The journey of QSO haloes from $z \sim 6$ to the present

R. E. Angulo,[1]⋆ V. Springel,[2,3] S. D. M. White,[1] S. Cole,[4] A. Jenkins,[4] C. M. Baugh[4] and C. S. Frenk[4]

[1]*Max-Planck-Institute for Astrophysics, Karl-Schwarzschild-Str. 1, 85740 Garching, Germany*
[2]*Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnenweg 35, 69118 Heidelberg, Germany*
[3]*ARI, Zentrum für Astronomie der Universität Heidelberg, Mönchhofstr. 12-14, 69120 Heidelberg, Germany*
[4]*Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE*

ABSTRACT

We apply a recently developed scaling technique to the Millennium-XXL, one of the largest cosmological N-body simulations carried out to date (3×10^{11} particles within a cube of volume $\sim 70 \, \mathrm{Gpc}^3$). This allows us to investigate the cosmological parameter dependence of the mass and evolution of haloes in the extreme high-mass tail of the $z = 6$ distribution. We assume these objects to be likely hosts for the population of rare but ultraluminous high-redshift quasars discovered by the Sloan Digital Sky Survey (SDSS). Haloes with a similar abundance to these quasars have a median mass of $9 \times 10^{12} \, \mathrm{M}_\odot$ in the currently preferred cosmology, but do not evolve into equally extreme objects at $z = 0$. Rather, their descendants span the full range, conventionally assigned to present-day clusters, 6×10^{13}–$2.5 \times 10^{15} \, \mathrm{M}_\odot$ for the same cosmology. The masses both at $z = 6$ and at $z = 0$ shift up or down by factors exceeding 2 if cosmological parameters are pushed to the boundaries of the range discussed in published interpretations of data from the *Wilkinson Microwave Anisotropy Probe* satellite. The main factor determining the future growth of a high-mass $z = 6$ halo is the mean overdensity of its environment on scales of 7–14 Mpc, and descendant masses can be predicted six to eight times more accurately if this density is known than if it is not. All these features are not unique to extreme high-z haloes, but are generic to hierarchical growth. Finally, we find that extreme haloes at $z = 6$ typically acquire about half of their total mass in the preceding 100 Myr, implying very large recent accretion rates which may be related to the large black hole masses and high luminosities of the SDSS quasars.

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

1 INTRODUCTION

Bright $z \sim 6$ quasars are extremely rare objects. Their luminosity is thought to be a result of supermassive black holes accreting gas at enormous rates (Fan et al. 2003; Kollmeier et al. 2006; Shen et al. 2008) at a time when only ~ 10 per cent of the mass in the Universe was in haloes where gas could cool efficiently and about half was still in diffuse form (Angulo & White 2010a). If these quasars are long lived, then their dark matter haloes belong to an equally extreme tail (presumably the most massive tail) of the halo distribution, a hypothesis supported by the strong observed clustering of high-z quasars (e.g. Shen et al. 2007). The observed QSO number density (about 1 Gpc^{-3}) combined with abundance matching implies that their host halo masses are well above $10^{12} \, \mathrm{M}_\odot$ at $z = 6$. These masses correspond to much more extreme peaks in the initial

Gaussian density fluctuation field than those associated with even the largest galaxy clusters today. Such quasars are clearly exceptional events in a Λ cold dark matter (ΛCDM) universe, but their properties, together with those of the intergalactic absorption seen in their spectra, encrypt key information about many astrophysical processes (e.g. Fan 2006) and perhaps also about the background cosmological model.

Studying the properties, environment and fate of the high-mass haloes in which the $z \sim 6$ quasars may live is a challenging task for cosmological N-body simulations. It requires a very large computational volume to obtain a representative sample of extremely rare objects, sufficient spatial and mass resolution to resolve their structure reliably, and sufficient time resolution to build the merger trees needed to trace evolution throughout cosmic time. Previous studies in this area have relied on analytic models calibrated using simulations of less extreme objects (e.g. Trenti, Santos & Stiavelli 2008) or resimulation techniques applied to a limited number of systems (Li et al. 2007; Sijacki, Springel & Haehnelt 2009; Romano-Diaz

⋆ E-mail: reangulo@gmail.com

et al. 2011). In this paper, we present an extremely large *N*-body calculation, which meets the computational challenges directly by simulating a very large volume at adequate resolution within the ΛCDM paradigm.

With this simulation in hand, we study a variety of topics associated with the assembly and the future of quasar haloes. For example, where are today's descendants of the massive black holes that powered quasars at $z \sim 6$? This question was explored by Springel et al. (2005) using the Millennium Simulation (MS). They concluded that these black holes would today lie at the centres of cD galaxies in massive galaxy clusters. However, the MS is too small to contain even one object like the Sloan Digital Sky Survey (SDSS) $z \sim 6$ quasars, so this conclusion was based on a small number of less rare systems. Trenti et al. (2008) used analytic tools to extend these MS-based results, arguing for a much greater diversity in the present-day descendants of SDSS QSO haloes. A similar conclusion was reached by Di Matteo et al. (2008) with cosmological smoothed particle hydrodynamic simulations. Here, we are able to use fully resolved *N*-body merger trees for haloes identified at the same space density as the brightest SDSS quasars. We confirm that these massive $z = 6$ haloes should evolve into a variety of systems today. Most of their associated black holes should end up in the central galaxies of haloes ranging from rich groups to clusters. We show that the median mass of descendants approaches $10^{15} \, \mathrm{M_\odot}$, but this value and that of early haloes themselves depend on the parameters assumed for the background cosmological model.

Another issue we investigate is the large-scale environment surrounding extreme high-redshift haloes. Such haloes are expected to be strongly biased towards overdense regions. Although we confirm this, we also show that there is considerable scatter, and many extreme haloes have environments of moderate overdensity. Indeed, some even have environments which are slightly underdense.

Finally, we explore the assembly histories predicted for extreme $z = 6$ haloes for three different assumptions about the parameters underlying the background ΛCDM cosmology. We measure the accretion rates of these haloes over the time period immediately preceding the epoch of observation and show that these can be extremely high. This may well be related to the fuelling of the extraordinarily high luminosities and masses measured for the $z \sim 6$ quasars.

This paper is set out as follows. In Section 2, we present technical details of our simulation and of the methods we use to identify dark matter haloes and their evolutionary paths. Section 3 then presents our results, including the properties of the $z = 0$ descendants of high-z quasar haloes and their assembly histories at higher redshift. Our conclusions are presented in Section 4.

2 NUMERICAL METHODS

In this section we describe the numerical tools used in this paper. We first present our *N*-body simulation (Section 2.1), including a description of the construction of halo and subhalo catalogues and of merger trees (Section 2.1.1). In Section 2.2 we then explain how these numerical data can be used to explore structure formation in cosmologies other than the one used to carry out the original simulation. The last subsection (Section 2.3) defines the samples of extreme haloes that we will study in the remainder of the paper.

2.1 The MXXL *N*-body simulation

Our simulation, named the Millennium-XXL or MXXL, represents the matter content of a ΛCDM universe using more than 303 bil-

Table 1. Parameters of the original MXXL simulation and of the scaled versions used to represent other cosmologies. The columns are as follows: (1) the name of the simulation; (2) the mass of a dark matter particle in units of $10^{10} \, \mathrm{M_\odot}$; (3) the side of the computational box in units of Mpc; (4) the total matter density; (5) the baryon density; (6) the linear fluctuation amplitude at $z = 0$. In all cases, the primordial spectral index is $n_s = 1$, the Hubble constant at $z = 0$ is $H_0 = 73 \, \mathrm{km \, s^{-1} \, Mpc^{-1}}$ and the dark energy is assumed to be a cosmological constant.

	m_{dm}	Box	Ω_m	Ω_b	σ_8
MXXL(M1)	0.95	4167	0.250	0.045	0.9
M3	1.64	5091	0.238	0.0416	0.761
M7	1.26	4467	0.272	0.0416	0.807

lion particles (6720^3) within a comoving cube of side 4.1 Gpc. Our choice of cosmological parameters is identical to that used in the other MSs (see Table 1; Springel et al. 2005; Boylan-Kolchin et al. 2009), implying a particle mass of $8.45 \times 10^9 \, \mathrm{M_\odot}$. This set of parameters is inconsistent with the most recent constraints from observations of the cosmic microwave background and low-redshift large-scale structure (Komatsu et al. 2011), but, as discussed in Section 2.2, our results can be scaled accurately to any nearby cosmology, including those that are now more favoured. The comoving Plummer-equivalent softening length of the simulation is $\epsilon = 13.7 \, \mathrm{kpc}$, implying $\sim 10^{15}$ effective resolution elements in the full simulated volume (the forces are exactly Newtonian only beyond $2.8 \times \epsilon$). The enormous statistical power and dynamical range that it implies are illustrated in Fig. 1, which shows the $z = 0$ density field on four different scales, starting from the whole box in the background, then zooming progressively on to the most massive halo in the insets.

The initial phase-space distribution of the particles was set up at $z = 63$ by perturbing a glass-like distribution (Baugh, Gaztanaga & Efstathiou 1995; White 1996) using second-order Lagrangian perturbation theory (2LPT; Scoccimarro 1998). The use of 2LPT has several advantages over the more common Zel'dovich approximation (Zel'dovich 1970), including small and rapidly decaying transients in matter clustering and a better representation of the perturbations that seed extremely large dark matter haloes (Crocce, Pueblas & Scoccimarro 2006; Knebe et al. 2009; Jenkins 2010). The latter is particularly relevant for this paper. The amplitude of individual Fourier modes was set hierarchically as described in Jenkins (in preparation). This allowed an efficient and consistent generation of initial conditions for the resimulation of any selected subregion at, in principle, arbitrarily high-mass resolution.

Note that the MXXL is the first simulation with a large enough volume and a small enough particle mass to contain a representative sample of well-resolved (≥ 1000 particles; Trenti et al. 2010) haloes which could represent the hosts of the $z \sim 6$ SDSS quasars. Performing a simulation with these characteristics posed severe computational challenges with respect to raw execution time, scalability of the simulation algorithms, memory consumption and I/O performance. Thanks to a highly optimized simulation code and one of the largest supercomputers in Europe, these challenges were successfully met. Specifically, the MXXL was carried out with a special version of the GADGET3 code (Springel 2005), which aggressively reduced peak memory consumption at runtime, incorporated a number of analysis tools on the fly and implemented a special
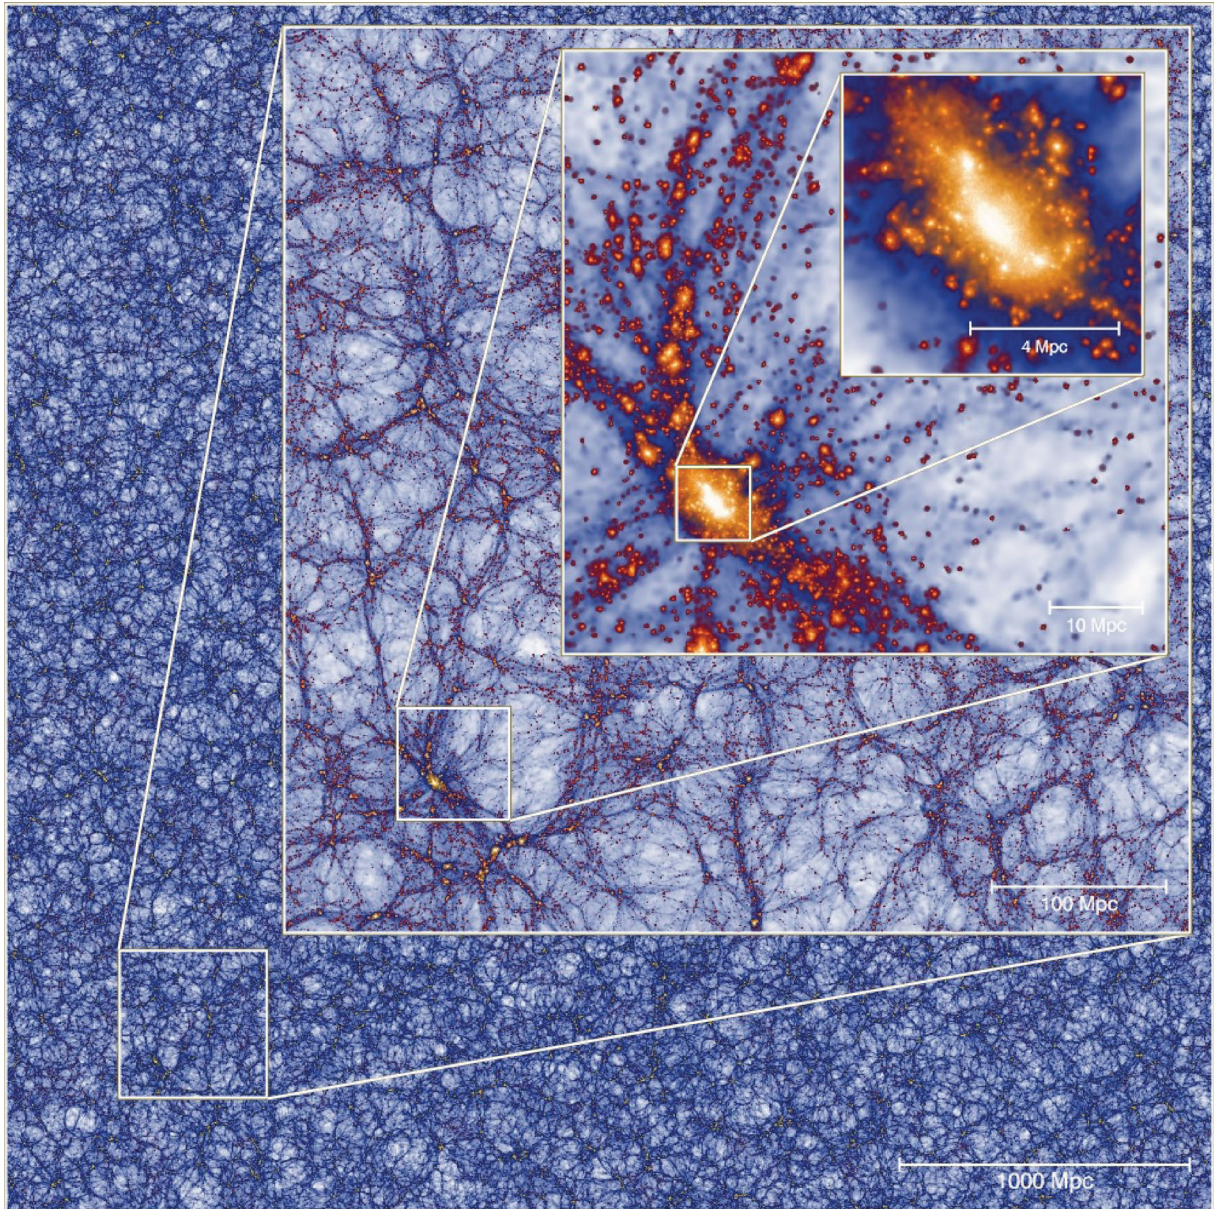
Figure 1. Images of the matter density field in the MXXL focusing on the most massive halo present in the simulation at $z = 0$. Each inset zooms by a factor of 8 from the previous one; the side length varies from 4.1 Gpc down to 8.1 Mpc. The intensity of each pixel is proportional to the logarithm of the dark matter density projected through a 25 Mpc thick slab. This simulation has a dynamic range of 10^5 on each spatial dimension, simultaneously resolving the internal structure of collapsed objects and the large-scale quasi-linear fluctuations in a ΛCDM universe.

compression of the output data. As a result, the MXXL was completed in late summer 2010 in less than 3 million CPU hours (including post-processing calculations), using 30 Tb of RAM and 12 228 cores of the Juropa cluster at the Jülich Supercomputer Center in Germany. We refer the readers to Angulo et al. (2012) for more details about the simulation.

2.1.1 Halo and subhalo catalogues

We identified self-bound halo/subhalo structures throughout the MXXL at the same 64 redshifts used in the MS and MS-II. This redshift frequency (roughly equally spaced in time by 300 Myr for $z < 2$ and by 100 Myr at $z \sim 6$) allowed us to build detailed merger trees. At each output time, we first applied a friends-of-friends (FOF) algorithm (Davis et al. 1985), with a linking length of 0.2 times

the mean interparticle separation to build an FOF group catalogue down to a limit of 20 particles. We then used a memory-efficient implementation of the SUBFIND algorithm (Springel et al. 2001) to identify self-bound substructures within each FOF group down to a limit of 15 particles. These calculations were performed *on the fly* during the *N*-body calculation, so that it was not necessary to store the particle data at all output times. This significantly reduced the I/O and storage requirements of the simulation.

Summing over all output times, there are 2.5×10^{10} FOF groups in the MXXL with more than 20 particles. At $z = 6$ there are 3.7×10^7 such groups and at $z = 0$, 6.5×10^8. The most massive FOF group at $z \sim 6$ contains 3285 particles and 4 substructures. This is about 300 times less massive than the biggest halo at the $z = 0$ snapshot, which contains 1062 232 particles and 688 substructures with more than 15 particles.

Finally, we built 'merger trees' similar to those described in Springel et al. (2005) in order to follow the evolution of a halo/subhalo structure in detail. For every subhalo in our catalogues we define a pointer to a unique descendant in the subsequent snapshot by locating the subhalo containing the greatest number of its 15 most bound particles.[1] These pointers are used to create a tree-like data structure, which represents the full assembly history, the current substructure and the future evolution of every halo. In particular, this allows us to map out the formation histories of our putative $z = 6$ quasar hosts and to follow their later evolution down to $z = 0$.

2.2 Exploring structure formation in other cosmologies

The MXXL is an extremely expensive numerical simulation and so can only be carried out once for a specific set of cosmological parameters. However, the rescaling method of Angulo & White (2010b) allows us to use the simulation to analyse any neighbouring cosmology with Gaussian initial fluctuations, and in this paper we will show results not only for the original MS cosmology but also for two other versions of the standard ΛCDM cosmology. The accuracy of the rescaling scheme is remarkably high if it is applied carefully; masses of individual objects are reproduced to better than 10 per cent and positioned to better than 100 kpc (Angulo & White 2010b; Ruiz et al. 2011). In the following, we briefly recap the main features of the method.

First, consider a 'target' cosmological model at $z = z_B$ which we seek to match using the results of an 'original' cosmological simulation of side length L_A (in units of h^{-1} Mpc). The heart of the method is to find length transformation, $L_A \rightarrow L_B = sL_A$, and relabelling of the time variable $z_A \rightarrow z_B$, by requiring that the variance of the linear density field in the target cosmology, $\sigma_B^2(R, z_B)$, over the range $[R_1, R_2]$ is as close as possible to that of the original cosmology, $\sigma_A^2(R, z_A)$, over the range $[s^{-1}R_1, s^{-1}R_2]$ at redshift z_A. Thus, we minimize

$$\int_{R_1}^{R_2} \frac{\mathrm{d}R}{R} \left[\sigma_B^2(R) D_B(z_B) - \sigma_A^2(s^{-1}R) D_A(z_A) \right]^2 \quad (1)$$

over s and z_A, where $D(z)$ is the linear growth factor in units of its present-day value.

In the Press–Schechter theory (Press & Schechter 1974) the halo mass function is determined by the linear variance of the underlying dark matter field; thus, equation (1) also minimizes the difference between the halo mass functions in the target and original cosmologies over the mass range $[M(R_2), M(R_1)]$. We usually take $M(R_2)$ to be the mass of the largest halo in the simulation at the lowest redshift of interest ($z = 0$ here) and $M(R_1)$ that of the least massive resolved halo.

As a result, the original box size will be expanded by a factor of s, and the output at redshift z_A will represent redshift z_B in the target cosmology. Redshifts in the target cosmology (z') corresponding to the redshifts (z) of stored data in the original simulation are then determined implicitly by $D_B(z') = [D_A(z)/D_A(z_A)]D_B(z_B)$. The mass of a simulation particle (in units of $M_\odot$) in the target cosmology is $m_B = (\Omega_{m,B} H_B^2 / \Omega_{m,A} H_A^2) s^3 m_A$, where $\Omega_{m,X}$ and H_X ($X = A$ or B) are the dimensionless total matter densities and Hubble parameters of the two cosmologies.

The second part of the algorithm of Angulo & White (2010b) corrects differences in power spectrum shape on large scales between the original and target cosmologies by altering the amplitude of quasi-linear modes using the Zel'dovich approximation. In this paper we do not make this correction since we are interested in the internal structure, abundance and evolution of massive haloes rather than their spatial distribution. We only use the original cosmology when looking at the overdensity around haloes, a quantity that is slightly affected by this correction.

With this technique we have created two additional halo catalogues in alternative cosmologies which we denote by M3 and M7. These have cosmological parameters motivated by the 3-yr (Spergel et al. 2007) and 7-yr (Komatsu et al. 2011) analyses of data from the *Wilkinson Microwave Anisotropy Probe* (*WMAP*) satellite. The main features of these models are values for σ_8 which are lower than those in the MXXL and the other MSs, and different values for Ω_m (see Table 1). The corresponding length scalings are $s = 1.222$ and 1.072 for M3 and M7, respectively. The $z = 0.623$ and 0.319 outputs of the MXXL represent $z' = 0$ in the M3 and M7 cosmologies.

2.3 QSO haloes

In this paper we will assume that the QSO luminosity is a monotonically increasing function of the FOF host halo mass at any given time and that there is a duty cycle that is independent of the halo mass. Models with these characteristics appear to be preferred by clustering analyses (e.g. White, Martini & Cohn 2008; Bonoli et al. 2010), but we note that they are not the only possibility. In physically motivated models, the QSO host depends on the details of the galaxy formation model and in particular on active galactic nucleus feedback (Marulli et al. 2008; Bonoli et al. 2009; Fanidakis et al. 2011, 2012).

Therefore, we consider halo samples limited by FOF mass at two different thresholds corresponding to two different abundances at $z \sim 6$ which we keep constant across our three cosmologies. 'Long-lived QSO haloes' (LLQ haloes) have a comoving number density of $n = 0.4 \, \mathrm{Gpc}^{-3}$, requiring FOF masses above $[15.4, 9, 5] \times 10^{12} \, M_\odot$ for the M1, M7 and M3 cosmologies, respectively. The number density of this sample (which is well below the limit that could be probed by the MS) matches that of the extremely bright quasars observed at $z > 5$ in the SDSS ($n = 0.6 \, \mathrm{Gpc}^{-3}$; Fan et al. 2003, 2006). Thus, they mimic the assembly history of SDSS QSOs if they have a 100 per cent duty cycle, i.e. if they shine constantly at their full brightness.

Our second sample, 'Short-lived QSO haloes' (SLQ haloes), is selected at 30 times higher abundance, i.e. $n = 11.6 \, \mathrm{Gpc}^{-3}$. This corresponds to minimum FOF halo masses of $[7.0, 4.2, 2.0] \times 10^{12} \, M_\odot$ for the M1, M7 and M3 cosmologies. This sample could be regarded as representing the hosts of the high-redshift SDSS quasars if these have a 1/30 duty cycle, i.e. each object shines at full brightness only 1/30 of the time.

Due to our length scaling, the total number of haloes in our two samples varies by a factor of s^3 between the three cosmologies; the SLQ samples contain 810, 997 and 1478 objects for the M1, M7 and M3 cases, respectively (note that we expect four objects above this mass threshold in the MS which, in fact, contained only two), whereas the LLQ samples contain 27, 34 and 50 haloes. The redshift of the samples also differs between cosmologies because the MXXL data were stored only at a discrete set of times; we use $z = 6.18$ (M1), 6.19 (M3) and 5.96 (M7) for the three cases.

---

[1] If two subhaloes contain the same number of these particles, we choose the one with the largest total binding energy.

Finally, we note that alternative sample definitions, for example, using virial masses or circular velocities rather than FOF masses do not produce significant differences in our results. 70 per cent of the 1000 most massive FOF haloes rank within the 1800 most massive haloes according to M_{200} and within the 4000 according to peak circular velocity.

3 RESULTS AND DISCUSSION

Using the numerical tools and halo samples presented in the previous section, we now study the evolution of the host haloes of $z \sim 6$ quasars. We look into the type of objects they turn into at the present day (Section 3.1) and how the masses of these descendants are related to the environments of their $z \sim 6$ progenitors (Section 3.3). In addition, we explore the growth rates of our quasar host haloes both prior (Section 3.4) and subsequent (Section 3.2) to their identification at $z \sim 6$.

3.1 The mass and fate of $z \sim 6$ quasar hosts

Since we are assuming that high-redshift quasars live in the most massive objects present at that time, one might naively expect that their descendants will lie at the centres of the most massive haloes at any later time, in particular, in the central galaxies of large galaxy clusters today. However, in this section we confirm the results of previous studies showing that this is not generally true (Trenti et al. 2008; Overzier et al. 2009). The descendants of QSOs can be found in haloes spanning a factor of 30 in mass, and in a few cases, they are not even located in the central object of this halo but in an orbiting subhalo.

The upper panel of Fig. 2 shows differential mass functions for the two $z \sim 6$ QSO host halo samples (defined in Section 2.3) and for the three cosmological models of Table 1, which differ primarily in their values of σ_8 and Ω_m. By construction, these samples consist of very massive haloes. Their mean mass ranges from 2.6×10^{12} to 8×10^{12} M$_\odot$ for the SLQ samples and from 5.2×10^{12} to 15.2×10^{12} M$_\odot$ for the sparser LLQ samples – for both samples these mean masses increase smoothly with σ_8, as expected. Within each sample, more than 99 per cent of all haloes lie within a factor of 2 in the mass of the threshold for inclusion. This is a consequence of the exponentially falling high-mass tail of the halo mass function.

The magnitude of the shifts between the three cosmologies is well described by the Jenkins et al. (2001) and Angulo et al. (2012) fitting formulae, which we display as solid or dashed lines in Fig. 2. However, at this redshift these formulae overpredict the number of haloes of a given mass by a factor of 2–3. In part, this reflects the fact that these formulae were calibrated using simulations and redshifts where the most massive haloes were much less extreme than those considered here, but the disagreement might also be in part a consequence of the non-universal behaviour of the halo mass function (see e.g. Tinker et al. 2008).

We identify the $z = 0$ descendant of each halo in our samples as the object that contains the majority of its 15 most bound particles. Although in principle it is possible that a different structure contains most of the mass of our high-z haloes, following the innermost particles should represent well the fate of a hypothetical black hole lying at the centre of the halo. The bottom panel of Fig. 2 displays the FOF mass distributions of these descendant haloes. For comparison, we also present predictions made using the analytic merger tree algorithm of Parkinson et al. (2008) which is based on the extended Press–Schechter (EPS) formalism and calibrated using merger trees extracted from the MS.
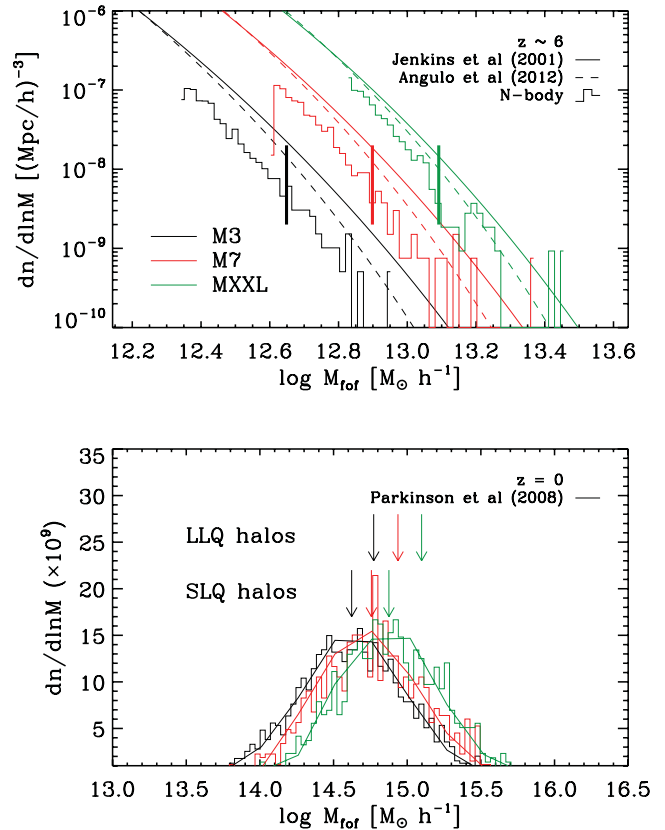
Figure 2. Comparison of the FOF mass distributions of our SLQ host halo samples at $z \sim 6$ to those of their descendants at $z = 0$. The top panel shows the differential number density of haloes for the three cosmologies of Table 1, corresponding to the parameters preferred by analyses of one, three and seven years of data from the *WMAP* satellite. On each histogram a thick vertical line indicates the mass threshold for the corresponding LLQ host halo sample. The bottom panel shows the distributions of FOF masses of the $z = 0$ descendants of these haloes. Predictions of the Jenkins et al. (2001) and Angulo et al. (2012) fitting formulae, and of merger trees constructed using the Parkinson, Cole & Helly (2008) algorithm are shown in the top and bottom panels, respectively. The vertical arrows in the bottom panel indicate the median masses of the descendants of SLQ haloes ($n = 30$ Gpc^{-3}) and LLQ haloes ($n = 1$ Gpc^{-3}). Note that haloes of similar masses at $z \sim 6$ end up in haloes with a wide range of masses at $z = 0$.

The median FOF mass of the descendants of our SLQ haloes varies by a factor of just 1.7 across our three cosmologies (see the lower set of coloured arrows in Fig. 2). This is significantly smaller than the spread of a factor of 3.1 in the initial median masses. The typical descendant mass is $M \sim 5.6 \times 10^{14}$ M$_\odot$, corresponding to a moderately rich $z = 0$ cluster. The median masses of the descendants of the sparser LLQ haloes are larger by a factor of about 1.5 and also vary slightly more with σ_8 (see the upper set of coloured arrows in Fig. 2). The median initial masses of the LLQ haloes are typically a factor of 1.9 larger than those of the SLQ haloes, so both within a single cosmology and between cosmologies the evolution is convergent in the sense that descendant masses are more similar than those of the original $z \sim 6$ objects. This is probably a result of QSO haloes, in all cases, descending into much less extreme peaks for which the differences among cosmologies are smaller. For example, the mass function for the M1 and M7 cosmologies is almost identical at $z = 0$ for masses below $M \sim 10^{14}$ M$_\odot$.

In all cases, there is a large spread in the masses of the descendants, as shown in the bottom panel of Fig. 2. This distribution at $z = 0$ can be well described by a log normal with $\sigma = 0.36$ dex. The most massive tail indeed corresponds to rare, high-mass clusters, but the least massive one corresponds to galaxy groups. The scatter is slightly smaller for the LLQ sample. We would also like to note that every descendant of an LLQ halo is the dominant structure of its $z = 0$ group, but in the M1 cosmology 39 out of the 810 descendants of SLQ haloes are satellite subhaloes. For the M3 and M7 cases the corresponding numbers are similar. As we will see in the next subsection, all this reflects the variety of formation histories among dark matter haloes of similar masses.

Our distributions are in good qualitative agreement with the EPS-based calculations of Trenti et al. (2008). Using the same M1 cosmology, these authors found 68 per cent of the descendants of a sample analogous to our LLQ to have masses in the range 2.5×10^{14}–$12.2 \times 10^{14}\,\mathrm{M}_\odot$ with a median of $5.6 \times 10^{14}\,\mathrm{M}_\odot$. For this cosmology, our results show a similar scatter, but with a median which is 35 per cent larger. As can be seen in Fig. 2, similar scatter and median masses are predicted by the merger tree algorithm of Parkinson et al. (2008) for all cosmologies, if we use a halo sample that matches the number density of the SLQ haloes. This confirms that our scaling algorithm captures the main features of structure growth in different background cosmologies.

3.2 The journey to $z = 0$

We now look in more detail at the paths connecting our samples of $z \sim 6$ quasar host haloes to their present-day descendants, concentrating on results in the unscaled MXXL, i.e. in the original M1 cosmology. In Fig. 3 we plot mass growth, defined as the ratio of descendant mass at each redshift to initial mass at $z \sim 6$. The light blue and brown regions indicate 86 per cent of the trajectories for SLQ and LLQ samples, respectively. Darker regions of each colour outline the regions containing 68 per cent of the trajectories. Growth histories seem to be remarkably similar in the two samples and to be faster than exponential, described approximately by $\log_{10} M(z)/M(z = 6.19) = 0.21 \times (6.19 - z)^{1.2}$. This seems to be independent of initial mass at $z = 6$, at least over the relatively restricted range of (high) masses considered here. Of course, this formula describes only the typical behaviour, and very different growth histories occur for haloes of similar initial masses. Some massive $z = 6$ haloes have grown only by a factor of 20–30 by $z = 0$ – much less than the change in the characteristic non-linear mass-scale M_* over the same redshift range – while others have increased their mass by factors of 200–300.

The spread in these trajectories increases substantially with decreasing redshift. At $z = 3$ the *rms* scatter in the log of the fractional growth is 0.168, at $z = 1$ it is 0.271 and at $z = 0$ it is 0.33, slightly larger than the initial separation in the median mass between our SLQ and LLQ samples. Although we do not display them, the mass accretion histories of haloes in other cosmologies are very similar. This is, of course, expected.

Fig. 3 also shows the mass growth for the most massive (dotted) and least massive (dashed) haloes in our SLQ sample. By chance, the least massive halo is one of the fastest growing, with a fractional growth rate faster than that of the most massive halo at all redshifts. Indeed, this 'low-mass' halo grows into a $2 \times 10^{15}\,\mathrm{M}_\odot$ object by $z = 0$, whereas the present-day descendant of the initially most massive halo is actually slightly smaller ($1.8 \times 10^{15}\,\mathrm{M}_\odot$). Neither of these haloes is in the extreme tail at $z = 0$, ranking 2224th and 3459th in mass among MXXL haloes and being four times less
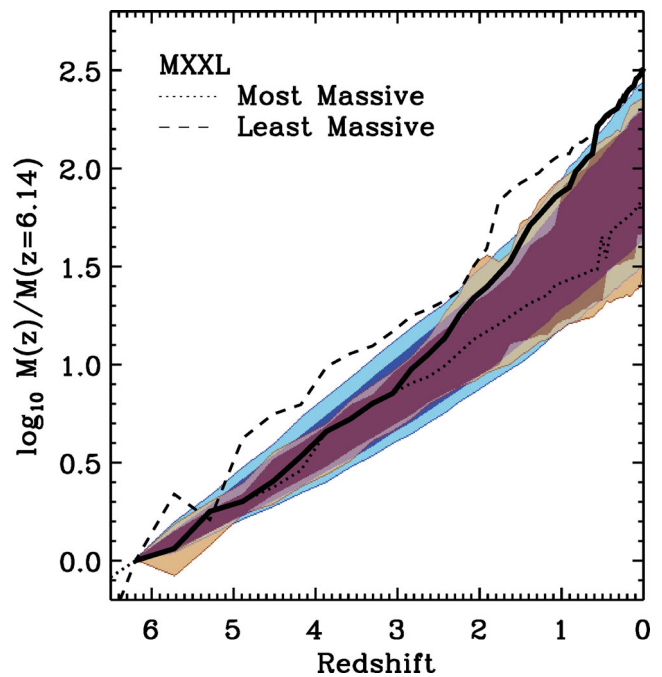
Figure 3. Mass growth of high-redshift quasar host haloes in the M1 cosmology (i.e. in the original MXXL). The blue regions show mass growth for the SLQ sample, whereas the brown regions show growth for the sparser and more massive LLQ sample. In each case, the dark and light regions enclose 68 and 86 per cent of the trajectories. Individual lines highlight particular trajectories. The thick solid line is the growth history that the largest halo at $z = 6$ would have if its mass equalled that of the most massive halo in the simulation at all later redshifts. In contrast, the dotted line shows the actual growth history of this halo. Finally, the dashed line displays the growth of the lowest mass halo in the SLQ sample which, by chance, is one of the fastest growing of all haloes.

massive than the most extreme object. Curiously, none of the 20 most massive $z = 0$ haloes in the MXXL has a progenitor in the SLQ sample.

Finally, the thick solid line in Fig. 3 indicates the mass growth that the most massive halo at $z = 6$ would have to have in order to be the most massive halo at all redshifts. This is close to the actual trajectory of the largest halo until $z \sim 3$, but at later times, other objects take over the top spots.

All these examples are not exclusive to extremely massive haloes at high redshifts, but are a generic illustration that the most massive haloes at any given epoch will no longer be the most massive haloes at a later time. This is a consequence of the diverse assembly histories of haloes in a hierarchical universe.

We explore this behaviour directly in Fig. 4. The top panel indicates the fraction of SLQ halo descendants whose masses place them above three abundance thresholds. The bottom panel shows a complementary picture, showing the fraction of SLQ haloes that rise above various mass thresholds at later times.

At redshift 3, all descendants of SLQ haloes have masses above $2 \times 10^{13}\,\mathrm{M}_\odot$ and rank among the 3000 most massive haloes per cubic gigaparsec. (Recall that initially these objects are *defined* as the high-mass tail of the halo distribution with an abundance of $30\,\mathrm{Gpc}^{-3}$.) In contrast, only 5 per cent are in haloes with $M > 1.4 \times 10^{14}\,\mathrm{M}_\odot$ and only 20 per cent still rank among the 30 most massive haloes per cubic gigaparsec. With time, the SLQ descendants fall further and further behind the most massive haloes in the MXXL. By the present day, only $2-5$ per cent (depending on cosmology) would
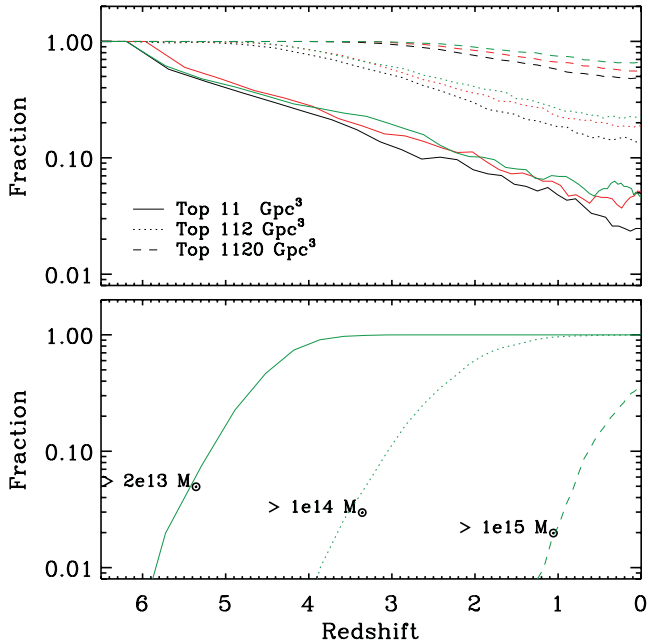
Figure 4. Fraction of descendants of SLQ haloes that are among the 11 (solid), 112 (dotted) or 1120 (dashed) most massive haloes per cubic giga-parsec at each later time (top panel), or that have a mass above 2.8×10^{13}, 1.4×10^{14} or 1.4×10^{15} M$_\odot$ (bottom panel). In the top panel, green, red and black lines correspond to the results for the M1, M7 and M3 cosmologies.

still be included in a mass-limited sample with $n = 11.6 \, \text{Gpc}^{-3}$, and about 50–70 per cent would be included in a sample with 100 times greater abundance.

3.3 The environment of QSO haloes

Why do some high-redshift haloes keep growing rapidly until the present, whereas others appear to shut down their accretion? Is this a random process or can it be related to some halo property at $z \sim 6$?

Fig. 5 shows how the masses of the $z = 0$ descendants of our SLQ and LLQ haloes correlate with their environment densities. The overdensity surrounding each high-redshift halo is computed on a variety of scales by mapping the underlying dark matter distribution on to a 2048^3 grid and then convolving this density field with a Gaussian filter of different sizes.

Clearly, most of the QSO haloes live in overdense regions, but there is considerable diversity in how extreme these regions are. On small scales, all live in at least a 3σ region, but 7σ is the typical overdensity of an SLQ halo and 10σ that of an LLQ halo. With increasing smoothing, our quasar host haloes are found in progressively less extreme regions and the separation between the SLQ and LLQ haloes decreases. For a smoothing of 28 Mpc the typical SLQ halo lives in a region of overdensity 1.4σ, and 10 per cent are located in regions of below-average density. Thus, our results suggest that it might be possible to find a quasar in the middle of a $14-28$ Mpc underdense region, even if quasars reside in the most massive haloes at $z \sim 6$.

Fig. 5 also shows that there is clearly a strong correlation between the overdensity around a high-z halo and the mass of its $z = 0$ descendant. In fact, this correlation is stronger than that between the mass of the $z = 6$ progenitor and that of its $z = 0$ descendant (or with any other property in our catalogues). For example, if we consider the environment at 14 Mpc, haloes that live in $<1\sigma$
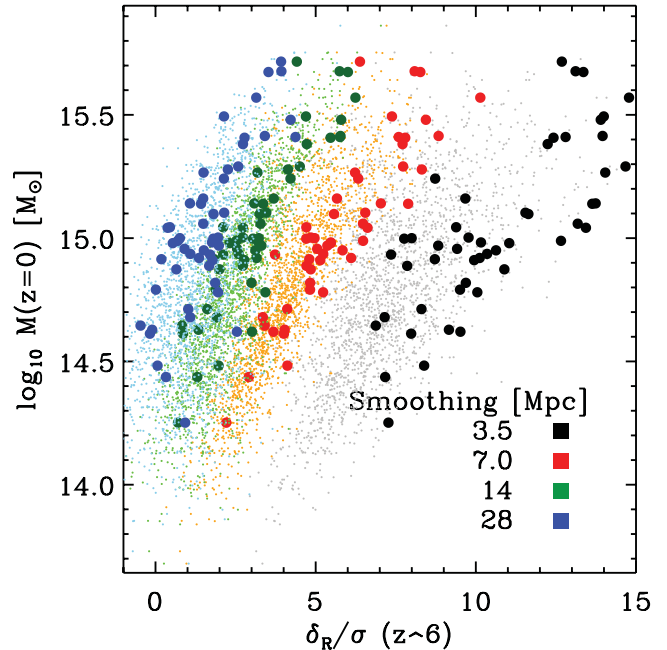
Figure 5. Correlation between the overdensity in which a QSO halo sits at $z = 6$ and the mass of its descendant at $z = 0$. The horizontal axis shows Gaussian-smoothed overdensity in units of its *rms* value $\sigma_{3.5} = 0.193$, $\sigma_7 = 0.125$, $\sigma_{14} = 0.071$ and $\sigma_{28} = 0.036$, where the suffix indicates the 1D smoothing radius of the Gaussian in units of Mpc comoving. Different colours refer to different smoothing scales as indicated in the figure. The small and large points show the overdensities of the environments of SLQ and LLQ haloes, respectively.

regions end up in haloes of $M \sim 2.8 \times 10^{14}$ M$_\odot$, whereas those found in $>4\sigma$ regions typically end up in 10 times more massive haloes (2.8×10^{15} M$_\odot$). The scatter in descendant mass at a fixed overdensity is $\sigma_{\log M} = [0.061, 0.044, 0.045, 0.063]$ for the fields smoothed on $[3.5, 7, 14, 28]$ Mpc scales, respectively. These figures are to be compared with a scatter of 0.36 for the sample as a whole. Thus, if the environment density surrounding a quasar is known, then the mass of its $z = 0$ descendant can be predicted six to eight times more accurately than if it is not.

These results are easy to understand, since the masses of the $z = 0$ descendants correspond to the mass contained within a sphere of radius about $7-14$ Mpc, and an object can only form by $z = 0$ if its material is already overdense by a factor of about $1.68/D_+(z=6)^2 = 0.32$ at $z = 6$. Our findings also warn against a naive connection between objects at different redshifts – the linkage depends not only on the actual properties of the objects, but also on their environment.

3.4 Prior accretion histories

Our $z \sim 6$ quasar host haloes are the most massive objects present at that time and so, by definition, are the objects which had the highest mean mass growth rates averaged over previous epochs. It is these extreme growth rates which must supply the baryons needed to build up the supermassive black holes and to fuel the extraordinarily luminous quasars which assume to lie in their cores. In this final section, we examine when and how fast our LLQ and SLQ haloes achieve their extreme masses.

---

[2] $D_+(z)$ is the growth factor in units of its present-day value.
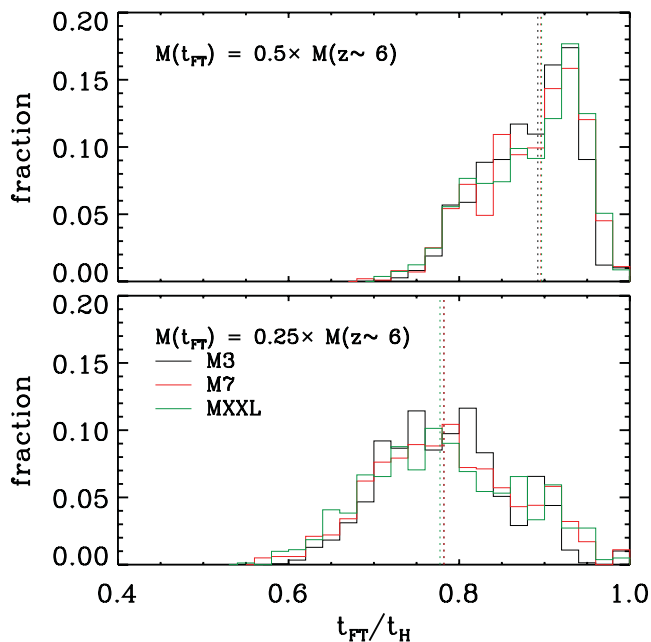
Figure 6. The distribution of formation times for haloes in our SLQ sample. We define the formation time of a halo as the epoch when the main progenitor of a halo had 50 or 25 per cent of the final halo mass, which we display in the top and bottom panels, respectively. These values are given in terms of the age of the Universe at the time the SQL sample is identified and correspond to 0.92, 0.93 and 0.94 Gyr for the M1, M3 and M7 cosmologies, respectively. The dotted lines in both panels indicate the median values for each sample.

We define an accretion history for each halo in our samples from the growth in mass along the main branch of its assembly tree. This branch is defined by stepping back in time from the $z \sim 6$ object, selecting the most massive progenitor of the current main branch object as the main branch object at the immediately preceding step. Note that with this algorithm the main progenitor at $z = 10$, say, is not necessarily the most massive of all the $z = 10$ progenitors. In fact, only about ~ 40 per cent of our SLQ haloes have an identifiable main progenitor at $z \sim 10$ (i.e. with a mass above the resolution limit, 20 particles or 1.2×10^{11} M$_\odot$) even though the MXXL contains 125 000 identifiable haloes at this time. Conversely, of the 100 most massive MXXL haloes at $z = 10$, only 30 have a descendant among our SLQ sample at $z \sim 6$.

A consequence of these statistics is that SLQ haloes typically accrete most of their mass in a relatively short period before they are identified. This is illustrated explicitly in Fig. 6, which shows histograms of the 25 and 50 per cent growth times of haloes in each of our three SLQ samples. These are defined for each halo as the times since it had a quarter and a half of its final mass, and are given in units of the age of the universe, t_H, at the time the samples were defined. Median values are $t_{FT}/t_H = 0.89$ and 0.78, respectively, for our two definitions of formation time, and they are almost independent of the cosmological model. These values correspond to roughly 100 and 200 Myr prior $z \sim 6$.

Recent accretion rates are clearly substantially larger than the mean value required to grow the halo in the Hubble time. Median accretion rates for the last half of the halo growth are $[9.7, 8.9, 7.6] \times 10^3$ M$_\odot$ yr^{-1} for the M1, M7 and M3 samples, respectively. The growth times of our LLQ haloes are similarly distributed to those shown in Fig. 6 and, as a result, their median accretion rates are two to three times higher. In all three cosmologies, the growth

rates we find appear to be comfortably large enough to fuel even quasars as bright as the SDSS objects at $z \sim 6$, provided, of course, that the associated baryons are able to shed most of their angular momentum and reach the central regions despite the tremendous luminosity being generated there.

4 CONCLUSIONS

In this paper we have combined the largest high-resolution cosmological simulation to date with a scaling technique which allows a simulation to represent structure growth in cosmologies other than that in which it was originally carried out. This allows us to explore the properties and the evolution of extremely massive haloes that might host $z \sim 6$ quasars in three cosmologies with parameters spanning the observationally allowed range.

We found significant differences in the growth of such haloes subsequent to their identification at $z \sim 6$. Some increase their masses by a factor of 200 by $z = 0$, whereas others grow only by a factor of 10. As a result, the descendants of bright high-redshift quasars are inferred to live in haloes with a wide range of halo masses today. The median descendant mass of haloes in a mass-limited sample with space density 11.6 Gpc^{-3} at $z \sim 6$ is 5.7×10^{14} M$_\odot$ for a *WMAP*7-like cosmology, while more massive objects with 30 times lower abundance, thus matching the directly observed number density of luminous SDSS quasars, end up in haloes with median mass about a factor of 2 higher. In both samples, descendants spread in mass by a factor of several above and below this median. Conversely, in the same cosmology, only 4 per cent of present-day haloes with mass above 2.8×10^{15} M$_\odot$, corresponding to a space density of 11.6 Gpc^{-3}, have a progenitor at $z \sim 6$ with mass above 7.1×10^{12} M$_\odot$ and so would be considered a potential quasar host at the same abundance. These figures change only slightly for the other two cosmologies we consider.

Another aspect of the same effect is that haloes ranked among the most massive at a given time will gradually occupy lower positions and other haloes, initially less massive, will take over the top positions. The dissimilar mass growth is also expected to influence the galaxies that would form in these haloes; two haloes of the same mass may thus host galaxies with different properties (Zhu et al. 2006). Since the large-scale clustering of haloes of given masses depends on assembly history (e.g. Gao, Springel & White 2005), this violates the core assumption of HOD modelling and simple abundance matching techniques, such that the galaxy population in a halo depends only on its mass and not on its large-scale environment.

We find that the best way to predict the later growth of $z = 6$ haloes is to look at their local environment on 14 Mpc scales, which correlates with the halo mass at $z = 0$ much more strongly than the actual mass at $z \sim 6$. We emphasize that the behaviour we described in the paper is not restricted to $z \sim 6$ haloes, but it is a general feature expected in hierarchical growth, where the initial amplitudes of different Fourier modes are independent of each other. This behaviour is an example of a general property of certain mathematical distributions known as 'regression to the mean', which describes the migration of an extreme sample to a less extreme one at a later time.[3] The 'regression to the mean' can only be avoided if there were a perfect and monotonically increasing relation between the mass of a halo and that of its descendant. In this

---

[3] This phenomenon was first pointed out by Sir Francis Galton, a cousin of Charles Darwin, in the 19th century.

case, the rank order of haloes by mass is perfectly preserved. Thus, extreme haloes at, for instance, $z = 6$ beget equally extreme haloes at $z = 0$. However, we have confirmed earlier results (De Lucia & Blaizot 2007; Trenti et al. 2008; Overzier et al. 2009) which showed that this situation does not apply to hierarchical structure formation from Gaussian initial conditions.

Finally, we explored the assembly of QSO haloes prior to their identification at $z \sim 6$, finding one of the fastest accretion rates ever seen in simulated objects: a median of about $8.6 \times 10^3 \, \mathrm{M_\odot \, yr^{-1}}$ (but up to a factor of 2 larger) for the 100 Myr preceding identification at $z \sim 6$, almost independent of the cosmological model. This appears sufficient to fuel the bright quasars observed in the SDSS.

REFERENCES

Angulo R. E., White S. D. M., 2010a, MNRAS, 401, 1796
Angulo R. E., White S. D. M., 2010b, MNRAS, 405, 143
Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, preprint (arXiv:1203.3216)
Baugh C. M., Gaztanaga E., Efstathiou G., 1995, MNRAS, 274, 1049
Bonoli S., Marulli F., Springel V., White S. D. M., Branchini E., Moscardini L., 2009, MNRAS, 396, 423
Bonoli S., Shankar F., White S. D. M., Springel V., Wyithe J. S. B., 2010, MNRAS, 404, 399
Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, MNRAS, 398, 1150
Crocce M., Pueblas S., Scoccimarro R., 2006, MNRAS, 373, 369
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
De Lucia G., Blaizot J., 2007, MNRAS, 375, 2

Di Matteo T., Colberg J., Springel V., Hernquist L., Sijacki D., 2008, ApJ, 676, 33
Fan X., 2006, New Astron. Rev., 50, 665
Fan X. et al., 2003, AJ, 125, 1649
Fan X. et al., 2006, AJ, 131, 1203
Fanidakis N., Baugh C. M., Benson A. J., Bower R. G., Cole S., Done C., Frenk C. S., 2011, MNRAS, 410, 53
Fanidakis N. et al., 2012, MNRAS, 419, 2797
Gao L., Springel V., White S. D. M., 2005, MNRAS, 363, L66
Jenkins A., 2010, MNRAS, 403, 1859
Jenkins A., Frenk C. S., White S. D. M., Colberg J. M., Cole S., Evrard A. E., Couchman H. M. P., Yoshida N., 2001, MNRAS, 321, 372
Knebe A., Wagner C., Knollmann S., Diekershoff T., Krause F., 2009, ApJ, 698, 266
Kollmeier J. A. et al., 2006, ApJ, 648, 128
Komatsu E. et al., 2011, ApJS, 192, 18
Li Y. et al., 2007, ApJ, 665, 187
Marulli F., Bonoli S., Branchini E., Moscardini L., Springel V., 2008, MNRAS, 385, 1846
Overzier R. A., Guo Q., Kauffmann G., De Lucia G., Bouwens R., Lemson G., 2009, MNRAS, 394, 577
Parkinson H., Cole S., Helly J., 2008, MNRAS, 383, 557
Press W. H., Schechter P., 1974, ApJ, 187, 425
Romano-Diaz E., Shlosman I., Trenti M., Hoffman Y., 2011, ApJ, 736, 66
Ruiz A. N., Padilla N. D., Domínguez M. J., Cora S. A., 2011, MNRAS, 418, 2422
Scoccimarro R., 1998, MNRAS, 299, 1097
Shen Y. et al., 2007, AJ, 133, 2222
Shen Y., Greene J. E., Strauss M. A., Richards G. T., Schneider D. P., 2008, ApJ, 680, 169
Sijacki D., Springel V., Haehnelt M. G., 2009, MNRAS, 400, 100
Spergel D. N. et al., 2007, ApJS, 170, 377
Springel V., 2005, MNRAS, 364, 1105
Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726
Springel V. et al., 2005, Nat, 435, 629
Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, ApJ, 688, 709
Trenti M., Santos M. R., Stiavelli M., 2008, ApJ, 687, 1
Trenti M., Smith B. D., Hallman E. J., Skillman S. W., Shull J. M., 2010, ApJ, 711, 1198
White S. D. M., 1996, in Schaefer R., Silk J., Spiro M., Zinn-Justin J., eds, Cosmology and Large-Scale Structure. Elsevier, Dordrecht
White M., Martini P., Cohn J. D., 2008, MNRAS, 390, 1179
Zel'Dovich Y. B., 1970, A&A, 5, 84
Zhu G., Zheng Z., Lin W. P., Jing Y. P., Kang X., Gao L., 2006, ApJ, 639, L5

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