The Abacus Cosmos: A Suite of Cosmological N -body Simulations

Lehman H. Garrison¹ , Daniel J. Eisenstein¹, Douglas Ferrer¹, Jeremy L. Tinker², Philip A. Pinto³, and David H. Weinberg^{4,5}

¹Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA 02138, USA

²Center for Cosmology and Particle Physics, New York University, 4 Washington Place, New York, NY 10003, USA

³Steward Observatory, University of Arizona, 933 N. Cherry Ave., Tucson, AZ 85121, USA

⁴Department of Astronomy, The Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA

⁵Center for Cosmology and AstroParticle Physics, The Ohio State University, Columbus, OH 43210, USA

Received 2017 December 14; revised 2018 March 27; accepted 2018 April 19; published 2018 June 5

Abstract

We present a public data release of halo catalogs from a suite of 125 cosmological N -body simulations from the ABACUS project. The simulations span 40 Λ CDM cosmologies centered on the Planck 2015 cosmology at two mass resolutions, $4 \times 10^{10} h^{-1} M_{\odot}$ and $1 \times 10^{10} h^{-1} M_{\odot}$, in $1.1 h^{-1}$ Gpc and $720 h^{-1}$ Mpc boxes, respectively. The boxes are phase-matched to suppress sample variance and isolate cosmology dependence. Additional volume is available via 16 boxes of fixed cosmology and varied phase; a few boxes of single-parameter excursions from Planck 2015 are also provided. Catalogs spanning $z = 1.5$ to 0.1 are available for friends-of-friends and ROCKSTAR halo finders and include particle subsamples. All data products are available at <https://lgarrison.github.io/AbacusCosmos>.

Key words: large-scale structure of universe – methods: numerical

1. Introduction

High-precision forward-modeling of large-scale structure is a necessary complement to upcoming galaxy surveys, including DESI (Levi et al. 2013), *WFIRST* (Spergel et al. 2015), *Euclid* (Laureijs et al. 2011), and LSST (LSST Science Collaboration et al. 2009), that will catalog tens of millions of galaxies in unprecedentedly large volumes. Mock catalogs must be available that allow designing of survey strategies and testing of cosmological parameter estimation and covariance. Although many fast approximate methods exist for generating such catalogs (e.g., COLA Tassev et al. 2013 and QPM White et al. 2014; see Monaco 2016 for a recent review), the highest-quality mocks remain those derived from N -body simulations. These gravity-only simulations directly evolve collisionless Monte Carlo tracers of the matter density field from the early, linear regime to late, nonlinear times when galaxies form. These simulations are fraught with challenges, from bias introduced by discretization at early times (e.g., Joyce & Marcos 2007; Garrison et al. 2016), to artificial relaxation at late times (e.g., Diemand et al. 2004; Power et al. 2016), to lack of hydrodynamical and baryonic physics (Mead et al. 2015; Schaller 2015; Schneider & Teyssier 2015, and references therein), to disagreements among N -body solvers (Schneider et al. 2016) and among halo finders (Knebe et al. 2013; Behroozi et al. 2015), to the sheer computational challenge of evolving $\mathcal{O}(10^{12})$ self-interacting particles, the scale that will be needed within a few years (Potter et al. 2016; Schneider et al. 2016). Nevertheless, N -body remains an invaluable tool in testing models of large-scale structure.

Given the computational expense of N -body, most public data products lie at one of two extremes: full halo catalogs available at a fixed cosmology, or reduced data products that allow variations in the cosmology. Some examples of the former are simulations like Bolshoi and MultiDark (Riebe et al. 2013), Las Damas (McBride et al. 2009), and Millennium (Lemson & Virgo Consortium 2006), which provide halo catalogs but focus on a single fiducial cosmology. Examples of the latter are tools like COSMICEMU (Lawrence et al. 2017) or

fitting formulae like those found in Tinker et al. (2008) or Comparat et al. (2017) for common data products like the power spectrum or halo mass function. These tools take cosmological parameters as inputs but do not provide access to the underlying halo catalogs or particles, so the science applications are limited. We aim to bridge these two extremes by providing a suite of many different cosmologies with access to the underlying halos and subsamples of the particles. In this sense, our approach is most similar to skiesanduniverses.org (Klypin et al. 2017) because the focus is access to raw halos and particles.

In the next section, we introduce the code used to run the simulations, and in Section 3 we discuss various parameter choices for the simulations and initial conditions. In Section 4, we introduce the design of the Latin hypercube grid of 40 cosmologies that form the core of our simulation effort, and in Section 5 we provide an overview of all the available sets of simulations. In Section 6, we detail the data products that we generate for each simulation, and in Section 7 we validate their accuracy. We summarize in Section 8.

2. ABACUS: Fast and Precise N -body Cosmology

2.1. Overview

The simulations in this work all employ ABACUS, an N -body code for cosmological simulations first described in Garrison et al. (2016). ABACUS is both extremely fast and accurate, capable of computing over 100 billion pairwise force interactions per second on a single computer node, with the option to compute forces to nearly machine precision while maintaining competitive speeds. It derives its performance from a combination of novel computational techniques and high-performance commodity hardware (in particular GPUs). ABACUS is not yet parallelized for distributed memory systems; each simulation here was run on a single node. ABACUS will be further described in D. Ferrer et al. (2018, in preparation) and M. Metchnik & P. A. Pinto (2018, in preparation); see also

Table 1
ABACUS Code Parameters, Described in Section 3.3

Parameter	Value
LagrangianPTOrder	2
Order	8
SofteningLength	63 or 41 kpc h^{-1}
TimeSliceRedshifts	1.5, 1.0, 0.7, 0.5, 0.3, [0.1]
TimeStepAccel	0.15
TimeStepDlna	0.03
ZD_PLT_target_z	5
ZD_qPLT	1
ZD_qPLT_rescale	1
Max (median) force error	1×10^{-4} (2×10^{-6})

Note. The force error is the maximum fractional error on the unsoftened forces on a set of 66 k randomly distributed particles compared to the true $1/r^2$ forces computed with an ewald summation in 256 bit precision. The Last Time-SliceRedshift is only present for the higher-resolution simulations.

Zhang et al. (2017) for a recent example of a massive, 130 billion particle simulation completed with ABACUS.

2.2. Algorithm

The computational domain in ABACUS is divided into a grid of CPD^3 cells, where CPD is the number of cells per dimension. The force computation is split into near-field and far-field components based on this cell decomposition. However, unlike most mesh-based N -body methods, such as P^3M , this decomposition is exact—the pairwise force between any two particles is always given exactly by either the near-field or far-field force.

Particles interact via the near-field force if they are within `NearFieldRadius` cells of one another. For example, `NearFieldRadius = 2` means a cell’s 124 nearest neighbor cells are included in the near-force calculation. We compute the force as a direct summation of $1/r^2$ forces (or some appropriately softened form; see below). This calculation can be accelerated with GPUs, which enables the impressive single-node performance.

Particles interact via the far-field force if they are separated by more than `NearFieldRadius` cells. ABACUS computes a multipole expansion of the particles in each cell using a variant of the fast multipole method (M. Metchnik & P. A. Pinto 2018, in preparation; see also Nitadori 2014 for a similar method). This multipole expansion is then convolved with a derivatives tensor to yield a set of coefficients for a Taylor-series expansion of the force in a given cell. The derivatives tensor is a fixed property of the grid for a given CPD, `NearFieldRadius`, and series expansion `Order`, and is thus precomputable. The multipole computation, derivatives convolution, and evaluation of the Taylor series happens for each timestep. Our performance tuning strategy is to change CPD to balance the cost of the near- and far-field computations, since the former can be offloaded to the GPUs while the CPUs are computing the latter.⁶ The far-field `Order` parameter largely sets the accuracy; we find `Order = 8` provides excellent accuracy and performance (see Table 1).

⁶ Although GPUs are well-suited to convolutions, we find that CPU convolutions are faster for our application when using modern CPUs with large SIMD vectors and a high-performance FFT library like Intel MKL. This is likely due to the data transfer overhead to the GPU and relatively low compute density (FLOPs per byte) of the operation.

The use of the fast multipole method on a discrete mesh enables our exact near-field/far-field force decomposition. Such an exact decomposition is also possible in tree codes—one can always choose to open a leaf (near-field) or leave it closed and use its multipoles (far-field)—but the key is that all possible vectors from one cell to another are precomputable if the multipoles are computed on a regular mesh instead of a tree. This allows one to solve the far-field as a convolution of mesh multipoles instead of explicit multipole interactions.

2.3. Hardware and Performance

Most of the simulations in this work were run on the El Gato cluster at the University of Arizona. The GPU nodes are dual-socket machines with two 8-core Intel Ivy Bridge E5-2650v2 processors (2.6 GHz) and 256 GB DDR3 RAM (1800 MHz). ABACUS is not designed to hold all particles in memory but rather stream them from a disk, so the I/O demands are fairly high. In many cases, we use hardware RAID, but for these simulations our strategy was to choose a problem size that would fit on a ramdisk (a filesystem hosted in RAM). A local scratch disk or network filesystem would typically not have the bandwidth to handle ABACUS’s I/O demands. Thus, the choice of 1440³ particles was set by the size of the system ramdisk.

Our typical speeds at the outset of this project were about 4 Mpart s^{-1} (million particle updates per second) per timestep, with about 20% of this time in the convolution step between each primary timestep; later code improvements increased this to about 12 Mpart s^{-1} . This higher speed is still a factor of two lower than typical speeds on our production machines (which were not used in this work), in part due to the extra load on the memory bandwidth from the ramdisk. The simulations took roughly 1000 timesteps to reach the final redshift from $z_{\text{init}} = 49$. The large number of timesteps is due to the lack of adaptive timestepping, so the global timestep is set by the shortest dynamical time in the whole box. This will be addressed with on-the-fly group finding and multi-stepping within those groups in future versions of ABACUS.

2.4. Force Softening

Several force softening options are available in ABACUS. The simplest is Plummer softening, where the $\mathbf{F}(\mathbf{r}) = \mathbf{r}/r^3$ force law is modified as

$$\mathbf{F}(\mathbf{r}) = \frac{\mathbf{r}}{(r^2 + \epsilon_p^2)^{3/2}}, \quad (1)$$

where ϵ_p is the softening length. This softening is very fast to compute but is not compact, meaning it never explicitly switches to the exact r^{-2} form at any radius (in contrast with spline softening). This modifies the growth of structure on large scales (Joyce & Marcos 2007). This is the softening we employ for the `emulator_1100box_planck` and `emulator_720box_planck` sets of simulations (see Section 5).

An alternative is spline softening, in which the force law is softened for small radii but explicitly changes to the unsoftened form at large radii. Traditional spline implementations split the force law into three or more piecewise segments (e.g., the cubic spline of Hernquist & Katz 1989); we split only once for computational efficiency⁷ and call this `single_spline`. We

⁷ Multiple splits can cause code-path branching, which incurs a significant performance penalty on CPUs and an even larger one on GPUs.

derive this form by considering a Taylor expansion in r of Plummer softening (Equation (1)) and requiring a smooth transition at the softening scale up to the second derivative.⁸ This gives

$$\mathbf{F}(\mathbf{r}) = \begin{cases} (10 - 15(r/\epsilon_s) + 6(r/\epsilon_s)^2)\mathbf{r}/\epsilon_s^3, & r < \epsilon_s; \\ \mathbf{r}/r^3, & r \geq \epsilon_s. \end{cases} \quad (2)$$

This is the softening we employ for the AbacusCosmos_1100box and AbacusCosmos_720box simulation sets.

The softening scales ϵ_s and ϵ_p imply different minimum dynamical times (an important property, as this sets the requisite temporal resolution to resolve orbits). We always set the softening length as if it were a Plummer softening and then internally convert to a softening length that gives the same minimum dynamical time for the chosen softening method. For single_spline, the conversion is $\epsilon_s = 2.16\epsilon_p$.

3. Simulation Details

We present technical details of the simulation configuration here. For an overview of the available sets of simulations, see Section 5.

3.1. Initial Conditions

The initial conditions were generated by the public zeldovich-PLT⁹ code of Garrison et al. (2016). We do not provide initial conditions files with the catalogs, but we do provide the input parameter file (info/abacus.par) for the IC code and the input power spectrum from CAMB (Lewis et al. 2000). The initial conditions can thus be generated by re-running the IC code with those inputs.

The simulations use second-order Lagrangian perturbation theory (2LPT) initial conditions, but zeldovich-PLT only outputs first-order displacements. The 2LPT corrections are generated by ABACUS on-the-fly using the configuration-space method of Garrison et al. (2016).

Two non-standard first-order corrections are implemented by zeldovich-PLT. The first is that the displacements use the particle lattice eigenmodes rather than the curl-free continuum eigenmodes. This eliminates transients that arise due to the discretization of the continuum dynamical system (the Vlasov-Boltzmann distribution function) into particles on small scales near k_{Nyquist} . The second correction is “rescaling,” in which initial mode amplitudes are adjusted to counteract the violation of linear theory that inevitably happens on small scales in particle systems. This violation usually takes the form of growth suppression; thus the initial adjustments are mostly amplitude increases. We choose $z_{\text{target}} = 5$ as the redshift at which the rescaled solution will match linear theory; this choice is tested in Garrison et al. (2016).

Later, we will refer to simulations that are “phase-matched” in the initial conditions. This refers to initial conditions with the same random number generator seed, called ZD_Seed in the abacus.par file. Matching this value (and ZD_NumBlock) between two simulations guarantees that the amplitudes and phases of the initial modes are identical between the

simulations (up to differences in the input power spectrum and cosmology).

3.2. Input Power Spectrum

We use CAMB (Lewis et al. 2000) to generate a linear $z = 0$ power spectrum for each cosmology in our grid. We then scale the power spectrum back to $z_{\text{init}} = 49$ by scaling σ_8 by the ratio of the growth factors $D(z = 49)/D(z = 0)$. This σ_8 is passed to zeldovich-PLT, which handles the re-normalization of the power spectrum. The computation of the growth factors is done by ABACUS’s cosmology module, so it is consistent by construction with the simulation’s cosmological evolution. We only use massless neutrinos and include no cosmological neutrino density. The exact CAMB inputs and outputs are available with each simulation.

In the computation of the power spectrum, baryons and CDM are treated as separate species; however, ABACUS is a gravity-only N -body solver (i.e., no hydrodynamics or baryonic physics), so we simply combine the baryonic and CDM density when computing the overall mass density of the universe. It is this combined mass density that sets our particle mass.

3.3. Code Parameters

The important ABACUS parameters (including those that might affect code accuracy) are given in Table 1 and described here. The simulations were run with a mix of force softening laws. The “ZD” parameter prefix indicates that this is an input to our zeldovich-PLT IC code.

LagrangianPTOrder: the order of Lagrangian perturbation theory corrections to compute at runtime (see Section 3.1).

Order: the multipole order used for computing the far-field force.

SofteningLength: the comoving Plummer-equivalent force softening length (see Section 2.4).

TimeSliceRedshifts: the output redshifts. The last slice ($z = 0.1$) is only available for the higher-resolution “720box” simulations.

TimeStepAccel: parameter to limit the timestep based on particle accelerations. The timestep is set such that $a_{\text{max}}\Delta t/v_{\text{rms}}$ is never larger than this parameter. This often sets the choice of timestep at late times. Various choices of this parameter are tested in D. Ferrer et al. (2018, in preparation); 0.15 is considered a conservative choice.

TimeStepDlna: maximum $\Delta(\ln a)$ allowed for a timestep. This often sets the choice of timestep at early times. For example, 0.03 means at least 33 steps per e -folding of the scale factor.

ZD_PLT_target_z: the target redshift for PLT rescaling of the initial conditions (see Section 3.1 for a description of PLT corrections).

ZD_qPLT: initialize the displacements and velocities in the eigenmodes of the particle system.

ZD_qPLT_rescale: do PLT rescaling.

4. Cosmology Grid Design

The 40 cosmologies listed in Table 2 are distributed in a six-dimensional Λ CDM parameter space according to a Latin hypercube algorithm (see Figure 1 for a visual representation).

⁸ A Taylor expansion in r^2 is also possible, but we discard that solution due to a large plateau of constant angular frequency near $r \sim 0$ that could excite dynamical instabilities.

⁹ <https://github.com/lgarrison/zeldovich-PLT>

Table 2

Cosmologies for the AbacusCosmos Sets of Simulations (Both Resolutions)

	H_0	Ω_{DE}	Ω_M	n_s	σ_8	w_0
00	69	0.698	0.302	0.93	0.854	-1.14
01	63	0.669	0.331	0.982	0.719	-0.765
02	72.3	0.739	0.261	0.97	0.851	-1.08
03	66.2	0.671	0.329	0.975	0.858	-0.982
04	74.2	0.733	0.267	0.954	0.889	-1.22
05	71.5	0.706	0.294	0.954	0.913	-1.29
06	68.6	0.717	0.283	0.96	0.768	-0.969
07	66.7	0.709	0.291	0.987	0.734	-0.788
08	65.8	0.695	0.305	0.957	0.692	-0.796
09	72.7	0.719	0.281	0.931	0.89	-1.24
10	67	0.677	0.323	0.961	0.83	-0.995
11	72.2	0.732	0.268	0.963	0.851	-1.16
12	65.2	0.687	0.313	0.99	0.778	-0.827
13	64.1	0.67	0.33	0.976	0.776	-0.792
14	74.8	0.731	0.269	0.971	0.999	-1.37
15	67.4	0.708	0.292	0.976	0.732	-0.89
16	70.9	0.727	0.273	0.969	0.835	-0.999
17	62.1	0.658	0.342	0.95	0.714	-0.745
18	67.9	0.693	0.307	0.97	0.831	-0.934
19	71	0.717	0.283	0.955	0.849	-1.04
<hr/>						
	H_0	Ω_{DE}	Ω_M	n_s	σ_8	w_0
20	64.5	0.675	0.325	0.967	0.74	-0.742
21	68.3	0.689	0.311	0.975	0.847	-1.02
22	73.3	0.723	0.277	0.937	0.923	-1.3
23	62.7	0.645	0.355	0.965	0.77	-0.796
24	73.9	0.731	0.269	0.966	0.938	-1.24
25	61.6	0.633	0.367	0.965	0.728	-0.754
26	69.3	0.705	0.295	0.963	0.831	-1.06
27	62.5	0.663	0.337	0.959	0.687	-0.655
28	65.5	0.676	0.324	0.937	0.795	-0.941
29	67.9	0.684	0.316	0.935	0.875	-1.1
30	64.6	0.674	0.326	0.956	0.735	-0.87
31	69.8	0.702	0.298	0.977	0.89	-1.14
32	63.3	0.681	0.319	0.984	0.647	-0.661
33	66.1	0.686	0.314	0.982	0.744	-0.866
34	70.1	0.692	0.308	0.941	0.906	-1.22
35	73.2	0.714	0.286	0.962	0.979	-1.35
36	63.8	0.659	0.341	0.963	0.738	-0.762
37	70.5	0.727	0.273	0.973	0.825	-1.04
38	74.5	0.747	0.253	0.951	0.886	-1.2
39	71.9	0.726	0.274	0.953	0.881	-1.16

Note. The cosmologies were chosen by a latin hypercube algorithm centered on the Planck 2013 cosmology; see Section 4.

Specifically, the algorithm samples a $(H_0, \Omega_M h^2, \Omega_b h^2, \sigma_8, n_s, w_0)$ space centered on the Planck 2013 cosmology (Planck Collaboration et al. 2014). We use $N_{\text{eff}} = 3.046$ in all cases. These 40 cosmologies are realized in the two AbacusCosmos simulation sets.

The goal of the cosmological parameter selection is to evenly span the parameter space with only a limited number of cosmologies. The procedure used here follows the Latin hypercube method described in Heitmann et al. (2009). In the Latin hypercube method, each of the N dimensions is divided into M bins, with M being the number of desired cosmologies. Each bin in each dimension is only sampled once, thus guaranteeing that the entire range of parameter space is covered in the set of cosmologies. The hypercube is optimized such that, for each cosmology, the distance to the nearest neighboring cosmology is maximized.

The optimized hypercube is then rotated into the parameter space defined by the union of WMAP9 year (Hinshaw et al. 2013) and Planck 2013 (Planck Collaboration et al. 2014) CMB results, combined with recent BAO and SN results. The results used to define this space are obtained from Anderson et al. (2014). The axes of the Latin hypercube then correspond to the eigenvectors of the CMB-defined parameter space. In the original hypercube design, each axis ranges from 0 to 1. In the CMB parameter space, these ranges correspond to -4 to 4 times the eigenvalue along each eigenvector, so we are sampling from the 4σ CMB constraints.

Although the Planck 2013 results were used in the hypercube design instead of the 2015 results, the resulting cosmologies span a larger space than allowed by either data set. Thus, derivatives measured from these simulations should be equally useful when assuming a fiducial cosmology of either Planck 2013 or 2015. Indeed, our fiducial cosmology for the emulator_planck simulations is Planck 2015, and it falls nearly at the center of the cosmology hypercube (the blue square in Figure 1).

5. Catalogs

We present four collections of simulations organized into “sets.” Sets have names like AbacusCosmos_1100box, while individual simulations have a two-digit number appended to them like AbacusCosmos_1100box_00. The two-digit number refers to the cosmology. Different phases of the same cosmology are indicated by a dash and number after the cosmology, as in emulator_1100box_planck_00-1.

The four sets of simulations are as follows:

AbacusCosmos_1100box: 41 phase-matched¹⁰ boxes: 40 spanning the six-dimensional Λ CDM cosmology parameter space (Section 4) and one with the fiducial Planck cosmology. Each box has a size of $1100 h^{-1}$ Mpc and particle mass $4 \times 10^{10} h^{-1} M_\odot$. The spline softening length is $63 h^{-1}$ Kpc.

AbacusCosmos_720box: same as the above, but at a higher mass resolution ($1 \times 10^{10} h^{-1} M_\odot$) and smaller box size ($720 h^{-1}$ Mpc). Note that these are not zoom-in simulations of the larger boxes, but independent realizations of the power spectrum. The spline softening length is $41 h^{-1}$ kpc.

emulator_1100box_planck: 21 boxes, 16 with identical Planck cosmologies but different IC phases, and 5 phase-matched boxes of single-parameter variations from the Planck cosmology ($\pm 5\%$ in σ_8 and H_0 , plus a central Planck box; see Table 3). All boxes have the same size as above $1100 h^{-1}$ Mpc and particle mass $4 \times 10^{10} h^{-1} M_\odot$. The Plummer softening length is $63 h^{-1}$ Kpc.

emulator_720box_planck: same as the above, but at a higher mass resolution ($1 \times 10^{10} h^{-1} M_\odot$) and smaller box size ($720 h^{-1}$ Mpc). As with the AbacusCosmos boxes, these are not zoom-in simulations. The Plummer softening length is $41 h^{-1}$ kpc.

The 40 AbacusCosmos simulations are designed to allow estimation of derivatives of cosmological measurables with respect to cosmology. They can either be used as an ensemble to construct an emulator/interpolator, or each individual box

¹⁰ “Phase-matched” means the initial conditions random number generator was given the same seed. This ensures differences between boxes are due to cosmology and not cosmic variance; see Section 3.1.

Figure 1. Corner plot representation of the cosmology space spanned by the AbacusCosmos simulations, where we have combined Ω_{CDM} and Ω_b into Ω_M . The blue square marks the fiducial central cosmology.

Table 3

Cosmologies for the emulator_1100box_planck and emulator_720box_planck Sets of Simulations, All with $n_s = 0.965$, $w_0 = -1$, and $N_{\text{eff}} = 3.04$

	H_0	Ω_{DE}	Ω_M	σ_8
00 and 00-0 to 00-15	67.3	0.686	0.314	0.83
01	67.3	0.686	0.314	0.78
02	67.3	0.686	0.314	0.88
03	64.3	0.656	0.344	0.83
04	70.3	0.712	0.288	0.83

Note. The first cosmology represents fiducial parameters for which 17 boxes with different phases were run. The latter four represent “derivative” boxes in which one parameter (in bold) is changed at a time. Note that the cosmologies are actually chosen in the space of physical densities $\Omega_i h^2$, which is why a change in H_0 results in change in Ω_i .

can be differenced with the central AbacusCosmos_planck box to provide an estimate of the derivative for that particular change in cosmology.

The boxes do not provide much cosmic volume of any individual cosmology ($1.7 h^{-3} \text{Gpc}^3$ if both resolutions are combined), but the 17 phase-varied emulator_planck boxes provide additional volume and a path toward suppressing cosmic variance and estimating covariance. The four emulator_planck single-parameter excursion boxes provide a more direct route to measuring derivatives but only for two parameters.

The motivation for two mass resolutions was first to provide a convergence test for large-scale structure properties, at least in the intermediate regime well-sampled by both resolutions (see Section 7.2). Second, the larger boxes provide the volume that is needed for BAO-type studies (Klypin & Prada 2017 argue that modes longer than $1 h^{-1} \text{Gpc}$ have very little impact on the matter power spectrum), while the smaller boxes provide the halo resolution that is needed by weak-lensing studies (see, e.g., Wibking et al. 2017).

Our softening lengths—41 and $63 h^{-1} \text{Kpc}$ in the two box resolutions—were chosen to support halos, not subhalos. Future enhancements to ABACUS should make it possible for

us to reach dramatically smaller softening lengths. In the near term, we are continuing to run simulations at these mass scales and resolutions to build volume. These simulations will be made available on the release website as they finish.

In total, these 125 simulations represent roughly 200K CPU-hours, or 25K GPU-hours, of computational effort. This is a relatively modest amount for a collection of 370 billion particles and is a testament to ABACUS’ computational efficiency.

6. Data Products: Halos and Power Spectra

We provide three data products at every redshift slice: friends-of-friends (FoF) halos, ROCKSTAR halos, and a high-resolution matter power spectrum. The redshift slices are $z = \{1.5, 1.0, 0.7, 0.5, 0.3, [0.1]\}$, where the $z = 0.1$ slice is only provided for the higher-resolution “720box” simulations. Particle subsamples are included with the catalogs. The data formats of each product are described on the data release website.

6.1. Friends-of-friends

The FoF algorithm links particles separated by less than a linking length b , equal to 0.186 in our catalogs (expressed as a fraction of the mean particle spacing). A halo is defined as a set of linked particles (Davis et al. 1985). Our implementation is based on the University of Washington N -body Shop’s halo finder, and we include halos down to 25 particles. The linking length 0.186 was chosen to correspond to an overdensity contour of 100 times the background density using the percolation theory results of More et al. (2011).

We compute a number of halo properties relative to the FoF centers of mass and velocity, including velocity dispersion, circular velocity profiles, and radial quantiles. However, the center of mass is susceptible to “barbell” pathologies, in which two largely distinct halos are connected by a chance alignment of a thin particle “bridge.” To guard against this, we also compute a second level of FoF with a smaller linking length (equal to 0.117 in our catalogs). This linking length was chosen to capture half of the mass of a singular isothermal sphere. We store the masses of the four most massive subhalos and additionally compute global halo properties centered on the most massive subhalo. These allow for a first-order defense against common FoF pathologies.

We provide a 10% subsample of particles both inside and outside halos (“halo particles” and “field particles”). Particles are assigned to be subsample particles as a pseudo-random function of the particle ID, so the same 10% of particles are output at every timestep. This allows for construction of crude merger trees and other simple box-to-box comparisons. A uniform sampling of the density field (that is, a 10% sample of all particles) can be formed via the union of the halo and field particle subsamples. Particle information includes positions, velocities, and IDs.

6.2. ROCKSTAR

ROCKSTAR (Behroozi et al. 2013) is a hierarchical halo finder that uses FoF in six-dimensional phase space at successively smaller linking lengths to find halos and substructure. The innermost substructure defines a halo seed to which particles are assigned based on their phase-space

proximity. ROCKSTAR also can use temporal information (multiple time slices) to improve structure tracking, but we do not use this mode as the time between our outputs is large. We run ROCKSTAR with mostly default settings, including the default mass definition of M_{vir} ; however, we report both regular ROCKSTAR masses and strict spherical overdensity masses. Strict spherical overdensity masses contain all particles, including those considered “unbound” and those not associated with the halo. Subhalos are reported and tagged with their parent halo as decided by ROCKSTAR. We also output a 10% subsample of particles in halos. The exact ROCKSTAR input file is available with each catalog (`rockstar.cfg`).

6.3. Power Spectra

We compute the matter power spectra by gridding the particles onto a mesh (2048^3 or finer) with triangle-shaped cloud (TSC) mass assignment. We then Fourier transform the density field, convert the result to a power spectrum, deconvolve the TSC-aliased window function from Jeong (2010), and bin in spherical annuli. The resulting 1D power spectrum is available for every simulation time slice.

6.4. Plummer versus Spline Data Products

The two AbacusCosmos sets were run with spline softening, while the `emulator_planck` boxes were run with Plummer softening (Section 2.4). The spline was developed partway through the simulation campaign and was thus only applied to the remaining sets of simulations. However, we also re-ran one of the Plummer boxes with spline softening to calibrate differences in the data products between spline and Plummer. Those results are presented here. The simulation is also available on the data release website as `emulator_1100box_planck_spline_00` if further calibrations are required.

In Figure 2, we show halo mass functions from the FoF halo finder for Plummer and spline. The differences are small ($<3\%$ for halos above 100 particles), but the number of Plummer halos steadily increases with increasing mass (except for the highest mass bin, which has very few halos). The Plummer softening may be inflating halos, causing large halos to come in contact with neighboring structures, increasing their mass. At the low-mass end, inflating small, tenuously bound halos (below 100 particles) may be causing them to unbind, resulting in fewer halos. These trends are almost completely insensitive to redshift. ROCKSTAR and ROCKSTAR SO halos behave very similarly at the low-mass end, but Plummer underpredicts the number of >1000 particle halos by 1%–3%. This is consistent with the picture that high-mass halos are inflated with Plummer softening, since ROCKSTAR is less susceptible than FoF to over-merging of inflated structures.

In Figure 3, we show matter power spectra for Plummer and spline softening. As expected, Plummer results in a loss of power at high k . At the Nyquist wavenumber of the particle lattice, Plummer misses 6% of power compared to spline. This loss of power due to non-compact gravitational softening is a known phenomenon; see, e.g., Joyce & Marcos (2007), Garrison et al. (2016). As with halo mass, there is almost no evolution of this relation with redshift.

Figure 2. Comparison of the FoF halo mass function of two identical simulations that differ only in the force softening technique (Plummer or spline). For halos larger than 100 particles, the differences are consistently small ($<3\%$), although the relative number of Plummer halos steadily increases with increasing mass. There is almost no evolution of this relation with redshift.

Figure 3. Comparison of the matter power spectra of two identical simulations that differ only in the force softening technique (Plummer vs. spline). As expected, the Plummer simulation is missing power at high k (6% at k_{Nyquist}) due to the long tail of the Plummer force softening law.

6.5. Example Python Interfaces

Example Python code is provided on the release website to load and manipulate the halo catalogs, including particle subsamples. The data formats are also documented on the website, allowing any user to write their own code to parse the catalogs. However, the provided Python code can at least serve as an implementation example of the data specifications.

Figure 4. Comparison of the $z = 0.5$ power spectra from the AbacusCosmos_1100box simulations to the COSMICEMU power spectrum emulator (Lawrence et al. 2017). The comparison is shown for the 28 cosmologies that fall within the COSMICEMU domain. The shaded bar shows the 1% error region. Each line represents a simulation; the colors have no meaning beyond distinguishing the lines. The overall agreement is very good; the low- k differences are due to cosmic variance, while the high- k differences are due to the finite resolution of the simulations.

A wrapper that loads the halo catalogs into HALOTOOLS¹¹ format (Hearin et al. 2017) is also provided.

Another halo occupation distribution (HOD) code that is well-integrated with the Abacus Cosmos simulation suite is GRAND-HOD.¹² In particular, GRAND-HOD can use the halo particle subsamples when populating a halo with satellite galaxies.

7. Validation

7.1. COSMICEMU and HALOFIT

We validate our power spectrum results against COSMICEMU (Lawrence et al. 2017) which produces a power spectrum as a function of cosmology. However, the ABACUSCOSMOS cosmologies span a slightly larger parameter space than COSMICEMU, so we are only able to do this test for 28 of our 40 sims.

Figure 4 shows that AbacusCosmos and COSMICEMU are in very good agreement (at the level of a few percent) for $k \sim 0.03$ to 4 at $z = 0.5$. At the low- k end, we see large deviations due to cosmic variance, or finite box size.¹³ At the high- k end, we see a downturn in the ABACUS power spectra past the Nyquist wavenumber of the particle lattice. This downturn is expected due to the failure of particle systems to reproduce linear theory near, and especially past, k_{Nyquist} (Joyce & Marcos 2007; Garrison et al. 2016).

¹¹ <http://halotools.readthedocs.io>

¹² <https://github.com/SandyYuan/GRAND-HOD>

¹³ For statistics that need more cosmological volume, this low- k scatter can be averaged down with the emulator_planck sims, which have 17 boxes of the same cosmology but different phases. See Figure 5.

Figure 5. Comparison of the $z=0.3$ power spectra from the `emulator_720box_planck` simulations to the COSMICEMU (Lawrence et al. 2017) and HALOFIT (Takahashi et al. 2012) power spectrum emulators. The inner shaded bar shows 1% agreement, while the outer bar shows the 4% accuracy quoted by COSMICEMU. Each ABACUS line represents a simulation, all of which have the same cosmology but different initial condition phases. The solid black line is the average of the ABACUS lines. The overall agreement is very good among ABACUS, COSMICEMU, and HALOFIT. The low- k scatter due to cosmic variance is suppressed when averaging over multiple ABACUS simulations, while the high- k differences are due to the finite resolution of the simulations and Plummer softening.

Figure 5 shows the agreement of the `emulator_planck` simulations with COSMICEMU and HALOFIT (Takahashi et al. 2012), as invoked through CAMB’s `do_nonlinear` feature at $z = 0.3$. In particular, it shows that the low- k scatter in the power spectrum seen in Figure 4 can be suppressed by averaging the results of many boxes (recall that these boxes have fixed cosmology but different IC phases). The agreement is within 4% (the accuracy quoted by COSMICEMU) for $k \sim 0.01$ to 3, with a downturn before k_{Nyquist} due to the Plummer softening (see also Figure 3). The `emulator_720-box_planck_00-0` simulation shows a small systematic deviation around $k = 1$ from the mean ABACUS result, but we find no evidence for misbehavior of this simulation in our diagnostics.

7.2. Convergence

We compare simulation data products in the intermediate regime well-sampled by both `emulator_720box_planck` and `emulator_1100box_planck`. We examine the power spectrum and halo mass function at three different redshifts. In all cases, we average over all 17 boxes at each resolution to suppress sample variance.

The matter power spectrum agreement in Figure 6 is excellent from $k = 0.03$ to $k = 6$ (the Nyquist wavenumber of the higher-resolution box), indicating that our results are stable with respect to box size and mass resolution. To compare the matter power spectra at incommensurate wavenumbers, we applied a cubic spline in log-log space to the reference spectrum. However, at the smallest k the binning breaks the

Figure 6. Comparison of the matter power spectrum at the two mass resolutions/box sizes. In the intermediate k regime well-sampled by both resolutions, we expect the results to be converged, since we are averaging over 17 boxes to suppress sample variance. Indeed, we find excellent agreement of $\sim 1\%$ over a wide range of k . The agreement changes by less than a percentage point from $z = 0.7$ to 0.3.

smoothness of the power spectrum, so the cubic spline appears to slightly overpredict the disagreement between the resolutions. Regardless, agreement in this regime is cosmic-variance-limited even after averaging over 17 boxes, so this discrepancy is not concerning.

The FoF halo mass function is converged to within 6% in the regime sampled by at least 100 particles at both mass resolutions (Figure 7). The lower mass resolution systematically overproduces small halos; this is a known effect in FoF that arises from the mismatch in spatial stochasticity in the particle sampling at the two resolutions (More et al. 2011). See Figure 10 of (Garrison et al. 2016) for a very similar test, in which downsampling the higher-resolution simulation before running FoF produces excellent agreement with the lower-resolution result. With ROCKSTAR, the downturn effect almost entirely disappears (Figure 8).

When considering ROCKSTAR SO masses, however, a substantial deviation is seen for all but the largest halos (Figure 9). At 100 particles (400 particles at the higher resolution), 30% more halos are found at the higher resolution. We speculate that the higher-resolution box finds more physically small halos (due to mass and softening resolution) near large halos; these small halos then have their SO masses inflated by the presence of nearby structure, since ROCKSTAR allows overlapping SO spheres and “double-counting” of mass.

8. Summary

We have presented a suite of cosmological N -body simulations produced by the new ABACUS code. The modest computational requirements of ABACUS (a single GPU node for a few days) enabled us to run one-hundred twenty-five ~ 1 Gpc boxes, each with 3 billion particles. These boxes span 40 cosmologies near Planck 2015, allowing for emulation/

Figure 7. A comparison of the FoF halo mass function at the two mass resolutions/box sizes. The lower mass resolution box overpredicts the number of 100-particle halos by 5% compared to 400-particle halos at the higher resolution. This is an effect of friends-of-friends and is not seen with ROCKSTAR; see Section 7.2.

Figure 8. Same as Figure 7, but for ROCKSTAR halos.

interpolation in this important parameter region. The accompanying halo catalogs include particle subsamples, allowing for detailed investigations of galaxy bias models and other effects sensitive to dark matter halo structure, like weak lensing. The data products are publicly available and include example code to manipulate the catalogs.

We thank the Research Computing Group in the FAS Division of Science at Harvard University for their assistance

Figure 9. Same as Figure 7, but for ROCKSTAR halos with spherical overdensity masses.

in hosting the catalogs. L.H.G. thanks Ben Wibking, Andres Salcedo, and Sihan Yuan for their feedback as “alpha” testers of these catalogs, Nina Maksimova for helpful discussions on the Abacus algorithm, and the referee for comments that helped improve the quality of the work. Our FoF implementation is based on the University of Washington *N*-body Shop friends-of-friends code. This work has been supported by grant AST-1313285 from the National Science Foundation and by grant DE-SC0013718 from the U.S. Department of Energy. Some of the computations used in this study were performed on the El Gato supercomputer at the University of Arizona, supported by grant 1228509 from the National Science Foundation. P.A.P. was supported by NSF AST-1312699.

ORCID iDs

Lehman H. Garrison <https://orcid.org/0000-0002-9853-5673>

David H. Weinberg <https://orcid.org/0000-0001-7775-7261>

References

- Anderson, L., Aubourg, É., Bailey, S., et al. 2014, *MNRAS*, **441**, 24
- Behroozi, P., Knebe, A., Pearce, F. R., et al. 2015, *MNRAS*, **454**, 3020
- Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013, *ApJ*, **762**, 109
- Comparat, J., Prada, F., Yepes, G., & Klypin, A. 2017, arXiv:1702.01628
- Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, *ApJ*, **292**, 371
- Diemand, J., Moore, B., Stadel, J., & Kazantzidis, S. 2004, *MNRAS*, **348**, 977
- Garrison, L. H., Eisenstein, D. J., Ferrer, D., Metchnik, M. V., & Pinto, P. A. 2016, *MNRAS*, **461**, 4125
- Hearin, A. P., Campbell, D., Tollerud, E., et al. 2017, *AJ*, **154**, 190
- Heitmann, K., Higdon, D., White, M., et al. 2009, *ApJ*, **705**, 156
- Hernquist, L., & Katz, N. 1989, *ApJS*, **70**, 419
- Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, *ApJS*, **208**, 19
- Jeong, D. 2010, PhD thesis, Univ. Texas at Austin
- Joyce, M., & Marcos, B. 2007, *PhRvD*, **76**, 103505
- Klypin, A., & Prada, F. 2017, arXiv:1701.05690
- Klypin, A., Prada, F., & Comparat, J. 2017, arXiv:1711.01453
- Knebe, A., Pearce, F. R., Lux, H., et al. 2013, *MNRAS*, **435**, 1618
- Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv:1110.3193
- Lawrence, E., Heitmann, K., Kwan, J., et al. 2017, arXiv:1705.03388

- Lemson, G. & Virgo Consortium 2006, arXiv:[astro-ph/0608019](#)
- Levi, M., Bebek, C., Beers, T., et al. 2013, arXiv:[1308.0847](#)
- Lewis, A., Challinor, A., & Lasenby, A. 2000, [ApJ](#), **538**, 473
- LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:[0912.0201](#)
- McBride, C., Berlind, A., Scoccimarro, R., et al. 2009, *BAAS*, **41**, 253
- Mead, A. J., Peacock, J. A., Heymans, C., Joudaki, S., & Heavens, A. F. 2015, [MNRAS](#), **454**, 1958
- Monaco, P. 2016, [Galax](#), **4**, 53
- More, S., Kravtsov, A. V., Dalal, N., & Gottlöber, S. 2011, [ApJS](#), **195**, 4
- Nitadori, K. 2014, arXiv:[1409.5981](#)
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, [A&A](#), **571**, A16
- Potter, D., Stadel, J., & Teyssier, R. 2016, arXiv:[1609.08621](#)
- Power, C., Robotham, A. S. G., Obreschkow, D., Hobbs, A., & Lewis, G. F. 2016, [MNRAS](#), **462**, 474
- Riebe, K., Partl, A. M., Enke, H., et al. 2013, [AN](#), **334**, 691
- Schaller, M. 2015, PhD thesis, Durham Univ.
- Schneider, A., & Teyssier, R. 2015, [JCAP](#), **12**, 049
- Schneider, A., Teyssier, R., Potter, D., et al. 2016, [JCAP](#), **4**, 047
- Spiegel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv:[1503.03757](#)
- Takahashi, R., Sato, M., Nishimichi, T., Taruya, A., & Oguri, M. 2012, [ApJ](#), **761**, 152
- Tassev, S., Zaldarriaga, M., & Eisenstein, D. J. 2013, [JCAP](#), **6**, 036
- Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, [ApJ](#), **688**, 709
- White, M., Tinker, J. L., & McBride, C. K. 2014, [MNRAS](#), **437**, 2594
- Wibking, B. D., Salcedo, A. N., Weinberg, D. H., et al. 2017, arXiv:[1709.07099](#)
- Zhang, H., Eisenstein, D. J., Garrison, L. H., & Ferrer, D. W. 2017, *ApJ*, submitted (arXiv:[1712.05787](#))