

The merger rate of galaxies in the Illustris simulation: a comparison with observations and semi-empirical models

Vicente Rodriguez-Gomez,^{1*} Shy Genel,^{1,2†} Mark Vogelsberger,³ Debora Sijacki,⁴ Annalisa Pillepich,¹ Laura V. Sales,^{1,5} Paul Torrey,^{3,6} Greg Snyder,⁷ Dylan Nelson,¹ Volker Springel,^{8,9} Chung-Pei Ma¹⁰ and Lars Hernquist¹

¹Harvard–Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

²Department of Astronomy, Columbia University, 550 West 120th Street, New York, NY 10027, USA

³Department of Physics, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁴Kavli Institute for Cosmology, Cambridge, and Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, UK

⁵Department of Physics & Astronomy, University of California, Riverside, 900 University Avenue, Riverside, CA 92521, USA

⁶TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA

⁷Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

⁸Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnengasse 35, D-69118 Heidelberg, Germany

⁹Zentrum für Astronomie der Universität Heidelberg, ARI, Mönchhofstr. 12-14, D-69120 Heidelberg, Germany

¹⁰Department of Astronomy, University of California, Berkeley, CA 94720, USA

Accepted 2015 February 5. Received 2015 February 4; in original form 2014 December 10

ABSTRACT

We have constructed merger trees for galaxies in the Illustris simulation by directly tracking the baryonic content of subhaloes. These merger trees are used to calculate the galaxy–galaxy merger rate as a function of descendant stellar mass, progenitor stellar mass ratio, and redshift. We demonstrate that the most appropriate definition for the mass ratio of a galaxy–galaxy merger consists in taking both progenitor masses at the time when the secondary progenitor reaches its maximum stellar mass. Additionally, we avoid effects from ‘orphaned’ galaxies by allowing some objects to ‘skip’ a snapshot when finding a descendant, and by only considering mergers which show a well-defined ‘infall’ moment. Adopting these definitions, we obtain well-converged predictions for the galaxy–galaxy merger rate with the following main features, which are qualitatively similar to the halo–halo merger rate except for the last one: a strong correlation with redshift that evolves as $\sim(1+z)^{2.4-2.8}$, a power law with respect to mass ratio, and an increasing dependence on descendant stellar mass, which steepens significantly for descendant stellar masses greater than $\sim 2 \times 10^{11} M_{\odot}$. These trends are consistent with observational constraints for medium-sized galaxies ($M_{*} \gtrsim 10^{10} M_{\odot}$), but in tension with some recent observations of the close pair fraction for massive galaxies ($M_{*} \gtrsim 10^{11} M_{\odot}$), which report a nearly constant or decreasing evolution with redshift. Finally, we provide a fitting function for the galaxy–galaxy merger rate which is accurate over a wide range of stellar masses, progenitor mass ratios, and redshifts.

Key words: methods: numerical – galaxies: interactions – cosmology: theory.

1 INTRODUCTION

Structure formation in Λ cold dark matter (Λ CDM) cosmological models is hierarchical in nature, which makes galaxy mergers an essential aspect of galaxy formation and evolution. In particular, it is important to quantify the galaxy–galaxy merger rate, namely, the

frequency of galaxy mergers as a function of the masses of the objects involved, redshift, and possibly other parameters such as gas fractions. A precise determination of this quantity is of fundamental interest for understanding the growth and assembly of galaxies, for bringing galaxy formation models into agreement with the observed distribution of galaxy morphologies, and for explaining the frequency of starburst galaxies and active galactic nuclei at high redshifts.

Although significant progress has been made in the determination of dark matter (DM) halo–halo merger rates using N -body

*E-mail: vrodiguez-gomez@cfa.harvard.edu

†Hubble fellow.

cosmological simulations (e.g. Fakhouri & Ma 2008; Genel et al. 2009, 2010; Fakhouri, Ma & Boylan-Kolchin 2010), with most theoretical predictions agreeing within a factor of ~ 2 , similar convergence has yet to be achieved in the determination of the galaxy–galaxy merger rate, in particular using theoretical models of galaxy formation and evolution.

There are three main approaches for making theoretical predictions of the galaxy–galaxy merger rate: (1) semi-empirical methods, which typically use an N -body cosmological simulation and ‘populate’ DM subhaloes with galaxies according to observational constraints, in particular by applying the halo occupation distribution or abundance matching formalisms, (2) semi-analytic models (SAMs), which use an N -body cosmological simulation as the ‘backbone’ of a galaxy formation model, which is implemented in post-processing, and (3) hydrodynamic simulations, which model the DM and baryonic components of a cosmological volume self-consistently. Therefore, the main difference between SAMs or hydrodynamic simulations with respect to semi-empirical methods is that the latter do not attempt to model galaxy formation processes from first principles (i.e. in an a priori fashion), therefore avoiding many of the associated difficulties.

Perhaps the best-known determination of the galaxy–galaxy merger rate using a SAM is the one by Guo & White (2008), although several other examples can be found in Hopkins et al. (2010b). On the other hand, there have been very few attempts to determine the galaxy–galaxy merger rate using hydrodynamic simulations (e.g. Maller et al. 2006; Kaviraj et al. 2014) due to the fact that until recent years it was not possible to produce statistically significant and sufficiently realistic populations of galaxies in cosmological hydrodynamic simulations.

In general, calculations of the galaxy–galaxy merger rate using semi-empirical methods (Stewart et al. 2009; Hopkins et al. 2010a) are in relatively good agreement with each other, while predictions of the galaxy–galaxy merger rate obtained from various SAMs and hydrodynamic simulations show discrepancies of about an order of magnitude between them, as demonstrated in Hopkins et al. (2010b). In order to resolve these discrepancies, further work on galaxy merger rates using a priori models of galaxy formation is required. This approach has several advantages, such as providing insight into the physical mechanisms included in the models, making predictions in situations where observational data are unavailable, and accounting for merger time-scales self-consistently.

Observational estimates of the galaxy–galaxy merger rate have also not converged yet, although significant progress has been made in this direction (Lotz et al. 2011). For instance, in the case of massive galaxies ($M_* \gtrsim 10^{11} M_\odot$), some studies find an increasing redshift dependence (Bluck et al. 2009, 2012; Bundy et al. 2009; Man et al. 2012), while others find a nearly constant or even decreasing redshift evolution (Williams, Quadri & Franx 2011; Newman et al. 2012). Recently, Man, Zirm & Toft (2014) compared the consequences of selecting major mergers by stellar mass and by flux ratio, concluding that the former approach leads to a decreasing redshift dependence, while the latter results in the opposite. This appears to reconcile the differences between the observations by Bluck et al. (2009, 2012) and Man et al. (2012) with those by Williams et al. (2011) and Newman et al. (2012), where major mergers were selected according to their flux and stellar mass ratios, respectively. However, this is in conflict with the increasing redshift evolution observed for medium-sized galaxies ($M_* \gtrsim 10^{10} M_\odot$), as demonstrated in Lotz et al. (2011), as well as with predictions from semi-empirical models (Stewart et al. 2009; Hopkins et al. 2010a). The dependence on stellar mass of the galaxy–galaxy merger rate

is also a subject of some discussion, with some studies finding an increasing mass dependence and others the opposite (see Casteels et al. 2014, for a review).

We point out that the galaxy–galaxy merger rate cannot be measured directly from observations. Instead, the merger *fraction* must be estimated first, typically from observations of close pairs or morphologically disturbed galaxies, and then converted into a merger rate by adopting some averaged ‘observability’ time-scale (Lotz et al. 2011). Nevertheless, the merger rate and the merger fraction have many common features, such as their evolution with redshift (assuming that the observability time-scales do not change significantly with redshift). For this reason, we will sometimes use the two terms interchangeably when comparing to observations.

In this work, we study the galaxy–galaxy merger rate using the Illustris simulation, a hydrodynamic cosmological simulation carried out in a periodic box of ~ 106.5 Mpc on a side, which has been shown to reproduce many important properties of galaxies at $z = 0$ (Vogelsberger et al. 2014a, 2014c) as well as at higher redshifts (Genel et al. 2014). Because of the large volume covered by the simulation, the self-consistent treatment of baryons, and the physically motivated galaxy formation model used (Vogelsberger et al. 2013), the Illustris simulation provides a unique opportunity to study the galaxy–galaxy merger rate with unprecedented precision and physical fidelity.

This paper is organized as follows. In Section 2, we briefly describe the suite of simulations from the Illustris Project, as well as the methods used to identify haloes and galaxies. Section 3 presents the methodology used to construct merger trees of galaxies and DM haloes. The merger rate of DM haloes is calculated and compared to previous theoretical work in Section 4. We present the definitions and methods used to calculate the galaxy–galaxy merger rate in Section 5.1; in Section 5.2, we compare different approaches for estimating the mass ratio of a merger. We furthermore explore the dependence of the galaxy–galaxy merger rate as a function of descendant mass, progenitor mass ratio, and redshift in Section 5.3, and compare our results with previous work based on observations and semi-empirical models in Section 5.4. We finally present a fitting function for the galaxy–galaxy merger rate in Section 5.5. We discuss our results and present our conclusions in Section 6.

2 THE SIMULATIONS

2.1 Overview

The Illustris Project (Genel et al. 2014; Vogelsberger et al. 2014a,c) is a suite of hydrodynamic cosmological simulations of a periodic box of $75 h^{-1}$ Mpc ≈ 106.5 Mpc on a side, carried out with the moving mesh code AREPO (Springel 2010). A fiducial physical model has been adopted in these simulations, which includes star formation and evolution, primordial and metal-line cooling with self-shielding corrections, gas recycling and chemical enrichment, stellar supernova feedback, and supermassive black holes with their associated feedback. This model has been described and shown to reproduce several key observables in Vogelsberger et al. (2013), while its implications for galaxies across different redshifts have been discussed in Torrey et al. (2014). This model has also been used in hydrodynamic simulations of Milky Way-sized haloes (Marinacci, Pakmor & Springel 2013) and dwarf galaxies (Vogelsberger et al. 2014b).

The largest simulation from the Illustris project, Illustris-1 (also referred to as *the* Illustris simulation), follows the dynamical evolution of 2×1820^3 resolution elements (1820^3 DM particles and

approximately 1820^3 gas cells or stellar/wind particles), in addition to 1820^3 passively evolved Monte Carlo tracer particles. Two lower resolution simulations, Illustris-2 and Illustris-3, follow the dynamical evolution of 2×910^3 and 2×455^3 resolution elements, respectively. There are also DM-only variants of the simulations, known as Illustris-Dark-1, Illustris-Dark-2, and Illustris-Dark-3, which can be used to study the effects of baryons on DM haloes and subhaloes. Each simulation produced 136 snapshots between $z = 46$ and 0. The 61 snapshots at $z > 3$ are spaced with $\Delta \log_{10}(1+z) \approx 0.02$, while the 75 snapshots at $z < 3$ are spaced with $\Delta t \approx 0.15$ Gyr.

The cosmological parameters used throughout this paper are $\Omega_m = 0.2726$, $\Omega_\Lambda = 0.7274$, $\Omega_b = 0.0456$, $\sigma_8 = 0.809$, $n_s = 0.963$, and $h = 0.704$, which are consistent with the 9-yr *Wilkinson Microwave Anisotropy Probe* (WMAP) measurements (Hinshaw et al. 2013). Unless otherwise noted, all results presented in this paper are derived from Illustris-1.

2.2 Identifying the substructure

DM haloes are identified using the standard friends-of-friends (FoF) approach (Davis et al. 1985) with a linking length equal to 0.2 times the mean interparticle separation. The algorithm is applied to the DM particles, keeping only haloes with at least 32 DM particles. After this step, baryonic resolution elements are assigned to the same FoF group as their nearest DM particle. Substructure within the FoF groups is identified using an extension of the `SUBFIND` algorithm (Springel et al. 2001; Dolag et al. 2009), which can be applied to hydrodynamic simulations.

The original version of `SUBFIND` (Springel et al. 2001) estimates the density field using adaptive kernel interpolation and then identifies subhalo candidates as locally overdense regions. The boundary of each subhalo candidate is determined by the first isodensity contour that passes through a saddle point of the density field. Each subhalo candidate is then subjected to a gravitational unbinding procedure, so that the remaining structures are self-bound. Particles from satellite subhaloes which are dropped during the unbinding procedure are tentatively added to the central subhalo (also known as *background* halo) from the same FoF group, which is checked again for gravitational boundness at the end of the process.

In the version of `SUBFIND` used with `AREPO`, the density field is calculated for all particles and gas cells using an adaptive smoothing length corresponding to the distribution of DM particles around each point. Subhalo candidates are defined in the same way as before, but during the unbinding procedure the gas thermal energy is also taken into account. We keep subhaloes with at least 20 resolution elements (including gas and stars).

We point out that the stellar masses used throughout this paper are the ones given by `SUBFIND`, *without* truncating the particles found outside a fiducial radius equal to twice the stellar half-mass radius (Genel et al. 2014; Vogelsberger et al. 2014a). We find that using this alternative definition does not change the galaxy merger rate by more than 10 per cent.

3 CONSTRUCTING MERGER TREES

In this section, we describe the algorithms used to construct merger trees. The code for creating subhalo merger trees has been featured in the *Sussing Merger Trees* comparison project (Srisawat et al. 2013; Avila et al. 2014; Lee et al. 2014), where it is referred to as `SUBLINK`. Essentially, merger trees are constructed at the subhalo level using a methodology similar to the one described in Springel

et al. (2005) and Boylan-Kolchin et al. (2009), with slight modifications in the merit function used to determine the descendants, a different definition for the *first progenitor* (also known in the literature as the *main progenitor*), and a new method for skipping snapshots. Furthermore, merger trees can be constructed for different particle types, such as DM, stars, and star-forming gas, as explained below.

We define two varieties of merger trees: (1) *DM-only*, which follow exclusively the DM particles of a simulation, and (2) *baryonic*, which follow the star particles *plus* the star-forming gas elements in the simulation. A gas cell is considered to be star forming if its hydrogen particle density is above 0.13 cm^{-3} (Springel & Hernquist 2003). We note that although following a gas cell is not entirely equivalent to following a collisionless stellar or DM particle, the hydrodynamic scheme implemented in `AREPO` is quasi-Lagrangian, which means that the cells of the moving mesh follow the gas flow to a large extent. Therefore, we assume that star-forming gas cells, which are typically found in the central, denser regions of subhaloes, are able to preserve their ‘identity’ for durations of at least a few snapshots, and can therefore add valuable information when determining the descendant of a given subhalo. A less approximate treatment is in principle possible – although not done here – by following Monte Carlo tracer particles instead of gas cells (Genel et al. 2013; Nelson et al. 2013). We find that including star-forming gas besides only stellar particles is very useful for constructing robust merger trees at high redshifts, where galaxies have relatively large gas contents.

If a subhalo does not contain any stars or star-forming gas, then it does not exist in the baryonic merger trees. Conversely, a subhalo without any DM particles does not exist in the DM-only trees (although this situation is extremely rare). The DM-only and baryonic merger trees of the Illustris-1 simulation contain approximately 5×10^8 and 7×10^7 objects, respectively, taking all 136 snapshots into account. All results in this paper were obtained using the baryonic merger trees, with the exception of Section 4, where we present results about the merger rate of DM haloes rather than galaxies.

3.1 Finding the descendants

Each subhalo is assigned a *unique* descendant (if any) from the next snapshot, an approximation which is consistent with the hierarchical buildup of structure in Λ CDM cosmologies. This is done in three steps. Firstly, descendant candidates are identified for each subhalo as those subhaloes in the following snapshot that have common particles with the subhalo in question. Secondly, each descendant candidate is given a score based on the following merit function:

$$\chi = \sum_j \mathcal{R}_j^{-1}, \quad (1)$$

where \mathcal{R}_j denotes the binding energy rank of particles from the subhalo in question which are also contained in the descendant candidate. In the case of the baryonic merger trees, equation (1) is modified to include the mass m_j (taken at the same time as the binding energy rank \mathcal{R}_j) of the resolution elements:

$$\chi = \sum_j m_j \mathcal{R}_j^{-1}. \quad (2)$$

Thirdly, the unique descendant of the subhalo in question is defined as the descendant candidate with the highest score.

It is worth mentioning that the merit function presented in Boylan-Kolchin et al. (2009) features an exponent of $-2/3$ instead of -1 .

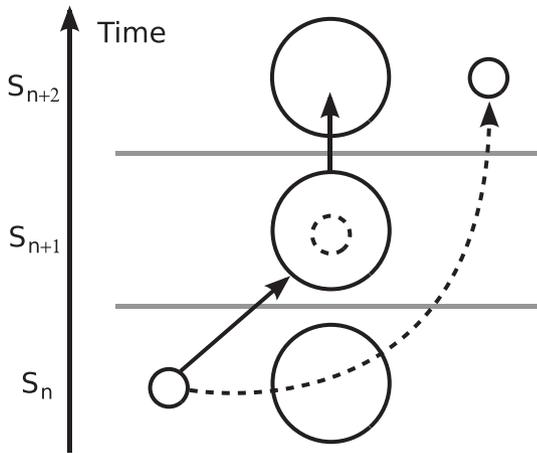


Figure 1. Illustration of snapshot ‘skipping,’ a simple approach for handling flyby events. The arrows indicate descendant links. A small subhalo identified at snapshot S_n is ‘lost’ during snapshot S_{n+1} because it is passing through a larger, denser object. In order to keep track of subhaloes in situations like this one, a descendant is also determined at snapshot S_{n+2} . If the ‘skipped’ descendant (dashed arrow) is not the same object as the ‘descendant of the descendant’ (solid arrows), then we define the ‘skipped’ one as the correct, unique descendant.

We find that an exponent of -1 allows the algorithm to follow subhaloes more robustly in major merger scenarios, particularly when three or more objects of comparable sizes and densities interact. Since the outer regions of subhaloes are subject to numerical truncation at the saddle points of the density field, as well as physical stripping, one should prioritize tracking the central parts of subhaloes, which are the ones that survive the longest. We find that the central few particles of a subhalo are remarkably stable over long periods of time. This is in agreement with previous work (Springel et al. 2001; Wetzel, Cohn & White 2009), in which reliable merger trees have been constructed by tracking only the 10–20 most bound particles of each subhalo.

Sometimes a halo finder cannot detect a small subhalo that is passing through a larger structure, simply because the density contrast is not high enough (see Fig. 1). We address this issue in the following way. For each subhalo from snapshot S_n , a ‘skipped descendant’ is identified at S_{n+2} , which is then compared to the ‘descendant of the descendant’ at the same snapshot. If the two possible descendants at S_{n+2} are not the same object, we keep the one obtained by skipping a snapshot since, by definition, it is the one with the largest score at S_{n+2} . This allows us to deal with flyby events, as long as the smaller subhalo is not ‘lost’ during more than one snapshot.

The similarity between the different merger tree algorithms compared in Srisawat et al. (2013) and Lee et al. (2014) suggests that allowing the search for descendants to extend over exactly two snapshots is enough for most cosmological simulations, which have relatively coarse snapshot spacings. However, a cosmological simulation with extremely high time resolution may require extending the search for descendants over more than two snapshots. This will be explored in future work using the small subboxes described in Vogelsberger et al. (2014a), which have 3976 snapshots each.

The validity of the single-descendant assumption can be investigated by quantifying ‘how hard’ it is to select the best descendant candidate. We did this in the following way. For each galaxy, we calculated the ratio between the ‘scores’ of the best and second-best descendant candidates, namely, $\xi = \text{score}(\text{second})/\text{score}(\text{first})$. We found that $\xi > 0.5$ ($\xi > 0.1$) in 1 per cent (4 per cent) of the

cases. This indicates that, although the approximation is certainly not perfect, in most cases the decision is an ‘easy’ one.

3.2 Merger trees of subhaloes and galaxies

We say that subhalo A is a progenitor (sometimes also called ‘direct’ progenitor, to distinguish it from earlier progenitors) of subhalo B if and only if subhalo B is the descendant of subhalo A . Note that a subhalo can have many progenitors, but at most a single descendant, an approximation motivated by the hierarchical buildup of structure in the Universe.

Once all the descendant connections have been made, as described in Section 3.1, the first progenitor of each subhalo is defined as the one with the ‘most massive history’ behind it (De Lucia & Blaizot 2007). This removes the arbitrariness in defining the first progenitor as simply the most massive one, which is subject to noise when the two largest progenitors have similar masses. As a result, the mass history of any particular galaxy or halo can be robustly compared across simulations carried out at different numerical resolutions or with variations in the physical model, as long as the initial conditions are the same.

Knowledge of all the subhalo descendants, along with the definition of the first progenitor, uniquely determines the merger trees. However, it is often convenient to rearrange this information into a more useful and physically motivated form. For example, one might be interested in retrieving the mass of a given object for all previous times, which would be a burdensome task if given the raw descendant information alone. We therefore construct merger trees in the following way. First, a linked-list structure is created for the whole simulation, so that each subhalo is assigned pointers to five ‘key’ subhaloes (Springel et al. 2005):

First progenitor: the progenitor of the subhalo in question, if any, which has the ‘most massive history’ behind it.

Next progenitor: the subhalo, if any, which shares the same descendant as the subhalo in question, and which has the next largest ‘mass history’ behind it.

Descendant: the unique descendant of the subhalo in question, if any.

First subhalo in FoF group: the main subhalo (defined as the one with the ‘most massive history’ behind it) from the same FoF group as the subhalo in question. Note that this link can point back to the subhalo under consideration.

Next subhalo in FoF group: the next subhalo from the same FoF group, if any, in order of decreasing ‘mass history.’

After this, the linked-list structure is stored in a depth-first fashion (Lemson & Springel 2006) into several files on a ‘per tree’ basis, where each tree is defined as a set of subhaloes that are connected by progenitor/descendant links or by belonging to the same FoF group. More specifically, two subhaloes belong to the same tree if and only if they can be reached by successively following the pointers described above. The resulting trees are completely independent from each other, which allows for easy parallelization of computationally expensive post-processing tasks, such as the construction of halo merger trees.

3.3 Merger trees of haloes (FoF groups)

Although most of the results in this paper were obtained using galaxy merger trees, we also construct halo (i.e. FoF group) merger trees in order to calculate the halo–halo merger rate. This quantity is relatively well constrained by theoretical models, so it can be used

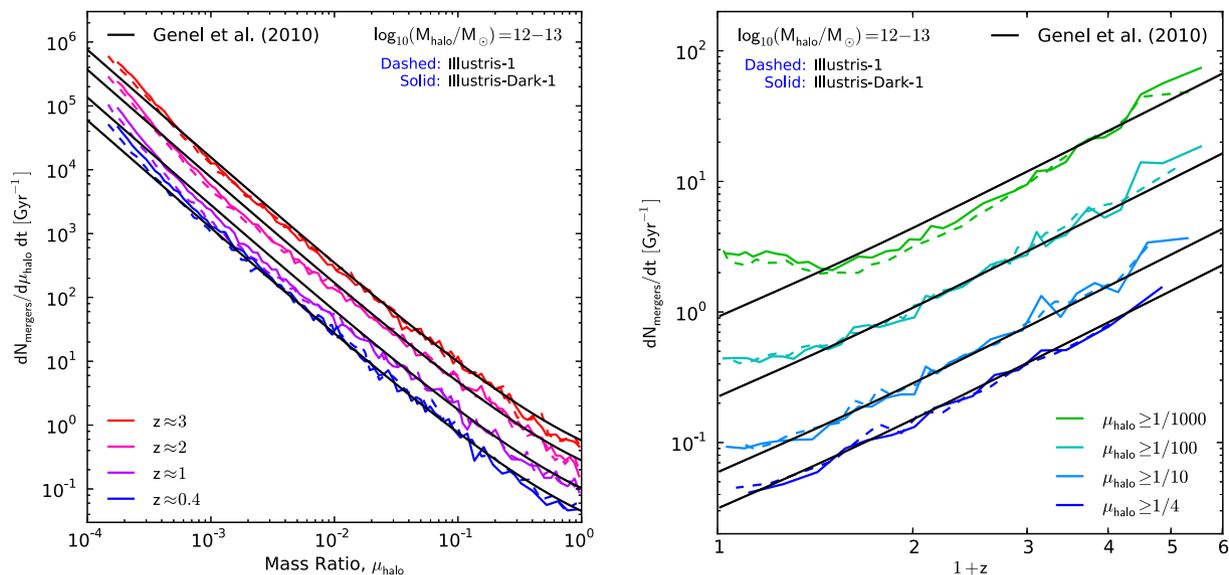


Figure 2. Left: the halo–halo merger rate as a function of the mass ratio μ_{halo} , shown for different redshifts. Right: the *cumulative* (with respect to mass ratio) halo–halo merger rate as a function of redshift, shown for different minimum mass ratios. Both panels correspond to mergers with descendant halo masses in the range $10^{12} \leq M_{\text{halo}}/M_{\odot} < 10^{13}$. The solid black lines are predictions from the fitting function given in Genel et al. (2010). The coloured dashed and solid lines correspond to the Illustris-1 (baryonic) and Illustris-Dark-1 (DM-only) simulations, respectively. The very good agreement between the dashed and solid lines indicates that baryons do not have a significant influence on the halo–halo merger rate. The increase in the merger rate seen at low redshifts is due to a limitation of the splitting algorithm as it approaches the final snapshots of the simulation, since spurious mergers can only be distinguished from real ones when there is a sufficient number of ‘future’ snapshots.

to validate some of our most basic results, as well as to assess the effects of cosmic variance on the cosmological volume used for this study, as discussed in Section 4.

Halo merger trees can contain fragmentation events in which a halo is split into two (or more) descendant haloes. These events arise because particles in a progenitor halo rarely end up in exactly one descendant halo; a decision therefore must be made to select a unique descendant halo. There is not a unique way to do this, and various algorithms have been proposed (see e.g. Fakhouri & Ma 2009, for a detailed comparison).

Here, we construct halo merger trees using the *splitting* algorithm (Fakhouri & Ma 2009, 2010; Genel et al. 2009, 2010). Instead of tracking the particles from each FoF group directly, this method takes the *subhalo* merger trees as input and constructs *halo* merger trees which are completely free of halo fragmentation events, as described below. The mass of each halo is defined as its *bound* mass, i.e. the combined mass of all particles gravitationally bound to its subhaloes, instead of the FoF group mass, which can contain a significant contribution from unbound particles.

Halo fragmentations are removed in the following way. For every tentative merger event between two haloes, the splitting algorithm checks whether the two haloes separate at a later time (as would happen in the case of a flyby), and, if that is the case, it then considers the two haloes as separate objects for all times. More specifically, for every halo at redshift z_{high} , the algorithm checks whether the halo contains at least one pair of subhaloes which at some lower redshift z_{low} do not belong to the same halo. Such halo would then be split in the following way: two subhaloes which belong to different haloes at z_{low} will also belong to different haloes at z_{high} , while subhaloes that stay together at z_{low} will also be together at z_{high} .

The splitting algorithm yields a new population of DM haloes and associated merger trees which are completely free from fragmentations, while leaving the DM halo mass function relatively unchanged (Genel et al. 2009).

4 THE HALO–HALO MERGER RATE

The merger rate of DM haloes has been studied extensively in previous work (e.g. Fakhouri & Ma 2008; Genel et al. 2009, 2010; Fakhouri et al. 2010, and references therein), with different theoretical predictions being similar within a factor of ~ 2 . Therefore, before calculating the galaxy–galaxy merger rate, we first verify that the halo–halo merger rate in Illustris is consistent with previous work.

Using the *splitting* method (Fakhouri & Ma 2009, 2010; Genel et al. 2009, 2010), *halo* merger trees were constructed by taking the DM-only *subhalo* merger trees as input (see Section 3.3). The resulting halo–halo merger rate is plotted in Fig. 2, both as a function of mass ratio for different redshifts (left) and as a function of redshift for different minimum mass ratios (right), for haloes with total *bound* masses (see Section 3.3) between 10^{12} and $10^{13} M_{\odot}$.¹ The solid black lines correspond to predictions from the fitting function provided by Genel et al. (2010). The coloured dashed and solid lines show the halo merger rate in the Illustris-1 (baryonic) and Illustris-Dark-1 (DM-only) simulations, respectively.

The very good agreement between the baryonic and DM-only Illustris runs in Fig. 2 indicates that baryons do not play an important

¹ The current implementation of the splitting algorithm supports a single particle type with a fixed mass. For this reason, the masses used to calculate the merger rate in Illustris-1 (dashed lines) actually correspond to the DM components rather than the total masses, which makes them smaller than their Illustris-Dark-1 counterparts (solid lines) by ~ 20 per cent (without taking baryonic effects into account). However, this difference is negligible for our purposes because of the weak mass dependence of the halo merger rate, $\sim M_{\text{halo}}^{0.15}$, which results in a change in the merger rate below 3 per cent. For comparison, the typical error bar size in both panels of Fig. 2, produced by Poisson noise in the number of mergers, is ~ 10 – 20 per cent. Thus, Fig. 2 would be essentially unchanged if we had used the total mass instead of the DM mass for Illustris-1 haloes.

role in the merger rate of haloes. Although not shown in this work, we have also calculated the halo merger rate for all the different feedback implementations described in Vogelsberger et al. (2013), as well as for the GADGET and AREPO runs described in Vogelsberger et al. (2012), which also resulted in no significant difference between any of them. This again shows that the halo merger rate is remarkably robust to different implementations of baryonic physics.

Fig. 2 shows that the halo–halo merger rate in the Illustris simulation is in excellent agreement with the formula provided by (Genel et al. 2010, except for redshifts $z \lesssim 0.4$, as discussed below). This is noteworthy given the fact that the best-fitting parameter values were obtained using the Millennium and Millennium II simulations (Springel et al. 2005; Boylan-Kolchin et al. 2009), which were carried out with cosmological parameters different from those in Illustris, and for which different subhalo merger trees were used as input for the splitting algorithm. In agreement with Genel et al. (2010), we find that the halo–halo merger rate scales with redshift as $\sim(1+z)^{2.3}$, with mass ratio as $\sim\mu_{\text{halo}}^{-1.7}$, and with descendant mass as $\sim M_{\text{halo}}^{0.15}$. These values are similar to the ones found by Fakhouri et al. (2010).

The good agreement between the halo merger rate in Illustris and the fit from Genel et al. (2010) also suggests that cosmic variance can be neglected in the 106.5 Mpc box used for this study, i.e. that the initial conditions used in the simulation are indeed representative of the large-scale density field. A detailed discussion about cosmic variance and the choice of initial conditions in Illustris can be found in Genel et al. (2014).

The increase in the merger rate seen at low redshifts is an unavoidable limitation of the splitting algorithm as it approaches the end of the simulation, since it becomes impossible to determine whether two recently merged haloes will ‘remain’ merged after $z = 0$, and therefore spurious mergers cannot be removed. For this reason, the calculated merger rate at $z \lesssim 0.4$ is overestimated and an extrapolation should be used instead. It is worth mentioning that analytic estimates of the halo merger rate based on the extended Press–Schechter formalism (Press & Schechter 1974; Bond et al. 1991; Lacey & Cole 1993; Neistein & Dekel 2008a,b) predict that the halo merger rate remains roughly a power law with respect to $(1+z)$ up to (and beyond) $z = 0$, which justifies the extrapolation used by Genel et al. (2009, 2010).

Finally, we point out that since we used the splitting algorithm to construct the halo merger trees (Section 3.3), the fitting formula from Genel et al. (2010) is the only analytical expression that can provide a meaningful comparison with previous work. If we had instead constructed halo merger trees using the *stitching* method (Fakhouri & Ma 2008), then the fit from Fakhouri et al. (2010) would be a better description of the resulting data. The halo merger rates obtained using these two methods can differ by up to a factor of 2 at $z \approx 0.4$ (Genel et al. 2009, figs 5 and 6).

5 THE GALAXY–GALAXY MERGER RATE

In this section, we describe how the galaxy–galaxy merger rate was calculated and explore its scaling as a function of descendant stellar mass, progenitor stellar mass ratio, and redshift. We also compare the merger rate with observations from the literature and provide a fitting formula which is reasonably accurate over a large range of masses, mass ratios, and redshifts.

We point out that the results about galaxy merger rates presented in this section were obtained directly from the *galaxy* merger trees (Section 3.2). Thus, they are independent from details about *halo* merger trees and rates (Sections 3.3 and 4).

5.1 Definitions

5.1.1 Merger

A merger takes place when a galaxy has more than one direct progenitor. Direct progenitors are usually found within the previous snapshot, but in some rare cases they are found two snapshots before, as discussed in Section 3.

We assume that all mergers are binary, which means that if a galaxy has N_p direct progenitors, we count $N_p - 1$ mergers, each between the first progenitor and each of the other ones. Fakhouri & Ma (2008) studied the effects produced by assuming binary versus multiple mergers and determined that for a low-redshift snapshot spacing of $\Delta z = 0.02$, the binary counting method was a good approximation for a wide range of halo masses and mass ratios. The low-redshift snapshot spacing in Illustris is $\Delta z \approx 0.01$, i.e. two times smaller than the value recommended by Fakhouri & Ma (2008), which means that Illustris is in the ‘safe’ regime with respect to the binary counting approximation.

Each merger is characterized by three parameters:

M_* : the stellar mass of the descendant immediately *after* the merger takes place.²

μ_* : the ratio between the stellar masses of the primary and secondary progenitors, taking both masses at t_{max} , defined as the moment when the secondary reaches its maximum stellar mass (see Section 5.2).

z : the redshift of the descendant snapshot.

Most halo finders have difficulty in correctly identifying subhaloes (or galaxies) during the final stages of a merger, which leads to ‘orphaned’ subhaloes during the construction of merger trees and a subsequent overestimation of the merger rate, since some merger events would be counted more than once. In order to avoid this, we only consider mergers which show a clear *infall* moment, that is, mergers for which both progenitors, followed back in time through their main branches in the merger trees, belonged to different FoF groups at some point in the past. This condition also becomes necessary in connection with the different definitions for the progenitor mass ratio discussed in Section 5.2.

5.1.2 Merger rate

The galaxy–galaxy merger rate describes the frequency of galaxy mergers as a function of descendant stellar mass M_* , progenitor stellar mass ratio μ_* , and redshift z . In this work, we focus on the merger rate *per galaxy*, which corresponds to the number of mergers per descendant galaxy, per unit time, per unit mass ratio. This quantity is typically given in units of Gyr^{-1} , and we denote it by

$$\frac{dN_{\text{mergers}}}{d\mu_* dt}(M_*, \mu_*, z). \quad (3)$$

In practice, equation (3) can be approximated in four steps: (i) defining bins in M_* , μ_* , and z , (ii) counting the number of mergers that fall into each bin, (iii) dividing by the average number of

² Although we could include star-forming gas in the mass of a galaxy, as we did when constructing the baryonic merger trees, for the rest of this paper we shall mostly be concerned with the stellar mass, since this quantity can be more directly compared to observations. This is consistent with our goal of quantifying the frequency of mergers, rather than the effects produced by them (in which case the gas content would indeed play an important role).

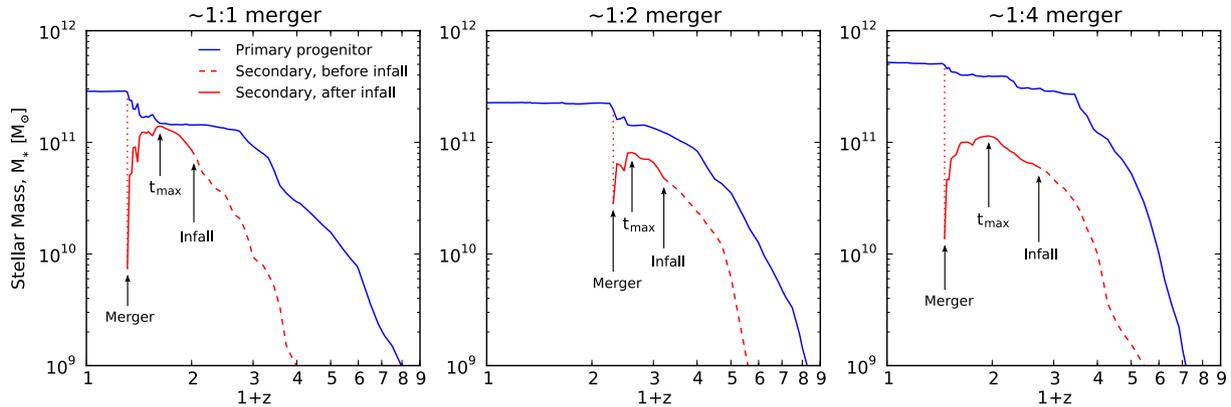


Figure 3. Stellar mass as a function of redshift, shown for galaxies undergoing mergers of different mass ratios (approximately 1:1, 1:2, and 1:4, from left to right). In each panel, the blue line corresponds to the main branch of a galaxy identified at $z = 0$, while the red line represents the main branch of a secondary galaxy that merges with the primary. The moment when the two galaxies merge is indicated with a vertical dotted line connecting the two progenitors. The secondary progenitor is drawn with a solid line when it is found inside the same FoF group as the primary, and with a dashed one when it is outside. In order to calculate the mass ratio of a merger, the masses of both progenitors are taken at t_{\max} , i.e. at the time when the *secondary* progenitor reaches its maximum stellar mass. Note that the mass ratio would be severely underestimated if the progenitor masses were taken right before the merger.

galaxies per snapshot for each corresponding bin, and (iv) dividing by the time interval, which is determined by the time difference between the snapshots that are located just before the edges of each redshift bin. Note that each redshift bin can contain more than one snapshot.

We make sure that each bin contains a minimum number of mergers (usually 5 or 10), so that bins are joined together when this is not the case. Additionally, we impose a resolution limit of at least 10 stellar particles for the smallest progenitor in each merger. The uncertainty in the calculated merger rate is determined by the Poisson noise from the number of mergers in each bin.³

Finally, since we are defining the first progenitor as the one with the ‘most massive history’ behind it, rather than as simply the most massive one, it is possible to have mass ratios greater than one. In these cases, we invert the mass ratio, so that we always have $\mu_* \leq 1$. We find that this minor correction has a negligible effect for all our results, with the resulting merger rate being indistinguishable from the one obtained by simply discarding mergers with $\mu_* > 1$ (the difference is much smaller than the uncertainty produced by the Poisson noise from the number of mergers).

5.2 The mass ratio of a merger

As mentioned above, the mass ratio of a merger is based on the stellar masses of the two progenitors taken at t_{\max} , i.e. at the time when the secondary progenitor reaches its maximum stellar mass. Here, we provide justification for this choice and explore other alternatives, such as taking the progenitor masses right before the merger and at virial infall.

Fig. 3 shows typical mass histories of galaxies that are undergoing mergers of different mass ratios, approximately 1:1, 1:2, and 1:4. Each panel shows (i) the moment when the two galaxies merge, (ii) t_{\max} , the time when the secondary progenitor reaches its maximum stellar mass, and (iii) the infall moment, i.e. the time when the secondary progenitor enters the same FoF group as the primary one.

³ In general, there are many more galaxies than mergers for any of the time-scales considered, so we neglect the error contribution from the number of galaxies in each bin.

In all panels, we observe that, shortly before the merger takes place, there appears to be an ‘exchange’ of mass between the primary and secondary progenitors. This is a consequence of how the halo finder imposes a distinction between centrals and satellites: even when two merging objects have nearly identical initial masses, one of them will be defined as the central subhalo and the other one as a satellite. Then, by construction, the central subhalo (or *background* halo) will be assigned most of the loosely bound matter residing in the FoF group, while the satellite will be ‘truncated’ by the saddle points in the density field. As a result, the central is typically much more massive than the satellite, even when the particle distribution of the two objects remains approximately symmetrical. This means that the mass ratio of a merger would be severely underestimated if we took the masses of the progenitors right before the merger.

Such effects are well known in the context of DM-only simulations. In particular, it has been observed that the DM mass of a satellite subhalo artificially correlates with its distance to the centre of the halo (e.g. Sales et al. 2007; Wetzel et al. 2009; Muldrew, Pearce & Power 2011). Here we show, however, that care must also be taken when considering the stellar content of a subhalo, despite the fact that stars are much more concentrated than DM and therefore less susceptible to numerical truncation.

Although phase-space halo finders seem to capture the masses of subhaloes more reliably during major mergers (see Avila et al. 2014, for a review), there is an additional reason why we avoid taking the progenitor masses immediately before a merger, which is to avoid effects from the merger itself, such as enhanced star formation and physical (as opposed to numerical) stripping, among other possible effects produced by mergers.

There are alternative definitions for the mass ratio of a merger. The two most relevant ones consist of taking the progenitor masses at t_{infall} , the time when the secondary progenitor enters the same FoF group as the primary one, and taking them at t_{\max} , the time when the secondary progenitor reaches its maximum stellar mass.

Fig. 4 shows the major merger rate as a function of descendant mass (left) and as a function of redshift (right), for the three mass ratio definitions mentioned so far, which are indicated with different colours. Additionally, we show with dashed lines the corresponding merger rates obtained by replacing each stellar mass with the

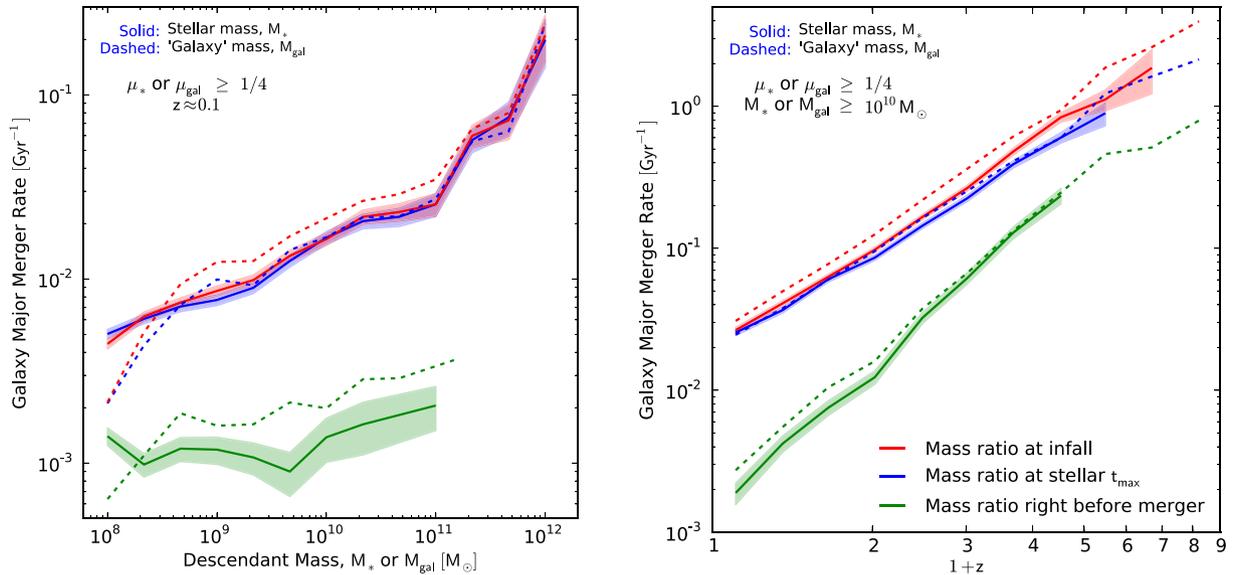


Figure 4. Left: major merger rate per galaxy as a function of descendant mass, for a redshift bin centred around $z \approx 0.1$. Right: major merger rate per galaxy as a function of redshift, for descendant galaxy masses greater than $10^{10} M_{\odot}$. The different colours show merger rates calculated by taking the mass ratio at different times, while the solid and dashed lines indicate merger rates calculated by using stellar and ‘galaxy’ (stars plus star-forming gas) masses, respectively. The shaded regions represent the Poisson noise from the number of mergers in each bin. We observe that taking the progenitor masses right before a merger can severely underestimate the major merger rate.

corresponding ‘galaxy’ (baryonic) mass, defined as the stellar *plus* star-forming gas mass of each galaxy.

Clearly, taking the galaxy masses right before a merger can underestimate the major merger rate by an order of magnitude or more. In fact, we find that both t_{infall} and t_{max} result in galaxy merger rates that are very well converged with resolution, while taking the progenitor masses right before a merger yields a merger rate that becomes smaller with increasing resolution. Indeed, as the resolution of a simulation is increased, two merging galaxies can be individually identified for a longer time before they finally merge, which means that they can get closer to each other, leading to a more extreme mass difference. This means that a major merger ($\mu_* \geq 1/4$) will appear to be a much more minor one by the time the merger actually takes place, which results in an underestimation of the major merger rate.

Another noticeable trend from Fig. 4 is that using baryonic masses instead of stellar masses results in slightly larger merger rates. This is a consequence of the decreasing fraction of cold gas as a function of stellar mass. Indeed, if we make the approximation $M_{\text{gal}} \propto M_*^{\alpha}$, where $\alpha < 1$, then a $\mu_* \gtrsim 1/4$ (major) merger in baryonic mass would correspond to a more minor one in stellar mass, which might not contribute to the major merger rate in this case.

In general, we observe that the merger rates obtained by taking the progenitor masses at t_{infall} and t_{max} are very similar, which is a consequence of the mass *ratio* being mostly unchanged between t_{infall} and t_{max} . However, even when the mass ratio is similar, the *masses* themselves can be very different across these two times, as a consequence of the large amount of star formation that can take place after infall. This can be seen in the three merger examples from Fig. 3, where the stellar mass grows by approximately a factor of 2 between t_{infall} and t_{max} (see also Sales et al. 2015, for a discussion about star formation in Illustris satellites after infall and their resulting colours). This suggests that taking the progenitor masses at t_{infall} is too early for making any meaningful comparison with observations of galaxy close pairs, which presumably involve

observations of galaxies that have already assembled most of their stellar mass.

To address the time delay between t_{infall} and t_{max} more generally, Fig. 5 shows the elapsed time since t_{max} and since t_{infall} for all merging (left) and surviving (right) satellites at $z = 0$. The bottom panels show the difference between the two times, which indicates that for the vast majority of satellites, t_{max} takes place a few Gyr after t_{infall} . The difference between these two time-scales is more pronounced and shows a smaller scatter in the case of merging satellites, which is partly explained by the fact that the surviving satellite population (right) includes galaxies which have been more recently accreted on to the halo, shifting Δt_{infall} downward (i.e. infall takes place at a later time) relative to the merging satellite population. Furthermore, it is less likely that newly accreted satellites have undergone increased star formation due to interactions with other galaxies, which shifts Δt_{max} upward (i.e. the maximum stellar mass was reached earlier) relative to merging satellites which, by definition, have already undergone such interactions.

All of this favours t_{max} over the other alternatives for the time when the merger mass ratio should be defined.⁴ By this time most of the stellar mass of the galaxy has already been formed, but it is also before numerical and physical effects from the merger itself begin to dominate. In other words, by taking the progenitor masses at t_{max} we are minimizing the bias from two different effects that tend to underestimate the stellar mass of the secondary progenitor, although for different reasons. We conclude that taking the progenitor masses at t_{max} is a reasonable choice for calculating both the merger rate *and* the stellar mass accretion rate, which will be the topic of upcoming work.

⁴ Perhaps another interesting alternative would consist of taking both progenitor masses at the time when the secondary enters the tidal radius of the interacting pair. However, such an alternative would be sensitive to the mass ratio between the primary and secondary progenitors, which, as we have seen, is largely influenced by details of the halo finding algorithm.

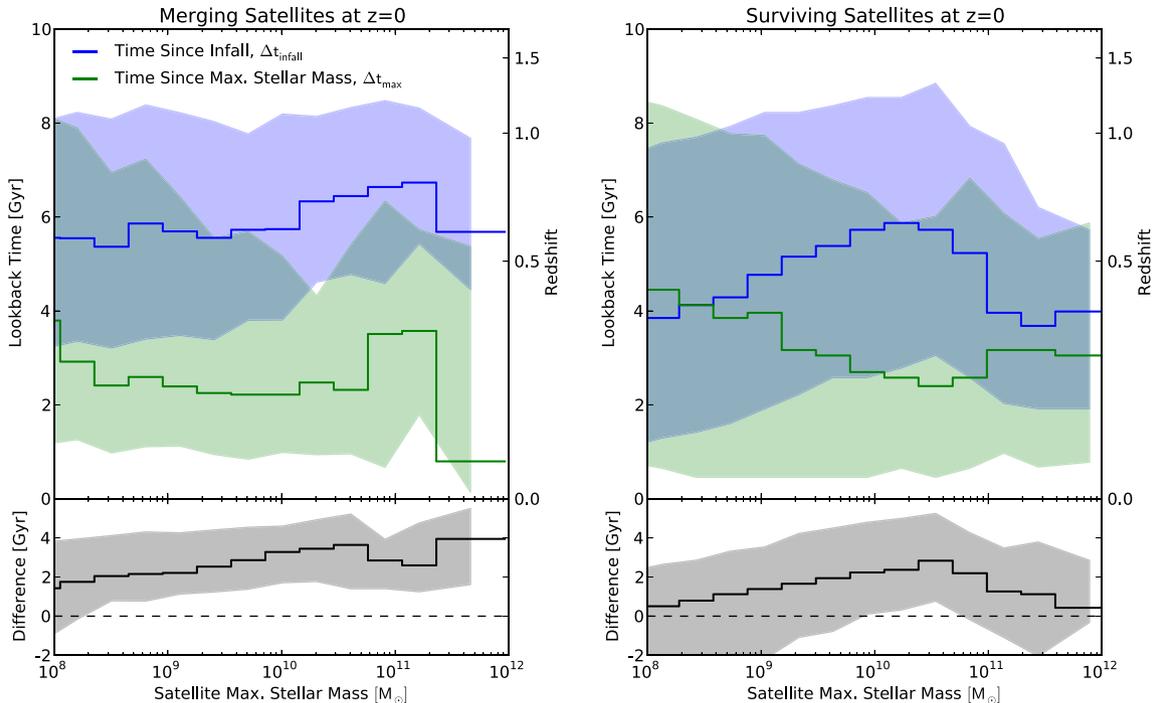


Figure 5. Left: the median elapsed times since virial infall (Δt_{infall} , blue) and since the moment of maximum stellar mass (Δt_{max} , green), shown for *merging* satellites at $z = 0$ as a function of maximum stellar mass. Right: the same for *surviving* satellites at $z = 0$. The bottom panels show the median of the difference between t_{infall} and t_{max} , calculated for each galaxy. The shaded regions indicate the range between the 16th and 84th percentiles, or approximately 1σ (note that the two shaded regions can overlap, which results in a darker colour). We observe that most satellites reach their maximum stellar mass a few Gyr *after* infall. The apparent sign reversal in the right-hand panels around $10^8 M_{\odot}$ happens simply because median values are not additive, so that median (t_{infall}) < median (t_{max}) does not necessarily imply that median ($t_{\text{infall}} - t_{\text{max}}$) < 0, and vice versa.

5.3 Results

In this section, we present the main features of the galaxy–galaxy merger rate as a function of descendant stellar mass M_* , progenitor stellar mass ratio μ_* , and redshift z . We consider both the ‘differential’ merger rate, which corresponds to mergers with mass ratios within a given interval, as well as the ‘cumulative’ merger rate, which includes all mergers with mass ratios greater than a given minimum value.

Fig. 6 shows the differential merger rate, given by equation (3), as a function of descendant mass (top) and as a function of mass ratio (bottom). The panels from left to right correspond to redshift bins centred around $z = 0.1, 1, \text{ and } 2$, respectively. The dashed black line corresponds to the fit from Table 1. We note that the merger rate has a relatively simple dependence on both M_* and μ_* . The dependence with respect to μ_* is well described by a power law, while the dependence on M_* can be modelled with a double power law with a break around $M_* \approx 2 \times 10^{11} M_{\odot}$. We note that the merger rate is always an increasing function of descendant galaxy mass: at low masses it grows as $\sim M_*^{0.2}$, which is very close to the mass dependence of the halo merger rate, $\sim M_{\text{halo}}^{0.13-0.15}$ (Fakhouri et al. 2010; Genel et al. 2010), and it steepens at $M_* \gtrsim 2 \times 10^{11} M_{\odot}$.

Close inspection of the bottom panels of Fig. 6 reveals that the lines corresponding to different descendant masses are not exactly parallel to each other. This feature is modelled by a ‘mixed’ term which includes both M_* and μ_* , and which is parametrized by γ (see Table 1). This means that, unlike with the halo merger rate, the galaxy merger rate is not separable with respect to descendant mass and mass ratio. This feature of the merger rate implies that more massive galaxies have a slightly larger relative contribution from more minor mergers, compared to less massive galaxies.

Fig. 7 is similar to Fig. 6, except that the merger rate is now ‘cumulative’ with respect to the mass ratio, i.e. it includes all mergers with mass ratios greater than a given μ_* , and the fitting function has been integrated accordingly (strictly speaking, it is not a fit anymore). We note that the enhancement in the galaxy merger rate above $M_* \approx 2 \times 10^{11} M_{\odot}$ becomes more noticeable after the merger rate has been integrated with respect to mass ratio. This feature is presumably a manifestation of the ‘turnover’ in the M_*-M_{halo} relationship, as explained in Hopkins et al. (2010b).

Fig. 7 can be useful for making quick assessments of the number of mergers that galaxies of a certain stellar mass are expected to undergo during a given time interval. For example, the major merger rate at $z \approx 0.1$ (blue line, upper-left panel) for Milky Way-like galaxies with $M_* \approx 6 \times 10^{10} M_{\odot}$ (McMillan 2011) is slightly larger than 0.02 Gyr^{-1} , which means that roughly one in every 50 Milky Way-like galaxies has undergone a major merger during the last Gyr.

Fig. 8 shows the redshift dependence of the cumulative (i.e. including all mergers with mass ratios larger than a given μ_*) galaxy merger rate. The left-hand panel shows the merger rate of galaxies with a fixed descendant mass $M_* \approx 10^{11} M_{\odot}$ for different mass ratio thresholds, while the right-hand panel shows the major ($\mu_* \geq 1/4$) merger rate for different descendant masses.

The right-hand panel from Fig. 8 demonstrates that the redshift dependence of the major merger rate becomes slightly weaker for more massive galaxies, as observed by Hopkins et al. (2010a) using semi-empirical methods. We find that the major merger rate of $M_* \approx 10^9 M_{\odot}$ galaxies has a redshift dependence proportional to $\sim (1+z)^{2.87}$, while the the major merger rate of $M_* \approx 10^{11} M_{\odot}$ galaxies evolves as $\sim (1+z)^{2.43}$. On the other hand, the left-hand panel

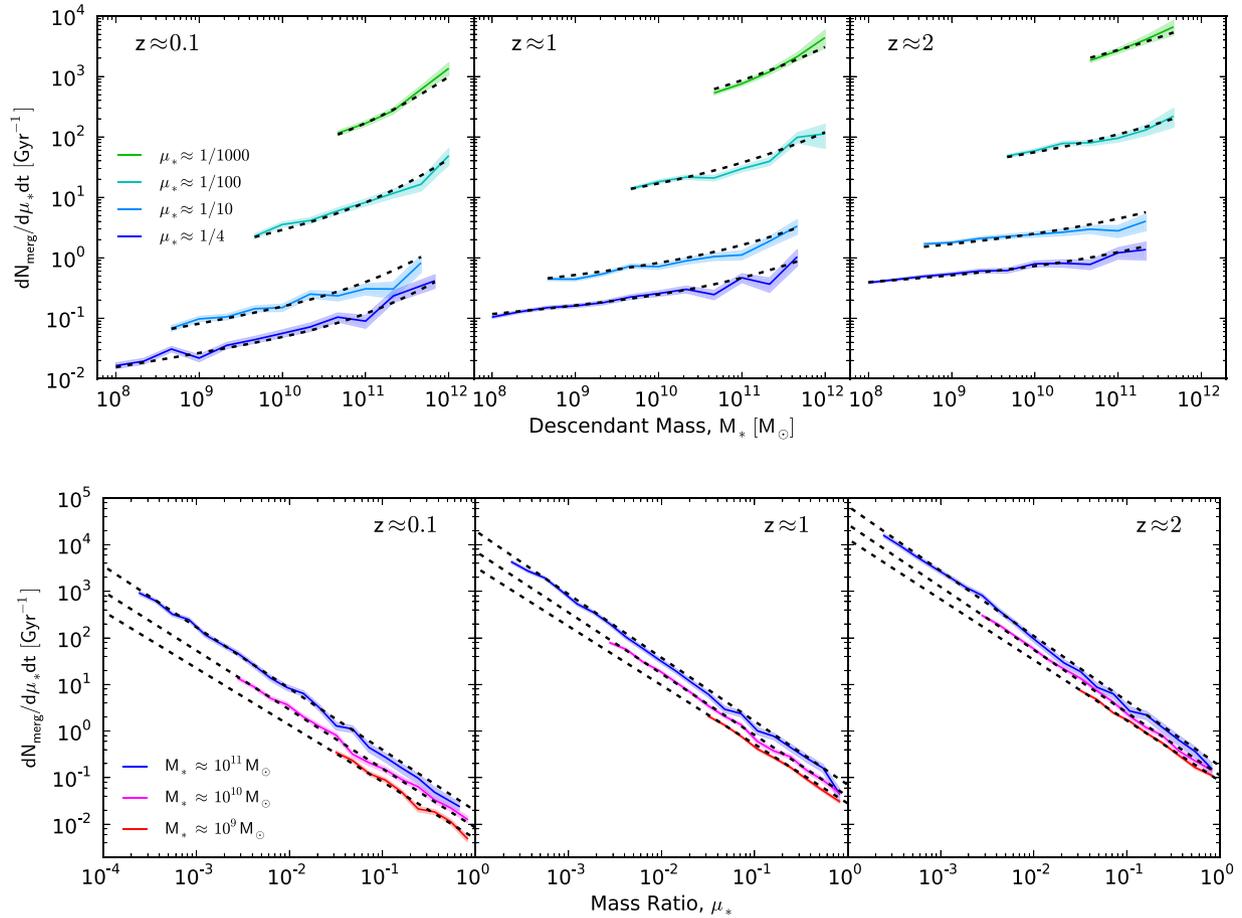


Figure 6. Top: the galaxy merger rate as a function of descendant mass M_* , for different mass ratios. Bottom: the galaxy merger rate as a function of mass ratio μ_* , for different descendant masses. The left-hand, centre, and right-hand panels correspond to redshift bins centred around 0.1, 1, and 2, respectively. The shaded regions indicate the Poisson noise in the number of mergers in each bin. The dashed black line represents the fitting function from Table 1.

Table 1. Fitting function and best-fitting parameters for the galaxy–galaxy merger rate (both M and μ correspond to stellar masses). See Section 5.5 for details.

Definition units	$\frac{dN_{\text{mergers}}}{d\mu dt}(M, \mu, z)$ Gyr $^{-1}$
	$A(z) \left(\frac{M}{10^{10} M_\odot} \right)^{\alpha(z)} \left[1 + \left(\frac{M}{M_0} \right)^{\delta(z)} \right] \mu^{\beta(z) + \gamma \log_{10} \left(\frac{M}{10^{10} M_\odot} \right)},$
	where
Fitting function	$A(z) = A_0(1+z)^\eta,$ $\alpha(z) = \alpha_0(1+z)^{\alpha_1},$ $\beta(z) = \beta_0(1+z)^{\beta_1},$ $\delta(z) = \delta_0(1+z)^{\delta_1},$ and $M_0 = 2 \times 10^{11} M_\odot$ is fixed.
$\log_{10}(A_0/\text{Gyr}^{-1})$	-2.2287 ± 0.0045
η	2.4644 ± 0.0128
α_0	0.2241 ± 0.0038
α_1	-1.1759 ± 0.0316
β_0	-1.2595 ± 0.0026
β_1	0.0611 ± 0.0021
γ	-0.0477 ± 0.0013
δ_0	0.7668 ± 0.0202
δ_1	-0.4695 ± 0.0440
χ_{red}^2	1.16

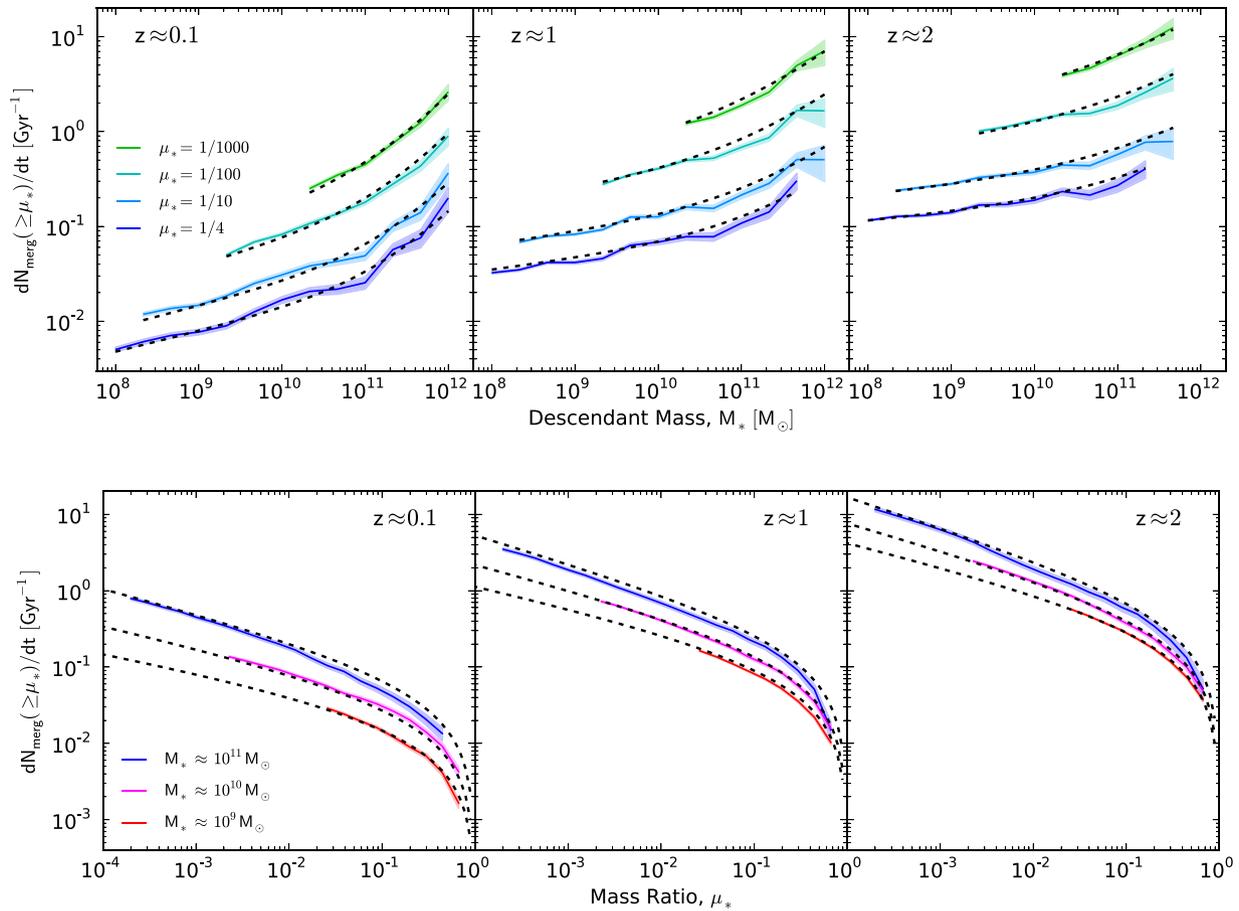


Figure 7. Top: the *cumulative* (with respect to mass ratio) merger rate as a function of the descendant mass M_* , shown for different minimum mass ratios. Bottom: the *cumulative* merger rate as a function of the mass ratio μ_* , shown for different descendant masses. The left-hand, centre, and right-hand panels correspond to redshift bins centred around 0.1, 1, and 2, respectively. The shaded regions indicate the Poisson noise in the number of mergers in each bin. The black dashed line is the *integral* of the fitting function from Table 1, integrated over the appropriate range of mass ratios μ_* (and therefore not a direct fit to the data shown in this figure).

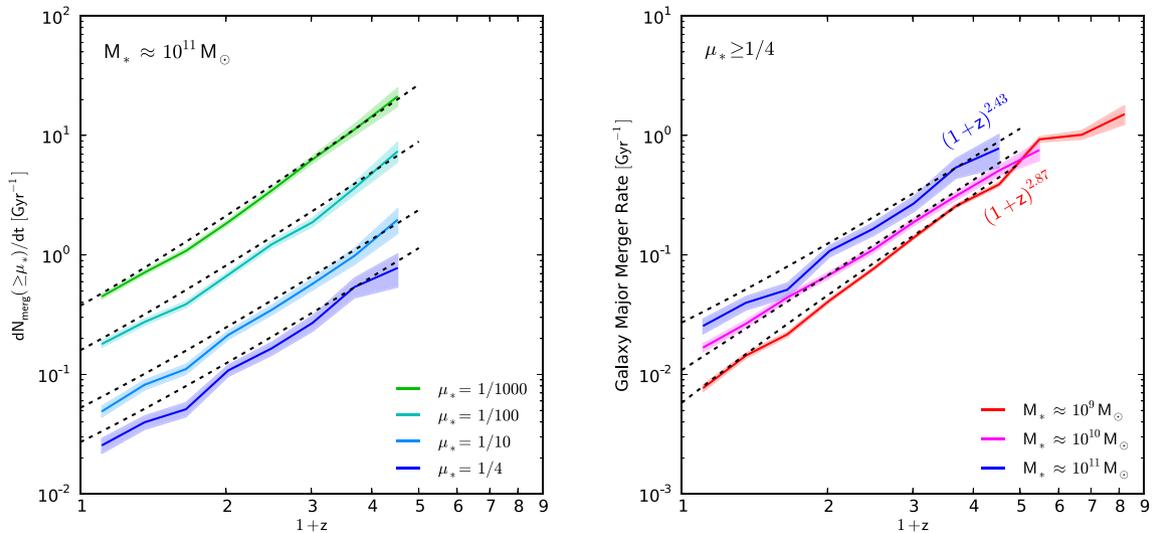


Figure 8. Left: the *cumulative* (with respect to mass ratio) merger rate as a function of redshift for descendant masses $M_* \approx 10^{11} M_\odot$, shown for a variety of minimum mass ratios. Right: the *cumulative* (with respect to mass ratio) merger rate as a function of redshift for mass ratios $\mu_* \geq 1/4$ (i.e. the major merger rate), shown for different descendant masses. The shaded regions indicate the Poisson noise in the number of mergers in each bin. The black dashed line represents the fitting function from Table 1, integrated over the appropriate mass ratio interval (it is therefore not a direct fit to the data shown in this figure).

from Fig. 8 shows that the slope of the merger rate with respect to redshift is practically independent of the mass ratio. In other words, the relative amount of major and minor mergers undergone by every galaxy (on average) is the same for all redshifts. In general, we find that the redshift dependence of the galaxy merger rate is very similar to the one of the halo merger rate, which evolves as $\sim(1+z)^{2.2-2.3}$ (Fakhouri et al. 2010; Genel et al. 2010).

5.4 Comparison to observations and semi-empirical models

In this section, we compare our main results with observational estimates of the galaxy merger rate, as well as with predictions from semi-empirical models. We do not include results from SAMs (e.g. Guo & White 2008) or hydrodynamic simulations (e.g. Maller et al. 2006) because their differences with respect to observations and semi-empirical models have already been studied in Hopkins et al. (2010b).

Fig. 9 shows the major ($\mu_* \geq 1/4$) merger rate of medium-sized ($M_* \geq 10^{10} M_\odot$, left) and massive ($M_* \geq 10^{11} M_\odot$, right) galaxies as a function of redshift. The blue, red, and green solid lines correspond to the different resolutions of Illustris, while the dot-dashed and solid black lines show predictions from the semi-empirical models of Stewart et al. (2009) and Hopkins et al. (2010a), respectively. These semi-empirical models disagree among themselves by factors of up to $\sim 2-3$, and our results from Illustris generally lie within this uncertainty range.

We point out that the galaxy merger rate in Fig. 9 is slightly different from the one in Fig. 8 because we now include all galaxies with stellar masses *larger than* a given value, rather than *around* a given value. This is done in order to have a more meaningful comparison with observations, which typically consider all galaxies with stellar masses (or luminosities) above a certain threshold. Additionally, the fitting functions from Stewart et al. (2009) and Hopkins et al. (2010a) represent slightly different quantities. The one from Hopkins et al. (2010a) describes the merger rate for all

galaxies with masses *larger than* a given value, while Stewart et al. (2009) provide three different versions of their fitting formula, with parameters corresponding to the mass ranges $10^{10} < M_*/M_\odot < 10^{10.5}$ (shown in the left-hand panel of Fig. 9), $10^{10.5} < M_*/M_\odot < 10^{11}$, and $M_* > 10^{11} M_\odot$ (shown in the right-hand panel of Fig. 9). Since the galaxy merger rate in both of these models (as well as in the current work) is an increasing function of descendant mass, the fit by Stewart et al. (2009) on the left-hand panel of Fig. 9 should be considered as a lower bound (although by less than 30 per cent, as a consequence of the weak mass dependence of the merger rate, and also because the number density of galaxies is dominated by less massive ones).

The left-hand panel of Fig. 9 also shows the range allowed by observations according to theoretical work by Lotz et al. (2011), where observational estimates of the major merger fraction from Kartaltepe et al. (2007), Lin et al. (2008), de Ravel et al. (2009), and Bundy et al. (2009) are converted into merger rates by means of ‘cosmologically averaged’ observability time-scales, which are determined from hydrodynamic merger simulations in combination with a galaxy formation model (Somerville et al. 2008). The corresponding galaxy merger rates predicted by Illustris are in good agreement with the predictions from Lotz et al. (2011), as well as with the semi-empirical models of Stewart et al. (2009) and Hopkins et al. (2010b), which are all allowed by the observational constraints.

The right-hand panel of Fig. 9, which corresponds to more massive galaxies ($M_* \geq 10^{11} M_\odot$), includes observational estimates of the merger rate based on merger fraction measurements by Bundy et al. (2009), Bluck et al. (2009, 2012), Williams et al. (2011), Man et al. (2012), and López-Sanjuan et al. (2012), which are shown as symbols with error bars. In all cases, we adopt the merger time-scales suggested by the authors, which are typically between 0.4 and 0.5 Gyr, except for the pair fraction observations of Williams et al. (2011) and López-Sanjuan et al. (2012), where we adopt a time-scale of 0.4 Gyr instead of the significantly larger suggested

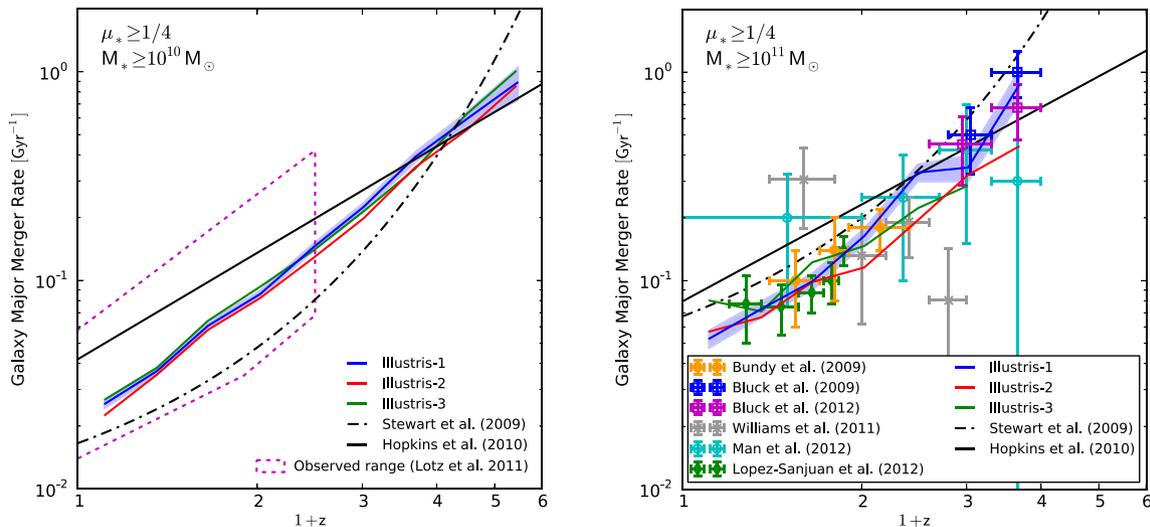


Figure 9. The galaxy major merger rate ($\mu_* \geq 1/4$) as a function of redshift, for descendant stellar masses greater than $10^{10} M_\odot$ (left) and $10^{11} M_\odot$ (right). The blue, red, and green lines correspond to the three resolution levels of Illustris. The shaded regions (Illustris-1 only) correspond to the Poisson noise from the number of mergers in each bin. Fitting functions from the semi-empirical models of Stewart et al. (2009) and Hopkins et al. (2010a) are indicated with dot-dashed and solid black lines, respectively. The magenta dashed range on the left-hand panel encapsulates the observational constraints for medium-sized galaxies ($M_* \gtrsim 10^{10} M_\odot$), determined from observations of the merger fraction by Kartaltepe et al. (2007), Lin et al. (2008), de Ravel et al. (2009), and Bundy et al. (2009), in combination with cosmologically averaged merger time-scales from Lotz et al. (2011). The right-hand panel includes different observational estimates of the merger rate for massive galaxies ($M_* \gtrsim 10^{11} M_\odot$), shown as symbols with error bars.

time-scales. Other observations of the merger fraction for massive galaxies which are not shown have been carried out by de Ravel et al. (2011), Newman et al. (2012), Xu et al. (2012), Ferreras et al. (2014), and Lackner et al. (2014).

In the case of massive galaxies, different authors find qualitatively different trends in the redshift evolution of the merger fraction: a decreasing redshift dependence (Williams, Quadri & Franx 2011; Ferreras et al. 2014), a nearly constant or mildly increasing redshift dependence (de Ravel et al. 2011; Man et al. 2012; Newman et al. 2012; Xu et al. 2012), or a strongly increasing redshift dependence (Bundy et al. 2009; Bluck et al. 2009, 2012; López-Sanjuan et al. 2012; Lackner et al. 2014). Recently, Man et al. (2014) attempted to resolve these differences by pointing out that studies in which major mergers are selected by flux ratio instead of stellar mass ratio tend to include very bright galaxies which nevertheless have very small masses, and should therefore not be counted as major mergers (using a stellar mass ratio definition). Therefore, Man et al. (2014) also support a decreasing redshift dependence, assuming that major mergers are selected by their stellar mass ratio. Yet, the Illustris simulation (as well as semi-empirical models) predict a strongly increasing redshift dependence, despite the fact that major mergers are also selected by stellar mass ratio.

The reason for this discrepancy is unclear at this stage. On the one hand, until observations converge to an agreed result better than a factor of ~ 2 , they will not be able to place significant constraints on modern theoretical models. On the other hand, considering that the halo–halo merger rate also exhibits a strong, positive correlation with redshift, we cannot envision any physical mechanism for which such trend should reverse in the case of galaxy mergers.

Fig. 10 shows the major ($\mu_* \geq 1/4$) merger rate of galaxies as a function of descendant stellar mass. As before, the three resolutions of Illustris are indicated with blue, red, and green solid lines, and predictions from the semi-empirical models of Stewart et al. (2009) and Hopkins et al. (2010a) are shown with dot–dashed and solid black lines, respectively. The black circles with error bars correspond to recent observational work on the mass-dependent merger rate by Casteels et al. (2014), based on observations of the fraction of highly asymmetric galaxies in the local Universe ($z \lesssim 0.2$), which are converted into merger rates by using the mass-dependent merger time-scales from Conselice (2006). The galaxy merger rate in Illustris is in good agreement with the observations by Casteels et al. (2014) for galaxies with stellar masses $M_* \gtrsim 10^{10} M_\odot$, although there is some disagreement below $\sim 10^{10} M_\odot$. We point out that the time-scales used by Casteels et al. (2014) require gas fraction measurements, which are only available for $M_* > 10^{10} M_\odot$. Therefore, an extrapolation has been used for $M_* < 10^{10} M_\odot$, which can introduce significant uncertainties into the corresponding observability time-scales.

The predictions from Hopkins et al. (2010a) appear to be larger than the ones from Illustris by a factor of ~ 2 – 5 . Part of this difference is explained by the fact that the model from Hopkins et al. (2010a) describes the merger rate for all galaxies with masses *larger than* a given value, while the other estimates in Fig. 10 correspond to galaxies with stellar masses *around* a given value. According to calculations with Illustris, this can account for a factor of ~ 2 at the low-mass end, $M_* \lesssim 10^9 M_\odot$, but the effects become less significant at higher masses. The remaining differences are possibly related to the merger time-scales involved, which are included self-consistently in Illustris (see Section 6).

Finally, we point out that different observational estimates for the mass dependence of the galaxy merger rate have also not converged yet, with some studies supporting an increasing mass dependence

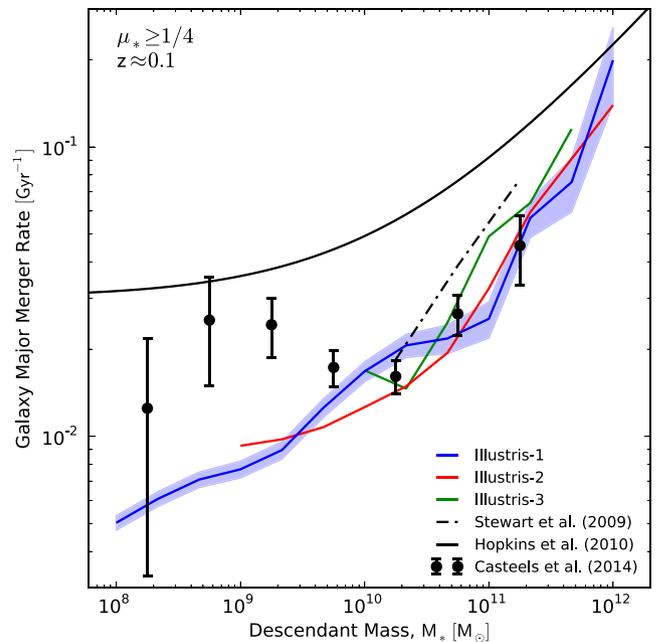


Figure 10. The galaxy major merger rate ($\mu_* \geq 1/4$) as a function of descendant stellar mass, for a redshift bin centred around $z = 0.1$. The blue, red, and green lines correspond to the three resolution levels of Illustris. The shaded regions (Illustris-1 only) correspond to the Poisson noise from the number of mergers in each bin. Fitting functions from the semi-empirical models of Stewart et al. (2009) and Hopkins et al. (2010a), evaluated at $z = 0.1$, are indicated with dot–dashed and solid black lines, respectively. The model from Hopkins et al. (2010a) has been scaled so that it corresponds to a major merger definition of $\mu_* \geq 1/4$ instead of $\mu_* \geq 1/3$. The black circles with error bars correspond to recent observations from Casteels et al. (2014).

and others suggesting the opposite (see Casteels et al. 2014, for a discussion). In fact, the observations by Casteels et al. (2014) are consistent with both an increasing and a decreasing mass dependence, depending on the stellar mass range considered. The Illustris simulation, on the other hand, always predicts an increasing mass dependence, which becomes steeper for larger galaxy masses.

5.5 A fitting formula

In Table 1, we provide a fitting formula for the galaxy–galaxy merger rate, along with the corresponding best-fitting parameters. For the sake of readability, we have dropped the asterisk subscript from the symbols M_* and μ_* . All masses and mass ratios in this section correspond to stellar masses.

We find that the galaxy–galaxy merger rate has a relatively simple dependence on the descendant mass M , the progenitor mass ratio μ , and the redshift z . The expression from Table 1 is qualitatively similar to the fitting function for DM halo merger rates presented in Fakhouri & Ma (2008), which is essentially a power law in M , μ , and $(1+z)$.

The main difference between the mathematical forms of the halo–halo and galaxy–galaxy merger rates is that the mass dependence steepens significantly at the high-mass end in the case of galaxies, such that it is better described by a double power law with a break around $2 \times 10^{11} M_\odot$. Furthermore, the exponents α and δ of the double power law exhibit some redshift dependence, which we parametrize as $\alpha(z) = \alpha_0(1+z)^{\alpha_1}$ and $\delta(z) = \delta_0(1+z)^{\delta_1}$. Both α_1 and δ_1 are negative, which means that the mass dependence of the

merger rate weakens with increasing redshift. This also means, as mentioned earlier, that the redshift dependence is stronger for lower mass galaxies. Additionally, the expression from Table 1 contains a ‘mixed’ term, parametrized by γ , that depends on both the stellar mass M and the mass ratio μ . This shows that the galaxy merger rate is not fully separable with respect to these variables, even for a fixed redshift.

The fits were carried out in log-space by minimizing a chi-squared merit function with a Markov chain Monte Carlo (MCMC) algorithm (Foreman-Mackey et al. 2013, <http://dan.iel.fm/emcee/current/>), considering all mergers which satisfy $M \geq 10^8 M_\odot$, $\mu \geq 1/1000$, and $z \leq 4$. The data points were obtained by creating bins in M , μ , and $(1+z)$ with widths corresponding to factors of 2, 1.2, and 1.1, respectively, and calculating the merger rate, along with the associated uncertainties, as explained in Section 5.1. In some cases the bins were rearranged so that there were at least five mergers per bin.

In all cases, the MCMC algorithm produced approximately Gaussian marginal distributions for each parameter. Therefore, we define the best-fitting value of each parameter as the mode of its marginal distribution, and the associated uncertainty as half the interval between the 16th and 84th percentiles, which corresponds to approximately 1σ . The resulting best-fitting parameters yield a reduced chi-squared statistic with a value of 1.16, which indicates that the model from Table 1 is a reasonably good fit to the data, without overfitting it.

6 DISCUSSION AND CONCLUSIONS

We have developed a theoretical framework for constructing and analysing merger trees of galaxies and DM haloes, which we apply to the Illustris simulation (Genel et al. 2014; Vogelsberger et al. 2014a,c) to make theoretical predictions for the merger rates of galaxies and DM haloes.

We find that the overall properties of DM halo merger trees and rates, which have been computed using the *splitting* method (Fakhouri & Ma 2009, 2010; Genel et al. 2009, 2010), are robust to baryonic effects and are also in very good agreement with previous theoretical work by Genel et al. (2010), who provided a fitting formula with parameters tuned to the Millennium and Millennium II simulations (Springel et al. 2005; Boylan-Kolchin et al. 2009). This agreement shows that the volume covered by the Illustris simulation can be considered to be ‘representative’ of the large-scale density field of the Universe.

The most novel aspect of this work pertains to the galaxy–galaxy merger rate, which we determine with unprecedented precision using a cosmological hydrodynamic simulation. We construct galaxy merger trees using an algorithm that has been shown to be reliable under a wide variety of circumstances (Srisawat et al. 2013; Avila et al. 2014; Lee et al. 2014). In particular, our merger trees are designed to track the innermost regions of subhaloes and galaxies, feature a robust definition of the first progenitor (i.e. the main progenitor), and avoid flyby events to some extent by allowing some objects to ‘skip’ a snapshot when finding a descendant.

When calculating galaxy merger rates, we argue that the most meaningful definition of the merger mass ratio consists in taking the two progenitor masses at the moment when the secondary progenitor reaches its maximum stellar mass. This happens, on average, a few Gyr after the secondary progenitor infalls into the same FoF group as the main progenitor. Additionally, we only consider mergers which have a well-defined *infall* moment, as explained in

Section 5.1. These definitions result in merger rates that are very well converged with resolution, as we show in Figs 9 and 10.

We find that the galaxy merger rate has a relatively simple dependence on descendant stellar mass, progenitor stellar mass ratio, and redshift, which is described by the fitting function given in Table 1. Essentially, this fit consists of a double power law with respect to stellar mass with a break around $\sim 2 \times 10^{11} M_\odot$, and single power laws for the mass ratio and redshift dependences. Some of the power-law exponents change with redshift, which results in a mass dependence that weakens with increasing redshift, or, equivalently, a redshift dependence that weakens with increasing mass. There is also a clear correlation between descendant mass and progenitor mass ratio, even at a fixed redshift, which implies that the galaxy–galaxy merger rate is not separable with respect to these variables, in contrast with the mathematical form of the halo–halo merger rate (e.g. Fakhouri & Ma 2008).

The strong, positive correlation with redshift found in this work is in disagreement with some observations of the major merger fraction for massive galaxies ($M_* \gtrsim 10^{11} M_\odot$), which find a nearly constant or decreasing evolution with redshift (Williams et al. 2011; Newman et al. 2012; Man et al. 2014). On the other hand, our results are in reasonable agreement with observations that suggest an increasing redshift evolution (Bluck et al. 2009, 2012; Bundy et al. 2009; Man et al. 2012).

For medium-sized galaxies ($M_* \gtrsim 10^{10} M_\odot$), the galaxy merger rate in Illustris is consistent with the general observational picture (e.g. Lotz et al. 2011). However, observational estimates of the merger rate must converge to a factor better than ~ 2 in order to distinguish predictions based on semi-empirical models (Stewart et al. 2009; Hopkins et al. 2010a) – which disagree among themselves by factors of up to ~ 2 – 3 – from those of Illustris, which typically lie inside this uncertainty range.

Observational work on the mass dependence of the merger rate has also not converged. We find good agreement with Casteels et al. (2014) for galaxies with stellar masses above $\sim 10^{10} M_\odot$, but find tension towards lower masses. This is possibly due to uncertainties in the observability time-scales assumed by Casteels et al. (2014), which require extrapolating the gas fraction for galaxies with stellar masses below $\sim 10^{10} M_\odot$, where observational data are unavailable.

As already mentioned, the galaxy merger rate in Illustris is in good qualitative agreement with predictions from semi-empirical models (Stewart et al. 2009; Hopkins et al. 2010b). However, it is worth noting that such models are designed to give reasonable agreement with observations by construction, without attempting to model galaxy formation from first principles. Because of this, they cannot be used to study the dependence of the galaxy merger rate and related quantities with respect to variations in physical models of galaxy formation. Additionally, semi-empirical models are generally not applicable in situations where observational data are scarce, such as for making predictions for the merger rate in the very minor merger regime ($\mu_* \lesssim 1/10$), at high redshifts ($z \gtrsim 3$), or when measurements of gas fractions are required.

State-of-the-art cosmological hydrodynamic simulations are better suited for such tasks. Furthermore, they have the advantage of handling merger time-scales self-consistently, which makes them ideal for measuring the galaxy–galaxy merger rate. For example, merger time-scales can be complicated by interactions with a third external object, which appears to be a fairly common occurrence (Moreno et al. 2013). Additionally, the final stages of a major merger are dominated by loss of angular momentum due to baryonic resonances and tidal torques (see Hayward et al. 2014, for a review),

which are difficult – or impossible – to describe using simple prescriptions for merger time-scales.

Previous attempts to measure the galaxy–galaxy merger rate using hydrodynamic simulations have yielded results which are significantly different from the ones presented in this work. In particular, the merger rate found by Maller et al. (2006) shows a much stronger dependence on descendant mass and redshift, which results in relatively poor agreement with observations and semi-empirical methods, as discussed in Hopkins et al. (2010b). More recently, Kaviraj et al. (2014) calculated the galaxy merger rate in the Horizon-AGN cosmological hydrodynamical simulation (Dubois et al. 2014) and found a nearly constant evolution with redshift, in disagreement with the results found by Maller et al. (2006), as well as with the ones from Illustris. These differences can be driven by various factors, including differences in star formation physics and AGN feedback (or lack thereof), details about the substructure finding algorithm and merger tree construction method, or different definitions when calculating the merger rate, most notably the progenitor mass ratio. In general, estimating the galaxy–galaxy merger rate using a priori models of galaxy formation is a non-trivial task.

The results presented in this paper are also in stark contrast with those found by Guo & White (2008), who applied the SAM proposed by De Lucia & Blaizot (2007) to the Millennium simulation (Springel et al. 2005) and found that the galaxy–galaxy merger rate has a strong dependence on stellar mass, but a weak one on redshift. In contrast, we find that it has a relatively weak dependence on stellar mass, but a strong one on redshift, which makes the galaxy–galaxy merger rate *qualitatively* similar to the halo–halo merger rate (except for the ‘knee’ in the mass dependence and the other features mentioned in Section 5.5). Interestingly, Guo & White (2008) also present an estimate of the halo–halo merger rate which is consistent with other theoretical calculations, including the one in this work. This implies that Guo & White (2008) find large *qualitative* differences between halo–halo and galaxy–galaxy merger rates, in disagreement with this work and with semi-empirical models. Some of these differences appear to be caused by satellite-specific prescriptions in the SAM of De Lucia & Blaizot (2007), in particular that galaxies cannot accrete gas after they have become satellites, and therefore cease to form stars once their supply of cold gas has been depleted.

The generally good agreement between the galaxy merger rate in Illustris and the one implied by observations comes with an important caveat: in order to convert a merger *fraction* into a merger *rate*, an observability time-scale has to be applied. This time-scale can shift the merger rate ‘vertically’ to larger or smaller values, introducing some arbitrariness in its normalization. Up to now, the most accurate observability time-scales have been determined from hydrodynamic merger simulations (e.g. Lotz et al. 2011), averaged in a ‘cosmological context’ by making several assumptions. Most importantly, a model of galaxy formation must be adopted in order to assign weights to the distribution of merger parameters at each redshift. Additionally, such merger simulations are usually considered to be in isolation, but, as mentioned above, Moreno et al. (2013) show that interactions between pairs of galaxies are often complicated by a third external object. These simplifying assumptions have a non-negligible effect on the merging time-scales of close pairs of galaxies, which affects estimates of the merger rate proportionately.

A more direct comparison with observations would consist of measuring the close pair fraction directly from the simulation, which could then be compared to observations without having to make assumptions about the merger time-scales involved. Unfortunately, this approach is complicated by the halo finder in situa-

tions where two large galaxies are found at very small separations ($\lesssim 20 h^{-1}$ kpc). Ultimately, the best approach may consist in creating synthetic images of galaxy surveys using the Illustris simulation (Torrey et al. 2015, Snyder et al., in preparation) and then applying the same source identification algorithms that are used with observational images. These topics will be explored in upcoming work.

Whereas it is reassuring that the *normalization* of the galaxy merger rate obtained in this work appears to agree well with observations, perhaps a more convincing indication of agreement is that the *slope* of the merger rate as a function of redshift follows the same trend as the range allowed by observational constraints for medium-sized galaxies ($M_* \gtrsim 10^{10} M_\odot$), which is proportional to $\sim(1+z)^{2.2-2.5}$. Although we cannot make a similar statement for more massive galaxies ($M_* \gtrsim 10^{11} M_\odot$) due to the *qualitative* disagreement between different observations of the merger fraction (Fig. 9, right-hand panel), the agreement with at least some of the sets of observations is encouraging. Additionally, the slope of the galaxy merger rate with respect to descendant mass is in good agreement with recent observations by Casteels et al. (2014) for galaxies with stellar masses above $\sim 10^{10} M_\odot$, where the observability time-scales used are more reliable. The body of these results shows that the Illustris simulation can be used to make realistic predictions about galaxy merger rates and related quantities. Further work on merger time-scales and mock galaxy surveys will lead to even more detailed comparisons between theoretical models of galaxy formation and observations of interacting and morphologically disturbed galaxies.

ACKNOWLEDGEMENTS

We are grateful to Asa Bluck, Gurtina Besla, Sarah Wellons, and Joshua Suresh for useful comments and discussions. We also thank Laura Blecha for carefully reading the manuscript. VS acknowledges support by the European Research Council through ERC-StG grant EXAGAL-308037. LH acknowledges support from NASA grant NNX12AC67G and NSF grant AST-1312095.

REFERENCES

- Avila S. et al., 2014, MNRAS, 441, 3488
 Bluck A. F. L., Conselice C. J., Bouwens R. J., Daddi E., Dickinson M., Papovich C., Yan H., 2009, MNRAS, 394, L51
 Bluck A. F. L., Conselice C. J., Buitrago F., Grützbauch R., Hoyos C., Mortlock A., Bauer A. E., 2012, ApJ, 747, 34
 Bond J. R., Cole S., Efstathiou G., Kaiser N., 1991, ApJ, 379, 440
 Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, MNRAS, 398, 1150
 Bundy K., Fukugita M., Ellis R. S., Targett T. A., Belli S., Kodama T., 2009, ApJ, 697, 1369
 Casteels K. R. V. et al., 2014, MNRAS, 445, 1157
 Conselice C. J., 2006, ApJ, 638, 686
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
 De Lucia G., Blaizot J., 2007, MNRAS, 375, 2
 de Ravel L. et al., 2009, A&A, 498, 379
 de Ravel L. et al., 2011, preprint (arXiv:1104.5470)
 Dolag K., Borgani S., Murante G., Springel V., 2009, MNRAS, 399, 497
 Dubois Y. et al., 2014, MNRAS, 444, 1453
 Fakhouri O., Ma C.-P., 2008, MNRAS, 386, 577
 Fakhouri O., Ma C.-P., 2009, MNRAS, 394, 1825
 Fakhouri O., Ma C.-P., 2010, MNRAS, 401, 2245
 Fakhouri O., Ma C.-P., Boylan-Kolchin M., 2010, MNRAS, 406, 2267
 Ferreras I. et al., 2014, MNRAS, 444, 906

- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Genel S., Genzel R., Bouché N., Naab T., Sternberg A., 2009, *ApJ*, 701, 2002
- Genel S., Bouché N., Naab T., Sternberg A., Genzel R., 2010, *ApJ*, 719, 229
- Genel S., Vogelsberger M., Nelson D., Sijacki D., Springel V., Hernquist L., 2013, *MNRAS*, 435, 1426
- Genel S. et al., 2014, *MNRAS*, 445, 175
- Guo Q., White S. D. M., 2008, *MNRAS*, 384, 2
- Hayward C. C., Torrey P., Springel V., Hernquist L., Vogelsberger M., 2014, *MNRAS*, 442, 1992
- Hinshaw G. et al., 2013, *ApJS*, 208, 19
- Hopkins P. F. et al., 2010a, *ApJ*, 715, 202
- Hopkins P. F. et al., 2010b, *ApJ*, 724, 915
- Kartaltepe J. S. et al., 2007, *ApJS*, 172, 320
- Kaviraj S., Devriendt J., Dubois Y., Slyz A., Welker C., Pichon C., Peirani S., Borgne D. L., 2014, preprint ([arXiv:1411.2595](https://arxiv.org/abs/1411.2595))
- Lacey C., Cole S., 1993, *MNRAS*, 262, 627
- Lackner C. N. et al., 2014, *AJ*, 148, 137
- Lee J. et al., 2014, *MNRAS*, 445, 4197
- Lemson G., Springel V., 2006, in Gabriel C., Arviset C., Ponz D., Solano E., eds, *ASP Conf. Ser. Vol. 351, Astronomical Data Analysis Software and Systems XV*. Astron. Soc. Pac., San Francisco, p. 212
- Lin L. et al., 2008, *ApJ*, 681, 232
- López-Sanjuan C. et al., 2012, *A&A*, 548, A7
- Lotz J. M., Jonsson P., Cox T. J., Croton D., Primack J. R., Somerville R. S., Stewart K., 2011, *ApJ*, 742, 103
- McMillan P. J., 2011, *MNRAS*, 414, 2446
- Maller A. H., Katz N., Kereš D., Dave R., Weinberg D. H., 2006, *ApJ*, 647, 763
- Man A. W. S., Toft S., Zirm A. W., Wuyts S., van der Wel A., 2012, *ApJ*, 744, 85
- Man A. W. S., Zirm A. W., Toft S., 2014, preprint ([arXiv:1410.3479](https://arxiv.org/abs/1410.3479))
- Marinacci F., Pakmor R., Springel V., 2013, *MNRAS*, 437, 1750
- Moreno J., Bluck a. F. L., Ellison S. L., Patton D. R., Torrey P., Moster B. P., 2013, *MNRAS*, 436, 1765
- Muldrew S. I., Pearce F. R., Power C., 2011, *MNRAS*, 410, 2617
- Neistein E., Dekel A., 2008a, *MNRAS*, 383, 615
- Neistein E., Dekel A., 2008b, *MNRAS*, 388, 1792
- Nelson D., Vogelsberger M., Genel S., Sijacki D., Keres D., Springel V., Hernquist L., 2013, *MNRAS*, 429, 3353
- Newman A. B., Ellis R. S., Bundy K., Treu T., 2012, *ApJ*, 746, 162
- Press W. H., Schechter P., 1974, *ApJ*, 187, 425
- Sales L. V., Navarro J. F., Abadi M. G., Steinmetz M., 2007, *MNRAS*, 379, 1464
- Sales L. V. et al., 2015, *MNRAS*, 447, L6
- Somerville R. S., Hopkins P. F., Cox T. J., Robertson B. E., Hernquist L., 2008, *MNRAS*, 391, 481
- Springel V., 2010, *MNRAS*, 401, 791
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
- Springel V. et al., 2005, *Nature*, 435, 629
- Srisawat C. et al., 2013, *MNRAS*, 436, 150
- Stewart K. R., Bullock J. S., Barton E. J., Wechsler R. H., 2009, *ApJ*, 702, 1005
- Torrey P., Vogelsberger M., Genel S., Sijacki D., Springel V., Hernquist L., 2014, *MNRAS*, 438, 1985
- Torrey P. et al., 2015, *MNRAS*, 447, 2753
- Vogelsberger M., Sijacki D., Kereš D., Springel V., Hernquist L., 2012, *MNRAS*, 425, 3024
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- Vogelsberger M. et al., 2014a, *MNRAS*, 444, 1518
- Vogelsberger M., Zavala J., Simpson C., Jenkins A., 2014b, *MNRAS*, 444, 3684
- Vogelsberger M. et al., 2014c, *Nature*, 509, 177
- Wetzel A. R., Cohn J. D., White M., 2009, *MNRAS*, 395, 1376
- Williams R. J., Quadri R. F., Franx M., 2011, *ApJ*, 738, L25
- Xu C. K., Zhao Y., Scoville N., Capak P., Drory N., Gao Y., 2012, *ApJ*, 747, 85

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