaminations of halo star selections by faint foreground stars and the extragalactic background.

3.3 Limitations

While the mock catalogues we have generated are based on sophisticated cosmological models, they are not without limitations. The biggest limitation is that our model only generates the accreted component of the stellar halo. It does not include processes that create *in situ* halo components in hydrodynamical simulations, such as the scattering of disc stars or star formation in gaseous streams. This means that the model is most suitable for studying the outer regions of a galaxy ($r > 20$ kpc). There are various possible ways to add in additional components to create a more complete model of the entire galaxy, and we will explore these in future papers.

Stars belonging to the galactic disc are not included in the mocks and the effect of having a disc present within the halo during its evolutionary history is also missing. Subhaloes that pass near or through a disc are likely to be disrupted more rapidly (D'Onghia et al. 2010). The importance of such effects is unclear because the disc has been growing as the halo grows; even today, it presents a small cross-section to infalling satellites. This effect is therefore most relevant to the small number of satellites that pass close to the centre of the halo at late times. The consequences of neglecting enhanced disruption by the galactic disc are likely to vary among our five simulations, according to their particular mass assembly histories and subhalo orbital distributions. While the properties of individual streams may change significantly, we expect that, in most cases, the overall properties of the halo should be robust to the inclusion of a realistic disc.

As discussed in Section 2.3, the C10 technique uses dark matter particles to trace the distribution and dynamics of stars. In reality, stars may form on phase-space trajectories that are not well sampled in collisionless simulations (for example, centrifugally supported discs). For this reason, the spatial configuration of stars within surviving satellites may not be reliable in detail. The importance of the initial configuration is greatly diminished once a satellite has been disrupted and the stars phase-mixed into the halo, but studies of chemical and kinematic gradients along coherent streams should keep this limitation in mind.

3.4 Example mocks

We now describe a set of example mock catalogues generated from the Aquarius simulations and use them to illustrate how particular samples of stars can be easily selected. Each sample is defined by a

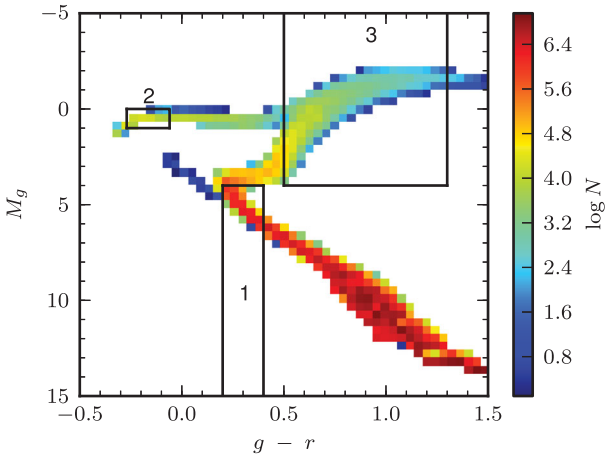

Figure 3. Colour–magnitude diagram for the overall stellar population of the accreted stellar halo of Aq-A. Stars still bound to satellites are not included. The colours encode the number of stars within a bin and the boxes illustrate the approximate regions covered in our three mock catalogue selection functions (see the text for details of the actual selection criteria, which make use of additional photometric bands and metallicity information). Box 1 marks our selection for MSTO stars, box 2 marks the Bell et al. (2010) selection for BHB stars and box 3 the Xue et al. (2014) selection for K giants. Photometric magnitudes are given in the SDSS *ugriz* filter system.

set of colour and magnitude cuts that selects a particular stellar type commonly used for surveying the stellar halo. The overall colour–magnitude diagram for all the stars in the stellar halo of Aq-A (with the contribution from stars still bound in satellites removed) is shown in Fig. 3. The selections for our three samples of halo tracers are marked by boxes.

The vast majority of stars are low mass stars that have not yet evolved off the main sequence and account for a significant fraction of the total stellar mass. However, as mentioned above, these are very faint. They can only be observed in the solar neighbourhood and are of little interest to studies of the distant halo or its global properties. A clear red giant branch and horizontal branch are visible; these contain the older, brighter stars which can be observed to large galactocentric distances and which are thus are most useful when studying the structure of the halo.

3.4.1 MSTO

Main-sequence turn-off (MSTO) stars can be unambiguously selected using the colour cut $0.2 < g - r < 0.4$ (e.g. Bell et al. 2010), and an absolute magnitude cut $M_g > 4$, in order to only select dwarf stars and remove giants. MSTO are relatively numerous and therefore serve as a good proxy for the general stellar content of the stellar halo. Approximately 40–50 per cent of the stellar mass brighter than $M_g = 7$ is found in MSTO stars and there are 81 506 853 such stars in the Aq-A halo. The volume currently accessible using MSTO stars is a relatively small fraction of the total halo, but their high surface density in this volume allows a great deal of information to be gathered by wide-field photometric surveys such as SDSS (e.g. Jurić et al. 2008).

3.4.2 BHB

Blue horizontal branch (BHB) stars are very often used as tracers of the distant stellar halo. They are luminous and have a very nearly constant magnitude, $M_g \sim 0.7$, so they can be located accurately to very large distances along the line of sight even with photometry alone. Unfortunately, the conditions under which BHB stars form are still not well understood. This makes their abundance in our mock catalogues highly sensitive to the assumptions made by the theoretical isochrones we use.

To select BHB stars in our mock catalogues we use the colour criteria proposed by Bell et al. (2010): $0.98 < u - g < 1.28$, $-0.27 < g - r < -0.06$ and excluding the region $([u - g - 0.98]/0.215)^2 + ([g - r + 0.06]/0.17)^2 < 1$. With this selection, BHB stars make up a small fraction of the total mass in stars brighter than $M_g = 7$, ranging from 1 to 3 per cent in the five mocks. There are 239 950 BHBs in the Aq-A halo.

The main source of stellar contamination in photometric BHB selection is fainter blue straggler (BS) stars. The formation of BS is also not well understood, and they are not included in the PARSEC isochrones. Hence, even halo BS (and white dwarfs) are not a source of contamination in the mocks. This should be kept in mind when comparing to observed photometric BHB counts, as should the absence of quasar contamination at faint magnitudes.

3.4.3 K giants

K giants are another stellar type frequently used to map the Milky Way's stellar halo (Morrison, Flynn & Freeman 1990; Morrison et al. 2000; Starkenburg et al. 2009). They have the advantage that they are found in predictable numbers in old populations of all metallicities. However, unlike other standard candles, their luminosities range by two orders of magnitude, which makes obtaining accurate distances more challenging. K giants are one of the types that have been targeted for spectroscopy by the Sloan Extension for Galactic Understanding and Exploration (SEGUE, Yanny et al. 2009) and recently, Xue et al. (2014) have estimated distances to a sample of 6036 K giants using a combination of observed $g - r$ colours and spectroscopic metallicity.

In order to isolate K giants in our stellar halo models we have used the colour cuts $0.5 < g - r < 1.3$ and $0.5 < u - g < 3.5$ from Xue et al. (2014), along with their proposed empirical polynomial relation between $(g - r)$ and [Fe/H] to remove red horizontal branch and red clump (RC) giants. All stars with $(g - r) > 0.087\,[\mathrm{Fe/H}]^2 + 0.39\,[\mathrm{Fe/H}] + 0.96$ is excluded from the selection. An additional cut of $M_g < 4$ removes faint dwarf stars of the same colours. Our selection results in the subset of stars contained within box 3 of Fig. 3, with a large fraction of the stars at the fainter end removed as being either red horizontal branch or RC giants. About 6–7 per cent of the total stellar halo mass in stars brighter than $M_g = 7$ is found as K giants and there are 956 246 K giants in the Aq-A halo.

4 DENSITY PROFILE OF THE STELLAR HALO

4.1 Spherically averaged radial profiles and breaks in the stellar halo

One of the most basic, but still unanswered, questions in galactic structure is the density profile and shape of the Milky Way stellar halo. While there is now general consensus that the stellar halo is oblate and follows a broken power-law density profile, the exact form is still ill constrained. Evidence for a break has been found by studies using RR Lyrae variables as standard candle tracers of the distant metal-poor halo, with measured break radii of ∼25 kpc (Watkins et al. 2009), ∼30 kpc, (Sesar et al. 2010) and ∼45 kpc

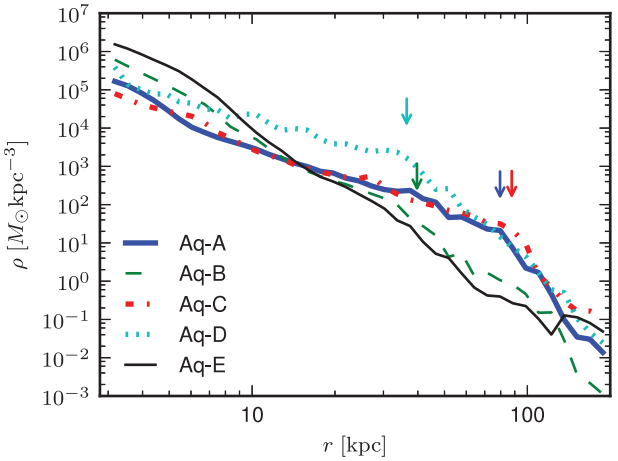

Figure 4. The spherically averaged density profiles of the five Aquarius stellar haloes. The arrows show the location of the break in a best-fitting broken power law. Three out of the five haloes have clear breaks, a fourth has a weak break while Aq-E is well fitted by an unbroken power law.

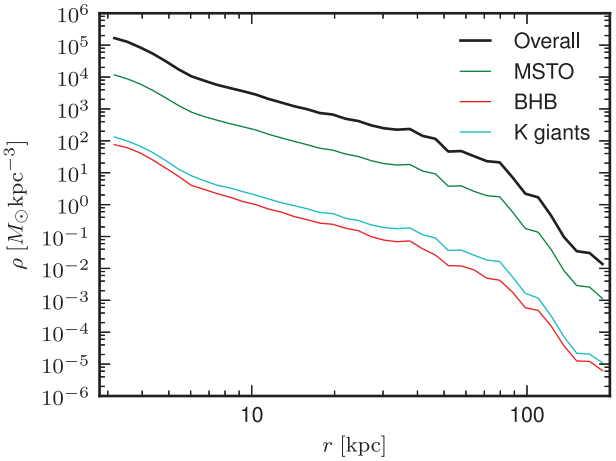

Figure 5. The spherically averaged density profile of the Aq-A stellar halo obtained using different types of tracer stars as indicated in the legend. The black line shows the density profile of all stars in the halo.

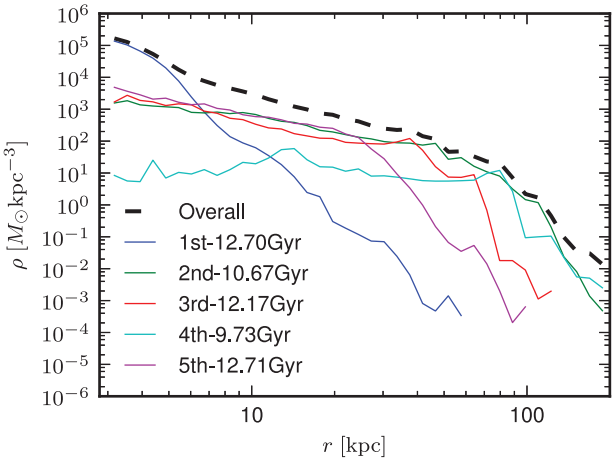

Figure 6. The density profile of the Aq-A stellar halo broken down into the five satellites which each contribute the largest amount of mass to the stellar halo. Numbers next to the legend show the look-back infall time of the five satellites.

(Keller et al. 2008; Akhter et al. 2012). Measurements using other types of tracers have also suggested the existence of a break. Deason et al. (2011), using a large sample of BHB and BS stars from SDSS Data Release 8 (DR8), again found the profile to follow a broken power law, with a break radius at ∼27 kpc. Similarly, from a sample of near MSTO stars from the Canada–France–Hawaii Telescope Legacy Survey (CFHTLS), Sesar et al. (2011) measured the break to be at ∼28 kpc. In contrast, no break or slope steepening was seen by De Propris, Harrison & Mares (2010), who analysed a sample of 666 BHB stars in two different directions separated by about 150° on the sky. Instead they found a smooth stellar distribution obeying a power-law density profile of index ∼−2.5. Akhter et al. (2012) suggested that an explanation for these differences in the measured location of the break could be that, at larger distances, the outer halo is not be smooth but dominated by clumps and debris. Other studies of the power-law index (or indices) of the Milky Way halo include Chiba & Beers (2000), Miceli et al. (2008) and Carollo et al. (2010).

In Fig. 4, we show the density profiles of the accreted stellar haloes in our five simulations, and we have excluded stars in surviving bound subhaloes. Three of them show clear breaks, although generally at larger distances than estimated for the Milky Way. To check whether the type of tracer makes a difference, we also show in Fig. 5 the stellar density profile of the Aquarius A halo measured from MSTO, BHB and K giant stars. The overall density profile is that of all stellar mass in the tags from the C10 method, while the tracer profiles are from samples taken from the mock catalogued described in Section 3.4. All of the curves have the same shape, with the break occurring at the same radius, but with differing normalizations, depending on the abundance of the stellar type.

It has been suggested that the presence of a break in the Milky Way's density profile occurs at the transition between an inner region where the halo is dominated by *in situ* stars and an outer region which is primarily composed of accreted halo stars. However, this cannot be the case for the breaks seen in the profiles of the mock catalogues, since they do not contain any *in situ* component. Recently, Deason et al. (2013) have proposed that the breaks in stellar halo profiles are related to the satellite accretion history of their galaxies and that a massive accretion event occurring 6–9 Gyr ago is likely to have caused the break in Milky Way density profile. Satellites that are just falling into the halo and are still in the process of being tidally disrupted, may have already lost a significant fraction of their stars through tidal stripping. Although these stripped stars are unbound and now form part of the halo, their tidal debris is still in the form of large coherent structures roughly tracing the orbits of their progenitor satellites. The merit of constructing a galactocentric, spherically averaged density profile for these debris streams and clouds is questionable, as they are far from being spherical or dynamically relaxed. Such profiles generally appear flat in the central regions and break to a very rapid falloff at their outer edge.

In Fig. 6, we show the density profiles of the stars that were brought in by the five largest contributors to the Aq-A stellar halo. Numbers beside the legends correspond to the look-back time of infall for these satellites. Except for the very largest contributor, the profiles of the others all have very strong breaks. A visual inspection reveals that the largest contributor was accreted earliest at about 12.7 Gyr ago. By the end of the simulation, it has been completely disrupted and mixed, with a correspondingly smooth profile. In contrast, the fourth largest contributor has fallen in more recently 9.73 Gyr ago and is still in the form of a coherent, massive stream. It has a sharp break around ∼80 kpc. The other three contributors

are in intermediate stages of disruption; all show clear caustic/shell structures (visible as localized density peaks). Notice the fifth most massive contributor shows these structures despite having been accreted at roughly the same time as the largest contributor, which is well mixed. The largest contributor has about 4.5 times the mass of the fifth most massive, and hence the decay of its orbit through dynamical friction may have been more rapid.

In Aq-A, breaks in the individual contributor profiles correspond to these caustic structures (this is true also for the other four haloes, which we do not plot here). This finding is consistent with the results of Deason et al. (2013), that breaks are related to the apocentres of accreted satellites. These few most massive contributors determine the overall profile – over certain ranges of radius, almost the entire stellar mass comes from a single progenitor. For example, around 50 kpc the overall density profile is almost entirely determined by the second largest contributor. It is clear that the break in the overall profile in this region is driven by stars from this object. It can be clearly seen that the number of breaks in the overall profile, and their sharpness, will depend strongly on the relative distribution of stars from the few satellites that contribute significant amounts of debris to the halo.

It can be argued that when measuring the spherically averaged stellar halo density profile, coherent structures such as those just described should be removed. However, while a massive stream such as the Sagittarius stream in the Milky Way might be easy to exclude, objects in later stages of disruption are more difficult to identify. In the simulations, we are able to use the SUBFIND halo finder to identify bound structures and follow the accretion histories of stars in order to find candidate streams but this cannot be done with observational data. In addition, even if structures have been identified, the question is then how to decide whether they should be excluded from the halo. For these reasons, as the overall profile is so sensitive to individual accretion events, and thus is likely to change rapidly with time, it is not clear how useful it is to measure the density profile of a stellar halo, other than for the purpose of gaining information about recent large accretion events.

4.2 Pencil-beam surveys

As shown in Fig. 5, the spherically averaged density profile is not sensitive to the type of tracer used. What is much more important is the area of the sky over which the profile is measured. In order to measure the profile to large radii, narrow pencil-beam surveys of bright tracers are often used, probing deep into the stellar halo over a limited area of sky. In order to quantify the intrinsic differences in the slope of the profile in different directions we employ a similar observational strategy to Sesar et al. (2011) and measure the density the simulated halo stars over in patches of small solid angle. Sesar et al. (2011) used four 'pencil beams' from the CFHTLS to select MSTO stars in order to measure the slope of the stellar halo density profile between 5 and 35 kpc in heliocentric radius. A direct comparison with Sesar et al. (2011) is limited by the lack of *in situ* halo stars in our model; it may be that these dominate the inner part of the halo. Moreover, Sesar et al. (2011) attempted to excise the most obvious accreted substructures from their analysis (the Sagittarius and Monoceros streams). This is a subjective procedure and implicitly assumes an underlying 'smooth' halo. We do not try to excise visible streams, or even bound satellites.[2] Our aim in this section is to demonstrate how substructure in the halo affects the determination of the density profile, rather than to compare directly with the results of Sesar et al. (2011).

We select stars in the same colour–apparent magnitude range as Sesar et al. (2011) ($0.2 < g - r < 0.3, g > 17, 17 < r < 22.5, i > 17$), and measure their galactocentric radial density profile in patches of 8×8 deg^2. Note that the Sesar et al. colour selection criteria is slightly different from our own (simpler) 'MSTO' criterion used in all other sections of this paper (shown in box 1 of Fig. 3), and also that the apparent magnitudes used for this selection are those for an observer at our fiducial solar position (see Section 3.1). We focus on the slope of the profile for stars in the range 5–35 kpc from the Sun. Although the selection of stars is based on heliocentric distance, $r_{\rm h}$, the density profiles are expressed in terms of galactocentric radius, $r_{\rm g}$, hence pencil beams in different directions sample different ranges of $r_{\rm g}$. We tile the entire sky with equal-area pencil beams and compute the parameters of the best-fitting power-law density profile, $\rho \propto r^n$, in each beam. The top panel of Fig. 7 shows a sky map of Aq-A with pixels colour-coded by the value of n along a given beam. There are large variations in n across the sky, with profiles ranging from very steep ($n \sim -5$) in some areas to flat ($n \sim 0$) in others.

The large-scale gradient in the value of n with galactic latitude is likely to result from the non-spherical overall distribution of halo stars. This is seen clearly in the projected number density of MSTO stars (including those bound to satellites) shown in the lower panel of Fig. 7. The surface density of accreted stars is high towards the galactic centre and falls off towards the galactic poles. This panel also shows that variations of n on smaller scales (hundreds of square degrees) correspond to localized density substructures in the halo. Some of the most dramatic fluctuations occur around surviving galaxies that are being tidally stripped; for example the density 'hot spots' in the lower panel, labelled 'A' and 'B', correspond to regions in the upper panel with $n \sim 0$. However, in general there are very few massive satellites in the range $5 < r_{\rm h} < 35$ kpc; we have overplotted black circles on Fig. 7 to highlight all those in Aq-A, with relative diameters scaled to 10 times the angular size of their subhaloes (as seen by our fiducial solar observer). Most small-scale fluctuations in the measured density slope, are due to inhomogeneities in the stellar halo itself rather than bound substructures.

In order to quantify how much the slope varies over the sky, the distributions of n are shown in Fig. 8 for all five haloes. These are 'area weighted' distributions, in that each bin shows the fraction of the total number of 8 deg $\times$ 8 deg tiles with a given slope. In each panel, a vertical dashed line marks the overall slope of the spherically averaged density profile in the same radial range (effectively the mass weighted average slope). For Aq-A, B and C, the overall slope is close to either the median or *mode* of the distribution, but for Aq-D and E the overall slope is significantly different, lying at the steep end of the distribution. This can be explained by the fact that, as we saw earlier, a substructure or stellar stream can completely dominate the overall profile, making it radically different from that measured in a typical patch of the halo. It should be noted that, so far, we have been considering the slope of the profile to no more than 35 kpc from the Sun. In the Milky Way, it has been suggested that an *in situ* component makes a significant contribution to the halo out to galactocentric distances of $r_{\rm g} \sim 20$ kpc. This component may also be much smoother than the accreted halo originating from disrupted satellites. A smooth *in situ* component may reduce the variation in the density profile.

To investigate the structure of the outer halo, we take all stars brighter than $M_g = 7$ in the mock catalogue with galactocentric radii in the range $20 < r_{\rm g} < 80$ kpc. This upper limit is chosen

[2] Note that we remove bound satellites from the analysis in all other sections.

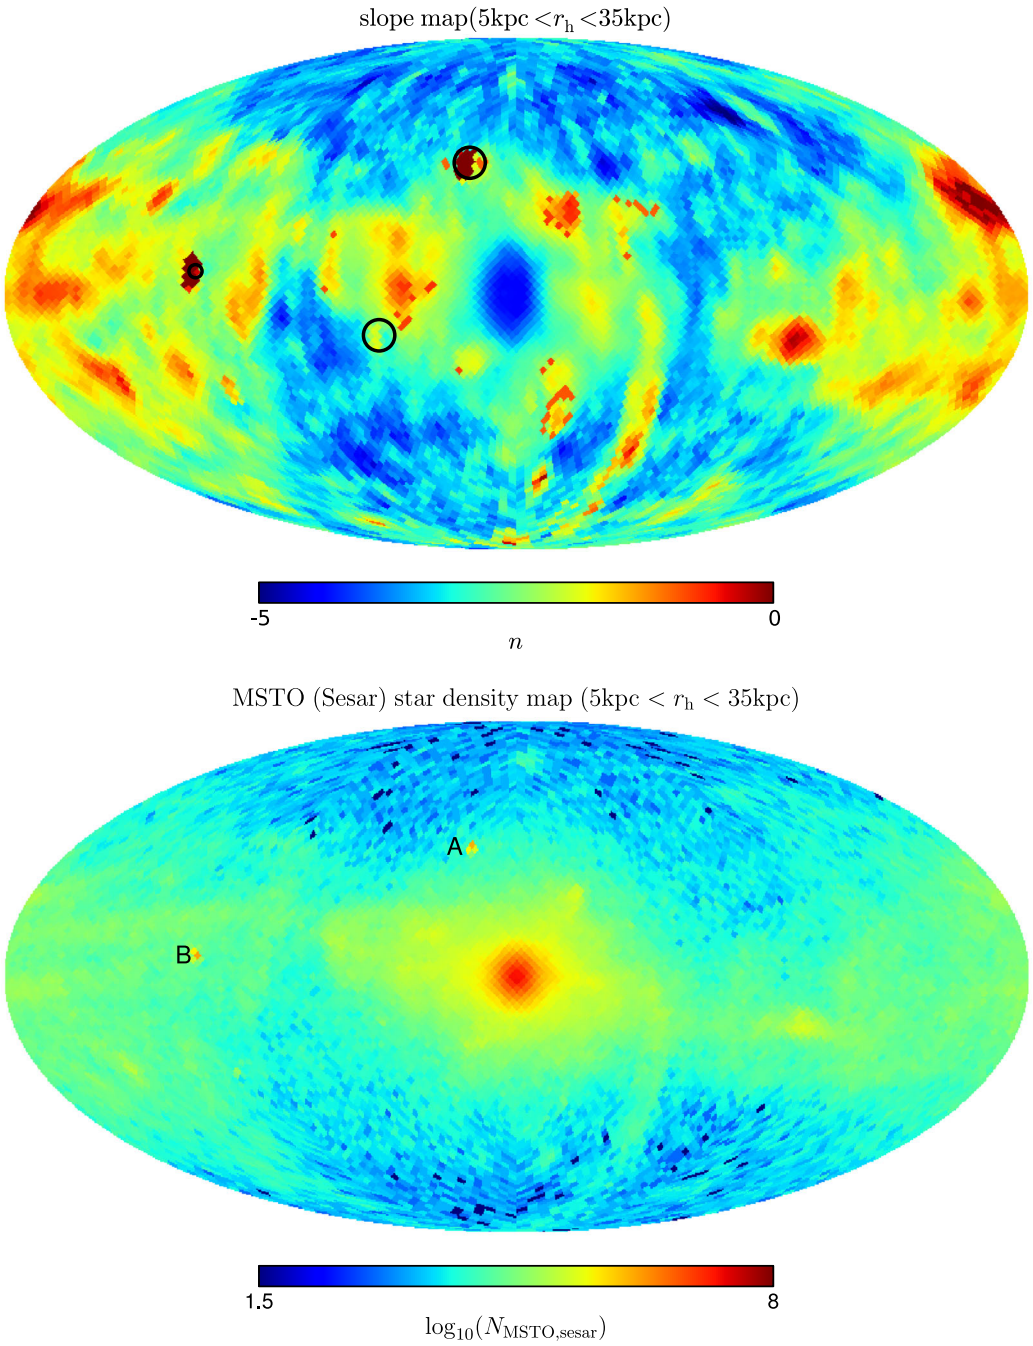

Figure 7. Top: a heliocentric all-sky map in which the colour of each 8 deg^2 pixel corresponds to a value of n, the best-fitting (dimensionless) power-law slope of the galactocentric density profile of stars from our mock catalogues, selected according to the MSTO colour–magnitude criteria of Sesar et al. (2011, see the text) in the heliocentric distance range $5 < r_h < 35$ kpc. The Galactic Centre is at the centre of this map and our fiducial disc plane is oriented along the equator. Unlike previous figures, bound satellites have not been excised in this analysis, although there are only three in this distance range (black circles, diameter equal to 10 times the apparent angular size of the corresponding subhalo). Bottom: the logarithm of the projected surface number density of stars selected and observed in the same way as the top panel. Labels A and B mark the location of satellites discussed in the text.

because the overall density profiles of Fig. 4 do not show breaks out to 80 kpc. We again construct pencil beams to a fixed distance from the Sun and measure the slopes of the density profiles from the galactic centre. The all sky map of slopes is presented in the top panel of Fig. 9 and the surface density of stars in the lower panel.

Compared with the top panel of Fig. 7, the fluctuations across the sky in Fig. 9 are more numerous and pronounced. The large-scale variations in n no longer trace large-scale variations in the surface density of stars (bottom plot). However, we can see many of the small-scale variations in these two maps correspond. Since stars within 20 kpc to the galactic centre are excluded, the central overdensity in Fig. 7 disappears.

The more complex fluctuations in Fig. 9 arise because there are more substructures at these larger galactocentric radii, and the local variations in the stellar halo are more pronounced. Again, we mark all subhaloes that are more massive than 10^8 M$_\odot$ with black

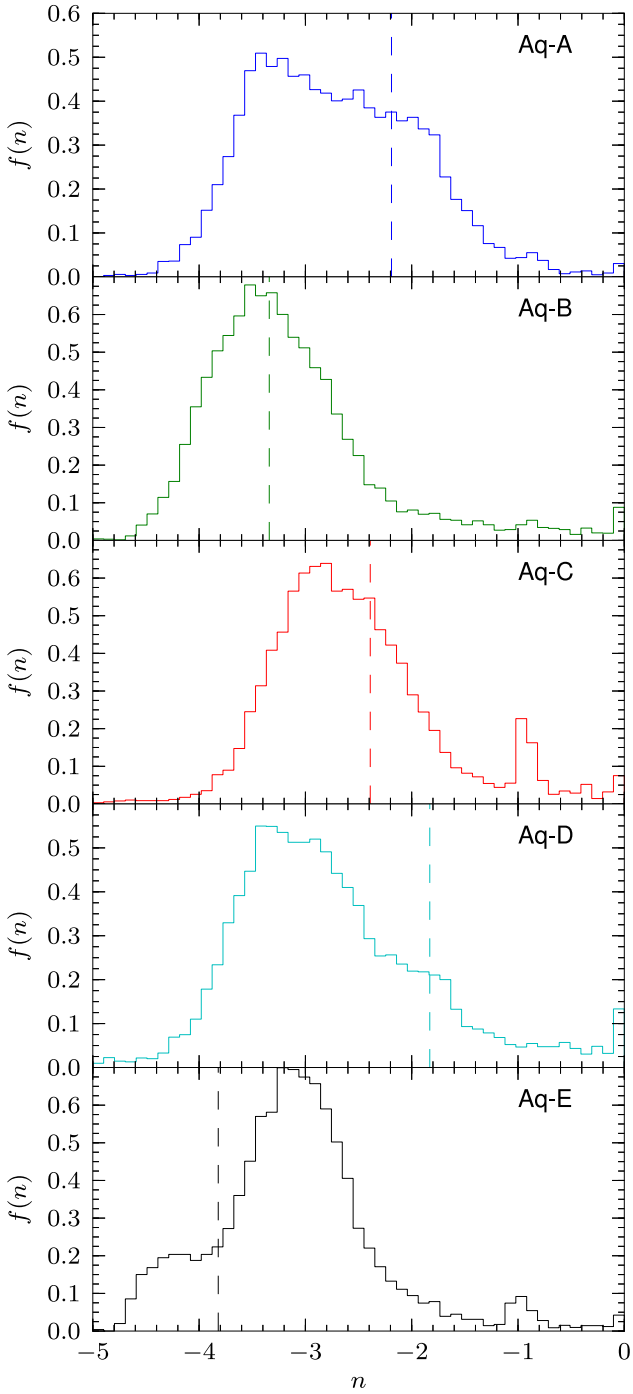

Figure 8. The distribution of volume density profile power-law slope measured for halo stars in our mock catalogue selected using the criteria of Sesar et al. (2011, see the text) in patches of area 8 deg^2 tiling the whole sky. Each panel shows a different mock catalogue. Bound subhaloes and other overdensities have not been removed; they contribute a small fraction of the pixels in these area-weighted distributions. The dashed vertical lines mark the overall slope from a power-law fit to the (mass-weighted) spherically averaged density profile in the same volume.

circles. Most of these subhaloes are distant and are thus small in their apparent size. The three most prominent local patches are associated with three massive and extended substructures. Interestingly, the two patches that are close to the North Galactic Pole and on the left-hand side of the map, which can also be identified in Fig. 7, now give very steep slopes in contrast to their flat slopes in Fig. 7. This is because the associated subhaloes are close to the solar observer. The radial range of 5–35 kpc from the Sun has a very flat density distribution, beyond which the density profile drops very quickly.

5 TRACER DISTRIBUTION

In Section 4.1, we showed that the shape of the spherically averaged density profiles of different stellar tracers in our mock catalogues are similar when averaged over the whole sky. In this section, we examine the sky distribution of different stellar tracers in more detail. We first look at the ellipsoidal shape of different populations, and then consider the degree of variation amongst different tracers in small patches of the sky.

5.1 The overall shape of tracers and dark matter

In this section, we aim to quantify the overall ellipsoidal shape of the stellar halo. The degree of 'flattening' (or triaxiality) of the stellar halo is of great interest, although to be meaningful, measurements based on ellipsoidal fits to the 3D distribution of stars require that a significant fraction of the halo be 'smooth' and symmetric.

Fig. 10 gives the ratio between the ellipsoidal minor and major axes, c/a, obtained by diagonalizing the standard inertia tensor of different stellar populations (coloured lines, see the legend) and dark matter particles (black lines) in each simulation. We consider stars and dark matter in four cumulative radial bins, the outer edges of which correspond to 1/4, 1/3, 1/2 and 1 $\times$ R_{200}. We exclude stars within 20 kpc of the galactic centre and we do not include stars in any bound satellites.

The overall shape of dark matter particles and stars are somewhat different from each other. In the outer halo, with the exception of Aq-E, the distribution of stars is significantly flatter than that of the dark matter. This reflects the tendency for stars even in the outer stellar halo to be concentrated towards the main symmetry plane of the halo, as noted by Cooper et al. (2011). As particles at larger radii are included in the measurement, c/a remains approximately constant for stars and dark matter. The large change in the outermost point for Aq-E is likely to be because the stellar distribution at large radius is dominated by only one coherent tidal stream. Interestingly, we find that BHB stars have a slightly more isotropic distribution than the other tracers.

Fig. 11 illustrates how the major axes of stellar tracers correlate with the major axes of dark matter particles.[3] For all the five haloes, we see the misalignment angle is typically small (less than $\sim$20 deg) particularly within $R_{200}/4$. Aq-E is distinguished by a sudden jump in this angle as the outer limit of the distribution is increased, which is related to the dominant stream mentioned above. Aq-C shows a notable misalignment between the BHB stars and other tracers, particularly in our innermost bin. A systematic offset between BHBs and the other tracers (such that BHBs are slightly more misaligned with the dark matter) is also evident in Aq-A, Aq-B and Aq-C. In Aq-D, on the other hand, all tracers of the stellar halo are very well aligned with the dark matter out to R_{200}, despite the significant differences in the flattening of these components seen in Fig. 10.

[3] We have checked the direction of the major axis of dark matter particles is almost constant with radius.

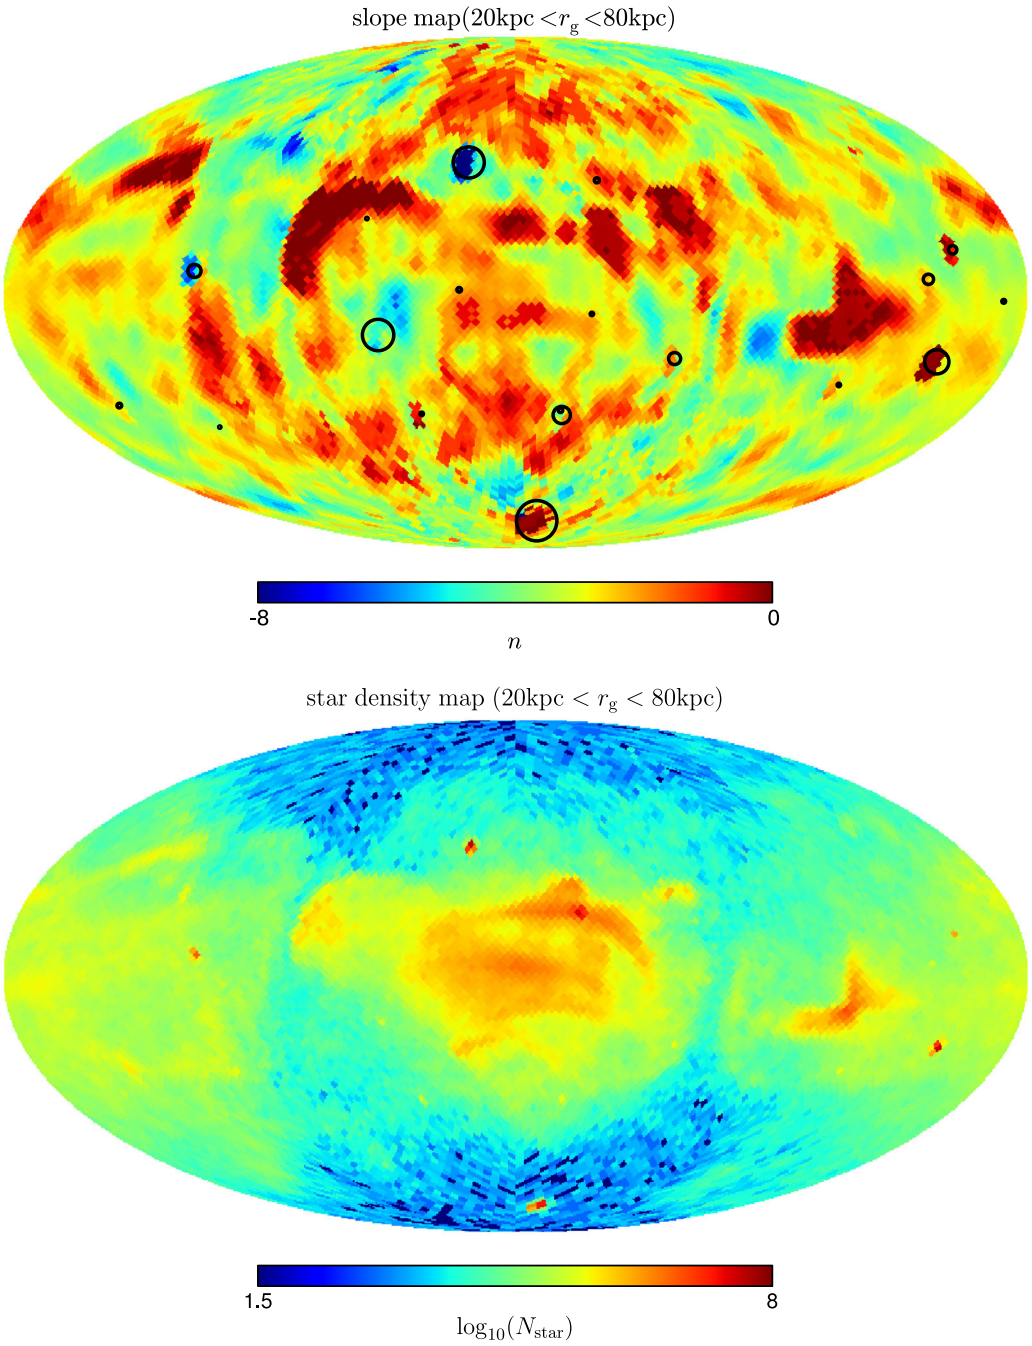

Figure 9. Heliocentric sky map similar to Fig. 7, but showing the power-law density slope (top) and surface number density (bottom) of all stars in our mock catalogue with galactocentric distance $20 < r_g < 80$ kpc. Bound satellites and other overdensities have not been excised. Black solid circles mark subhaloes that host satellites, with the angular size of circle equal to 10 times the half-mass radius of the subhalo as viewed from our fiducial solar position.

5.2 Tracer distribution over the sky

It has been proposed that localized variations in the mix of stellar populations may provide a method for discovering new stellar structures in the Milky Way halo (e.g. Ibata et al. 2007; McConnachie et al. 2009). Significant fluctuations in the ratio of BHB to MSTO stars have been found in the stellar halo (Bell et al. 2010). Some of these, such as the 'low-latitude stream' are almost devoid of BHB stars, while other structures are rich in them. The Sagittarius tidal stream even shows a variation in the BHB/MSTO ratio along its extent (Bell et al. 2010), possibly related to a variation in mean metallicity that has been interpreted as the result of a population gradient in the progenitor dwarf galaxy (e.g. Chou et al. 2007). In this section, we perform a similar test with the BHB, MSTO and K giant samples from our mock catalogues and attempt to understand the origin of the differences.

The simplest test is to search for variations by comparing sky maps of the projected map for our different tracers as seen by a hypothetical observer located on the solar circle. Fig. 12 (left-hand column) shows a number of such maps for the Aq-C halo. (Aq-D and Aq-B show similar features; to avoid repetition we focus on Aq-C only.) It contains many clearly defined streams that are visible in all the tracers. What is most interesting, however, is that there

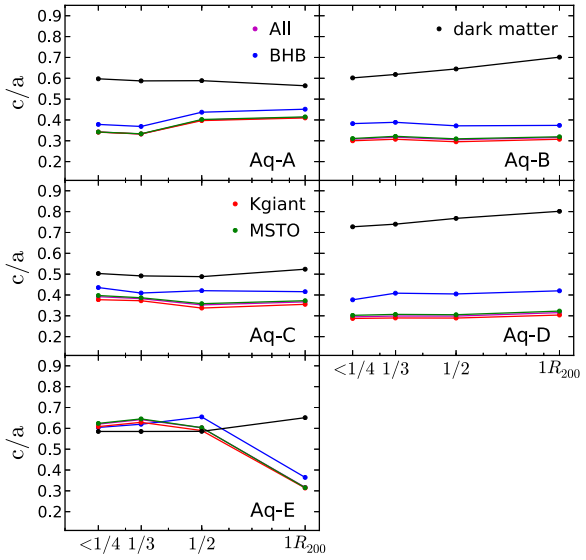

Figure 10. The overall axis ratio (minor axis versus major axis) of different stellar tracers and the underlying dark matter, reported for four shells centred on the galaxy with inner radius 20 kpc in all cases and outer radii of 1/4,1/3, 1/2 and 1 × R_{200}, respectively.

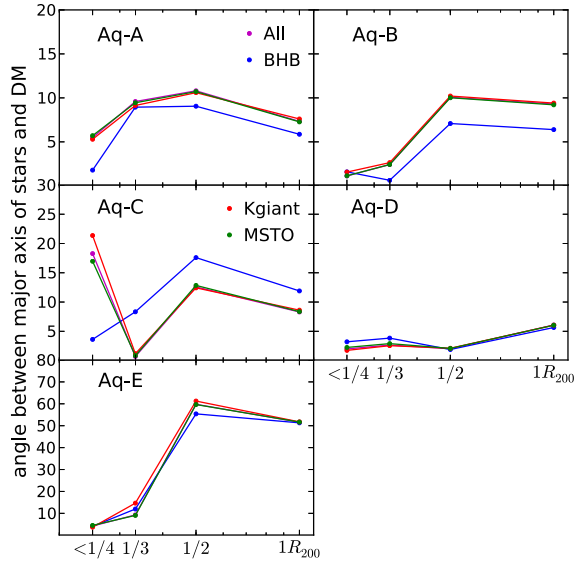

Figure 11. The angle (in units of degrees) between the major axis of stellar tracers and the major axis of dark matter particles reported for four shells centred on the galaxy with inner radius 20 kpc in all cases and outer radii of 1/4,1/3, 1/2 and 1 × R_{200}, respectively.

are two streams that do not appear in the BHB map but appear in the two other maps. In particular, the large continuous horizontal stream lying close to the equatorial plane that are clearly seen in the MSTO and K giant map is partially missing in the BHB map, while the shorter stream located near the top of the MSTO map, towards the galactic north pole, is totally absent in the BHB map. The absence of these streams in the BHB map helps to explain why BHB stars show distinct features in the behaviour of axis ratio and major axis at large radii.

Comparing the ratio of BHBs/MSTO stars (Fig. 12 right-hand column), the two missing streams are easily identifiable. While the large, equatorial stream is at least partially present in the BHB map, it has a lower than average abundance resulting in a higher MSTO/BHB ratio. In contrast, the stream near the top of the MSTO map is almost completely absent in the BHB map and has a huge MSTO/BHB ratio. Other streams with more standard stellar populations have BHB to MSTO ratios similar to that of the overall halo and thus do not show up in the ratio maps. There are also differences in the MSTO/K giant ratio, but in the opposite sense. From the MSTO/K giant map it can be seen that the two streams have a higher than average abundance of K giants. One other feature is that the BHB stars are slightly more abundant in the galactic centre.

These differences in the relative abundance of the three types of tracer can be explained by differences in the age and metallicity of the stellar populations that make up the features. In Fig. 13, we show a map of the average age (left) and the average metallicity (right) of the Aq-C stellar halo. These have been generated from the mean age/metallicity of all stars within each equal-area patch of the sky. It is immediately obvious that the streams missing from the BHB map are composed of the very youngest stars in the halo. As well as being a few gigayears younger than the mean, they are also slightly more metal rich. In these populations BHB stars have not yet had time to form, which explains their deficit. The actual age of these structures may be even younger, but the average is increased by the other non-associated halo stars along the same line of sight. Interestingly, there is a patch of intermediate-age stars that stands out in the lower-right corner of the age map, but is not visible in any of the tracer maps. This suggests that while differences in the make-up of stellar populations will help identify some structures, not all will be found with this technique.

6 CONCLUSIONS

In this paper, we have described a new way to construct mock catalogues of stellar haloes based on cosmological *N*-body simulations. We have adopted the particle tagging method of C10 and extended it to turn the output into a catalogue of individual stars that can be directly compared to observations.

Using the Aquarius *N*-body simulations of galactic haloes as the basic framework, we employed the semi-analytical galaxy formation model, GALFORM, to calculate the evolution of the baryonic component of the Universe. This predicts the amount and properties of the stars forming in each halo throughout the simulation. These stellar populations were used to tag dark matter particles in the *N*-body simulation. Finally, these tags are turned into a mock catalogue using theoretical isochrones to convert the fundamental properties of the stellar populations, such as age and metallicities, into a set of observables such as colour and magnitudes. A phase-space sampling technique was applied to 'explode' the massive particles into numerous individual stars while maintaining the phase-space structure of the original simulation.

The output of our method represents, within the limitations of the model, a perfect data set containing precise information about the entire galaxy. Actual observations usually provide an incomplete, limited view of a galaxy. They may be subject to selection criteria, be limited in coverage, or, in the case of surveys of the stellar halo, be obstructed by other structures such as the galactic disc. To make a fair comparison between simulations and observations, it is necessary to convert the simulation data into a form that also reproduces the limitations and biases of the observational data. The selection criteria of an actual survey, encoded in a mask or selection function, can be readily applied to our mocks. Observational errors, foreground dust extinction and contaminating sources can also be included easily. Thus, in principle, the mock catalogues can be viewed and analysed in a manner that mimics the observational

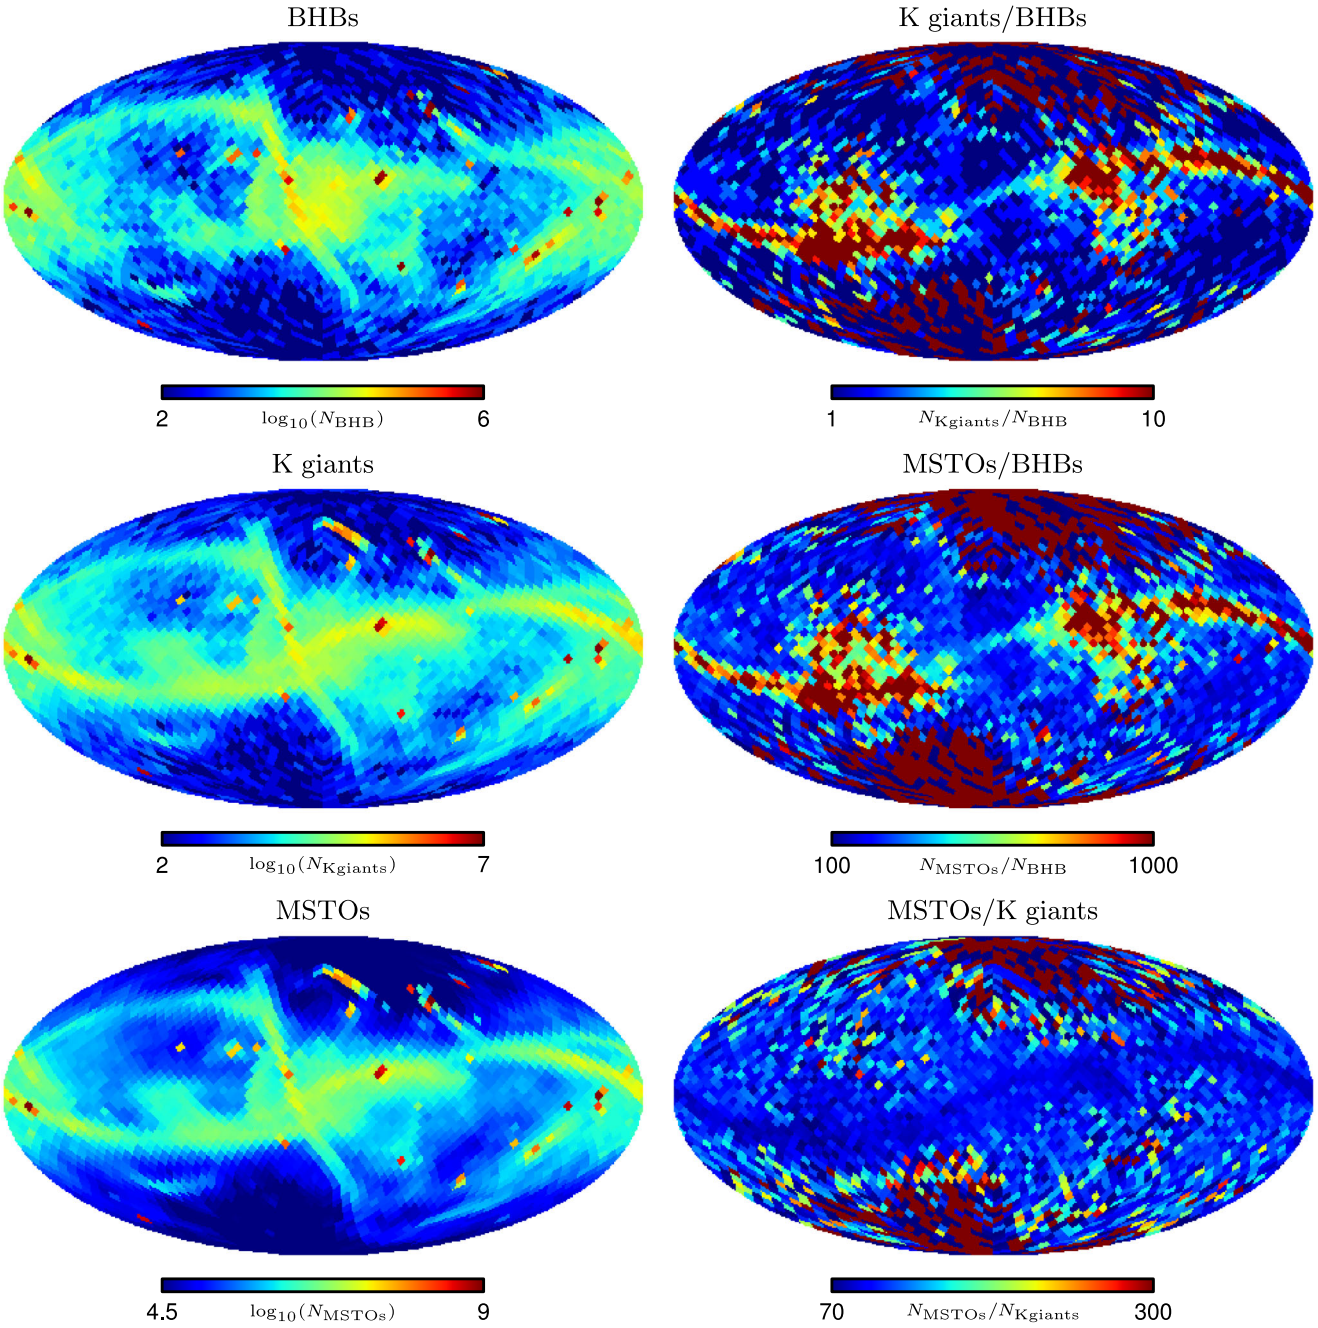

Figure 12. The distribution of different stellar tracers over the sky for the Aq-C stellar halo as seen by our fiducial heliocentric observer. The Galactic Centre is at the centre of the map and the disc plane is oriented along the equator. Stars within 20 kpc of the galactic centre have been excluded. Left: the projected logarithmic number of BHBs, K giants and MSTO stars. Right: the ratio of K giants/BHBs, MSTO/BHBs and MSTO/K giant stars.

data. In this paper, we have used the Aquarius simulations to make mock catalogues of galactic stellar haloes that can be used to help interpret data from upcoming surveys such as those to be carried out by the *Gaia* mission. However, there is no reason why this technique could not be applied to other simulations of both smaller haloes, such as dwarfs, or larger haloes, such as clusters.

Using our mock catalogues we carried out simple analyses to explore whether current observations of the Milky Way halo provide an accurate picture. We have found that:

(i) The accreted stellar halo is mainly built up from a few massive objects. Those that are in early stages of disruption still maintain a coherent structure and are locally concentrated. They dominate the spherically averaged density profile at their corresponding radii. Since the overall density profile is so sensitive to these few substructures it provides limited information about the structure of the underlying dark matter halo. Instead, it tells us more about the recent accretion history (Deason et al. 2013). Three of the five haloes we have investigated show clear breaks in their spherically averaged density profiles while two are well described by a power law. In our model, these breaks have nothing to do with the distinction between accreted and *in situ* components, as has often been claimed for the Milky Way halo. Rather, they are produced by the accretion of satellites, although not necessarily by a single one.

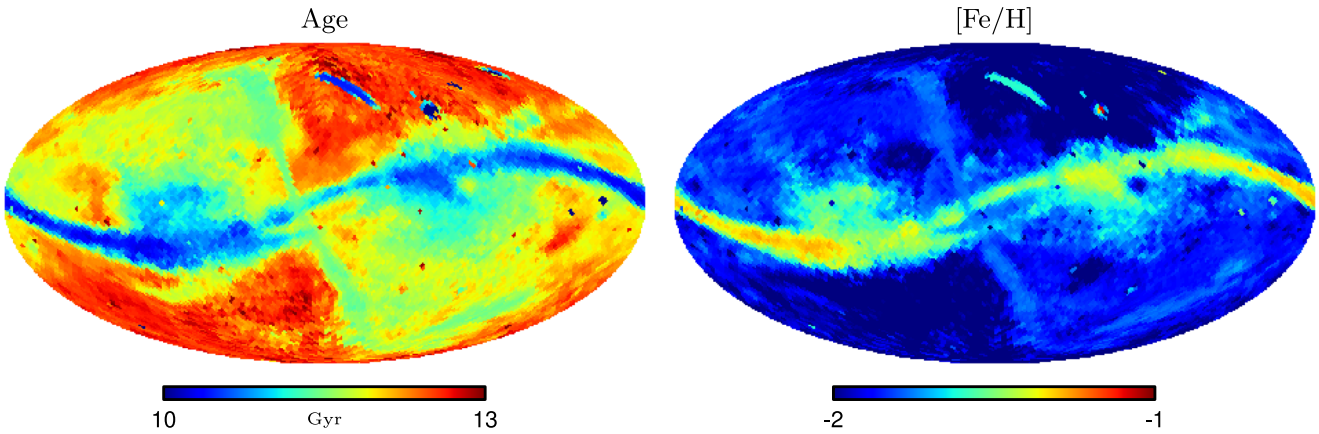

Figure 13. All-sky maps of the mean properties of the Aq-C stellar halo when averaged along the line of sight. Left map: average age. Right map: average metallicity map. Stars within 20 kpc of the galactic centre have been excluded.

(ii) The structure of the accreted stellar halo is clumpy and thus varies considerably over the sky. Attempts to measure properties of the halo, such as its density profile, based on limited regions of the sky, provide biased estimates. For example, if the profile is modelled as a power law, the slopes measured in different pencil-beam surveys through our mock stellar haloes vary from very steep to flat depending on the direction of the survey.

(iii) Beyond 20 kpc, the flattening of dark matter and halo stars (as measured by c/a, the major to minor axis ratio of an ellipsoidal fit) is approximately constant out to R_{200}. The stellar distribution is significantly flatter than the dark matter. The minor axis of the stellar distribution is approximately aligned with that of the inner dark matter halo, which we have also assumed to be the normal vector to the main symmetry plane.

(iv) Differences in the ages and metallicities of the satellites that build up the stellar halo are visible as local variations in the composition of stellar halo populations. These are reflected in differences in the abundance of different stellar types. For example young, more metal rich objects tend to lack BHB stars. Searching for such fluctuations in halo star surveys offers a way to identify streams and structure in the halo.

In this paper, we have developed a method of converting massive stellar particles into stars based on the output of the C10 *N*-body particle tagging method. However, with only minor modification it should be possible to use the same conversion technique on hydrodynamical galaxy formation simulations. By combining these two approaches, components that are missing from our treatment, such as the galactic disc or bulge, could be included, while other components, such as the stellar halo, could be followed with much higher resolution than is attainable with current hydrodynamic simulations.

The five complete mock catalogues for the Aquarius haloes A-E, are publicly available at http://galaxy-catalogue.dur.ac.uk:8080/StellarHalo, in a data base that can be queried via SQL. The properties that are available in each catalogue are described in the Table 1.

ACKNOWLEDGEMENTS

We thank the anonymous referee for their careful reading of our paper and helpful suggestions. WW is grateful for useful discussions with Jiaxin Han and Yanchuan Cai. This work was supported by the European Research Council [GA 267291] COSMIWAY and Science and Technology Facilities Council [ST/F001166/1,ST/L00075X/1]. WW is supported by the Durham International Junior Research Fellowship [RF040353]. APC acknowledges a Chinese Academy of Sciences International Research Fellowship and NSFC grant no. 11350110323. The simulations for the Aquarius Project were carried out at the Leibniz Computing Centre, Garching, Germany, at the Computing Centre of the Max-Planck-Society in Garching, at the Institute for Computational Cosmology in Durham, and on the STELLA supercomputer of the LOFAR experiment at the University of Groningen. Sky maps were produced with the HEALPY implementation of the HEALPIX algorithms (http://healpix.jpl.nasa.gov; Górski et al. 2005).

REFERENCES

Abadi M. G., Navarro J. F., Steinmetz M., 2006, MNRAS, 365, 747
Akhter S., Da Costa G. S., Keller S. C., Schmidt B. P., 2012, ApJ, 756, 23
Arad I., Dekel A., Klypin A., 2004, MNRAS, 353, 15
Bailin J., Bell E. F., Valluri M., Stinson Greg S., Debattista Victor P., Couchman H. M. P., Wadsley J. 2014, ApJ, 783, 95
Baugh C. M., 2006, Rep. Prog. Phys., 69, 3101
Beers T. C. et al., 2012, ApJ, 746, 34
Bell E. F., Xue X. X., Rix H.-W., Ruhland C., Hogg D. W., 2010, AJ, 140, 1850
Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, MNRAS, 370, 645
Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, MNRAS, 427, 127
Bullock J. S., Johnston K. V., 2005, ApJ, 635, 931
Bullock J. S., Kravtsov A. V., Weinberg D. H., 2001, ApJ, 548, 33
Carollo D. et al., 2010, ApJ, 712, 692
Chiba M., Beers T. C., 2000, AJ, 119, 2843
Chou M.-Y. et al., 2007, ApJ, 670, 346
Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, MNRAS, 319, 168
Colless M. et al., 2001, MNRAS, 328, 1039
Cooper A. P. et al., 2010, MNRAS, 406, 744 (C10)
Cooper A. P., Cole S., Frenk C. S., Helmi A., 2011, MNRAS, 417, 2206
Cooper A. P., D'Souza R., Kauffmann G., Wang J., Boylan-Kolchin M., Guo Q., Frenk C. S., White S. D. M., 2013, MNRAS, 434, 3348
Cooper A. P., Gao L., Guo Q., Frenk C. S., Jenkins A., Springel V., White S. D. M., 2014, preprint (arXiv:1407.5627)
Crain R. A. et al., 2009, MNRAS, 399, 1773
D'Onghia E., Springel V., Hernquist L., Keres D., 2010, ApJ, 709, 1138
de Bruijne J. H. J., 2012, Ap&SS, 341, 31
De Lucia G., Helmi A., 2008, MNRAS, 391, 14
De Propris R., Harrison C. D., Mares P. J., 2010, ApJ, 719, 1582
Deason A. J., Belokurov V., Evans N. W., 2011, MNRAS, 416, 2903

Deason A. J., Belokurov V., Evans N. W., Johnston K. V., 2013, ApJ, 763, 113
Font A. S., Johnston K. V., Bullock J. S., Robertson B. E., 2006, ApJ, 638, 585
Font A. S., McCarthy I. G., Crain R. A., Theuns T., Schaye J., Wiersma R. P. C., Dalla Vecchia C., 2011a, MNRAS, 416, 2802
Font A. S. et al., 2011b, MNRAS, 417, 1260
Gao L., Navarro J. F., Cole S., Frenk C. S., White S. D. M., Springel V., Jenkins A., Neto A. F., 2008, MNRAS, 387, 536
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Ibata R., Martin N. F., Irwin M., Chapman S., Ferguson A. M. N., Lewis G. F., McConnachie A. W., 2007, ApJ, 671, 1591
Ibata R. A. et al., 2014, ApJ, 780, 128
Johnston K. V., Bullock J. S., Sharma S., Font A., Robertson B. E., Leitner S. N., 2008, ApJ, 689, 936
Jordi C. et al., 2010, A&A, 523, A48
Jurić M. et al., 2008, ApJ, 673, 864
Kafle P. R., Sharma S., Lewis G. F., Bland-Hawthorn J. F., 2013, MNRAS, 430, 2973
Kaiser N. et al., 2010, in Stepp L. M., Filmozzi R., Hall H. I., eds, Proc. SPIE Conf. Ser. Vol. 7733, Ground-based and Airborne Telescopes III. SPIE, Bellingham, p. 14
Kazantzidis S., Kravtsov A. V., Zentner A. R., Allgood B., Nagai D., Moore B., 2004, ApJ, 611, L73
Keller S. C., Murphy S., Prior S., Da Costa G., Schmidt B., 2008, ApJ, 678, 851
Kennicutt R. C., Jr, 1983, ApJ, 272, 54
Libeskind N. I., Knebe A., Hoffman Y., Gottlöber S., Yepes G., 2011, MNRAS, 418, 336
McConnachie A. W. et al., 2009, Nature, 461, 66
Martínez-Delgado D. et al., 2010, AJ, 140, 962
Miceli A. et al., 2008, ApJ, 678, 865
Morrison H. L., Flynn C., Freeman K. C., 1990, AJ, 100, 1191
Morrison H. L., Mateo M., Olszewski E. W., Harding P., Dohm-Palmer R. C., Freeman K. C., Norris J. E., Morita M., 2000, AJ, 119, 2254
Navarro J. F., Eke V. R., Frenk C. S., 1996, MNRAS, 283, L72
Navarro J. F. et al., 2010, MNRAS, 402, 21
Pasetto S., Chiosi C., Kawata D., 2012, A&A, 545, A14
Peñarrubia J., Benson A. J., Walker M. G., Gilmore G., McConnachie A. W., Mayer L., 2010, MNRAS, 406, 1290
Pillepich A. et al., 2014, MNRAS, 444, 237
Pontzen A., Governato F., 2013, MNRAS, 430, 121
Robertson B., Bullock J. S., Font A. S., Johnston K. V., Hernquist L., 2005, ApJ, 632, 872
Sawala T., Frenk C. S., Crain R. A., Jenkins A., Schaye J., Theuns T., Zavala J., 2013, MNRAS, 431, 1366
Schönrich R., Binney J., Dehnen W., 2010, MNRAS, 403, 1829
Schönrich R., Asplund M., Casagrande L., 2014, ApJ, 786, 7
Sesar B., Jurić M., Ivezić Ž., 2011, ApJ, 731, 4
Sesar B. et al., 2010, ApJ, 708, 717
Sharma S., Steinmetz M., 2006, MNRAS, 373, 1293
Sharma S., Bland-Hawthorn J., Johnston K. V., Binney J., 2011, ApJ, 730, 3
Spergel D. N. et al., 2003, ApJS, 148, 175
Springel V. et al., 2008a, MNRAS, 391, 1685
Springel V. et al., 2008b, Nature, 456, 73
Starkenburg E. et al., 2009, ApJ, 698, 567
Tissera P. B., White S. D. M., Scannapieco C., 2012, MNRAS, 420, 255
Tissera P. B., Scannapieco C., Beers T. C., Carollo D., 2013, MNRAS, 432, 3391
Tumlinson J., 2010, ApJ, 708, 139
Walker M. G., Mateo M., Olszewski E. W., Gnedin O. Y., Wang X., Sen B., Woodroofe M., 2007, ApJ, 667, L53
Watkins L. L. et al., 2009, MNRAS, 398, 1757
Xue X.-X. et al., 2014, ApJ, 784, 170
Yanny B. et al., 2009, AJ, 137, 4377
Zolotov A., Willman B., Brooks A. M., Governato F., Hogg D. W., Shen S., Wadsley J., 2010, ApJ, 721, 738

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