

 Open access • Journal Article • DOI:10.1093/MNRAS/STZ361

Disruption of satellite galaxies in simulated groups and clusters: the roles of accretion time, baryons, and pre-processing — Source link

Yannick M. Bahé, Joop Schaye, David J. Barnes, Claudio Dalla Vecchia ...+6 more authors

Institutions: Leiden University, Massachusetts Institute of Technology, Spanish National Research Council, University of La Laguna ...+2 more institutions

Published on: 11 May 2019 - Monthly Notices of the Royal Astronomical Society (Oxford University Press)

Topics: Satellite galaxy, Galaxy and Dark matter

Related papers:

- [Galaxy evolution in groups and clusters: satellite star formation histories and quenching time-scales in a hierarchical Universe](#)
- [The Hydrangea simulations: galaxy formation in and around massive clusters](#)
- [Populating a cluster of galaxies - I. Results at z=0](#)
- [On the Infall of Matter into Clusters of Galaxies and Some Effects on Their Evolution](#)
- [The evolution of large-scale structure in a universe dominated by cold dark matter](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/disruption-of-satellite-galaxies-in-simulated-groups-and-14or409ydp>

Durham Research Online

Deposited in DRO:

12 March 2019

Version of attached file:

Published Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Bahé, Yannick M. and Schaye, Joop and Barnes, David J. and DallaVecchia, Claudio and Kay, Scott T. and Bower, Richard G. and Hoekstra, Henk and McGee, Sean L. and Theuns, Tom (2019) 'Disruption of satellite galaxies in simulated groups and clusters : the roles of accretion time, baryons, and pre-processing.', *Monthly notices of the Royal Astronomical Society.*, 485 (2). pp. 2287-2311.

Further information on publisher's website:

<https://doi.org/10.1093/mnras/stz361>

Publisher's copyright statement:

© 2019 The Author(s) Published by Oxford University Press on behalf of the Royal Astronomical Society.

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

Disruption of satellite galaxies in simulated groups and clusters: the roles of accretion time, baryons, and pre-processing

Yannick M. Bahé¹*, Joop Schaye¹, David J. Barnes², Claudio Dalla Vecchia^{3,4}, Scott T. Kay⁵, Richard G. Bower^{1,6}, Henk Hoekstra¹, Sean L. McGee⁷ and Tom Theuns⁶

¹Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands

²Department of Physics, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

³Instituto de Astrofísica de Canarias, C/Vía Láctea s/n, E-38205 La Laguna, Tenerife, Spain

⁴Departamento de Astrofísica, Universidad de La Laguna, Av. del Astrofísico Francisco Sánchez s/n, E-38206 La Laguna, Tenerife, Spain

⁵Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK

⁶Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK

⁷School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

Accepted 2019 January 31. Received 2019 January 10; in original form 2018 October 30

ABSTRACT

We investigate the disruption of group and cluster satellite galaxies with total mass (dark matter plus baryons) above $10^{10} M_{\odot}$ in the Hydrangea simulations, a suite of 24 high-resolution cosmological hydrodynamical zoom-in simulations based on the EAGLE model. The simulations predict that ~ 50 per cent of satellites survive to redshift $z = 0$, with higher survival fractions in massive clusters than in groups and only small differences between baryonic and pure N -body simulations. For clusters, up to 90 per cent of galaxy disruption occurs in lower-mass subgroups (i.e. during pre-processing); 96 per cent of satellites in massive clusters that were accreted at $z < 2$ and have not been pre-processed survive. Of those satellites that are disrupted, only a few per cent merge with other satellites, even in low-mass groups. The survival fraction changes rapidly from less than 10 per cent of those accreted at high z to more than 90 per cent at low z . This shift, which reflects faster disruption of satellites accreted at higher z , happens at lower z for more massive galaxies and those accreted on to less massive haloes. The disruption of satellite galaxies is found to correlate only weakly with their pre-accretion baryon content, star formation rate, and size, so that surviving galaxies are nearly unbiased in these properties. These results suggest that satellite disruption in massive haloes is uncommon, and that it is predominantly the result of gravitational rather than baryonic processes.

Key words: methods: numerical – galaxies: clusters: general – galaxies: evolution – galaxies: stellar content.

1 INTRODUCTION

A key prediction of the concordance Λ cold dark matter (Λ CDM) cosmology is that dark matter structures form hierarchically: small objects collapsed first and then built up successively more massive structures through mergers (e.g. Press & Schechter 1974; Searle & Zinn 1978; White & Rees 1978). Galaxy groups and clusters represent the highest level of this hierarchy at the present day,

built up from the largest number of individual accreted galaxies.¹ Once accreted, galaxies are subject to mass loss due to tidal forces and ram pressure stripping, while dynamical friction can drive them towards the centre of their host halo and therefore enhance the mass loss yet further (see e.g. Binney & Tremaine 2008). In this way, the galaxy may be reduced to a mass below a given detection threshold, or even disrupted completely (e.g. Hayashi et al. 2003).

¹We here use the term ‘galaxy’ to refer to distinct self-bound objects, irrespective of their mass or composition. A galaxy therefore includes the dark matter halo as well as stellar component and gas reservoir, where they exist.

* E-mail: bahe@strw.leidenuniv.nl

Understanding the extent to which satellite galaxies survive this mass loss is desirable for a number of reasons. It allows measuring the halo mass from the abundance of galaxies (see e.g. Rozo et al. 2009; Budzynski et al. 2012; Rykoff et al. 2014; Andreon 2015; Saro et al. 2015) or kinematics (e.g. Zhang et al. 2011; Bocquet et al. 2015; Sereno & Etori 2015; see also Armitage et al. 2018). Detailed characterization of substructure is one of the most promising avenues to constrain the nature of dark matter (e.g. Randall et al. 2008; Lovell et al. 2012; Vegetti et al. 2012; Harvey et al. 2015; Robertson et al. 2018). Finally, satellite galaxies differ from isolated galaxies of the same stellar mass in key aspects, such as their colour (e.g. Peng et al. 2010), star formation rate (e.g. Kauffmann et al. 2004; Wetzel, Tinker & Conroy 2012), and morphology (e.g. Dressler 1980). The detailed origins of these differences are still unsolved puzzles, which also requires understanding to what extent satellites survive at all: if, for example, survival correlates with galaxy properties prior to infall, this may (partly) explain the aforementioned differences.

Because of its complexity, this problem needs to be addressed with numerical simulations (see e.g. van den Bosch & Ogiya 2018). Since the late 1990s, these have achieved sufficiently high resolution to avoid ubiquitous numerical dissolution of satellite galaxies (or ‘subhaloes’; e.g. Ghigna et al. 1998; Moore et al. 1999; Springel et al. 2001, 2008; Gao et al. 2012), which prompted a multitude of studies that analysed their evolution and survival in detail (e.g. Tormen, Diaferio & Syer 1998; De Lucia et al. 2004; Gao et al. 2004; Weinberg et al. 2008; Dolag et al. 2009; Xie & Gao 2015; Chua et al. 2017; van den Bosch 2017). The qualitatively consistent conclusion from these studies is that subhaloes survive for a limited amount of time, with the lowest survival rate (i.e. fastest disruption) at both the highest and lowest ends of the subhalo mass range. The majority of surviving subhaloes in massive clusters were therefore accreted relatively recently, at $z < 1$ (De Lucia et al. 2004; Gao et al. 2004). Of those that were accreted earlier, only a small fraction was typically predicted to survive to $z = 0$: both Gao et al. (2004) and Jiang & van den Bosch (2017), for instance, found that only 10 per cent of simulated subhaloes accreted at $z = 2$ could still be identified at $z = 0$.

An inherent limitation in all numerical studies is that limited resolution precludes the identification of subhaloes below a limiting mass, even if they are physically not completely disrupted. If survival is defined as the subhalo retaining at least a given number of particles (e.g. Gao et al. 2004; Xie & Gao 2015; van den Bosch 2017) or a minimum mass set by the resolution of the simulation (e.g. Chua et al. 2017), simulations with higher resolution predict higher survival fractions: for example, Xie & Gao (2015) found that in the Phoenix dark matter only galaxy cluster simulations (Gao et al. 2012), which resolve each cluster with $\sim 10^8$ particles, more than half of all subhaloes with mass above $10^{10} M_{\odot}$ accreted at $z = 2$ survive to the present day.

A more subtle consequence of numerical resolution has been pointed out in a recent series of papers by van den Bosch (2017), van den Bosch et al. (2018), and van den Bosch & Ogiya (2018): they found that the complete disruption of subhaloes should, physically, be extremely rare and that numerical artefacts can occur even well above the nominal resolution limit of a simulation. Through a suite of idealized N -body experiments, van den Bosch & Ogiya (2018) demonstrated that inadequate force softening – i.e. spatial resolution – and particle numbers – i.e. mass resolution – both act to accelerate the tidal disruption of subhaloes, even when they are ‘well resolved’ with $\gtrsim 100$ particles. Due to the extremely demanding resolution requirements found to be necessary to prevent

such numerical disruption, van den Bosch & Ogiya (2018) argued that this constitutes a serious roadblock on the path to understanding the evolution of satellite galaxies.

Another limitation in many of the aforementioned simulations is the neglect of baryons. Ram pressure can efficiently remove gas from infalling galaxies (Gunn & Gott 1972), making them more susceptible to disruption (e.g. Saro et al. 2008), while gas cooling and star formation may have a stabilizing effect through the formation of dense cores, which are more difficult to disrupt. Non-radiative hydrodynamical simulations have given discrepant answers about the impact of gas removal on subhalo survival, with some finding it to be more relevant (Saro et al. 2008; Dolag et al. 2009) than others (Tormen, Moscardini & Yoshida 2004).

The modelling of additional baryonic effects, such as gas cooling, star formation, and its associated energy feedback, remains uncertain (see e.g. Scannapieco et al. 2012 and the discussion in Schaye et al. 2015) and cosmological hydrodynamical simulations accounting for them have long struggled to produce even realistic *isolated* galaxies. They have therefore, perhaps unsurprisingly, led to a variety of contradictory conclusions about the net effect of baryons on satellite survival: Weinberg et al. (2008) found that their inclusion increases survival, particularly in low-mass galaxies, while Dolag et al. (2009) concluded that the effect of gas cooling and star formation is largely cancelled by the disruptive effect of gas stripping. The Illustris simulation (Vogelsberger et al. 2014) predicts a net disruptive effect of baryons (Chua et al. 2017).

With an improved implementation of energy feedback that largely overcomes numerical cooling losses (Dalla Vecchia & Schaye 2012), and by calibrating the uncertain subgrid prescriptions against observational relations in the local Universe, the EAGLE project (Schaye et al. 2015) has produced a population of galaxies that match not only these calibration diagnostics, but also their evolution to high redshift (Furlong et al. 2015, 2017) and a wide range of other observables including galaxy colours (Trayford et al. 2015, 2017), star formation rates (Schaye et al. 2015), and neutral gas content (Lagos et al. 2015; Bahé et al. 2016; Marasco et al. 2016; Crain et al. 2017). This model therefore provides realistic initial conditions to study the evolution of satellite galaxies.

The Hydrangea simulation suite applies this successful model to the scale of galaxy clusters by combining it with the zoomed initial conditions technique (e.g. Katz & White 1993). Despite some tensions in the mass of their simulated central cluster galaxies (Bahé et al. 2017b) and hot gas fractions (Barnes et al. 2017b), the $z = 0.1$ satellite stellar mass function agrees remarkably well with observations, down to stellar masses far below that of the Milky Way (Bahé et al. 2017b). This suggests that the fraction of satellites that survive to the present day is modelled correctly. The Hydrangea suite therefore allows us to study the evolution of satellites in a realistic way, over a wide range of host and galaxy masses.

With this tool, we revisit the question of satellite survival in massive haloes. We aim to address in particular the following three questions: (i) What fraction of accreted satellites survive to $z = 0$, and how does this depend on accretion time, galaxy mass, and host mass? How important, therefore, is satellite disruption² in a simulation suite that is characteristic of the current state of the art

²Throughout this paper, we use ‘disruption’ as antonym to ‘survival’. It therefore refers to the dispersal of galaxies into their host halo as well as to mergers with another galaxy.

in cosmological hydrodynamical simulations that include massive clusters (see also e.g. Pillepich et al. 2018 and Tremmel et al. 2019)? (ii) What is the predicted effect of baryons on galaxy survival? (iii) What is the role of environmental effects on galaxies prior to accretion on to their (final) halo? This ‘pre-processing’ step (e.g. Fujita 2004; Berrier et al. 2009; McGee et al. 2009; Balogh & McGee 2010) has been identified as a key stage in the evolution of cluster galaxies (e.g. Zabludoff & Mulchaey 1998; Berrier et al. 2009; McGee et al. 2009; Bahé et al. 2013; Wetzel et al. 2013; Han et al. 2018), but to our knowledge, no study has so far examined its role in satellite disruption.

The remainder of this paper is structured as follows. Section 2 summarizes the key aspects of the Hydrangea simulations and the relevant post-processing steps, including an overview of our new method to trace simulated galaxies through time. The predicted survival fractions are presented in Section 3, followed by an analysis of the roles of pre-processing, satellite–satellite mergers, and galaxy accretion time in Section 4. We investigate the influence of galaxy properties prior to accretion on their survival in Section 5, and summarize our conclusions in Section 6. In appendices, we provide a detailed description of our new tracing method (Appendix A), a verification of the robustness of our results against numerical limitations (Appendix B), and a comparison to the numerical experiments of van den Bosch & Ogiya (2018, Appendix C). A companion study (Paper II; Bahé et al. in preparation) examines the mechanisms of galaxy disruption and its role in building central group/cluster galaxies and their extended haloes.

Throughout, we assume the same flat Λ CDM cosmology as the EAGLE project, with parameters as determined by Planck Collaboration XVI (2014): Hubble parameter $h \equiv H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1}) = 0.6777$, dark energy density parameter $\Omega_\Lambda = 0.693$ (dark energy equation of state parameter $w = -1$), matter density parameter $\Omega_M = 0.307$, and baryon density parameter $\Omega_b = 0.04825$. All galaxy stellar, dark matter, and total masses are computed as the sum of all gravitationally bound particles of the respective type as identified by the SUBFIND code (see Section 2.2).

2 SIMULATIONS AND POST-PROCESSING

2.1 The Hydrangea simulations

The Hydrangea simulations are part of the C-EAGLE project, a suite of cosmological hydrodynamical zoom-in smoothed particle hydrodynamics simulations of 30 massive galaxy clusters (Bahé et al. 2017b; Barnes et al. 2017b). They were run with the ‘AGNdT9’ variant of the EAGLE model (see table 3 of Schaye et al. 2015), with initial particle masses $m_{\text{DM}} = 9.7 \times 10^6 M_\odot$ and $m_{\text{gas}} = 1.8 \times 10^6 M_\odot$ for dark matter and gas, respectively. The (spatially constant, Plummer-equivalent) gravitational softening length of the simulations is $\epsilon = 0.7$ proper kpc at $z < 2.8$. Here, we provide a succinct summary of their key features and refer to Bahé et al. (2017b) and Barnes et al. (2017b) for more details.

The 30 clusters of the C-EAGLE project were chosen from a low-resolution N -body simulation (Barnes et al. 2017a), in the mass range³ $14.0 > \log_{10}(M_{200c}^z/M_\odot) > 15.4$ at $z = 0$ and without a more massive halo closer than $\max(20r_{200c}, 30 \text{ Mpc})$ at $z = 0$. 24 clusters – the Hydrangea suite – were simulated with a high-resolution

region extending to at least $10r_{200c}$ from the centre of the target cluster (defined as the location of its potential minimum). Within these large zoom-in regions, they contain a multitude of additional lower-mass groups and clusters on the outskirts of the main target cluster.

The EAGLE code (Schaye et al. 2015), which was used for the zoom-in resimulations, is a substantially modified version of the GADGET-3 code (last described in Springel 2005). The changes include updates to the hydrodynamics scheme collectively referred to as ‘ANARCHY’ (Schaller et al. 2015; Schaye et al. 2015) and a large number of subgrid physics models to simulate unresolved astrophysical processes, which are described in detail by Schaye et al. (2015). They include models for radiative cooling, photoheating, and reionization (Wiersma, Schaye & Smith 2009a); star formation based on the Kennicutt–Schmidt relation cast as a pressure law (Schaye & Dalla Vecchia 2008) but with a metallicity-dependent star formation threshold (Schaye 2004); a pressure floor corresponding to $P \propto \rho^{4/3}$ imposed on gas with $n_{\text{H}} \geq 10^{-1} \text{ cm}^{-3}$ to prevent the formation of an inadequately modelled cold gas phase; mass and metal enrichment of gas due to stellar outflows based on Wiersma et al. (2009b); energy feedback from star formation in thermal stochastic form based on Dalla Vecchia & Schaye (2012); and seeding, growth of, and energy feedback from supermassive black holes based on Springel, Di Matteo & Hernquist (2005), Rosas-Guevara et al. (2015), and Schaye et al. (2015).

Particularly relevant to this study is that those subgrid parameters that are not well-constrained by observations – primarily the efficiency scaling of star formation feedback – were calibrated so that the simulated *field* galaxy population matches low-redshift observations in terms of the stellar mass function and stellar sizes (as described by Crain et al. 2015). These are crucial prerequisites for meaningful predictions about the survival of cluster galaxies, because an overly massive or overly compact stellar component may make the simulated galaxies artificially resilient against disruption (and vice versa).

In addition to the main simulation with hydrodynamics and baryon physics, each volume was also simulated in N -body-only mode, i.e. starting from the same initial conditions but assuming that all matter is dark. These ‘DM-only’ simulations allow us to directly quantify the net impact of baryons (see also Armitage et al. 2018).

2.2 Structure identification

The primary output from each simulation consists of 30 snapshots, which are mostly spaced equidistant in time between $z = 14.0$ and $z = 0$ with $\Delta t = 500 \text{ Myr}$. In each of these outputs, structures were identified with the SUBFIND code (Springel et al. 2001; Dolag et al. 2009) in a two-step process.

First, spatially disjoint groups of particles were found with a friends-of-friends (FoF) algorithm with a linking length of $b = 0.2$ times the mean inter-particle separation. As shown by More et al. (2011), this linking length corresponds approximately (within a factor of ≈ 2) to a limiting isodensity contour of $\delta \equiv \rho/\rho_{\text{mean}} = 82$. The FoF algorithm is applied only to DM particles; baryon particles are attached to the FoF group (if any) of their nearest DM neighbour particle (Dolag et al. 2009). Groups with less than $N_{\text{FoF}} = 32$ DM particles are deemed unresolved and discarded.

Within each FoF group, SUBFIND then identifies gravitationally self-bound ‘subhaloes’. This procedure is described in detail by Springel et al. (2001) and Dolag et al. (2009). Candidate subhaloes are identified as locally overdense regions, limited by the isodensity

³ $M_{200c}^z=0$ denotes the total mass within a sphere of radius r_{200c} , centred on the potential minimum of the cluster, within which the average density equals 200 times the critical density.

contour at the density saddle point that separates the candidate subhalo from the local background. Within each candidate, gravitationally unbound particles are iteratively removed and candidates retaining more than 20 particles (excluding gas) are identified as genuine subhaloes. Finally, all particles in the FoF group that are not part of any subhalo are collected into the ‘background’ subhalo, provided that they are gravitationally bound to it.

In the following, we will refer to all subhaloes as ‘galaxies’, including the background subhalo (which is typically the most massive one in any FoF group). The latter will be referred to as ‘central’ and all others as ‘satellites’. This nomenclature is independent of the stellar content of a subhalo (which may be zero); unless specifically stated otherwise, we define galaxies as including all particle types, including their gaseous and dark matter haloes.

Previous work has shown that the subhalo identification step of SUBFIND tends to incorrectly assign particles near the edge of satellites to the central subhalo (e.g. Muldrew, Pearce & Power 2011). In idealized tests, Muldrew et al. (2011) have shown that this can artificially suppress the mass of even massive subhaloes ($M = 10^{12} M_{\odot}$) by as much as 90 per cent near the centre of a galaxy cluster; in extreme cases, it may be lost altogether. Our tracing procedure accounts for this spurious, temporary ‘disruption’ where possible (see below), and we have verified that only a minute fraction of galaxies missing from the $z = 0$ SUBFIND catalogue still exist as self-bound structures (see Appendix B2). One must, however, bear in mind that the masses of satellite subhaloes calculated by SUBFIND may be (substantially) underestimated.

2.3 Tracing galaxies through time

The subhalo catalogues returned by SUBFIND describe the simulated structures at one point in time. In order to follow individual simulated galaxies – physical objects that appear at some point in time and potentially disappear later – these outputs must be linked together as an additional post-processing step. We accomplish this with the ‘SPIDERWEB’ algorithm, a substantially modified version of the procedure outlined in Bahé et al. (2017b). A full description of the code elements and their physical motivation is provided in Appendix A; the following is a brief summary of its main aspects.

SPIDERWEB follows a galaxy through time by identifying the sequence of subhaloes in subsequent snapshots that share the highest fraction of particles. Although this is conceptually straightforward, subtleties arise due to interactions between galaxies, particularly in the dense environments of groups and clusters. We therefore consider multiple candidate descendants for each subhalo in a given snapshot (i), namely all those in the subsequent snapshot (j) that are ‘linked’ to the original subhalo by sharing at least one particle. In the case of multiple links from one subhalo in i , the highest priority is given to the one that contains the largest number of its 5 per cent most bound collisionless particles (its ‘core’). The other links are reserved as backup in case this highest priority link leads to a subhalo in j that already overlaps more closely with another subhalo in i : this may, for example, happen if the galaxy is undergoing severe stripping so that most of its (core) particles are transferred to another galaxy between two snapshots. We note that this approach differs from other ‘merger tree’ algorithms (e.g. Rodriguez-Gomez et al. 2015; Qu et al. 2017), which only consider one possible descendant for each subhalo.

To account for instances of a galaxy temporarily not being identified at all by SUBFIND, SPIDERWEB attempts to re-connect lost galaxies after up to 5 snapshots (corresponding to a maximum gap of 2.5 Gyr at our standard snapshot spacing). Our code also

gives special consideration to the treatment of mergers, by explicitly accounting for prior mass transfers between galaxies when selecting the main progenitor of a subhalo in j that is linked to multiple subhaloes in i .

If no descendant can be found for a subhalo in i , its galaxy is treated as disrupted and ‘merged’ on to the galaxy that contains the largest number of its core particles. By following these target galaxies (possibly over multiple mergers), SPIDERWEB identifies a unique ‘carrier’ galaxy at $z = 0$ as the endpoint of every galaxy that has ever existed in the simulation. For a comprehensive description and justification of these methods, the interested reader is referred to Appendix A.

2.4 Sample selection

Galaxies are characterized by the peak (total) subhalo mass they have ever attained, which we denote as $M_{\text{tot}}^{\text{peak}}$. In contrast to the equivalent mass at $z = 0$ ($M_{\text{tot}}^{z=0}$), this can be homogeneously computed for both surviving and disrupted galaxies, and compared to the stellar peak mass M_{\star}^{peak} , it allows a direct comparison between hydrodynamical and DM-only simulations. There is a fairly tight relation between $M_{\text{tot}}^{\text{peak}}$ and M_{\star}^{peak} (see also Moster, Naab & White 2013 and Behroozi et al. 2018), with a 1σ scatter of typically only ≈ 0.5 dex: $M_{\text{tot}}^{\text{peak}} = 10^{10} (10^{11.5}, 10^{12.5}) M_{\odot}$ corresponds approximately to $M_{\star}^{\text{peak}} = 10^{7.7} (10^{10.1}, 10^{11}) M_{\odot}$.

Here, we analyse galaxies with $M_{\text{tot}}^{\text{peak}} > 10^{10} M_{\odot}$ ($M_{\star}^{\text{peak}} \gtrsim 5 \times 10^7 M_{\odot}$), i.e. those that have at some point been resolved by > 1000 particles. Many baryonic $z = 0$ properties of our simulated galaxies are already unconverged or in tension with observations at $M_{\text{tot}}^{\text{peak}} < 10^{11.5} M_{\odot}$, including stellar masses (at $M_{\text{tot}}^{\text{peak}} < 5 \times 10^{10} M_{\odot}$), sizes, quenched fractions (both at $M_{\text{tot}}^{\text{peak}} < 10^{11} M_{\odot}$), metallicities (Schaye et al. 2015), and neutral gas content (Crain et al. 2017). We include these low-mass galaxies here to test the predicted survival fractions in this poorly converged regime, but emphasize that they should be interpreted with caution, at least to the extent that they deviate between hydrodynamical and DM-only simulations.

We exclude a small number of galaxies ($\ll 1$ per cent at $M_{\text{tot}}^{\text{peak}} > 10^{10} M_{\odot}$) that are formed predominantly from particles that were previously associated with another galaxy. These ‘spectres’ typically correspond to substructures within a more massive galaxy (e.g. a dense part of a spiral arm) that are temporarily identified as a separate subhalo (see Appendix A for further details).

Because the Hydrangea simulations use the zoom-in technique, some subhaloes in each snapshot lie close to the edge of the high-resolution region and may be subject to numerical artefacts. We therefore exclude all galaxies from our analysis whose potential minimum lies closer than 5 comoving Mpc from a low-resolution boundary particle in any snapshot. We also exclude a very small population of low-mass galaxies (< 0.1 per cent at $M_{\text{tot}}^{\text{peak}} < 10^{11} M_{\odot}$) that have no identifiable carrier at $z = 0$ because all their particles became unbound when they were disrupted.

2.5 Satellite accretion times

As a final step, we need to identify galaxies that have been accreted by a group or cluster at some point in their lives. Not all of these are satellites at $z = 0$: some may have been disrupted completely, and others may have temporarily or permanently escaped as ‘backsplash’ galaxies (see e.g. Gill, Knebe & Gibson 2005). For each galaxy, we therefore first identify the snapshots in which it

is a satellite⁴; if there are none, the galaxy is discarded. In each of these snapshots, we then find the corresponding central galaxy. The FoF group containing this central (or its carrier, in case the central itself has merged) at $z = 0$ is a candidate host of the galaxy under consideration. If there are multiple candidates (from different snapshots), we select the one that is a candidate in the largest number of snapshots and, in the event of a tie, the one from the earliest snapshot. By definition, all hosts correspond to FoF groups at $z = 0$ and can therefore be classified by their present-day $M_{200c}^{z=0}$.

An illustration of our host assignment scheme is provided in Fig. 1. This follows one galaxy (represented by purple circles) through six consecutive snapshots at times t_0 – t_5 (different rows from top to bottom), with the last row at t_5 corresponding to $z = 0$. Circles in other colours represent other galaxies. The purple galaxy is a satellite in four snapshots (t_1 – t_4), during which it is a member of the FoF groups indicated with dotted ellipses in the colour of their centrals (which are denoted with a ‘C’). Because one of these (blue) is itself a satellite (of the green one) at $z = 0$, there are only two candidate host groups, indicated with the green and grey dashed ellipses in the bottom row. The galaxy under consideration (purple) was associated to the green candidate in three snapshots (t_2 – t_4) and to the grey candidate in only one (t_1). The former is therefore selected as its host, even though the purple galaxy is, in this example, not actually part of it⁵ at $z = 0$.

We exclude galaxies that are the central galaxy of their own host at $z = 0$, which can occur as a result of satellite–central swaps. This only affects 0.1 per cent of our galaxies, but because these all have⁶ $M_{\text{tot}}^{\text{peak}} \approx M_{200c}^{z=0}$, the fraction is almost 50 per cent within the most extreme combination of high $M_{\text{tot}}^{\text{peak}}$ ($> 10^{12.5} M_{\odot}$) and low $M_{200c}^{z=0}$ ($= 10^{12.5} - 10^{13.5} M_{\odot}$). Our final sample contains 165 566 galaxies with $M_{\text{tot}}^{\text{peak}} > 10^{10} M_{\odot}$ that are associated with a host of $M_{200c}^{z=0} > 10^{12.5} M_{\odot}$, including 3 433 with $M_{\text{tot}}^{\text{peak}} > 10^{12} M_{\odot}$.

With a host halo selected for each galaxy, we next find their accretion times. We consider two alternative definitions, but note that a plethora of others have been used in the literature (see e.g. Gao et al. 2004; Xie & Gao 2015; Chua et al. 2017). The ‘branch accretion time’ (t_{branch}) is the middle of the snapshot interval before the galaxy first became a satellite in any progenitor branch of its host halo (in other words, in a halo whose central – or its carrier – at $z = 0$ is in the same group as the galaxy’s host). The ‘main accretion time’ (t_{main}) is taken as the analogous point when the galaxy became a satellite in its actual host halo. Galaxies that never reach their host halo, for example because they disrupted in a side branch (see Section 4.1), are assigned $t_{\text{main}} = \infty$. When a galaxy became a satellite and then merged before the next snapshot was written (so that it is never recorded as a satellite), we assign an




Figure 1. An illustration of our host assignment scheme. Shown are six consecutive snapshots at times t_0 – t_5 (the latter corresponding to $z = 0$). Coloured circles represent four individual galaxies, of which the purple one is currently under consideration. Although it is a central at redshift $z = 0$, it was a satellite in four previous snapshots (t_1 – t_4), with the respective groups outlined by dotted ellipses. Their centrals lie in two FoF groups at $z = 0$, indicated by the grey and green dashed ellipses in the bottom row. These are the two candidate hosts of the galaxy, and the shading behind each snapshot label (on the far left) indicates to which one it was associated at this point. Because it is associated most often to the green candidate, this is selected as the galaxy’s host. The two horizontal red lines indicate the two accretion times used in this paper, corresponding to first infall into the host itself (t_{main}) and one of its progenitor branches (t_{branch}).

accretion time half-way between the last snapshot in which the galaxy was detected, and the first in which it was not.

For the situation depicted in Fig. 1, these two definitions of accretion time are indicated by red horizontal lines. We highlight that t_{branch} is, in this example, not equivalent to the first time at which the purple galaxy became a satellite, because its (brief) association with the grey group in t_1 is not yet part of its accretion into its final host (green).

In Fig. 2, we show the cumulative distribution of both branch (top) and main (bottom) accretion times for galaxies with different peak total galaxy masses ($M_{\text{tot}}^{\text{peak}}$; solid lines in shades of green and blue) and host halo masses ($M_{200c}^{z=0}$; dashed lines in shades of yellow

⁴As noted in Section 2.2, we define satellite status and accretion times in terms of a galaxy’s membership to an FoF group: it is a satellite if it is not the central subhalo of the FoF group to which it belongs. Not all of these satellites are necessarily within r_{200c} from the central, particularly in highly aspherical groups.

⁵Fig. 1 deliberately depicts the non-standard situation of a galaxy that has escaped from its host at $z = 0$, to highlight that our host assignment scheme does not depend (exclusively) on $z = 0$ group membership. The choice of host and accretion times would be exactly the same in the (more typical) situation of the purple galaxy being part of the green group at $z = 0$, or having merged with one of its members.

⁶There are small differences between $M_{\text{tot}}^{\text{peak}}$ and $M_{200c}^{z=0}$ even for galaxies that are their own host, because the latter excludes particles beyond r_{200c} , but also includes unbound particles and those in satellites within this radius.

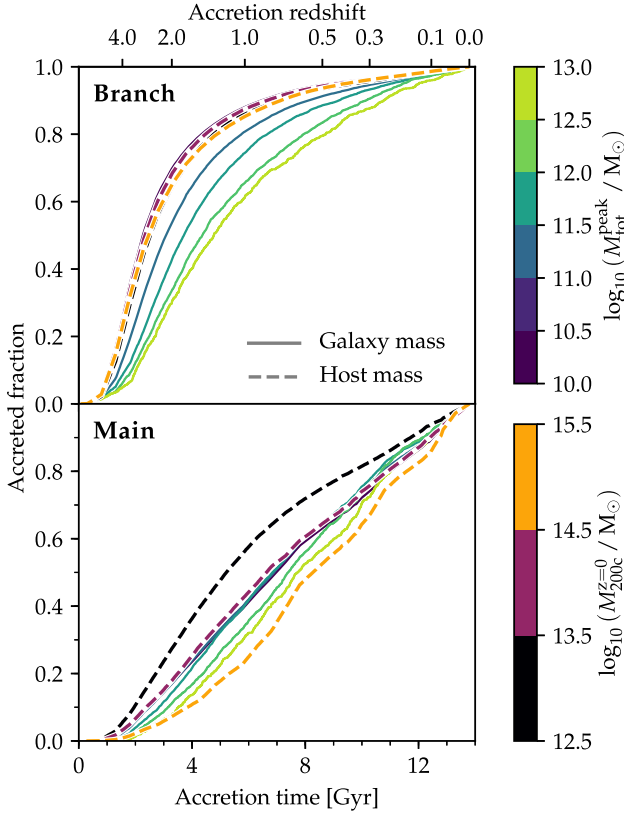


Figure 2. Cumulative distribution of accretion times for different peak total galaxy masses (solid lines) and different host halo masses (dashed lines) in the hydrodynamical simulations. *Top*: accretion on to any branch of the host group (t_{branch}), *bottom*: accretion on to its main progenitor branch (t_{main}), for galaxies with $t_{\text{main}} < \infty$. Values of t_{branch} are predominantly early (around $z \approx 2$), while t_{main} is more evenly spread out. The former depend mostly on galaxy mass, the latter on host mass.

and red). For the former we divide galaxies into six equal bins in log space, from $M_{\text{tot}}^{\text{peak}} = 10^{10}$ to $10^{13} M_{\odot}$. For the hosts, we distinguish between ‘massive clusters’ ($M_{200c}^{z=0} > 10^{14.5} M_{\odot}$; orange), ‘low-mass clusters’ ($M_{200c}^{z=0} = 10^{13.5} - 10^{14.5} M_{\odot}$; lilac), and ‘groups’ ($M_{200c}^{z=0} = 10^{12.5} - 10^{13.5} M_{\odot}$; black).

Due to the set-up of our simulations, the latter two bins are dominated by objects at the periphery of a more massive cluster and therefore not necessarily representative of all haloes in these mass bins. However, we found that the survival fractions shown below only vary by ≈ 5 per cent between galaxies with a host at < 5 and $5 - 10 r_{200c}$ from the central cluster of their simulation volume, respectively.⁷ We are therefore confident that the large-scale environment does not induce a significant bias in our conclusions for lower-mass haloes. For display purposes, all times are offset by a random value of up to ± 250 Myr to suppress artificial discreteness due to the finite number of snapshots.

Galaxies are accreted over a wide redshift range, $4 \gtrsim z \geq 0$. The distribution of t_{branch} (top; median at $z \approx 1.5 - 3$) is more concentrated towards high z than that of t_{main} (bottom; median at $z \approx 0.5 - 1$). In addition, t_{branch} depends strongly on $M_{\text{tot}}^{\text{peak}}$ – more massive galaxies are accreted later (compare the dark blue and yellow-green solid lines) – but hardly on $M_{200c}^{z=0}$ (the orange and black dotted lines lie

⁷The survival fraction is, in general, slightly higher for galaxies whose host lies closer to the central cluster.




Figure 3. Fraction of satellite galaxies that are pre-processed (first accreted by a subgroup, rather than their final host) as a function of their peak total mass. Hydrodynamical simulations are represented by solid lines (shaded bands indicating their binomial 1σ uncertainties), the DM-only runs by dotted lines. Different colours represent galaxies in hosts of different mass, as indicated by the colour bar along the right edge. Pre-processing is ubiquitous, especially for low-mass galaxies and those associated with massive clusters (orange).

almost on top of each other).⁸ The main accretion time shows the opposite behaviour, with a clear difference between different hosts – galaxies in clusters (orange dashed) are accreted later than those in groups (black dashed) – but a much weaker dependence on galaxy mass. There is hardly any difference between hydrodynamical and DM-only simulations (omitted for clarity).

The gap between the median t_{branch} and t_{main} implies a significant role of pre-processing, i.e. that many galaxies first fall into a subgroup that is later accreted by their final host (e.g. Berrier et al. 2009; McGee et al. 2009; Balogh & McGee 2010). This is shown directly in Fig. 3, where we plot the fraction of galaxies with $t_{\text{branch}} < t_{\text{main}}$ as a function of $M_{\text{tot}}^{\text{peak}}$ for the three host mass bins in the hydrodynamical simulations (solid lines; shaded bands indicate binomial 1σ uncertainties following Cameron 2011 both here and in subsequent figures) and the corresponding DM-only runs (dotted lines).

The pre-processed fraction is very high: 87 (73) percent of galaxies in massive clusters with $M_{\text{tot}}^{\text{peak}} > 10^{10}$ (10^{12}) M_{\odot} , and still ≈ 35 per cent of Milky Way analogues ($M_{\text{tot}}^{\text{peak}} \sim 10^{12} M_{\odot}$) in groups (black) were first a satellite in a subgroup that was later accreted by their main host. This is notably higher than what previous authors have found for *surviving* galaxies (only ≈ 50 per cent even in massive clusters; e.g. Bahé et al. 2013; Wetzel et al. 2013; Han et al. 2018). As we show below, this discrepancy arises because many galaxies do not survive the pre-processing stage.

At $M_{\text{tot}}^{\text{peak}} > 2 \times 10^{11} M_{\odot}$, the pre-processed fraction in the hydrodynamical simulations agrees closely with the DM-only runs. Only at lower masses is there a small, but consistent, tendency towards slightly higher pre-processing fractions in the hydrodynamical

⁸There is a slight dependence on $M_{200c}^{z=0}$ when only considering more massive galaxies ($M_{\text{tot}}^{\text{peak}} > 10^{11.5} M_{\odot}$), in the sense that t_{branch} is ≈ 1 Gyr later for low-mass groups than clusters (for clarity not shown in Fig. 2).




Figure 4. The fraction of all accreted galaxies (including pre-processed ones) that survive with $M_{\text{tot}}^{z=0} > 5 \times 10^8 M_{\odot}$. Different host masses at $z = 0$ are indicated by different colours. Both the hydrodynamical simulations (solid lines, shaded bands indicate 1σ binomial uncertainties) and the DM-only counterparts (dotted lines) predict a survival fraction of ~ 50 per cent. At fixed galaxy mass, survival is slightly more common in more massive hosts.

simulations (by < 6 per cent). This could be caused by subtle differences in the halo finder between the two simulation types, or reflect a small impact of baryons on the actual accretion paths of low-mass galaxies.

3 SURVIVAL FRACTIONS OF SATELLITES

We begin by investigating the survival of all galaxies from their point of first accretion (t_{branch}). Our fiducial definition of survival requires that the galaxy is identified by SUBFIND at $z = 0$ and has a mass of at least $M_{\text{tot}}^{z=0} = 5 \times 10^8 M_{\odot}$ (corresponding to ≈ 50 DM or ≈ 270 baryon particles); the effect of varying this threshold is explored below. The survival fraction of all galaxies ever accreted is plotted in Fig. 4 as a function of peak total galaxy mass $M_{\text{tot}}^{\text{peak}}$, in three halo mass bins. Solid lines represent the hydrodynamical simulations (with shaded bands representing binomial 1σ uncertainties, as in Fig. 3), while the corresponding fractions from the DM-only simulations are shown by dotted lines. We have not matched individual galaxy pairs in the two simulation sets, because these may follow significantly different orbits due to amplifications of small differences in the cluster environment (Prins 2018).

The survival fraction is ~ 50 per cent, with only a moderate dependence on galaxy or host mass. Perhaps surprisingly, it is slightly higher in massive clusters than groups (51 versus 44 per cent when averaged over all $M_{\text{tot}}^{\text{peak}} > 10^{10} M_{\odot}$). It is also mildly higher around $M_{\text{tot}}^{\text{peak}} = 10^{12} M_{\odot}$ than at the highest and lowest galaxy masses, at least in clusters (up to 67 per cent). Averaged over our entire sample, 47 per cent of satellites with $M_{\text{tot}}^{\text{peak}} > 10^{10} M_{\odot}$ and $M_{200c}^{z=0} > 10^{12.5} M_{\odot}$ survive at $z = 0$. In Appendix B1, we demonstrate that these numbers are insensitive to an increase in mass resolution by a factor of 8, at least in low-mass groups and for $M_{\text{tot}}^{\text{peak}} \lesssim 3 \times 10^{11} M_{\odot}$ (more massive objects are not sampled well by our high-resolution runs due to their smaller volumes).

A second key feature of Fig. 4 is that the survival fractions in the hydrodynamical simulations closely follow those in their DM-

only counterparts. There are some minor differences, for example at $M_{\text{tot}}^{\text{peak}} \sim 10^{12} M_{\odot}$ in massive clusters – where the inclusion of baryons increases the survival fraction by a few per cent – and at the low-mass end ($M_{\text{tot}}^{\text{peak}} \lesssim 10^{11} M_{\odot}$), where the baryonic galaxies are slightly more susceptible to disruption at fixed $M_{\text{tot}}^{\text{peak}}$, possibly as a consequence of poor resolution (see above). Overall, however, *the effect of baryons on galaxy survival is small*: if star formation and gas stripping separately have non-negligible impact, they happen to cancel each other almost exactly.

The close agreement between the survival fractions in the hydrodynamical and DM-only simulations implies that the former should not contain many remnants that are (almost) completely devoid of dark matter and only survive because of their baryon content. To verify this, we have also computed the survival fractions above a *dark matter* mass threshold of $5 \times 10^8 M_{\odot}$ in the hydrodynamical simulations,⁹ which agree almost exactly with those from the equivalent threshold in total mass (not shown).

This absence of (almost) purely baryonic remnants appears to be in tension with semi-analytic models, which typically require a large fraction of (baryonic) galaxies to survive the disruption of their dark matter subhalo in the form of ‘orphan’ or ‘type 2’ satellites (e.g. Somerville et al. 2008; Guo et al. 2011; Henriques et al. 2015). In the Guo et al. (2011) model applied to the Millennium-II simulation, for example – which has almost exactly the same resolution as the Hydrangea DM-only runs – 25 per cent of all satellite galaxies with $M_{\star}^{z=0} = 10^{9.5} M_{\odot}$ are orphans, and still almost 20 per cent at $M_{\star}^{z=0} = 10^{10.5} M_{\odot}$.

The small net influence of baryons is also at odds with the recent study of Chua et al. (2017), who found that, in the Illustris simulation, the inclusion of baryons *reduces* the survival fraction by ≈ 5 –20 per cent, at all masses. It is plausible that these differences reflect different subgrid physics implementations, so that a destabilizing effect of gas stripping dominates in Illustris, while it is approximately cancelled by the cohesive effect of star formation in Hydrangea.¹⁰

3.1 Influence of the detection threshold

3.1.1 Thresholds in total galaxy mass

In Fig. 4, we counted any galaxy as ‘surviving’ that was identified by SUBFIND at $z = 0$ and had a total mass of at least $5 \times 10^8 M_{\odot}$. To elucidate the sensitivity of our predictions to this threshold, we plot in Fig. 5 the survival fractions with a number of other definitions; for clarity, only the massive cluster bin is shown, but we have verified that the qualitative conclusions also apply to lower-mass hosts.

The top panel compares the survival fractions at our fiducial mass threshold of $M_{\text{tot}}^{z=0} = 5 \times 10^8 M_{\odot}$ (dark blue, identical to the orange lines in Fig. 4) to both those obtained from considering all SUBFIND detections at $z = 0$ as surviving (grey), and two stricter mass thresholds of $M_{\text{tot}}^{z=0} = 3 \times 10^9$ and $10^{10} M_{\odot}$ (medium and light blue, respectively). As in Fig. 4, we show results from

⁹This threshold is not fully equivalent to $M_{\text{tot}}^{z=0} > 5 \times 10^8 M_{\odot}$ in the DM-only (DMO) version, because the DM particles in the DMO simulations also account for the mass contributed by baryons and are therefore more massive, by a factor of $(1 - \Omega_b/\Omega_m)^{-1} = 1.19$.

¹⁰Note that the absolute survival fractions in the DM-only simulation of Chua et al. (2017) are significantly higher than in our Fig. 4, because they do not explicitly include the pre-processing phase. We have verified that this does not account for the different impact of baryon physics in the two simulations.

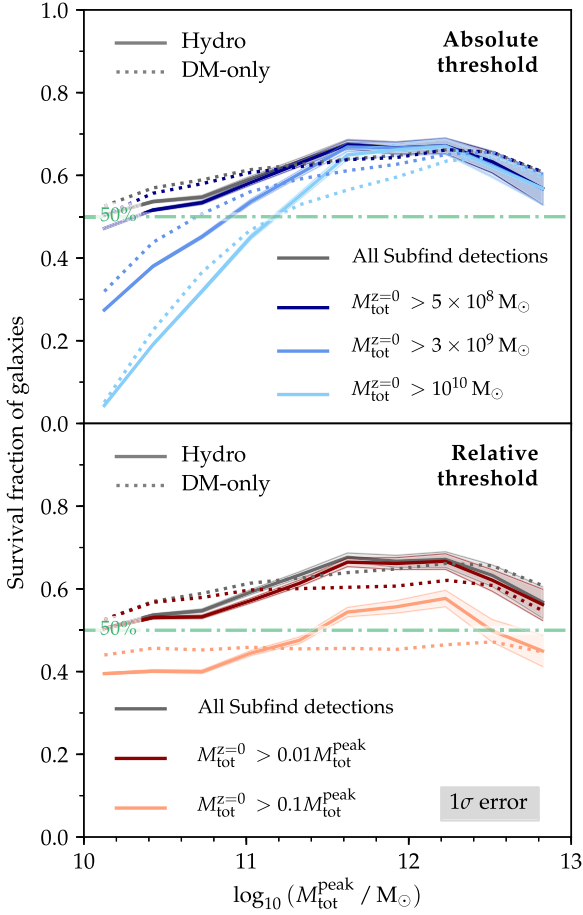


Figure 5. Dependence of satellite survival fractions on the imposed detection threshold in the hydrodynamical (solid lines) and DM-only simulations (dotted lines) of massive clusters. Grey lines show the total survival fraction, i.e. all galaxies detected by SUBFIND at $z = 0$. In the *top panel*, the dark, medium, and light blue lines show, respectively, the fraction of galaxies retaining a total mass of at least 5×10^8 , 3×10^9 , and $10^{10} M_{\odot}$, respectively, at $z = 0$. The *bottom panel* gives the fraction of galaxies that retain at least 1 per cent (dark red) and 10 per cent (light red) of their total peak mass at $z = 0$. All thresholds apart from this last one converge in the hydrodynamic simulations at $M_{\text{tot}}^{\text{peak}} > 3 \times 10^{11} M_{\odot}$. In contrast, many lower-mass galaxies – and in the DM-only simulations even some Milky Way analogues – only survive as low-mass remnants below $10^{10} M_{\odot}$.

the hydrodynamical simulations as solid, and from the DM-only runs as dotted lines. The lower panel shows the survival fractions above two *relative* mass thresholds, requiring the galaxy to retain at least 1 per cent (dark red) or at least 10 per cent (light red) of their peak total mass. Recall that SUBFIND-derived satellite masses may be biased low, so that these lines should more accurately be interpreted as representing lower limits on the true surviving fractions.

Compared to our fiducial threshold of $M_{\text{tot}}^{z=0} = 5 \times 10^8 M_{\odot}$ (dark blue lines in the top panel), the survival fractions hardly increase when including all SUBFIND detections (grey), in both the hydrodynamical and DM-only simulations; only at $M_{\text{tot}}^{\text{peak}} \lesssim 10^{11} M_{\odot}$ is there a difference of a few per cent. This indicates that $M_{\text{tot}}^{z=0} < 5 \times 10^8 M_{\odot}$ remnants can, in principle, be resolved by our simulations, but also that they are very uncommon in the (peak) mass range considered here. This is confirmed in Appendix B1, where we show that the survival fractions of satellites

with $M_{\text{tot}}^{\text{peak}} \gtrsim 3 \times 10^{10} M_{\odot}$ in low-mass groups are unchanged when the mass resolution is increased, and the mass threshold for survival lowered, by a factor of 8. The more restrictive thresholds, on the other hand (medium and light blue), remove a successively larger fraction of galaxies with $M_{\text{tot}}^{\text{peak}} < 3 \times 10^{11} M_{\odot}$ that have a remnant in the $z = 0$ SUBFIND catalogue (69 per cent with $M_{\text{tot}}^{z=0} < 10^{10} M_{\odot}$), indicating a continuous distribution of remnant masses between a lower limit ($\gtrsim 5 \times 10^8 M_{\odot}$) and $M_{\text{tot}}^{\text{peak}}$. Our fiducial limit of $M_{\text{tot}}^{z=0} = 5 \times 10^8 M_{\odot}$ is therefore a physically and numerically meaningful definition of galaxy survival in our simulations.¹¹ At lower resolution, it may not be possible to identify remnants with $M_{\text{tot}}^{z=0} \lesssim 10^{10} M_{\odot}$, which could plausibly account for the higher disruption rates reported by e.g. Jiang & van den Bosch (2017).

An alternative criterion to distinguish between surviving and disrupted galaxies is the fraction of their peak mass retained at $z = 0$. As the bottom panel shows, a relative threshold of 1 per cent of the peak mass (dark red line) agrees to per cent level with the survival fraction from the entire SUBFIND catalogue in the hydrodynamical simulations. This implies a near-total absence of galaxies that lose more than 99 per cent of their mass but still survive as self-bound objects that can be detected at the resolution of our simulations. This is true even amongst the most massive galaxies ($M_{\text{tot}}^{\text{peak}} > 10^{12} M_{\odot}$) for which a remnant with one per cent of its peak mass would be well above the resolution limit of the simulations. In Paper II, we show that this is because massive galaxies predominantly merge with the core of the central group/cluster galaxy, rather than gradually dispersing into its halo.

In contrast, a significant (but nevertheless minor) fraction of galaxies – around 10 per cent in the hydrodynamical simulations, almost independent of mass – are identified by SUBFIND at $z = 0$ but only retain less than one tenth of their peak mass (the difference between the light red and grey lines). These galaxies experienced strong mass loss (plausibly due to tidal stripping), but are nevertheless not disrupted completely.

Although the DM-only versions (dotted lines in Fig. 5) yield broadly the same result as the hydrodynamical simulations discussed so far, there is an interesting second-order difference, especially at $M_{\text{tot}}^{\text{peak}} \sim 10^{12} M_{\odot}$. In this regime, the DM-only runs do produce a (small) population of galaxies that survive only as a very small remnant with mass below $10^{10} M_{\odot}$ or 1 per cent of their peak mass. This offset is particularly evident in the bottom panel, where the DM-only trends for both thresholds are almost flat, while they show a ≈ 50 per cent variation with $M_{\text{tot}}^{\text{peak}}$ in the hydrodynamical simulations. This suggests that baryons do have a non-negligible impact on mass stripping from satellites, but not on whether they ultimately survive as a (potentially very small) remnant.

3.1.2 Thresholds in stellar mass

In Fig. 6, we test similar thresholds in *stellar* mass in the hydrodynamical simulations, and also classify galaxies by their stellar peak mass M_{\star}^{peak} . In terms of absolute thresholds (green lines), the result is qualitatively consistent with our findings for total mass: surviving galaxies with $M_{\star}^{\text{peak}} \gtrsim 3 \times 10^9 M_{\odot}$ almost always retain a significant stellar remnant ($M_{\star}^{z=0} > 10^9 M_{\odot}$ or $> 0.5 M_{\star}^{\text{peak}}$ at

¹¹ A much lower threshold (e.g. $10^6 M_{\odot}$) would be numerically meaningless because our simulations could not possibly resolve such a small remnant. A higher threshold would not do justice to the resolution of our simulations.

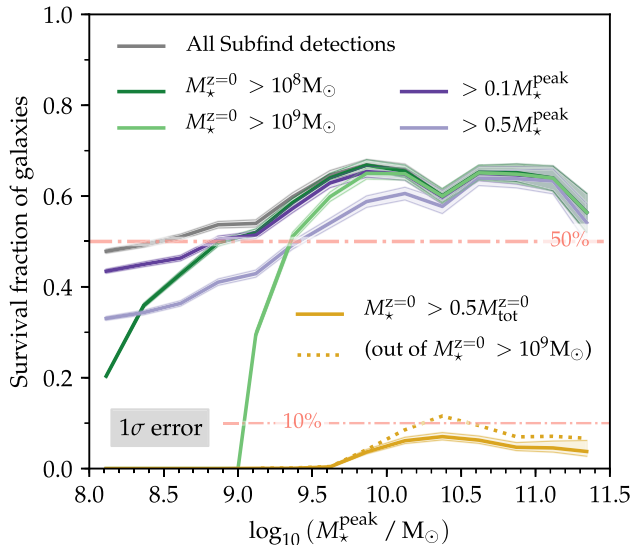


Figure 6. Dependence of the satellite survival fraction on *stellar* peak mass (M_{\star}^{peak}) and detection threshold in the hydrodynamical simulations. The grey line shows the fraction of galaxies at a given M_{\star}^{peak} that are detected by SUBFIND at $z = 0$, with binomial 1σ uncertainties marked by the shaded band. Dark and light green (purple) lines show the fraction whose stellar mass at $z = 0$ exceeds 10^8 and $10^9 M_{\odot}$ (10 and 50 per cent of M_{\star}^{peak}), respectively. The yellow lines, near the bottom of the plot, give the fraction of star-dominated survivors ($M_{\star}^{\text{peak}} > 0.5 M_{\text{tot}}^z$ at $z = 0$) out of all galaxies (solid) and only those surviving with $M_{\star}^z > 10^9 M_{\odot}$ (dotted). In the range plotted, virtually all surviving galaxies retain a resolved stellar remnant within 1 dex of their peak mass, but only a small subset are dominated by stars.

$z = 0$), but many lower-mass galaxies drop¹² below a threshold of 10^9 (and to a lesser extent also 10^8) M_{\odot} .

When considering relative thresholds, however (purple lines), it becomes clear that stellar mass loss from surviving satellites is considerably less severe than loss of total mass: even at $M_{\star}^{\text{peak}} = 10^8 M_{\odot}$, only a few per cent are reduced to less than one tenth of their peak stellar mass (compare the grey and dark purple lines), and such strong loss hardly occurs at all above $10^9 M_{\odot}$. Even only 50 per cent stellar mass loss is almost non-existent at the high-mass end ($M_{\star}^{\text{peak}} > 2 \times 10^{10} M_{\odot}$) and only affects less than half the surviving lowest-mass galaxies (compare the grey and light purple lines). This is consistent with the findings of Bahé et al. (2017a), who found a median stripped stellar mass fraction from surviving galaxies in groups and low-mass clusters of < 10 per cent, and with the works of Barber et al. (2016) and van Son et al. (2019), who demonstrate that (massive) galaxies that lost around 90 per cent of their initial stellar mass are extreme outliers from the relations between stellar mass and black hole mass or stellar size. *In terms of their stellar mass, satellite galaxy survival is therefore almost binary: they either retain a large part of it, or are lost completely.*

Also shown in Fig. 6 is the fraction of galaxies that survive in stellar mass dominated form (i.e. with $M_{\star}^{z=0} > 0.5 M_{\text{tot}}^{z=0}$; yellow solid line) and the analogous fraction out of only those that survive with $M_{\star}^{z=0} > 10^9 M_{\odot}$ (yellow dotted line). Both are small, with only the latter reaching ≈ 10 per cent at $M_{\star}^{\text{peak}} \sim 3 \times 10^{10} M_{\odot}$. Despite the much weaker loss of stellar than total mass, our simulations

¹²We note that this mass loss includes a contribution from stellar winds, in addition to stripping of stars through e.g. tidal forces.

therefore predict that the vast majority of surviving galaxies, at any mass, remain dominated by their non-stellar component. Qualitatively, this agrees with the conclusions of Dolag et al. (2009) based on lower-resolution simulations.

To summarize, the Hydrangea simulations predict that baryons have some impact on the mass loss of satellite galaxies, but are negligible with respect to their survival. The survival fraction is higher in more massive haloes – up to 67 per cent for Milky Way analogue galaxies in massive clusters – but still 44 per cent in low-mass groups. While many low-mass galaxies only survive as a small remnant with $M_{\text{tot}}^{z=0} < 10^{10} M_{\odot}$ – but often still within a factor of > 0.1 of their peak value in stellar mass – at $z = 0$, more massive galaxies with $M_{\text{tot}}^z > 3 \times 10^{11} M_{\odot}$ either disrupt completely, or retain a substantial core with $M_{\text{tot}}^{z=0} > 10^{10} M_{\odot}$ and $M_{\star}^{z=0} > 0.5 M_{\star}^{\text{peak}}$ at $z = 0$. Galaxies rarely survive in purely (or even mostly) stellar form.

4 INFLUENCE OF PRE-PROCESSING, OTHER SATELLITES, AND ACCRETION TIME

We now investigate different factors contributing to satellite disruption in more detail. The role of pre-processing (i.e. accretion on to a subgroup that is later accreted by their final host) is tested in Section 4.1, and that of satellite–satellite mergers in Section 4.2. We then show how the survival fraction depends on accretion redshift (Section 4.3) and time elapsed since accretion (Section 4.4), and conclude by investigating the distribution of galaxy disruption events over cosmic history (Section 4.5).

4.1 Role of pre-processing

4.1.1 Survival fractions of directly accreted and pre-processed galaxies

In Fig. 7, we repeat the survival analysis from Section 3 (Fig. 4), but this time we only consider galaxies that were not pre-processed, i.e. with $t_{\text{main}} = t_{\text{branch}}$. Different colours represent different host mass bins, and results from the hydrodynamical (DM-only) simulations are shown as solid (dashed) lines.

It is evident that the survival fraction amongst these ‘directly accreted’ galaxies is considerably higher than in the total population (cf. Fig. 4): in massive clusters (orange), it reaches ≈ 85 per cent even at the low-mass galaxy end ($M_{\text{tot}}^{\text{peak}} \sim 10^{10} M_{\odot}$), and peaks above 90 per cent at $M_{\text{tot}}^{\text{peak}} \sim 3 \times 10^{11} M_{\odot}$. Even for groups (black), the survival fraction of directly accreted galaxies exceeds 60 per cent, albeit only at $M_{\text{tot}}^{\text{peak}} < 10^{11} M_{\odot}$. This contrasts starkly with the survival fractions for pre-processed galaxies, which are shown – for clarity only for massive clusters in the hydrodynamical simulations – as dashed lines in Fig. 7 and lie in the range of ≈ 40 –60 per cent. Pre-processing is evidently much more disruptive than the final host environment, consistent with the trend towards lower survival fractions in lower-mass (final) haloes.

Similar to the total satellite population, the survival fractions of directly accreted galaxies agree closely between DM-only and hydrodynamical simulations. The survival fraction of *all* galaxies accreted by their final host is only ≈ 10 per cent lower than for their directly accreted subset, as shown for massive clusters by the orange dash-dotted line in Fig. 4. Even higher is the survival fraction of only those galaxies that were a central immediately prior to t_{main} (irrespective of whether they were previously pre-processed, not shown). As we demonstrate below, this is because most disruption of pre-processed galaxies occurs outside of their final host halo.

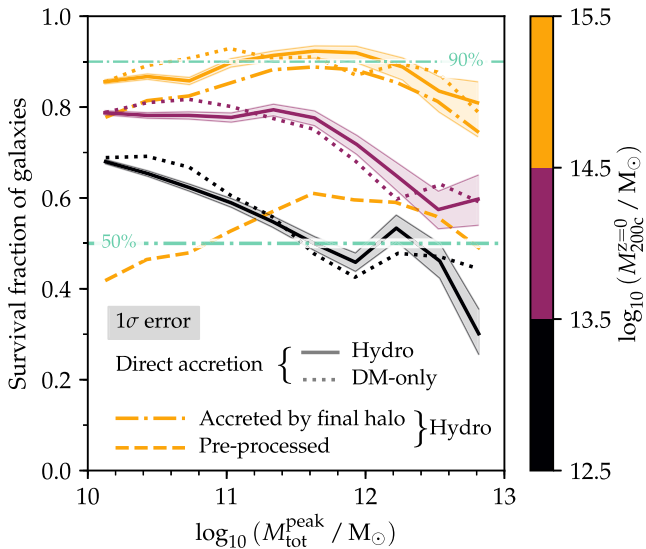


Figure 7. Survival fraction of galaxies ($M_{\text{tot}}^{z=0} > 5 \times 10^8 M_{\odot}$) that were directly accreted on to their final host, in the hydrodynamical (solid lines, shaded bands indicate binomial 1σ errors) and DM-only simulations (dotted). For comparison, the survival fractions of all galaxies that reach their main host, and of pre-processed galaxies, are shown as dash-dotted and dashed lines, respectively; for clarity, we only show these for massive clusters ($M_{200c}^{z=0} > 10^{14.5} M_{\odot}$) in the hydrodynamical simulations. Survival is more common for galaxies that are not pre-processed.

Massive clusters in particular therefore preserve a near-complete ‘fossil record’ of all galaxies with $M_{\text{tot}}^{\text{peak}} > 10^{10} M_{\odot}$ that have ever orbited within them. Keeping in mind that simulations may also disrupt satellite galaxies for numerical, rather than physical, reasons (van den Bosch & Ogiya 2018), the true survival fractions may, in principle, be even higher than what is shown in Fig. 7. To test this, we compute in Appendix C the fraction of surviving remnants that are numerically unreliable according to the criteria of van den Bosch & Ogiya (2018). Amongst massive galaxies ($M_{\text{tot}}^{\text{peak}} > 3 \times 10^{11} M_{\odot}$), numerically unreliable remnants are rare ($\lesssim 1$ per cent) in our simulations, but at $M_{\text{tot}}^{\text{peak}} \sim 10^{10} M_{\odot}$, up to one-third of remnants may be unreliable. The survival fractions shown in Fig. 7, however, are not consistent with significant numerical disruption of low-mass satellites (e.g. they depend only weakly on galaxy mass). This suggests that *numerical disruption of satellites is not common in our simulations*, at least at $M_{\text{tot}}^{\text{peak}} > 10^{10} M_{\odot}$.

4.1.2 Where are pre-processed galaxies disrupted?

Pre-processed galaxies can be disrupted either in their subgroup (prior to t_{main}), or later in their final host. In Fig. 8, we disentangle these two scenarios, for simplicity combining all hosts with $M_{200c}^{z=0} > 10^{12.5} M_{\odot}$ into a single bin (we have verified that differences between different host masses are small). Different lines show the fractional contribution of different merger types to the disruption of pre-processed galaxies. Clearly dominant (≈ 50 – 80 per cent, highest at lowest $M_{\text{tot}}^{\text{peak}}$) are mergers with the pre-processing host (black solid line), i.e. those that merged prior to t_{main} with a galaxy that was previously the disrupted galaxy’s central.

In addition, the next most common disruption route is also due to the pre-processing host, but only after it became itself a satellite of the (final) host halo (purple dash-dotted line). Although these are technically mergers between two satellites in the final halo, it

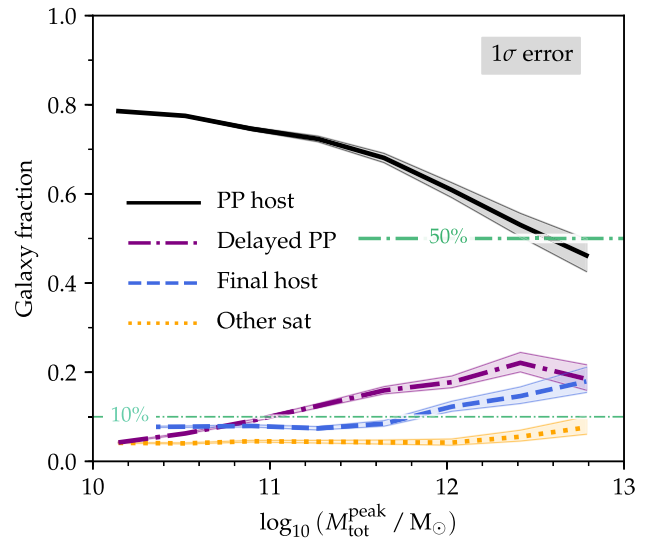


Figure 8. Merger routes of disrupted pre-processed satellite galaxies as a function of $M_{\text{tot}}^{\text{peak}}$ (all three host mass bins combined). The black solid line shows the fraction that merges with their pre-processing host before infall into their final halo (t_{main}). The purple dash-dotted and blue dashed lines show the fractions that undergo, after t_{main} , a ‘delayed’ merger with their pre-processing host, and a merger with their final host, respectively. The orange dotted line represents mergers with another satellite, almost all of which occur during pre-processing. Shaded bands give binomial 1σ uncertainties. The vast majority merge with their pre-processing host, either before or after reaching the final halo.

is more appropriate to consider them as a case of ‘delayed pre-processing’, since the infalling subgroup may retain its physical identity for some time after having been subsumed into its host. Including these, pre-processing hosts account for ≈ 70 per cent of all disruption of pre-processed galaxies, at all masses we probe. The remaining galaxies merge with their final host (< 10 per cent at $M_{\text{tot}}^{\text{peak}} < 10^{12} M_{\odot}$, dashed blue line) or, even less commonly, with another unrelated satellite (orange dotted line), mostly during pre-processing. This is consistent with the recent study of Han et al. (2018), who inferred from a different set of simulations that pre-processing has a decisive impact on mass stripping from infalling galaxies, in particular when the mass ratio between galaxy and pre-processing host is low.

4.1.3 The contribution of pre-processing to galaxy disruption

To conclude our investigation of pre-processing, we show in Fig. 9 the fraction of all satellite disruption that is due to pre-processing (including delayed mergers and mergers with other satellites prior to t_{main}), as a function of $M_{\text{tot}}^{\text{peak}}$ and $M_{200c}^{z=0}$. The combination of a higher pre-processed fraction at lower $M_{\text{tot}}^{\text{peak}}$ and higher $M_{200c}^{z=0}$ (Fig. 3), and their much lower survival fraction compared to directly accreted galaxies (Fig. 7) implies that the vast majority, ≈ 80 – 90 per cent, of all disruption at $M_{\text{tot}}^{\text{peak}} < 10^{12} M_{\odot}$ in massive clusters is the result of pre-processing. The fraction decreases somewhat towards higher masses, but pre-processing still accounts for ≈ 70 per cent of all disruption even at $M_{\text{tot}}^{\text{peak}} = 10^{13} M_{\odot}$. In lower-mass haloes, pre-processing is overall much less important, and only accounts for ≈ 20 per cent of the disruption of Milky Way analogues in groups.

The DM-only simulations broadly agree with the hydrodynamical runs, but generally predict a slightly lower fraction of disruption that is due to pre-processing (by ≈ 5 per cent) and a slightly smoother




Figure 9. The fraction of all satellite disruption that is due to pre-processing (including delayed mergers), in the hydrodynamical simulations (solid lines) and their DM-only counterparts (dotted). Different host masses are represented by different colours (see the colour bar on the right). Pre-processing is by far the dominant cause of disruption in cluster galaxies with $M_{\text{tot}}^{\text{peak}} < 10^{12} M_{\odot}$, but it becomes much less relevant for more massive galaxies and those in lower-mass hosts.

transition from the flat part at low $M_{\text{tot}}^{\text{peak}}$ to the decline at high mass (especially in clusters). This suggests that baryons have a (small) disruptive effect in situations where the mass contrast between the satellite and host is not too large; we investigate this further in Paper II.

To summarize, we have found that *pre-processing plays a crucial role in disrupting galaxies*, particularly in clusters where it accounts for the vast majority of all disruption (≈ 90 per cent at $M_{\text{tot}}^{\text{peak}} \sim 10^{12} M_{\odot}$ and $M_{200c}^{z=0} > 10^{14.5} M_{\odot}$). Galaxies accreted directly on to their final host survive to ≈ 85 per cent in massive clusters, and still ≈ 80 per cent in lower-mass clusters at $M_{\text{tot}}^{\text{peak}} \leq 3 \times 10^{11} M_{\odot}$. Pre-processing disruption mostly involves mergers with the central galaxy of the subgroup. The lowest-mass haloes are therefore the most efficient in disrupting satellites at fixed $M_{\text{tot}}^{\text{peak}}$, plausibly as a consequence of dynamical friction, while massive galaxy clusters should preserve a near-complete record of all galaxies (at least with $M_{\text{tot}}^{\text{peak}} > 10^{10} M_{\odot}$) that they have ever accreted.

4.2 Role of satellite–satellite mergers

We had noted above that satellite–satellite mergers are rather uncommon for pre-processed galaxies. Their role in the (final) host haloes themselves is explored in Fig. 10, where we show the fraction of all disruption events amongst directly accreted galaxies ($t_{\text{branch}} = t_{\text{main}}$) that are due to mergers with other satellite galaxies. We exclude cases where this other satellite was previously the galaxy’s central (due to central–satellite swaps, which is only relevant for massive galaxies in groups).

The key feature is that satellite–satellite mergers in massive haloes are extremely rare; note that the y-axis range is reduced to $[0, 0.1]$ in order to highlight any deviations from zero at all. At $M_{\text{tot}}^{\text{peak}} < 10^{12} M_{\odot}$, they account for less than 1 per cent of disruption events in massive clusters, and still $\lesssim 3$ per cent in groups. Only amongst the most massive galaxies are they slightly more




Figure 10. The fraction of non-surviving directly accreted satellites that are disrupted by mergers with other satellites. Hydrodynamical simulations are represented by solid lines (with shaded bands indicating binomial 1σ uncertainties), DM-only runs by dotted lines. For clarity, the y-axis range is reduced compared to the other plots. Satellite–satellite mergers are very uncommon (particularly in massive clusters): only at the highest masses ($M_{\text{tot}}^{\text{peak}} \gtrsim 10^{12} M_{\odot}$) do they account for ≈ 10 per cent of disruption events.

relevant, with fractions of up to 10 per cent in low-mass clusters at $M_{\text{tot}}^{\text{peak}} \approx 10^{13} M_{\odot}$. What disruption occurs in massive haloes (see above) is therefore almost exclusively due to mergers with the central galaxy (including dispersal into its halo, as we test in Paper II). We note that interactions between satellites may nevertheless contribute significantly to their mass loss and ultimate dispersal (see e.g. Moore et al. 1996; Marasco et al. 2016). At all galaxy and host masses that we consider, the predictions from hydrodynamical and DM-only simulations agree to within the statistical uncertainties, which rules out a significant impact of baryon physics on this merger channel.

4.3 Evolution of surviving fraction with accretion time

We now examine the influence of accretion time on galaxy survival. For ease of interpretation, we focus here on directly accreted galaxies. In Fig. 11, galaxies are split into three host mass bins (three different panels, $M_{200c}^{z=0}$ increasing from left to right) and six bins in galaxy peak mass $M_{\text{tot}}^{\text{peak}}$ (different coloured lines, increasing from purple to green). Each line traces the fraction of galaxies that survive (with $M_{\text{tot}}^{z=0} > 5 \times 10^8 M_{\odot}$) as a function of accretion time (t_{acc} , or equivalently redshift (z_{acc})). For clarity, only the hydrodynamical simulations are shown, but we have verified that the DM-only runs give very similar results.

The dominant trend of all lines in Fig. 11 is that galaxies that were accreted later are more likely to survive to $z = 0$, in agreement with previous work (e.g. De Lucia et al. 2004; Gao et al. 2004). At $z_{\text{acc}} \approx 0$ the survival fraction approaches unity, as should be expected. The few per cent of galaxies that were accreted very early, on the other hand ($z_{\text{acc}} \gtrsim 4$, see Fig. 2), almost never survive to $z = 0$.

Within each bin of host and galaxy mass (individual lines in Fig. 11), the survival fraction always transitions quite rapidly from ~ 0 to ~ 1 , over a period of typically only a few Gyr. The accretion time (measured from the big bang) at which the survival fraction reaches 50 per cent (t_{trans}) depends in general on both $M_{\text{tot}}^{\text{peak}}$ and




Figure 11. The fraction of non-pre-processed galaxies accreted at a given redshift (z_{acc}) that survive to $z=0$ (with $M_{\text{tot}}^{z=0} > 5 \times 10^8 M_{\odot}$), in the hydrodynamical simulations. Only bins with at least 10 galaxies are shown, which is why not all lines extend to the earliest accretion times. Different panels show different host mass ranges, as indicated in the bottom-right corners. In all three panels, the survival fraction of low-mass galaxies approaches unity for galaxies accreted at $z_{\text{acc}} \lesssim 1$, and then drops rapidly towards earlier accretion times. More massive galaxies, and those in less massive groups, are still disrupted at lower z_{acc} .

$M_{200c}^{z=0}$. In low-mass groups (left-hand panel), $t_{\text{trans}} \approx 2.5$ Gyr ($z \approx 2$) for the lowest-mass galaxies (purple) and then increases fairly gradually to $t_{\text{trans}} \approx 7.5$ Gyr ($z \approx 0.6$) at $M_{\text{tot}}^{\text{peak}} > 10^{12.5} M_{\odot}$ (yellow-green). While the lowest-mass galaxies therefore already survive to 90 per cent at $z_{\text{acc}} = 1.3$, those with the highest masses only reach this point at $z_{\text{acc}} = 0.3$.

The dependence of t_{trans} on galaxy mass is noticeably less strong in more massive hosts. In low-mass clusters ($M_{200c}^{z=0} \sim 10^{14} M_{\odot}$; middle panel of Fig. 11) the lowest-mass galaxies follow almost exactly the same trend as in groups, but not until $M_{\text{tot}}^{\text{peak}} = 10^{12} M_{\odot}$ is there a noticeable shift towards later t_{trans} . Consequently, even the most massive galaxies reach 50 (90) per cent survival already at $z_{\text{acc}} = 0.9$ ($z_{\text{acc}} = 0.6$).

In massive clusters (right-hand panel), any differences with $M_{\text{tot}}^{\text{peak}}$ are very small, but there is a slight shift towards *even earlier* t_{trans} with increasing galaxy mass, at least for those bins where our simulations contain enough galaxies to identify t_{trans} . This shift may reflect the enhanced ability of more massive galaxies to withstand tidal stripping, while their mass is still so far below that of the host cluster that e.g. dynamical friction does not cause accelerated disruption in the same way as in lower-mass hosts. Milky Way analogues ($M_{\text{tot}}^{\text{peak}} \sim 10^{12} M_{\odot}$) therefore reach 90 per cent survival already at $z_{\text{acc}} = 2.0$ and 96 per cent of all galaxies with $M_{\text{tot}}^{\text{peak}} > 10^{10} M_{\odot}$ and $z_{\text{acc}} < 2$ survive at $z=0$. The small fraction of galaxies that are disrupted in massive clusters are therefore predominantly those that were accreted the earliest.

4.4 From accretion to disruption: rapid, delayed, or continuous?

A natural question to ask is whether the relatively rapid transition from disruption- to survival-dominated accretion redshifts is indicative of a long, mass-dependent delay between accretion and disruption. In other words, galaxies accreted just after t_{trans} may survive at $z=0$ because they have (just) not been a satellite for long enough, while those accreted just before could have been disrupted very recently. We now demonstrate that such a delay time argument

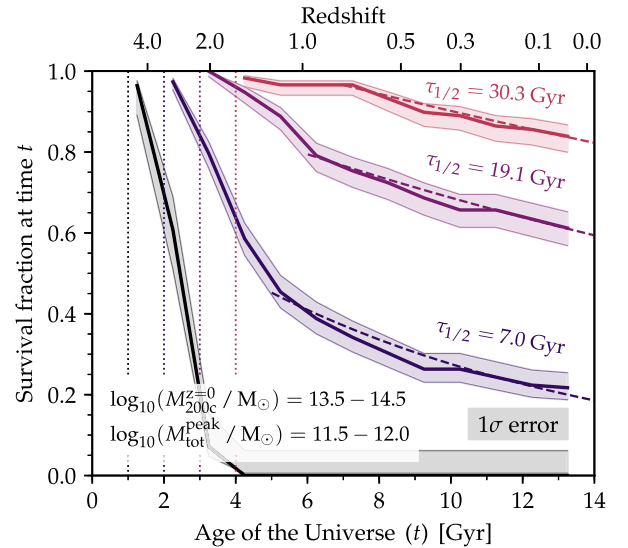


Figure 12. Connection between accretion and disruption time for one bin in galaxy and host mass (see bottom-left corner) in the hydrodynamical simulations. We select galaxies that were accreted within four intervals of $\Delta t = 500$ Myr, centred on the times indicated by vertical dotted lines, and show their survival fraction (above a total mass threshold of $5 \times 10^8 M_{\odot}$) as a function of time. Thin dashed lines represent the best-fitting late-time exponential decay model, with corresponding half-lifetime $\tau_{1/2}$ as indicated on the right. The earliest accreted galaxies have all disrupted rapidly after accretion. Later generations show a progressively shallower decline in survival fraction, with $\tau_{1/2} \gg t_{\text{Hubble}}$ at $z_{\text{acc}} < 2$.

cannot be invoked as the reason for the lower survival fraction of early accreted galaxies.

For this purpose, Fig. 12 shows the survival fraction of galaxies as a function of cosmic time t , i.e. the fraction with $M_{\text{tot}}(t) > 5 \times 10^8 M_{\odot}$. We select galaxies that were not pre-processed in four bins of $\Delta t_{\text{acc}} = 500$ Myr, with centres indicated by the vertical dotted lines. For clarity, we focus on only one bin in galaxy mass ($M_{\text{tot}}^{\text{peak}} = 10^{11.5} - 10^{12.0} M_{\odot}$) and host mass (low-mass clusters) in

the hydrodynamical simulations. For each bin in accretion time, the correspondingly coloured solid line shows the fraction of galaxies still alive at time t , and the bands the corresponding 1σ binomial uncertainties.

It is immediately evident that there is no universally long delay between accretion and disruption, particularly at high z_{acc} (black/indigo). The disruption rate (i.e. the line slope) is greatest within the first few Gyr after accretion and then flattens off. In the earliest accretion bin ($z_{\text{acc}} > 4$; black), all galaxies are disrupted within 3 Gyr of accretion, while a successively higher fraction of galaxies accreted later survive at least this long. At $t > t_{\text{acc}} + 3$ Gyr, the survival fraction decays approximately exponentially with t . The best fits are given by the dashed lines, with a systematically increasing half-life time $\tau_{1/2}$ for lower z_{acc} . At $z_{\text{acc}} < 2$, $\tau_{1/2}$ exceeds (significantly) the available time until $z = 0$, which naturally explains why most of these galaxies survive until today.

The strong dependence of the survival fraction on accretion redshift (Fig. 11) is therefore the result of the disruption efficiency decreasing (strongly) with time. It is conceivable that this reflects the lower host halo masses at higher z_{acc} , but we have verified that our results are not markedly changed when galaxies are instead binned by their host mass *at accretion*, as long as it remains¹³ $\gtrsim 1$ dex above $M_{\text{tot}}^{\text{peak}}$. Instead, the fact that the half-life times shown in Fig. 12 scale with accretion redshift approximately as $(1 + z_{\text{acc}})^{-3/2}$ – the expected scaling of the dynamical time with redshift (McGee, Bower & Balogh 2014) – suggests that the low survival fraction of early accreted galaxies is due to different orbital conditions imprinted at accretion. In Paper II, we show that early accreted galaxies lose mass more rapidly because they have (much) shorter orbital periods, while massive galaxies are more strongly dragged towards the host centre at high z_{acc} and can therefore merge more efficiently with the (growing) central cluster galaxy.

4.5 Galaxy disruption times

We have so far only distinguished galaxies by their accretion times, but a related question of interest – particularly for connection with observational work – is when galaxies actually disrupt. This is shown in Fig. 13, which gives the cumulative fraction of (non-surviving) galaxies that were disrupted (i.e. fell below our mass threshold of $5 \times 10^8 M_{\odot}$) prior to a given time t_{disrupt} . The three bins in host mass are represented by differently coloured lines; two bins in galaxy mass are distinguished by different line styles. As in Fig. 2, all times are offset by a random value of up to ± 250 Myr to suppress artificial discreteness due to the limited number of snapshots.

The distribution is qualitatively similar to that of the branch accretion times shown in the top panel of Fig. 2, consistent with the picture that most galaxies are disrupted during pre-processing, soon after first accretion. In line with their lower accretion redshifts, more massive galaxies (solid lines) are disrupted slightly later, with median disruption redshifts of $z = 2$ and ≈ 1 for $M_{\text{tot}}^{\text{peak}} = 10^{10}$ – 10^{11} and $10^{11.5}$ – $10^{12.5} M_{\odot}$, respectively. The lower-mass bin shows no dependence of t_{disrupt} on $M_{200c}^{z=0}$ at all, but there is a slight tendency towards later disruption in groups than clusters for more massive galaxies (by $\lesssim 2$ Gyr), consistent with the equivalent trends in t_{branch} . Due to the typically long delay between accretion and disruption for most galaxies accreted at intermediate and low redshifts (Fig. 12),

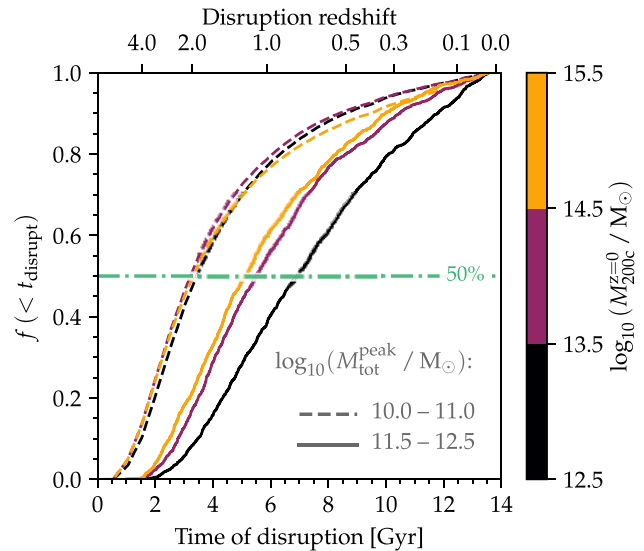


Figure 13. Distribution of galaxy disruption times. Shown is the fraction of (ultimately disrupted) galaxies that fall below the survival threshold of $5 \times 10^8 M_{\odot}$ prior to time t_{disrupt} , in two bins of peak galaxy mass (different line styles) and three host mass bins (different colours). No selection is made with regard to pre-processing. Disruption was most prevalent at redshift $z \approx 1$ – 3 , but with a broad tail extending to $z = 0$. Massive galaxies were disrupted later, but the influence of host mass is small.

disruption is still prevalent in the low-redshift Universe, in particular amongst massive galaxies ($M_{\text{tot}}^{\text{peak}} \sim 10^{12} M_{\odot}$), for which ≈ 10 per cent of all disruption events occur at $z < 0.3$.

5 BIASES BETWEEN SURVIVING AND DISRUPTED GALAXIES

For the final part of our analysis we test whether there are any differences in pre-infall properties between galaxies that are disrupted and those that survive, at fixed (total) $M_{\text{tot}}^{\text{peak}}$. Such differences could cause subtle biases between (surviving) cluster and field galaxies without any actual galaxy transformation process. To pre-empt the answer, we did not find any strong differences of this kind in terms of either the baryonic or dark matter properties of galaxies – at least those accreted around $z \approx 2$ – and can therefore rule out such ‘differential disruption’ as a significant contributor to the observed differences between field and cluster galaxies in the local Universe.

A complication in comparing the pre-infall properties of disrupted and surviving galaxies is that, as we have found above (Fig. 11), disrupted galaxies were preferentially accreted earlier than survivors. Because the relations of e.g. stellar mass and star formation rate with halo mass evolve with redshift (see e.g. Furlong et al. 2015 and references therein), a comparison between *all* disrupted and surviving galaxies would show strong differences that are purely the result of this redshift bias. A meaningful comparison is therefore only possible between galaxies with similar accretion redshift z_{acc} and furthermore – due to the finite number of galaxies in our simulation – only around the z_{acc} where the survival and disruption fractions are comparable (i.e. $z_{\text{acc}} \approx 2$).

In the top panel of Fig. 14, we show the gas mass (green), star formation rate (blue), and stellar mass (purple) in the snapshot before the main accretion time for disrupted galaxies that were directly accreted between $z_{\text{acc}} = 1.5$ and 2.5 . The DM half-mass radius (black), stellar half-mass radius (red), and maximum circular

¹³In other words, excluding situations better described as minor or major galaxy mergers, rather than accretion of satellites.

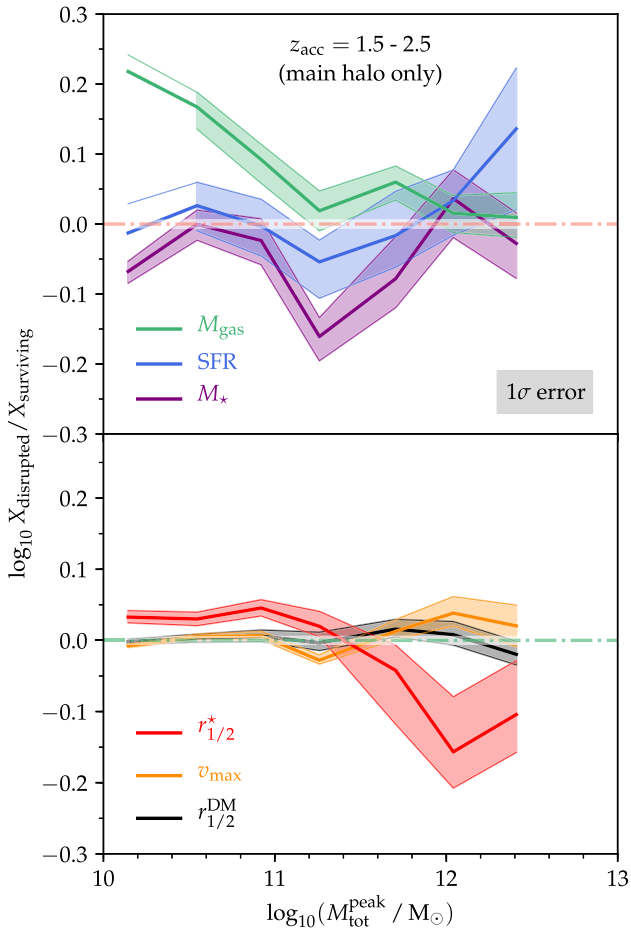


Figure 14. *Top panel:* bias of disrupted galaxies in baryonic properties prior to accretion, compared to surviving galaxies (at a mass threshold of $5 \times 10^8 M_\odot$ for survival). Stellar mass is shown in purple, star formation rate in blue, and gas mass in green. Only galaxies directly falling into their final host between $z = 2.5$ and 1.5 are shown to avoid indirect biases due to different accretion times. Only bins with at least 10 surviving and disrupted galaxies are shown. At fixed $M_{\text{tot}}^{\text{peak}}$, disrupted galaxies had slightly higher gas mass, and slightly lower stellar mass, than surviving galaxies. *Bottom panel:* the same for DM half-mass radius (black), stellar half-mass radius (red), and maximum circular velocity (orange). Out of these, only the stellar half-mass radius shows a significant (but relatively small) difference between disrupted and surviving galaxies.

velocity (orange) are compared in the bottom panel. All values are normalized to their analogues for surviving galaxies: we compute the median and 1σ uncertainty for disrupted and surviving galaxies as a function of $M_{\text{tot}}^{\text{peak}}$ and then plot their logarithmic ratio. The 1σ errors shown as shaded bands are here computed as the difference between the median and the 16th/84th percentiles, divided by \sqrt{N} where N is the number of galaxies per bin. Note that for SFR and M_{gas} , the 16th percentile is equal to zero in the lowest-mass bin, so that we cannot compute a meaningful (logarithmic) lower error boundary.

The key feature of Fig. 14 is the absence of any clear, strong differences between disrupted and surviving galaxies. There is a mildly significant negative bias in stellar mass, i.e. in the sense that disrupted galaxies contained less stellar mass prior to accretion than equally massive surviving galaxies, but only by $\lesssim 0.1$ dex. Similarly, there is a mild positive bias in gas mass, at least for low-mass galaxies ($M_{\text{tot}}^{\text{peak}} < 10^{11} M_\odot$). This is consistent with a picture

in which there are (small) individual effects of gas stripping (e.g. Saro et al. 2008) and (past) star formation (e.g. Weinberg et al. 2008), which largely cancel each other on average.

The pre-accretion stellar half-mass radius of disrupted low-mass galaxies ($M_{\text{tot}}^{\text{peak}} \lesssim 3 \times 10^{11} M_\odot$) is marginally (but significantly) larger than for surviving galaxies, consistent with the expectation that less compact galaxies are more susceptible to tidal stripping. Interestingly, our simulations predict the opposite trend for more massive galaxies, where disrupted galaxies were, on average, slightly more compact prior to accretion. This is further evidence for two different disruption channels for low- and high-mass galaxies, as we discuss in Paper II. No consistent and significant difference is seen for SFR, DM half-mass radius ($r_{1/2}^{\text{DM}}$), or maximum circular velocity (v_{max}).

The key implication is that *whether a galaxy survives or not depends at best weakly on its internal properties*. This fits in with our earlier conclusion that the survival fraction is similar between DM-only and hydrodynamical simulations, and does not depend strongly on galaxy mass. The almost an order of magnitude higher stellar mass fractions of (surviving) satellites in clusters (Bahé et al. 2017b), in particular those accreted early (Armitage et al. 2018), are therefore predominantly a consequence of dark matter being stripped more efficiently than stars from satellite galaxies (see Section 3.1), rather than star formation enhancing the likelihood of survival. We caution, however, that we could only do this test for galaxies in a relatively narrow and high range of accretion redshifts. A significantly larger cluster sample would be required to check whether the same conclusion holds for galaxies that were accreted later.

Finally, we note that we have also considered the equivalent biases for pre-processed galaxies (with quantities calculated at t_{branch} , not shown). Most features are qualitatively consistent with Fig. 14, but there appears to be a stronger negative bias in stellar mass (approximately -0.15 dex), and a small but significant negative bias in $r_{1/2}^{\text{DM}}$ (approximately -0.06 dex) for disrupted pre-processed galaxies. This may, however, simply be a manifestation of indirect bias due to large-scale environmental influence of the host.¹⁴ Again, we would require a larger simulation volume to control for this indirect effect and test whether internal galaxy properties are causally connected to survival in the pre-processing phase.

6 SUMMARY AND DISCUSSION

We have investigated the disruption of galaxies in groups and clusters with the aid of the Hydrangea simulations, a suite of cosmological, hydrodynamical/ N -body zoom-in simulations of 24 galaxy clusters and their large-scale environments. From the evolutionary histories of individual simulated galaxies with a peak (i.e. maximum past) total (baryons plus dark matter) mass of $M_{\text{tot}}^{\text{peak}} > 10^{10} M_\odot$ – corresponding to a peak stellar mass $M_\star^{\text{peak}} \gtrsim 5 \times 10^7 M_\odot$ – that we have computed with an updated tracing procedure, we have searched for galaxies that were accreted by a group/cluster in the past and identified those as ‘surviving’ that still correspond to

¹⁴Galaxies whose pre-processing begins closer to their final host have a higher chance of survival (due to the shorter time before their pre-processing host is itself accreted) and are more strongly affected by large-scale environmental influence of their final host. Directly accreted galaxies are not subject to this bias, because their hosts affected all of them approximately equally at the point of accretion.

distinct subhaloes with total mass above $5 \times 10^8 M_\odot$ at $z = 0$. Our main conclusions may be summarized as follows:

(i) Averaged over the entire history of the Universe, our simulations predict that 47 per cent of all satellite galaxies with peak total mass $M_{\text{tot}}^{\text{peak}} \geq 10^{10} M_\odot$ that were accreted on to groups or clusters ($M_{200c}^{z=0} > 10^{12.5} M_\odot$) survive to the present day. The survival fraction increases somewhat with host halo mass and is rather insensitive to galaxy mass. The fraction is highest (67 per cent) for galaxies with $M_{\text{tot}}^{\text{peak}} \sim 10^{12} M_\odot$ in massive clusters, and differs only marginally ($\lesssim 5$ per cent) between simulations with and without baryons (Fig. 4).

(ii) Many surviving galaxies have lost a large fraction of their $M_{\text{tot}}^{\text{peak}}$ by $z = 0$, and may therefore not be counted as surviving with higher mass thresholds and/or in lower-resolution simulations. However, hardly any galaxies in the hydrodynamical simulations survive with less than 1 per cent of their $M_{\text{tot}}^{\text{peak}}$, even where such a remnant would be well resolved (Fig. 5). Hence, once a galaxy has lost $\gg 90$ per cent of its peak total mass, its chance of survival is very small.

(iii) Stellar mass loss from surviving galaxies is less severe than total mass loss. Even at very low peak stellar masses ($M_\star^{\text{peak}} \sim 10^8 M_\odot$) and including mass loss from stellar evolution, only a few per cent of galaxies survive with less than one tenth of their peak stellar mass. At $M_\star^{\text{peak}} > 10^{10} M_\odot$, even 50 per cent stellar mass loss is rare. In terms of stellar mass, survival is therefore almost binary: either a significant fraction is retained, or the galaxy is lost completely. Nevertheless, only $\lesssim 10$ per cent of surviving galaxies are stellar-mass dominated at $z = 0$, even at the most favourable $M_\star^{\text{peak}} \approx 2 \times 10^{10} M_\odot$ (Fig. 6).

(iv) Most galaxy disruption in clusters, and at $10^{10} M_\odot \leq M_{\text{tot}}^{\text{peak}} < 10^{11} M_\odot$ also in groups, occurs during pre-processing. At $M_{\text{tot}}^{\text{peak}} \sim 10^{12} M_\odot$, 90 per cent of all disrupted galaxies in massive clusters ($M_{200c}^{z=0} > 10^{14.5} M_\odot$) were pre-processed (Figs 8 and 9). The survival fraction of galaxies that were directly accreted by their final host is as high as ≈ 90 per cent (at $M_{\text{tot}}^{\text{peak}} = 10^{11.5} M_\odot$ in a massive cluster), with only per cent-level variations between hydrodynamical and DM-only simulations (Fig. 7). The most massive host haloes are therefore the least efficient in disrupting satellites of a given mass, and vice versa.

(v) The survival fraction of satellite galaxies depends strongly and non-linearly on their accretion redshift (z_{acc}). In massive clusters, and at $M_{\text{tot}}^{\text{peak}} < 10^{11} M_\odot$ even in low-mass groups, ≈ 95 per cent of non-pre-processed galaxies with $z_{\text{acc}} \leq 1$ survive to $z = 0$. Towards higher z_{acc} , the survival fraction drops steeply and universally becomes negligible for $z_{\text{acc}} \gtrsim 4$. Below the scale of massive clusters ($M_{200c}^{z=0} < 10^{14.5} M_\odot$), this transition from low (< 10 per cent) to high (> 90 per cent) survival fractions occurs at lower z_{acc} for galaxies with higher $M_{\text{tot}}^{\text{peak}}$ and those in lower-mass hosts: at fixed z_{acc} and $M_{200c}^{z=0}$, the lowest-mass galaxies are therefore the most likely to survive (Fig. 11). This redshift dependence is the result of a strong evolution in the disruption efficiency with z_{acc} , rather than reflecting a uniformly long delay time between accretion and disruption (Fig. 12).

(vi) The disruption of galaxies continues until $z = 0$. Half of all non-surviving galaxies with $M_{\text{tot}}^{\text{peak}} \sim 10^{12} M_\odot$ are disrupted after $z = 1$ (including during pre-processing); for clusters ($M_{200c}^{z=0} > 10^{13.5} M_\odot$), 10 per cent of these disruption events occur after $z = 0.3$ (Fig. 13).

(vii) The survival of galaxies is not strongly correlated with their internal properties before accretion, at least for those accreted in the interval $1.5 \leq z_{\text{acc}} \leq 2.5$. Compared to survivors with the same

$M_{\text{tot}}^{\text{peak}}$, disrupted galaxies contained only slightly more gas (≈ 0.2 dex) and slightly less stellar mass (≈ 0.1 dex). Stellar half-mass radii show a slight, mass-dependent bias between disrupted and surviving galaxies; star formation rate, maximum circular velocity, and dark matter half-mass radius display no significant offsets. The observed differences between cluster and field galaxies at $z \approx 0$ are therefore unlikely the result of biased survival amongst the former (Fig. 14).

According to these findings, the disruption of satellite galaxies is not a ubiquitous feature of cosmological galaxy cluster simulations, at least not at $M_{\text{tot}}^{\text{peak}} > 10^{10} M_\odot$ and at the relatively high mass resolution of Hydrangea ($\sim 10^6$ and $10^7 M_\odot$ for baryons and DM, respectively). In contrast to recent predictions from idealized N -body experiments (van den Bosch & Ogiya 2018), galaxies with the lowest (peak) mass are in fact the most likely ones to survive to $z = 0$ at any z_{acc} and $M_{200c}^{z=0}$ (with the possible exception of the most massive clusters, where galaxies of all masses we have considered display a similarly high survival fraction). The mass range that we have probed extends well below the scale at which baryonic properties of galaxies become affected by poor resolution ($M_{\text{tot}}^{\text{peak}} \sim 10^{11} M_\odot$; Schaye et al. 2015). This suggests that artificial disruption of satellites is not a major roadblock for cosmological hydrodynamical simulations.

Although massive galaxy clusters give rise to strong tidal and ram pressure forces, our simulations predict that this is in fact the environment in which the *smallest* fraction of satellite galaxies are destroyed. Instead, they should contain a near-complete ‘fossil record’ of all galaxies that have ever orbited within them, whereas $\approx 1/3$ of satellites in low-mass groups are disrupted before $z = 0$. Despite their rarity, massive clusters therefore constitute a valuable laboratory to study the effect of environmentally induced galaxy transformations over time. These findings are consistent with the observational detection of an upturn in the satellite luminosity function at the faint end in clusters (e.g. Lan, Ménard & Mo 2016), which suggests that low-mass satellites are indeed able to survive and accumulate in massive haloes.

There are two regimes where our simulations do predict a significant fraction of satellites to be disrupted: pre-processing in lower-mass groups, which then later assemble into a more massive group or cluster, and satellites accreted at high redshift, where disruption was evidently much more widespread, and more swift, than in the present-day Universe. This agrees with the observational evidence for widespread (dwarf) galaxy disruption during the early stages of cluster formation (López-Cruz et al. 1997). We defer further exploration of these trends to a follow-up paper, where we show that they are the consequence of enhanced mergers between satellite and central galaxies, and a strong evolution of the orbital time-scale of galaxies, with increasing z_{acc} (see also Han et al. 2018). Both effects highlight the impact of the cosmological environment of groups and clusters on the predicted evolution of their member galaxies.

Our simulations suggest that the role of baryons in determining the survival of satellite galaxies – but not the degree of stripping they experience – is small, which is important in two ways. First, it rules out ‘biased survival’ as a significant contributor to the environmental differences that are observed in the local Universe: in principle, e.g. the relative overabundance of red, quenched galaxies could also have stemmed from a preferential disruption of their blue, star-forming cousins. Our findings therefore corroborate the hypothesis that these differences are the result of individual galaxies being transformed by their environment, through processes such as ram-pressure stripping, strangulation, or tidal stripping.

Secondly, the small impact of baryons implies that pure N -body simulations can, at least in principle, predict the survival of galaxies with reasonable accuracy.

We finally emphasize that negligible total disruption of satellites in massive clusters, as predicted by our study, is not incompatible with (significant) mass loss from surviving satellites. Indeed, we have shown that many low-mass galaxies only survive as small remnants with total (but typically not stellar) mass well below their peak values. In future work, we will investigate in more detail how this mass loss is connected to the build-up and growth of central group and cluster galaxies, and of their extended dark matter and stellar haloes.

ACKNOWLEDGEMENTS

We thank the reviewer for their report, which improved the presentation of the results in this paper. We thank Lydia Heck for expert computational support with the Cosma machine in Durham, which was used for part of the work presented here. YMB acknowledges funding from the EU Horizon 2020 research and innovation programme under Marie Skłodowska-Curie grant agreement 747645 (ClusterGal) and the Netherlands Organisation for Scientific Research (NWO) through VENI grant 016.183.011. CDV acknowledges financial support from the Spanish Ministry of Economy and Competitiveness (MINECO) through grants AYA2014-58308 and RYC-2015-1807. The Hydrangea simulations were in part performed on the German federal maximum performance computer ‘HazelHen’ at the maximum performance computing centre Stuttgart (HLRS), under project GCS-HYDA / ID 44067 financed through the large-scale project ‘Hydrangea’ of the Gauss Center for Supercomputing. Further simulations were performed at the Max Planck Computing and Data Facility in Garching, Germany. This work also used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grant ST/H008519/1, and STFC DiRAC Operations grant ST/K003267/1 and Durham University. DiRAC is part of the National E-Infrastructure.

REFERENCES

- Andreon S., 2015, *A&A*, 582, A100
 Armitage T. J., Barnes D. J., Kay S. T., Bahé Y. M., Dalla Vecchia C., Crain R. A., Theuns T., 2018, *MNRAS*, 474, 3746
 Bahé Y. M., McCarthy I. G., 2015, *MNRAS*, 447, 969
 Bahé Y. M., McCarthy I. G., Balogh M. L., Font A. S., 2013, *MNRAS*, 430, 3017
 Bahé Y. M., Schaye J., Crain R. A., McCarthy I. G., Bower R. G., Theuns T., McGee S. L., Trayford J. W., 2017a, *MNRAS*, 464, 508
 Bahé Y. M. et al., 2016, *MNRAS*, 456, 1115
 Bahé Y. M. et al., 2017b, *MNRAS*, 470, 4186
 Balogh M. L., McGee S. L., 2010, *MNRAS*, 402, L59
 Barber C., Schaye J., Bower R. G., Crain R. A., Schaller M., Theuns T., 2016, *MNRAS*, 460, 1147
 Barnes D. J., Kay S. T., Henson M. A., McCarthy I. G., Schaye J., Jenkins A., 2017a, *MNRAS*, 465, 213
 Barnes D. J. et al., 2017b, *MNRAS*, 471, 1088
 Behroozi P., Wechsler R., Hearin A., Conroy C., 2018, preprint ([arXiv:1806.07893](https://arxiv.org/abs/1806.07893))
 Berrier J. C., Stewart K. R., Bullock J. S., Purcell C. W., Barton E. J., Wechsler R. H., 2009, *ApJ*, 690, 1292
 Binney J., Tremaine S., 2008, *Galactic Dynamics*, 2nd edn. Princeton Univ. Press, Princeton, NJ
 Bocquet S. et al., 2015, *ApJ*, 799, 214
 Budzynski J. M., Kogosov S. E., McCarthy I. G., McGee S. L., Belokurov V., 2012, *MNRAS*, 423, 104
 Cameron E., 2011, *PASA*, 28, 128
 Chua K. T. E., Pillepich A., Rodriguez-Gomez V., Vogelsberger M., Bird S., Hernquist L., 2017, *MNRAS*, 472, 4343
 Correa C. A., Wyithe J. S. B., Schaye J., Duffy A. R., 2015, *MNRAS*, 452, 1217
 Crain R. A. et al., 2015, *MNRAS*, 450, 1937
 Crain R. A. et al., 2017, *MNRAS*, 464, 4204
 Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
 De Lucia G., Blaizot J., 2007, *MNRAS*, 375, 2
 De Lucia G., Kauffmann G., Springel V., White S. D. M., Lanzoni B., Stoehr F., Tormen G., Yoshida N., 2004, *MNRAS*, 348, 333
 Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497
 Dressler A., 1980, *ApJ*, 236, 351
 Dutton A. A., Macciò A. V., 2014, *MNRAS*, 441, 3359
 Fujita Y., 2004, *PASJ*, 56, 29
 Furlong M. et al., 2015, *MNRAS*, 450, 4486
 Furlong M. et al., 2017, *MNRAS*, 465, 722
 Gao L., White S. D. M., Jenkins A., Stoehr F., Springel V., 2004, *MNRAS*, 355, 819
 Gao L., Navarro J. F., Frenk C. S., Jenkins A., Springel V., White S. D. M., 2012, *MNRAS*, 425, 2169
 Ghigna S., Moore B., Governato F., Lake G., Quinn T., Stadel J., 1998, *MNRAS*, 300, 146
 Gill S. P. D., Knebe A., Gibson B. K., 2005, *MNRAS*, 356, 1327
 Gunn J. E., Gott J. R., III, 1972, *ApJ*, 176, 1
 Guo Q. et al., 2011, *MNRAS*, 413, 101
 Han S., Smith R., Choi H., Cortese L., Catinella B., Contini E., Yi S. K., 2018, *ApJ*, 866, 78
 Harvey D., Massey R., Kitching T., Taylor A., Tittley E., 2015, *Science*, 347, 1462
 Hayashi E., Navarro J. F., Taylor J. E., Stadel J., Quinn T., 2003, *ApJ*, 584, 541
 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
 Jiang F., van den Bosch F. C., 2017, *MNRAS*, 472, 657
 Katz N., White S. D. M., 1993, *ApJ*, 412, 455
 Kauffmann G., White S. D. M., Heckman T. M., Ménard B., Brinchmann J., Charlot S., Tremonti C., Brinkmann J., 2004, *MNRAS*, 353, 713
 Lagos C. del P. et al., 2015, *MNRAS*, 452, 3815
 Lan T.-W., Ménard B., Mo H., 2016, *MNRAS*, 459, 3998
 López-Cruz O., Yee H. K. C., Brown J. P., Jones C., Forman W., 1997, *ApJ*, 475, L97
 Lovell M. R. et al., 2012, *MNRAS*, 420, 2318
 Marasco A., Crain R. A., Schaye J., Bahé Y. M., van der Hulst T., Theuns T., Bower R. G., 2016, *MNRAS*, 461, 2630
 McGee S. L., Balogh M. L., Bower R. G., Font A. S., McCarthy I. G., 2009, *MNRAS*, 400, 937
 McGee S. L., Bower R. G., Balogh M. L., 2014, *MNRAS*, 442, L105
 Moore B., Katz N., Lake G., Dressler A., Oemler A., 1996, *Nature*, 379, 613
 Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, *ApJ*, 524, L19
 More S., Kravtsov A. V., Dalal N., Gottlöber S., 2011, *ApJS*, 195, 4
 Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
 Muldrew S. I., Pearce F. R., Power C., 2011, *MNRAS*, 410, 2617
 Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
 Peng Y.-j. et al., 2010, *ApJ*, 721, 193
 Pillepich A. et al., 2018, *MNRAS*, 475, 648
 Planck Collaboration XVI, 2014, *A&A*, 571, A16
 Press W. H., Schechter P., 1974, *ApJ*, 187, 425
 Prins L., 2018, Master’s thesis, Universiteit Leiden
 Qu Y. et al., 2017, *MNRAS*, 464, 1659
 Randall S. W., Markevitch M., Clowe D., Gonzalez A. H., Bradač M., 2008, *ApJ*, 679, 1173
 Robertson A. et al., 2018, *MNRAS*, 476, L20
 Rodriguez-Gomez V. et al., 2015, *MNRAS*, 449, 49

- Rosas-Guevara Y. M. et al., 2015, *MNRAS*, 454, 1038
 Rozo E. et al., 2009, *ApJ*, 703, 601
 Rykoff E. S. et al., 2014, *ApJ*, 785, 104
 Saro A. et al., 2015, *MNRAS*, 454, 2305
 Saro A., De Lucia G., Dolag K., Borgani S., 2008, *MNRAS*, 391, 565
 Scannapieco C. et al., 2012, *MNRAS*, 423, 1726
 Schaller M., Dalla Vecchia C., Schaye J., Bower R. G., Theuns T., Crain R. A., Furlong M., McCarthy I. G., 2015, *MNRAS*, 454, 2277
 Schaye J. et al., 2015, *MNRAS*, 446, 521
 Schaye J., 2004, *ApJ*, 609, 667
 Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
 Searle L., Zinn R., 1978, *ApJ*, 225, 357
 Sereno M., Ettori S., 2015, *MNRAS*, 450, 3675
 Somerville R. S., Hopkins P. F., Cox T. J., Robertson B. E., Hernquist L., 2008, *MNRAS*, 391, 481
 Springel V., 2005, *MNRAS*, 364, 1105
 Springel V., Di Matteo T., Hernquist L., 2005, *MNRAS*, 361, 776
 Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
 Springel V. et al., 2008, *MNRAS*, 391, 1685
 Tormen G., Diaferio A., Syer D., 1998, *MNRAS*, 299, 728
 Tormen G., Moscardini L., Yoshida N., 2004, *MNRAS*, 350, 1397
 Trayford J. W. et al., 2015, *MNRAS*, 452, 2879
 Trayford J. W. et al., 2017, *MNRAS*, 470, 771
 Tremmel M. et al., 2019, *MNRAS*, 483, 3336
 van den Bosch F. C., 2017, *MNRAS*, 468, 885
 van den Bosch F. C., Ogiya G., 2018, *MNRAS*, 475, 4066
 van den Bosch F. C., Ogiya G., Hahn O., Burkert A., 2018, *MNRAS*, 474, 3043
 van Son L. A. C. et al., 2019, *MNRAS*, 485, 396
 Vegetti S., Lagattuta D. J., McKean J. P., Auger M. W., Fassnacht C. D., Koopmans L. V. E., 2012, *Nature*, 481, 341
 Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
 Weinberg D. H., Colombi S., Davé R., Katz N., 2008, *ApJ*, 678, 6
 Wetzel A. R., Tinker J. L., Conroy C., 2012, *MNRAS*, 424, 232
 Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2013, *MNRAS*, 432, 336
 White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341
 Wiersma R. P. C., Schaye J., Smith B. D., 2009a, *MNRAS*, 393, 99
 Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, *MNRAS*, 399, 574
 Xie L., Gao L., 2015, *MNRAS*, 454, 1697
 Zabludoff A. I., Mulchaey J. S., 1998, *ApJ*, 496, 39
 Zhang Y.-Y., Andernach H., Caretta C. A., Reiprich T. H., Böhringer H., Puchwein E., Sijacki D., Girardi M., 2011, *A&A*, 526, A105

APPENDIX A: THE SPIDERWEB TRACING ALGORITHM

In this Appendix, we provide a detailed description of the SPIDERWEB algorithm that we have used to trace simulated galaxies through time, a significantly updated version of the procedure described in Bahé & McCarthy (2015) and Bahé et al. (2017b). The fundamental assumption is that simulated galaxies are physical structures that persist through time and are therefore, in general, present in multiple snapshots. In each snapshot, an (identified) galaxy corresponds to exactly one subhalo in the SUBFIND catalogue. Tracing galaxies through time therefore equates to identifying those subhaloes in successive snapshots that represent the same galaxy (see Fig. A1 for a schematic illustration).

A1 Extraction of links from subhalo catalogues

We make use of the Lagrangian nature of the Hydrangea simulations, which allows us to identify the same particle in successive snapshots i and j . Any subhaloes in i and j that have particles in

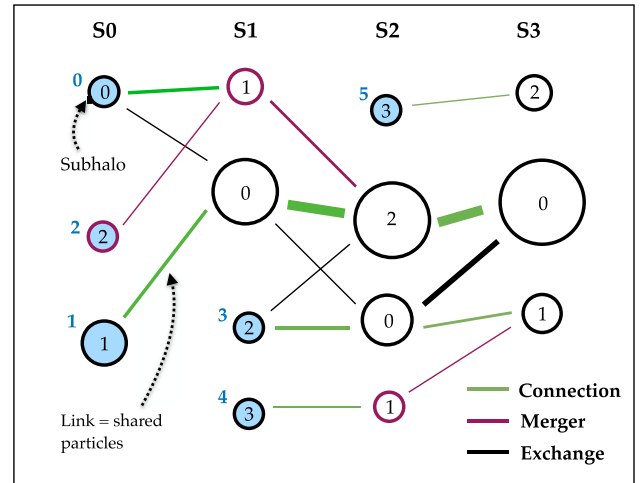


Figure A1. Example tracing situation with 14 subhaloes (circles with an inscribed number representing their arbitrary ID and size indicating their mass) in four snapshots S0–S3 (different columns). Lines represent links, i.e. particle overlaps between subhaloes, whose width scales with link particle number (for simplicity assuming that all particles have equal mass). Green lines represent links that connect successive subhaloes of the same galaxy. The first subhalo of each galaxy is shaded blue, with a number to the left indicating the (arbitrary) galaxy ID. Purple circles indicate the last subhalo of a galaxy about to disrupt/merge, with the corresponding purple link pointing to the galaxy’s carrier in the next snapshot. Black lines represent exchange links between galaxies.

common may, in principle, represent the same galaxy. We therefore begin by extracting all such ‘links’ (i.e. particle overlaps) between subhaloes in i and j , including particles of all types (not just dark matter).

Intuitively, it may be more natural to only consider one link per subhalo in i (as is done in the schemes of e.g. Rodriguez-Gomez et al. 2015 and Qu et al. 2017). However, our more general choice is justified in the regime of groups and clusters, where galaxies may not only grow, but also lose mass through tidal and/or hydrodynamic stripping. As illustrated in Fig. A1 (galaxy 3 between S2 and S3), this can lead to the majority of particles from one galaxy being transferred to another (e.g. the central cluster galaxy), and so would require predicting which particles are least likely to be transferred. Our approach is, instead, to test whether the main link (with the largest particle overlap, see below) leads to a viable descendant, and consider alternative links if this is not the case. In this way, we aim to trace individual galaxies for as long as possible.

Nevertheless, we also give special consideration to a small set of ‘core’ particles in each subhalo, defined as the 5 per cent most bound collisionless particles (i.e. excluding gas), limited to a maximum number of 10^5 . We have found that this is necessary to correctly trace galaxies in situations where a large fraction of particles are transferred from one galaxy to another as a result of a swap in the central/satellite classification between the two.¹⁵ As illustrated in the top panel of Fig. A2, this could lead to a transfer of galaxy ID from one object to the other ($i_1 \rightarrow j_0$), leaving one subhalo without descendant (i_0) and one without progenitor (j_1). Under the plausible assumption that the core particles are least likely to be affected by

¹⁵Due to the way in which SUBFIND associates particles to satellite galaxies, there can be a large population of ‘ambiguous’ particles that are preferentially assigned to the central, and therefore change subhalo membership if the central/satellite classification is swapped.

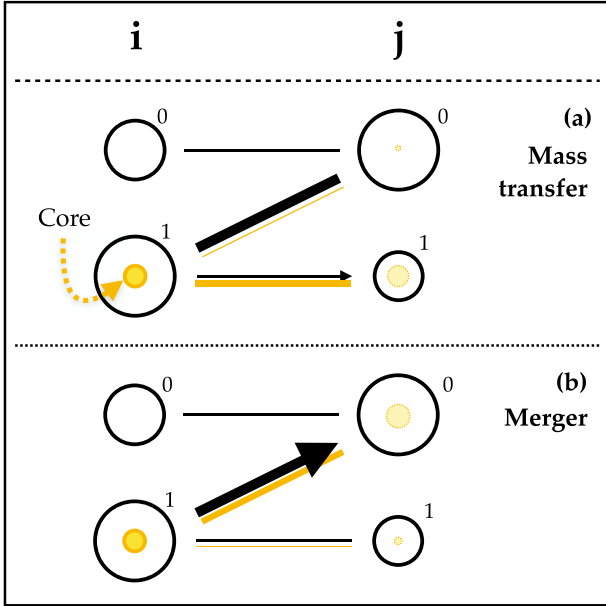


Figure A2. Identification of mass transfer by considering subhalo cores. Black circles represent two subhaloes each in two successive snapshots i and j . The total particle overlap between them, as indicated by the width of the connecting black lines, is identical in the top and bottom configurations and could be interpreted as either mass transfer or a merger. The small yellow circle in subhalo i_1 represents its most bound ‘core’ particles. As indicated by the yellow lines, these transfer differently in the two scenarios: under mass transfer (top), they remain mostly within their own galaxy and end up in j_1 (as indicated by the faint yellow circles in the j subhaloes), but in a merger they end up (mostly) in subhalo j_0 .

such an artificial particle transfer, they represent a robust tracer for their galaxy (yellow lines in Fig. A2) and can therefore distinguish this case from a similar situation in which two galaxies merge (bottom panel of Fig. A2).

Although this reasoning would suggest that the core should be as small as possible – ideally containing only the few most-bound particles – there are two arguments against a very small core. First, small regions of a galaxy, such as a spiral arm or a central clump near a massive black hole, can occasionally become self-bound and dense enough that they appear as a separate entry in the subhalo catalogue. With a very small core, there is a risk that the majority of core particles become members of such a spurious subhalo, which would then be (wrongly) identified as the descendant. Secondly, we have found that within the most bound few per cent of particles the ordering in terms of binding energy fluctuates noticeably between snapshots. In other words, only a small fraction of e.g. the 0.01 per cent most bound particles in i are also the 0.01 per cent most bound in j , whereas at a threshold of ≈ 1 –5 per cent, this fraction approaches unity. We have found that our choice of core fraction, a compromise between these competing constraints, produces stable tracing results across the full range of subhalo masses encountered in our simulations.

For each link, we record the subhalo to which it is connected in i (which we call its ‘sender’) and in j (its ‘receiver’), as well as its total number of particles (N), their total mass (M), and number of particles that form the core of its sender (N_{core} , which may be zero). This information is then used in the following steps to deduce which links connect the same galaxy between snapshots, and which represent interactions between two different galaxies. We note that

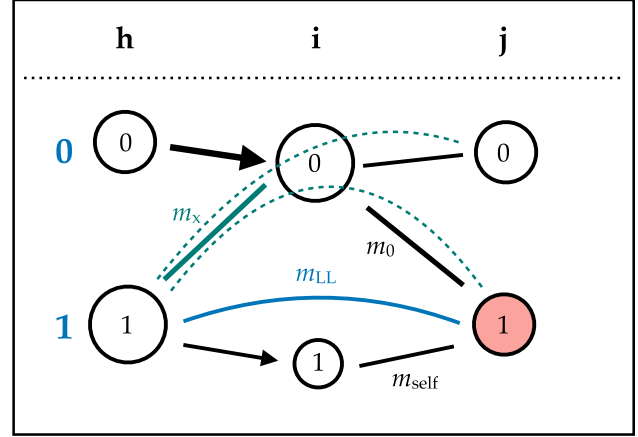


Figure A3. Accounting for prior mass exchanges between interacting galaxies. Subhalo j_1 (red) may be the descendant of either i_0 (galaxy 0) or i_1 (galaxy 1). Prior to the snapshot interval i – j , these galaxies have already exchanged mass (turquoise link h_1 – i_0). These particles may continue along either of the two curved dashed turquoise lines, i.e. to j_0 or j_1 , but only the latter case is relevant in determining the progenitor of j_1 . As described in the text, the fraction following this path is estimated from the three link masses indicated as m_x , m_{self} , and m_{LL} . Lines with arrowheads represent connections following the same galaxy.

there are typically only slightly more links than subhaloes (within ~ 50 per cent).

A2 Compensation of prior mass exchanges

In the simplest scenario, each subhalo in j would receive only one link, in which case it could be unambiguously identified as representing the same galaxy as that link’s sender in i . In reality, however, there will frequently be situations where a subhalo in j receives several links, because it amalgamates matter from several subhaloes in i (as a result of mergers or mass exchange). In this situation, illustrated in Fig. A3, it is less obvious to decide which link to select as the one leading back to the progenitor of the target galaxy (j_1 , highlighted in red).

Physically, it is desirable to rank candidate progenitors in order of the mass that they contribute to the target galaxy in j . A complication with this approach is that galaxies may exchange mass – both physically and numerically – over an extended period of time. Neither the total subhalo masses in i , nor the mass of their links to the target in j , may therefore be a fair proxy of what fraction of the target galaxy is actually contributed by each progenitor candidate. A common way to correct for this, first suggested by De Lucia & Blaizot (2007), is to rank the candidate progenitors by their ‘branch mass’, i.e. the total mass of their progenitors in all previous snapshots.

Here, we follow a different strategy, exploiting the fact that the link network provides us with a complete record of all galaxy interactions prior to snapshot i . We can therefore reconstruct the net prior mass exchange between each pair of candidate progenitors, and then adjust the masses of their links to the target appropriately.

However, as illustrated in Fig. A3, the link network does not provide full information about the correlation between links in different snapshot intervals. For instance, the particles in the exchange link between snapshots h and i (solid turquoise line) may, in the interval i – j , either be carried on towards j_0 (the subhalo at the top), or to the target currently under consideration (j_1), as indicated

by the two dashed lines. In principle, it is possible to distinguish between these cases by explicitly comparing particle IDs in different links, but this would add unjustified complexity to the code.

Instead, we estimate the connection between links by comparing the particle IDs between subhaloes in *non-adjacent* snapshots, i.e. h and j in the current example. These ‘long links’ directly measure how many particles have been transferred from subhalo h_1 to the target (j_1), but contain no information about which subhaloes (if any) these particles were associated with in i . We therefore estimate the total mass that has been ‘bypassed’ around a subhalo (i_1 in Fig. A3) via another (here, i_0) as $m_{\text{by}} = m_{\text{LL}} - m_{\text{self}}$ (limited to the interval $[0, m_x]$), where m_{LL} , m_x , and m_{self} are, respectively, the masses of the long link (h_1-j_1), the exchange link (h_1-i_0), and the direct link i_1-j_1 (see Fig. A3). If there is more than one subhalo in i along which mass could be routed from h_1 to j_1 , we compute the ratio $f_{\text{by}} = m_{\text{by}}/\Sigma_m$ (where Σ_m is the sum of link masses from h_1 to all such subhaloes in i), limited to the interval $[0, 1]$, and assume that a fraction f_{by} of each individual exchange links is routed towards the target j_1 .

This direct accounting scheme is only performed for up to four snapshot intervals prior to $i-j$ (i.e. typically 2.5 Gyr prior to j). In situations where galaxies have already interacted before this point, we count the full mass of these ‘old’ exchange links. This is justified because such long-lasting interactions typically affect satellites orbiting a much more massive host, where there is almost guaranteed to be no confusion about the correct progenitor–descendant identification. Furthermore, the weighting factors f_{by} computed for recent interactions are typically close to unity, with a (mass-weighted) average of ≈ 0.8 .

As a result, we obtain a $N \times N$ matrix (\mathbf{X}) that contains for any pair of the N candidate progenitors in i ($N = 2$ in the example of Fig. A3) an estimate of the total mass that was previously transferred from one to the other, and is now transferred to the target j_1 . In other words, $X_{1,0}$ contains (an estimate of) the mass in the link i_0-j_1 that should actually be counted towards the mass of link i_1-j_1 . We then check whether this reassignment of mass is physically justified: galaxies that are currently undergoing significant stripping are unlikely to simultaneously re-accrete mass from other galaxies. As illustrated in Fig. A4, we therefore test for each candidate link to the target (j_1) whether it carries at least $2/3$ the number of core particles (n_C) in the link with the highest n_C from the same sender (which may be the link itself). If this is not the case, we judge that any galaxy connected along this link would be unlikely to re-gain particles and therefore set the corresponding entry in the exchange matrix \mathbf{X} to zero.

As a final consistency check, we compute for each candidate progenitor in i the total mass that it needs to return to the other progenitors. If this sum exceeds the total mass of its link to the target (m_0 in Fig. A3), all exchanges are scaled down such that their sum is equal to this link mass. At last, we then define a ‘compensated link mass’ m_{comp} equal to the original link mass, reduced by the total mass returned to other links and increased by returns from other links. These compensated masses represent an estimate of the true contribution of each candidate galaxy to the target.

A3 Ranking and filtering of links

To determine the order in which each link should be considered when finding the descendant of its sender, and the progenitor of its receiver, it is necessary to rank all links sent, and likewise all those received, by one subhalo according to some priority criterion.

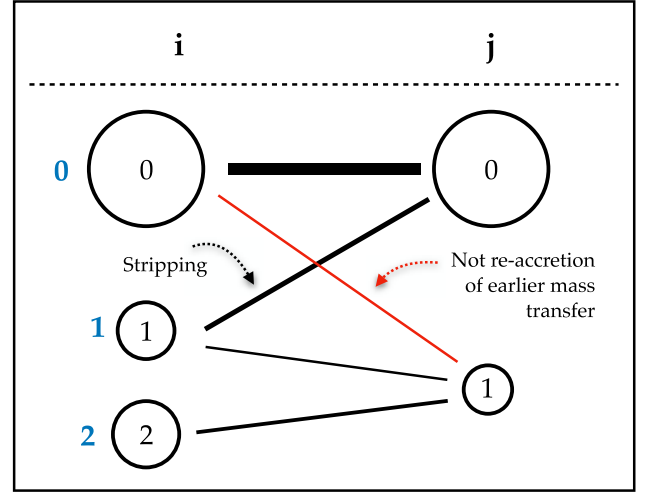


Figure A4. A situation in which transfer compensation is not allowed. Galaxy 1 transfers the majority of its mass to galaxy 0 (i_1-j_0), so it is unphysical to assume that it regains previously transferred mass at the same time (i_0-j_1). This can affect the choice between subhaloes i_1 and i_2 as progenitor of j_1 .

For sender ranking, we order links by their *number of core particles* (n_C). As explained above, this ensures that the most bound particles, which are least affected by numerical mass transfers, are given the highest weight in determining the descendant of a galaxy. Note that we do not weight particles by mass here. This is because some core particles (notably BHs) can be orders of magnitude more massive than others, but we are here treating the particles as tracers to determine the most plausible descendant subhalo. To limit the effect of small-number statistics, we group all links with $n_C < 3$ and rank them at the bottom according to their *total* number of particles. Because link selection proceeds from the highest ranked downwards (see next section), those links are typically irrelevant for determining the evolution of galaxies, except in very low-mass systems close to the resolution limit.

Analogously, all links to the same receiver subhalo are ranked according to their *compensated mass* (see previous section). The reason for employing mass weighting here, and including gas particles, is that we are now interested in determining the galaxy that has contributed the most to the subhalo under consideration, so that all particles should be included and more massive particles *should* carry more weight. Note that, for particle species that can change their mass over the course of the simulation, all masses are consistently determined at the later of the two snapshots (j). We first select those links with $n_C \geq 3$ and rank them in inverse order of their compensated mass, giving highest priority to those with the highest m_{comp} . All links with $n_C < 3$ are then ranked in a second group at the bottom, again in inverse order of m_{comp} . We also assign an analogous receiver rank based on the original, uncompensated link masses.

As a final step before selecting the links that connect subhaloes belonging to the same galaxy, we need to filter the links to exclude those that would lead to physically questionable connections. As discussed above, a key feature of our approach is to include the possibility of connecting subhaloes along links that carry only a minority of the (core) particles from the sender subhalo. This can be physically motivated in the case of strong stripping. It can also, however, lead to physically undesirable situations in which a small part of a galaxy that is temporarily identified as an independent

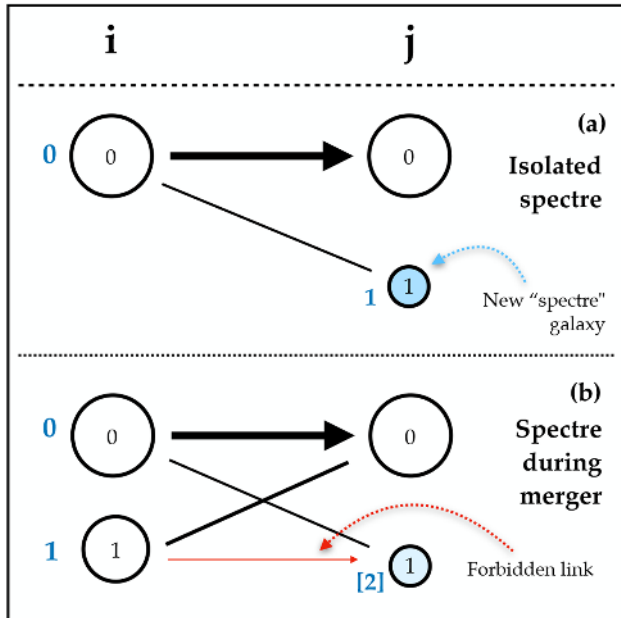


Figure A5. Exclusion of links to prevent connections to spurious ‘spectre’ galaxies. *Top:* formation of a spectre in isolation, which is easily identified as a new galaxy. *Bottom:* when a spectre forms at the same time as a merger, it may be mis-identified as the descendant of the merging galaxy (1). The link i_1-j_1 is therefore marked as forbidden.

subhalo – which we refer to as a ‘spectre’ in the following – is selected as a descendant.

To illustrate this possibility, consider the situation depicted in Fig. A5. The top panel shows a simple scenario in which a spectre is formed from part of an existing galaxy. Because the link to the spectre (blue) carries only a small fraction of the core particles, it is ranked below the link to the main subhalo (indicated by an arrow). The latter is therefore identified as the galaxy’s progenitor, while the spectre becomes a new galaxy (see below).

The situation becomes more complex in the bottom panel. Here, the spectre is generated during a galaxy merger, and contains matter from both merging galaxies. Both subhaloes in i send their highest-ranked link towards subhalo j_0 , but because link i_0-j_0 contains more mass than i_1-j_0 , i_0 is selected as the progenitor of j_0 . However, i_1 has a second link to j_1 (the spectre), which could therefore undesirably be selected as its descendant.

To exclude such mis-identifications, we mark a link as ‘forbidden’ if it carries neither an appreciable fraction of its sender subhalo’s core nor of its total particles (i.e. $n_c < 2/3n_c^{\max}$ and $n < 2/3n^{\max}$, where m and n are the link mass and number of particles, respectively, and n_c^{\max} and n^{\max} the maximum number of core and total particles sent from subhalo i_1 along any link, respectively), and additionally satisfies either $m < 2/3m_{\text{recv}}^{\max}$ or $n < 2/3n_{\text{recv}}^{\max}$ (with m_{recv}^{\max} and n_{recv}^{\max} the maximum mass and particle number received along any link at subhalo j_1 , respectively). These criteria capture both the situation depicted in the bottom panel of Fig. A5 – in which all four conditions are satisfied – and more subtle situations in which subhalo i_1 contributes the majority of mass, but not particles to the spectre (typically, this occurs if a massive BH particle originally belonging to i_1 becomes part of the spectre). At the same time, it does not exclude physically plausible scenarios, for instance if galaxy 1 were severely stripped (m and n close to m_{recv}^{\max} and n_{recv}^{\max} , respectively) or re-accreted mass that was (physically or numerically) temporarily ascribed to i_0 (n_c or n close to n_c^{\max} or n^{\max}).

A4 Select connecting links

Once all links are ranked, and those that are not physically plausible excluded, those links that connect the progenitor and descendant subhaloes of the same galaxy are selected. As discussed above, our approach is to consider several possible links for each subhalo and attempt connections first along the highest-ranked (most plausible) one, and then successively lower-ranked alternatives if the former were unsuccessful. The highest-priority class of links are clearly those with the highest sender- and receiver-rank (which we denote as 0). For subsequent levels, there is an ambiguity between sender-rank 0, but receiver-rank 1, and sender-rank 1, but receiver-rank 0 (and analogous for lower levels). We here prioritize the former, which effectively prefers connecting galaxies in such a way that they retain the largest possible fraction of their core particles, rather than accrete the smallest possible fraction of mass from other objects. In practice, we have found that there is hardly any difference between these two ordering options.

We therefore iterate through successively lower sender ranks and consider all those links whose sender- and receiver-subhaloes have both not yet been connected. All links that are the only ones leading to their respective receiver subhalo can be selected to connect its two associated subhaloes as part of the same galaxy.

For subhaloes in j that receive multiple links in the current iteration, the most straightforward solution would be to select the one with the highest receiver rank. However, to increase the robustness of the tracing results, we first test whether the link with the highest receiver rank based on compensated mass is the same as that obtained with the uncompensated, original masses. If so, the situation is unambiguous and the highest receiver-rank link is connected.

If the two estimates differ – as depicted in Fig. A6 – we proceed to the *next* snapshot interval ($j-k$) and identify the ‘likely descendant’ of the subhalo currently under consideration. The motivation behind this is to select the progenitor that maximizes the long-term particle overlap between different subhaloes of the galaxy. The links between j and k are analysed in the same way as between i and j , but without mass compensation. We next test whether there are long links between the two respective sender subhaloes in i and the likely descendant in k (i.e. whether they share any particles), and if so, whether the long-link corresponding to the $i-j$ link with the higher original receiver rank (i_0 in the example of Fig. A6) has a higher sender rank than its alternative (i_1). If this is the case, disregarding the mass compensation leads to a better long-term particle consistency, so the corresponding link is selected (i_0-j_0). In all other situations (including if j is the last snapshot of the simulation), we select the link with the higher compensated receiver rank. Reassuringly, this covers the vast majority of cases: only in ≈ 10 per cent of ambiguous situations (i.e. with different links labelled as highest-ranked by compensated and original mass) is the selected one that which is highest-ranked by original mass. Overall, the receiver ranks computed from compensated and original mass differ in only ≈ 0.5 per cent of selected links in each snapshot interval.

A5 Connect temporary non-identifications

At this point, all subhaloes in i and j that can plausibly be identified as representing the same galaxy are connected. However, it is still possible that a subhalo in j could not be connected although it represents an existing galaxy: subhaloes are occasionally missed by SUBFIND, especially against the dense background of a

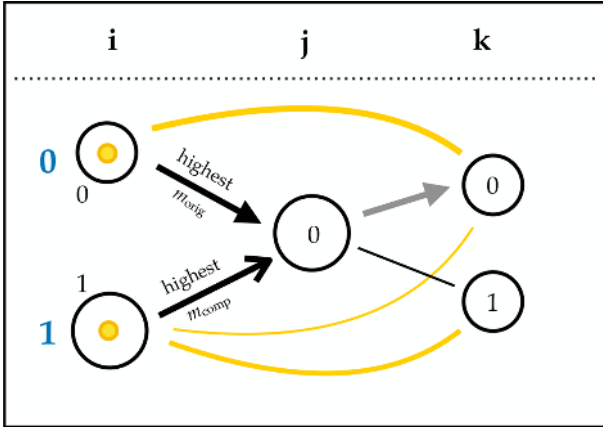


Figure A6. Decision between ambiguous progenitors. Subhalo j_0 could be the descendant of i_0 (which contributes the largest amount of mass, m_{orig}) or i_1 (whose mass contribution adjusted for prior exchanges, m_{comp} , is greatest). In the following snapshot (k), k_0 is the likely descendant of j_0 , but only i_0 sends most of its core particles to this subhalo. The progenitor of j_0 is therefore chosen to be i_0 (continuing galaxy 0), while galaxy 1 is merged (and will likely be re-connected to k_1 in the next snapshot, as described in Section A5). Under all other circumstances, i_1 would be selected as the progenitor of j_0 .

massive galaxy cluster. An orbiting satellite galaxy may therefore be (temporarily) left without a counterpart in the subhalo catalogue in snapshot i . Uncorrected for, such galaxies would appear to spuriously disrupt and then form as new galaxies a short time later. This is clearly not an appropriate description of their actual evolution.

To prevent such mis-classifications, we make use of the already-mentioned long links and retrospectively connect galaxies that were left without a descendant in an earlier snapshot ($z > z_i$) to a subhalo in j (see Fig. A7). Such long-link connections are enabled over up to five snapshot intervals – if a galaxy can still not be connected after this period, it is marked as disrupted. The selection of long-links is performed in analogy to the steps for direct links described above.¹⁶ Because connections along long links should be an exception, rather than the rule, a number of additional constraints on their eligibility are imposed. All of them aim to limit the selection of long links to cases where they are clearly required:

(i) The sender subhalo must not currently have a descendant, or if it does, this descendant subhalo could in turn not be connected. The first case covers the standard situation of a galaxy temporarily disappearing from the catalogue. The second, less common, case arises in situations where a small part of a disappearing galaxy (e.g. a spectre) is still identified as a separate subhalo, but is not strongly enough linked to the galaxy when it re-appears to establish a connection (bottom panel of Fig. A7). In this case, we have the option of re-establishing a link from the last snapshot in which the galaxy was properly identified. The original descendant (h_1) is then disconnected and turned into an ‘orphan’ galaxy that only exists in one snapshot.

(ii) The link must contain at least three core particles. This is to exclude connections that represent only marginal (core) particle

¹⁶The only difference is that, in computing their compensated masses, all links are allowed to re-gain particles. This is because galaxies may transition from stripping to re-accretion over the longer time intervals probed by long links.

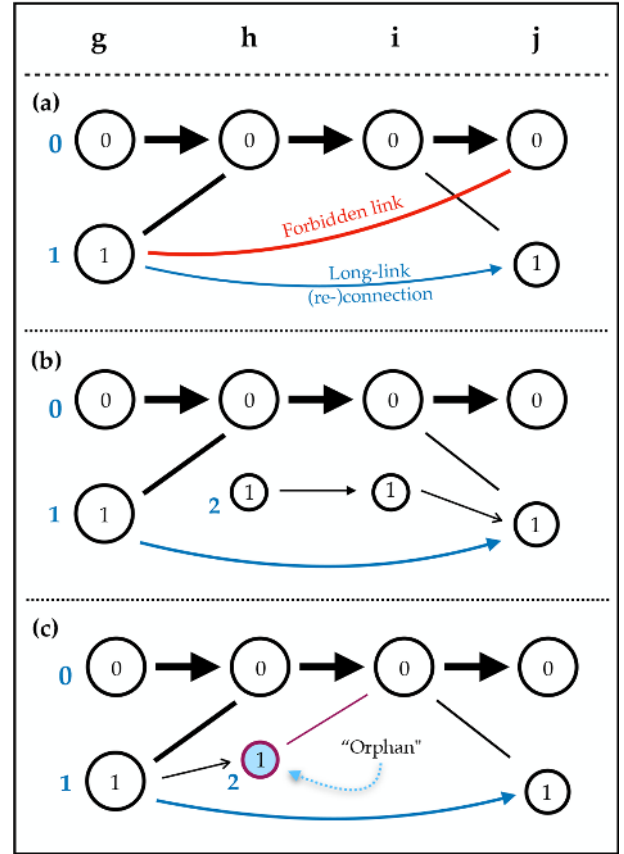


Figure A7. Three example situations involving the re-connection of subhaloes via long-links. Panel (a) (top) depicts the simplest case in which a subhalo without descendant (g_1) is re-connected to a subhalo without progenitor (j_1). As described in the text, link g_1-j_0 is forbidden because it would invert the survival order established in snapshot h . Panel (b) (middle) illustrates a long-link connection that ‘overrides’ a (weaker) direct link from galaxy 2 (subhalo i_1). Panel (c) (bottom) shows a case in which an originally identified descendant (h_1) is disconnected and turned into an ‘orphan’ that is only alive in one snapshot, because the long-link g_1-j_1 allows the continued tracing of galaxy 1 to snapshot j .

overlap, which is not justified in the exceptional situation of linking across multiple snapshots.

(iii) The link must have a higher compensated mass than a (potential) currently connected shorter link to the same receiver subhalo. This is because we want to permit re-connections also in cases where a galaxy has, during its absence from the subhalo catalogue, accreted a smaller galaxy. Naively, the latter would be identified as its progenitor, but if the long link contributes more mass, it should be connected instead (middle panel of Fig. A7).

(iv) The link must not be received by any subhalo that is (backwards) connected to a subhalo that the sender subhalo already sends a link to (see the top panel of Fig. A7 for a schematic illustration, in which the upper long-link, coloured in red, satisfies this criterion). This condition imposes that once a galaxy has been mis-classified as merged with another (as would happen if it has been missed by the subhalo finder against a dense group/cluster background), it cannot at a later point be identified as the progenitor of that galaxy. We have found that this is necessary to prevent unintended situations in which two galaxies of similar mass that have physically merged both ‘survive’ by alternately skipping snapshots, often for many Gyr.

(v) If more than one of the long links satisfy the above constraints per sender subhalo – i.e. if there is more than one option to re-connect a galaxy that has temporarily disappeared – only that with the highest sender rank is allowed, or others that contain at least 2/3 of the number of core particles of that link (and which therefore offer a comparably strong connection).

A6 New galaxies

Once eligible long links are connected, each subhalo in j that can be identified as continuing a pre-existing galaxy is connected with its progenitor, from which they inherit a unique galaxy ID (an identifier that is the same for all, and only those, subhaloes that represent the same galaxy; see Fig. A1). Typically, some subhaloes are still not connected at this stage, because they represent newly formed galaxies. They are therefore assigned new galaxy IDs, which may be passed on to their descendants in subsequent snapshots.¹⁷

While the majority of these new galaxies are relatively small, isolated objects that have just emerged above the detection threshold of SUBFIND, a subset of them are typically spectres, anti-hierarchically formed transient substructures within larger galaxies. Since they typically form in baryon-dominated regions, they can reach up to $\sim 10^9 M_\odot$ in stellar mass and could therefore be confused with genuine galaxies in stellar-mass selected samples. To avoid such contamination, we flag new galaxies as (likely) spectres if they receive at least one link from another galaxy and less than half their particles (by number or mass) were unbound in the previous snapshot. At high z , almost all newly emerging galaxies are genuine, but the fraction of spectres increases steadily with time and reaches ≈ 25 per cent at $z = 0$. In this paper, we have consistently excluded galaxies that were flagged as spectres.¹⁸

A7 Carrier list

The steps described above are repeated for all snapshots in the simulation. At the end of this process, every subhalo in any of the 30 snapshots corresponds to exactly one galaxy, and each galaxy to at most one subhalo in each snapshot. Typically, there are a factor of a few more galaxies than there are subhaloes at the final snapshot, because the majority does not survive¹⁹ to $z = 0$.

As a final step, SPIDERWEB identifies the ‘carrier’ of each disrupted galaxy, i.e. the galaxy that inherits the largest fraction of its (core) particles in the first snapshot after the galaxy has been lost. Note that such a carrier may not exist for all galaxies: if e.g. gradual mass loss brings them below the SUBFIND detection threshold, its particles may all be unbound (not assigned to any subhalo) in the next snapshot. In many other cases, however, including those of interest in this paper, a galaxy is lost because it dissolves (merges) into a more massive galaxy. In this situation, most of its particles are still members of a galaxy, which can therefore be identified as the dissolved galaxy’s ‘carrier’.

To keep track of these mergers, we define a ‘carrier ID’ for each galaxy, which is initially equal to its galaxy ID. Once a galaxy

¹⁷In the first snapshot, no subhalo can have a descendant and so each is assigned a new galaxy ID.

¹⁸Because most spectres are dominated by stars, they constitute an appreciable fraction in stellar mass limited galaxy samples, approximately 15 per cent at $M_\star^{\text{peak}} > 10^9 M_\odot$.

¹⁹We note that this is not in contradiction to our findings in the main part of this paper: the majority of galaxies have lower peak masses than those that we have analysed, and/or disrupt in lower-mass haloes.

Table A1. Example carrier list for the situation depicted in Fig. A1 in each of the four snapshots shown.

Snapshot	Gal. 0	Gal. 1	Gal. 2	Gal. 3	Gal. 4	Gal. 5
S0	0	1	2	–	–	–
S1	0	1	0	3	4	–
S2	1	1	1	3	4	5
S3	1	1	1	3	3	5

is disrupted, its carrier ID is updated to that of the galaxy which receives its highest (sender-)rank link, i.e. which carries the largest share of its (core) particles and is therefore the most plausible merger target. Any other galaxies that it had itself accreted in the past, and whose carrier IDs were therefore equal to its own, are likewise updated. On the one hand, this enables an easy identification of ‘where a galaxy ends up’ at a given snapshot after its disruption. On the other hand, it also provides a simple method of determining all progenitors of a galaxy: those are simply all galaxies whose carrier ID is equal to the galaxy’s own ID.

As an illustration, Table A1 shows the carrier list for the scenario depicted in Fig. A1 in each of the four snapshots. Note that galaxy 2 undergoes two mergers, so its carrier ID is first changed to 0 (in S1) and then to 1 (in S2). The three galaxies that end up in subhalo 0 in S3 (which represents galaxy 1) all have the same carrier ID (1). Galaxies 1, 3, and 5 retain its own galaxy ID as carrier ID because they are still alive in S3.

APPENDIX B: ROBUSTNESS OF SUBHALO IDENTIFICATION

In Fig. 5, we had shown that only a small fraction of galaxies with $M_{\text{tot}}^{\text{peak}} > 10^{10} M_\odot$ that are detected in the SUBFIND catalogue at $z = 0$ have a total mass below $5 \times 10^8 M_\odot$ at $z = 0$. It is conceivable that there are additional galaxies that do (physically) survive – possibly with lower mass – but which are missed in the SUBFIND catalogue for numerical reasons (either the limited resolution of our simulations, or shortcomings of the SUBFIND algorithm).

B1 Possibility of insufficient resolution

To test the possibility that surviving subhaloes may be missed due to the finite resolution of the Hydrangea simulations, we have repeated our analysis on three simulations from the EAGLE project that model the evolution of a $(25 \text{ cMpc})^3$ cube at two different levels of resolution. One, L0025N0376/Ref, uses the same mass and spatial resolution as the Hydrangea simulations analysed in the main part of this paper.²⁰ Two others (L0025N0752/Ref and L0025N0752/Recal) have a mass (spatial) resolution that is higher by a factor of 8 (2), where the latter also uses recalibrated simulation parameters to achieve a similarly good match to the galaxy stellar mass function as the lower-resolution counterpart (see Schaye et al. 2015). Due to their limited volume, these simulations only contain around a dozen low-mass groups ($M_{200c}^{z=0} = 10^{12.5} - 10^{13.5} M_\odot$), with correspondingly larger statistical uncertainties than in the Hydrangea analysis.

In Fig. B1, we show the predicted survival fractions of low-mass group satellites as a function of their peak mass in these

²⁰The Ref model uses slightly different parameters for subgrid feedback from active galactic nuclei than the AGNdT9 model used for Hydrangea, but this is of no significance to the resolution test presented here.




Figure B1. Survival fraction of satellites in low-mass groups in three (25 cMpc)³ simulations from the EAGLE project. L0025N0376/Ref (black solid line) uses the same resolution as the Hydrangea simulations analysed in the main part of this paper, while the other two (blue dashed and purple dotted lines, corresponding to different choices of subgrid parameters) have eight times better mass resolution. In the *top* panel, the same survival threshold ($5 \times 10^8 M_{\odot}$) is applied to all three, while the *bottom* panel uses an eight times lower threshold for the higher-resolution simulations. In the first case, the higher resolution has no impact, but with a lower mass threshold the fraction of surviving galaxies increases slightly at the low-mass end.

three simulations, in analogy to Fig. 4 for Hydrangea. The default-resolution simulation L0025N0376/Ref is shown as a solid black line, while the two higher-resolution versions L0025N0752/Ref and L0025N0752/Recal are represented by blue dashed and purple dotted lines, respectively. The top panel applies the same survival threshold of $5 \times 10^8 M_{\odot}$ to all three and reveals near-perfect agreement between the two resolution levels (irrespective of whether the subgrid parameters in the high-resolution version are re-calibrated or not). At least within the relatively low galaxy and halo masses accessible with these (25 cMpc)³ simulations, the survival fractions above this threshold are therefore insensitive to the finite resolution of our simulations.

In the bottom panel, we explore the effect of lowering the survival mass threshold by a factor of 8 for the two high-resolution simulations. Only at the lowest galaxy masses that we probe ($M_{\text{tot}}^{\text{peak}} \sim 10^{10} M_{\odot}$) does this increase the survival fraction, by up to 20 per cent (in a relative sense) from 35 ± 2 to 40 ± 2 or 43 ± 3 per

cent (in the Ref and Recal models, respectively).²¹ At higher masses – i.e. in the regime where baryonic processes are approximately converged (Schaye et al. 2015; Crain et al. 2017) – there is no indication that a significant population of galaxy remnants is missed because of the finite resolution of the Hydrangea simulations.

In summary, we conclude that our fiducial survival fractions (above a threshold of $5 \times 10^8 M_{\odot}$) are insensitive to an increase in resolution, and that they represent the total survival fractions for satellite galaxies with $M_{\text{tot}}^{\text{peak}} \gtrsim 3 \times 10^{10} M_{\odot}$.

B2 Possibility of missed subhaloes

To test the possibility that SUBFIND may have missed (resolved) subhaloes, we have iteratively recomputed the bound mass of all those galaxies that are *not* present in the SUBFIND catalogue. Starting from the dark matter, star, and black hole particles that constitute each galaxy in the last snapshot in which it is identified,²² we compute the gravitational potential ϕ_i and kinetic energy K_i of each particle i (with respect to the mass-weighted average velocity of the particles in the most negative decile in gravitational potential). Any particle whose binding energy $\epsilon_i = \phi_i + K_i$ is positive is removed, and the iteration continued until less than 0.05 per cent of particles are removed in any one iteration.

The result is shown as the light green line in Fig. B2, which shows the fraction of galaxies in massive clusters ($M_{200c}^{z=0} > 10^{14.5} M_{\odot}$) that are either present in the original SUBFIND catalogue at $z = 0$ or for which the recomputation yielded a remnant with at least 10 bound particles. For comparison, the grey line shows the fraction of only those galaxies identified by SUBFIND (as in Fig. 5). It is evident that the re-computation only increases the survival fraction by at most ≈ 5 per cent, implying that most physically surviving subhaloes are indeed detected by SUBFIND.

In addition, the recomputation described above *considered* each galaxy individually and therefore represents an upper limit to the fraction of galaxies surviving as *independent* self-bound structures. Some of them, while self-bound, may in fact be an indistinguishable part of a more massive galaxy, in the same way as a random selection of stars from the Milky Way’s bulge may be self-bound to each other without constituting a separate galaxy. To estimate the impact of this effect, the dark green line in Fig. B2 includes only those galaxies with a remnant from the re-computation that lies outside of $\min(30 \text{ kpc}, R_{1/2}^{\text{star}})$, where $R_{1/2}^{\text{star}}$ is the stellar half-mass radius from any subhalo in the SUBFIND catalogue. In effect, this limits new detections to the outer halo of more massive galaxies. With this stricter definition, the difference between the recomputed and SUBFIND catalogue disappears almost entirely. Although this may in turn be overly restrictive – some of the galaxies recovered in the recomputed catalogue within $\min(30 \text{ kpc}, R_{1/2}^{\text{star}})$ may in fact be genuine, independent, survivors – we conclude from Fig. B2 that SUBFIND robustly identifies the vast majority of surviving galaxies in a group/cluster environment, and that our results in this paper are therefore not an artefact of our particular subhalo finder.

²¹Fig. 5 already indicated that a few per cent of these low-mass galaxies survive with a mass below our fiducial threshold at the intermediate resolution level of Hydrangea. We have verified, however, that the low-mass survival fraction is still raised in the higher-resolution simulations even when a uniform threshold of $6 \times 10^7 M_{\odot}$ is applied to all three simulations.

²²We do not include gas because this component is expected to be efficiently removed through ram pressure.

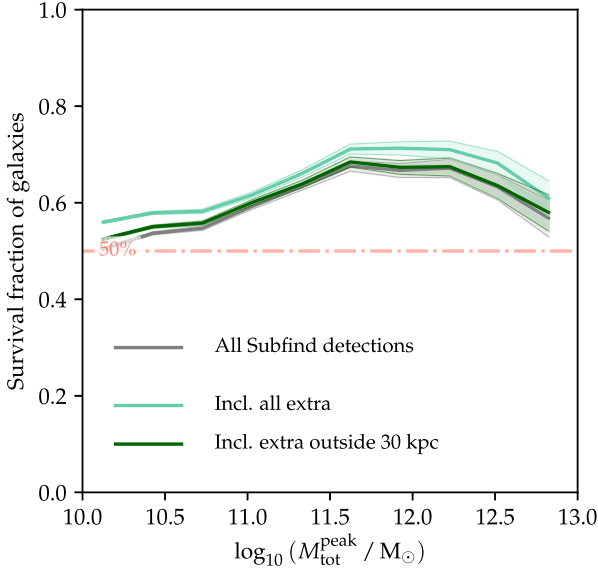


Figure B2. Fraction of galaxies in massive clusters ($M_{200c}^{z=0} > 10^{14.5} M_{\odot}$) that survive to $z = 0$. The grey line includes all galaxies present in the SUBFIND catalogue. The light green line adds galaxies that retain a self-bound remnant with at least 10 particles when starting from all particles that are part of the galaxy in its last snapshot. The dark green line limits those recovered detections to only those that lie away from subhaloes in the SUBFIND catalogue (see the text for details). There is no significant population of surviving galaxies that is missed by SUBFIND, particularly with the more restrictive definition (dark green).

APPENDIX C: COMPARISON TO IDEALIZED EXPERIMENTS

Our simulations predict substantial survival of even low-mass galaxies ($M_{\text{tot}}^{\text{peak}} \sim 10^{10} M_{\odot}$) in massive haloes, especially if they were accreted at $z < 2$. This appears to be in tension with the idealized N -body experiments of van den Bosch & Ogiya (2018), in which numerical disruption occurs even for satellites that are initially resolved by $\geq 10^5$ particles (corresponding to $M_{\text{tot}}^{\text{peak}} \gtrsim 10^{12} M_{\odot}$ at our resolution). For a quantitative comparison to their work, we use the criteria in their equations (21) and (22): these specify the minimum mass fraction that a satellite must retain to avoid numerical artefacts from inadequate force softening and particle discreteness noise, respectively, as

$$f_{\text{bound, min}}^{\text{softening}} = 1.12 \frac{c^{1.26}}{f^2(c)} \left(\frac{\epsilon}{r_{s,0}} \right)^2 \quad \text{and} \quad (\text{C1})$$

$$f_{\text{bound, min}}^{\text{discreteness}} = 0.32 \left(\frac{N_{\text{acc}}}{1000} \right)^{-0.8}, \quad (\text{C2})$$

where $c = r_{s,0}/r_{200c}$ is the concentration parameter of the galaxy's DM halo, $r_{s,0}$ its NFW scale radius (Navarro, Frenk & White 1996), $f(c) = \ln(1+c) - c/(1+c)$, ϵ is the force softening length of the simulation (here equal to 0.7 proper kpc), and N_{acc} the number of particles bound to the galaxy at the time of accretion.

In our simulations, we take N_{acc} from the SUBFIND catalogue at the last snapshot before accretion (t_{branch}). Instead of fitting NFW profiles, we estimate c (and hence $r_{s,0}$) from the redshift-dependent M_{200c} - c relation of Correa et al. (2015, their appendix B1), including lognormal scatter with $\sigma = 0.11$ dex. For galaxies that were a central prior to accretion, we use the M_{200c} of its FOF group, for others, we estimate c from $M_{\text{tot}}^{\text{peak}}$.

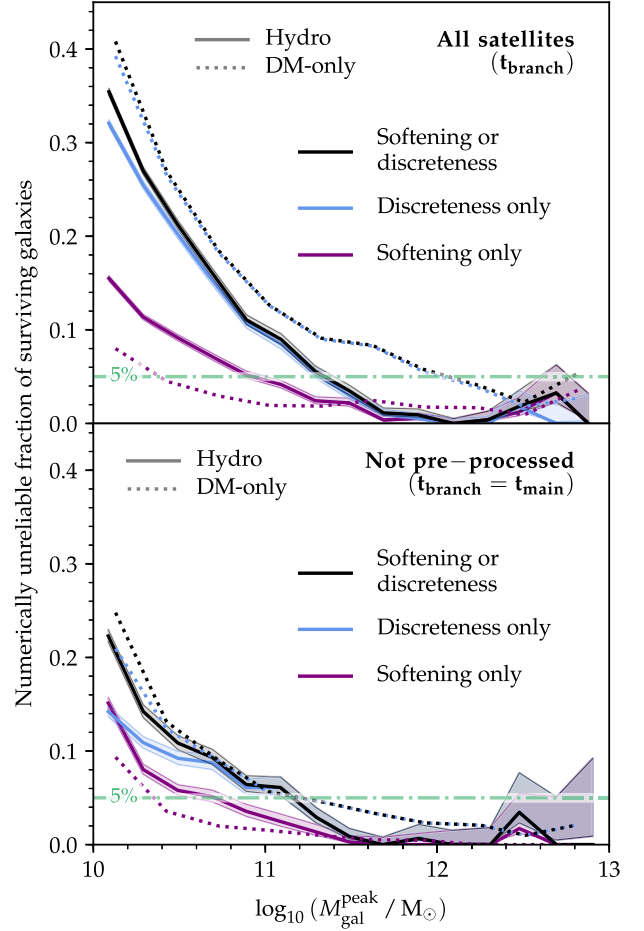


Figure C1. The fraction of surviving galaxies ($M_{\text{tot}}^{z=0} > 5 \times 10^8 M_{\odot}$) whose $z = 0$ remnants are affected by numerical artefacts, according to the criteria of van den Bosch & Ogiya (2018). In the *top panel* we show all surviving satellites (including those that were pre-processed) and calculate pre-accretion properties at t_{branch} . In the *bottom panel* only galaxies that were not pre-processed are shown, with pre-accretion properties calculated at t_{main} . In both cases, the blue, purple, and black lines show, respectively, the fraction of survivors that violate the discreteness noise criterion, the softening criterion, or either of them. Results from the hydrodynamical simulations are shown as solid lines (with binomial 1σ errors as shaded bands), while DM-only simulations are represented by dotted lines. At $M_{\text{tot}}^{\text{peak}} < 3 \times 10^{11} M_{\odot}$, numerical artefacts should play a non-negligible role, but this is not reflected in the survival fractions.

Fig. C1 shows the fraction of surviving ($M_{\text{tot}}^{z=0} > 5 \times 10^8 M_{\odot}$) galaxies in massive clusters ($M_{200c}^{z=0} > 10^{14.5} M_{\odot}$) whose remaining bound fraction ($f_{\text{bound}} = M_{\text{tot}}^{z=0}/M_{\text{tot}}^{\text{peak}}$) is below the requirements in equations (C1) and (C2), i.e. those that are expected to be susceptible to numerical artefacts. Galaxies affected by inadequate force softening are shown by purple lines, those subject to particle discreteness noise in blue, and black lines give the fraction of galaxies violating either constraint. In the top panel, all satellites are included, while the bottom panel only shows galaxies that were not pre-processed.

At $M_{\text{tot}}^{\text{peak}} \gtrsim 3 \times 10^{11} M_{\odot}$, the fraction of remnants violating either numerical reliability constraint is close to zero in the hydrodynamical simulations (black solid lines). For these, numerical artefacts would only occur for bound fractions below 1 per cent, which are extremely rare (Fig. 5). We therefore conclude that

massive galaxies are unaffected by numerical disruption, because they are *physically* disrupted before losing sufficient mass to become numerically unreliable. We note that DM-only simulations do predict a small fraction (≈ 10 per cent) of survivors even at $M_{\text{tot}}^{\text{peak}} = 10^{12} M_{\odot}$ that were so severely stripped that they should have become numerically unreliable, at least when pre-processed galaxies are included (dotted black/blue lines in the top panel; see also Fig. 5).

In lower-mass galaxies, the fraction of numerically unreliable remnants increases rapidly, and reaches 35 (23) per cent of all (non-pre-processed) satellites at $M_{\text{tot}}^{\text{peak}} = 10^{10} M_{\odot}$. The dominant driver is susceptibility to discreteness noise, with softening by itself only affecting a few per cent of galaxies except at $M_{\text{tot}}^{\text{peak}} < 3 \times 10^{10} M_{\odot}$. Because the discreteness noise threshold is independent of concentration, this means that the extent of (overall) numerical unreliability (black lines) is not significantly affected by our simplified approach of estimating c from the Correa et al. (2015) relation.²³ In the DM-only runs, the fraction of unreliable remnants of low-mass galaxies is even slightly higher.

These numerically unreliable survivors should be accompanied by similarly massive galaxies that were numerically disrupted. In our simulations, we instead find near-complete survival of low-mass galaxies in massive clusters (Figs 7 and 11), and only small differences with satellite mass in all host mass bins (Fig. 4) that also tend to be shallower at the lowest masses, rather than steepening as seen in the fraction of unreliable survivors. Finally, the survival fractions of low-mass galaxies are (slightly) lower in the

hydrodynamical simulations compared to the DM-only runs (Figs 5 and 7), although they have a slightly lower fraction of unreliable remnants. All this suggests that numerical disruption of satellites is rare even when the bound fraction falls below the van den Bosch & Ogiya (2018) thresholds.

There are (at least) two possible explanations for this. First, it might be that many low-mass galaxies are actually affected by numerical artefacts in our simulations, but that these are mild and only cause some unphysical mass loss, rather than any appreciable number of extra disruption events (even at a threshold of $0.1 M_{\text{tot}}^{\text{peak}}$ as explored in Fig. 5). Secondly, the criteria of van den Bosch & Ogiya (2018) may be overly conservative in the more realistic situations produced by our simulations. It is conceivable, for instance, that most of the mass loss is due to transient events, such as encounters with other satellites ('galaxy harassment'; Moore et al. 1996) or pre-processing, while the tidal field of the host itself is too weak to induce numerical inaccuracies. More work would be required to test these scenarios in detail.

²³We have verified that using the relation of Dutton & Macciò (2014), or even shifting the $M_{200c}-c$ relation upwards by $1\sigma = 0.11$ dex does not cause an appreciable difference in the overall fraction of numerically unreliable galaxies.

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