

Simulating the joint evolution of quasars, galaxies and their large-scale distribution

Supplementary Information

V. Springel¹, S. D. M. White¹, A. Jenkins², C. S. Frenk², N. Yoshida³, L. Gao¹, J. Navarro⁴, R. Thacker⁵, D. Croton¹, J. Helly², J. A. Peacock⁶, S. Cole², P. Thomas⁷, H. Couchman⁵, A. Evrard⁸, J. Colberg⁹ & F. Pearce¹⁰

This document provides supplementary information for the above article in Nature. We detail the physical model used to compute the galaxy population, and give a short summary of our simulation method. Where appropriate, we give further references to relevant literature for our methodology.

Characteristics of the simulation

Numerical simulations are a primary theoretical tool to study the nonlinear gravitational growth of structure in the Universe, and to link the initial conditions of cold dark matter (CDM) cosmogonies to observations of galaxies at the present day. Without direct numerical simulation, the hierarchical build-up of structure with its three-dimensional dynamics would be largely inaccessible.

Since the dominant mass component, the dark matter, is assumed to consist of weakly interacting elementary particles that interact only gravitationally, such simulations use a set of discrete point particles to represent the collisionless dark matter fluid. This representation as an N-body system is obviously only a coarse approximation, and im-

proving its fidelity requires the use of as many particles as possible while remaining computationally tractable. Cosmological simulations have therefore always striven to increase the size (and hence resolution) of N-body computations, taking advantage of every advance in numerical algorithms and computer hardware. As a result, the size of simulations has grown continually over the last four decades. Fig. 1 shows the progress since 1970. The number of particles has increased exponentially, doubling roughly every 16.5 months. Interestingly, this growth parallels the empirical ‘Moore’s Law’ used to describe the growth of computer performance in general. Our new simulation discussed in this paper uses an unprecedentedly large number of 2160^3 particles, more than 10^{10} . We were able to finish this computation in 2004, significantly ahead of a simple extrapolation of the past growth rate of simulation sizes. The simulation represented a substantial computational challenge that required novel approaches both for the simulation itself, as well as for its analysis. We describe the most important of these aspects in the following. As an aside, we note that extrapolating the remarkable progress since the 1970s for another three decades, we may expect cosmological simulations with $\sim 10^{20}$ particles some time around 2035. This would be sufficient to represent all stars in a region as large as the Millennium volume with individual particles.

Initial conditions. We used the Boltzmann code CMBFAST²⁴ to compute a linear theory power spectrum of a Λ CDM model with cosmological parameters consistent with recent constraints from WMAP and large-scale structure data^{25,26}. We then constructed a random realization of the model in Fourier space, sampling modes in a sphere up to the Nyquist frequency of our 2160^3 particle load. Mode amplitudes $|\delta_{\mathbf{k}}|$ were determined by random sampling from a Rayleigh distribution with second moment equal to $P(k) = \langle |\delta_{\mathbf{k}}|^2 \rangle$, while phases were chosen randomly. A high quality random number generator with period $\sim 10^{171}$ was used for this purpose. We em-

¹Max-Planck-Institute for Astrophysics, Karl-Schwarzschild-Str. 1, 85740 Garching, Germany

²Institute for Computational Cosmology, Dep. of Physics, Univ. of Durham, South Road, Durham DH1 3LE, UK

³Department of Physics, Nagoya University, Chikusa-ku, Nagoya 464-8602, Japan

⁴Dep. of Physics & Astron., University of Victoria, Victoria, BC, V8P 5C2, Canada

⁵Dep. of Physics & Astron., McMaster Univ., 1280 Main St. West, Hamilton, Ontario, L8S 4M1, Canada

⁶Institute of Astronomy, University of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, UK

⁷Dep. of Physics & Astron., University of Sussex, Falmer, Brighton BN1 9QH, UK

⁸Dep. of Physics & Astron., Univ. of Michigan, Ann Arbor, MI 48109-1120, USA

⁹Dep. of Physics & Astron., Univ. of Pittsburgh, 3941 O’Hara Street, Pittsburgh PA 15260, USA

¹⁰Physics and Astronomy Department, Univ. of Nottingham, Nottingham NG7 2RD, UK

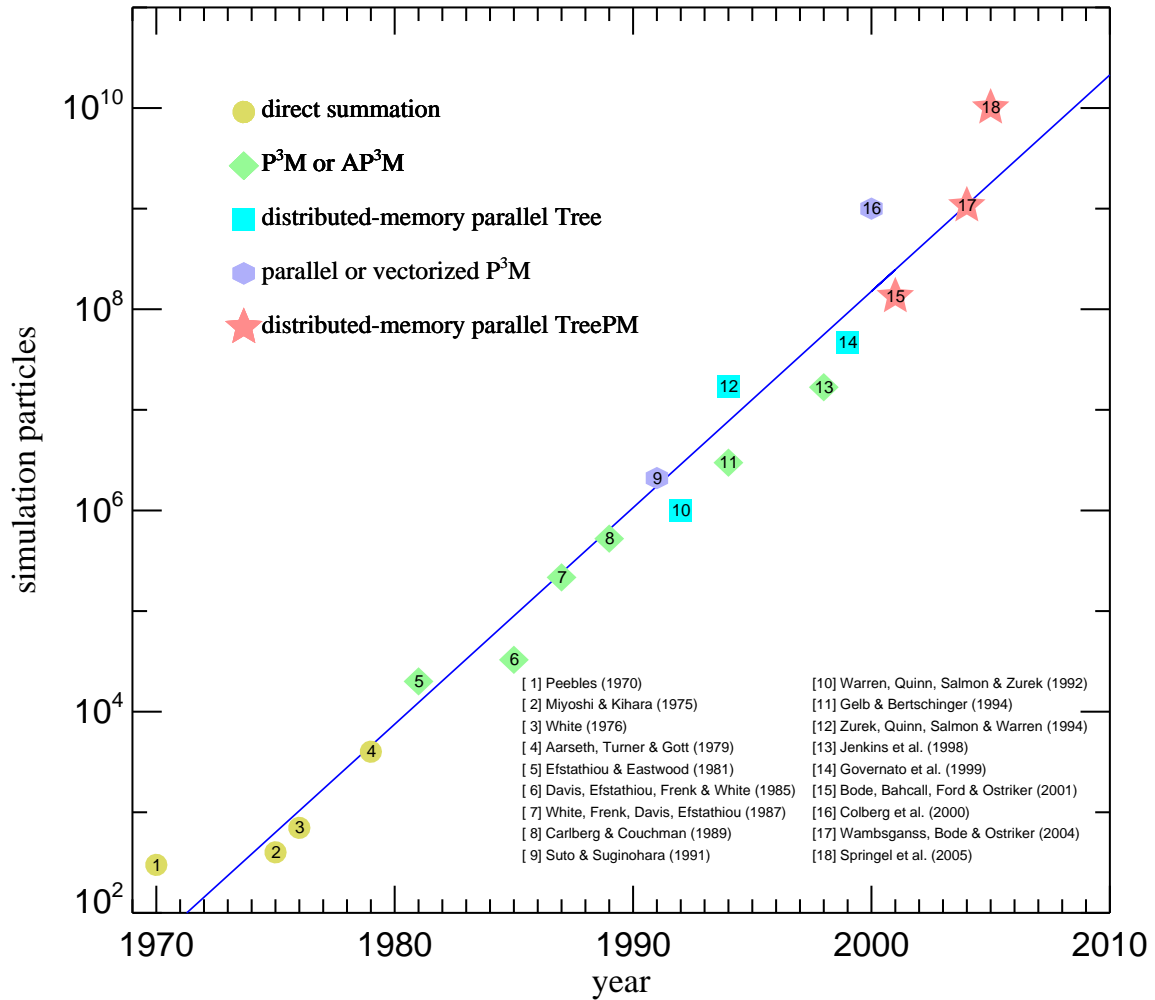


Figure 1: Particle number in high resolution N-body simulations of cosmic structure formation as a function of publication date¹⁻¹⁷. Over the last three decades, the growth in simulation size has been exponential, doubling approximately every ~ 16.5 months (blue line). Different symbols are used for different classes of computational algorithms. The particle-mesh (PM) method combined with direct particle-particle (PP) summation on sub-grid scales has long provided the primary path towards higher resolution. However, due to their large dynamic range and flexibility, tree algorithms have recently become competitive with these traditional P³M schemes, particularly if combined with PM methods to calculate the long-range forces. Plain PM simulations¹⁸⁻²² have not been included in this overview because of their much lower spatial resolution for a given particle number. Note also that we focus on the largest simulations at a given time, so our selection of simulations does not represent a complete account of past work on cosmological simulations.




Figure 2: Different realizations of the initial power spectrum. The top and bottom panels show measured power-spectra for 20 realizations of initial conditions with different random number seeds, together with the mean spectrum (red symbols). The latter lies close to the input linear power spectrum (black solid line). In the bottom panel, the measurements have been divided by a smooth CDM-only power spectrum²³ to highlight the acoustic oscillations. One of the realizations has been drawn in blue; it shows a fluctuation pattern that superficially resembles the pattern around the second acoustic peak. However, this is a chance effect; the fluctuations of each bin are independent.

employed a massively parallel complex-to-real Fourier transform (which requires some care to satisfy all reality constraints) to directly obtain the resulting displacement field in each dimension. The initial displacement at a given particle coordinate of the unperturbed density field was obtained by tri-linear interpolation of the resulting displacement field, with the initial velocity obtained from the Zel’dovich approximation. The latter is very accurate for our starting redshift of $z = 127$. For the initial unperturbed

density field of 2160^3 particles we used a *glass-like* particle distribution. Such a glass is formed when a Poisson particle distribution in a periodic box is evolved with the sign of gravity reversed until residual forces have dropped to negligible levels²⁷. For reasons of efficiency, we replicated a 270^3 glass file 8 times in each dimension to generate the initial particle load. The Fast Fourier Transforms (FFT) required to compute the displacement fields were carried out on a 2560^3 mesh using 512 processors and a distributed-memory code. We deconvolved the input power spectrum for smoothing effects due to the interpolation off this grid.

We note that the initial random number seed was picked in an unconstrained fashion. Due to the finite number of modes on large scales and the Rayleigh-distribution of mode amplitudes, the mean power of the actual realization in each bin is expected to scatter around the linear input power spectrum. Also, while the expectation value $\langle |\delta_{\mathbf{k}}|^2 \rangle$ is equal to the input power spectrum, the median power per mode is biased low due to the skew-negative distribution of the mode amplitudes. Hence, in a given realization there are typically more points lying below the input power spectrum than above it, an effect that quickly becomes negligible as the number of independent modes in each bin becomes large. We illustrate this in the top panel of Figure 2, where 20 realizations for different random number seeds of the power spectrum on large scales are shown, together with the average power in each bin. Our particular realization for the Millennium Simulation corresponds to a slightly unlucky choice of random number seed in the sense that the fluctuations around the mean input power in the region of the second peak seem to resemble the pattern of the acoustic oscillations (see the top left panel of Figure 6 in our Nature article). However, we stress that the fluctuations in these bins are random and uncorrelated, and that this impression is only a chance effect. In the bottom panel of Figure 2, we redraw the measured power spectra for the 20 random realizations, this time normalised to a smooth CDM power spectrum without acoustic oscillations in order to highlight the baryonic ‘wiggles’. We have drawn one of the 20 realizations in blue. It is one that resembles the pattern of fluctuations seen in the Millennium realization quite closely while others scatter quite differently, showing that such deviations are consistent with the expected statistical distribution.

Dynamical evolution. The evolution of the simulation particles under gravity in an expanding background is governed by the Hamiltonian

$$H = \sum_i \frac{\mathbf{p}_i^2}{2m_i a(t)^2} + \frac{1}{2} \sum_{ij} \frac{m_i m_j \varphi(\mathbf{x}_i - \mathbf{x}_j)}{a(t)}, \quad (1)$$

where $H = H(\mathbf{p}_1, \dots, \mathbf{p}_N, \mathbf{x}_1, \dots, \mathbf{x}_N, t)$. The \mathbf{x}_i are comoving coordinate vectors, and the corresponding canonical momenta are given by $\mathbf{p}_i = a^2 m_i \dot{\mathbf{x}}_i$. The explicit time

dependence of the Hamiltonian arises from the evolution $a(t)$ of the scale factor, which is given by the Friedman-Lemaître model that describes the background cosmology. Due to our assumption of periodic boundary conditions for a cube of size L^3 , the interaction potential $\varphi(\mathbf{x})$ is the solution of

$$\nabla^2 \varphi(\mathbf{x}) = 4\pi G \left[-\frac{1}{L^3} + \sum_{\mathbf{n}} \delta_{\varepsilon}(\mathbf{x} - \mathbf{n}L) \right], \quad (2)$$

where the sum over $\mathbf{n} = (n_1, n_2, n_3)$ extends over all integer triplets. The density distribution function $\delta_{\varepsilon}(\mathbf{x})$ of a single particle is spread over a finite scale ε , the gravitational softening length. The softening is necessary to make it impossible for hard binaries to form and to allow the integration of close particle encounters with low-order integrators. We use a spline kernel to soften the point mass, given by $\delta_{\varepsilon}(\mathbf{x}) = W(|\mathbf{x}|/2.8\varepsilon)$, where $W(r) = 8(1 - 6r^2 + 6r^3)/\pi$ for $0 \leq r < 1/2$, $W(r) = 16(1 - r)^3/\pi$ for $1/2 \leq r < 1$, and $W(r) = 0$ otherwise. For this choice, the Newtonian potential of a point mass at zero lag in non-periodic space is $-Gm/\varepsilon$, the same as for a ‘Plummer-sphere’ of size ε , and the force becomes fully Newtonian for separations larger than 2.8ε . We took $\varepsilon = 5 h^{-1} \text{kpc}$, about 46.3 times smaller than the mean particle separation. Note that the mean density is subtracted in equation (2), so the solution of the Poisson equation corresponds to the *peculiar potential*, where the dynamics of the system is governed by $\nabla^2 \phi(\mathbf{x}) = 4\pi G[\rho(\mathbf{x}) - \bar{\rho}]$.

The equations of motion corresponding to equation (1) are $\sim 10^{10}$ simple differential equations, which are however coupled tightly by the mutual gravitational forces between the particles. An accurate evaluation of these forces (the ‘right hand side’ of the equations) is computationally very expensive, even when force errors up to $\sim 1\%$ can be tolerated, which is usually the case in collisionless dynamics²⁸. We have written a completely new version of the cosmological simulation code GADGET²⁹ for this purpose. Our principal computational technique for the gravitational force calculation is a variant of the ‘TreePM’ method^{30–32}, which uses a hierarchical multipole expansion³³ (a ‘tree’ algorithm) to compute short-range gravitational forces and combines this with a more traditional particle-mesh (PM) method³⁴ to determine long-range gravitational forces. This combination allows for a very large dynamic range and high computational speed even in situations where the clustering becomes strong. We use an explicit force-split³² in Fourier-space, which produces a highly isotropic force law and negligible force errors at the force matching scale. The algorithms in our code are specially designed for massively parallel operation and contain explicit communication instructions such that the code can work on computers with distributed physical memory, a prerequisite for a simulation of the size and computational cost of the Millennium Run.

For the tree-algorithm, we first decompose the simulation volume spatially into compact *domains*, each served by one processor. This domain decomposition is done by dividing a space filling Peano-Hilbert curve into segments. This fractal curve visits each cell of a fiducial grid of 1024^3 cells overlaid over the simulation exactly once. The decomposition tries to achieve a work-load balance for each processor, and evolves over time as clustering progresses. Using the Peano-Hilbert curve guarantees that domain boundaries are always parallel to natural tree-node boundaries, and thanks to its fractal nature provides for a small surface-to-volume ratio for all domains, such that communication with neighbouring processors during the short-range tree force computation can be minimised. Our tree is fully threaded (i.e. its leaves are single particles), and implements an oct-tree structure with monopole moments only. The cell-opening criterion was relative³⁵; a multipole approximation was accepted if its conservatively estimated error was below 0.5% of the total force from the last timestep. In addition, nodes were always opened when the particle under consideration lay inside a 10% enlarged outer node boundary. This procedure gives forces with typical errors well below 0.1%.

For the PM algorithm, we use a parallel Fast Fourier Transform (FFT)* to solve Poisson’s equation. We used a FFT mesh with 2560^3 cells, distributed into 512 slabs of dimension $5 \times 2560 \times 2560$ for the parallel transforms. After clouds-in-cells (CIC) mass assignment to construct a density field, we invoke a real-to-complex transform to convert to Fourier space. We then multiplied by the Greens function of the Poisson equation, deconvolved for the effects of the CIC and the trilinear interpolation that is needed later, and applied the short-range filtering factor used in our TreePM formulation (the short range forces suppressed here are exactly those supplied by the tree-algorithm). Upon transforming back we obtained the gravitational potential. We then applied a four-point finite differencing formula to compute the gravitational force field for each of the three coordinate directions. Finally, the forces at each particle’s coordinate were obtained by trilinear interpolation from these fields.

A particular challenge arises due to the different data layouts needed for the PM and tree algorithms. In order to keep the required communication and memory overhead low, we do not swap the particle data between the domain and slab decompositions. Instead, the particles stay in the domain decomposition needed by the tree, and each processor constructs patches of the density field for all the slabs on other processors which overlap its local domain. In this way, each processor communicates only with a small number of other processors to establish the binned density field on the slabs. Likewise, the slab-decomposed potential field is transferred back to processors so that a local region is formed covering the local domain, in addition

*Based on the www.fftw.org libraries of MIT.




Figure 3: The power spectrum of the dark matter distribution in the Millennium Simulation at various epochs (blue lines). The gray lines show the power spectrum predicted for linear growth, while the dashed line denotes the shot-noise limit expected if the simulation particles are a Poisson sampling from a smooth underlying density field. In practice, the sampling is significantly sub-Poisson at early times and in low density regions, but approaches the Poisson limit in nonlinear structures. Shot-noise subtraction allows us to probe the spectrum slightly beyond the Poisson limit. Fluctuations around the linear input spectrum on the largest scales are due to the small number of modes sampled at these wavelengths and the Rayleigh distribution of individual mode amplitudes assumed in setting up the initial conditions. To indicate the bin sizes and expected sample variance on these large scales, we have included symbols and error bars in the $z = 0$ estimates. On smaller scales, the statistical error bars are negligibly small.

to a few ghost cells around it, such that the finite differencing of the potential can be carried out for all interior points.

Timestepping was achieved with a symplectic leap-frog scheme based on a split of the potential energy into a short-range and long-range component. The short-range dynamics was then integrated by subcycling the long-range step³⁶. Hence, while the short-range force had to be computed frequently, the long-range FFT force was needed only comparatively infrequently. More than 11000 timesteps in total were carried out for the simulation, using individual and adaptive timesteps[†] for the particles. A timestep of a particle was restricted to be smaller than $\Delta t = \sqrt{2\eta\epsilon/|\mathbf{a}|}$, where \mathbf{a} is a particle’s acceleration and $\eta = 0.02$ controls the integration accuracy. We used a binary hierarchy of timesteps to generate a grouping of particles onto timebins.

The memory requirement of the code had to be aggressively optimised in order to make the simulation possible on the IBM p690 supercomputer available to us. The total aggregated memory on the 512 processors was 1 TB, of which about 950 GB could be used freely by an application program. In our code *Lean-GADGET-2* produced for the Millennium Simulation, we needed about 400 GB for particle storage and 300 GB for the fully threaded tree in the final clustered particle state, while the PM algorithm consumed in total about 450 GB in the final state (due to growing variations in the volume of domains as a result of our work-load balancing strategy, the PM memory requirements increase somewhat with time). Note that the memory for tree and PM computations is not needed concurrently, and this made the simulation feasible. The peak memory consumption per processor reached 1850 MB at the end of our simulation, rather close to the maximum possible of 1900 MB.

On the fly analysis. With a simulation of the size of the Millennium Run, any non-trivial analysis step is demanding. For example, measuring the dark matter mass power spectrum over the full dynamic range of the simulation volume would require a 3D FFT with $\sim 10^5$ cells per dimension, which is unfeasible at present. In order to circumvent this problem, we employed a two stage procedure for measuring the power spectrum where a “large-scale” and a “small-scale” measurement were combined. The former was computed with a Fourier transform of the whole simulation box, while the latter was constructed by folding the density field back onto itself¹³, assuming periodicity for a fraction of the box. The self-folding procedure leads to a sparser sampling of Fourier space on small scales, but since the number of modes there is large, an accurate small-scale measurement is still achieved. Since

[†] Allowing adaptive changes of timesteps formally breaks the symplectic nature of our integration scheme, which is however not a problem for the dynamics we follow here.




Figure 4: Measured distribution of mode amplitudes in the Millennium Simulation at redshift $z = 4.9$. Only modes in the k -range $0.03 h/\text{Mpc} < k < 0.07 h/\text{Mpc}$ are included (in total 341 modes), with their amplitude normalised to the square root of the expected linear power spectrum at that redshift. The distribution of modes follows the expected Rayleigh distribution very well. The bottom panel shows the relative deviations of the measurements from this distribution, which are in line with the expected statistical scatter.

the PM-step of the simulation code already computes an FFT of the whole density field, we took advantage of this and embedded a measurement of the power spectrum directly into the code. The self-folded spectrum was computed for a 32 times smaller periodic box-size, also using a 2560^3 mesh, so that the power spectrum measurement effectively corresponded to a 81920^3 mesh. We have carried out a measurement each time a simulation snapshot was generated and saved on disk. In Figure 3, we show the resulting time evolution of the *dark matter* power spectrum in the Millennium Simulation. On large scales and at early times, the mode amplitudes grow linearly, roughly in proportion to the cosmological expansion factor. Non-linear evolution accelerates the growth on small scales when the dimensionless power $\Delta^2(k) = k^3 P(k)/(2\pi^2)$ approaches unity; this regime can only be studied accurately using numerical simulations. In the Millennium Simulation, we are able to determine the nonlinear power spectrum over a larger range of scales than was possible in earlier work¹³, almost five orders of magnitude in wavenumber k .

On the largest scales, the periodic simulation volume encompasses only a relatively small number of modes and, as a result of the Rayleigh amplitude sampling that we used, these (linear) scales show substantial random fluctuations around the mean expected power.

This also explains why the mean power in the k -range $0.03 h/\text{Mpc} < k < 0.07 h/\text{Mpc}$ lies below the linear input power. In Figure 4, we show the actual distribution of normalised mode amplitudes, $\sqrt{|\delta_{\mathbf{k}}|^2/P(k)}$, measured directly for this range of wavevectors in the Millennium Simulation at redshift $z = 4.9$. We see that the distribution of mode amplitudes is perfectly consistent with the expected underlying Rayleigh distribution.

Useful complementary information about the clustering of matter in real space is provided by the two-point correlation function of dark matter particles. Measuring it involves, in principle, simply counting the number of particle pairs found in spherical shells around a random subset of all particles. Naive approaches to determine these counts involve an N^2 -scaling of the operation count and are prohibitive for our large simulation. We have therefore implemented novel parallel methods to measure the two-point function accurately, which we again embedded directly into the simulation code, generating a measurement automatically at every output. Our primary approach to speeding up the pair-count lies in using the hierarchical grouping provided by the tree to search for particles around a randomly selected particle. Since we use logarithmic radial bins for the pair counts, the volume corresponding to bins at large radii is substantial. We use the tree for finding neighbours with a range-searching technique. In carrying out the tree-walk, we check whether a node falls fully within the volume corresponding to a bin. In this case, we terminate the walk along this branch of the tree and simply count all the particles represented by the node at once, leading to a significant speed-up of the measurement.

Finally, the exceptionally large size of the simulation prompted us to develop new methods for computing friends-of-friends (FOF) group catalogues in parallel and on the fly. The FOF groups are defined as equivalence classes in which any pair of particles belongs to the same group if their separation is less than 0.2 of the mean particle separation. This criterion combines particles into groups with a mean overdensity that corresponds approximately to the expected density of virialised groups. Operationally, one can construct the groups by starting from a situation in which each particle is first in its own single group, and then testing all possible particle pairs; if a close enough pair is found whose particles lie in different groups already present, the groups are linked into a common group. Our algorithm represents groups as link-lists, with auxiliary pointers to a list's head, tail, and length. In this way we can make sure that, when groups are joined, the smaller of two groups is always attached to the tail of the larger one. Since each element of the attached group must be visited only once, this procedure avoids a quadratic contribution to the operation count proportional to the group size when large groups are built up. Our parallel algorithm works by first determining the FOF groups on local domains, again exploiting the tree

for range searching techniques, allowing us to find neighbouring particles quickly. Once this first step of group finding for each domain is finished, we merge groups that are split by the domain decomposition across two or several processors. As groups may in principle percolate across several processors, special care is required in this step as well. Finally, we save a group catalogue to disk at each output, keeping only groups with at least 20 particles.

In summary, the simulation code evolved the particle set for more than 11000 timesteps, producing 64 output time slices each of about 300 GB. Using parallel I/O techniques, each snapshot could be written to disk in about 300 seconds. Along with each particle snapshot, the simulation code produced a FOF group catalogue, a power spectrum measurement, and a two-point correlation function measurement. Together, over ~ 20 TB of data were generated by the simulation. The raw particle data of each output was stored in a special way (making use of a space-filling curve), which allows rapid direct access to subvolumes of the particle data. The granularity of these subvolumes corresponds to a fiducial 256^3 mesh overlaid over the simulation volume, such that the data can be accessed randomly in pieces of ~ 600 particles on average. This storage scheme is important to allow efficient post-processing, which cannot make use of an equally powerful supercomputer as the simulation itself.

Postprocessing of the simulation data

Substructure analysis. High-resolution simulations like the present one exhibit a rich substructure of gravitationally bound dark matter subhalos orbiting within larger virialised structures³⁷. The FOF group finder built into the simulation code is able to identify the latter, but not the ‘subhalos’. In order to follow the fate of infalling halos and galaxies more reliably, we therefore determine dark matter substructures for all identified FOF halos. We accomplish this with an improved and extended version of the SUBFIND algorithm³⁸. This computes an adaptively smoothed dark matter density field using a kernel-interpolation technique, and then exploits the topological connectivity of excursion sets above a density threshold to identify substructure candidates. Each substructure candidate is subjected to a gravitational unbinding procedure. If the remaining bound part has more than 20 particles, the subhalo is kept for further analysis and some basic physical properties (angular momentum, maximum of its rotation curve, velocity dispersion, etc.) are determined. An identified subhalo was extracted from the FOF halo, so that the remainder formed a featureless ‘background’ halo which was also subjected to an unbinding procedure. The required computation of the gravitational potential for the unbinding was carried out with a tree algorithm similar to the one used in the simulation code itself.

Finally, we also compute a virial mass estimate for

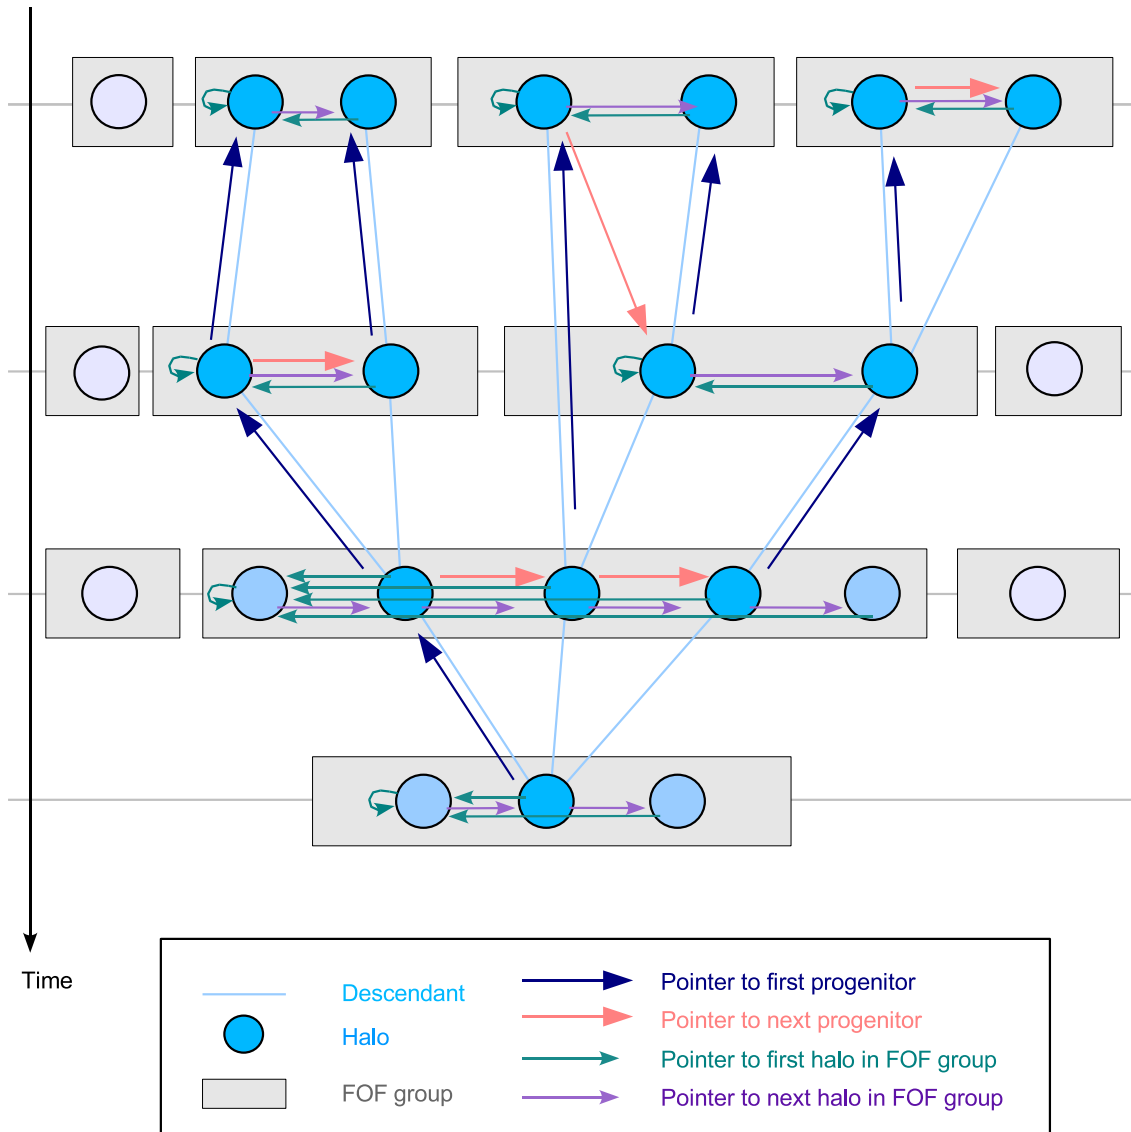


Figure 5: Schematic organisation of the merger tree in the Millennium Run. At each output time, FOF groups are identified which contain one or several (sub)halos. The merger tree connects these halos. The FOF groups play no direct role, except that the largest halo in a given FOF group is the one which may develop a cooling flow according to the physical model for galaxy formation implemented for the trees. To facilitate the latter, a number of pointers for each halo are defined. Each halo knows its descendant, and its most massive progenitor. Possible further progenitors can be retrieved by following the chain of ‘next progenitors’. In a similar fashion, all halos in a given FOF group are linked together.

each FOF halo in this analysis step, using the spherical-overdensity approach and the minimum of the gravitational potential within the group as the central point. We identified 17.7×10^6 FOF groups at $z = 0$, down from a maximum of 19.8×10^6 at $z = 1.4$, where the groups are more abundant yet smaller on average. At $z = 0$, we found a total of 18.2×10^6 subhalos, and the largest FOF group contained 2328 of them.

Merger tree definition and construction. Having determined all halos and subhalos at all output times, we tracked these structures over time, i.e. we determined the hierarchical merging trees that describe in detail how structures build up over cosmic time. These trees are the key information needed to compute physical models for the properties of the associated galaxy population.

Because structures merge hierarchically in CDM universes, a given halo can have several progenitors but, in general, it has only one descendant because the cores of virialised dark matter structures do not split up into two or more objects. We therefore based our merger tree construction on the determination of a unique descendant for any given halo. This is, in fact, already sufficient to define the merger tree construction, since the progenitor information then follows implicitly.

To determine the appropriate descendant, we use the unique IDs that label each particle and track them between outputs. For a given halo, we find all halos in the subsequent output that contain some of its particles. We count these particles in a weighted fashion, giving higher weight to particles that are more tightly bound in the halo under consideration. In this way, we give preference to tracking the fate of the inner parts of a structure, which may survive for a long time upon infall into a bigger halo, even though much of the mass in the outer parts can be quickly stripped. The weighting is facilitated by the fact that the results of the SUBFIND analysis are stored in order of increasing total binding energy, i.e. the most bound particle of a halo is always stored first. Once these weighted counts are determined for each potential descendant, we select the one with the highest count as the descendant. As an additional refinement (which is not important for any of our results), we have allowed some small halos to skip one snapshot in finding a descendant. This deals with cases where we would otherwise lose track of a structure that temporarily fluctuates below our detection threshold.

In Figure 5, we show a schematic representation of the merger tree constructed in this way. The FOF groups are represented at different times with boxes each of which contains one or more (sub)halos. For each halo, a unique descendant is known, and there are link-list structures that allow the retrieval of all progenitors of a halo, or of all other halos in the same FOF group. Not all the trees in the simulation volume are connected with each other. Instead, there are 14.4×10^6 separate trees, each essentially

describing the formation history of the galaxies contained in a FOF halo at the present time. The correspondence between trees and FOF halos is not exactly one-to-one because some small FOF halos did not contain a bound subhalo and were dropped, or because some FOF halos can be occasionally linked by feeble temporary particle bridges which then also combines their corresponding trees. We have stored the resulting tree data on a per-tree basis, so that the physical model for galaxy formation can be computed sequentially for all the trees individually, instead of having to apply the model in one single step. The latter would have been impossible, given that the trees contain a total of around 800 million halos.

Physical model for galaxy formation

‘Semi-analytic’ models of galaxy formation were first proposed more than a decade ago³⁹. They have proven to be a very powerful tool for advancing the theory of galaxy formation^{40–47}, even though much of the detailed physics of star formation and its regulation by feedback processes has remained poorly understood. The term ‘semi-analytic’ conveys the notion that while in this approach the physics is parameterised in terms of simple analytic models, following the dark matter merger trees over time can only be carried out numerically. Semi-analytic models are hence best viewed as simplified simulations of the galaxy formation process. While the early work employed Monte-Carlo realizations of dark matter trees^{48,49}, more recent work is able to measure the merging trees directly from numerical dark matter simulations⁵⁰. In the most sophisticated version of this technique, the approach is extended to include dark matter substructure information as well³⁸. This offers substantially improved tracking of the orbits of infalling substructure and of their lifetimes. In the Millennium Simulation, we have advanced this method further still, using improved substructure finding and tracking methods, allowing us fully to exploit the superior statistics and information content offered by the underlying high-resolution simulation.

Our semi-analytic model integrates a number of differential equations for the time evolution of the galaxies that populate each hierarchical merging tree. In brief, these equations describe radiative cooling of gas, star formation, the growth of supermassive black holes, feedback processes by supernovae and AGN, and effects due to a reionising UV background. Morphological transformation of galaxies and processes of metal enrichment are modelled as well. Full details of the scheme used to produce specific models shown in our Nature article will be provided in a forthcoming publication⁵¹, but we here include a very brief summary of the most important aspects of the model. Note, however, that this is just one model among many that can be implemented in post-processing on our stored Millennium Run data-structures. A prime goal of our project is to evaluate such schemes against

each other and against the observational data in order to understand which processes determine the various observational properties of the galaxy population.

Radiative cooling and star formation. We assume that each virialised dark matter halo contains (initially) a baryonic fraction equal to the universal fraction of baryons, $f_b = 0.17$, which is consistent with WMAP and Big-Bang nucleosynthesis constraints. A halo may lose some gas temporarily due to heating by a UV background or other feedback processes, but this gas is assumed to be reaccreted once the halo has grown in mass sufficiently. The influence of the UV background is directly taken into account as a reduction of the baryon fraction for small halos, following fitting functions obtained from detailed hydrodynamical models^{52,53}.

We distinguish between cold condensed gas in the centre of halos (forming the interstellar medium), and hot gas in the diffuse atmospheres of halos. The latter has a temperature equal to the virial temperature of the halo, and emits bremsstrahlung and line radiation. The corresponding cooling rate is estimated following standard parameterisations^{38,39}, which have been shown to provide accurate matches to direct hydrodynamical simulations of halo formation including radiative cooling^{54,55}. We note that, following the procedures established already in reference³⁹, the cooling model accounts for a distinction between a cold infall regime and cooling out of a hot atmosphere. The transition is at a mass scale close to that found in detailed analytic calculations of the cooling process^{56,57} and in recent hydrodynamical simulations⁵⁸.

The cooling gas is assumed to settle into a disk supported by rotation. We directly estimate the disk size based on the spin parameter of the hosting dark matter halo⁵⁹. Once the gas surface density exceeds a critical threshold motivated by observations⁶⁰, we assume that star formation proceeds in the disk, with an efficiency of order $\simeq 10\%$ on a disk dynamical time. This parameterisation reproduces the phenomenological laws of star formation in observed disk galaxies^{61,62} and the observed gas fractions at low redshift.

Supernova explosions associated with short-lived massive stars are believed to regulate star formation in galaxies, particularly in small systems with shallow potential wells⁶³. Observations suggest that supernovae blow gas out of star-forming disks, with a rate that is roughly proportional to the total amount of stars formed⁶⁴. We adopt this observational scaling, and estimate how much of this gas can join the hot halo of the galaxy given the total amount of energy released by the supernovae, and how much may be blown out of the halo entirely. The efficiency of such mass-loss is a strong function of the potential well depth of the galaxy. In our model, small galaxies may blow away their remaining gas entirely in an intense burst of star formation, while large galaxies do not exhibit

any outflows.

Morphological evolution. We characterise galaxy morphology by a simple bulge-to-disk ratio which can be transformed into an approximate Hubble type according to observational trends⁶⁵. While the generic mode of gas cooling leads to disk formation, we consider two possible channels for the formation of bulges: secular evolution due to disk instabilities, or as a consequence of galaxy merger events.

Secular evolution can trigger bar and bulge formation in disk galaxies. We invoke simple stability arguments for self-gravitating stellar disks⁵⁹ to determine the mass of stars that needs to be put into a nuclear bulge component to render the stellar disk stable.

Galaxy mergers are described by the halo merger tree constructed from the simulation, augmented with a timescale for the final stages of a merger whenever we lose track of a substructure due to finite spatial and time resolution. We then estimate the remaining survival time in a standard fashion based on the dynamical friction timescale. We use the mass ratio of two merging galaxies to distinguish between two classes of mergers. *Minor mergers* involve galaxies with mass ratio less than 0.3. In this case, we assume that the disk of the larger galaxy survives, while the merging satellite becomes part of the bulge component. For larger mass ratios, we assume a *major merger* takes place, leading to destruction of both disks, and reassembly of all stars in a common spheroid. Such an event is the channel through which pure elliptical galaxies can form. The cold gas in the satellite of a minor merger, or the cold gas in both galaxies of a major merger, is assumed to be partially or fully consumed in a nuclear starburst in which additional bulge stars are formed. The detailed parameterisation of such induced starbursts follows results obtained from systematic parameter studies of hydrodynamical galaxy collision simulations⁶⁶⁻⁶⁸.

Spectrophotometric modelling. To make direct contact with observational data, it is essential to compute spectra and magnitudes for the model galaxies, in the passbands commonly used in observations. Modern population synthesis models allow an accurate prediction of spectrophotometric properties of stellar populations as a function of age and metallicity^{69,70}. We apply such a model⁷⁰ and compute magnitudes in a number of passbands separately for both bulge and disk components and in both rest- and observer-frames. Dust obscuration effects are difficult to model in general and present a major source of uncertainty, especially for simulated galaxies at high redshift. We apply a rather simple plane-parallel slab dust model⁵⁰, as a first-order approximation to dust extinction.

Metal enrichment of the diffuse gas component can also be important, because it affects both cooling rates in mod-

erately sized halos and galaxy colours through the population synthesis models. Our treatment of metal enrichment and transport is close to an earlier semi-analytic model⁷¹. In it, metals produced and released by massive stars are placed first into the cold star forming gas, from which they can be transported into the diffuse hot halo or into the intergalactic medium by supernova feedback. We assume a homogenous metallicity (i.e. perfect mixing) within each of the gas components, although the hot and cold gas components can have different metallicities.

Active galactic nuclei. Supermassive black holes are believed to reside at the centre of most, if not all, spheroidal galaxies, and during their active phases they power luminous quasars and active galactic nuclei. There is substantial observational evidence that suggests a connection between the formation of galaxies and the build-up of supermassive black holes (BH). In fact, the energy input provided by BHs may play an important role in shaping the properties of galaxies^{72–74}, and in reionising the universe^{75,76}.

Our theoretical model for galaxy and AGN formation extends an earlier semi-analytic model for the joint build-up of the stellar and supermassive black hole components⁷⁷. This adopts the hypothesis that quasar phases are triggered by galaxy mergers. In these events, cold gas is tidally forced into the centre of a galaxy where it can both fuel a nuclear starburst and be available for central AGN accretion. We parameterise the efficiency of the feeding process of the BHs as in the earlier work⁷⁷, and normalise it to reproduce the observed scaling relation between the bulge mass and the BH mass at the present epoch^{78,79}. This ‘quasar mode’ of BH evolution provides the dominant mass growth of the BH population, with a total cumulative accretion rate that peaks at $z \simeq 3$, similar to the observed population of quasars.

A new aspect of our model is the addition of a ‘radio mode’ of BH activity, motivated by the observational phenomenology of nuclear activity in groups and clusters of galaxies. Here, accretion onto nuclear supermassive BHs is accompanied by powerful relativistic jets which can inflate large radio bubbles in clusters, and trigger sound waves in the intracluster medium (ICM). The buoyant rise of the bubbles^{80,81} together with viscous dissipation of the sound waves⁸² is capable of providing a large-scale heating of the ICM, thereby offsetting cooling losses⁸³. These physical processes are arguably the most likely explanation of the ‘cooling-flow puzzle’: the observed absence of the high mass dropout rate expected due to the observed radiative cooling in clusters of galaxies. We parameterise the radio mode as a mean heating rate into the hot gas proportional to the mass of the black hole and to the $3/2$ power of the temperature of the hot gas. The prefactor is set by requiring a good match to the bright end of the observed present-day luminosity function of galaxies. The

latter is affected strongly by the radio mode, which reduces the supply of cold gas to massive central galaxies and thus shuts off their star formation. Without the radio mode, central cluster galaxies invariably become too bright and too blue due to excessive cooling flows. The total BH accretion rate in this radio mode becomes significant only at very low redshift, but it does not contribute significantly to the cumulative BH mass density at the present epoch.

References

1. Peebles, P. J. E. Structure of the Coma Cluster of Galaxies. *Astron. J.* **75**, 13–20 (1970).
2. Miyoshi, K. & Kihara, T. Development of the correlation of galaxies in an expanding universe. *Pub. Astron. Soc. Jap.* **27**, 333–346 (1975).
3. White, S. D. M. The dynamics of rich clusters of galaxies. *Mon. Not. R. Astron. Soc.* **177**, 717–733 (1976).
4. Aarseth, S. J., Turner, E. L. & Gott, J. R. N-body simulations of galaxy clustering. I - Initial conditions and galaxy collapse times. *Astrophys. J.* **228**, 664–683 (1979).
5. Efstathiou, G. & Eastwood, J. W. On the clustering of particles in an expanding universe. *Mon. Not. R. Astron. Soc.* **194**, 503–525 (1981).
6. Davis, M., Efstathiou, G., Frenk, C. S. & White, S. D. M. The evolution of large-scale structure in a universe dominated by cold dark matter. *Astrophys. J.* **292**, 371–394 (1985).
7. White, S. D. M., Frenk, C. S., Davis, M. & Efstathiou, G. Clusters, filaments, and voids in a universe dominated by cold dark matter. *Astrophys. J.* **313**, 505–516 (1987).
8. Carlberg, R. G. & Couchman, H. M. P. Mergers and bias in a cold dark matter cosmology. *Astrophys. J.* **340**, 47–68 (1989).
9. Suto, Y. & Sugimoto, T. Redshift-space correlation functions in the cold dark matter scenario. *Astrophys. J. Let.* **370**, L15–L18 (1991).
10. Warren, M. S., Quinn, P. J., Salmon, J. K. & Zurek, W. H. Dark halos formed via dissipationless collapse. I - Shapes and alignment of angular momentum. *Astrophys. J.* **399**, 405–425 (1992).
11. Gelb, J. M. & Bertschinger, E. Cold dark matter. 1: The formation of dark halos. *Astrophys. J.* **436**, 467–490 (1994).

12. Zurek, W. H., Quinn, P. J., Salmon, J. K. & Warren, M. S. Large-scale structure after COBE: Peculiar velocities and correlations of cold dark matter halos. *Astrophys. J.* **431**, 559–568 (1994).
13. Jenkins, A. *et al.* Evolution of Structure in Cold Dark Matter Universes. *Astrophys. J.* **499**, 20–40 (1998).
14. Governato, F. *et al.* Properties of galaxy clusters: mass and correlation functions. *Mon. Not. R. Astron. Soc.* **307**, 949–966 (1999).
15. Bode, P., Bahcall, N. A., Ford, E. B. & Ostriker, J. P. Evolution of the Cluster Mass Function: GPC³ Dark Matter Simulations. *Astrophys. J.* **551**, 15–22 (2001).
16. Colberg, J. M. *et al.* Clustering of galaxy clusters in cold dark matter universes. *Mon. Not. R. Astron. Soc.* **319**, 209–214 (2000).
17. Wambsganss, J., Bode, P. & Ostriker, J. P. Giant Arc Statistics in Concord with a Concordance Lambda Cold Dark Matter Universe. *Astrophys. J. Let.* **606**, L93–L96 (2004).
18. Klypin, A. A. & Shandarin, S. F. Three-dimensional numerical model of the formation of large-scale structure in the Universe. *Mon. Not. R. Astron. Soc.* **204**, 891–907 (1983).
19. White, S. D. M., Frenk, C. S. & Davis, M. Clustering in a neutrino-dominated universe. *Astrophys. J. Let.* **274**, L1–L5 (1983).
20. Centrella, J. & Melott, A. L. Three-dimensional simulation of large-scale structure in the universe. *Nature* **305**, 196–198 (1983).
21. Park, C. Large N-body simulations of a universe dominated by cold dark matter. *Mon. Not. R. Astron. Soc.* **242**, 59P–61P (1990).
22. Bertschinger, E. & Gelb, J. M. Cosmological N-body simulations. *Computers in Physics* **5**, 164–175 (1991).
23. Bardeen, J. M., Bond, J. R., Kaiser, N. & Szalay, A. S. The statistics of peaks of Gaussian random fields. *Astrophys. J.* **304**, 15–61 (1986).
24. Seljak, U. & Zaldarriaga, M. A Line-of-Sight Integration Approach to Cosmic Microwave Background Anisotropies. *Astrophys. J.* **469**, 437–444 (1996).
25. Spergel, D. N. *et al.* First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters. *Astrophys. J. Suppl.* **148**, 175–194 (2003).
26. Seljak, U. *et al.* Cosmological parameter analysis including SDSS Ly-alpha forest and galaxy bias: constraints on the primordial spectrum of fluctuations, neutrino mass, and dark energy. *PRD*, submitted, astro-ph/0407372 (2004).
27. White, S. D. M. Formation and evolution of galaxies: Les houches lectures. In Schaefer, R., Silk, J., Spiro, M. & Zinn-Justin, J. (eds.) *Cosmology and Large-Scale Structure* (Dordrecht: Elsevier, astro-ph/9410043, 1996).
28. Hernquist, L., Hut, P. & Makino, J. Discreteness Noise versus Force Errors in N-Body Simulations. *Astrophys. J. Let.* **402**, L85–L88 (1993).
29. Springel, V., Yoshida, N. & White, S. D. M. GADGET: a code for collisionless and gasdynamical cosmological simulations. *New Astronomy* **6**, 79–117 (2001).
30. Xu, G. A New Parallel N-Body Gravity Solver: TPM. *Astrophys. J. Suppl.* **98**, 355–366 (1995).
31. Bode, P., Ostriker, J. P. & Xu, G. The Tree Particle-Mesh N-Body Gravity Solver. *Astrophys. J. Suppl.* **128**, 561–569 (2000).
32. Bagla, J. S. TreePM: A Code for Cosmological N-Body Simulations. *Journal of Astrophysics and Astronomy* **23**, 185–196 (2002).
33. Barnes, J. & Hut, P. A Hierarchical O(NlogN) Force-Calculation Algorithm. *Nature* **324**, 446–449 (1986).
34. Hockney, R. W. & Eastwood, J. W. *Computer Simulation Using Particles* (New York: McGraw-Hill, 1981, 1981).
35. Salmon, J. K. & Warren, M. S. Skeletons from the treecode closet. *J. Comp. Phys.* **111**, 136 (1994).
36. Duncan, M. J., Levison, H. F. & Lee, M. H. A Multiple Time Step Symplectic Algorithm for Integrating Close Encounters. *Astron. J.* **116**, 2067–2077 (1998).
37. Ghigna, S. *et al.* Dark matter haloes within clusters. *Mon. Not. R. Astron. Soc.* **300**, 146–162 (1998).
38. Springel, V., White, S. D. M., Tormen, G. & Kauffmann, G. Populating a cluster of galaxies - I. Results at z=0. *Mon. Not. R. Astron. Soc.* **328**, 726–750 (2001).
39. White, S. D. M. & Frenk, C. S. Galaxy formation through hierarchical clustering. *Astrophys. J.* **379**, 52–79 (1991).
40. Kauffmann, G., White, S. D. M. & Guiderdoni, B. The Formation and Evolution of Galaxies Within Merging Dark Matter Haloes. *Mon. Not. R. Astron. Soc.* **264**, 201–218 (1993).

41. Cole, S., Aragon-Salamanca, A., Frenk, C. S., Navarro, J. F. & Zepf, S. E. A Recipe for Galaxy Formation. *Mon. Not. R. Astron. Soc.* **271**, 781–806 (1994).
42. Kauffmann, G. The age of elliptical galaxies and bulges in a merger model. *Mon. Not. R. Astron. Soc.* **281**, 487–492 (1996).
43. Kauffmann, G., Nusser, A. & Steinmetz, M. Galaxy formation and large-scale bias. *Mon. Not. R. Astron. Soc.* **286**, 795–811 (1997).
44. Baugh, C. M., Cole, S., Frenk, C. S. & Lacey, C. G. The Epoch of Galaxy Formation. *Astrophys. J.* **498**, 504–521 (1998).
45. Somerville, R. S. & Primack, J. R. Semi-analytic modelling of galaxy formation: the local Universe. *Mon. Not. R. Astron. Soc.* **310**, 1087–1110 (1999).
46. Cole, S., Lacey, C. G., Baugh, C. M. & Frenk, C. S. Hierarchical galaxy formation. *Mon. Not. R. Astron. Soc.* **319**, 168–204 (2000).
47. Benson, A. J., Lacey, C. G., Baugh, C. M., Cole, S. & Frenk, C. S. The effects of photoionization on galaxy formation - I. Model and results at $z=0$. *Mon. Not. R. Astron. Soc.* **333**, 156–176 (2002).
48. Kauffmann, G. & White, S. D. M. The merging history of dark matter haloes in a hierarchical universe. *Mon. Not. R. Astron. Soc.* **261**, 921–928 (1993).
49. Somerville, R. S. & Kolatt, T. S. How to plant a merger tree. *Mon. Not. R. Astron. Soc.* **305**, 1–14 (1999).
50. Kauffmann, G., Colberg, J. M., Diaferio, A. & White, S. D. M. Clustering of galaxies in a hierarchical universe - I. Methods and results at $z=0$. *Mon. Not. R. Astron. Soc.* **303**, 188–206 (1999).
51. Croton, D. J., White, S. D. M., Springel, V. & et al. The many lives of AGN: super-massive black holes, cooling flows, galaxy colours and luminosities. *Mon. Not. R. Astron. Soc.* in preparation (2005).
52. Gnedin, N. Y. Cosmological Reionization by Stellar Sources. *Astrophys. J.* **535**, 530–554 (2000).
53. Kravtsov, A. V., Gnedin, O. Y. & Klypin, A. A. The Tumultuous Lives of Galactic Dwarfs and the Missing Satellites Problem. *Astrophys. J.* **609**, 482–497 (2004).
54. Yoshida, N., Stoehr, F., Springel, V. & White, S. D. M. Gas cooling in simulations of the formation of the galaxy population. *Mon. Not. R. Astron. Soc.* **335**, 762–772 (2002).
55. Helly, J. C. *et al.* A comparison of gas dynamics in smooth particle hydrodynamics and semi-analytic models of galaxy formation. *Mon. Not. R. Astron. Soc.* **338**, 913–925 (2003).
56. Forcado-Miro, M. I. & White, S. D. M. Radiative shocks in galaxy formation. I: Cooling of a primordial plasma with no sources of heating. astro-ph/9712204 (1997).
57. Birnboim, Y. & Dekel, A. Virial shocks in galactic haloes? *Mon. Not. R. Astron. Soc.* **345**, 349–364 (2003).
58. Keres, D., Katz, N., Weinberg, D. H. & Dave, R. How Do Galaxies Get Their Gas? *Mon. Not. R. Astron. Soc.* astro-ph/0407095 (2004).
59. Mo, H. J., Mao, S. & White, S. D. M. The formation of galactic discs. *Mon. Not. R. Astron. Soc.* **295**, 319–336 (1998).
60. Kauffmann, G. Disc galaxies at $z=0$ and at high redshift: an explanation of the observed evolution of damped Ly α absorption systems. *Mon. Not. R. Astron. Soc.* **281**, 475–486 (1996).
61. Kennicutt, R. C. The star formation law in galactic disks. *Astrophys. J.* **344**, 685–703 (1989).
62. Kennicutt, R. C. The Global Schmidt Law in Star-forming Galaxies. *Astrophys. J.* **498**, 541–552 (1998).
63. Dekel, A. & Silk, J. The origin of dwarf galaxies, cold dark matter, and biased galaxy formation. *Astrophys. J.* **303**, 39–55 (1986).
64. Martin, C. L. Properties of Galactic Outflows: Measurements of the Feedback from Star Formation. *Astrophys. J.* **513**, 156–160 (1999).
65. Simien, F. & de Vaucouleurs, G. Systematics of bulge-to-disk ratios. *Astrophys. J.* **302**, 564–578 (1986).
66. Mihos, J. C. & Hernquist, L. Triggering of starbursts in galaxies by minor mergers. *Astrophys. J. Let.* **425**, L13–L16 (1994).
67. Mihos, J. C. & Hernquist, L. Gasdynamics and Starbursts in Major Mergers. *Astrophys. J.* **464**, 641–663 (1996).
68. Cox, T. J., Primack, J., Jonsson, P. & Somerville, R. S. Generating Hot Gas in Simulations of Disk-Galaxy Major Mergers. *Astrophys. J. Let.* **607**, L87–L90 (2004).

69. Bruzual A., G. & Charlot, S. Spectral evolution of stellar populations using isochrone synthesis. *Astrophys. J.* **405**, 538–553 (1993).
70. Bruzual, G. & Charlot, S. Stellar population synthesis at the resolution of 2003. *Mon. Not. R. Astron. Soc.* **344**, 1000–1028 (2003).
71. De Lucia, G., Kauffmann, G. & White, S. D. M. Chemical enrichment of the intracluster and intergalactic medium in a hierarchical galaxy formation model. *Mon. Not. R. Astron. Soc.* **349**, 1101–1116 (2004).
72. Silk, J. & Rees, M. J. Quasars and galaxy formation. *Astr. & Astrop.* **331**, L1–L4 (1998).
73. Di Matteo, T., Springel, V. & Hernquist, L. Energy input from quasars regulates the growth and activity of black holes and their host galaxies. *Nature* **433**, 604–607 (2005).
74. Springel, V., Di Matteo, T. & Hernquist, L. Black Holes in Galaxy Mergers: The Formation of Red Elliptical Galaxies. *Astrophys. J. Let.* **620**, L79–L82 (2005).
75. Haardt, F. & Madau, P. Radiative Transfer in a Clumpy Universe. II. The Ultraviolet Extragalactic Background. *Astrophys. J.* **461**, 20–37 (1996).
76. Madau, P., Rees, M. J., Volonteri, M., Haardt, F. & Oh, S. P. Early Reionization by Mini-quasars. *Astrophys. J.* **604**, 484–494 (2004).
77. Kauffmann, G. & Haehnelt, M. A unified model for the evolution of galaxies and quasars. *Mon. Not. R. Astron. Soc.* **311**, 576–588 (2000).
78. Magorrian, J. *et al.* The Demography of Massive Dark Objects in Galaxy Centers. *Astron. J.* **115**, 2285–2305 (1998).
79. Ferrarese, L. & Merritt, D. A Fundamental Relation between Supermassive Black Holes and Their Host Galaxies. *Astrophys. J. Let.* **539**, L9–L12 (2000).
80. Churazov, E., Brüggén, M., Kaiser, C. R., Böhringer, H. & Forman, W. Evolution of Buoyant Bubbles in M87. *Astrophys. J.* **554**, 261–273 (2001).
81. Brüggén, M. & Kaiser, C. R. Hot bubbles from active galactic nuclei as a heat source in cooling-flow clusters. *Nature* **418**, 301–303 (2002).
82. Fabian, A. C. *et al.* A deep Chandra observation of the Perseus cluster: shocks and ripples. *Mon. Not. R. Astron. Soc.* **344**, L43–L47 (2003).
83. Vecchia, C. D. *et al.* Quenching cluster cooling flows with recurrent hot plasma bubbles. *Mon. Not. R. Astron. Soc.* **355**, 995–1004 (2004).