

Extending the halo mass resolution of N -body simulations

Raul E. Angulo,^{1,2,3}★ Carlton M. Baugh,³ Carlos S. Frenk³ and Cedric G. Lacey³

¹*Centro de Estudios de Física del Cosmos de Aragón, Plaza San Juan 1, Planta-2, 44001 Teruel, Spain*

²*Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA*

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

Accepted 2014 May 30. Received 2014 May 29; in original form 2013 October 15

ABSTRACT

We present a scheme to extend the halo mass resolution of dark matter N -body simulations. The method uses the simulated density field to predict the number of sub-resolution haloes expected in different regions, taking as input the abundance and the bias factors of haloes of a given mass. These quantities can be computed analytically or measured from higher resolution simulations. We show that the method recovers the abundance and clustering in real- and redshift-space of haloes with mass below $\sim 7.5 \times 10^{13} h^{-1} M_{\odot}$ at $z = 0$ to better than 10 per cent. By applying the method to an ensemble of 50 low-resolution, large-volume simulations, we compute the expected correlation function and covariance matrix of luminous red galaxies (LRGs), which we compare to state-of-the-art baryonic acoustic oscillation measurements. The original simulations resolve just two-thirds of the LRG population, so we extend their resolution by a factor of 30 in halo mass in order to recover all LRGs. Using our method, it is now feasible to build the large numbers of high-resolution large volume mock galaxy catalogues required to compute the covariance matrices necessary to analyse upcoming galaxy surveys designed to probe dark energy.

Key words: cosmology: theory – large-scale structure of Universe.

1 INTRODUCTION

The spatial distribution of galaxies is an important resource in physical cosmology, encoding information about the physics of galaxy formation and the values of the basic cosmological parameters (Guzzo et al. 2008; Cabré & Gaztañaga 2009; Sánchez et al. 2009, 2012; Beutler et al. 2011; Zehavi et al. 2011; Reid et al. 2012). A number of galaxy surveys are underway or planned which share the primary science goal of using the large-scale structure of the Universe to constrain the nature of dark energy (e.g. Laureijs et al. 2011). To achieve this, these surveys will map galaxies over many tens of cubic gigaparsecs. As the clustering signals predicted by competing cosmological models are often very similar, the scientific exploitation of the surveys will be limited by how well we are able to understand the systematic errors which may affect statistical measures of the large-scale structure of the Universe.

A complete understanding of the systematic and sampling errors associated with clustering measurements requires many effects to be modelled, including cosmic variance, non-linear evolution of density fluctuations, scale-dependent bias, redshift space distortions, discreteness effects and survey geometry. To meet the challenge of providing the best possible theoretical predictions, the most accurate techniques have to be employed. Currently, this means using N -body

simulations of the hierarchical clustering of the dark matter (DM; see Springel, Frenk & White 2006). The need to model clustering accurately on scales beyond $100 h^{-1}$ Mpc requires computational boxes in excess of $1 h^{-1}$ Gpc on a side (Angulo et al. 2008a). Resolving Milky Way mass haloes or smaller in such calculations is expensive but has been achieved in a small number of cases (for a summary of the state of the art, see Kuhlen, Vogelsberger & Angulo 2012). Such calculations are currently one-offs and the computational resources are not available to generate the large numbers of such runs which are required to compute covariance matrices for large-scale structure statistics.

The principal way to study errors on clustering measurements from galaxy surveys is through an accurate model of the experiment itself (Baugh 2008). For the case of relevance here (the spatial distribution of galaxies), this is optimally achieved in a three step process. First, the halo clustering is predicted by following the evolution of particles in an N -body simulation (see the recent reviews of Springel et al. 2006; Kuhlen et al. 2012). Secondly, the properties of galaxies within these haloes are predicted using a semi-analytical model of galaxy formation (for reviews see Baugh 2006; Benson 2010). And finally, the appropriate flux limit, sample selection, redshift completeness and the geometry of the survey need to be applied to the catalogues (e.g. Merson et al. 2013). Some of these steps may be modified. For example, in the case of a low-resolution simulation, ‘galaxies’ may be added using an empirical rule based on the smoothed density of the DM (White et al. 1987; Cole et al.

★ E-mail: reangulo@gmail.com

1998). The predictions of the semi-analytical model may be substituted by empirical techniques tuned to match observational data, such as halo occupation distribution (HOD) modelling or sub-halo abundance matching (Zehavi et al. 2011; Simha et al. 2012).

The direct approach to modelling the errors on clustering statistics, namely populating many large volume, high-resolution N -body simulations with galaxies using semi-analytical galaxy formation models, is computationally expensive for two reasons. (i) The large number of independent N -body simulations needed to make robust estimates of the errors. An adequate estimate of the variance requires several dozen realizations of the density field (e.g. a 10 per cent error on the variance for a Gaussian distribution requires ~ 50 realizations). An order of magnitude more simulations is needed to robustly compute the full covariance matrix (Takahashi et al. 2009). (ii) The huge dynamic range required to resolve the haloes which are likely to host the galaxies observed. For instance, in N -body simulations with box sizes of a few gigaparsecs, only Milky Way sized haloes can be identified robustly, even in the highest mass resolution simulation of this type carried out to date (Angulo et al. 2012). The mass limit is much larger in typical large-volume simulations (Fosalba et al. 2008; Kim et al. 2009; Teyssier et al. 2009; Alimi et al. 2012). In fact, currently there is no simulation which can simultaneously model, for instance, the faintest galaxies and the volume to be probed by the Dark Energy Survey (The Dark Energy Survey Collaboration 2005). Although algorithms and computer hardware are constantly improving, finite computational resources impose a limit on N -body simulations: carrying out a single simulation, not to mention an ensemble of them, meeting the desired requirements is currently prohibitively expensive computationally.

Several authors have proposed algorithms to predict galaxy clustering efficiently and to overcome the difficulties stated above. Amongst the simplest are realizations of Gaussian or log-normal density fields (e.g. Percival et al. 2001; Mesinger & Furlanetto 2007). More sophisticated ideas have been implemented using second order perturbation theory (Monaco, Theuns & Taffoni 2002; Scoccimarro & Sheth 2002; Kitaura & Heß 2013; Manera et al. 2013; Monaco et al. 2013; Kitaura, Yepes & Prada 2014). In a different approach, the use of simulation particles to mimic galaxy clustering has been adopted in several studies by invoking a prescription based on the local DM density (Cole et al. 1998; Cabré & Gaztañaga 2009). White, Tinker & McBride (2014) recently proposed an extension to these ideas, using low-resolution particle-mesh N -body simulations to generate large numbers of realisations of the DM density field, which is sampled to mimic the clustering of different samples of DM haloes. However, none of these approaches has fully achieved the combination of simplicity and accuracy desirable when modelling a given cosmological experiment.

The objective of this paper is to present and test a scheme to create mock catalogues of the large-scale distribution of galaxies in a computationally inexpensive way.¹ Our approach uses the DM density field extracted from N -body simulations to predict a halo population whose properties can be derived from a higher resolution simulation or analytically. This effectively extends the halo mass resolution of N -body simulations down to an arbitrarily low limit. Note, a very similar method has been developed independently by de la Torre & Peacock (2013). Here, we present an alternative formulation of the method together with an enhanced suite of tests, which focus on larger scales, specifically the creation of covariance matrices for the

two-point correlation function on the baryonic acoustic oscillation (BAO) scale.

The structure of this paper is as follows. In Section 2, we provide details of our method along with its theoretical motivation. In Section 3, we apply our algorithm to an ensemble of N -body simulations to investigate the limitations and range of applicability of the procedure. The haloes created by our algorithm can be combined with higher mass haloes which are identified directly in the simulations to extend the range of halo masses in the simulation box. The resulting hybrid halo catalogue can be fed into a semi-analytic galaxy formation model or combined with an HOD model. To illustrate the feasibility of the idea, our procedure is shown in action in Section 4, where we use an HOD to predict the errors on the clustering of luminous red galaxies (LRGs). Finally, in Section 5, we present a summary and discussion of our findings.

2 METHOD

In this section, we present the algorithm used to generate a halo population from the density field in DM simulations. We start by giving the motivation and main ideas behind the method (Section 2.1) and then we outline the steps to be followed in a practical implementation of the technique (Section 2.2).

2.1 Theoretical motivation

Assuming that the abundance of haloes at a given position, \mathbf{x} , a function of the local underlying non-linear DM density alone, we can write the number density field of haloes of mass, M , as

$$\delta_{h,R}(\mathbf{x}, M) = f_M(\delta_{dm,R}(\mathbf{x})), \quad (1)$$

where f_M is a smooth and arbitrary function (which could, in principle, be different for haloes of different mass), $\delta(\mathbf{x})$ is the density contrast, defined as $\rho(\mathbf{x})/\langle\rho(\mathbf{x})\rangle - 1$, where $\rho(\mathbf{x})$ is the density at \mathbf{x} and $\langle\rho(\mathbf{x})\rangle$ is the mean density, and the subscripts h and dm refer to the density field of haloes and DM, respectively. R is the scale on which both density fields are smoothed and is set by the smallest scale on which equation (1) holds.

On sufficiently large scales the DM density approaches the mean value, $|\delta_{dm}(\mathbf{x}|R)| \ll 1$, which allows us to express equation (1) as a Taylor series expansion in δ_{dm} (see e.g. Fry & Gaztanaga 1993):

$$\delta_{h,R}(\mathbf{x}, M) = \sum_{k=0}^{\infty} \frac{b_k(M)}{k!} \delta_{dm,R}^k(\mathbf{x}), \quad (2)$$

where the subscript R denotes the smoothing scale. The coefficients b_k are usually referred to as the bias parameters. In particular, b_1 is known as the linear bias. These parameters can be derived analytically from collapse models (Mo, Jing & White 1997) or measured directly from N -body simulations (Angulo, Baugh & Lacey 2008b). Note that the functional form adopted in equation (2) is not the only possibility; for instance, de la Torre & Peacock (2013) invoke an exponential model. However, both models converge asymptotically on large scales.

It is straightforward to write down an expression for the expected number density of haloes of a given mass in a region in which the DM density field has been smoothed,

$$\begin{aligned} N_{h,R}(\mathbf{x}, M) &= \langle N_{h,R}(\mathbf{x}, M) \rangle \\ &\times \left[b_0 + b_1 \delta_{dm,R}(\mathbf{x}) + \frac{b_2}{2} \delta_{dm,R}^2(\mathbf{x}) + \mathcal{O}(\delta_{dm,R}^3) - 1 \right]. \end{aligned} \quad (3)$$

¹ The method was introduced in the PhD thesis of the lead author (Angulo 2008).

Here, the brackets $\langle \rangle$ denote an average over all smoothing regions, and so $\langle N_{h,R}(\mathbf{x}, M) \rangle$ is the standard halo mass function. Note that b_0 is set by requiring that the expression inside square brackets is equal to the unity when averaged over all regions. As we discuss below, it is possible to use this expression to construct a halo density field which displays the halo abundance and clustering properties expected in an N -body simulation.

2.2 Implementation

It can be seen clearly that, under our assumptions, the expected abundance of haloes at a given location (equation 3), depends on three quantities: (i) the DM density field at the location, (ii) the mean number density of haloes of a given mass and (iii) the bias parameters as a function of halo mass. The core of our method is that it is possible to recover the underlying DM density field directly from simulations with high fidelity (even in the case of low-resolution simulations), and also that both the bias parameters and the mean number of haloes can be calculated easily, either analytically or from high-resolution N -body simulations (which will typically be of much smaller volume than the simulations we wish to populate with haloes). As a consequence of bringing these ingredients together, a population of DM haloes, which spans an *arbitrarily wide range of masses*, can be created.

Subject to the validity of our assumptions as discussed below, the population of haloes generated using our method has, by construction, the correct abundance and clustering on scales larger than the chosen smoothing scale. In fact, not only are the two-point statistics reproduced for the halo distribution but, in principle, the correct volume-averaged higher order statistics are also recovered (as can be seen from equation 2). We call the haloes generated using our technique ‘sub-resolution’ haloes. In the next section, we will test our method by applying it to generate all of the haloes in a simulation volume, in order to assess the validity of the approach. In practice, we will use a hybrid halo catalogue made up of haloes which are resolved directly in the simulation, and lower mass haloes which are added using our technique, hence the name ‘sub-resolution’.

There are, inevitably, limitations in the sub-resolution halo catalogues which arise from our simplified treatment of halo formation. First, our expressions are only strictly valid when the density contrast is small, $\delta_{\text{dm}} \ll 1$. This sets a minimum smoothing scale that can be used which in turn determines the smallest scale on which the halo clustering can be reproduced. Secondly, in a practical implementation, equations (2) and (3) have to be truncated at a given order which creates two problems: (i) The clustering statistics of orders higher than the truncation cannot be reproduced accurately. There will be some information about the higher order clustering of haloes since we are applying our technique to the evolved density field in the DM simulation. (ii) In underdense regions, equation (3) can predict a negative number of massive haloes. This would happen in an expansion truncated at first order if $b(M) > 1$ and $\delta_{\text{dm}} < -1/b$, implying $\delta_{\text{nh}} < -1$. Consequently, we expect our procedure to break down for haloes more massive than M_* . These restrictions are not prohibitive though, since our algorithm is primarily designed to add low-mass and therefore low-bias haloes. Moreover, as discussed by de la Torre & Peacock (2013), the small-scale clustering in magnitude limited galaxy samples tends to be dominated by satellite galaxies hosted by massive haloes. Since these are typically resolved directly in N -body simulations, a relatively large smoothing scale does not introduce noticeable artefacts even in the small-scale clustering of catalogues constructed using our algorithm. We

investigate and quantify these restrictions in the following sections where we present our algorithm in action.

We note that while we were preparing this manuscript, de la Torre & Peacock (2013) independently developed and explored essentially the same idea as the one presented here. There are, however, differences in the implementation and in the applications in which we focus. Here, we confirm the high accuracy of the method in general, but we extend it and apply it to the modelling of large-scale clustering with particular attention on BAO analyses and relevant covariance matrices.

3 TESTING THE METHOD

We now apply and test the procedure outlined in the previous section. In Section 3.1, we provide details of the implementation of the method and present some general characteristics of the resulting halo catalogues. In Section 3.2, we show the results of three basic tests and a comparison with haloes identified directly in a high-resolution N -body simulation. The sub-resolution catalogues we generate in this section cover a wide range of halo masses, including those of haloes that are resolved in the N -body simulations. The goal in this section is to establish the range of validity of our method in view of the assumptions and approximations which underpin it. As we pointed out in the previous section, the actual implementation of the method (Section 4) will make use of ‘hybrid’ halo catalogues in which the higher mass haloes are those directly resolved in the simulation and the lower mass ones are the ‘sub-resolution’ population generated by our algorithm.

3.1 The sub-resolution halo catalogue

To characterize the performance of our method, we use the simulations described in Angulo et al. (2008a). These include a suite of 50 low-resolution simulations, referred to as the L-BASICC ensemble. Each of these modelled the gravitational interactions between 448^3 particles of mass $1.85 \times 10^{12} h^{-1} M_\odot$ in a periodic box of side $1340 h^{-1} M_\odot$. We also employ a higher resolution run, dubbed BASICC, which used 1448^3 particles of mass $5.49 \times 10^{10} h^{-1} M_\odot$, also in a periodic box of side $1340 h^{-1} M_\odot$. Note that one of the L-BASICC simulations has exactly the same initial density field as used in the BASICC run. Haloes are identified in the simulation outputs using a Friends-of-Friends (FoF) percolation algorithm (Davis et al. 1985). We stress that it is computationally inexpensive to carry out such a set of low-resolution simulations. Each of the L-BASICC runs would only take approximately 150 CPU-hours on modern supercomputers.

Following the algorithm described in Section 2, we computed a sub-resolution halo catalogue for the three outputs ($z = 0, 0.5$ and 1) of each of the 50 simulations in the L-BASICC ensemble. This process is made up of three steps. The first is the construction of the DM density field in the simulations. This is performed by placing particles on to a grid using the nearest grid point mass assignment scheme (Hockney & Eastwood 1981). We use a grid of 256^3 cells (the cell size is $5.2 h^{-1}$ Mpc) which is set so that $\langle \delta^2 \rangle \sim 1$. We therefore expect to obtain an inaccurate estimation of the halo clustering on scales smaller than a few times the size of the grid cell. Note that de la Torre & Peacock (2013) followed an alternative path and constructed the DM density field from the resolved halo population. As there are fewer haloes than particles, there is a larger amount of noise in the reconstructed density field. Dealing directly with simulation particles would also allow the use of Lagrangian smoothing techniques (Abel, Hahn & Kaehler 2012;

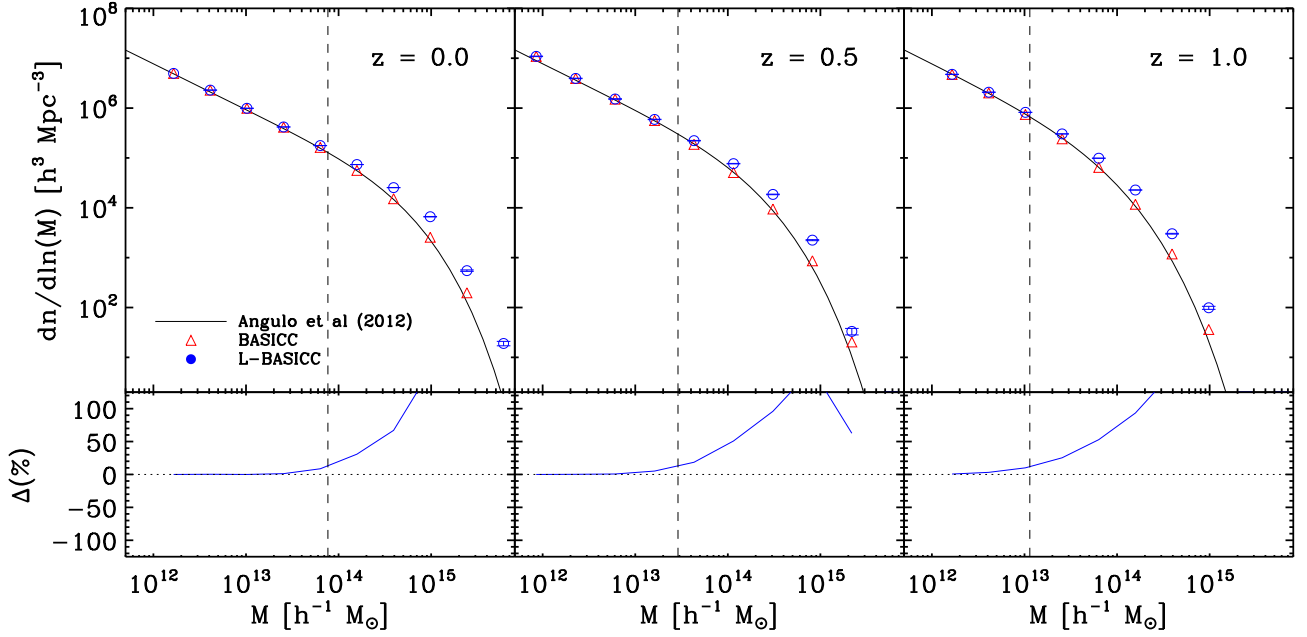


Figure 1. A test of the method showing the comparison between the mass function of FoF haloes resolved in the BASICC simulation (red triangles) and the mean of sub-resolution halo catalogues built from the L-BASICC ensemble (blue circles). Note that in this test case sub-resolution haloes are generated across the whole mass range plotted to assess the range of validity of the technique. Top: the mean number of haloes per logarithmic mass bin, as a function of their mass. The error bars show the dispersion from applying our algorithm to 50 simulations (L-BASICC). Each column shows a different redshift and the empirical fit to various N -body results from Angulo et al. (2012, solid lines). The vertical dashed lines indicate the halo mass at which the number of haloes resolved directly in the simulations and those created by our algorithm first differ by 10 per cent moving in the direction of increasing mass. The relative difference between these catalogues is shown by the curves in the bottom panels.

Shandarin, Habib & Heitmann 2012). Although we do not use such techniques here, they have been recently shown to have extremely low discreteness noise (Angulo et al. 2013a; Angulo, Hahn & Abel 2013b) which could improve the accuracy of our method in the future.

The next step is to tabulate the halo bias parameters and the number density of haloes as a function of mass. We extract these relationships from the higher resolution BASICC simulation in logarithmic mass bins of width $\Delta \log_{10} M = 0.426$. Both quantities are computed by smoothing the haloes and DM field in 256^3 cells and then averaging the values across the grid.

Finally, these three quantities are brought together to compute the expectation value for the number density of haloes on every point of the grid. There are several points regarding the placement of haloes that are worth noting. (i) The actual number of haloes in each cell is generated from a Poisson distribution with the expectation value as the mean. In doing this, we have also neglected the covariance between halo mass bins, which is justified given the box size of our simulations (Smith & Marian 2011). (ii) The haloes are placed randomly within each of the smoothing volumes. (iii) Each of these haloes is given a peculiar velocity equal to the mean velocity of the DM particles within the same cell. Alternatively, one could use some sort of interpolation scheme such as that used by de la Torre & Peacock (2013). (iv) Equation (3) is truncated at linear order.

As a result of following this procedure, we obtained 50 independent sub-resolution halo catalogues at the three redshifts mentioned above. Each contains approximately 17 million haloes with mass between 5.48×10^{11} and $1 \times 10^{16} h^{-1} M_{\odot}$ at $z = 0$. In the following subsection, we will explore the properties of these catalogues.

3.2 Abundance and clustering

In this subsection, we compare the abundance and clustering strength in our sub-resolution halo catalogues with the same quantities measured using haloes directly identified by an FoF algorithm in a high-resolution simulation (for details of the FoF catalogues see Angulo et al. 2008a).

The upper panels of Fig. 1 show the differential halo mass function from our catalogues (blue filled circles) and that from FoF haloes identified in the BASICC simulation. In the lower panels, we can see the differences between the two populations more clearly on a linear scale. This figure shows that there is excellent agreement between the number of haloes generated using our algorithm and that obtained directly in the higher resolution N -body simulation. This represents an initial validation of the ideas and their implementation presented in this paper. Our method predicts an abundance of haloes that agrees with the direct simulation results to better than 10 per cent for objects of mass $M < 7.51 \times 10^{13} h^{-1} M_{\odot}$ at $z = 0$, $M < 2.7 \times 10^{13} h^{-1} M_{\odot}$ at $z = 0.5$ and $M < 1.14 \times 10^{13} h^{-1} M_{\odot}$ at $z = 1$. There is a strong disagreement between the numbers of sub-resolution and FoF haloes at the high-mass end. This is caused by the fact that equation (2) is inconsistent for highly biased haloes in low-density regions where $\delta_h < -1$ (the problem is alleviated in the low-mass regime where $b \lesssim 1$) and because we have truncated equation (3) at the linear bias term. As a consequence of haloes of a fixed mass becoming more biased with increasing redshift, the mass function of sub-resolution haloes provides an acceptable match to the simulation results (i.e. better than 10 per cent agreement) over a reduced range of masses at high redshift.

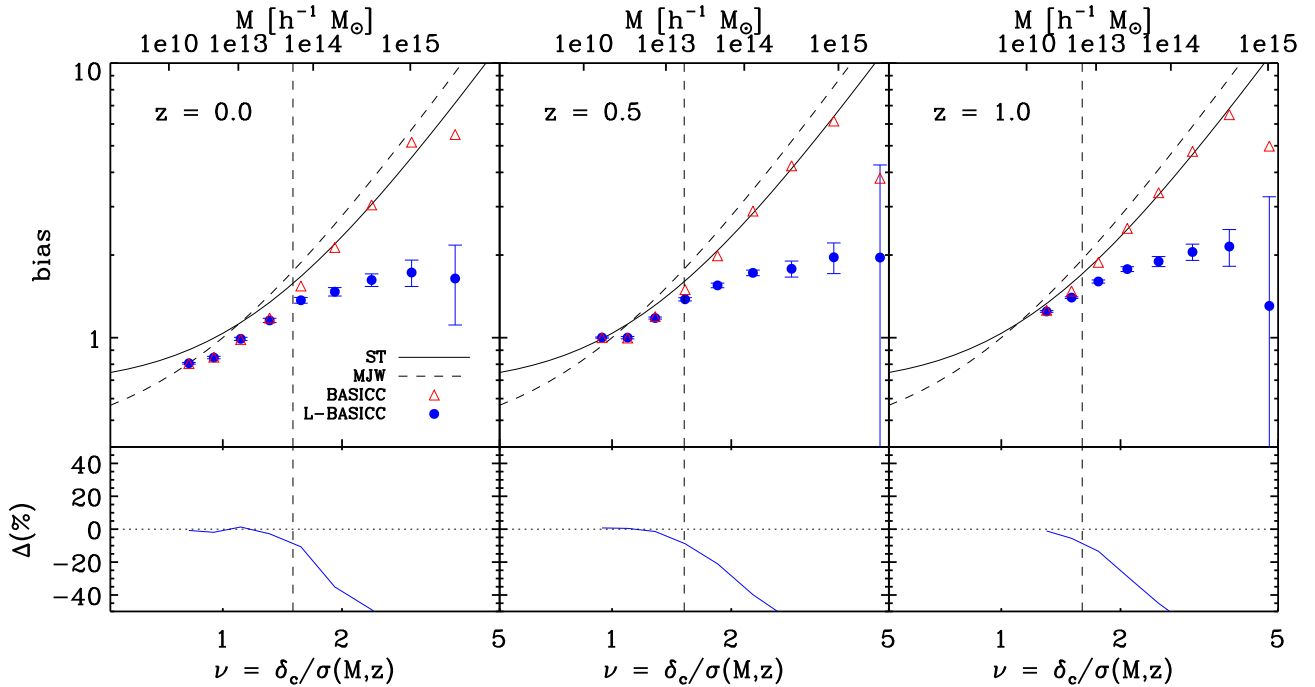


Figure 2. A test of the clustering predicted by our method. In this test case, sub-resolution haloes are generated for the full mass range plotted to assess the range of validity of the technique. Top: the linear bias parameter as a function of halo mass (top axis) and peak height $\delta_c/\sigma(M, z)$ (bottom axis). Each panel shows a different redshift as labelled. The blue filled circles with error bars show the mean bias and the dispersion for the sub-resolution haloes generated by our algorithm applied to the 50 low-resolution L-BASICC simulations. The bias measured from FoF haloes in the high-resolution BASICC simulation is shown as red triangles. Theoretical predictions from Mo et al. (1997) and Sheth, Mo & Tormen (2001) are also shown using different line styles as labelled. The lower panels show the relative difference between the halo bias from the two different catalogues. The vertical dashed lines show where the clustering of haloes in the sub-resolution catalogue first differs from that measured from the FoF haloes by 10 per cent, moving in the direction of increasing mass.

We extend the comparison by investigating the clustering strength in the sub-resolution catalogues. Each column of Fig. 2 displays the linear bias parameter as a function of the peak height,² $\delta_c/\sigma(M, z)$ on the bottom axis and as a function of mass on the top axis. Note that we compute the linear bias, b , by smoothing the halo and DM density fields in cells of size $167 h^{-1}$ Mpc, and taking the ratio i.e. $b^2 = \langle \delta_{\text{hh}}^2 / \delta_{\text{mm}}^2 \rangle$. As in the previous plot, the vertical lines indicate the maximum halo mass at which the result from the sub-resolution haloes agrees to within 10 per cent with that of the resolved haloes. Similar to the behaviour seen in Fig. 1, at the high-mass end, the sub-resolution haloes fail to reproduce the clustering measured from the resolved FoF catalogues, which suggests a common origin for the discrepancies seen in the abundance and clustering of sub-resolution haloes at high masses. Note that the 10 per cent-difference mass limit derived from the clustering comparison is slightly smaller than that derived from the mass function at; $z = 0$ $M_{\text{max}} = 5.23 \times 10^{13} h^{-1} M_{\odot}$ while at $z = 0.5$ and $z = 1$ $4.3 \times 10^{13} h^{-1} M_{\odot}$ and $6.73 \times 10^{12} h^{-1} M_{\odot}$, respectively. Again, the very good agreement apparent at low masses validates our approach.

Finally, we explore the spherically averaged clustering of the halo catalogues in redshift space. Fig. 3 shows the ratio between the linear bias parameter measured in redshift space and that measured in real space for the sub-resolution haloes and for the FoF haloes. In

linear perturbation theory, this quantity is equivalent to the square root of the Kaiser ‘boost factor’ (Kaiser 1987):

$$f = \left(1 + \frac{2}{3}\beta + \frac{1}{5}\beta^2 \right), \quad (4)$$

where $\beta = \Omega_m(z)^{0.55}/b$, with Ω_m denoting the matter density parameter and b the linear bias parameter. This expression is overplotted in Fig. 3 for comparison. Note that, in practice, the Kaiser factor is only attained asymptotically (Jennings, Baugh & Pascoli 2011), so we again measure the bias in redshift space by comparing densities in grid cells of size $167 h^{-1}$ Mpc. Furthermore, there is no reason to expect this relation to hold for highly non-linear objects corresponding to high peaks (e.g. Angulo et al. 2005).

Despite the scatter among the sub-resolution halo catalogues, we find reasonably good agreement between the theoretical expectations and the measurements from the BASICC FoF haloes. Given the comparisons presented in previous figures, it is not surprising to see the differences for haloes corresponding to high peaks. Nevertheless, our scheme to assign peculiar velocities to haloes performs satisfactorily in the regime where the abundance and clustering in real space are properly imprinted on the sub-resolution catalogues. This is a remarkable success, extending the usability of the method to modelling redshift-space distortions.

Now that we have established the range of mass scales over which the sub-resolution halo catalogues give an accurate reproduction of the results seen in high-resolution N -body simulations, in the application of the method presented in the next section, we will use a hybrid halo catalogue, made up of directly resolved haloes and lower-mass, sub-resolution haloes.

² Here, δ_c is the threshold for collapse in linear perturbation theory and $\sigma(M, z)$ is the linear theory rms variance in the density field smoothed on a scale enclosing mass M at redshift z .

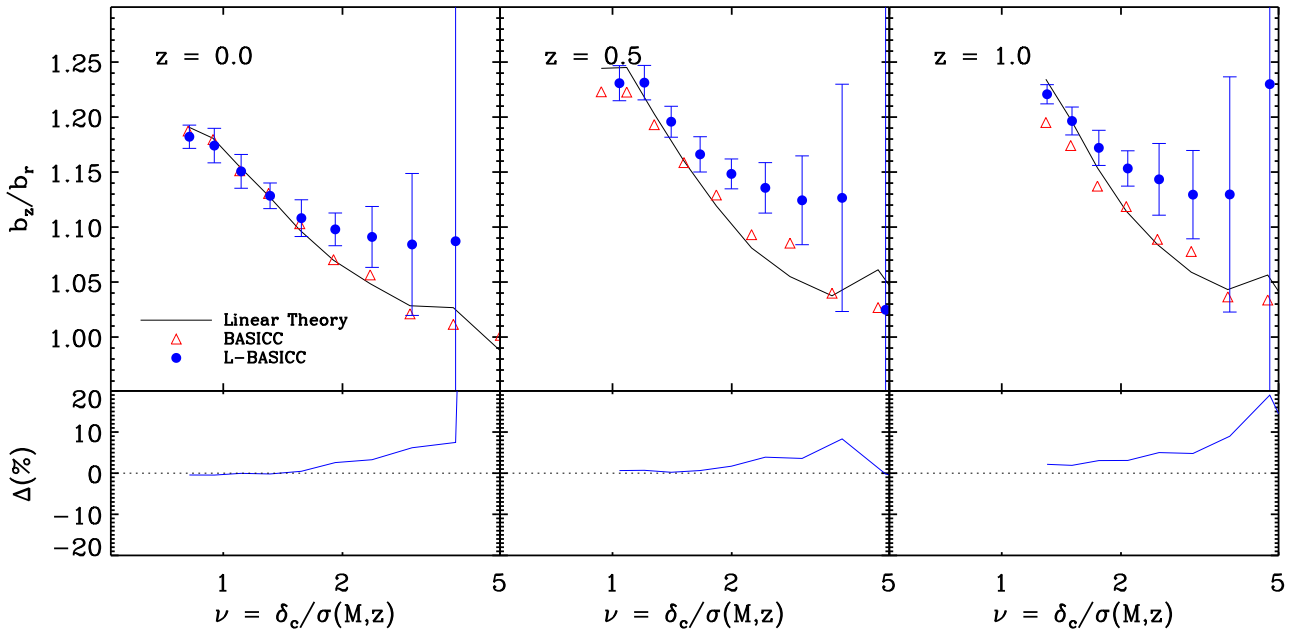


Figure 3. A test of the performance of the model predictions for the clustering measured in redshift space. In the test case, sub-resolution haloes are generated across the full range of halo masses plotted to evaluate the performance of the technique. The linear bias parameter for haloes measured in redshift space, b_z , divided by that measured in real space, b_r , as a function of the peak height (bottom axis). The mean and dispersion of this quantity, measured from our sub-resolution haloes in an ensemble of low-resolution simulations, is displayed using blue symbols with error bars. We display the results measured from FoF haloes in the BASICC simulation using red triangles. For comparison, we have also included the prediction based on linear theory (solid line, see equation 4). The lower panels show the relative difference between the sub-resolution results and those obtained from the BASICC run.

4 APPLICATION: LARGE-SCALE CLUSTERING OF LRGs

Recently, the clustering of LRGs has been of great importance in probing different cosmological scenarios. The low number density but strong clustering of these galaxies means that the spatial distribution of LRGs can be mapped over vast regions of the sky at relatively low-observational cost. A large survey volume enables tight constraints to be placed on cosmological parameters, in particular by measuring the BAO feature (Eisenstein et al. 2005; Cabré & Gaztañaga 2009; Gaztañaga, Miquel & Sánchez 2009; Sánchez et al. 2009). Unfortunately, there is still an incomplete understanding of the errors associated with clustering measurements on large scales. A realistic model for the uncertainties, including systematic errors, is crucial to extract cosmological constraints from the data, since the determination of the best-fitting model, together with the allowed regions in cosmological parameter space, depend sensitively on the availability of an accurate covariance matrix.

Semi-analytical modelling and observational evidence suggest that LRGs not only populate very massive haloes but they can also be found in haloes with masses as small as $10^{11} h^{-1} M_\odot$ albeit with a low probability (Almeida et al. 2008; Wake et al. 2008). Therefore, the modelling of LRG clustering, and the BAO feature imprinted on it, requires huge simulations with a considerable dynamic range in mass. Although such extremes can be achieved in modern supercomputers these tend to be one-off runs and the computational cost is enhanced to inaccessible levels when studying uncertainties or subtle features present in the clustering which require many realizations.

In this section, we approach this problem using the algorithm we described above. Specifically, we generate 50 LRG catalogues to compute the mean and variance of the two-point correlation

function. Details of the creation of the LRG catalogues as well as the clustering measurements are presented in the following subsections.

4.1 The haloes and LRG catalogues

The starting point in the creation of the LRG mock catalogues is to predict the abundance and spatial distribution of the DM haloes that are likely to host such galaxies. For this purpose, we created 50 hybrid halo catalogues, each one spanning 4 orders of magnitude in mass within a volume of $2.4 h^{-3} \text{ Gpc}^3$ at $z = 0.5$.

The halo catalogues are hybrid in the sense that they consist of two types of haloes. The high mass ones ($M > 1.85 \times 10^{13} h^{-1} M_\odot$) correspond to objects identified directly using an FoF algorithm, with at least 10 particles, in each of the L-BASICC simulations. Then, smaller mass sub-resolution haloes ($5.48 \times 10^{11} < M/(h^{-1} M_\odot) < 1.85 \times 10^{13}$) were created using the algorithm described in Section 2. We recall that our method is accurate to better than the 10 per cent level for this mass range. In this way, we are effectively extending the dynamic range of the L-BASICC simulations towards lower masses. Combining the two types of haloes also eliminates the need to reproduce high-mass haloes in the sub-resolution catalogues, which proved to be troublesome (see Section 3.2).

Once we have generated the catalogues that contain all the haloes that are expected to host LRGs, we use an HOD model to determine how many LRGs on average populate each DM halo (for a review of the halo model see Cooray & Sheth 2002). Following Wake et al. (2008), we can express the mean number of central LRGs, N_c as a function of the host halo mass, M_{halo} as

$$\langle N_c | M_{\text{halo}} \rangle = \exp(-M_{\text{min}}/M_{\text{halo}}), \quad (5)$$

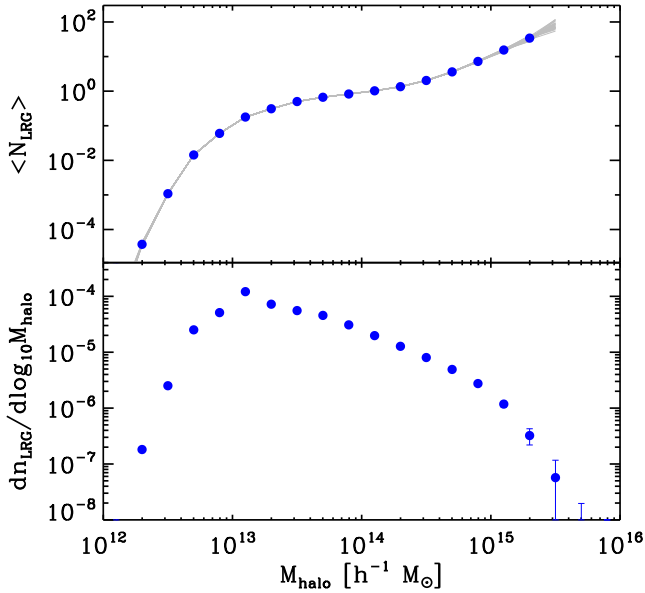


Figure 4. The application of the method to the construction of LRG catalogues. Here, we use a hybrid catalogue comprised of sub-resolution haloes and haloes which are directly resolved in the simulations. Top: the mean number of LRGs per halo as a function of the host halo mass in our simulations. Bottom: the number density of LRGs per decade in host halo mass. The grey lines in the top panel show the quantities in each of our simulations while the blue symbols show the mean and dispersion.

and the mean number of satellite LRGs, N_s as

$$\langle N_s | M_{\text{halo}} \rangle = (M_{\text{halo}} / M_1)^\alpha. \quad (6)$$

Consequently, the total number of LRGs has an expected value of

$$\langle N_{\text{LRG}} | M_{\text{halo}} \rangle = \langle N_c | M_{\text{halo}} \rangle [1 + \langle N_s | M_{\text{halo}} \rangle], \quad (7)$$

where α , M_{min} and M_1 are, in principle, free parameters that can be constrained either by comparing to observational estimates of clustering observations or through semi-analytical galaxy formation modelling. Indeed, Wake et al. (2008), using the measured clustering of 2SLAQ LRGs (Cannon et al. 2006), found that the best-fitting values of M_{min} , M_1 and α are 2.19×10^{13} , $2.82 \times 10^{13} h^{-1} \text{M}_\odot$ and 1.86, respectively. In the second step, we assume that N_c and N_s follow Poisson distributions with means given by equations (5) and (6), which, combined with the values from Wake et al., allows us to place LRGs within our hybrid halo catalogues. Note that the alternative approach of applying full semi-analytic modelling to the hybrid halo catalogues could also have been taken.

Each of our final catalogues contain 398 963 galaxies, or equivalently, a number density of $1.66 \times 10^{-4} h^3 \text{Mpc}^{-3}$. Even though, on average, there is less than one LRG per sub-resolution halo, together the sub-resolution haloes host a total of 114 243 galaxies which represents 28 per cent of whole LRG sample. Fig. 4 shows the resulting mean number of LRG in our catalogues per halo (top panel) as well as the total number LRGs (bottom panel), in both cases as a function of the mass of the host halo.

4.1.1 Correlation function of LRGs

At this point, we are now in a position to investigate the clustering of LRGs. We measure the correlation function using fast Fourier transforms. This approach is considerably more efficient than computations carried out in configuration space, when one is interested

in the correlation function on large scales measured from catalogues containing a large number of objects.

In brief, the method uses a pixelization of the density field from which the (real- and redshift-space) spherically averaged correlation function can be estimated from the amplitude of Fourier modes as

$$\xi(\mathbf{r}) = \mathcal{F}^{-1} \{ |\mathcal{F}[\delta(\mathbf{x})]|^2 \}, \quad (8)$$

where $\delta = (n(\mathbf{x}) - \langle n \rangle) / \langle n \rangle$ is the density fluctuation (in real or redshift space) on a grid, and $\mathcal{F}[\delta]$ is its Fourier transform. Vertical bars denote the modulus of a complex field, and \mathcal{F}^{-1} an inverse Fourier Transform. We carry out this operation using a fast Fourier transform with a grid of dimensions $N_{\text{grid}} = 1024$, which corresponds to $1.3 h^{-1} \text{Mpc}$ for the L-BASICC simulation box size. This method gives an accurate estimation of the correlation function for scales larger than a few grid cells.

Fig. 5 shows the result of applying this procedure to compute the correlation function for LRGs in each of our 50 catalogues.

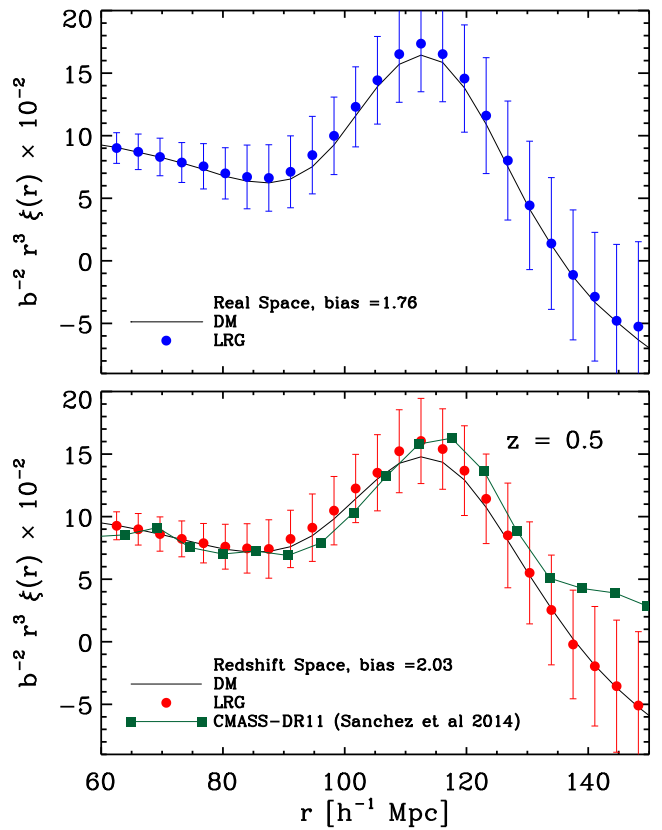


Figure 5. The mean and variance of the correlation function measured from LRG samples constructed from our hybrid catalogues of resolved and sub-resolution haloes at $z = 0.5$. The top panel shows the real space correlation function and the bottom panel shows redshift space. In each case, filled circles indicate the clustering measured from the LRG catalogues created by populating hybrid halo catalogues (i.e. the mixture of sub-resolution and FoF haloes) with the LRG HOD inferred by Wake et al. (2008). The solid lines show the mean correlation function measured from the DM particles of all 50 L-BASICC realisations. To allow a full comparison, we have divided each measurement (and the respective variance) by a constant bias, measured in the range $r = [60-70] h^{-1} \text{Mpc}$, and by the expected Kaiser boost factor in the case of redshift-space measurements. Note that we display $\xi(r) \times r^3$ on the y-axis to enhance the appearance of the BAO peak. For comparison, in the bottom panel, we also display the measurements of Sánchez et al. (2014) for the correlation function of BOSS-CMASS sample at $z = 0.56$ as solid green squares.

The top panel displays the measurements in real space while the bottom panel shows redshift space. In both cases, the mean and variance of the measurements are indicated by the filled circles and error bars. In order to assess our results, we have measured the correlation function for subsets of DM particles at $z = 0.5$ from the L-BASICC simulations. We display the mean of all 50 simulations in real and redshift space as a solid line in the top and bottom panels, respectively. This allows us to compare the form of the correlation function measured from our LRG catalogues with that of the underlying DM distribution. Note that the y-axis shows $\xi \times r^3$ instead of ξ , as in this way the acoustic peak is highlighted. In addition, the results (including the errors) in both real and redshift space have been renormalized as described in the figure caption.

By comparing the correlation function of LRGs with that of the DM, we can see the effects of galaxy bias. Fig. 5 shows that the respective correlation functions, after applying a scaling in amplitude, agree fairly well with one another, implying that the LRG bias is approximately scale independent over the range of pair separations plotted. There is a small residual dependence of the bias on scale in real space which seems to be accentuated in redshift space. Although the discrepancy is not significant given the size of the errors associated with the simulation volume, using a simulation with 10 times larger volume and 3375 times more particles, Angulo et al. (2013c) recently showed that distortions of this type are expected in biased tracers of the DM field (see also Padmanabhan & White 2009; Mehta et al. 2011). This scale-dependent bias, absent in approaches that simply apply a biasing scheme on top of the DM field, is an example of the benefits of an hybrid approach like the one proposed here.

In addition, in the bottom panel of Fig. 5, we show the correlation function of the BOSS-CMASS sample at $z = 0.56$, as measured by Sánchez et al. (2014). Our predictions are in very good agreement with the data, and all but the last two points located at $r > 140 h^{-1} \text{ Mpc}$ agree at the 1σ level. The residual difference is likely to be caused by the differences between the fiducial cosmology employed in our simulations and that preferred by the data. This supports the idea that our method is accurate enough for modelling and interpretation of the BAO signal.

In Fig. 6, we compare the variance measured from our ensemble of LRG catalogues (filled circles) with that measured from the DM samples (triangles). By comparing both measurements, we illustrate the importance of shot-noise in the expected variance. The dotted line shows a theoretical prediction for the variance based on power-spectrum measurements which include the effects of a finite number of modes, discreteness noise, bias and binning (see Sánchez, Baugh & Angulo 2008, for more details). The theoretical predictions by Sánchez et al. (2008) provide a fairly good match to the variance in our LRG samples, showing that our catalogues have the expected variance.

We extend this comparison in Fig. 7 in which we display the normalized covariance matrix (Cohn 2006; Smith, Scoccimarro & Sheth 2008), $C_\xi(r, r') \equiv ((\xi(r) - \bar{\xi})(\xi(r') - \bar{\xi}))/(\sigma(r)\sigma(r'))$, in real space (left-hand plot) and in redshift space (right-hand plot). The above diagonal part of the plot shows the expected covariance as computed following Sánchez et al. (2008). The below diagonal part shows the covariance for the LRG catalogues. The non-diagonal parts of the covariance matrix show a reasonably good agreement between the mock LRG catalogues and the theoretical expectations, similar to the case of the comparison of the variances. The agreement is not perfect and our LRG catalogues show slightly stronger off-diagonal correlations than the expectation and also show more structure, in particular an excess correlation at the BAO location.

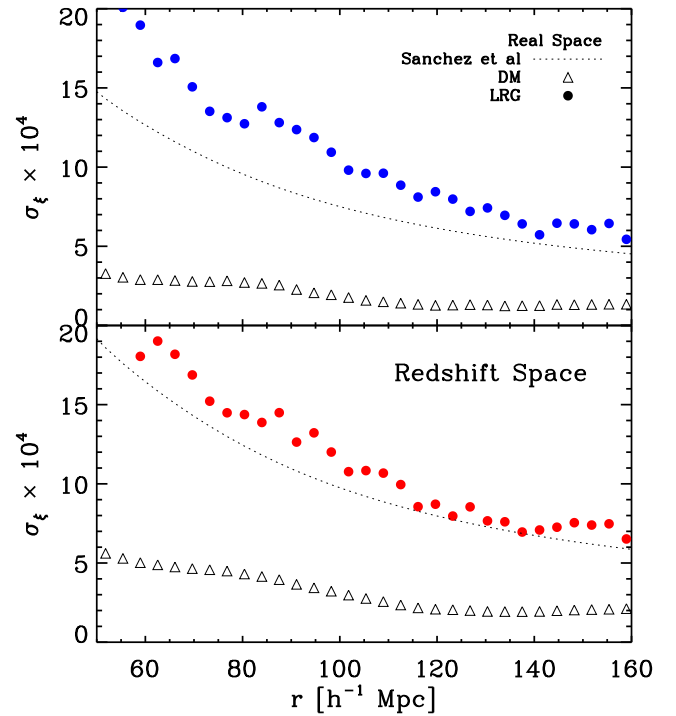


Figure 6. The variance in the two-point correlation function measured from 50 LRGs catalogues in real space (top) and redshift space (bottom) at $z = 0.5$, constructed using the hybrid catalogues of resolved and sub-resolution haloes. We also plot a theoretical estimate for the variance from Sánchez et al. (2008, dotted lines)

One possible explanation could be the contribution of the higher order moments of the halo density field, which are present in our LRG samples but absent in the Sanchez et al. predictions. As shown by Angulo et al. (2008b), the higher order moments of haloes differ considerably from those of the DM. As an example, recall that even if the DM density field is Gaussian (i.e. the higher moments are zero), then haloes will have non-zero higher order correlations which contribute to the covariance matrix. Nevertheless, the results are still noisy given the small number of simulations in our ensemble and further investigation is required. In any case, the performance of our catalogues is remarkable and illustrates the feasibility of constructing detailed covariance matrices from computationally cheap N -body simulations that have the correct diagonal terms.

5 SUMMARY

Due to the large volumes that future surveys are expected to map, the resulting measurements of galaxy clustering will be of exquisite accuracy, with the target of distinguishing between different models for the acceleration of the cosmic expansion. The clustering signals predicted by competing models often differ by small amounts. It is therefore essential to understand the systematic and sampling errors associated with the measurements. Only in this way will it be possible to extract robust conclusions from the data. In practice, this challenge can only be met by techniques which make use of cosmological N -body simulations, since this approach gives the best estimate of the contribution of various non-linear effects to the measured clustering.

We have devised and illustrated the feasibility of a scheme that allows the rapid and efficient creation of large numbers of galaxy mock catalogues which are able to resolve all of the galaxies selected

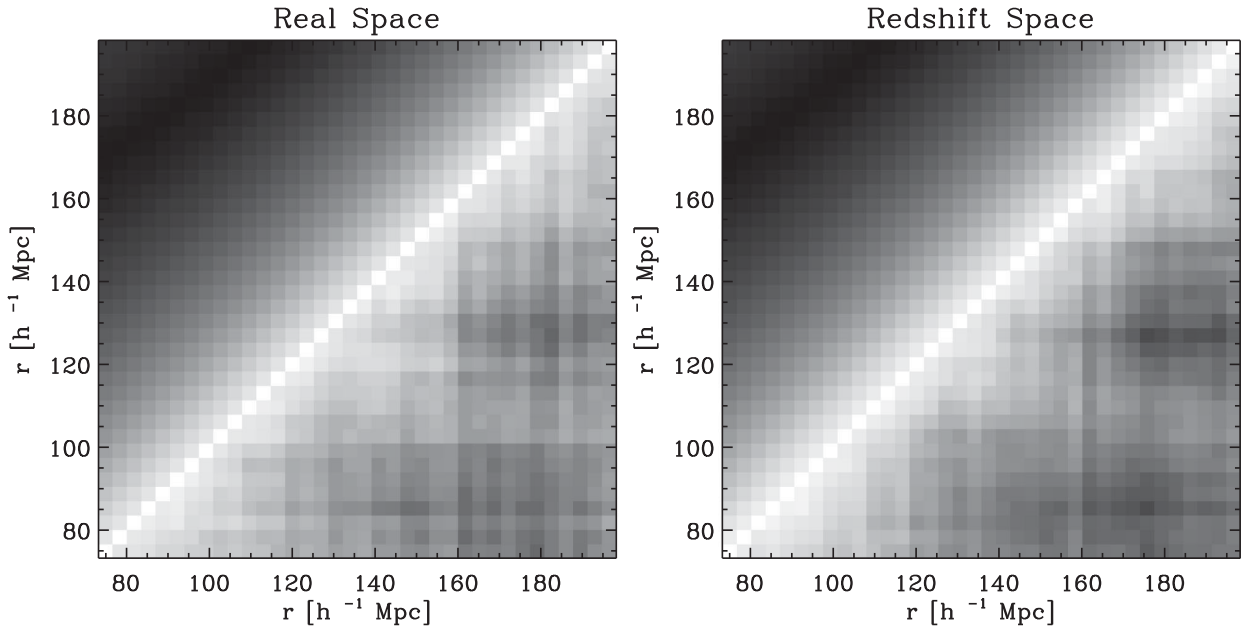


Figure 7. The normalized covariance matrix from our ensemble of LRG mock catalogues, constructed from the hybrid catalogues of resolved and sub-resolution haloes (lower triangular region) and from an analytical prediction from Sánchez et al. (2008) that incorporates the correct volume, bias, number density and binning (upper triangular region). The left-hand plot displays the results in real space while the right displays the covariance matrix in redshift space.

in upcoming surveys. This is done by taking moderate-resolution simulations and effectively extending their dynamic range in halo mass to mimic running a simulation with a substantially larger number of particles. Our method uses the density field extracted from the moderate-resolution N -body simulation and combines it with the bias parameters and mass functions extracted from a higher resolution simulation. In this way, it is possible to predict statistically the expected density field of DM haloes in the moderate-resolution simulation volume. Since low-resolution simulations are relatively easy to generate, our procedure allows the investigation of uncertainties in both the measurements themselves and in the procedures employed to extract robust information from the data.

We have shown that, on large scales, the generated halo population agrees with the population seen directly in a high-resolution simulation over a considerable range of masses. At $z = 0$, in particular, the abundance and clustering strength, in both real and redshift space, of haloes less massive than $7.51 \times 10^{13} h^{-1} M_{\odot}$ agree to within 10 per cent with those computed directly from FoF haloes identified in a high-resolution simulation. For high-mass haloes or at higher redshifts, our procedure performs less satisfactorily.

An interesting application of our scheme is to the creation of hybrid halo catalogues. High-mass haloes can be extracted directly from cosmological N -body simulations, whilst low-mass haloes which lie beyond the grasp of the simulation can be generated using our technique. In this way, we can employ our algorithm in the regime where it works best. As an example, we have created 50 such catalogues from the L-BASICC simulations which are combined with an HOD for LRGs. From the resulting galaxy catalogues, we are successfully able to predict their mean correlation function along with the full covariance matrix. We found that the variance in a sample of DM particles drawn from the simulations and analytical estimates are in agreement with measurements from the LRG catalogues. In spite of this, differences in the off-diagonal terms of the covariance matrix were found.

In the LRG example presented, we extended the halo mass resolution of the L-BASICC runs by a factor of ≈ 30 , since this was all that was demanded by this application. The specifications of the available high-resolution N -body simulation set the limit on the boost attainable in the resolution of the moderate resolution runs. For example, if we had instead chosen to augment the L-BASICC runs using the Millennium-II simulation of Boylan-Kolchin et al. (2009), which modelled the growth of structure using 2160^3 particles in a volume of $(100)^3 h^{-3} \text{Mpc}^3$, then the resulting halo catalogue would be the equivalent of that expected from a simulation employing $28\,944^3$ or more than 24 trillion particles. This is around 50 times larger than the largest number of particles used in an N -body simulation to date. Our approach will allow the production of halo catalogues equivalent to running large numbers of such simulations.

The algorithm presented here is already useful for generating mock observations and in creating covariance matrices, particularly if combined with novel techniques to mimic running very large ensembles of simulations (e.g. Schneider et al. 2011). Nevertheless, it could be enhanced in a number of ways, including the following.

- (i) The placement of haloes within the smoothing volume could be improved by distributing haloes following a given correlation function.
- (ii) Haloes could be placed recursively using different smoothing scales, starting with the whole box and stopping at any desired scale. Here, a parent cell puts constraints on all their child cells, which could be used to include an arbitrary scale-dependent biasing scheme.
- (iii) A more complex biasing scheme could be implemented, which can be calibrated directly using N -body simulations, and could be different for haloes of different mass. Dependences in addition to the density, such as the tidal field, could be taken into account.

(iv) The form of the probability distribution function of haloes given a DM overdensity can be calibrated directly with N -body simulations instead of making the assumption that this has a standard form such as a Poisson distribution.

(v) Similarly, the covariance matrix among different halo mass bins can be used in sampling the DM–halo relationship.

(vi) Extended features can be incorporated such as exclusion effects between haloes, an additional density-dependent velocity dispersion, and the substructure content of haloes.

(vii) The merger history tree associated to each DM halo can be build using, for instance, Sheth & Lemson (1999) or Parkinson, Cole & Helly (2008), which enables more complex and realistic modelling of the galaxy hosted by the sub-resolution haloes. This is another advantage of our method over those that simply sample the DM density field.

Nevertheless, even without these improvements, we expect that the simple technique presented in this paper will improve the understanding and treatment of uncertainties in observations and, therefore, will allow the full potential of measurements of the large-scale distribution of galaxies to be reached.

ACKNOWLEDGEMENTS

The calculations for this paper were performed on the ICC Cosmology Machine, which is part of the DiRAC Facility jointly funded by STFC, the Large Facilities Capital Fund of BIS, and Durham University. The authors acknowledge the support of the STFC rolling grant in theoretical cosmology ST/F001166/1. CSF acknowledges support from ERC Advanced Investigator grant, COSMIWAY. For questions about accessing the simulation data used in this paper, please contact the first author.

REFERENCES

Abel T., Hahn O., Kaehler R., 2012, *MNRAS*, 427, 61
 Alimi J.-M. et al., 2012, preprint ([arXiv:1206.2838](https://arxiv.org/abs/1206.2838))
 Almeida C., Baugh C. M., Wake D. A., Lacey C. G., Benson A. J., Bower R. G., Pimbblet K., 2008, *MNRAS*, 386, 2145
 Angulo R. E., 2008, PhD thesis, Durham Univ.
 Angulo R., Baugh C. M., Frenk C. S., Bower R. G., Jenkins A., Morris S. L., 2005, *MNRAS*, 362, L25
 Angulo R. E., Baugh C. M., Frenk C. S., Lacey C. G., 2008a, *MNRAS*, 383, 755
 Angulo R. E., Baugh C. M., Lacey C. G., 2008b, *MNRAS*, 387, 921
 Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, *MNRAS*, 426, 2046
 Angulo R. E., Chen R., Hilbert S., Abel T., 2013a, preprint ([arXiv:1309.1161](https://arxiv.org/abs/1309.1161))
 Angulo R. E., Hahn O., Abel T., 2013b, *MNRAS*, 434, 3337
 Angulo R. E., White S. D. M., Springel V., Henriques B., 2013c, preprint ([arXiv:e-prints](https://arxiv.org/abs/1309.1161))
 Baugh C. M., 2006, *Rep. Prog. Phys.*, 69, 3101
 Baugh C. M., 2008, *Phil. Trans. R. Soc. A*, 366, 4381
 Benson A. J., 2010, *Phys. Rep.*, 495, 33
 Beutler F. et al., 2011, *MNRAS*, 416, 3017
 Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, *MNRAS*, 398, 1150
 Cabré A., Gaztañaga E., 2009, *MNRAS*, 393, 1183

Cannon R. et al., 2006, *MNRAS*, 372, 425
 Cohn J. D., 2006, *New. Astron.*, 11, 226
 Cole S., Hatton S., Weinberg D. H., Frenk C. S., 1998, *MNRAS*, 300, 945
 Cooray A., Sheth R., 2002, *Phys. Rep.*, 372, 1
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
 de la Torre S., Peacock J. A., 2013, *MNRAS*, 435, 743
 Eisenstein D. J. et al., 2005, *ApJ*, 633, 560
 Fosalba P., Gaztañaga E., Castander F. J., Manera M., 2008, *MNRAS*, 391, 435
 Fry J. N., Gaztanaga E., 1993, *ApJ*, 413, 447
 Gaztañaga E., Miquel R., Sánchez E., 2009, *Phys. Rev. Lett.*, 103, 091302
 Guzzo L. et al., 2008, *Nature*, 451, 541
 Hockney R. W., Eastwood J. W., 1981, *Computer Simulation Using Particles*. McGraw-Hill, New York
 Jennings E., Baugh C. M., Pascoli S., 2011, *MNRAS*, 410, 2081
 Kaiser N., 1987, *MNRAS*, 227, 1
 Kim J., Park C., Gott J. R., Dubinski J., 2009, *ApJ*, 701, 1547
 Kitaura F.-S., Heß S., 2013, *MNRAS*, 435, L78
 Kitaura F.-S., Yepes G., Prada F., 2014, *MNRAS*, 439, L21
 Kuhlen M., Vogelsberger M., Angulo R., 2012, *Phys. Dark Universe*, 1, 50
 Laureijs R. et al., 2011, preprint ([arXiv:1110.3193](https://arxiv.org/abs/1110.3193))
 Manera M. et al., 2013, *MNRAS*, 428, 1036
 Mehta K. T., Seo H.-J., Eckel J., Eisenstein D. J., Metchnik M., Pinto P., Xu X., 2011, *ApJ*, 734, 94
 Merson A. I. et al., 2013, *MNRAS*, 429, 556
 Mesinger A., Furlanetto S., 2007, *ApJ*, 669, 663
 Mo H. J., Jing Y. P., White S. D. M., 1997, *MNRAS*, 284, 189
 Monaco P., Theuns T., Taffoni G., 2002, *MNRAS*, 331, 587
 Monaco P., Sefusatti E., Borgani S., Crocce M., Fosalba P., Sheth R. K., Theuns T., 2013, *MNRAS*, 433, 2389
 Padmanabhan N., White M., 2009, *Phys. Rev. D*, 80, 063508
 Parkinson H., Cole S., Helly J., 2008, *MNRAS*, 383, 557
 Percival W. J. et al., 2001, *MNRAS*, 327, 1297
 Reid B. A. et al., 2012, *MNRAS*, 426, 2719
 Sánchez A. G., Baugh C. M., Angulo R. E., 2008, *MNRAS*, 390, 1470
 Sánchez A. G., Crocce M., Cabré A., Baugh C. M., Gaztañaga E., 2009, *MNRAS*, 400, 1643
 Sánchez A. G. et al., 2012, *MNRAS*, 425, 415
 Sánchez A. G. et al., 2014, *MNRAS*, 440, 2692
 Schneider M. D., Cole S., Frenk C. S., Szapudi I., 2011, *ApJ*, 737, 11
 Scoccimarro R., Sheth R. K., 2002, *MNRAS*, 329, 629
 Shandarin S., Habib S., Heitmann K., 2012, *Phys. Rev. D*, 85, 083005
 Sheth R. K., Lemson G., 1999, *MNRAS*, 305, 946
 Sheth R. K., Mo H. J., Tormen G., 2001, *MNRAS*, 323, 1
 Simha V., Weinberg D. H., Davé R., Fardal M., Katz N., Oppenheimer B. D., 2012, *MNRAS*, 423, 3458
 Smith R. E., Marian L., 2011, *MNRAS*, 418, 729
 Smith R. E., Scoccimarro R., Sheth R. K., 2008, *Phys. Rev. D*, 77, 043525
 Springel V., Frenk C. S., White S. D. M., 2006, *Nature*, 440, 1137
 Takahashi R. et al., 2009, *ApJ*, 700, 479
 Teyssier R. et al., 2009, *A&A*, 497, 335
 The Dark Energy Survey Collaboration, 2005, preprint ([arXiv:astro-ph/0510346](https://arxiv.org/abs/astro-ph/0510346))
 Wake D. A. et al., 2008, *MNRAS*, 387, 1045
 White S. D. M., Davis M., Efstathiou G., Frenk C. S., 1987, *Nature*, 330, 451
 White M., Tinker J. L., McBride C. K., 2014, *MNRAS*, 437, 2594
 Zehavi I. et al., 2011, *ApJ*, 736, 59

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