

Indra: a public computationally accessible suite of cosmological N -body simulations

Bridget Falck,¹★ Jie Wang,^{2,3} Adrian Jenkins^{id},⁴ Gerard Lemson,¹ Dmitry Medvedev,¹ Mark C. Neyrinck^{id},^{5,6} and Alex S. Szalay¹

¹*Department of Physics and Astronomy, Johns Hopkins University, 3400 N Charles St, Baltimore, MD 21218, USA*

²*National Astronomical Observatories, Chinese Academy of Sciences, Beijing 100012, China*

³*University of Chinese Academy of Sciences, School of Astronomy and Space Science, Beijing 100049, China*

⁴*Department of Physics, Institute for Computational Cosmology, Durham University, Durham DH1 3LE, UK*

⁵*Ikerbasque, the Basque Foundation for Science, E-48009, Bilbao, Spain*

⁶*Department of Physics, University of the Basque Country UPV/EHU, E-48080 Bilbao, Spain*

Accepted 2021 June 21. Received 2021 June 18; in original form 2021 January 10

ABSTRACT

Indra is a suite of large-volume cosmological N -body simulations with the goal of providing excellent statistics of the large-scale features of the distribution of dark matter. Each of the 384 simulations is computed with the same cosmological parameters and different initial phases, with 1024^3 dark matter particles in a box of length $1 h^{-1}$ Gpc, 64 snapshots of particle data and halo catalogues, and 505 time-steps of the Fourier modes of the density field, amounting to almost a petabyte of data. All of the Indra data are immediately available for analysis via the SciServer science platform, which provides interactive and batch computing modes, personal data storage, and other hosted data sets such as the Millennium simulations and many astronomical surveys. We present the Indra simulations, describe the data products and how to access them, and measure ensemble averages, variances, and covariances of the matter power spectrum, the matter correlation function, and the halo mass function to demonstrate the types of computations that Indra enables. We hope that Indra will be both a resource for large-scale structure research and a demonstration of how to make very large data sets public and computationally accessible.

Key words: methods: numerical – dark matter – large-scale structure of Universe.

1 INTRODUCTION

Understanding the observations made by large-scale structure surveys such as *Euclid*, Roman, DESI, and the VRO requires input from theoretical predictions of structure formation in the form of numerical simulations. Cosmological N -body simulations solve for the non-linear gravitational collapse of matter and are indispensable tools both for testing cosmological models and for planning for, and analysing, observations. For example, simulations are necessary to predict the large-scale structure observables of dark energy and modified gravity theories in the non-linear regime (e.g. Winther et al. 2015; Joyce, Lombriser & Schmidt 2016), which can be compared to simulations of the current Λ CDM (cold dark matter with a cosmological constant, Λ) cosmological paradigm. Simulations have also become an integral part of the analysis of measurements of baryon acoustic oscillations (BAO) from galaxy redshift surveys (e.g. Percival et al. 2014).

One of the pressing challenges when making inferences from large observational surveys is that we can only observe the Universe from one vantage point in a finite volume. This ‘cosmic variance’ places limits on the precision with which we can measure statistical fluctuations on very large scales. This also presents a problem for

simulations: we cannot simulate exactly the positions of the galaxies in this observational volume. We can, however, simulate an ensemble of universes, increasing the precision with which we can numerically predict the statistical properties (e.g. the covariance matrix) of the matter distribution on large scales for a given cosmological model.

Ideally, one would want to simulate the full volume and resolution that the next generation of galaxy surveys will attain, repeat this thousands of times for different realizations of the same cosmological model and parameters, run this ensemble for thousands of combinations of cosmological parameters in a given model, and repeat the entire process for any number of cosmological models one expects the observations to test. Such a scenario has obvious drawbacks: it would require a prohibitively expensive amount of computing power and produce an excessive amount of data which would be impossible to analyse without equally massive computing resources. Because of these immense technical challenges, workarounds are being devised: cosmological emulators (Heitmann et al. 2014, 2016; Garrison et al. 2018; DeRose et al. 2019; Nishimichi et al. 2019) attempt to span cosmological parameter space with a few full simulations and interpolate between them, low-resolution particle mesh simulations attempt to capture features only at large scales (Takahashi et al. 2009; Blot et al. 2015), approximate simulations combine numerical shortcuts with analytic solutions to produce fast and cheap ensembles of realizations with a limited range of accuracy (see Chuang et al. 2015; Lippich et al. 2019, and references therein), and analytic

* E-mail: bridget.falck@jhu.edu

methods attempt to model the statistics in order to reduce the number of full realizations required to obtain the same precision (Pope & Szapudi 2008; Dodelson & Schneider 2013; Taylor, Joachimi & Kitching 2013; Heavens et al. 2017).

This paper aims to address the technical challenges head-on by producing a suite of cosmological N -body simulations called Indra¹ hosted on SciServer,² a science platform that allows the community to perform analysis where the data are stored. This method of *server-side analysis* was pioneered in astronomy by the Sloan Digital Sky Survey SkyServer (Szalay et al. 2000), which enabled astronomers to interactively explore a multiterabyte data base by leveraging state-of-the-art archiving technologies. Modelled on the success of the SkyServer, the Millennium Simulation Data base (Lemson & Virgo Consortium 2006) made available the post-processing outputs of the Millennium simulation (Springel et al. 2005), including halo catalogues, merger trees, and later, semi-analytic galaxy catalogues (see also Riebe et al. 2011; Loebman et al. 2014; Bernyk et al. 2016). The dark matter particle positions and velocities, however, were not in the data base, as their large size (roughly 20 TB) and usage patterns (being point clouds instead of object catalogues) presented significant technical challenges.

Though data base technology has seen significant advances in the past decade, our early efforts to build a relational data base for Indra were motivated by reducing the number of rows in a table – the planned full suite of 512 simulations amounts to 35 trillion output particles. We considered storing particles in spatially sorted array-like ‘chunks’ (Dobos et al. 2012) and developed an inverted indexing scheme (Crankshaw et al. 2013) to access individual particles in these arrays (e.g. to track their movement across many snapshots). However, many technical advances opened up the capability to release data in their native binary file format, without needing to first load them into relational data bases, by mounting the data to contained compute environments running on virtual machines. Indra takes advantage of the SciServer science platform, built at Johns Hopkins University where the data are stored, which provides interactive and batch mode server-side analysis in various environments, public and private file storage in both flat files and data bases, and many collaboration tools (Taghizadeh-Popp et al. 2020).

Recent advances have seen some tera- and peta-scale cosmological simulation suites that provide their data for download, but with no server-side computation (Blot et al. 2015; Heitmann et al. 2019; Villaescusa-Navarro et al. 2019). Additionally, the latest version of the Illustris hydrodynamical simulations (Nelson et al. 2019) provides both data for download and an interactive compute environment similar to SciServer. Modern astronomical surveys are developing infrastructures and tools for the analysis and cross-matching of data sets that are hundreds of terabytes to petabytes in size (Jurić et al. 2017). The server-side analysis of both astronomical surveys and the theoretical predictions from simulations will soon become common practice for science involving very large data sets.

The Indra simulations are motivated by the goal of obtaining excellent statistics of the large-scale matter density field on $100 h^{-1}$ Mpc scales. We designed a suite of 512 (of which 384 are or will be currently available) realizations of $1 h^{-1}$ Gpc-sided cosmological N -body simulations, each with 1024^3 dark matter particles. We saved 64 snapshots per simulation, many more than other simulation suites, so that ensemble statistics may be calculated for a large number of

redshifts and merger histories of dark matter haloes are able to be built. We describe the simulations and their data products in Section 2, including details on how the initial conditions were created and an issue that affected the first 128 simulations. In Section 3, we measure ensemble averages, variances, and covariances from the three Indra data products: the dark matter particles, the halo catalogues, and the Fourier modes of the density field. In the appendix, we describe how the full suite of simulations can be accessed and analysed by the community.

2 THE INDRA SIMULATIONS

2.1 Overview

The Indra suite of simulations were run using the code L-GADGET2 (Springel 2005), each with a box length of $1 h^{-1}$ Gpc, 1024^3 particles, and WMAP7 cosmological parameters ($\Omega_M = 0.272$, $\Omega_\Lambda = 0.728$, $h = 0.704$, $\sigma_8 = 0.81$, $n_s = 0.967$; Komatsu et al. 2011). They are purely gravitational dark-matter-only cosmological simulations; each dark matter particle has a mass of $m_p = 7.031 \times 10^{10} h^{-1} M_\odot$. The force softening length is $40 h^{-1}$ kpc, and the parameters used for the force accuracy, time integration, and tree criteria are the same as used for the Millennium Run (Springel et al. 2005), which used the same lean version of GADGET. Initial conditions were generated using second-order Lagrangian perturbation theory (2LPT; Scoccimarro 1998; Jenkins 2010), with pre-initial conditions on a regular grid and a starting redshift of $z = 127$.³ A total of 64 snapshots are saved down to $z = 0$ at the same redshifts as the Millennium Run. In general the box size, particle number, and other parameters of Indra are geared toward resolving haloes more massive than that of the Milky Way and obtaining good statistics on BAO scales, while saving snapshots at a large number of redshifts and not exceeding a petabyte of data. Since Indra is intended to be a public resource, we hope it can serve a variety of cosmological applications.

The Indra simulations can be identified using three digits that range from 0 to 7, which we denote as $x.y.z$, corresponding to the coordinates of an $8 \times 8 \times 8$ cube. Note, however, that each simulation is fully independent and obeys periodic boundary conditions in a $1 (h^{-1} \text{Gpc})^3$ volume; see Section 2.3 for details about how the initial phases were set up. Due to a bug in the initial conditions of the first 128 simulations, denoted $0.y.z$ and $1.y.z$, these are not available, though we may re-run them in the future.

Halo catalogues were generated on-the-fly using a standard friends-of-friends (FOF) algorithm (Davis et al. 1985), with a linking length of $b = 0.2$ times the interparticle separation, L/N , and a minimum of 20 particles per halo. After the simulations were finished, we ran SUBFIND (Springel et al. 2001) to identify subhaloes in phase-space and calculate halo properties such as R_{200} and M_{200} .

The Indra simulations may be accessed via SciServer, a science platform that allows users to perform analysis where the data are stored without the need to download terabytes to a local compute cluster (Taghizadeh-Popp et al. 2020). Further details of data access

³Since the inception of this project over a decade ago, there has been much progress in understanding the scale- and resolution-dependent effects of initial conditions generation and starting redshift in cosmological simulations. In particular, Michaux et al. (2021) quantify the discreteness effects of the initial lattice in simulations with the same box size and resolution as Indra. The difference between their reference ‘FCC’ simulation and Indra’s choice of $z_{\text{init}} = 127$ with 2LPT varies with redshift, scale, and choice of statistic: for the power spectrum (their fig. 7), from a few per cent at $z = 3$ to 1 per cent at $z = 0$ and $k = 1 h \text{Mpc}^{-1}$.

¹‘Indra’ refers to the Buddhist metaphor of Indra’s net, in which each of the infinite jewels reflects every other jewel.

²<http://www.sciserver.org>

are given in Appendix A. Below, we describe the three types of data output for each of the Indra simulations: A total of 64 snapshots of particle positions and velocities, 64 snapshots of halo catalogues, and 505 snapshots of the Fourier modes of the density field. We then describe in detail the initial conditions of the simulations, which were created using Panphasia (Jenkins 2013) to allow higher resolution re-simulations of interesting regions.

At the time of this writing, we are finishing up the last series of 64 simulation runs and post-processing, and we will make these available as they are ready. Therefore, the analysis performed in this work is done on 320 Indra simulations: 2_0_0 to 6_7_7.

2.2 Data products

The output of the Indra simulations consists of three types of data products: the positions, velocities, and identifiers (IDs) of dark matter particles; halo and subhalo catalogues, along with the IDs of their constituent particles; and the Fourier modes of the coarse-gridded density field. The data are permanently stored on the DataScope, developed by the Institute for Data Intensive Engineering and Science⁴ at the Johns Hopkins University. The DataScope is a multipetabyte file server with high throughput capability that hosts large data sets across many domains; Indra data are accessed from SciServer compute containers over a fast network connection. A subset of the data, including all halo catalogue files and FFT outputs and some of the particle snapshots, is also stored on a distributed file system called the FileDB, which allows for more computationally heavy analysis than the default SciServer Compute environment. Read times are somewhat faster for the data on FileDB, and a system for parallel computation using Dask⁵ is in development. All data are accessed through SciServer using the INDRA-TOOLS software package.⁶ Both interactive Compute and batch Jobs run Docker containers on virtual machines, and the INDRA-TOOLS software, whether run from a PYTHON script or Jupyter notebook, hides the filesystem from the user, choosing to read from FileDB if the requested data are there and from DataScope if not. Instructions on how to access the full suite of Indra data, as well as a description of the example Jupyter notebooks included with INDRA-TOOLS, are provided in Appendix A.

2.2.1 Dark matter particles

The bulk of the petabyte of Indra data are in the form of particle positions, velocities, and IDs for 64 snapshots of each simulation. These are stored as binary files in the standard GADGET format. The entirety of the particle snapshots are stored on the DataScope, and a subset of the particle data, consisting of a few full simulations and several full sets of snapshots, are also stored on FileDB servers. This subset might be slightly faster to work with and currently includes five full simulations (2_0_0, 3_0_0, 4_0_0, 5_0_0, and 6_0_0) and nine full sets of snapshots at redshifts 0, 0.1, 0.5, 1, 1.4, 1.7, 2.0, 2.4, and 127.

The snapshots of particle data are immediately available for analysis in the SciServer Compute environment. As an example, we present a visualization of the particle data in Fig. 1, which shows the density field of slices through two of the 512 Indra simulations at $z = 1.7$ and $z = 0$. The density fields were measured from interpolations

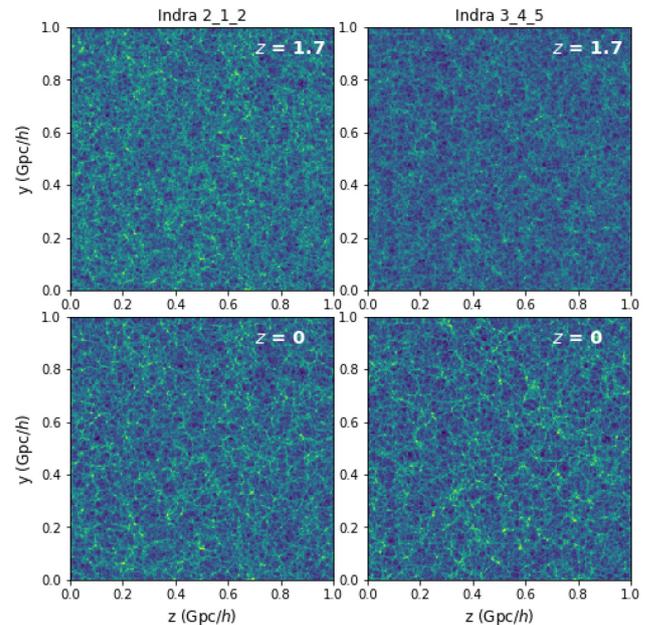


Figure 1. Slices of the logarithmically scaled density field from 2 of the Indra simulations at 2 of the 64 snapshots, $z = 1.7$ and $z = 0$, measured from CIC-interpolations of the 1024^3 particle positions. Slices are 1 cell-length thick with $3.9 h^{-1}$ Mpc grid cells.

of particle positions on to a grid using the cloud-in-cell (CIC) assignment scheme. One CIC grid of a 1024^3 -particle snapshot can be computed in a few minutes using the package PMESH (Feng, Hand & biweidai 2017). We discuss preliminary analysis of the particle data, including power spectrum and correlation function covariances, in Section 3.1.

2.2.2 Dark matter haloes

The halo data consist of FOF group tables computed on-the-fly, SUBFIND halo catalogues computed after the runs have finished, and the IDs of the constituent particles of both the FOF and SUBFIND haloes. The FOF data are minimal and contain only the number of particles in each halo, along with indexing information with which to pick out the correct particle IDs from the group ID data. Post-processing of the haloes, including calculation of the mass, radius, position, velocity, and other halo properties is performed by SUBFIND and stored in the subhalo tables. SUBFIND identifies one subhalo as the ‘main’ subhalo of that parent FOF group and measures quantities such as mass and radius for these haloes, so in this way the subhalo catalogue contains properties both of parent FOF groups and of their subhaloes.

For every FOF group, SUBFIND measures the radius within which the FOF group has an overdensity 200 times the *mean* density of the simulation, the mass within this $R_{200,\text{mean}}$, the radius within which the FOF group has an overdensity 200 times the *critical* density of the simulation, the mass within this $R_{200,\text{crit}}$, the radius within which the FOF group has an overdensity corresponding to the value at virialization in the top-hat collapse model for this cosmology, and the mass within this $R_{\text{top-hat}}$. For every subhalo, including the ‘main’ subhalo, SUBFIND measures the 3D coordinates of position as determined by the particle with the minimum gravitational potential, the 3D coordinates of peculiar velocity, the 1D velocity dispersion, the maximum of the circular velocity curve, the 3D coordinates of

⁴<http://idies.jhu.edu/what-we-offer/sciserver/datascope/>

⁵<http://dask.org/>

⁶<http://github.com/bfalck/indra-tools>

halo spin, and the radius containing half of the mass. In the binary catalogue data, all mass units are in $h^{-1} M_{\odot}$, radius and position units in $h^{-1} \text{Mpc}$, velocity units in km s^{-1} , and spin units in $h^{-1} \text{Mpc km s}^{-1}$. In the data base, the units are the same except the masses are in units of $10^{10} h^{-1} M_{\odot}$.

The FOF and SUBFIND catalogues of the full suite of Indra simulations, for every snapshot, are stored in their binary file format in both the DataScope and FileDB filesystems. Additionally, they are loaded in to relational data base tables, which can be queried directly from compute containers on SciServer. The usefulness of hosting halo catalogues in a relational data base was successfully demonstrated by the Millennium Simulation data base (Lemson & Virgo Consortium 2006). They provide the capability of speeding up analysis in many cases, for example, selecting haloes by mass from many snapshots and hundreds of Indra volumes without the need to read all of the binary catalogue files. The Indra data base also exploits a spatial indexing library developed for numerical simulations (Lemson, Budavari & Szalay 2011) that enables efficient selection of haloes or particles within shapes such as spheres or cones.

The suite of Indra subhalo data, saved both in binary files and in relational data base tables, is a data-rich resource for studies of large-scale structure. For example, having 64 snapshots for each simulation allows the construction of halo merger histories, which is not possible when only a few snapshots are stored. In addition, since all particle positions and velocities are saved, any other halo finder may be run to complement the existing FOF and SUBFIND catalogues, and new halo properties may also be measured from the constituent particles. We discuss preliminary analysis of the halo data, including the covariance of the mass function, in Section 3.2.

2.2.3 Fourier modes of the density

The Fourier-space density field is output as the simulation runs, more frequently than are the particle snapshots, for a total of 505 time-steps. The real and complex modes are defined on a $129 \times 129 \times 65$ grid of (k_x, k_y, k_z) , where k_x and k_y range from 0 to $\pm 0.4 h \text{Mpc}^{-1}$ and $k_z \geq 0$, which is sometimes denoted the ‘upper half-sphere’ of k -space. The largest wavenumber is thus $|\mathbf{k}|_{\text{max}} = 0.7 h \text{Mpc}^{-1}$ and correspondingly, the smallest length-scale in the FFT data are $9 h^{-1} \text{Mpc}$. All of the FFT data, for every run and all 505 time-steps, are stored on both the DataScope and FileDB filesystems and amount to about 2 TB. In Section 3.3, we measure the distribution functions of the Fourier amplitudes of the density field and of the mode-dependent growth function, $D(a, k)$.

2.3 Initial conditions

The initial conditions are generated using the IC_2lpt_Gen code (Jenkins 2010). Each Indra volume is a realization of a LCDM universe with cubic periodic boundary conditions. To make the realizations independent and representative of the set of all possible volumes, we set the phases using the Panphasia multiscale realization of a Gaussian white noise field (Jenkins 2013; Jenkins & Booth 2013).⁷ Each Indra volume is assigned its own small cubic white noise field patch within the very large Panphasia field. This individual white noise field is then convolved in k -space with an appropriate real non-negative filter to produce a realization of the LCDM power spectrum. The so-called k -space ‘corner modes’, i.e. the power at

⁷<http://icc.dur.ac.uk/Panphasia.php>

modes larger than the 1D particle Nyquist frequency, are not set to zero (Falck et al. 2017).

The Panphasia realization is completely defined on all scales, which means once the cubic patch in the Panphasia field has been chosen, the phases for that Indra volume are known on all scales below the box scale – even down to scales that could only be accessed by future zoom simulations in these volumes.

For reference, we will give the locations of the phase information within the Panphasia field in this section. The Panphasia field has an octree structure, so to describe a cubic region it is necessary to define five numbers: the level in the octree, three coordinates for the location of one corner of the cube, and the side-length of that cube. Jenkins (2013) defined a fixed format text descriptor to hold these five pieces of information, in addition to a sixth check digit, for error checking purposes only, and the name for the volume. This text descriptor is used as an input by the IC_2lpt_Gen code to set the phases for a particular periodic volume. The descriptor, together with the cosmological parameters, box size, and the linear power spectrum, is all that is needed to set-up the Indra volumes themselves. Any future zoom simulations of any region in any of the Indra volumes, down to the putative CDM free streaming scale as need be, use the Panphasia descriptor for the parent volume.

The phases of all of the Indra volumes can be given as a set of descriptors that are functions of the three-digit identifiers x_y_z :

[Panph1, L15, (X, Y, Z), S75, CH – 999, INDRA_xyz],

where x, y, z are octal integers, each in the range of 0–7, which specify an Indra volume. The symbols, X, Y, Z, give the locations of one corner of each cubic patch, where

$$\begin{aligned} X &= 31248 + 100 * x, \\ Y &= 31376 + 100 * y, \\ Z &= 31504 + 100 * z. \end{aligned} \quad (1)$$

We have set the check digits always to be –999 (which instructs the IC_2lpt_Gen code not to do the error checking). So, for example, the descriptor for the Indra_555 volume is

[Panph1, L15, (31748, 31876, 32004), S75, CH–999, INDRA_555].

Although the Indra volumes are all close together in the Panphasia field, they are disjoint, which for a Gaussian white noise field means they are completely independent. Because of this, they cannot be mapped directly to an $(8 h^{-1} \text{Gpc})^3$ volume; in principle it would be possible to set up a realization of a roughly $10.7 h^{-1} \text{Gpc}$ -sided cube that includes all of the Indra volume phase information. While individual patches of such a simulation would resemble individual Indra volumes to some extent, there would be clear differences on the largest scales as the periodic constraints are very different. In addition, for individual Indra volumes, the value of the integral of the white noise field over the entire simulation volume is ignored, as all Indra volumes must have exactly the mean density of the universe. In a larger volume this information would be retained, contributing to the largest scale modes of the periodic volume.

3 ENSEMBLE AVERAGES, VARIANCES, AND COVARIANCES

One major motivation for creating the Indra simulations is to enable precise measurements of ensemble averages, variances, and covariances at very large scales from full, not approximate, cosmological N -body simulations. In this section, we demonstrate the types of

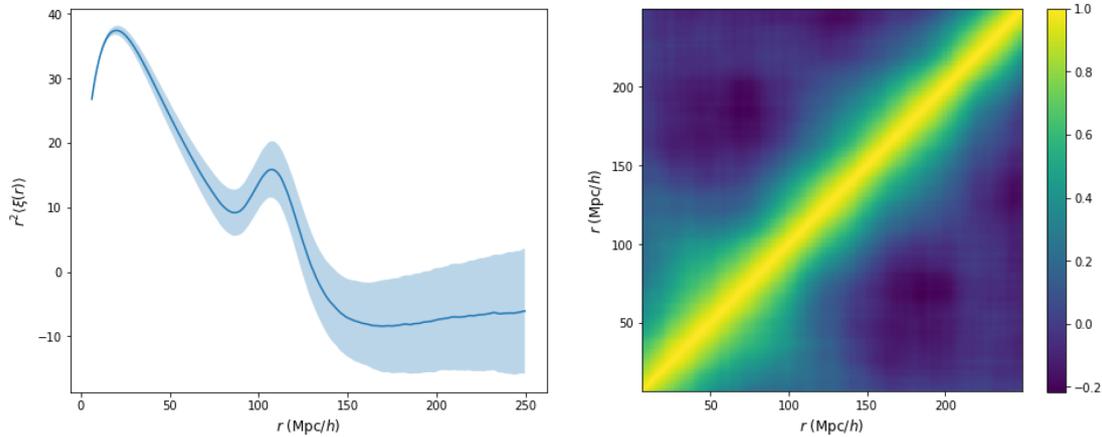


Figure 2. *Left:* The $z = 0$ ensemble average correlation function measured from 320 Indra simulations. The shaded region indicates the standard deviation of the mean. *Right:* The correlation matrix of the same 320 correlation functions.

computations that Indra enables on each of the three main data products. Section 3.1 measures the mean, variance, and covariance of the matter correlation function and power spectrum from the dark matter particle positions. Section 3.2 focuses on the mass function of the dark matter haloes. In Section 3.3, we study the Fourier modes of the density field and measure the distribution function, and its evolution, of the Fourier amplitudes as well as the mode-dependent growth function $D(a, k)$.

3.1 Matter power spectrum and correlation function

The matter density contrast, $\delta = (\rho - \bar{\rho})/\bar{\rho}$, is a function of both space and time. Its spatial distribution can be summarized in terms of its two-point statistics in the form of the autocorrelation function, $\xi(\mathbf{r}) = \langle \delta(\mathbf{r}')\delta(\mathbf{r}' + \mathbf{r}) \rangle$, or its Fourier dual, the power spectrum, $P(k) \equiv \langle |\delta_k|^2 \rangle$. These statistical descriptions completely specify a Gaussian random field, such as the density field in the very early Universe (Bardeen et al. 1986). The covariance matrix, \mathbf{C} , of a statistic y is given by

$$C_{ij} \equiv \langle (y_i - \langle y_i \rangle)(y_j - \langle y_j \rangle) \rangle, \quad (2)$$

where i and j are bins of either r (for the correlation function covariance) or k (for the power spectrum covariance).

Constraining cosmological parameters with large-scale structure surveys requires measuring the covariance matrix from a large number of mock galaxy catalogues (see e.g. Percival et al. 2014; Lippich et al. 2019), perhaps in combination with analytic methods (e.g. Pope & Szapudi 2008). The error in covariance estimation due to the finite number of realizations propagates to errors in parameter constraints (Dodelson & Schneider 2013; Taylor et al. 2013; Blot et al. 2016; Sellentin & Heavens 2016). Additionally, the mass resolution of the N -body simulation has a systematic affect on the covariance (Blot et al. 2015). The distribution of power spectra departs from Gaussianity even at large scales (Takahashi et al. 2009; Blot et al. 2015), necessitating the use of N -body simulations to study these effects.

We measure the correlation function and power spectrum of the Indra simulations by first interpolating the particle positions on to a grid using a cloud-in-cell (CIC) assignment scheme, which results in a measure of the density field in voxels of constant volume. This is then Fourier-transformed to obtain the Fourier modes of the density field, δ_k , from which we compute the power spectrum and correlation function.

The correlation functions are measured from a CIC density grid of 1024^3 cells (with a Nyquist frequency of $3.2 h \text{ Mpc}^{-1}$). Fig. 2 shows the $z = 0$ mean correlation function, and its variance, from 320 Indra simulations. The covariance of the correlation function is shown in the right-hand panel of the figure in terms of its Pearson correlation matrix:

$$R_{ij} = \frac{C_{ij}}{\sqrt{C_{ii}C_{jj}}}. \quad (3)$$

Fig. 3 shows the $z = 0$ mean matter power spectrum and its variance, as well as the correlation matrix, of the dark matter particles of 320 Indra simulations, measured from a CIC density grid of 512^3 cells (corresponding to a Nyquist frequency of $1.6 h \text{ Mpc}^{-1}$). The CIC window function was deconvolved in the power spectrum calculation, but no shot-noise subtraction or anti-aliasing scheme was performed (see e.g. Jing 2005).

Because we measure the correlation functions and power spectra over the full periodic simulation volumes, there are no modes present larger than our ‘survey window’, so the covariance matrices do not contain the effects of beat-coupling (Hamilton, Rimes & Scoccimarro 2006) or local averaging (de Putter et al. 2012). An observational galaxy survey would measure these effects in its covariance matrix since modes are present in the Universe on scales larger than any survey volume. However, de Putter et al. (2012) found that these effects mostly cancel each other, and additionally for a $(1 h^{-1} \text{ Gpc})^3$ volume, the correction to the power spectrum due to the local average effect is small, $\lesssim 10^{-4}$.

The Poisson error on the power spectra for 320 Indra simulations, given by $1/\sqrt{N(k)} * 320$ for $N(k)$ modes at wavenumber k , is 2.5 per cent at the largest scales where $k \sim 0.01 h \text{ Mpc}^{-1}$ and 0.3 per cent at $0.06 h \text{ Mpc}^{-1}$, roughly the BAO scale. The Indra simulations provide an excellent N -body benchmark for detailed studies of the matter power spectrum and correlation function covariance using analytic and approximate methods.

3.2 Halo mass function

In the context of understanding or analysing large-scale galaxy surveys, the halo catalogues are of primary importance. Methods that aim to speed up N -body simulations by introducing approximations or simplifications of the fully non-linear gravitational collapse focus on how well they can reproduce the spatial distribution, masses, and key characteristics of the dark matter haloes identified in full

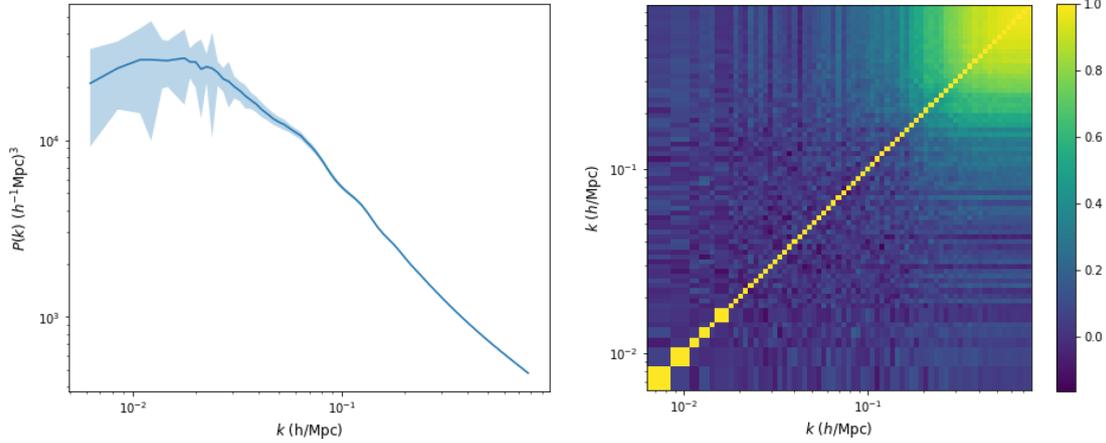


Figure 3. *Left:* The $z = 0$ ensemble average power spectrum measured from 320 Indra simulations. The shaded region indicates the standard deviation of the mean. *Right:* The correlation matrix of the same 320 power spectra.

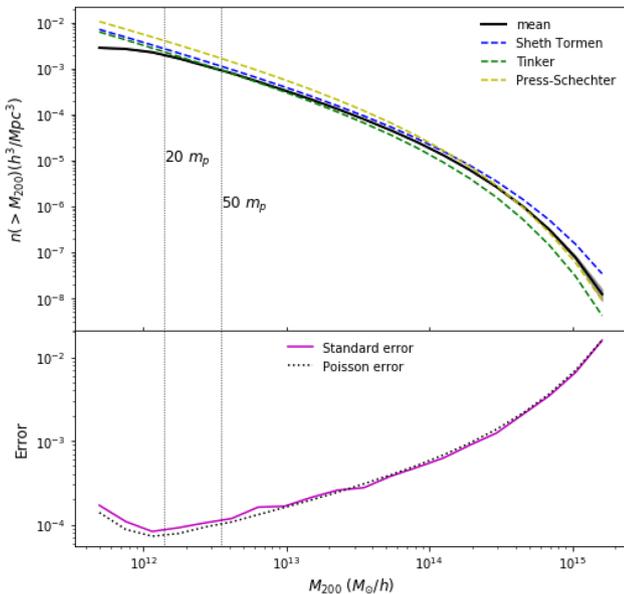


Figure 4. *Upper:* Average M_{200} mass function at $z = 0$ of 320 simulations. The grey-shaded region shows the standard deviation in each mass bin, and vertical lines show where the M_{200} mass is equivalent to 20 and 50 times the particle mass. Theoretical mass functions from Sheth, Mo & Tormen (2001), Press & Schechter (1974), and Tinker et al. (2010) are also shown. *Lower:* Standard deviation of the $z = 0$ mass functions of Fig. 4, along with the Poisson error. The Poisson error is a very good approximation of the standard deviation of the mass function for larger haloes, but underestimates it for smaller haloes.

N -body simulations (Blot et al. 2019; Lippich et al. 2019). In this data release paper, we focus on the mass distribution of the Indra haloes and its ensemble average, variance, and covariance. Specifically, we measure the M_{200} mass function of the main subhaloes identified by SUBFIND, excluding the children of the main subhaloes. M_{200} is defined as the mass within R_{200} , the radius within which the density is greater than 200 times the critical density, ρ_c .

The upper panel of Fig. 4 shows the average M_{200} mass function at $z = 0$ of 320 Indra simulations, along with theoretical mass functions

(Press & Schechter 1974; Sheth et al. 2001; Tinker et al. 2010) derived using HMFcalc (Murray, Power & Robotham 2013).⁸ The standard deviation of these mass functions is given by the grey-shaded area around the mean and is only visible for the highest masses. As is standard, the mass function begins to flatten at masses for which haloes are not well resolved due to the resolution limit of the simulation. Vertical lines show where the halo mass is equivalent to 20 and 50 times the particle mass, though note that these masses do not correspond to the number of particles in the haloes (all haloes have a minimum of 20 particles) since we have measured the M_{200} mass function; the lines are merely a guide to suggest a point along the mass function at which haloes may be considered well resolved, which is a subjective judgment in any case.

Though the mass function is generally plotted as the cumulative number density of haloes greater than a given mass as a function of mass, the variance of the mass function is better characterized by the mean number of haloes (not the cumulative number density) in each mass bin, \bar{N} . The bottom panel of Fig. 4 shows the relative standard deviation of the mean number of haloes in each mass bin, $(\sigma_{\bar{N}}/\bar{N})/\sqrt{320}$, along with the Poisson error, $1/\sqrt{\bar{N}} * 320$. The Poisson error is a very good approximation to the standard deviation of the mass function for haloes above $\sim 10^{13} h^{-1} M_{\odot}$ and slightly underestimates the standard deviation for lower mass haloes. The errors generally increase as a function of mass, as the number of haloes at a given mass decreases, but also increase at low mass, as fewer haloes are found due to the resolution limit of the simulation.

Fig. 5 shows the correlation matrix of 320 $z = 0$ mass functions. As a guide, the horizontal and vertical lines show where the M_{200} mass is equivalent to 50 times the particle mass, or $3.5 \times 10^{12} h^{-1} M_{\odot}$. The off-diagonal elements are greatest for the lowest mass haloes, echoing the small-scale (large k) behaviour of the power spectrum covariance in Fig. 3. The covariance matrix has a rather rich structure that to our knowledge we are the first to measure to this precision. There is a substantial anticorrelation far from the diagonal, i.e. between large and small haloes. We suspect this is a result of conservation of mass in these simulations, i.e. larger haloes grow at the expense of smaller haloes.

⁸<http://hmf.icrar.org/>

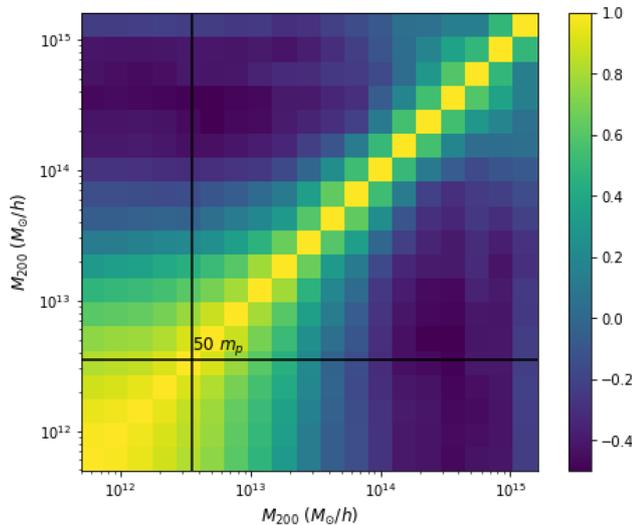


Figure 5. Correlation matrix of the mass functions of 320 Indra simulations at $z = 0$. The horizontal and vertical lines show where the M_{200} mass is equivalent to 50 times the particle mass.

3.3 Fourier modes of the density field

In this section, we analyse the Fourier modes of the density field that are output as the simulations run. These data enable the study of the mode-by-mode evolution of the density field without the need to measure computationally expensive density grids from the particle positions, though only down to scales of $k \sim 0.7 h \text{ Mpc}^{-1}$ ($9 h^{-1} \text{ Mpc}$). We first measure the mean and variance of the $z = 0$ power spectra and compare this to those derived from the particle positions. We then look in detail at the one-point distribution functions of the density, $\delta_k(a)$, and the mode-dependent growth function, $D(a, k)$, in four bins of wavenumber.

3.3.1 Power spectra

Fig. 6 shows the mean and variance of the power spectra of 320 Indra realizations at $z = 1$, measured both from the output Fourier modes of the density field and from the CIC density grid derived from the particle positions. Both the mean and the variance of the FFT and particle-based power spectra agree very well with each other up to the Nyquist frequency of the FFT output, with only a slight deviation at higher frequencies (at the corner modes; see Section 2.3). Note that the FFT and CIC power spectra were not computed in the same k -space bins, which explains some of the disagreement at low frequencies. The FFT outputs are therefore an exceptional resource for studies of the Fourier-space density field from the largest scales down to the mildly non-linear regime of $10 h^{-1} \text{ Mpc}$, at very high temporal resolution, without the need for expensive interpolations of the density field from the particle data. These data can also be an excellent training set for neural network architectures.

3.3.2 One-point distribution function

The properties of the distribution of density fluctuations at early times are well described by a Gaussian random field, as extensively studied by Bardeen et al. (1986). In particular, the Fourier transform of a Gaussian random field allows analytical study of the decomposition of the field into its Fourier amplitudes and phases. For a given scale-factor a and wavenumber k , the Fourier amplitudes of the

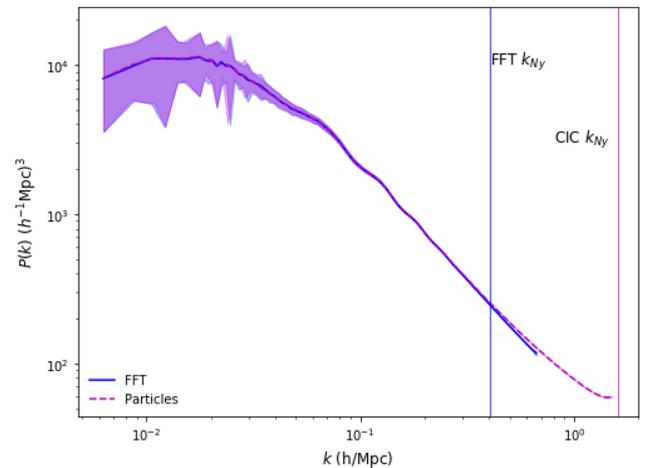


Figure 6. Ensemble average power spectra of 320 simulations at $z = 1$. In blue is the mean and standard deviation of power spectra derived from the FFT output, and in magenta is the mean and standard deviation of power spectra from CIC density fields of the particle positions (as in Section 3.1). The vertical lines show the Nyquist frequencies of the FFT output (blue) and of the CIC grid (magenta).

density field can vary for different random initial conditions and for different wavevectors within a bin of wavenumber. We can thus measure the one-point distribution function of Fourier modes from an ensemble of simulations such as Indra and study its evolution and scale dependence.

In Matsubara (2007), the complex Fourier modes are normalized by the power spectrum, and a distribution function is obtained for

$$\alpha_k = \frac{\delta_k(a)}{\sqrt{P(a, k)}} = \frac{\delta_k(a)}{\sqrt{(|\delta_k(a)|^2)}} = A_k e^{i\theta_k}. \quad (4)$$

Matsubara (2007) finds that the distribution function of a Fourier mode does not depend on the phase θ_k , with the result that, for a random field (not necessarily Gaussian) in a spatially homogeneous space, the one-point distribution function of Fourier phase, $P(\theta_k)$, is always homogeneous, that is, the phases are uniformly distributed.

Fig. 7 shows the evolution of individual Fourier modes for 20 random wavevectors having nearly the same wavenumber in four k bins, defined in Table 1: the largest modes (upper left), modes around the BAO scale (upper right), the regime where $L \sim 21 h^{-1} \text{ Mpc}$ (lower left), and the mildly non-linear regime where $L \sim 10.4 h^{-1} \text{ Mpc}$ (lower right). As k increases and the length-scale gets smaller, the evolution of a given Fourier mode becomes more non-linear, but at all scales, the range of normalized mode amplitudes within the k bin is similarly spread out. These trajectories look rather random, but are deterministic, driven by the evolving phase and amplitude couplings between each mode and all other modes (i.e. the rest of the density field; e.g. Bernardeau et al. 2002). The distribution of mode amplitudes at a given time is the one-point distribution function studied by Bardeen et al. (1986) and Matsubara (2007).

The Fourier modes of a Gaussian random field have a distribution function that is Rayleigh-distributed (Bardeen et al. 1986):

$$P_G(A_k, \theta_k) dA_k d\theta_k = 2A_k e^{-A_k^2} dA_k \frac{d\theta_k}{2\pi} \quad (5)$$

(see equation 68 in Matsubara 2007). In the upper panels of Fig. 8, we show the one-point distribution function of Fourier modes from 320 Indra simulations for the four k bins given in Table 1 and three values of the scale factor, $a = 0.2, 0.5$, and 0.8 , corresponding to redshifts of $z = 4, 1$, and 0.25 . The lower panels of Fig. 8 show the

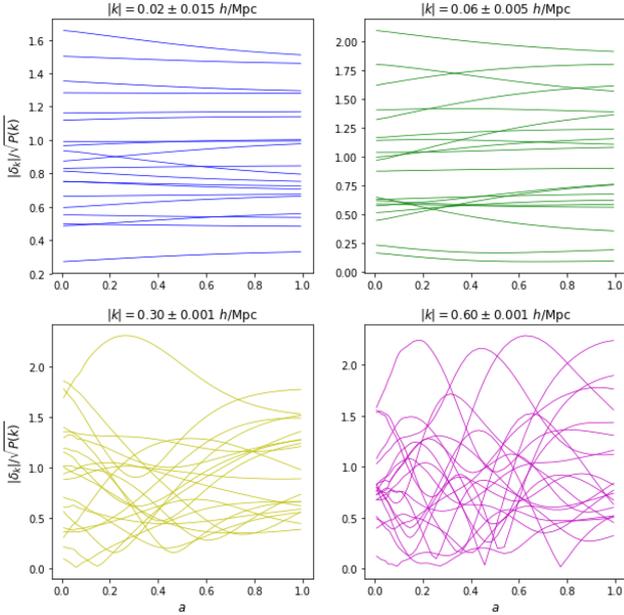


Figure 7. Evolution of normalized Fourier modes of the density field for a random selection of wavevectors in a bin of wavenumber in one Indra simulation.

Table 1. The wavenumber bins, their corresponding range of scales, and the number of wavevectors (for one simulation) in each bin. Thus, the total number of modes in each bin is N_m multiplied by the number of simulations used, which is 320 in this paper.

$ k $ ($h \text{ Mpc}^{-1}$)	δk	l ($h^{-1} \text{ Mpc}$)	(l_{\min}, l_{\max})	N_m
0.02	0.015	314	(180, 1256)	417
0.06	0.005	105	(97, 114)	998
0.30	0.001	20.9	(20.87, 21.01)	4548
0.60	0.001	10.47	(10.45, 10.49)	1224

difference between these distribution functions and that of a Gaussian random field. All of the distribution functions appear to be relatively consistent with a Rayleigh-distributed Gaussian random field.

As Fig. 8 shows, there is very little difference between the distribution functions at different times and scales. However, Fig. 7 shows differences in the time-evolution of individual Fourier modes at different scales. In Fig. 9, we show summary statistics of the one-point distributions of Fourier modes as a function of scale factor. The upper panel shows the ratio of the mean to the standard deviation, μ/σ , which has a constant value of $\sqrt{\pi/(4-\pi)} \simeq 1.913$ for a Rayleigh-distributed Gaussian random field. The lower panel shows the evolution of the skewness, which is equal to $(2\sqrt{\pi}(\pi-3))/(4-\pi)^{3/2} \simeq 0.631$ for a Gaussian random field. The error bands are given by the standard errors of each statistic, assuming that the statistics are Gaussian distributed (Harding, Tremblay & Cousineau 2014). The largest scale mode, at $k = 0.02 h \text{ Mpc}^{-1}$, is the only one for which the statistics lie outside those of a Rayleigh distribution. This may be because there is a factor of 2 fewer modes in this largest bin than the $k = 0.06 h \text{ Mpc}^{-1}$ bin. It is also interesting to note the oscillation of the distribution function statistics of the smallest scale mode, $k = 0.6 h \text{ Mpc}^{-1}$, but further investigations are outside the scope of this paper.

3.3.3 Scale-dependent evolution

In linear theory, the matter density contrast δ evolves independently of scale. The *linear growth function* $D_L(a)$ can be defined via

$$\delta(a) = D_L(a)\delta(a=1). \quad (6)$$

For a flat universe where $\Omega_M + \Omega_\Lambda = 1$,

$$D_L(a) = 5/2 \Omega_M \sqrt{\Omega_M/a^3 + \Omega_\Lambda} \int_0^1 \frac{da}{(a\sqrt{\Omega_M/a^3 + \Omega_\Lambda})^3}, \quad (7)$$

which reduces to $D_L(a) = a$ when $\Omega_M = 1$.

When $\delta \gtrsim 1$, the evolution becomes non-linear, and the k -modes no longer evolve independently, an effect known as mode-coupling (Meiksin & White 1999; Scoccimarro, Zaldarriaga & Hui 1999). In this case, one can solve for the evolution of the density field perturbatively:

$$\delta(a, k) \approx D(a)\delta^{(1)}(k) + D^2(a)\delta^{(2)}(k) + \dots \quad (8)$$

Taking the first order only, and approximating $\delta^{(1)}(k)$ as the density field of the first output of the simulations at $z = 127$, we measure the *mode-dependent growth function* $D(a, k)$ as the amplitude of the ratio of the complex Fourier modes of the density field:

$$D(a, k) = |\delta_k(a)/\delta_k(z=127)|. \quad (9)$$

We then divide this by the mode-independent linear growth function, calculated from the analytic solution in equation (7) and scaled by the growth function at $z = 127$:

$$D(a) = D_L(a)/D_L(z=127). \quad (10)$$

The ratio is then

$$D(a, k)/D(a) = \frac{|\delta_k(a)/\delta_k(z=127)|}{D_L(a)/D_L(z=127)}. \quad (11)$$

When this ratio deviates from unity, that is a sign that the evolution is non-linear and the mode-coupling term is important.

The collection of $D(a, k)$'s that contribute to a k bin can exhibit a large amount of variation for one simulation and across many realizations of random initial phases. Fig. 10 shows the evolution of the mode-dependent growth function for 20 random wavevectors in the four k bins defined in Table 1. Note that the y-axis ranges change for each wavenumber. Also shown are the median, mean, and inner 68 per cent percentiles of the distributions of all wavevectors within each wavenumber range over 320 Indra simulations, shown in black.

At a given a , the distribution of mode-dependent growth functions for the largest mode is very nearly Gaussian, while the BAO ($105 h^{-1} \text{ Mpc}$), mildly non-linear ($20.9 h^{-1} \text{ Mpc}$), and more non-linear ($10.5 h^{-1} \text{ Mpc}$) distributions are positively skewed. It turns out that these distributions, for a given scale factor, are very nearly lognormal. In Fig. 11, we plot the sigma-normalized mean of the logarithm of these distributions, or explicitly, $\langle \ln(D(a, k)/D(a)) \rangle$, as a function of the scale factor. Here, the average $\langle \rangle$ is taken over the simulations in the Indra ensemble and the modes in each k bin. Notably, the means of the two largest scale bins are below zero, meaning the means of the mode-dependent growth functions are less than the linear growth rate, at all scale factors, while the two smaller scale modes are above zero and grow over time. This shows how energy cascades from large to small scales as a result of mode-mode coupling.

As with the Fourier mode distribution functions in Fig. 8, it is illustrative to look at the full distributions of the mode-dependent growth function at a given scale-factor. In Fig. 12, we show the distributions of $\ln(D(a, k)/D(a))$ in four k bins, normalized by

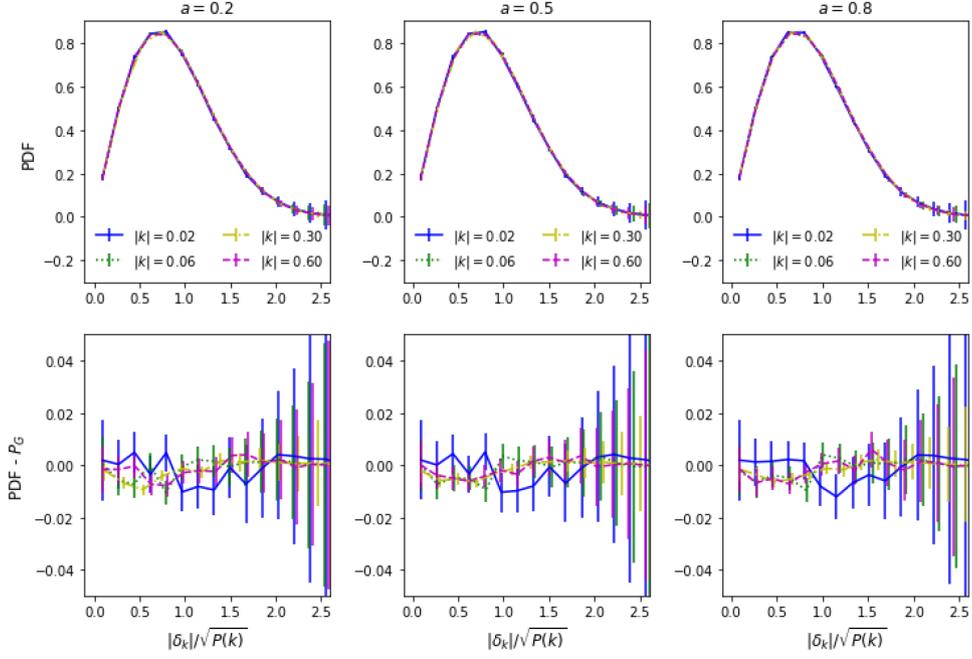


Figure 8. *Upper panels:* One-point distribution functions of normalized Fourier amplitudes of the density field in four k bins and at three values of the scale factor, measured for 320 Indra simulations. *Lower panels:* Difference between the Indra distribution functions and those of a Gaussian random field. The error bars correspond to Poisson errors.

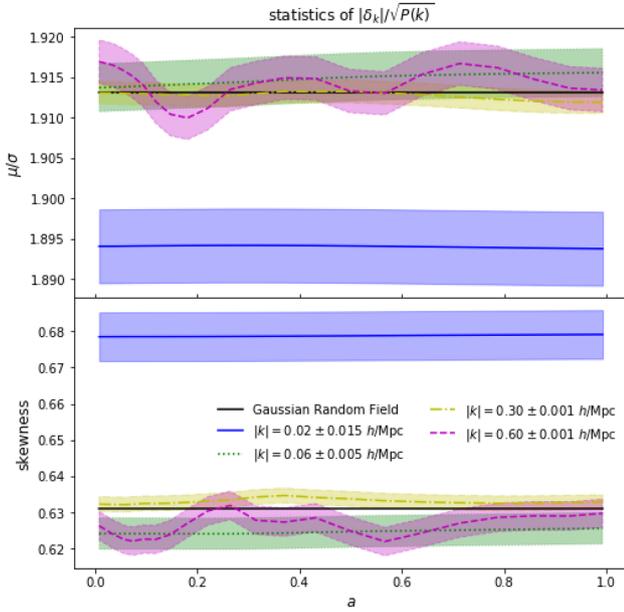


Figure 9. Statistics of Fourier mode distribution functions as a function of scale-factor in four k bins. The upper panel shows the ratio of the mean to the standard deviation, along with its value for a Rayleigh-distributed Gaussian random field, while the lower panel shows the skewness and its Gaussian value. The bands show the standard errors of each statistic.

their standard deviations, for three different scale factors. Recall from equation (9) that we have defined the mode-dependent growth function as the amplitude of the complex Fourier modes of the density field, $\delta_k(a)$, normalized by $\delta_k(z = 127)$ as a proxy for the linear density field, and that the ratio $D(a, k)/D(a)$ is near unity when the density field evolves linearly. Even after normalizing the width of the histograms by their standard deviations, which is much larger for the

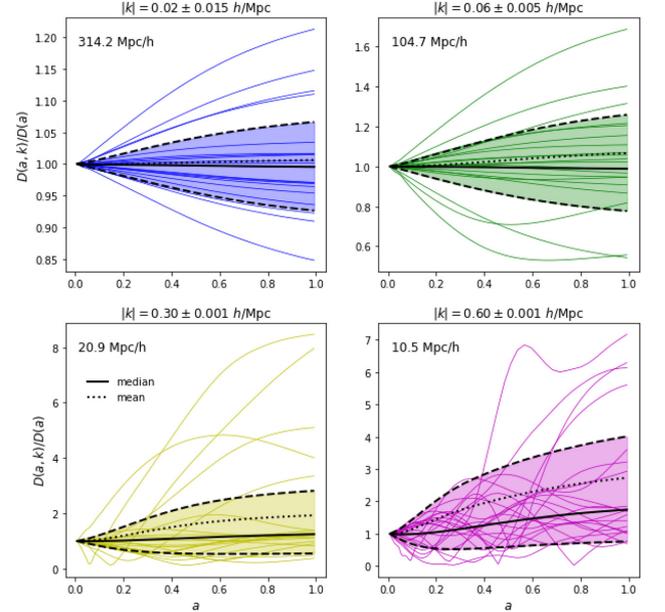


Figure 10. Mode-dependent growth function evolution for a random selection of wavevectors having the same wavenumber, as well as the median (solid), mean (dotted), and 16th and 84th percentiles (dashed) of the distributions of all wavevectors having within the given range of wavenumber over 320 Indra simulations. Wavelengths of modes in the middle of each bin appear in the upper left corner of each panel.

smaller scale modes, the distributions become wider as k increases and the growth function tends to deviate far from the linear theory value. The distribution of the largest-scale mode ($|k| = 0.02 \text{ h Mpc}^{-1}$) is sharply peaked around 0, which is expected for linear evolution.

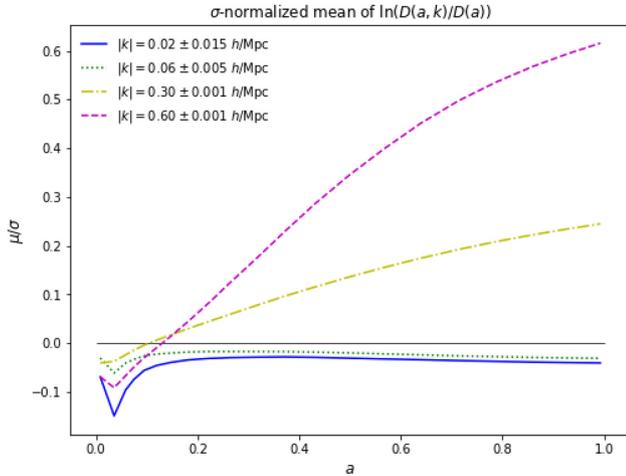


Figure 11. Evolution of the sigma-normalized mean of the logarithm of the mode-dependent growth function, $D(a, k)$.

4 CONCLUSIONS

Indra is a suite of cosmological N -body simulations made available and computationally accessible through the SciServer science platform. Each of the 384 (h^{-1} Gpc) 3 volumes was run with the same cosmological parameters and different realizations of the initial conditions, providing the means to calculate precise averages and covariances on large scales. The large number of saved snapshots (64) for each simulation provides excellent redshift coverage and enables the calculation of halo merger trees, semi-analytic catalogues, and other derived data products.

In this paper, we have presented ensemble averages, variances, and covariances for each of Indra’s three main data products: dark matter particles, halo catalogues, and Fourier-space density grids. We also studied in detail the one-point distribution function of the Fourier-space density field and the evolution and mode dependence of the growth function. We anticipate that the suite of Indra simulations will enable a variety of scientific investigations, such as

(i) conditional statistics: mass functions, correlation functions, and power spectra conditional on, e.g. a density criterion (see e.g. Massara et al. 2020);

(ii) extreme statistics, e.g. of very large or very rare structures (see e.g. Watson et al. 2014);

(iii) re-simulations of interesting structures at higher resolution [enabled by Panphasia (Jenkins 2013) initial phases];

(iv) new data bases of structures: haloes from other halo finders, voids, or filaments defined by haloes or by the dark matter particles, etc.;

(v) ‘event data base’ of interesting things from the full suite, e.g. merging clusters;

(vi) suites of mock galaxy catalogues and light-cone catalogues, optimized by fast spatial searches and well sampled over a large range of redshifts;

(vii) studies where a full N -body reference ensemble would benchmark tests of, or eliminate the need for, suites of approximate simulation methods; and

(viii) the development of machine learning methods requiring large training sets of structure formation on 10–100 h^{-1} Mpc scales.

The ability to ask these and other scientific questions of the petabyte-scale Indra simulations is supported by the data storage and computational infrastructure of the DataScope and SciServer (Taghizadeh-Popp et al. 2020). For modern data sets of the size as the Indra suite, it becomes unfeasible to require users to download the data for local analysis. Instead, a science platform such as SciServer allows users to bring their analysis to the data. In this paper, we have presented patterns of storage (data base and distributed file systems) and compute (Jupyter notebooks in Docker containers) developed for SciServer that we found useful for large simulation data and that can in principle be replicated by other teams. SciServer itself has been deployed in a number of other institutes and various groups are building their own science platforms. Current and future large-scale structure surveys such as *Euclid*, Roman, DESI, and the VRO also require large suites of simulations to perform their analysis, and we would strongly encourage those groups to publish any simulation data catalogues together with their observational data.

We believe that making Indra available in an environment such as SciServer is worth the considerable effort and resources. Making the data explorable and interactive, without the need to download terabytes to a local compute environment, opens up discovery space to the community of simulation experts as well as students and the public.

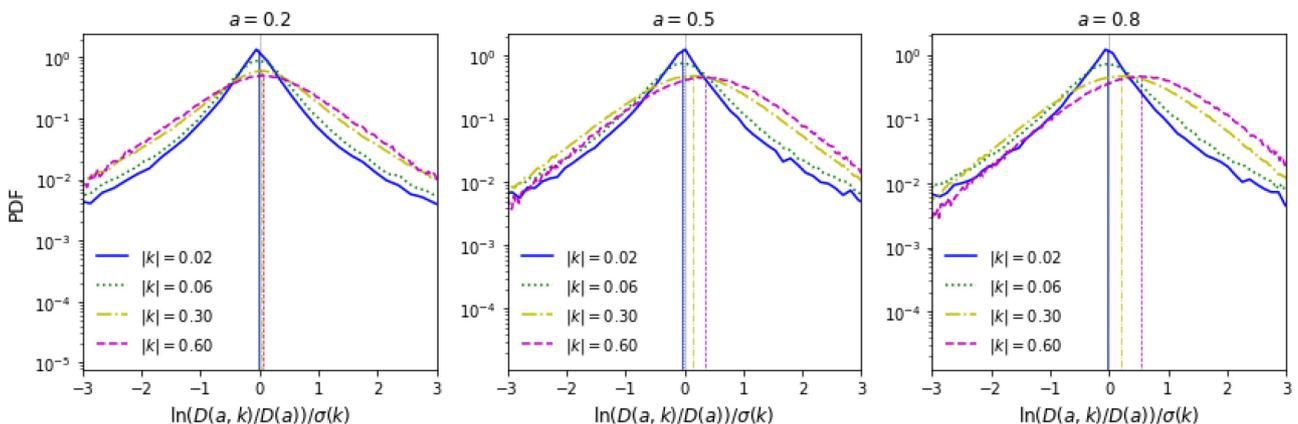


Figure 12. Distribution functions of σ -normalized $\ln(D(a, k)/D(a))$ in four k bins and at three values of the scale factor, measured for 320 Indra simulations. The σ -normalized mean values of the distributions are given as vertical lines under the histograms, showing that the two larger scale modes have means less than 0, and the two smaller-scale modes have means greater than 0. See Fig. 10 for a sense of distribution widths without σ -normalizing; the distributions are narrow at low k .

ACKNOWLEDGEMENTS

We could not have made Indra publicly available without the support of IDIES and SciServer developers and engineers, including Brian Mohr, Jan Vandenberg, Manu Taghizadeh-Popp, Sue Werner, and Lance Joseph. We are thankful for the work done on earlier ideas for a public Indra data base by Daniel Crankshaw, Tamás Budavári, and László Dobos and for preliminary scientific analysis performed by Nuala McCullagh. We are grateful to Volker Springel for providing us with the L-GADGET2 and SUBFIND codes, and our analysis extensively used the following code packages: iPYTHON (Perez & Granger 2007), MATPLOTLIB (Hunter 2007), NUMPY (van der Walt, Colbert & Varoquaux 2011), and SCIPY (<http://www.scipy.org>). This work could not have been carried out without support from the administrative and custodial staff at several institutions. MCN is grateful for funding from Basque Government grant IT956-16.

Computational resources for the Indra simulations were provided by the Homewood High-Performance Cluster (HHPC) and the Maryland Advanced Computing Center (MARCC) through the Institute for Data-Intensive Engineering and Science (IDIES) at JHU. This research used resources of the National Energy Research Scientific Computing Center (NERSC, ROR <https://ror.org/05v3mvq14>, GRID ID grid.484489.d), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. Work relating to the initial conditions used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the Science and Technology Facilities Council (STFC) DiRAC HPC Facility (<http://www.dirac.ac.uk>). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grants ST/H008519/1 and ST/K00087X/1, STFC DiRAC Operations grant ST/K003267/1 and Durham University. DiRAC is part of the UK National E-Infrastructure.

AUTHOR CONTRIBUTIONS

BF did the bulk of the analysis and wrote the majority of this manuscript. BF coordinated the data storage, kept track of the simulation and post-processing progress, and worked with GL and DM to set-up and test the public access framework on SciServer. BF wrote the PYTHON package that reads Indra data and the example notebooks that show how to interact with Indra on SciServer. JW maintained the production GADGET and SUBFIND codes, coordinated supercomputer time, and ran the bulk of the simulation and the post-processing codes. ARJ helped with planning and setting up the initial conditions and contributed to the writing of part of Section 2 of the manuscript. GL is the projects manager of SciServer and provided scripts for loading the data bases and for the 3D spatial query libraries. DM developed critical infrastructure to enable public access and computational support of Indra on SciServer. MCN helped to plan the simulations and analysis, coordinated supercomputer time for the simulations, and ran several of them. AS helped to procure funding, plan the project, and analyse the results.

DATA AVAILABILITY

The data underlying this article are available from SciServer at <http://sciserver.org>. The derived data generated in this research will be shared on reasonable request to the corresponding author.

REFERENCES

- Bardeen J. M., Bond J. R., Kaiser N., Szalay A. S., 1986, *ApJ*, 304, 15
 Bernardeau F., Colombi S., Gaztañaga E., Scoccimarro R., 2002, *Phys. Rep.*, 367, 1
 Bernyk M. et al., 2016, *ApJS*, 223, 9
 Blot L., Corasaniti P. S., Alimi J. M., Reverdy V., Rasera Y., 2015, *MNRAS*, 446, 1756
 Blot L., Corasaniti P. S., Amendola L., Kitching T. D., 2016, *MNRAS*, 458, 4462
 Blot L. et al., 2019, *MNRAS*, 485, 2806
 Chuang C.-H. et al., 2015, *MNRAS*, 452, 686
 Crankshaw D., Falck B., Szalay A. S., Burns R., Budavári T., Wang J., 2013, in Proceedings of the 25th International Conference on Scientific and Statistical Data base Management, SSDBM. Association for Computing Machinery, New York, NY
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
 de Putter R., Wagner C., Mena O., Verde L., Percival W. J., 2012, *J. Cosmol. Astropart. Phys.*, 2012, 019
 DeRose J. et al., 2019, *ApJ*, 875, 69
 Dobos L., Szalay A. S., Blakeley J., Falck B., Budavári T., Csabai I., 2012, in Ballester P., Egret D., Lorente N. P. F., eds, ASP Conf. Ser. Vol. 461, Astronomical Data Analysis Software and Systems XXI. Astron. Soc. Pac., San Francisco, p.323
 Dodelson S., Schneider M. D., 2013, *Phys. Rev. D*, 88, 063537
 Falck B., McCullagh N., Neyrinck M. C., Wang J., Szalay A. S., 2017, *ApJ*, 837, 181
 Feng Y., Hand N., biweidai, 2017, rainwoodman/pmesh 0.1.33 (Zenodo). Available at: <http://doi.org/10.5281/zenodo.1051254>
 Garrison L. H., Eisenstein D. J., Ferrer D., Tinker J. L., Pinto P. A., Weinberg D. H., 2018, *ApJS*, 236, 43
 Hamilton A. J. S., Rimes C. D., Scoccimarro R., 2006, *MNRAS*, 371, 1188
 Harding B., Tremblay C., Cousineau D., 2014, *Quant. Methods Psychol.*, 10, 107
 Heavens A. F., Sellentin E., de Mijolla D., Vianello A., 2017, *MNRAS*, 472, 4244
 Heitmann K., Lawrence E., Kwan J., Habib S., Higdon D., 2014, *ApJ*, 780, 111
 Heitmann K. et al., 2016, *ApJ*, 820, 108
 Heitmann K. et al., 2019, *ApJS*, 244, 17
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Jenkins A., 2010, *MNRAS*, 403, 1859
 Jenkins A., 2013, *MNRAS*, 434, 2094
 Jenkins A., Booth S., 2013, preprint ([arXiv:1306.5771](https://arxiv.org/abs/1306.5771))
 Jing Y. P., 2005, *ApJ*, 620, 559
 Joyce A., Lombriser L., Schmidt F., 2016, *Annu. Rev. Nucl. Part. Sci.*, 66, 95
 Jurić M. et al., 2017, in Lorente N. P. F., Shorridge K., Wayth R., eds, ASP Conf. Ser. Vol. 512, Astronomical Data Analysis Software and Systems XXV. Astron. Soc. Pac., San Francisco, p. 279
 Komatsu E. et al., 2011, *ApJS*, 192, 18
 Lemson G., Virgo Consortium T., 2006, preprint ([arXiv:astro-ph/0608019](https://arxiv.org/abs/astro-ph/0608019))
 Lemson G., Budavari T., Szalay A. S., 2011, in Cushing J. B., French J. C., Bowers S., eds, Lecture Notes in Computer Science, Vol. 6809, Proceedings of Scientific and Statistical Data base Management - 23rd International Conference, SSDBM 2011, Portland, OR, USA, July 20-22, 2011. Springer, Berlin, p. 509
 Lippich M. et al., 2019, *MNRAS*, 482, 1786
 Loebman S. et al., 2014, Proceedings of Workshop on Data Analytics in the Cloud, DanaC'14. Association for Computing Machinery, New York, NY, p. 1
 Massara E., Villaescusa-Navarro F., Ho S., Dalal N., Spergel D. N., 2021, *Phys. Rev. Lett.*, 126, 011311
 Matsubara T., 2007, *ApJS*, 170, 1
 Meiksin A., White M., 1999, *MNRAS*, 308, 1179
 Michaux M., Hahn O., Rampf C., Angulo R. E., 2021, *MNRAS*, 500, 663
 Murray S. G., Power C., Robotham A. S. G., 2013, *Astron. Comput.*, 3, 23

- Nelson D. et al., 2019, *Comput. Astrophys. Cosmol.*, 6, 2
 Nishimichi T. et al., 2019, *ApJ*, 884, 29
 Percival W. J. et al., 2014, *MNRAS*, 439, 2531
 Perez F., Granger B. E., 2007, *Comput. Sci. Eng.*, 9, 21
 Pope A. C., Szapudi I., 2008, *MNRAS*, 389, 766
 Press W. H., Schechter P., 1974, *ApJ*, 187, 425
 Riebe K. et al., 2011, *New Astron.*, preprint (arXiv:1109.0003)
 Scoccimarro R., 1998, *MNRAS*, 299, 1097
 Scoccimarro R., Zaldarriaga M., Hui L., 1999, *ApJ*, 527, 1
 Sellentin E., Heavens A. F., 2016, *MNRAS*, 456, L132
 Sheth R. K., Mo H. J., Tormen G., 2001, *MNRAS*, 323, 1
 Springel V., 2005, *MNRAS*, 364, 1105
 Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
 Springel V. et al., 2005, *Nature*, 435, 629
 Szalay A. S., Kunszt P. Z., Thakar A. R., Gray J., Slutz D., 2000, in Manset N., Veillet C., Crabtree D., eds, *ASP Conf. Ser. Vol. 216, Astronomical Data Analysis Software and Systems IX*. Astron. Soc. Pac., San Francisco, p. 405
 Taghizadeh-Popp M. et al., 2020, *Astron. Comput.*, 33, 100412
 Takahashi R. et al., 2009, *ApJ*, 700, 479
 Taylor A., Joachimi B., Kitching T., 2013, *MNRAS*, 432, 1928
 Tinker J. L., Robertson B. E., Kravtsov A. V., Klypin A., Warren M. S., Yepes G., Gottlöber S., 2010, *ApJ*, 724, 878
 van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
 Villaescusa-Navarro F. et al., 2020, *ApJS*, 250, 2
 Watson W. A., Iliev I. T., Diego J. M., Gottlöber S., Knebe A., Martínez-González E., Yepes G., 2014, *MNRAS*, 437, 3776
 Winther H. A. et al., 2015, *MNRAS*, 454, 4208

APPENDIX: ACCESSING AND ANALYSING THE INDRA SIMULATIONS

Indra is accessed by creating an account on the SciServer science platform (Taghizadeh-Popp et al. 2020) at <http://www.sciserver.org>. SciServer users have access to public data sets and SQL data bases, private data storage, file-sharing and group collaborative tools, and the interactive and batch-mode compute environments. Users effectively opt-in to data sets they are interested in by choosing the appropriate *Science Domain*; Indra is in the *Cosmological Simulations* Science Domain, which also contains data from the Millennium simulations. Once a user joins this Domain, they will be able to query the Indra data base in *CasJobs* and load the Indra data and software in a *Compute* container, both of which are reached from SciServer *Dashboard* after you log in. General SciServer tutorials can be found at <http://www.sciserver.org/support/help/>.

All of the Indra data, including the data base of halo catalogues, can be analysed in a *Compute* container. When creating a new container, choose the *Cosmological Simulations* compute image, which has INDRA-TOOLS⁹ and some other simulation software pre-installed, and check all of the Indra Data Volumes. The INDRA-TOOLS GITHUB repository includes several example Jupyter notebooks to help get users started with accessing the data and performing some analysis:

⁹<http://github.com/bfalck/indra-tools>

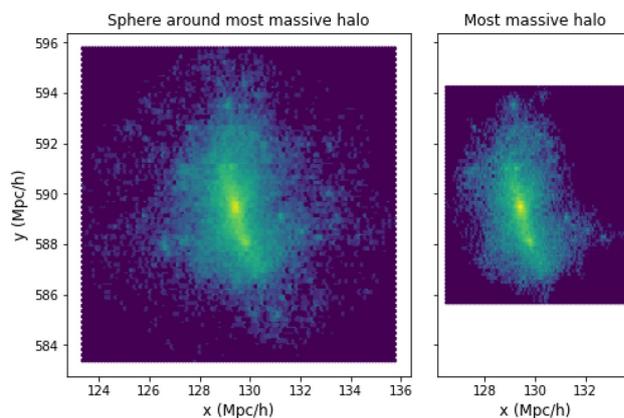


Figure A1. *Left:* Particles in a sphere around the most massive halo in run 2_0_0 at $z = 0$ whose radius is three times the radius of the halo. *Right:* Particles identified as being in this halo.

(i) *read_examples:* How to read all of the data products: snapshots of particle positions and velocities, plus pre-computed power spectra at select snapshots; Fourier modes of the coarse-gridded density field; and the halo and subhalo catalogues, including how to index the halo catalogues and retrieve IDs of particles in haloes.

(ii) *database_examples:* How to query the halo data base tables, including sample queries that demonstrate how to select from one run and snapshot, one run and multiple snapshots, and one snapshot and multiple runs.

(iii) *density_field_examples:* How to compute real-space density fields from the Fourier-space density fields and from the snapshots of particle positions as well as how to create quick slices for plots using the Shape3D functionality.

(iv) *Shape3D_examples:* How to use Shape3D objects to efficiently read subsets of particles contained in spheres, boxes, cones, and cone segments to, e.g. grab all particles around (or in) a given halo or create light-cones.

Since reading all 1024^3 particles from one snapshot can take some time, the Shape3D functionality can significantly speed up analyses that are only concerned with portions of the full volume. This functionality uses the spatial indexing library developed for numerical simulations (Lemson et al. 2011) to define shapes within a periodic volume and selectively read only those particles contained within the shape. Fig. A1 demonstrates the result of querying the most massive halo for run 2_0_0 at $z = 0$ from the data base and reading all particles in a sphere around that halo, which takes a few seconds to execute. The right-hand panel shows only those particles identified as being within the halo, which can take 1–2 min because it requires reading the entire binary halo catalogue and the halo particle IDs for that run and snapshot. This is one of the examples in the *Shape3D_examples* notebook.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.