

Constructing high-fidelity halo merger trees in ABACUSSUMMIT

Sownak Bose ^{1,2}★ Daniel J. Eisenstein,¹ Boryana Hadzhiyska ¹, Lehman H. Garrison ³
and Sihan Yuan ^{1,4}

¹Center for Astrophysics | Harvard & Smithsonian, 60 Garden St, Cambridge, MA 02138, USA

²Institute for Computational Cosmology, Department of Physics, Durham University, Durham DH1 3LE, UK

³Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA

⁴Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Stanford, CA 94305, USA

Accepted 2022 February 24. Received 2022 February 23; in original form 2021 September 21

ABSTRACT

Tracking the formation and evolution of dark matter haloes is a critical aspect of any analysis of cosmological N -body simulations. In particular, the mass assembly of a halo and its progenitors, encapsulated in the form of its merger tree, serves as a fundamental input for constructing semi-analytic models of galaxy formation and, more generally, for building mock catalogues that emulate galaxy surveys. We present an algorithm for constructing halo merger trees from ABACUSSUMMIT, the largest suite of cosmological N -body simulations performed to date consisting of nearly 60 trillion particles, and which has been designed to meet the Cosmological Simulation Requirements of the Dark Energy Spectroscopic Instrument (DESI) survey. Our method tracks the cores of haloes to determine associations between objects across multiple time slices, yielding lists of halo progenitors and descendants for the several tens of billions of haloes identified across the entire suite. We present an application of these merger trees as a means to enhance the fidelity of ABACUSSUMMIT halo catalogues by flagging and ‘merging’ haloes deemed to exhibit non-monotonic past merger histories. We show that this cleaning technique identifies portions of the halo population that have been deblended due to choices made by the halo finder, but which could have feasibly been part of larger aggregate systems. We demonstrate that by cleaning halo catalogues in this post-processing step, we remove potentially unphysical features in the default halo catalogues, leaving behind a more robust halo population that can be used to create highly accurate mock galaxy realizations from ABACUSSUMMIT.

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

1 INTRODUCTION

In our present paradigm of structure formation, galaxies are thought to form within potential wells of dark matter that have collapsed out of density fluctuations in the primordial Universe due to gravitational instability. These potential wells – so-called ‘haloes’ of dark matter – are self-gravitating, virialized structures that have decoupled from the background expansion of the Universe. As gas shock heats and, subsequently, cools and condenses within dark matter haloes, it sparks star formation that leads to the eventual formation of galaxies.

Modelling the hierarchical formation and evolution of dark matter haloes is therefore a fundamental component of any theoretical framework of galaxy formation. A well-known, early approach to this problem was presented in the analytic model of Press & Schechter (1974), who, assuming initial fluctuations seeded by a Gaussian random field, derived a methodology for computing the multiplicity function of dark matter haloes as a function of mass and redshift. Despite the relative simplicity of the model, the number density of haloes predicted by the Press–Schechter approach has been shown to be in reasonable agreement with early numerical experiments of cosmic structure formation (e.g. Efstathiou et al. 1988; Lacey & Cole 1994).

A refinement of this model, often dubbed the ‘extended’ Press–Schechter (EPS) theory, was later introduced by Bond et al. (1991), who described a framework for computing the distribution of haloes with mass M at some redshift z by following trajectories of individual mass elements through the linear overdensity field. In particular, EPS enables the calculation of the quantity $f(M_1, z_1 | M_2, z_2)$, or the fraction of mass from haloes with mass M_2 at redshift z_2 that are contained in *progenitor* haloes of mass M_1 identified at an earlier redshift z_1 . The EPS method therefore enables the computation of halo formation and merger rates over multiple epochs (e.g. Lacey & Cole 1993; Somerville & Kolatt 1999; Parkinson, Cole & Helly 2008) – in other words, the formation and evolution of progenitor and descendant haloes, the ensemble of which defines a *halo merger tree*. Once constructed, the merger tree acts as the backbone upon which semi-analytic models of galaxy formation may be built (e.g. Kauffmann, White & Guiderdoni 1993; Cole et al. 2000; Croton et al. 2006; Somerville et al. 2008; Benson 2012; Henriques et al. 2015; Lagos et al. 2018).

While analytic methods like EPS are instructive, modelling the formation and evolution of cosmic structures self-consistently in their large-scale environment is the order of the day if we wish to compare theoretical models to the observed Universe. Over the last several decades, cosmological N -body simulations have become near indispensable tools for augmenting our understanding of structure formation. In terms of the computational volume of

* E-mail: sownak.bose@cfa.harvard.edu

their calculations, simulations that model the dark matter only are currently the industry standard (e.g. Springel et al. 2005; Klypin, Trujillo-Gomez & Primack 2011; Heitmann et al. 2015; Potter, Stadel & Teyssier 2017; Elahi et al. 2018; Garrison et al. 2018; Baugh et al. 2019; Ishiyama et al. 2020). Halo catalogues and merger trees extracted from these simulations may then be augmented with models of the galaxy-halo connection to build mock catalogues that can then be compared to galaxy surveys (see Wechsler & Tinker 2018, for a review). Although hydrodynamical simulations of galaxy formation are continuously improving in both their scale and their sophistication (see Vogelsberger et al. 2020, for a review), at present they fall short of the demands of modern galaxy surveys in terms of box size and the number of simulations available.

The identification of dark matter haloes in N -body simulations is by no means a ‘solved’ problem. At some level, any halo finder has to ultimately resort to making a series of choices when defining the existence and extent of haloes. For example, the first choice comes in deciding whether halo particle memberships are to be determined in configuration space, as in the ‘friends-of-friends’ (Davis et al. 1985) or spherical overdensity-based methods (e.g. Press & Schechter 1974; Klypin & Holtzman 1997), in phase space (e.g. Behroozi, Wechsler & Wu 2013a), or overall characterizations of the cosmic density field (e.g. Neyrinck, Gnedin & Hamilton 2005; Falck, Neyrinck & Szalay 2012). A subsequent step may (or may not) involve determining which subsets of ‘halo’ particles are in fact gravitationally bound (e.g. Springel et al. 2001; Knollmann & Knebe 2009). Each choice can lead to slightly different populations of haloes realized from the same particle field (see Knebe et al. 2011, for an exhaustive study of the effects of halo finders on the properties of the final catalogue).

Given these uncertainties, halo merger trees may then also be used as useful diagnostic tools for assessing the fidelity of the halo population itself. The dynamical history of these objects introduces further complications as they may ‘split’, fly-by one another, merge partially, or simply flicker above and below the mass threshold for recording a halo. Each of these processes can significantly alter the mass associated with the halo over time. This is particularly significant for applications that use the properties of haloes to build mock galaxy catalogues. A prominent example is that of the Halo Occupation Distribution (HOD; e.g. Peacock & Smith 2000; Benson et al. 2000; Berling & Weinberg 2002; Zheng et al. 2005), in which the probability that a halo contains a certain population of galaxies is dependent on properties like its mass or maximum circular velocity. The dynamical histories of haloes may be stochastic enough to make their final mass very different from the mass they would have had when a galaxy in the real Universe would have formed in them, in which case a model like the HOD may result in a biased mock galaxy population. In this regard, a working definition of a ‘well-defined halo’ may simply be a persistent entity – i.e. that which can be *cleanly* associated from one time-step of the simulation to the next. By traversing the merger trees of individual haloes, one can then flag and remove objects that fail to pass this persistence test in post-processing.

There are now several approaches in the literature for constructing merger trees from cosmological simulations, differing in their complexity and the information content that acts as input for building these trees. While some algorithms attempt to establish associations between haloes across time-steps based simply on their mass/position/velocity in the box (Onions et al. 2012), other methods aim at connecting haloes (or subhaloes) that maximize a merit function based on overlapping particle sets across two time slices (Klimontowski et al. 2010; Elahi et al. 2018). Particle-based

correlator methods are the more established flavour in recent times, seeing as they display greater stability in the properties of their merger trees (see, e.g. Srisawat et al. 2013). Some algorithms assign greater weight in the merit function to those particles that are more tightly bound (based on e.g. their binding energy as in Springel et al. 2005; Jiang et al. 2014). Finally, there are also those algorithms that, in addition to using particles to associate haloes, also make use of the past history of individual objects (such as their bulk motion, or indeed their particle membership) to construct the final merger tree (e.g. Behroozi et al. 2013b; Han et al. 2018).

As modern galaxy surveys like the Dark Energy Spectroscopic Instrument (DESI; Levi et al. 2013), *Euclid* (Laureijs et al. 2011), *The Nancy Grace Roman Space Telescope* (Spergel et al. 2015), and the Vera Rubin Observatory’s Legacy Survey of Space and Time (LSST; Ivezić et al. 2019) continue to ramp up the statistical power of observational data, there is an increasingly pressing demand on theoretical models to keep up to pace. The computational challenge lies in the fact that the relevant numerical simulations not only require the several Gpc³ in volume to emulate these surveys, but they need to simultaneously resolve the mass scales of interest. For example, the largest galaxy sample in DESI will target emission-line galaxies (ELGs) in the redshift range between $z \sim 0.6$ and 1.6. Resolving the host haloes of ELGs (typically of the order of $10^{11} M_{\odot} h^{-1}$ in halo mass, Gonzalez-Perez et al. 2018; Hadzhiyska et al. 2021) with at least a modest number of particles within Gpc-sized boxes requires tens of billions of resolution elements. In this paper, we will describe our methodology for constructing and applying halo merger trees from the ABACUSSUMMIT simulations (Maksimova et al. 2021), which have been designed to meet exactly these specifications.

The layout of this paper is as follows. In Section 2, we describe the details of the ABACUSSUMMIT suite, and the algorithms used to identify haloes and construct merger trees from them. In Section 3, we present tests of the robustness of some typical merger tree statistics against the cadence of simulation outputs used to construct them. Section 4 presents our main results, and displays some of the primary applications of the merger trees, including the methodology we use to ‘clean’ halo catalogues by coalescing haloes that have been oversplit by the halo finder, or undergone any number of the dynamical processes described above. We then demonstrate the power of this cleaning technique by applying ‘cleaned’ ABACUSSUMMIT haloes to a realistic application of constructing mock galaxy catalogues to observed galaxy clustering data. Finally, a summary of our main findings is presented in Section 5.

2 METHODOLOGY

In this section, we describe ABACUSSUMMIT, an extremely large set of numerical simulations that are designed to match DESI’s cosmological simulation requirements (Section 2.1), and COMPASO (Hadzhiyska et al. 2022), the on-the-fly halo-finding algorithm that has been specially designed for this project (Section 2.2). We then discuss the design philosophy underpinning our merger tree algorithm (Section 2.3), before discussing the methodology itself in detail (Section 2.4). Finally, we conclude by discussing some of the optimizations that have been implemented to speed-up the construction of merger trees in ABACUSSUMMIT (Section 2.5).

2.1 The ABACUSSUMMIT simulation suite

ABACUSSUMMIT is a suite of very large, high-accuracy cosmological simulations of structure formation that have been designed to match – and exceed – the Cosmological Simulation Requirements of

the DESI survey. In total, we have performed simulations totalling almost 60 trillion particles spanning 97 cosmological models, centred around the Planck Collaboration VI (2020) parameters, which acts as our fiducial cosmological model. A subset of these runs also model the evolution of structure in the presence of massive neutrinos. Individual simulations within the ABACUSSUMMIT suite are labelled according to their cosmology, phase (which are set while generating the corresponding initial conditions), box size, and number of resolution elements. For further details, we refer the reader to the main ABACUSSUMMIT release paper by Maksimova et al. (2021) and the accompanying online documentation.¹ The different numerical setups adopted in each one of the ABACUSSUMMIT simulations are broadly classified as follows:

- (i) *Base*: a 2 Gpc h^{-1} box with 6912³ particles;
- (ii) *Highbase*: a 1 Gpc h^{-1} box with 3456³ particles;
- (iii) *High*: a 1 Gpc h^{-1} box with 6300³ particles;
- (iv) *Huge*: a 7.5 Gpc h^{-1} box with 8640³ particles;
- (v) *Hugebase*: a 2 Gpc h^{-1} box with 2304³ particles; and
- (vi) *Fixedbase*: 1.185 Gpc h^{-1} box with 4096³ particles with fixed amplitude initial conditions.

Note that *Base*, *Highbase*, and *Fixedbase* all have the same mass resolution (around $2 \times 10^9 M_{\odot} h^{-1}$ per particle). On the other hand, *High* has six times better mass resolution, while both *Huge* and *Hugebase* have 27 times worse mass resolution. For the majority of this paper, we will only consider the *Highbase* and *High* simulations. In both instances, we will present results from simulations run at the base cosmology (*c000* in our notation) and the same initial phase (*ph100* in our notation). Collectively, we refer to these simulations as *c000_ph100* throughout the remainder of this paper. The only exception to this is our investigation in Section 4.7, where we make use of one *Base* simulation at *c000* cosmology, but with phase *ph000*. A comprehensive tabulation of all simulation variants is listed in the online documentation for ABACUSSUMMIT.²

All simulations were performed at the Summit supercomputer at the Oak Ridge Leadership Computing Facility, using the ABACUS code (Garrison et al. 2021, see also Garrison, Eisenstein & Pinto 2019). ABACUS is a high-performance, high-accuracy cosmological N -body code that has been optimized for GPU architectures and for large-volume, moderately clustered simulations. By harnessing the power of GPUs, ABACUS is able to clock over 30 million particle updates per second on commodity dual-Xeon, dual-GPU computers and nearly 70 million particle updates per second on each node of the Summit supercomputer. Importantly, the speedup achieved does not compromise on accuracy: ABACUS reports a median force accuracy below 10^{-5} . In-depth descriptions of the near- and far-field force computations, optimization routines, and code tests are presented in Garrison et al. (2021).

2.2 Halo finding with COMPASO

To identify haloes in the ABACUSSUMMIT simulations, we have developed a new, on-the-fly halo finder, COMPASO, which is a hybrid friends-of-friends/spherical overdensity-based halo finder. The algorithm is described in detail in Hadzhiyska et al. (2022); here, we provide a short summary of its basic operation:

(i) For each particle in the simulation, we compute its ‘local density’, Δ , using a weighting kernel defined as

$$W(r; b) = 1 - r^2/b^2,$$

where b is set to be equal to 0.4 times the mean interparticle spacing, l_{mean} . Δ here is defined in units of the cosmic mean density.

(ii) We then mark particles as being eligible to be in groups if they satisfy the condition $\Delta > 60$. These particles are then segmented into what we call ‘L0 haloes’ using a standard friends-of-friends algorithm with linking length $l_{\text{FoF}} = 0.25l_{\text{mean}}$.

(iii) Next, within each L0 halo, we define ‘L1 haloes’ using *Competitive Assignment to Spherical Overdensities* (COMPASO). The particle with the largest Δ is selected to be the first halo nucleus. We then search outward to find the innermost radius within which the enclosed density drops below the L1 threshold density, Δ_{L1} . Particles interior to this radius, R_{L1} , are tentatively assigned to the L1 group defined by the first nucleus. The set of particles that are interior to 80 per cent of R_{L1} are marked as ‘ineligible’ to be future halo nuclei; the remaining particles may yet be eligible.

(iv) The remaining ‘eligible’ particles are then searched to find the particle with the next highest Δ that is also a density maximum. This condition is met when a particle is denser than *all* other particles – eligible or not – that are within a radius of $0.4l_{\text{mean}}$. If the condition is met, we spawn a new nucleus and once again search outwards for its L1 radius, R_{L1} , using *all* L0 particles.

(v) A particle is assigned to the newly created group if it was previously unassigned *or* if its enclosed density with respect to the new group is at least *twice* that of its enclosed density with respect to the group it is currently assigned to. In practice, these enclosed densities are estimated by scaling from R_{L1} assuming an inverse square density profile.

(vi) We continue searching for new nucleation centres until we reach the minimum density threshold, defined as the central density of a singular isothermal sphere with 35 particles within a radius enclosing 200 times the average background density of the universe.

(vii) Finally, within each L1 halo, we repeat steps (iii)–(v) to identify ‘L2 haloes’; the centre of mass of the largest L2 halo is used to define the centre relative to which all L1 halo statistics are outputted.

As part of the ABACUSSUMMIT data products, we output properties of all L1 haloes containing at least 35 particles. Conversely, we only store masses for the five largest L2 haloes within each L1 halo. Haloes are arranged across 34 separate files (or ‘superslabs’) that themselves are comprised of sets of 50 individual slabs which the computational domain in ABACUS is split into (see Garrison et al. 2021, for further details). Hadzhiyska et al. (2022) present a number of comparisons of the halo populations identified by COMPASO with those of haloes identified using the ROCKSTAR halo finder (Behroozi et al. 2013a) run on the same particle sets.

Note that other configuration space-based halo finders perform an additional ‘unbinding’ step that is used to discard high-velocity particles from the set assigned to a given halo. In the case of COMPASO, we have opted to forego such a step: both due to constraints this places on the runtime efficiency of the ABACUS code, and also because energy-based unbinding algorithms have problems inherently associated with the method. Instead, we apply a ‘cleaning’ procedure in post-processing (described in Section 4.2) that is very effective at identifying and filtering haloes with a large fraction of unbound particles. Section 3.3 in Hadzhiyska et al. (2022)

¹<https://abacussummit.readthedocs.io/en/latest/index.html>.

²<https://abacussummit.readthedocs.io/en/latest/cosmologies.html>.

presents a series of detailed tests that compares the outcome of our cleaning method with a halo catalogue produced using COMPASO and unbinding.

2.3 Design philosophy of ABACUS merger trees

Before describing the algorithm in detail, we begin by framing the major concepts and philosophies of our approach to merger trees.

The ABACUSSUMMIT suite is designed to support mock galaxy catalogues based on catalogues and properties of haloes, with rather limited attention to subhaloes. We seek to provide merger trees to support computation of halo properties such as mass-quantile formation times, to detect major merger events, and to track associations of progenitors and descendants across long stretches of cosmic time. In principle, one could seek to modulate the properties of central galaxies or quench the growth of infalling satellites based on such associations, such as is done in semi-analytic galaxy formation models. However, the focus on haloes rather than subhaloes does limit the ability to attempt to model individual satellite galaxies in more physical detail.

Because of this focus on halo scales and associations over longer time, we build the ABACUSSUMMIT merger trees from a moderately dense set of epochs. We use 33, with a separation in redshift of $\Delta z = 0.05$ at $z \leq 0.5$, by $\Delta z = 0.075$ for $0.5 < z < 1.7$, and somewhat coarser spacing at higher redshift. This corresponds to a typical time interval of $\Delta t = \Delta \ln a/H(a) \approx (0.04 \text{ to } -0.05)/H(a)$, where a is the cosmological scalefactor and $1/H(a)$ is the Hubble time. The choice of 33 is lower than the 60–200 that is commonly adopted in other simulation suites (e.g. Srisawat et al. 2013; Wang et al. 2016), but for halo applications we find it to be sufficient (as will be tested in Section 3). We stress that the typical halo crossing time (diameter divided by circular velocity) at the L1 radius, given its associated enclosed overdensity, is roughly $t_{\text{cross}} \approx 0.2/H(a)$. Hence, ABACUSSUMMIT provide four to five output epochs per halo crossing time, giving an intuitive reason as to why our merger-tree results are adequately converged against cadence. We caution that associating importance to halo events on time-scales shorter than a crossing time is probably not physically motivated, e.g. the intrahalo medium cannot equilibrate faster than the sound crossing time. Higher cadences are motivated by the desire to track subhalo mergers, which utilize a higher density threshold and hence smaller dynamical time. Of course, limiting the number of output epochs is also important to economize on the data volume of the simulations, which was 2 PB even for this adopted set.

Our methodology of building merger trees is based on match particle membership, utilizing the fact that ABACUSSUMMIT did track and output unique particle identification numbers. This allows for a direct association, without the need to use dynamical predictions of halo orbits. This methodology is common to several algorithms in the literature and has been shown to result in merger trees with greater stability (see, e.g. Srisawat et al. 2013, for details). In this sense, the algorithm presented in Section 2.4 is readily generalizable to any halo catalogues that record the IDs and densities of the particle sets defining haloes and is not limited to use in conjunction with COMPASO.

One application of the association of haloes between epochs is to detect inconsistencies between single-epoch halo catalogues, in particular in regard to the delicate and inexact problem of deblending of nearby density peaks. Different halo finders will produce different results on such situations, and there is no perfect answer. As described in Section 4.2, we use the merger trees to investigate such issues in compaso, resulting in a data product of cleaned halo catalogues at

a series of epochs. We recommend that these cleaned catalogues are an improved starting point for the construction of mock galaxy catalogues such as from the HOD model.

2.4 Merger tree algorithm

We now describe in detail the algorithm employed for constructing halo merger trees in ABACUSSUMMIT. Note that a ‘halo’ in this context refers to L1 haloes as defined in Section 2.2.

To be able to accurately match haloes across output times, we make use of subsampled particles that are output alongside the halo catalogues themselves. In ABACUSSUMMIT, the subsampled particle list is split into 3 and 7 per cent (yielding a total of 10 per cent) sets, which, respectively, are referred to as subsamples ‘A’ and ‘B’. These subsamples are selected based on a hash of their (unique) particle ID number, and are consistent across redshifts. Each particle is assigned a 64-bit integer, PID, which stores both the ID number as well as its kernel density. The ID is simply the (i, j, k) index location of the particle in the initial grid; these numbers are stored as the lower three 16-bit integers in PID. The kernel density is stored as the square root of the density (in units of the cosmic density) in bits 1..12 of the upper 16-bit integer in PID. The sets of particles associated with L1 haloes are stored contiguously in the subsampled PID files; for the purpose of merger tree construction, we only consider those particles that are marked as being part of a halo (i.e. excluding subsampled particles in the ‘field’).

With the list of L1 haloes and their corresponding particles at hand, the merger tree construction, in *reverse time order*, proceeds as follows:

(i) Start with the first L1 halo, `halo_now`, identified at snapshot `i`, and retrieve its centre (defined by the centre of mass of its largest L2 halo) and its 10 per cent particle subsample.

(ii) Identify the list of all haloes in the *two preceding* snapshots, `i-1` and `i-2`, that could have *plausibly* been associated with `halo_now`. A plausible association is defined as a halo that is located at most at a distance of r_{max} from the centre of `halo_now`. Throughout, we assume $r_{\text{max}} = 4 \text{ Mpc } h^{-1}$.

(iii) From the list of plausible associations, we identify *candidate associations* as those objects that share a non-zero fraction of their unique particle IDs with `halo_now`. The fraction of particles donated by these candidate associations to `halo_now` is labelled as f_{donate} . We also record the fraction of subsampled particles in `halo_now`, *weighted by kernel density*, that is donated to it by its candidate associations in snapshots `i-1` and `i-2`. We denote these values by f_{match} .

(iv) We then mark a candidate association as a Progenitor if $f_{\text{donate}} \geq f_{\text{thresh}}$. The remaining candidate associations are discarded. We assume $f_{\text{thresh}} = 0.5$ throughout.

(v) We identify the MainProgenitor as the halo association at snapshot `i-1` that contributes the largest f_{match} . We further record MainProgenitorPrec (‘preceding’), which is the MainProgenitor identified in snapshot `i-2`. While we keep track of the full Progenitor list from snapshot `i-1`, we only track the value of MainProgenitorPrec from snapshot `i-2`.

(vi) Repeat steps (i)–(v) for the remaining haloes at snapshot `i`.

(vii) Repeat steps (i)–(vi) for all subsequent triples of output times in a given simulation.

A cartoon illustration of this halo association procedure across any two ABACUSSUMMIT snapshots is shown in Fig. 1. For any given

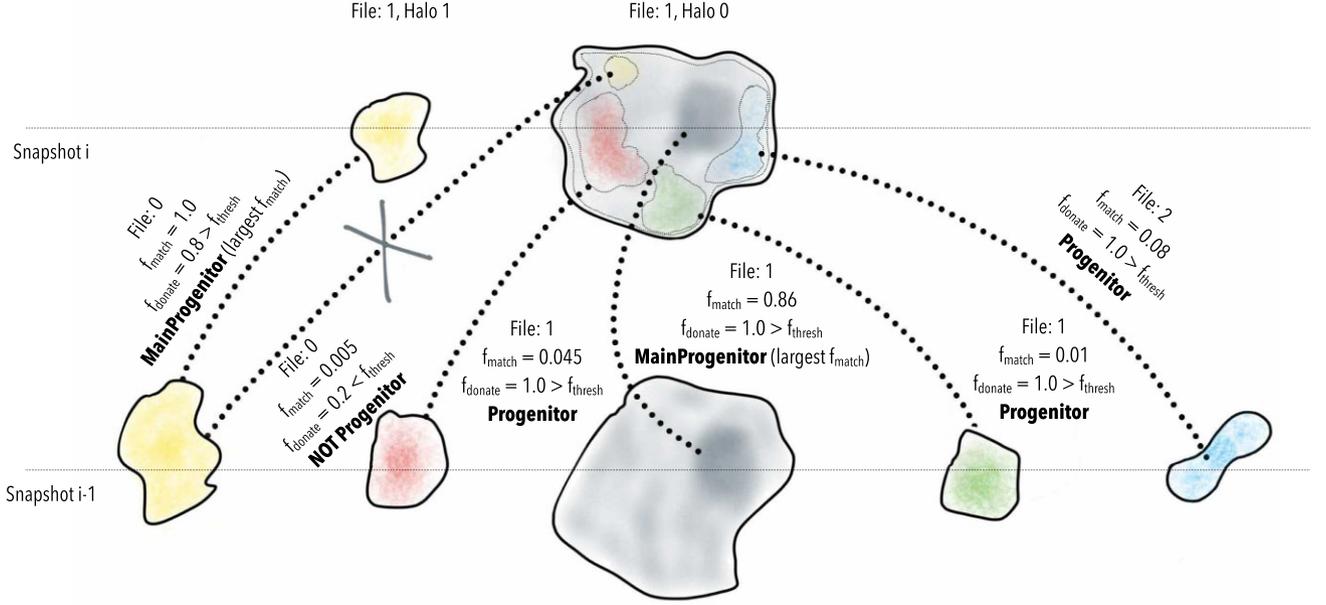


Figure 1. A schematic illustration of the merger tree algorithm. We identify associations between haloes across multiple time slices by tracking subsamples of halo particles that have been tagged. A candidate association is flagged when a halo in snapshot $i-1$ donates some fraction of its subsampled particles, f_{donate} , to a halo in snapshot i ; these are marked by the dotted lines. The candidate is then marked as a **Progenitor** if $f_{\text{donate}} \geq f_{\text{thresh}}$; throughout the ABACUSSUMMIT suite, we assume $f_{\text{thresh}} = 0.5$. The quantity f_{match} is defined as the fraction of subsampled particles in a halo in snapshot i that was found in a candidate association in snapshot $i-1$. The association with the largest f_{match} is marked as the **MainProgenitor**. Note that associations may be identified across slabs (labelled by File number in this illustration), although typically by never more than one superslab on either side.

halo, the joint list of its **Progenitors** and **MainProgenitors** across all output times defines its merger tree. To conveniently pinpoint a halo in the simulation, we define a unique halo identifier, **HaloIndex**, defined as

$$\begin{aligned} \text{HaloIndex} = & 1e12 * \text{FullStepNumber} \\ & + 1e9 * \text{SuperslabNumber} \\ & + \text{IndexInSuperslabFile}, \end{aligned}$$

where **FullStepNumber** is the number of the time-step at which the halo catalogue has been output, and **IndexInSuperslabFile** is the array index of the halo in the corresponding halo catalogue file. The **Progenitor** and **MainProgenitor** lists for a given halo are then simply collections of **HaloIndex** values appended over multiple output times.

As described in the methodology above, haloes identified at any given output time are matched across *two* preceding snapshots simultaneously. The primary advantage of this is that it allows us to keep track of fly-bys between pairs of haloes that might occur between one output time and the next. For example, consider the case of two haloes that were previously separate objects at snapshot $i-2$, ‘merge’ momentarily at snapshot $i-1$, and then again split away at snapshot i . In this instance, for at least one halo in this pair,

$$\begin{aligned} \text{MainProgenitor}(\text{MainProgenitor}(\text{halo at snapshot } i)) \\ \neq \text{MainProgenitorPrec}(\text{halo at snapshot } i) \end{aligned}$$

In other words, tracking **MainProgenitorPrec** allows us to ‘switch’ to the correct branch of the merger tree on the occasions where a transient fly-by occurs.

One subtlety that arises in our method pertains to step (ii) in the algorithm, where the list of plausible halo associations is narrowed down by first filtering out haloes that are located more than a

distance r_{max} at the preceding output time. Since haloes are split across multiple files in superslabs (see Section 2.2), it may be the case (particularly for haloes located close to the boundary of a superslab file) that several plausible halo associations may be found in other superslabs, although typically by never more than one superslab on either side (in the Base simulations, each superslab has a width of around $60 \text{ Mpc } h^{-1}$). For example, a halo located in superslab file 1 at snapshot i may have candidate association haloes (and, therefore, progenitors) located in superslab files 0, 1, or 2 at snapshot $i-1$ (see Fig. 1). For this reason, at any given step of our merger tree algorithm, we are required to simultaneously hold in memory halo and PID catalogues for nine superslabs at any given time: three per snapshot, and three snapshots in total ($i, i-1, i-2$).

Finally, in some rare circumstances, we identify associations where one halo is marked as the **MainProgenitor** of another, but does not also appear in the **Progenitor** list of this halo. While this may at first seem logically inconsistent, such a situation can occur when one halo overwhelmingly contributes the largest f_{match} , but does not actually donate the majority of its particles to the same halo (i.e. $f_{\text{donate}} < f_{\text{thresh}}$). This happens in the case of ‘splits’: When a low-mass ‘halo’ that once belonged to a larger aggregate unit breaks off at a later time, either due to the object’s orbital trajectory, which ejects it from its parent halo, or due to a deblending decision made by COMPASO. In this sense, a split may be physical (in the former case) or numerical (in the latter). In the outputs produced by the merger tree calculation, these incidences are marked with a boolean flag, **IsPotentialSplit**, which is set to 1 when the condition above is met. Note that these cases are infrequent: occurring in ~ 2 per cent of 150-particle haloes, and in 0.001 per cent of 15 000-particle haloes (see section 3.4 and table 2 of Hadzhiyska et al. 2022).

2.5 Code optimization routines

Our code for building halo merger trees from ABACUSSUMMIT has been developed to be used as a post-processing tool, and is written in PYTHON (although some of the packages it uses are PYTHON wrappers around C code). The task of determining halo associations for hundreds of millions of haloes across multiple output times for each simulation in the ABACUSSUMMIT suite is demanding: For this reason, we employ a few optimization routines that dramatically improved the performance of our code:

(i) Step (ii) of the algorithm, which first performs a neighbour search around the halo of interest to narrow down the list of all plausible halo associations from the previous output time significantly reduces the number of objects whose particle list have to be cross-matched. The neighbour search is performed using a multi-threaded tree search, for which we use the SCIPY package’s `cKDTree` implementation (Virtanen et al. 2020).

(ii) Steps (iii)–(v), which are performed on a per-halo basis, is accelerated by a factor of ~ 15 using the Numba just-in-time compiler (Lam, Pitrou & Seibert 2015), which translates PYTHON functions to optimized machine code at runtime.

(iii) The loop over all haloes in a single step is a classic example of an ‘embarrassingly parallel’ problem, whose speed scales linearly with the number of CPU processes employed for the task. We address this using the `joblib` library,³ which uses PYTHON’s multiprocessing functionality to achieve such a speedup. On Rhea,⁴ we use 16 cores for processing ABACUSSUMMIT merger trees, which yields another 16 times speedup.

Finally, we note that that merger trees are not constructed for every single halo in the ABACUSSUMMIT suite, but only for those in which particle tracking can be carried out reliably. We determine this to be the case when a halo contains at least 50 particles in total, with at least five subsampled particles available for tracking. The 50-particle limit corresponds to a lower limit of $1.05 \times 10^{11} M_{\odot} h^{-1}$ in halo mass at the Base mass resolution.

3 CADENCE OF OUTPUTS FOR BUILDING MERGER TREES

In this section, we present statistics of merger trees constructed from our algorithm for different cadences in the number of snapshots used to build halo trees. The aim of this section is to explore the influence of our choice of 33 intervals (cf. Section 2.3) for typical ABACUSSUMMIT simulations.

To this end, we have run a dedicated simulation set that we refer to as HIGHCADENCE. This simulation has a box size of $296 \text{ Mpc } h^{-1}$ with 1024^3 particles, resulting in an effective mass resolution that is identical to ABACUSSUMMIT Base resolution. We output 186 snapshots separated in cosmological scale factor by $\Delta \ln a \approx 0.007$, corresponding to roughly six times the density of outputs of a typical ABACUSSUMMIT simulation. To perform our tests, we subsequently build trees using all 186 (highest cadence), 64 (approximately three times coarser, medium cadence), and 33 outputs (approximately six times coarser, Summit-like cadence).

We begin by computing perhaps the simplest diagnostic tool for assessing halo merger trees that is their mass accretion histories, i.e. quantifying the mass growth of a halo as a function of redshift. The most common way of expressing this information is to traverse

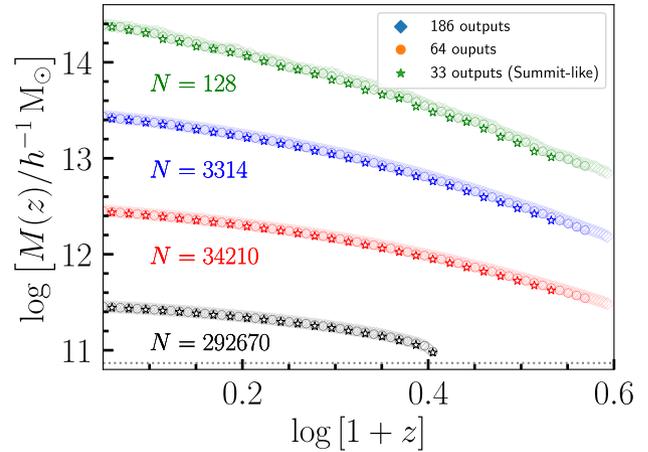


Figure 2. Median mass accretion histories of haloes identified at $z = 0.1$ in the HIGHCADENCE simulation from merger trees built using the fiducial spacing of output times used in ABACUSSUMMIT (stars, 33 outputs), and with approximately two and six times denser outputs (circles and diamonds, respectively). The grey horizontal dotted line marks the 50-particle limit, which corresponds to the smallest haloes for which merger tree associations are tracked. In each case, we follow haloes along the MainProgenitor branch of their merger tree. We find excellent agreement in the accretion histories irrespective of the density of outputs used to build the merger tree.

the main progenitor branch of individual halo merger trees, and to parametrize the accretion history as simply the mass evolution of the main progenitor.

Fig. 2 shows the median mass accretion histories associated with haloes in bins of mass centred at $\log [M_0/M_{\odot} h^{-1}] = [11.5, 12.5, 13.5, 14.5]$, each with a width of ± 0.2 dex. Here, M_0 is the mass at $z = 0.1$. The results for trees with 186, 64, and 33 outputs, respectively, are represented by diamonds, circles, and stars. It is clear from this figure that mass accretion histories are recovered equally well irrespective of the temporal density of the trees, with the agreement at $z \lesssim 1.5$ in particular being excellent. At higher redshift, there is a noticeable deviation in the agreement between the Summit-like accretion histories and the highest cadence trees, particularly in the most massive bin, where we average over only 128 haloes. The lower cadence merger trees tend to slightly underestimate the mass of the main progenitor at high redshift, which could potentially be due to the algorithm choosing a different main progenitor branch when there are two similar mass progenitors from which to choose. This can happen when a halo’s MainProgenitor in the previous time-step is identified as a singular entity in fine-grained trees (which looks over shorter time intervals), but may be split between two objects when looking back over the longer time interval in the coarser trees. However, the effect is small: In the most massive bin, the typical agreement (across all redshifts) in the mass of the MainProgenitor is of the order of ~ 4 per cent between the highest cadence and the Summit-like cadence trees. In lower mass bins, the agreement improves to within 1 per cent.

Fig. 3 shows a summary statistic that can be used to quantify the accretion histories of haloes, the so-called redshift of formation, z_{form} , as a function of halo mass. This parameter can be used to augment mass-only HOD models to e.g. incorporate the dependence of galaxy occupation and/or properties on halo age (e.g. Hearin & Watson 2013). We measure z_{form} by traversing the main progenitor branch of the merger tree and identifying the formation redshift as the epoch at which the halo achieves 50 per cent of its final-day mass.

³<https://joblib.readthedocs.io/>.

⁴<https://www.olcf.ornl.gov/olcf-resources/compute-systems/rhea/>.

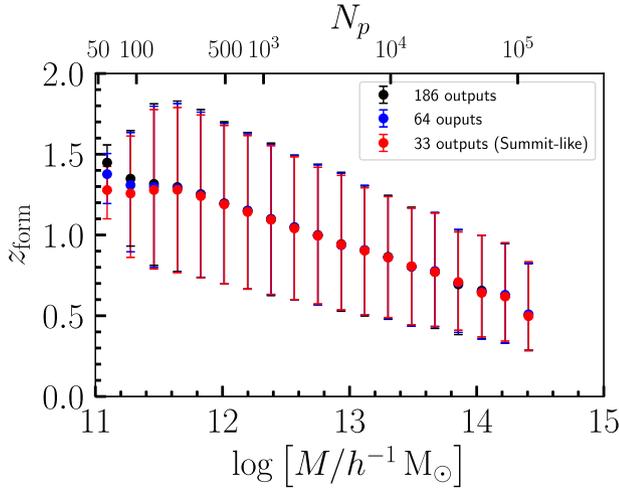


Figure 3. The redshift of formation, z_{form} , of haloes in the HIGHCADENCE simulation for objects identified at $z = 0.1$. The epoch of formation is defined as the redshift at which the halo attains 50 per cent of its final mass. The circles mark the median z_{form} , while the error bars indicate the 16th–84th percentile of this distribution. Both the median and the scatter in the measurement of formation time are near identical irrespective of whether 186 outputs or 33 outputs are used to construct the merger tree.

The mass dependence of halo formation time is shown in black, blue, and red, respectively, for the high, medium, and Summit-like cadence merger trees. The symbols represent the median value of z_{form} , while the error bars span the 16th–84th percentile of the distribution. As expected, there is a negative correlation between z_{form} and halo mass, indicating that low-mass haloes are older than high-mass haloes, which is the standard expectation from hierarchical structure formation. There is excellent agreement between all three versions of the merger tree calculation, both in the median and in the scatter. Encouragingly, the consistency between the three calculations is preserved all the way down to the 50-particle limit. This test demonstrates that for estimating the formation epoch of haloes, which is a common use case for merger trees generally, the Summit-like output cadence of 33 snapshots is sufficient.

Finally, we investigate how the density of outputs used to construct a merger tree affects the identification of progenitors beyond the primary branch. Fig. 4 shows the *conditional* mass function, which represents the fraction of mass contained in progenitor haloes of mass M_1 at $z = 1.1$ that merge into haloes of mass $10^{13} M_{\odot} h^{-1}$ by $z = 0.1$. We construct this quantity by stepping through every time slice in the merger tree, and recording the progenitors (and progenitors of progenitors) whose final descendant points to a $10^{13} M_{\odot} h^{-1}$ halo at $z = 0.1$ (of which there are 5912 examples in the HIGHCADENCE box).

The measurements from the high, medium, and Summit-like cadence merger trees, respectively, are represented by the black, blue, and red histograms. We find that the conditional mass function is reproduced near identically in all three cases, demonstrating the fact the complete progenitor histories of haloes are recovered even using the coarser spacing of snapshots used to build merger trees in ABACUSSUMMIT. The three versions of the trees agree down to progenitor mass ratios of $M_1/M_0 = 10^{-2}$, corresponding to haloes of mass $10^{11} M_{\odot} h^{-1}$ at $z = 1.1$, which contribute to ≈ 10 per cent of the final mass of the descendant halo. The dashed green line shows the conditional mass function predicted by the Monte Carlo algorithm of Parkinson et al. (2008), which is based on an extension of EPS,

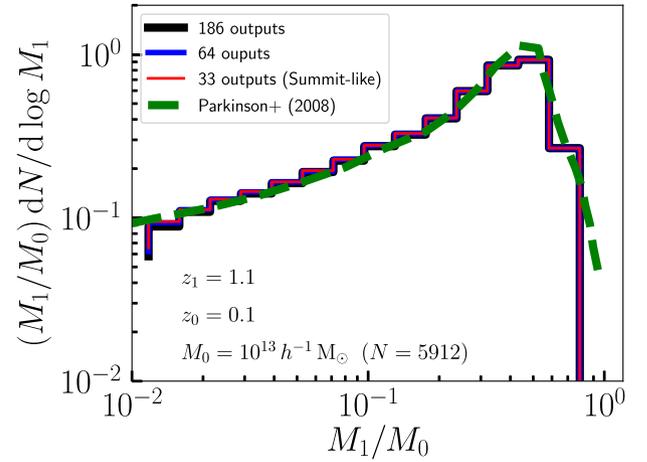


Figure 4. The fraction of mass contained in haloes of mass M_1 at $z = 1.1$ that merge into haloes of mass $10^{13} h^{-1} M_{\odot}$ by $z = 0.1$ measured in the HIGHCADENCE simulation. The so-called ‘conditional mass function’ is nearly identically irrespective of how many outputs are used to build the merger tree. The dashed green curve represents the conditional mass function predicted by the algorithm of Parkinson et al. (2008), which is based on a modified version of EPS theory.

calibrated to reproduce the results of the Millennium simulation (Springel et al. 2005). Reassuringly, we find very good agreement between the conditional mass function predicted by this model, to that extracted from the merger tree algorithm presented in this work.

The results in this section demonstrate that the snapshot spacing we have chosen for constructing trees in ABACUSSUMMIT is sufficient for halo-level merger tree statistics, particularly for the kinds of metrics that would typically be used in building mock galaxy catalogues. In the subsequent section, we apply our algorithm to the ABACUSSUMMIT suite, and present some initial applications of the resulting data products.

4 APPLICATIONS OF THE MERGER TREE

In the following subsections, we present some initial applications of the halo merger trees constructed from ABACUSSUMMIT using the algorithm that has been described in Section 2.4. While merger tree associations have been constructed for all simulations in the ABACUSSUMMIT suite, here we report results from the Highbase and High simulations only (both at the `c000_ph100` cosmology).

4.1 Halo mass accretion histories

In Fig. 5, we present the median mass accretion histories of haloes identified at $z = 0.5$ in the `c000_ph100` Highbase (solid lines and shaded regions) and High (symbols with error bars) simulations. Recall that the latter has six times better mass resolution in the same $1 \text{ Gpc } h^{-1}$ box. The symbols/lines represent the median accretion history while the error bars/shaded regions encompass the 16th and 84th percentiles of the scatter around the mean. Different colours represent accretion histories from different halo mass ranges, where we have set the bin centres to be at $\log [M_0/M_{\odot} h^{-1}] = [11.5, 12.5, 13.5, 14.5]$, each with a width of ± 0.2 dex. In Highbase, this yields 11 574 539, 1251 068, 96 822, and 1501 haloes, respectively.

We notice immediately that the two simulations show excellent agreement in the regions where they overlap, despite the factor of 6 difference in mass resolution. This provides a simple, yet reassuring,

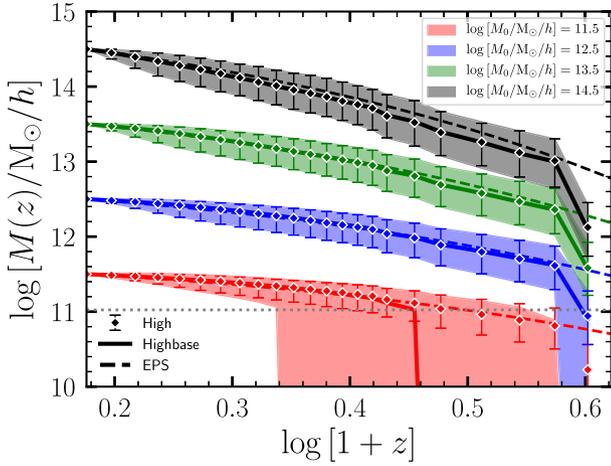


Figure 5. Median mass accretion histories of haloes identified at $z = 0.5$ in the `c000_ph100` cosmology at the `Highbase` (solid lines) and `High` (symbols) resolution, where the latter has six times better mass resolution than the former. The mass bins are centred at $\log [M_0/M_\odot h^{-1}] = [11.5, 12.5, 13.5, 14.5]$, each with a bin width of ± 0.2 dex. The shaded bands (error bars) encompass the 16th and 84th percentiles of the scatter in the accretion histories for `Highbase` (`High`). The dashed curves show predictions for the accretion histories using the semi-analytic model described in Correa et al. (2015). The grey horizontal dotted line marks the 50-particle limit in `Highbase`, which corresponds to the smallest haloes for which merger tree associations are tracked.

test of numerical convergence in two of our `ABACUSSUMMIT` simulations, at least from the perspective of this diagnostic for halo merger trees. This is seen most clearly in the lowest mass bin, where we see a sharp break in the accretion histories extracted from `Highbase` below the 50-particle limit, which is indicated using the horizontal dashed line in grey. Beyond this limit, we see a smooth transition into the regime that is now resolved in the `High` simulation.

For comparison with expected trends, Fig. 5 also shows predictions for the mass accretion history of haloes using the semi-analytic model developed by Correa et al. (2015), which is itself inspired by the extended Press–Schechter (EPS) formalism. In general, there is good agreement between this model and our measurements from the simulations, although there are slight differences in the most massive mass bin. This is perhaps unexpected, as the model by Correa et al. (2015) has been calibrated against a different set of numerical simulations. Furthermore, the comparison in the largest mass bin may be compromised by the relatively small number of objects that fall under this mass range (624).

4.2 Cleaning COMPASO catalogues

As we have discussed in the Introduction and at the end of Section 2.4, there is often ambiguity when it comes to reliably tracking the halo population in a cosmological simulation. A wide variety of processes including fly-bys, partial mergers, splits, or simply limited numerical resolution can undermine the robustness of halo properties output at any given time. Much like in the construction of the trees themselves, there are a variety of ways in which these fractures in the merger tree may be addressed. Examples include algorithms that try to ‘patch’ trees across gaps in a halo’s history by inserting additional haloes into the catalogue that guarantee smoother mass evolution (e.g. Tweed et al. 2009; Knebe et al. 2010; Behroozi et al. 2013b; Jiang et al. 2014). Halo splits (into multiple descendants) may well be treated

Table 1. A summary of the effects of the cleaning procedure described in Section 4.2 for three choices of $\kappa = 1.2, 2.0,$ and 20.0 .

$\log [M_{\text{halo}}/M_\odot h^{-1}]$	κ	f_{merged}			f_{inc}		
		1.2	2.0	20.0	1.2	2.0	20.0
11.5 ± 0.1		0.10	0.05	0.02	$5e-3$	$4e-3$	$4e-5$
12.5 ± 0.1		0.06	0.02	$6e-3$	0.34	0.16	0.04
13.5 ± 0.1		0.03	0.01	$8e-5$	0.97	0.89	0.81
14.5 ± 0.1		0.01	0.00	0.00	1.00	1.00	1.00

Notes. f_{merged} denotes the fraction of haloes within a specific mass bin (as listed in the first column) that are deleted as separate entities in the halo catalogue and merged on to a larger object. f_{inc} denotes the fraction of haloes in each mass bin whose mass is increased (relative to the default halo catalogue) after they are merged with the cleaned haloes. The numbers listed here are taken from the `c000_ph100_Highbase` simulation at $z = 0.5$.

as physically allowed scenarios, such as when two haloes that are in reality distinct dynamically but have at some point been treated as a single object by the halo finder. Characterizing the split objects appropriately requires careful consideration of their energy and phase space properties; recent works that describe the treatment of these splits include those by Han et al. (2018) and Roper, Thomas & Srisawat (2020).

In `ABACUSSUMMIT`, another instance when pathologies in the halo catalogue may arise is when the `COMPASO` algorithm has been overly aggressive in deblending single haloes into two or more components. This may happen due to the strict spherical overdensity criterion in `COMPASO` that draws hard edges in the ellipsoidal halo particle distribution, or indeed `COMPASO`’s eligibility condition that is adept at finding new halo nucleation points on the outskirts of haloes (see Hadzhiyska et al. 2022, for details). A general effect of these processes is to detach particles from haloes, leading to a reduction in their mass relative to the time when a potential galaxy hosted by this halo may have formed. A model like the HOD, which determines the galaxy occupancy of haloes based on their mass at any given time may therefore incorrectly predict the galaxy content of such objects.

To overcome these issues in `ABACUSSUMMIT`, we take a conservative approach and simply ‘clean’ the `COMPASO` halo catalogues of objects that may have been compromised by the processes listed above. One way to do this is to simply flag a halo that at some epoch passes through (or fly-by) another, sharing some of its particles in the process, before re-emerging as two separate objects at some later time. We then ‘merge’ the particles of the two haloes into a singular entity and, for all subsequent output times, assume that this agglomerated object remains as a persistent entity (unless it itself undergoes a future fly-by/splashback event). In general, this method of cleaning halo catalogues in post-processing removes ~ 1 – 5 per cent of objects (see Table 1), predominantly composed of a population of low-mass haloes floating around the boundaries of more massive haloes, and appends their particle lists to these neighbours. Note that the same technique will also result in re-merging overly deblended haloes.

The task of identifying and re-merging the offending haloes is simplified with the help of halo merger trees. In detail, our cleaning procedure proceeds as follows:

- (i) For a halo identified at redshift z_i , record its `COMPASO` mass, $M_{\text{halo}}(z_i)$.
- (ii) Traverse its merger tree along the `MainProgenitor` branch, and identify the epoch, z_{max} , at which it achieves its maximum mass, $M_{\text{halo}}(z_{\text{max}})$.
- (iii) If the ratio $M_{\text{halo}}(z_{\text{max}})/M_{\text{halo}}(z_i) > \kappa$, flag this halo for cleaning.

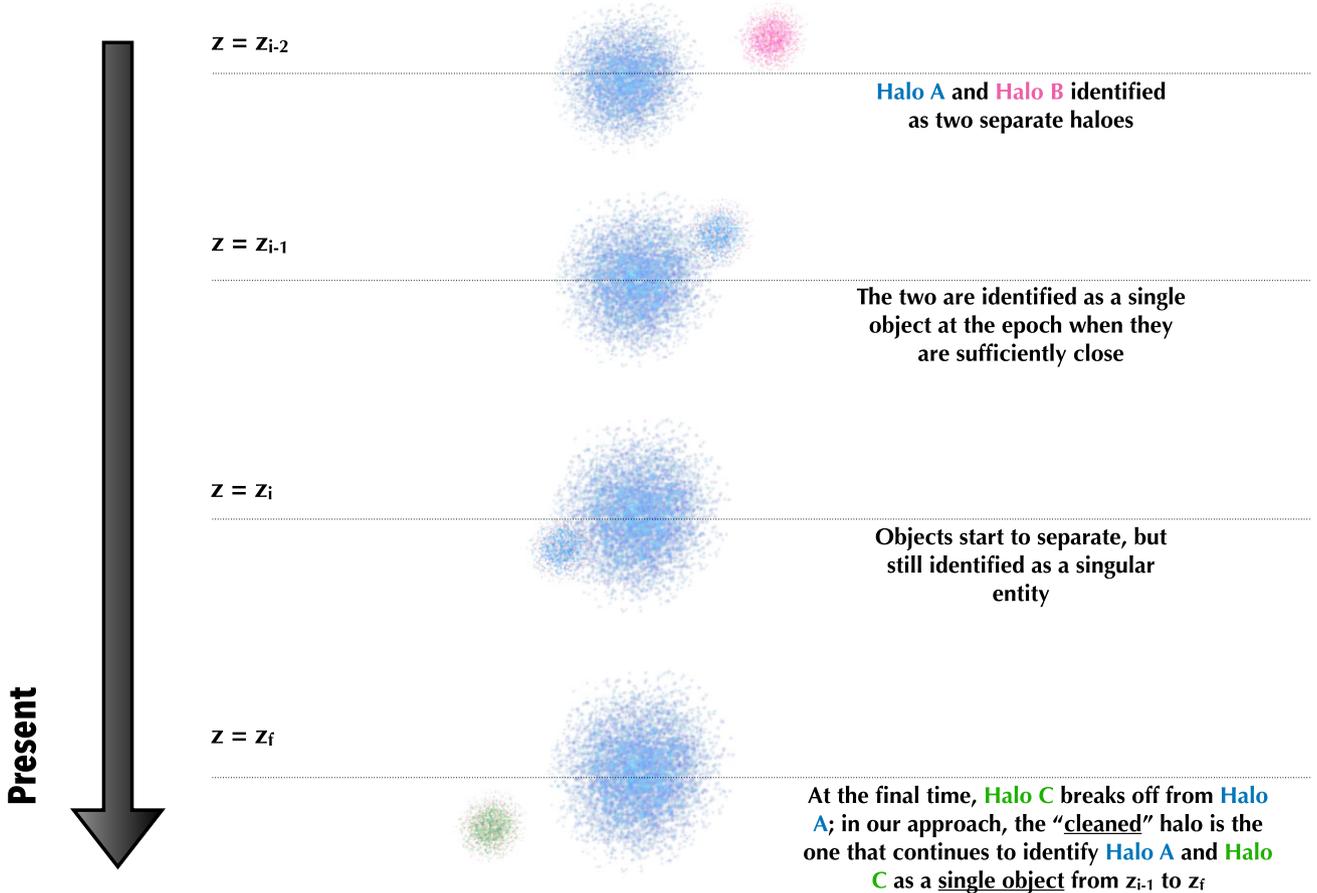


Figure 6. A schematic diagram showing the effect of cleaning as described in Section 4.2. At some early epoch, $z = z_{i-2}$, objects marked as Halo A (blue) and Halo B (pink) are identified as two distinct haloes. They subsequently interact and share particles such that they are then labelled as belonging to the same aggregate halo. At some later time, z_f , particles that once belonged to Halo B (or perhaps a subset of particles from B and some from A) may break off from this aggregate body, and be identified as some separate set C (green). In some halo catalogues, the green particles may be identified as a separate halo in its own right. In our ‘cleaning’ approach, we continue to treat the green and blue particles as one object.

(iv) At redshift z_{\max} , find the most massive halo whose `HaloIndex` is equal to the `MainProgenitor` of the flagged halo at this time. Once this entity has been found, add the mass of the flagged halo to the mass of its massive neighbour.

(v) The aggregated halo now remains merged for all subsequent outputs. The flagged halo is eliminated from the halo catalogue at each of the subsequent outputs. The particle list of the aggregate halo is the union of the particles of the original massive halo and *all* haloes that have been merged with it.

We display an illustration of the effect of cleaning in Fig. 6. Here, we start with having two sets of particles labelled as distinct haloes A (blue) and B (pink) at some early epoch, z_{i-2} . As these objects come together and interact, at some point, they will share enough particles that the halo finder will label them as belonging to the same aggregate halo (between z_{i-1} and z_i). As the erstwhile Halo B emerges out of this aggregate at z_f , it may now be composed of some combination of particles from A and B: we label these in green. Different halo catalogues will treat the green particles differently, typically identifying it as some new Halo C, whose mass and other properties will be the outcome of the complex interaction between Halo A and Halo B. In the cleaning procedure we have described in this section, the particles in the green and blue sets continue to be treated as a single, aggregate halo (i.e. all the way from z_{i-1}

to z_f . In other words, the particle list in green is appended to the particle list in blue. A concrete example of this procedure applied to the `c000_ph100` Highbase simulation is shown in Fig. 7, where in the different colours we highlight some of the haloes that are ‘merged’ on to some nearby massive neighbour following the cleaning strategy.

In very rare instances (~ 0.1 per cent of the time), a halo may be flagged for cleaning but no distinct host halo is identified in step (iv) of the procedure outlined above. This is typically the case for two initially separate objects that just come into contact for the first time, where COMPASO assigns more than half the particles associated with the smaller halo to its more massive neighbour. The larger halo will not be marked as the `MainProgenitor` of the smaller one, but may appear in the complete list of `Progenitors`. We then merge the smaller object with this halo. If no match is found even in the `Progenitors` list, we then simply merge the flagged halo with the nearest object that contains at least as many particles as the difference in the particle number between the peak and present-day mass of the flagged halo.

The method outlined in this subsection identifies objects whose present-day mass may be under-counted due to fly-bys, splits, partial merging, etc. This algorithm preferentially flags (and removes) low-mass objects, and append their particle lists to heavier, neighbouring haloes that they were once attached to. Note that for objects with

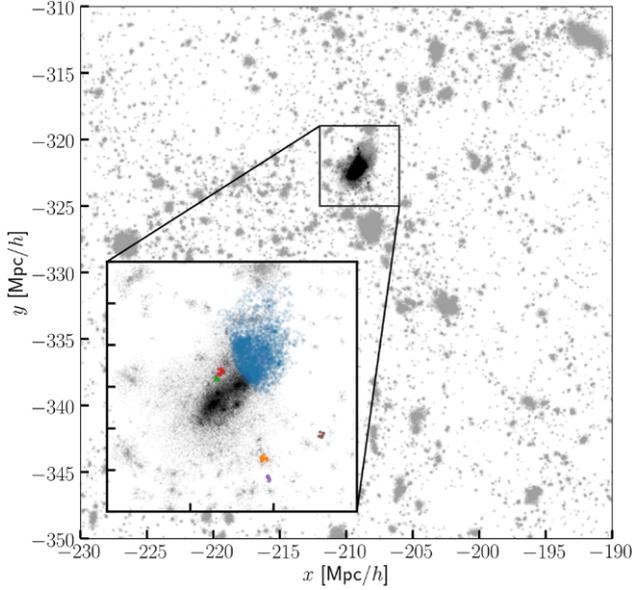


Figure 7. A visual depiction of the cleaning procedure outlined in Section 4.2. The background panel shows a $(40 \text{ Mpc } h^{-1})^3$ region where the most massive halo ($M_{\text{halo}} = 10^{15} M_{\odot} h^{-1}$) is represented in black. The remaining haloes are displayed in grey where, for clarity, we only show haloes more massive than $10^{11} M_{\odot} h^{-1}$. In each case, we only display the 10 per cent subsampled particles associated with each halo. The inset panel zooms into region of size $6 \times 6 \times 40 \text{ (Mpc } h^{-1})^3$, in which we have now highlighted the haloes (six in total, represented by the different colours) that are flagged by our cleaning algorithm and merged with the $10^{15} M_{\odot} h^{-1}$ halo. While most of the merged objects are of relatively low mass ($\sim 10^{11} M_{\odot} h^{-1}$), the halo represented by the blue points is a rather massive object, with $M_{\text{halo}} = 1.02 \times 10^{14} M_{\odot} h^{-1}$. This is a clear example where the COMPASO algorithm has debanded two haloes based on the density criteria described in Section 2.2. The net effect of merging these haloes with the neighbouring cluster is to increase the mass of the central host by roughly 10 per cent.

‘well-behaved’, smooth merger histories, $M_{\text{halo}}(z_i) = M_{\text{halo}}(z_{\text{max}})$, so they will not be flagged by this procedure. Our method has a single free parameter, κ , for which our recommended choice is $\kappa = 2.0$. Based on our tests, however, we find that the properties of the cleaned halo catalogues do not change significantly for small variations in κ . In the following subsections, we present comparisons of the statistics of the halo population in the cleaned and default COMPASO catalogues.

4.3 Halo mass functions

Fig. 8 compares the halo mass function (measured at $z = 0.5$) for the `c000_ph100` Highbase simulation for the default (red) and cleaned ($\kappa = 1.2$: brown; $\kappa = 2.0$: blue; $\kappa = 20.0$: green) versions of the COMPASO catalogues. In black, we also show the halo mass function measured by the ROCKSTAR halo finder (Behroozi et al. 2013a) run on the same simulation. The most significant difference in the way in which ROCKSTAR operates is that is a phase-space halo finder, utilizing both the positions and velocities of dark matter particles to identify haloes. This provides a very useful, independent method to compare the predictions of COMPASO and its variations in the form of the cleaned catalogues.

Considering just ROCKSTAR and the default COMPASO (in red) first, we find that there are noticeable differences on both extremities of the halo mass function. In particular, COMPASO shows an ~ 10 per cent excess in the abundance of $10^{11} M_{\odot} h^{-1}$ haloes. On the other hand,

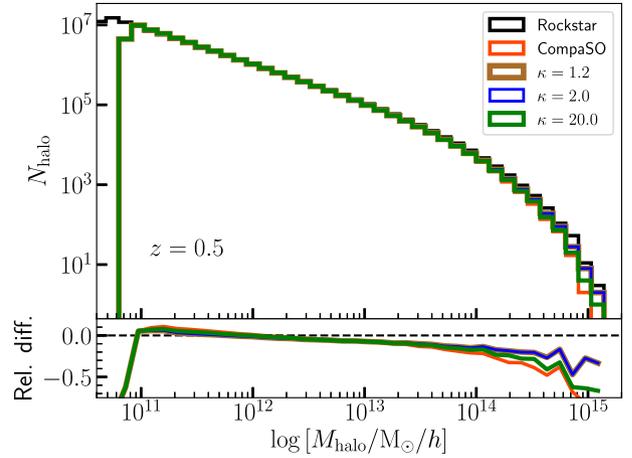


Figure 8. A comparison of the halo mass functions extracted from the Highbase `c000_ph100` simulation at $z = 0.5$. Halo finding has been performed using both ROCKSTAR (black) and COMPASO (red). The ‘cleaned’ versions of the COMPASO catalogues obtained after assuming cleaning parameters, $\kappa = 1.2, 2.0$, and 20.0 are shown, respectively in brown, blue, and green (see Section 4.2 for details of the cleaning procedure). The standard COMPASO mass function deviates from ROCKSTAR for haloes more massive than $\sim 10^{13} M_{\odot} h^{-1}$, being suppressed relative to it by 50 per cent at $M_{\text{halo}} \approx 10^{15} M_{\odot} h^{-1}$. The ‘cleaned’ catalogues bring the COMPASO mass functions into better agreement with ROCKSTAR: $\kappa = 2.0$, which is our fiducial choice, reduces the discrepancy to 25 per cent; there is no discernible difference in the $\kappa = 1.2$ case. For $\kappa = 20.0$, which is the least aggressive cleaning choice, there is only a marginal difference relative to the default COMPASO catalogue.

at the high-mass end, COMPASO shows an ~ 50 per cent *deficiency* in the abundance of $10^{15} M_{\odot} h^{-1}$ haloes. While some of this difference stems inevitably from the different methods used to define haloes in ROCKSTAR and COMPASO, the full extent of the discrepancy cannot be attributed to differences in halo finding.

The green curve shows the modification to the COMPASO halo mass function after applying the cleaning strategy described in Section 4.2 assuming $\kappa = 20.0$. This is the most conservative of three parameter variations we examine in this paper, flagging only those haloes for cleaning that were once part of haloes with mass $\geq 20 \times$ their mass at $z = 0.5$. The net result is a very small change with respect to the default COMPASO case, only marginally shifting the mass function to higher masses on scales greater than $\sim 10^{14} M_{\odot} h^{-1}$. The $\kappa = 2.0$ version (in blue) shows a much more dramatic change, particularly in the regime of low-mass groups and clusters, reducing the discrepancy relative to ROCKSTAR to the 25 per cent level. Finally, we see no noticeable change in transitioning from $\kappa = 2.0$ to 1.2 (in brown), where the two curves overlap almost exactly.

In Table 1, we contrast the fraction of haloes (as a function of mass) that are affected by the cleaning method for each choice of κ , either because they are removed from the COMPASO catalogue and merged, or because their reported mass has been augmented by the addition of merged neighbouring haloes. We see that, as expected, low-mass haloes are the ones that are more likely to be flagged for removal, while a larger fraction of the more massive haloes end up with increased masses after being merged with several smaller haloes. Furthermore, we see that the $\kappa = 1.2$ and 2.0 cases yield very similar results, consistent with what we see in Fig. 8; on the other hand, the choice of $\kappa = 20.0$ tends to be far less effective at removing intermediate-mass haloes.

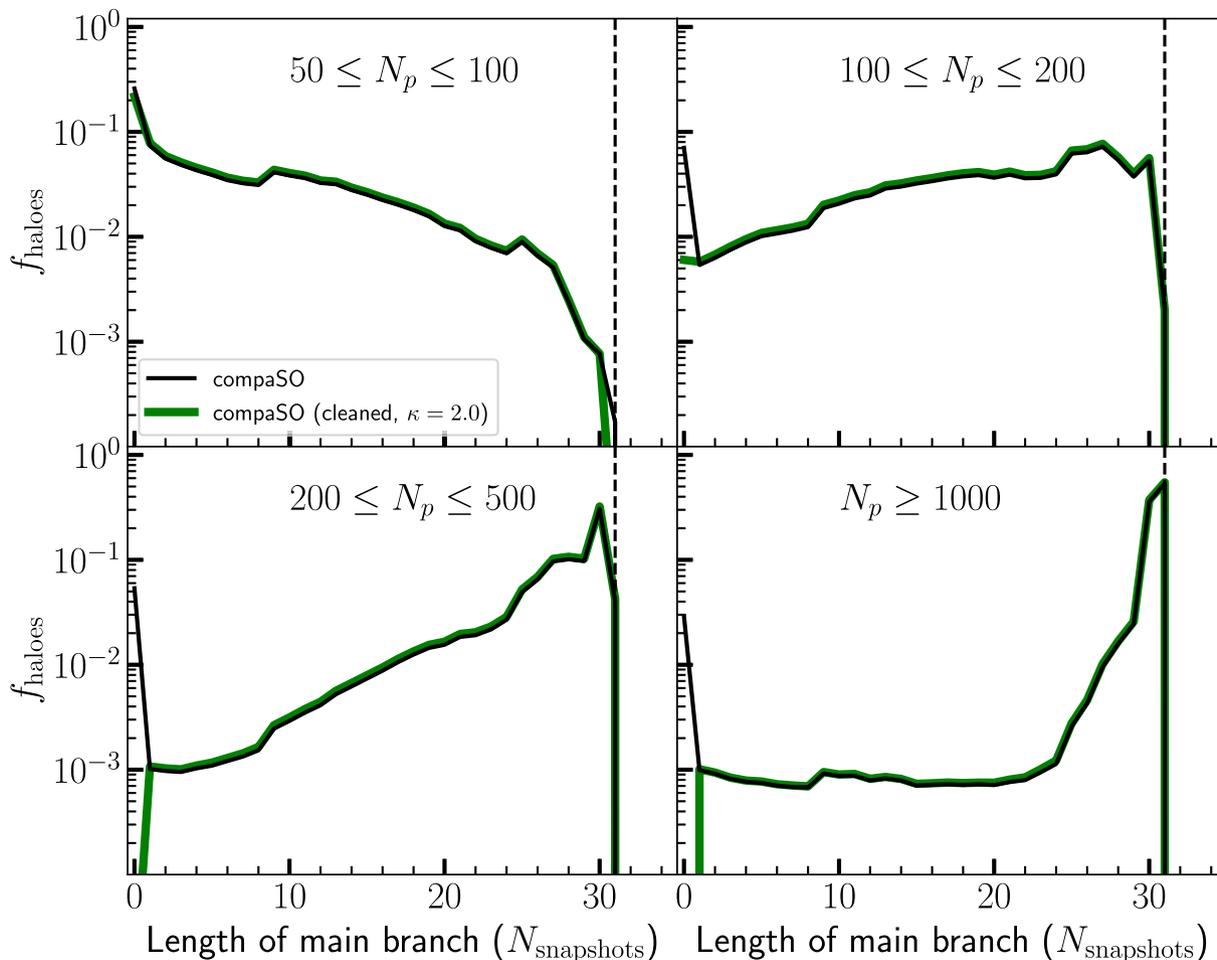


Figure 9. A comparison of the length of the main progenitor branch of haloes extracted from the Highbase c000_ph100 simulation at $z = 0.1$. The length of the main progenitor branch is defined as the number of snapshots through which the `MainProgenitor` of any given halo can be tracked back in time. Results for the standard COMPASO halo catalogue are shown in black, while ‘cleaned’ COMPASO haloes are shown in green. The four panels show the results from four different mass bins selected according to the number of particles in the halo, N_p , at $z = 0.1$. f_{haloes} denotes the fraction of haloes in each mass bin that can be tracked for the corresponding number of snapshots. The dashed vertical line marks the total number of snapshots that have been used to create the merger tree (excluding the final output). The standout feature is that in the cleaned version, there is a sharp decrease in the number of haloes that can only be tracked for zero or one snapshots. These correspond predominantly to the population of filtered haloes (i.e. those that have been flagged for cleaning). It is also clear that haloes resolved with more particles are easier to track back in time.

4.4 Length of the main progenitor branch

In the merger tree comparison project presented in Srisawat et al. (2013), a particularly useful metric for comparing different algorithms turns out to be a measure of the length of the main progenitor branch of haloes – essentially, the number of snapshots into the past that one is able to reliably track a halo’s main progenitor.

Fig. 9 shows the results of measuring the length of the main progenitor branch for the default (black) and cleaned (green) COMPASO catalogues. To perform the cleaning, we have assumed our fiducial value of $\kappa = 2.0$. The two curves largely overlap, suggesting that cleaning does not change this measure significantly, which is to be expected since the objects that are affected by cleaning comprise a small fraction of the full population of haloes (see Table 1). The most prominent difference is the sharp reduction haloes that can only be tracked for zero or one snapshots in the cleaned catalogue. This is dominated by objects in the default COMPASO catalogue that have been flagged and identified for merging by the algorithm in Section 4.2. This fraction is larger in the lower mass bins ($N_p \geq 500$), as expected.

From Fig. 9, we also note that haloes that are better resolved are more easily tracked through multiple output times, with the vast majority of haloes ($\gtrsim 70$ per cent) with $N_p \geq 1000$ being tracked across the entire length of the merger tree (denoted by the dashed vertical line). Finally, we find that in the lowest mass bin considered ($50 \leq N_p \leq 100$), it becomes increasingly difficult to track haloes through several snapshots. The reason for this is that reliably associating haloes across snapshots becomes unstable when only a few particles are available for matching. It is also limited by the choice we make in the construction of the merger tree in which we require that a halo contain at least five subsampled particles in order to be considered for tracking. An increasing fraction of haloes fail this test as we proceed to earlier snapshots.

An interesting way to further assess the fidelity of the merger trees for cleaned versus uncleaned haloes would be through visual characterization, such as the dendrogram diagnostic proposed by Poulton et al. (2018). We leave such an exercise for future study.

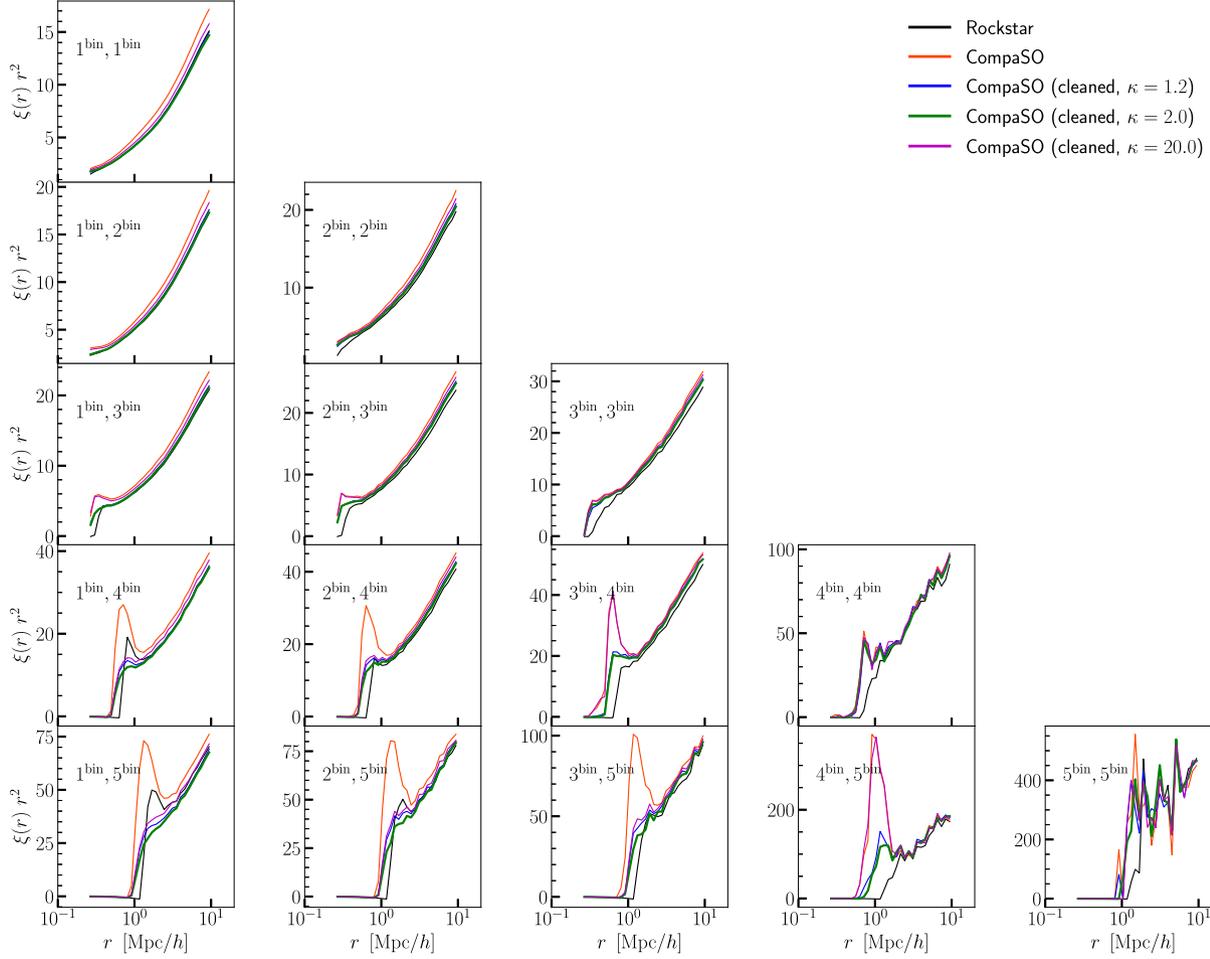


Figure 10. The two-point clustering of haloes at $z = 0.5$ measured using ROCKSTAR (black), COMPASO (red), and the ‘cleaned’ COMPASO (with $\kappa = 1.2$, blue; $\kappa = 2.0$, green; $\kappa = 20.0$, magenta) catalogues for the Highbase c000_ph100 simulation. Haloes are split into five distinct mass bins: $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = 11.5 \pm 0.1$ (first bin); $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = 12.0 \pm 0.1$ (second bin); $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = 12.5 \pm 0.1$ (third bin); $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = 13.5 \pm 0.1$ (fourth bin); and $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = 14.5 \pm 0.3$ (fifth bin). Individual panels show the auto- and cross-correlation functions of haloes identified between pairs of mass bins. The unprocessed COMPASO catalogue shows a strong excess of low-mass haloes in the outskirts of galaxy clusters, which manifests as a sharp peak in the cross-correlation function on scales $\sim 1 \text{ Mpc } h^{-1}$, well in excess of what is measured for ROCKSTAR haloes. Cleaning the halo catalogue using the strategy described in Section 4.2 brings COMPASO into much closer agreement with ROCKSTAR, although there are residual differences in the highest mass bins. We also find that, as in the case of the halo mass function (Fig. 8), the difference between the $\kappa = 1.2$ and 2.0 catalogues is almost negligible. The $\kappa = 20.0$ case, which results in the least aggressive cleaning, shows marked differences and yields results that are more comparable to the default COMPASO haloes.

4.5 2-point clustering: the auto- and cross-correlation functions of haloes

In the previous subsections, we have seen that applying our COMPASO catalogue cleaning procedure has the effect of flagging low-mass objects ($\log [M_{\text{halo}}/M_{\odot} h^{-1}] \lesssim 14.5$) on to larger, neighbouring haloes. This suggests that a large number of ‘offending’ haloes (i.e. those that are flagged and merged by our cleaning algorithm) corresponds to low-mass haloes that exist in the peripheries of more massive haloes. This assertion can be tested by examining the clustering of haloes, measured using the two-point correlation function.

We calculate the correlation function, $\xi(r)$, using the Peebles (1980) estimator:

$$\xi(r) = \frac{\langle DD(r) \rangle}{\langle RR(r) \rangle} - 1,$$

where $\langle DD(r) \rangle$ and $\langle RR(r) \rangle$, respectively, are the normalized data–data and random–random pair counts of objects separated by a

distance r . To examine the clustering as a function of mass, we split the halo population into five mass bins: $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = 11.5 \pm 0.1$, 12.0 ± 0.1 , 12.5 ± 0.1 , 13.5 ± 0.1 , and 14.5 ± 0.3 . In practice, we compute $\xi(r)$ using the CORRFUNC package (Sinha & Garrison 2020).

Fig. 10 shows the bin-by-bin auto-correlation (diagonal panels) and cross-correlation (off-diagonal panels) functions of haloes identified by ROCKSTAR (black), COMPASO (red), as well as the cleaned COMPASO catalogue obtained using our recommended choice of $\kappa = 2.0$ (green), as well as $\kappa = 1.2$ (blue) and 20.0 (magenta). The clustering is measured at $z = 0.5$ in the Highbase c000_ph100 simulation. In each case, we have taken care to ensure that the clustering is computed at the same halo abundance; this guarantees that any systematic differences measured in $\xi(r)$ are as a result of physical differences in the three halo catalogues, rather than being due to a different number of objects being included based on the mass definition used in each case.

Upon comparing the clustering in the default COMPASO catalogue with ROCKSTAR, we immediately notice significant differences. It is clear that COMPASO has a greater tendency to find low-mass dark matter haloes and identify them as independent objects in their own right. This manifests first as a larger amplitude in the auto-correlation function of $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = 11.5$ and more prominently as the sharp spike in the cross-correlations of low-mass haloes with heavier objects. In particular, the spike appears around the transition region between the one- and two-halo terms in the cross-correlation functions ($r \approx 1\text{--}3 \text{ Mpc } h^{-1}$), which corresponds roughly to the typical virial radii of the most massive clusters in our sample. ROCKSTAR, on the other hand, does not show such a pronounced feature in the same regime. In other words, the COMPASO method finds an excess of low-mass dark matter structures floating around the boundaries of massive haloes; ROCKSTAR seemingly does not deblend these sets into populations of independent objects.

Fig. 10 also shows that much of the discrepancy between ROCKSTAR and COMPASO disappears after the cleaning procedure is applied to the COMPASO haloes. By merging haloes that our method flagged for cleaning with heavier neighbours, we eradicate the vast majority of the population of haloes that were originally found at the boundaries of massive clusters. This results in a dramatic suppression in the prominent spike that previously punctuated the measurement of $\xi(r)$ for COMPASO haloes. That said, the auto-correlation functions of haloes heavier than $\log [M_{\text{halo}}/M_{\odot} h^{-1}] \geq 12.5$ remain largely unaffected by cleaning, and they show appreciable differences with respect to ROCKSTAR on scales smaller than $r \lesssim 2 \text{ Mpc } h^{-1}$. We note, however, that matching ROCKSTAR exactly is not our intention – nor, indeed, is it necessary – as some differences between the two catalogues will inevitably persist due to the different philosophies used for identifying dark matter haloes.

We also find that, consistent with our observations of the halo mass function in Fig. 8, $\kappa = 1.2$ and 2.0 yield identical results in all panels. Both choices lead to an ~ 15 per cent reduction in the auto-correlation of $\log [M_{\text{halo}}/M_{\odot} h^{-1}] \leq 12.0$ haloes. The $\kappa = 20.0$ case results in a more modest change (5–10 per cent in the auto-correlation function). The biggest contrast to the lower κ (i.e. more aggressive) cleaning choices is seen in the cross-correlations between massive haloes. This is seen most prominently in bins (3,4) and (4,5) of Fig. 10, where we see that the $\kappa = 20.0$ choice has no real cleaning effect relative to the default COMPASO halo catalogue.

4.6 Radial infall velocity profiles

Thus far, we have seen that the population of haloes that is preferentially removed by our cleaning method corresponds to those that are typically located at the boundaries of massive clusters. Their presence manifests as a sharp bump in the cross-correlation function on the scale of $1\text{--}3 \text{ Mpc } h^{-1}$, which is comparable to the typical virial radius of low-mass clusters. Next, we examine the dynamics of the ‘cleaned’ objects as measured by their infall velocities in the vicinity of the neighbouring clusters.

Fig. 11 shows the mean (peculiar) radial velocity profile of a tracer population of haloes, measured relative to clusters in the mass range $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = [14.0, 14.5]$. For the tracer population, we select haloes in the mass range $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = [11.0, 13.0]$. The profiles are shown for the default COMPASO (black) and cleaned COMPASO (with $\kappa = 2.0$, green) catalogues, both measured at $z = 0.5$ in the `Highbase c000_ph100` simulation. The peculiar radial velocity, $\langle v_{\text{rad}} \rangle$, is measured as a function of distance from the cluster centre. The upper axis denotes this distance in units of r_{95} , which is

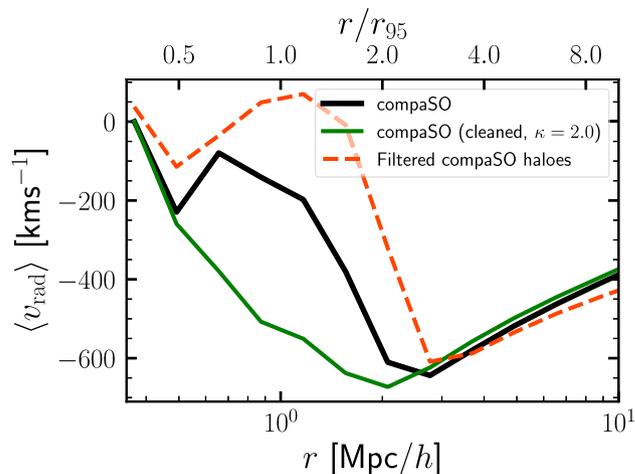


Figure 11. The mean peculiar radial velocity, $\langle v_{\text{rad}} \rangle$, of haloes in the mass range $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = [11.0, 13.0]$, relative to a sample of low-mass cluster haloes in the mass range $\log [M_{\text{halo}}/M_{\odot} h^{-1}] = [14.0, 14.5]$. The curves in black and green, respectively, show results from the COMPASO and cleaned COMPASO catalogues, where we assume $\kappa = 2.0$. The dashed red curve shows the radial profile for only those haloes that have been flagged and removed due to cleaning (i.e. merged into a larger halo) according to the procedure described in Section 4.2. The upper x-axis denotes the radial range in units of r_{95} from the cluster centre, where r_{95} is the radius encompassing 95 per cent of the mass of the halo. The cleaning tends to remove haloes with more positive radial velocities in and around the outskirts of clusters ($1\text{--}3 \text{ Mpc } h^{-1}$), with a sizeable population with $\langle v_{\text{rad}} \rangle > 0$. This is likely representative of population of haloes that have exited their host clusters following at least one orbital passage.

the radius that encompasses 95 per cent of the mass of the cluster measured from its L2 centre of mass.

First, we see that the radial velocity profiles measured in both the default and cleaned COMPASO cases show similar overall trends. As haloes get increasingly close to the cluster centre from $r \sim 10 \text{ Mpc } h^{-1}$ to about $r \sim 3 \text{ Mpc } h^{-1}$, $\langle v_{\text{rad}} \rangle$ becomes increasingly negative, signifying the fact that these objects are infalling into the cluster. As these objects reach the vicinity of the virial radius of the cluster (comparable to the value of r_{95}), there is a sharp upturn in the profile, signifying that the velocities of the tracer haloes begin to virialize within the cluster. The system then reaches equilibrium in the innermost portion of the halo, where $\langle v_{\text{rad}} \rangle = 0$.

The behaviour of the COMPASO haloes before and after cleaning are identical on scales larger than about $3 \text{ Mpc } h^{-1}$. Within this radius, however, there is a systematic deviation in the profile, with the haloes in the cleaned catalogue showing a more negative infall velocity on average than in the default COMPASO case. Notably, this change is seen in and around the virial radius of the cluster haloes, which is indeed the scale where the cross-correlation function of haloes in each catalogue differed the most.

We can understand the origin of this shift by considering the mean radial infall velocity profile of those objects that have been flagged by our cleaning procedure and removed from the COMPASO catalogue; in Fig. 11, this is shown as the dashed red curve. The profile for the filtered objects is shifted systematically to more positive values of $\langle v_{\text{rad}} \rangle$. It is especially noticeable that $\langle v_{\text{rad}} \rangle > 0$ in the regime between 0.8 and $2 r_{95}$, which is indicative of a population of haloes whose radial motions point away from the cluster centre. These are likely examples of ‘splashback’ haloes – those that were once part of a larger object, but have since exited following at least one orbital

passage within their former hosts. The cleaning method we have applied guarantees that the past merger tree of the filtered objects once shared a main progenitor with their neighbouring massive halo, which adds further confidence to our notion that the red curve in Fig. 11 receives a large contribution from these splashback objects. By removing the contribution of these objects from the mean radial infall velocity profiles, we more easily isolate the population of objects in the boundaries of massive clusters that are actually infalling at present. This also has a small effect in changing the location of the upturn in the velocity profile, shifting from $\approx 3 \text{ Mpc } h^{-1}$ to around $2 \text{ Mpc } h^{-1}$ between the default and cleaned COMPASO catalogues.

4.7 From cleaned haloes to mock galaxy catalogues

As a final test of the value of our cleaning method, we apply a set of ‘cleaned’ ABACUSSUMMIT simulations in a real use-case scenario: building mock galaxy catalogues to match the number density and redshift-space clustering of an observed galaxy sample.

Specifically, we utilize galaxies from the CMASS sample of the Baryon Oscillation Spectroscopic Survey (BOSS; Bolton et al. 2012; Dawson et al. 2013) Data Release 12, part of the Sloan Digital Sky Survey (SDSS) III programme (Eisenstein et al. 2011). In particular, we focus on galaxies in the redshift range $0.46 < z < 0.6$, yielding a sample of approximately 600 000 galaxies (predominantly Luminous Red Galaxies) at an average number density of $n_{\text{gal}} = (3.01 \pm 0.3) \times 10^{-4} h^3 \text{ Mpc}^{-3}$. The redshift-space clustering of the sample is determined by the anisotropic two-point correlation function, $\xi(r_p, \pi)$. r_p and π , respectively, are the transverse and line-of-sight galaxy separation in comoving units. We utilize the full-shape anisotropic clustering instead of the projected clustering as it provides stronger constraints on the HOD (Yuan et al. 2021). We determine the anisotropic two-point correlation function $\xi(r_p, \pi)$ using the Landy & Szalay (1993) estimator:

$$\xi(r_p, \pi) = \frac{DD - 2DR + RR}{RR},$$

where DD , DR , and RR are the data–data, data–random, and random–random pair counts in bins of (r_p, π) . For the present analysis, we choose eight logarithmically spaced bins between 0.169 and $30 \text{ Mpc } h^{-1}$ in the transverse direction, and six linearly spaced bins between 0 and $30 \text{ Mpc } h^{-1}$ bins along the line-of-sight direction.

To produce the galaxy mock, we make use of the ABACUSSUMMIT `Base_c000_ph000` simulation, which provides an eight times larger volume than the `Highbase` box we have considered until now. The mass resolution ($\sim 2 \times 10^9 M_\odot h^{-1}$ per particle) is more than sufficient for the CMASS galaxies, which typically live in haloes of $> 10^{12} M_\odot h^{-1}$. We populate COMPASO dark matter haloes with galaxies using the ABACUSHOD⁵ code (Yuan et al. 2022), which determines the galaxy–halo connection using sophisticated particle-based HOD models. Specifically, the HOD we use for this analysis is based on the Zheng & Weinberg (2007) model where, respectively, the mean central and satellite occupations in haloes with mass M_h are defined as

$$\langle N_{\text{cen}}(M_h) \rangle = \frac{1}{2} \text{erf} \left[\frac{\log(M_{\text{cut}}/M_h)}{\sqrt{2}\sigma} \right], \quad (1)$$

and

$$\langle N_{\text{sat}}(M_h) \rangle = \left(\frac{M_h - \kappa_0 M_{\text{cut}}}{M_1} \right)^\alpha \langle N_{\text{cen}}(M_h) \rangle. \quad (2)$$

⁵<https://abacusutils.readthedocs.io/en/latest/hod.html>.

Here, M_{cut} , σ , κ_0 , M_1 , and α are model parameters determined through fitting. M_{cut} is the minimum mass of haloes that host a central galaxy, and σ is the width of the transition from non-zero to zero central occupancy. Furthermore, M_1 is the characteristic mass of haloes that host satellites, while $\kappa_0 M_{\text{cut}}$ represents the minimum halo mass for hosting at least one satellite. α is the power-law slope of the satellite occupation. Galaxy positions in ABACUSHOD are assigned such that the central galaxy is located at the host halo centre, while satellites are placed on halo particles. Note that a more common notation for κ_0 is simply κ ; we have chosen the former to avoid confusion with the cleaning factor, κ , we have adopted until now.

Because of redshift-space distortions, the line-of-sight velocities of the galaxies are translated into biases in the measurements of galaxy positions. It is thus important to have a model of the galaxy velocities, particularly their line-of-sight component. In this analysis, we allow for velocity bias effects, which introduces flexibilities in the velocities of the central and satellite galaxies (Guo et al. 2015). Specifically, the central galaxy velocity along the line of sight is given by

$$v_{\text{cent},z} = v_{L2,z} + \alpha_c \delta v(\sigma_{\text{LOS}}), \quad (3)$$

where $v_{L2,z}$ denotes the line-of-sight component of the L2 centre-of-mass velocity, $\delta v(\sigma_{\text{LOS}})$ denotes the Gaussian scatter, and α_c is the central velocity bias parameter, which modulates the amplitude of the central velocity bias effect. By definition, α_c is non-negative, and $\alpha_c = 0$ corresponds to no central velocity bias, in which case the central galaxy simply assumes the centre-of-mass velocity of the L2 subhalo. The satellite velocity is given by

$$v_{\text{sat},z} = v_{L2,z} + \alpha_s (v_{p,z} - v_{L2,z}), \quad (4)$$

where $v_{p,z}$ denotes the line-of-sight component of particle velocity, and α_s is the satellite velocity bias parameter. $\alpha_s = 1$ corresponds to no satellite velocity bias, in which case the satellite simply assumes the velocity of the particle. In summary, our HOD model consists of seven parameters: five in the baseline HOD model (M_{cut} , M_1 , σ , α , κ_0), and two velocity bias parameters (α_c , α_s). ABACUSHOD is also capable of implementing multi-tracer occupation, and presents additional important extensions to the baseline HOD model including assembly bias and satellite profile flexibility (see Yuan, Eisenstein & Garrison 2018, for details). We do not incorporate these extensions in this work. A full analysis enabling all relevant extensions will be described in Yuan et al. (2022).

In order to determine the best-fitting values of the free parameters in equations (1) and (2), we perform a joint fit of the CMASS DR12 galaxy number density, n_{gal} , and the redshift-space two-point correlation function, $\xi(r_p, \pi)$. The quality of the fit is determined using a χ^2 statistic, defined as

$$\chi^2 = \chi_\xi^2 + \chi_{n_{\text{gal}}}^2,$$

where

$$\chi_\xi^2 = (\boldsymbol{\xi}_{\text{mock}} - \boldsymbol{\xi}_{\text{data}})^T \mathbf{C}^{-1} (\boldsymbol{\xi}_{\text{mock}} - \boldsymbol{\xi}_{\text{data}}),$$

and

$$\chi_{n_g}^2 = \begin{cases} \left(\frac{n_{\text{mock}} - n_{\text{data}}}{\sigma_n} \right)^2 & (n_{\text{mock}} < n_{\text{data}}) \\ 0 & (n_{\text{mock}} \geq n_{\text{data}}). \end{cases}$$

Here, \mathbf{C} is the jackknife covariance on ξ , and σ_n is the jackknife uncertainty on the galaxy number density. f_{ic} is an incompleteness factor that uniformly downsamples our mocks to match the desired number density. We sample the posterior space with the DYNESTY nested sampling code (Speagle 2020). Full details of this fitting procedure are described in section 4.1 of Yuan et al. (2021).

Table 2. A summary of the best-fitting HOD parameters (first seven rows) derived from mocks built on the uncleaned (left-hand column) and cleaned (right-hand column) versions of the `Base_c000_ph000` simulation, assuming $\kappa = 2.0$.

	COMPASO	COMPASO (cleaned, $\kappa = 2.0$)
M_{cut}	12.887	12.857
M_1	14.387	14.055
$\log \sigma$	-2.651	-1.807
α	0.813	1.189
κ_0	0.335	0.343
α_c	0.226	0.227
α_s	1.217	0.913
$\log \mathcal{Z}$	-76	-62
$\log \mathcal{L}$	-57	-43

Notes. The last two rows, respectively, list the Bayesian evidence, $\log \mathcal{Z}$, and log-likelihood, $\log \mathcal{L}$, of the best-fitting mocks. The corresponding values suggest a significant preference for the cleaned catalogue over the uncleaned version.

Table 2 demonstrates the significant improvement in the fit to $\xi(r_p, \pi)$ due to cleaning. Specifically, the two columns summarize the best-fitting HODs obtained using mocks constructed from the default (uncleaned) COMPASO halo catalogues (left-hand column) and from the ‘cleaned’ version of this catalogue assuming $\kappa = 2.0$ (right-hand column). The bottom two rows compare the the log-likelihood, $\log \mathcal{L}$, and Bayesian evidence, $\log \mathcal{Z}$, resulting from each of the cleaned and uncleaned catalogues. The Bayesian evidence, which describes the consistency between the model and the data, for the uncleaned mock catalogue is $\log \mathcal{Z} = -76$, compared to $\log \mathcal{Z} = -62$ for the cleaned version. In other words, the BOSS redshift-space clustering exhibits a roughly 14 e -fold preference for the cleaned mocks over the uncleaned ones. Similarly, we see a 14 e -fold improvement to the best-fitting likelihood.

The full marginalized posterior distributions of our seven-parameter HOD mocks fit to $\xi(r_p, \pi)$ are shown in Fig. 12, in which the parameter constraints for the uncleaned catalogue are shown in black, while for the cleaned catalogue these are in green. The maximum likelihood values of the HOD parameters are listed in the first seven rows of Table 2. Unsurprisingly, the switch from uncleaned to cleaned COMPASO haloes results in a shift in the best-fitting HOD parameters, particularly M_1 , α , and α_s . The decrease in M_1 and the increase in α suggest a larger satellite population in the cleaned mock, consistent with the fact that the halo cleaning re-merges smaller haloes on the fringe of larger haloes, converting centrals in these fringe haloes to satellites in the re-merged halo. The decrease in the satellite velocity bias parameter, α_s , pulls it back into agreement with no satellite velocity bias ($\alpha_s = 1$). This is consistent with prior velocity bias fits by e.g. Guo et al. (2015). Judging from the marginalized posterior contours, it is also worth noting that the cleaning also generally yields better-constrained HOD parameter posteriors and the removal of several parameter degeneracies. This suggests that the cleaned haloes are indeed more consistent with the data, and thus results in a preferred redshift-space two-point correlation function, $\xi(r_p, \pi)$ obtained using mocks constructed from the uncleaned halo catalogues (left-hand panel) and from the cleaned catalogue (right-hand panel). In particular, there is a noticeable reduction in the residual on scales $r_p \sim 3\text{--}10 \text{ Mpc } h^{-1}$, where the uncleaned COMPASO mocks predict an excess in

clustering relative to the data. The removal of low-mass, high velocity haloes at the boundaries of groups and clusters (see Fig. 11) in the cleaned version results in an improved fit. The shift from uncleaned to cleaned haloes also results in an appreciable improvement in the residuals on sub-Mpc projected separations. While it is difficult to conduct ‘ χ by eye’, the right-hand panel sums up to significantly lower χ^2 compared to the left-hand panel (see $\log \mathcal{L}$ comparison in Table 2).

5 CONCLUSIONS

Tracking the formation and growth of dark matter haloes is a critical step in the analysis of cosmological simulations of structure formation. In this paper, we have introduced an algorithm for constructing high-fidelity halo merger trees from ABACUSSUMMIT (Maksimova et al. 2021), an extremely large suite of cosmological simulations that has been designed to meet the Cosmological Simulation Requirements of the DESI survey. The flagship ABACUSSUMMIT data set consists of 139 simulations of size $2 \text{ Gpc } h^{-1}$, each resolved with 6912^3 particles, yielding an effective particle mass of $2 \times 10^9 M_\odot h^{-1}$, making it an ideal data set for theoretical applications in the mass scale relevant to emission-line galaxies (ELGs), a primary target for DESI. Combined with the several smaller and larger boxes that we have also simulated as part of this project, ABACUSSUMMIT totals nearly 60 trillion particles, with simulations spanning a wide range of cosmological parameter variations, with the majority also including the effect of massive neutrinos on structure growth.

To identify haloes in ABACUSSUMMIT, we use the *Competitive Assignment to Spherical Overdensities*, or COMPASO, algorithm (Hadzhiyska et al. 2022), which is a highly efficient, on-the-fly halo finder that has been purpose-built for ABACUSSUMMIT. COMPASO identifies haloes from a given particle set by first applying a density-bounded friends-of-friends linking algorithm to identify groups, followed by a hierarchical sequence of particle assignment steps that competitively link together particles that satisfy certain criteria based on pre-defined density thresholds. A brief outline of how COMPASO operates is provided in Section 2.2.

Like several other merger tree algorithms that have been published in the literature (e.g. Behroozi et al. 2013b; Jiang et al. 2014; Poole et al. 2017; Han et al. 2018; Elahi et al. 2019; Rangel et al. 2020), in our method, which is described in detail in Section 2.4, the ability to build associations between sets of haloes relies on the ability to accurately track the particle content of these entities across multiple snapshots. In performing these matches, particular preference is given to particles with the highest kernel densities (as computed by COMPASO). This tends to trace material in the cores of haloes, which has been shown to be more reliable choice for the construction of merger trees (e.g. Srisawat et al. 2013). A potential halo association occurs when some fraction of these particles overlaps between any two haloes across two adjacent output times. Once these associations are established, we define halo progenitors as only those objects at some snapshot $i-1$ that contribute a sizeable fraction of matched particles (i.e. exceeding a threshold fraction) to a recipient halo at snapshot i . A visualization of our algorithm is shown in Fig. 1. In Section 3, we perform a series of tests to show that our algorithm is robust to the density of snapshots used to construct halo merger trees, particularly when measuring the kinds of statistics that are commonly used in the construction of mock catalogues (Figs 2–4).

The most direct application of a merger tree is to compute the mass accretion histories of dark matter haloes. Given the mass resolution and volume of the ABACUSSUMMIT simulations presented in this

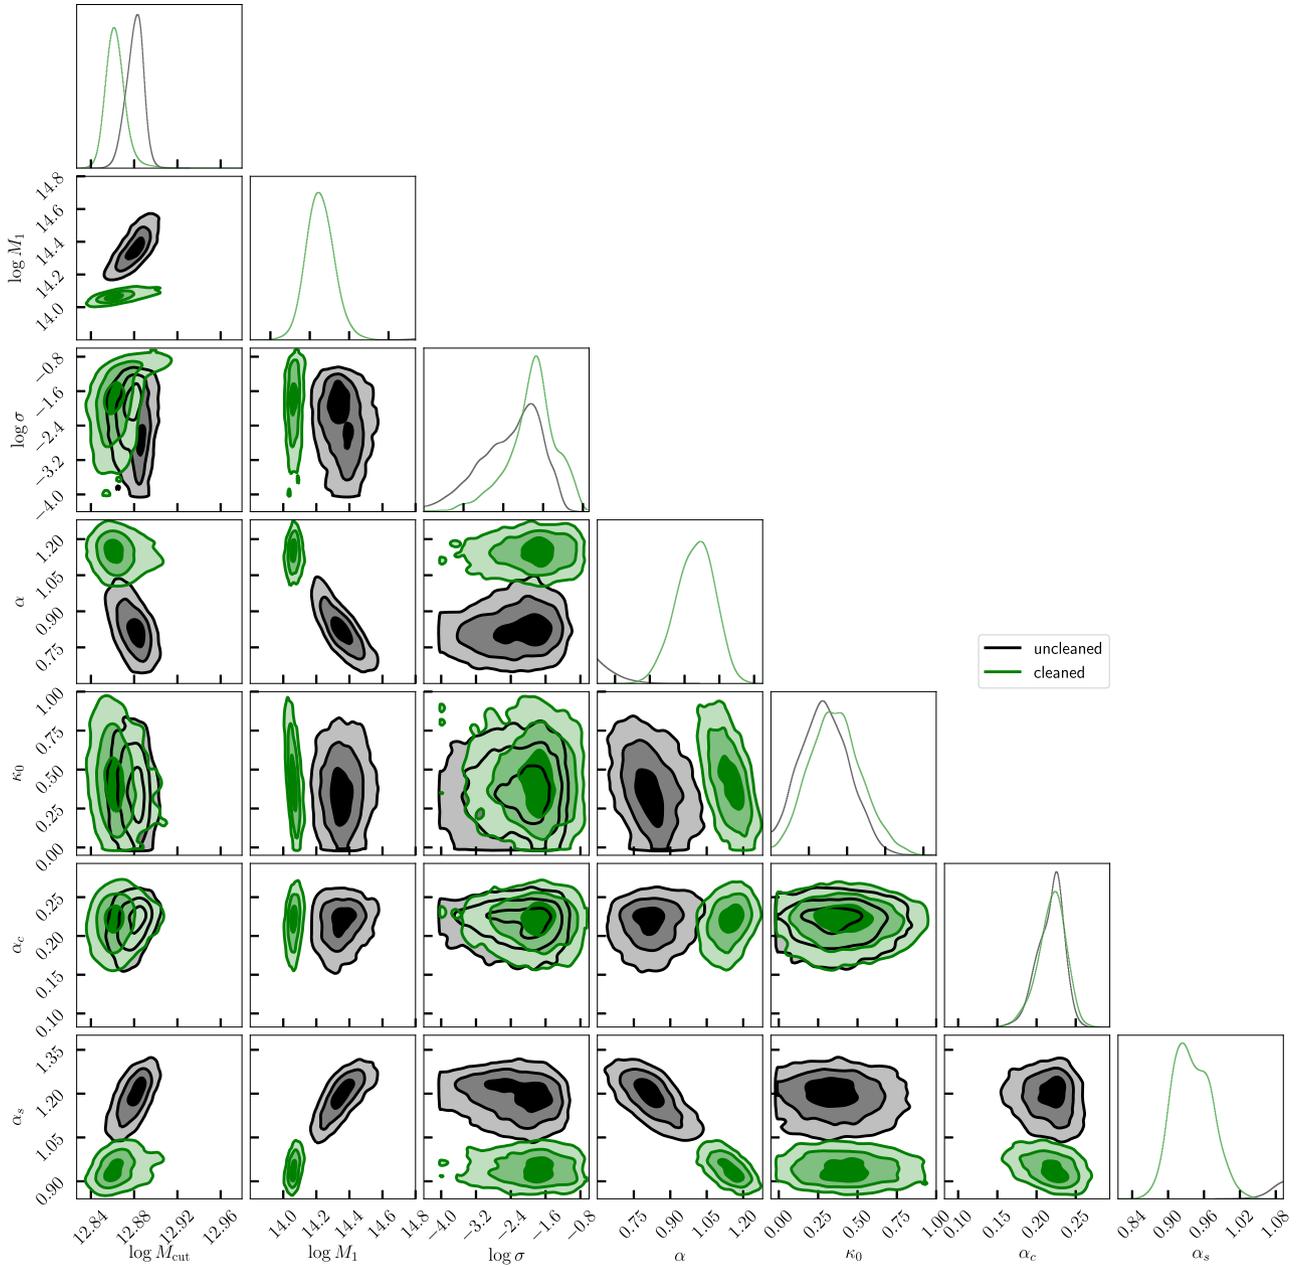


Figure 12. The 1D and 2D marginalized posterior constraints on our seven-parameter HOD fits to the BOSS CMASS DR12 redshift-space two-point correlation function, $\xi(r_p, \pi)$. The distributions for the uncleaned and cleaned halo catalogues, respectively, are represented by the black and green contours. The three levels correspond to 1σ , 2σ , and 3σ level constraints. The posterior medians on each of these parameters are listed in the first seven rows of Table 2.

paper, we are able to accurately track the assembly history of objects ranging from the ELG-mass hosts to rich clusters of galaxies (Fig. 5). We also find good convergence between pairs of simulations that have been run using the same initial phases, but that vary in resolution by a factor of 6.

We then describe a procedure that uses merger trees to ‘clean’ the default COMPASO halo catalogues of objects whose masses are deemed to be unreliable due to physical and/or numerical effects during their past evolution. In particular, haloes identified at the present time may have undergone fly-bys, been ‘split’ off larger haloes by the COMPASO algorithm, or simply passed through the interiors of more massive objects. Each of these processes can hamper the accuracy with which the particle content (and, therefore, the present-

day properties) of the halo are tracked, which can consequently bias models of the galaxy–halo connection that are applied to our default halo outputs. In ABACUS, we take on a conservative approach and identify all such ‘problematic’ objects and eliminate these as independent haloes using the procedure described in Section 4.2. In short, this method flags those objects whose maximum mass (measured along the main progenitor branch) exceeds its present-day mass by some factor, κ . After identifying the more massive, neighbouring halo they were once part of, the objects are ‘merged’ on to this host and the newly aggregated system is treated as a single entity at all subsequent output times (see Figs 6 and 7).

The net effect of this cleaning method is to predominantly flag low-mass haloes that have especially truncated merger trees (Fig. 9),

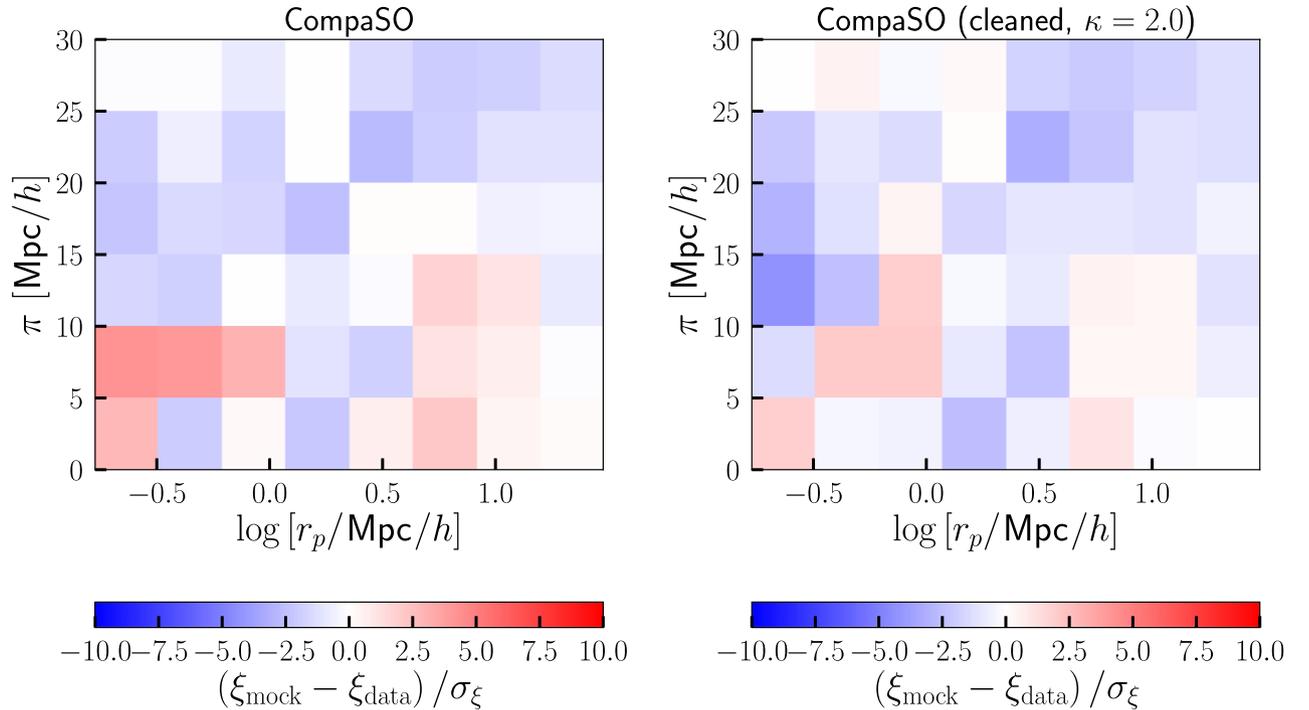


Figure 13. The best-fitting two-point correlation function fits to $\xi(r_p, \pi)$ using the uncleaned (left-hand panel) and cleaned (right-hand panel) versions of the `Base_c000_ph000` simulation. The colour maps show the *residual* on the best-fitting mocks with respect to the data. The errors, σ_ξ , are computed using the diagonal elements of the inverse covariance matrix (see the main text). The mocks constructed using the cleaned COMPASO catalogues yield a better fit to the clustering, particularly in the lower left and lower right quadrants.

and those that are found preferentially near the boundaries of cluster-mass haloes. Their removal results in a suppression in the two-point cross-correlation on scales of $1\text{--}3 \text{ Mpc } h^{-1}$ (Fig. 10), and we find that the response of this effect is largely insensitive to small changes in κ . On examining the detailed dynamics of the ‘cleaned’ haloes, we find that our method preferentially removes a large population of objects with net positive radial velocities directed away from the centres of clusters (Fig. 11).

As a final test, we apply our cleaned catalogues in a real use-case application of building HOD mocks to fit the BOSS CMASS DR12 redshift-space two-point correlation function, $\xi(r_p, \pi)$. To this end, we use the ABACUSHOD code (Yuan et al. 2022) that allows for several generalized extensions to the baseline halo-mass-only HOD model. We find that shifting from uncleaned to cleaned catalogues results in an enhanced satellite population in the latter, consistent with the idea low-mass haloes that was originally identified as individual entities are now merged with larger neighbours and become part of their satellite population (Fig. 12). We also find that the resulting fit to $\xi(r_p, \pi)$ is much improved in the cleaned version than in the uncleaned version, with Bayesian evidence and log-likelihood that suggest a 14 e -fold preference for the former over the latter (Table 2 and Fig. 13).

While in this paper we have presented results from only three sets of ABACUSSUMMIT simulations, our merger tree construction has been applied across the entire suite. This comprises an extensive repertoire of information pertaining to the redshift evolution of halo properties spanning a generous range in cosmology variations. This makes the ABACUSSUMMIT suite a potentially invaluable resource for constructing and testing models of the galaxy–halo connection, covariance estimation, and emu-

lator building in anticipation of next-generation galaxy redshift surveys.

ACKNOWLEDGEMENTS

We thank the referee for providing a number of useful comments and suggestions that have strengthened quality of this work. We are grateful to Alex Smith for testing an early version of the halo catalogue cleaning method, and for identifying areas of improvement. This work has been supported by NSF AST-1313285, DOE-SC0013718, and NASA ROSES grant 12-EUCLID12-0004. SB is supported by the UK Research and Innovation (UKRI) Future Leaders Fellowship [grant number MR/V0233]. DJE is supported in part as a Simons Foundation investigator. NAM was supported in part as an NSF Graduate Research Fellow. LHG is supported by the Center for Computational Astrophysics at the Flatiron Institute, which is supported by the Simons Foundation.

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. Computation of the merger trees used resources of the National Energy Research Scientific Computing Center (NERSC), a US Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. The ABACUSSUMMIT simulations have been supported by OLCF projects AST135 and AST145, the latter through the Department of Energy ALCC program.

We would like to thank the OLCF and NERSC staff for their highly responsive and expert assistance, both scientific and administrative, during the course of this project.

DATA AVAILABILITY

The ABACUSSUMMIT simulations have been placed into the public domain, subject to the academic citations described at <https://abacussummit.readthedocs.io/en/latest/citation.html>. A user interface for reading and analysing ABACUSSUMMIT data, as well as the ABACUSHOD module, are available as part of the ABACUSUTILS package (<https://abacusutils.readthedocs.io/en/latest/>).

Data access is available through OLCF's Constellation portal. The initial data release includes cleaned halo catalogues for six ABACUSSUMMIT simulations. The persistent DOI describing these data is 10.13139/OLCF/1811689. Cleaned halo catalogues and particle subsamples for all other ABACUSSUMMIT boxes are available through the DOI 10.13139/OLCF/1828535. Instructions for accessing the data are given at <https://abacussummit.readthedocs.io/en/latest/data-access.html>.

REFERENCES

- Baugh C. M. et al., 2019, *MNRAS*, 483, 4922
- Behroozi P. S., Wechsler R. H., Wu H.-Y., 2013a, *ApJ*, 762, 109
- Behroozi P. S., Wechsler R. H., Wu H.-Y., Busha M. T., Klypin A. A., Primack J. R., 2013b, *ApJ*, 763, 18
- Benson A. J., 2012, *New A*, 17, 175
- Benson A. J., Cole S., Frenk C. S., Baugh C. M., Lacey C. G., 2000, *MNRAS*, 311, 793
- Berlind A. A., Weinberg D. H., 2002, *ApJ*, 575, 587
- Bolton A. S. et al., 2012, *AJ*, 144, 144
- Bond J. R., Cole S., Efstathiou G., Kaiser N., 1991, *ApJ*, 379, 440
- Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, *MNRAS*, 319, 168
- Correa C. A., Wyithe J. S. B., Schaye J., Duffy A. R., 2015, *MNRAS*, 450, 1521
- Croton D. J. et al., 2006, *MNRAS*, 365, 11
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- Dawson K. S. et al., 2013, *AJ*, 145, 10
- Efstathiou G., Frenk C. S., White S. D. M., Davis M., 1988, *MNRAS*, 235, 715
- Eisenstein D. J. et al., 2011, *AJ*, 142, 72
- Elahi P. J., Welker C., Power C., Lagos C. d. P., Robotham A. S. G., Cañas R., Poulton R., 2018, *MNRAS*, 475, 5338
- Elahi P. J., Poulton R. J. J., Tobar R. J., Cañas R., Lagos C. d. P., Power C., Robotham A. S. G., 2019, *Publ. Astron. Soc. Aust.*, 36, e028
- Falck B. L., Neyrinck M. C., Szalay A. S., 2012, *ApJ*, 754, 126
- Garrison L. H., Eisenstein D. J., Ferrer D., Tinker J. L., Pinto P. A., Weinberg D. H., 2018, *ApJS*, 236, 43
- Garrison L. H., Eisenstein D. J., Pinto P. A., 2019, *MNRAS*, 485, 3370
- Garrison L. H., Eisenstein D. J., Ferrer D., Maksimova N. A., Pinto P. A., 2021, *MNRAS*, 508, 575
- Gonzalez-Perez V. et al., 2018, *MNRAS*, 474, 4024
- Guo H. et al., 2015, *MNRAS*, 446, 578
- Hadzhiyska B., Tacchella S., Bose S., Eisenstein D. J., 2021, *MNRAS*, 502, 3599
- Hadzhiyska B., Eisenstein D., Bose S., Garrison L. H., Maksimova N., 2022, *MNRAS*, 509, 501
- Han J., Cole S., Frenk C. S., Benitez-Llambay A., Helly J., 2018, *MNRAS*, 474, 604
- Hearin A. P., Watson D. F., 2013, *MNRAS*, 435, 1313
- Heitmann K. et al., 2015, *ApJS*, 219, 34
- Henriques B. M. B., White S. D. M., Thomas P. A., Angulo R., Guo Q., Lemson G., Springel V., Overzier R., 2015, *MNRAS*, 451, 2663
- Ishiyama T. et al., 2020, *MNRAS*, 506, 4210
- Ivezić Ž. et al., 2019, *ApJ*, 873, 111
- Jiang L., Helly J. C., Cole S., Frenk C. S., 2014, *MNRAS*, 440, 2115
- Kauffmann G., White S. D. M., Guiderdoni B., 1993, *MNRAS*, 264, 201
- Klimontowski J., Lokas E. L., Knebe A., Gottlöber S., Martínez-Vaquero L. A., Yepes G., Hoffman Y., 2010, *MNRAS*, 402, 1899
- Klypin A., Holtzman J., 1997, preprint (astro-ph/9712217)
- Klypin A. A., Trujillo-Gomez S., Primack J., 2011, *ApJ*, 740, 102
- Knebe A., Libeskind N. I., Knollmann S. R., Yepes G., Gottlöber S., Hoffman Y., 2010, *MNRAS*, 405, 1119
- Knebe A. et al., 2011, *MNRAS*, 415, 2293
- Knollmann S. R., Knebe A., 2009, *ApJS*, 182, 608
- Lacey C., Cole S., 1993, *MNRAS*, 262, 627
- Lacey C., Cole S., 1994, *MNRAS*, 271, 676
- Lagos C. d. P., Tobar R. J., Robotham A. S. G., Obreschkow D., Mitchell P. D., Power C., Elahi P. J., 2018, *MNRAS*, 481, 3573
- Lam S. K., Pitrou A., Seibert S., 2015, Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM '15, Association for Computing Machinery, New York, NY
- Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
- Laureijs R. et al., 2011, preprint (arXiv:1110.3193)
- Levi M. et al., 2013, preprint (arXiv:1308.0847)
- Maksimova N. A., Garrison L. H., Eisenstein D. J., Hadzhiyska B., Bose S., Satterthwaite T. P., 2021, *MNRAS*, 508, 4017
- Neyrinck M. C., Gnedin N. Y., Hamilton A. J. S., 2005, *MNRAS*, 356, 1222
- Onions J. et al., 2012, *MNRAS*, 423, 1200
- Parkinson H., Cole S., Helly J., 2008, *MNRAS*, 383, 557
- Peacock J. A., Smith R. E., 2000, *MNRAS*, 318, 1144
- Peebles P. J. E., 1980, The Large-Scale Structure of the Universe.
- Planck Collaboration et al., 2020, *A&A*, 641, A6
- Poole G. B., Mutch S. J., Croton D. J., Wyithe S., 2017, *MNRAS*, 472, 3659
- Potter D., Stadel J., Teyssier R., 2017, *Comput. Astrophys. Cosmol.*, 4, 2
- Poulton R. J. J., Robotham A. S. G., Power C., Elahi P. J., 2018, *PASA*, 35, e042
- Press W. H., Schechter P., 1974, *ApJ*, 187, 425
- Rangel E., Frontiere N., Habib S., Heitmann K., Liao W.-k., Agrawal A., Choudhary A., 2020, preprint (arXiv:2008.08519)
- Roper W. J., Thomas P. A., Srisawat C., 2020, *MNRAS*, 494, 4509
- Sinha M., Garrison L. H., 2020, *MNRAS*, 491, 3022
- Somerville R. S., Kolatt T. S., 1999, *MNRAS*, 305, 1
- Somerville R. S., Hopkins P. F., Cox T. J., Robertson B. E., Hernquist L., 2008, *MNRAS*, 391, 481
- Speagle J. S., 2020, *MNRAS*, 493, 3132
- Spergel D. et al., 2015, preprint (arXiv:1503.03757)
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
- Springel V. et al., 2005, *Nature*, 435, 629
- Srisawat C. et al., 2013, *MNRAS*, 436, 150
- Tweed D., Devriendt J., Blaizot J., Colombi S., Slyz A., 2009, *A&A*, 506, 647
- Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
- Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020, *Nat. Rev. Phys.*, 2, 42
- Wang Y. et al., 2016, *MNRAS*, 459, 1554
- Wechsler R. H., Tinker J. L., 2018, *ARA&A*, 56, 435
- Yuan S., Eisenstein D. J., Garrison L. H., 2018, *MNRAS*, 478, 2019
- Yuan S., Hadzhiyska B., Bose S., Eisenstein D. J., Guo H., 2021, *MNRAS*, 502, 3582
- Yuan S., Garrison L. H., Hadzhiyska B., Bose S., Eisenstein D. J., 2022, *MNRAS*, 510, 3301
- Zheng Z., Weinberg D. H., 2007, *ApJ*, 659, 1
- Zheng Z. et al., 2005, *ApJ*, 633, 791

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