Genetically modified haloes: towards controlled experiments in ΛCDM galaxy formation

Nina Roth,[★] Andrew Pontzen and Hiranya V. Peiris

Department of Physics and Astronomy, University College London, Gower Street, London, WC1E 6BT, UK

ABSTRACT

We propose a method to generate 'genetically modified' (GM) initial conditions for high-resolution simulations of galaxy formation in a cosmological context. Building on the Hoffman–Ribak algorithm, we start from a reference simulation with fully random initial conditions, then make controlled changes to specific properties of a single halo (such as its mass and merger history). The algorithm demonstrably makes minimal changes to other properties of the halo and its environment, allowing us to isolate the impact of a given modification. As a significant improvement over previous work, we are able to calculate the abundance of the resulting objects relative to the reference simulation. Our approach can be applied to a wide range of cosmic structures and epochs; here we study two problems as a proof of concept. First, we investigate the change in density profile and concentration as the collapse times of three individual haloes are varied at fixed final mass, showing good agreement with previous statistical studies using large simulation suites. Secondly, we modify the $z = 0$ mass of haloes to show that our theoretical abundance calculations correctly recover the halo mass function. The results demonstrate that the technique is robust, opening the way to controlled experiments in galaxy formation using hydrodynamic zoom simulations.

Key words: galaxies: evolution – galaxies: formation – dark matter.

1 INTRODUCTION

Understanding galaxy formation requires us to take account of the variety of halo assembly histories that build today's population. Many pressing questions – such as the origin of varying morphologies (e.g. van Dokkum et al. 2013; Papovich et al. 2015) and bulge sizes (Kormendy 2015) – will be answered by understanding the interplay between complex, non-linear physics and the various histories for mass accretion. The fundamental difficulty is that these histories are in turn determined by the random initial conditions (ICs) seeded in the early universe.

This paper is the first in a series to directly tackle that problem using a novel approach. The most typical solution is to simulate large numbers of galaxies in a representative volume (e.g. Genel et al. 2014; Codis et al. 2015; Schaye et al. 2015). However, this is computationally expensive and limits the resolution that can be achieved for any single object. Conversely, zoom-in simulations achieve the maximum level of physical detail for a given computational time. They have been used to establish that qualitatively different processes come into play at sub-kpc resolutions, where processes within the interstellar medium begin to be resolved (e.g. Governato et al. 2007; Brook et al. 2011; Guedes et al. 2011;

Hopkins et al. 2013; Pontzen & Governato 2014). But such approaches only sample over a small and potentially biased range of merger histories. A third tactic is to use isolated, idealized set-ups to test particular hypotheses (e.g. Naab, Burkert & Hernquist 1999; Robertson et al. 2006; Hopkins et al. 2009), but these by definition lack a full cosmological environment. Thus, it is difficult to quantitatively connect the results of isolated and zoom simulations to the observed galaxy population.

Our aim is to combine the best aspects of these three types of numerical study. We proceed by systematically changing aspects of individual galaxies' histories (such as mass and merger history) within a cosmological simulation, while keeping track of the statistical likelihood of the changes to understand the relative abundance of objects of different types. This can be achieved by using the Hoffman–Ribak algorithm (Hoffman & Ribak 1991, hereafter HR91; see also Bardeen et al. 1986; Bertschinger 1987; van de Weygaert & Bertschinger 1996 for further theoretical background). A more common use for HR91 is to obtain simulations resembling the local universe by turning a given observational data set (e.g. the local distribution of galaxies) into a prescription for the ICs of a numerical simulation (Bistolas & Hoffman 1998; Kravtsov, Klypin & Hoffman 2002; Mathis et al. 2002; Klypin et al. 2003; Heß, Kitaura & Gottlöber 2013; Jasche & Wandelt 2013; Sorce et al. 2014). There is a significant amount of literature that uses this technique to study the formation history of the Local Group (Zavala et al. 2009;

★ E-mail: n.roth@ucl.ac.uk

Klimentowski et al. 2010; Libeskind et al. 2010; Forero-Romero et al. 2011; Iliev et al. 2011; Dayal, Libeskind & Dunlop 2013; Doumler et al. 2013; Kitaura 2013; Brook et al. 2014; Nuza et al. 2014).

Instead, we propose to use constrained ICs as an experimentation toolkit for the formation of a particular halo embedded in a cosmological volume. This approach has precedent: for example, Frenk et al. (1999) used the HR91 method to create galaxy cluster initial conditions for a comparative study of numerical simulation codes. More recently, Romano-Díaz et al. (2006, 2007) and Hoffman et al. (2007) simulated a single dark matter object of $\sim 10^{12}\,\mathrm{M}_\odot$ with different substructures to understand the impact of quiescent and violent accretion phases on the inner properties of the halo, and the origin of the universal halo density profile. By including baryonic physics, Romano-Díaz et al. (2011a,b, 2014) studied galactic properties in extremely overdense regions which may host the early precursors of QSOs. In a similar vein, Dubois et al. (2012) use the numerical implementation from Prunet et al. (2008) to investigate the accretion of material in the cores of very massive haloes to shed light on the formation of black holes at high redshifts.

In all of the above cases, a simulated object was created by constraining the properties of a region defined by an analytical profile (typically a Gaussian peak). Constrained properties included the height of the density peak at the origin and its first- and second-order derivatives (see e.g. Prunet et al. 2008 and the appendix of Romano-Díaz et al. 2011a). This creates objects that are well defined in a theoretical sense (e.g. one can predict their collapse time reasonably well), but that represent configurations which may or may not be common in fully random ICs.

What sets our work apart from these previous efforts is that we always start with a 'reference' halo from a simulation based on fully random ICs. We are able to impose constraints on volumes of completely arbitrary shape, using the particles that make up a single dark matter halo embedded in a cosmological volume. Once the constraints are applied, we re-run the simulation and compare the results to the original reference run.

This has two immediate benefits. First, we can fine-tune selected properties of the halo while demonstrably ensuring that the constrained object is as similar as possible to the reference run – a controlled 'genetic modification' (GM) of the halo. Secondly, we can calculate the change in the likelihood of the field after the modification; in other words, we can assess the relative abundance of the genetically modified systems compared to the original. This will allow us to test whether connections between merger history and morphology quantitatively account for observed population statistics.

The current work provides a first illustration of both these aspects of the technique. Specifically, we study the properties of several haloes as their total mass and merger history are systematically changed. We investigate the concentration at $z = 0$ for different mass accretion histories and find overall excellent agreement of our constrained haloes with relations derived from statistical averages over large simulations. There are many studies that connect the concentration parameter to other halo properties like the mass, collapse time or mass accretion history, halo shape and angular momentum (e.g. Bullock et al. 2001; van den Bosch 2002; Wechsler et al. 2002; Zhao et al. 2003; Reed et al. 2005; Bett et al. 2007; Macciò et al. 2007; Neto et al. 2007; Duffy et al. 2008; Macciò, Dutton & van den Bosch 2008; Zhao et al. 2009; Ragone-Figueroa et al. 2010; Prada et al. 2012; Ludlow et al. 2013, 2014; Correa et al. 2015a,b,c; Klypin et al. 2014). Often, these studies operate by considering a statistical sample from a large volume simulation to find

correlations and provide fitting functions. Even though the statistical power in recent simulations is excellent, there is still considerable scatter around the median relations. Since the density profile of dark matter haloes is an important ingredient in theoretical models, it is important to understand these correlations and the scatter. Given the large number of parameters that could influence the evolution of a halo, principal component analysis has been used to investigate correlations between them (Jeeson-Daniel et al. 2011; Skibba & Macciò 2011; Wong & Taylor 2012). Our approach of designing 'experiments' in galaxy formation provides a complementary approach to computationally expensive statistical studies.

This paper is organized as follows: in Section 2 we give a brief outline of the HR91 technique and our specific implementation. Section 3 contains details of the numerical simulations that are used to obtain the results in the rest of the paper. In Section 4 we provide a brief illustration of some of the constraints we have applied to the reference initial density field, focusing on influencing a single halo traced by its particles. Next, we study the results of designing different merger histories for a set of haloes in Section 5, focusing on their collapse–concentration relation. In Section 6, we discuss a method for assessing the relative abundance of the modified haloes by defining a χ^2 measure, and show that our results are consistent with the cosmological halo mass function. We summarize in Section 7. Finally, Appendices A and B contain the mathematical details of our reformulation of the HR91 technique including a translation between our notation and theirs.

2 OUTLINE OF THE METHOD

We now present a brief outline of the mathematical technique by which ICs can be generated that satisfy certain constraints, while remaining consistent with a Λ cold dark matter (ΛCDM) power spectrum. This technique is described in a slightly different formulation by HR91. The full derivation can be found in Appendix A.

By assumption, the density field in the early universe is linearly perturbed around a background density ρ_0, so that

$$\rho(\boldsymbol{x}) = \rho_0\left(1 + \delta(\boldsymbol{x})\right), \tag{1}$$

where $\delta(\boldsymbol{x})$ is a Gaussian random field with statistical properties specified by the ΛCDM transfer function and inflationary tilt.

Generating ICs involves sampling the Gaussian random field at a list of discrete points $\boldsymbol{x}_\nu$, where the integer value ν decides which point we are discussing. In particular, when running a uniform-resolution cosmological simulation with N particles on a box side, ν runs from 1 to N^3. The sampled field then consists of an N^3-length vector $\boldsymbol{\delta}$, where an element is given by $\delta_\nu \equiv \delta(\boldsymbol{x}_\nu)$. The values of $\boldsymbol{\delta}$ are drawn from a multivariate Gaussian probability distribution with mean $\langle\boldsymbol{\delta}\rangle = \boldsymbol{\mu}_0$ and covariance matrix $\mathbf{C}_0 = \langle(\boldsymbol{\delta} - \boldsymbol{\mu}_0)^\dagger(\boldsymbol{\delta} - \boldsymbol{\mu}_0)\rangle$. Cosmological ICs have zero mean, $\boldsymbol{\mu}_0 = 0$, but the HR91 technique is not limited to this case.

A constrained field is defined by requiring $\boldsymbol{\alpha}^\dagger\boldsymbol{\delta} = d$ for some constraint vector $\boldsymbol{\alpha}$, which also contains N^3 elements. In general, d is real, though the formalism also extends to the case where it is complex-valued. A simple example would be to fix the density contrast to 0 at position $\boldsymbol{x}_1$. This requires $\alpha_\nu = 1$ for $\nu = 1$ (before normalization, see below) and 0 otherwise, and $d = 0$. Throughout, we will use Greek indices in this way to denote values of either $\boldsymbol{\delta}$ or $\boldsymbol{\alpha}$ at a specific grid position $\boldsymbol{x}_\nu$.

To actually create a field satisfying any given constraint, one could sample repeatedly from the underlying population until obtaining a realization that satisfies (or is close to satisfying) the requirement. However, such an accept–reject algorithm would be

computationally expensive to implement in practice; instead, the HR91 technique makes a mathematical rearrangement that requires only one set of random numbers to be generated. As detailed in Appendix A, this rearrangement also shows that a constrained Gaussian random field remains Gaussian. This allows us to apply a large number of constraints independently, with the final result obeying $\boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta} = d_i$ for each i, where the Roman index i denotes the n different constraints. The properties of the constrained field are then determined by a new mean and covariance

$$\boldsymbol{\mu}_n = \boldsymbol{\mu}_0 + \sum_{i=1}^n \mathbf{C}_0 \boldsymbol{\alpha}_i \left(d_i - \boldsymbol{\alpha}_i^\dagger \boldsymbol{\mu}_0 \right) \tag{2}$$

$$\mathbf{C}_n = \mathbf{C}_0 - \sum_{i=1}^n \mathbf{C}_0 \boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^\dagger \mathbf{C}_0, \tag{3}$$

provided that the $\{\boldsymbol{\alpha}_i\}$ have been orthonormalized[1] in the sense that $\boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_j = \delta_{ij}$.

A realization of the constrained field could therefore be obtained by calculating $\boldsymbol{\mu}_n$ and $\mathbf{C}_n$ and drawing random numbers accordingly. However, in practice, dealing directly with $\mathbf{C}_n$ from equation (3) becomes prohibitively expensive for large N. The problem is that, whereas $\mathbf{C}_0$ is the ΛCDM power spectrum and therefore diagonal in Fourier space, $\mathbf{C}_n$ is generally not sparse in either pixel or Fourier space. Instead, one can make the ansatz that a realization obeying n constraints, $\boldsymbol{\delta}_n$, can be obtained starting from a realization of the unconstrained field, $\boldsymbol{\delta}_0$, via

$$\boldsymbol{\delta}_n = \mathbf{P}_n \left(\boldsymbol{\delta}_0 - \boldsymbol{\mu}_0 \right) + \boldsymbol{\mu}_n, \tag{4}$$

where $\mathbf{P}_n$ is a matrix that depends on $\mathbf{C}_0$ and $\{\boldsymbol{\alpha}_i\}$.

By requiring that $\boldsymbol{\delta}_n$ obeys the correct statistics and, additionally, requiring that the changes made to the field are minimal, one can uniquely derive the HR91 solution for $\mathbf{P}_n$. The details are given in Appendix A, with the result that

$$\boldsymbol{\delta}_n = \boldsymbol{\delta}_0 + \sum_{i=1}^n \mathbf{C}_0 \boldsymbol{\alpha}_i \left(d_i - d_{i0} \right), \tag{5}$$

where we have defined $d_{i0} = \boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta}_0$ to represent the value of constrained quantities in the unconstrained realization, and again require that the $\{\boldsymbol{\alpha}_i\}$ are orthonormalized. This reduces the actual calculation to a series of vector multiplications and summations in Fourier space (since $\mathbf{C}_0$ is diagonal there). The memory requirements are manageable since we need only store vectors of length N^3, instead of the $N^3 \times N^3$ matrix $\mathbf{C}_n$.

Consequently, a continuum of constrained realizations can be generated from a single realization of the original ensemble. We select a single dark matter halo from a reference run at $z=0$, and then return to the ICs and place constraints on the particles that make up this object. In this way, our constrained regions are defined directly via the halo particles, without requiring any assumptions about the properties of density peaks in the ICs, or any type of direct smoothing of the density field (only indirectly through the halo finding at $z=0$).

3 SIMULATION SET-UP

All our simulations were run with P-GADGET-3 (Springel 2005; Springel et al. 2008). The ICs have been set up at redshift $z=99$ and

[1] Note that this orthonormalization can always be arranged for any set of non-conflicting original constraints.

Table 1. Overview of the simulations used in this paper. For more details on the individual runs, see the text in the sections mentioned in the last column. MH and MH* stand for the two different ways of constraining the merger history of the halo. The simulations marked in bold are not actually used because they are not in equilibrium at $z=0$ (see Section 5 for a discussion).

Name	Halo	Constraint	Section
Reference	N/A	None	3
H24-MH	Halo 24	$d_{10}/d_{\mathrm{ref}} = \{\mathbf{0.5}, 1.5\}$	4, 5
H24-MH*	Halo 24	$d(z=1)/d_{\mathrm{ref}} = \{0.9, 1.1\}$	4, 5
H37-MH	Halo 37	$d_{10}/d_{\mathrm{ref}} = \{\mathbf{0.5}, 1.5\}$	4, 5
H37-MH*	Halo 37	$d(z=1)/d_{\mathrm{ref}} = \{0.9, 1.1\}$	4, 5
H40-MH	Halo 40	$d_{10}/d_{\mathrm{ref}} = \{0.5, 1.5, 2\}$	4, 5
H40-MH*	Halo 40	$d(z=1)/d_{\mathrm{ref}} = \{\mathbf{0.9}, 1.1, \mathbf{1.2}\}$	4, 5
H24-mass	Halo 24	$d/d_{\mathrm{ref}} = \{0.5, 0.8, 1.2, 1.5\}$	4, 6
H37-mass	Halo 37	$d/d_{\mathrm{ref}} = \{0.5, 0.8, 1.2, 1.5\}$	4, 6
H40-mass	Halo 40	$d/d_{\mathrm{ref}} = \{0.5, 0.8, 1.2, 1.5\}$	4, 6

evolved to $z=0$, saving 100 snapshots from $z=9$ equally spaced in scale factor. The cosmological model is $\Omega_{\mathrm{m}} = 0.279$, $\Omega_{\mathrm{b}} = 0.045$, $\Omega_\Lambda = 0.721$, $\sigma_8 = 0.817$, $h = 0.701$, $n_{\mathrm{s}} = 0.96$, corresponding to a *Wilkinson Microwave Anisotropy Probe* 5 cosmology (Dunkley et al. 2009). While these cosmological parameters have been revised in more recent data sets, they allow an easier comparison of our results with the literature. All simulations have a (comoving) box size of $L = 50\,h^{-1}\,\mathrm{Mpc} \sim 71.3\,\mathrm{Mpc}$, and $N_{\mathrm{part}} = 256^3$ dark matter particles, resulting in a particle mass of $8.24 \times 10^8\,\mathrm{M}_\odot$. The Plummer equivalent force softening length (which limits the smallest accessible scales) is $\epsilon = 25.6\,\mathrm{kpc}$ in comoving units.

We use the SUBFIND code (Springel et al. 2001), which finds haloes with the friends-of-friends (FoF) method. SUBFIND also identifies subhaloes inside the top level FoF groups, but we always use the whole group for particle tracking. Each FoF group is assigned a unique number, sorted in descending order by mass. SUBFIND also provides a list of halo particle IDs, which allows us to track the haloes between snapshots and across different simulations (by selecting those objects which have the most particles in common at $z=0$). We choose the standard FoF linking length of 0.2 times the mean interparticle distance.

Our analysis makes use of the PYTHON module PYNBODY (Pontzen et al. 2013). We select haloes with mass $M_{200} \sim 10^{13}\,\mathrm{M}_\odot$ at $z=0$, which are well resolved but not the most massive (and therefore rare) objects in the box. Throughout this paper, M_{200} refers to the halo mass contained in r_{200}, the radius within which the mass density is 200 times the critical density of the universe at that time. We construct halo merger trees by tracing halo particles from $z=0$ backwards in time through each simulation snapshot. This allows us to determine the mass accretion history and other internal properties of each object as a function of time. We take advantage of the fact that the density field is first set up by assigning one particle to each grid node (the displacements are applied later). This means that any particle at $z=0$ can be traced back to a grid position $\boldsymbol{x}_\nu$ in the ICs, and no additional interpolation is necessary to calculate $\delta(\boldsymbol{x}_\nu)$.

Table 1 gives an overview of the different runs that are used in this paper; the last column provides the section where the constraints are described and the results are discussed. In order to show that the technique is robust, we have selected three haloes of similar mass in the reference run, and constrained their properties in different ways. Therefore, each simulation name contains a halo number and constraint type.
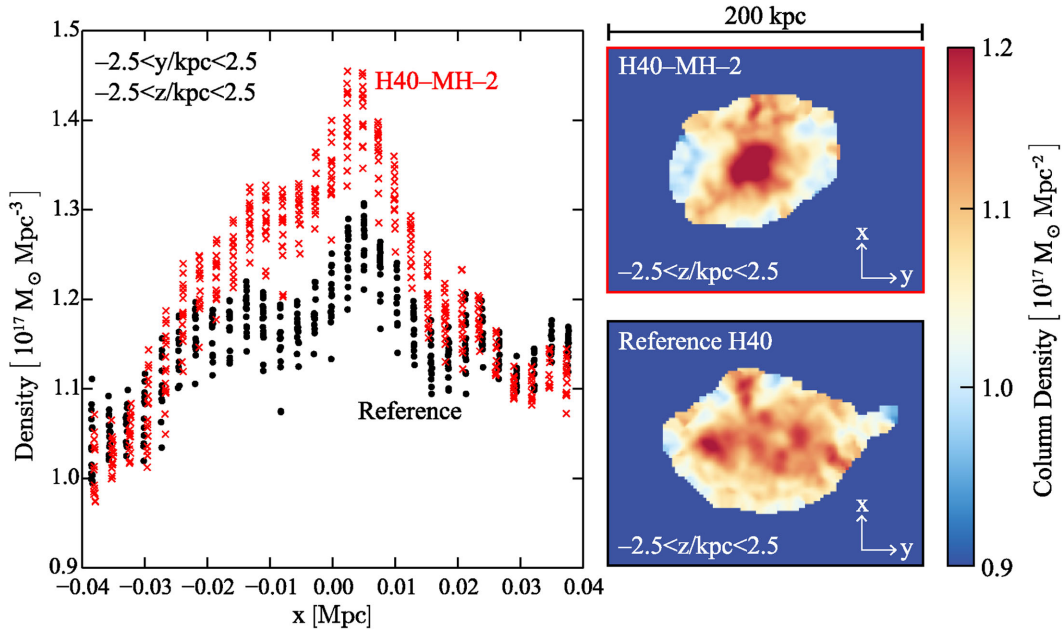
Figure 1. Left-hand panel: the density of the reference ICs (black circles) and modified H40-MH-2 ICs (red crosses) for the early collapse constraint, where the density of the 10 per cent innermost particles is increased by a factor of 2. The slice is 5 kpc wide in the y- and z-coordinates, to give an impression of the 3D structure. Each symbol corresponds to a single particle/initial grid point. The constrained density field maintains the complicated (sub-)structure that was present in the reference run. Right-hand panels: the same two ICs as a 2D projection in the x–y plane. Only those particles that form each halo at $z = 0$ are shown here; it is these particles that are used for generating the constraint in our algorithm. The higher central density is clearly visible in the constrained case. The results of these simulations will be discussed in detail in Section 5.

4 ILLUSTRATION OF CONSTRAINTS

We now present a simple illustration of the technique with which we generate a *density constraint*. We will discuss the results obtained from running simulations with these constraints in the following sections.

Our approach constrains the actual Lagrangian region that collapses into a halo at $z = 0$; by contrast, in previous work the constraints were typically chosen to follow some analytical form, in order to connect the constraints to theoretical models such as Press–Schechter theory (e.g. van de Weygaert & Bertschinger 1996; Romano-Díaz et al. 2006). Since we know which particles are going to collapse in the reference run, we do not need to assume a specific form for the peak or a smoothing scale. The resulting constrained halo will be very similar to the reference object, unless the constraint radically changes the collapsing region, e.g. by introducing a large overdensity in a region which only forms an intermediate mass halo in the reference run.

Designing the constraints for a given modification to the final halo requires a physical understanding of the evolution. Ultimately, a proposed constraint must be tested by trialling the changes and testing that they have the desired effect and that they are statistically consistent with the modified halo existing in the unconstrained universe. We will demonstrate both of these properties over the remainder of the paper.

Changing the mass can be achieved by changing the density contrast of the halo particles in the ICs. By creating a larger or smaller overdensity, we influence the final mass by increasing or decreasing the overall size of the region which has the average threshold density to collapse by a specified redshift (Press & Schechter 1974). We

term this a *density constraint*: in the ICs, we calculate the average mass overdensity of all N_{part} particles in the reference halo

$$\frac{1}{N_{\text{part}}} \sum_{\nu=1}^{N_{\text{part}}} \delta(\boldsymbol{x}_\nu) \equiv d, \qquad (6)$$

where we again use the fact that each particle corresponds to a grid position $\boldsymbol{x}_\nu$. Before orthonormalization, the value of the constraint vector $\boldsymbol{\alpha}$ is then $1/N_{\text{part}}$ for each particle which belongs to the halo and 0 otherwise. The density can now be increased or decreased by enforcing the value of d. Results of simulations with different choices for d that produce haloes with higher or lower mass at $z = 0$ will be used in Section 6.

More specifically, according to the Press–Schechter argument, the collapse time of a halo is related to its peak height $\nu = \delta/\sigma(R)$, where $\sigma(R)$ is the variance of the density field smoothed on a scale R. Therefore, by fine-tuning the overdensity on different scales within the ICs, we can modify the accretion history. In particular, by increasing (or decreasing) the density contrast in an *inner region* of the halo, while requiring that the overall density contrast of the halo particles stays the same, we are able to generate a halo with very similar mass at $z = 0$, but a faster (or slower) accretion history.

Fig. 1 illustrates an example of such a constraint acting on the ICs. In the left-hand panel, we show a slice through the dark matter density field in the ICs, centred on the halo's centre of mass. Each black circle corresponds to a density value in the reference run, and the red crosses show the same position in the constrained run. The slice is 5 kpc wide in the y- and z-coordinates, to give an
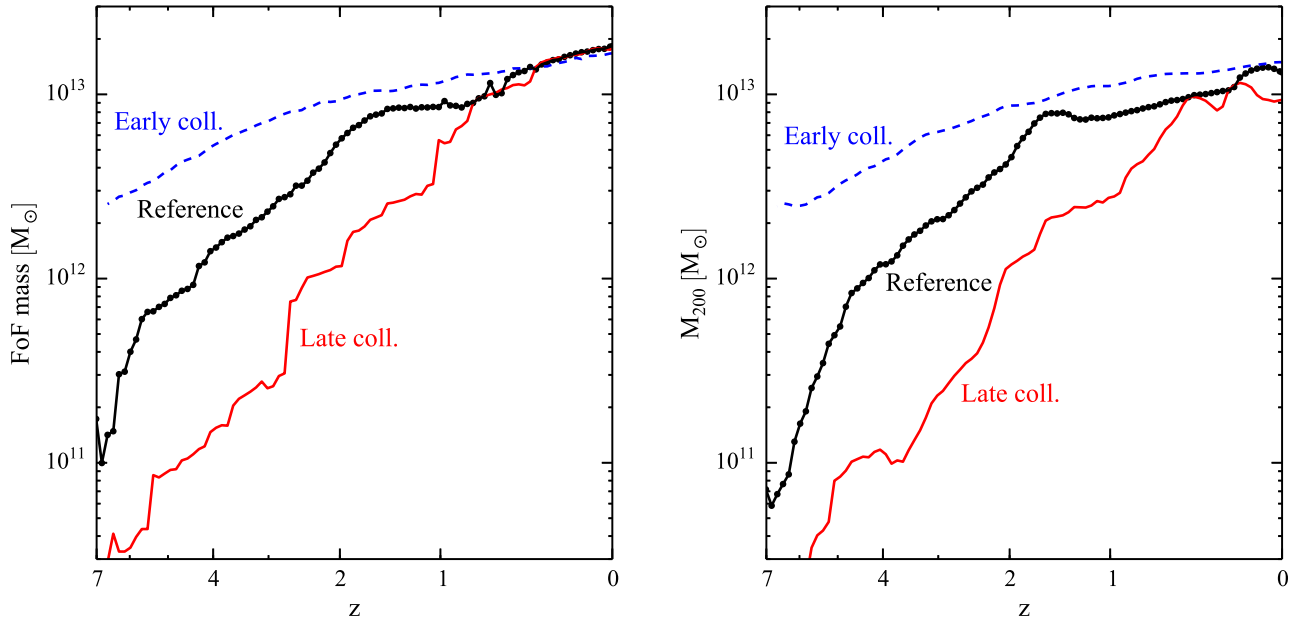
Figure 2. Left-hand panel: mass accretion history for early (red solid) and late collapse (blue dashed) runs, expressed by the FoF mass (all particles assigned by the halo finder). The black solid line with points shows the same halo in the reference run; each point is one snapshot, illustrating the time resolution of our simulations. Right-hand panel: same but for the virial mass M_{200}, which does not converge to a common value at late times because M_{200} probes the inner regions of the halo (see Fig. 3), which are affected by the collapse time.

impression of the 3D structure. Here, we show the constraint that will be used in the 'early' run (discussed in detail in Section 5), where we have increased the density of the innermost 10 per cent of halo particles to be a factor of 2 higher, while keeping the overall density the same. An alternative approach is to actually identify substructures at an intermediate redshift and apply the 'inner' constraints to those specific particles. We have tried this second approach for the present work, selecting the particles which have already collapsed around $z \sim 1$ by constructing a merger tree from the SUBFIND output in the reference run. Table 1 contains an overview of all simulations used in this paper; we denote the first method of constraining the *inner region* by d_{10} and the second by $d(z = 1)$ there. The two modification methods give near-identical results (Section 5), making the outcomes reassuringly insensitive to the intuition guiding the modifications. We have found that it is also possible to add further constraints to modify the build-up in different subhaloes and so fine-tune the accretion history to any required degree.

As explained in Section 2, the modified field is constructed to follow the peaks and troughs of the underlying density field, thereby maintaining the same substructure as much as possible. In the right-hand panel, we show the 2D projection (*x–y* plane) of the density of halo particles in the ICs, again for both the reference run and constrained run. In both cases, the density is calculated for all particles that are part of the halo at $z = 0$, which have been traced back to $z = 99$. The effect of increasing the density in the innermost region can be clearly seen in the constrained run (red border). In addition, the second constraint, which keeps the overall mass the same, leads to a compensation effect in the ICs, removing some particles in the outer regions which fall into the reference halo but not the constrained one. We will discuss the results of simulations with these modified ICs in the next section.

5 MERGER HISTORY AND CONCENTRATION–COLLAPSE RELATION

So far we have looked at how applying various constraints modifies the initial linear overdensity field. We will now consider the changes that result when the new ICs are used in a numerical simulation, starting with our modified merger history.

The mass accretion histories for simulations with the H40-MH-2 'early' (circles) and H40-MH-0.5 'late' collapse constraint (crosses) are shown in Fig. 2. Here, we chose the innermost 10 per cent of particles and changed their overdensity by a factor of 2 (0.5) for the early (late) collapse cases. In the left-hand panel, we show the time evolution of the total mass (including all substructures) for the two constrained runs and the reference halo. It is clear that the accretion rates differ quite significantly at early times, but are compensated at late times, leaving the overall mass of the objects the same. In the right-hand panel, we show the time evolution of M_{200}. This quantity only measures the mass up to r_{200} instead of the total mass of linked FoF particles (which can extend out to several r_{200}). As in the former case, the accretion rate follows the expected behaviour in the early and late collapse cases. However, at late times, M_{200} differs between the constrained runs and the reference runs. This is due to a change in the halo density profile related to the collapse time, as we will now explore.

The halo radial density profiles for these three simulations at $z = 0$ are illustrated in Fig. 3, with inset panels showing projected density maps. As the collapse is delayed, the slope in the inner regions becomes less steep. The location of the virial radius r_{200} is indicated by an arrow in each case; as expected given our discussion above, this is displaced inwards by the relative shallowness of the late collapse H40-MH-0.5 case.
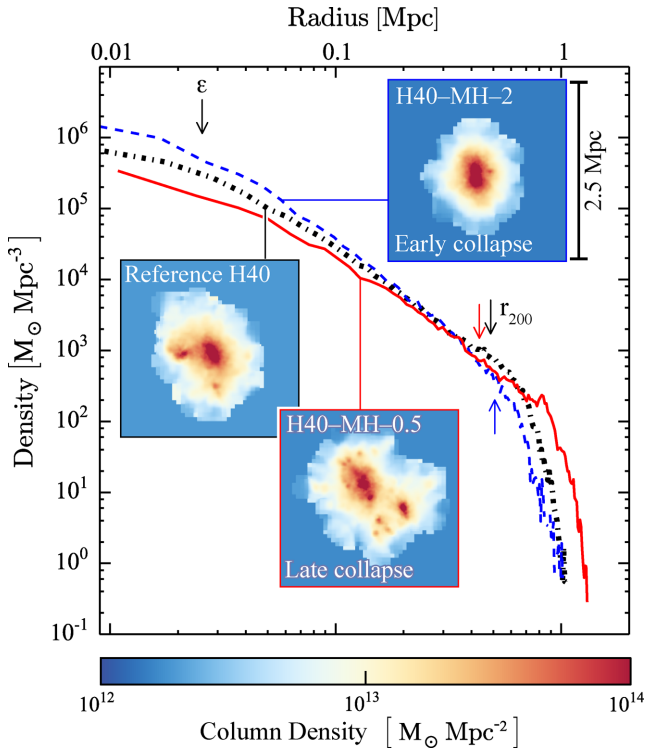
Figure 3. Density profile of the reference halo (black dot–dashed) and the 'early' (blue dashed) and 'late' (red solid) constrained runs at $z = 0$. The leftmost arrow indicates the softening length of the simulation, and the other arrows indicate the virial radius of each halo. Inset panels: density projection (x–y plane) of the resulting haloes at $z = 0$. All panels show a region 2.5 Mpc across, include only the FoF group particles, and use the same colour scale for the column density.

The difference between the density profiles can be encapsulated in the concentration parameter

$$c = \frac{r_{200}}{r_s}, \tag{7}$$

where r_{200} is the virial radius of the halo (defined in Section 3) and r_s is the scale radius in the NFW density profile (Navarro, Frenk & White 1997)

$$\rho(r) = \frac{4\rho_s}{(r/r_s)(1 + r/r_s)^2}. \tag{8}$$

We fit an NFW profile to each of our haloes at $z = 0$, after determining its centre using a shrinking-sphere method and estimating the density in 100 radial bins of equal size. In order not to contaminate the fits with numerical artefacts introduced by the finite particle resolution, we exclude from the fit the innermost regions (two times the softening length, ϵ), which are affected by the force softening (e.g. Power et al. 2003), and regions with $r > 0.6\, r_{200}$, which may not be relaxed. We tested that the choice of the minimum and maximum radius has negligible impact on the estimate of r_s. The measured values of the concentration are 4.2, 6.1 and 11.1 for late, reference and early simulations, respectively.

The GM method allows us to study the relationship between the collapse time of a halo (defined below) and its concentration as measured at $z = 0$. Using a large statistical sample, Wechsler et al. (2002) found that $c \propto a_{\mathrm{coll}}^{-1}$, the scale factor at collapse time. We

follow their procedure to obtain the collapse scale factor by fitting the mass accretion history of each halo with

$$M(z) = M_0 \times \exp[-\alpha\, z], \tag{9}$$

with $M_0 \equiv M_{200}(z = 0)$ and $\alpha = 2a_{\mathrm{coll}}$. Extensions to this simple function have been proposed by e.g. Tasitsiomi et al. (2004), McBride, Fakhouri & Ma (2009) and Correa et al. (2015a), but for our purposes these are not necessary: the refined formulae are designed to accurately represent the median mass accretion histories for many haloes, and the corrections are smaller than the scatter between individual haloes.

There are some differences between the conventions of Wechsler et al. (2002) and the present work which we need to understand before proceeding. In the left-hand panel of Fig. 4, the grey band and black dashed line show the average concentration and scatter from measuring the properties of $\sim$120 haloes of mass $M \sim 10^{13}\,\mathrm{M_\odot}$ in our unmodified box. Wechsler et al. (2002) used a slightly different definition of halo mass from the one used in our study. Instead of defining M_{200} with respect to the critical density of the universe, they define M_{200}^{mean} relative to the mean density.

To show how this affects the measured concentration, the green points show a sample of concentration parameters estimated using M_{200}^{mean}, using the same haloes that were used to generate the grey shaded region. There is an overall upward offset of these points relative to the grey band because r_{200}^{mean} is correspondingly larger. The black solid line and the green band show the median relation and scatter predicted by Wechsler et al. (2002, taken from their fig. 7), corrected by a factor of 0.8 following Duffy et al. (2008) to account for their different σ_8 (1, instead of 0.817 in our simulations).

In summary, once the differences in conventions and cosmological parameters are taken into account, we can reproduce the median results of Wechsler et al. (2002) in our unmodified boxes. The scatter from our simulation is also compatible with the much larger Wechsler et al. (2002) sample; outliers are likely due to the fact that we do not pre-select relaxed haloes as they do. For consistency with the rest of the paper, we will use all quantities derived with respect to *critical density* in the following analysis.

We are now ready to see how this relationship emerges when using GM haloes instead of a statistical sample. For this study, we have selected three different haloes in the reference run, which are all of similar mass ($M \sim 10^{13}\,\mathrm{M_\odot}$). The right-hand panel of Fig. 4 shows the result for 13 simulations (four each for haloes 24 and 37, and five for halo 40) with different collapse times. We call each set a 'halo family', including the reference run. The right-hand panel of Fig. 4 shows the results for the three families illustrated by red diamonds, green circles and blue squares for family 24, 37 and 40, respectively.

The slopes of the three families appear to be consistent with, but scattered around, the population average (grey band). We find that each halo family is well described by a linear relation

$$c = \frac{\mathrm{const}_1}{a_{\mathrm{coll}}} + \mathrm{const}_2, \tag{10}$$

which contains the offset as an additional parameter compared to the fit used in Wechsler et al. (2002). These fits are shown in the right-hand panel of Fig. 4 as coloured solid lines; we also fit all 13 points together (black dashed line). In addition, the grey band shows the scatter expected for haloes in our selected mass bin, as in the left-hand panel. The fit to all 13 simulations is very similar to the median relation from our unconstrained box which is consistent with the larger sample from Wechsler et al. (2002).
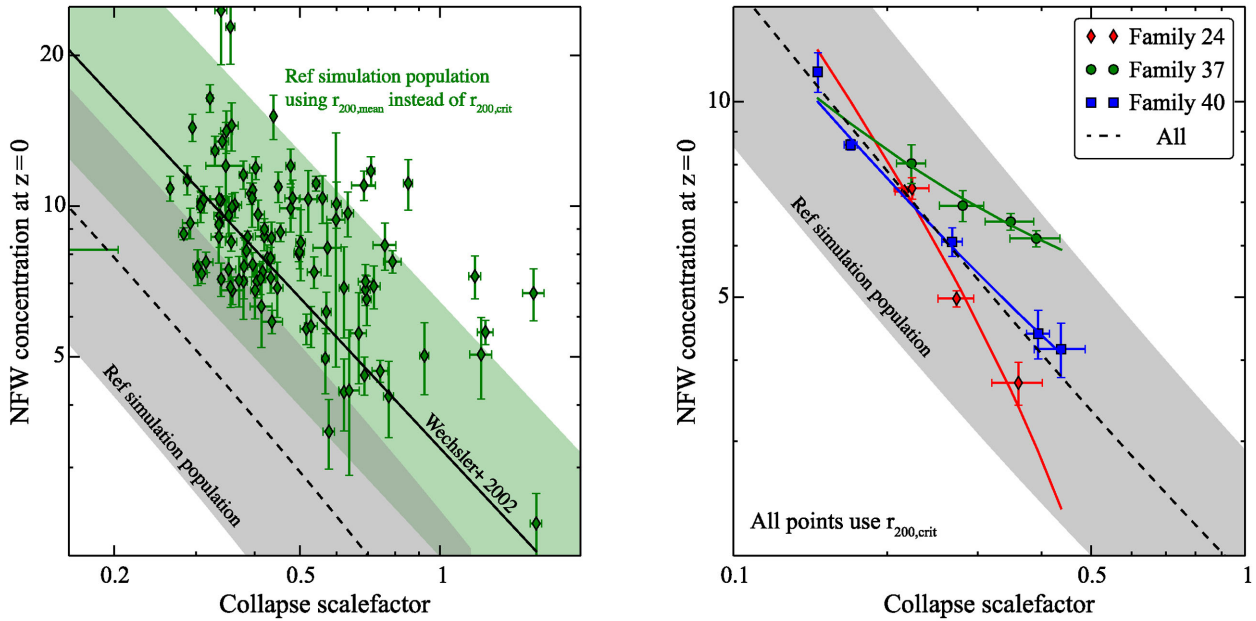
Figure 4. Halo concentration parameter as a function of the collapse time. Left-hand panel: our reference simulation gives a volume to probe the relationship using the traditional statistical technique. Taking 120 haloes from our simulation, this results in a scatter of points in the region of the grey band. To compare with the existing literature, we need to redefine r_{200} (and hence c) relative to the mean (rather than critical) density, after which these haloes are represented by the green points with error bars. The black solid line and green band show the average relation and scatter as predicted by Wechsler et al. (2002), multiplied by a factor of 0.8 to account for their different choice of σ_8 (Duffy et al. 2008). Right-hand panel: our three constrained haloes (24, 37 and 40), showing fits to each constrained family individually (colours) and all of them together (black dashed). The grey band is the same as in the left-hand panel. Together the panels establish that (left) haloes in our reference volume recover the known relationship between concentration and collapse scale factor, and (right) relationships consistent with this relation are also recovered individually by each GM family. The scatter of slopes between different families is expected (see the text).

The shift of each family member along its line is dictated by the direction and amplitude of the density constraint in the ICs. Higher values of the density in the inner region shift a point towards the top left, and lower values to the bottom right with respect to the reference run. This gives us considerable insight into the kind of results that can be expected from GM compared to large population studies. The scatter of individual simulations within a GM family is very small – in other words, the concentration is highly predictable from a single variable. This is because, as we have previously emphasized, the history of each halo within a single family is as similar as possible to all the others. The normal scatter in the concentration–collapse relation is then seen to be due to factors that are not being constrained within a single family (such as more detailed aspects of the merger history or other variables such as halo spin). The GM technique allows for a detailed exploration of results from specific, precise changes.

For halo 24 (red diamonds), two of the results are nearly identical: the point with the highest concentration value is actually two points nearly on top of each other. These points have been obtained using the two methods for setting mass accretion history constraints discussed previously, emphasizing that they can lead to very similar results. Indeed, for each halo we have performed constrained runs using both methods, and a mixture of the resulting data points are shown in Fig. 5. For a list of all the simulations used in this paper, see Table 1.

The results of four additional runs (one each for haloes 24 and 37, two for halo 40) are excluded from this figure. In each case, the estimated density profile was not well described by an NFW profile due to the halo undergoing a merger or the presence of large substructures. This is in agreement with Zhao et al. (2003) who find that at least part of the scatter around the Wechsler et al.

(2002) relation is due to poor fits to the NFW profiles and the mass accretion history.

6 LIKELIHOOD OF THE MODIFIED FIELD

As explained previously, the HR91 algorithm constructs a constrained realization which is equivalent to (but much more efficient than) rejection sampling, i.e. repeatedly drawing from an ensemble of ΛCDM universes until one obtains a realization that satisfies a given number of constraints. However, a naive choice of constraints can easily result in extreme configurations which are very unlikely to occur within the Hubble volume of the real universe. Depending on context, this could even be intentional (e.g. when investigating rare objects; Romano-Díaz et al. 2011a,b, 2014; Dubois et al. 2012); but nevertheless it is important to understand how likely it is for a given constrained configuration to arise, relative to the reference realization. We now derive a general expression for evaluating this likelihood and show how it is related to the abundance of haloes when changing the mass.

We can compare the unconstrained and constrained fields with respect to the *unmodified* ΛCDM covariance matrix $\mathbf{C}_0$ by evaluating the change in χ^2, defined as

$$\Delta\chi^2 = \boldsymbol{\delta}_n^{\dagger} \mathbf{C}_0^{-1} \boldsymbol{\delta}_n - \boldsymbol{\delta}_0^{\dagger} \mathbf{C}_0^{-1} \boldsymbol{\delta}_0, \tag{11}$$

where $\boldsymbol{\delta}_n$ is a field with n constraints. This constrained field has a relative abundance in the universe of $e^{-\Delta\chi^2/2}$ compared to the original, unconstrained field $\boldsymbol{\delta}_0$. Since this is only a relative abundance, applying a constraint to a halo that is rare in the reference simulation will in general also generate a rare object in the constrained run. We therefore modify several haloes in a similar way, in order to reduce the impact that picking a rare object may have on any of our results.
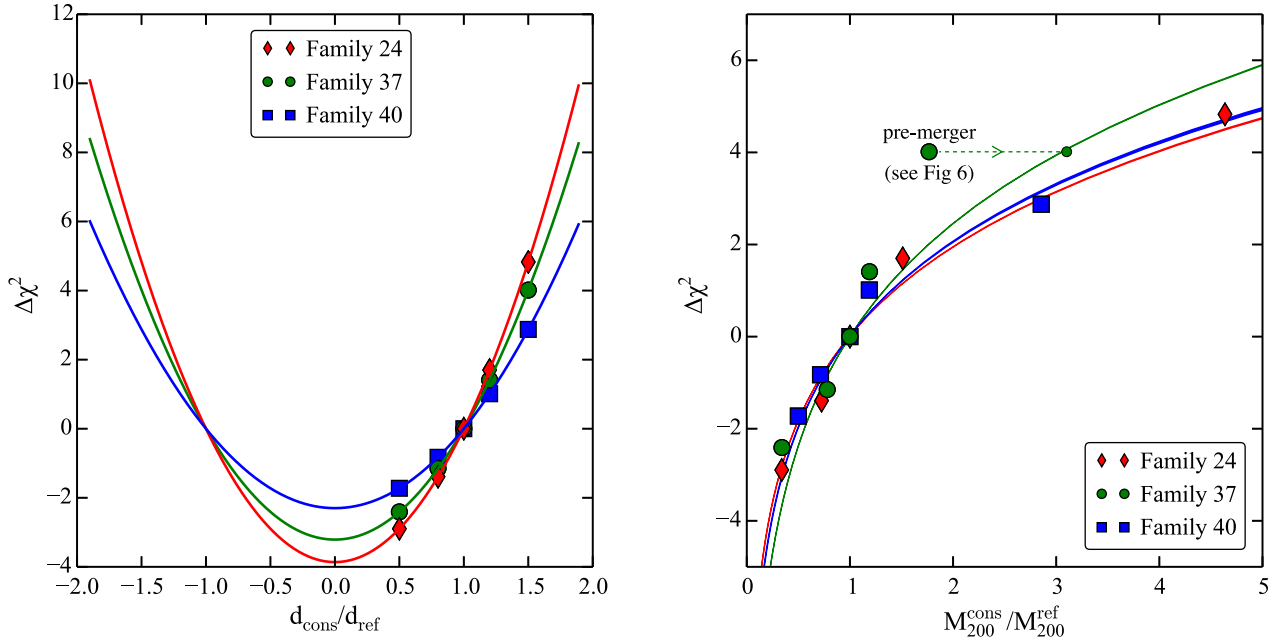
Figure 5. Left-hand panel: the relationship between $\Delta\chi^2$ and the initial overdensity for different halo families (H24-mass, H37-mass and H40-mass; see Table 1). Lines show the theoretical prediction from equation (13), whereas points give the actual change measured from the IC generator output, confirming that the algorithm is operating as expected. Right-hand panel: $\Delta\chi^2$ values (points) can be interpreted as giving the relative abundance of the haloes within each genetically modified family, and therefore should agree with estimates from a halo mass function (lines). The agreement is indeed good except for the H37-mass-1.5 point which appears to have too small a mass at $z = 0$ compared to expectations. This is because the halo mass function is based on an average mass build-up rate, whereas in this specific case the mass will only be acquired after a major merger in around 3 Gyr (see Fig. 6 and discussion in the text).

One can calculate $\Delta\chi^2$ directly from the density field, but we can also expand the above equation analytically by inserting equation (5) and making use of $\boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_j = \delta_{ij}$. This leads to a series of cancellations, with the final result

$$\Delta\chi^2 = \sum_{i=1}^{n} \left(|d_i|^2 - |d_{i0}|^2 \right), \qquad (12)$$

where we again use $d_{i0} = \boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta}_0$ to express the value of the constraints in the underlying realization.

Crucially, the details of the original realization $\boldsymbol{\delta}_0$ have disappeared except in the initial values of the constrained quantities, d_{i0}. In other words, the relative likelihood of the constrained simulation compared to the unconstrained case is dependent only on the choice of constraints. It is therefore specifically related to properties of the individual halo, not to details of its surroundings. This is another very desirable property of the HR91 formalism and reflects the minimality of the changes made to the field going from $\boldsymbol{\delta}_0$ to $\boldsymbol{\delta}_n$.

For a single constraint, equation (12) has a particularly transparent interpretation. Because of the normalization condition, the variance of d_{i0} for $i = n = 1$ in unconstrained realizations is

$$\langle d_0^* d_0 \rangle = \langle \boldsymbol{\delta}_0^\dagger \boldsymbol{\alpha}_1 \boldsymbol{\alpha}_1^\dagger \boldsymbol{\delta}_0 \rangle = \boldsymbol{\alpha}_1^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_1 = 1. \qquad (13)$$

Thus, for a single constraint, a change in $\Delta\chi^2$ of 1 corresponds to a 1σ variation in the property measured in the population at large.

The left-hand panel of Fig. 5 shows $\Delta\chi^2$ for a single constraint as a function of $d_{\mathrm{cons}}/d_{\mathrm{ref}}$, the ratio between a halo's average density contrast after the constraint and its value in the reference run (see Section 4). The values are calculated directly from the fields (points) and using equation (12) (lines); the two methods agree to within numerical accuracy, which is a useful verification of the algorithm. As before, the results for the three families are illustrated by red

diamonds, green circles and blue squares for haloes 24, 37 and 40, respectively. The minimum at $d_{\mathrm{cons}} = 0$, as well as the symmetry, is expected for a zero-mean Gaussian random field.

6.1 Connecting ICs and non-linear structure

In this section, we work towards establishing a quantitative connection between the degree of change in the ICs and in the final, non-linear structure. Using a single constraint, we investigate how a change in density contrast in the ICs is related to the resulting halo mass at late times. Qualitatively, an increase in overdensity should lead to a more massive object at $z = 0$ as explained in Section 4.

Quantitatively, the probability of finding a halo with mass M at $z = 0$ is given by the halo mass function, $n(M)\mathrm{d}M$, which depends on the cosmological power spectrum and growth function. Given two haloes of masses M_0 and M_1, their relative abundance is given by the ratio of the halo mass function at those masses, $n(M_1)/n(M_0)$. Assuming that the statistical properties of constrained and unconstrained simulations can be related by the change in mass of the target halo alone, this ratio also gives the relative probability of the structure in the two simulations.

Additionally, we can calculate the relative probability of the two ICs using equation (11); specifically

$$\frac{p(d_1)}{p(d_0)} = \exp\left[-\Delta\chi^2/2\right], \qquad (14)$$

where $p(d)$ is shorthand for the probability of a constrained field using one 'density constraint' with value d. Since we only consider probability ratios and $\Delta\chi^2$, any terms which only change the normalization of $p(d)$ have dropped out.

Now we have two methods of calculating the relative probabilities, and we can check whether they agree. We rewrite $p(d)$ in terms of halo mass using the conservation of probability

$$p(M) = p(d) \, d'(M), \tag{15}$$

where M is the halo mass at $z = 0$ and $d'(M)$ is the derivative of the constraint d with respect to the mass, evaluated at M. By combining equations (14) and (15), we can find a relationship between the $p(M_1)/p(M_0)$ and $\Delta\chi^2$, namely

$$\frac{p(M_1)}{p(M_0)} = \exp\left[-\Delta\chi^2/2\right] d'(M_1) \left[d'(M_0)\right]^{-1}. \tag{16}$$

If our assumptions are correct, this expression for $p(M_1)/p(M_0)$ should be equal to the mass function ratio $n(M_1)/n(M_0)$. We have used HMFcalc (Murray, Power & Robotham 2013) to generate a halo mass function at $z = 0$ for our cosmological model, and confirmed that it provides a good fit to our simulations.[2]

Evaluating equation (16) also requires an estimate of the Jacobian factors on the right-hand side. This can be obtained either from a physical model underlying the mass function or by using an empirically calibrated $M(d)$ relationship from the simulations. We chose the latter approach by fitting a power-law relation between d and M, which allows us to obtain values for $d'(M)$ at different halo masses separately for each halo family (24, 37 and 40; introduced in the previous section). This leaves us with a 'semi-analytical' prediction: theoretical halo mass function plus fit to the Jacobian.

The right-hand panel of Fig. 5 shows the results of the calculation (lines), as well as points evaluated directly from the simulations. Overall, most points show good agreement: there is consistency between the population statistics and the abundance calculated from the GM $\Delta\chi^2$ values. The broad agreement justifies our set of assumptions for calculating abundances in this specific case of a single density constraint. However, the individual haloes do scatter around the relation and there is one point that clearly does not fit the expectations. This arises from our H37-mass-1.5 run where we increase the initial overdensity of the proto-halo 37 region by a factor of 1.5.

The mismatch can be understood by considering the discrete nature of merger histories. Specifically, Fig. 6 shows the projected density at the last output ($z = 0$) in a region around halo 37 in the original run (upper panel) and the ×1.5 run (lower panel). In the latter case, a major merger (mass ratio ~2) will occur in around 3 Gyr. After this merger, the anomalous point will shift significantly rightwards in Fig. 5 to the correct mass ratio ($\simeq 3.1$) according to the $\Delta\chi^2(M)$ derived from the halo mass function.[3] We can frame this in another, more general way: the halo mass function is a statistical construction that corresponds to averaging over all possible histories, but the individual points in Fig. 5 represent modifications to specific haloes which have a discretized accretion history. Therefore, they scatter away from the line, especially when seen at special times (such as shortly before a major merger).

The main conclusion from Fig. 5 is therefore that the changes in the χ^2 give us a good quantitative handle on the relative abundance of haloes of different types, at least in this case where we have only changed the mass. However, the complex non-linear connection

---

[2] This is preferable to obtaining a (noisy) estimate of the mass function directly from our limited volume; the fitting functions included in the HM-Fcalc tool were validated using detailed studies of large simulation suites (e.g. Tinker et al. 2008).

[3] Note that the nearest massive halo in the original run corresponds to the same particles, but is considerably further away from halo 37 and is not on a trajectory that will lead to a merger within a Hubble time.
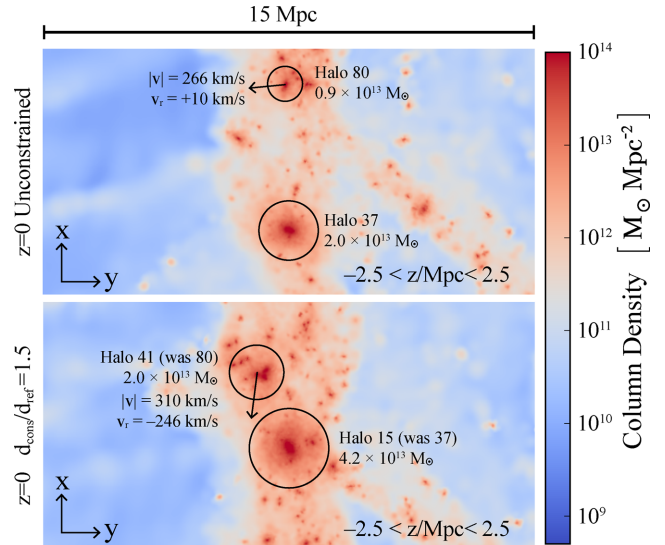
Figure 6. Slices from the original simulation (upper panel) and H37-mass-1.5 simulation (lower panel) illustrate how the target halo is, at $z = 0$, seen at a time where it is about to undergo a major merger in the latter case. For that reason, its mass undershoots the expectation from the halo mass function (Fig. 5) which averages over all the possible discrete realizations of the accretion history. Black circles show the size of the virial radius.

between initial and final states means that $\Delta\chi^2$ will always need to be interpreted with care.

7 DISCUSSION AND CONCLUSIONS

In this paper, we have demonstrated an extension of the HR91 technique to modify the ICs of a numerical simulation. For this modification, we selected regions of arbitrary shape, defined solely by the particles that form a halo in our reference simulation. This is a different approach than that used in previous works, which relied on imposing constraints of a given analytic profile. By applying our constraints only to the halo particles, we showed that we can 'genetically modify' a single object, changing its properties in a smooth and continuous way.

Using constraints on the density averaged over all halo particles controls the total mass, whereas adding additional constraints allows us to change the halo's collapse time. This serves as a demonstration of the technique and is the basis of the creation of further constraint types to study the impact of other halo characteristics on a halo's evolution.

Using the collapse time constraint, we investigated the density profiles of the resulting haloes at $z = 0$ and found that the distribution of their concentration parameter is consistent with the results of statistical analyses such as Wechsler et al. (2002). However, we also find that different haloes occupy different regions in the parameter space, and have different trajectories when their collapse time is changed. We plan to study this behaviour in future work, in order to determine which other halo parameters have changed. This should be complementary to the principal component analysis carried out by Skibba & Macciò (2011), Jeeson-Daniel et al. (2011) and Wong & Taylor (2012), which revealed somewhat inconclusive correlations between additional internal halo parameters. With our constrained simulations, we will not only be able to find correlations but to explicitly test their significance. Since we can directly compare the constrained halo to its reference in the unconstrained

run, we can establish exactly which changes in the halo parameters have a physical impact.

We have provided a way of quantifying the likelihood of modified ICs via a χ^2 difference between the constrained and unconstrained fields. In general, this statistic can be used to assess how compatible the modified object is with the underlying cosmology. Similar expressions were obtained by van de Weygaert & Bertschinger (1996); however, our orthonormalization procedure allows for the derivation of the considerably simpler equation (12). The statistic can be used to quantify the rarity of genetically modified objects relative to the unconstrained realization. As an example, we showed that modifying the mass produces abundance constraints that are quantitatively consistent with the traditionally measured halo mass function at $z = 0$. Individual haloes have discrete accretion histories and scatter around the mean relation predicted by the mass function; the strongest outlier in our study is about to undergo a merger at $z = 0$, which significantly lowers its current mass. In principle, the $\Delta\chi^2$ measure could also be used to specifically create objects that are 'rare' in a ΛCDM universe (similar to e.g. Romano-Díaz et al. 2011a,b, 2014; Dubois et al. 2012); we leave such a study to future work.

In this paper, we have used a uniform resolution over a box size of $50\ h^{-1}$ Mpc; having a sufficiently large box is important to ensure that haloes are embedded in the correct large-scale environment. Our code also produces ICs for 'zoom' simulations with varying resolution (see also Prunet et al. 2008; Romano-Díaz et al. 2014); the only major difference when generating these is the extra computational complexity introduced in the transformation between real space and Fourier space on an irregularly spaced grid, which has been tackled elsewhere in the literature (e.g. Bertschinger 2001; Hahn & Abel 2011). The real power of the approach to generate insight into a population from a handful of runs will become more apparent as we begin to use these zoom ICs in tandem with high-resolution baryonic physics.

While our focus here has been on basic properties such as the formation time and mass of a system, many other interesting aspects of evolution can be changed by constraining different properties. One example with which we are experimenting is the specific angular momentum, which can be controlled because tidal torque theory describes the connection between ICs and final spin (e.g. White 1984; Catelan & Theuns 1996; Porciani, Dekel & Hoffman 2002a,b); the internal properties of the galaxy forming inside the dark matter halo will naturally depend on the spin parameter of the halo. We are able to generate ICs that modify the spin parameter of the halo, but leave the mass and merger history untouched. Studies of halo spin constraints for dark matter and hydrodynamic simulations will be presented in future work.

ACKNOWLEDGEMENTS

REFERENCES

Bardeen J. M., Bond J. R., Kaiser N., Szalay A. S., 1986, ApJ, 304, 15
Bertschinger E., 1987, ApJ, 323, L103
Bertschinger E., 2001, ApJS, 137, 1
Bett P., Eke V., Frenk C. S., Jenkins A., Helly J., Navarro J., 2007, MNRAS, 376, 215
Bistolas V., Hoffman Y., 1998, ApJ, 492, 439
Brook C. B. et al., 2011, MNRAS, 415, 1051
Brook C. B., Di Cintio A., Knebe A., Gottlöber S., Hoffman Y., Yepes G., Garrison-Kimmel S., 2014, ApJ, 784, L14
Bullock J. S., Kolatt T. S., Sigad Y., Somerville R. S., Kravtsov A. V., Klypin A. A., Primack J. R., Dekel A., 2001, MNRAS, 321, 559
Catelan P., Theuns T., 1996, MNRAS, 282, 436
Codis S. et al., 2015, MNRAS, 448, 3391
Correa C. A., Wyithe J. S. B., Schaye J., Duffy A. R., 2015a, MNRAS, 450, 1514
Correa C. A., Wyithe J. S. B., Schaye J., Duffy A. R., 2015b, MNRAS, 450, 1521
Correa C. A., Wyithe J. S. B., Schaye J., Duffy A. R., 2015c, MNRAS, 452, 1217
Dayal P., Libeskind N. I., Dunlop J. S., 2013, MNRAS, 431, 3618
Doumler T., Gottlöber S., Hoffman Y., Courtois H., 2013, MNRAS, 430, 912
Dubois Y., Pichon C., Haehnelt M., Kimm T., Slyz A., Devriendt J., Pogosyan D., 2012, MNRAS, 423, 3616
Duffy A. R., Schaye J., Kay S. T., Dalla Vecchia C., 2008, MNRAS, 390, L64
Dunkley J. et al., 2009, ApJS, 180, 306
Forero-Romero J. E., Hoffman Y., Yepes G., Gottlöber S., Piontek R., Klypin A., Steinmetz M., 2011, MNRAS, 417, 1434
Frenk C. S. et al., 1999, ApJ, 525, 554
Genel S. et al., 2014, MNRAS, 445, 175
Governato F., Willman B., Mayer L., Brooks A., Stinson G., Valenzuela O., Wadsley J., Quinn T., 2007, MNRAS, 374, 1479
Guedes J., Callegari S., Madau P., Mayer L., 2011, ApJ, 742, 76
Hahn O., Abel T., 2011, MNRAS, 415, 2101
Heß S., Kitaura F.-S., Gottlöber S., 2013, MNRAS, 435, 2065
Hoffman Y., Ribak E., 1991, ApJ, 380, L5 (HR91)
Hoffman Y., Romano-Díaz E., Shlosman I., Heller C., 2007, ApJ, 671, 1108
Hopkins P. F., Cox T. J., Younger J. D., Hernquist L., 2009, ApJ, 691, 1168
Hopkins P. F., Cox T. J., Hernquist L., Narayanan D., Hayward C. C., Murray N., 2013, MNRAS, 430, 1901
Iliev I. T., Moore B., Gottlöber S., Yepes G., Hoffman Y., Mellema G., 2011, MNRAS, 413, 2093
Jasche J., Wandelt B. D., 2013, MNRAS, 432, 894
Jeeson-Daniel A., Dalla Vecchia C., Haas M. R., Schaye J., 2011, MNRAS, 415, L69
Kitaura F.-S., 2013, MNRAS, 429, L84
Klimentowski J., Łokas E. L., Knebe A., Gottlöber S., Martinez-Vaquero L. A., Yepes G., Hoffman Y., 2010, MNRAS, 402, 1899
Klypin A., Hoffman Y., Kravtsov A. V., Gottlöber S., 2003, ApJ, 596, 19
Klypin A., Yepes G., Gottlober S., Prada F., Hess S., 2014, preprint (arXiv:1411.4001)
Kormendy J., 2015, in Laurikainen E., ed., Galactic Bulges. Springer-Verlag, New York, p. 431
Kravtsov A. V., Klypin A., Hoffman Y., 2002, ApJ, 571, 563
Libeskind N. I., Yepes G., Knebe A., Gottlöber S., Hoffman Y., Knollmann S. R., 2010, MNRAS, 401, 1889
Ludlow A. D. et al., 2013, MNRAS, 432, 1103
Ludlow A. D., Navarro J. F., Angulo R. E., Boylan-Kolchin M., Springel V., Frenk C., White S. D. M., 2014, MNRAS, 441, 378
McBride J., Fakhouri O., Ma C.-P., 2009, MNRAS, 398, 1858

Macciò A. V., Dutton A. A., van den Bosch F. C., Moore B., Potter D., Stadel J., 2007, MNRAS, 378, 55

Macciò A. V., Dutton A. A., van den Bosch F. C., 2008, MNRAS, 391, 1940

Mathis H., Lemson G., Springel V., Kauffmann G., White S. D. M., Eldar A., Dekel A., 2002, MNRAS, 333, 739

Murray S. G., Power C., Robotham A. S. G., 2013, Astron. Comput., 3, 23

Naab T., Burkert A., Hernquist L., 1999, ApJ, 523, L133

Navarro J. F., Frenk C. S., White S. D. M., 1997, ApJ, 490, 493

Neto A. F. et al., 2007, MNRAS, 381, 1450

Nuza S. E., Parisi F., Scannapieco C., Richter P., Gottlöber S., Steinmetz M., 2014, MNRAS, 441, 2593

Papovich C. et al., 2015, ApJ, 803, 26

Pontzen A., Governato F., 2014, Nature, 506, 171

Pontzen A., Roškar R., Stinson G., Woods R., 2013, Astrophysics Source Code Library, record ascl:1305.002

Porciani C., Dekel A., Hoffman Y., 2002a, MNRAS, 332, 325

Porciani C., Dekel A., Hoffman Y., 2002b, MNRAS, 332, 339

Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, MNRAS, 338, 14

Prada F., Klypin A. A., Cuesta A. J., Betancort-Rijo J. E., Primack J., 2012, MNRAS, 423, 3018

Press W. H., Schechter P., 1974, ApJ, 187, 425

Prunet S., Pichon C., Aubert D., Pogosyan D., Teyssier R., Gottloeber S., 2008, ApJS, 178, 179

Ragone-Figueroa C., Plionis M., Merchán M., Gottlöber S., Yepes G., 2010, MNRAS, 407, 581

Reed D., Governato F., Verde L., Gardner J., Quinn T., Stadel J., Merritt D., Lake G., 2005, MNRAS, 357, 82

Robertson B., Bullock J. S., Cox T. J., Di Matteo T., Hernquist L., Springel V., Yoshida N., 2006, ApJ, 645, 986

Romano-Díaz E., Faltenbacher A., Jones D., Heller C., Hoffman Y., Shlosman I., 2006, ApJ, 637, L93

Romano-Díaz E., Hoffman Y., Heller C., Faltenbacher A., Jones D., Shlosman I., 2007, ApJ, 657, 56

Romano-Díaz E., Shlosman I., Trenti M., Hoffman Y., 2011a, ApJ, 736, 66

Romano-Díaz E., Choi J.-H., Shlosman I., Trenti M., 2011b, ApJ, 738, L19

Romano-Díaz E., Shlosman I., Choi J.-H., Sadoun R., 2014, ApJ, 790, L32

Schaye J. et al., 2015, MNRAS, 446, 521

Sherman J., Morrison W. J., 1950, Ann. Math. Stat., 21, 124

Skibba R. A., Macciò A. V., 2011, MNRAS, 416, 2388

Sorce J. G., Courtois H. M., Gottlöber S., Hoffman Y., Tully R. B., 2014, MNRAS, 437, 3586

Springel V., 2005, MNRAS, 364, 1105

Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726

Springel V. et al., 2008, MNRAS, 391, 1685

Tasitsiomi A., Kravtsov A. V., Gottlöber S., Klypin A. A., 2004, ApJ, 607, 125

Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, ApJ, 688, 709

van de Weygaert R., Bertschinger E., 1996, MNRAS, 281, 84

van den Bosch F. C., 2002, MNRAS, 331, 98

van Dokkum P. G. et al., 2013, ApJ, 771, L35

Wechsler R. H., Bullock J. S., Primack J. R., Kravtsov A. V., Dekel A., 2002, ApJ, 568, 52

White S. D. M., 1984, ApJ, 286, 38

Wong A. W. C., Taylor J. E., 2012, ApJ, 757, 102

Zaroubi S., Hoffman Y., Fisher K. B., Lahav O., 1995, ApJ, 449, 446

Zavala J., Jing Y. P., Faltenbacher A., Yepes G., Hoffman Y., Gottlöber S., Catinella B., 2009, ApJ, 700, 1779

Zhao D. H., Mo H. J., Jing Y. P., Börner G., 2003, MNRAS, 339, 12

Zhao D. H., Jing Y. P., Mo H. J., Börner G., 2009, ApJ, 707, 354

APPENDIX A: TECHNIQUE

Here, we will give a detailed description of the derivation of the HR91 operator and our equation (5).

The values of the three-dimensional overdensity field $\delta \equiv \delta(\boldsymbol{x})$ in the ICs of our cosmological simulations are distributed according to a multivariate Gaussian

$$p_0(\boldsymbol{\delta}) \propto \exp\left(-\frac{1}{2}(\boldsymbol{\delta} - \boldsymbol{\mu}_0)^\dagger \mathbf{C}_0^{-1}(\boldsymbol{\delta} - \boldsymbol{\mu}_0)\right), \tag{A1}$$

with mean $\langle\boldsymbol{\delta}\rangle = \boldsymbol{\mu}_0$ and $\mathbf{C}_0 = \langle(\boldsymbol{\delta} - \boldsymbol{\mu}_0)^\dagger(\boldsymbol{\delta} - \boldsymbol{\mu}_0)\rangle$ (the covariance or the power spectrum in Fourier space). Cosmological ICs have zero mean, $\boldsymbol{\mu}_0 = 0$, but we will consider the fully general case.

We will build the general procedure by induction. Suppose we have a $p_{i-1}(\boldsymbol{\delta})$ that describes the probability distribution function for $i - 1$ constraints; we now want to add the ith constraint, ensuring that $\boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta} = d_i$ for some constraint vector $\boldsymbol{\alpha}_i$ and constant d_i. To gain samples from the constrained distribution, one could sample from the original distribution and reject all those trials which lie too far away from $\left|\boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta} - d_i\right|^2 = 0$. Mathematically, this can be expressed by multiplying the original probability distribution by a penalty function, e.g.

$$p_i(\boldsymbol{\delta}) \propto \lim_{\beta\to\infty} p_{i-1}(\boldsymbol{\delta}) \exp\left(-\frac{\beta}{2}\left|\boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta} - d_i\right|^2\right), \tag{A2}$$

where the constant of proportionality renormalizes the probability distribution function and is dependent on β. In the limit $\beta \to \infty$, the penalty function becomes a Dirac-delta distribution and the constraint is satisfied exactly.

Under the assumption that p_{i-1} is Gaussian, the new probability function is the product of two Gaussians, and so remains Gaussian itself; consequently after imposing i constraints, we must be able to write

$$p_i(\boldsymbol{\delta}) \propto \exp\left(-\frac{1}{2}(\boldsymbol{\delta} - \boldsymbol{\mu}_i)^\dagger \mathbf{C}_i^{-1}(\boldsymbol{\delta} - \boldsymbol{\mu}_i)\right) \tag{A3}$$

for some mean $\boldsymbol{\mu}_i$ and covariance $\mathbf{C}_i$ which we will now derive. By multiplying out equation (A2), we obtain

$$p_i(\boldsymbol{\delta}) \propto \lim_{\beta\to\infty} \exp\left[-\frac{1}{2}(\boldsymbol{\delta} - \boldsymbol{\mu}_i)^\dagger\left(\mathbf{C}_{i-1}^{-1} + \beta\boldsymbol{\alpha}_i\boldsymbol{\alpha}_i^\dagger\right)(\boldsymbol{\delta} - \boldsymbol{\mu}_i)\right.$$
$$\left. - \boldsymbol{\mu}_i^\dagger\left(\mathbf{C}_{i-1}^{-1} + \beta\alpha_i\alpha_i^\dagger\right)\boldsymbol{\delta} + \boldsymbol{\mu}_{i-1}^\dagger\mathbf{C}_{i-1}^{-1}\boldsymbol{\delta} + \beta d_i\boldsymbol{\alpha}_i^\dagger\boldsymbol{\delta}\right], \tag{A4}$$

where we have already thrown away several terms which are zero order in $\boldsymbol{\delta}$ since they just change the normalization. By comparing terms in equations (A3) and (A4), we can first read off $\mathbf{C}_i^{-1} = \mathbf{C}_{i-1}^{-1} + \beta\boldsymbol{\alpha}_i\boldsymbol{\alpha}_i^\dagger$. We will also need a normalization for the $\boldsymbol{\alpha}_i$, which conveniently can be chosen[4] as

$$\boldsymbol{\alpha}_i^\dagger \mathbf{C}_{i-1}\boldsymbol{\alpha}_i = 1. \tag{A5}$$

Next, we apply the Sherman–Morrison formula (Sherman & Morrison 1950)

$$(\mathbf{C}_{i-1}^{-1} + \beta\boldsymbol{\alpha}_1\boldsymbol{\alpha}_1^\dagger)^{-1} = \mathbf{C}_{i-1} - \beta\frac{\mathbf{C}_{i-1}\boldsymbol{\alpha}_i\boldsymbol{\alpha}_i^\dagger\mathbf{C}_{i-1}}{1 + \beta\boldsymbol{\alpha}_i^\dagger\mathbf{C}_{i-1}\boldsymbol{\alpha}_i}$$
$$\simeq \mathbf{C}_{i-1}\left[1 - (1 - \beta^{-1})\boldsymbol{\alpha}_i\boldsymbol{\alpha}_i^\dagger\mathbf{C}_{i-1}\right], \tag{A6}$$

where we have used $\beta > 1$ and the normalization condition (A5) in the second step.

---

[4] Unless $\boldsymbol{\alpha}_i$ is a null direction of $\mathbf{C}_{i-1}$, but then there would be zero probability of our constraint in the original distribution.

The terms in the second line of equation (A4) have to cancel exactly. Plugging equation (A6) into this expression leads to

$$\boldsymbol{\mu}_i = \lim_{\beta \to \infty} \boldsymbol{\mu}_{i-1} - (1 - \beta^{-1}) \mathbf{C}_{i-1} \boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^\dagger \boldsymbol{\mu}_{i-1} + d_i \boldsymbol{\alpha}_i \mathbf{C}_{i-1} \boldsymbol{\alpha}_i^\dagger$$

$$= \boldsymbol{\mu}_{i-1} + \mathbf{C}_{i-1} \boldsymbol{\alpha}_i \left(d_i - \boldsymbol{\alpha}_i^\dagger \boldsymbol{\mu}_{i-1} \right), \tag{A7}$$

and finally taking the limit in equation (A6) yields

$$\mathbf{C}_i = \mathbf{C}_{i-1} - \mathbf{C}_{i-1} \boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^\dagger \mathbf{C}_{i-1}. \tag{A8}$$

This result allows us to apply as many constraints as desired analytically – by looping over the constraints and updating the covariance matrix and mean at each step, then drawing a constrained realization – but this would be computationally expensive. Instead, the constrained realization can be constructed from the unconstrained field using a projection operator, which we will now derive.

For notational simplicity, in addition to normalizing the constraints, it is also helpful to make them orthogonal (e.g. through a Gram–Schmidt procedure) in the sense that $\boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_j = 0$ for $i \neq j$. One can then verify by substitution (see Appendix A1) that the constrained field has mean

$$\boldsymbol{\mu}_n = \boldsymbol{\mu}_0 + \sum_{i=1}^n \mathbf{C}_0 \boldsymbol{\alpha}_i \left(d_i - \boldsymbol{\alpha}_i^\dagger \boldsymbol{\mu}_0 \right) \tag{A9}$$

and covariance

$$\mathbf{C}_n = \mathbf{C}_0 - \sum_{i=1}^n \mathbf{C}_0 \boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \tag{A10}$$

for orthonormalized $\{\boldsymbol{\alpha}_i\}$.

Efficiently drawing from the distribution implied by the above mean and covariance is made possible by any operator $\mathbf{P}_n$ that takes a realization from the unconstrained field $\boldsymbol{\delta}_0$ and forms a new realization under n constraints via the ansatz

$$\boldsymbol{\delta}_n = \mathbf{P}_n \left(\boldsymbol{\delta}_0 - \boldsymbol{\mu}_0 \right) + \boldsymbol{\mu}_n, \tag{A11}$$

where to gain the correct covariance $\mathbf{C}_n = \langle (\boldsymbol{\delta}_n - \boldsymbol{\mu}_n)(\boldsymbol{\delta}_n - \boldsymbol{\mu}_n)^\dagger \rangle$ one must demand

$$\mathbf{P}_n \mathbf{C}_0 \mathbf{P}_n^\dagger = \mathbf{C}_n. \tag{A12}$$

There are an infinity of operators $\mathbf{P}_n$ with this property. Given any specific $\mathbf{P}_n$, one can form $\mathbf{P}_n' = \mathbf{U} \mathbf{P}_n$, where $\mathbf{U}^\dagger \mathbf{C}_0 \mathbf{U} = \mathbb{1}$, and the new $\mathbf{P}_n'$ satisfies the required identity (A12). To obtain the unique HR91 operator, we additionally require $\mathbf{P}_n$ to make *minimal* changes to the field. This implies $\mathbf{P}_n \boldsymbol{\delta}_n = \boldsymbol{\delta}_n$ – in other words, that no changes are made if the field already satisfies the constraints. Using equation (A11), it immediately follows that $\mathbf{P}_n^2 = \mathbf{P}_n$ and $\mathbf{P}_n \boldsymbol{\mu}_n = \boldsymbol{\mu}_n$. The first of these conditions implies that all eigenvalues of $\mathbf{P}_n$ are either 1 or 0.

One can verify by substitution that all these requirements are satisfied by

$$\mathbf{P}_n = \mathbb{1} - \sum_{i=1}^n \mathbf{C}_0 \boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^\dagger \text{(for orthonormalized } \{\boldsymbol{\alpha}_i\}), \tag{A13}$$

Note that the HR91 form given in their equations (2–4) builds the orthonormalization procedure into the projection operator (appearing as ξ_{ij}^{-1} in their notation). However, as stated above, we found it notationally simpler to pre-condition the constraints into orthonormal form using the Gram–Schmidt procedure. Both formulations are mathematically equivalent (see Appendix B).

Inserting equations (A13) and (A9) into equation (A11) then leads to the final expression

$$\boldsymbol{\delta}_n = \boldsymbol{\delta}_0 + \sum_{i=1}^n \mathbf{C}_0 \boldsymbol{\alpha}_i \left(d_i - d_{i0} \right), \tag{A14}$$

where $d_{i0} = \boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta}_0$.

In practice, most of the necessary calculations are performed in Fourier space, because there $\mathbf{C}_0$ is the ΛCDM power spectrum which is diagonal. Any constraint vector $\boldsymbol{\alpha}_i$ and density field $\boldsymbol{\delta}$ can be easily converted using numerical fast Fourier transformations.

Note that the algorithm in its current form only takes into account the contribution from the power spectrum. If one wanted to generate constrained ICs based on an observational data set, the associated uncertainties would introduce extra contributions in the new covariance matrix (e.g. Zaroubi et al. 1995; van de Weygaert & Bertschinger 1996), which is not included in the current implementation.

A1 Comments on normalization

Throughout this paper, we use the same notation for the normalized and unnormalized constraints (expressed by $\boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta} = d_i$). In practice, these quantities are affected by the normalization condition in the following way: if $\boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_i = \kappa_i$ before normalization, then we immediately find $\boldsymbol{\alpha}_i \to \boldsymbol{\alpha}_i / \sqrt{\kappa_i}$ to satisfy $\boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_i = 1$. Accordingly, the constant d_i transforms as $d_i \to d_i / \sqrt{\kappa_i}$ as well, such that $\boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta} = d_i$ is still obeyed after normalization.

The consistency of the Gram–Schmidt condition $\boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_j = \delta_{ij}$ and our normalization $\boldsymbol{\alpha}_i^\dagger \mathbf{C}_{i-1} \boldsymbol{\alpha}_i = 1$ can be shown as follows: multiplying (A10) by $\boldsymbol{\alpha}_n$ on both sides yields

$$\boldsymbol{\alpha}_n^\dagger \mathbf{C}_{n-1} \boldsymbol{\alpha}_n = \boldsymbol{\alpha}_n^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_n$$

$$- \sum_{i=1}^{n-1} \boldsymbol{\alpha}_n^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_n, \tag{A15}$$

where we have shifted the index n by 1 compared to equation (A10) for ease of notation. Inserting $\boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_j = \delta_{ij}$ into the right-hand side leads to

$$\boldsymbol{\alpha}_n^\dagger \mathbf{C}_{n-1} \boldsymbol{\alpha}_n = 1 - \sum_{i=1}^{n-1} \delta_{ni} \delta_{in} = 1, \tag{A16}$$

which proves the equality.

APPENDIX B: EQUIVALENCE TO HR91

In this appendix, we translate our notation into the one used by HR91 to show that our final expression (A14) is equivalent to their equation (4), which states

$$\boldsymbol{\delta}_n = \boldsymbol{\delta}_0 + \sum_{i,j=1}^n \boldsymbol{\xi}_i \xi_{ij}^{-1} (d_i - d_{i0}), \tag{B1}$$

where we have already translated their notation for the constrained and unconstrained field, and the values of the constraints ($f \to \delta$, $c_i \to d_i$). The two additional functions are

$$\boldsymbol{\xi}_i = \langle \boldsymbol{\delta}_0 \mathcal{C}_i^\dagger \rangle \tag{B2}$$

and

$$\xi_{ij} = \langle \mathcal{C}_i \mathcal{C}_j^\dagger \rangle, \tag{B3}$$

where $\mathcal{C}_i = \boldsymbol{\alpha}_i^\dagger \boldsymbol{\delta}_0$ (not to be confused with our covariance matrix $\mathbf{C}_i$), and we have added daggers to the HR91 notation to allow for complex-valued quantities. Inserting these expressions into equation (B1) leads to

$$\boldsymbol{\delta}_n = \boldsymbol{\delta}_0 + \sum_{i,j=1}^n \mathbf{C}_0 \boldsymbol{\alpha}_i \left[\boldsymbol{\alpha}_i^\dagger \mathbf{C}_0 \boldsymbol{\alpha}_j \right]^{-1} (d_i - d_{i0}) \qquad (B4)$$

$$= \boldsymbol{\delta}_0 + \sum_{i=1}^n \mathbf{C}_0 \boldsymbol{\alpha}_i (d_i - d_{i0}), \qquad (B5)$$

where in the first step we used that $\mathbf{C}_0 = \langle \boldsymbol{\delta}_0 \boldsymbol{\delta}_0^\dagger \rangle$ because $\boldsymbol{\mu}_0 = 0$ in HR91, and our orthonormalization condition in the second step. This shows that the Gram–Schmidt approach is equivalent to performing the matrix inversion ξ_{ij}^{-1}.

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