



STRAWBERRY: Finding haloes in the gravitational potential

Tamara R. G. Richardson^{1,*} , Jens Stücker², and Raul E. Angulo^{1,3} 

¹ Donostia International Physics Center (DIPC), Paseo Manuel de Lardizabal, 4, 20018 Donostia-San Sebastián, Gipuzkoa, Spain

² Institute for Astronomy, University of Vienna, Türkenschanzstraße 17, Vienna 1180, Austria

³ IKERBASQUE, Basque Foundation for Science, E-48013 Bilbao, Spain

Received 13 August 2025 / Accepted 17 January 2026

ABSTRACT

We present a novel algorithm that discriminates between bound and unbound particles by consideration of the gravitational potential from an accelerated reference frame, also referred to as the ‘boosted potential’. Particles are considered to be bound if their energy does not exceed the escape energy of a potential well, given by the closest saddle point that connects to a deeper potential minimum. This approach has core benefits over previous approaches, since it does not require any ad hoc thresholds (such as overdensity criteria), while it does include the gravitational effect of all particles in the binding criterion (improving over widely used self-potential binding checks) and it only operates with instantaneous information (making it simpler than approaches based on dynamical histories). We show that particles typically become bound between their first peri- and apo-centric passage and that bound and unbound populations show very distinct characteristics through their distribution in phase space, as well as in their density profiles, virial ratios, and redshift evolution. Our findings suggest that it is possible to understand haloes as two-component systems, with one component being bound, virialised, of a finite extent, and evolving slowly in quasi-equilibrium and the other component being unbound, un-virialised, and evolving rapidly.

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

1. Introduction

In the standard Λ cold dark matter (Λ CDM) cosmological model, dark matter haloes are one of the building blocks of non-linear structure. As such, a great deal of research is founded upon our understanding of these structures. A non-exhaustive list of examples includes the halo model of the non-linear power spectrum and its extensions (e.g. [Seljak 2000](#); [Mead et al. 2020](#); [Asgari et al. 2023](#); [Ayçoberry et al. 2025](#); [Salazar et al. 2025](#)), which rely on knowing how abundant haloes are and how matter is distributed within them; galaxy clustering studies (e.g. [Alam et al. 2021](#); [Abbott et al. 2018, 2022](#); [Adame et al. 2025](#); [Euclid Collaboration: Mellier et al. 2025](#)), which are heavily reliant on the definition and properties of haloes in order to model the clustering properties of galaxies (e.g. [Jing et al. 1998](#); [Benson et al. 2000](#); [Peacock & Smith 2000](#); [Vale & Ostriker 2006](#); [Conroy et al. 2006](#); [Contreras et al. 2021](#); [Ortega-Martinez et al. 2024](#)); galaxy cluster number count analyses, which are heavily dependent on the calibration of scaling relations between observable properties and halo properties (e.g. [Planck Collaboration XXVII 2016](#); [Heymans et al. 2021](#); [Lesci et al. 2022](#); [Bleem et al. 2024](#); [Bulbul et al. 2024](#); [Serenio et al. 2025](#)); and dedicated ‘zoom-in’ simulation runs of galaxy and structure formation (e.g. [Hahn et al. 2017](#); [Cui et al. 2018, 2022](#); [Wang et al. 2020](#); [Nadler et al. 2023](#); [Pellissier et al. 2023](#); [Nelson et al. 2024](#)), which require a robust understanding of how haloes form to accurately select the Lagrangian patch that is re-simulated. Therefore, it is evident that it is important to have a clear definition of haloes and their boundaries.

The simplest, and arguably the most widely used boundary defines haloes as spherical regions, which are denser than the cosmological background. This results in a set of characteristic mass scales known as spherical overdensity masses, typically denoted $M_{\Delta c}$ or $M_{\Delta b}$, which represent the total mass contained within a spherical region enclosing an average density that is Δ times the critical or background matter density of the Universe. This definition is motivated by the theoretical description of virialised haloes within the spherical collapse framework. This definition has many benefits that have made it popular. First, it is intuitive and easy to implement. Second, it is flexible, with the possibility of varying the value of Δ depending on the application. For example, studies that focus on the innermost regions of haloes, such as the study of hot gas in the intracluster medium (e.g. [Planck Collaboration XXVII 2016](#); [Hilton et al. 2021](#); [Bleem et al. 2024](#); [Bulbul et al. 2024](#); [Serenio et al. 2025](#)), will tend to use mass definitions corresponding to these regions (e.g. M_{500c} , M_{2500c}), while studies interested in larger radii, such as weak lensing mass estimates (e.g. [Serenio 2015](#); [Bellagamba et al. 2019](#); [Umetsu et al. 2020](#); [Lesci et al. 2022](#)), will tend to use masses defined at lower overdensities (e.g. M_{200c} or M_{200b}).

Nonetheless, this halo definition has several drawbacks. Notably, it enforces spherical symmetry, whereas haloes are typically ellipsoidal and present many small-scale features such as sub-haloes or caustics (e.g. [Eisenstein & Loeb 1995](#); [Springel et al. 2004](#); [Vera-Ciro et al. 2011](#); [Despali et al. 2014](#); [Bonnet et al. 2022](#)). Moreover, spherical overdensities are not well suited to studying the evolution of haloes, due to the fact that the definition itself changes depending on the cosmological background, through an effect known as pseudo-evolution ([Diemer et al. 2013](#)). For example, this effect makes it more

* Corresponding author: tamara.richardson@dipc.org

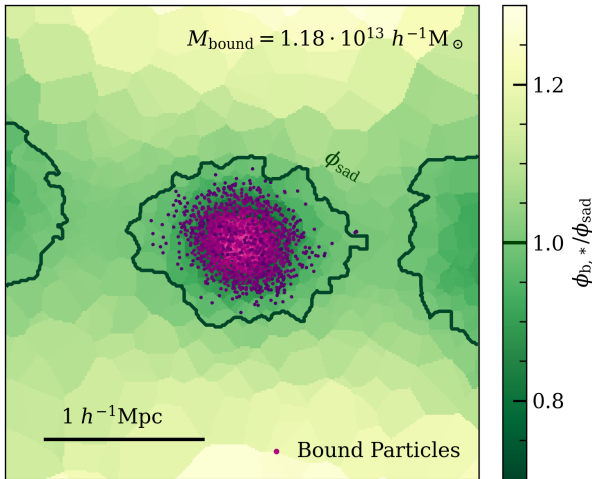


Fig. 1. Slice of the boosted gravitational potential field centred on a halo. Dark green contours mark the saddle point energy, ϕ_{sad} , while the bound population of particles is shown as magenta points.

difficult to disentangle how much matter has been accreted over a given time period, as the spherical overdensity mass will increase over time even in the absence of accretion, simply because the background density is decreasing. This definition is also ill-suited to describe the abundance of haloes, as expressed through the halo mass function (HMF). Indeed, excursion-set theory predicts that the HMF should have no explicit cosmology dependence when expressed in the correct units (Press & Schechter 1974; Bond et al. 1991); in other words, the HMF should be universal. In practice, simulations have revealed this to not be the case (e.g. Tinker et al. 2008; Despali et al. 2016; Castro et al. 2021; Ondaro-Mallea et al. 2022). However, an open question remains as to the origin of the observed non-universality: whether it is generated through a physical process or simply emerges from the mass definition.

In N -body simulations, another popular choice for the definition is the mass of friends-of-friends (FoF) groups, (e.g. Press & Davis 1982; Springel et al. 2001; Roy et al. 2014), which depend on a linking length between individual particles that serves as a proxy for the local matter density. This has the advantage of removing the assumption of spherical symmetry and it is numerically efficient to implement. Moreover, this definition provides valuable insight as to the clumpy nature of haloes but comes at the expense of being more difficult to link to theory and not being directly applicable to observations. In addition, FoF has the tendency to detect chance alignment of particles, marking the presence of a halo where there is none. For example, in Angulo et al. (2013) and Stücker et al. (2020, 2021), caustics and numerical fragments were detected by the FoF approach as haloes in warm dark matter simulations. Furthermore, because FoF relies on the local density of particles, it may create unphysically large structures by bridging haloes (see e.g. Knebe et al. 2011; Leroy et al. 2021). This particular issue stems from the fact that haloes lack a sharp boundary in terms of density.

Nevertheless, in simulations, the physical properties of dark matter, such as the density, velocity, and acceleration fields, are known, allowing us to construct a more in-depth picture of haloes by using more advanced definitions. While the edges of haloes are difficult to define from the density field alone (see e.g. Eisenstein & Hut 1998; Springel et al. 2001; Aubert et al. 2004; Codis et al. 2015, 2018), it is possible to do so in terms

of dynamical properties. For instance, particles can be selected based on their orbital status, considering only particles that have undergone their first apocentric passage (Adhikari et al. 2014; Diemer 2017) or that have undergone shell crossing in three dimensions (Falck et al. 2012; Stücker et al. 2020) as part of the halo. Alternatively, the boundary can be defined in energy space by introducing cuts in terms of kinetic energy (García et al. 2023; Salazar et al. 2025). These approaches are, however, much more computationally expensive, with the first methods noted above requiring the motion of every particle to be tracked, while the second requires an internal fitting procedure to optimise the kinetic energy cut.

One quantity that is often overlooked when studying haloes is the gravitational potential, ϕ . As the scalar quantity sourcing the motion, the gravitational potential is the natural quantity defining the area of influence of a given halo. Until recently, using the gravitational potential has been impossible in practice since it is dominated by large-scale gradients that mask the potential wells of smaller structures. However, Stücker et al. (2021) showed that it is possible to use the gravitational potential to define structures by adding a ‘boost’,

$$\phi_b(\mathbf{x}) = \phi(\mathbf{x}) - (\mathbf{x} - \mathbf{x}_0) \cdot \nabla_{\mathbf{x}}\phi(\mathbf{x})|_{\mathbf{x}=\mathbf{x}_0} = \phi(\mathbf{x}) + (\mathbf{x} - \mathbf{x}_0) \cdot \mathbf{a}_0, \quad (1)$$

which removes the locally-averaged large-scale gradient, $\nabla_{\mathbf{x}}\phi(\mathbf{x})|_{\mathbf{x}=\mathbf{x}_0}$. This operation is equivalent to studying a halo in a reference frame free-falling with an acceleration of \mathbf{a}_0 , with the equivalence principle implying that this type of transformation does not impact the internal dynamics of the system. In this context, haloes can be defined as host potential wells, for which the boundary is set by a saddle point in the field which leads to a deeper potential well. Particles within the well that have energies below the saddle-point energy are bound to the structure, whereas particles with higher energies could potentially escape. This definition of boundness can be understood as a generalisation of commonly used self-binding checks, (e.g. Springel et al. 2001; Behroozi et al. 2013) since it naturally includes the gravitational effect of surrounding structures. This is illustrated in Fig. 1, where we show a slice of the boosted gravitational potential field centred on a $10^{13} h^{-1} M_{\odot}$ halo. Here, we outline the saddle point energy, ϕ_{sad} , marking the dynamical edge of the halo with a thick green line and showing the particles that are bound to this potential well in magenta.

In this work, we propose expanding the study of the properties of haloes as defined using the boosted potential in Stücker et al. (2021). As such, we have developed a particle assignment algorithm, known as STRAWBERRY¹, to efficiently find the saddle point energy and check whether particles are bound to a given potential well. This work is structured as follows. After discussing the theoretical framework in Sect. 2, we present our particle assignment algorithm in Sect. 3 for which several convergence tests can be found in Appendix A. In Sect. 4, we study the properties of particle distributions assigned to haloes, examining how particle energies evolve and how they are then distributed in the halo. In addition, we study their virialisation state, and how their bound masses compare to the more commonly used M_{200b} . Finally, in Sect. 5, we summarise this work and present our conclusions and future prospects.

2. Theory

The goal of the STRAWBERRY algorithm is to find bound structures in the gravitational potential. However, notions of

¹ STRucture Assignment Within BoostED Reference fRames in cYthon.

boundedness and the potential energy depend on the choice of a coordinate frame. Here, we discuss our notion of boundedness, along with notions of how it operates in the idealised case of a static potential landscape, how the expansion of the universe can be accounted for, and how to choose an optimal frame of reference.

2.1. Notion of boundedness

It is difficult to give a clear definition of what ‘boundedness’ signifies in the case of cosmological simulations. For isolated systems, binding notions tend to reference the gravitational energy normalised to $\phi(\rightarrow \infty) = 0$ and performing checks on whether a given particle’s energy is sufficient to escape to infinity. However, this assumption shows weakness when the additional potential of other structures is considered. For example, considering the gravitational potential of the Earth, in principle, it would be possible to bind particles at arbitrary large distances if they were to have low enough relative velocities. However, due to the gravitational field of the Sun, any particles at a distance greater than the Lagrangian point L1 will orbit around the Sun, rather than the Earth. While the self-potential binding check is a powerful heuristic in many situations, we probably do not want to use it as a baseline of what it means to be ‘bound’.

Intuitively, what we understand by particles being ‘bound’ to a structure is that they are restricted to orbiting around it. Approximately, we might say that a particle, i , is bound to a given set of particles at a time, t , if for all future times it will not leave a certain area of influence. However, his notion of boundedness is not very precise, as it has several ambiguities: (1) the coordinate frame that we have used to define distances needs to be specified. For example, in cosmology, under this definition, a structure may be bound in comoving space, but unbound in physical space; (2) it is not clear what reference sets of particles and area of influence should be considered; and (3) it is also a somewhat impractical notion overall, as it requires predicting the full future of the particle. For instance, if a particle orbits in the vicinity of a structure for a long time, but then gets expelled (e.g. through some event that strongly changes its energy), it is not obvious whether or not it should be called ‘bound’ prior to the ejection event.

Despite these ambiguities, we argue that the goal of any physical definition of boundedness is to approximate statements of this form. In other words, heuristics based on energy arguments that could be used to make approximate predictions about the future of particles. For a given heuristic, a particle then becomes bound when it becomes clear that it is restricted in this sense for all future times. Our goal is to define a heuristic based on the ‘full’ gravitational potential that allows us to detect bound structures in this regard.

Before we go on to discuss the full cosmological scenario, it is insightful to consider a simplified thought experiment. Let us assume that we have a Hamiltonian of the form,

$$H = \phi(\mathbf{x}) + \frac{1}{2}\mathbf{v}^2, \quad (2)$$

with a complicated (non-monotonic), but static potential, ϕ . As energy is conserved in this scenario, it is easy to find bound regions: if a particle has an energy level $E = H$, then it may only orbit in the space restricted by $\phi(\mathbf{x}) < E$. Furthermore, it is restricted to the connected space (that contains its starting position) that fulfils this criterion. We can associate a connected region $\phi(\mathbf{x}) < E$ with the deepest minimum, ϕ_{\min} , inside of it. The maximal extent of space that may be associated with




Fig. 2. Illustration of our binding notion for a fixed potential landscape. The region where particles may be bound to the minimum extends up to the first saddle point connecting to a deeper minimum.

a given minimum is then limited by the first saddle point, ϕ_{sad} , which connects to a deeper minimum. The difference in energy $\Delta\phi = \phi_{\text{sad}} - \phi_{\text{min}}$ is called the ‘persistence’. Therefore, in the described scenario, we may say that a particle is bound to the deepest minimum in the connected space $\phi(\mathbf{x}) \leq E$. This is illustrated in Figure 2.

While in the former scenario, we can provide a very clear notion of boundedness, the cosmological scenario is considerably more complicated. This is because the potential notion depends on the adopted reference frame and is generally highly time-dependent. However, it is possible to recover a scenario that is close to the one described above by transforming it into an appropriate reference frame that accounts for the expansion of the universe and for large-scale accelerations.

2.2. Potential energy

In physical space, it is common to define the potential through the Poisson equation,

$$\nabla_r^2 \phi_r = 4\pi G \rho(\mathbf{r}), \quad (3)$$

where ∇_r denotes a derivative with respect to physical coordinates \mathbf{r} and $\rho(\mathbf{r})$ is the physical density. The physical energy is given by

$$E = \phi_r(\mathbf{r}) + \frac{1}{2}v^2 - \frac{1}{6}\Lambda c^2 r^2, \quad (4)$$

$$= \phi_r(\mathbf{r}) + \frac{1}{2}v^2 - \frac{1}{2}H^2 \Omega_\Lambda(a)r^2, \quad (5)$$

where Λ is the cosmological constant. The dark energy term is often omitted, as it is usually sufficiently small at the small distances that are of interest in physical space. However, since it has an effect on larger scales, it is included for consistency.

In the comoving frame, we can write

$$\mathbf{x} = \frac{\mathbf{r}}{a}, \quad (6)$$

$$\mathbf{v} = \dot{\mathbf{x}}a^2 = a(\dot{\mathbf{r}} - \mathbf{r}H), \quad (7)$$

$$\nabla_x^2 \phi = 4\pi G a^2 [\rho(\mathbf{r}) - \rho_m] = \frac{3}{2}a^{-1} \Omega_{m,0} H_0^2 \delta, \quad (8)$$

$$\dot{\mathbf{v}} = -\nabla_x \phi, \quad (9)$$

where x is the comoving coordinate, a is the scale factor, \mathbf{v} is the canonical momentum which relates to peculiar velocities as $\mathbf{v}_{\text{pec}} = \mathbf{v}/a$, ρ_m is the mean matter density of the universe, H_0 the Hubble parameter at $a = 1$, δ the relative matter overdensity, ϕ is the peculiar potential, and the dot marks derivatives with respect

to the proper time for a comoving observer. We note that conventions on the definition of the peculiar potential may differ by a factor, a , for example the cosmological simulation code *Gadget* (Springel et al. 2021) outputs $\phi_{\text{gadget}} = a \cdot \phi$ in snapshots. Importantly, the physical potential and the peculiar potential differ by a quadratic term,

$$\phi = \phi_r - \frac{1}{4}H^2 r^2 \Omega_m, \quad (10)$$

which accounts for the effect of the mean density. We can express the physical energy from the comoving frame,

$$\begin{aligned} E &= \frac{1}{4}H^2 r^2 [\Omega_m(a) - 2\Omega_\Lambda(a)] + \phi + \frac{\dot{r}^2}{2}, \\ &= \frac{1}{4}H^2 a^2 x^2 [2 + \Omega_m(a) - 2\Omega_\Lambda(a)] + H\mathbf{v} \cdot \mathbf{x} + \phi + \frac{v^2}{2a^2}, \\ &= \frac{3}{4}H^2 a^2 x^2 \Omega_m(a) + H\mathbf{v} \cdot \mathbf{x} + \phi + \frac{v^2}{2a^2}, \end{aligned} \quad (11)$$

where the third line is assuming a flat universe without radiation, so $\Omega_\Lambda = 1 - \Omega_m$. We can calculate E from the two different frames and get consistent results for the same particle. However, the split between which part of the energy is considered as potential energy and kinetic energy is frame-dependent. For a binding check, this implies that we may define a procedure which is consistent between resting and expanding reference frames by requiring $E < E_\infty$, where E_∞ is the reference energy above which a particle may escape the potential well. The motivation of using the physical energy, E , is that it is closest to a conserved quantity for static physical systems that have decoupled from the expansion of the Universe.

The remaining challenge is to define E_∞ . In the scenario considered in Section 2.1, it appeared natural to consider saddle-points of the potential as the largest energies and furthest points where particles may be bound. However, due to the quadratic term in Eq. (10), we have different saddle-points in the comoving and the physical frame. In the absence of dark energy, saddle points of the physical potential signify points that would stay at a fixed physical distance if $\dot{r} = 0$, whereas saddle-points of the peculiar potential signify points that stay at a fixed comoving distance if $v = 0$. This ambiguity is similar to the ambiguity in the notion of boundedness, as pointed out above.

A further necessary consideration is that it is easy to find examples where (for both the comoving and physical potential) there is no saddle point. For instance, we can take an overdensity $\delta(r)$ for which the spherical average within a radius, r , is monotonically decreasing and positive, $\bar{\delta}(r) \geq 0$, for all r , and $\bar{\delta}(r \rightarrow \infty) \rightarrow 0$. If we compute the physical potential,

$$\partial_r \phi_r = \frac{4\pi G}{3} r \rho_m [1 + \bar{\delta}(r)], \quad (12)$$

we find that it is monotonically increasing ($\partial_r \phi_r > 0$), and, as such, it has no saddle point. In contrast, for the peculiar potential, it is

$$\partial_r \phi = \frac{4\pi G}{3} r \rho_m \bar{\delta}(r), \quad (13)$$

which reaches $\partial_r \phi \rightarrow 0$ only at infinity. In practice, neither case is convenient.

Pragmatically, we must therefore refine our definition and ask what is the furthest point that we may want to consider in a binding check. To answer this question, we can call upon the spherical collapse formalism. In this context, we may propose

that we do not want to consider in our binding check any points beyond the turn-around radius, r_{ta} , where spherically collapsing shells are characterised by $\dot{r} = 0$. Within this formalism, this radius corresponds to an average enclosed density $\bar{\delta}$, with $\bar{\delta} \approx 4.55$ for an Einstein-de-Sitter (EdS) Universe (see e.g. Mo et al. 2010). From these ingredients, we can construct a potential of

$$\partial_r \phi_* = \frac{4\pi G \rho_m}{3} r [\bar{\delta}(r) - \bar{\delta}], \quad (14)$$

explicitly forcing the existence of a saddle point at $\bar{\delta}(r) = \bar{\delta}$. We refer to this potential as the turn-around potential, which can further be expressed as

$$\phi_* = \phi - \frac{1}{4}H^2 \Omega_m r^2 \bar{\delta}, \quad (15)$$

$$= \phi - \frac{4\pi G}{6} \rho_m r^2 \bar{\delta}. \quad (16)$$

While this potential only guarantees, for a spherically symmetric system, that the saddle will be located at the turn-around radius, it does guarantee that at sufficiently large radii, a saddle-point and a deeper minimum exist even for rare pathological cases. In terms of this potential, the physical energy is given by

$$E = \frac{1}{4}H^2 r^2 [\Omega_m(a)(1 + \bar{\delta}) - 2\Omega_\Lambda(a)] + \phi_* + \frac{\dot{r}^2}{2}, \quad (17)$$

and we define the reference energy, E_∞ , as the energy level of a particle that is physically at rest $\dot{r} = 0$ (therefore, $\mathbf{v} = -a^2 H \mathbf{x}$) at the saddle-point, $\mathbf{x}_{\text{sad},*}$,

$$E_\infty = \frac{1}{4}H^2 a^2 x_{\text{sad},*}^2 [\Omega_m(a)(1 + \bar{\delta}) - 2\Omega_\Lambda(a)] + \phi_*(\mathbf{x}_{\text{sad},*}). \quad (18)$$

We note that the choice of the potential that is used to define the saddle-point energy level leaves a degree of freedom that may seem somewhat arbitrary. However, in practice the difference between the different potential notions is almost always negligible, since they only tend to differ at much larger radii than the location of the saddle-point. We will show this in more detail in Appendix A.3.

2.3. The boosted potential

Definitions of energy depend on the adopted reference frame. For example, a Galilean transformation changes the kinetic energy term and it modifies the time-dependence of the potential. Furthermore, a transformation into an accelerated reference system with an arbitrary time-dependent offset, $\mathbf{x}_0(t)$,

$$\mathbf{x}' = \mathbf{x} - \mathbf{x}_0(t), \quad (19)$$

$$\mathbf{v}' = \mathbf{v} - \mathbf{v}_0(t), \quad (20)$$

with $\mathbf{v}_0 = a^2 \dot{\mathbf{x}}_0$, introduces an additional fictitious force in the transformed coordinates,

$$\dot{\mathbf{v}}' = \dot{\mathbf{v}} - \mathbf{a}_0(t), \quad (21)$$

with $\mathbf{a}_0(t) = \dot{\mathbf{v}}_0(t)$. With the equation of motion $\dot{\mathbf{v}} = -\nabla_x \phi$, we can interpret this as a modification of the potential,

$$\phi_b(\mathbf{x}', t) = \phi(\mathbf{x}(\mathbf{x}', t), t) + \mathbf{x}' \cdot \mathbf{a}_0(t), \quad (22)$$

which we call a ‘boosted potential’ (Stücker et al. 2021). According to the equivalence principle, transformations of this type do not affect the internal dynamics of systems, but they do

change the energies that we would define for particles. Therefore, we need to choose a unique reference frame to define energies. Here, we argue that the preferred reference frame is the one that minimises the explicit time dependence of the potential in a region of interest. In such a frame, the conservation of energies is the least violated and a binding check is the most meaningful.

Since the potential arises from the mass distribution, the explicit time-dependence of the mass distribution, $\rho(\mathbf{x}', t)$, is expected to be small in an optimal reference frame. As such, given a set of particles $\{\mathbf{x}_i, \mathbf{v}_i\}$ (e.g. all particles in a region of interest), we chose our reference frame so that the centre of mass, the centre of mass velocity, and the centre of mass acceleration are fixed at 0:

$$\langle \mathbf{x}' \rangle(t) = 0, \quad (23)$$

$$\langle \mathbf{v}' \rangle(t) = 0, \quad (24)$$

$$\langle \mathbf{a}' \rangle(t) = 0, \quad (25)$$

where the expectation values go over all considered particles. We can then evaluate the transformations in terms of our original coordinates:

$$\mathbf{x}' = \mathbf{x} - \mathbf{x}_0(t) = \mathbf{x} - \langle \mathbf{x} \rangle(t), \quad (26)$$

$$\mathbf{v}' = \mathbf{v} - \mathbf{v}_0(t) = \mathbf{v} - \langle \mathbf{v} \rangle(t), \quad (27)$$

$$\mathbf{a}' = \mathbf{a} - \mathbf{a}_0(t) = \mathbf{a} + \langle \nabla \phi(\mathbf{x}) \rangle(t). \quad (28)$$

Combining the definition of the boosted potential with that of the turn around potential, we define

$$\phi_{b,*} = \phi - \langle \nabla \phi \rangle \cdot \mathbf{x} - \frac{1}{4} H^2 \Omega_m r^2 \tilde{\delta}, \quad (29)$$

in the accelerated reference frame. This potential is the field that will be used in the following section to define the dynamical boundary of haloes, and energies of particles by replacing ϕ_* with $\phi_{b,*}$ in Eqs. (17) and (18). Indeed, this definition grants us all of the benefits discussed above by both minimising the explicit time dependence of the potential, while also guaranteeing that there will always exist a saddle point to which we can define the boundary of the potential well.

3. Algorithm

In the previous section, we lay down the theoretical framework upon which we constructed our particle assignment scheme. In this section, we go over more specifically how this framework is adapted to sets of particles and how it is implemented to produce groups of particles bound to potential wells. To summarise this section, the STRAWBERRY halo assignment algorithm (hereafter simply referred to as the algorithm or STRAWBERRY) proceeds in four main steps:

1. Propose a region of interest;
2. Apply the transformation to an accelerated frame;
3. Find the saddle-point contour of the boosted turn-around potential, $\phi_{b,*}$;
4. Perform the binding check by evaluating whether $E < E_\infty$ for all particles within the saddlepoint contour.

3.1. Moving about the simulation

Before getting into the core of the algorithm, one crucial step is required; namely, to translate the theoretical framework of Sect. 2 (expressed in terms of continuous fields) into a form applicable to a discrete set of particles. In practice, this step is

important because we rely on features of these fields to define structures, namely minima, gradients, and saddle points, making them an integral part of the numerical scheme. As discussed further in Sect. 3.3, we constructed structures by iteratively adding neighbouring particles until we fill the potential well. This requires two elements: a means of computing the boosted potential at the location of each visited particle and then a definition of neighbouring particles.

As such, STRAWBERRY begins by establishing a connectivity between particles. In this context, connectivity can be seen as a set of edges which connect particles to each other, assigning to each one a set of neighbours. There are many ways that one could define this connectivity. For instance, a common choice is to use a Delaunay tessellation (Forman 1998, 2002; Sousbie 2011; Tierny et al. 2018; Stücker et al. 2021), which has many desirable properties, such as guaranteeing that any particle is accessible from any other particle in the simulation, while using a minimal number of edges. This kind of connectivity is computationally expensive, making cheaper alternatives more desirable given the large number of particles used in modern cosmological simulations.

Here, we have instead opted for a simpler connectivity definition. For each particle, the algorithm saves a list of the N_{ngb} nearest neighbours, with STRAWBERRY tracking the ten nearest neighbours of each particle by default. This choice is both simple to implement, given the fixed size of the list of neighbours, and is efficient to compute, through the use of the k -d tree algorithm (Bentley 1975) for example. Because this particular connectivity does not have the guarantees of more advanced choices, it can be prone to forming clumps and islands, prohibiting the exploration of certain parts of the simulation. However this can be easily avoided if the number of saved neighbours is sufficiently high. This is explained in Appendix A.2, where we show that for the number of neighbours used by default, the recovered physical properties of haloes are converged. Alternatively, additional information about the field that is being explored can be provided; for example, SUBFIND (Springel et al. 2001), which establishes an effective connectivity based on the density field by connecting each particle to its two neighbours with highest density within a preselected list of nearest neighbours. While this may seem more efficient, in the case at hand, this would mean recomputing the connectivity each time we switch reference frame, resulting in a significantly increased computational cost.

3.2. Initialisation and reference frame

As stated earlier in this work, the definition of the boosted potential relies on local averages in order to pass into the accelerated reference frame of the structure. Without an initial estimate of the positions and sizes of structures, it is difficult to define these quantities (see e.g. Stücker et al. 2021). As a first guess, our algorithm uses an initial catalogue of particles split into FoF groups. It is from these groups that we compute mean acceleration, $\langle \mathbf{a} \rangle$, along with the mean position of particles, $\langle \mathbf{x} \rangle$, which provides an initial guess as to the location of the minimum of the potential well. We note that these FoF groups only serve as a selection of the zones of interest that the STRAWBERRY algorithm will visit, and are not used during the assignment and binding procedures described in Sect. 3.3.

Using this initial guess we compute the boosted potential for all particles within the FoF group and select the particle which has the lowest potential value. In the majority of cases,

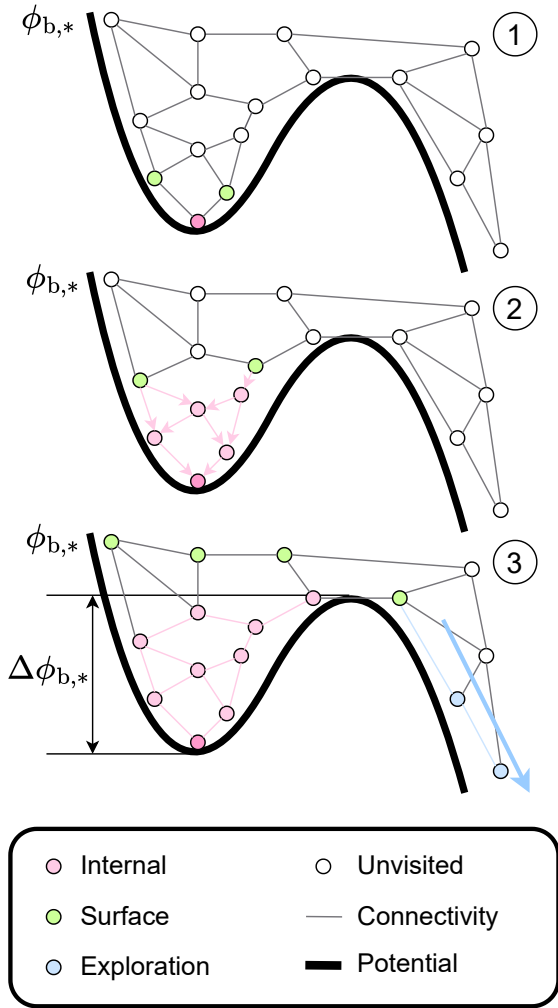


Fig. 3. Illustration of main particle assignment scheme implemented in STRAWBERRY. We start by defining the turn around boosted potential, $\phi_{b,*}$, shown as a thick black line, and assigning the connectivity, thin grey lines, between the particles, drawn as circles. We then place ourselves at the bottom of the potential well and defining the group within the well, pink particle, and tracking particles adjacent to this group, in green, also known as surface particles. In the second panel, we grow the pink group by including surface particles if all of their neighbours with lower potentials are also part of the group. In the third panel, we stop growing the group as we have found a surface particle, to the right, which is connected to particles, in blue, which lead to a deeper potential valley. This final particle sets the persistence of the group, $\Delta\phi_{b,*}$.

over 99 percent, the resulting particle has no neighbours with lower potential values, meaning that it resides at the bottom of a potential well. In the remaining cases, which are primarily highly perturbed haloes or numerical artefacts, the selected particle does not fulfil this condition, typically residing on a slope in the potential at the boundary of the FoF group. In these cases, we performed an approximate gradient descent starting from the mean position of particles in the FoF group, which finds the minimum of the well. In the rare cases where this second procedure did not find a minimum and instead led to particles that were not within the FoF group (i.e. in fewer than 0.01 percent of haloes), then the FoF halo was discarded.

3.3. Filling the potential well

With all of this in hand, we can go on to fill the boosted potential well and locate the saddle point. We achieve this by defining a connected group of ‘internal particles’ and continuously keeping track of a set of ‘surface particles’, which are all the neighbours of the internal group that are not marked as internal particles. We initialised these two groups, respectively, as containing only the potential minimum and all of its neighbours. To grow the group and fill the potential well, the assignment procedure transfers surface particles to the internal group until the saddle point is found, after which the binding procedure is performed on the internal particles. We illustrate this process in Fig. 3, where we take a simple one-dimensional example for ease of comprehension.

In the second phase (see panel 2), we grow the internal group by iteratively including surface particles. To do so, we select the surface particle which has the lowest potential value and include it in the group if all its neighbours with lower potential values are also part of the internal group. If the particle is added to the growing structure, all of its unassigned neighbours are then considered surface particles.

We repeat this step until we find a surface particle that does not satisfy this condition; namely, which has neighbours with lower potentials which are not part of the internal group. This configuration can occur in two cases: either we have found the saddle point that leads to a deeper potential valley, or we have found a sub-structure, as is illustrated in Fig. 2. To check whether we have found the saddle point leading to a deeper potential valley, we create a new exploratory sub-group (see the blue particles in Fig. 3). We explore this region by repeatedly connecting the lowest potential exploration-group-adjacent unvisited particle to this sub-group. This procedure continues until we either reach a particle that has a boosted potential value lower than the minimum or we run out of particles with potential values lower than the candidate saddle point. If we encounter a particle that has a lower potential than the minimum of our well, then we consider that we have found the desired saddle point and, thus, filled the well. In this case, we can define the persistence, $\Delta\phi_{b,*}$, as the difference in the potential between this saddle point and the lowest point in the well. If not, we are dealing with a sub-minimum. All the selected particles are then included into the internal group and all of their unassigned neighbours tracked as surface particles. We note that we keep track of these sub-groups of particles as they require special attention during the binding check.

3.4. Binding check

The final step of this assignment pipeline is the binding check. The energy of every selected particle is computed using Eq. (17) and then compared to the energy at the saddle point given by Eq. (18), after recentring the reference frame around the minimum of the potential well. Particles that have an energy lower than the saddle point energy are considered bound, while those above are considered unbound, thereby creating two populations of particles (bound and unbound) inhabiting the potential well. From this selection, we define the bound mass, M_{bound} , of the halo as the total mass of particles assigned to the bound group.

In most dark matter models, haloes present a certain amount of sub-structure in the form of sub-haloes. These objects, just as their hosts, are gravitationally bound, and influence the potential landscape by creating local minima, which are in motion within the potential of the host halo. When it comes to the binding procedure, these local minima are problematic because if we do

not account for their motion, they allow certain particles to pass the binding check that would otherwise have failed. An intuitive example of this is a sub-structure that is being accreted onto the central halo. Let us consider a particle within the sub-structure which is, by chance, moving in exactly the opposite direction to the bulk motion of the sub-structure. In the reference frame of the central halo, its peculiar velocity is then $\mathbf{v} = \mathbf{0}$. In this case, the energy of the particle is simply given by the sum of the potential of the central halo and infalling halo, along with the Hubble term. Given this configuration and provided that the infalling structure is sufficiently close to the potential minimum, the lack of apparent motion will allow the particle to easily pass the binding check. To avoid this issue, we introduce a secondary bulk binding check which is applied to sub-structures.

Before doing so, we must first decide which particles should be considered part of a sub-group. In practice, this equates to simply analysing the local potential minimum in the same manner as the main group. We recall that during the group assignment procedure, we assigned certain particles to candidate sub-groups for being in the vicinity of a local dip in the boosted potential. In the case where the sub-group is sufficiently large², it is likely that the local minimum in the potential is generated by the presence of a sub-structure.

After the central halo's binding energy is found, each sub-group is re-analysed, in the same way we analysed the main group and using each candidate sub-group as a seed for a sub-structure in the same way we used a FoF group for the main halo. Using these sub-group particles, regardless of whether they are bound to the main structure or not, we measure the local average acceleration, finding the particle that has the lowest potential in this sub-well and repeating the assignment and binding procedures. From this, we obtain a new sub-group of particles that are gravitationally bound to the local minimum.

We treat the binding of the whole sub-group to the host halo as a single object, with its kinetic energy being defined by the bulk velocity of the structure and using its maximum potential, in the reference frame of the main structure, to estimate its potential energy. Given these quantities, we either bind or discard the entire structure, in the same fashion as individual particles in the main group.

4. The bound population of haloes

The STRAWBERRY algorithm presented in the previous sections allows us to study large quantities of haloes through the rapid analysis of simulation snapshots, and producing catalogues of halo properties. In this section, we analyse the resulting particle distributions to better understand the bound population. Specifically, we first focus on where and when particles become bound to their host, how these particles are then distributed within the halo in terms of energies, and how this selection translates to positions and velocities. Here, we focus on halo dynamics in a single Λ CDM cosmology and leave a larger analysis of any possible cosmological dependence for future work.

To perform our study, we used a gravity-only N-body Λ CDM simulation, for which the cosmological parameters were set to the values of the Illustris TNG300 simulation (Springel et al. 2018); namely, $\Omega_{m,0} = 0.3089$, $\Omega_{\Lambda,0} = 0.6911$, $\Omega_{b,0} = 0.045$, $n_s = 0.9667$, and $h = 0.6774$. The simulation was run using the GADGET 4 (Springel et al. 2021) code and represents a comoving periodic box of side length, $L_{\text{box}} = 205 h^{-1}$ Mpc, sampled by $N = 625^3$ particles, leading to a particle mass of $3 \cdot 10^9 h^{-1} M_{\odot}$.

² By default 20 particles.




Fig. 4. Median energy histories of three sets of particles from a sample of 100 haloes with masses $M_{\text{bound}} \in [0.5 \cdot 10^{14} h^{-1} M_{\odot}, 1.5 \cdot 10^{14} h^{-1} M_{\odot}]$, selected according to their respective time of binding. The individual energy histories of particles from different haloes are normalised by the persistence $\Delta\phi_b(a=1)$ of the host halo before the median is computed. The median persistence history of the sample of haloes is shown in black.

The initial conditions are generated at $z_{\text{init}} = 99$ using second-order Lagrangian perturbation theory as implemented in the built-in initial condition generator, NGENIC (Springel 2015), using the Eisenstein & Hu (1998) approximation for the linear power spectrum.

4.1. When and where particles become bound to a halo

The question of when and where particles become bound to haloes echoes dynamical halo boundary definitions, where the boundary of haloes is set not in terms of density but the dynamics of particles such as particle orbits or energies (Diemer 2017; García et al. 2023; Salazar et al. 2025). To answer this question, we selected a sample of 100 haloes with bound masses between $5 \cdot 10^{13} h^{-1} M_{\odot}$ and $1.5 \cdot 10^{14} h^{-1} M_{\odot}$ for which we tracked the assembly history from $a = 0.5$ to $a = 1$, and simultaneously tracked all particles which enter within a spherical region of radius $R < 1.5\sqrt{3} \simeq 2.6 h^{-1}$ Mpc centred on the position of the main progenitors. During the tracking procedure, we travelled backwards in time, defining the main progenitor group as the FoF halo within the tracked region that contains the most particles that will be bound to the main descendant halo in the following snapshot. Once the main progenitor was detected, we recentred the region on its location before proceeding with the particle assignment procedure described in Sect. 3. Within each region we tracked the individual positions, velocities, boosted potential values, and energies of all particles.

From the resulting dataset, we estimated the energy histories of all particles that enter these 100 regions. Among these energy tracks, we selected three sub-samples of particles that become bound to their host haloes at scale factors of $a \in [0.6, 0.7]$, $a \in [0.7, 0.8]$, and $a \in [0.8, 0.9]$. To stack energy histories originating from different haloes, we normalised the individual histories by the persistence of their host halo at $a = 1$. In Fig. 4, we show the evolution of the median energies of all three sub-samples along with the evolution of the median persistence, normalised at $a = 1$. Over time, the energy of particles continuously grows slowly by about 20% to 30% between $a = 0.5$ and

$a = 1.0$ up to the time when they are accreted, after which their energy stays constant. On the other hand, the persistence grows at a significantly faster rate, almost doubling since $a = 0.5$. Since particles become bound when they cross the persistence threshold, the core reason that particles become bound here is that the persistence grows – primarily by additional mass being accreted outside of the orbital radius of the particles of interest.

We note that this is a quite different picture to what would be observed if the self-potential was used to define energies. A notable difference is that the self-potential is normalised to zero at infinity, whereas our potential notion is normalised to be zero at the minimum. In the self-potential, particles become bound because they decrease their self-energy and cross the zero-energy level. For instance, if material is accreted at far larger radii than the orbital radius of a particle of interest, then the self-energy decreases, whereas the energy defined here would be unaffected (while the persistence of the potential increases). This can be seen in Fig. 4, where the energy history of particles flattens out after they become bound. We argue that this is a desirable property and suggest that normalising energies to zero at the minimum may make comparisons of energy levels across time more meaningful. Indeed, we can see that the energies remain clearly separated and maintain their ordering after accretion, indicating that once the halo is virialised it maintains a memory of the accretion process. We note that such a perspective also simplifies understanding other processes, such as tidal stripping (e.g. Drakos et al. 2020; Errani et al. 2022; Stücker et al. 2023).

To gauge at what point of their orbit particles become bound, in Fig. 5, we compute the fraction of bound particles after a given number of dynamical times (Jiang & van den Bosch 2016), expressed as

$$N_{\text{dyn}}(a|a_{\text{peri}}) = \frac{\sqrt{2}}{\pi} \int_{a_{\text{peri}}}^a \frac{1}{a} \sqrt{\frac{\bar{\rho}}{\rho_c}} da, \quad (30)$$

since the first pericentric passage, with $\bar{\rho}$, the characteristic density of haloes which we choose to be 200 times the mean cosmological density, ρ_c , the critical density, a , the scale factor, and, a_{peri} , the scale factor at first pericentric passage. The pericentre is chosen as the first minimum radius with respect to the position of the potential minimum, and maximum velocity norm with respect to the average velocity of the bound population. This specific choice of variable allows us to stack the binding histories of all particles, regardless of the time of their first pericentric passage.

We find that the fraction of bound particles begins to increase as soon as particles enter the saddle point contour of the halo, we mark this moment with a green circle, with the error bars around this point marking the standard deviation of N_{dyn} between the particle entering the halo and its first pericentric passage. Most particles that become bound at this stage are particles that are accreted with lower-than-average energies. As we approach the pericentre, the fraction of bound particles increases rapidly, with 50% of particles being bound to the structure at the first pericentric passage, marked by an orange star³. At the first apocentre, marked by a red square, we find that 75% of particles are considered bound to the structure. Finally, we observe that after two dynamical times, 90% of the particles are considered bound. From these considerations, we confirm that gravitational binding occurs on a relatively short timescale, typically only a




Fig. 5. Fraction of bound particles as a function of the number of dynamical times since the first pericentric passage, shown as the median, (solid black line) and 1- σ region (shaded area) as estimated over 100 haloes. The coloured points represent different events in a particle’s history, green, the particle enters the potential well, orange, the particle passes the pericentre of its orbit for the first time, and red, the particle reaches its first apocentre.




Fig. 6. Stacked radial phase space density distribution of 100 massive haloes in the simulation. In magenta, we show the bound population and in green the infalling population. From darkest to brightest, both contours jointly encompass, 50%, 84%, and 95% of all particles. We note that the splashback contour only displays one shade.

few dynamical times. These observations support the motivations behind particle selection criteria based on orbital status. For instance, a particle may be selected based on whether it has passed its first pericentre or apocentre (e.g. Adhikari et al. 2014; Diemer 2017), meaning parallels can be drawn between these types of mass definition and the one presented in this work.

The binding of particles can further be explored in phase space. Indeed, it is generally accepted that in phase space, CDM haloes can be represented as ravelled three-dimensional manifolds (Shandarin & Zeldovich 1989). When projected in terms of radius, r , and radial velocity, v_r , haloes appear as a spiral pattern, with the most recently accreted material being deposited on the outskirts of the halo and older material being concentrated in the centre. The halo is then connected to the background medium through a stream of infalling material (see e.g. Wang et al. 2011; Adhikari et al. 2014; Zavala & Frenk 2019).

In Fig. 6, we show the stacked radial phase space density of the same 100 haloes, The phase space density is split into

³ This point is by definition at $N_{\text{dyn}} = 0$ and hence does not have any dispersion.

bound (magenta), unbound and infalling (green), and unbound and splashback (orange) populations. We select bound particles according to the boosted potential criterion presented above and separate the unbound sets using the orbits of particles. The brightness of each contour is set to enclose 50%, 84%, and 95%, respectively from darkest to lightest, of the total population, including both bound and unbound population. Here, the bound population occupies a triangular region around the centre of the halo. This population is relatively symmetric in terms of radial velocity. As such it is reasonable to think that this population is virialised, we discuss this quantitatively in Sect. 4.3. This region of phase space, is flanked at larger radii and radial velocities by material classified as unbound. We observe that the second component can be split into two parts, one with negative radial velocities, which we refer to as the first infall stream and a second with positive radial velocities, which we refer to as the first splashback stream. We note that the first splashback stream is less dense than the first infall stream. This is in line with our previous observation, where we find that by the first pericentric passage approximately half of particles are already bound to the halo, as such they will not be included here. This configuration of bound and unbound particles is consistent with the kinetic energy-radius cuts proposed by Salazar et al. (2025) to filter out recently accreted material. Indeed, when plotted in terms of kinetic energy and radius, the distinction between bound and unbound particles (see Fig. 6) follows the valley features used in Salazar et al. (2025) to define the orbiting and infalling populations.

4.2. Matter density profiles

Given the previous observations, we go on to investigate in greater detail how the bound population is distributed by studying the matter density profile of the resulting structures. Through simulations, it has been shown that haloes exhibit the same density profile over at least 20 orders of magnitude in terms of halo mass (Wang et al. 2020). In particular, it has been shown that the density profile is well described by a Navarro-Frenk-White (NFW) profile (Navarro et al. 1997), which depends solely on a single scale radius and amplitude. This profile, however, has the significant shortcoming of not having a convergent mass, meaning that integrating the profile to infinitely large radii does not converge to a finite value. Because of this ambiguity over where haloes ‘end’, it has become common practice to define mass with respect to an average enclosed overdensity Δ , which has the advantage of being independent of the specific profile of the halo, thus standardising the definition of ‘halo mass’. In recent years, this standardisation has been put into question. Indeed, using dynamical information (Adhikari et al. 2014), it has been shown that haloes, contrary to the most popular profiles used to describe them, have a finite extent as shown by Diemer (2017, 2022), García et al. (2023), and Salazar et al. (2025), amongst others.

In the upper panel of Fig. 7, we show the average matter density profile of bound and unbound particles, respectively, as pink and green lines. Here, the unbound population is defined as all particles around the halo that are not considered bound, meaning that we also include particles ‘outside’ the potential well. Each line style corresponds to the median profile over different mass ranges, as indicated by the legend. Both axes are normalised with respect to the radius, $r_{70\%}$, containing 70% of the bound mass.

We see in Fig. 7 that all bound profiles exhibit a similar behaviour. Indeed, for all mass bins, we observe that up to roughly $r_{70\%}$, the slope of the profile varies slowly as a function




Fig. 7. Upper panel: Median matter density profiles of haloes split into bound (pink), unbound (green), and all (grey) particles using STRAWBERRY for all haloes in three mass bins. Lower panel: Median logarithmic derivative of the matter density profile split into bound (pink), unbound (green), and all (grey) particles using STRAWBERRY for three mass bins. The axes are normalised with respect to the radius, $r_{70\%}$, containing 70% of the bound mass. The bound and total profiles are cut at six times the softening length of the simulation. The unbound profiles are only shown for bins containing more than ten particles.

of radius. Beyond this point, the profile steepens rapidly, exhibiting an exponential cut-off. As such, we find that haloes possess a clear-cut edge, beyond which particles are not bound to the structure. This boundary results in a convergent halo mass, justifying the definition of a bound mass. This boundary is, however, not directly visible in the density profile and can only be obtained in terms of particle dynamics. Indeed, this result is analogous to that of Diemer (2022), where the author obtains an exponential cut off in the profile when selecting only particles that have passed their first pericentre. Considered in tandem with Fig. 5, this result is unsurprising, as the fraction of bound particles rises sharply as it approaches its first pericentric passage, meaning that in practice both approaches result in a similar selection of particles.

To further quantify the evolution of the slope of the density profile, in the lower panel of Fig. 7, we show the median logarithmic slope, $\gamma = d \log_{10} \rho(r) / d \log_{10} r$, for each mass bin. For the bound population, the slope decays slowly from $\gamma = -1$ at $r = 0$ to $\gamma \approx -3$ at $r = r_{70\%}$, which is consistent with the NFW or Einasto profile. Further from the centre, on the other hand, we observe an exponential decay of the profile, with the logarithmic slope plunging rapidly. Beyond this point, the grey curves representing the logarithmic slope of the total profile (i.e. including both bound and unbound particles) detach from the bound population. We observe that the density of the unbound population decreases at a slower rate than its bound counterpart, becoming dominant farther from the centre, with the total profile beginning to significantly depart from the bound population around $2 r_{70\%}$. Beyond the exponential cut-off in the bound profile, the logarithmic slope of the total population remains above $\gamma \gtrsim -3$, before presenting a dip, also known as the splashback feature




Fig. 8. Median comoving matter density profile of haloes with bound masses, $9 \cdot 10^{12} h^{-1} M_{\odot} < M_{\text{bound}} < 2 \cdot 10^{13} h^{-1} M_{\odot}$. For each redshift, we separate the total profile, in grey, into three components: the bound population, in magenta, the unbound population, in green, and particles outside the potential well, in blue. For each line colour, higher redshifts are marked by darker colours. The dotted pink lines represent the best fit Diemer (2023) model B profiles to the bound population.

(Adhikari et al. 2014), between 2 and 3 times $r_{70\%}$ and increasing again.

In Fig. 8 we plot the same decomposition of the halo profile, but for a single tighter mass range, $9 \cdot 10^{12} h^{-1} M_{\odot} < M_{\text{bound}} < 2 \cdot 10^{13} h^{-1} M_{\odot}$, and multiple redshifts, $z = 0, 0.5, 1$, and $z = 2$. Interestingly, we observe that the individual components of the profile evolve differently. We find that the outskirts of the bound population change very little with time. This is despite an apparent change to the inner slope and hints at a tight correlation between the bound mass and size of the bound region. Conversely, we find the evolution of the total profile is substantial in this region and is primarily dominated by the evolution of the unbound population. We observe that as the background density decreases with decreasing redshift, so does the density of infalling material surrounding the halo.

In this same figure, we fit the Diemer (2023) model B profile,

$$\rho(r) = \rho_0 \exp[S(r, r_s, r_t, \alpha, \beta)],$$

$$S = -\frac{2}{\alpha} \left[\left(\frac{r}{r_s} \right)^{\alpha} - 1 \right] - \frac{1}{\beta} \left[\left(\frac{r}{r_t} \right)^{\beta} - \left(\frac{r_s}{r_t} \right)^{\beta} \right] + \frac{1}{\eta} \left(\frac{r_s}{r_t} \right)^{\beta} \left[\left(\frac{r}{r_s} \right)^{\eta} + 1 \right], \quad (31)$$

finding a good agreement between the parametric form and the density profile of the bound populations. Moreover, we find that the resulting best fit parameters are consistent with the ranges quoted in Diemer (2025), indicating that it is possible to distinguish between the orbiting and infalling populations, as defined in these works, using purely instantaneous information, meaning without the need to reconstruct particle orbits. For completeness, in Appendix B we fit this model to the density profiles of individual haloes and present the resulting best fit parameters.

4.3. Virialisation

In theoretical settings, haloes are considered to be stable, gravitationally bound structures. To have a system that remains finite




Fig. 9. Median virial ratios as a function of mass, in purple for STRAWBERRY selected particle distributions, in green for STRAWBERRY unbound particles, and in orange for SUBFIND selected particle distributions. The brightness of each curve corresponds to the redshift of the sample, with darker tones representing higher redshifts.

over long time spans it is required that it fulfils the virial theorem:

$$2T + G = \langle \mathbf{v} \cdot \mathbf{v} \rangle + \langle \mathbf{r} \cdot \mathbf{a} \rangle = 0, \quad (32)$$

where G is the virial, T is the mean kinetic energy, and \mathbf{r} , \mathbf{v} , and \mathbf{a} are all defined as having vanishing means. Here we use the more fundamental ‘force’ virial theorem, as it does not require any additional assumptions, which are needed to define the potential energy used in the more common ‘potential’ virial theorem. In its original form, the brackets represent time averages performed over the trajectory of single particles. However, under the assumption that the population is fixed and ergodic, these time averages can be replaced by averages over a sufficiently large number of N -body particles sampling different orbits inside the system.

In simulations, it has been shown that particle distributions recovered by classical halo finding algorithms do not fulfil the virial condition, typically underpredicting the ratio $-G/T$ (e.g. Shaw et al. 2006; Poole et al. 2006; Davis et al. 2011). This effect is understood to be caused by the unintended inclusion of unbound particles in the outskirts of haloes (Shaw et al. 2006). Indeed, as discussed in Sect. 4, these particles have not yet mixed with the bound population and tend to have larger kinetic energies while being located at larger radii. As such their inclusion increases the average kinetic energy, T , and thus decreases the virial ratio.

The result of this effect can be seen in Fig. 9 where the orange curves represent the median virial ratio measured from SUBFIND particle distributions. Here the aforementioned offset is clearly visible, with the median virial ratio being approximately 10% to 25% below the expected value, depending on the redshift. We note that the redshift evolution of these curves is in line with our previous observations that haloes, on average, tend to be surrounded by more infalling material at higher redshifts, thus resulting in a larger bias towards high kinetic energies.

In Stücker et al. (2021), the authors show that the bound population at $z = 0$ is virialised with the resulting virial ratio being in agreement with the theoretical prediction. Here, we expand on these previous results and also show the redshift evolution. Indeed, the STRAWBERRY selected particles, shown in purple,




Fig. 10. Median mass ratios as a function of M_{200b} . The colour of the lines correspond to the ratio in question, in grey $M_{200b}/M_{\text{bound}}$, and in pink $M_{\text{bound}}(<r_{200b})/M_{\text{bound}}$. The shade of each line corresponds to the redshift, specifically, from brightest to darkest, $z = 0, 0.5, 1.0,$ and 2.0 .

yield a median in agreement with the theoretical expectation. We note a residual shift towards lower values at high masses and high redshifts which we attribute to recently formed, accretion-dominated haloes that have not had time to fully relax. For the purpose of visual clarity, we have not shown the dispersion around the median, which is roughly constant of the order of 10%. Furthermore, if we plot the median virial ratios estimated only using the particles considered unbound by STRAWBERRY (green lines in Fig. 9) we observe a very strong bias toward lower values, and also recover the redshift trend seen in the SUBFIND selected particles. This confirms that it is indeed the inclusion of unbound particles which produces the decrease seen in classical halo finders.

4.4. Bound masses

In this final section, we investigate how our previously introduced mass definition, M_{bound} , relates to the wider ecosystem of masses used in cosmology. Indeed, Stücker et al. (2021), observe that, at $z = 0$, the bound mass appears to closely track the SO mass M_{200b} . Here, we wish to see if this is also the case at higher redshift. To begin, we estimate the median mass ratio $M_{200b}/M_{\text{bound}}$ as a function of M_{200b} . This ratio, shown as grey lines in Fig. 10, is computed for four redshifts, $z = 0, 0.5, 1.0,$ and 2.0 , corresponding with the shade of each line, respectively from brightest to darkest. We plot the resulting curves omitting the scatter around these measurements for visual clarity, the latter being roughly constant, $\sigma \approx 0.2$, for all masses and redshift.

At redshift $z = 0$, we find that the bound mass is comparable to M_{200b} . At higher redshifts, we see that while both masses remain qualitatively similar, the ratio does shift to a lower value. This trend is driven by the bound region being more extended at higher redshifts, which can also be seen in Fig. 8. We can further observe this trend by plotting the ratio of the amount of bound mass within r_{200b} with respect to the total bound mass, in pink, using the same brightness convention as previously. In doing so, we can clearly observe that this ratio increases with redshift. This indicates that at higher redshifts r_{200b} contains less bound material than at lower redshift, reinforcing the observation that the bound region is more extended at higher redshifts.

To further understand the origin of this trend, we plot the median fraction of unbound material within r_{200b} , in green, following the same convention for higher redshifts. In this case, the redshift dependence appears to be less visible. This indicates that while M_{200b} will primarily contain bound material, the physical region it will be probing will change as a function of redshift. A possible explanation for the observed trend is that r_{200b} is defined with respect to the total matter density and cosmological background. As such, it will not only depend on how matter is distributed inside the halo but also on the cosmological background evolution. This creates a compounding effect where changes in the density profile, as seen in Fig. 8, and cosmological background both change the definition of r_{200b} (see e.g. Diemer & Kravtsov 2014).

5. Discussion

In this work, we present a novel algorithm for the assignment of particles to haloes in cosmological simulations. This algorithm, operating on the boosted gravitational potential (Stücker et al. 2021), defines a halo as a set of particles that is energetically confined to a potential well. In this context, gravitational binding is determined by directly comparing the energy of particles to the escape energy of the halo given by the saddle point in the potential that leads to a lower minimum. As such, this can be seen as a self-consistent and parameter-free definition of haloes. Particles with lower energies are retained and considered bound to the structure, while particles with higher energies are considered unbound as they would theoretically be able to exit through the saddle point. This algorithm presents a considerable improvement with respect to that used by Stücker et al. (2021), being both more numerically efficient and taking into account additional effects; for instance, in the handling of energies on large scales and at redshifts of $z \neq 0$, allowing for a more in-depth analysis.

Equipped with a clear binding-notion, we studied how particles become bound to haloes. In this formalism, we find that despite the fact that the energy of particles is continuously increasing, particles become bound because the saddle point energy increases faster. Moreover, we find that once a particle is accreted, binding takes place on a short time scale, typically only in one dynamical time. This observation echoes the particle selection criteria of Adhikari et al. (2014) and Diemer (2017), where orbiting particles were selected on the basis of whether they have passed a given point in their first orbit.

Similarly to previous studies (e.g. Wang et al. 2011), we observe that newly accreted particles are deposited on large orbits in the outskirts of the halo. In addition, we find that the bound population is virialised and exhibits a well defined edge. On the other hand, the unbound population, which inhabits the outskirts, is not virialised and does not exhibit such an edge and even obscures the edge of the bound population when looking at the density profile of all particles surrounding the halo. In phase space and energy space, we find that this edge separates the bound population from the infalling and first splashback streams, qualitatively corresponding to a feature previously used to define the orbiting population (García et al. 2023; Salazar et al. 2025).

At higher redshifts, as seen in Fig. 8, we identified distinct evolution mechanisms for both populations. At similar bound masses, the bound component does not exhibit significant differences at the edge, with the extent of the bound region remaining relatively stable. In contrast, the evolution of the total profile is substantial in this region, being mainly driven by the unbound population. As such, in the outskirts the evolution is driven by

external factors such as the accretion rate, the background density and external tidal fields.

This work paves the way for deeper studies of dark matter haloes, particularly in cases where a clear definition in terms of energies is advantageous. For instance, it is useful when creating idealised initial conditions for haloes, where the relation between the density and gravitational potential is used to perform the Eddington inversion and obtain the particle distribution function (e.g. Binney & Tremaine 2008). Alternatively, the definition of haloes outlined here can be of interest when a distinction between a virialised and non-virialised state is needed. This might be the case in studies of how gas settles in the potential well or how tidal stripping affects haloes. Finally, Ondaro-Mallea et al. (2022) associated deviations from universality in the halo mass function defined at M_{200b} to the mass accretion history of haloes. In this context and given the stability of the bound profile described in Sect. 4.2, it would be of interest to further investigate whether M_{bound} provides a more universal alternative to standard spherical overdensity masses.

Data availability

The codes designed for and used throughout this work along with additional examples are publicly available on GitHub: <https://github.com/trg-richardson/strawberry>

Acknowledgements. The authors would like to thank Eduardo Rozo, Edgar Salazar, Lurdes Ondaro-Mallea, and the anonymous referee for their insightful comments. This project has received funding from the Spanish Government's grant program "Proyecto de Generación de Conocimiento" under grant number PID2021-128338NB-I00. JS acknowledges funding by the Austrian Science Fund (FWF) [10.55776/ESP705]. Many of the numerical computations carried out, and plots presented in this work were made possible by the SciPy (Virtanen et al. 2020), NumPy (Harris et al. 2020), Matplotlib (Hunter 2007), and Bacco (Angulo et al. 2021) packages.

References

- Abbott, T. M. C., Abdalla, F. B., Alarcon, A., et al. 2018, *Phys. Rev. D*, **98**, 043526
- Abbott, T. M. C., Aguena, M., Alarcon, A., et al. 2022, *Phys. Rev. D*, **105**, 023520
- Adame, A. G., Aguilar, J., Ahlen, S., et al. 2025, *JCAP*, **2025**, 028
- Adhikari, S., Dalal, N., & Chamberlain, R. T. 2014, *JCAP*, **2014**, 019
- Alam, S., Aubert, M., Avila, S., et al. 2021, *Phys. Rev. D*, **103**, 083533
- Angulo, R. E., Hahn, O., & Abel, T. 2013, *MNRAS*, **434**, 3337
- Angulo, R. E., Zennaro, M., Contreras, S., et al. 2021, *MNRAS*, **507**, 5869
- Asgari, M., Mead, A. J., & Heymans, C. 2023, *Open J. Astrophys.*, **6**, 39
- Aubert, D., Pichon, C., & Colombi, S. 2004, *MNRAS*, **352**, 376
- Ayçoberry, E., Pranjali, R. S., Benabed, K., et al. 2025, *A&A*, **693**, A182
- Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013, *ApJ*, **762**, 109
- Bellagamba, F., Sereno, M., Roncarelli, M., et al. 2019, *MNRAS*, **484**, 1598
- Benson, A. J., Cole, S., Frenk, C. S., Baugh, C. M., & Lacey, C. G. 2000, *MNRAS*, **311**, 793
- Bentley, J. L. 1975, *Commun. ACM*, **18**, 509
- Binney, J., & Tremaine, S. 2008, *Galactic Dynamics*, 2nd edn. (Princeton: Princeton University Press)
- Bleem, L. E., Klein, M., Abbot, T. M. C., et al. 2024, *Open J. Astrophys.*, **7**, 13
- Bond, J. R., Cole, S., Efstathiou, G., & Kaiser, N. 1991, *ApJ*, **379**, 440
- Bonnet, G., Nezri, E., Kraljic, K., & Schimd, C. 2022, *MNRAS*, **513**, 4929
- Bulbul, E., Liu, A., Kluge, M., et al. 2024, *A&A*, **685**, A106
- Castro, T., Borgani, S., Dolag, K., et al. 2021, *MNRAS*, **500**, 2316
- Codis, S., Gavazzi, R., Dubois, Y., et al. 2015, *MNRAS*, **448**, 3391
- Codis, S., Jindal, A., Chisari, N. E., et al. 2018, *MNRAS*, **481**, 4753
- Conroy, C., Wechsler, R. H., & Kravtsov, A. V. 2006, *ApJ*, **647**, 201
- Contreras, S., Angulo, R. E., & Zennaro, M. 2021, *MNRAS*, **508**, 175
- Cui, W., Knebe, A., Yepes, G., et al. 2018, *MNRAS*, **480**, 2898
- Cui, W., Dave, R., Knebe, A., et al. 2022, *MNRAS*, **514**, 977
- Davis, A. J., D'Aloisio, A., & Natarajan, P. 2011, *MNRAS*, **416**, 242
- Despali, G., Giocoli, C., & Tormen, G. 2014, *MNRAS*, **443**, 3208
- Despali, G., Giocoli, C., Angulo, R. E., et al. 2016, *MNRAS*, **456**, 2486
- Diemer, B. 2017, *ApJS*, **231**, 5
- Diemer, B. 2022, *MNRAS*, **513**, 573
- Diemer, B. 2023, *MNRAS*, **519**, 3292
- Diemer, B. 2025, *MNRAS*, **536**, 1718
- Diemer, B., & Kravtsov, A. V. 2014, *ApJ*, **789**, 1
- Diemer, B., More, S., & Kravtsov, A. V. 2013, *ApJ*, **766**, 25
- Drakos, N. E., Taylor, J. E., & Benson, A. J. 2020, *MNRAS*, **494**, 378
- Eisenstein, D. J., & Hu, W. 1998, *ApJ*, **496**, 605
- Eisenstein, D. J., & Hut, P. 1998, *ApJ*, **498**, 137
- Eisenstein, D. J., & Loeb, A. 1995, *ApJ*, **439**, 520
- Errani, R., Navarro, J. F., Ibata, R., & Peñarrubia, J. 2022, *MNRAS*, **511**, 6001
- Euclid Collaboration (Mellier, Y., et al.) 2025, *A&A*, **697**, A1
- Falck, B. L., Neyrinck, M. C., & Szalay, A. S. 2012, *ApJ*, **754**, 126
- Forman, R. 1998, *Adv. Math.*, **134**, 90
- Forman, R. 2002, *Sémin. Lothar. Comb.*, **48**, 35
- García, R., Salazar, E., Rozo, E., et al. 2023, *MNRAS*, **521**, 2464
- Hahn, O., Martizzi, D., Wu, H.-Y., et al. 2017, *MNRAS*, **470**, 166
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, **585**, 357
- Heymans, C., Tröster, T., Asgari, M., et al. 2021, *A&A*, **646**, A140
- Hilton, M., Sifón, C., Naess, S., et al. 2021, *ApJS*, **253**, 3
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Jiang, F., & van den Bosch, F. C. 2016, *MNRAS*, **458**, 2848
- Jing, Y. P., Mo, H. J., & Börner, G. 1998, *ApJ*, **494**, 1
- Knebe, A., Knollmann, S. R., Muldrew, S. I., et al. 2011, *MNRAS*, **415**, 2293
- Leroy, M., Garrison, L., Eisenstein, D., Joyce, M., & Maleubre, S. 2021, *MNRAS*, **501**, 5064
- Lesci, G. F., Marulli, F., Moscardini, L., et al. 2022, *A&A*, **659**, A88
- Mead, A. J., Tröster, T., Heymans, C., Van Waerbeke, L., & McCarthy, I. G. 2020, *A&A*, **641**, A130
- Mo, H., van den Bosch, F. C., & White, S. 2010, *Galaxy Formation and Evolution* (Cambridge: Cambridge University Press)
- Nadler, E. O., Mansfield, P., Wang, Y., et al. 2023, *ApJ*, **945**, 159
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, *ApJ*, **490**, 493
- Nelson, D., Pillepich, A., Ayromlou, M., et al. 2024, *A&A*, **686**, A157
- Ondaro-Mallea, L., Angulo, R. E., Zennaro, M., Contreras, S., & Aricó, G. 2022, *MNRAS*, **509**, 6077
- Ortega-Martinez, S., Contreras, S., & Angulo, R. 2024, *A&A*, **689**, A66
- Peacock, J. A., & Smith, R. E. 2000, *MNRAS*, **318**, 1144
- Pellissier, A., Hahn, O., & Ferrari, C. 2023, *MNRAS*, **522**, 721
- Planck Collaboration XXVII. 2016, *A&A*, **594**, A27
- Poole, G. B., Fardal, M. A., Babul, A., et al. 2006, *MNRAS*, **373**, 881
- Press, W. H., & Davis, M. 1982, *ApJ*, **259**, 449
- Press, W. H., & Schechter, P. 1974, *ApJ*, **187**, 425
- Roy, F., Bouillot, V. R., & Rasera, Y. 2014, *A&A*, **564**, A13
- Salazar, E. M., Rozo, E., García, R., et al. 2025, *Phys. Rev. D*, **111**, 043527
- Seljak, U. 2000, *MNRAS*, **318**, 203
- Sereno, M. 2015, *MNRAS*, **450**, 3665
- Sereno, M., Maurogordato, S., Cappi, A., et al. 2025, *A&A*, **693**, A2
- Shandarin, S. F., & Zeldovich, Y. B. 1989, *Rev. Mod. Phys.*, **61**, 185
- Shaw, L. D., Weller, J., Ostriker, J. P., & Bode, P. 2006, *ApJ*, **646**, 815
- Sousbie, T. 2011, *MNRAS*, **414**, 350
- Springel, V. 2015, *Astrophysics Source Code Library* [record ascl:1502.003]
- Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001, *MNRAS*, **328**, 726
- Springel, V., White, S. D. M., & Hernquist, L. 2004, *IAU Symp.*, **220**, 421
- Springel, V., Pakmor, R., Pillepich, A., et al. 2018, *MNRAS*, **475**, 676
- Springel, V., Pakmor, R., Zier, O., & Reinecke, M. 2021, *MNRAS*, **506**, 2871
- Stücker, J., Hahn, O., Angulo, R. E., & White, S. D. M. 2020, *MNRAS*, **495**, 4943
- Stücker, J., Angulo, R. E., & Busch, P. 2021, *MNRAS*, **508**, 5196
- Stücker, J., Ogiya, G., Angulo, R. E., Aguirre-Santaella, A., & Sánchez-Conde, M. A. 2023, *MNRAS*, **521**, 4432
- Tiery, J., Favelier, G., Levine, J. A., Gueunet, C., & Michaux, M. 2018, *IEEE Trans. Visual Comput. Graphics*, **24**, 832
- Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, *ApJ*, **688**, 709
- Umetsu, K., Sereno, M., Lieu, M., et al. 2020, *ApJ*, **890**, 148
- Vale, A., & Ostriker, J. P. 2006, *MNRAS*, **371**, 1173
- Vera-Ciro, C. A., Sales, L. V., Helmi, A., et al. 2011, *MNRAS*, **416**, 1377
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Meth.*, **17**, 261
- Wang, J., Navarro, J. F., Frenk, C. S., et al. 2011, *MNRAS*, **413**, 1373
- Wang, J., Bose, S., Frenk, C. S., et al. 2020, *Nature*, **585**, 39
- Zavala, J., & Frenk, C. S. 2019, *Galaxies*, **7**, 81

Appendix A: Convergence

This appendix is devoted to validating the various numerical aspects of the algorithm presented in Sect. 3, and provide additional specifics regarding the implementation of the latter. Most notably, we first test the effect of numerical resolution on the quantities presented throughout this work before investigating the effect of the two main design choices, namely the number of neighbours assigned to each particle and the impact of the large scale potential dampening factor, $\tilde{\delta}$.

A.1. Simulation resolution

Here, we first test the effect of resolution. To visualise this, we measure the HMF and median persistence at redshift, $z = 0$, as estimated in three simulations which differ solely by the number of particles they contain. To do so, in addition to the simulation used throughout the main text we run two lower resolution simulations, respectively with $N = 156^3$, and $N = 312^3$ particles, using the same realisation of the Gaussian random noise field.

In Fig. A.1, we plot the HMF estimated from the measured bound masses. We observe that the HMF converges for haloes containing more than a few hundred particles. Below this mass, we see a significant difference between the higher and lower resolution simulations. This is due to two effects, one intrinsic to the resolution and a second selection effect. The first effect is caused by the loss of spatial resolution in the inner regions of haloes. This reduces the finesse with which the algorithm can locate the minimum in the potential, effectively smoothing out the inner cusp. This smoothing, results in a lower persistence and as such a lower mass. The second effect mentioned above comes from the choice of which seeds are or are not considered starting points for the algorithm. Indeed, here, we only consider haloes detected using SUBFIND as potential seeds if they contain more than 50 particles. These haloes will have a range of bound masses, meaning that the resulting selection in terms of bound halo masses introduces a systematic bias at the low mass end, and is in fact dominant here.

The first effects becomes more apparent when plotting the median persistence as a function mass as shown in Fig. A.2. Indeed, as stated above, this is because the central cusp of the potential is sampled with fewer particles. This results in an underestimation of the potential depth when using lower resolution simulations, but also for smaller haloes which are sampled by fewer particles. In Fig. A.2, this is visualised by the slow convergence of the curves at high masses and high resolution. While this does not significantly affect the recovered masses, as evidenced by the absence of bias in Fig. A.1, it is clear that resolution effects should be taken into account when analysing these persistences, with an observed relative difference of the order 20% between the median persistence of estimated using the highest resolution simulation and both lower simulations.

A.2. Number of neighbours

We go on to examine the algorithmic choices that are made in Sect. 3, starting with the number of neighbours we assign to each particle. To do so, we repeat the same exercise as above but changing the number of neighbours assigned to each particle. In the context of this work, we have chosen to limit ourselves to considering a fixed number of nearest neighbours, $N = 10$, for all particles. To justify this decision, in Fig. A.3 we estimate the HMF for four different numbers of nearest neighbours, specifically $N = 5, 10, 15$, and 20. At high masses, we observe a gen-




Fig. A.1. Upper panel: Bound halo mass function as estimated in three simulations with the same cosmology but differing particle numbers. Lower panel: Relative error with respect to the highest resolution ($N = 625^3$) simulation. The shaded region corresponds to a 5% deviation.




Fig. A.2. Upper panel: Median persistence as a function of halo mass for three particle resolutions. For each curve, the shaded region contains 68% of haloes in the corresponding mass bin. Lower panel: Relative difference with respect to the median persistence in the highest resolution simulation.

eral agreement between all the curves. However, at low masses we observe a significant divergence of the $N = 5$ HMF, with this particular choice of neighbours under-predicting the abundance of haloes in this mass range. The main systematic in this case, is a clumping effect, which hinders exploration of the medium in low density regions. We can see this in Fig. A.3 where there is good convergence passed $N = 10$ but for $N = 5$ we overestimate the abundance of intermediate mass haloes and underestimate the abundance of the low mass counterparts. Beyond $N = 10$, we do not observe any significant changes that would justify the additional computational resources needed for these connectivity definitions. Given these observations, we decide to use the




Fig. A.3. Upper panel: Bound halo mass function as estimated using STRAWBERRY outputs on a single simulation but changing the number of nearest neighbours used by the algorithm. Lower panel: Relative error with respect to the nominal value ($N = 10$) of nearest neighbours. The shaded region corresponds to a 5% deviation.

$N = 10$ nearest neighbours of each particle to define the connectivity.

A.3. Large scale potential dampening

As discussed in the main-text, the definition of the gravitational potential in an expanding (and uniformly accelerated) universe inevitably leaves a degree of freedom associated to what should be considered ‘at rest’. While we have proposed the pragmatic choice of $\tilde{\delta} = 4.55$, targetting the turn-around radius in spherical collapse scenarios in EdS universes as a saddle-point, we show here that the precise value of this has a very minor impact on the haloes that are identified.

It is however important to note that this decision does have an influence on the estimated physical quantities and as such is worth analysing in further detail. To do so, we repeat the measurements presented above for ten values of $\tilde{\delta}$ ranging from 0 to 10.939⁴. Starting with the HMF shown in Fig. A.4. Here, we observe that depending on the value of the parameter, the abundance of haloes of a given mass can be either underestimated or overestimated up to the order of 5% to 10%. This is indicative that while the large scale correction is designed to operate only on the largest haloes, it impacts haloes at all masses. This is because for small and intermediate mass haloes, the correction still lowers the level of the saddle point, even if its contribution is sub-dominant. As a result, certain particles are considered unbound because their mechanical energy falls between the physical potential depth and the artificially lowered value. In turn, this shifts the HMF towards lower masses, as observed here.

Regarding the effect $\tilde{\delta}$ has on the persistence, we observe in Fig. A.5 a 5% shift of the $\Delta\phi_{b,*} - M_{\text{bound}}$ relation. As the




Fig. A.4. Upper panel: Bound halo mass function as estimated using STRAWBERRY outputs on a single simulation using ten different values of the large scale parameter, $\tilde{\delta}$, ranging from 0 to 10.939. Lower panel: Relative difference with respect to halo mass function estimated using the nominal value ($\tilde{\delta} = 4.55165$) used in this work. The shaded region corresponds to a 5% deviation.




Fig. A.5. Upper panel: Median persistence as a function of bound halo mass for ten values of $\tilde{\delta}$ ranging from 0 to 10.939. Lower panel: Relative difference with respect to the median persistence estimated for $\tilde{\delta} = 4.55$.

persistence is corrected for $\tilde{\delta}$, this shift is indicative of higher persistence haloes being assigned lower masses for larger values of large scale parameter. While this deviation is smaller than the impact of the resolution on the persistence it is nonetheless important to keep in mind when studying the persistences.

Appendix B: Profile fitting parameters

As stated in the main text we find that, on average, the Diemer (2023) model B profile provides a good fit to the density profile of the bound population. In this appendix, we provide an

⁴ This specific value corresponding to the turn-around density contrast at $z = 0$ for the cosmology of the simulation.




Fig. B.1. Median and $1\text{-}\sigma$ dispersion of best fit Diemer (2023) model B profile parameters as a function of bound mass, respectively solid lines and shaded regions. Dotted lines represent the best fit power law to the median.

expanded view where we fit this model to the individual log profiles of haloes using a standard χ^2 minimisation procedure.

In Fig. B.1 we plot the distributions of resulting fitting parameters as a function of the bound mass. We observe that both the scale radius and truncation radius vary significantly with bound mass. We quantify this trend by fitting a power-law to the median fitted values, excluding the largest mass bins which are dominated by only a few high mass haloes. The best-fit power laws,

$$r_s = (0.05 \pm 0.01) \left[M_{\text{bound}} / 10^{13} h^{-1} M_{\odot} \right]^{0.403 \pm 0.007} h^{-1} \text{Mpc}, \quad (\text{B.1})$$

$$r_t = (0.46 \pm 0.01) \left[M_{\text{bound}} / 10^{13} h^{-1} M_{\odot} \right]^{0.352 \pm 0.007} h^{-1} \text{Mpc}, \quad (\text{B.2})$$

are shown as dotted lines in the upper panel of Fig. B.1. In contrast, and as also observed in Diemer (2025), we find that the two shape parameters,

$$\alpha = (0.20 \pm 0.01) \left[M_{\text{bound}} / 10^{13} h^{-1} M_{\odot} \right]^{-0.041 \pm 0.008} \quad \text{and} \quad (\text{B.3})$$

$$\beta = (4.15 \pm 0.03) \left[M_{\text{bound}} / 10^{13} h^{-1} M_{\odot} \right]^{-0.05 \pm 0.01}, \quad (\text{B.4})$$

vary only slowly as a function of the bound mass.

A full quantification of the behaviour of these parameters as a function of redshift and cosmology, is out of the scope of this work and requires a separate dedicated study. Nonetheless, the ranges of parameters is compatible with those found in Diemer (2025), which can be attributed to the similarity between the SPARTA orbital selection and our equivalent effective selection as illustrated in Fig. 5.