

Universal structure of dark matter haloes over a mass range of 20 orders of magnitude

<https://doi.org/10.1038/s41586-020-2642-9>

J. Wang^{1,2}✉, S. Bose³, C. S. Frenk⁴✉, L. Gao^{1,2}, A. Jenkins⁴, V. Springel⁵ & S. D. M. White⁵✉

Received: 21 November 2019

Accepted: 23 June 2020

Published online: 2 September 2020

 Check for updates

Cosmological models in which dark matter consists of cold elementary particles predict that the dark halo population should extend to masses many orders of magnitude below those at which galaxies can form^{1–3}. Here we report a cosmological simulation of the formation of present-day haloes over the full range of observed halo masses (20 orders of magnitude) when dark matter is assumed to be in the form of weakly interacting massive particles of mass approximately 100 gigaelectronvolts. The simulation has a full dynamic range of 30 orders of magnitude in mass and resolves the internal structure of hundreds of Earth-mass haloes in as much detail as it does for hundreds of rich galaxy clusters. We find that halo density profiles are universal over the entire mass range and are well described by simple two-parameter fitting formulae^{4,5}. Halo mass and concentration are tightly related in a way that depends on cosmology and on the nature of the dark matter. For a fixed mass, the concentration is independent of the local environment for haloes less massive than those of typical galaxies. Haloes over the mass range of 10^{-3} to 10^{11} solar masses contribute about equally (per logarithmic interval) to the luminosity produced by dark matter annihilation, which we find to be smaller than all previous estimates by factors ranging up to one thousand³.

Figure 1 illustrates our simulation scheme. The top left panel shows the present-day distribution of dark matter in a slab cut from a large cosmological simulation (L0) identical to the 2005 Millennium Simulation⁶, except that cosmological parameters are updated to reflect recent analyses of cosmic microwave background data from the Planck satellite. The total mass in this simulation is about 10^{19} solar masses (M_{\odot}). The circle outlines a spherical region chosen to avoid any of the more massive structures. The material in this region was traced back to the initial time and used to define a Lagrangian volume within which the particle count was increased by a factor of about 2,000, the particle mass was decreased by the same factor, and the representation of the linear cosmological fluctuation field was extended to ~ 10 times smaller scale while retaining all structure present in the original simulation. The mass outside this ‘zoomed’ region was consolidated into a smaller number of particles whose individual masses increased with distance from its centre. These new initial conditions were then integrated down to the present day. The top middle panel of Fig. 1 shows a projection of the mass within the largest sphere enclosed in the high-resolution region. It has resolution 2,000 times better in mass and ~ 10 times better in length than the first panel but contains a comparable number of well resolved haloes (that is, made up of 10^4 or more simulation particles).

The small circle in this panel outlines a spherical subregion of this first level zoom (L1) that avoids any larger structures. It was again traced back to the initial conditions, refined by another factor of 500 in mass, and re-simulated to give a second-level zoom (L2) for which the final structure within the high-resolution region is shown in the top right

panel of Fig. 1. This whole process was repeated eight times, each revealing ever-smaller structures, to give a final simulation (L8c) with eight levels of refinement and a high-resolution particle mass of $\sim 10^{-11} M_{\odot}$, hence a dynamic range of 30 orders of magnitude. The final mass distributions in the high-resolution regions at each stage are shown in the remaining panels of Fig. 1. Their initial conditions were set by using second-order Lagrangian perturbation theory with an initial power spectrum with Planck parameters together with a free-streaming cut-off at small spatial scales corresponding to a thermal weakly interacting massive particle (WIMP), which, for illustrative purposes, we assume to have mass 100 GeV. One of the zooms (L7c) was repeated without this cut-off (giving L7) in order to understand its effects on halo structure (see the Methods section for further details).

Considerable effort was needed to ensure that the initial conditions procedure, the force calculation accuracy and the time integration scheme of the simulation code were adequate to give reliable results over such a large dynamic range. In the Methods section we describe some of these improvements, and we present convergence tests that demonstrate that they were successful. The more massive haloes in the high-resolution region at each level can all be individually identified in the ‘parent’ level, making it possible to check that the masses agree in the two cases. For the most massive haloes, the resolution of the parent level is sufficient to test that their radial density profiles also agree. The plots in the Methods section show that both these tests are passed for all adjacent level pairs, giving us confidence that our results for the internal structure of dark matter haloes are reliable for

¹Key Laboratory for Computational Astrophysics, National Astronomical Observatories, Chinese Academy of Sciences, Beijing, China. ²School of Astronomy and Space Science, University of Chinese Academy of Sciences, Beijing, China. ³Center for Astrophysics | Harvard and Smithsonian, Cambridge, MA, USA. ⁴Institute for Computational Cosmology, Department of Physics, Durham University, Durham, UK. ⁵Max-Planck-Institut für Astrophysik, Garching, Germany. ✉e-mail: jie.wang@nao.cas.cn; c.s.frenk@durham.ac.uk; swhite@mpa-garching.mpg.de




Fig. 1 | Projected dark matter density maps at each simulation level. Images of the dark matter distribution in a slice 30 Mpc thick through the base level of our simulation (L0) and in spheres almost entirely contained within the higher-resolution region of each of the eight successive levels of zoom (L1 to L8c). The zoom sequence is indicated by arrows between the panels, and a

circle in each of the first eight panels indicates the zoom region shown in the next panel. Bars give a length scale for each plot. In the first panel, the largest haloes have a mass similar to that of a rich galaxy cluster, whereas in the last panel the smallest clearly visible haloes have a mass comparable to that of the Earth.

$10^{-6} < M_{\text{halo}}/M_{\odot} < 10^{15}$, the entire halo mass range that should be populated if the dark matter consists of 100-GeV WIMPs.

Figure 2 shows the first major result of this article. At each level of our simulation we identify a sample of a few tens of well resolved, quasi-equilibrium haloes of similar mass. For these, we construct a mean, spherically averaged mass density profile which we compare with two well known two-parameter fitting formulae, the NFW profile⁷,

$$\rho(r) = \rho_s r_s^3 / r(r + r_s)^2, \quad (1)$$

where ρ_s and r_s are the characteristic density and scale radius respectively, and the Einasto profile^{5,8},

$$\rho(r) = \rho_2 \exp[-2\alpha^{-1}((r/r_{-2})^\alpha - 1)], \quad (2)$$

where r_{-2} is the radius at which the logarithmic slope is -2 , and α is a shape parameter that we fix to $\alpha = 0.16$. These formulae were fitted to the mean profiles at each level over the radial range where these are numerically robust. Figure 2 shows differences between the measured profiles and these best fits in two different ways. Figure 2a gives the logarithmic slope of the profiles as a function of r/r_{-2} , where r_{-2} is the characteristic radius of the best Einasto fit. In such a plot, each fitting formula predicts a universal curve, a Z-shaped transition between values of -1 and -3 on the vertical axis in the NFW case, and a smoother, more gradual change of slope in the Einasto case. Over 20 orders of magnitude in mass, the mean profiles of the simulated haloes are all

very similar and are closer to the Einasto case than to the NFW. The only clear exception is that the curve for L0, representing haloes of moderately rich galaxy clusters, is noticeably steeper than the others. Larger values of α have previously been shown to give a better fit to such objects, but the trend in α does not continue to the much lower masses that we have now simulated. Figure 2b and c shows that over the factor of about 10^4 in density for which these profiles are robustly measured, NFW fits are almost everywhere accurate to better than about 10% and Einasto fits to a few per cent. This universality over 20 orders of magnitude in halo mass is remarkable, not least because reliable simulation data for the nearby Universe (redshift $z = 0$) have not previously been available for most of this range.

The mass of a dark matter halo is conventionally taken as that within the virial radius, defined here as r_{200} , the radius enclosing a mean density 200 times the critical value. Mass and concentration, $c = r_{200}/r_{\text{ch}}$, can then be used as alternative parameters for the above fitting functions, with $r_{\text{ch}} = r_s$ and r_{-2} , respectively, for the NFW and Einasto cases. Figure 3 shows the mass–concentration relation in our simulation, considering only haloes with enough particles for a reliable concentration measurement ($>10^4$ at the higher levels, somewhat fewer in L0, L1 and L2). Each coloured band gives the [10%, 90%] range for haloes at a given level, with a white line indicating the median concentration at each mass. Over the mass range $10^{15} > M_{200}/M_{\odot} > 10^{10}$ relevant for galaxy clusters and for all but the very faintest galaxies, concentration rises fairly rapidly with decreasing mass. The relation becomes shallower for lower-mass haloes, however, and eventually turns down as the free-streaming mass

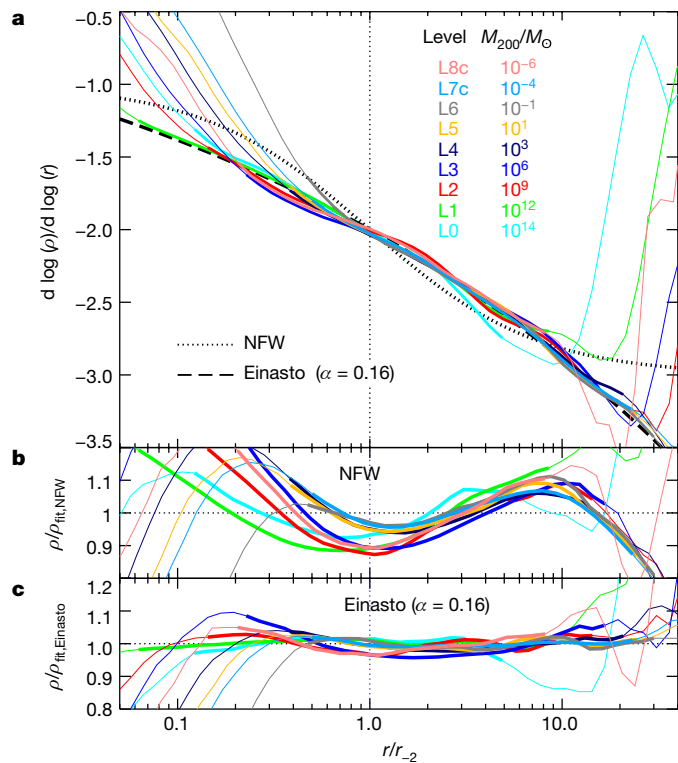


Fig. 2 | Density profiles for haloes with mass between that of the Earth and that of a rich galaxy cluster. As described in detail in the Methods section, results for all well resolved equilibrium haloes in a narrow mass bin at each level are averaged together. **a**, The logarithmic slope $d \log(\rho)/d \log(r)$ is shown as a function of radius normalized by r_{-2} . The result for each level is represented by a different colour, as indicated in the key. A thicker line is used over the most reliable range between the convergence radius r_{conv} and r_{200} . The number of haloes in each stack is listed in Extended Data Table 1. Predictions for NFW and Einasto profiles are shown as dotted and dashed black curves, respectively. **b**, The ratio of each stacked profile to the best-fit NFW profile is shown as a function of r/r_{-2} . **c**, As in **b**, for the Einasto profile (with α fixed at 0.16).

is approached. This turn-down is most clearly seen by comparing results for L7c and L8c, where the initial conditions included a free-streaming cut-off, with those for L7, where they did not. In the Methods section we compare matched objects in L7 and L7c, showing that the cut-off reduces the concentration of individual haloes by an increasing amount as the free-streaming mass (about Earth mass for a 100-GeV WIMP) is approached. Like all N -body simulations of structure formation with a free-streaming cut-off, both L7c and L8c form spurious small-scale clumps. As discussed in the Methods section, this does not affect the results of this article.

Other points of interest in Fig. 3 are that the scatter in concentration depends very little on halo mass, being about 0.15 dex over the full halo mass range plotted, and that previously published mass–concentration relations, although agreeing roughly for galaxy- and cluster-mass haloes, give wildly divergent results when extrapolated down to the lower halo masses that are simulated here. Only one model^{9,10} represents our results relatively well, both with and without a free-streaming cut-off. In the Methods section we give a simpler fitting formula that fits our numerical data even better and follows their approach to predict the effects of varying the free-streaming scale.

The concentration–mass relation is of critical importance for predicting WIMP annihilation radiation signals. Previous work implied that these should be dominated by haloes with mass relatively close to the free-streaming limit, but this changes substantially for the reduced concentrations that we find (see the Methods section for details).




Fig. 3 | Halo concentration as a function of mass over a mass range of 20 orders of magnitude. The median values of the concentration, $c_{\text{Einasto}} = r_{200}/r_{-2}$ (from the best-fitting Einasto profile), in each mass bin are shown as white curves, with coloured regions showing the range [10%, 90%]. As before, each zoom level is shown with a different colour, and we give results both for L7c, which has a free-streaming cut-off, and for the otherwise identical L7, which does not. Mass–concentration relations from five published models^{10,18–21} are shown as smooth dashed lines in different colours: black, ref.¹⁰; dark blue, ref.¹⁸; red, ref.¹⁹; green, ref.²⁰; light blue, ref.²¹. Solid black lines show the fitting formulae given in the Methods section. The pairs of solid and dashed black lines give predictions for the cases with and without a free-streaming cut-off. The vertical dotted lines indicate the limits below which spurious haloes are expected to affect L7c and L8c (see Methods).

Structures down to very small scales should also be present in the outer regions of much more massive haloes, resulting in a substantial boost in the total amount (and a flattening of the radial profile) of their annihilation luminosity. Our simulation cannot address these issues directly, but it can be used to inform the further modelling required¹¹.

The high-resolution region of L8c is only about 300 pc across at the final time and contains a total mass that is only about 1% that of the Sun, implying a mean density about 0.3% that of L0. This low value is a consequence of repeatedly choosing to refine regions that avoid any massive nonlinear structure. It is still rather larger than the median $z=0$ density of a Universe with Planck cosmology dominated by a 100-GeV WIMP¹². One may nevertheless question whether the haloes that we have simulated can be considered representative of the general population of similar-mass objects. In the Methods section, we test this by investigating how the concentration of our haloes depends on the density of their immediate environment, measured in a spherical shell between $5r_{200}$ and $10r_{200}$. Remarkably, despite the low mean density of the higher refinement levels, the distribution of this environment density is centred just below the cosmic mean for all haloes less massive than about $10^{10} M_{\odot}$, with a spread of at least an order of magnitude. In addition, such haloes show no systematic trend of concentration with local density. This encourages us to believe that the concentration–mass relation of Fig. 3 should be representative of the full halo population. Previous attempts^{13–16} to simulate the structure of very-low-mass haloes have failed precisely because they did not take account of the low-density larger-scale environment in which such haloes live at $z=0$.

A final related issue is that our simulation follows dark matter only, neglecting the effects of the 16% of cosmic matter which is baryonic. Both relative velocity and pressure effects¹⁷ are expected to prevent the gas from following the dark matter on the very small scales that we have simulated. Although accurate treatment of these effects is beyond present capabilities given the dynamic range that we are considering, we may expect that at the higher refinement levels they would increase the mean density (because on average the baryons will be less underdense than the dark matter) but reduce the growth rate of haloes (because this is driven by the dark matter density only, rather than by the total density). Given that halo concentration depends weakly on halo mass and not at all on local environment density, we expect these effects to shift our results by at most small factors, but this will require further work for confirmation.

The universal halo structure that we have demonstrated across 20 orders of magnitude in halo mass and the associated mass–concentration relation differ substantially from previously proposed extrapolations. This affects predictions not only for annihilation signals, which depend strongly on the concentration of the lowest-mass haloes, but also for perturbations of image structure in strong gravitational lenses and for structure in stellar streams in the Galaxy’s halo, both of which aim to constrain the nature of dark matter using haloes of mass $10^6 M_\odot$ to $10^9 M_\odot$.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-020-2642-9>.

1. Blumenthal, G. R., Faber, S. M., Primack, J. R. & Rees, M. J. Formation of galaxies and large-scale structure with cold dark matter. *Nature* **311**, 517–525 (1984).
2. White, S. D. M. & Frenk, C. S. Galaxy formation through hierarchical clustering. *Astrophys. J.* **379**, 52–79 (1991).

3. Bertone, G., Hooper, D. & Silk, J. Particle dark matter: evidence, candidates and constraints. *Phys. Rep.* **405**, 279–390 (2005).
4. Navarro, J. F., Frenk, C. S. & White, S. D. M. A universal density profile from hierarchical clustering. *Astrophys. J.* **490**, 493–508 (1997).
5. Navarro, J. F. et al. The inner structure of Λ CDM haloes—III. Universality and asymptotic slopes. *Mon. Not. R. Astron. Soc.* **349**, 1039–1051 (2004).
6. Springel, V. et al. Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature* **435**, 629–636 (2005).
7. Navarro, J. F., Frenk, C. S. & White, S. D. M. The structure of cold dark matter halos. *Astrophys. J.* **462**, 563 (1996).
8. Einasto, J. On the construction of a composite model for the galaxy and on the determination of the system of galactic parameters. *Trudy Astrofizicheskogo Instituta Alma-Ata* **5**, 87–100 (1965).
9. Ludlow, A. D. et al. The mass–concentration–redshift relation of cold dark matter haloes. *Mon. Not. R. Astron. Soc.* **441**, 378–388 (2014).
10. Ludlow, A. D. et al. The mass–concentration–redshift relation of cold and warm dark matter haloes. *Mon. Not. R. Astron. Soc.* **460**, 1214–1232 (2016).
11. Springel, V. et al. The Aquarius project: the subhaloes of galactic haloes. *Mon. Not. R. Astron. Soc.* **391**, 1685–1711 (2008).
12. Stücker, J., Busch, P. & White, S. D. M. The median density of the Universe. *Mon. Not. R. Astron. Soc.* **477**, 3230–3246 (2018).
13. Diemand, J., Moore, B. & Stadel, J. Earth-mass dark-matter haloes as the first structures in the early Universe. *Nature* **433**, 389–391 (2005).
14. Ishiyama, T., Makino, J. & Ebisuzaki, T. Gamma-ray signal from earth-mass dark matter microhalos. *Astrophys. J. Lett.* **723**, L195–L200 (2010).
15. Anderhalden, D. & Diemand, J. Density profiles of CDM microhalos and their implications for annihilation boost factors. *J. Cosmol. Astropart. Phys.* **04**, 009 (2013).
16. Angulo, R. E., Hahn, O., Ludlow, A. D. & Bonoli, S. Earth-mass haloes and the emergence of NFW density profiles. *Mon. Not. R. Astron. Soc.* **471**, 4687–4701 (2017).
17. Tselikhovich, D. & Hirata, C. Relative velocity of dark matter and baryonic fluids and the formation of the first structures. *Phys. Rev. D* **82**, 083520 (2010).
18. Neto, A. F. et al. The statistics of Λ CDM halo concentrations. *Mon. Not. R. Astron. Soc.* **381**, 1450–1462 (2007).
19. Sánchez-Conde, M. A. & Prada, F. The flattening of the concentration–mass relation towards low halo masses and its implications for the annihilation signal boost. *Mon. Not. R. Astron. Soc.* **442**, 2271–2277 (2014).
20. Dutton, A. A. & Macciò, A. V. Cold dark matter haloes in the Planck era: evolution of structural parameters for Einasto and NFW profiles. *Mon. Not. R. Astron. Soc.* **441**, 3359–3374 (2014).
21. Diemer, B. & Joyce, M. An accurate physical model for halo concentrations. *Astrophys. J.* **871**, 168 (2019).

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2020

Methods

Simulations

The hierarchical resimulation strategy that allows us to follow the evolution of haloes over 20 orders of magnitude in mass was described in the main body of this article. The base level (level 0 or L0) is a cube of length 738 Mpc and particle mass $1.55 \times 10^9 M_\odot$. At subsequent levels (L1–L8c) the mass resolution increases by factors between a few hundred and 2,000, and the volume decreases by similar factors until the particle mass reaches $1.6 \times 10^{-11} M_\odot$ in L8c. At each level, well resolved haloes (that is, with $>10^4$ particles within the virial radius) span two to three orders of magnitude in mass, ranging from $M_{200} = 10^{15} M_\odot$ in L0 to $M_{200} = 10^{-6} M_\odot$ (the Earth’s mass) in L8c. Here M_{200} is defined as the mass within a sphere enclosing a mean density 200 times the critical value. The parameters of the various levels of our simulation are listed in Extended Data Table 1. The simulation assumes a Λ CDM cosmology with Planck 2014 parameters²². Specifically, the mean matter density, mean baryon density and cosmological constant, in units of the critical density, have values $\Omega_m = 0.307$, $\Omega_b = 0.48$ and $\Omega_\Lambda = 0.693$; the present-day Hubble parameter is $H_0 = 67.77 \text{ km s}^{-1} \text{ Mpc}^{-1}$; the power-law index of the power spectrum of primordial adiabatic perturbations is $n_s = 0.961$; and the normalization of the linear power spectrum is $\sigma_8 = 0.829$.

Linear power spectrum

To create the displacement field for the initial conditions, a linear power spectrum is required that covers more than nine orders of magnitude from the fundamental modes of the L0 volume to the Nyquist cut-off of the L8c simulation. For relatively large scales, we use the same linear power spectrum as the EAGLE project²³. This was computed using the CAMB code²⁴ assuming Λ CDM with the values of the cosmological parameters given previously. As we are only modelling the dominant dark matter component, we use the BBKS fitting formula (equation G3 of ref. ²⁵) to extrapolate the power spectrum to very small scales²⁵. We adopt a similar approach to previous work^{11,26}, creating a composite matter power spectrum that smoothly transitions from the EAGLE matter power spectrum to the BBKS form over wavenumbers 7–70 Mpc^{-1} . We determined by inspection that setting the parameter $\Gamma = 0.1673$ in the BBKS transfer function and adopting an effective normalization, $\sigma_8 = 0.8811$, matches the BBKS power spectrum accurately to the EAGLE power spectrum over the transition wavenumber range; we interpolate linearly in log wavenumber over this range to produce a smooth power spectrum for pure CDM.

We then used the formulae of Green et al.²⁷ to represent the free-streaming cut-off expected for a 100-GeV WIMP. This corresponds to a specific particle physics model, but the corresponding mass scale is close to the peak of the posterior probability for the Constrained Minimal Supersymmetric Standard Model (CMSSM) parameter space²⁸ and can still be considered representative in view of more recent constraints²⁹. We show below how to adapt our results to alternative particle physics models that would predict a free-streaming cut-off on a different scale.

Making the initial conditions

Although the setting up and running of zoom simulations has become commonplace in the field of numerical cosmology, the initial conditions required for the present project are much more extreme in terms of the range of mass and length scales modelled than in previously published simulations. These exceptional demands have driven developments that go beyond the techniques described in previous work^{11,30,31}.

The initial conditions for levels L1 to L8c were created and evolved sequentially in order of increasing mass resolution. After each level was completed, a region avoiding any massive halo was selected from its high-resolution region, and this then became the next level (see Fig. 1). The amplitude and phase of the initial fluctuations present in the initial conditions of all lower levels were retained, but the amplitude and

phase of all higher-frequency fluctuations added at the new level were set independently and at random according to the power spectrum. In principle, we could make initial conditions at the resolution of L8c for any Lagrangian region within L0 without running any intermediate levels, but in the great majority of cases this would result in all of the mass being incorporated into a single halo of mass larger than that of the entire high-resolution region of L8c.

The specific features that emerge at any redshift—for example, the positions, masses and orientations of individual haloes or filaments—are a consequence of our particular realization of the linear initial conditions: that is, of our adopted power spectrum together with the specific phases and amplitudes chosen for each wave in a Fourier space representation of the initial Gaussian random field. Our phase information was taken from the Panphasia white-noise field^{31,32}, an extremely large single realization of a Gaussian white-noise field with a hierarchical octree structure. Because the Panphasia field is completely specified ahead of time, all of the structure uncovered at all resolutions is essentially predetermined, as is the similar structure that would be uncovered by a different hierarchical zoom into any other region of the L0 cube.

The creation of initial conditions at each zoom level can be divided into three stages: stage 1 is to specify the region of interest; stage 2 is to build a particle load focusing most of the particles, and therefore most of the computational effort, in the small region of interest, while aggregating particles for lower levels so that the computational time for these regions is reduced while maintaining accurate tidal forces in the high-resolution region; stage 3 is to generate and apply the displacement field to the particle load, and assign velocities to each particle.

For stage 1, we start by selecting a spherical region of interest at redshift zero from a previously completed simulation of the lower level. For L1, this is the cosmological simulation L0, but for all higher levels it is itself a zoom simulation. The region is selected by eye using projections of the density field to avoid large haloes that were previously simulated with good resolution and would be prohibitive to simulate at much higher resolution. At the same time, we avoid regions that are more underdense than necessary, as these would yield few new haloes. The region size is dictated by the cost of resimulating at the resolution desired for the next level, given that we can afford simulations with a few billion high-resolution particles. Having selected a sphere, we then use its particles to determine the location and shape of the corresponding Lagrangian region by binning their high-redshift positions onto a 40^3 cubic grid just large enough to enclose them all. Within this cube, we define a simply connected region by selecting grid cells that either contain a particle or are adjacent to one that does.

For constructing the particle load in stage 2, we use a set of cubes with a variety of sizes that tessellate the entire simulation volume. In each cube, we place one or more particles of identical mass in an arrangement that ensures that the centre of mass of the particles within every cube is at the cell centre, and we choose the total particle mass in each cell so that it has precisely the mean density of the Universe. We also place the particles as evenly as possible within each cube in the sense that if that cube were tessellated over all space, the gravitational forces on each particle due to all other particles would be essentially zero. For the region outside the high-resolution cube, we lay down a set of ‘tidal’ particles arranged within a series of concentric cubic shells centred on the high-resolution cube. For the initial conditions of the highest level, L8c, more than 47 million tidal cells, each containing a single particle, are used to represent the mass distribution in the lower levels.

The particles within the high-resolution region are the ones that form the structures analysed at each level. In L8c, for example, we place a pre-prepared set of about 50,000 particles with a glass-like structure in each of the retained cells from our 40^3 mesh. This glass-like arrangement is created in a small periodic cube and results in the net gravitational force on every particle being extremely small. We also enforce the condition that the centre of mass of the glass be exactly at

the cell centre. Because the glass is generated using periodic boundary conditions, it is simple to tile the entire high-resolution region with multiple replicas. The number of particles in the glass determines the mass resolution in the high-resolution region.

In stage 3, we generate and apply the displacement field following exactly the method described in previous work³¹. The displacement field is computed using Fourier methods for a series of concentric meshes centred on the high-resolution cube. The top-level mesh covers the entire domain, and the smallest mesh just covers the high-resolution cube. Each successive mesh is exactly half the linear size of the one above and adds independent information taken from the Panphasia field so as to be able to double the linear resolution of the displacement field. The L8c simulation required 23 levels in total with the smallest mesh being approximately 180 pc on a side.

We have used second-order Lagrangian perturbation theory (2LPT) to create the displacement and velocity fields for most of the initial conditions. In practice, however, we have found that using first-order (Zeldovich) initial conditions instead of 2LPT makes no important difference to the results provided that the starting redshift is high enough. In particular, for our chosen starting redshift of 127 for levels 0–2, and 255 for levels 3–8, there was no important difference in the halo density profile or its concentration between runs using the Zeldovich and 2LPT initial conditions. Nonetheless, we used the 2LPT initial conditions for all but levels 4 and 6.

Simulation code

The simulations were run with GADGET-4, a new version of the well tested GADGET³³ cosmological N -body code. Several improvements were implemented in this code to allow the extreme zooms considered here to be executed with the required accuracy. The most relevant is an extension of the hierarchical multipole force computation algorithm to higher expansion order, yielding better force accuracy for given computational cost. A further efficiency gain comes from replacing the one-sided Barnes and Hut tree algorithm³⁴ with a fast multipole method³⁵, where the multipole expansion is carried out symmetrically both at the source and the sink side of two interacting particle groups.

The extreme dynamic range of our zooms revealed two problems that had not shown up in more conventional cosmological simulations with uniform mass resolution. Because the magnitude of the peculiar acceleration vector, \mathbf{a} , of particles in the small structures targeted here is typically dominated by matter perturbations on much larger scales, the local timestep criterion most commonly used in cosmological N -body CDM simulations, $\Delta t \propto (\varepsilon/|\mathbf{a}|)^{1/2}$, where ε is the gravitational softening length, often fails to provide a reasonable proxy for the local dynamical time in our smallest dark matter haloes. Rather, it tends to become unrealistically small because $|\mathbf{a}|$ remains at the large values characteristic of the resolved cosmic large-scale structures in our 738-Mpc periodic box, whereas ε shrinks to the tiny scales resolved in our calculations. We address this problem by applying a hierarchical time integration algorithm³⁶ in which the Hamiltonian describing the system is recursively split into parts that evolve sufficiently slowly to be treated with a relatively long timestep, and faster parts that require shorter timesteps. This procedure effectively decouples the small-scale dynamics from the large-scale forces. The above canonical timestep criterion then yields a reasonable timestep for the smallest forming structures once it is applied (some steps down the hierarchy) only to the partial accelerations created by the high-resolution region itself.

A more subtle issue that becomes apparent with our very high dynamic range arises from the fact that force errors in our hierarchical multipole algorithm are spatially correlated. As a result, neighbouring particles normally have very similar node interaction lists. Formally, this creates force discontinuities across boundaries of the hierarchically nested cubes of the global octree geometry because the interaction lists and the field expansions (in the case of the fast multipole

method) change there. Small haloes, for which internal peculiar accelerations are small compared with that induced by large-scale structure, can be greatly affected by such errors if they are cut by an octree boundary corresponding to a geometrically large node. In such cases, the force error discontinuity can be appreciable relative to the peculiar acceleration. At high redshift, this error can build up over many timesteps if the halo is nearly at rest relative to the octree pattern. To alleviate such effects, we decorrelate these errors in time by translating the whole particle set by a random vector (drawn uniformly from the cubic volume) after every timestep. Physically, this does not change anything as the periodic system is translationally invariant. Numerically, it causes the above errors to average out in time, thereby preventing the build-up of sizeable momentum errors over many steps.

Convergence

A critical test of our numerical techniques is convergence in the properties of our simulations. We first examine maps of the mass distribution in common regions of adjacent zoom levels. As an example, in Extended Data Fig. 1 we compare projected density distributions in L1, L3 and L8c with the corresponding distributions in the same region of the parent level. It is clear that large-scale structure in the simulations is converged.

We next check the convergence properties of the halo mass function, again by comparing results for common regions of adjacent levels. Mass functions of haloes in spherical volumes of radius approximately 90% that of the entire high-resolution region are shown as solid curves in Extended Data Fig. 2, with different levels indicated by different colours. The mass functions of haloes in the same region in the parent simulations are shown as dotted curves.

The convergence of the halo mass functions in adjacent zoom levels is remarkable. Small differences appear at low masses when comparing simulations that resolve the free-streaming cut-off (L7c and L8c). These stem from the presence of spurious haloes that form due to discreteness effects when a cut-off in the power spectrum is resolved³⁷. The two vertical dotted lines in Extended Data Fig. 2 indicate the masses below which the abundance of these spurious haloes becomes important in the high-resolution regions of levels L7c and L8c³⁷. For these cases, convergence can be tested only to the right of the dotted lines and, as the figure shows, in this regime convergence is very good. Here we have only considered haloes in L7c and L8c with mass above these limits.

A convergence test of the internal structure of haloes is shown in Extended Data Fig. 3. Here we compare the density profile of one of the most massive haloes in a given level (solid lines) with the same halo in its parent level (dashed lines). The profiles of haloes from the parent simulation are plotted as thick solid lines in the radial range between the ‘convergence radius’³⁸ and r_{200} . The ratio of the density profiles of the matched pairs is plotted in the bottom panel of the figure. At radii larger than the convergence radius, the profiles agree to within a few per cent.

In summary, Extended Data Figs. 1, 2 and 3 show that in the regime where convergence can be tested, the spatial distribution, the abundance and the density profiles of haloes converge remarkably well over a factor of several hundred in mass resolution for all adjacent pairs of levels in our simulation.

Global properties of our simulation levels

Extended Data Table 1 lists a number of properties of the high-resolution regions at each level of our simulation: n_p is the number of high-resolution particles and m_p is the mass of each one, so that $M_{\text{tot}} = n_p m_p$ is the total mass of the high-resolution region. The r.m.s. linear fluctuation (extrapolated to $z=0$) expected in a spherical region that on average contains this mass is given as $\sigma(M_{\text{tot}}, z=0)$. For levels 2 and higher, this number exceeds unity, reaching 17.6 in L8c. As a result, the actual $z=0$ densities of the high-resolution regions, given as $\langle \rho \rangle / \rho_{\text{mean}}$, are small, typically a few per cent. This is expected since most

of the mass of the Universe is contained in ‘high-mass’ haloes that have, by construction, been excluded from the higher simulation levels¹². The fact that typical present-day haloes of very low mass (for example, Earth mass) form from Lagrangian regions with atypically low (linear) overdensity on larger scales (for example, $1-10^6 M_\odot$) explains why previous work has been unable to follow their evolution to low redshift¹⁶.

Density profiles

We selected only ‘relaxed’ or ‘equilibrium’ haloes defined as those that satisfy the following two criteria¹⁸: (i) the mass fraction in subhaloes within the virial radius is less than 0.1, and (ii) the offset between the centre of mass and the minimum of the potential is less than $0.07r_{200}$. As listed in Extended Data Table 1, more than 90% of our well resolved haloes satisfy these conditions at all levels (more than 95% at the highest levels). This can be understood as a result of the relatively high typical formation redshifts of the haloes (also listed in the table), although these are considerably lower than found in earlier work which was unable to follow such low-mass haloes to low redshift. Interestingly, we find the lowest-mass haloes to have lower typical formation redshifts than slightly more massive objects, consistent with the lower concentrations that we find below. We note also that we found no higher-mass particle within 0.95 of the radius of the high-resolution region at any level, and that extremely few well resolved haloes were contaminated by such particles. We excluded from our analysis any halo with such a particle within twice its virial radius, r_{200} .

To make mean mass-density profiles, we averaged binned mass densities in the radial range $(0.001-10)r_{200}$ for haloes lying between 0.8 and 1.2 times the central mass values listed in Extended Data Table 1. We then fitted NFW⁷ and Einasto^{5,8} formulae (equations (1) and (2)) to the stacked profiles using the bins between the ‘convergence’ radius³⁸ and r_{200} by minimizing the expression¹⁸:

$$\psi^2 = \frac{1}{N_{\text{bin}}} \sum_{i=1}^{N_{\text{bin}}} [\ln \rho_{\text{sim},i} - \ln \rho_{\text{fit},i}]^2, \quad (3)$$

where $\rho_{\text{sim},i}$ and $\rho_{\text{fit},i}$ are the simulation data and the fitted density profile in radial bin i . For the Einasto fits, the shape parameter α was set to 0.16 (ref. ³⁹) so that only two parameters are varied in both the NFW and Einasto fits.

The logarithmic slopes of the stacked radial density profiles of haloes are plotted in Fig. 2 out to large radii, $10r_{200}$. The ratios of the stacked profiles to the best-fit NFW and Einasto formulae are plotted in the lower panels of this figure and show that the NFW fits are almost everywhere accurate to better than 10%, and the Einasto fits to a few per cent.

As may be seen in Fig. 2, although Einasto profiles with $\alpha = 0.16$ fit the data well overall, for haloes with $M_{200} > 10^{12} M_\odot$ the fits have relatively large residuals. We therefore carried out Einasto fits to individual haloes with all three parameters free. The resulting median dependence of α on halo mass is well described by:

$$\alpha = 0.16 + 0.0238 \times (M_{200}/M_*)^{1/3} \quad (4)$$

where M_* is defined by $\sigma(M_*) = 1.68$, where $\sigma(M)$ is the r.m.s. linear fluctuation within a sphere that on average contains mass M . This extends previously published formulae³⁹ to much lower halo mass. For the Planck cosmology that we use here, $M_* = 1.14 \times 10^{14} M_\odot$. We now refit all haloes using for each an α value given by equation (4) and adjusting only the two remaining parameters. In this way, we obtain a robust estimate of the concentration–mass relation over the full halo mass range accessed by our simulation.

The resulting relation between M_{200} and $c_{\text{Einasto}} = r_{200}/r_{-2}$ is shown in Fig. 3. Simple extrapolations of empirical formulae derived for haloes of mass $M \geq 10^{10} M_\odot$ overestimate the concentrations of low-mass haloes ($M_{200} < 10^6 M_\odot$) by large factors. On the other hand, formulae derived from halo mass accretion histories^{9,10} match our data better over

the entire halo mass range, both with and without a free-streaming cut-off. We fit a simple parametrized form used previously^{19,40} to the median concentration–mass relation for levels L0 to L7 (that is, with no free-streaming cut-off), namely:

$$c_{\text{Einasto}}(M_{200}) = \sum_{i=0}^5 c_i \left[\ln \frac{M_{200}}{h^{-1} M_\odot} \right]^i \quad (5)$$

When the free-streaming cut-off is important (that is, for L7c and L8c), the concentration drops exponentially at the low-mass end and the relation is well fitted by:

$$c_{\text{Einasto}}(M_{200}) = \exp \left[c_6 \times \left(\frac{M_{\text{fs}}}{M_{200}} \right)^{1/3} \right] \times \sum_{i=0}^5 c_i \left[\ln \frac{M_{200}}{h^{-1} M_\odot} \right]^i. \quad (6)$$

In these relations, the c_i are dimensionless constants and the free-streaming mass scale is given by $M_{\text{fs}} = \frac{4\pi}{3} \times \left(\frac{2\pi}{k_{\text{fs}}} \right)^3 \times \rho_{\text{mean}}$, where k_{fs} is the free-streaming wave number defined by equation (3) of ref. ²⁷. For a thermal WIMP of mass 100 GeV, $M_{\text{fs}} = 7.3 \times 10^{-6} M_\odot$ and $k_{\text{fs}} = 1.77 \text{ pc}^{-1}$. We find the following best-fit values for the other parameters: $c_i = [27.112, -0.381, -1.853 \times 10^{-3}, -4.141 \times 10^{-4}, -4.334 \times 10^{-6}, 3.208 \times 10^{-7}, -0.529]$ for $i \in \{0, \dots, 6\}$.

Environmental dependence

Our strategy for simulating haloes over the entire mass range expected in a Λ CDM universe relies on successive resimulation of low-density regions. An important question is then whether the structure of these haloes is typical of the overall population. We can address this by examining how the concentration of haloes of a given mass varies with environment. We characterize the local environment of each halo by the mean density, $\langle \rho \rangle$, averaged over a surrounding shell with inner and outer radii 5 and 10 times the halo’s virial radius, r_{200} .

In Extended Data Fig. 4 we plot halo concentration as a function of $\langle \rho \rangle / \rho_{\text{crit}}$ for haloes averaged over mass bins in the range $(0.5-2) M_{\text{char}}$, where M_{char} , listed in Extended Data Table 1, is the typical mass of equilibrium haloes at each level. The white curves show median values and the surrounding shaded regions the r.m.s. scatter. Even though we focus on underdense regions, the density around haloes of mass below $10^{10} M_\odot$ is centred just below the mean density and that around more massive haloes is centred just above the mean density. Furthermore, the value of the environment density spans at least an order of magnitude, two orders of magnitude in the case of smaller mass haloes. Concentrations show no monotonic trend over this range, suggesting that the concentration–mass relation of Fig. 3 is representative of the halo population as a whole.

Impact of the free-streaming cut-off on halo structure

We can assess the effect of the free-streaming cut-off on the internal structure of individual haloes by comparing levels 7 and 7c. We do this in Extended Data Fig. 5, which shows density profiles for matched halo pairs in the two simulations. The haloes are matched by mass (mass difference less than 10%) and separation (offset less than 10% of the radius of the high-resolution region). Matched pairs were stacked in four different bins of L7 mass: $m_{200} = 5 \times 10^{-5} M_\odot$; $10^{-4} M_\odot$; $5 \times 10^{-4} M_\odot$; and $10^{-3} M_\odot$. The numbers of halo pairs in these bins are 152, 132, 40 and 24, respectively.

The effect of the free-streaming cut-off is to reduce the inner density (with a corresponding slight increase in the outer density) by an amount that grows as the free-streaming mass is approached. The net result is a progressive reduction in the concentration of haloes with decreasing mass, as may also be seen by comparing concentrations for haloes in levels L7 and L7c in Fig. 3. This effect reflects the later formation of haloes in L7c relative to their counterparts in L7.

Impact of the concentration–mass relation on annihilation luminosities

The annihilation luminosity from a dark matter halo scales as the square of the local dark matter density integrated over its volume. For smooth profiles of the kind that we fit to our simulation, almost all the luminosity comes from well inside the characteristic radius (r_s or r_{-2}) and scales with the characteristic parameters of the halo as V_{\max}^4/r_{\max} where V_{\max} and r_{\max} are the maximum circular velocity of the halo and the radius at which this is attained. The luminosity per unit mass of a halo thus scales as $V_{\max}^4/(r_{\max}M_{200})$ which, at given redshift, depends only on halo concentration, $c = r_{-2}/r_{200}$, for Einasto haloes with constant α ($= 0.16$ in our case).

For the well resolved haloes at each level of our simulation, we can measure V_{\max} and r_{\max} directly and so estimate their light-to-mass ratios. Averaging over all haloes of given mass and multiplying by their contribution to the mass density of the Universe according to some halo mass function (for illustration we here use the analytic Sheth–Tormen function⁴¹), we can construct Extended Data Fig. 6 which shows how the total annihilation luminosity of the present Universe is distributed over halo mass. Remarkably, we find the contribution per unit logarithmic halo mass interval to be almost constant over the range $-3 \leq \ln M_{200}/M_\odot \leq 11$. This is considerably less weighted towards low-mass haloes than estimates based on previously published extrapolations of the concentration–mass relation, and the total luminosity density is lower by factors ranging up to 10^3 . Thus, the importance of very small structures for annihilation luminosities has been overestimated in the past, often by substantial factors. Note that this observation is likely to apply to substructure within haloes as well as to emission from the main (smooth) halo, even though we have not considered such subhalo contributions here.

Data availability

The data generated, analysed and presented in this study are available from the corresponding authors on request, although the requester will be responsible for providing the very considerable resources needed for transferring and storing these data.

Code availability

The parent code GADGET-2 has been publicly available for some time. It is expected that most of the extensions and modifications made to meet the extreme requirements of this project will be made available in the future release of GADGET-4; those interested can contact V.S. (vspringel@MPA-Garching.mpg.de) for further information. The very large white-noise field, Panphasia, used to code our initial conditions, is also publicly available; those interested can contact A.J. (A.R.Jenkins@durham.ac.uk).

22. Ade, P. A. R. et al. (Planck Collaboration) Planck 2013 results. I. Overview of products and scientific results. *Astron. Astrophys.* **571**, A1 (2014).
23. Schaye, J. et al. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *Mon. Not. R. Astron. Soc.* **446**, 521–554 (2015).

24. Lewis, A., Challinor, A. & Lasenby, A. Efficient computation of cosmic microwave background anisotropies in closed Friedmann–Robertson–Walker Models. *Astrophys. J.* **538**, 473–476 (2000).
25. Bardeen, J. M., Bond, J. R., Kaiser, N. & Szalay, A. S. The statistics of peaks of Gaussian random fields. *Astrophys. J.* **304**, 15 (1986).
26. Springel, V. et al. Prospects for detecting supersymmetric dark matter in the Galactic halo. *Nature* **456**, 73–76 (2008).
27. Green, A. M., Hofmann, S. & Schwarz, D. J. The power spectrum of SUSY-CDM on subgalactic scales. *Mon. Not. R. Astron. Soc.* **353**, L23–L27 (2004).
28. Martínez, G. D. et al. Indirect dark matter detection from dwarf satellites: joint expectations from astrophysics and supersymmetry. *J. Cosmol. Astropart. Phys.* **06**, 014 (2009).
29. Bertone, G. & Tait, T. M. P. A new era in the search for dark matter. *Nature* **562**, 51–56 (2018).
30. Jenkins, A. Second-order Lagrangian perturbation theory initial conditions for resimulations. *Mon. Not. R. Astron. Soc.* **403**, 1859–1872 (2010).
31. Jenkins, A. A new way of setting the phases for cosmological multiscale Gaussian initial conditions. *Mon. Not. R. Astron. Soc.* **434**, 2094–2120 (2013).
32. Jenkins, A. & Booth, S. Panphasia: a user guide. Preprint at <https://arXiv.org/abs/1306.5771> (2013).
33. Springel, V. The cosmological simulation code GADGET-2. *Mon. Not. R. Astron. Soc.* **364**, 1105–1134 (2005).
34. Barnes, J. & Hut, P. A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature* **324**, 446–449 (1986).
35. Dehnen, W. A very fast and momentum-conserving tree code. *Astrophys. J. Lett.* **536**, L39–L42 (2000).
36. Pelupessy, F. I., Jänes, J. & Portegies Zwart, S. N -body integrators with individual time steps from hierarchical splitting. *New Astron.* **17**, 711–719 (2012).
37. Wang, J. & White, S. D. M. Discreteness effects in simulations of hot/warm dark matter. *Mon. Not. R. Astron. Soc.* **380**, 93–103 (2007).
38. Power, C. et al. The inner structure of Λ CDM haloes—I. A numerical convergence study. *Mon. Not. R. Astron. Soc.* **338**, 14–34 (2003).
39. Gao, L. et al. The redshift dependence of the structure of massive Λ cold dark matter haloes. *Mon. Not. R. Astron. Soc.* **387**, 536–544 (2008).
40. Lavalle, J., Yuan, Q., Maurin, D. & Bi, X. J. Full calculation of clumpiness boost factors for antimatter cosmic rays in the light of Λ CDM N -body simulation results. Abandoning hope in clumpiness enhancement? *Astron. Astrophys.* **479**, 427–452 (2008).
41. Sheth, R. K. & Tormen, G. An excursion set model of hierarchical clustering: ellipsoidal collapse and the moving barrier. *Mon. Not. R. Astron. Soc.* **329**, 61–75 (2002).

Acknowledgements This work was supported by the National Natural Science Foundation of China (NSFC) under grant 11988101, the National Key Program for Science and Technology Research Development (2017YFB0203300), and the UK Science and Technology Facilities Council (STFC) consolidated grant ST/P000541/1 to Durham University. J.W. and L.G. acknowledge support by NSFC grants 11373029, 11873051 and 11851301. S.B. acknowledges support from Harvard University through the ITC Fellowship. C.S.F. acknowledges support by the European Research Council (ERC) through Advanced Investigator grant DMIDAS (GA 786910). The visualizations in Fig. 1 were produced with the py-sphviewer code (<http://github.com/alejandrobl/py-sphviewer>). This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (<http://www.dirac.ac.uk>). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1 and ST/R002371/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

Author contributions J.W. contributed to the design of the project, carried out simulations, performed analysis and wrote the first draft of the paper. S.B. contributed to design of project, carried out simulations and performed analysis. C.S.F. contributed to the concept of project and its design, analysis of results and writing of final draft. A.J. contributed to the design of the project, upgraded initial conditions code and generated initial conditions, and contributed to analysis of results. L.G. contributed to the design of the project and analysis of results. V.S. wrote and upgraded simulation code, carried out simulations and contributed to analysis of results. S.D.M.W. initially conceived the project, and contributed to its design, analysis of results and writing of final draft.


Competing interests The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to J.W., C.S.F. or S.D.M.W.


Peer review information *Nature* thanks Mike Boylan-Kolchin and James Bullock for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>.




Extended Data Fig. 1 | Projected density maps for different zoom levels. a–f, Levels L1, L3 and L8c (a, c, e) are compared with maps of the same regions in their parent levels L0, L2 and L7 respectively (b, d, f). The regions shown are

the largest spheres that fit almost entirely within the high-resolution region of the higher level. Only high-resolution particles are used to make the images.




Extended Data Fig. 2 | The cumulative halo number as a function of mass, M_{200} , in the high-resolution region of each zoom level compared with that in the same region of the parent level. Different colours denote different levels as indicated in the legend. Results from the parent levels are shown as dotted curves. The two vertical black dotted lines indicate the upper mass limit

for spurious haloes in L7c and L8c, calculated as described in ref. ³⁷. Note the excellent agreement between the solid and dotted curves above the resolution limit of the latter (and above the L7c mass limit for spurious haloes in the case of L7c versus L8c).




Extended Data Fig. 3 | Comparison of the density profile of a massive halo at each level with that of its counterpart in the parent level. a, The density profile of one of the most massive haloes in the high-resolution region of each zoom level is compared with that of the same halo at the parent level. Results from different levels are shown with different colours, as indicated by the legend, which also gives the masses of the haloes concerned. Higher-resolution profiles are shown as dashed curves, while those from the parent levels are

shown as solid curves. The range between the convergence radius and r_{200} is plotted as a thick line in the lower-resolution case. **b,** The ratio of the density profiles of each pair in the upper panel. Again, results in the range between the convergence radius in the lower-resolution case and r_{200} are shown as thick lines. Note the excellent convergence between simulation pairs over this radial range, which typically differ in mass resolution by a factor of several hundred.




Extended Data Fig. 4 | Dependence of halo concentration on local environment in the high-resolution region at each zoom level. Results are shown for haloes in the mass range $[0.5, 2]M_{\text{char}}$; the key gives the characteristic mass, M_{char} , for each level and also defines the colour key. Each white curve gives the median concentration for the best-fit Einasto profile, while the surrounding coloured region gives the r.m.s. scatter. Local environment density is defined

here as the mean in a thick spherical shell, $5 < r/r_{200} < 10$, surrounding each halo, and is given in units of the critical density. All haloes are used for this plot. A vertical dotted line shows the cosmic mean density. Note that although concentration depends strongly on mass, any dependence on local environment density is weak.



Extended Data Fig. 5 | Stacked density profiles of matched haloes in two L7 simulations. The result for the L7 and the L7c are presented as solid and dashed curves individually. The densities are multiplied by r^2 to increase the dynamical range of the figure. Different colours correspond to different mass bins with central values quoted in the legend. The profiles are shown as thick lines over

the range where they are most reliable, between the convergence radius, r_{conv} and r_{200} . The vertical dotted line indicates the softening length in the high-resolution region at this level. The effect of the free-streaming cut-off is to reduce the density in the inner parts, and therefore the concentration, by an increasing amount as the halo mass approaches the free-streaming mass.



Extended Data Fig. 6 | The mean luminosity density produced by dark matter annihilation in today's Universe as a function of halo mass. Solid coloured lines indicate estimates obtained by multiplying the mean $\langle L/M \rangle$ obtained from well resolved, simulated haloes by the halo number density predicted for each bin by the Sheth-Tormen halo mass function⁴¹. Shaded regions indicate the estimated 1σ uncertainty. Dashed curves indicate results

found if we instead use $\langle L/M \rangle$ values predicted by previously published mass-concentration relations (light blue²¹, yellow¹⁰, red¹⁹, green²⁰, dark blue¹⁸). Corresponding predictions for the total integrated luminosity density are given in the key by values preceding each model's name. Note that the units of annihilation luminosity are arbitrary, so only ratios of values are meaningful.

Extended Data Table 1 | Parameters of the simulation levels

level	R_{high} [Mpc]	n_p	ϵ [kpc]	m_p [M_\odot]	$\sigma(M_{\text{tot}}, z = 0)$	$\langle \rho \rangle / \rho_{\text{mean}}$	M_{char} [M_\odot]	N_{char}	z_{form}	f_{vir}
L0	738	1.0×10^{10}	7.4	1.5×10^9		1.0	10^{14}	127	0.94	0.92
L1	52	1.0×10^{10}	4.4×10^{-1}	7.4×10^5	0.34	0.39	10^{12}	59	1.66	0.91
L2	8.8	5.4×10^9	5.6×10^{-2}	1.5×10^3	1.66	0.082	10^9	29	1.91	0.93
L3	1.0	1.8×10^9	8.3×10^{-3}	2.8	4.22	0.036	10^6	27	2.61	0.94
L4	0.27	2.0×10^9	1.0×10^{-3}	5.5×10^{-3}	6.96	0.026	10^3	59	4.44	0.94
L5	0.035	1.5×10^9	2.2×10^{-4}	5.8×10^{-5}	9.36	0.024	10	30	4.68	0.94
L6	0.0066	1.7×10^9	3.8×10^{-5}	2.6×10^{-7}	12.12	0.014	10^{-1}	35	4.84	0.94
L7	0.0011	2.5×10^9	5.3×10^{-6}	8.6×10^{-10}	15.06	0.016	10^{-4}	201	5.21	0.96
L7c	0.0011	2.5×10^9	5.3×10^{-6}	8.6×10^{-10}	15.06	0.016	10^{-4}	202	4.83	0.97
L8c	0.00024	1.5×10^9	1.4×10^{-6}	1.6×10^{-11}	17.60	0.028	10^{-6}	24	1.96	0.94

Column 1: name of the level; column 2: R_{high} , the radius of the high-resolution region; column 3: n_p , the total number of high-resolution particles; column 4: ϵ , the softening length of the high-resolution particles; column 5: m_p , the mass of the high-resolution particles; column 6: $\sigma(M_{\text{tot}}, z = 0)$, the r.m.s. linear overdensity, extrapolated to $z = 0$, within spheres that on average contain mass $M_{\text{tot}} = n_p m_p$; column 7: $\langle \rho \rangle / \rho_{\text{mean}}$, the mean mass density at $z = 0$ in the high-resolution region in units of the cosmic mean; column 8: M_{char} , the typical mass of the equilibrium haloes for which profiles were stacked in Fig. 2; column 9: N_{char} , the number of haloes in the mass bin $[0.8M_{\text{char}}, 1.2M_{\text{char}}]$ used in the stacks; column 10: z_{form} , the median formation redshift of equilibrium haloes of the characteristic mass; column 11: f_{vir} , the fraction of haloes with more than 3,000 particles that are in equilibrium according to the criteria given in the text.